

Physical properties of simulated galaxy populations at $z = 2$ – I. Effect of metal-line cooling and feedback from star formation and AGN

Marcel R. Haas,^{1,2,3*} Joop Schaye,² C. M. Booth,^{2,4,5} Claudio Dalla Vecchia,⁶
Volker Springel,^{7,8} Tom Theuns^{9,10} and Robert P. C. Wiersma^{2†}

¹Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Rd., Piscataway, NJ 08854, USA

²Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA, Leiden, the Netherlands

³Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

⁴Department of Astronomy and Astrophysics, The University of Chicago, Chicago, IL 60637, USA

⁵Kavli Institute for Cosmological Physics and Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

⁶Max Planck Institute for Extraterrestrial Physics, Gießenbachstraße, D-85748 Garching, Germany

⁷Heidelberger Institut für Theoretische Studien, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁸Zentrum für Astronomie der Universität Heidelberg, Astronomisches Recheninstitut, Mönchhofstr. 12-14, D-69120 Heidelberg, Germany

⁹Institute for Computational Cosmology, Department of Physics, University of Durham, Science Laboratories, South Road, Durham DH1 3LE, UK

¹⁰Department of Physics, University of Antwerp, Campus Groenenborger, Groenenborgerlaan 171, B-2020 Antwerp, Belgium

Accepted 2013 August 7. Received 2013 July 22; in original form 2012 November 5

ABSTRACT

We use hydrodynamical simulations from the OverWhelmingly Large Simulations (*OWLS*) project to investigate the dependence of the physical properties of galaxy populations at redshift 2 on metal-line cooling and feedback from star formation and active galactic nuclei (AGN). We find that if the sub-grid feedback from star formation is implemented kinetically, the feedback is only efficient if the initial wind velocity exceeds a critical value. This critical velocity increases with galaxy mass and also if metal-line cooling is included. This suggests that radiative losses quench the winds if their initial velocity is too low. If the feedback is efficient, then the star formation rate is inversely proportional to the amount of energy injected per unit stellar mass formed (which is proportional to the initial mass loading for a fixed wind velocity). This can be understood if the star formation is self-regulating, i.e. if the star formation rate (and thus the gas fraction) increases until the outflow rate balances the inflow rate. Feedback from AGN is efficient at high masses, while increasing the initial wind velocity with gas pressure or halo mass allows one to generate galaxy-wide outflows at all masses. Matching the observed galaxy mass function requires efficient feedback. In particular, the predicted faint-end slope is too steep unless we resort to highly mass loaded winds for low-mass objects. Such efficient feedback from low-mass galaxies ($M_* \ll 10^{10} M_\odot$) also reduces the discrepancy with the observed specific star formation rates, which are higher than predicted unless the feedback transitions from highly efficient to inefficient just below $M_* \sim 5 \times 10^9 M_\odot$.

Key words: methods: numerical – galaxies: evolution – galaxies: formation – galaxies: fundamental parameters – cosmology: theory.

1 INTRODUCTION

Simulating the growth of dark matter (DM) haloes from initially small density perturbations through to the present day has become well established. Even the complex, non-linear stage of structure

formation can be predicted by means of high-resolution gravitational N -body simulations. The distribution of DM haloes derived from these simulations agrees very well with observations. The formation and evolution of galaxies is, however, much less well understood. Modelling the baryonic component is much more difficult than simulating the DM due to the collisional nature of the gas and the wealth of phenomena that need to be taken into account (cooling, star formation, feedback, etc.).

There are two popular approaches to tackle this challenging task. In semi-analytic models, simple descriptions of the behaviour of the

* E-mail: mhaas@physics.rutgers.edu

† Present address: Atomic Energy of Canada Limited, Chalk River Laboratories, Chalk River, Ontario K0J1J0, Canada.

baryonic component, as a function of the DM halo mass, merging history and environment, describe the evolution of gas and stars (e.g. Kauffmann et al. 1999; Somerville & Primack 1999; Cole et al. 2000). The freedom to choose functional forms and parameter values combined with the ability to run large numbers of models, ensure that reproducing observations is usually within reach. While this approach has great advantages, such as the ability to make mock galaxy surveys that are sufficiently realistic to reveal observational biases, there are also significant drawbacks. The large number of parameters can make it difficult to identify the key physical processes. More importantly, the ability to reproduce observations with a model that uses unphysical functional forms or unrealistic parameter values to describe physical processes can easily result in erroneous conclusions and misplaced confidence.

Complimentary to the semi-analytic models are models in which no explicit functional forms are assumed for physical processes. The so-called ‘abundance matching’ techniques link theoretically constructed DM halo populations to observed galaxy populations. These techniques rely on few assumptions, such as a correspondence between galaxy luminosity and DM halo mass (e.g. Kravtsov et al. 2004; Vale & Ostriker 2004). Although such models are relatively easy to construct and give a potentially robust match between DM properties and galaxy observables, all baryonic processes that govern galaxy formation are hidden in the assumed correspondence, so even though the consequences of galaxy formation physics can be characterized, it cannot be probed directly.

The other approach is to follow both the DM and the baryonic components by hydrodynamic simulation. While the DM is nearly always simulated using particles, baryons can either be modelled with Eulerian methods [discretizing the volume in an (adaptive) grid, e.g. Cen et al. 1990; Ryu, Vishniac & Chiang 1990 or using the Lagrangian approach also used for the DM (discretizing the mass using particles, e.g. Evrard 1988; Hernquist & Katz 1989; Thomas & Couchman 1992). Here, the freedom is limited to the parametrization of unresolved sub-grid processes. The high computational expense associated with full numerical simulations and the reduced level of freedom in the sub-grid modelling mean that, thus far, numerical simulations have been much less successful in reproducing observations of galaxy populations than semi-analytic models. Compared with the semi-analytic method, the advantages of the simulation approach include the much reduced (though still present) risk of getting the right answers for the wrong reasons, the ability to ask more detailed questions due to the tremendous increase in resolution, and the fact that not only galaxies, but also the intergalactic medium (IGM) is modelled.

Following the baryonic component in direct simulation is numerically expensive. There is therefore a trade-off between simulated volume and numerical resolution. Whereas, in this work we will simulate a representative cosmological volume, in order to get a large sample of galaxies, other works use a zoom-in technique to obtain high resolution in a couple of regions of interest (typically a handful of galaxies), e.g. Agertz et al. (2009); Ceverino et al. (2010); Ceverino et al. (2012). Such simulations, through their higher resolution, justify the use of more detailed physical models, at the price of a much smaller sample of simulated galaxies.

As many processes related to the baryons are not (well) resolved by even the highest resolution simulations, they are dealt with in the so-called sub-grid models. Among these are radiative cooling, the temperature and pressure of the multiphase gas at high densities [in the rest of the paper loosely called ‘the interstellar medium (ISM)’] and the formation of stars, the energy and momentum feedback by these stars into the ISM/intracluster medium/IGM, stellar mass-loss

and the growth of supermassive black holes (BHs) and associated feedback processes.

In this work, we employ cosmological, hydrodynamical simulations to investigate a number of basic baryonic properties of haloes, including: the (specific) star formation rate (SFR); the stellar, gas and baryon fractions; the gas-consumption time-scale and the galaxy stellar mass function. Reproducing observations in detail is not the main goal of this paper and so we have not attempted to tune our models or to optimize the sub-grid implementations to match any particular data set. Instead, we focus on understanding *how* different physical mechanisms shape the galaxy population.

We make use of a sub-set of the simulations from the *Overwhelmingly Large Simulations* project *OWLS* (Schaye et al. 2010), a large (~ 50 simulations) set of cosmological, smoothed particle hydrodynamics (SPH) simulations, each of the same volume of the universe, run with a wide variety of different prescriptions for sub-grid physical processes. The large variety of input physics in the *OWLS* runs allows us to investigate properties of haloes and their relation to the physical and numerical parameters. In this paper, we investigate the effect of feedback from star formation and from accreting supermassive BHs, as well as metal-line cooling. In a companion paper (Haas et al. 2013, hereafter Paper II), we investigate the effects of the assumed cosmology, the reionization history, the treatment of the unresolved, multiphase ISM, the assumed star formation law and the stellar initial mass function (IMF). These ingredients are all present in the simulations presented here, but the parameters are not varied. A sub-set of these simulations have been used by Sales et al. (2009, 2010) to investigate the dependence of the angular momentum and sizes of galaxies on the various feedback ingredients.

This work (together with Paper II) complements that of Schaye et al. (2010), where the cosmic star formation histories predicted by the *OWLS* models were analysed. The global SFR density can be decomposed into a DM halo mass function, which is determined by the cosmology, and the statistical distribution of the SFR as a function of halo mass. Here we will study the latter, which is astrophysically more relevant than the global SFR density as it removes the main effect of cosmology (the mass function) and allows us to investigate how the effects of the various baryonic processes vary with mass. Whilst we will add a dimension to the work of Schaye et al. (2010) by investigating the dependence on mass, we will remove another one in order to keep the scope of the study manageable. Thus, we limit ourselves to $z = 2$ and to the high-resolution series presented in Schaye et al. (2010) (these runs were halted at this redshift).

This paper is structured as follows: In Section 2, we describe the reference simulation used in this study (Section 2.1), how we define and extract haloes (Section 2.2) and how we compare the simulations to observations (Section 2.3). In Section 3, we describe how the properties of the galaxies in our reference simulation depend on both halo mass (Section 3.1) and galaxy stellar mass (Section 3.2). In Section 4, we describe how galaxy properties depend upon the physics included in the simulation. In this paper, we consider the effects of: metal-line cooling (Section 4.1), constant-energy supernova (SN) winds (Section 4.2), the effect of a top-heavy IMF in high-density gas (Section 4.3), ‘momentum-driven’ winds (Section 4.4), winds decoupled from the hydrodynamics (Section 4.5), thermally coupled SN-driven winds (Section 4.6) and the effect of feedback from active galactic nuclei (AGN) (Section 4.7). Finally, in Section 5 we summarize our findings and conclude. In Appendix A, we study the numerical convergence of our results using simulations of different volumes and with different numerical resolution

while in Appendix B we show that our results are insensitive to our particular choice of halo finder.

2 NUMERICAL TECHNIQUES

For a detailed discussion of the full set of *OWLS* models we refer the reader to Schaye et al. (2010). Here, we briefly summarize the reference model. Throughout this paper, we refer to this reference simulation as ‘*REF*’.

2.1 The reference simulation

We ran our simulations in periodic boxes of 25 comoving $h^{-1}\text{Mpc}$ with 512^3 DM and baryonic particles (which originally are collisional ‘gas’ particles, but can be converted into collisionless ‘star’ particles in the course of the simulation). We evolved the particles using an extended version of the *N*-Body Tree/SPH code *GADGET3* (last described in Springel 2005). The simulation particles have masses of $8.68 \times 10^6 h^{-1} M_{\odot}$ for DM and $1.85 \times 10^6 h^{-1} M_{\odot}$ for baryons (initially, the baryonic particle masses change in the course of the simulation due to mass transfer from stars back to the gaseous phase). In Appendix A, we show that our results are reasonably well converged with respect to the resolution and box size of the simulation. The gravitational softening length is initially fixed at 1/25 the inter-particle spacing in comoving coordinates (1.95 comoving h^{-1} kpc), but below $z = 2.91$ it is fixed at $0.5 h^{-1}$ kpc in proper units.

CMBFAST (Seljak & Zaldarriaga 1996) was used to generate initial conditions, that were evolved forward in time using the Zel’Dovich (1970) approximation from an initial glass-like state. The simulation is started at $z = 127$. The value of the cosmological parameters are $\Omega_m = 0.238$, $\Omega_{\Lambda} = 0.762$, $\Omega_b = 0.0418$, $h = 0.73$, $\sigma_8 = 0.74$ and $n_s = 0.951$. These values are derived from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) 3-yr data and largely consistent¹ with the 7-yr *WMAP* data (Komatsu et al. 2011). Because we will study the relative differences between different sub-grid models, the actual values of the cosmological parameters are not of major importance. The variation in sub-grid physics implementations is the main power of the *OWLS* set of simulations. The rest of this paper deals with variations of the sub-grid models for SN and AGN feedback, and their influence on the galaxy population. Therefore, we will first describe the parameters and sub-grid models used in the reference simulation.

The simulation explicitly follows the 11 elements H, He, C, N, O, Ne, Mg, Si, S, Ca and Fe. Radiative cooling and heating are calculated element-by-element in the presence of the Cosmic Microwave Background and the Haardt & Madau (2001) model for the UV/X-ray background radiation from quasars and galaxies, as described in Wiersma, Schaye & Smith (2009a). In these calculations, the gas is assumed to be in photoionization equilibrium and optically thin.

In the centres of haloes the pressure is so high that the gas is expected to be in multiple phases, with cold and dense molecular clouds embedded in a warmer, more tenuous gas. This multi-phase structure is not resolved by our simulations (and the simulations lack the physics to describe these phases), so we impose a polytropic effective equation of state for particles with densities $n_{\text{H}} > 10^{-1} \text{cm}^{-3}$. These particles are also assumed to be star forming, as this is the

density required to form a cold interstellar gas phase in the disc plane (under the assumption of hydrostatic equilibrium in the disc, this central density corresponds to a column density in Hydrogen of $\sim 10^{21} \text{cm}^{-2}$ Schaye 2004). We set the pressure of these particles to $P \propto \rho^{\gamma_{\text{eff}}}$, where γ_{eff} is the polytropic index and ρ is the physical proper mass density of the gas. In order to prevent spurious fragmentation due to a lack of numerical resolution we set $\gamma_{\text{eff}} = 4/3$, as then the ratio of the Jeans length to the SPH kernel and the Jeans mass are independent of density (Schaye & Dalla Vecchia 2008). The normalization of the polytropic equation of state is such that for atomic gas with primordial composition, the energy per unit mass corresponds to 10^4 K, namely ($P/k = 1.08 \times 10^3 \text{K cm}^{-3}$ for $n_{\text{H}} = 10^{-1} \text{cm}^{-3}$). The implementation of star formation is stochastic, as described in Schaye & Dalla Vecchia (2008), with a pressure-dependent SFR, obtained from local hydrostatic equilibrium and the observed Kennicutt–Schmidt law (KS-law Kennicutt 1998).

For the same 11 elements that we use for the cooling, we follow the production by asymptotic giant branch stars and by Type Ia and Type II (including Type Ib,c) SNe, as described in Wiersma et al. (2009b). The star particles are assumed to be simple stellar populations (SSPs) with a Chabrier (2003) IMF. SN feedback is implemented kinetically. After a short delay of 30 Myr, corresponding to the maximum lifetime of stars that end their lives as core-collapse SNe, newly formed star particles inject kinetic energy into their surroundings by kicking a fraction of their SPH neighbours in random directions. Each SPH neighbour i of a newly formed star particle j has a probability of $\eta m_j / \sum_{i=1}^{N_{\text{ngb}}} m_i$ of receiving a kick with a velocity v_w . Our reference model uses $\eta = 2$ and $v_w = 600 \text{km s}^{-1}$, which for our assumed IMF corresponds to ~ 40 per cent of the available energy from SNe being injected as winds (where all stars more massive than $6 M_{\odot}$ are assumed to explode in an SN with 10^{51} erg of energy). See Table 1 for the parameters that are varied in this paper.

2.2 Halo identification

Haloes are identified using a Friends-of-Friends (FoF) algorithm, which links together all DM particles which are closer to each other than the linking parameter ($b = 0.2$ times the mean interparticle distance). FoF identifies iso-overdensity contours of $\delta \equiv (\rho - \bar{\rho})/\bar{\rho} \simeq 3/(2\pi b^3) \simeq 60$ (Davis et al. 1985; Lacey & Cole 1994). Baryonic particles are linked to their nearest DM particle and belong to the same group, if any.

Following the convergence tests presented in Appendix A, we only include haloes that contain at least 100 star particles when considering halo properties as a function of stellar mass and we use a minimum of 2000 DM particles when we plot properties against halo mass. These two cuts produce nearly identical halo samples in the reference simulation and ensure that only well-resolved haloes are considered. In Appendix B, we confirm that our results are robust to changes in the definition of halo mass used. Physical properties of haloes (e.g. stellar mass and SFR) are just the sum of the properties of all constituent particles (in case of fractions like baryon fraction and specific SFR, it is the ratio of the sums of numerator and denominator).

Whenever we show the correlation between two halo properties, the plot consists of lines that connect the medians of bins evenly spaced in the quantity plotted along the horizontal axis. Each bin contains at least 30 data points. If there are fewer than 30 points in a given bin, it is extended until it includes 30 objects. The last bin may

¹ The only significant discrepancies are in σ_8 , which is 8 per cent, or 2.3σ , lower than the value favoured by the *WMAP* 7-yr data and the Hubble parameter, which is 1σ below the 7-yr value.

Table 1. Overview of the simulations and the input physics that is varied in this paper. Bold face indicates departures from the reference model. The first column gives the name of the simulation; the second column denotes the type of SN feedback, either kinetic, thermal or none. The third column shows the wind velocity for kinetic feedback models (the circular velocity is defined as $v_c = \sqrt{GM_{\text{vir}}/R_{\text{vir}}}$ and the velocity dispersion σ is related to the gravitational potential Φ : $\sigma = \sqrt{-\Phi/2}$); the fourth column indicates, for kinetic feedback models, the wind mass loading. The fifth column indicates whether the winds are decoupled from the hydrodynamics, the sixth column indicates whether metal-dependent cooling is followed in the simulation and the seventh column indicates the stellar IMF(s). The eighth column shows which simulations include AGN feedback and the last column specifies the section in which each simulation is discussed.

Name	Kinetic/thermal SN feedback	v_{wind} (km s ⁻¹)	$\eta = \frac{\dot{M}_{\text{wind}}}{M_*}$	Winds Decoupled?	Z Cooling	IMF	AGN Feedback?	Section
<i>REF</i>	Kinetic	600	2	no	yes	Chabrier	no	All
<i>NOSN_NOZCOOL</i>	None	n.a.	n.a.	n.a.	no	Chabrier	no	4.1
<i>NOZCOOL</i>	Kinetic	600	2	no	no	Chabrier	no	4.1
<i>WML4</i>	Kinetic	600	4	no	yes	Chabrier	no	4.2
<i>WML8V300</i>	Kinetic	300	8	no	yes	Chabrier	no	4.2
<i>WML4V424</i>	Kinetic	424	4	no	yes	Chabrier	no	4.2
<i>WML1V848</i>	Kinetic	848	1	no	yes	Chabrier	no	4.2
<i>WDENS</i>	Kinetic	$\sim c_s \propto \rho^{1/6}$	$\sim \rho^{-1/3}$	no	yes	Chabrier	no	4.2
<i>DBLIMFML14</i>	Kinetic	600	2, 14	no	yes	Chabrier, top-heavy	no	4.3
<i>DBLIMFV1618</i>	Kinetic	600, 1618	2	no	yes	Chabrier, top-heavy	no	4.3
<i>DBLIMFCONTSFV1618</i>	Kinetic	600, 1618	2	no	yes	Chabrier, top-heavy	no	4.3
<i>WVCIRC</i>	Kinetic	$5v_c/\sqrt{2}$	$150/(\sqrt{2}v_c)$	no	yes	Chabrier	no	4.4
<i>WPOTNOKICK</i>	Kinetic	3σ	$150/(\sqrt{2}v_c)$	no	yes	Chabrier	no	4.4
<i>WPOT</i>	Kinetic	5σ	$150/(\sqrt{2}v_c)$	no	yes	Chabrier	no	4.4
<i>WHYDRODEC</i>	Kinetic	600	2	yes	yes	Chabrier	no	4.5
<i>WTHERMAL</i>	Thermal	n.a.	n.a.	n.a.	yes	Chabrier	no	4.6
<i>AGN</i>	Kinetic	600	2	no	yes	Chabrier	yes	4.7

contain between 0 and 30 objects. We bin the data starting from the high-mass end. There, the difference in mass for two consecutive haloes is much larger than at the low-mass end, and in this way we are sure that the value of the mass at the high-mass end of the plots is always the mean of the mass of the 15th and 16th most massive systems.

2.3 Comparing simulations to observations

In this section, we describe how we compare our simulated galaxy population to observed stellar masses and SFRs. To convert observationally inferred stellar masses and SFRs from the cosmology assumed in the literature to our cosmology, we multiply them by the square of the ratio of luminosity distances, $[d_{L, \text{our cosm}}(z)/d_{L, \text{obs cosm}}(z)]^2$. The subscripts ‘our cosm’ and ‘obs cosm’ denote our cosmology and the cosmology under which the observations are transformed into masses/SFRs, respectively. Using the observational data sets compared to in this paper, this ratio varies from very close to 1 to ~ 2.6 .

Note that we are using FoF haloes, so all satellites are added to the central galaxy. We show in Appendix B that it makes very little difference to treat the satellites separately.

We compare our simulated SFRs to those observed by Daddi et al. (2007), who measured obscured and unobscured star formation by taking SFRs from both the UV and IR for *K*-selected *sBzK* galaxies (star forming, see Daddi et al. 2004) in the GOODS fields at $z \sim 2$. The median of the observed SFR as a function of stellar mass is well fit by $\text{SFR} = 250 \times (M_*/10^{11} M_\odot)^{0.9}$. The IMF assumed in the observations is the Salpeter (1955) IMF, whereas our stellar masses and SFRs are based on the Chabrier (2003) IMF. We therefore divide the observationally inferred SFRs by a factor of 1.65, which is the asymptotic (reached after only 10^8 yr) ratio of the number of ionizing photons per unit stellar mass predicted by Bruzual & Charlot (2003) for a constant SFR.

For stellar masses, the IMF conversion factor is more sensitive to the age of the population and the observed rest-frame wavelength. As the light in most wavelength bands is dominated by massive stars and the high-mass ends of both the Salpeter and Chabrier IMFs are power laws with very similar power-law indices, we use the same factor of 1.65 as we used for the SFRs. For very old populations observed in red wavelength bands (tracing stellar continua rather than dust emission) the conversion factor should be different. We verified that the K-band mass-to-light ratio is about a factor of 1.65 smaller for a Chabrier than for a Salpeter IMF for SSPs and constantly star-forming populations, for the full range of ages and metallicities available in the Bruzual & Charlot (2003) population synthesis package. We therefore also divide by a factor of 1.65 to convert stellar masses from the Salpeter to the Chabrier IMF.

We compare to the galaxy stellar mass function of Marchesini et al. (2009), a combined sample, using the deep near-infrared Multiwavelength Survey by Yale-Chile, the Faint Infrared Extragalactic Survey and the Great Observatories Origins Deep Survey-Chandra Deep Field South surveys at $z \sim 2$. Specifically, we compare to the $1/V_{\text{max}}$ results of Marchesini et al. (2009), including all of their uncertainties, with the exception of bottom-light IMFs. The reason for this choice is that models with bottom-light IMFs dominate the systematic errors and represent a more extreme assumption than the variations in the other quantities. Additionally, theoretical justifications for using bottom-light IMFs exist, thus far, only at high redshift (Davé 2008; van Dokkum 2008; Wilkins et al. 2008). We obtain the error bars on the observed data points by considering all of the sources of random errors included in the observational data points, which include: Poisson errors on the number counts, cosmic variance and the random errors from the use of photometric redshifts. We add the sources of random error in quadrature and, in addition, linearly add the maximum of the systematic errors in the same mass bins, just as Marchesini et al. (2009). The systematic errors include the systematic component in the errors from

photometric redshifts, errors arising from different population synthesis packages tested (Bruzual & Charlot 2003; Maraston 2005; Marchesini et al. 2009; Charlot & Bruzual, in preparation), varying the metallicities of the stellar populations, and the use of different extinction curves (Milky Way from Allen 1976, SMC from Prevot et al. 1984, Bouchet et al. 1985 and Calzetti et al. 2000).

Because our redshift of interest ($z = 2$) corresponds to the boundary between two of the redshift bins of Marchesini et al. (2009), which are $1.3 < z < 2$ and $2 < z < 3$, we weigh the averaging to the sizes of the redshift intervals (weights 1.2 and 0.8 Gyr, respectively), which gives results consistent with the $z = 2$ results of the Newfirm Medium-Band Survey (Marchesini et al. 2010). Additionally, the observed mass bins are not exactly the same size in both redshift intervals, although the differences are very small. Observed mass bins are of size 0.3 dex (at $1.3 < z < 2$) and 0.29 dex (at $2 < z < 3$). We interpolate the observed mass bins to a constant size of 0.3 dex. The resulting $1/V_{\max}$ estimate of the analysis of Marchesini et al. (2009) is shown as the yellow shaded regions in the bottom-right panels of Fig. 1, and Figs 3–9, labelled (I).

Marchesini et al. (2009) use a diet Kroupa IMF, for which the correction factor to our Chabrier IMF is very small (~ 1.03 , diet Kroupa being slightly more massive for the same flux). As this number is also derived from population synthesis packages, which come along with their own uncertainties, we chose not to convert masses for the small difference in IMFs. We do correct the masses for the difference in luminosity distances as described earlier. Number densities also need to be converted, as the volume at a given redshift is different for different angular diameter and comoving distances. Therefore, the number density (ϕ_*) is corrected for the ratio of volume elements (which is a function of the assumed cosmology).

When we compare our simulated galaxies to observations of the molecular gas mass in galaxies, we use a sub-set of the compilation used in Genzel et al. (2010). As their total range of redshifts is very large, we chose to use only the two sub-sets directly above and below $z = 2$. The near-IR long-slit $H\alpha$ sample from Erb et al. (2006) contains 11 galaxies with a mean redshift of 2.3 and is originally drawn from the BX selected sample of Steidel et al. (2004) and Reddy et al. (2005). The four galaxies with a mean redshift of 1.5 that we include are star-forming BzK galaxies from Daddi et al. (2008, 2010). For both samples, the molecular gas masses come from CO measurements with the Plateau de Bure interferometer and the stellar masses come from optical/UV SED fitting, under the assumption of a Chabrier IMF. We therefore only have to correct for the difference in luminosity distance, as explained above. Our simulations do not explicitly calculate the molecular fraction in the high-density gas. Nevertheless, the galaxies we compare to are at the very highest stellar mass end of our simulated galaxies. For those galaxies, the molecular gas fraction (by mass) is thought to be very high due to high pressures in the ISM, so we assume that all gas on the imposed equation of state is molecular. A recent study by Narayanan, Bothwell & Davé (2012) suggests that with an improved conversion factor from CO to H_2 , gas fractions of these galaxies are likely somewhat lower than quoted by Genzel et al. (2010).

3 PHYSICAL PROPERTIES OF SIMULATED GALAXIES AS A FUNCTION OF DARK MATTER HALO MASS

Fig. 1 shows, as a function of total halo mass the nine different galaxy properties we consider in this paper: medians of stellar

mass fraction ($f_{\text{star}} = M_{\text{star}}/M_{\text{tot}}$, panel A), SFR (panel B), baryon fraction ($f_{\text{baryon}} = M_{\text{baryon}}/M_{\text{tot}}$, panel C), fraction of star-forming gas ($f_{\text{ISM}} = M_{\text{ISM}}/M_{\text{tot}}$, panel D) and gas mass fraction ($f_{\text{gas,halo}} = (M_{\text{gas,total}} - M_{\text{ISM}})/M_{\text{tot}}$, panel E). Then, as a function of stellar mass: medians of the molecular gas mass in the ISM (panel F), specific SFR (sSFR = SFR/M_* , panel G), the inverse of the gas consumption time-scale ($\text{SFR}/M_{\text{ISM}}$, panel H) and galaxy number density (the galaxy stellar mass function, panel I).

In each panel of Fig. 1, we show the properties of the galaxies in the reference simulation with a black curve. The blue curves show results from the simulations analysed in Paper II, which include changes in the cosmology, ISM physics, star formation law and IMF. The red curves show the simulations discussed in this paper (see Section 4), which comprise changes in the feedback processes that are modelled. The red and blue lines are included in Fig. 1 to give a visual impression of the magnitude of the effects that we are considering. Here, we note that the magnitude of the differences between simulated galaxy properties due to uncertainties in the sub-grid physics are larger than the differences that arise from changing the hydrodynamics method from SPH to a (moving-) mesh code (Springel 2010), as e.g. shown by Vogelsberger et al. (2012); Kereš et al. (2012) and Scannapieco et al. (2012). The relative accuracy of different numerical techniques can nevertheless be important, but an investigation of this issue falls outside the scope of this paper.

The error bar in the lower-right corner of each panel (except the stellar mass function in panel H) indicates the typical scatter of galaxies around the median in the reference simulation. The error bar is a mass-weighted (by halo mass in panels A–E and by stellar mass in panels F–H) mean σ_{tot} of the standard deviation of all galaxies in all bins $\sigma_{\text{bin}}: \sigma_{\text{tot}} = \sum_{\text{bins}} \sigma_{\text{bin}} M_{\text{bin}} / \sum_{\text{bins}} M_{\text{bin}}$. We note that in all panels this scatter is dominated by the lower mass bins, where the scatter is typically much larger (and the bins are smaller and thus more numerous), and possibly partly due to resolution effects. Also, at the high-mass end the ‘scatter’ in the data is partially caused by the (almost exclusively positively sloped) relation between the dependent variable and mass. The scatter of galaxies around the median is typically smaller than the difference between the median relations in the most extreme simulations. A difference between two simulations that is smaller than the scatter is nevertheless still meaningful.

In the remainder of this section, we describe and explain the trends seen in the properties of the galaxies in the reference simulation, and comment briefly on the effects of the physics variations. Note that the first five panels (galaxy properties as a function of DM halo mass) and the last four panels (galaxy properties as a function of stellar mass) are resolved down to slightly different resolution limits, as discussed in Appendix A.

3.1 Properties as a function of halo mass

In panel (A) of Fig. 1, we plot the stellar mass fraction as a function of total halo mass. Stellar mass and total mass are tightly (and almost linearly) correlated, so differences between the models are emphasized by plotting the ratio of the masses. In the reference model, the fraction of the total mass locked into stars increases smoothly as a function of halo mass, from ~ 1 per cent in the smallest resolved haloes ($\log_{10}(M_{\text{tot}}/M_{\odot}) \sim 10.5$), to 4 per cent in the highest mass haloes. It is difficult to make a quantitative comparison with galaxy formation efficiencies quoted in much of the literature (e.g. Guo et al. 2010), as the definition of halo mass used in this study is different from the spherical overdensity halo definitions used in these works, but the qualitative trend is the same. As we move

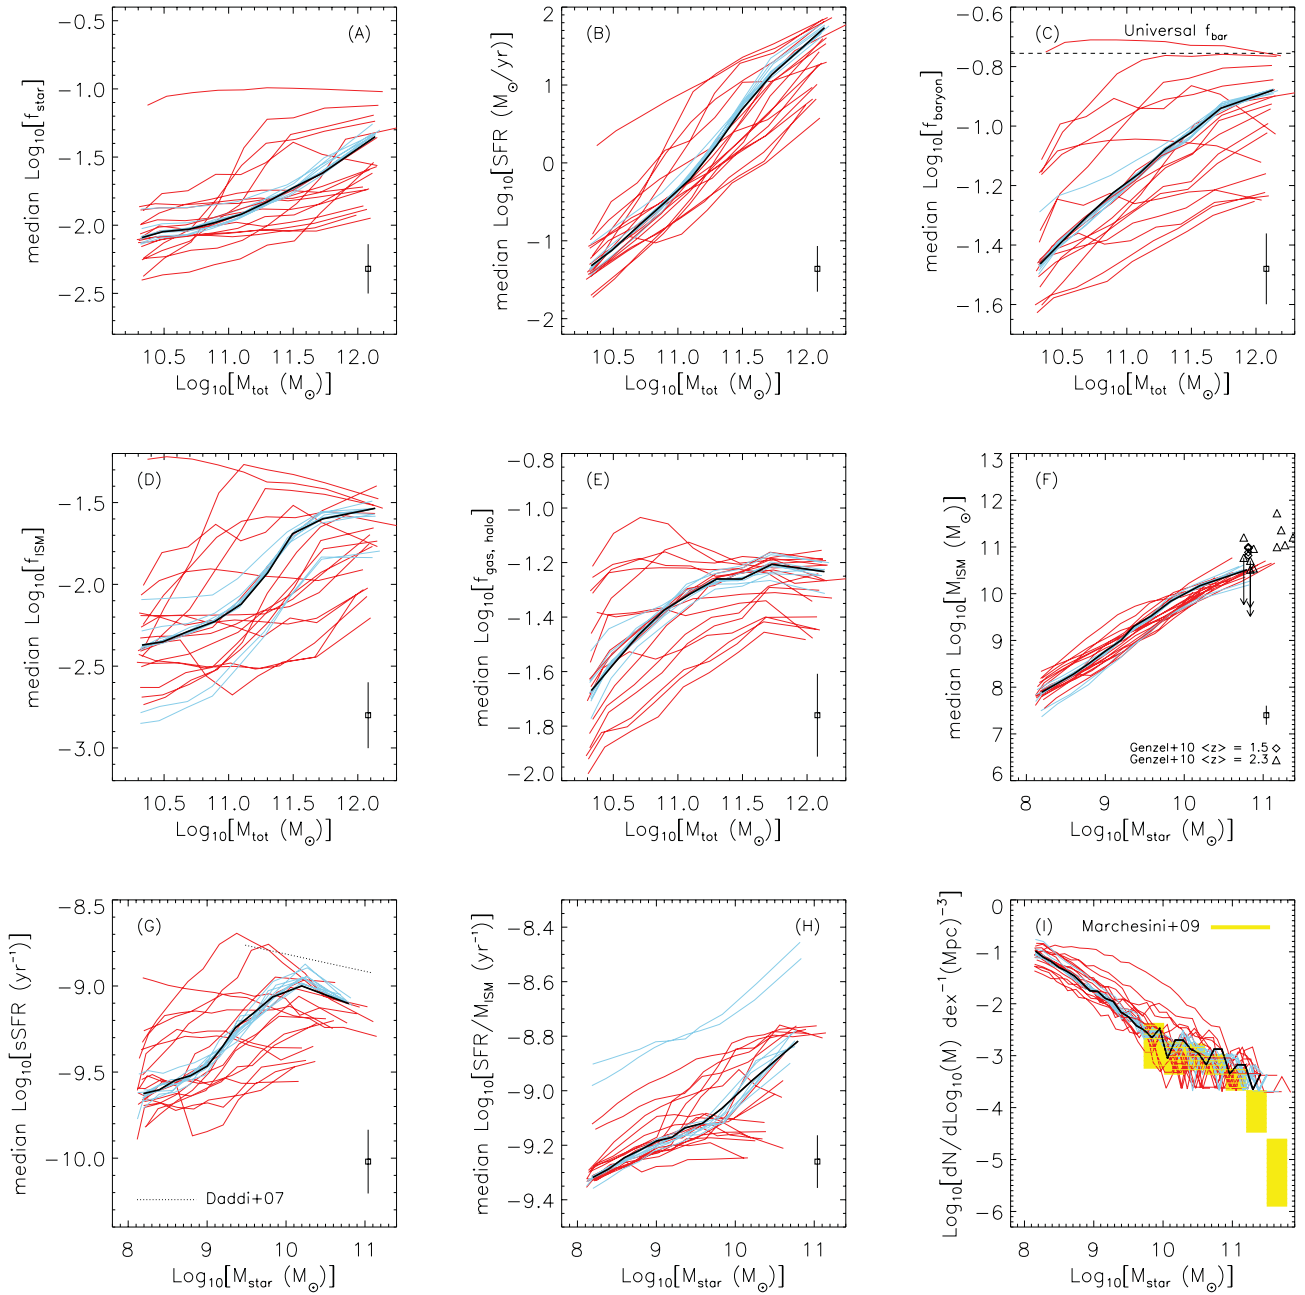


Figure 1. Median relations between halo properties in all the simulations used in this work (red lines) and in Paper II (blue lines). The reference model is shown as a black curve in each panel. In Section 4, we consider sub-sets of simulations in more detail. In the first five panels we show, as function of total halo mass, medians of stellar mass fraction ($f_{\text{star}} = M_{\text{star}}/M_{\text{tot}}$, panel A), SFR (panel B), baryon fraction ($f_{\text{baryon}} = M_{\text{baryon}}/M_{\text{tot}}$, panel C), fraction of star-forming gas ($f_{\text{ISM}} = M_{\text{ISM}}/M_{\text{tot}}$, panel D) and gas fraction ($f_{\text{gas,halo}} = (M_{\text{gas,total}} - M_{\text{ISM}})/M_{\text{tot}}$, panel E). The next four panels show, as a function of stellar mass, medians of the molecular gas mass in the ISM (panel F), specific SFR (sSFR = SFR/M_* , panel G), the inverse of the gas consumption time-scale ($\text{SFR}/M_{\text{ISM}}$, panel H) and galaxy number density (the galaxy stellar mass function, panel I). As described in the text, we show medians in bins along the horizontal axes for all haloes that satisfy the convergence criteria that apply to that specific panel. The horizontal, dashed line in panel (C) shows the universal baryon fraction for our chosen cosmology, the data points in panel (F) show a sub-set of the compilation studied in Genzel et al. (2010), the dotted black line in panel (G) shows the stellar mass – sSFR relation from the GOODS field (Daddi et al. 2007) and the shaded yellow region in panel (I) shows the galaxy stellar mass function of Marchesini et al. (2009). The error bars in the lower-right corners show the mass-weighted mean scatter of galaxies in the reference simulation, as explained in the text.

towards higher halo masses, a fractionally larger proportion of the mass is locked into stars. It is clear from inspection of panel (A) that the changes in physics considered in this paper (red curves) have a large effect on the $M_{\text{halo}}-M_*$ relation (much bigger than the simulations discussed in Paper II; blue curves). We discuss the reasons for this sensitivity in Section 4. The one simulation that is

a strong outlier, with much higher stellar mass at fixed halo mass, is the simulation that neglects SN feedback and, as such, contains no mechanism to prevent the catastrophic cooling of gas into stars. Most observational studies do not go out as far as $z = 2$ and rely on extrapolations beyond the observed galaxy population, especially at the low-mass end. Yang et al. (2012) study the stellar mass – halo

mass relation around $z = 2$, but their lowest observed point is at $M_{\text{halo}} \sim 10^{12} M_{\odot}$, which is our very highest mass bin. At that mass, the reference simulation slightly overpredicts the mass in stars (a factor of ~ 2), as expected from the inefficient SN feedback that will be discussed at length below. Some simulations with very efficient feedback underproduce stars at $M_{\text{halo}} \sim 10^{12} M_{\odot}$, as compared to these sub-halo abundance matching techniques.

In panel (B), we show the integrated SFR inside haloes as a function of their total mass. In all simulations, the SFR is an increasing function of halo mass in the mass range probed here ($10.5 < \log_{10}(M_{\text{tot}}/M_{\odot}) < 12.0$). As in panel (A), the vast majority of the physics variations considered in Paper II have very little effect, but the variations discussed here are very important. This implies that the sub-grid physics employed to model the ISM, other than that related to feedback, does not affect the large-scale, stellar properties of the galaxies. The outlier among the red curves in this panel is again the simulation that neglects SN feedback, which shows that feedback is an important component of the sub-grid models.

Panel (C) shows the baryon fractions of the haloes as a function of halo mass. The universal baryon fraction appropriate for our default cosmology ($\Omega_b/\Omega_m = 0.18$) is overplotted as a horizontal, dashed line. In the default simulation, feedback processes are able to eject baryons from the halo very easily in low-mass haloes, leading to low baryon fractions (0.6 dex below universal) for $M_{\text{tot}} \sim 10^{10.5} M_{\odot}$, whereas they become gradually less efficient as the halo mass increases. Again, the physics changes presented in this paper have a much more significant effect on this relation than those discussed in Paper II. Panels (D) and (E) show the fraction of the mass that is in gas in the ISM (which is sensitive to the physics variations in Paper II) and the rest of the gas in the haloes, respectively. In general, both are increasing functions of the total mass.

3.2 Properties as a function of stellar mass

The mass in the ISM as a function of the galaxy stellar mass is slightly sub-linear, as depicted in panel (F). This reflects the fact that at higher masses, gas is more efficiently transformed into stars. Our simulations slightly underpredict the molecular gas mass as a function of stellar mass when compared to the sub-sample of the galaxies studied by Genzel et al. (2010).

Both the stellar mass (panel A) and SFR (panel B) of galaxies are correlated almost linearly with halo mass, suggesting, that there is a linear relation between the stellar mass and the SFR. This linear trend masks important differences between the simulations, so to examine galaxy properties in more detail, we use the galaxy specific SFR ($s\text{SFR} = \text{SFR}/M_*$).

Panel (G) shows the median $s\text{SFR}$ as a function of galaxy stellar mass. In this panel, the black, dotted line shows the observations of Daddi et al. (2007). The scatter in the data is not shown, but is constant at approximately 0.2 dex. The scatter in the simulation data points is similar, although somewhat smaller ($\sim 0.1 - 0.15$ dex). We plot the observed relation over the mass range that is observed, and do not extrapolate beyond there.

The trend seen in the simulations is such that at low masses the $s\text{SFR}$ is an increasing function of mass, but at $\sim 10^{10} M_{\odot}$, the trend reverses and the $s\text{SFR}$ decreases with increasing stellar mass. At the highest masses the relation is declining more steeply than observed. As discussed in Section 3.1, most of the physics variations discussed in this paper have little effect on the integrated stellar properties of the haloes. The underestimate of the SFR is likely connected to the underestimate of molecular mass shown in panel (F).

Because SFRs are expected to be more strongly influenced by the mass of available star-forming gas than by the amount of stars already formed, we define a second normalized SFR, the ratio of the galaxy SFR to its star-forming gas mass. This is the inverse of the time needed to convert the present reservoir of star-forming gas into stars, assuming that the present SFR is maintained. We plot the inverse of the gas consumption time-scale as a function of stellar mass in panel (H). The gas consumption time is a monotonically decreasing function of stellar mass in the simulations. In the cores of the more massive haloes, gas reaches higher densities and pressure, so gas is converted more efficiently into stars.

In panel (I), we examine the galaxy stellar mass function: the number density of galaxies as a function of their stellar mass. The yellow, shaded regions show the data from Marchesini et al. (2009). When the observational uncertainties are taken into account, most of the simulated mass functions fall well within the observed range. We note that, at low masses, our simulated stellar mass function is steeper than most faint-end slopes of derived Schechter function parametrizations at similar redshifts (e.g. Conselice, Blackburne & Papovich 2005; Marchesini et al. 2009, 2010), but that this discrepancy exists largely outside the range of stellar masses where the mass function has been observed. Newer results are complete down to only slightly lower masses (e.g. Mortlock et al. 2011, who are complete down to $10^{9.5} M_{\odot}$). Our simulation box size is too small to form the rare objects that populate the exponential cutoff of the stellar mass function.

Comparing the behaviour of the reference simulation between panels (G) and (I) reveals an interesting behaviour: although the galaxy $s\text{SFR}$ is a factor of a few too low at $z = 2$ over the observed range, the simulations form enough galaxies of all masses, and possibly too many low-mass systems (a similar behaviour was also noted by Choi & Nagamine 2012). This may indicate that the observations are not internally consistent, such that the integral of observed SFRs do not add up to the observed current stellar mass after correcting for stellar mass-loss. A perhaps more likely explanation is that galaxy formation in the simulation is too efficient in low-mass systems at high redshift and that other physics variations are required to flatten the stellar mass function and to get the SFRs at $z = 2$ to agree with observed values.

4 ISOLATING THE EFFECTS OF THE INPUT PHYSICS

In this section, we discuss each of the variations of the input physics. Table 1 summarizes the simulations discussed in this paper and indicates in which sub-section each one is discussed. Bold face values indicate differences between each simulation and ‘REF’. In this paper, we discuss how metal-line cooling and the inclusion of kinetic SN feedback affect galaxy properties (Section 4.1); how the parameters of the kinetic SN feedback model affect the results (Section 4.2); the effect of a top-heavy IMF at high pressures (Section 4.3); winds with ‘momentum-driven’ scalings that depends on galaxy properties (Section 4.4); the effects of decoupling the winds from the hydrodynamics (Section 4.5); the effect of injecting SN energy thermally (Section 4.6); and the effects of strong, AGN feedback (Section 4.7).

A graphical representation of the gas density of a galaxy formed in a representative set of models is shown in Fig. 2. The galaxy resides in a halo of total mass $\sim 10^{12.5} M_{\odot}$. It was first identified in the ‘REF’ simulation, where its position (defined as the centre of mass of all particles within 10 per cent of the virial radius) is

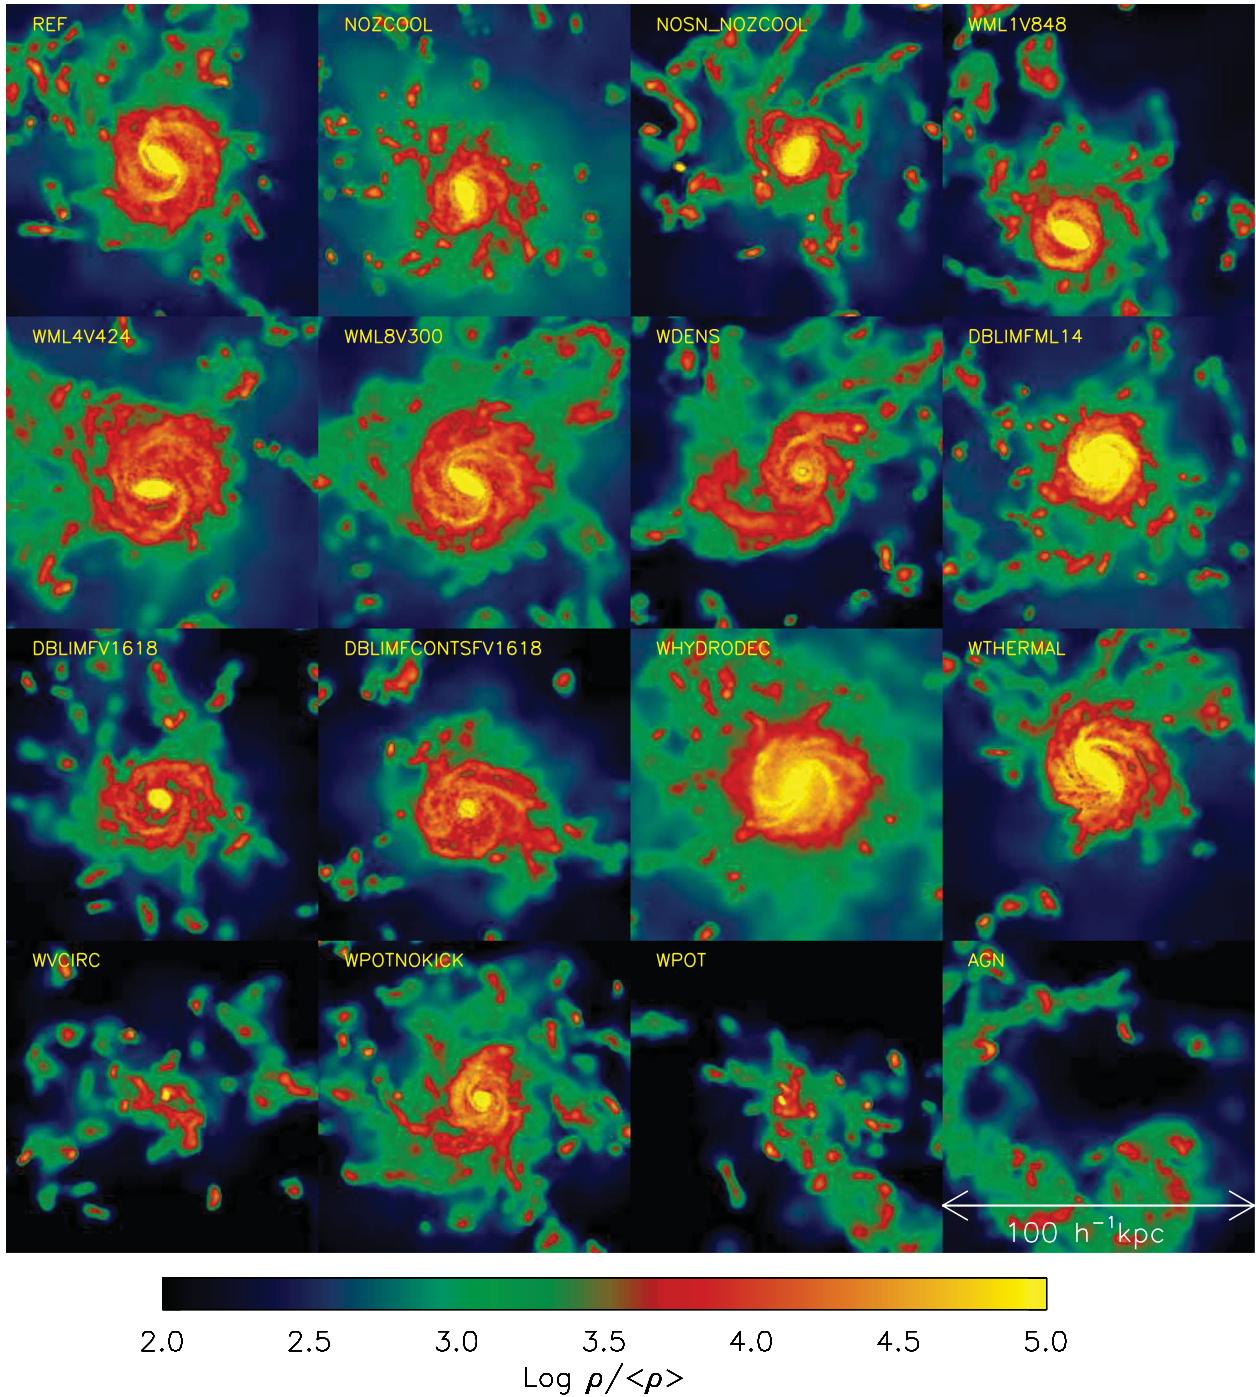


Figure 2. A graphical representation of a galaxy in a halo of $10^{12.5} M_{\odot}$ in 15 of our simulations at redshift 2. The colour coding denotes the gas density in a slice of $100 h^{-1}$ kpc thickness, divided by the mean density of the universe. All frames are 100 comoving h^{-1} kpc on a side and are centred on the position of the galaxy in the ‘REF’ simulation. All frames have a thickness of 100 comoving h^{-1} kpc. The orientation of the line of sight is along the z -axis, which is almost perfectly aligned with the angular momentum vector of all material inside 10 per cent of the virial radius of this galaxy in the ‘REF’ simulation. It is clear that different feedback prescriptions can have a large effect on the size of the galactic disc. Note that ‘WML4’ is not included here. That simulation is also discussed in Paper II, and it is included there.

determined. The line of sight is along the z -axis, which is almost perfectly aligned with the angular momentum vector of the gas within 10 per cent of the virial radius ($\cos(\phi) = 0.994$). For the other simulations, the image is centred on the same position, showing the remarkable similarity in the positions and orientations of the galaxies. It is immediately clear that changing the feedback model can have a large effect on the galaxy morphology. Going so far as

almost completely destroying the galaxy disc in the strongest cases (‘AGN’ and ‘DBLIMFV1618’).

4.1 Metal-line cooling

In this set of simulations, we investigate how gas cooling through metal lines affects the galaxy population. Metal-line cooling is the

dominant mechanism by which enriched gas can cool in the temperature range $10^5 < T < 10^6$ K (e.g. Wiersma et al. 2009a). Metals are also the main coolants for gas at temperatures well below 10^4 K, but we do not include those phases of the ISM in these simulations. If gas shock heats to high temperatures while accreting on to galaxies, neglecting metal-line cooling will greatly decrease the supply of gas that can cool out of haloes to fuel the galaxy (see van de Voort et al. 2011, for a comprehensive discussion of galaxy fuelling in these simulations). In order to isolate these effects, we compare the ‘REF’ simulation to one in which cooling through metal-lines was switched off (‘NOZCOOL’), and to a further simulation in which both cooling through metal-lines and SN feedback were switched

off (‘NOZCOOL_NOSN’). The morphology of a typical massive galaxy in these simulations can be seen in Fig. 2. Turning off metal-line cooling reduces the extent of the gaseous disc in this massive system as less gas cools out of the halo into the galaxy. Neglecting SN feedback leads to a hugely centrally concentrated galaxy as there is no mechanism to eject low-angular momentum material or to stop the fragmentation and associated angular momentum losses.

To isolate the effect of metal-line cooling, we can compare the ‘REF’ simulation (solid, black curve) to the ‘NOZCOOL’ simulation (dotted, red curve) in Fig. 3. Panel (B) of Fig. 3 shows the effect of metal-line cooling on the SFRs of galaxies. In general, metal-line cooling increases SFRs, because cooling rates increase with

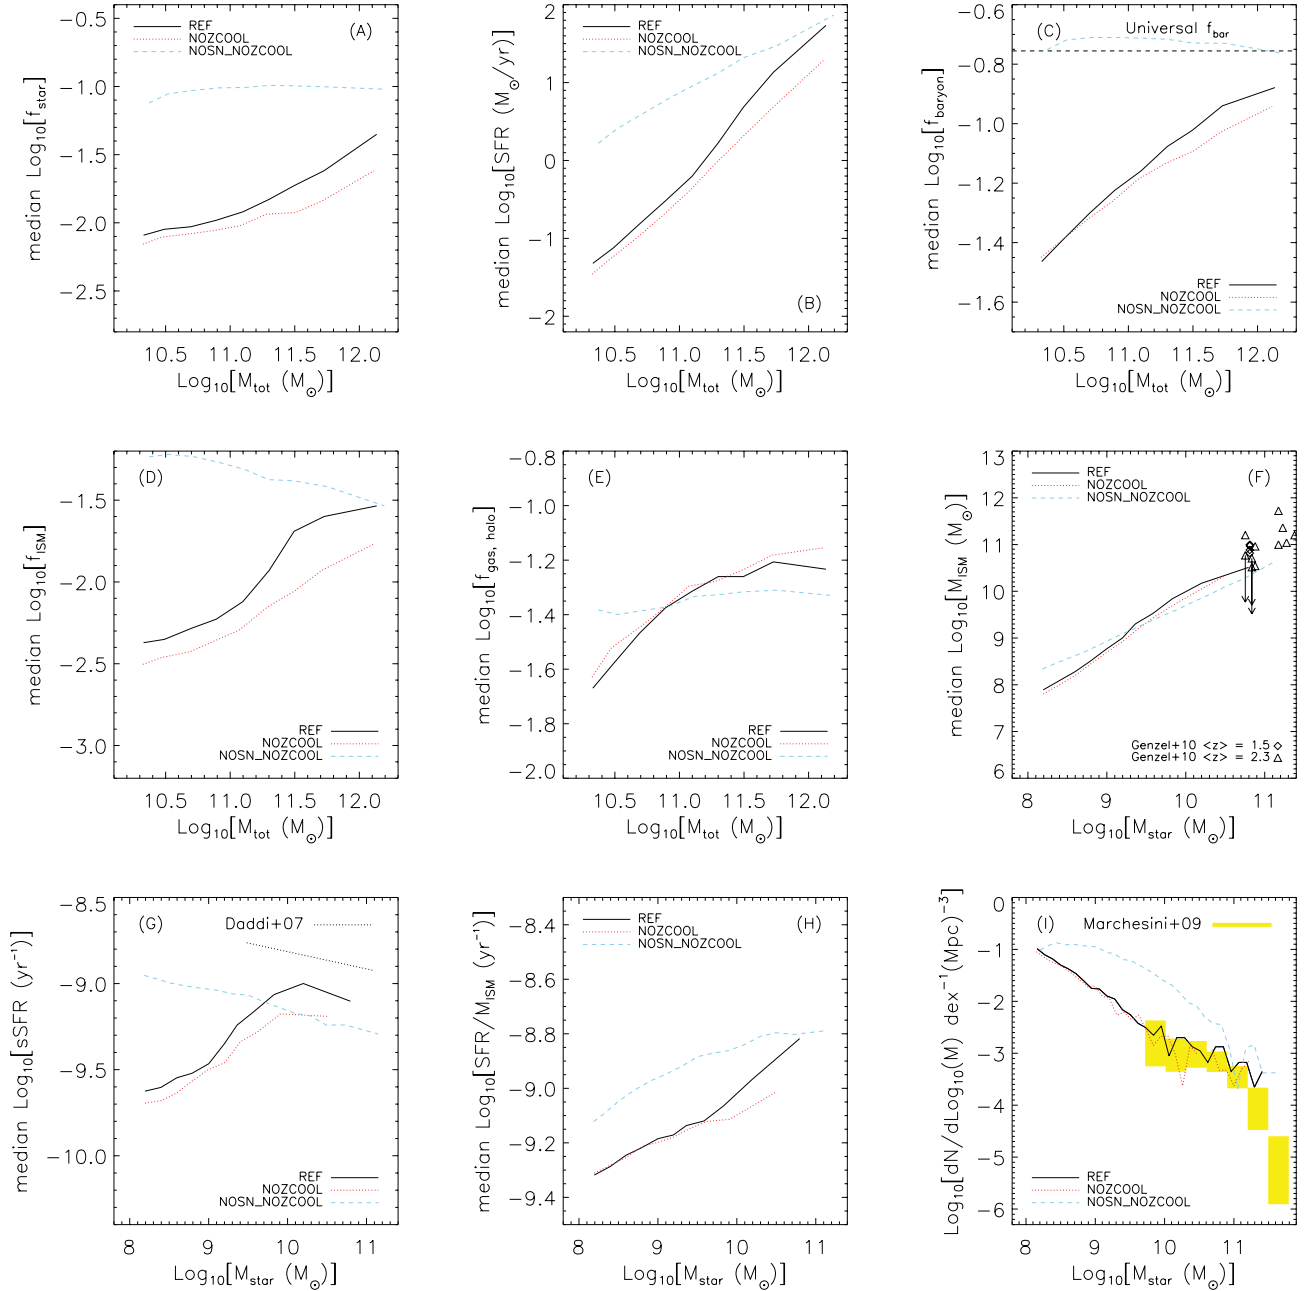


Figure 3. As Fig. 1, but showing only the sub-set of simulations in which the metal-line cooling and/or kinetic SN feedback are turned off. The solid, black curve shows the ‘REF’ simulation. The red, dotted line shows the effect of turning off only metal-line cooling (‘NOZCOOL’). The blue, dashed curve shows the effect of switching off both metal-line cooling and SN feedback (‘NOZCOOL_NOSN’). The effect of the SN feedback can thus be isolated by comparing the blue, dashed and red, dotted curves.

increasing metallicities. The magnitude of the differences between the simulations increases with halo mass because both the fraction of gas that is shock heated when it enters the galaxy (e.g. Birnboim & Dekel 2003; Kereš et al. 2005; Ocvirk, Pichon & Teyssier 2008; van de Voort et al. 2011) and the halo virial temperature increase with mass, making the effect of metal-cooling more pronounced. The same trends are evident in other gas properties as a function of halo mass and stellar mass. Without metal-line cooling, less gas is able to cool on to the galaxy (panel D), and hence stellar masses (panel A) are lower.

As a function of stellar mass, we see that sSFRs are higher (panel G) and gas consumption time-scales are shorter (panel H) in the ‘REF’ than in the ‘NOZCOOL’ simulation, as gas can cool more quickly into the galaxy when metal-line cooling is allowed, leading to higher gas densities and more efficient star formation. At high masses, winds are inefficient and more metal-rich gas piles up, increasing the difference between the simulations. The stellar mass functions of the two simulations (panel I) are almost identical up to $M_{\text{star}} = 10^{9.5} M_{\odot}$, but above this mass the ‘NOZCOOL’ simulation lies systematically 0.1–0.2 dex below ‘REF’. This occurs because the normalization of the $M_{\text{star}}-M_{\text{tot}}$ relation is decreased in ‘NOZCOOL’, and so a given stellar mass corresponds to a larger, rarer halo.

Turning our attention now to the effect of SN feedback by comparing the ‘NOZCOOL’ simulation (dotted, red curve) to the ‘NOZCOOL_NOSN’ simulation (dashed, blue curve), we see that neglecting SN feedback leads to dramatic changes in the galaxy population as with no energy injection there is no process that can prevent gas from cooling directly into the galaxy.

Notably, panel (C) demonstrates that, unlike all of the other simulations we consider, ‘NOZCOOL_NOSN’ has baryon fractions *above* the universal value. This is due to the fact that star formation and cooling allow more gas to be pulled into the potential well of the galaxy. Non-radiative simulations predict baryon fractions of ~ 0.8 times the universal value (Crain et al. 2007, but see Kravtsov, Nagai & Vikhlinin 2005 who find a value much closer to one in a non-radiative mesh-based simulation), as pressure support forces more gas outside of haloes.

Interestingly, we note that (unlike all of the other simulations considered here) the slope of the $M_{\text{star}}-\text{sSFR}$ relation (panel G) in the ‘NOZCOOL_NOSN’ has an almost identical slope to that observed, although with a normalization that is lower by a factor of approximately 2. That the slope of the observed $M_{\text{star}}-\text{sSFR}$ relation is so similar to that of the simulation without feedback may indicate that perhaps feedback is ineffective at high galaxy masses, as is the case for the reference simulation above $10^{10} M_{\odot}$ (see also Paper II). The similarity to the observed relation is striking. The galaxies in the observations are selected to be star forming, so one would naively expect the star formation in those galaxies to be regulated by feedback.

4.2 Winds with constant energy per unit stellar mass formed

Winds driven by feedback from star formation have become a fundamental ingredient of the galaxy formation picture by preventing too much gas from being locked into stars. In the ‘REF’ simulation, such winds are assumed to be driven by SNe and are implemented kinetically. That is, the effect of SNe is to ‘kick’ some amount of gas. This model contains two free parameters: the initial mass-loading $\eta = \dot{M}_{\text{wind}}/\dot{M}_{\text{s}}$ and the initial wind velocity v_w , which remain essentially unconstrained. We therefore compare a series of four simulations that use the same SN energy per unit stellar mass

formed (i.e. ηv_w^2 is the same as for ‘REF’), but distribute it differently between the velocity and mass loading. The mass loadings in the four simulations are 1, 2 (i.e. ‘REF’), 4 and 8, with corresponding velocities of 848, 600, 424 and 300 km s⁻¹, respectively. The parameters of the wind model are contained in the simulation name. For example, the simulation ‘WMLIV848’ uses a mass-loading of 1 and a wind velocity of 848 km s⁻¹. A fifth simulation with kinetic feedback with the same feedback energy, ‘WDENS’, uses a wind velocity that scales in proportion to the local gas sound-speed ($\propto \rho^{1/6}$) and a mass loading that scales such that the total injected energy is constant ($\propto \rho^{-1/3}$). The normalization of the parameters in the ‘WDENS’ simulation is such that the wind velocity and mass loading are the same as in the reference model if the gas density equals the star formation threshold, i.e. $n_{\text{H}} = 0.1 \text{ cm}^{-3}$. For illustrative purposes, we also consider a simulation in which the available SN energy relative to ‘REF’ is doubled (by doubling the mass loading to 4 and keeping the wind velocity the same). This simulation is termed ‘WML4’.

We first compare the set of simulations in which we hold the amount of SN energy that is injected per unit stellar mass constant, but vary how the energy is distributed between the initial mass loading and wind velocity. We show in Fig. 4 that the simulation with the highest wind velocity (‘WMLIV848’; red, dotted curve) suppresses SF much more effectively in high-mass objects than the simulations with lower wind velocities (e.g. ‘WML8V300’). However, at low galaxy masses, this trend reverses and it is the simulations with the lowest wind velocities that suppress SF most efficiently (panels A, B and D). At the low-mass end, the amount of suppression is set by the amount of gas that is kicked, where models with large mass-loadings are able to inject large amounts of gas into the wind, much of which escapes into the halo or beyond, leading to these models being most effective. However, in massive haloes a large wind velocity is required to drive gas out of the galaxy, so the models with large mass loadings become incapable of driving winds and form more stars in high-mass objects. There exists a clear transition mass above which, for a given wind velocity, SN feedback is no longer able to drive gas from the galaxy (see also Springel & Hernquist 2003b). This is visible both in the properties of the galaxies (panels F–H), and the properties of the galaxy haloes (panels A–E), where above a mass (that depends on the initial wind velocity) the results tend towards those in a simulation that includes no feedback (Fig. 3; blue, dashed curve).

Dalla Vecchia & Schaye (2008) showed that the reason why effective feedback in more massive galaxies requires higher velocity winds is not directly related to the requirement that the winds overcome the gravitational potential. Instead, they found that the outflows are already quenched in the ISM and only if gas pressure forces are taken into account. Because the pressure in the ISM increases with the depth of the potential well, more massive galaxies require higher velocity winds. A deeper understanding was provided by Dalla Vecchia & Schaye (2012), who demonstrated analytically, and confirmed numerically, that the outflows are quenched due to radiative losses in the shock-heated ISM. Higher velocities imply higher post-shock temperatures and hence longer cooling times. As the ISM of more massive galaxies is denser, higher velocities are required for effective feedback in more massive galaxies. They pointed out that cosmological simulations overestimate the radiative losses due to their finite resolution and showed that models with different wind parameters, but the same amount of energy per unit stellar mass, make converging predictions if the resolution is sufficiently high. The resolution of our simulations is low compared to the analytic estimates of the required resolution provided by Dalla

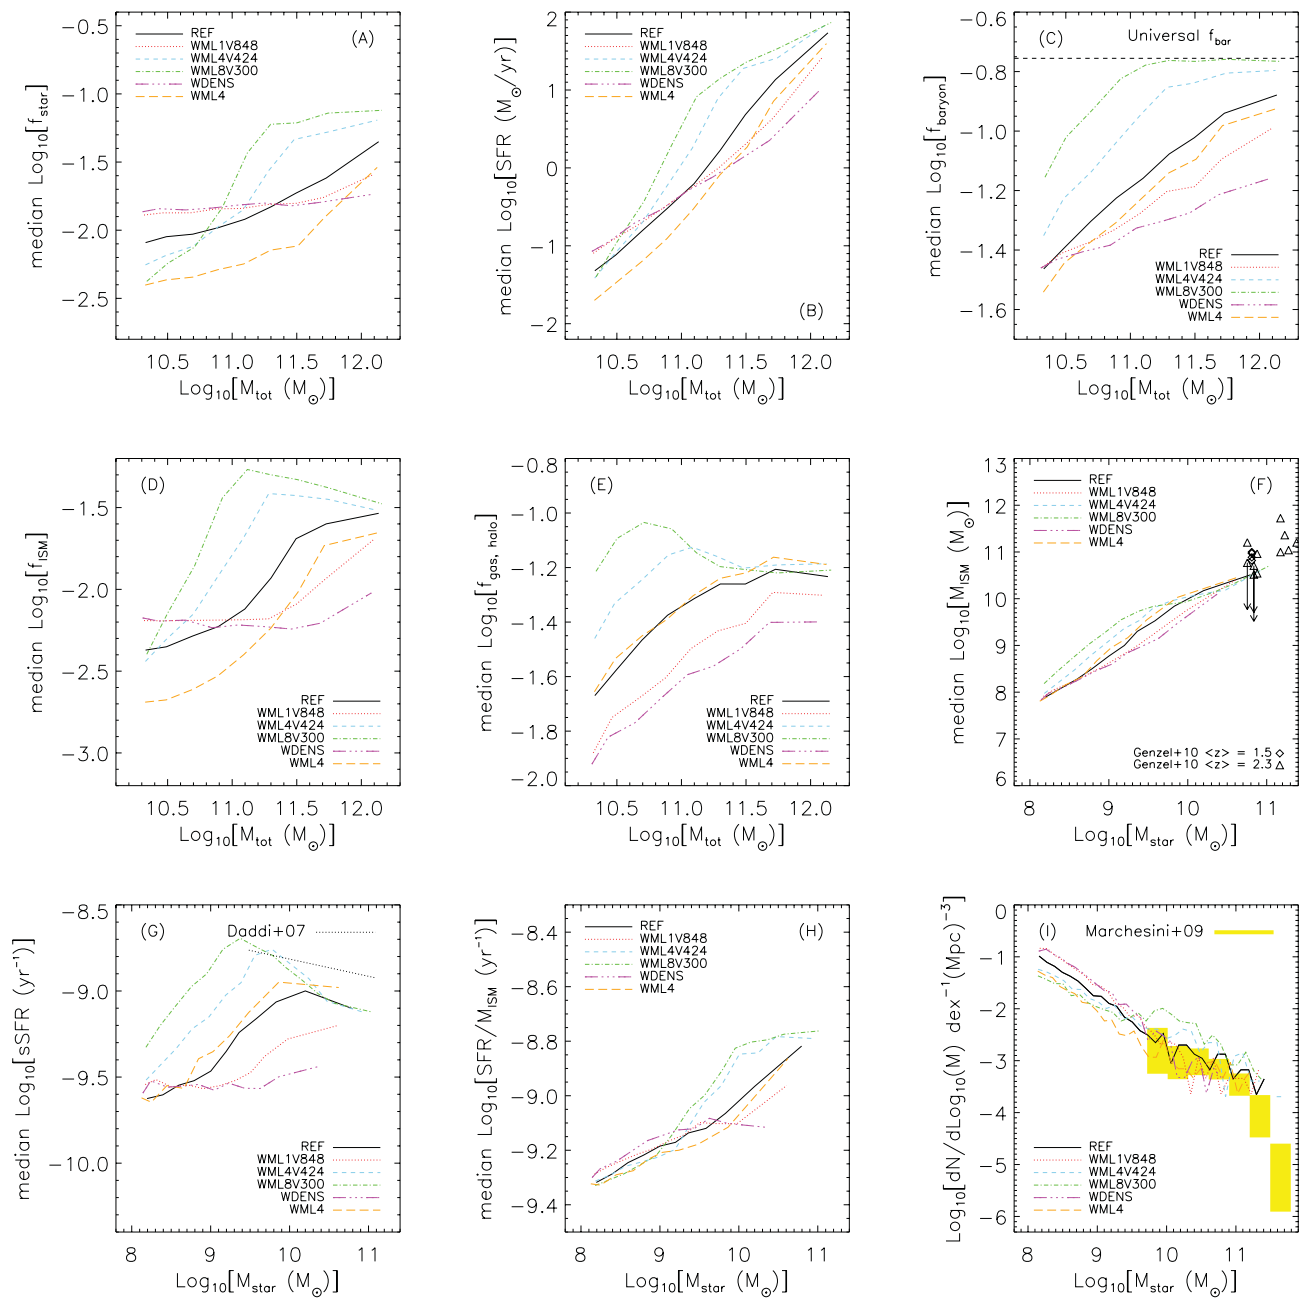


Figure 4. As Fig. 1, but showing only a set of simulations in which the initial wind velocity and the initial wind mass loading are varied. The reference simulation (black, solid curve) has a mass loading of $\eta = 2$ and a wind velocity of 600 km s^{-1} . The simulations shown by the red-dotted, blue-dashed and green-dot-dashed lines show simulations that use the same total energy per unit stellar mass but different combinations of the mass loading and wind velocity. The simulations shown have mass-loadings of 1 (*WML1V848*; red, dotted curve), 4 (*WML4V424*; blue, dashed curve) and 8 (*WML8V300*; green, dot-dashed curve). The magenta triple dot-dashed line represents a simulation (*WDENS*) which has a mass loading and velocity dependent on the local density, such that the energy in the wind is still the same and the initial velocity is proportional to the local sound speed. Changing the parameters of the kinetic feedback model can change the galaxy stellar properties by up to an order of magnitude, even when the energy that is injected is kept constant. The halo mass above which SN feedback becomes inefficient at removing gas from the galaxy (as seen by a steeper rise in the galaxy SFR with halo mass) depends primarily on the SN-driven wind velocity.

Vecchia & Schaye (2012), which explains our finding that galaxy properties are sensitive to the choice of initial wind velocity.

A second instructive comparison is between *REF* and the simulation that injects twice as much SN energy per unit stellar mass, *WML4* (orange, long-dashed curves). This simulation uses the same wind velocity as *REF* and so the SN-driven winds become ineffective at the same mass. This can be seen by comparing

REF (black, solid curve) to *WML4* in panel (B), where above $M_{\text{tot}} = 10^{11.25} M_{\odot}$, the two simulations converge towards each other with increasing mass, but below this mass, *WML4* can suppress SF significantly more efficiently than *REF*. We can see this more quantitatively in panel (A), where we see that at low masses (where the feedback is effective), f_{star} is almost exactly a factor of 2 lower than in the *REF* simulation. This implies that the star formation

is self-regulating: a galaxy's SFR increases until the rate of energy injection into the ISM reaches a critical value. This critical rate of energy injection depends on the halo mass and is presumably the rate for which outflows balance inflows (e.g. Schaye et al. 2010; Davé, Finlator & Oppenheimer 2012; see also Booth & Schaye 2010 for an analogous discussion of self-regulated BH growth). Because the outflow rate depends not only on the rate of energy injection, but also on the initial wind velocity (e.g. Dalla Vecchia & Schaye 2008), the same will be true for the critical SFR. Doubling the SN energy that is injected per unit stellar mass, while keeping the initial wind velocity constant, will half the SFR required to produce the same outflow rate and will thus cause the galaxy to form half as many stars. This represents one of the fundamental conclusions of this paper, star formation is regulated by the interplay between the available fuel supply and feedback processes.

Because a given galaxy requires a fixed rate of energy injection to achieve a balance between inflows and outflows, increasing the efficiency of SN feedback causes the galaxies to decrease their gas reservoirs (panel D). Because the star formation law is non-linear and because the ISM of more massive galaxies tends to be denser, the decrease in the mass in the ISM is larger for lower mass galaxies.

The gas consumption time-scales of all models are very similar in the regime where SN feedback is effective (panel H) because the ratio between the amount of gas available for star formation and the rate at which stars are formed is mainly determined by the star formation law. At the point where feedback becomes incapable of lifting gas out of the galaxy, the gas in the galaxy is used up very efficiently.

The kinetic feedback models discussed above introduce a slope in the relation between stellar mass and sSFR, that does not agree with observations. Models that do not use a constant wind velocity and mass loading result in different relations between these quantities as we will show below.

For the 'WDENS' model, the energy injected in the wind per unit stellar mass is also the same as in the reference model, but the initial wind velocity scales with the local sound speed. The relation between halo mass and SFR is even shallower than it is for the run with $v_w = 848 \text{ km s}^{-1}$, indicating that the feedback is effective for all haloes. At high masses, this is our most effective wind model with constant energy. This indicates that choosing the mass-loading and wind velocity carefully can make SN-driven winds capable of driving winds in all objects resolved in these simulations.

As is clear from panel I of Fig. 4, decreasing the slope of the low-mass end of the stellar mass function can be attained by increasing the mass loading factor in constant energy winds. The highest mass loading still gives a low-mass end slope that is somewhat steeper than power-law fits to the low-mass end in the observations, although the discrepancy only becomes severe at masses lower than those observed. Such efficient feedback in low-mass objects also boosts the sSFRs in the higher-mass objects for which feedback from star formation is just becoming inefficient, as appears to be required by the observations (panel G).

These simulations suggest that the fit to the observed mass function could be much improved by varying the mass loading factor with galaxy mass (see also Oppenheimer et al. 2010; Okamoto et al. 2010; Bower, Benson & Crain 2012; Puchwein & Springel 2012).

4.3 Simulations with a top-heavy IMF at high pressures

Determining the stellar IMF represents a significant challenge outside of the solar neighbourhood. The IMF determined locally is usually taken to be universal and is applied to all galaxies, at all

redshifts. There is, however, some evidence suggesting that the IMF may be 'top-heavy' in extreme regions such as starbursts (e.g. Baugh et al. 2005; Haberland, Anderson & James 2010; Weidner, Kroupa & Pflamm-Altenburg 2011). We performed a series of runs in which a different IMF is assumed in stars that formed in high-pressure environments.

In these simulations, stars form with an IMF $dN/dM \propto M^{-1}$ (compared with $M^{-2.3}$ for the high-mass end of the Chabrier IMF) if the gas pressure exceeds $P/k = 2.0 \times 10^6 \text{ cm}^{-3} \text{ K}$ (evaluated at the resolution limit of the simulations). This process mimics the observation that stars that form in extreme environments (such as starbursts) appear to have an IMF that is flatter than the Chabrier IMF we assume in 'REF' (e.g. McCrady, Gilbert & Graham 2003; Stolte et al. 2005; Maness et al. 2007).

For a top-heavy IMF, the available energy from SNe per unit stellar mass formed is higher by a factor of ~ 7 than for the Chabrier IMF. Numerically, we have some freedom in how this energy is distributed. It can be used either to increase the wind mass loading or the wind velocity. We therefore ran two additional simulations with a top-heavy IMF at high pressures. In the first simulation, the wind velocity remains fixed at 600 km s^{-1} (the 'REF' value), but the mass loading is increased to $\eta = 14$ ('DBLIMFML14') for stars that form in high-pressure environments. In a second simulation, the mass loading is kept fixed at $\eta = 2$ (the 'REF' value) and the wind velocity is increased to $v_w = 1618 \text{ km s}^{-1}$ ('DBLIMFV1618') for stars that form at high pressures.

When changing the IMF suddenly as a function of the pressure, it is not immediately clear how to treat the star formation law. The Kennicutt–Schmidt law is inferred from observations that probe only massive stars. The actual SFR, therefore, depends on how many low-mass stars are formed per unit mass of massive stars. When changing the IMF, the star formation law can be changed in two ways.

(i) From observations there is no indication of a discontinuity in the formation rate of massive stars with pressure. Although this is most likely the result of the IMF being a continuous function of SFR or pressure (if there is a relation at all), we nevertheless implemented a model that changes the normalization of the KS-law such that the formation of massive stars is continuous, resulting in a discontinuous SFR as a function of pressure (the KS-law normalization drops at the pressure above which the IMF is top-heavy). This is the procedure we follow in most of our simulations, and is the assumption made by both the simulations 'DBLIMFML14' and 'DBLIMFV1618'.

(ii) If the (total) SFR as a function of pressure is continuous, the formation of massive stars must be discontinuous, given that we assumed the IMF to change suddenly above some critical pressure. We run one additional model under this assumption, termed 'DBLIMFCONTSFV1618', which assumes that the total SFR as a function of pressure is continuous and the extra energy injected by a top-heavy IMF goes into increasing the wind velocity relative to the 'REF' simulation.

When comparing simulations to observations, we do not correct the stellar mass of simulations with a double IMF. On average, only ~ 10 per cent of the star particles in the simulation box formed with a top-heavy IMF (this depends slightly on resolution and weakly on whether the rate of formation of massive stars is a continuous function of the pressure). In Schaye et al. (2010), we showed that at late times, this correction can be significant, but that at $z = 2$ the integrated SFR of the universe is the same, regardless of whether or not the SFRs of particles at pressures higher than the threshold

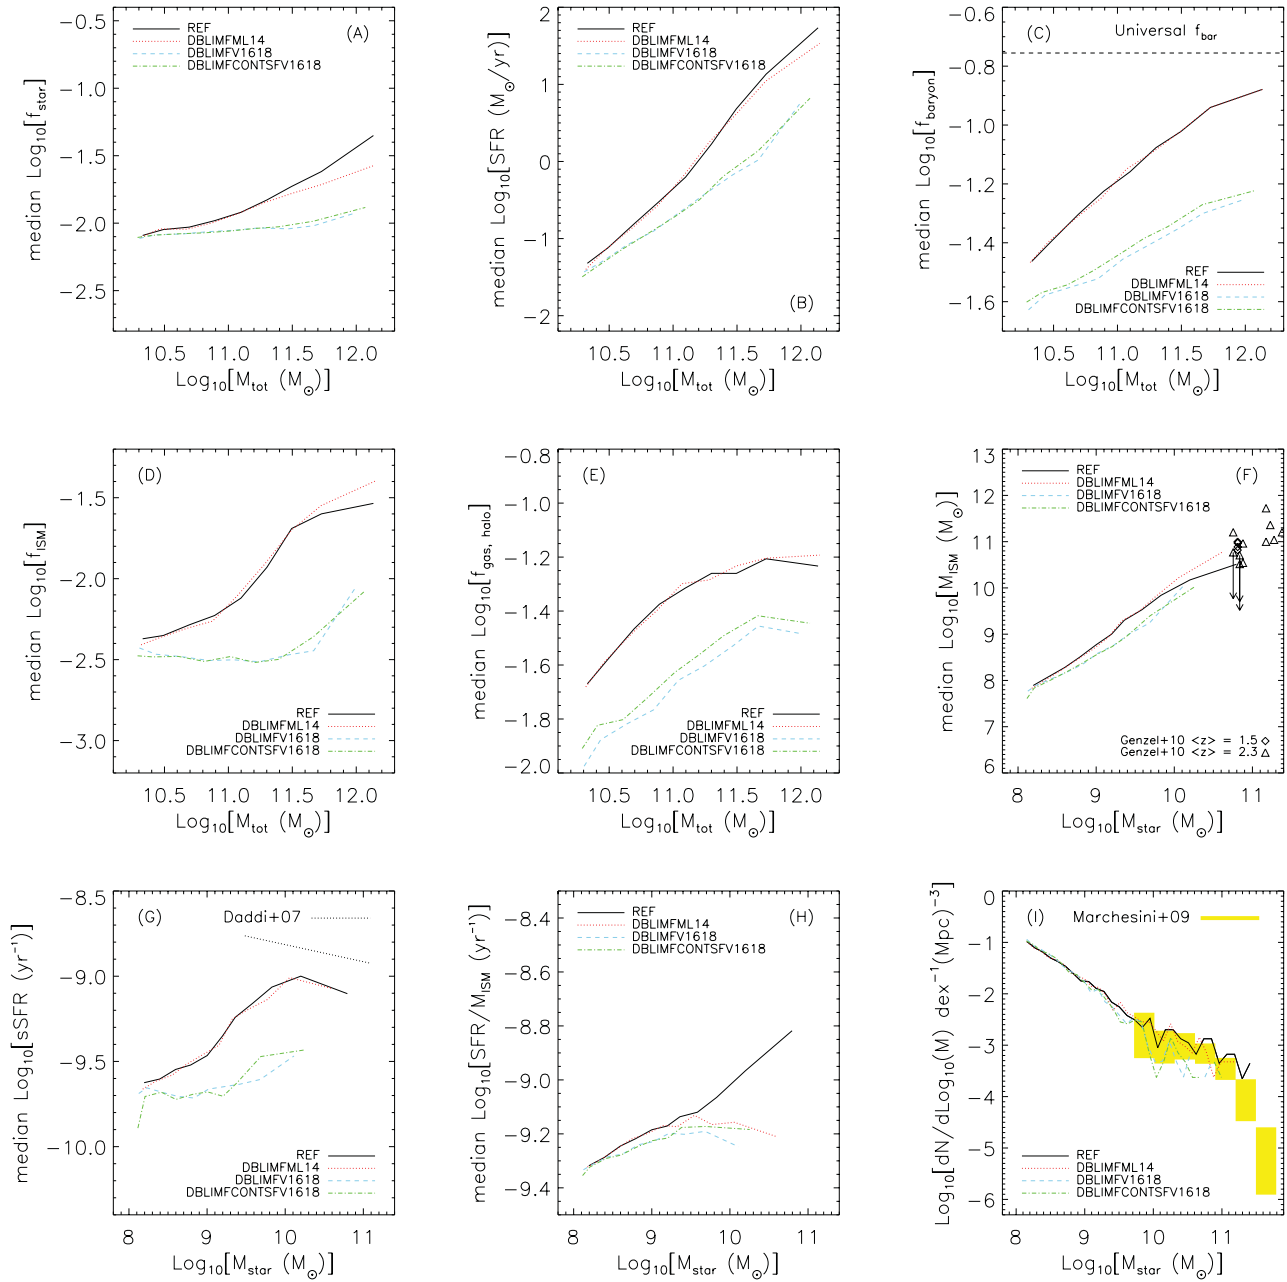


Figure 5. As Fig. 1, but showing only the set of simulations in which a top-heavy IMF is used for star formation at high pressure. The ‘REF’ simulation (without top-heavy IMF) is shown by the black, solid curve. The extra SN energy available from a top-heavy IMF can either go towards increasing the mass loading (‘DBLIMFML14’; red, dotted curve) or the initial velocity of the winds (‘DBLIMFV1618’; blue, dashed curve). In the presence of a sudden change in the IMF at some pressure, one can either allow the SFR or the rate of formation of massive stars (which is what is observed) to be a continuous function of the pressure (‘DBLIMFCONSTSFV1618’; green, dot-dashed curve). Putting the extra SN energy into the mass-loading (‘DBLIMFML14’) does not increase the level of suppression of the SFR, whereas increasing the wind velocity does allow more gas to escape from galaxies.

pressure for the top-heavy IMF are corrected for another assumed IMF. Also, we demonstrate in Paper II, by considering simulations in which the normalization and slope of the star formation law are varied, that the form of the assumed SF law does not strongly influence the galaxy mass function or the SFR distributions in galaxies (it mainly affects the ISM fraction). At $z = 2$, we therefore expect any differences between the runs that use a top-heavy IMF and the ‘REF’ simulation to be mostly due to the extra energy input from SN feedback and/or the increased rate of production of metals that results from a top-heavy IMF.

We show results from these simulations in Fig. 5. In all of the plots there are two ‘families’ of curves. The simulations ‘REF’ (black, solid curves) and ‘DBLIMFML14’ are, in almost all plots, similar to one another, and the simulations ‘DBLIMFV1618’ (blue, dashed curves) and ‘DBLIMVCONSTSFV1618’ (green, dot-dashed curves) are also very similar to one another. The fact that ‘DBLIMFV1618’ and ‘DBLIMVCONSTSFV1618’ are always very similar tells us that, if feedback is effective, whether or not we have a continuous SFR at the threshold pressure is only of minor importance when it comes to galaxy properties, and as such we will not consider ‘DBLIMV-CONSTSFV1618’ any further.

Next we compare ‘REF’ and ‘DBLIMFML14’. As discussed in Section 4.2, SNe are only capable of driving winds out of a galaxy when they are launched at sufficiently high velocities. Therefore, ‘DBLIMFML14’ behaves similarly to ‘REF’ because the two models are capable of driving winds out of the same objects. It is apparent that ‘DBLIMFML14’ does form fewer stars than ‘REF’ at high halo masses ($M_{\text{tot}} \gtrsim 10^{11.5} M_{\odot}$; panels A and B), even though the halo baryon fractions of the two simulations are very similar (panels C and E). That the two models differ only at high halo masses can be explained by the fact that the pressure in the ISM increases with the depth of the potential. Hence, only in massive haloes does a significant fraction of the ISM have pressures above the critical value above which we assume a top-heavy IMF.

The reason why ‘DBLIMFML14’ has lower stellar mass fraction in massive haloes than ‘REF’ can be seen by comparing the gas consumption time-scale (panel H) and the mass of gas in the ISM (panels D and F) for the same two simulations. Although there is significantly more gas in the ISM in ‘DBLIMFML14’ than in ‘REF’, the gas consumption time-scale is much longer. This may reflect the drop in the normalization of the KS-law in high-pressure gas in ‘DBLIMFML14’. Because the gas consumption time-scale becomes extremely long in the high-pressure gas ($\sim 10^{10}$ yr) and because the initial wind velocity is too low to be effective, self-regulation may be difficult to achieve.

Finally, we compare ‘REF’ and the top-heavy IMF simulation in which the initial wind velocity is increased (‘DBLIMFV1618’). It is immediately obvious that the effect of the top-heavy IMF is to strongly suppress star formation (panels A, B, F, G and H) as well as to effectively remove gas from galaxies (panel D) and from their host haloes (panels C and E). The effects increase with the mass, because the ISM pressure, and thus the fraction of stars forming with a top-heavy IMF, increase with the depth of the potential well. The effect of the top-heavy IMF is very similar to increasing the wind velocity (see Section 4.2). In both simulations, the SN-driven winds are able to drive out gas from all but the most massive objects, but the stellar fractions are lower for ‘DBLIMFV1618’ than for ‘WMLV848’. This is a result of the increase in the available SN energy per unit stellar mass, so a smaller SFR is required to drive winds from a galaxy and to effectively regulate star formation. Finally, we note that the simulations that include strong feedback from a top-heavy IMF suppress the amplitude of the galaxy stellar mass function below the observed points (panel I). This behaviour is seen in the other models that contain very strong feedback (e.g. ‘WPOT’ and ‘WVCIRC’ in Section 4.4 and ‘AGN’ in Section 4.7).

We conclude that the effect of a top-heavy IMF at high pressures mainly reflects the increased efficiency of winds driven by massive stars. The extra energy available for driving winds is more important than the increase in the metal yields associated with a top-heavy IMF.

4.4 ‘Momentum-driven’ wind models

Galactic winds can be driven by radiation pressure on dust grains in the wind, which drag along the gas (e.g. Murray, Quataert & Thompson 2005). Here, the driving force of the wind is the radiation pressure which injects momentum into the outflow. If the cooling time in the shock-heated gas is sufficiently short, then the wind will conserve its momentum and the wind is said to be ‘momentum driven’. Although cosmological, hydrodynamical simulations have not yet tried to model this process directly by including radiation transport, simulations have been run that include the same type of kinetic feedback as used for SN-driven winds, but in which the initial

wind velocity varies with local physical conditions or galaxy properties, and the initial mass loading varies as $\eta \propto v_w^{-1}$, such that the rate with which momentum is injected per unit stellar mass, $\propto \eta v_w$, is constant. This scaling should be contrasted with that used for the ‘constant energy’ winds considered earlier, for which $\eta \propto v_w^{-2}$.

We implemented several ‘momentum-driven wind’ models that are similar to those used by Oppenheimer & Davé (2006, 2008). Here, the wind parameters vary with either the local potential (‘WPOT’ and ‘WPOTNOKICK’; Oppenheimer & Davé 2006) or the circular velocity, $v_c = \sqrt{GM_{\text{vir}}/R_{\text{vir}}}$ of the halo from which the wind is launched (‘WVCIRC’; Oppenheimer & Davé 2008). In ‘WVCIRC’, the wind velocity and mass loading are given by $v_w = (3+n)v_c/\sqrt{2}$ and $\eta = \frac{1}{\sqrt{2}} \times (v_c/v_{\text{crit}})^{-1}$, where n and v_{crit} are parameters, set to 2 and 150 km s⁻¹, respectively. In ‘WPOTNOKICK’, the wind velocity is given by $v_w = 3\sigma$, where σ is the velocity dispersion, calculated from the gravitational potential: $\sigma = \sqrt{-\Phi/2}$, and $\eta = 150 \text{ km s}^{-1} \sigma^{-1}$. In ‘WPOT’ an extra velocity kick of $2 \times \sigma$ is added to each event. These models are not directly comparable to those of Oppenheimer & Davé (2006, 2008, 2009) because these authors used different parameter values and different models in different papers (see section 4.9 of Schaye et al. 2010). Moreover, contrary to Oppenheimer & Davé, we do not decouple the wind particles from the hydrodynamics, which likely makes our implementations less efficient. As was shown by Dalla Vecchia & Schaye (2008) and Schaye et al. (2010), and as we will show in Section 4.5, temporarily turning off pressure forces for wind particles makes the winds more efficient and has a large effect on galaxy properties.

One important difference between the models discussed here and those in Section 4.2 is that here the total amount of energy put into the wind per unit stellar mass formed is not constant, but depends on the mass of the halo in which it takes place ($E \propto M_{\text{tot}}^{1/3}$ in ‘WVCIRC’), and we caution that the total energy used in feedback exceeds the total available energy from SNe for the most massive galaxies.² Most worryingly, the amount of momentum in the winds also exceeds the momentum available from SNe and radiation, assuming that every photon is absorbed once to drive the outflow (so no boost from optical thickness in the IR), by a factor of about 7 (Haas 2010, see also the discussion in Davé, Oppenheimer & Finlator 2011). It should therefore be kept in mind that these models may inject an unphysically large amount of momentum. Inspection of Fig. 2 reveals that these wind prescriptions are capable of almost completely disrupting the galaxy disc.

We compare the effects of the different momentum-driven wind models in Fig. 6. All the simulations studied in this section predict relatively flat relations between M_{tot} and f_{star} (panel A), yielding a flatter galaxy mass function at the low-mass end (panel I). None of the models show a characteristic halo mass above which feedback becomes ineffective (panels B–F), indicating that the feedback remains effective at all masses. In all of these simulations, the combination of a large mass-loading in low-mass haloes and a high wind-velocity in high-mass haloes means that feedback from star formation is always able to drive strong winds and suppress SF.

4.5 Hydrodynamically decoupled winds

Many SPH simulations that use kinetic SN feedback employ a method known as decoupling that allows SN-driven winds to

² This occurs for $M_{\text{tot}} \gtrsim 10^{12.5} M_{\odot}$, although the exact mass of equality is redshift dependent, due to the redshift dependence of the virial radius.

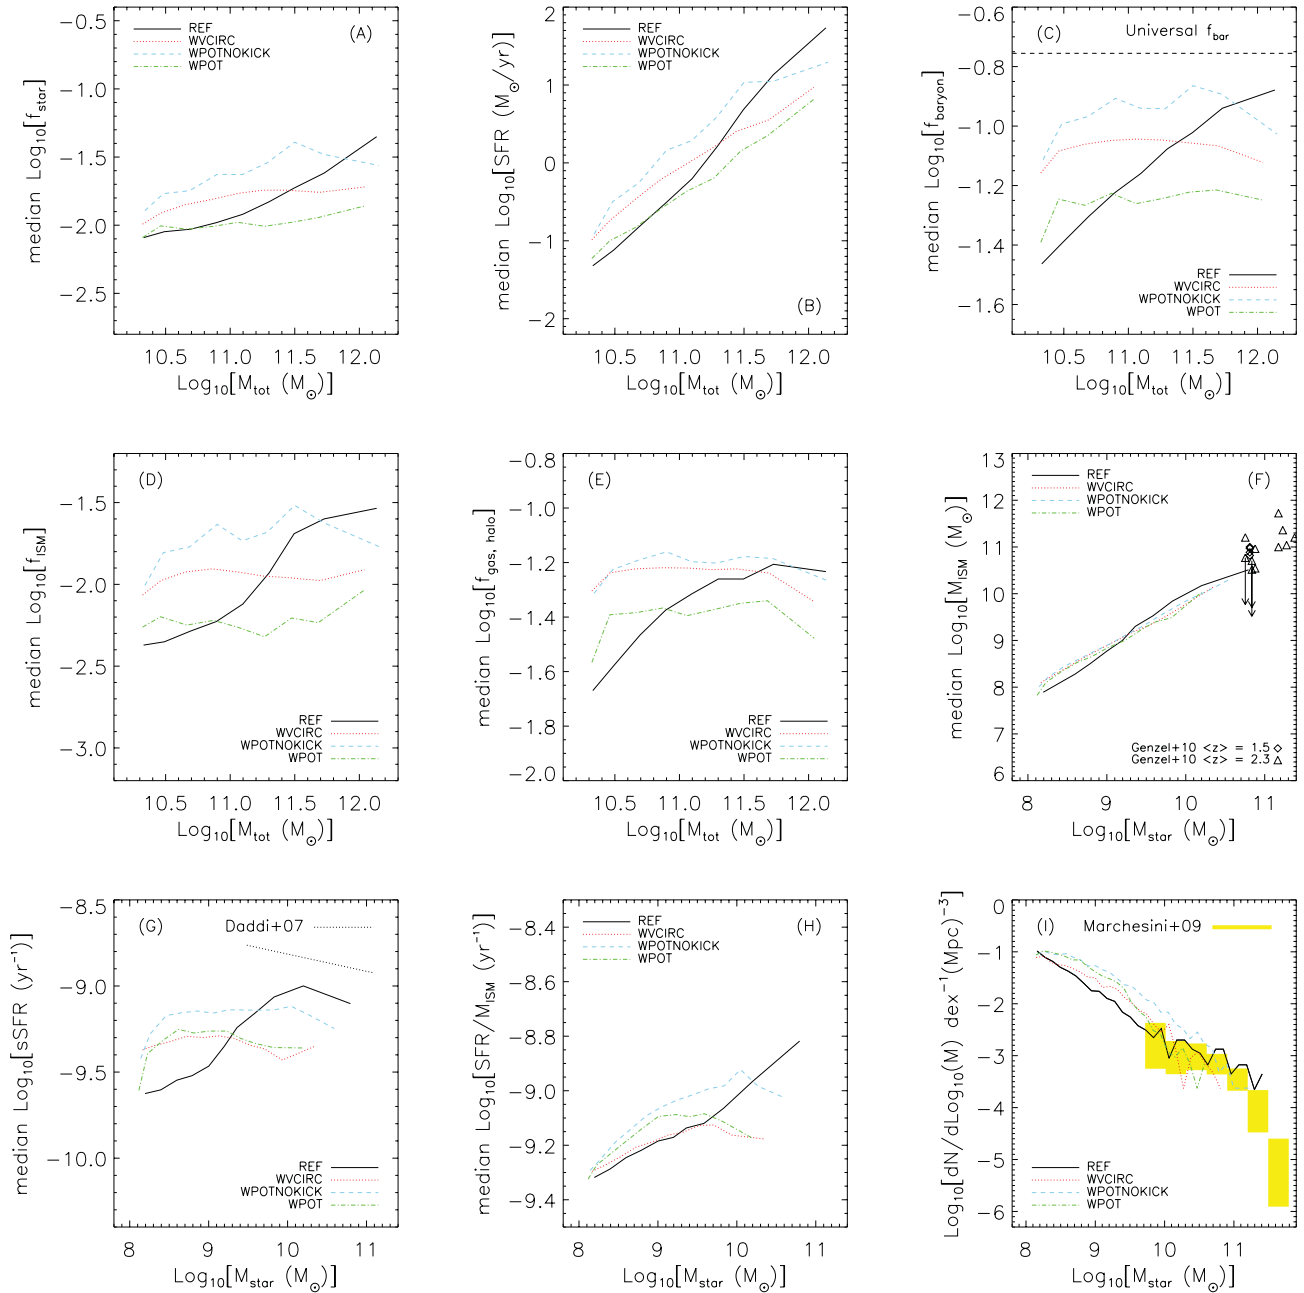


Figure 6. As Fig. 1, but showing only the reference simulation (*REF*; black, solid curve) and the set of simulations in which ‘momentum-driven wind’ scalings are employed. In *WVCIRC* (red, dotted curve), the initial wind velocity scales with on the circular velocity of the halo the wind is launched from, while in *WPOTNOKICK* and *WPOT* the wind velocity is set by the local gravitational potential (without and with an extra kick, respectively, shown by the blue, dashed and green, dot–dashed curves, respectively). In these models, the energy injected into the wind per unit stellar mass is not constant and generally exceeds the energy in the reference simulation, as is evident from the increased suppression of star formation.

effectively escape galaxies before they begin to interact with the halo gas. In these simulations, wind particles, once launched, are temporarily decoupled from the hydrodynamics until they escape the ISM. During decoupling a gas particle experiences gravity, but feels no hydrodynamic drag. Decoupling prevents shock-heating and hence the radiative losses which may otherwise quench the outflow in high-pressure gas (Dalla Vecchia & Schaye 2012), and it prevents the wind from entraining any surrounding ISM. To some extent, decoupling wind particles from the hydrodynamics mimics the existence of unresolved ‘chimneys’ of low density in the ISM, through which winds can easily escape without entraining much

other gas. However, decoupling does not account for the energy required to create such chimneys. For a detailed study of the effect of decoupling for the case of isolated disc galaxy simulations, see Dalla Vecchia & Schaye (2008).

To explore the effect of decoupling, one of the *OWLS* simulations, *WHYDRODEC*, employs decoupling, following the recipe of Springel & Hernquist (2003a). Every particle that is kicked into the wind feels no hydrodynamic forces for either a fixed amount of time (50 Myr), or until the density of the wind particle falls below some value (10 per cent of the star formation density threshold, i.e. when $n_{\text{H}} < 10^{-2} \text{ cm}^{-3}$). If the wind would retain its original velocity

of 600 km s^{-1} , 50 Myr would correspond to a travelling distance of roughly 30 kpc, so it ensures that SN-driven winds escape the galaxy.

Inspection of Fig. 2 reveals that decoupling the SN-driven winds dramatically increases the density of gas around this galaxy as every wind particle is able to escape from the galaxy (see also Dalla Vecchia & Schaye 2008). Decoupling the winds means that SN feedback remains capable of suppressing SF to significantly higher halo masses, until the wind velocity falls below the escape speed from the halo (see also Springel & Hernquist 2003b). This is visible in Fig. 7 where the galaxies in the ‘REF’ simulation show sharp upturns in M_{star} (panel A), SFR (panel B), f_{ISM} (panel D) at a halo mass $\sim 10^{11.25} M_{\odot}$, which are not present in the ‘WHYDRODEC’

simulation, indicating that the SN feedback remains capable of quenching SF if winds are artificially decoupled from the hydrodynamics. The decoupled winds are capable of ejecting more gas entirely from both the galaxy (panels D and F) and the halo (panels C and E).

Interestingly, for the lowest halo masses the SFR is higher for ‘WHYDRODEC’ than for ‘REF’ (panel B). This is expected, because wind particles will drag other particles into the wind, thereby increasing the effective mass loading, provided the initial wind velocity is sufficiently high for the winds to leave the galaxy.

The differences between ‘WHYDRODEC’ and ‘REF’ are very significant, up to 0.5 dex in the ISM fraction (panel D). From this, we can conclude that gravity (which acts on the winds in both

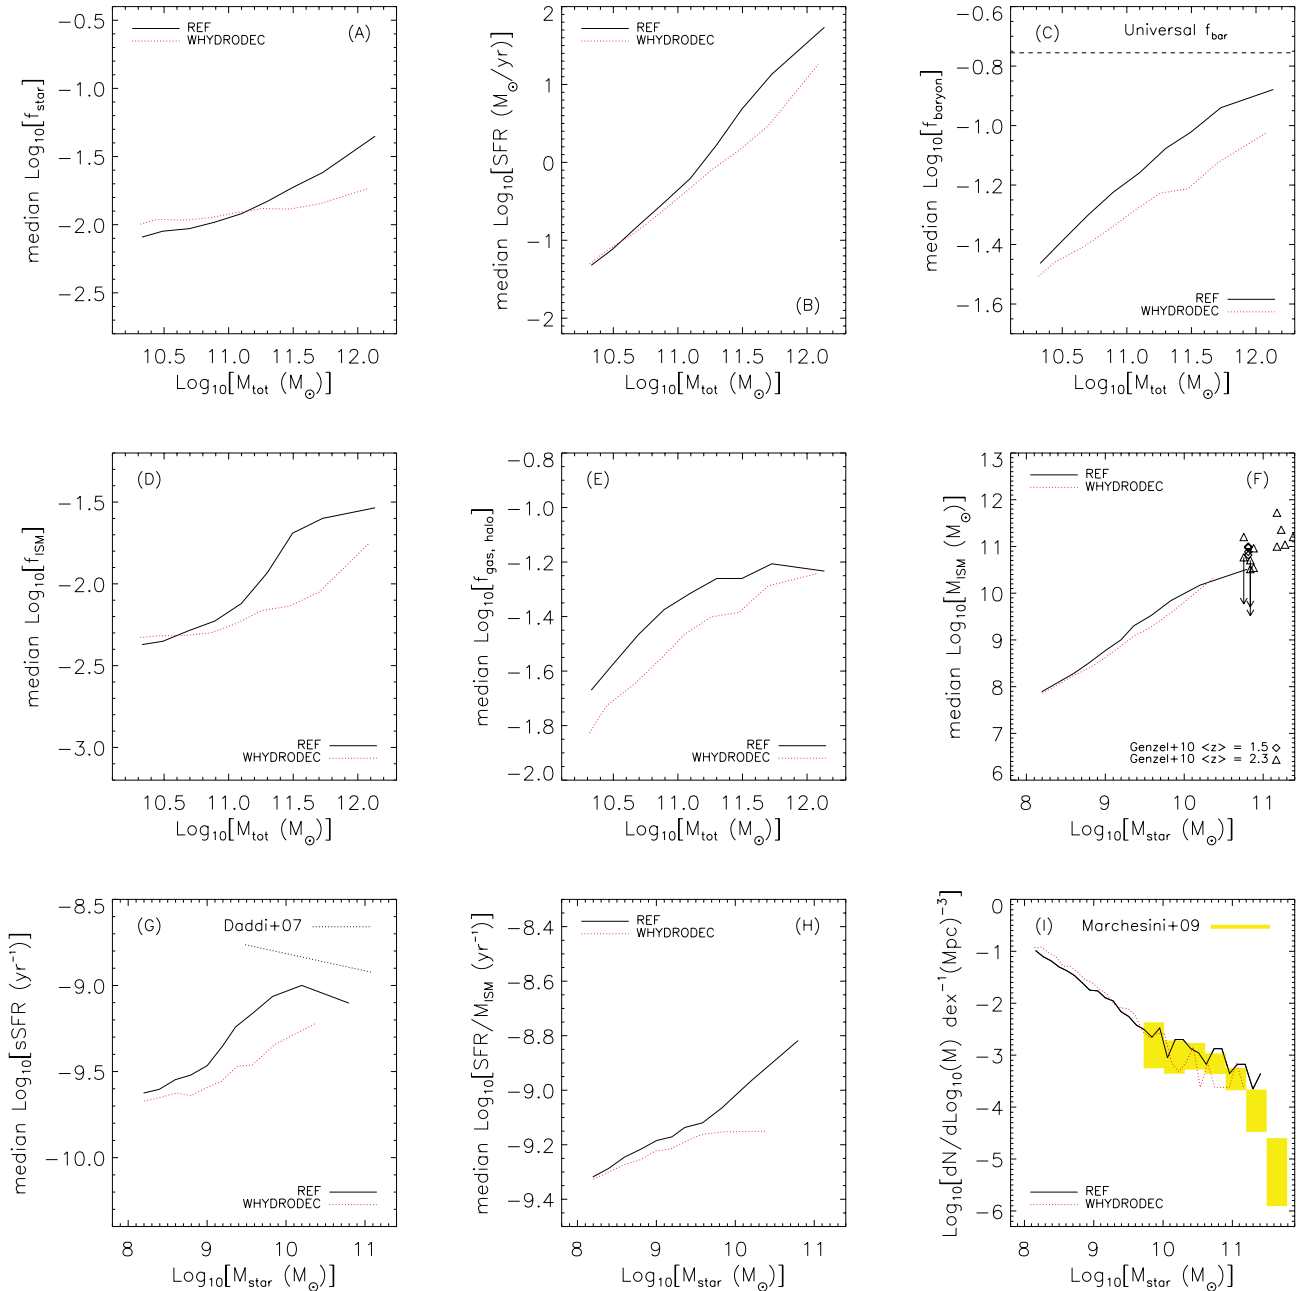


Figure 7. As Fig. 1, but showing only the ‘REF’ (black, solid curve) and the simulation in which the wind particles are temporarily decoupled from the hydrodynamics (‘WHYDRODEC’; red, dotted curve). Despite using the same wind parameters in both of these simulations, artificially decoupling the SN-driven winds from the galaxies makes the winds very efficient at suppressing star formation.

simulations) is not the process that makes the winds ineffective. It is rather the hydrodynamic drag and the associated radiative losses that make the winds less able to escape from high-mass haloes (see Dalla Vecchia & Schaye 2008, 2012; Creasey et al. 2011).

4.6 Thermal SN feedback

All the models that we considered up till now, implemented feedback from star formation kinetically. In this section, we investigate what happens if, instead of launching the wind by injecting kinetic energy, we inject thermal energy into the gas surrounding each newly formed star particle. As described in Dalla Vecchia & Schaye (2012), the ‘thermal feedback’ is implemented stochastically. If the SN energy is distributed amongst all of a star particle’s neighbours, then the rise in the temperature of each particle is so low that cooling times remain very short and the particle immediately re-radiates all of the energy. In this case, the feedback will have little effect, unless the cooling is temporarily suppressed (Mori et al. 1997; Thacker & Couchman 2000; Kay et al. 2002; Sommer-Larsen, Götz & Portinari 2003; Brook et al. 2004; Stinson et al. 2006). Therefore, we choose to inject the thermal energy into neighbouring gas particles stochastically by specifying a temperature jump, $\Delta T = 10^{7.5}$ K, and then calculating for each neighbouring gas particle the probability that it is heated such that the expectation value for the total injected energy agrees with the amount of feedback energy that is available. This method has some similarity with the promotion feedback model of Scannapieco et al. (2006). We do not turn off radiative cooling at any time. The simulation that employs this model is termed ‘*W THERMAL*’. Given our choices for ΔT and IMF, the expectation value for the number of heated particles per star particle is about 1.34 for a fully ionized plasma (equation 8 of Dalla Vecchia & Schaye 2012). We inject ~ 40 per cent of the available SN energy in order to facilitate comparison with the other models.

‘*W THERMAL*’ is compared to ‘*REF*’ in Fig. 8. Although the two models inject the same amount of SN energy per unit stellar mass formed, ‘*W THERMAL*’ suppresses star formation less effectively than ‘*REF*’ at most masses. This is visible in panels A–D, where the thermal feedback simulation lies systematically above the kinetic feedback model. The differences are small at high masses, because there both types of feedback are effective. However, at the lowest halo masses ($M_{\text{tot}} \lesssim 10^{10.5} M_{\odot}$, close to our resolution limit), ‘*W THERMAL*’ predicts lower stellar masses, but still slightly higher SFRs, than ‘*REF*’. This suggests that the thermal feedback is more efficient than the kinetic feedback in the poorly resolved galaxies that are not plotted here.

Our finding that the thermal feedback is less efficient than kinetic feedback is consistent with Dalla Vecchia & Schaye (2012), who predict that for $\Delta T = 10^{7.5}$ K, radiative losses should become significant for densities $n_{\text{H}} \gtrsim 1 \text{ cm}^{-3}$ (see their equation 18), which are certainly reached in all but the lowest mass galaxies in our simulations.

4.7 AGN feedback

AGN feedback is implemented using the method of Booth & Schaye (2009) which is, in turn, a modification of that of Springel, Di Matteo & Hernquist (2005). We frequently run a halo finder and insert a seed mass BH ($m_{\text{seed}} = 10^5 h^{-1} M_{\odot}$) into every halo with mass $> 4 \times 10^{10} h^{-1} M_{\odot}$ that does not yet contain a BH. These seed BHs then grow both through merging and gas accretion (which is limited to the Eddington rate). Accretion rates in low-density

gas ($n_{\text{H}} < 10^{-1} \text{ cm}^{-3}$) are assumed to be equal to the Bondi–Hoyle rate. For higher density star-forming gas, the Bondi–Hoyle rate is boosted by a factor of $(n_{\text{H}}/10^{-1} \text{ cm}^{-3})^2$ to compensate for the lack of a cold, interstellar gas phase and the finite resolution (see Booth & Schaye 2009, for a full discussion). The BH growth rate is related to its accretion rate by $\dot{m}_{\text{BH}} = (1 - \epsilon_{\text{r}})\dot{m}_{\text{accr}}$, where $\epsilon_{\text{r}} = 0.1$ is the radiative efficiency of the BH. The amount of energy coupled to the surroundings of the BH is then given by $\dot{E} = \epsilon_{\text{f}}\epsilon_{\text{r}}\dot{m}_{\text{accr}}c^2$, where c is the speed of light and ϵ_{f} is a free parameter, the ‘feedback efficiency’, that is tuned to reproduce the global BH density at $z = 0$. In the fiducial runs $\epsilon_{\text{f}} = 0.15$. This model reproduces the observed BH scaling relations (Booth & Schaye 2009), and the observed X-ray and optical properties of low-redshift galaxy groups (McCarthy et al. 2010). Note, however, that our simulations have a substantially higher resolution than the $100h^{-1}$ Mpc, 512^3 runs used to compare with low-redshift observations. As shown by Booth & Schaye (2009), the model is not fully converged for intermediate masses with higher resolution, resulting in more efficient BH growth until the growth becomes self-regulating. This model does not include ‘radio-mode’ feedback (which may be necessary to get a sharp exponential drop-off in stellar mass function) and we have not varied the parameters of this model in this paper.

Fig. 2 shows that AGN feedback is the most destructive form of feedback, and in a massive galaxy at $z = 2$ the AGN destroys the gaseous disc. This effect is also visible in panels C–E of Fig. 9, where it is apparent that the ISM and halo gas content (and thus the baryonic content, which is dominated by gas in the halo) of galaxies is strongly suppressed relative to the ‘*REF*’ simulation.

In the very lowest mass resolved haloes ($M_{\text{tot}} < 10^{10.6} M_{\odot}$), ‘*AGN*’ and ‘*REF*’ show virtually identical results because in these haloes BHs have not yet been seeded. However, even slightly above these halo masses, the BH accretion prescription used in these simulations is very efficient at this resolution and the BHs grow and effectively suppress star formation (panels A, B, G and H of Fig. 9). It has been argued that AGN feedback is necessary to suppress rapid cooling of hot halo gas and suppress star formation in high-mass haloes (e.g. Di Matteo, Springel & Hernquist 2005; Croton et al. 2006; Bower, McCarthy & Benson 2008; Booth & Schaye 2009; McCarthy et al. 2010). In panel (B) of Fig. 9, we show the halo SFR as a function of mass for ‘*AGN*’. The effect of AGN is indeed very strong at high masses. Although the galaxy sSFRs are almost an order of magnitude below the observations (panel G) this may not be a problem. The observations are of galaxies that are actively star forming. While in the ‘*REF*’ simulation all galaxies are active, in the ‘*AGN*’ simulation there are both active and passive galaxies at all masses, leading to a median sSFR that lies below the relation obtained for active galaxies alone. The stellar mass function (panel I), slightly undershoots the observed stellar mass function in the range where both observations and simulations exist. This may occur because the SN feedback in the ‘*REF*’ simulation was tuned to reproduce the peak in the global SFR density of the Universe. Including AGN feedback will then underproduce the amount of stars.

5 CONCLUSIONS

We have analysed a large set of high-resolution cosmological simulations from the *OWLS* project (Schaye et al. 2010), focusing on the baryonic properties of the galaxy population at redshift 2. In particular, we studied the effects of variations in the input physics on the stellar mass, SFR, baryon fraction, ISM fraction and gas fraction as a function of halo mass, on the sSFR, ISM mass and the

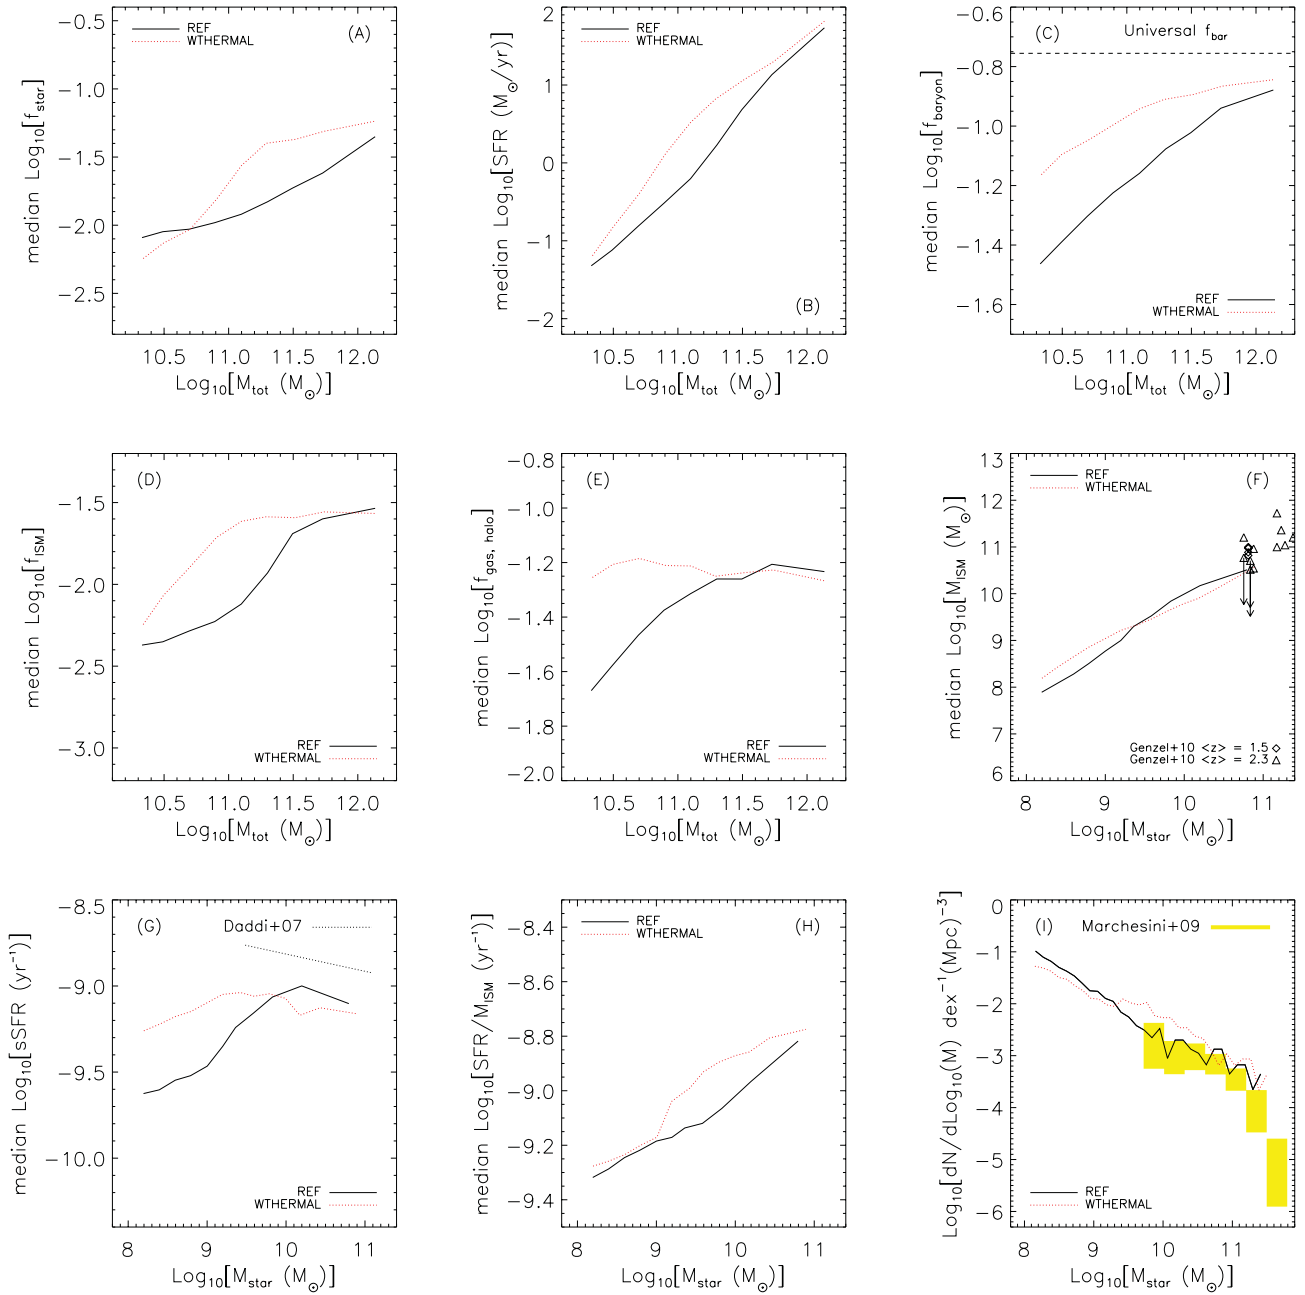


Figure 8. As Fig. 1, but showing only the simulation in which SN feedback is implemented thermally (*W THERM*; red, dotted curve) and *REF* (black, solid curve). Both thermal and kinetic feedback suppress star formation, but for the models presented here, thermal feedback is somewhat weaker, except for the lowest, and in these simulations poorly resolved, halo masses.

gas consumption time-scale as a function of stellar mass, and on the galaxy stellar mass function. In this paper, we focused on variations of the parameters of the sub-grid models for radiative cooling, feedback from star formation and AGN, as well as the box size and the resolution, as shown in Appendix A. In Paper II, we concentrate on the remaining sub-grid prescriptions (e.g. those for star formation), which tend to be less important for galaxy properties other than the mass and density distribution of the ISM.

A central conclusion from this work is that the SFR in star-forming galaxies is self-regulated by galactic winds driven by massive stars. Many of our results can be understood if the SFR adjusts so that the (time averaged) rate at which energy and

momentum are injected is sufficient to drive a large-scale outflow that, on average, balances the gas accretion rate (e.g. Schaye et al. 2010; Davé et al. 2012). Chemical feedback is also important, because, for a fixed redshift and halo mass, the accretion rate is sensitive to the radiative cooling rate and hence to the metallicity.

Our results for variations of the kinetic feedback parameters and radiative cooling rates can be summarized as follows.

(i) Feedback from star formation is only effective if the initial wind velocity used for kinetic feedback is sufficiently high. The required wind velocity increases with the mass of the galaxy and also if metal-line cooling is included. If wind particles are temporarily

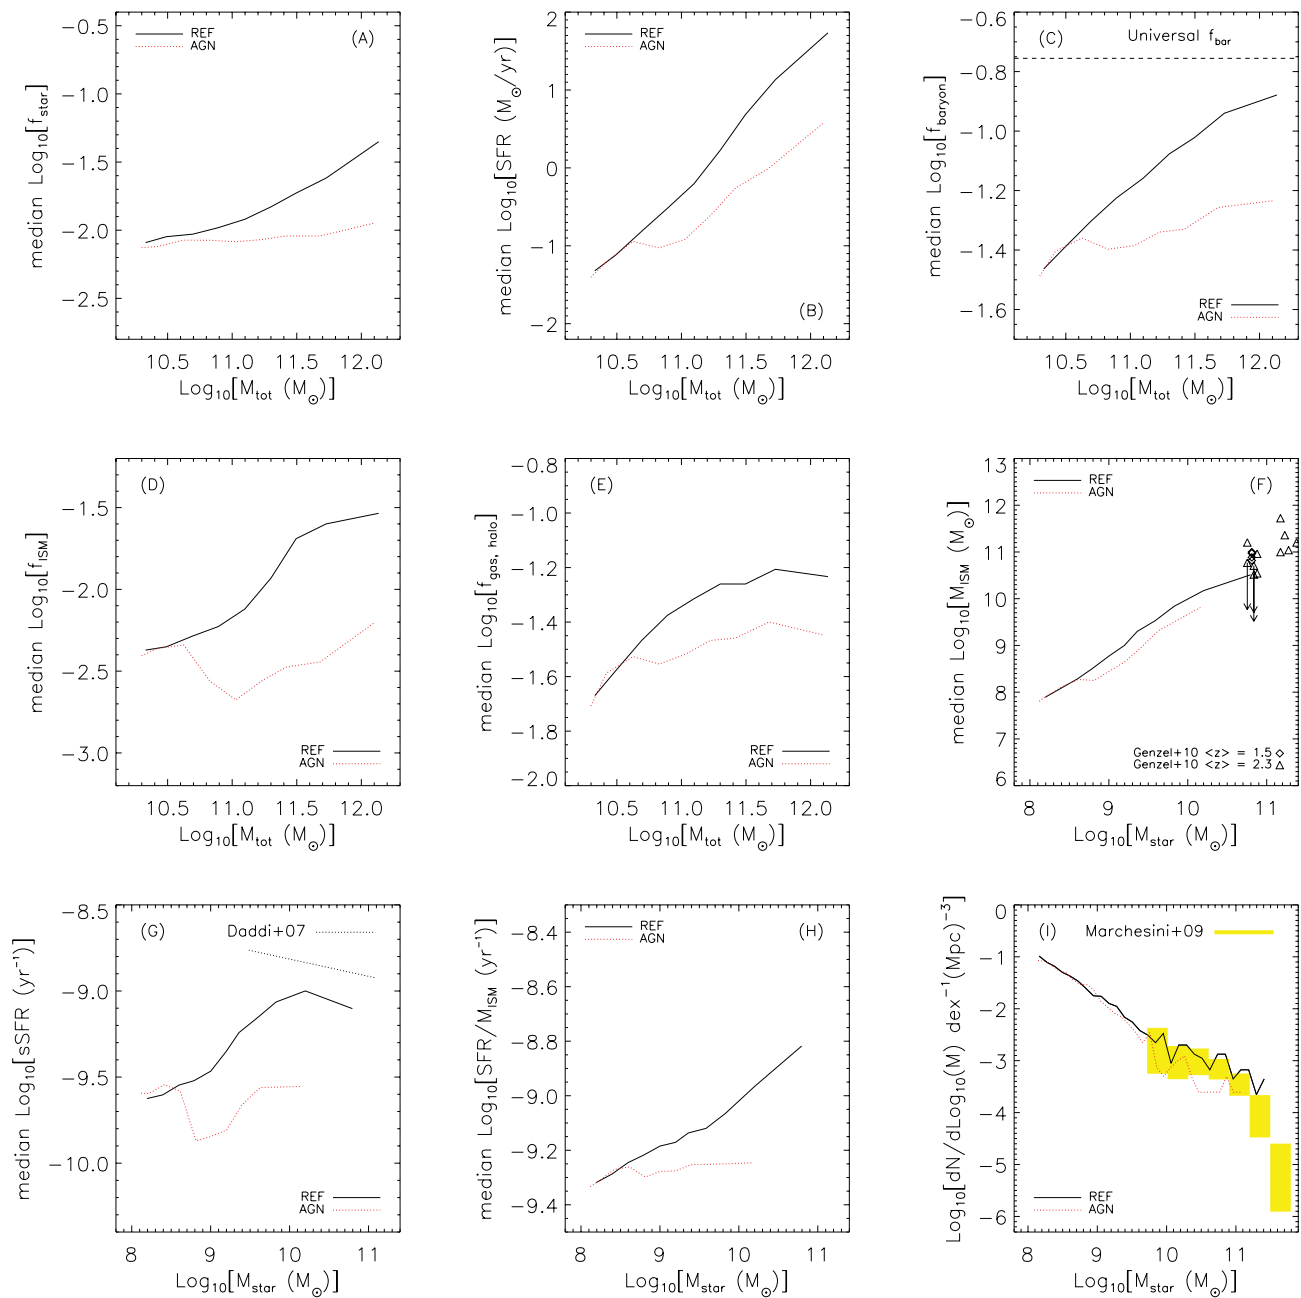


Figure 9. As Fig. 1, but comparing only the simulation with AGN feedback ('AGN'; red, dotted curve) to 'REF' (black solid line). AGN feedback effectively ejects gas from galaxies in all haloes where BH seeds have been placed ($M_{\text{tot}} > 4 \times 10^{10} h^{-1} M_{\odot}$), thus strongly suppressing the SFR.

decoupled from the hydrodynamics, then the feedback remains effective up to much higher masses. All these results can be understood if radiative losses inside the ISM are responsible for the quenching of the winds (Dalla Vecchia & Schaye 2012, see also Creasey et al. 2011).

(ii) If the winds are not quenched, then the SFR is inversely proportional to the amount of energy that is injected per unit stellar mass. For a fixed initial wind velocity, this rate of energy injection is determined by the initial wind mass loading.

(iii) If, on the other hand, the initial wind velocity is too low for the winds to escape the galaxies, then the behaviour of the simulations tends to that of a simulation without any feedback at all: catastrophic cooling resulting in excessive star formation.

The main conclusions drawn from other variations are as follows.

(i) Using a top-heavy IMF for star formation at high pressure mainly influences the simulated haloes through the extra amount of SN energy available per unit stellar mass.

(ii) Feedback can remain highly efficient at all masses if the initial wind velocity is increased with the galaxy mass while keeping the momentum that is injected per unit stellar mass constant, as motivated by models of winds driven by radiation pressure (although the momentum-driven winds used here and in the literature may use much more momentum than is actually available).

(iii) AGN feedback is very efficient at reducing the SFR and gas content, especially at high masses.

We compared our predictions to three different observational results: the molecular gas mass as a function of stellar mass, the sSFR as a function of stellar mass and the galaxy stellar mass function. The latter can be thought of as a convolution of the halo mass function and the stellar mass as a function of halo mass. The comparison with observations revealed the following.

(i) For most simulations, the mass in the ISM in simulated galaxies is slightly lower than the observed molecular gas masses as a function of stellar mass. Note though, that the observed gas masses may have been overestimated (Narayanan et al. 2012).

(ii) Except for models with inefficient feedback, the galaxy stellar mass function is close to the observed one over much of the observed mass range. The shape is different though, with most simulations predicting a steeper low-mass end than (Schechter-like extrapolations of) observational results. Models with higher initial wind mass loading factors predict shallower faint-end slopes, as appears to be required by the observations (see also Bower et al. 2012). None of the simulations predict a clear exponential cut-off at the high-mass end, but this could be due to our limited box size or lack of ‘radio-mode’ AGN feedback.

(iii) The predicted sSFRs as a function of stellar mass tend to be lower than observed. For the observed mass range, the discrepancy is worst if the feedback is efficient. The high observed sSFR can only be matched if the feedback is inefficient at the observed mass, but highly efficient at lower masses. The observed negative slope in the relation between the sSFR and stellar mass is only reproduced for galaxies for which feedback is inefficient.

Thus, there is tension between the comparison of simulated and observed galaxy stellar mass functions and the comparison between simulated and observed sSFRs. The high observed sSFRs are difficult to match unless feedback suddenly becomes inefficient at the stellar masses for which observations are available ($M_* > 10^{9.5} M_\odot$). It can, however, not remain inefficient as the stellar mass increases or else the galaxy mass function would become too high at the high-mass end. Even though there may be other explanations, one possibility is that our simulations form too many stars in low-mass systems at higher redshift. That would result in a steep low-mass end of the mass function, a lower gas reservoir in galaxies at $z = 2$ and consequently, too low sSFRs.

We have shown that winds driven by feedback from star formation determine the main properties of galaxies residing in haloes of a given mass (the scatter among the red lines in Fig. 1 is much larger than for the blue lines that represent the simulations appearing in Paper II). Even for a fixed amount of feedback energy per unit stellar mass, variations in the sub-grid implementation, e.g. different wind velocities and mass loadings (that can be either constant or fixed functions of local physical conditions), provide us with considerable freedom to ‘tune’ galaxy properties. This freedom can be exploited to match observations spanning a wide range of masses, which would provide the simulations with some of the attractions of semi-analytic models. However, this potential success comes also with the disadvantages of such models: the underlying physics may remain poorly understood. As higher resolution simulations become feasible, the freedom provided by sub-grid models to generate galactic outflows in cosmological simulations will be reduced, as the results become less sensitive to the manner in which the energy is injected (Dalla Vecchia & Schaye 2012).

Furthermore, improvement in our understanding of the physics that determines the global properties of galaxies will likely come from theoretical models and observations focusing on galactic winds. The physics of star formation is less crucial as self-regulation

implies that the time-averaged, galaxy-wide SFRs are determined by the large-scale inflows and the efficiency with which star formation drives galactic winds.

ACKNOWLEDGEMENTS

The authors kindly thank the anonymous referee for a thoughtful report that helped improve the paper. Furthermore, we thank Danilo Marchesini for useful discussions about the observations of the stellar mass function and for providing his results in digital form. Furthermore, we are grateful to the members of the *OWLS* collaboration for useful discussions about, and comments on, this work. The simulations presented here were run on Stella, the LOFAR BlueGene/L system in Groningen, on the Cosmology Machine at the Institute for Computational Cosmology in Durham (which is part of the DiRAC Facility jointly funded by STFC, the Large Facilities Capital Fund of BIS, and Durham University) as part of the Virgo Consortium research programme, and on Darwin in Cambridge. This work was sponsored by the National Computing Facilities Foundation (NCF) for the use of supercomputer facilities, with financial support from the Netherlands Organization for Scientific Research (NWO), also through a VIDI grant. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement 278594-GasAroundGalaxies and from the Marie Curie Training Network CosmoComp (PITN-GA-2009-238356). VS acknowledges support through SFB 881, The Milky Way System, of the DFG.

REFERENCES

- Agertz O., Teyssier R., Moore B., 2009, *MNRAS*, 397, L64
 Allen C. W., 1976, in Allen C. W., ed., *Astrophysical Quantities*. Athlone, London
 Baugh C. M., Lacey C. G., Frenk C. S., Granato G. L., Silva L., Bressan A., Benson A. J., Cole S., 2005, *MNRAS*, 356, 1191
 Birnboim Y., Dekel A., 2003, *MNRAS*, 345, 349
 Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
 Booth C. M., Schaye J., 2010, *MNRAS*, 405, L1
 Bouchet P., Lequeux J., Maurice E., Prevot L., Prevot-Burnichon M. L., 1985, *A&A*, 149, 330
 Bower R. G., McCarthy I. G., Benson A. J., 2008, *MNRAS*, 390, 1399
 Bower R. G., Benson A. J., Crain R. A., 2012, *MNRAS*, 422, 2816
 Brook C. B., Kawata D., Gibson B. K., Flynn C., 2004, *MNRAS*, 349, 52
 Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000 (BC03)
 Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *ApJ*, 533, 682
 Cen R. Y., Ostriker J. P., Jameson A., Liu F., 1990, *ApJ*, 362, L41
 Ceverino D., Dekel A., Bournaud F., 2010, *MNRAS*, 404, 2151
 Ceverino D., Dekel A., Mandelker N., Bournaud F., Burkert A., Genzel R., Primack J., 2012, *MNRAS*, 420, 3490
 Chabrier G., 2003, *PASP*, 115, 763
 Choi J.-H., Nagamine K., 2012, *MNRAS*, 419, 1280
 Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, *MNRAS*, 319, 168
 Conselice C. J., Blackburne J. A., Papovich C., 2005, *ApJ*, 620, 564
 Crain R. A., Eke V. R., Frenk C. S., Jenkins A., McCarthy I. G., Navarro J. F., Pearce F. R., 2007, *MNRAS*, 377, 41
 Creasey P., Theuns T., Bower R. G., Lacey C. G., 2011, *MNRAS*, 415, 3706
 Croton D. J. et al., 2006, *MNRAS*, 365, 11
 Daddi E., Cimatti A., Renzini A., Fontana A., Mignoli M., Pozzetti L., Tozzi P., Zamorani G., 2004, *ApJ*, 617, 746
 Daddi E. et al., 2007, *ApJ*, 670, 156
 Daddi E., Dannerbauer H., Elbaz D., Dickinson M., Morrison G., Stern D., Ravindranath S., 2008, *ApJ*, 673, L21
 Daddi E. et al., 2010, *ApJ*, 713, 686

- Dalla Vecchia C., Schaye J., 2008, MNRAS, 387, 1431
 Dalla Vecchia C., Schaye J., 2012, MNRAS, 426, 140
 Davé R., 2008, MNRAS, 385, 147
 Davé R., Oppenheimer B. D., Finlator K., 2011, MNRAS, 415, 11
 Davé R., Finlator K., Oppenheimer B. D., 2012, MNRAS, 421, 98
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
 Di Matteo T., Springel V., Hernquist L., 2005, Nat, 433, 604
 Dolag K., Borgani S., Murante G., Springel V., 2009, MNRAS, 399, 497
 Erb D. K., Steidel C. C., Shapley A. E., Pettini M., Reddy N. A., Adelberger K. L., 2006, ApJ, 647, 128
 Evrard A. E., 1988, MNRAS, 235, 911
 Genzel R. et al., 2010, MNRAS, 407, 2091
 Guo Q., White S., Li C., Boylan-Kolchin M., 2010, MNRAS, 404, 1111
 Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays, CEA, Saclay
 Haas M. R., 2010, PhD thesis, Univ. Leiden, Leiden
 Haas M. R., Schaye J., Booth C. M., Dalla Vecchia C., Springel V., Theuns T., Wiersma R. P. C., 2013, MNRAS, 435, 2955 (Paper II)
 Haberman S. M., Anderson J. P., James P. A., 2010, ApJ, 717, 342
 Hernquist L., Katz N., 1989, ApJS, 70, 419
 Kauffmann G., Colberg J. M., Diaferio A., White S. D. M., 1999, MNRAS, 303, 188
 Kay S. T., Pearce F. R., Frenk C. S., Jenkins A., 2002, MNRAS, 330, 113
 Kennicutt R. C., Jr, 1998, ARA&A, 36, 189
 Kereš D., Katz N., Weinberg D. H., Davé R., 2005, MNRAS, 363, 2
 Kereš D., Vogelsberger M., Sijacki D., Springel V., Hernquist L., 2012, MNRAS, 425, 2027
 Komatsu E. et al., 2011, ApJS, 192, 18
 Kravtsov A. V., Berlind A. A., Wechsler R. H., Klypin A. A., Gottlöber S., Allgood B., Primack J. R., 2004, ApJ, 609, 35
 Kravtsov A. V., Nagai D., Vikhlinin A. A., 2005, ApJ, 625, 588
 Lacey C., Cole S., 1994, MNRAS, 271, 676
 Maness H. et al., 2007, ApJ, 669, 1024
 Maraston C., 2005, MNRAS, 362, 799
 Marchesini D., van Dokkum P. G., Förster Schreiber N. M., Franx M., Labbé I., Wuyts S., 2009, ApJ, 701, 1765
 Marchesini D. et al., 2010, ApJ, 725, 1277
 McCarthy I. G. et al., 2010, MNRAS, 406, 822
 McCrady N., Gilbert A. M., Graham J. R., 2003, ApJ, 596, 240
 Mori M., Yoshii Y., Tsujimoto T., Nomoto K., 1997, ApJ, 478, L21
 Mortlock A., Conselice C. J., Bluck A. F. L., Bauer A. E., Grützbauch R., Buitrago F., Ownsworth J., 2011, MNRAS, 413, 2845
 Murray N., Quataert E., Thompson T. A., 2005, ApJ, 618, 569
 Narayanan D., Bothwell M., Davé R., 2012, MNRAS, 426, 1178
 Ocvirk P., Pichon C., Teysier R., 2008, MNRAS, 390, 1326
 Okamoto T., Frenk C. S., Jenkins A., Theuns T., 2010, MNRAS, 406, 208
 Oppenheimer B. D., Davé R., 2006, MNRAS, 373, 1265
 Oppenheimer B. D., Davé R., 2008, MNRAS, 387, 577
 Oppenheimer B. D., Davé R., 2009, MNRAS, 395, 1875
 Oppenheimer B. D., Davé R., Kereš D., Fardal M., Katz N., Kollmeier J. A., Weinberg D. H., 2010, MNRAS, 406, 2325
 Prevot M. L., Lequeux J., Prevot L., Maurice E., Rocca-Volmerange B., 1984, A&A, 132, 389
 Puchwein E., Springel V., 2012, MNRAS, 428, 2966
 Reddy N. A., Erb D. K., Steidel C. C., Shapley A. E., Adelberger K. L., Pettini M., 2005, ApJ, 633, 748
 Ryu D., Vishniac E. T., Chiang W.-H., 1990, ApJ, 354, 389
 Sales L. V., Navarro J. F., Schaye J., Dalla Vecchia C., Springel V., Haas M. R., Helmi A., 2009, MNRAS, 399, L64
 Sales L. V., Navarro J. F., Schaye J., Dalla Vecchia C., Springel V., Booth C. M., 2010, MNRAS, 409, 1541
 Salpeter E. E., 1955, ApJ, 121, 161
 Scannapieco C., Tissera P. B., White S. D. M., Springel V., 2006, MNRAS, 371, 1125
 Scannapieco C. et al., 2012, MNRAS, 423, 1726
 Schaye J., 2004, ApJ, 609, 667
 Schaye J., Dalla Vecchia C., 2008, MNRAS, 383, 1210
 Schaye J. et al., 2010, MNRAS, 402, 1536
 Seljak U., Zaldarriaga M., 1996, ApJ, 469, 437
 Somerville R. S., Primack J. R., 1999, MNRAS, 310, 1087
 Sommer-Larsen J., Götz M., Portinari L., 2003, ApJ, 596, 47
 Springel V., 2005, MNRAS, 364, 1105
 Springel V., 2010, MNRAS, 401, 791
 Springel V., Hernquist L., 2003a, MNRAS, 339, 289
 Springel V., Hernquist L., 2003b, MNRAS, 339, 312
 Springel V., Di Matteo T., Hernquist L., 2005, MNRAS, 361, 776
 Steidel C. C., Shapley A. E., Pettini M., Adelberger K. L., Erb D. K., Reddy N. A., Hunt M. P., 2004, ApJ, 604, 534
 Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, MNRAS, 373, 1074
 Stolte A., Brandner W., Grebel E. K., Lenzen R., Lagrange A.-M., 2005, ApJ, 628, L113
 Thacker R. J., Couchman H. M. P., 2000, ApJ, 545, 728
 Thomas P. A., Couchman H. M. P., 1992, MNRAS, 257, 11
 Vale A., Ostriker J. P., 2004, MNRAS, 353, 189
 van de Voort F., Schaye J., Booth C. M., Haas M. R., Dalla Vecchia C., 2011, MNRAS, 414, 2458
 van Dokkum P. G., 2008, ApJ, 674, 29
 Vogelsberger M., Sijacki D., Kereš D., Springel V., Hernquist L., 2012, MNRAS, 425, 3024
 Weidner C., Kroupa P., Pflamm-Altenburg J., 2011, MNRAS, 412, 979
 Wiersma R. P. C., Schaye J., Smith B. D., 2009a, MNRAS, 393, 99
 Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, MNRAS, 399, 574
 Wilkins S. M., Hopkins A. M., Trentham N., Tojeiro R., 2008, MNRAS, 391, 363
 Yang X., Mo H. J., van den Bosch F. C., Zhang Y., Han J., 2012, ApJ, 752, 41
 Zel'Dovich Y. B., 1970, A&A, 5, 84

APPENDIX A: NUMERICAL CONVERGENCE

Here, we explore the sensitivity of our results to numerical resolution and simulation box size, and assess down to what mass limits we can consider our results numerically converged. All simulations used in this section employ the same physical model as the ‘REF’ simulation, but are run with different particle numbers and box sizes. We denote the simulations with ‘LXXXYYYY’, where XXX is the size of the simulation box in comoving h^{-1} Mpc and YYY is the number of particles per spatial dimension (for both DM and baryons we use YYY^3 particles). In this nomenclature, the ‘REF’ simulation is ‘L025N512’. In order to independently examine the effects of box size and resolution, we compare two sets of simulations:

(i) *L025N512*, *L012N256* and *L006N128*: These simulations have the same resolution, but the box size is varied in steps of a factor of 2 from 25 to 6.25 $\text{Mpc } h^{-1}$. These three runs are shown as black curves with different line styles (solid, dashed and dot-dashed, respectively) in all figures in this appendix

(ii) *L025N512*, *L025N256* and *L025N128*: These simulations have the same box size, but different resolutions. The DM particle mass in these simulations is 6.3×10^6 , 5.1×10^7 and $4.1 \times 10^8 M_{\odot} h^{-1}$, respectively and the maximum proper gravitational softening is 0.5, 1 and 2 $\text{kpc } h^{-1}$, respectively. In all figures in this appendix, these three simulations are shown as solid black, blue and red curves, respectively.

All plots in this appendix include all haloes identified by the FoF algorithm (the lowest mass haloes contain 20 DM particles). Fig. A1 shows the median SFR as a function of halo mass. It is immediately clear that simulation box size has no influence on the SFRs, as the lines with different line styles (*L025N512*, *L012N256* and *L006N128*) show almost perfect agreement. The only effect of

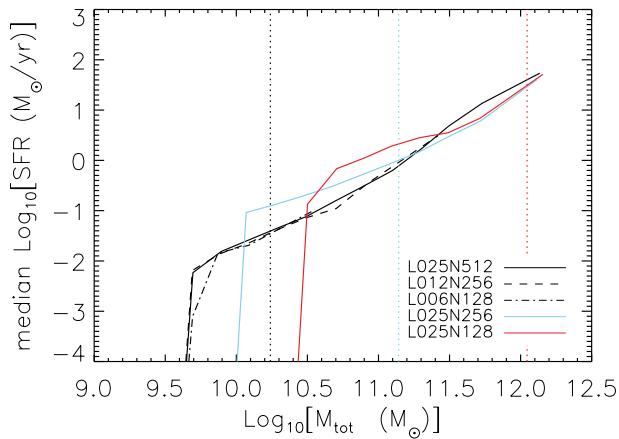


Figure A1. Median SFR as a function of halo mass at $z = 2$ for five simulations with different particle numbers and/or box sizes as indicated in the legend. The vertical dotted lines indicate the mass of 2000 DM particles in the simulations shown by the curves in the corresponding colours. At the low-mass end, the median SFR falls to zero, as more than half of the haloes in a bin do not have gas particles with a density above the star formation threshold. Above a mass corresponding to 2000 DM particles per halo, the SFR as a function of halo mass is reasonably well resolved.

box size is that a larger box allows one to sample more massive, rarer objects.

The simulations are less well converged with respect to resolution (compare the solid black, red and blue curves in Fig. A1). The vertical, dotted lines denote 2000 times the DM particle mass in the simulations of the same colour. The halo SFRs are reasonably converged above these halo masses, so we employ 2000 DM particles as our resolution limit when comparing galaxy properties as a function of halo mass. The difference between the ‘L025N512’ and ‘L025N256’ runs are $\lesssim 0.5$ dex. It is clear from Fig. A1 that at low masses the SFR is always slightly overpredicted in lower resolution simulations, so the true convergence level is likely better than 0.5 dex. Note that the difference between the high- and intermediate-resolution simulations just above the convergence criterium for the intermediate-resolution simulation (the vertical blue dotted line) is much smaller, so one expects that the difference between our high-resolution simulations and even higher resolution simulations at our imposed particle limit will be smaller than 0.5 dex. Observe that the halo mass regime where the median SFR is zero is effectively removed when demanding a minimum number of 2000 DM particles per halo, because more than half of the haloes with a low number of particles do not have any gas particles with densities above the star formation threshold.

The build-up of stellar mass is influenced by the SFR at all epochs prior to the epoch at which it is measured. As all haloes were initially small and thus poorly resolved, the early build-up of stellar mass is underestimated. We therefore expect the convergence of the sSFR as a function of stellar mass to be worse than that of the SFR as a function of total halo mass.

Fig. A2 shows the sSFR as a function of stellar mass. The vertical cut-off at the low-mass end corresponds again to haloes for which the median SFR is zero. The mass range over which the sSFR appears approximately converged with respect to resolution corresponds to ~ 100 star particles (indicated by the vertical dotted lines). Convergence can be seen by comparing the solid black and blue lines rightwards of the blue, dotted line and by comparing the solid blue and red curves rightwards of the red, dotted line. We

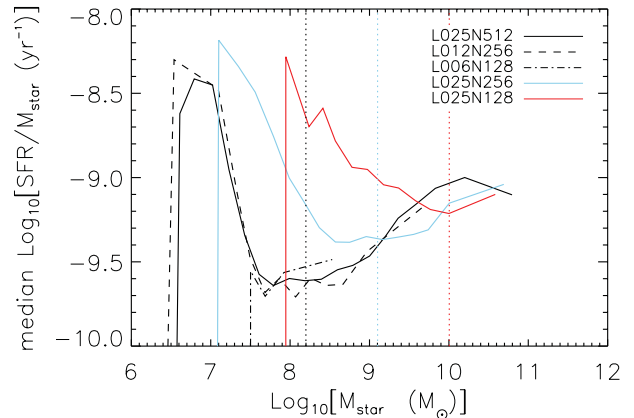


Figure A2. The median sSFRs of haloes as a function of their stellar mass at $z = 2$ for five simulations with different particle numbers and/or box sizes as indicated in the legend. The vertical dotted lines indicate the mass corresponding to 100 star particles in the simulations shown by the curves in the corresponding colours. The sharp cut-off at low masses again stems from the fact that there is a minimum to the (non-zero) SFR. Right of the vertical dotted lines the sSFRs are reasonably well converged.

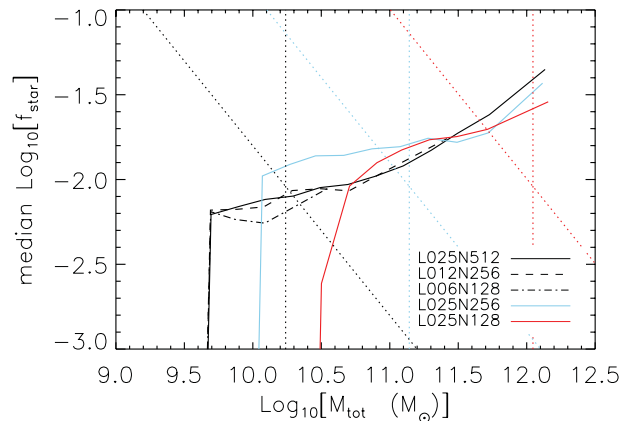


Figure A3. Median stellar mass fraction as a function of halo mass at $z = 2$ for five simulations with different particle numbers and/or box sizes as indicated in the legend. The vertical dotted lines indicate the mass corresponding to 2000 DM particles in the simulations shown by the curves in the corresponding colours. The diagonal black dotted line indicates the relation for haloes with 100 star particles (the cut that is made in the rest of the paper where relations with stellar mass are shown). As can be seen, in the highest resolution simulation, the cuts made throughout this paper in minimum number of DM particles and the minimum number of star particles roughly correspond to each other. At lower resolutions, the cut in DM particle number is more stringent.

note that, as expected, the same trends are found for $\text{SFR}/M_{\text{gas}}$ (not shown).

Finally, in Fig. A3 we show the median stellar mass fraction as a function of halo mass. The vertical dotted lines indicate our adopted resolution limit of 2000 DM particles. The diagonal dotted lines indicate the stellar mass fraction for haloes consisting of 100 star particles, which is our resolution limit for plots with stellar mass on the horizontal axis. The fact that for a given resolution (i.e. line colour), the solid curve intersects the two dotted lines in about the same place, implies that the cuts of 100 star particles and 2000 DM particles are comparable for these simulations. Above the limit of 2000 DM particles, the stellar mass fractions are nearly converged. At lower resolution, the cut in DM particle number is more stringent

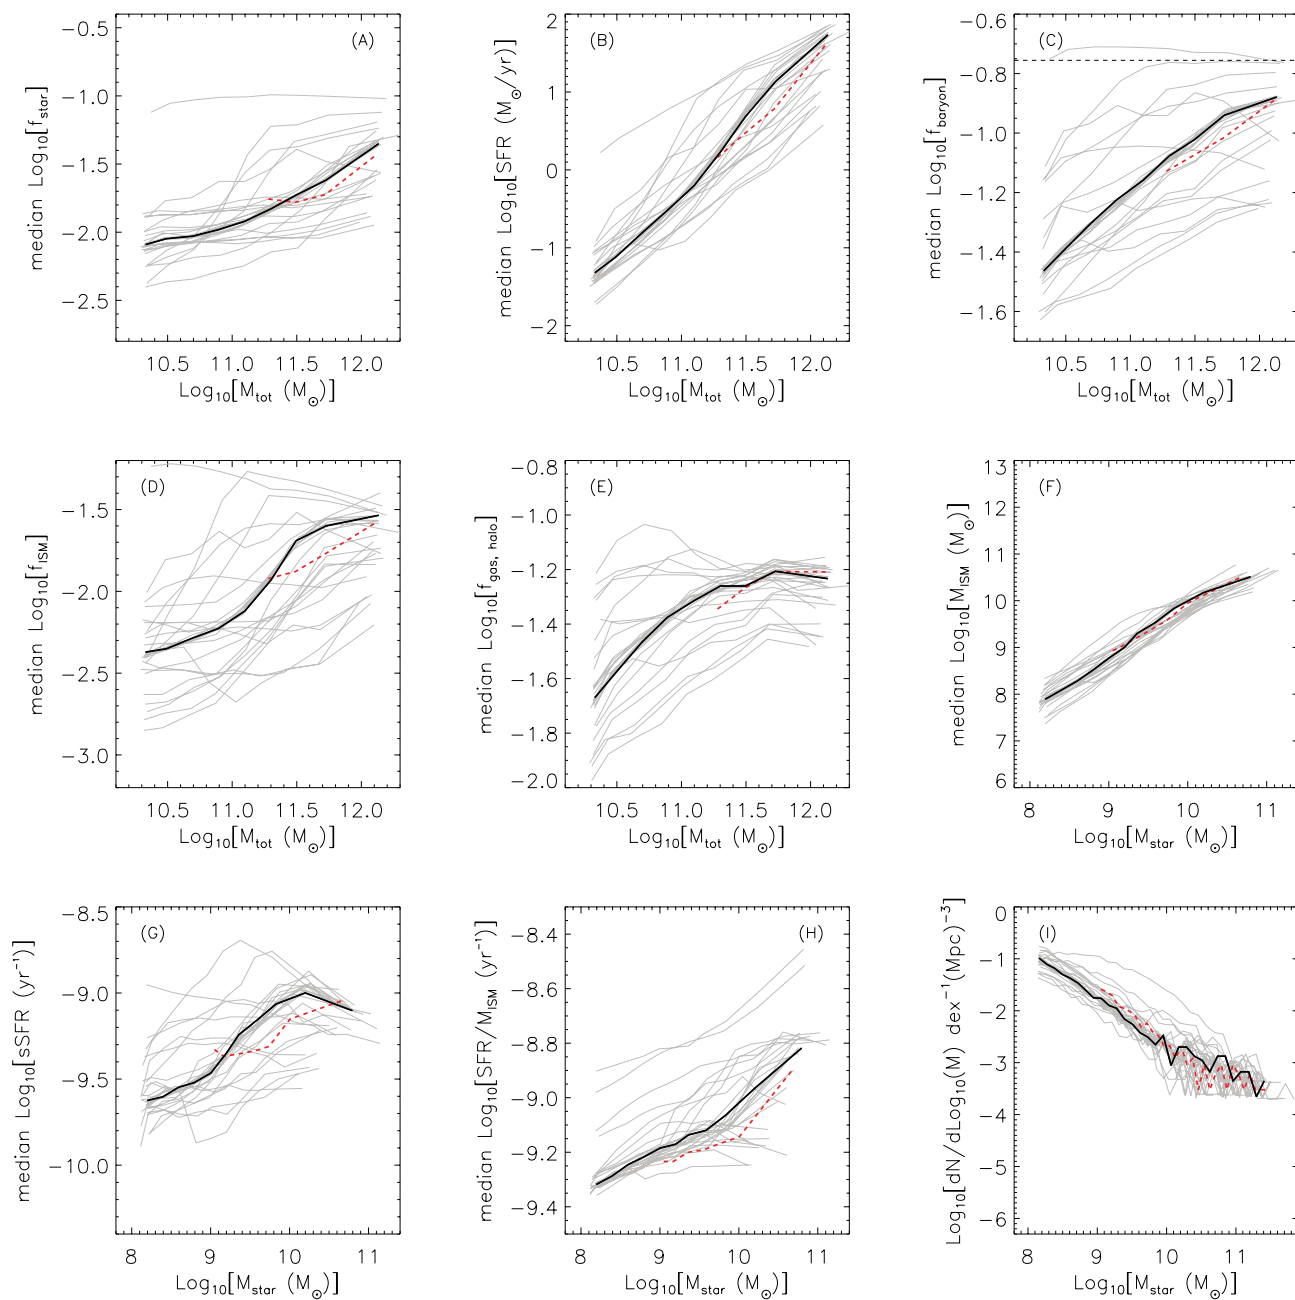


Figure A4. As Fig. 1, but for the reference input physics at two different resolutions: *L025N512* (black solid lines) and *L025N256* (red dashed lines), which has an eight times lower mass resolution. For both sets of simulations, the convergence criteria as used in the main body of the paper are used: a minimum of 2000 DM particles for panels (A) through (E) and a minimum of 100 star particles for panels (F) through (I). The thin grey lines are all physics variations at the same resolution as the black solid line. The spread due to physics variations, at fixed resolution, is comparable to or larger than the difference between simulation results at different resolution.

than the cut on the minimum number of star particles. Therefore, throughout the paper we use a minimum of 2000 DM particles when we plot quantities as a function of halo mass.

In summary, our resolution requirements are as follows: To avoid biasing the results due to resolution effects, we impose a cut of 2000 DM particles when looking at relations with total halo mass and of 100 stars particles when investigating correlations with stellar mass. In Fig. A4 we show all simulations used in this paper and paper II in grey lines in the background, as well as the ‘REF’ simulation at the resolution used in the paper (*L025N512*, black solid lines) and at lower resolution (*L025N256*, red dashed lines). The spread

in simulation results due to physics variations at fixed resolution is comparable to or larger than the difference between simulation results using the resolution that is used in the main body of the paper and a mass resolution that is eight times lower than that. Even for results where the difference between two simulations with different physics at the same resolution is smaller than the difference between the simulations with the same physics but at different resolution, the difference between the simulations at the same resolution can still be significant. In such cases, it becomes hard to tell whether the differences are due to a difference in resolution dependence of the two physics implementations or really only due to the physics.

However, because they are run at the same resolution, it is unlikely that resolution affects the qualitative conclusions we draw from the comparison of the different models.

APPENDIX B: THE EFFECT OF HALO DEFINITION

Throughout this paper, all of the halo masses that we quote are FoF (Davis et al. 1985) masses, meaning that all particles associated with a given FoF halo contribute to its mass. This particular definition of mass is somewhat arbitrary, so we check in this appendix whether using a different halo mass definition would significantly affect our results. In particular, we compare FoF masses to those returned by SUBFIND (Dolag et al. 2009), for both the main haloes and the sub-haloes. SUBFIND iteratively calculates, for each FoF halo, the mass of both baryonic and collisionless matter that is gravitationally bound to the same structure.

In Fig. B1, we show the galaxy stellar mass function in FoF (solid curve) and SUBFIND (dotted curve) haloes. The two stellar mass functions are virtually identical. This is largely because the stars live preferentially in the centres of haloes, and changing the

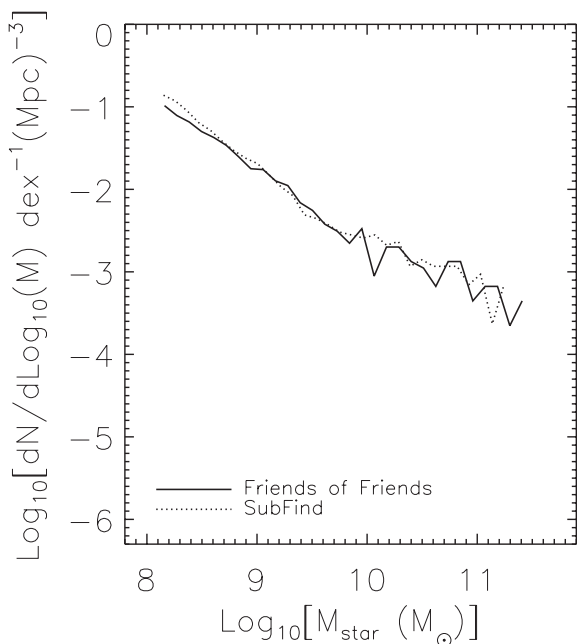


Figure B1. Galaxy stellar mass function for the ‘REF’ simulation with two different halo mass definitions. The solid (dotted) curve shows the stellar mass function when FoF (SUBFIND) haloes are used. Results for mass functions are independent of the halo mass definition used, suggesting that any results plotted as a function of stellar mass are robust to changes in the halo mass definition.

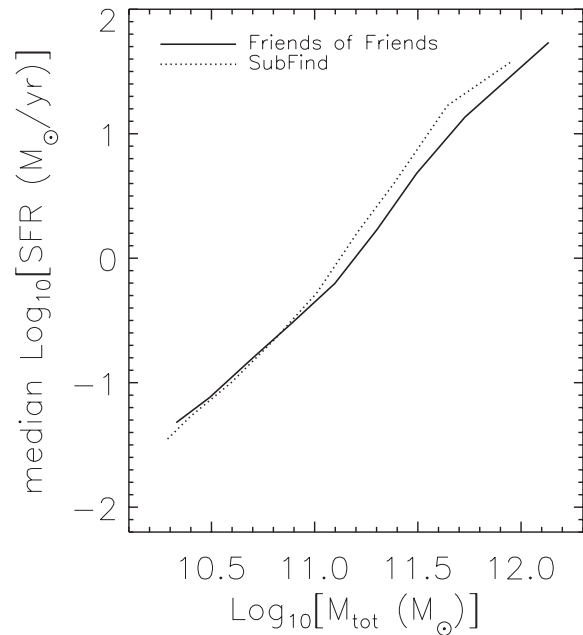


Figure B2. Relation between halo mass and galaxy SFR for the ‘REF’ simulation with two different halo mass definitions. The solid (dotted) curve shows the relation for FoF (SUBFIND) haloes. Although at the highest masses, the different mass definitions give results that differ by up to 0.1 dex, this difference is far smaller than the magnitudes of the effects we are probing, and also affects all of the different simulations equally.

halo mass definition affects mainly particles at the edges of the halo. This plot suggests that all of our results that are plotted as a function of stellar mass are robust to the definition of halo used.

The halo mass is somewhat more affected by the mass definition. In Fig. B2, we show the median SFR as a function of halo mass. The solid curve shows the FoF haloes, and the dotted curve shows the SUBFIND haloes. Here, we see almost perfect agreement at the low-mass end, but the overall effect of SUBFIND at the high-mass end is to unbind haloes that are artificially linked together by the FoF algorithm, so at the highest masses the median halo mass can decrease by up to 0.1 dex. We note, however, that this same effect will exist in all of the simulations analysed here, so although these curves may shift by a small amount, the main focus of this paper is on the relative differences between the simulations, and these remain totally unaffected by the halo mass definition. Note also that the SUBFIND algorithm only removes mass from FoF haloes and never adds any. Many studies use the mass of a sphere containing an overdensity of ~ 200 as total mass for the main sub-haloes, which results in halo masses closer to the FoF halo masses used in this paper.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.