

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

Semi-Analytic Modeling of Galaxy Formation for the Future: The Challenge of Emulating Cosmological Hydrodynamical Simulations

Permalink

<https://escholarship.org/uc/item/9xc1v7c9>

Author

Pandya, Viraj

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

**SEMI-ANALYTIC MODELING OF GALAXY FORMATION FOR
THE FUTURE: THE CHALLENGE OF EMULATING
COSMOLOGICAL HYDRODYNAMICAL SIMULATIONS**

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ASTRONOMY AND ASTROPHYSICS

by

Viraj Pandya

September 2021

The Dissertation of Viraj Pandya
is approved:

Professor Piero Madau, Chair

Professor Kevin Bundy

Dr. Rachel Somerville

Peter Biehl
Vice Provost and Dean of Graduate Studies

Copyright © by

Viraj Pandya

2021

Table of Contents

List of Figures	vi
List of Tables	xx
Abstract	xxi
Acknowledgments	xxiii
1 Introduction	1
1.1 History of comparing SAMs and hydrodynamical simulations	4
1.2 A more complete picture of stellar feedback in simulations	8
1.3 The need for a more flexible SAM with explicit preventative stellar feedback	12
1.4 Outline of this work	14
2 Comparing galaxy formation in semi-analytic models and hydrodynamical simulations	16
2.1 Introduction	16
2.2 Model Descriptions	19
2.2.1 FIRE-2 Simulations	19
2.2.2 Santa Cruz Semi-Analytic Model	21
2.3 Analysis	28
2.3.1 Generating merger trees and SAM predictions	29
2.3.2 Computing bulk and flow quantities in the simulations	30
2.4 Results	33
2.4.1 Stellar, ISM and CGM mass scalings at $z = 0$	35
2.4.2 Stellar mass histories	37
2.4.3 Star formation stochasticity	40
2.4.4 ISM mass histories	40
2.4.5 CGM mass histories	42
2.4.6 Halo baryon fraction evolution	43
2.4.7 Halo mass inflow rates	45
2.4.8 ISM inflow rates	47

2.4.9	ISM outflow rates	49
2.4.10	Halo outflow rates	51
2.5	Discussion	52
2.5.1	Interpreting the model discrepancies	54
2.5.2	A simple preventative stellar feedback model	63
2.5.3	Results using dark matter only simulations	70
2.6	Summary	73
3	Characterizing simulated multi-phase supernova-driven winds	80
3.1	Introduction	80
3.2	Simulation Description	82
3.3	Analysis	85
3.3.1	Accurately defining outflows	85
3.3.2	Multi-phase outflow selection criteria	92
3.3.3	Computing wind loading factors	96
3.4	ISM wind loading factors	106
3.4.1	Comparison to FIRE-1	109
3.4.2	Global halo circular velocity	110
3.4.3	Multi-phase ISM mass loadings	113
3.4.4	Multi-phase ISM momentum loadings	116
3.4.5	Multi-phase ISM energy loadings	119
3.4.6	Multi-phase ISM metal loadings	122
3.4.7	Wind radial velocities	125
3.4.8	Trends with SFR and ISM gas mass surface densities	130
3.4.9	SF burstiness, dense ISM gas fractions and inner CGM virialization	131
3.5	Halo wind loading factors	134
3.5.1	Bernoulli velocity versus potential depth	135
3.5.2	Halo mass loading	138
3.5.3	Halo momentum loading	139
3.5.4	Halo energy loading	139
3.5.5	Halo metal loading	140
3.6	Discussion	140
3.6.1	Comparison to theoretical expectations and other simulations . .	144
3.6.2	Systematic uncertainties	152
3.7	Summary	155
4	A new semi-analytic model with explicit preventative stellar feedback	160
4.1	Introduction	160
4.2	FIRE-2 Simulations: Sample and analysis	161
4.2.1	Simulation sample	161
4.2.2	Particle tracking fluxes	161
4.3	A Prototype Simulation-Calibrated Semi-Analytic Model	165
4.3.1	Guiding philosophy	165

4.3.2	Halo inflow	167
4.3.3	ISM inflow (CGM cooling)	169
4.3.4	Star formation	173
4.3.5	Stellar feedback (ejective)	174
4.3.6	Stellar feedback (preventative)	177
4.3.7	Chemical evolution	181
4.4	Case study of modeling a low-mass dwarf	183
4.4.1	Impact of new rapid CGM cooling model	183
4.4.2	Impact of SN heating of the IGM	183
4.4.3	Impact of SN heating of the CGM	185
4.4.4	Stellar mass assembly history	186
4.4.5	Wind energy coupling efficiency	187
4.5	Results for the entire FIRE-2 ensemble	190
4.5.1	Evolution of individual halos	190
4.5.2	IGM heating efficiency function	197
4.5.3	CGM heating efficiency function	198
4.5.4	CGM depletion time as a heating diagnostic	198
4.6	Discussion	201
4.6.1	Significance and implications of preventative SN feedback	201
4.6.2	Remaining uncertainties in the prototype SAM	204
4.7	Summary	213
5	Conclusions and Future Work	217
5.1	Summary	217
5.2	Future Work	219
5.2.1	Observable consequences of preventative stellar feedback	219
5.2.2	Wind recycling	222
5.2.3	More faithful modeling of CGM physics	223
5.2.4	Extension to satellites, higher mass halos and SMBHs	225
5.2.5	Towards an open-source hierarchical Bayesian SAM	227

List of Figures

- 2.1 An illustration of our zone definitions for analyzing bulk and flow quantities in the FIRE-2 simulations. The background image shows the projected gas density distribution of the MW-mass halo m12m at $z = 0$ (purple is low density, with green and yellow representing progressively higher densities, respectively). The solid yellow circles represent the virial radius (outer circle) and the “edge” of the ISM (inner circle). The dashed white circles demarcate the outer limits of the virial and ISM shells ($1.0 - 1.1 \times R_{\text{vir}}$ and $0.1 - 0.2 \times R_{\text{vir}}$, respectively) through which mass inflow and outflow rates are computed. Stellar and ISM masses are computed using all star and gas particles within $0.1 \times R_{\text{vir}}$, respectively, whereas CGM masses are computed using all gas particles between $0.1 - 1.0 \times R_{\text{vir}}$. These definitions are well-matched for comparison to the Santa Cruz SAM. 34

2.2	Our FIRE-2 measurements (circles) and SAM predictions (crosses) for stellar-to-halo mass ratios (left), ISM-to-stellar mass ratios (middle), and CGM-to-halo mass ratios (right) at $z = 0$. Halos are colored according to their mass bin (low-mass dwarfs in purple, intermediate-mass dwarfs in green and MW-mass halos in red). We also show observationally inferred scaling relations for median stellar-to-halo mass ratios (Rodríguez-Puebla et al., 2017; Behroozi et al., 2019) and ISM-to-stellar mass ratios (Boselli et al., 2014; Peeples et al., 2014; Calette et al., 2018). We do not show observational estimates of CGM-to-halo mass ratios since they are highly uncertain and the SAM is not calibrated to match observed CGM properties. The SAM and FIRE-2 agree relatively well with each other and with these observations in terms of stellar-to-halo and ISM-to-stellar mass ratios at a fixed mass (the ISM-to-stellar mass ratios predicted by the SAM for low-mass dwarfs would agree better with observations if we only included the cold atomic and molecular phases; Popping et al., 2014; Somerville et al., 2015). By comparison, the SAM and FIRE-2 predictions for CGM-to-halo mass ratios disagree dramatically with each other, especially for the dwarfs where the SAM predictions are generally lower by ~ 3 orders of magnitude.	38
2.3	Comparison of the stellar mass assembly history measured in FIRE-2 and as predicted by the SAM. Top: logarithmic ratio of the SAM and FIRE-2 time series color-coded by mass bin (m10 halos in purple, m11 halos in green, and m12 halos in red). Bottom: individual stellar mass assembly histories for one representative halo from each mass bin (m10q left, m11c middle, m12f right). All time series are smoothed over ~ 1 Gyr for easier visual comparison. With one exception (m10v), the SAM generally reproduces the FIRE-2 stellar mass assembly histories within a factor of two.	39
2.4	Normalized SFHs for FIRE-2 measurements (colored curves) and SAM predictions (black). The time series are the instantaneous SFHs divided by the corresponding 100 Myr boxcar-smoothed SFHs. The halos are ordered based on increasing $z = 0$ halo mass from top to bottom (m10 halos in purple, m11 in green, m12 in red). The m10 and m11 FIRE-2 halos are bursty at all times, and the m12 halos are bursty at early times but not late times. In contrast, the SAM predicts much less SFH stochasticity.	41
2.5	Similar to Figure 2.3 but now for ISM mass as a function of cosmic age. The SAM agrees relatively well with FIRE-2 for the m12 halos except at very early times, and this is also true for the m11 halos, albeit with more scatter. But the systematic discrepancy for the m10 halos remains at the order of magnitude level at all times, consistent with the $z = 0$ SAM excess in Figure 2.2.	42

- 2.12 The ratio of the cumulative mass ejected from the halo versus from the ISM as a function of cosmic time for FIRE-2 (left panel) and the SAM (right panel). The lines are color-coded according to $z = 0$ halo mass. The SAM qualitatively follows the FIRE-2 trends for the m11 and m12 halos at $z \lesssim 2$ (i.e., cosmic ages $\gtrsim 3$ Gyr): an increasingly larger fraction of winds are able to leave the halo in progressively lower mass halos. But in the FIRE-2 simulations, the m10 halos strikingly tend to have expelled more material through R_{vir} than has ever left the ISM boundary ($0.1R_{\text{vir}}$), implying significant entrainment of ambient CGM material by the outflows (this is also true for the progenitors of all halos at very early times $z \gtrsim 6$). Since the ratio cannot exceed 1 in the SAM by construction, it asymptotes to 1 for the low-mass dwarfs (all of their winds leave the halo). 53
- 2.13 A cartoon that schematically illustrates the results of our comparison between the SAM and FIRE-2 (restricted to $z = 0$ for simplicity). On the left we show a representative dwarf (m10q) and on the right a representative MW-mass halo (m12f). For each halo, the left half portrays the SAM and the right half depicts FIRE-2. From inside out, we show the bulk masses of stars (black), ISM (blue), CGM (red) and the ejected gas reservoir (magenta; restricted to the SAM since there is no clear definition of this component for FIRE-2). The opacity can be used to compare the mass of a single component between the two models or the mass of different components within a single model. The arrows illustrate inflows and outflows between the different bulk mass components (note that the purple arrows show the total halo gas accretion rate, not the recycling rate). Larger size arrows convey higher flow rates. Note how the stellar masses agree very well between the two models for both galaxies despite significant differences in the other bulk components. 63

2.14	<p>The halo gas accretion efficiency ($\dot{M}_{\text{in,gas}}/\dot{M}_{\text{in,DM}}$ in the virial shell) as a function of the halo virial mass for the FIRE-2 simulations at $z = 0$ (black points) and $z = 2$ (magenta points). If gas accretion perfectly tracked DM accretion at R_{vir} with the universal baryon fraction ($f_b = 0.158$ from Planck Collaboration et al., 2016) as commonly assumed, then the halos would all lie along the horizontal solid gray line. The Okamoto et al. (2008) model describing the suppression of halo gas accretion due to the ionizing UV background is shown with the black dashed line for $z = 0$ and the magenta dashed line for $z = 2$. The FIRE-2 halo gas accretion efficiencies fall below the expectation from UV background heating alone. The dotted lines show the behavior of our simple preventative stellar feedback model if we assume the halo energy loading factor is a constant 0.1; we see that it is incapable of explaining the data points, in fact showing the opposite trend at low masses. However, the solid lines show that our model can explain the data points remarkably well if we assume that the halo energy loading factor is preferentially higher for dwarfs ($\eta_E = \varepsilon_{\text{heat}}\eta_M$ with $\varepsilon_{\text{heat}} = 0.01$, implying that η_E is of order unity for low-mass dwarfs). The agreement is better at $z = 0$ than $z = 2$, suggesting either a redshift dependence for $\varepsilon_{\text{heat}}$ or that our simple model is breaking down.</p>	71
2.15	<p>A few relevant halo properties measured in the full hydrodynamical simulations (magenta lines) and the corresponding DM-only simulations (black lines). These are the 5 halos for which corresponding DM-only FIRE-2 runs exist. The DM-only halo properties are very similar to the hydro-based halo properties, with the M_{vir} and R_{vir} being lower in the hydro version by only 10-20% on average. The major systematic difference is in the halo concentration which tends to be lower in the dwarfs in the hydro version (presumably due to stellar feedback adiabatically expanding the halo center) whereas it is higher for the MW-mass halos in the hydro run (presumably due to the significant stellar mass adiabatically contracting the halo center).</p>	74

2.16	Time series of the main properties considered for the five individual halos with corresponding DM-only simulations. From left to right: m10q, m10v, m11q, m12f and m12m. From top to bottom: stellar mass, ISM mass, CGM mass, halo gas mass inflow rate, ISM mass inflow rate, ISM mass outflow rate and halo mass outflow rate. In every panel, the solid black line shows the prediction of the SAM when run on the DM-only simulation merger trees. The other two curves follow the same convention as the individual halo panels in the figures from the main body of the text: solid colored curves for the FIRE-2 measurements and dotted colored curves for the SAM predictions using the full hydrodynamical simulation merger trees. The colors show the mass bin that each halo belongs to (purple for m10, green for m11 and red for m12). The main takeaway is that our conclusions do not change if we use the SAM results from the DM-only simulation merger trees: the new SAM predictions agree with our fiducial ones relatively well, and hence the DM-only-based SAM trends relative to FIRE-2 remain qualitatively the same.	75
3.1	Visualizing and quantifying a strong outflow at $z \sim 0$ in a MW-mass halo (m12f). This is a single frame from a movie that is available for download.	93
3.2	Similar to Figure 3.1 but now for a low-mass dwarf (m10q). This movie frame is during a major outflow episode at high redshift $z \sim 3.5$, a regime where dwarfs are often characterized as having mass loadings of ~ 100 or more. If we divide the values of the individual ISM mass outflow rate peaks by their associated preceding SFR spikes (bottom-left panel), we would indeed infer instantaneous mass loadings of ~ 100 . The halo-scale mass loadings (magenta) are even larger due to entrainment of CGM gas by outflows. Note how there is a hot bubble in the projection panels created by the strong outflows.	94
3.3	Temperature distribution of ISM-scale winds. These distributions are based on \dot{M}_{out} -weighted averages over three broad redshift bins. Solid gray vertical lines are cooling physics-based temperature cuts at 10^3 and 10^5 K. Dashed vertical colored lines illustrate example virial temperatures of the inner halo (at $0.1R_{\text{vir}}$) for representative halos from each mass bin. The virial temperatures roughly align with the temperature distribution peak for hot (virialized) outflows. Cold outflows are more prominent at high redshift. The two lowest mass halos generally do not have multiphase outflows. Note that the m13 halos were only run down to $z = 1$ and so are absent from the bottom panel.	97

3.4	<p>Illustration of our automated algorithm for measuring instantaneous loading factors. These are three representative 1 Gyr time chunks showing m12f at high-redshift when SF is bursty (top), m12f at lower redshift when SF is more continuous (middle), and m10q at the same lower redshift when its SF is still bursty (bottom). In each panel, the SFH is shown in black, the original unshifted mass outflow rate history as transparent magenta, and the time-shifted mass outflow rate history as opaque magenta. The cyan dots and horizontal lines identify peaks and their baselines, respectively. Above each detected peak, we write the burst-integrated wind mass, stellar mass, and instantaneous mass loading factor. Note how the instantaneous mass loading is $\sim 5 - 10$ times weaker in m12f at lower redshift compared to high redshift. Note also how highly mass-loaded the two bursts in m10q are despite being at the same lower redshift, and also how far apart these two bursts are in time (~ 600 Myr with zero SF in between).</p>	107
3.5	<p>Comparison of our ISM-scale mass and metal loading factors as a function of stellar mass to previous FIRE-1 work (Muratov et al., 2015, 2017; Anglés-Alcázar et al., 2017a). <i>Top</i>: Our fiducial mass loadings are roughly $\sim 2\times$ lower than Muratov et al. (2015), who found $\eta_{\text{M,ISM}} \propto (M_*/M_\odot)^{-0.35}$ (excluding the low-redshift m12 halos from their fit). This is due to our Bernoulli velocity wind criterion excluding slower, turbulent flows which would otherwise lead to larger mass loadings as in Muratov et al. (2015); the transparent symbols show that we would agree with FIRE-1 remarkably well if we use the same wind selection criteria. We also plot the particle tracking-based measurements of mass loadings in FIRE-1 from Anglés-Alcázar et al. (2017a), which are even higher since they track outflows directly out of the ISM (much of which recycles back). <i>Bottom</i>: Despite our stricter wind selection criterion, our metal loadings agree with Muratov et al. (2017), suggesting that the switch from FIRE-1 to FIRE-2 subgrid physics affected outflow metallicities. The transparent symbols show that we would predict even larger metal loadings than FIRE-1 if we use the same wind selection criteria as Muratov et al. (2017). Nevertheless, our overall conclusion is similar: nearly all metals produced by Type II SNe are ejected from the ISM of dwarfs but retained within the ISM of more massive galaxies.</p>	111

- 3.6 ISM mass loading as a function of halo circular velocity. The big markers show our fiducial redshift-averaged measurements whereas the small dots show our instantaneous mass loading measurements color-coded by redshift (V_{vir} for the latter is an \dot{M} weighted average over individual burst intervals). Muratov et al. (2015) found a shallower slope of ISM mass loading with halo circular velocity for FIRE-1 halos with $V_{\text{vir}} > 60$ km/s and interpreted it to mean a transition from energy-driven winds in dwarfs to momentum-driven winds in higher mass halos (dotted colored lines). We do not see this flattening with our more stringent outflow selection criteria in FIRE-2 except possibly at the very massive end ($V_{\text{vir}} \gtrsim 300$ km s $^{-1}$) starting at intermediate redshift. Instead, our measurements are roughly consistent with a single power law that goes as V_{vir}^{-2} at high-redshift (red line), with a steepening at later times (green and blue lines). 114
- 3.7 Evolution of multi-phase ISM mass loading factors with stellar mass. The bigger markers show our fiducial redshift-averaged measurements and the smaller dots are our instantaneous burst measurements color-coded by redshift (M_* for the latter is an \dot{M} weighted average over each individual burst interval). The horizontal gray line denotes order unity and the other lines show approximate fits (see text). *Top*: For all outflow phases combined, η_{M} is less than 1 for massive halos and rises to ~ 100 for dwarfs. *Bottom-left*: The fraction of mass loading in the cold phase is far less than one but correlates strongly with redshift. High redshift halos can have $\sim 10\%$ of their mass loading in the cold phase, but this drops to less than 1% for most halos at later times. *Bottom-middle*: the warm outflow phase dominates by mass fraction in the dwarfs. Note also the much tighter correlation and lack of redshift dependence compared to the cold phase. *Bottom-right*: the hot phase dominates in the massive halos but steadily drops off toward lower mass halos. Interestingly, when the total $\eta_{\text{M,ISM}} < 1$, the hot phase dominates, whereas in the dwarfs with $\eta_{\text{M,ISM}} \gg 1$ the warm phase dominates. 117
- 3.8 Similar to Figure 3.7 but now for the ISM momentum loading factor. *Top*: The overall momentum loading factor is of order unity for higher redshift dwarfs, dropping to ~ 0.1 for some dwarfs at $z \sim 0$. For massive halos, overall momentum loadings are generally less than 0.1. *Bottom-left*: The fraction of momentum loading in the cold phase is negligible except for high redshift dwarfs where it can approach $\sim 10\%$. *Bottom-middle*: The warm phase carries nearly all of the momentum in the lowest mass dwarfs, gradually dropping to $\lesssim 10\%$ for more massive halos. *Bottom-right*: In contrast, the hot phase carries nearly all of the momentum in massive halos, gradually dropping to $\sim 10\%$ for the lowest mass dwarfs. 120

- 3.9 Similar to Figure 3.7 but now for the evolution of ISM energy loading factor. *Top*: Overall energy loadings are of order unity for dwarfs at high-redshift, and between 0.1 and 1 for lower redshift dwarfs. MW and m13 halos have $\eta_E \sim 0.1$ especially at lower redshift. *Bottom-left*: The fraction of energy loading carried by the cold phase is negligible, except in some low-mass dwarfs at high-redshift where it is almost $\sim 10\%$ (signifying large kinetic energy). *Bottom-middle*: The warm phase carries only about $\sim 10\%$ of the energy loadings in massive halos but becomes increasingly important for low-mass dwarfs, where its fractional contribution approaches $\sim 50\%$ or higher. *Bottom-right*: The hot phase carries $\sim 50\%$ of the energy in the lowest mass dwarfs and effectively becomes the sole carrier of energy in massive halos ($f_{\text{hot}} \sim 100\%$). . . . 123
- 3.10 Similar to Figure 3.7 but now showing evolution of the ISM metal loading factor. *Top*: Overall metal loadings are of order unity for dwarfs, i.e., nearly all metals produced by SNe are ejected from the ISM of dwarfs. In contrast, the SN metal yield is mostly retained within the ISM of massive halos. *Bottom-left*: The fraction of metals carried by the cold phase is generally negligible for all halos except in high-redshift dwarfs and intermediate-redshift massive halos (for which $f_{\text{cold}} \sim 10\%$). *Bottom-middle*: The warm phase carries nearly all of the metals in the lowest mass dwarfs, and becomes progressively less important for more massive halos, dropping to $\sim 10\%$. *Bottom-right*: In contrast, the hot phase carries nearly all of the metals in massive halos compared to only $\sim 10\%$ in the lowest mass halos. 126
- 3.11 Similar to Figure 3.3 but now showing the distributions of ISM mass outflow rate in bins of radial velocity for wind particles classified as cold (left column), warm (middle column) or hot (right column). The different rows show the average distributions over our fiducial large redshift bins, where snapshots with higher total $\dot{M}_{\text{out,ISM}}$ are given higher weight in the average. 128
- 3.12 Analogous to Figure 3.15 but now, following common practice, we plot the mass-flux-weighted average radial velocity versus halo virial velocity. From left to right: we plot this for the cold, warm and hot ISM outflows. The solid gray line is the one-to-one mapping between radial velocity halo circular velocity; the dashed and dotted lines are twice and five times the circular velocity, respectively. We see that generally cold and warm outflows cluster around $\approx 2V_{\text{vir}}$, with slightly lower radial velocities in the m13 halos. The hot outflows tend to be faster, approaching $\approx 5 \times V_{\text{vir}}$ on average for some dwarfs. Interestingly the radial velocity of hot outflows in some halos can be less than V_{vir} , which either suggests deceleration due to interactions or that the slower component of the hot wind dominates (as illustrated in the bottom-right panel of Figure 3.1 and Figure 3.2). . . 129

- 3.13 Instantaneous ISM mass loading factor versus \dot{M} -weighted average gas mass surface density (top) and SFR surface density (bottom) within individual burst windows. There is a clear negative correlation with gas mass surface density such that $\eta_M \propto \Sigma_{\text{gas}}^{-1.2}$ (diagonal black line; excluding the m13 halos from the fit). The trend with SFR surface density follows $\eta_M \propto \Sigma_{\text{SFR}}^{-0.5}$ but the deviation from a simple power law is more apparent. 132
- 3.14 Instantaneous ISM mass loading factors versus three “derived” physical properties. From left to right: (1) the maximum SFR in the burst interval divided by the 1 Gyr-averaged SFR, (2) the dense ISM gas fraction, and (3) the cooling time to free-fall time ratio at $0.1R_{\text{vir}}$, which is a measure of inner halo virialization. Each of these physical properties are \dot{M} -weighted averages within individual burst windows. The large black symbols and errorbars denote binned medians with 25th and 75th percentiles. We see that instantaneous mass loadings are higher for more powerful starbursts (i.e., when the SFH is locally bursty). There is also a weak trend where instantaneous η_M tends to be higher when the dense ISM gas fraction is lower, although there is a lot of scatter. Finally, the instantaneous η_M steadily declines as $t_{\text{cool}}/t_{\text{ff}}$ increases, with the mass loading becoming $\ll 1$ when the inner halo is virialized as indicated by $t_{\text{cool}}/t_{\text{ff}} > 2$ (this is the case for massive halos whereas dwarfs have a non-virialized inner halo and higher mass loadings). 134
- 3.15 Average mass-flux-weighted Bernoulli velocity (i.e., specific kinetic energy plus enthalpy) of multi-phase ISM outflows versus the difference in escape velocity between $0.1R_{\text{vir}}$ and R_{vir} (a proxy for the potential difference). This gives a sense of whether outflows can be expected to reach R_{vir} in the absence of interactions (given our ISM wind selection criteria, outflows should make it to $\sim 0.5R_{\text{vir}}$ at minimum). Cold and warm outflows (left and middle panels) in some dwarfs have up to twice the energy needed to make it to R_{vir} , but in the more massive halos the cold/warm outflow energy is comparable to the potential difference. In contrast, hot outflows (right) have Bernoulli velocities that are far in excess of the energy needed to make it to R_{vir} . This is obvious for lower mass halos where the hot outflows contain up to $\sim 5\times$ more energy than needed to escape the halo (hence these outflows can be expected to travel very large distances, probably becoming unbound). Hot outflows in low-redshift MW halos have just enough energy to reach R_{vir} . Hot SN-driven outflows can also escape from m13 halos at high redshift but not necessarily at intermediate redshift. 137

3.16	Comparing “escaping” wind loading factors for the ISM and virial shells. The top row is ISM loading factors using only the subset of ISM outflows that have enough energy to reach $2R_{\text{vir}}$ instead of just $0.5R_{\text{vir}}$. The middle row is halo-scale loading factors using only outflows at R_{vir} that have enough energy to get to at least $2R_{\text{vir}}$. The bottom row is the ratio of these halo-scale and ISM-scale loading factors. From left to right we show mass, momentum, energy and metal loading factors. <i>Mass</i> : dwarf halo mass loadings are a few times larger than the ISM mass loadings, perhaps indicative of additional swept up material. Low-redshift MW halos have a ratio close to ~ 1 (recall that our ISM loadings exclude slower outflows, which may be substantial but unlikely to reach R_{vir}). m13 halos show a larger ratio at intermediate redshift than at high redshift. <i>Momentum</i> : Dwarf outflows often have more momentum at the halo scale than at the ISM scale, in contrast to low-redshift MW halos whose outflows have comparable momentum at R_{vir} and $0.1R_{\text{vir}}$. <i>Energy</i> : In dwarfs, halo energy loadings are comparable to ISM energy loadings. In contrast, for MW halos at low redshift, the halo-scale energy loadings are ~ 0.1 smaller than their ISM-scale energy loadings. <i>Metals</i> : In dwarf halos, most metals leaving the ISM also leave the halo. In low-redshift MW halos, only $\sim 10\%$ of ISM metal outflows leave the halo (surprisingly, intermediate redshift m13 halos have high metal loadings).	141
4.1	ISM gas depletion time as a function of V_{vir} and redshift for the FIRE-2 halos. We can use our continuous fit to estimate SFRs in the prototype SAM.	175
4.2	ISM mass loading factor (left) and halo mass loading factor (right) as a function of V_{vir} and redshift. Note that these are based on particle tracking fluxes without any cut on minimum Bernoulli velocity (i.e., most closely corresponding to a zero-radial-velocity threshold between outflows and inflows). We can use our continuous fits in the SAM, although for some individual dwarf halos the scatter can be substantial and should be taken into account.	177
4.3	Bernoulli velocity of ISM outflows (left) and halo outflows (right). We can assign energy outflow rates to the SAM using the continuous parameterizations as a function of both V_{vir} and redshift (see text). This can then serve as the basis for preventative heating of the CGM and IGM by SN-driven winds (more important in dwarfs).	181
4.4	Inflow and outflow metallicities ($\dot{Z}/\dot{M}/Z_{\odot}$) for the ISM shell and halo shell. Our parameterizations that depend on V_{vir} and redshift enable easy implementation into the prototype SAM.	184

4.5	Time evolution of CGM mass (top), halo gas inflow rate (middle) and halo gas outflow rate (bottom) for the low-mass dwarf m10q. FIRE measurements are in black, Santa Cruz SAM in cyan. There are 3 lines for the prototype SAM: without any IGM/CGM heating (magenta), with only IGM heating (yellow), and with both IGM and CGM heating (red). We see that our new rapid CGM cooling model is capable of roughly reproducing the FIRE CGM mass whereas the SC-SAM is low by orders of magnitude. However, this agreement is achieved still with very high halo inflow/outflow rates, whereas enabling SN heating of the IGM and CGM leads to good agreement in both the mass and flux time series. . .	188
4.6	Analogous to Figure 4.5 but now the time evolution of ISM mass and associated ISM inflow/outflow rates for the low-mass dwarf m10q. . . .	189
4.7	Analogous to Figure 4.5 but now the time evolution of stellar mass. . .	190
4.8	Time evolution of the accretion suppression factor (top) and wind energy coupling efficiency (bottom) for the low-mass dwarf m10q. The IGM-heating-only scenario for the prototype SAM is shown with the yellow lines, and the CGM plus IGM heating scenario in red. When $f_{\text{SN}} < 1$, that means there was otherwise excess accretion $\Delta\dot{M}$ that needed to be offset. We see that this low-mass dwarf requires both CGM and IGM heating. Furthermore, we see that nearly all of the available wind energy needs to couple to the CGM to offset excess ISM accretion, and there needs to be substantial wind energy coupling with the IGM as well. . . .	191
4.9	Time evolution of CGM mass and associated inflow/outflow rates (left column), ISM mass and corresponding inflow/outflow rates (middle column), and stellar mass and underlying SFR (right column) for the low-mass dwarf m10y. Line styles are the same as in previous figures.	192
4.10	Similar to Figure 4.9 but for the low-mass dwarf m10z.	192
4.11	Similar to Figure 4.9 but for the intermediate-mass dwarf m11a.	193
4.12	Similar to Figure 4.9 but for the intermediate-mass dwarf m11b.	193
4.13	Similar to Figure 4.9 but for the intermediate-mass dwarf m11c.	194
4.14	Similar to Figure 4.9 but for the intermediate-mass dwarf m11q.	194
4.15	Similar to Figure 4.9 but for the intermediate-mass dwarf m11v.	195
4.16	Similar to Figure 4.9 but for the intermediate-mass dwarf m11f.	195
4.17	Similar to Figure 4.9 but for the MW-mass halo m12i.	196
4.18	Similar to Figure 4.9 but for the MW-mass halo m12f.	196
4.19	Similar to Figure 4.9 but for the MW-mass halo m12m.	197

4.20	The halo accretion suppression factor, $f_{\text{SN,IGM}}$ as a function of halo virial velocity in three broad redshift bins: $z = 0.0 - 0.5$ (purple), $z = 0.5 - 2.0$ (green) and $z = 2.0 - 4.0$ (red). A low value of $f_{\text{SN,IGM}}$ indicates greater suppression whereas $f_{\text{SN,IGM}} = 1$ means no heating was assumed by the prototype SAM. There is an overall trend for dwarfs requiring greater suppression albeit with significant scatter. Higher-mass halos do not require significant halo accretion suppression (as expected, since at high-redshift they are feed by cold filaments, and at later times, their virial temperature is higher than the ambient IGM gas enabling efficient accretion).	199
4.21	Analogous to Figure 4.20 but now for CGM heating. There is much more heating assumed in the prototype SAM overall for the CGM as compared to the IGM, with some dwarfs requiring the “cold/rapid mode” CGM cooling rate $M_{\text{CGM}}/t_{\text{dyn}}$ to be scaled down by as much as $f_{\text{SN,CGM}} \sim 0.1$. The MW-mass halos at late times do not require as much CGM heating, though that may be due to the lack of a fountain flow term and underproduction of CGM metallicity in the prototype SAM (leading to lower cooling rates versus FIRE to begin with).	200
4.22	The effective CGM depletion time ($M_{\text{CGM}}/\dot{M}_{\text{in,ISM}}$ as a function of halo virial velocity in three broad redshift bins. The FIRE-2 data are shown on the left, the prototype SAM without any SN heating of the IGM/CGM in the middle, and the prototype SAM with full SN heating of both the IGM/CGM on the right. The horizontal dashed lines illustrate the halo dynamical time (at R_{vir}) at the midpoint of the three redshift bins. While there is a lot of scatter at lower redshifts, the high-redshift halos would naively be expected to undergo “cold/rapid mode” accretion where $t_{\text{dep}} = t_{\text{dyn}}$ since the ISM inflow rate is limited by the free-fall time of the CGM (middle panel). However, in FIRE, the high-redshift dwarfs have preferentially longer CGM depletion times suggestive of heating. Indeed, the prototype SAM is able to qualitatively reproduce the negative-sloped high-redshift trend when SN heating of the CGM/IGM is enabled. . . .	202
4.23	Illustration of the ISM extending to $\sim 0.2R_{\text{vir}}$ for an intermediate-mass dwarf due to cycles of expansion and re-contraction driven by strong SN feedback (top). This leads to systematically higher measurements of the ISM inflow rate in FIRE than would be predicted from CGM cooling alone by the prototype SAM (bottom panel). With multi-snapshot particle tracking in the future, we will distinguish first-time ISM accretion from recycled ISM accretion to guide the implementation of a new rapid fountain flow reservoir in the prototype SAM.	206

4.24 Representative example showing that we can roughly reproduce the overall masses of the CGM, ISM and stars as well as the metallicity of the ISM and stars. However, the CGM metallicity is significantly underestimated. The gray curves are bulk measurements from FIRE. The green curves are from simply integrating the \dot{M} and \dot{Z} terms as measured via our particle tracking. 208

List of Tables

3.1	Average properties and loading factors of the FIRE-2 halos in our low redshift bin ($z = 0.0 - 0.5$).	101
3.2	Identical to Table 3.1 but now for our intermediate redshift bin ($z = 0.5 - 2.0$).	102
3.3	Identical to Table 3.1 but now for our high redshift bin ($z = 2.0 - 4.0$).	103
3.4	Catalog of instantaneous burst properties for all halos at $z < 4$	108

Abstract

Semi-Analytic Modeling of Galaxy Formation for the Future: The Challenge of
Emulating Cosmological Hydrodynamical Simulations

by

Viraj Pandya

It is a grand challenge to develop a simple physical model that can faithfully emulate the results of high-resolution galaxy formation simulations and scale up their predictions to cosmologically representative volumes. Semi-analytic models (SAMs) are a promising means of tracking the physical processes associated with galaxy formation, but they are based on many approximations that have not been rigorously tested. In this dissertation, I demonstrate that SAMs can be re-tooled to roughly reproduce the overall baryon cycle of simulated galaxies.

I start by comparing predictions between the Santa Cruz SAM and the Feedback In Realistic Environments (FIRE) simulation suite. I show that both models predict remarkably similar stellar and interstellar medium (ISM) mass, but differ dramatically in terms of the underlying mass fluxes and circumgalactic medium (CGM) mass. The SAM predicts much higher gas accretion rates for dwarfs than FIRE-2, and compensates by requiring higher mass outflow rates. The SAM also predicts orders of magnitude lower CGM mass for dwarfs than FIRE-2. I argue that these model discrepancies are caused by the lack of preventative stellar feedback in the SAM and by its assumptions for halo gas cooling and recycling.

I then uniquely characterize the mass, momentum, energy and metallicity of multi-phase galactic winds in the same simulations. Among other results, I find that winds from dwarfs conserve nearly all of the available supernova energy, momentum and metal mass out to beyond the halo radius. I leverage these insights to prototype a new SAM in which supernova-driven winds can shock-heat gas within and beyond the CGM, thereby suppressing cooling and accretion. With this simple preventative feedback model, my SAM is able to simultaneously reproduce the bulk masses and mass fluxes of the simulated halos. Rapid inner halo recycling and the chemical enrichment of the CGM remain poorly understood, but extracting additional data from the simulations can pin down these and other uncertainties for SAMs. I discuss some observable consequences of my work and implications for the quest to develop a more complete, realistic and standard physical model of galaxy formation.

Acknowledgments

This thesis would not have been possible without the encouragement, guidance and support of many people. First and foremost, I am deeply grateful to my advisors Rachel Somerville and Kevin Bundy, and my committee chair Piero Madau. Thank you to Rachel for introducing me to the problem of modeling galaxy formation in a cosmological context when I was an undergraduate math and economics major at Rutgers University taking your “Galaxies and the Milky Way” class. I have learned so much from working with you to develop a next-generation semi-analytic model of galaxy formation (building on your pioneering thesis work here at UC Santa Cruz many years ago) and I feel tremendously lucky to have been afforded the opportunity to spend so much time working with you and others at the Flatiron Institute. Thank you to Kevin for taking me on as a theory student: I have learned so much about observational and instrumental astronomy as a member of your group, and I greatly appreciate all of your support, trust and independence-nurturing advising style. I am beyond excited to work with you in the future to apply my new model to interpret data from current and next-generation facilities. Thank you to Piero for chairing my thesis committee, keeping me on track and helping me brainstorm follow-up project ideas motivated by my thesis that we can work on together in the future.

It has been a privilege to pursue a PhD in the Department of Astronomy and Astrophysics at UC Santa Cruz while being surrounded by so many inspiring and accomplished faculty, postdocs and fellow grad students. Special thanks go to Joel Primack,

Sandy Faber, David Koo, Aaron Romanowsky, Jean Brodie and Francesco Belfiore for sharing their wisdom and mentoring me on many of my early projects. Thank you also to my many collaborators at the Flatiron Institute without whose guidance this thesis would not have been possible – special thanks go to Daniel Anglés-Alcázar, Drummond Fielding, John Forbes, Greg Bryan, Chris Hayward and Yakov Faerman. Double thanks to Greg for sponsoring my post-doc applications and teaching me so much in such a short amount of time – I cannot wait to work with you as a Hubble Fellow at Columbia. I also thank my many collaborators in the SMAUG, FIRE and CANDELS teams – working on joint projects together has made astronomy more fun, memorable and thought-provoking for me. I am also grateful to the UCSC Osterbrock Leadership Program through which I got to invest a small fraction of my time on various leadership-building activities and gain valuable advice from many mentors including Sandy, Andy Skemer, Brant Robertson, Bob Williams, Steve Mandel, Genevieve Graves and Steve Finkelstein.

I am grateful to the Princeton Post-Baccalaureate Program in Astrophysics, and especially Jenny Greene, for taking a chance on me and making it possible for me to pursue a PhD in Astrophysics. Thank you to Saurabh Jha for letting me start my first-ever astronomy research project during the same semester that I would graduate from Rutgers. Thank you to John Mulchaey for teaching me X-ray astronomy and for all the encouragement. Thank you to Roderich Tumulka for showing me that there's more to math and academia than just taking classes by advising me on an independent study of and research project on the foundations of quantum mechanics, which led to

my first publication. Thank you to the astronomy faculty at Rutgers for being sources of inspiration early on. Thank you to Carl Sagan whose writings I discovered one winter break in college – this serendipitous event inspired me to found the Rutgers Astronomical Society and ultimately changed my life.

Finally, thanks to all the friends that I have made throughout the years while working on this thesis. Special thanks to my amazing housemates in the two Santa Cruz “Astro Grad Houses” that I lived in: Asher, Kat, Brittany and Enia as well as Tiffany, Eric, Mickey and César. Thanks also to Asher and Felipe for the many inspiring jam sessions, Mickey for getting me into intramural basketball, Grecco for tennis and co-advising undergraduates with me, and all of the awesome graduate students at UCSC and the Flatiron Institute that I had the pleasure of getting to know. Lastly, thanks to my parents, sister and Emma for their support, especially throughout this last year while I was wrapping up my thesis during the pandemic.

Published Material: This dissertation includes reprints of the following previously published material with permission from the co-authors. Chapter 2 was published in *The Astrophysical Journal* as Pandya et al. (2020) and the American Astronomical Society holds the copyright; here I reproduce the accepted manuscript version. I was responsible for all of the analysis, text and figures with the exception of Figure 2.13 which was created by my co-author John Forbes. I benefited from numerous discussions with my co-authors, and the two models compared in that paper were developed originally by my advisor Rachel Somerville and the FIRE Collaboration. Chapter 3 is a preprint available as Pandya et al. (2021) and has received a favorable referee report (it

will be published in the Monthly Notices of the Royal Astronomical Society soon after this thesis with the copyright held by Oxford Journals). Again, I was responsible for all of the analysis, text and figures, and my co-authors provided the simulation data and significant input. Chapter 4 is written up as a draft but has not yet been submitted to a journal (I expect to do so after this thesis is published). Again, I was responsible for all of the work in this chapter. For the purpose of this dissertation, I have combined and streamlined the introductions from the three individual papers into a single non-redundant introduction.

Funding: I was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1339067. I was also supported by the Flatiron Institute, a division of the Simons Foundation, in the form of a Summer Research Associate position, a Pre-Doctoral Fellowship, and access to their supercomputing facilities which I utilized to do the work presented in this thesis. I thank the Scientific Computing Core at the Flatiron Institute for their help with my supercomputing questions.

Chapter 1

Introduction

In the Λ CDM paradigm of galaxy formation, the growth of dark matter halos is paralleled by the accretion of gas from the intergalactic medium (IGM; e.g., White & Rees, 1978; Blumenthal et al., 1984). The accreted gas is thought to reside within the circumgalactic medium (CGM), which acts as a buffer between the interstellar medium (ISM) and the IGM. Radiative cooling of this CGM gas leads to the build-up of the ISM and eventually star formation. The resulting feedback from stars and supernovae (SNe) is capable of heating and ejecting gas from the ISM back into the CGM or IGM, and the energy and momentum carried by these stellar-driven winds can also suppress future gas cooling and accretion (and hence star formation). Gas that has been previously ejected from the ISM can re-accrete, which together with the other gas flow processes gives rise to the “baryon cycle” of galaxies (e.g., Oppenheimer et al., 2010; Christensen et al., 2016; Anglés-Alcázar et al., 2017a). These and other physical processes ultimately shape the evolutionary histories of individual galaxies, with the statistical properties of galaxy

populations (e.g., the stellar mass function and galaxy scaling relations) emerging as a result. This is the modern high-level picture of galaxy formation gleaned from both observations and interpretive models, but many uncertainties remain in our detailed understanding of the relevant physics (see the recent reviews by Somerville & Davé, 2015; Naab & Ostriker, 2017).

Models of galaxy formation span a continuum in terms of volume and resolution. To thoroughly understand galaxy formation in a cosmological context requires modeling populations of galaxies, which in turn requires modeling large volumes (several 100^3 Mpc^3). Such large-volume population studies are important to: (1) ensure a robust sampling of the scatter in halo growth histories at a fixed mass, (2) explore the range of physical processes at play across different large-scale environments, (3) enable comparisons to observations from large-volume surveys, and (4) allow galaxies to ultimately be used as reliable cosmological probes. However, owing to resolution limits, all currently existing large-volume models contain a “phenomenological” component, which is to say that: (1) physics occurring below the resolution limit is parameterized, often in an ad hoc way, and (2) the free parameters of the model are adjusted to match a limited set of observations. This is generally true for modern large-volume hydrodynamical simulations, which solve the equations of gravity and fluid dynamics along with “subgrid recipes” (e.g., Vogelsberger et al., 2014; Genel et al., 2014; Schaye et al., 2015; Pillepich et al., 2018; Davé et al., 2019, and references therein). It is also true for semi-analytic models (SAMs), which attempt to distill the key insights from more sophisticated simulations using a set of coupled ordinary differential equations

that track the flow of mass between different galactic components (e.g., White & Frenk, 1991; Kauffmann et al., 1993; Cole et al., 1994; Somerville & Primack, 1999). Both phenomenological approaches have their advantages, disadvantages and simplifying assumptions, but ultimately they are complementary and inform each other.

In this dissertation, I critically examine key assumptions made in SAMs of galaxy formation with an emphasis on the overall baryon cycle. While SAMs have had tremendous success in interpreting and indeed reproducing various observations at a range of redshifts, their approximations have not been rigorously tested. It is timely to do so now because of recent successes in developing high-resolution but small-volume cosmological simulations where more of the relevant small-scale physics (such as star formation and supernova feedback) can be implemented more self-consistently. Such simulations can serve as benchmarks for more approximate SAMs, with the latter having the potential to emulate the computationally expensive simulations and scale up their predictions to statistically large volumes. On the other hand, since any simulation is incomplete and subject to code and physics uncertainties, the flexibility and efficiency of SAMs can be leveraged to explore the role of alternative or additional physics beyond what is included in the simulations. While the model comparison may reveal many areas of disagreement and hence future work, I hone in on the lack of explicit modeling of SN-driven heating of the CGM and IGM in SAMs. I demonstrate how to more completely characterize the thermodynamic properties of SN-driven winds in simulations, and how to use that resulting rich dataset to pin down many otherwise uncertain free parameters of SAMs. Before continuing further, it is important that I provide some historical

background.

1.1 History of comparing SAMs and hydrodynamical simulations

There is a long history of comparing the predictions of SAMs to hydrodynamical simulations. Benson et al. (2001) first demonstrated how the parameters controlling halo gas cooling and merger rates in a simplified SAM could be adjusted to better match predictions from a cosmological hydrodynamical simulation. They focused on the cosmic number and mass densities of hot halo gas and dense ISM gas. Overall, their study showed remarkable consistency between the two very different approaches for modeling halo gas cooling in cosmological volumes (in an average statistical sense). Subsequently, Yoshida et al. (2002) and Helly et al. (2003) each ran their own simplified SAM on halo merger trees extracted directly from hydrodynamical simulations and compared predictions for gas cooling and accretion on an individual halo-by-halo basis. Both of these studies demonstrated the striking correspondence, with minimal systematic offsets, between their SAM and hydrodynamical predictions. In the years since, there have been a number of studies that compared the predictions of SAMs and hydrodynamical simulations (using both statistical and individual halo-by-halo approaches). Owing to the ever-increasing sophistication of the simulations, the comparisons have expanded to include a wider range of physical processes beyond just halo gas cooling: evolution of dark matter subhalos (Jiang & van den Bosch, 2016), UV background heating due

to spatiotemporally inhomogeneous reionization (Mutch et al., 2016), cold/rapid versus hot/slow mode accretion (Cattaneo et al., 2007; Lu et al., 2011a; Hirschmann et al., 2012), relating halo and galaxy angular momentum (Guo et al., 2016; Stevens et al., 2017; Mitchell et al., 2018), multi-phase ISM and dust modeling (e.g., Popping et al., 2019), feedback processes (Weinmann et al., 2012; Qin et al., 2018; Ayromlou et al., 2020) and baryonic effects on dark matter halo concentrations (e.g., Dutton et al., 2016).

Among the many SAM versus hydrodynamical simulation comparisons, the studies by Stringer et al. (2010) and Neistein et al. (2012) are particularly informative. Stringer et al. (2010) modified several aspects of an existing SAM to ask how closely it could reproduce the evolution of a single Milky Way (MW)-mass halo simulated at high resolution. They found remarkable potential in the ability of their SAM to match the predictions of the more sophisticated simulation as a function of time, including the evolution of shocked versus unshocked halo gas accretion, halo gas scale length, disk gas scale length, disk circular velocity, stellar mass, cold gas mass, hot gas mass, hot disk gas mass and outflow gas mass. They further showed that their fiducial, previously published SAM (used for observational comparisons) predicted a very different evolution for the same simulated MW-mass halo, primarily due to its assumptions of much lower star formation efficiency and much stronger supernova feedback. Neistein et al. (2012) went a step further and characterized the efficiencies of various processes in a large-volume hydrodynamical simulation using a novel particle phase tracking approach. They derived mass- and redshift-dependent functions that summarized accretion, cooling, star

formation and feedback in the simulation. They emphasized that these functions were significantly different than the assumptions built into traditional SAMs, but that the functions represented a common language for connecting SAMs and simulations.

It is clear from the many previous studies that SAMs show the potential to transparently summarize the complicated physics of and emergent predictions from more sophisticated cosmological hydrodynamical simulations. However, an outstanding question that still remains today is whether modifications made to SAMs to bring them into better agreement with simulations must also necessarily come at the expense of no longer matching observations (Cattaneo et al., 2007; Stringer et al., 2010). Primarily, this puzzle must be driven by the fact that SAMs include only a limited description of the full range of phenomena found in simulations. However, a secondary cause is the somewhat circular logic of comparing to reference simulations that are themselves phenomenologically calibrated and hence effectively SAM-like in nature (e.g., Crain et al., 2015; Pillepich et al., 2018). Even if such simulations agree with a plethora of observations, the choice of subgrid model and associated free parameters carry degeneracies that propagate as largely unknown systematic uncertainties on predictions for galaxy populations. These uncertainties make it difficult to firmly interpret observations, but this has motivated important recent work on improving the flexibility and computational efficiency of SAMs (e.g., Henriques et al., 2009; Lu et al., 2011b; Henriques et al., 2013; Lagos et al., 2018; Forbes et al., 2019).

Chapter 2, which is published as Pandya et al. (2020), presents a compelling new take on this long problem. Instead of comparing SAMs to other phenomenological,

effectively SAM-like, large-volume simulations, or to a single galaxy simulated at high-resolution, I compare to a suite of several high-resolution state-of-the-art simulated halos spanning a range of masses. The models I use are the “Santa Cruz” SAM (the most recent version: Somerville et al., 2021) and the Feedback In Realistic Environments (FIRE) Project¹ (the second generation FIRE-2 suite; Hopkins et al., 2014, 2018a). The FIRE-2 simulations represent a good comparison suite because their stellar feedback model deposits mass, energy, momentum and metals locally without any explicit “tuning” to match observations; the resulting large-scale effects are hence emergent phenomena (e.g., outflows, CGM heating and recycling; Muratov et al., 2015; Anglés-Alcázar et al., 2017a). I focus not just on comparing overall bulk masses of the CGM, ISM and stars but also the corresponding mass inflow and outflow rates that build up those different galactic components. The emphasis on mass flow rates is crucial because fundamentally the system of ODEs governing SAM predictions deals with mass fluxes, and it is possible to get to the same mass in different ways. I will show that the lack of explicit preventative stellar feedback in SAMs, which is otherwise naturally captured in the simulations, leads to dramatic differences between the two models. I will also present a simple model for heating by SN-driven winds that serves as motivation for a deeper analysis of the thermodynamic properties of outflows in the simulations.

¹<http://fire.northwestern.edu>

1.2 A more complete picture of stellar feedback in simulations

SN-driven winds play a fundamental role in modern models of galaxy formation by helping to regulate star formation. Without SN-driven winds, models would predict an overabundance of dwarf galaxies compared to observations (e.g., White & Frenk, 1991; Benson et al., 2003; Kereš et al., 2009), overestimate the average stellar masses formed within dwarf halos (e.g., Dekel & Silk, 1986; Springel & Hernquist, 2003), fail to match the redshift evolution of several observed scaling relations (e.g., Somerville et al., 2001), and produce overly compact galaxies. In addition to regulating star formation, galactic winds are thought to affect the thermodynamic state and metal content of the circumgalactic medium (CGM; e.g., see the review by Tumlinson et al., 2017) as well as chemically enrich the intergalactic medium (IGM; e.g., Oppenheimer & Davé, 2006). Winds may also fuel a significant fraction of late-time star formation in more massive halos by recycling back into the interstellar medium (ISM; e.g., Oppenheimer et al., 2010; Henriques et al., 2013; White et al., 2015; Anglés-Alcázar et al., 2017a). In lower mass halos, SN-driven winds may more easily escape and heat the CGM/IGM, causing preventative feedback effects by suppressing gas accretion in the first place (e.g., Davé et al., 2012; Lu et al., 2015; Pandya et al., 2020) and decreasing the metal and dust content of dwarfs (e.g., Davé et al., 2011; Feldmann, 2015).

Despite their central importance, a complete characterization of galactic winds in a cosmological context and their implications for galaxy evolution has remained elu-

sive. In the current landscape of models, genuinely emergent wind properties have been predicted by “resolved” ISM simulations but these only represent a relatively small sub-galactic region and generally assume $z = 0$ Milky Way-like global conditions (e.g., Walch et al., 2015; Martizzi et al., 2016; Li et al., 2017; Kim & Ostriker, 2018; Kim et al., 2020a). Extending SN-driven wind predictions to global galaxy scales has been challenging, but much progress has been made using idealized high-resolution simulations of dwarfs and more massive galaxies (e.g., Hopkins et al., 2012; Fielding et al., 2017b; Smith et al., 2018; Hu, 2019; Li & Tonnesen, 2020). On cosmological scales, all large-volume models are effectively phenomenological: they must either implement wind scalings “by hand” (e.g., White & Frenk, 1991; Davé et al., 2016) or rely on subgrid approaches that require tunable free parameters such as hydrodynamically decoupled wind particles or temporary shut off of cooling (e.g., Springel & Hernquist, 2003; Stinson et al., 2006). In between these approaches sit a relatively new generation of cosmological “zoom-in” simulations such as the Feedback In Realistic Environments project² (FIRE; Hopkins et al., 2014, 2018a), where in some cases SN remnants can be resolved. When SN remnants are unresolved, the subgrid approach is to deposit the additional momentum expected from the unresolved energy-conserving phase of SN remnants using even higher resolution simulations for calibration instead of observational tuning (Hopkins et al., 2018b). In addition, a variety of physical processes are accounted for in such zoom-in simulations that may not otherwise be captured in small-scale simulations (e.g., self-consistent clustering of star formation, cosmological gas accretion and

²<http://fire.northwestern.edu>

mergers with complex global geometry, etc.). It is timely to ask how the emergent wind properties from such zoom-in simulations compare to those of higher resolution subgalactic simulations (as studied by, e.g., Gurvich et al., 2020), and to derive new wind scalings that can be implemented into large-volume simulations and semi-analytic models (SAMs; as presented by, e.g., Muratov et al., 2015).

When analyzing galactic winds, it is common practice to focus on “mass loading factors” and “metal loading factors,” which respectively describe gas mass outflow rates and metal outflow rates conveniently normalized by reference star formation rates and supernova metal injection rates. It has long been appreciated that dwarf halos preferentially have higher mass and metal loading factors (e.g., Dekel & Silk, 1986; Mac Low & Ferrara, 1999; Efstathiou, 2000). The common interpretation of this is that dwarfs have shallower potential wells and hence SN ejecta can more easily escape. Simple arguments suggest that we should expect a power law relation between the mass loading factor and global halo circular velocity whose slope will be steeper if winds are “energy-conserving” and shallower if they are “momentum-conserving” (Murray et al., 2005). Much work has gone into testing this simple energy- and momentum-driven dichotomy using hydrodynamical simulations (e.g., Hopkins et al., 2012; Muratov et al., 2015; Christensen et al., 2016), and the language of this framework is commonly used to justify assumed wind scalings in SAMs and simulations with insufficient resolution to capture SN remnant evolution (e.g., Somerville et al., 2008a; Oppenheimer et al., 2010; Anglés-Alcázar et al., 2014). While characterizing winds in this way has provided useful insights, a more detailed analysis of the thermodynamic properties of multi-phase winds

(i.e., temperature and velocity distributions) provides additional clues about whether their driving energy source is kinetic or thermal, and enables more careful consistency checks between different simulations (and against observations).

In addition to characterizing the mass and metal loading of galactic winds, it is also crucial to explicitly measure their multi-phase energy and momentum loading factors: how much of the energy and momentum input by SNe also make it out of the ISM? The explicit calculation of energy and momentum loading factors can help test whether winds are energy-driven or momentum-driven in a simple way, and help to interpret any secondary heating or “pushing” effects on the CGM/IGM. In recent years, small-scale, high-resolution idealized simulations have made quantitative predictions for energy and momentum loadings, with a common finding that the cold phase carries most of the mass whereas the hot phase carries most of the energy (Kim & Ostriker, 2018; Fielding et al., 2018; Hu, 2019; Li & Bryan, 2020; Kim et al., 2020a). These idealized numerical experiments have also been able to correlate their loading factors against the granular conditions of the ISM in which SNe go off rather than just the global halo circular velocity (e.g., Creasey et al., 2013; Fielding et al., 2017a; Li & Bryan, 2020; Kim et al., 2020a). A similarly comprehensive analysis of multi-phase galactic winds in cosmological simulations would provide major insights on how “ejective feedback” (quantified by mass and metal loading) and “preventative feedback” (quantified by energy and momentum loading) may act in concert to regulate galaxy evolution.

Chapter 3 presents a novel analysis of galactic outflows in the FIRE simulations (Hopkins et al., 2018a). This chapter is taken from Pandya et al. (2021), which

has received a favorable journal referee review and will be published soon after this thesis. In addition to measuring the overall wind mass loading factor and its scaling with various galaxy/halo properties, I go considerably beyond previous work by also measuring momentum, energy and metal loading factors split by phase. Motivated by analysis procedures for small-scale idealized simulations, I present a multi-dimensional characterization of the outflow properties with both global halo properties and several “quasi-local” ISM-scale properties. The insights from this chapter serve as the basis for implementing and calibrating preventative stellar feedback in SAMs.

1.3 The need for a more flexible SAM with explicit preventative stellar feedback

While SAMs of galaxy formation have seen tremendous success in reproducing properties of galaxies both near and far (see reviews by Benson, 2010; Somerville & Davé, 2015; Wechsler & Tinker, 2018), it remains unclear how much of this agreement comes at the expense of physically viable assumptions. For example, the implementation of continuous “disk mode” star formation is inconsistent with the apparent bursty nature of SF in observed and simulated dwarfs, and the lack of preventative feedback modeling requires assuming very high mass loading of SN-driven winds. The potentially incomplete description of all relevant physics in SAMs manifests in long-standing problems such as dwarfs forming their stars too early (Henriques et al., 2013; White et al., 2015), the inability to simultaneously match the observed redshift evolution of both the

stellar and gas-phase mass-metallicity relations (Somerville et al., 2015), and orders of magnitude uncertainty in CGM masses, metallicities, flow rates and overall multi-phase partitioning (Pandya et al., 2020).

SAMs contain on the order of tens of free parameters. Many of these free parameters control crucial mass flow rates that are observationally unconstrained and may effectively never be directly observed for the full galaxy population. It is common practice to benchmark SAMs by comparing to observable properties, namely bulk stellar and ISM masses, but this alone is inadequate since it is possible to reach the same final mass in different ways. One way around this is to reproduce the redshift evolution of, e.g., the stellar mass function, but that still does not directly constrain the mass flow terms operating under the hood in SAMs. Another idea is to look to more sophisticated hydrodynamical simulations where mass flow rates can explicitly be measured and serve to constrain SAMs. Ideally, the simulations would be of sufficiently high resolution to minimize the impact of uncertain subgrid assumptions so that the SAM would mainly be used to interpret the simulation. On the other hand, since any cosmological simulation is subject to such physics- and code-related uncertainties, it is still of great interest to develop a flexible SAM whose parameters can be cross-matched to phenomenological parameters of the simulation. Such a SAM can then be used to both distill the complexity of the simulation and to assess the potential impact of alternative subgrid prescriptions for unresolved physics (e.g., IMF, CGM cooling, SMBH formation and feedback, etc.).

In Chapter 4, we present the prototype of a new SAM in which nearly all of the

free parameters and associated functional forms are constrained from the core FIRE-2 simulations (Hopkins et al., 2018a). We combine insights from Chapters 2 and 3 to implement an alternative CGM cooling model for dwarfs and allow SN-driven winds to heat the CGM and IGM, thereby suppressing accretion. While previous studies have argued for preventative feedback (e.g., Lu et al., 2015; Hirschmann et al., 2016; Lu et al., 2017; Wright et al., 2020), here I show for the first time that SN-driven winds have enough energy to provide the necessary heating and estimate the energy coupling efficiency. I show that the prototype SAM gets remarkably close in reproducing the overall evolution of the individual FIRE halos provided that we allow for this SN-driven heating of the CGM and IGM around dwarfs. However, uncertainties remain in modeling rapid recycling of “fountain flows” (weak winds) in the inner halo, and the chemical enrichment of the CGM from both the central galaxy and its satellites. Despite these and related “missing physics” issues, this work provides a promising path forward for the challenging goal of re-tooling SAMs to emulate simulations. We discuss how we plan to address remaining model uncertainties, implement additional physics modules, eventually publicly release a flexible open-source SAM code framework, and explore anticipated observational implications in the near future.

1.4 Outline of this work

Chapter 2 compares predictions between two state-of-the-art galaxy formation models and pinpoints ways in which the two models could be made to better agree with

each other. Chapter 3 analyzes the hydrodynamical simulations in more detail to achieve a more complete characterization of SN-driven winds. Chapter 4 combines insights from the previous two chapters and presents the prototype of a new SAM that can roughly reproduce the evolution of FIRE halos, provided preventative feedback is accounted for. Finally, in Chapter 5, we summarize our results and discuss remaining modeling uncertainties, anticipated observational implications, and how this dissertation sets the foundation for significant future work.

Chapter 2

Comparing galaxy formation in semi-analytic models and hydrodynamical simulations

2.1 Introduction

In this chapter, we compare two state-of-the-art galaxy formation models: the Santa Cruz SAM (Somerville et al., 2015, and references therein) and the FIRE-2 cosmological zoom-in simulations (Hopkins et al., 2018a). Both of these models have been shown to reproduce a large range of observations, but they are built very differently. The SAM includes an approximate treatment of many of the physics relevant to galaxy formation whereas the FIRE-2 simulations rely on a more direct approach of solving the equations of fluid dynamics along with subgrid models. The FIRE-2 simulations represent a good comparison suite because their stellar feedback model deposits mass,

energy, momentum and metals locally without any explicit “tuning” to match observations; the resulting large-scale effects are hence emergent phenomena (e.g., outflows, CGM heating and recycling; Muratov et al., 2015; Anglés-Alcázar et al., 2017a).

As with any simulation, caution is warranted regarding the absolute correctness and completeness of the FIRE-2 simulations (improvements can always be made to the numerical algorithms, the range of physical processes implemented, and the diversity of halo mass accretion histories and large-scale environments probed). However, for the purposes of improving physical prescriptions for SAMs, we can confidently use FIRE-2 as a baseline for comparison, identify systematic discrepancies, and develop plausible solutions to guide future work. With the FIRE-2 suite, we will study the time evolution of 13 individual halos across a broad range in mass: low-mass dwarfs ($M_{\text{vir}} \sim 10^{10} M_{\odot}$ at $z = 0$), intermediate-mass dwarfs ($\sim 10^{11} M_{\odot}$), and MW-mass galaxies ($\sim 10^{12} M_{\odot}$). We will also restrict the scope of our comparison to a few bulk quantities that characterize the overall baryon cycle of galaxies (the foundation of any SAM): stellar, ISM and CGM masses, and the corresponding mass inflow and outflow rates for the ISM and CGM. The inclusion of flow *rates* in addition to global bulk quantities is, to our knowledge, a novel feature of this work which has not been widely studied in the past (but see Hirschmann et al., 2012).

The work in this chapter was published as (Pandya et al., 2020) and advanced one of the key goals of the SMAUG Collaboration¹ which is to ask: is it possible to develop a model that faithfully captures the essential physics of galaxy formation in

¹Simulating Multiscale Astrophysics to Understand Galaxies

a more computationally efficient way than fully numerical large-volume simulations? Given that the physical processes involved in galaxy formation are not fully understood and also span a vast range in scale, it is not feasible to develop a single “ab initio” simulation that is capable of making credible predictions on the scale of galaxy populations. Instead, SMAUG aims to carefully design a suite of high-resolution numerical experiments whose results can be coarse grained to develop realistic subgrid prescriptions for cosmological simulations. As part of the first results from SMAUG², the resolved ISM simulations by Kim et al. (2020a) and resolved black hole accretion simulations by Angles-Alcazar et al. (2020) take the first step towards this goal. The complementary parameter space study of simulated star-forming regions by Motwani et al. (2020) is designed to provide the initial conditions for a future suite of resolved ISM simulations building on Kim et al. (2020a). In the present work, we take the first step towards re-tooling and calibrating SAMs using physically self-consistent simulations instead of observations so that SAMs may become more predictive rather than descriptive in nature. Our emphasis on the need to improve phenomenological modeling of stellar feedback and gas flows in the CGM underscores the work of Fielding et al. (2020b), who find that the properties of the multi-phase CGM depend strongly on the nature of feedback, cosmological accretion and simulation methodology.

This chapter is organized as follows. We describe the FIRE-2 simulations and the Santa Cruz SAM in Section 2.2, and our analysis methods in Section 2.3. We present the results of our comparison in Section 2.4, while Section 2.5 is devoted

²<https://www.simonsfoundation.org/flatiron/center-for-computational-astrophysics/galaxy-formation/smaug/papersplash1>

to interpreting the model discrepancies and presenting a preventative stellar feedback model for inclusion in future SAMs. A summary is provided in Section 2.6.

2.2 Model Descriptions

Here we describe the FIRE-2 cosmological hydrodynamical “zoom-in” simulations and the Santa Cruz SAM. Note that the FIRE-2 simulations assume $h = 0.70$, $\Omega_{m,0} = 0.27$, $\Omega_{\Lambda,0} = 0.73$ and $\Omega_b = 0.045$ (see section 2.8 of Hopkins et al., 2018a). The Santa Cruz SAM assumes the Planck Collaboration et al. (2016) cosmology, with $h = 0.678$, $\Omega_{m,0} = 0.308$, $\Omega_{\Lambda,0} = 0.692$ and $\Omega_b = 0.0486$. The main differences between the SAM and FIRE-2 that we focus on in this chapter are unlikely to be driven by the small differences in assumed cosmology.

2.2.1 FIRE-2 Simulations

We use the FIRE-2 suite of cosmological hydrodynamical “zoom-in” simulations described in Hopkins et al. (2018a). The simulations were run with the Gizmo³ code (Hopkins, 2015) using the Lagrangian meshless finite-mass method and fully-adaptive gravitational force softening lengths for gas. Briefly, a large DM-only box was evolved to $z = 0$, and relatively isolated halos were chosen to be re-simulated at much higher resolution with baryons included. The initial zoom region is defined to be $\sim 5R_{\text{vir}}$ around the halo at $z = 0$, but in practice only the zoom region within $\sim 2R_{\text{vir}}$ is guaranteed to avoid contamination from low resolution DM particles.

³<http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>

The FIRE-2 simulations account for gas heating and cooling between temperatures of 10K and 10^{10} K, including free-free, photoionization/recombination, Compton, photoelectric, metal-line, molecular, fine-structure, dust collisional and cosmic ray processes (the corresponding cooling tables are given in Appendix B of Hopkins et al., 2018a). A spatially uniform but redshift dependent UV background is imposed based on Faucher-Giguère et al. (2009). Star formation occurs stochastically in self-gravitating, molecular, self-shielding gas that has hydrogen number density $n_H \geq 1000 \text{ cm}^{-3}$. Owing to the high spatial and mass resolution (see below), stellar feedback is modeled via local deposition of mass, momentum, energy and metal mass from star particles to neighboring gas particles. The feedback accounts for both Type Ia and Type II supernovae, stellar winds, momentum from radiation pressure, photo-ionization and photo-electric heating. In this way, the generation, propagation and recycling of large-scale galactic winds are emergent phenomena rather than being put in “by hand” via delayed cooling, thermal bombs or decoupled winds (e.g., Muratov et al., 2015; Anglés-Alcázar et al., 2017a).

Of the 27 high resolution FIRE-2 halos listed in Table 1 of Hopkins et al. (2018a), we use the 13 halos for which particle data were output for the full set of 600 snapshots: m10q, m10v, m10y, m10z, m11a, m11b, m11q, m11c, m11v, m11f, m12i, m12f and m12m (these specific halos were originally presented in Wetzel et al., 2016; Garrison-Kimmel et al., 2017; Chan et al., 2018; Hopkins et al., 2018a). Our sample of FIRE-2 halos is identical to those of Hafen et al. (2019, see their Table 1), with the addition of m11f but excluding their metal diffusion runs. The halos are grouped into

three mass bins based on their $z = 0$ virial mass: the four m10 halos have $M_{\text{vir}} \sim 10^{10} M_{\odot}$ (low-mass dwarf bin), the six m11 halos have $M_{\text{vir}} \sim 10^{11} M_{\odot}$ (intermediate-mass dwarf bin), and the three m12 halos have $M_{\text{vir}} \sim 10^{12} M_{\odot}$ (MW-mass halo bin). With this sample, we will be able to study systematic trends with halo mass for discrepancies between the SAM and FIRE-2. The mass and spatial (gravitational force softening) resolution vary with halo mass, and are systematically higher for the dwarfs. The star/gas particle masses are $250 M_{\odot}$ for the m10 halos, $880 M_{\odot}$ for m11q, $2100 M_{\odot}$ for m11a, m11b and m11c, $7100 M_{\odot}$ for m11v, m12i, m12f and m12m, and $12000 M_{\odot}$ for m11f. The DM particle masses are roughly $\sim 5\times$ higher. The minimum adaptive gravitational force softening lengths for the gas are on the order of ~ 1 pc (see Hopkins et al., 2018a, for more details). In addition, the typical snapshot spacing is ~ 20 Myr, which allows us to accurately track variability in halo mass accretion and star formation for comparison to the SAM.

2.2.2 Santa Cruz Semi-Analytic Model

The Santa Cruz SAM was first presented in Somerville & Primack (1999), with significant updates described in Somerville et al. (2008a), Somerville et al. (2012), Porter et al. (2014), Popping et al. (2014) and Somerville et al. (2015). Here, we use the latest Somerville et al. (2015) version, which includes recipes for multi-phase partitioning of ISM gas. We adopt the same calibration of free parameters for this version as used in Popping et al. (2019); we will report the adopted parameter values for each of the relevant equations that we review below. We will not review the details of satellite-

specific processes since our comparison to FIRE-2 only involves central halos (section 2.8 and section 2.1, respectively, of Somerville & Primack, 1999; Somerville et al., 2008a, describe how subhalos are modeled in the SAM). In addition, since active galactic nucleus (AGN) feedback is not implemented in the FIRE-2 simulations employed here, we have disabled it in the SAM for a more consistent comparison and will not review the corresponding equations here.⁴

We emphasize that we have not made any other changes to the Santa Cruz SAM used in previously published works. The initial mass function assumed in the SAM (Chabrier, 2003) and FIRE-2 (Kroupa, 2001) are similar enough that they are unlikely to drive any significant model differences. For the cooling function, the SAM assumes Sutherland & Dopita (1993) whereas FIRE-2 has an implicit cooling algorithm based on more recent calculations for a wide range of physical processes. It is possible that if we implement the FIRE-2 cooling function approximations listed in Appendix B of Hopkins et al. (2018a), the SAM predictions for CGM cooling rates could change dramatically. However, we do not think the different cooling functions can explain the more fundamental qualitative differences between the two models (e.g., regarding halo gas accretion and recycling). On the other hand, the metallicity modeling and calibration are quite different between the SAM and FIRE-2, with the former using the instantaneous recycling approximation and only considering Type II supernovae with an assumed chemical yield $y = 1.6$ (section 2.8 of Somerville et al., 2008a). In contrast,

⁴Note that AGN feedback can have appreciable effects for MW halos but not dwarfs in the SAM (Somerville et al., 2008a). Nevertheless, for a consistent comparison with the FIRE-2 simulations we must disable it in the SAM.

FIRE-2 self-consistently tracks chemical yields of different species during various stages of star particle evolution (Appendix A of Hopkins et al., 2018a). Although we do not focus on comparing metallicities between the two models in this chapter, a more sophisticated treatment of metals in the SAM could affect cooling-related processes and have observable consequences (we defer an investigation of this to future work).

2.2.2.1 Halo gas accretion

For any given halo, the SAM begins by computing the DM accretion rate via finite differencing of the virial mass time series provided by the halo merger tree. Before the Universe is reionized (reionization is assumed to occur instantaneously at a specified redshift), it is assumed that gas accretion into the halo perfectly tracks DM accretion with the universal baryon fraction, i.e., $\dot{M}_{\text{gas}} = f_{\text{b}}\dot{M}_{\text{vir}}$ where $f_{\text{b}} = 0.158$ according to Planck Collaboration et al. (2016). After reionization, the pristine gas accretion rate is suppressed due to photoheating from the UV background:

$$\dot{M}_{\text{CGM,in,pristine}} = f_{\text{coll}}f_{\text{b}}\dot{M}_{\text{vir}}. \quad (2.1)$$

The factor f_{coll} gives the fraction of infalling baryonic mass that is able to collapse into the halo despite heating by the UV background. It depends on halo mass and redshift, and is taken from Okamoto et al. (2008) who characterized the suppressive effects of the Haardt & Madau (2001) UV background in their hydrodynamical simulations. In practice, the formula for f_{coll} involves computing a “characteristic filtering mass” at which the gas accretion rate drops to half of the universal f_{b} ; above this characteristic halo mass, the accretion rate approaches f_{b} , and below it the accretion rate drops

steeply such that UV background heating is more effective in lower mass halos. The filtering halo mass is computed according to Appendix B of Kravtsov et al. (2004); it is $M_{\text{filt}} \lesssim 10^8 M_{\odot}$ before reionization is complete, and rises to $M_{\text{filt}} \approx 10^{10} M_{\odot}$ by $z = 0$. We assume the IGM is fully reionized by $z \sim 8$, consistent with Planck Collaboration et al. (2016). All of our FIRE-2 halos have virial masses above the characteristic filtering mass at all times $z \lesssim 10$, except for m10v which becomes larger starting only at $z \sim 2$ (see also Figure 11 of Fitts et al., 2017). We have experimented with changing the filtering mass normalization to mimic using different UV background models, and find that our results are insensitive for reasonable changes.

On top of the pristine IGM gas accretion, the SAM adds the “re-accretion” of gas that was previously ejected from the halo due to stellar feedback:

$$\dot{M}_{\text{CGM,in,recycled}} = \chi_{\text{re-infall}} \left(\frac{M_{\text{ejected}}}{t_{\text{dyn}}} \right). \quad (2.2)$$

M_{ejected} is the total mass of the ejected gas reservoir (its growth rate is set by the product of Equation 2.6 and Equation 2.7 described below) and $\chi_{\text{re-infall}}$ is a free parameter that sets what fraction of the ejected gas reservoir can cool back into the halo at each time step. We assume $\chi_{\text{re-infall}} = 0.1$ as in previous Santa Cruz SAM papers; this implies that the ejected gas will re-accrete back into the halo on ten dynamical times $t_{\text{dyn}} \equiv \frac{R_{\text{vir}}}{V_{\text{vir}}}$, where $V_{\text{vir}} = \sqrt{\frac{GM_{\text{vir}}}{R_{\text{vir}}}}$ is the circular velocity of the halo at the virial radius (note that $t_{\text{dyn}} \approx 0.1 t_{\text{Hubble}}$, so the gas will effectively re-accrete over a Hubble time).

There are two additional sources of CGM gas from within the halo itself. The first is outflows from the ISM that get deposited into the CGM; we defer this to

the discussion of the relevant stellar feedback equations below. The second source is transfer from subhalos: the SAM assumes that once a halo becomes a subhalo, the CGM of the subhalo is instantaneously transferred to the CGM of the host halo. Although physical processes associated with satellite galaxies (i.e., subhalos) can indirectly affect the evolution of the central galaxy, we do not expect these processes to be the dominant ones in the simulations we are considering.

2.2.2.2 CGM gas cooling

Gas that has accreted into the halo as described above builds up the CGM mass. The cooling rate of this CGM gas into the ISM is computed according to White & Frenk (1991), which is also the basis for most, if not all, other SAMs. First, the CGM is assumed to uniformly be at the virial temperature of the halo at each time step. Then, the “radiative cooling time” is computed, which is the characteristic timescale for the gas to cool by radiating away its current thermal energy:

$$t_{\text{cool}}(r) = \frac{(3/2)\mu m_p k T_{\text{vir}}}{\rho_g(r) \Lambda(T_{\text{vir}}, Z_h)}. \quad (2.3)$$

μm_p is the mean molecular weight of the halo gas and $\Lambda(T_{\text{vir}}, Z_h)$ is the Sutherland & Dopita (1993) cooling function, which takes into account the metallicity of the halo gas Z_h . As is common practice, the gas mass density radial profile is assumed to be a singular isothermal sphere: $\rho_g(r) = \frac{M_{\text{CGM}}}{4\pi R_{\text{vir}} r^2}$. Plugging this into the equation for t_{cool} , one can solve for R_{cool} , the radius within which all of the gas can radiatively cool within t_{cool} (heating is neglected). Then, integrating to compute the total cooled mass within

R_{cool} and differentiating with respect to time gives the ISM mass accretion rate:

$$\dot{M}_{\text{ISM,in}} = \frac{1}{2} M_{\text{CGM}} \frac{R_{\text{cool}}}{R_{\text{vir}}} \frac{1}{t_{\text{cool}}}. \quad (2.4)$$

Note that although different choices have been adopted in the literature, it is common practice to assume that the cooling time is equal to the halo dynamical time at R_{vir} , i.e., $t_{\text{cool}} = t_{\text{dyn}}$. It is possible to have $R_{\text{cool}} > R_{\text{vir}}$ (this generally occurs for low mass halos), and these instances are assumed to represent “cold/fast/filamentary” mode accretion. Since $R_{\text{cool}} > R_{\text{vir}}$ implies that the cooling time is shorter than the dynamical time, the SAM ignores the radiative cooling prediction during these timesteps and instead sets the ISM accretion rate equal to the halo gas accretion rate (see also, e.g., Croton et al., 2006). Otherwise, the interpretation is that gas has been gravitationally shock-heated to the virial temperature upon first accreting into the halo, and is now radiatively cooling via the assumed “hot/slow/spherical” mode. As mentioned in Somerville et al. (2008a), reasonable variations within the framework of this particular gas cooling model (e.g., changing the definition of t_{cool} or assuming a different form for $\rho_g(r)$) can lead to variations in the ISM accretion rate by a factor of at most $\sim 2 - 3$.

2.2.2.3 Star formation and stellar feedback

Gas that has accreted into the ISM is partitioned into HI, H₂, HII and metals (details given in Popping et al., 2014; Somerville et al., 2015). The SAM keeps track of the mass surface density for these different gas phases in radial disk annuli (see Somerville et al., 2008b, for details about the SAM disk model). The default recipe for predicting the star formation rate (SFR) surface density is based on the molecular

hydrogen gas phase alone, accounting for a higher conversion efficiency above a critical H_2 surface density (Bigiel et al., 2008, 2011; Narayanan et al., 2012):

$$\Sigma_{\text{SFR}} = A_{\text{SF}} \left(\frac{\Sigma_{\text{H}_2}}{10 M_{\odot} \text{pc}^{-2}} \right) \left(1 + \frac{\Sigma_{\text{H}_2}}{\Sigma_{\text{H}_2, \text{crit}}} \right)^{N_{\text{SF}}}. \quad (2.5)$$

A_{SF} , N_{SF} and $\Sigma_{\text{H}_2, \text{crit}}$ are free parameters of this two-part scaling relation. We assume $A_{\text{SF}} = 5.98 \times 10^{-3} M_{\odot} \text{yr}^{-1} \text{kpc}^{-2}$, $N_{\text{SF}} = 1.0$ and $\Sigma_{\text{H}_2, \text{crit}} = 70 M_{\odot} \text{pc}^{-2}$ (following Popping et al., 2014, 2019). There are various ways to estimate the molecular hydrogen gas density Σ_{H_2} . Here we use the metallicity-dependent partitioning approach of Gnedin & Kravtsov (2011) that is the default in the Santa Cruz SAM.

On top of the continuous “disk mode” star formation, the SAM also superimposes “starbursts” due to galaxy mergers. The SFR spikes are modeled using a Laplace distribution (i.e., double exponential distribution) whose two parameters, the total starburst mass M_{burst} and the associated timescale τ_{burst} , are a function of progenitor properties and calibrated to binary galaxy merger simulations (Somerville et al., 2008a; Porter et al., 2014, and references therein). Note that while starbursts will contribute some variability to the overall star formation history (SFH), the disk star formation can exhibit its own stochasticity due to changes in the H_2 gas fraction (driven by changes in gas metallicity and galaxy size) and changes in the overall gas fraction (driven by stellar feedback and CGM gas cooling).

All stellar feedback in the SAM (aside from heating by the UV background) is ejective. At every timestep, the mass outflow rate from the ISM due to stellar feedback

is computed as:

$$\dot{M}_{\text{ISM,out}} = \epsilon_{\text{SN}} \left(\frac{V_{\text{max}}}{V_0} \right)^{\alpha_{\text{SN}}} \dot{M}_{\text{SFR}}. \quad (2.6)$$

Here, ϵ_{SN} and α_{SN} are free parameters, V_0 is an arbitrary normalization constant and V_{max} is the maximum circular velocity of the halo taken from the merger tree. We assume $\epsilon_{\text{SN}} = 1.5$ and $\alpha_{\text{SN}} = -2.6$ following Popping et al. (2019). The total mass blown out of the ISM is either transferred into the CGM or driven out of the halo completely (i.e., deposited into the ejected reservoir). The fraction of outflow mass that gets ejected from the halo is computed via

$$f_{\text{eject}} = [1.0 + (V_{\text{vir}}/V_{\text{eject}})^{\alpha_{\text{eject}}}]^{-1}, \quad (2.7)$$

where α_{eject} and V_{eject} are free parameters, with the latter representing a “threshold” halo virial velocity below which most ISM wind mass will leave the halo. We assume $\alpha_{\text{eject}} = 6$ and $V_{\text{eject}} = 110 \text{ km s}^{-1}$ following Somerville et al. (2008a, and more recent Santa Cruz SAM studies). Hence, $f_{\text{eject}} \times \dot{M}_{\text{ISM,out}}$ gives the mass addition rate for the ejected reservoir and the remainder $(1-f_{\text{eject}}) \times \dot{M}_{\text{ISM,out}}$ is deposited into the CGM. The ejected gas can re-accrete into the halo on a Hubble time and become re-eligible for cooling as described earlier.

2.3 Analysis

Here we describe how we analyze the hydrodynamical simulations and generate semi-analytic predictions for comparison.

2.3.1 Generating merger trees and SAM predictions

We run the Rockstar halo finder (Behroozi et al., 2013b) to generate halo catalogs at each snapshot for the full hydrodynamical FIRE-2 simulations which include both baryonic and DM particles. But since Rockstar will only use DM particles to define virial overdensities and hence halo boundaries, we enable its option to up-weight the DM density field. We adopt the Bryan & Norman (1998) definition of halo virial mass and radius. We only output properties of halos that have at least 100 DM particles associated with them (i.e., within their virial radius); this is the default threshold below which Rockstar discards halos as noise. Next, we run the companion consistent-trees code (Behroozi et al., 2013c) to generate gravitationally-consistent merger trees. This code corrects inconsistencies in the default Rockstar-based merger trees by: (1) removing spurious detections of halos, (2) inserting “phantom” halos at snapshots where a descendant halo is not identified but should obviously exist due to re-appearance of the halo in a subsequent snapshot, and (3) slightly modifying the positions and velocities of halo centers by comparing to the expected evolution between snapshots based on simple gravitational force calculations. In the end, our halo virial masses and radii agree with those reported in Table 1 of Hopkins et al. (2018a) to within 0.1 dex.

With the halo merger trees in hand, we run the Somerville et al. (2015) version of the Santa Cruz SAM with the same observational calibration as used in Popping et al. (2019, with AGN feedback disabled; see our subsection 2.2.2 above for details). Since the SAM has its own model for generating subhalos and predicting their evolution (section

2.1 of Somerville et al., 2008a), we have discarded all subhalos from the merger trees. This is appropriate for our study since we are only focusing on the evolution of the central halo in each of the FIRE-2 zoom simulations (along the most massive progenitor branch) and comparing subhalo modeling is deferred to future work. Note that we are running the SAM on merger trees extracted from the full hydrodynamical simulations, whereas it would be more faithful to use merger trees extracted from corresponding DM-only simulations. However, such DM-only simulations only exist for a subset of the FIRE-2 suite and hence we use the full hydrodynamical suite to increase our sample size (13 halos). In subsection 2.5.3, we show that none of our conclusions change when we use only the limited DM-only simulation suite. Nevertheless, we emphasize that the only input for the SAM is the dark matter halo merger trees: the SAM is not provided any information about the baryonic properties of the halos.

2.3.2 Computing bulk and flow quantities in the simulations

Our merger trees tell us the center position and radius of the central halo in every snapshot, as well as many other halo properties. With this information, we can use the simulation particle data to compute the baryonic properties of the central halo along the most massive progenitor branch. In Figure 2.1, we illustrate how we compute bulk masses and differential mass flow quantities in different “zones” for the hydrodynamical data. The definitions of these zones are well-matched to the SAM for comparison. We define the stellar mass as the sum of the masses of all star particles within $0.1R_{\text{vir}}$. We also define the ISM mass as the sum of all gas particle masses within $0.1R_{\text{vir}}$. The

CGM mass is defined as the sum of all gas particle masses between $0.1R_{\text{vir}}$ and $1.0R_{\text{vir}}$, irrespective of temperature, density, etc. We already have the dark matter halo mass from the merger trees, which is based on the sum of all DM particle masses within $1.0R_{\text{vir}}$. These constitute our main integrated mass measurements. We also compute instantaneous global galaxy SFRs by summing up the predicted instantaneous SFRs of all individual gas particles within $0.1R_{\text{vir}}$. We have also computed time-averaged SFRs based on adding up stars with ages younger than 20 Myr, 100 Myr, and 1 Gyr, and find good agreement with the instantaneous gas-based measurements after boxcar smoothing. By default, we use the instantaneous gas-based measurements since these are closer in definition to what the Santa Cruz SAM predicts.

We adopt the approach of Muratov et al. (2015) to measure instantaneous mass flow rates within radial shells. Specifically, for all particles within a given radial shell, we compute their radial velocities including the contribution from the Hubble flow (this is generally minor but it can have a differentially larger effect in halo outskirts). We define all particles with negative halo-centric radial velocities as inflowing, and similarly all particles with positive radial velocities as outflowing. Then, the mass inflow rate for a given radial shell is the weighted sum of the individual particle mass fluxes using only the particles with negative radial velocities:

$$\dot{M} = \sum_i \frac{m_i |v_{r,i}|}{dL}. \quad (2.8)$$

Here, m_i is the mass of particle i in the shell, $|v_{r,i}|$ is the absolute value of its radial velocity and dL is the shell width. An analogous calculation is done separately for the

mass outflow rate using only particles in the shell with positive radial velocities. In this way, particles with slower velocities contribute less mass flux than those with higher velocities (for a given particle mass), and the dependence of the mass flux measurement on the shell width is accounted for.

We make mass inflow and outflow rate measurements in two spherical shells at every snapshot. We define a “virial shell” that extends from $1.0 - 1.1R_{\text{vir}}$ and an “ISM shell” that extends from $0.1 - 0.2R_{\text{vir}}$. The widths of both shells are thus $0.1R_{\text{vir}}$. We have carried out extensive convergence tests for the location and width of each shell. In short, the definition of the virial shell is robust to reasonable changes in the centering and width, especially since we take the halo virial radius as a given from the merger tree. On the other hand, the definition of the ISM shell is more arbitrary since there is no obvious ISM “edge” in either the simulation or the SAM. The ISM shell width represents a good compromise between mitigating Poisson noise, systematically missing the fastest moving particles, and accurately capturing the bulk flow of mass as a function of radius across snapshots. The ISM shell is located at a considerable distance from the ISM which means that there can be contamination from ambient inner CGM material or fountain flows. However, if the shell is placed too close to the ISM, then the flow measurements can also be contaminated by the dense ISM. Without imposing more sophisticated criteria to select truly escaping or accreting ISM particles and accounting for the complicated geometrical evolution of galaxies, there thus needs to be a compromise. Overall, we find that our choice of shell definitions are sensible for comparison to the SAM and for measuring flow rates out to $z \sim 10$ (and they are also

standard in the literature; Faucher-Giguère et al., 2011; Muratov et al., 2015).

Another way to derive mass flow rates is via particle tracking, which also has the advantage of providing information about recycling distances and timescales. Anglés-Alcázar et al. (2017a), Hafen et al. (2019) and Hafen et al. (2020) have already performed this particle tracking analysis for both the FIRE-1 and FIRE-2 simulations, and we will discuss their results in the context of our work. Note that throughout this chapter we will use the “pure” inflow and outflow rates separately instead of the net inflow rate (i.e., inflow minus outflow). We do not attempt to excise satellites whose own orbits and outflows can contaminate our flow measurements for the central halo. Leaving satellites in may also bias our computed CGM masses a bit high, although it does make for a more consistent comparison to the SAM (which transfers the CGM of subhalos to that of the parent halo; subsection 2.2.2).

2.4 Results

Here we present results from our comparison between the SAM and FIRE-2 predictions. We will first present the bulk mass quantities and then the mass flow quantities to better pinpoint any discrepancies.⁵

⁵Unless otherwise noted, we boxcar smooth all time series by ± 1 Gyr to facilitate visual comparison. While this smoothing does wash out features on much smaller timescales, our conclusions would broadly remain the same had we not smoothed the data.

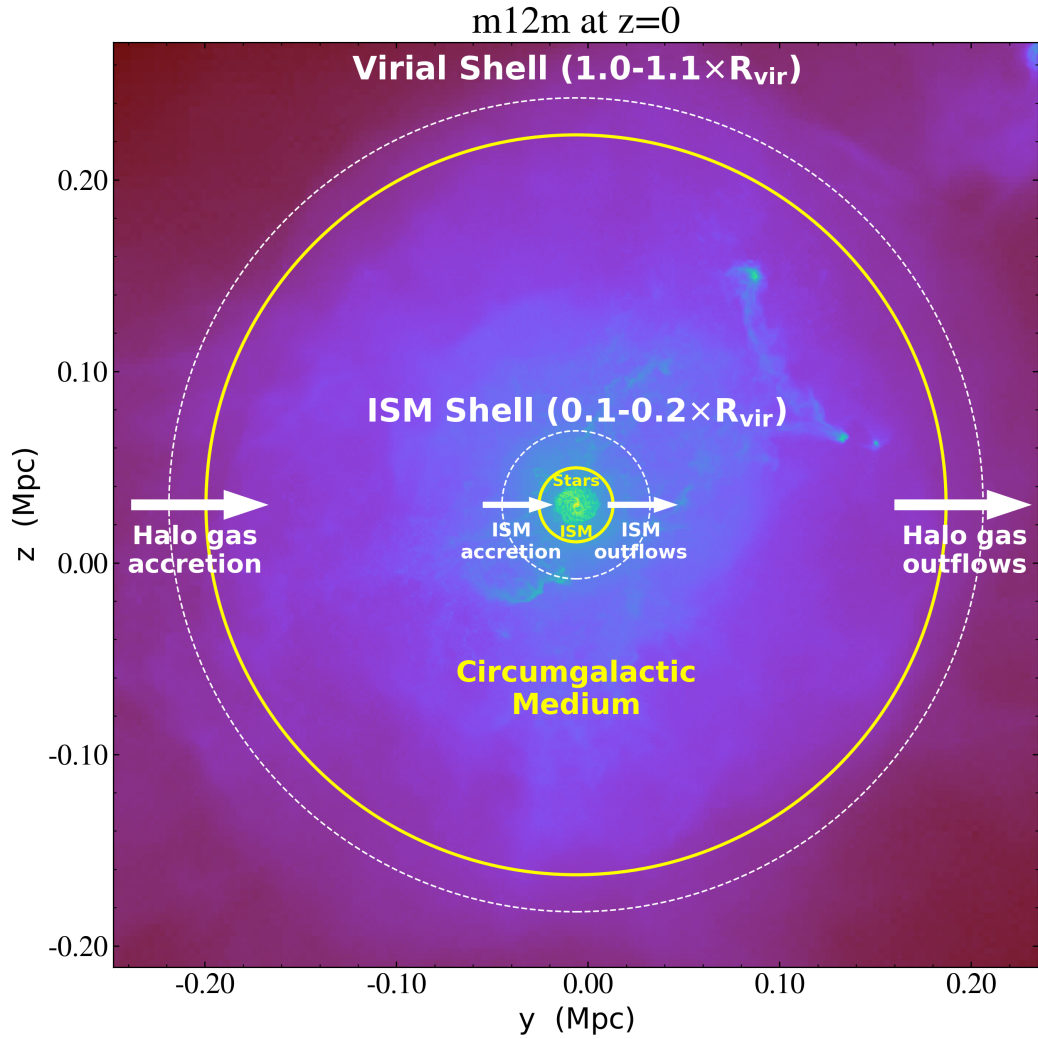


Figure 2.1: An illustration of our zone definitions for analyzing bulk and flow quantities in the FIRE-2 simulations. The background image shows the projected gas density distribution of the MW-mass halo m12m at $z = 0$ (purple is low density, with green and yellow representing progressively higher densities, respectively). The solid yellow circles represent the virial radius (outer circle) and the “edge” of the ISM (inner circle). The dashed white circles demarcate the outer limits of the virial and ISM shells ($1.0 - 1.1 \times R_{\text{vir}}$ and $0.1 - 0.2 \times R_{\text{vir}}$, respectively) through which mass inflow and outflow rates are computed. Stellar and ISM masses are computed using all star and gas particles within $0.1 \times R_{\text{vir}}$, respectively, whereas CGM masses are computed using all gas particles between $0.1 - 1.0 \times R_{\text{vir}}$. These definitions are well-matched for comparison to the Santa Cruz SAM.

2.4.1 Stellar, ISM and CGM mass scalings at $z = 0$

We begin by showing mass-dependent scaling relations at $z = 0$ for the SAM, FIRE-2 and observations in Figure 2.2. We focus on stellar-to-halo, ISM-to-stellar and CGM-to-halo mass ratios (as a function of the denominator mass; no boxcar smoothing). We include comparisons to median stellar-to-halo mass relations derived from halo abundance matching (from Rodríguez-Puebla et al., 2017; Behroozi et al., 2019) and ISM-to-stellar mass relations from observations (from Boselli et al., 2014; Peeples et al., 2014; Calette et al., 2018). We do not include observationally inferred CGM-to-halo mass ratios because there are large systematic uncertainties associated with measuring the total CGM masses of galaxies. Versions of the stellar-to-halo and (atomic plus molecular) ISM-to-stellar mass ratio relations are used to calibrate the SAM.⁶ We emphasize that these observational comparisons are purely illustrative: we have not made an effort to properly generate mock observables and there are a few caveats. First, the stellar-to-halo mass relations based on subhalo abundance matching are only valid at $M_{\text{vir}} \gtrsim 10^{10.5} M_{\odot}$ due to the resolution of the DM simulations used so we cannot comment on the low-mass dwarfs (but see Wheeler et al., 2019). In addition, we do not make any cuts on ISM gas phase for the SAM and FIRE-2 predictions even though the observationally inferred ISM-to-stellar mass ratios plotted in Figure 2.2 account for

⁶As mentioned in subsection 2.2.2, we have disabled AGN feedback in the SAM for the sake of a fair comparison with FIRE-2, even though the SAM relies on AGN feedback to agree with observations for MW and higher mass halos (there are no appreciable effects for dwarfs in the SAM). We find that enabling AGN feedback in the SAM decreases the stellar and ISM masses of MW halos by a couple tenths of a dex, and increases their CGM masses by more than a dex (owing to quasar winds ejecting ISM mass and radio jet heating offsetting CGM cooling; Somerville et al., 2008a). As expected, this brings the SAM MW halos into even better qualitative agreement with the observations shown.

only the cold atomic and molecular gas phases (i.e., HI and H₂). This is done to prevent confusion throughout the rest of this chapter where we will simply want to compare the total ISM masses between the SAM and FIRE-2 (neglecting the physics of multi-phase gas partitioning, which is beyond the scope of this chapter). Note, however, that the SAM predictions for the cold atomic and molecular ISM gas masses alone (excluding HII) have been shown agree well with observations at $z \sim 0$ (Popping et al., 2014; Somerville et al., 2015).

Overall we find that the SAM and FIRE-2 predictions agree relatively well with each other and with observations for the stellar-to-halo and ISM-to-stellar mass ratios at fixed mass, but disagree dramatically on CGM-to-halo mass ratios. In detail, the stellar-to-halo mass ratios generally agree with the abundance matching relations within a factor of two for both the SAM and FIRE-2.⁷ We do not attempt to extend the abundance matching relations to low-mass dwarfs. As for the ISM-to-stellar mass ratio, the SAM and FIRE-2 agree relatively well with each other and with the observations for the m11 and m12 halos. This is remarkable since no attempt was made to force the SAM to reproduce FIRE-2, and FIRE-2 itself was not calibrated to match observations. However, for the m10 halos, the SAM is higher than FIRE-2 by up to a factor of ten. This order of magnitude disagreement persists if we separately compare just the cold ISM mass (atomic plus molecular; defined crudely in FIRE-2 as all gas particles at

⁷Figure 7 of Hopkins et al. (2018a) shows even better agreement for the m12 halos. Our virial and stellar masses agree with those of Hopkins et al. (2018a, Table 1) within 0.1 dex, but our stellar masses are slightly larger whereas our virial masses are slightly smaller. Hence, our estimate of the stellar-to-halo mass ratio itself will be biased higher than theirs. The virial mass disagreement can likely be attributed to our different halo finders whereas the stellar mass difference is likely due to our different assumed integration radius. We use $0.1R_{\text{vir}}$ for simplicity but they use a more refined, slightly smaller definition (three times the iteratively computed 3D stellar half-mass radius).

$< 0.1R_{\text{vir}}$ with $T < 10^4$ K) or the warm ionized gas mass (HII; defined crudely in FIRE-2 using gas particles at $< 0.1R_{\text{vir}}$ that have $T = 10^4 - 10^5$ K).

Strikingly, all of these differences are eclipsed by discrepancies in the CGM-to-halo mass ratios: the SAM predictions are orders of magnitude lower than FIRE-2, with the deficit being systematically larger for lower mass halos (~ 3 orders of magnitude). The ability to agree relatively well on stellar and ISM mass but disagree by orders of magnitude on CGM mass reflects the flexibility allowed in phenomenological models for the baryon cycle. In our case, this flexibility arises because the SAM is not calibrated to match the observed CGM masses of galaxies (which are highly uncertain; it is not clear whether the bulk of extragalactic, non-ISM gas bound to halos is located within or outside of those halos).

2.4.2 Stellar mass histories

While the previous comparison of mass-dependent scaling relations at $z = 0$ is already suggestive of significant model discrepancies, it is insightful to compare the full time evolution of various properties. We start with stellar mass in Figure 2.3. Overall, the SAM and FIRE-2 agree roughly within a factor of two. Although the SAM was tuned to reproduce the $z = 0$ stellar mass function, it is *not* tuned to reproduce observations at earlier cosmic epochs, although its predictions have been shown to be in reasonably good agreement with observations such as luminosity and stellar mass functions out to $z \sim 10$ (Somerville et al., 2015; Yung et al., 2019a,b). Two trends are evident: the SAM tends to predict higher stellar masses than FIRE-2 at early times in

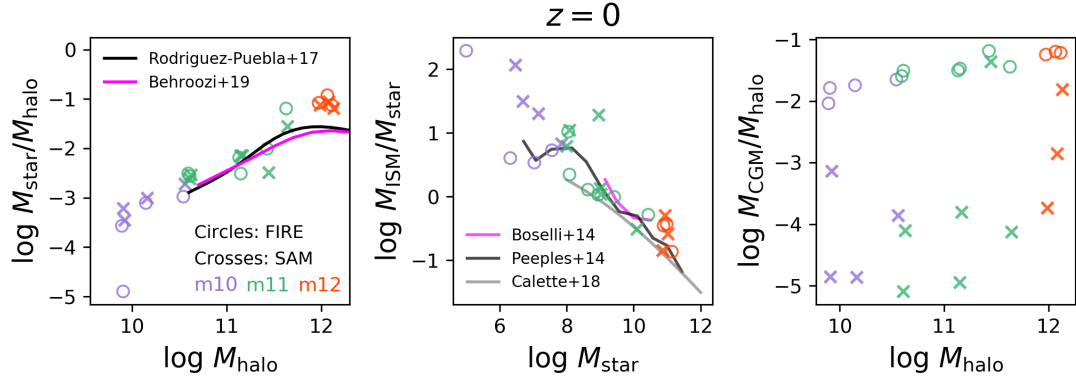


Figure 2.2: Our FIRE-2 measurements (circles) and SAM predictions (crosses) for stellar-to-halo mass ratios (left), ISM-to-stellar mass ratios (middle), and CGM-to-halo mass ratios (right) at $z = 0$. Halos are colored according to their mass bin (low-mass dwarfs in purple, intermediate-mass dwarfs in green and MW-mass halos in red). We also show observationally inferred scaling relations for median stellar-to-halo mass ratios (Rodríguez-Puebla et al., 2017; Behroozi et al., 2019) and ISM-to-stellar mass ratios (Boselli et al., 2014; Peeples et al., 2014; Calette et al., 2018). We do not show observational estimates of CGM-to-halo mass ratios since they are highly uncertain and the SAM is not calibrated to match observed CGM properties. The SAM and FIRE-2 agree relatively well with each other and with these observations in terms of stellar-to-halo and ISM-to-stellar mass ratios at a fixed mass (the ISM-to-stellar mass ratios predicted by the SAM for low-mass dwarfs would agree better with observations if we only included the cold atomic and molecular phases; Popping et al., 2014; Somerville et al., 2015). By comparison, the SAM and FIRE-2 predictions for CGM-to-halo mass ratios disagree dramatically with each other, especially for the dwarfs where the SAM predictions are generally lower by ~ 3 orders of magnitude.

MW-mass halos (by up to a factor of 10) and to a lesser extent in the m11 halos; and it also predicts higher stellar masses than FIRE-2 in the low-mass dwarfs at late times (but by less than a factor of two, except for the remarkably late-forming halo m10v, which we will discuss later).




Figure 2.3: Comparison of the stellar mass assembly history measured in FIRE-2 and as predicted by the SAM. Top: logarithmic ratio of the SAM and FIRE-2 time series color-coded by mass bin (m10 halos in purple, m11 halos in green, and m12 halos in red). Bottom: individual stellar mass assembly histories for one representative halo from each mass bin (m10q left, m11c middle, m12f right). All time series are smoothed over ~ 1 Gyr for easier visual comparison. With one exception (m10v), the SAM generally reproduces the FIRE-2 stellar mass assembly histories within a factor of two.

2.4.3 Star formation stochasticity

That the overall stellar mass assembly histories agree already suggests that the star formation histories (SFHs) must also agree when averaged over sufficiently long timescales. Indeed, we find that this is generally the case. However, on shorter timescales (~ 100 Myr), the behavior of the SAM and FIRE-2 SFHs are very different. In Figure 2.4, we show the normalized SFHs of all 13 FIRE-2 halos and include the SAM predictions. As already shown by Sparre et al. (2017) and Faucher-Giguère (2018), the FIRE-2 m10 and m11 halos have bursty SFHs at all times, whereas the more massive m12 halos are only bursty at early times ($z \gtrsim 1$ corresponding to cosmic ages $\lesssim 6$ Gyr) and settle into a more steady state at later times (see also Muratov et al., 2015; Anglés-Alcázar et al., 2017a; Ma et al., 2017). These trends are not predicted by the SAM, in which there is systematically much lower SFH variability compared to FIRE-2.

2.4.4 ISM mass histories

Figure 2.5 now compares the ISM mass histories between FIRE-2 and the SAM. Overall we see more disagreement here. The SAM predicts higher ISM masses than FIRE-2 in halos of all masses at very early times (up to a factor of ten). The SAM ISM masses are higher by at least a factor of ~ 5 -10 in nearly all the m10 halos over all of cosmic time (as discussed in subsection 2.4.1, these differences persist if we only consider the cold or warm ionized components). The m11 ISM masses predicted by the SAM tend to be higher than FIRE-2 by about a factor of 2-3 over most of cosmic time. The MW mass halos mostly show good agreement between the two methods (within a

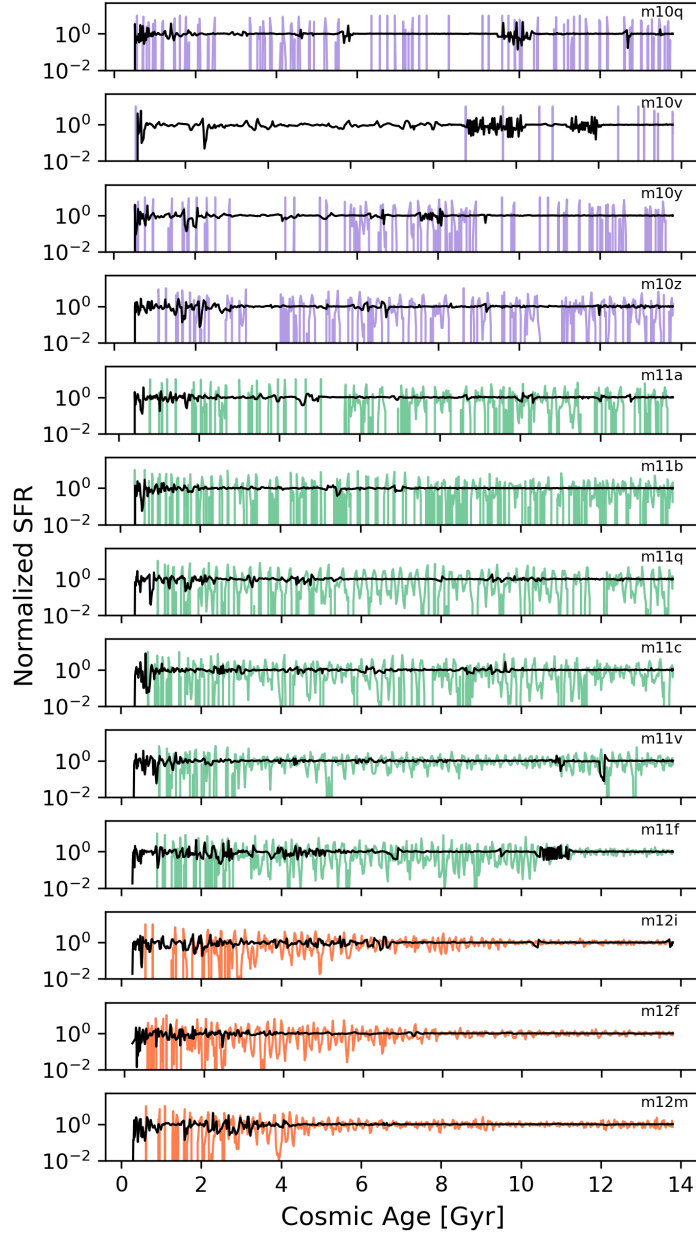


Figure 2.4: Normalized SFHs for FIRE-2 measurements (colored curves) and SAM predictions (black). The time series are the instantaneous SFHs divided by the corresponding 100 Myr boxcar-smoothed SFHs. The halos are ordered based on increasing $z = 0$ halo mass from top to bottom (m10 halos in purple, m11 in green, m12 in red). The m10 and m11 FIRE-2 halos are bursty at all times, and the m12 halos are bursty at early times but not late times. In contrast, the SAM predicts much less SFH stochasticity.

factor of 2) after a cosmic age of about 6 Gyr.




Figure 2.5: Similar to Figure 2.3 but now for ISM mass as a function of cosmic age. The SAM agrees relatively well with FIRE-2 for the m12 halos except at very early times, and this is also true for the m11 halos, albeit with more scatter. But the systematic discrepancy for the m10 halos remains at the order of magnitude level at all times, consistent with the $z = 0$ SAM excess in Figure 2.2.

2.4.5 CGM mass histories

Next we will compare the “CGM” mass predicted by the SAM and as measured in FIRE-2. In Figure 2.6, we plot the time evolution of the CGM mass in FIRE-2 and in the SAM. It is immediately obvious that the CGM mass is much lower in the SAM than in FIRE-2 in all halos at all times. The CGM mass is $\sim 3 - 4$ orders of magnitude

lower in the SAM than in FIRE-2 for the m10 and m11 halos. While the discrepancy is smaller for the m12s, the SAM still has lower CGM masses than FIRE-2 by ~ 1 order of magnitude. The “boxy” trajectories for CGM mass in the dwarfs are likely an artifact of the SAM CGM cooling model (the CGM mass may be constant when $R_{\text{cool}} > R_{\text{vir}}$ and the halo gas inflow rate equals the ISM inflow rate, assuming outflows and subhalo accretion are a negligible source of CGM mass growth; subsection 2.2.2.2).

For context, we also plot the time evolution of the “ejected” gas mass reservoir for the individual example SAM halos, and see that it dominates over the CGM mass. Most of this extragalactic (i.e., non-ISM but still bound) gas resides outside of the halo in the SAM, and its mass alone agrees better with the FIRE-2 CGM mass (especially for the MW halos).

2.4.6 Halo baryon fraction evolution

Finally, it is useful to combine the three previous bulk mass quantities and define the bulk halo baryon fraction:

$$f_{\text{b,halo}} = \frac{M_{\text{stars}} + M_{\text{ISM}} + M_{\text{CGM}}}{M_{\text{stars}} + M_{\text{ISM}} + M_{\text{CGM}} + M_{\text{DM}}} . \quad (2.9)$$

Consistent with Muratov et al. (2015), Fitts et al. (2017) and Hafen et al. (2019), in Figure 2.7 we show that lower mass FIRE-2 halos are more depleted of baryons than higher mass halos, relative to the universal baryon fraction ($f_{\text{b}} = 0.158$ according to Planck Collaboration et al., 2016). The SAM reproduces this overall trend. In more detail, the SAM predictions relative to FIRE-2 are systematically lower for the m11

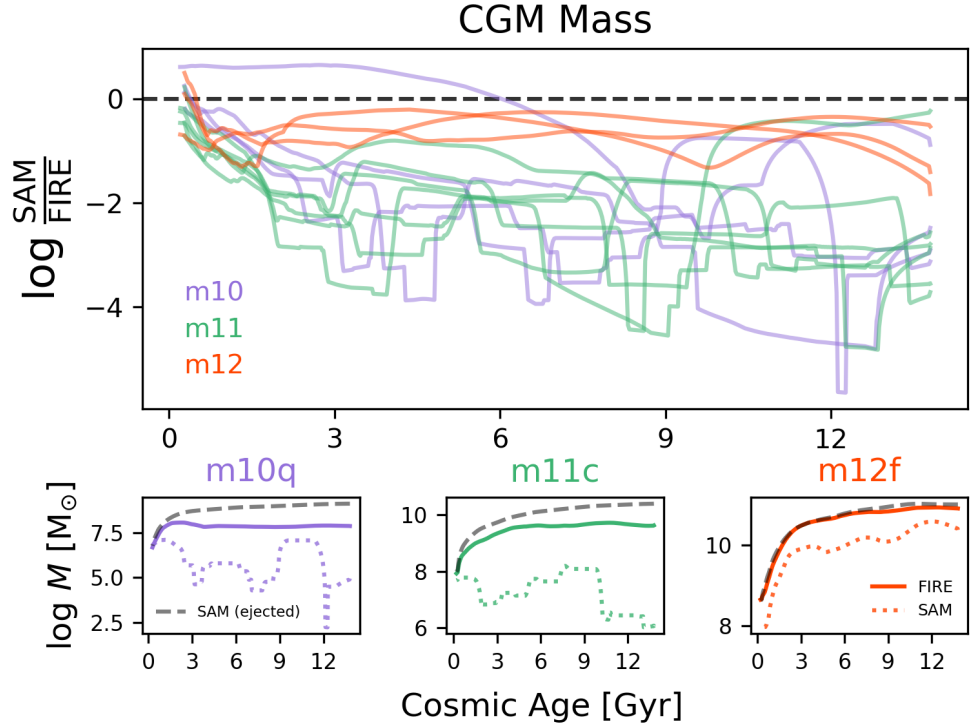


Figure 2.6: Similar to Figure 2.3 but now for CGM mass. The SAM predicts much lower CGM masses than FIRE-2 for all halos, with the deficit being worse for the m10 and m11 halos ($\sim 3 - 4$ orders of magnitude). In the individual example halo panels (bottom), we overplot the mass of the “ejected gas” reservoir (dashed gray lines) and see that it alone is larger than the FIRE-2 CGM mass. The SAM CGM masses are likely very low because most of the gas resides in this ejected reservoir (i.e., the SAM predicts that most of the extragalactic yet bound gas resides outside of the halo).

and m12 halos and similar or higher for the m10 halos. However, the differences are roughly at the factor of $\sim 2 - 3$ level at most and primarily driven by the CGM mass deficit in the SAM (which predicts that most of the extragalactic/non-ISM bound gas resides outside of the halos rather than in the CGM). The main reason why the m10 halos tend to have somewhat similar (or higher at late times) baryon fractions in the SAM than FIRE-2 is because their CGM mass deficit is somewhat offset by their ISM mass excess. It is interesting that any order of magnitude discrepancies in the individual mass components (namely ISM and CGM mass) manifest as relatively inconsequential differences for the halo baryon fraction, suggesting that this is an ambiguous quantity to interpret on its own.

2.4.7 Halo mass inflow rates

In order to better pinpoint what is driving the trends in the bulk quantities above, we now turn to a comparison of differential quantities, namely the corresponding mass inflow and outflow rates for the ISM and CGM. We begin with the halo mass inflow rate in Figure 2.8. For the MW-mass halos, the SAM agrees relatively well with FIRE-2 effectively at all times. But for progressively lower mass halos, the SAM predicts systematically higher halo gas accretion rates than measured in FIRE-2, with the discrepancy getting somewhat worse with time. For the m11 halos, the SAM is higher than FIRE-2 by more than a factor of two, and for the m10 halos, the SAM is higher than FIRE-2 by more than a factor of ten.

We can gain further insight by splitting the SAM halo gas accretion into pris-




Figure 2.7: Similar to Figure 2.3 but now for the bulk halo baryon fractions. The SAM reproduces the general trend in FIRE-2: lower mass halos are more depleted of baryons than higher mass halos, relative to the universal baryon fraction (horizontal gray lines in the bottom panels; $f_b = 0.158$ according to Planck Collaboration et al., 2016). There is relatively good agreement between the SAM and FIRE-2, with differences at the factor of $\sim 2 - 3$ level at most (despite order of magnitude differences in ISM and CGM mass; note the much smaller y -axis range in this figure compared to Figure 2.5 and Figure 2.6).

tine accretion versus re-accretion of gas that was previously ejected from the halo due to stellar feedback. It then becomes obvious that the re-accretion rate dominates over the pristine accretion rate in the dwarfs (see gray lines in the bottom panels of Figure 2.8; these halos are representative). Hence, the trend that the overall halo mass inflow rate is higher in the SAM than FIRE-2 for dwarfs is primarily driven by the high re-accretion rates. However, the pristine SAM accretion rate itself can still be significantly higher than FIRE-2 for the low-mass dwarfs, which may reflect preventative feedback not modeled by the SAM. Finally, for the MW-mass halos, the pristine accretion generally dominates over re-accretion, which is sensible since most stellar-driven winds cannot escape the potential well of these more massive halos (Equation 2.7). However, there can be dips in the pristine accretion that reflect the underlying DM halo merger history.⁸ Coincidentally, these dips are generally compensated for by the re-accretion rate, leading to overall agreement with the FIRE-2 halo inflow rates for the MW-mass halos (as seen for the example m12f halo).

2.4.8 ISM inflow rates

Figure 2.9 compares the ISM accretion rate between the SAM and FIRE-2. The SAM predicts much higher ISM accretion rates compared to FIRE-2. For the m10 halos, the SAM is higher than FIRE-2 by more than a factor of 10 whereas for the m11 and m12 halos the SAM is larger than FIRE-2 by more than a factor of two (especially

⁸For a central halo that experiences a merger, the halo mass will generally show a sharp jump because the halo finder suddenly assigns to the central halo all the particles belonging to the recently accreted subhalo. The subsequent DM accretion can be lower by comparison, especially while the halos have not fully coalesced.




Figure 2.8: Similar to Figure 2.3 but now for halo gas accretion rate. In the bottom panels, we also plot the SAM halo gas accretion rate split into pristine accretion (solid gray) and re-accretion of previously ejected gas (dashed gray). The SAM matches the MW-mass halo gas accretion rates relatively well, but predicts significantly higher values for the dwarfs (by $\sim 1 - 2$ orders of magnitude). This excess accretion in the SAM is primarily driven by its high ejected gas re-accretion rate, but the pristine accretion by itself is still higher than FIRE-2 for the low-mass dwarfs (m10q is representative).

at late times).




Figure 2.9: Similar to Figure 2.3 but now for ISM gas accretion rate. The SAM generally has higher ISM gas accretion rates than measured in FIRE-2, and this discrepancy is preferentially worse for the lower mass halos (up to two orders of magnitude). Even the m12 halos at late times have about a factor of two higher ISM accretion rates in the SAM than in FIRE-2.

2.4.9 ISM outflow rates

Next, we turn to the ISM mass outflow rate in Figure 2.10. The SAM ejects much more gas from the ISM than FIRE-2, with the discrepancy being more than a factor of 2 for most halos at most times. This is expected because if the SAM is to match the SMHM relation at $z = 0$, then it must remove the excess accreted ISM gas

via more efficient stellar feedback. Indeed, we verified that, on average, the net ISM inflow rates (inflow minus outflow) agree relatively well between the SAM and FIRE-2, with some slight discrepancies for the dwarfs (related to their excess ISM masses in Figure 2.5). However, the issue is that the SAM and FIRE-2 are achieving their similar net inflow rates in different ways.




Figure 2.10: Similar to Figure 2.3 but now for the ISM mass outflow rate. The SAM has higher ISM outflow rates than FIRE-2, with the discrepancy becoming larger than a factor of two for most halos by late times. This is necessary in the SAM to prevent excess star formation and match the observed stellar mass function.

2.4.10 Halo outflow rates

Lastly, we compare halo mass outflow rates in Figure 2.11. Again, the halo outflow rates are higher in the SAM than FIRE-2 for the m11 and m10 halos, and for the m12s at very early and at late times. This is somewhat expected given that the ISM outflow rates were higher as well, and the halo outflow rate is simply a halo circular velocity dependent re-scaling of the ISM outflow rate (Equation 2.7). However, comparing the cumulative mass ejected from the halo versus from the ISM (obtained via integration of the respective mass outflow rate histories without boxcar smoothing) as a function of time between the two models reveals a striking phenomenon. In Figure 2.12, we see that the ratio of halo outflow mass divided by ISM outflow mass is generally less than one in FIRE-2 for the m11 and m12 halos, except at very early times when the progenitor halos are in the dwarf phase. The SAM shows a qualitatively similar trend for these intermediate-mass dwarf and MW-mass halos at $z \lesssim 2$: an increasingly greater fraction of wind mass is able to leave the halo in progressively lower mass halos. However, it is striking that in FIRE-2, this ratio can exceed 1 for the m10 dwarfs. The ratio reaches a factor of ~ 1.5 for m10q and, incredibly, a factor of ~ 10 for the late-forming m10v (and even higher ratios are reached for the progenitors of all halos at very early times $z \gtrsim 6$). This implies that more mass has left the halo than has ever left the ISM (cumulatively), and is suggestive of entrainment of ambient CGM material by outflows (see also Muratov et al., 2015; Hafen et al., 2019, 2020). In contrast, the ratio can never exceed 1 by construction in the SC SAM. Hence, the SAM predicts

that nearly all winds will leave the halo in low-mass dwarfs as specified by the function Equation 2.7, but any potential effects resulting from entrainment are not captured by the SAM.




Figure 2.11: Similar to Figure 2.3 but now for the halo mass outflow rates. The SAM has higher halo mass outflow rates than the FIRE-2 measurements. This is not surprising because the halo mass outflow rate in the SAM is a halo circular velocity-dependent re-scaling of the ISM outflow rate, and the latter was already shown to be much higher than FIRE-2.

2.5 Discussion

Here we interpret the results from our comparison, discuss possible solutions to the model discrepancies with an emphasis on developing ways to modify SAMs to

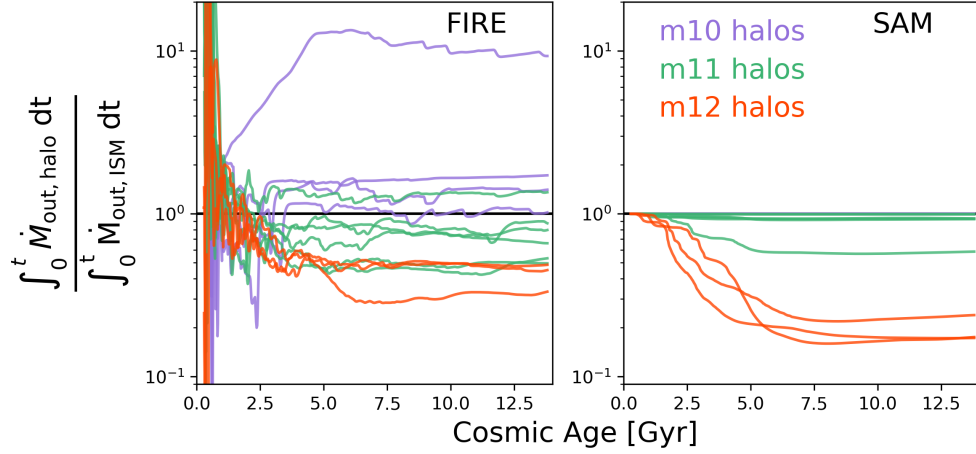


Figure 2.12: The ratio of the cumulative mass ejected from the halo versus from the ISM as a function of cosmic time for FIRE-2 (left panel) and the SAM (right panel). The lines are color-coded according to $z = 0$ halo mass. The SAM qualitatively follows the FIRE-2 trends for the m11 and m12 halos at $z \lesssim 2$ (i.e., cosmic ages $\gtrsim 3$ Gyr): an increasingly larger fraction of winds are able to leave the halo in progressively lower mass halos. But in the FIRE-2 simulations, the m10 halos strikingly tend to have expelled more material through R_{vir} than has ever left the ISM boundary ($0.1R_{\text{vir}}$), implying significant entrainment of ambient CGM material by the outflows (this is also true for the progenitors of all halos at very early times $z \gtrsim 6$). Since the ratio cannot exceed 1 in the SAM by construction, it asymptotes to 1 for the low-mass dwarfs (all of their winds leave the halo).

produce better agreement with FIRE, and present a new preventative stellar feedback model to help interpret the suppressed dwarf halo gas accretion efficiencies in FIRE-2.

2.5.1 Interpreting the model discrepancies

In this subsection, we will step through each of the four mass flow rates in turn and discuss the possible causes and solutions for the SAM versus FIRE-2 model discrepancies. But first we provide a high-level summary of the basic story, which is also summarized with a cartoon schematic in Figure 2.13. We showed that two very reasonable models of galaxy formation – the Santa Cruz SAM and the FIRE-2 simulations – agree relatively well with each other in terms of their stellar and ISM mass histories (Figure 2.2, Figure 2.3 and Figure 2.5). However, the two models disagree dramatically in terms of their CGM mass histories, with the SAM remarkably predicting ~ 3 orders of magnitude lower CGM mass than FIRE-2 for the dwarf halos (Figure 2.2 and Figure 2.6). The SAM assumes that most of the “missing” extragalactic gas resides outside of the halo in a so-called “ejected” reservoir (owing partially to observational uncertainties about the total CGM masses of galaxies). To better understand the discrepancies, we turned to the actual mass flow rates for the ISM and CGM. The fundamental discrepancy between the SAM and FIRE-2 arises in the halo gas accretion rate (Figure 2.8). While there is reasonable agreement for the MW-mass halos, the SAM predicts much higher halo gas accretion rates for the dwarfs than FIRE-2 (exceeding a factor of ten for the low-mass dwarfs by $z = 0$). The ISM gas accretion rates are also higher in the SAM than FIRE-2 by more than a factor of two for the m11 and m12 halos, and by more

than a factor of ten for the low-mass dwarfs (Figure 2.9). These higher inflow rates in the SAM are compensated for by higher ISM and halo outflow rates in the SAM compared to FIRE-2 (Figure 2.10 and Figure 2.11), making it possible to understand why the SAM and FIRE-2 predict similar stellar mass and ISM mass histories. In addition to these discrepancies, the SAM also does not capture star formation stochasticity (Figure 2.4) and the entrainment of ambient CGM material by outflows from low-mass dwarfs (Figure 2.12).

We begin by diagnosing the higher halo gas accretion rates of dwarf halos in the SAM compared to FIRE-2. We showed in Figure 2.8 that re-accretion of previously ejected gas dominates over pristine accretion for the SAM dwarf halos. Hence, to first order the halo gas ejection and recycling model must be updated, but this is an area of uncertainty that has long plagued SAMs. Previous works have shown that the way in which halo gas ejection and re-accretion is implemented in SAMs can significantly affect results. Early models were split between allowing no re-accretion at all versus assuming a single re-accretion timescale (e.g., see section 2.6 of Somerville & Primack, 1999, and references therein). Somerville et al. (2008a) claim that some re-accretion is necessary to match the observed baryon fractions of galaxy clusters (which would otherwise be predicted to be too low), but simultaneously reproducing the late formation times and mass functions of dwarfs has presented challenges. Henriques et al. (2013) proposed that the re-accretion timescale should depend inversely on halo mass with no dependence on redshift because that allowed their SAM to better match the observed evolution of dwarfs. White et al. (2015) re-visited this issue with the Santa Cruz SAM

and tested three alternative solutions for de-coupling the star formation and halo gas accretion histories of dwarfs: adding a redshift dependence for the mass loading factor of stellar-driven winds, changing the gas depletion timescale for star formation, and changing the re-accretion timescale as in Henriques et al. (2013). Their comprehensive observationally-driven study concluded, in qualitative agreement with Henriques et al. (2013), that preferentially increasing the re-accretion timescale for dwarfs may be the most promising solution.

Another approach for guiding SAMs is to explicitly track halo gas recycling in high-resolution simulations. Anglés-Alcázar et al. (2017a) showed that recycling is ubiquitous and occurs over a broad range of timescales in the FIRE-1 simulations, although the recycling events generally happen in the inner halo (“fountain flows”) leading to median recycling timescales of only $\sim 100 - 350$ Myr (see also Hafen et al., 2019, 2020, for the FIRE-2 suite). Interestingly, Christensen et al. (2016) and Tollet et al. (2019) both find longer median recycling timescales of ~ 1 Gyr in their respective zoom-in simulation suites, with little or no dependence on halo mass, supporting the use of a single recycling time as adopted in some SAMs. Indeed, our SAM assumes a single recycling time (roughly on the order of a Hubble time), with the caveat that our recycling refers to gas already ejected from the halo whereas many of the previous simulation analyses define recycling within the halo. Tollet et al. (2019) and Hafen et al. (2020) also emphasize the inherently multi-phase nature of outflows in their simulations, with the hot component more easily able to leave the halo and the cooler component likely to be recycled at the inner halo via fountain flows. This general multi-phase picture for

outflows is in agreement with even higher resolution but smaller scale simulations (e.g., Kim & Ostriker, 2018; Fielding et al., 2018; Li & Bryan, 2019; Li & Tonnesen, 2019; Kim et al., 2020a). This suggests that decreasing the high halo gas accretion rates of dwarfs involves more than just preferentially increasing their recycling timescales, namely improving how the ejected SAM component is modeled in the first place. Some fraction of gas in the ejected reservoir should be allowed to become unbound permanently (especially for the hottest, fastest winds in dwarfs), and a distinct bound wind component should be modeled which does not have enough energy to escape beyond R_{vir} but instead may recycle rapidly in the inner halo. In addition, it may not be necessary to appeal solely to moving gas outside of the halo to reduce CGM cooling rates and hence SFRs; instead, some fraction of the “ejected” reservoir could actually be placed within the CGM but in a thermal state that simply does not cool efficiently.

Figure 2.8 also implies that updating halo gas ejection and recycling may not be enough: the pristine gas accretion rates alone can still be significantly higher in the SAM than FIRE-2 for the low-mass dwarfs (m10q is representative). It is tempting to attribute this to the different UV background model assumed in the SAM (taken from the simulations of Okamoto et al., 2008, who themselves adopted the UV background model of Haardt & Madau 2001) versus in FIRE-2 (Faucher-Giguère et al., 2009). According to the Okamoto et al. (2008) prescription, the characteristic mass at which the bulk halo baryon fraction drops to half of the universal value at $z = 0$ is $\approx 10^{10} M_{\odot}$. It is a factor of a few higher, $\approx 5 \times 10^{10} M_{\odot}$ according to the simulations of Faucher-Giguère et al. (2011, their Figure 7), who implemented the Faucher-Giguère et al. (2009) UV

background model.⁹ Using the SAM, we have experimented with increasing the characteristic mass and/or changing the redshift of reionization (within a reasonable range of values) but find that this cannot satisfactorily explain the suppressed halo accretion, especially for the intermediate-mass dwarfs.¹⁰ The main exception is m10v, the late-forming low-mass dwarf for which our SAM predictions disagree dramatically with FIRE-2: its virial mass does not exceed the SAM fiducial characteristic mass until $z \sim 2$, compared to $z \sim 10$ for the other m10 halos. Interestingly, previous authors have argued that the low halo baryon fractions and accretion rates of the FIRE-2 dwarfs can at least partially be attributed to the UV background (Fitts et al., 2017; El-Badry et al., 2018; Hafen et al., 2019), in contrast to the weaker effects predicted by the SAM.

What else could possibly suppress the halo gas accretion rates of the FIRE-2 dwarfs?¹¹ There is an emerging consensus that some form of preventative feedback is needed in SAMs beyond UV background heating alone. Hirschmann et al. (2012) already showed that the Santa Cruz SAM predicts much higher halo gas accretion rates compared to their reference suite of cosmological zoom-in simulations (see their Figure 11). Interestingly, Lu et al. (2011a) found the opposite when comparing their SAM to the cosmological simulations of Kereš et al. (2009). Nevertheless, both of these

⁹But note that the Faucher-Giguère et al. (2011) simulations pre-date the FIRE-2 subgrid models, hydrodynamic solver, etc. To properly assess the redshift evolution of the characteristic mass in FIRE-2, we would need FIRE-2 simulations with all stellar feedback turned off, such that only the UV background and gravitational shock heating can systematically suppress gas inflows at R_{vir} .

¹⁰Note also that significantly changing the redshift of reionization or characteristic mass normalization would cause other predictions of the SAM to disagree with FIRE-2 and observations, and possibly make the SAM assumptions inconsistent with cosmology constraints from Planck Collaboration et al. (2016).

¹¹The large-scale cosmic web environment of a halo can be very relevant, but we think this is unimportant for our sample of FIRE-2 halos which are relatively isolated and have “typical” accretion histories for their mass (Hopkins et al., 2018a).

authors later assumed general “pre-heating” to suppress halo gas accretion rates for their SAMs (Lu et al., 2015; Hirschmann et al., 2016; Lu et al., 2017). More recently, Tollet et al. (2019) characterized the baryon cycle in the NIHAO simulations (Wang et al., 2015) and also argued that SAMs would need a new “maintenance feedback” channel to achieve lower cooling rates. They showed that in the NIHAO simulations, stellar-driven outflows from dwarf halos divert otherwise inflowing gas supplied by cosmic web filaments on scales as large as $6R_{\text{vir}}$, resulting in suppressed accretion. Furthermore, the entrainment of outflows implied by our Figure 2.12 may have additional preventative feedback effects that need to be better understood (this phenomenon is also seen in the FIRE-1, NIHAO and EAGLE simulations, respectively, by Muratov et al., 2015; Tollet et al., 2019; Mitchell et al., 2019). In the next section, we will present a simple but physically-motivated model for preventative stellar feedback that agrees remarkably well with the reduced halo gas accretion efficiencies in FIRE-2.

Now we turn to the ISM inflow rate: the ISM accretion rates may be higher in the SAM than in FIRE-2 in part because the same is already true for the halo gas accretion rates. However, subtle details of the SAM halo gas cooling model (based on White & Frenk, 1991) may also cause the ISM inflow rates to disagree. It is notable that although the CGM mass in the SAM is much lower than it is in FIRE-2 dwarfs, the ISM accretion rates are *higher*. In the regime where $R_{\text{cool}} < R_{\text{vir}}$ (“hot/slow” mode accretion), a higher overall CGM mass would likely correspond to higher cooling rates in the SAM (subsubsection 2.2.2.2). This implies that if the SAM CGM masses were somehow made to agree better with FIRE-2, the ISM inflow discrepancy would

presumably become worse with the existing SAM cooling model. However, the simple SAM assumption that gas accretes into the ISM on a dynamical time when $R_{\text{cool}} > R_{\text{vir}}$ (“cold/fast” mode accretion) could also be a factor. If gas accretes into the ISM too quickly, without spending enough time in the CGM, this would be consistent with both the lower CGM masses and higher ISM inflow rates of SAM dwarfs compared to FIRE-2. Since it is likely that most of the dwarf halos spend most of their lifetime experiencing this so-called “cold/fast” mode accretion in the SAM, this is an important regime to study in the future.

A critical point is that the SAM does not include a heating term due to stellar-driven winds that can offset the predicted halo gas cooling rate. In the FIRE-2 dwarfs, it is almost certainly the case that the energy and momentum of stellar-driven winds are suppressing accretion on the scale of the ISM (this may even have an effect in the MW-mass halos at late times, where the SAM ISM accretion rates are higher than FIRE-2 by more than a factor of two; Figure 2.9). In addition, the calculation of a “cooling radius” and ad hoc treatment of the case when it is greater than the virial radius can lead to unphysical looking behavior for the dwarfs (e.g., the “boxy” trajectories for CGM mass in Figure 2.6). Even for the radiative cooling timescale calculation itself, the assumed singular isothermal CGM mass density profile is likely an oversimplification for the simulated halos since bursty inflows and outflows may cause the CGM to have a more dynamic structure. SAMs do not generally explicitly model the structure and dynamics of the CGM, but this is slowly changing with work on new cooling flow solutions (e.g., Lu et al., 2011a; Stern et al., 2019) and explicit CGM substructure

models (e.g., Maller & Bullock, 2004; Voit et al., 2015; Faerman et al., 2019; Lan & Mo, 2019). Explicitly modeling the CGM with SAMs is important given that modern cosmological hydrodynamical zoom-in simulations, including the FIRE suite, might lack the resolution requirements to capture some of the relevant cooling and shock heating microphysics (e.g., see the recent enhanced halo resolution studies by Hummels et al., 2019; van de Voort et al., 2019; Peeples et al., 2019).

Finally, switching to the outflow side: it is again not surprising that the ISM outflow rates are much higher in the SAM than in FIRE-2 given the agreement between their stellar mass histories. The only plausible way to decrease the SAM ISM outflow rates is to implement preventative feedback that suppresses the high gas accretion rates in the first place. Improvements in this area may fundamentally require changing how we model “disk mode” star formation and what we assume about variations in the local star formation efficiency (e.g., Khullar et al., 2019, and references therein). Indeed, the order of magnitude ISM mass excess but factor of two agreement on stellar mass for low-mass dwarfs predicted by the SAM compared to FIRE-2 suggests that the assumptions for how gas forms stars are different in the two models. In addition, small-scale simulations suggest that preventative feedback effects may be stronger during bursty star formation episodes since those result in clustered supernovae that drive faster, more energetic winds (e.g., Gentry et al., 2017; Fielding et al., 2017a). To achieve local star formation efficiency variations and stochasticity in a physically self-consistent way, the SAM outflow model itself may need to be replaced with one that depends exclusively on local properties. Ideally, on the ISM scale, the mass, energy, momentum and metal

mass from stellar feedback should be deposited locally, e.g., within annuli of a radially-resolved disk (e.g. Forbes et al., 2019). For halo outflows, while the traditional approach of setting a wind escape fraction that depends on the global halo circular velocity may still be viable, Figure 2.12 suggests a need for additional variables that account for entrainment and ejection of ambient CGM material by multi-phase outflows (see also Guo et al., 2011; Muratov et al., 2015; Hu, 2019; Li & Bryan, 2019; Tollet et al., 2019; Mitchell et al., 2019; Hafen et al., 2019, 2020).

The discrepancies between the Santa Cruz SAM and FIRE-2 have implications for other models of galaxy formation. That two models with very different underlying baryon cycles can still match the observed evolution of the stellar mass function, and by extension the low-mass end of the stellar-to-halo mass relation, emphasizes ambiguities for interpreting observations with phenomenological models. These ambiguities are amplified even more with subhalo abundance matching and “semi-empirical models” that make even simpler assumptions for how halo mass accretion rates relate to galaxy star formation rates (e.g., Behroozi et al., 2013a; Moster et al., 2013; Rodríguez-Puebla et al., 2016, 2017; Moster et al., 2018; Tacchella et al., 2018; Behroozi et al., 2019). For example, it is common practice in these models to define the star formation efficiency as $\text{SFE} = \frac{\text{SFR}}{f_b \dot{M}_{\text{vir}}}$. If indeed the halo gas accretion rates of dwarfs follow the universal baryon fraction, then this would imply low SFEs in dwarfs. But our study suggests that it is also possible for the reverse interpretation to be true: for a given SFE, if less gas is flowing into the halo in the first place, then this can also explain the lower SFRs of dwarfs. With a SAM coupled to high resolution simulations, we can explicitly isolate

and model these preventative physical processes (as in the next section) and ultimately study the implications for the evolution of the galaxy–halo connection.

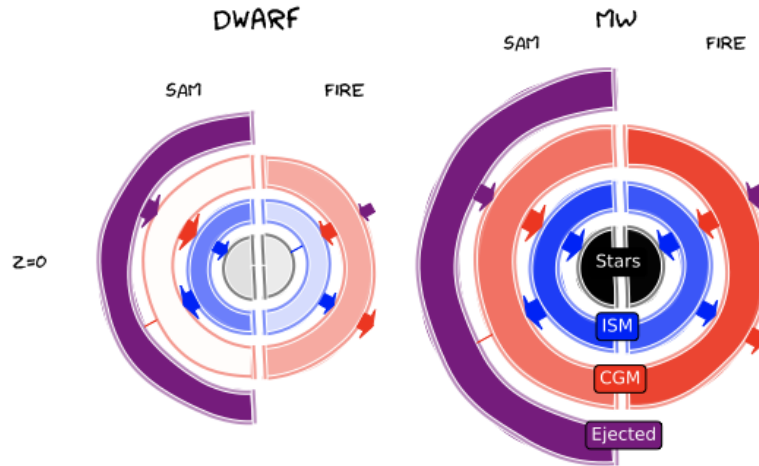


Figure 2.13: A cartoon that schematically illustrates the results of our comparison between the SAM and FIRE-2 (restricted to $z = 0$ for simplicity). On the left we show a representative dwarf (m10q) and on the right a representative MW-mass halo (m12f). For each halo, the left half portrays the SAM and the right half depicts FIRE-2. From inside out, we show the bulk masses of stars (black), ISM (blue), CGM (red) and the ejected gas reservoir (magenta; restricted to the SAM since there is no clear definition of this component for FIRE-2). The opacity can be used to compare the mass of a single component between the two models or the mass of different components within a single model. The arrows illustrate inflows and outflows between the different bulk mass components (note that the purple arrows show the total halo gas accretion rate, not the recycling rate). Larger size arrows convey higher flow rates. Note how the stellar masses agree very well between the two models for both galaxies despite significant differences in the other bulk components.

2.5.2 A simple preventative stellar feedback model

Here we present a simple physical model for how preventative stellar feedback can suppress gas accretion rates preferentially for dwarf halos on the scale of R_{vir} . We

deliberately keep the model simple as the goal here is to demonstrate that a reasonable model can approximately match the FIRE-2 inflow results, rather than trying to develop a detailed prescription for inclusion in semi-analytic models, a task we defer to future work. The essence of our model is that the energy from SN-driven winds can heat some fraction of the gas beyond R_{vir} to the virial temperature (or higher). Since the virial temperature is a measure of the gravitational potential depth, this would then imply that the heated gas becomes unbound from the halo and hence is unable to accrete.¹² Note that preventative feedback in this context refers to preventing gas from accreting into the halo in the first place (as in Lu et al., 2017), rather than preventing halo gas from accreting into the ISM (e.g., Mitra et al., 2015).

First, we define

$$f_{\text{in}} = \frac{\dot{M}_{\text{in,baryons}}}{f_{\text{b}}\dot{M}_{\text{in,DM}}} \quad (2.10)$$

as the ratio of the actual baryonic mass inflow rate ($\dot{M}_{\text{in,baryons}}$) to the baryon fraction-adjusted DM mass inflow rate ($f_{\text{b}}\dot{M}_{\text{in,DM}}$) at the virial radius.¹³

We can obtain an expression for the amount of gas mass that must be heated to suppress the accretion rate by first writing down an expression for halo gas binding energy:

$$E_{\text{b}} = \frac{1}{2}f_{\text{b}}M_{\text{vir}}V_{\text{vir}}^2. \quad (2.11)$$

Next, we take the time derivative of this expression and equate it to the heating rate

¹²The unbound, low-density hot gas may then travel outwards before eventually turning around and recooling on to the halo (e.g., as illustrated in Figure 1 of Noh & McQuinn, 2014). More complicated models may predict the detailed evolution of this gas, but here we restrict ourselves to simply deriving an effective suppression fraction for the initially accreting gas.

¹³Note that we should also multiply by an additional factor f_{coll} to account for UV background heating (Equation 2.1), but this is likely negligible for most of our halo masses, as we will show later.

(\dot{E}_{heat}); if we assume V_{vir} is constant and we isolate the gas mass term ($f_{\text{b}}\dot{M}_{\text{vir}}$), we get the needed mass heating/unbinding rate:

$$f_{\text{b}}\dot{M}_{\text{in,DM}}(1 - f_{\text{in}}) = \frac{2\dot{E}_{\text{heat}}}{V_{\text{vir}}^2}, \quad (2.12)$$

where we have used our definition of f_{in} to replace $\dot{M}_{\text{in,baryons}}$. The factor $(1 - f_{\text{in}})$ comes in because we want to equate the heating rate with the fraction of gas that does not accrete.¹⁴

Assuming the heating is provided by energy from star formation, we can write $\dot{E}_{\text{heat}} = \eta_{\text{E}}e_{\text{SN}}\text{SFR}$, where SFR is the star formation rate, $e_{\text{SN}} = 10^{51}\text{erg}/(100M_{\odot})$ is the specific energy produced by SNe per $100M_{\odot}$ of stars formed (this is approximate at the order of magnitude level given a reasonable assumption for the IMF), and η_{E} is the efficiency in transporting that energy from the SN site to the virial radius. Doing this allows us to solve for f_{in} :

$$f_{\text{in}} = 1 - 2\eta_{\text{E}}\frac{e_{\text{SN}}}{V_{\text{vir}}^2}\frac{\text{SFR}}{f_{\text{b}}\dot{M}_{\text{in,DM}}}. \quad (2.13)$$

Note that the ratio $e_{\text{SN}}/V_{\text{vir}}^2$ will be higher for dwarfs owing to their lower V_{vir} . The other important term is the star formation efficiency ratio $\text{SFR}/(f_{\text{b}}\dot{M}_{\text{in,DM}})$. Since the SFE defined in this way is generally lower for dwarfs, this new term acts in the opposite direction of the $e_{\text{SN}}/V_{\text{vir}}^2$ trend. To make further progress, we therefore need a prediction for SFR (or equivalently SFE). One option is to take this from the SAM or simulation itself (perhaps suitably time-shifted to allow for a delay as the energy

¹⁴An alternative derivation is to directly balance the heating rate with the specific gravitational potential energy of the fraction of gas beyond R_{vir} that was heated to at least T_{vir} and hence unable to accrete: $\dot{E}_{\text{heat}} = f_{\text{b}}\dot{M}_{\text{in,DM}}(1 - f_{\text{in}})\frac{k_{\text{B}}T_{\text{vir}}}{\mu m_{\text{H}}}$.

flows from the ISM to the virial radius). However, here we assume a simple equilibrium “bathtub” model (e.g., Davé et al., 2012; Mitra et al., 2015) in which the amount of gas in the ISM is fixed (at least over short periods of time) such that the amount of inflowing gas is balanced by the outflowing gas. In this case, we can write

$$\text{SFR} = \frac{f_{\text{in}} f_{\text{b}} \dot{M}_{\text{in,DM}}}{1 + \eta_{\text{M}}}, \quad (2.14)$$

where η_{M} is the mass-loading factor, or ratio of the mass outflow rate (near the ISM) to the star formation rate. Using this relation in Equation 2.13 allows us to solve for f_{in} :

$$f_{\text{in}} = (1 + \psi)^{-1}, \quad (2.15)$$

where

$$\psi \equiv \frac{2\eta_{\text{E}} e_{\text{SN}}}{(1 + \eta_{\text{M}}) V_{\text{vir}}^2} \quad (2.16)$$

is the ratio of specific SN energy to the specific halo gravitational potential, accounting for our mass and energy loading efficiencies. If the ratio $\eta_{\text{E}}/(1 + \eta_{\text{M}})$ was a constant, then ψ will be larger and hence f_{in} will be smaller for lower mass halos. This would give the expected qualitative behavior that a lower fraction of gas is able to accrete into dwarf halos. However, as a last step, we need to explicitly consider how η_{M} and η_{E} may evolve with halo mass and/or redshift. For η_{M} , we directly take the broken power law relation for the FIRE-1 simulations from Muratov et al. (2015). According to their equations 4 and 5, η_{M} follows a steeper power law for halos with $V_{\text{vir}} < 60 \text{ km s}^{-1}$ and there is a slight redshift dependence. For the energy loading factor η_{E} on the scale of R_{vir} , there is less precedent. We therefore consider two simple possibilities. First, we

assume a constant $\eta_E = 0.1$ motivated by the ISM wind breakout condition study of Li & Bryan (2019). Alternatively, we hypothesize that lower mass halos have preferentially higher energy loading factors (which is plausible given their preferentially higher mass loading factors and the apparently energy-conserving nature of their winds; Muratov et al., 2015). Specifically, we assume $\eta_E = \varepsilon_{\text{heat}}(1 + \eta_M)$ where $\varepsilon_{\text{heat}}$ is a constant that parameterizes our ignorance about the conversion from ISM mass loading to ISM energy loading and then to halo energy loading. With this simple parameterization, the strong halo mass (and slight redshift) dependence of η_M from Muratov et al. (2015) is canceled out, allowing us to see how our model behaves if indeed the ratio $\eta_E/(1 + \eta_M)$, rather than η_E alone, is constant.

In Figure 2.14, we plot the halo gas accretion efficiency as a function of halo mass for the FIRE-2 halos at $z = 0$ and $z = 2$, where we define the accretion efficiency to be the ratio of the gas accretion rate to the DM accretion rate in the virial shell (without any boxcar time smoothing). If gas accretion perfectly tracked DM accretion at R_{vir} as commonly assumed in halo models, then the halos should all lie along the mass-independent universal baryon fraction line ($f_b = 0.158$; Planck Collaboration et al., 2016). We know that heating from the UV background can preferentially suppress gas accretion into low-mass halos, so this cannot be strictly true. However, as we have already discussed above, the accretion efficiencies of the FIRE-2 halos fall below the expected suppression due to the UV background alone (comparing to Okamoto et al., 2008, which is the relation assumed in the Santa Cruz SAM). This is perhaps not so surprising because UV background heating is thought to primarily affect halos with

much lower masses than ours. Turning to the version of our model with a constant $\eta_E = 0.1$, we see that it is incapable of describing the data points; in fact, at low halo masses, this version of the model shows an upturn in f_ϵ . However, if we adopt the second version of the model with $\eta_E = \epsilon_{\text{heat}}(1 + \eta_M)$, and set $\epsilon_{\text{heat}} = 0.01$ (implying that η_E is preferentially higher, i.e., of order unity, in the low-mass dwarfs), then the prediction from our simple model matches the data points remarkably well, especially at $z = 0$. The predicted suppression at $z = 2$ is somewhat stronger than the data points, which may suggest that ϵ_{heat} should have a redshift dependence and/or that our simple equilibrium bathtub model is breaking down.

We again stress that our preventative stellar feedback model is very simple and although it is promising, there are several unknowns that should be addressed in the future. First and foremost, we started by assuming that SN-driven winds can reach R_{vir} and heat a fraction of the surrounding gas to the virial temperature or higher. This is certainly a plausible assumption for the FIRE-2 dwarfs given their high halo-to-ISM cumulative outflow mass ratios (Figure 2.12). It should be less the case in the MW-mass halos since winds would need a higher velocity to escape the potential well of these more massive halos; however, in detail this depends on the relative fraction of hot, fast-moving wind versus cooler, slower moving wind, and the rate at which the thermal energy of the wind is lost to the ambient CGM due to interactions/mixing (we have not distinguished between kinetic and thermal energy for the SN winds). Directly characterizing η_E and η_M for the FIRE-2 halos would be of great interest for testing and calibrating our model in the future. In addition, there will be degeneracies between

preventative stellar feedback, ejective stellar feedback, and gravitational shock heating of gas accreting onto the more massive halos. The implications of these degeneracies for the galaxy–halo connection can be explored with a SAM in the future, provided that the physical processes have been modeled and calibrated to faithfully represent the hydrodynamical simulations.

Many previous works have already suggested that preventative stellar feedback is important in dwarfs. Dekel & Silk (1986) derived the equations for SN-driven heating of halo gas and the implications for ejecting gas (based on comparing the specific SN energy with the halo virial temperature; see their section 4). Here we are explicitly considering suppression of gas accretion rather than gas ejection alone. Oppenheimer & Davé (2009) and Oppenheimer et al. (2010) used hydrodynamical simulations to infer that SN-driven winds must have an additional heating/preventative effect to offset gas cooling, but their results were not parameterized and easily translatable to SAMs (see also Pawlik & Schaye, 2009; van de Voort et al., 2011; Christensen et al., 2016; El-Badry et al., 2018; Wright et al., 2020). Lu et al. (2015) and Lu et al. (2017) explicitly implemented preventative feedback in their SAM and found that it is required (along with ejective feedback) to simultaneously explain the observed stellar mass–metallicity relation and the stellar mass function. However, their preventative feedback equation is more schematic in nature, and can be ascribed to “pre-heating” by a multitude of processes in a more general sense (see also Hirschmann et al., 2016). In contrast, we have explicitly constructed a model that isolates one potentially important preventative effect of SN-driven winds alone. There are likely additional preventative stellar feedback

effects such as an energy input rate into the ambient CGM that can offset the predicted radiative cooling rate and possibly even eject ambient CGM material (e.g., Guo et al., 2011).

2.5.3 Results using dark matter only simulations

Throughout this chapter, we have run the SAM on merger trees extracted for dark matter halos from the main hydrodynamical FIRE-2 simulations, i.e., baryons have affected the properties of DM halos in the merger trees. This was done to increase our sample size of halos (13) because DM-only N -body simulations do not exist for all of the FIRE-2 halos. Here, we re-run the SAM on a subset of the FIRE-2 halos that have corresponding DM-only simulations available (with the same initial conditions, resolution, snapshot output times, etc.). We run Rockstar and consistent-trees on these DM-only simulations in the same way as described in Section 2.3 except we no longer force Rockstar to up-weight the DM particle mass since there are no baryonic particles to account for.

Before comparing the SAM results, it is useful to compare a few relevant halo properties measured with Rockstar in the hydrodynamical versus corresponding DM-only N -body simulations. Figure 2.15 overplots the time series of the halo DM mass accretion rate, M_{vir} , R_{vir} , halo spin parameter and halo concentration from the two matching runs for each of the five halos. The baryonic effects on these DM halo properties are generally not significant. The DM mass accretion histories have the same normalization on average, except that some spikes in the halo mass accretion rate are




Figure 2.14: The halo gas accretion efficiency ($\dot{M}_{\text{in,gas}}/\dot{M}_{\text{in,DM}}$ in the virial shell) as a function of the halo virial mass for the FIRE-2 simulations at $z = 0$ (black points) and $z = 2$ (magenta points). If gas accretion perfectly tracked DM accretion at R_{vir} with the universal baryon fraction ($f_b = 0.158$ from Planck Collaboration et al., 2016) as commonly assumed, then the halos would all lie along the horizontal solid gray line. The Okamoto et al. (2008) model describing the suppression of halo gas accretion due to the ionizing UV background is shown with the black dashed line for $z = 0$ and the magenta dashed line for $z = 2$. The FIRE-2 halo gas accretion efficiencies fall below the expectation from UV background heating alone. The dotted lines show the behavior of our simple preventative stellar feedback model if we assume the halo energy loading factor is a constant 0.1; we see that it is incapable of explaining the data points, in fact showing the opposite trend at low masses. However, the solid lines show that our model can explain the data points remarkably well if we assume that the halo energy loading factor is preferentially higher for dwarfs ($\eta_E = \epsilon_{\text{heat}}\eta_M$ with $\epsilon_{\text{heat}} = 0.01$, implying that η_E is of order unity for low-mass dwarfs). The agreement is better at $z = 0$ than $z = 2$, suggesting either a redshift dependence for ϵ_{heat} or that our simple model is breaking down.

suppressed in the hydrodynamical version of the merger trees. This suppression of spikes in the mass accretion history is likely related to the virial mass and virial radius being smaller in the hydrodynamical run compared to the pure DM-only run, although the difference is only at the 10 – 20% level. The halo spin parameters are nearly identical as a function of time. The main systematic trend is in the halo concentration parameter: lower mass halos have lower concentrations in the hydrodynamical run, presumably due to adiabatic expansion of the halo due to strong baryonic feedback (Fitts et al., 2017). In contrast, the halo concentration parameter is larger in the hydrodynamical simulation for the MW-mass halos, presumably due to the greater central mass of baryons leading to adiabatic contraction of the halos (as is analytically expected, e.g., Mo et al., 1998; Dutton et al., 2016).

In Figure 2.16 we compare the SAM predictions when run on the FIRE-2 hydrodynamical and DM-only merger trees for the same five halos. It is immediately apparent that the two sets of SAM results agree relatively well with each other. As a consequence, the DM-only-based SAM trends relative to FIRE-2 remain qualitatively the same, and our conclusions would not have changed if we used the DM-only simulation merger trees instead of the fiducial hydrodynamical simulation merger trees. For example, the DM-only SAM still predicts similarly higher halo gas inflow rates for low-mass dwarfs than in FIRE-2 (by $\sim 2 - 3$ orders of magnitude). The main systematic difference between the two sets of SAM predictions is that the DM-only-based SAM predicts somewhat higher stellar masses for dwarfs than the hydrodynamic merger tree-based SAM. This might be due to the higher halo concentrations for dwarfs in the

DM-only simulations (no adiabatic expansion due to baryons) leading to smaller predicted disk sizes, which in turn causes gas surface densities and hence higher SFRs. In addition, the $z = 0$ CGM masses of the two MW-mass halos also agree better with FIRE-2 using the DM-only-based SAM, although the dwarfs continue to show similarly low CGM masses by orders of magnitude in the SAM compared to FIRE-2. Hence, while there are some relatively minor discrepancies that suggest a deeper look at how the SAM treats baryonic effects on DM halos, in the context of the global baryon cycle that is the main focus of this chapter, our conclusions remain the same overall.

2.6 Summary

We have used the FIRE-2 cosmological hydrodynamical “zoom-in” simulations (Hopkins et al., 2018a) to test some of the fundamental assumptions in the Santa Cruz SAM (Somerville et al., 2015) related to the global baryon cycle. We ran the Santa Cruz SAM on the FIRE-2 merger trees and compared, on an individual halo-by-halo basis, the time evolution of the masses of various components (stars, ISM, CGM) and the corresponding mass flow rates into and out of the ISM and CGM. We did not change anything in the SAM (which has been shown capable of matching many observations at $z = 0$ and higher redshift) except to turn off AGN feedback since that is not included in the FIRE-2 simulations we use. Our sample spans 13 halos grouped into three mass bins, with at least 3 halos per bin: low-mass dwarfs ($M_{\text{vir}} \sim 10^{10} M_{\odot}$ at $z = 0$), intermediate-mass dwarfs ($M_{\text{vir}} \sim 10^{11} M_{\odot}$), and MW-mass galaxies ($M_{\text{vir}} \sim 10^{12} M_{\odot}$). We also

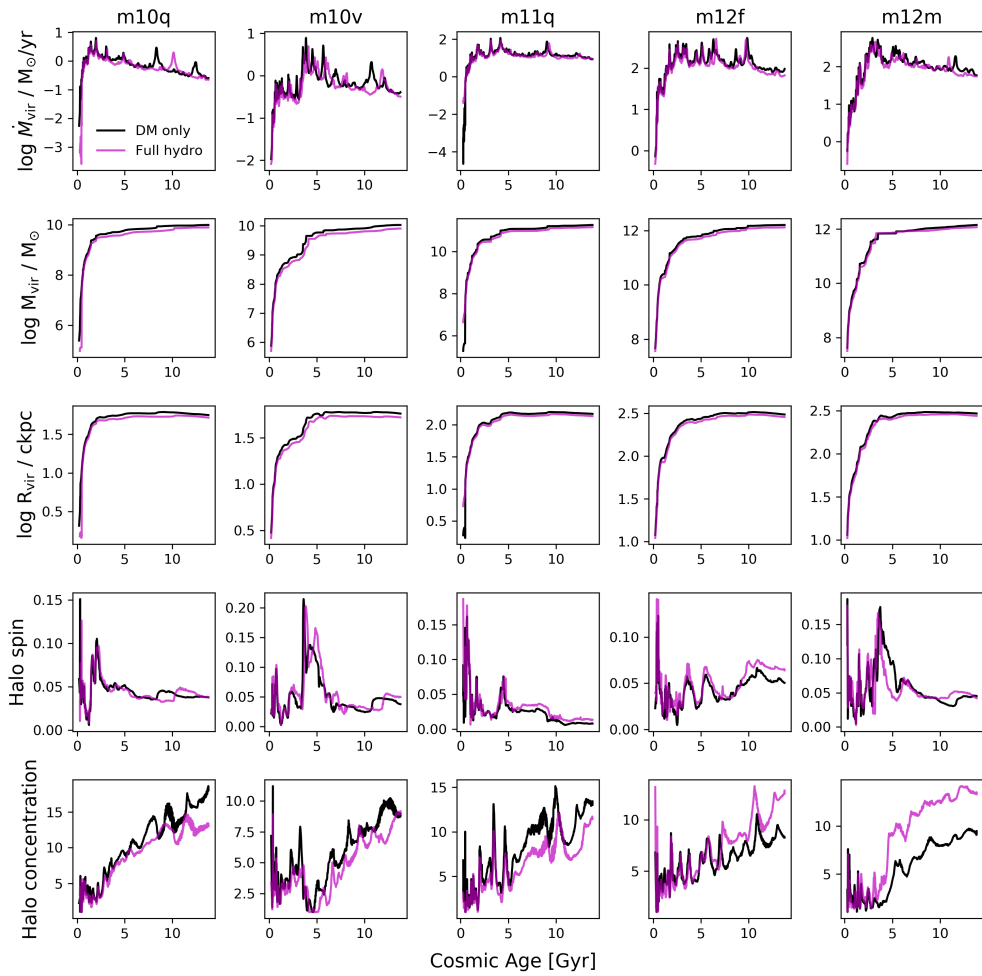


Figure 2.15: A few relevant halo properties measured in the full hydrodynamical simulations (magenta lines) and the corresponding DM-only simulations (black lines). These are the 5 halos for which corresponding DM-only FIRE-2 runs exist. The DM-only halo properties are very similar to the hydro-based halo properties, with the M_{vir} and R_{vir} being lower in the hydro version by only 10-20% on average. The major systematic difference is in the halo concentration which tends to be lower in the dwarfs in the hydro version (presumably due to stellar feedback adiabatically expanding the halo center) whereas it is higher for the MW-mass halos in the hydro run (presumably due to the significant stellar mass adiabatically contracting the halo center).

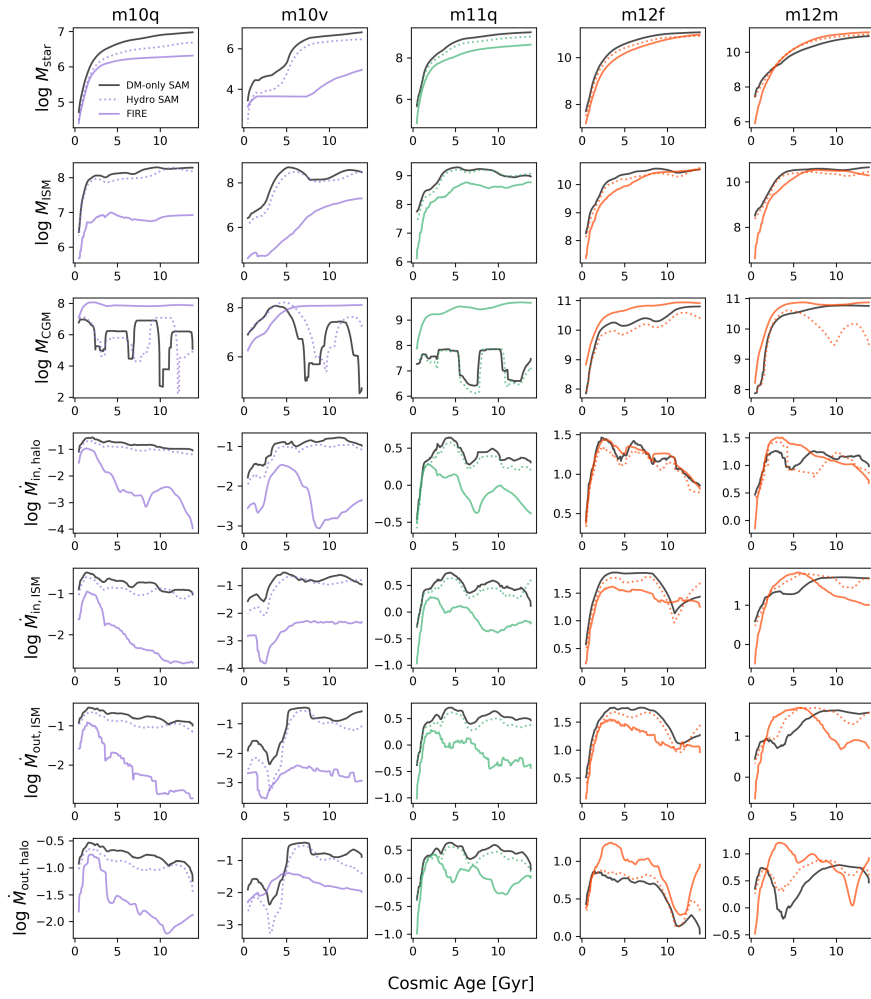


Figure 2.16: Time series of the main properties considered for the five individual halos with corresponding DM-only simulations. From left to right: m10q, m10v, m11q, m12f and m12m. From top to bottom: stellar mass, ISM mass, CGM mass, halo gas mass inflow rate, ISM mass inflow rate, ISM mass outflow rate and halo mass outflow rate. In every panel, the solid black line shows the prediction of the SAM when run on the DM-only simulation merger trees. The other two curves follow the same convention as the individual halo panels in the figures from the main body of the text: solid colored curves for the FIRE-2 measurements and dotted colored curves for the SAM predictions using the full hydrodynamical simulation merger trees. The colors show the mass bin that each halo belongs to (purple for m10, green for m11 and red for m12). The main takeaway is that our conclusions do not change if we use the SAM results from the DM-only simulation merger trees: the new SAM predictions agree with our fiducial ones relatively well, and hence the DM-only-based SAM trends relative to FIRE-2 remain qualitatively the same.

presented a simple physical model for how preventative stellar feedback can suppress halo gas accretion on the scale of R_{vir} preferentially for dwarfs.

Our main takeaways are as follows:

1. At $z = 0$, the SAM agrees relatively well with FIRE-2 and empirical constraints on the stellar-to-halo mass relation. The SAM and FIRE-2 also agree relatively well with each other and with observations for the ISM-to-stellar mass ratio at $z = 0$ (as a function of stellar mass). However, they disagree dramatically with each other in terms of CGM mass: the CGM mass of dwarfs is $\sim 3 - 4$ orders of magnitude lower in the SAM than in FIRE-2. This reflects the flexibility allowed in galaxy formation models to match observations of stars and the ISM while at the same time disagreeing greatly on the total CGM mass (owing partially to the observational uncertainty about whether most extragalactic gas resides within or outside of halos).
2. As a function of time, the SAM reproduces the stellar mass assembly histories of the FIRE-2 galaxies generally within a factor of two (with the exception of one late-forming dwarf m10v). However, despite the overall agreement on the stellar mass assembly history, the two models disagree on the star formation history on shorter timescales of ~ 100 Myr. The SAM does not demonstrate stochasticity in its SFHs whereas it is an ubiquitous phenomenon in the FIRE-2 dwarfs at all times and in the FIRE-2 MW-mass halo progenitors at early times (as also shown by Muratov et al., 2015; Anglés-Alcázar et al., 2017a; Ma et al., 2017; Sparre et al.,

2017; Faucher-Giguère, 2018).

3. The time series of ISM mass agrees relatively well between the SAM and FIRE-2, although the SAM tends to be higher than FIRE-2 for the low-mass dwarfs. The CGM mass discrepancy between the SAM and FIRE-2 at $z = 0$ (at fixed mass) also extends over all time. The mass of the “ejected” gas reservoir in the SAM dominates over the CGM mass at all times, even in the MW-mass halos; this previously ejected gas is assumed to re-accrete back into the CGM on roughly a Hubble time in the SAM. Despite these dramatic differences in the individual bulk components, the halo baryon fractions tend to agree within a factor of $\sim 2 - 3$ at all times, and both the SAM and FIRE-2 show the same qualitative trend: lower mass halos are more depleted of baryons than higher mass halos, relative to the universal baryon fraction.
4. Comparing the mass flow rates as a function of time gives clues to the discrepancies in the integrated masses. The fundamental mismatch is that the SAM has significantly higher halo gas accretion rates compared to FIRE-2, with the discrepancy being systematically worse for dwarfs by a factor of $\sim 10 - 100$. We argue that this is due to a combination of high re-accretion rates of gas previously ejected from the halo and the lack of preventative stellar feedback to suppress pristine halo gas accretion. The ISM accretion rates are also higher in the SAM than in FIRE-2, owing primarily to the halo gas cooling model and lack of preventative stellar feedback in the SAM. Correspondingly, to match the stellar assembly

histories and the observed $z = 0$ stellar mass function, the SAM has higher mass outflow rates than FIRE-2 from both the ISM and the halo. But, the low-mass dwarfs in FIRE-2 have cumulatively ejected more mass from their halo than has ever left their ISM (between a factor of ~ 1.5 up to ~ 10 by $z = 0$; even larger ratios are measured for the progenitors of all halos at very early times $z \gtrsim 6$). This implies significant entrainment of ambient CGM material and may have important preventative feedback effects that are not currently allowed for in the SAM by construction.

5. We propose a simple physical model for how stellar-driven winds can suppress halo gas accretion on the scale of R_{vir} for dwarfs. The essence of the model is that SN-driven winds can shock heat some fraction of gas beyond R_{vir} to the virial temperature or higher such that it can no longer accrete into the halo. We show that this simple model is capable of reproducing the reduced halo gas accretion efficiencies of the FIRE-2 dwarfs remarkably well, provided that the energy loading factor at R_{vir} is preferentially higher for dwarfs. Characterizing the mass and energy loading factors from the simulations in the future will help test and calibrate our preventative stellar feedback model.

Given all of the model discrepancies and potential improvements discussed herein, it is natural to ask whether a SAM that is calibrated to match FIRE-2 (or any zoom-in simulation suite) can also still match observations. This is one of our ultimate driving questions, but our work demonstrates that the overall foundational structure and

perhaps philosophy underlying SAMs may first need to be updated. Explicitly adding preventative feedback is arguably the most crucial step because the current Santa Cruz SAM does not contain the relevant physics to capture the low halo gas accretion rates of FIRE-2 dwarfs. The apparent success of our new preventative stellar feedback model suggests a path forward, but automated parameter space exploration techniques will be needed to map out degeneracies with existing SAM assumptions. Beyond that, we will need to improve (among other things) how we model halo gas cooling, the multi-phase structure of the CGM, the stochastic nature of star formation and associated outflows, and implement new channels for halo gas ejection and recycling. In parallel, it will be important to consider the statistical challenges associated with calibrating a SAM using the relatively small sample size of halos that can be provided by modern zoom-in simulation suites, and then scaling up predictions to the level of galaxy populations. In particular, it is presently unclear if the diversity in halo growth histories at a fixed halo mass is enough on its own to reproduce the scatter in galaxy properties at a fixed stellar mass, or if there is additional scatter on the baryonic physics side (e.g., from smaller sub-galactic scales) that needs to be modeled. These and related issues will be the subject of our future work.

Chapter 3

Characterizing simulated multi-phase supernova-driven winds

3.1 Introduction

In this chapter, we build on the analysis of winds in the FIRE-1 zoom-in simulations (Hopkins et al., 2014) by Muratov et al. (2015, 2017); Anglés-Alcázar et al. (2017a); Hafen et al. (2019, 2020). Whereas Muratov et al. (2015, 2017) focused on instantaneous mass and metal loadings in FIRE-1 and their scalings with stellar mass, halo mass and halo circular velocity, we use a suite of simulations run using the updated FIRE-2 code (Hopkins et al., 2018a). The FIRE-2 simulations model the same stellar processes as the FIRE-1 simulations but use a more accurate hydrodynamic solver. Thus, we expect many of the overall predictions to be similar between FIRE-1 and FIRE-2. Motivated by analysis procedures for small-scale idealized simulations, we

implement a sophisticated method for identifying galactic winds by considering their bulk kinetic, thermal and potential energies. Instead of focusing only on the total mass and metal loading factors as is common practice, our multi-dimensional analysis focuses on the temperature dependence of all four loading factors (mass, momentum, energy and metals) and how this varies as a function of galaxy mass. With the FIRE-2 simulation suite, we will comment on the nature of SN-driven galactic winds across a wide range of halo masses (low-mass dwarfs, intermediate-mass dwarfs, MW halos and their high-redshift dwarf progenitors, and more massive halos at high redshift). We also present scaling relations for the loading factors not just with the global halo circular velocity (as is commonly done), but also with several “quasi-local” ISM properties as a first step toward connecting the larger-scale emergent loadings with the smaller-scale conditions of the ISM in which the winds are launched.

The results described herein were presented in Pandya et al. (2021) and will be published soon after this dissertation. This chapter was directly motivated by the work in the preceding chapter (published as Pandya et al., 2020) wherein we argued that the lack of preventative stellar feedback in SAMs leads to much higher mass inflow and outflow rates than measured in FIRE. By measuring the energy and momentum of SN-driven winds, we will be extracting crucial data from the simulations that will later inform our modeling of preventative stellar feedback in next-generation SAMs (Chapter 4).

This chapter is organized as follows. Section 3.2 describes the FIRE-2 simulations and Section 3.3 details our analysis methods. In Section 3.4, we present wind

loading factors near the ISM and in Section 3.5 we describe results for winds leaving the halo at R_{vir} . We discuss our results in Section 3.6 and summarize in Section 3.7. We assume a standard flat Λ CDM cosmology consistent with the FIRE-2 code and Planck Collaboration et al. (2014); i.e., $h = 0.7$, $\Omega_M = 0.27$, $\Omega_\Lambda = 0.73$ and $f_b \equiv \Omega_b/\Omega_M \approx 0.16$.

3.2 Simulation Description

We use a suite of cosmological “zoom-in” simulations run using the FIRE-2 code (Hopkins et al., 2018a). Our analysis focuses on a “core” suite of 13 FIRE-2 halos: 4 low-mass dwarfs with $M_{\text{vir}} \sim 10^{10} M_\odot$ at $z = 0$ (m10q, m10v, m10y, m10z), 6 intermediate-mass dwarfs with $M_{\text{vir}} \sim 10^{11} M_\odot$ by $z = 0$ (m11a, m11b, m11q, m11c, m11v, m11f), and 3 MW-mass halos with $M_{\text{vir}} \sim 10^{12} M_\odot$ by $z = 0$ (m12i, m12f, m12m). These halos were first presented in Wetzel et al. (2016); Garrison-Kimmel et al. (2017); Chan et al. (2018); Hopkins et al. (2018a). To this core suite, we also add the four FIRE-2 massive halos (A1, A2, A4, and A8 with $M_{\text{vir}} \sim 10^{12.5} - 10^{13} M_\odot$ at $z = 1$) presented by Anglés-Alcázar et al. (2017b) and further analyzed in Cochrane et al. (2019); Wellons et al. (2020); Stern et al. (2020). These halos are denoted as “m13” throughout the chapter, were only run down to $z = 1$, and were previously simulated with the FIRE-1 model as part of the MassiveFIRE suite (Feldmann et al., 2016, 2017). While the m10, m11 and m12 halos agree well with empirical stellar-to-halo-mass relations, the m13 halos have unrealistically high stellar masses and central densities by $z = 1$ (e.g., Parsotan et al., 2021), and hence should not be taken as representative of the observed

population (this is a regime where feedback from supermassive black holes may have an appreciable effect but this is not included in these simulations).

We refer the reader to Hopkins et al. (2018a) for a detailed description of the simulations and methodology. Here we only briefly review the most relevant aspects, with a particular emphasis on the explicit stellar feedback model. The core FIRE-2 simulations model the same physical processes as in FIRE-1 but use the more accurate Lagrangian “meshless finite-mass” hydrodynamic solver as opposed to the “pressure–entropy” formulation of smoothed particle hydrodynamics (Hopkins, 2015). The FIRE-2 simulations implement a broad range of physics, including deposition of mass, momentum, energy and metals due to both Type Ia and Type II SNe, stellar winds, radiation pressure, and photo-ionization and photo-electric heating. There is a spatially-uniform but redshift-dependent UV background based on Faucher-Giguère et al. (2009).

The relatively high resolution of the FIRE-2 simulations (Lagrangian particle masses of $\sim 250M_{\odot}$ in the low-mass dwarfs, up to $\sim 7100M_{\odot}$ for the MW halos and $\sim 33000M_{\odot}$ for the m13 runs) allows stellar feedback to be modeled locally and explicitly. In particular, the generation and propagation of winds is not explicitly dependent on global halo properties and does not require subgrid approaches of limited predictive power (e.g., hydrodynamically-decoupled winds, shut-off of cooling, thermal bombs). Of course, not all SN remnants will be resolved, especially in the more massive halos which have comparatively worse resolution. As detailed in Hopkins et al. (2018b), this is “corrected” for in FIRE-2 by depositing onto nearby gas particles the additional momentum expected from the unresolved energy-conserving Sedov-Taylor phase (due

to PdV work). The thermal energy output by the unresolved SN remnant is also self-consistently reduced to account for radiative cooling after the energy-conserving phase. In cases where the SN remnant is resolved, the FIRE-2 subgrid model explicitly deposits the thermal energy expected from the energy-conserving phase, and allows the hydrodynamic solver to explicitly calculate the heating and momentum generation (PdV work). Note that while some small-scale simulations suggest that a resolution of $\lesssim 100M_{\odot}$ may be necessary to properly capture the evolution of SN remnants (e.g., Kim & Ostriker, 2015; Steinwandel et al., 2020), the combination of multiple stellar feedback effects (e.g., early radiative feedback) with self-consistent clustering of star formation in FIRE-2 may act to alleviate this resolution requirement. Hopkins et al. (2018b, Figure 9) showed that the FIRE subgrid model remains converged to the high-resolution result up to resolutions of $2000M_{\odot}$ for an m10 halo (see also Wheeler et al., 2019, who re-simulated a few FIRE-2 dwarfs with $30M_{\odot}$ resolution).

As our work builds on the analysis of FIRE-2 presented in Pandya et al. (2020), here we use the same halo catalogs and merger trees generated using the Rockstar and consistent-trees codes (Behroozi et al., 2013b,c). For halo masses and radii, we adopt the Bryan & Norman (1998) virial overdensity definition. We only focus on the main central halo in each of these simulations and do not analyze satellites.

3.3 Analysis

In this section, we describe how we select outflowing gas, define multi-phase outflows, and compute loading factors.

3.3.1 Accurately defining outflows

3.3.1.1 Selecting outflowing particles

It is common practice to define outflows in cosmological simulations using a single cut on the halo-centric radial velocity of particles (regardless of using the shell/Eulerian or particle-tracking/Lagrangian methods). The simplest cut often adopted is $v_{\text{rad}} > 0$ km/s which would select all particles that are traveling radially away from the halo center (as done by, e.g., Faucher-Giguère et al., 2011; Muratov et al., 2015). This can confuse slow random motions with galactic outflows. The other extreme is to select only particles at a given radius whose $v_{\text{rad}} > v_{\text{esc}}(r)$ where $v_{\text{esc}}(r)$ is the local escape velocity at that radius. This cut is often used to define the subset of fastest moving “wind” particles among the whole distribution of outflowing particles. There are variations on this radial velocity cut method in the literature: Muratov et al. (2015) use the velocity dispersion of the underlying virialized DM halo particles, Mitchell et al. (2020) use $0.25V_{\text{max}}$ where V_{max} is the maximum circular velocity of the halo, and Nelson et al. (2019) compute the cumulative mass fraction of outflowing particles with radial velocities above sequentially increasing velocity thresholds.

However, using a single cut on v_{rad} alone is sub-optimal for defining winds for

the following reason.¹ Consider that every gas particle possesses three forms of energy: kinetic, thermal and potential energy. A single cut on v_{rad} alone assumes the extreme case of “ballistic motion” in which thermal energy is not a source of acceleration. But since we are dealing with gas, we must account for the fact that: (1) the transverse velocities also carry kinetic energy, and (2) the thermal energy of gas particles can serve as a source of acceleration assuming adiabatic expansion, i.e., no external heating/cooling/interactions. This has long been realized in the literature for small-scale resolved ISM/CGM simulations (e.g., Martizzi et al., 2016; Kim et al., 2020a; Schneider et al., 2020) but has not been fully leveraged for cosmological simulations (though see Hopkins et al., 2012).

Here we introduce a slightly more sophisticated methodology to accurately define outflowing particles. First, we make a simple cut on $v_{\text{rad}} > 0$ km/s. This selects all particles that are flowing radially outwards. However, a large fraction of these particles may have relatively small radial velocities arising from underlying random velocity fluctuations. We only want to select particles that will be able to travel a significant distance. Hence, for every gas particle, we calculate the radial component of the total Bernoulli velocity $v_{\text{B,total}}$, which is a measure of the total specific energy (e.g., Hopkins et al., 2012; Li et al., 2017; Kim & Ostriker, 2018; Fielding et al., 2018):

$$v_{\text{B,total}}^2 \equiv \frac{1}{2}v_r^2 + \frac{c_s^2}{\gamma - 1} - \frac{1}{2}v_{\text{esc}}^2. \quad (3.1)$$

¹From an observational perspective, the simplest $v_{\text{rad}} > 0$ km/s cut might be justified and desirable (especially if detailed kinematics, phase information and gravitational potential constraints are unavailable), but here we are interested in robustly identifying and characterizing winds from a simulation perspective.

The first term is the specific radial kinetic energy quantified by the halo-centric particle radial velocity squared. The second term is the specific enthalpy assuming an ideal gas whose equation of state has adiabatic index γ and sound speed $c_s = \sqrt{\gamma \frac{kT}{\mu m_p}}$. We assume a monatomic ideal gas, hence $\gamma = \frac{5}{3}$.

The third term is equivalent to the specific gravitational potential energy, Φ . The simulation code internally keeps track of Φ for each particle to compute its gravitational acceleration, but unfortunately Φ is not one of the properties output in the particle snapshot files. Computing Φ in post-processing is tricky because the mass distribution is heterogeneous and that can disproportionately affect the potential for some particles, even if they have the same halo-centric distance. However, for simplicity, we assume the mass distribution can be approximated as spherically symmetric, which allows us to relate the potential to the enclosed mass profile in a simple way:

$$\Phi(r) = - \int_r^{r_\infty} \frac{GM(< r)}{r^2} dr \quad (3.2)$$

where r_∞ is an arbitrarily large radius. It is important to include the contribution of mass out to several virial radii since that can affect the Φ profile within the halo. Given that we are working with cosmological zoom-in simulations, we adopt the following strategy. Within $2R_{\text{vir}}$, the enclosed mass profile is based on all star, gas and high-resolution DM particles using spherical shells of width $0.01R_{\text{vir}}$. From $2 - 10R_{\text{vir}}$, we “stitch on” the enclosed mass profile accounting for only high-resolution and low-resolution DM particles using shell widths of $0.1R_{\text{vir}}$; baryons would have contributed only $\sim 15\%$ by mass so we safely ignore those at these large distances.² Finally, in

²The masses of high-resolution DM particles are five times larger than those of baryonic particles.

logarithmically spaced bins from $10 - 100R_{\text{vir}}$ (far beyond the zoom region), we again stitch on the enclosed mass profile using only low-resolution DM particles to capture the large-scale DM matter field. Note that $100R_{\text{vir}}$ is only ~ 5 Mpc for the m10 dwarfs but ~ 30 Mpc for the MW halos at $z = 0$.

We re-write Φ as an escape velocity using the energy conservation equation and assuming the gas particle at r is already maximally cold (i.e., ignore any changes in enthalpy):

$$\frac{1}{2}v_{\text{esc}}^2(r) + \Phi(r) = \Phi(r_{\text{zp}}) . \quad (3.3)$$

We set the zeropoint of the potential at $r_{\text{zp}} = 2R_{\text{vir}}$ since that is the turnaround radius for a virialized system and particles traveling beyond $2R_{\text{vir}}$ are likely unbound from the halo anyway (also, our zoom regions can start to become contaminated by low-resolution DM beyond $2R_{\text{vir}}$). In this way, we can derive the radial profile of escape velocity, which tells us how fast a particle must be going initially (at minimum) to fully climb out of the halo potential:

$$v_{\text{esc}}(r) = \sqrt{2(\Phi(2R_{\text{vir}}) - \Phi(r))} . \quad (3.4)$$

The quantity $v_{\text{B,total}}^2$ represents the radial component of the total specific energy of a gas particle at its current position. Note that $v_{\text{B,total}}^2$ can be negative, which means that a particle is bound (i.e., its kinetic energy plus enthalpy is less than its potential energy). By comparing this initial Bernoulli velocity to a hypothetical final Bernoulli velocity at some other larger halo-centric distance, we can assess whether a

The masses of low-resolution DM particles are larger by another factor of eight immediately outside of the zoom region, and continue to increase by additional factors of eight progressively farther from the zoom region (see footnote 37 of Hopkins et al., 2014).

given gas particle has enough starting energy to make it to that larger distance (neglecting interactions). For a particle to be able to travel from its current radius r_1 to some secondary radius r_2 , its initial Bernoulli velocity must be larger than the potential energy at that secondary radius.³ We use this to impose an additional criterion that selects only gas with sufficiently large $v_{\text{B,total}}$ relative to the escape velocity at some target distance (defined in the next section):

$$v_{\text{B,total}}^2(r_1) > -\frac{1}{2}v_{\text{esc}}^2(r_2). \quad (3.5)$$

This criterion along with $v_{\text{rad}} > 0$ km/s is a more physically meaningful and robust way to select wind particles compared to either $v_{\text{rad}} > 0$ km/s or $v_{\text{rad}} > v_{\text{esc}}(r)$ alone. It is effectively an intermediate case that avoids the very slow moving turbulent motions while still selecting the hotter and slower components of the wind. This definition is also a natural way to quantitatively distinguish between genuinely escaping winds and outflows expected to remain bound out to some larger radius.

3.3.1.2 Computing outflow fluxes

We compute outflow fluxes in two characteristic spherical shells:

1. $0.1 - 0.2R_{\text{vir}}$ (ISM boundary shell)
2. $1.0 - 1.1R_{\text{vir}}$ (virial boundary shell)

³This neglects the effect of heating by the UV background that prevents gas from cooling to arbitrarily low temperature. Thus, in principle for the secondary radius we should add the $\frac{c_s^2}{\gamma-1}$ term assuming the sound speed for gas in thermal equilibrium with the UV background at $\sim 10^4$ K, roughly 15 km/s. In practice, this makes a negligible difference for outflow selection (most of the gas tends to be escaping in low-mass halos anyway, and for MW-mass halos this 10^4 K gas sound speed term is an order of magnitude lower than the escape velocity term).

In each of these two shells, we must select particles that have enough energy to make it to some secondary radius, r_2 , if not farther (assuming an adiabatic flow). There is inevitably a large range of arbitrary choices that could be made for r_2 . For the ISM shell, we adopt a secondary radius of $r_2 = 0.5R_{\text{vir}}$, which we take to represent the “middle” of the CGM. Choosing a smaller target radius would pick up additional cooler/slower outflows, but we note that our ISM shell is already quite far out ($0.1 - 0.2R_{\text{vir}}$). For the virial shell, we adopt a secondary radius of $r_2 = 2.0R_{\text{vir}}$. This lets us select particles at $1.0 - 1.1R_{\text{vir}}$ that have at least enough energy to make it to $2.0R_{\text{vir}}$, if not farther. Since particles can be considered unbound if they travel beyond the turnaround radius of $2R_{\text{vir}}$, this is a natural way to estimate genuinely escaping outflows from the halo. Finally, since we will compare the halo outflow rate to the preceding ISM outflow rate, we also define a second more restrictive ISM outflow criterion by choosing $r_2 = 2R_{\text{vir}}$. This lets us additionally estimate the subset of ISM outflows that have enough energy to get not just to $0.5R_{\text{vir}}$ but rather escape to $2R_{\text{vir}}$ or beyond.

Finally, with outflowing particles identified for each of the two shells above, we compute their total mass, momentum, energy and metal mass outflow rates as follows:

$$\dot{M}_{\text{out}} = \sum_i \frac{m_i v_{r,i}}{\Delta L} \quad (3.6)$$

$$\dot{p}_{\text{out}} = \sum_i \dot{M}_{\text{out},i} v_{r,i} \left(1 + \frac{1}{\gamma \mathcal{M}_i^2} \right) \quad (3.7)$$

$$\dot{E}_{\text{out}} = \sum_i \dot{M}_{\text{out},i} v_{\text{B},i}^2 \quad (3.8)$$

$$\dot{M}_{\text{Z,out}} = \sum_i \dot{M}_{\text{out},i} Z_i \quad (3.9)$$

Here, the subscript i runs over all the selected outflowing particles in the shell, $\Delta L = 0.1R_{\text{vir}}$ is the width of our ISM and virial shells, v_r is the radial velocity, $\mathcal{M} \equiv v_r/c_s$ is the Mach number, and Z is the metal mass fraction of the particle. Note that the second term in the momentum flux accounts for the thermal pressure component (defined as $P = \rho c_s^2/\gamma$), which can be substantial for hot outflows or more generally when \mathcal{M} is small. v_B is the Bernoulli velocity neglecting the gravitational term and including the transverse kinetic energy component (as opposed to $v_{B,\text{total}}$ in Equation 3.1):

$$v_B^2 = \frac{1}{2}v^2 + \frac{3}{2}c_s^2, \quad (3.10)$$

where v is the magnitude of the total halo-centric particle velocity vector instead of just v_r . We neglect the gravitational term for \dot{E}_{out} because we want to quantify how much specific kinetic energy and enthalpy are being transported by outflows (these quantities, including the transverse velocity components, will be responsible for any heating and pushing of ambient gas). The gravitational term comes in earlier when we first want to identify escaping and bound outflows.

Figure 3.1 and Figure 3.2 respectively show examples of strong outflows in a MW halo at $z \sim 0$ and a dwarf halo at $z \sim 3$. The phase diagram of temperature versus radial velocity shows that our Bernoulli velocity wind criterion successfully captures the slower but hot wind component, which would otherwise be missed by requiring simply $v_{\text{rad}} > v_{\text{esc}}$. At the same time, our method excludes the slow and cold component of outflows, which is likely driven by random motions of gas particles and would otherwise complicate the interpretation of mass outflow rates (much of this slower gas may recycle

back into the ISM via fountain flows). There is generally a time lag between peaks in the star formation history (SFH) and subsequent spikes in the mass outflow rate time series. As outflows propagate from the inner halo to the outer halo, they can either deposit or sweep up mass in the CGM. This can be inferred qualitatively from the time evolution of the radial profile of \dot{M}_{out} since the amplitude and width of individual outflow spikes may change as they move to larger radius.

3.3.2 Multi-phase outflow selection criteria

It is important to distinguish between outflows of different temperatures since that can clarify whether the driving energy source is kinetic or thermal. The simplest way to do this is based on atomic cooling physics. We can divide the temperature distribution into roughly three phases:

1. $T < 10^3 K$ (cold outflows)
2. $10^3 < T < 10^5 K$ (warm outflows)
3. $T > 10^5 K$ (hot outflows).

These temperature cuts correspond to physically distinct regimes.⁴ The cut at 10^5 K corresponds to the peak of the cooling curve, so material is expected to separate naturally about this temperature. Likewise, the cut at 10^3 K corresponds to the unstable part of the cooling curve at the usual pressures and photoelectric heating rates found in the ISM/inner CGM, so gas is also expected to naturally separate about this

⁴Our warm gas is also termed cold gas in some CGM studies since it is much colder than the virial temperature of MW-mass halos.

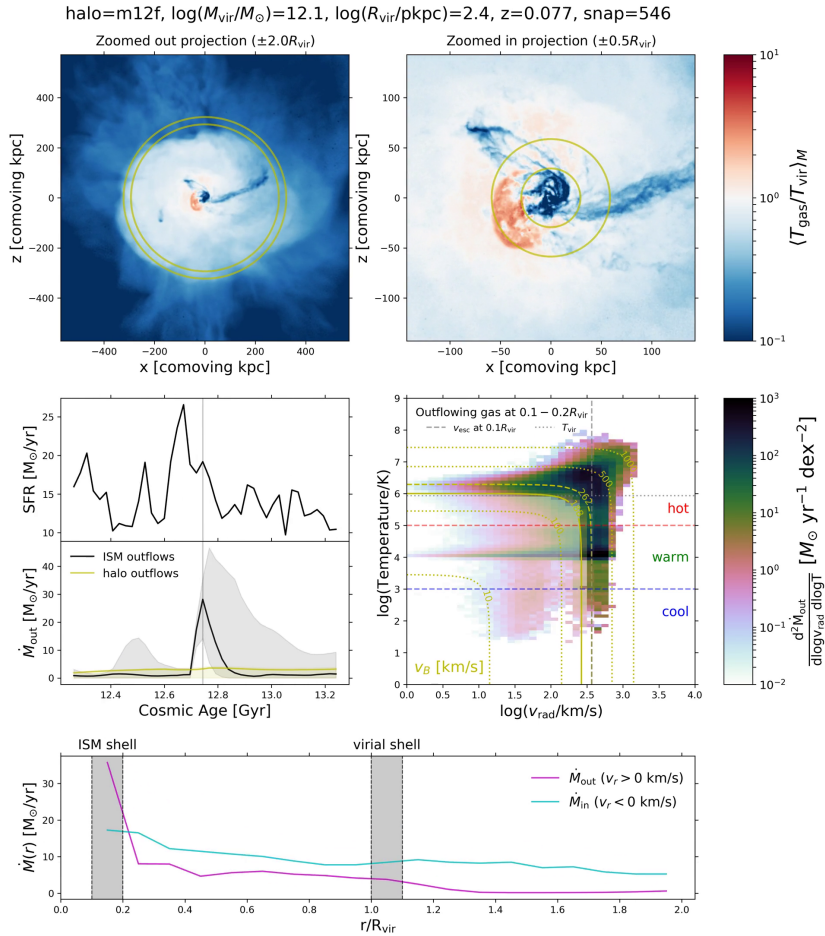


Figure 3.1: Visualizing and quantifying a strong outflow at $z \sim 0$ in a MW-mass halo (m12f). This is a single frame from a movie that is available for download. *Top-left*: Zoomed-out projection ($\pm 2R_{\text{vir}}$) of the mass-weighted average gas temperature. The colorbar has been normalized by the halo virial temperature. The two white circles demarcate our virial shell ($1.0 - 1.1R_{\text{vir}}$). *Top-right*: Similar to the top-left panel but now a zoomed-in projection ($\pm 0.5R_{\text{vir}}$). The two white circles mark our ISM shell ($0.1 - 0.2R_{\text{vir}}$). *Middle-left*: Time series of the SFR (top panel) and the mass outflow rate in the ISM and virial shells (bottom panel). The lines in the lower part of this panel show mass outflow rate measurements based on our fiducial Bernoulli velocity criterion. The shaded regions show how more extreme cuts would lead to different estimates: $v_{\text{rad}} > 0 \text{ km/s}$ gives an upper bound to the mass outflow rate whereas $v_{\text{rad}} > v_{\text{esc}}$ picks up only the fastest material and hence leads to a lower bound. The vertical gray line marks the current snapshot time. *Middle-right*: Phase diagram of temperature and radial velocity for the multi-phase ISM outflows identified using our Bernoulli velocity method. The colorbar shows the mass outflow rate in logarithmic bins of temperature and radial velocity. The horizontal red and blue lines demarcate our cool, warm and hot outflow temperature regimes. The horizontal dotted gray line indicates the halo virial temperature (computed at R_{vir}) and the vertical gray dashed line denotes v_{esc} at $0.1R_{\text{vir}}$. Yellow dotted contours show lines of constant Bernoulli velocity, with the potential difference between $0.1R_{\text{vir}}$ and either $0.5R_{\text{vir}}$ or $2R_{\text{vir}}$ shown as the solid and dashed yellow contours, respectively. The transparent histogram below the solid yellow contour shows what is excluded from our v_B cut. Selecting only outflows with $v_{\text{rad}} > v_{\text{esc}}$ would miss the slower but still hot wind component, which our Bernoulli velocity method successfully captures. *Bottom*: Radial profile of instantaneous mass flux for both outflows (magenta) and inflows (cyan) between $0.1 - 2R_{\text{vir}}$ in spherical shells of width $0.1R_{\text{vir}}$. For simplicity, we use $v_{\text{rad}} = 0 \text{ km/s}$ as the dividing point between outflows and inflows. Our ISM and virial shells are marked as the gray bands. This panel can be used to follow the radial evolution of individual outflow episodes and qualitatively infer CGM entrainment or wind mass losses. This movie and others in their entirety are available for download in the online journal.

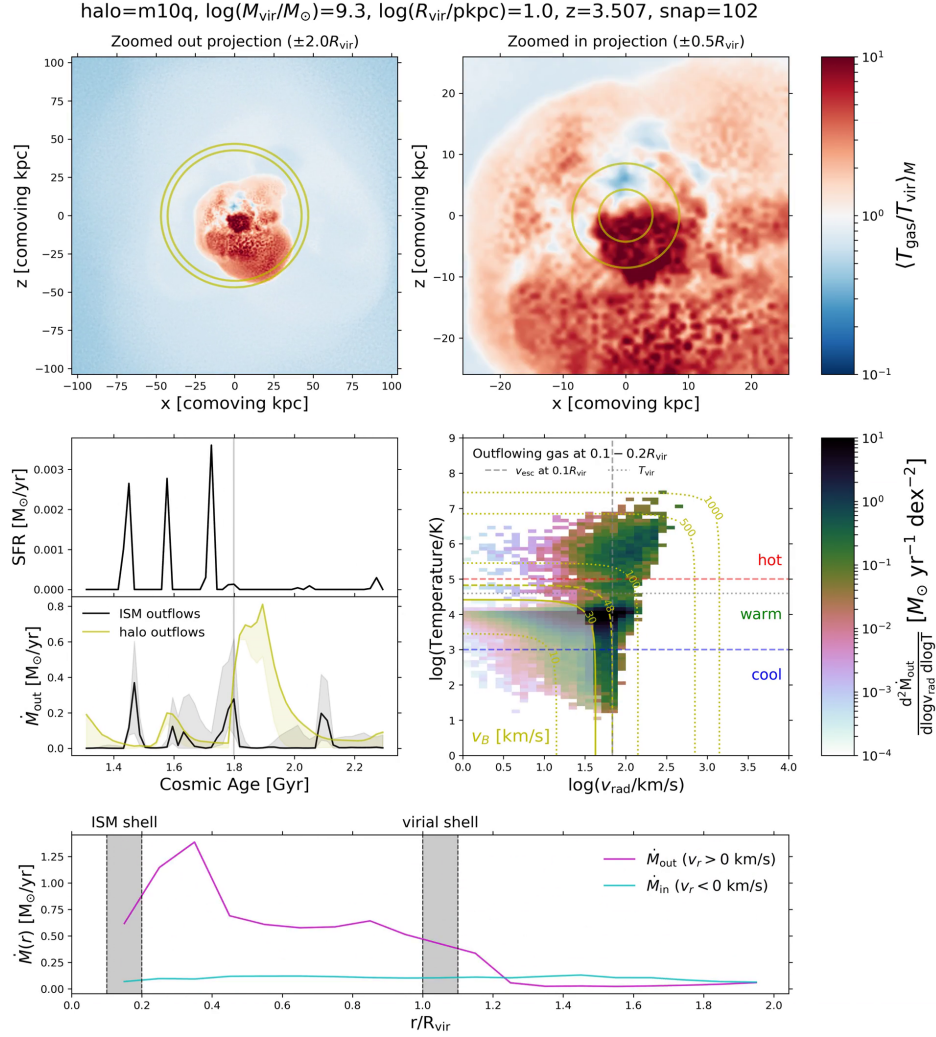


Figure 3.2: Similar to Figure 3.1 but now for a low-mass dwarf (m10q). This movie frame is during a major outflow episode at high redshift $z \sim 3.5$, a regime where dwarfs are often characterized as having mass loadings of ~ 100 or more. If we divide the values of the individual ISM mass outflow rate peaks by their associated preceding SFR spikes (bottom-left panel), we would indeed infer instantaneous mass loadings of ~ 100 . The halo-scale mass loadings (magenta) are even larger due to entrainment of CGM gas by outflows. Note how there is a hot bubble in the projection panels created by the strong outflows. This movie and others in their entirety are available for download in the online journal.

temperature. Lastly, a significant amount of gas can be expected to have $T \sim 10^4$ K since that is roughly the equilibrium temperature between photoionization from the UV background and recombination cooling. These cuts, therefore, mirror the delineations that are expected to arise naturally in and around galaxies. These temperature bins also trace what observers can measure: the cold phase corresponds to molecular/atomic outflows, the warm phase traces partially ionized gas that will produce H α emission and absorption from singly and doubly ionized metals, and the hot phase traces highly ionized gas that produces X-ray emission.

In Figure 3.3, we plot the average temperature distribution of the mass outflow rate of our halos through the ISM shell. The distribution is averaged over three broad redshift bins using the $\dot{M}_{\text{out,ISM}}$ in each snapshot as the weight. We see that outflows in our simulations are inherently multi-phase, except in the two lowest mass halos. The cold phase is more pronounced at higher redshift. The peak in the warm regime at $\sim 10^4$ K likely reflects the equilibrium temperature between heating and cooling, and the broad peak in the hot regime corresponds to the virial temperature in the inner halo, computed as

$$T_{\text{vir}} = 35.9 \frac{V_{\text{circ}}^2}{\text{km s}^{-1}} \text{ K} \quad (3.11)$$

where V_{circ} is the circular velocity at $0.1R_{\text{vir}}$ (as opposed to the common practice of using V_{circ} at the virial radius). The virial temperatures of the lowest mass halos are themselves below 10^5 K, so there is no pronounced peak in their hot outflow rates. The two lowest mass dwarfs, in particular, show a cut-off in their cold outflow rates at $\sim 10^4$ K. We will see later that this means the warm phase is remarkably important for

outflows in dwarfs.

3.3.3 Computing wind loading factors

3.3.3.1 Reference fluxes

Lastly, it is useful to compare the wind fluxes to reference fluxes at the ISM scale. By dividing the two, we can estimate the loading factor η and get a sense of how much mass, energy, momentum and metal mass is being ejected versus what was input from star formation and SNe. Computing the reference fluxes is non-trivial in cosmological simulations because of the wide range of processes that are simultaneously at play. We therefore limit ourselves to considering Type II SNe (we expect these to dominate over Type Ia SNe, radiative heating and mass loss from normal stellar evolution, and other processes, but see our discussion of caveats in subsection 3.6.2). In line with Kim et al. (2020a), we adopt the following reference fluxes:

1. $\dot{M}_{\text{ref}} = \text{SFR}$
2. $\dot{p}_{\text{ref}} = \dot{N}_{\text{SN}} \frac{E_{\text{SN}}}{v_{\text{cool}}} = \frac{\text{SFR}}{100M_{\odot}} \frac{E_{\text{SN}}}{v_{\text{cool}}}$
3. $\dot{E}_{\text{ref}} = \dot{N}_{\text{SN}} E_{\text{SN}} = \frac{\text{SFR}}{100M_{\odot}} E_{\text{SN}}$
4. $\dot{M}_{\text{Z,ref}} = \dot{N}_{\text{SN}} M_{\text{ej}} Z_{\text{SN}} = \frac{\text{SFR}}{100M_{\odot}} M_{\text{ej}} Z_{\text{SN}}$

Here, the total instantaneous galaxy SFR is computed by summing over the individual SFRs predicted by all gas particles⁵ within $0.1R_{\text{vir}}$. Then $\dot{N}_{\text{SN}} = \frac{\text{SFR}}{100M_{\odot}}$ is the supernova

⁵Alternatively, we could have summed the masses of star particles younger than, say, 20 Myr and then divided by that timescale. We do not expect our conclusions to change had we used this different SFR definition.




Figure 3.3: Temperature distribution of ISM-scale winds. These distributions are based on \dot{M}_{out} -weighted averages over three broad redshift bins. Solid gray vertical lines are cooling physics-based temperature cuts at 10^3 and 10^5 K. Dashed vertical colored lines illustrate example virial temperatures of the inner halo (at $0.1R_{\text{vir}}$) for representative halos from each mass bin. The virial temperatures roughly align with the temperature distribution peak for hot (virialized) outflows. Cold outflows are more prominent at high redshift. The two lowest mass halos generally do not have multiphase outflows. Note that the m13 halos were only run down to $z = 1$ and so are absent from the bottom panel.

rate; we adopt the common assumption that one SN occurs for every $100M_{\odot}$ of stars formed under reasonable assumptions for the IMF. This is consistent with the FIRE-2 assumptions of a Kroupa (2001) IMF and the STARBURST99 stellar population models (Leitherer et al., 1999); see section 2.5 of Hopkins et al. (2018a) for more details. $E_{\text{SN}} = 10^{51}$ erg is the total mechanical energy assumed to be released by a single Type II SN. $v_{\text{cool}} = 200$ km/s is the terminal velocity of the supernova remnant after it has shocked and swept up ambient ISM material (note that this is lower than the actual injection velocity of ≈ 2000 km/s). We assume the mean SN ejecta mass is $M_{\text{ej}} = 10M_{\odot}$ of which $2M_{\odot}$ is metal mass (so that the mean SN ejecta metallicity is $Z_{\text{SN}} = 0.2$). This is equivalent to the Muratov et al. (2017) approach of defining $\dot{M}_{\text{Z,ref}} = y \times \text{SFR}$, where they use $y = 0.02$ for the chemical yield of one SN per $100M_{\odot}$ stars formed (see their Footnote 4).

3.3.3.2 Redshift-averaged loading factors

To compute loading factors η , we cannot simply divide the wind fluxes by their corresponding reference fluxes in a given snapshot because of the time lag between generation and propagation of outflows (see again the time series in the bottom left of Figure 3.1 and Figure 3.2). The bursty nature of SF in dwarfs means that there will be extended periods of zero SF, which can lead to artificially high instantaneous loading factor estimates (if the time delay and burst integration are not properly accounted for). On the other hand, continuous, steady-state SF in more massive halos at late times also makes it challenging to derive accurate delay times and detect local maxima

in the SFH (Muratov et al., 2015; Hung et al., 2019; Kim et al., 2020a). Given the small-scale complexity of our time series and the fact that we are analyzing outflows over ~ 10 Gyr, we adopt the redshift-averaging approach of Muratov et al. (2015). We compute time integrals of the wind and reference fluxes over sufficiently long timescales so as to encompass multiple stellar feedback episodes. This avoids dependence of loading factors on averaging timescale (for sufficiently long timescales). As an example, for mass loading factors, we integrate the total mass of stars formed and the total mass of wind blown out over a large redshift interval, and divide the latter by the former. We do a similar calculation for cumulatively summed momentum, energy and metal loading factors.

We define the same three redshift bins as Muratov et al. (2015): low-redshift ($z = 0.0 - 0.5$, 5 Gyr), intermediate-redshift ($z = 0.5 - 2.0$, 5 Gyr), and high-redshift ($z = 2.0 - 4.0$, 2 Gyr). Although these redshift bins are extremely long, they have the advantage of giving us a robust estimate of the average loading factor for both the ISM shell and the virial shell, effectively marginalizing over the difference in delay times for $0.1R_{\text{vir}}$ and R_{vir} . This will allow us to more confidently compare our halo loading factors to our ISM loading factors and constrain any losses/gains in mass, energy, momentum and metals as outflows transit the CGM.

Table 3.1, Table 3.2 and Table 3.3 respectively give the average properties and multi-phase loading factors of the FIRE-2 halos in our low-redshift ($z = 0.0 - 0.5$), intermediate-redshift ($z = 0.5 - 2.0$) and high-redshift ($z = 2.0 - 4.0$) bins. Note that the tables printed in this dissertation are only a subset of the much longer set of tables that

give multi-phase loading factors in both the ISM and virial shells for mass, momentum, energy and metals. The full set of supplementary tables is available for download online. We note that we have provided the average global SFR which can be used to convert the dimensionless loading factors back into raw mass, momentum, energy and metal outflow rates in physical units (following subsection 3.3.3). The mass-weighted average Bernoulli velocity (excluding the gravitational term) can be approximated as $\sqrt{\dot{E}/\dot{M}}$ and the mass-weighted average radial velocity (including the thermal momentum component) can be approximated as \dot{p}/\dot{M} .

3.3.3.3 Instantaneous loading factors

In addition to our fiducial redshift-averaged loading factors, we compute instantaneous ISM loading factors for individual outflow episodes.⁶ Any instantaneous loading factor algorithm must involve three steps: (1) time-shifting, (2) peak detection, and (3) burst integration. We now describe our approach for each of these in turn.

Given that our time series span most of the history of the universe, we adopt the following strategy. For each halo, we first split the whole time series from $z = 4$ to $z = 0$ (~ 12 Gyr) into twelve 1 Gyr chunks. Then we cross-correlate the SFH and mass outflow rate history in each chunk to derive a single time lag for that chunk (using `numpy.correlate`). We define the time lag as the value at which the cross-correlation function peaks. Since it is unlikely for the time delay to exceed ~ 200 Myr, we limit the

⁶We will not attempt to derive instantaneous halo-scale loading factors in this chapter because the delay time and halo outflow duration can be substantially longer. There may also be significant variation in how different outflow episodes evolve as they transit the CGM.

Halo	$\log M_{\text{vir}}$	$\log M_*$	V_{vir}	$\log \text{SFR}$	$V_{\text{esc,ISM}}$	$V_{\text{esc,halo}}$	$\eta_{\text{M,ISM}}$	$f_{\text{M,ISM,cold}}$	$f_{\text{M,ISM,warm}}$
m10q	9.8546	6.2894	26.4429	-4.3402	65.9813	26.7645	35.7331	0.0000	0.9202
m10v	9.8418	4.5565	26.2860	-4.3757	62.1723	27.0999	53.3407	0.0000	0.9689
m10y	10.1182	6.9964	32.5870	-3.2242	78.8185	33.9036	37.3902	0.0029	0.7528
m10z	10.5395	7.4673	44.3759	-2.3722	105.8595	46.0907	15.9710	0.0034	0.6545
m11a	10.5537	7.9441	46.0388	-1.5776	113.1263	48.0332	5.9448	0.0003	0.6306
m11b	10.5742	7.9563	46.7756	-1.8808	120.7780	47.6691	2.4817	0.0000	0.4792
m11c	11.1002	8.8442	70.1262	-1.0490	172.2317	72.7490	5.1707	0.0012	0.3836
m11f	11.5972	10.2756	106.2536	0.7062	271.1231	109.1207	0.2519	0.0068	0.3691
m11q	11.1147	8.5553	70.5149	-1.4055	169.3726	73.2528	6.6084	0.0004	0.2552
m11v	11.1405	9.3015	73.6367	-0.4062	174.0561	79.0680	0.9904	0.0020	0.5341
m12f	12.0816	10.8553	156.5903	1.1966	395.9141	160.9539	0.1605	0.0026	0.1211
m12m	12.0184	11.0586	152.7088	1.3188	399.2297	156.7147	0.1790	0.0029	0.0810
m12i	11.9294	10.7769	139.7645	1.0767	360.6893	143.3412	0.1088	0.0003	0.0478

Table 3.1: Average properties and loading factors of the FIRE-2 halos in our low redshift bin ($z = 0.0 - 0.5$). We provide some basic global properties: halo virial mass (M_{\odot}), stellar mass (M_{\odot}), virial velocity (km s^{-1}), SFR ($M_{\odot} \text{ yr}^{-1}$), escape velocity from the ISM at $0.1R_{\text{vir}}$ (km s^{-1}) and escape velocity from the halo at R_{vir} (km s^{-1}). For the loading factors, we provide the average total loading factor (dimensionless) and the corresponding cold, warm and hot phase fractions (η_{phase}/η). As the full table is much longer, in the text here we only show a limited set of columns for the ISM-scale mass loading factors; there are additional columns giving the loading factors and their corresponding phase fractions for momentum, energy and metal outflows (for both the ISM and virial shells). The m13 halos are not shown since they are only run down to $z = 1$ (they appear in the subsequent two tables). **This supplementary table is available to download in its entirety in the online journal.**

Halo	$\log M_{\text{vir}}$	$\log M_*$	V_{vir}	$\log \text{SFR}$	$V_{\text{esc,ISM}}$	$V_{\text{esc,halo}}$	$\eta_{\text{M,ISM}}$	$f_{\text{M,ISM,cold}}$	$f_{\text{M,ISM,warm}}$
m10q	9.6679	6.1910	28.3840	-3.9136	68.7514	29.8457	42.2899	0.0000	0.9007
m10v	9.6041	3.6395	26.8659	-inf	57.1216	29.1336	inf	nan	nan
m10y	9.9107	6.5727	34.4039	-2.9452	79.6717	37.1819	25.3753	0.0041	0.8080
m10z	10.2583	7.0819	44.9341	-2.4946	103.5751	47.9005	17.6529	0.0013	0.6614
m11a	10.3552	7.3467	48.9326	-2.0557	112.5202	51.5762	15.8302	0.0068	0.6865
m11b	10.3977	7.6180	49.8646	-2.1166	118.1264	52.8621	20.0579	0.0211	0.7451
m11c	10.8999	8.4603	74.6748	-0.9751	173.6646	78.3523	7.5931	0.0141	0.5623
m11f	11.3949	9.6310	111.3398	0.4287	264.9936	116.8989	1.5574	0.0145	0.4897
m11q	10.9105	8.1635	73.8962	-1.3263	169.5094	78.1010	7.5609	0.0130	0.4647
m11v	10.8051	8.8465	70.3905	-0.5639	170.5799	74.1870	1.9865	0.0063	0.5575
m12f	11.8129	10.2684	155.6512	0.9783	370.1414	163.9375	0.8151	0.0093	0.3378
m12m	11.8957	10.2896	159.2550	1.2817	393.2780	170.7198	0.3144	0.0150	0.2311
m12i	11.7536	10.0605	147.9378	0.9407	344.9060	155.8567	0.7816	0.0052	0.2826
A1	12.4241	11.3586	276.6363	1.8137	695.7254	287.7614	0.2196	0.0000	0.0000
A2	12.6153	11.5553	324.3633	1.9258	786.7923	340.2941	0.4423	0.0014	0.0037
A4	12.4939	11.2806	288.3921	1.8442	687.2517	306.7733	0.4405	0.0022	0.0083
A8	12.7956	11.3537	366.9799	2.2786	831.0475	389.0615	0.3216	0.0139	0.0366

Table 3.2: Identical to Table 3.1 but now for our intermediate redshift bin ($z = 0.5-2.0$). This supplementary table is available to download in its entirety in the online journal.

Halo	$\log M_{\text{vir}}$	$\log M_*$	V_{vir}	$\log \text{SFR}$	$V_{\text{esc,ISM}}$	$V_{\text{esc,halo}}$	$\eta_{\text{M,ISM}}$	$f_{\text{M,ISM,cold}}$	$f_{\text{M,ISM,warm}}$
m10q	9.4611	5.8281	32.5904	-3.1300	71.2804	34.8458	44.1721	0.0258	0.8701
m10v	8.7289	3.6450	18.4540	-inf	40.5637	20.3351	inf	nan	nan
m10y	9.3987	5.8869	31.3700	-2.7648	67.8910	34.8602	53.5149	0.0693	0.8780
m10z	9.6626	6.0321	38.4115	-2.0593	84.1822	42.0751	29.4567	0.0850	0.7992
m11a	9.8915	6.7998	45.7891	-2.0118	102.0091	48.5946	24.2457	0.0685	0.8138
m11b	10.0150	7.1604	50.6737	-1.8223	113.8998	55.1179	20.0025	0.0460	0.7567
m11c	10.4254	7.8025	69.1876	-1.2756	159.2299	74.5761	10.3179	0.0314	0.6073
m11f	10.8506	8.2414	96.9847	-0.0968	218.5064	102.2394	4.1509	0.0704	0.5978
m11q	10.3770	7.3952	66.5144	-1.1906	146.0821	71.5669	15.1848	0.1659	0.6763
m11v	10.5013	7.9483	74.4011	-0.8721	166.8453	80.9191	7.8286	0.0552	0.6506
m12f	11.3000	9.0861	138.5344	0.6098	311.4655	145.5008	2.3737	0.0233	0.3868
m12m	11.0576	8.3822	115.4186	0.0730	255.1170	132.6813	3.6423	0.0574	0.4311
m12i	11.1531	8.3699	121.9358	-0.0764	258.3681	133.6676	4.8712	0.0358	0.5038
A1	12.1771	11.0075	278.5722	2.0537	682.5353	294.1316	0.2029	0.0046	0.0566
A2	12.2491	11.1264	299.9680	2.3505	736.2042	311.7156	0.1585	0.0074	0.1516
A4	12.0143	10.4359	244.9935	1.9061	567.1888	260.2105	0.3228	0.0354	0.2052
A8	11.8331	9.4767	212.0698	1.5120	475.8216	237.4192	0.6018	0.0463	0.2846

Table 3.3: Identical to Table 3.1 but now for our high redshift bin ($z = 2.0 - 4.0$). This supplementary table is available to download in its entirety in the online journal.

cross-correlation window length to ± 6 snapshots, which roughly translates to a ± 120 Myr window (given the typical snapshot spacing of ~ 20 Myr). Our chunking approach allows for the possibility that the time lag can systematically increase towards later times. Indeed, we find that the time lag roughly increases from $\sim 20 - 40$ Myr at $z \sim 4$ to ~ 50 Myr at $z \sim 1$, and finally to ~ 100 Myr at $z \sim 0$. Muratov et al. (2015, their Appendix B) found a similar systematic increase in the time lag towards low redshift and suggested it could be because halo radii grow with time while outflow velocities do not increase as dramatically (so outflows take longer to get to $0.1R_{\text{vir}}$). Based on visual inspection, we find our simple cross-correlation algorithm to work remarkably well. In dwarfs, SFHs are bursty so the time lag is most easily constrained. In more massive halos, although the SFH is continuous, there can still be peaks (often broad) in the mass outflow rate history that help constrain the cross-correlation. As we will show below, even when the time-shifting is imperfect for individual episodes, our burst integration baselines are usually wide enough to smooth over this error.

Next, in a given chunk, we detect peaks in the shifted mass outflow rate time series using the automated `scipy.signal.find_peaks` routine. This is a powerful algorithm that identifies local maxima based on their “topographic prominence” (i.e., how the amplitude of a peak compares to the amplitude of its direct neighbors). The function also estimates the peak baseline by extending a horizontal line on both sides of the peak until intersection with part of an even higher peak. For efficiency, we limit the extent of this baseline search window to a total of 8 snapshots (for a total possible burst duration of ~ 160 Myr). Although the width of an ISM outflow spike is unlikely

to exceed ~ 100 Myr, we find that allowing for this slightly larger max baseline helps correct for any imperfect time shifts due to the single-lag cross-correlation described earlier.

With the peak centers and baselines for outflow spikes in hand, we numerically integrate the references fluxes and (time-shifted) outflow fluxes within each burst window. While our adopted peak detection algorithm performs well (based on visual inspection), any time series analysis is fraught with uncertainty and some filtering criteria must be applied to remove unwanted, noisy detections. For simplicity, we only have two selection criteria for bursts.⁷ First, we remove outflow episodes where the corresponding burst-integrated stellar mass is zero; these scenarios likely reflect mergers and other inner halo activity. Second, adapting Muratov et al. (2015, their Appendix B), we only keep bursts whose integrated wind mass is at least 10% of the wind mass of the most powerful burst within their 1 Gyr time chunk. This choice is inevitably arbitrary but it is designed to pick up the clearer, well-defined and more interesting outflow episodes. While this does mean we have a floor on our instantaneous loading factors, in the case of mass loading, we can recover values as low as ~ 0.1 in the low-redshift MW halos. As our results will show, our instantaneous loading factors also agree remarkably well with our fiducial redshift-averaged measurements. This serves to validate the two very different approaches while also allowing us to get a sense of the instantaneous scatter in wind loadings.

In Figure 3.4, we illustrate our time-shifting, peak detection and burst inte-

⁷For the m13 halos, we further choose to only include bursts at $z = 2 - 4$ since both the SF and \dot{M}_{out} history are continuous at $z < 2$ and it is not clear that the derived time lags are meaningful.

gration results for three representative 1 Gyr time chunks using m12f and m10q as examples. Instantaneous mass loading factors are found to be $\sim 5 - 10$ times higher in m12f at high-redshift ($z = 1.9 - 1.4$) than at a lower redshift ($z = 0.7 - 0.5$). In the lower redshift chunk, SF is continuous with a non-zero baseline unlike at high-redshift for m12f. However, broad outflow peaks are still apparent and the cross-correlation result seems sensible. In the same lower redshift chunk, m10q only has two starbursts that are spaced far apart (by ~ 600 Myr) and the outflows are highly mass-loaded with $\eta_M \sim 500$ and 80. To better characterize and understand these trends, we will later correlate all individual detected outflow episodes using their associated burst-averaged physical properties.

Table 3.4 provides a catalog of instantaneous burst properties for all halos at $z < 4$. This catalog includes individual integrated burst stellar masses, wind masses and mass loading factors (both combined and split into cold, warm and hot phases). Similarly, integrated multi-phase momentum, energy and metal loadings are also provided for each individual burst. Burst interval averaged SFR and gas surface densities, dense ISM gas fractions, global stellar mass and halo virial velocity, etc. are also provided as discussed above.

3.4 ISM wind loading factors

Here we present our ISM loading factors as a function of a few galaxy/halo properties. It is beyond the scope of this chapter to identify a “universal” halo prop-

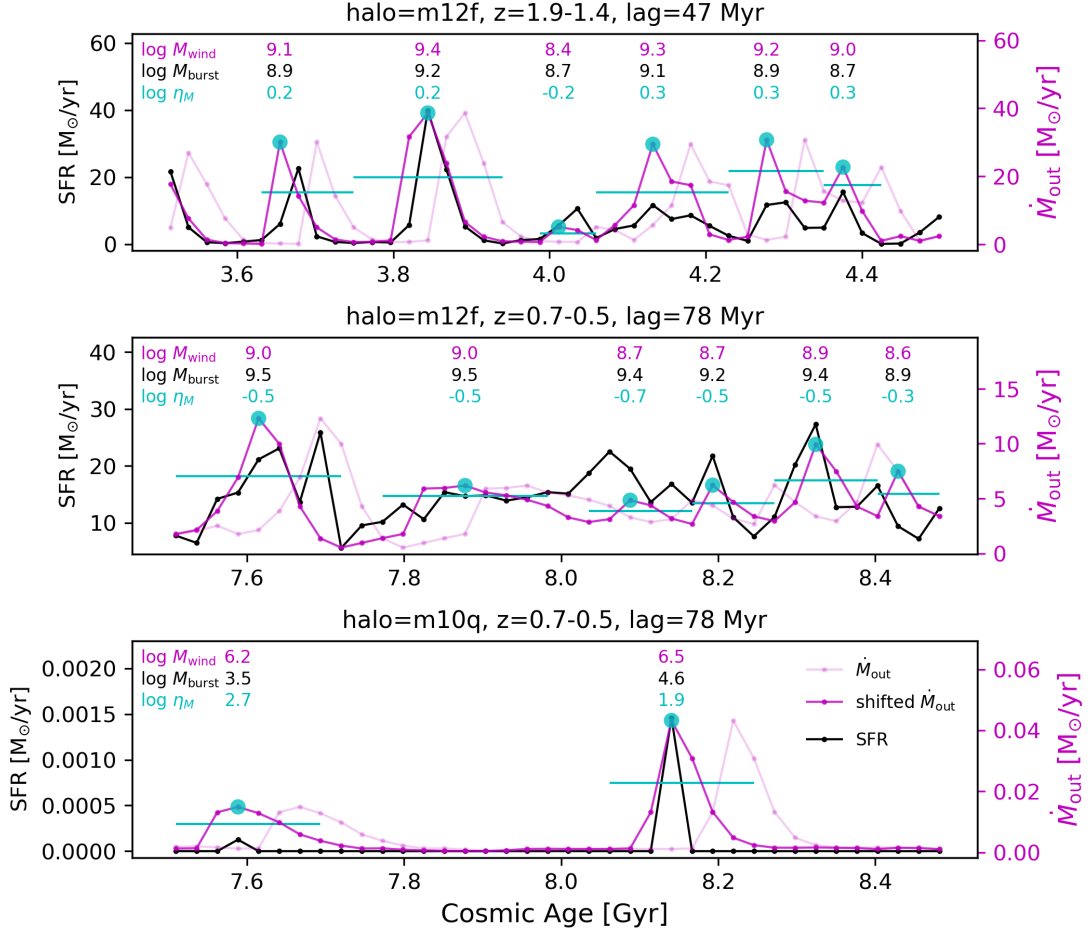


Figure 3.4: Illustration of our automated algorithm for measuring instantaneous loading factors. These are three representative 1 Gyr time chunks showing m12f at high-redshift when SF is bursty (top), m12f at lower redshift when SF is more continuous (middle), and m10q at the same lower redshift when its SF is still bursty (bottom). In each panel, the SFH is shown in black, the original unshifted mass outflow rate history as transparent magenta, and the time-shifted mass outflow rate history as opaque magenta. The cyan dots and horizontal lines identify peaks and their baselines, respectively. Above each detected peak, we write the burst-integrated wind mass, stellar mass, and instantaneous mass loading factor. Note how the instantaneous mass loading is $\sim 5 - 10$ times weaker in m12f at lower redshift compared to high redshift. Note also how highly mass-loaded the two bursts in m10q are despite being at the same lower redshift, and also how far apart these two bursts are in time (~ 600 Myr with zero SF in between).

Halo	Redshift	dt _{lag}	log $M_{*,burst}$	log M_{wind}	log η_M	log $M_{wind,cold}$	log $M_{wind,warm}$	log $M_{wind,hot}$
m10q	3.8854	0.0000	4.7019	6.7398	2.0379	5.0895	6.7139	5.2920
m10q	3.8121	0.0000	3.3653	6.5009	3.1357	4.0429	6.4945	4.5534
m10q	3.5071	0.0000	4.4318	7.0831	2.6513	5.0506	6.9918	6.3392
m10q	3.0765	0.0000	3.2909	6.8781	3.5872	5.4926	6.8177	5.8258
m10q	2.6924	0.0000	6.0332	7.4752	1.4420	5.9537	7.4178	6.4467
m10q	2.0145	67.2997	4.6990	5.5042	0.8052	-inf	5.4945	3.5448
m10q	1.9858	67.2997	3.6776	5.1833	1.5057	-inf	5.1631	3.6249

Table 3.4: A subset of columns and rows from our full instantaneous burst catalog for all halos at $z < 4$. Integrated starburst and wind masses are in M_\odot , momentum in M_\odot km/s, energy in erg, metal mass in M_\odot . Gas mass surface densities are in $M_\odot/\text{yr}/\text{pc}^2$ and SFR surface densities in $M_\odot/\text{yr}/\text{kpc}^2$. Global stellar mass and virial mass are in M_\odot , virial radius in proper kpc, and virial velocity at R_{vir} and $0.1R_{\text{vir}}$ in km/s. The time lag and burst baseline are in Myr (note that because of our algorithm, all bursts in the same 1 Gyr time slice have the same time lag). Unless noted otherwise, all other quantities are dimensionless (see also main text).

This supplementary table is available to download in its entirety in the online journal.

erty (or combination of properties) with which to unambiguously correlate the loading factors. We start with a brief comparison to previous work on mass and metal loading for FIRE-1 halos as a function of both stellar mass and halo virial velocity. We then investigate how all four multi-phase loading factors (mass, momentum, energy and metals) vary with stellar mass and redshift for the FIRE-2 halos. Finally, we correlate our instantaneous loading factors versus burst interval-averaged gas and SFR surface densities and a few other interesting physical properties.

For purely illustrative purposes, we provide fitting functions to approximate many of the trends. However, we caution that the scatter is often large and the optimal functional form is not always obvious. For simplicity, we fit (sometimes broken) power laws in log-log space. We use the `scipy.optimize.curve_fit` implementation of the

Levenberg-Marquardt damped least-squares method (without weighting). We report one standard deviation uncertainties on fitted parameters using the square root of the diagonal entries of the covariance matrix. Unless indicated otherwise, our fits are only done to the broad redshift-averaged measurements and generally include the overly massive m13 halos. In a future work, we will present scalings and quantify scatter in a form that can be implemented into SAMs.

3.4.1 Comparison to FIRE-1

In Figure 3.5, we compare our FIRE-2 measurements of mass and metal loading factors vs. stellar mass to the FIRE-1 results of Muratov et al. (2015) and Muratov et al. (2017), respectively.

Our mass loading factors are roughly a factor of two lower than Muratov et al. (2015), who found a redshift-independent relation with stellar mass: $\eta_{\text{M,ISM}} \propto (M_*/M_\odot)^{-0.35}$. Similar to FIRE-1, our mass loading factors drop off more steeply at $M_* \gtrsim 10^9 M_\odot$ (note that the low-redshift m12 halos were not used to fit the FIRE-1 relation; Muratov et al., 2015). Our lower normalization relative to FIRE-1 is driven by our different particle selection schemes: our Bernoulli velocity wind criterion excludes slower-moving, turbulent flows whereas the simpler $v_{\text{rad}} > 0$ km/s selection of Muratov et al. (2015) includes this slow component (and hence leads to upper limits). We have verified that if we use all particles with $v_{\text{rad}} > 0$ km/s and place the ISM shell at $0.2 - 0.3R_{\text{vir}}$ instead of $0.1 - 0.2R_{\text{vir}}$ (to be even more consistent with Muratov et al., 2015), then our mass loading factors increase and become remarkably similar to FIRE-

1. We also compare to the particle tracking-based measurements of mass loadings in FIRE-1 from Anglés-Alcázar et al. (2017a), which are even higher since they tracked outflows directly out of the ISM (much of which recycles back).

Our metal loading factors agree with FIRE-1 from Muratov et al. (2017) despite our more stringent wind selection criteria. Had we selected outflows with $v_{\text{rad}} > 0$ km/s at $0.2 - 0.3R_{\text{vir}}$ instead of $0.1 - 0.2R_{\text{vir}}$ (like Muratov et al., 2017), we would predict about a factor of two higher metal loading factors than our fiducial measurements. This suggests that although the subgrid physics change from FIRE-1 to FIRE-2 did not greatly affect the overall mass loading factors, there was an appreciable effect on the metal loading and hence metallicity of winds. Nevertheless, our conclusions remain broadly similar to Muratov et al. (2017): ISM metal outflows in dwarfs are comparable to the yield of type II SNe (i.e., $\eta_{\text{Z,ISM}} \sim 1$), with relatively lower ISM metal outflows in the more massive halos.

3.4.2 Global halo circular velocity

Next we plot the mass loading as a function of virial velocity in Figure 3.6. We follow the common practice of plotting ISM mass loading versus global halo circular velocity⁸; we find similar scalings when plotting ISM mass loading versus circular velocity at $0.1R_{\text{vir}}$ or halo-scale mass loading versus circular velocity at R_{vir} . The theoretical motivation for comparing mass loading to circular velocity is that the power

⁸Note that we define V_{vir} as the circular velocity at R_{vir} using our own calculated enclosed mass profile accounting for stars, gas and high-resolution dark matter (see Section 3.3). Some studies take V_{vir} directly from a halo finder, but this may not account for the reduced baryon fractions of dwarfs if only the dark matter particles are used.




Figure 3.5: Comparison of our ISM-scale mass and metal loading factors as a function of stellar mass to previous FIRE-1 work (Muratov et al., 2015, 2017; Anglés-Alcázar et al., 2017a). *Top*: Our fiducial mass loadings are roughly $\sim 2\times$ lower than Muratov et al. (2015), who found $\eta_{M, \text{ISM}} \propto (M_*/M_\odot)^{-0.35}$ (excluding the low-redshift m12 halos from their fit). This is due to our Bernoulli velocity wind criterion excluding slower, turbulent flows which would otherwise lead to larger mass loadings as in Muratov et al. (2015); the transparent symbols show that we would agree with FIRE-1 remarkably well if we use the same wind selection criteria. We also plot the particle tracking-based measurements of mass loadings in FIRE-1 from Anglés-Alcázar et al. (2017a), which are even higher since they track outflows directly out of the ISM (much of which recycles back). *Bottom*: Despite our stricter wind selection criterion, our metal loadings agree with Muratov et al. (2017), suggesting that the switch from FIRE-1 to FIRE-2 subgrid physics affected outflow metallicities. The transparent symbols show that we would predict even larger metal loadings than FIRE-1 if we use the same wind selection criteria as Muratov et al. (2017). Nevertheless, our overall conclusion is similar: nearly all metals produced by Type II SNe are ejected from the ISM of dwarfs but retained within the ISM of more massive galaxies.

law is expected to be steeper for energy-driven winds ($\eta_M \propto V^{-2}$) and shallower for momentum-driven winds ($\eta_M \propto V^{-1}$); see Murray et al., 2005). Muratov et al. (2015) found a very steep slope for the FIRE-1 dwarfs ($\propto V^{-3.2}$), and then a transition to a shallower slope ($\propto V^{-1}$) for more massive halos with $V_{\text{vir}} > 60$ km/s.

At high-redshift, we find that our measurements follow

$$\eta_{\text{M,ISM}} = 10^{4.6} \left(\frac{V_{\text{vir}}}{\text{km s}^{-1}} \right)^{-2.0} \quad \text{for } z = 2.0 - 4.0 \quad (3.12)$$

with a coefficient error of ± 0.2 dex and power law exponent error of ± 0.1 (the m13 halos are excluded from the fit). This is consistent with the expectation for energy-conserving winds. We do not see the need for a broken power law with a shallower slope for more massive halos at high-redshift. If anything, the m13 halos at high-redshift fall off more steeply than expected for a $\propto V_{\text{vir}}^{-2}$ scaling, perhaps suggesting they retain more of their outflows in the ISM as fuel for rapid early star formation.

At intermediate redshift, the relation steepens:

$$\eta_{\text{M,ISM}} = 10^{5.6} \left(\frac{V_{\text{vir}}}{\text{km s}^{-1}} \right)^{-2.6} \quad \text{for } z = 0.5 - 2.0 \quad (3.13)$$

with a coefficient error of ± 0.5 dex and power law exponent error of ± 0.2 (again excluding the m13 halos). The relation steepens even further by low redshift:

$$\eta_{\text{M,ISM}} = 10^{6.5} \left(\frac{V_{\text{vir}}}{\text{km s}^{-1}} \right)^{-3.4} \quad \text{for } z = 0.0 - 0.5 \quad (3.14)$$

with errors of ± 0.6 dex and ± 0.3 for the coefficient and power law exponent, respectively. There is a hint that a broken power law may be appropriate at intermediate-redshift given the elevated $\eta_{\text{M,ISM}}$ of the m13 halos, but this would be at a much higher pivot

point ($\gtrsim 200 \text{ km s}^{-1}$) than the 60 km s^{-1} found by Muratov et al. (2015). The MW halos follow our simple unbroken scalings at both intermediate- and low-redshift. As we will discuss later, the stronger redshift dependence when plotting against halo virial velocity instead of stellar mass may reflect the fact that the stellar-to-halo-mass ratio, at fixed halo mass, gets larger at later times whereas V_{vir} does not evolve as dramatically (this is particularly true for the massive halos).

3.4.3 Multi-phase ISM mass loadings

In Figure 3.7, we plot multi-phase ISM mass loading factors versus stellar mass. Our fiducial measurements in the broad redshift bins are shown with the symbols, but we also plot continuous “tracks” of the loading factors based on ± 1 Gyr boxcar smoothing to get a sense of the scatter. For total loadings (all phases combined), we plot the actual loading factors whereas for the individual phases we plot fractions for clarity (i.e., $\eta_{\text{phase}}/\eta_{\text{total}}$).

The topmost panel is similar to Figure 3.5: when combining all phases, dwarfs have significantly higher mass loading factors than more massive halos. The total mass loadings in low-mass dwarfs are of order ~ 100 , steadily dropping towards ~ 10 for intermediate-mass dwarfs and becoming less than unity for the m12 and m13 halos (despite the latter being at high redshift). The total mass loadings approximately follow

$$\eta_{\text{M,ISM}} = 10^{4.4} (M_*/M_{\odot})^{-0.45} \quad (3.15)$$

with errors of ± 0.2 dex and ± 0.02 for the coefficient and power law exponent, respec-

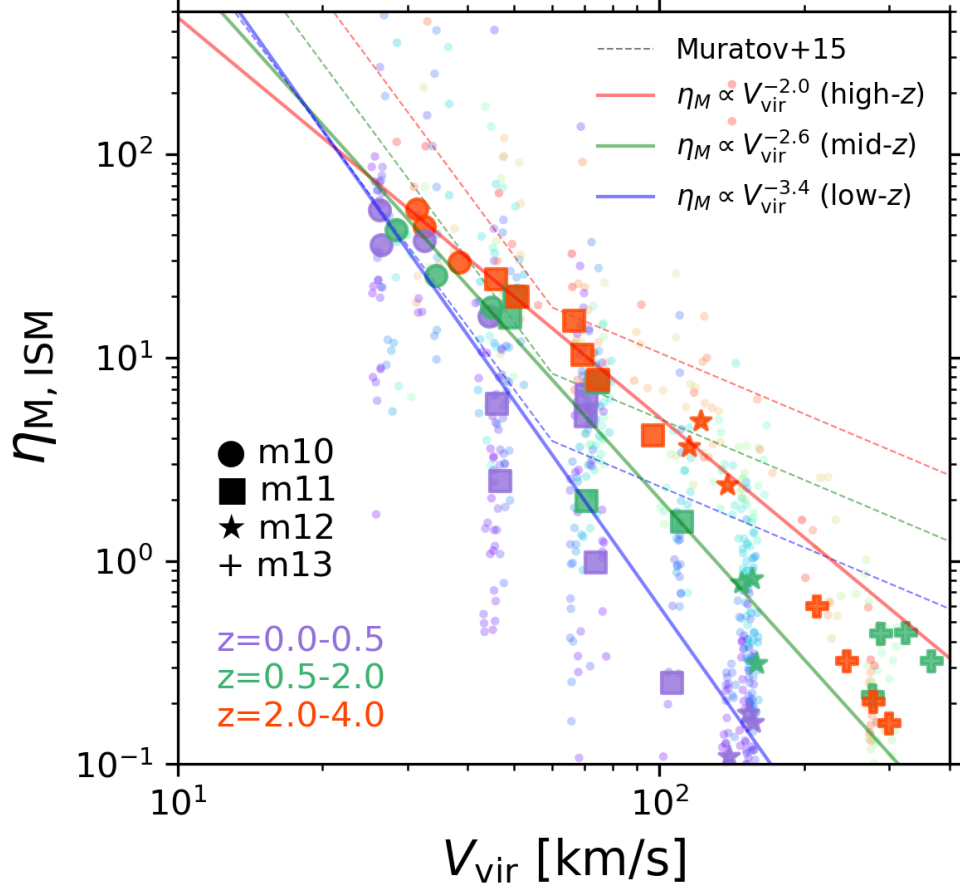


Figure 3.6: ISM mass loading as a function of halo circular velocity. The big markers show our fiducial redshift-averaged measurements whereas the small dots show our instantaneous mass loading measurements color-coded by redshift (V_{vir} for the latter is an \dot{M} weighted average over individual burst intervals). Muratov et al. (2015) found a shallower slope of ISM mass loading with halo circular velocity for FIRE-1 halos with $V_{\text{vir}} > 60$ km/s and interpreted it to mean a transition from energy-driven winds in dwarfs to momentum-driven winds in higher mass halos (dotted colored lines). We do not see this flattening with our more stringent outflow selection criteria in FIRE-2 except possibly at the very massive end ($V_{\text{vir}} \gtrsim 300$ km s $^{-1}$) starting at intermediate redshift. Instead, our measurements are roughly consistent with a single power law that goes as V_{vir}^{-2} at high-redshift (red line), with a steepening at later times (green and blue lines).

tively. At low redshift, a few of the m11 halos and all three m12 halos are a factor of a few below our approximate fit (see also FIRE-1; Muratov et al., 2015; Hayward & Hopkins, 2017). This may reflect the deeper potential wells at later times as well as the changing structure of the ISM and inner CGM over time, as we will discuss in Section 3.6.

The cold mass loading fractions correlate strongly with redshift but are generally flat with stellar mass (at a given redshift). High redshift dwarfs (including the progenitors of the MW halos) have cold mass loading fractions approaching ~ 0.1 . At lower redshifts, the cold mass loading fractions drop to 0.01 or less. We find that the cold mass loading fractions can be approximated simply as

$$f_{\text{M,cold}} \equiv \frac{\eta_{\text{M,ISM,cold}}}{\eta_{\text{M,ISM}}} = 10^{-3.1} (1+z)^{3.1} \quad (3.16)$$

with errors of ± 0.1 dex and ± 0.3 for the coefficient and power law exponent, respectively. We have excluded the m13 halos from the fits.

By comparison, the warm and hot phases show much less scatter. In the m13 and $z \sim 0$ MW halos, the hot phase carries most of the outflowing mass. In contrast, the warm phase carries most of the outflows in dwarfs, including the high-redshift MW progenitor dwarfs. As we will discuss later, this may reflect the fact that the virial temperatures of the dwarfs are below our hot phase cut-off of 10^5 K, meaning that the warm phase is still “dynamically hot” in the dwarfs. In the high-redshift dwarf progenitors of MW halos, the cold mass loadings are comparable to the hot mass loadings, yet the warm mass loadings dominate over the other two phases by a factor

of ~ 10 . There is no strong redshift dependence for the warm and hot mass loading fractions, but the warm mass loading fractions drop significantly at high stellar masses. We fit a broken power law to the warm mass loading fractions with the break fixed at $10^{10.5} M_\odot$:

$$f_{\text{M,warm}} \equiv \frac{\eta_{\text{M,ISM,warm}}}{\eta_{\text{M,ISM}}} = \begin{cases} 10^{-0.5} (M_*/10^{10.5} M_\odot)^{-0.09} & \text{for } M_* \lesssim 10^{10.5} M_\odot \\ 10^{-0.5} (M_*/10^{10.5} M_\odot)^{-2.0} & \text{for } M_* \gtrsim 10^{10.5} M_\odot \end{cases} \quad (3.17)$$

with errors of ± 0.1 dex for the coefficient, ± 0.04 for the low-mass exponent and ± 0.3 for the high-mass exponent. For the hot mass loading fractions, we assume a single power law:

$$f_{\text{M,hot}} \equiv \frac{\eta_{\text{M,ISM,hot}}}{\eta_{\text{M,ISM}}} = 10^{-1.9} (M_*/M_\odot)^{0.18} \quad (3.18)$$

with errors of ± 0.1 dex for the coefficient and ± 0.01 for the power law exponent.

3.4.4 Multi-phase ISM momentum loadings

In Figure 3.8, we show the multi-phase ISM momentum loading factors versus stellar mass. With all phases combined, the total momentum loadings are $\lesssim 0.1$ for the MW halos at $z \sim 0$ as well as the m13 halos at high-redshift. In contrast, the momentum loadings are of order unity in the dwarfs, including the high-redshift progenitors of MW halos. For the lowest mass halos, the momentum loading exceeds one, which may reflect some SN/superbubble clustering phenomenon (e.g., Gentry et al., 2017; Fielding et al., 2018; Faucher-Giguère, 2018). There is some scatter in the momentum loadings of dwarfs but averaged over long timescales their values exceed those of the more massive

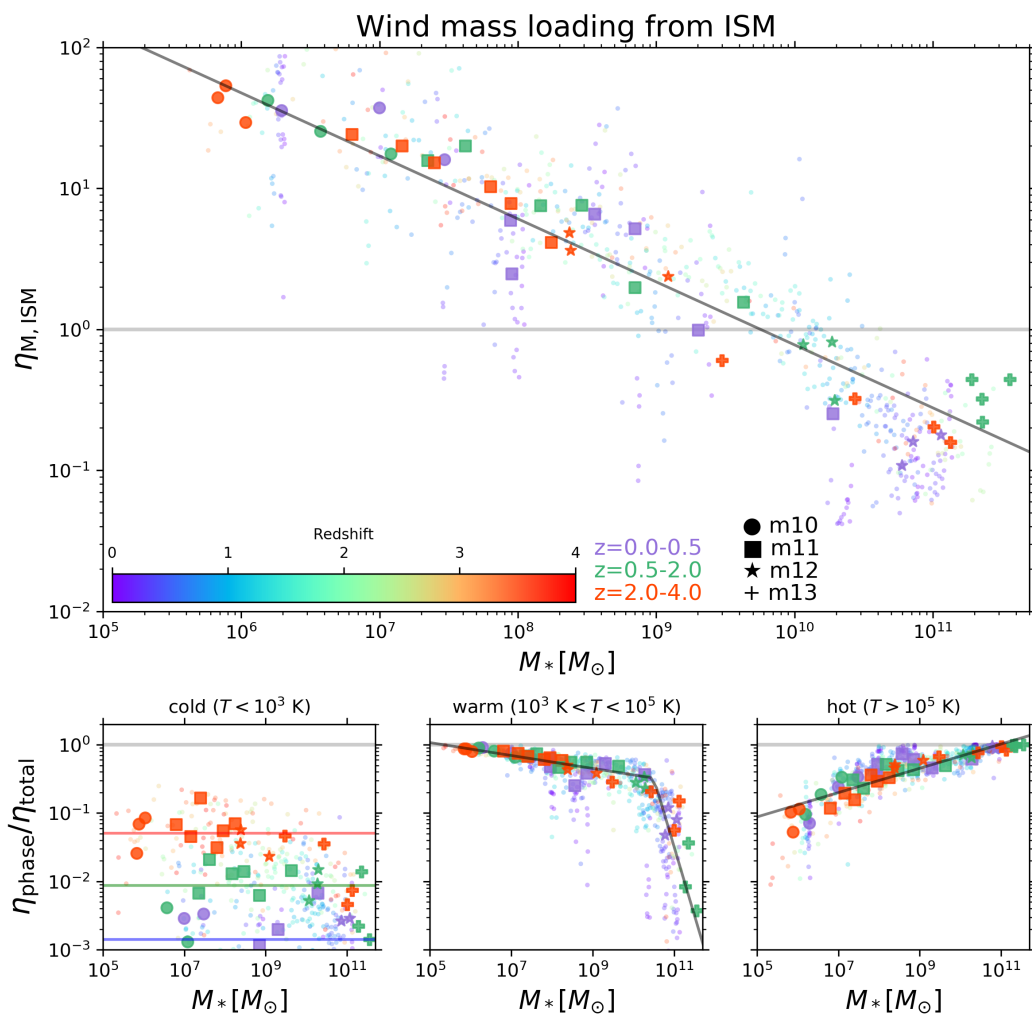


Figure 3.7: Evolution of multi-phase ISM mass loading factors with stellar mass. The bigger markers show our fiducial redshift-averaged measurements and the smaller dots are our instantaneous burst measurements color-coded by redshift (M_* for the latter is an \dot{M} weighted average over each individual burst interval). The horizontal gray line denotes order unity and the other lines show approximate fits (see text). *Top*: For all outflow phases combined, η_M is less than 1 for massive halos and rises to ~ 100 for dwarfs. *Bottom-left*: The fraction of mass loading in the cold phase is far less than one but correlates strongly with redshift. High redshift halos can have $\sim 10\%$ of their mass loading in the cold phase, but this drops to less than 1% for most halos at later times. *Bottom-middle*: the warm outflow phase dominates by mass fraction in the dwarfs. Note also the much tighter correlation and lack of redshift dependence compared to the cold phase. *Bottom-right*: the hot phase dominates in the massive halos but steadily drops off toward lower mass halos. Interestingly, when the total $\eta_{M, \text{ISM}} < 1$, the hot phase dominates, whereas in the dwarfs with $\eta_{M, \text{ISM}} \gg 1$ the warm phase dominates.

halos by about a factor of ten. We approximate the scaling between total momentum loading and stellar mass as

$$\eta_{\text{p,ISM}} = 10^{2.1} (M_*/M_\odot)^{-0.29} \quad (3.19)$$

with errors of ± 0.2 dex for the coefficient and ± 0.02 for the power law exponent.

Splitting by phase, the cold momentum loadings are negligible in low-redshift dwarfs, MW halos at $z \sim 0$ and the m13 halos. However, cold momentum loading fractions are more substantial in high redshift dwarfs: the progenitors of MW halos have values of a few percent whereas some lower mass dwarfs exceed 10%. A simple redshift-dependent formula can reasonably approximate the cold momentum loading fractions (excluding the m13 halos):

$$f_{\text{p,cold}} = 10^{-3.2} (1+z)^{2.9} \quad (3.20)$$

with errors of ± 0.1 dex for the coefficient and ± 0.3 for the power law exponent.

The warm and hot momentum loading fractions are significantly higher than the cold momentum loadings for all galaxies considered. For the more massive halos (including the MW halos at $z \sim 0$), the hot momentum loading fractions are much larger than the warm momentum loading fractions and approach order unity. In intermediate-mass dwarfs, the hot and warm momentum loading fractions are comparable, and in the lowest mass dwarfs the warm phase carries nearly all of the momentum. The importance of warm momentum loading in low-mass dwarfs may be related to their virial temperatures being lower than 10^5 K: their outflows may not satisfy our lower limit for hot temperatures but may still be “dynamically hot.” We approximate our trends for

the warm phase using a broken power law with the break fixed at $10^{10.5}M_{\odot}$:

$$f_{\text{p,warm}} = \begin{cases} 10^{-0.6}(M_*/10^{10.5}M_{\odot})^{-0.10} & \text{for } M_* \lesssim 10^{10.5}M_{\odot} \\ 10^{-0.6}(M_*/10^{10.5}M_{\odot})^{-2.1} & \text{for } M_* \gtrsim 10^{10.5}M_{\odot} \end{cases} \quad (3.21)$$

with errors of ± 0.1 dex for the coefficient, ± 0.05 for the low-mass exponent and ± 0.3 for the high-mass exponent. For the hot phase, we find simply

$$f_{\text{p,hot}} = 10^{-1.3}(M_*/M_{\odot})^{0.12} \quad (3.22)$$

with a coefficient error of ± 0.1 dex and power law exponent error of ± 0.01 .

3.4.5 Multi-phase ISM energy loadings

In Figure 3.9, we plot multi-phase ISM energy loading factors versus stellar mass. When we consider all phases combined, the total ISM energy loadings are less than 0.1 in the MW halos at low-redshift. The same is true for the m13 halos at both intermediate and high redshift, which have similarly low energy loadings. In contrast, dwarfs at high redshift have energy loadings of order unity. At lower redshifts, dwarfs show more scatter in their total energy loadings, but still maintain preferentially higher energy loadings compared to the massive halos (generally $\eta_{\text{E,ISM}} \gtrsim 0.2$). Taking into account this complicated redshift and mass dependence, we parameterize the energy loadings as a broken power law at high-redshift (with the break point fixed at 10^9M_{\odot})

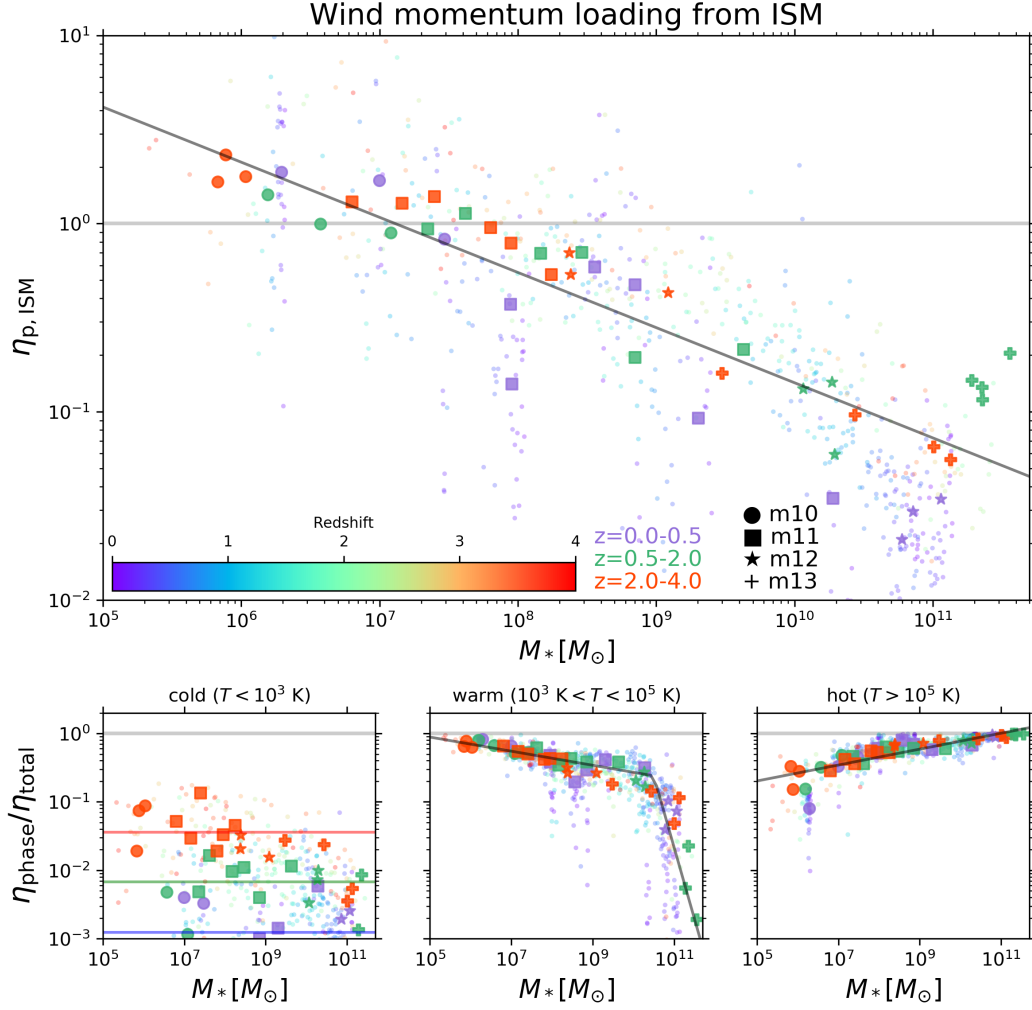


Figure 3.8: Similar to Figure 3.7 but now for the ISM momentum loading factor. *Top*: The overall momentum loading factor is of order unity for higher redshift dwarfs, dropping to ~ 0.1 for some dwarfs at $z \sim 0$. For massive halos, overall momentum loadings are generally less than 0.1. *Bottom-left*: The fraction of momentum loading in the cold phase is negligible except for high redshift dwarfs where it can approach $\sim 10\%$. *Bottom-middle*: The warm phase carries nearly all of the momentum in the lowest mass dwarfs, gradually dropping to $\lesssim 10\%$ for more massive halos. *Bottom-right*: In contrast, the hot phase carries nearly all of the momentum in massive halos, gradually dropping to $\sim 10\%$ for the lowest mass dwarfs.

and two distinct power laws for the other two redshift bins:

$$\eta_{\text{E,ISM}} = \begin{cases} 10^{1.3} (M_*/M_\odot)^{-0.25} & \text{for } z = 0.0 - 0.5 \\ 10^{0.5} (M_*/M_\odot)^{-0.11} & \text{for } z = 0.5 - 2.0 \\ 10^{0.001} (M_*/10^9 M_\odot)^{-0.04} & \text{for } z = 2.0 - 4.0 \text{ \& } M_* \lesssim 10^9 M_\odot \\ 10^{0.001} (M_*/10^9 M_\odot)^{-0.44} & \text{for } z = 2.0 - 4.0 \text{ \& } M_* \gtrsim 10^9 M_\odot . \end{cases} \quad (3.23)$$

The errors for the low-redshift power law are ± 0.6 dex for the coefficient and ± 0.07 for the exponent. The errors for the intermediate-redshift power law are ± 0.3 dex for the coefficient and ± 0.03 for the exponent. As for the high-redshift broken power law, the errors are ± 0.05 dex (coefficient), ± 0.03 (low-mass exponent), and ± 0.04 (high-mass exponent).

Splitting by phase, the cold energy loading fractions are negligible in all halos compared to the warm and hot energy loading fractions, although the scatter in cold energy loading fractions correlates positively with redshift. Just as for the cold mass and momentum loading fractions, we can approximate the redshift dependence in a simple way (again, excluding the m13 halos):

$$f_{\text{E,cold}} = 10^{-3.3} (1+z)^{2.4} \quad (3.24)$$

with errors of ± 0.1 dex (coefficient) and ± 0.3 (exponent).

The hot energy loading fractions dominate over the warm energy loading fractions by about an order of magnitude for the MW halos, their high redshift dwarf progenitors, and the m13 high-redshift halos. In contrast, a substantial fraction of energy is carried by the warm phase in lower mass halos. We approximate the warm

energy loading fractions as a broken power law with the break fixed at $10^{10.5}M_\odot$:

$$f_{\text{E,warm}} = \begin{cases} 10^{-0.9}(M_*/10^{10.5}M_\odot)^{-0.11} & \text{for } M_* \lesssim 10^{10.5}M_\odot \\ 10^{-0.9}(M_*/10^{10.5}M_\odot)^{-1.5} & \text{for } M_* \gtrsim 10^{10.5}M_\odot \end{cases} \quad (3.25)$$

where the errors are ± 0.1 dex (coefficient), ± 0.05 (low-mass exponent) and ± 0.3 (high-mass exponent). For the hot energy loading fractions, we find

$$f_{\text{E,hot}} = 10^{-0.60}(M_*/M_\odot)^{0.054} \quad (3.26)$$

where the errors are ± 0.07 dex (coefficient) and ± 0.008 (power law exponent).

3.4.6 Multi-phase ISM metal loadings

In Figure 3.10, we plot multi-phase ISM metal loading factors versus stellar mass. Similar to Figure 3.5, with all phases combined the total ISM metal loadings are of order unity in dwarfs at all redshifts (i.e., including progenitors of MW halos). However, in more massive halos, the ISM metal loadings drop steadily to ~ 0.1 when averaged over long timescales. There is no strong redshift dependence for the total metal loadings, allowing us to simply parameterize the trends with halo mass using a broken power law (with the break point fixed at 10^9M_\odot):

$$\eta_{\text{Z,ISM}} = \begin{cases} 10^{0.001}(M_*/10^9M_\odot)^{-0.04} & \text{for } M_* \lesssim 10^9M_\odot \\ 10^{0.001}(M_*/10^9M_\odot)^{-0.44} & \text{for } M_* \gtrsim 10^9M_\odot . \end{cases} \quad (3.27)$$

The errors are ± 0.05 dex (coefficient), ± 0.03 (low-mass exponent) and ± 0.04 (high-mass exponent).

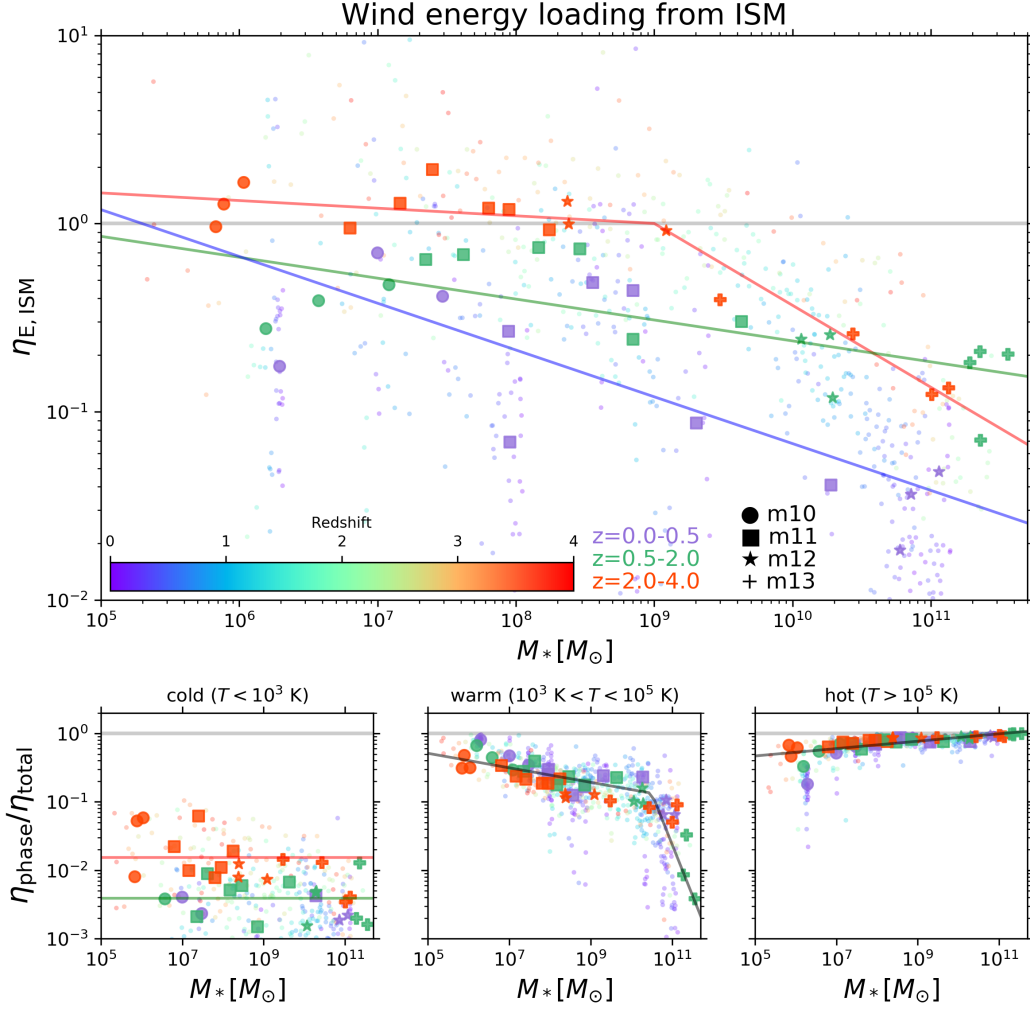


Figure 3.9: Similar to Figure 3.7 but now for the evolution of ISM energy loading factor. *Top*: Overall energy loadings are of order unity for dwarfs at high-redshift, and between 0.1 and 1 for lower redshift dwarfs. MW and m13 halos have $\eta_E \sim 0.1$ especially at lower redshift. *Bottom-left*: The fraction of energy loading carried by the cold phase is negligible, except in some low-mass dwarfs at high-redshift where it is almost $\sim 10\%$ (signifying large kinetic energy). *Bottom-middle*: The warm phase carries only about $\sim 10\%$ of the energy loadings in massive halos but becomes increasingly important for low-mass dwarfs, where its fractional contribution approaches $\sim 50\%$ or higher. *Bottom-right*: The hot phase carries $\sim 50\%$ of the energy in the lowest mass dwarfs and effectively becomes the sole carrier of energy in massive halos ($f_{\text{hot}} \sim 100\%$).

Splitting by phase, metals carried by the cold phase are negligible overall but there is a strong redshift dependence. At high redshift, all halos have roughly constant cold metal loading fractions of ≈ 0.05 . At later times, the cold metal loading fractions decrease but there seems to be a positive correlation with stellar mass. Excluding the m13 halos, we parameterize the cold metal loading fractions as

$$f_{Z,\text{cold}} = \begin{cases} 10^{-3.1}(M_*/M_\odot)^{0.10} & \text{for } z = 0.0 - 0.5 \\ 10^{-2.8}(M_*/M_\odot)^{0.10} & \text{for } z = 0.5 - 2.0 \\ 10^{-0.9}(M_*/M_\odot)^{-0.07} & \text{for } z = 2.0 - 4.0 . \end{cases} \quad (3.28)$$

The errors for the low-redshift power law are ± 0.8 dex (coefficient) and ± 0.08 (exponent). The errors for both the intermediate- and high-redshift power laws are the same: ± 0.5 dex (coefficient) and ± 0.06 (exponent).

The hot metal loading fraction is of order unity and the warm metal loading fraction is of order 0.1 in more massive halos. In contrast, for the lowest mass halos, the warm phase carries nearly all of the metals (the hot metal loading fraction drops to order 0.1). We fit a broken power law to the warm metal loading fractions assuming a fixed break at $10^{11}M_\odot$:

$$f_{Z,\text{warm}} = \begin{cases} 10^{-0.5}(M_*/10^{11}M_\odot)^{-0.08} & \text{for } M_* \lesssim 10^{11}M_\odot \\ 10^{-0.5}(M_*/10^{11}M_\odot)^{-2.7} & \text{for } M_* \gtrsim 10^{11}M_\odot \end{cases} \quad (3.29)$$

where the errors are ± 0.1 dex (coefficient), ± 0.03 (low-mass exponent), and ± 0.5 (high-mass exponent). For the hot metal loading fractions, we find:

$$f_{Z,\text{hot}} = 10^{-1.2}(M_*/M_\odot)^{0.10} \quad (3.30)$$

where the errors are ± 0.1 dex (coefficient) and ± 0.01 (power law exponent).

3.4.7 Wind radial velocities

Wind radial velocity distributions are another fundamental quantity for characterizing the thermodynamics of outflows, and are also more readily constrained by observations rather than total masses, momenta, energies, Bernoulli velocities, etc.

Figure 3.11 shows the distributions of ISM mass outflow rate in bins of radial velocity for wind particles classified as cold, warm or hot. The distributions are averaged over our fiducial redshift ranges weighting by the overall $\dot{M}_{\text{out,ISM}}$ in each contributing snapshot. As can also be seen in the velocity panel of our movies (Figure 3.1 and Figure 3.2), the full radial velocity distributions generally extend to very low values for hot and even warm outflows. These slow moving particles have enough energy from their temperature, compared to the halo potential, to still be classified as wind particles. However cold outflows generally do not extend to very low velocities since their thermal energy is negligible and cannot get them classified as winds according to our Bernoulli velocity criterion. All three phases show a sharp cutoff above a few thousand km/s since these are likely the fastest that stellar-driven winds can propagate (AGN feedback may lead to even faster winds but that is not implemented here).

Figure 3.12 collapses the full multi-phase radial velocity distributions into mass-flux-weighted average radial velocities in each phase as a function of halo circular velocity. This is a simpler, more traditional plot compared to Figure 3.15, where we compared the average mass-flux-weighted Bernoulli velocity to the difference in escape

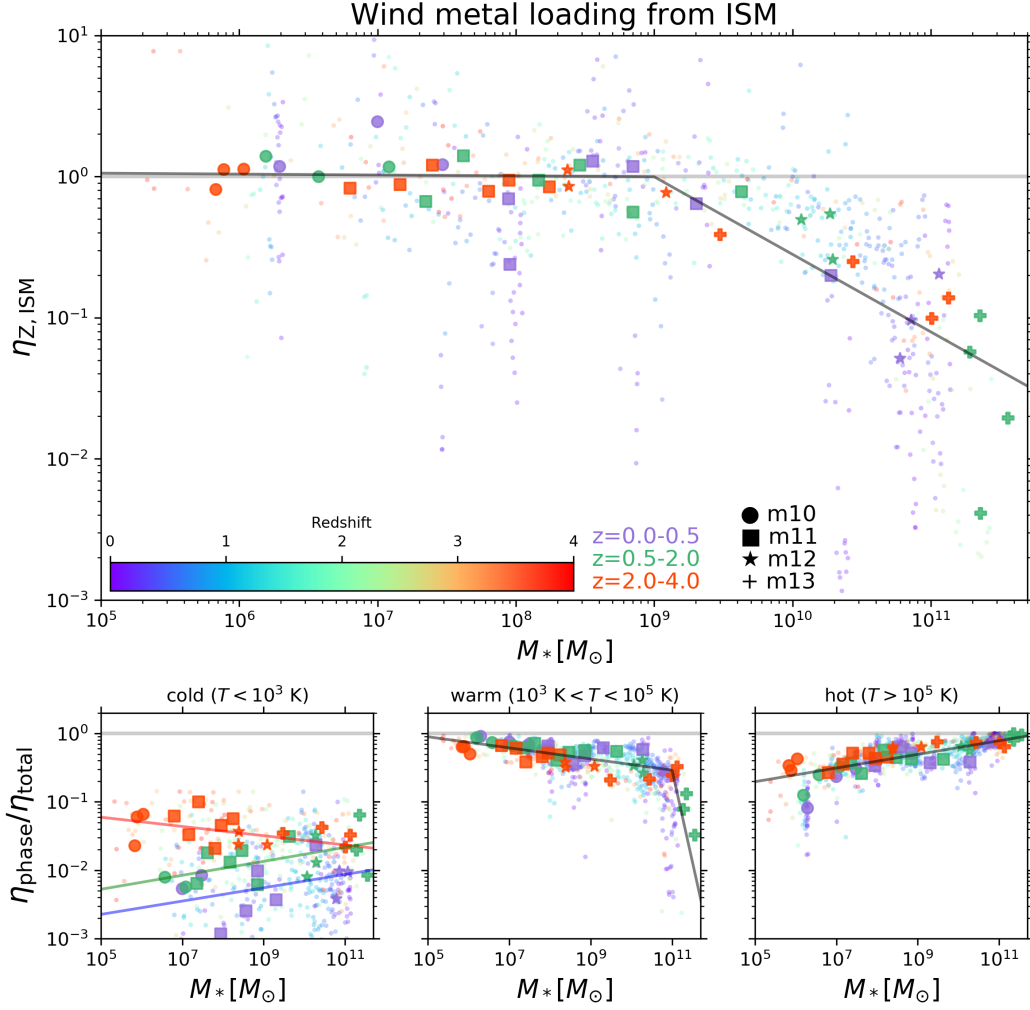


Figure 3.10: Similar to Figure 3.7 but now showing evolution of the ISM metal loading factor. *Top*: Overall metal loadings are of order unity for dwarfs, i.e., nearly all metals produced by SNe are ejected from the ISM of dwarfs. In contrast, the SN metal yield is mostly retained within the ISM of massive halos. *Bottom-left*: The fraction of metals carried by the cold phase is generally negligible for all halos except in high-redshift dwarfs and intermediate-redshift massive halos (for which $f_{\text{cold}} \sim 10\%$). *Bottom-middle*: The warm phase carries nearly all of the metals in the lowest mass dwarfs, and becomes progressively less important for more massive halos, dropping to $\sim 10\%$. *Bottom-right*: In contrast, the hot phase carries nearly all of the metals in massive halos compared to only $\sim 10\%$ in the lowest mass halos.

velocity between $0.1R_{\text{vir}}$ and $1.0R_{\text{vir}}$. Consistent with previous work, we find a positive correlation between average outflow radial velocity and halo circular velocity, with the cold and warm outflows generally clustering around $v_{\text{rad}} \approx 2 \times V_{\text{vir}}$. The relatively large radial velocities of cold and warm outflows is likely driven by the fact that most of their Bernoulli velocity must come from the kinetic term, which alone needs to be sufficiently large for traveling to $\gtrsim 0.5R_{\text{vir}}$. The hot outflows can have substantially larger radial velocities, especially in some high-redshift dwarfs where $v_r \approx 5 \times V_{\text{vir}}$. Interestingly, in some halos, the average hot outflow radial velocities can be lower than V_{vir} ; this may reflect deceleration of outflows due to high density gas within the ISM and inner CGM. However, these slower moving outflows still have enough energy to travel deep into the CGM owing to their hot temperatures and hence sufficiently high Bernoulli velocities.

It is beyond the scope of this work to present a detailed two-dimensional analysis of loading factors simultaneously in temperature and radial velocity bins (but see the bottom-right panel of Figure 3.1 and Figure 3.2). Such an analysis would provide useful constraints for launching of galactic winds in SAMs and lower resolution cosmological simulations (see discussion in Kim et al., 2020b). It is also beyond the scope of our work to investigate full radial velocity profiles and compare to observations. However, our analysis can be adapted in the future to study outflows closer to the ISM and make predictions for observables based on the trajectories and intrinsic evolution of wind particles (following, e.g., Anglés-Alcázar et al., 2017a; Hafen et al., 2020).

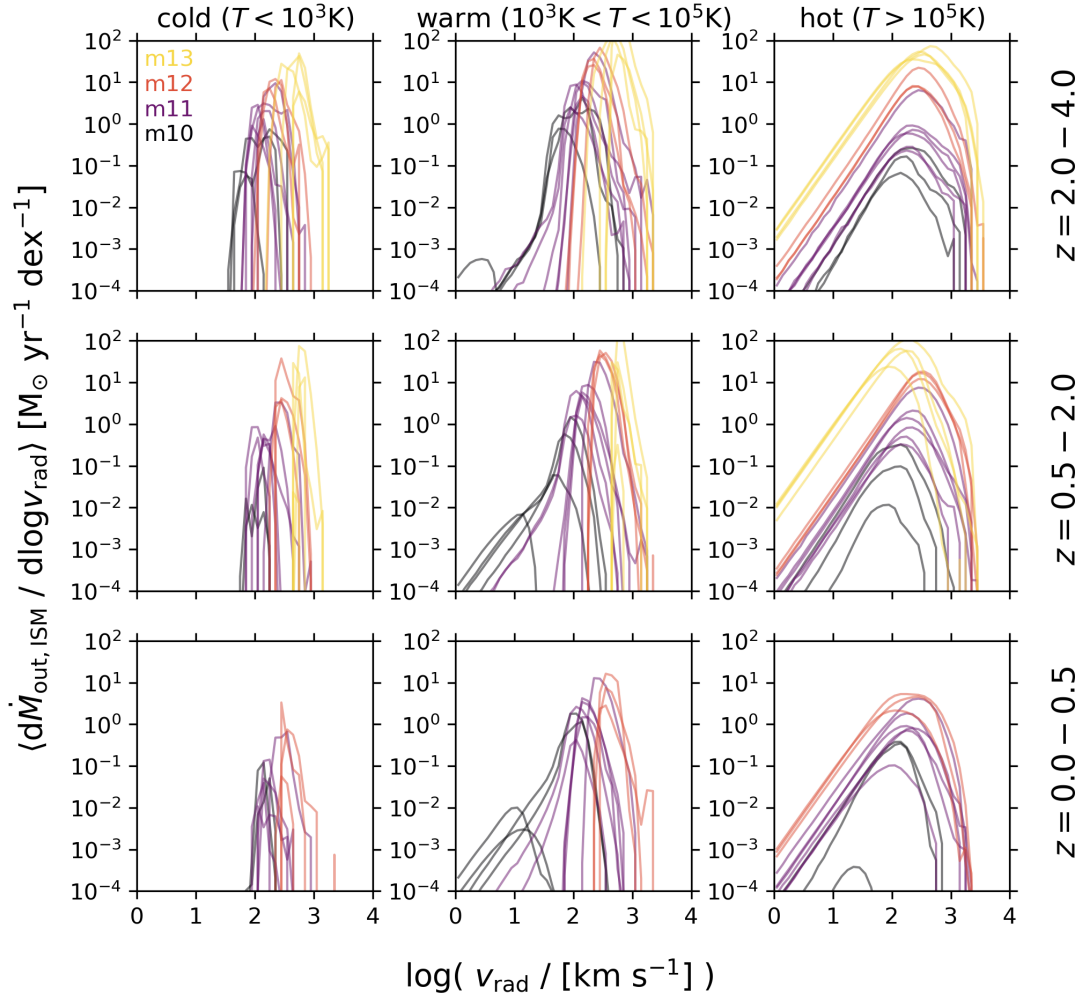


Figure 3.11: Similar to Figure 3.3 but now showing the distributions of ISM mass outflow rate in bins of radial velocity for wind particles classified as cold (left column), warm (middle column) or hot (right column). The different rows show the average distributions over our fiducial large redshift bins, where snapshots with higher total $\dot{M}_{\text{out,ISM}}$ are given higher weight in the average.




Figure 3.12: Analogous to Figure 3.15 but now, following common practice, we plot the mass-flux-weighted average radial velocity versus halo virial velocity. From left to right: we plot this for the cold, warm and hot ISM outflows. The solid gray line is the one-to-one mapping between radial velocity halo circular velocity; the dashed and dotted lines are twice and five times the circular velocity, respectively. We see that generally cold and warm outflows cluster around $\approx 2V_{\text{vir}}$, with slightly lower radial velocities in the m13 halos. The hot outflows tend to be faster, approaching $\approx 5 \times V_{\text{vir}}$ on average for some dwarfs. Interestingly the radial velocity of hot outflows in some halos can be less than V_{vir} , which either suggests deceleration due to interactions or that the slower component of the hot wind dominates (as illustrated in the bottom-right panel of Figure 3.1 and Figure 3.2).

3.4.8 Trends with SFR and ISM gas mass surface densities

In Figure 3.13, we plot our instantaneous burst-integrated mass loading factor as a function of the \dot{M} -weighted average Σ_{gas} and Σ_{SFR} within individual burst windows. Σ_{gas} is defined as $M_{\text{gas}}/(\pi R_{3\text{D}}^2)$ where M_{gas} is the mass of all gas particles within $0.1R_{\text{vir}}$ and $R_{3\text{D}}$ is the 3D stellar half-mass radius (a commonly used definition of galaxy size) computed using star particles within $0.1R_{\text{vir}}$. Σ_{SFR} is defined similarly except the numerator is the instantaneous SFR as described in subsection 3.3.3. Hence we are assuming, for simplicity, that the ISM gas and star formation within $0.1R_{\text{vir}}$ are mostly confined to a flat disk of radius $R_{3\text{D}}$.

Instantaneous mass loadings drop off as Σ_{gas} increases. They also drop off with increasing Σ_{SFR} although there is more scatter, especially at low Σ_{SFR} . The bursts in the m12 halos (red points) show a clear evolution from high mass loadings at low Σ_{gas} (i.e., in their dwarf progenitors) to low mass loadings of ~ 0.1 at high Σ_{gas} (i.e., at low redshift). Most of the bursts in the m13 halos occur at rather large surface densities since these halos were already quite massive by $z = 2 - 4$. For purely illustrative purposes, we parameterize the trends as

$$\eta_{\text{M}} = 10^{2.71} \left(\frac{\Sigma_{\text{gas}}}{M_{\odot} \text{ pc}^{-2}} \right)^{-1.17} \quad (3.31)$$

and

$$\eta_{\text{M}} = 10^{-0.47} \left(\frac{\Sigma_{\text{SFR}}}{M_{\odot} \text{ yr}^{-1} \text{ kpc}^{-2}} \right)^{-0.53} . \quad (3.32)$$

The errors for the Σ_{gas} scaling are ± 0.08 dex (coefficient) and ± 0.04 (power law exponent). The errors for the Σ_{SFR} scaling are ± 0.05 dex (coefficient) and ± 0.02 (power law

exponent).

3.4.9 SF burstiness, dense ISM gas fractions and inner CGM virialization

The previous global correlations with M_* and V_{vir} are not satisfying in terms of painting a physical picture because they do not address how small-scale ISM conditions may influence the initial properties of winds during breakout. This interpretation-related ambiguity remains even in cases where the global correlations appear statistically strong with minimal scatter. On the other hand, the instantaneous loading factor trends (or lack thereof) with Σ_{gas} and Σ_{SFR} are also not sufficiently informative because they lack a proper normalization and hence physical context. Although we cannot establish causality with the FIRE-2 dataset (that would require controlled numerical experiments), we can at least correlate our instantaneous loading factors against a few relevant “derived” physical properties. In this first attempt, we choose to only focus on the following three for simplicity. Do more powerful starbursts (relative to the average SFR over a longer time window) drive winds that are more highly mass-loaded? Are instantaneous loading factors higher when dense ISM gas fractions are lower since that may enable winds to break out without as much impedance? Does the virialization of the inner halo correlate with the strength of ISM winds? The following is a brief heuristic and empirical exploration of these three questions.

To quantify starburst strength (or “local burstiness”) for each instantaneous outflow episode, we divide the maximum SFR within the burst window by the 1 Gyr-

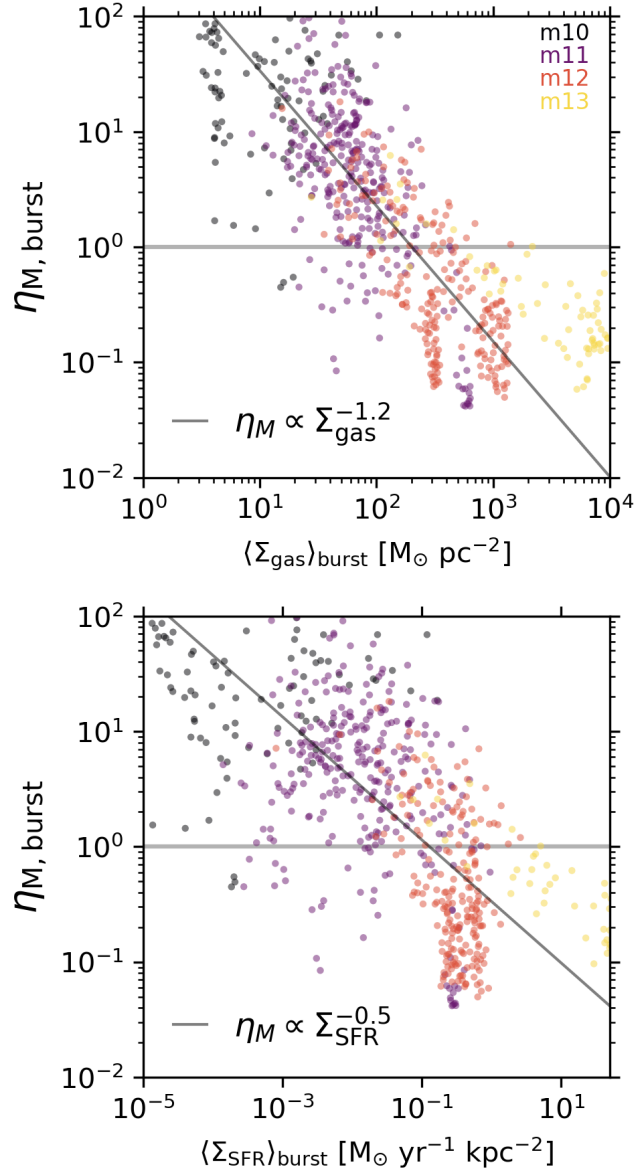


Figure 3.13: Instantaneous ISM mass loading factor versus \dot{M} -weighted average gas mass surface density (top) and SFR surface density (bottom) within individual burst windows. There is a clear negative correlation with gas mass surface density such that $\eta_M \propto \Sigma_{\text{gas}}^{-1.2}$ (diagonal black line; excluding the m13 halos from the fit). The trend with SFR surface density follows $\eta_M \propto \Sigma_{\text{SFR}}^{-0.5}$ but the deviation from a simple power law is more apparent.

averaged SFR (i.e., within each burst’s overall time chunk). The dense ISM gas fraction is computed as $f_{\text{dense}} = M_{\text{gas,dense}}/M_{\text{gas}}$ where $M_{\text{gas,dense}}$ is the mass of all gas particles within $0.1R_{\text{vir}}$ that have density $n > 1000 \text{ cm}^{-3}$ (this is the SF density threshold in FIRE-2).⁹ We take the \dot{M} -weighted average f_{dense} within each individual burst window. Finally, we take the $t_{\text{cool}}/t_{\text{ff}}$ ratio at $0.1R_{\text{vir}}$ from Stern et al. (2020), who analyzed the same simulations. This cooling time to free-fall time ratio is a measure of virialization in the inner CGM (specifically when $t_{\text{cool}}/t_{\text{ff}} \gtrsim 2$, the halo is virialized all the way down to the central galaxy). Following Stern et al. (2020), we do not include the low-mass (m10) dwarfs since they have $T_{\text{vir}} \sim 10^4 \text{ K}$ and the distinction between the dynamically hot and cool phases breaks down. As with f_{dense} , we estimate the \dot{M} -weighted average $t_{\text{cool}}/t_{\text{ff}}$ within each individual burst window.

Figure 3.14 shows our instantaneous mass loading factors as a function of the aforementioned three physical properties. The instantaneous mass loadings are clearly larger when starbursts are stronger (i.e., when the peak SFR is more prominent relative to the 1 Gyr-averaged SFR). In contrast, there is a lot of scatter and effectively no trend with f_{dense} , especially if we neglect the m13 halos which have large f_{dense} but low η_{M} (these halos are so massive that SN-driven winds cannot easily escape). Finally, the instantaneous mass loading steadily declines as the $t_{\text{cool}}/t_{\text{ff}}$ ratio gets larger, with $\eta_{\text{M}} \ll 1$ when the inner halo is virialized (i.e., when $t_{\text{cool}}/t_{\text{ff}} > 2$). This condition is met

⁹Our f_{dense} statistic is almost certainly too simplistic to capture the full complexity of the multi-phase ISM. A more robust measure of wind breakout conditions would take into account the full temperature–density distribution of the ISM to identify the warmer volume-filling phase fraction (e.g., Li & Bryan, 2020). However, it may be challenging to account for the complicated redshift and halo mass dependence of ISM geometry, multi-phase partitioning, and “contamination” of hot gas from the inner virialized CGM.




Figure 3.14: Instantaneous ISM mass loading factors versus three “derived” physical properties. From left to right: (1) the maximum SFR in the burst interval divided by the 1 Gyr-averaged SFR, (2) the dense ISM gas fraction, and (3) the cooling time to free-fall time ratio at $0.1R_{\text{vir}}$, which is a measure of inner halo virialization. Each of these physical properties are \dot{M} -weighted averages within individual burst windows. The large black symbols and errorbars denote binned medians with 25th and 75th percentiles. We see that instantaneous mass loadings are higher for more powerful starbursts (i.e., when the SFH is locally bursty). There is also a weak trend where instantaneous η_{M} tends to be higher when the dense ISM gas fraction is lower, although there is a lot of scatter. Finally, the instantaneous η_{M} steadily declines as $t_{\text{cool}}/t_{\text{ff}}$ increases, with the mass loading becoming $\ll 1$ when the inner halo is virialized as indicated by $t_{\text{cool}}/t_{\text{ff}} > 2$ (this is the case for massive halos whereas dwarfs have a non-virialized inner halo and higher mass loadings).

in the massive halos but not in the dwarfs (including the high-redshift dwarf progenitors of the m12 halos), which instead have high mass loadings and a non-virialized inner halo.

These trends will help inform our discussion later.

3.5 Halo wind loading factors

We now turn to halo-scale loading factors at R_{vir} . The driver of halo-scale outflows is more difficult to disentangle because there can be other input sources for mass, momentum, energy and metals in addition to the ISM outflows (e.g., CGM turbulence

stirred by satellite motions and their own outflows; Faucher-Giguère et al., 2016; Hafen et al., 2019, 2020). As a result, there may be ambiguities in interpreting “halo loading factors” which are computed as outflow fluxes in the virial shell normalized by reference fluxes on the ISM scale for type II SN inputs. However, we have verified through animations of the projected particle data that hot outflows generated by the central galaxy do often have enough energy to make it to R_{vir} , even in the MW halos at low redshift. Hence, it is informative to compare our broad redshift-averaged measurements of outflows at R_{vir} to those at $0.1R_{\text{vir}}$ (the large integration timescale means we are effectively marginalizing over complicated propagation and delay time physics).

3.5.1 Bernoulli velocity versus potential depth

Before presenting the halo loading factors, in Figure 3.15 we first compare the average mass-flux-weighted Bernoulli velocities ($v_B = \sqrt{\dot{E}/\dot{M}}$ following Equation 3.10) of multi-phase ISM outflows to the difference in escape velocity between $0.1R_{\text{vir}}$ and R_{vir} (which quantifies the halo potential depth). As outflows propagate outwards, they gain potential energy at the expense of kinetic and thermal energy; hence in the limiting case of adiabatic outflows, the decrease in Bernoulli velocity should mirror the decrease in escape velocity. Note that the upper limit on the Bernoulli velocity of SN-driven outflows is $\sqrt{\frac{10^{51}\text{erg}}{100M_{\odot}}} \approx 700 \text{ km s}^{-1}$; comparing this to the potential difference gives a simple estimate of whether SN-driven outflows can escape from halos of a given mass.

We see that cold and warm outflows contain just enough energy to make it to R_{vir} in the dwarfs and even the massive halos. On the other hand, the hot outflows

contain much more energy than needed to get to R_{vir} ; for high-redshift dwarfs, the energy of hot outflows is $\sim 5\times$ higher than the escape velocity difference, hence many of these outflows may become unbound from the halo. In the MW halos at low redshift, the hot outflows have just enough energy to make it to R_{vir} . This is also true for the m13 halos at high redshift but not at low redshift (where again, outflows may only reach $\sim 0.5R_{\text{vir}}$ in accordance with our wind selection criteria).

This exercise demonstrates that we should expect to see significant halo wind loading (especially for hot outflows in dwarfs) and that comparing characteristic outflow rates at R_{vir} to those at $0.1R_{\text{vir}}$ can help constrain average losses/gains in mass, momentum, energy and metals while winds transit the CGM. In Figure 3.16, we compare total mass, momentum, energy and metal loading factors in our ISM shell ($0.1 - 0.2R_{\text{vir}}$) to those in our virial shell ($1.0 - 1.1R_{\text{vir}}$). Both the ISM and halo loading factors in this figure only include outflows that have enough energy to get to at least $2R_{\text{vir}}$ if not farther.¹⁰ By defining this subset of “escaping” ISM and halo outflows, we can constrain what fraction of mass, momentum, energy and metals predicted to escape from the ISM to $2R_{\text{vir}}$ may actually do so. We will now describe each of these outflow quantities in turn.

¹⁰Note that these “escaping” ISM outflows are somewhat lower than our fiducial measurements to go from $0.1 \rightarrow 0.5R_{\text{vir}}$ since only a subset of ISM outflows will have the greater required initial energy to travel all the way to $2R_{\text{vir}}$. However, the overall trends are similar to our results above.




Figure 3.15: Average mass-flux-weighted Bernoulli velocity (i.e., specific kinetic energy plus enthalpy) of multi-phase ISM outflows versus the difference in escape velocity between $0.1R_{\text{vir}}$ and R_{vir} (a proxy for the potential difference). This gives a sense of whether outflows can be expected to reach R_{vir} in the absence of interactions (given our ISM wind selection criteria, outflows should make it to $\sim 0.5R_{\text{vir}}$ at minimum). Cold and warm outflows (left and middle panels) in some dwarfs have up to twice the energy needed to make it to R_{vir} , but in the more massive halos the cold/warm outflow energy is comparable to the potential difference. In contrast, hot outflows (right) have Bernoulli velocities that are far in excess of the energy needed to make it to R_{vir} . This is obvious for lower mass halos where the hot outflows contain up to $\sim 5\times$ more energy than needed to escape the halo (hence these outflows can be expected to travel very large distances, probably becoming unbound). Hot outflows in low-redshift MW halos have just enough energy to reach R_{vir} . Hot SN-driven outflows can also escape from m13 halos at high redshift but not necessarily at intermediate redshift.

3.5.2 Halo mass loading

We see that even for the low-redshift MW halos, the actual halo mass loading is comparable to, in fact even slightly larger than, the ISM mass loading defined using particles with enough energy to make it to $2R_{\text{vir}}$. If we had included slower moving, likely cold and turbulent, ISM outflows – which never had a chance of getting to $2R_{\text{vir}}$ anyway – then this ratio would be closer to 0.3 – 0.4 (Figure 12 of Pandya et al., 2020). While we do not know whether the identity of the gas leaving the virial shell is the same as the gas that was previously ejected from the ISM (e.g., much of the ISM outflows could have stalled in the CGM while still pushing ambient halo gas outwards), our finding that $\eta_{\text{halo}}/\eta_{\text{ISM}} \sim 1$ in the low-redshift MW halos combined with their relatively large Bernoulli velocities of hot outflows in Figure 3.15 suggests that outflows can have substantial effects in MW halos (see also our supplementary movies, e.g., Figure 3.1). This agrees with the conclusions drawn from the comparative CGM analysis of diverse simulations by Fielding et al. (2020b). For dwarfs, the halo mass loading is generally many times higher than the ISM mass loading expected to make it to $2R_{\text{vir}}$. Since the dwarfs are generally very isolated, this is likely due to ISM outflows sweeping up CGM gas (see also Muratov et al., 2015; Pandya et al., 2020). In contrast, for the m13 halos, the much higher halo mass loadings than expected are likely due to their rich satellite systems, which can stir up the CGM and have substantial outflows of their own (e.g., Anglés-Alcázar et al., 2017a; Hafen et al., 2020). While quantifying these entrainment and satellite effects is beyond the scope of this chapter, the time evolution of the radial

profile of \dot{M}_{out} and \dot{M}_{in} in our supplementary movies (as in Figure 3.1 and Figure 3.2) can qualitatively reveal these effects. For example, the amplitude and/or width of an outflow spike may increase as it propagates to larger radius, which would be indicative of CGM entrainment.

3.5.3 Halo momentum loading

In the dwarfs, the halo momentum loadings are larger than the ISM momentum loadings, which is expected for energy-conserving outflows (if $\dot{E} \sim \dot{M}v^2$ is roughly constant, then $\dot{p} \sim \dot{M}v$ will increase as the outflow decelerates due to sweeping up mass). Interestingly, the MW halos at low redshift have roughly similar outflow momentum at the halo and ISM boundaries. The m13 halos have anomalously high halo momentum loading factors, which may suggest additional momentum input sources (e.g., their rich satellite systems).

3.5.4 Halo energy loading

The halo energy loadings are comparable to ISM energy loadings in the dwarfs (including the high redshift progenitors of MW halos). This suggests that the relatively higher ISM energy loadings of dwarfs (Figure 3.9) are conserved to at least R_{vir} , which is consistent with the large Bernoulli velocities relative to their potential depth (Figure 3.15). The m13 halos also have high halo energy loadings but there is likely significant contamination from their rich satellite systems (which can introduce additional kinetic and thermal energy from stirring turbulence in the CGM, heating from their

own energy-rich outflows, etc.). In contrast, for the low redshift MW halos, the halo energy loading factors are only ~ 0.25 times their ISM energy loadings (i.e., 4 times lower).

3.5.5 Halo metal loading

In the MW halos at low redshift, the metal loading at R_{vir} is roughly $\sim 20\%$ of the ISM-scale metal loading. In contrast, for their dwarf progenitors and all dwarfs more generally, the halo metal loadings are comparable to the ISM metal loadings. This is consistent with the interpretation by Muratov et al. (2017) for FIRE-1 that metals escape from dwarf halos but are retained within MW halos (perhaps due to substantial interactions in the latter, although we also do not know the level of mixing with pristine gas). The m13 halos also have relatively low halo metal loadings perhaps owing to their deeper potential wells, although again their halo-scale measurements are likely contaminated due to metal-rich outflows from their large satellite systems (see also Hafen et al., 2019).

3.6 Discussion

Here we summarize the overall story suggested by our results, discuss our findings in the context of previous work, and list some possible systematic uncertainties in our analysis.

The results of our analysis tell a seemingly simple story. We have found that dwarfs have preferentially much higher ISM mass, momentum, energy and metal load-

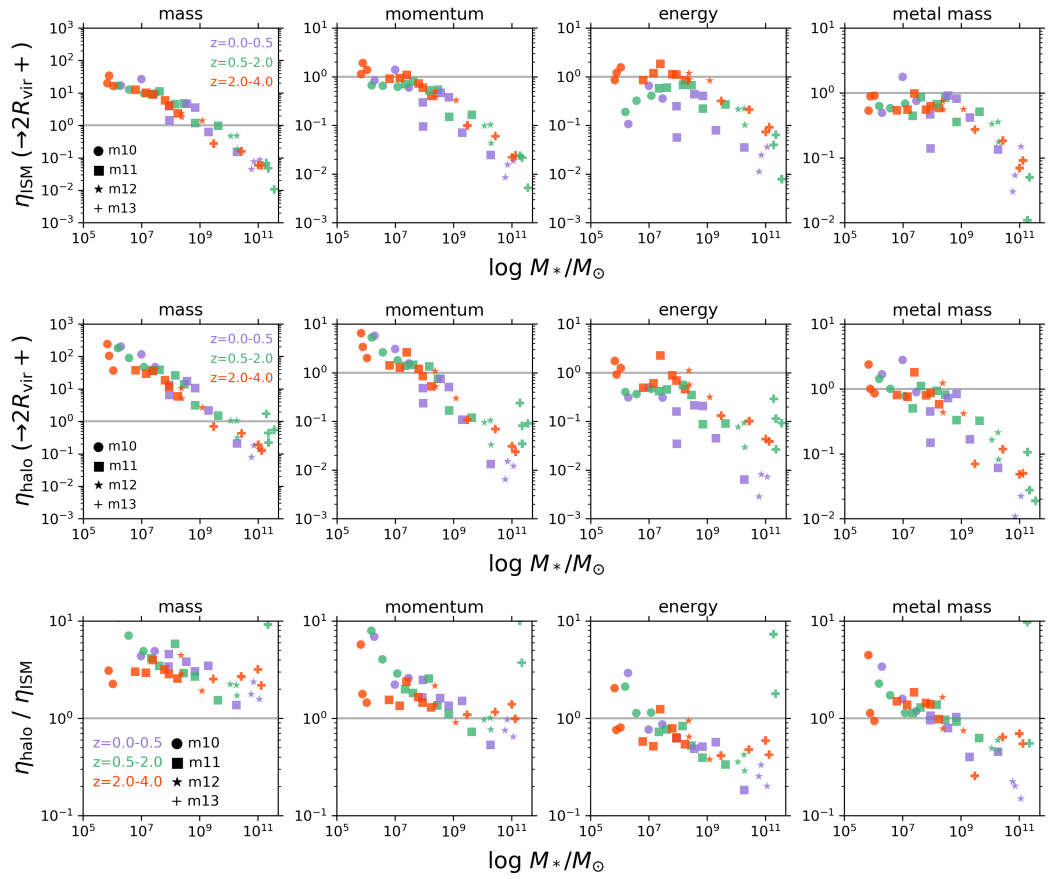


Figure 3.16: Comparing “escaping” wind loading factors for the ISM and virial shells. The top row is ISM loading factors using only the subset of ISM outflows that have enough energy to reach $2R_{\text{vir}}$ instead of just $0.5R_{\text{vir}}$. The middle row is halo-scale loading factors using only outflows at R_{vir} that have enough energy to get to at least $2R_{\text{vir}}$. The bottom row is the ratio of these halo-scale and ISM-scale loading factors. From left to right we show mass, momentum, energy and metal loading factors. *Mass*: dwarf halo mass loadings are a few times larger than the ISM mass loadings, perhaps indicative of additional swept up material. Low-redshift MW halos have a ratio close to ~ 1 (recall that our ISM loadings exclude slower outflows, which may be substantial but unlikely to reach R_{vir}). m13 halos show a larger ratio at intermediate redshift than at high redshift. *Momentum*: Dwarf outflows often have more momentum at the halo scale than at the ISM scale, in contrast to low-redshift MW halos whose outflows have comparable momentum at R_{vir} and $0.1R_{\text{vir}}$. *Energy*: In dwarfs, halo energy loadings are comparable to ISM energy loadings. In contrast, for MW halos at low redshift, the halo-scale energy loadings are ~ 0.1 smaller than their ISM-scale energy loadings. *Metals*: In dwarf halos, most metals leaving the ISM also leave the halo. In low-redshift MW halos, only $\sim 10\%$ of ISM metal outflows leave the halo (surprisingly, intermediate redshift m13 halos have high metal loadings).

ings than MW-mass halos at late times (and even the m13 halos at early times). The cold outflow phase is generally negligible for dwarfs except at high redshift where the cold phase can account for $\sim 10\%$ of each of the total loading factors.¹¹ The importance of the warm phase gradually increases toward lower stellar masses (for which the warm phase approaches $\sim 100\%$ by mass fraction). The suppression of multi-phase outflows in the lowest mass dwarfs may be a clue that the UV background together with the global thermodynamics of the halo (the virial temperatures of these dwarfs is much lower than our threshold of 10^5K for the hot phase) either prevents thermal instabilities or rapidly heats up cold outflows due to CGM mixing and/or shocks. In addition, much of the ISM of dwarfs may already be at a warm temperature, so significant cold mass loading may not be expected. In any case, it is remarkable that the overall momentum, energy and metal loadings are of order unity in the lowest mass dwarfs, implying that most of the SN-driven energy, momentum and metals make it quite far out of the ISM; the mass loadings being of order 100 also suggests that the outflows sweep up significant amounts of ambient material. The metal loadings being of order unity suggests that the ISM metallicity of dwarfs is in equilibrium (Forbes et al., 2014) since most of the metals produced as SN ejecta escape via metal-enriched, energy-conserving outflows (hence the ISM metallicity should be roughly constant with time). Note that the FIRE simulations have been shown to agree reasonably well with observed mass–metallicity relations for both gas and stars in the mass ranges that we examine here (Ma et al.,

¹¹Observed outflows driven by active galactic nuclei (AGN) can have high cold mass loading factors (e.g., Cicone et al., 2014; Fiore et al., 2017; Fluetsch et al., 2019, and references therein), but the FIRE-2 simulations do not include AGN feedback.

2016; Wetzel et al., 2016; Escala et al., 2018). The ratio $\eta_{\text{halo}}/\eta_{\text{ISM}} \gtrsim 1$ for the dwarfs, further suggesting that ISM outflows escape to quite large distances ($\gtrsim R_{\text{vir}}$) on average with their energy, momentum and metals intact.

In contrast, for low-redshift MW halos and high-redshift massive (m13) halos, winds are weaker and the hot phase generally carries most of the mass, momentum, energy and metals.¹² The warm phase is subdominant (though it can carry a substantial fraction of metals in the low-redshift MW halos; see purple stars in Figure 3.10) and the cold phase is generally negligible (a few percent by mass fraction). The loading factors for the low-redshift m12 halos are below unity (η_M is of order ~ 0.1 on average, and possibly smaller for individual weak outflow episodes), which means that only a fraction of the SN-driven mass, energy, momentum and metals make it out of the ISM (unlike for the dwarfs). Nevertheless, the ratio $\eta_{M,\text{halo}}/\eta_{M,\text{ISM}} \sim 1$ for the MW halos at low-redshift, suggesting that whenever there is a large breakout of wind from the ISM, there is subsequently also a large outflow from the halo. However, the $\eta_{\text{halo}}/\eta_{\text{ISM}}$ ratio is far below unity for energy and metals, meaning that a large fraction of wind energy is dissipated while metals are mixed into the CGM due to interactions (or the outflow metallicities are diluted due to sweeping up of metal-poor CGM gas).¹³ Interestingly, the $\eta_{p,\text{halo}}/\eta_{p,\text{ISM}}$ is closer to 1 for the low-redshift MW halos, and may be driven by the thermal pressure term which would be substantial for their predominantly hot outflows.

¹²Had we only used the simpler $v_{\text{rad}} > 0$ km/s cut, we would select substantially more warm outflows. However, some fraction of these may not travel far beyond $\sim 0.1R_{\text{vir}}$ and may represent random motions near the ISM edge.

¹³Many of the m13 halos have $\eta_{\text{halo}}/\eta_{\text{ISM}}$ ratios greater than unity, which is unexpected given their deep potential wells. We think this is likely due to additional input sources of mass, momentum, energy and metals at large radii. Possible sources include outflows and turbulence stirred by their numerous satellites as well as accretion shocks of infalling gas near R_{vir} .

3.6.1 Comparison to theoretical expectations and other simulations

3.6.1.1 Comparison to simple theoretical arguments

Traditionally, mass loading factors are correlated against the global halo circular velocity since that is a proxy for the potential depth and because the inferred power law slope may encode whether the winds are “energy-driven” ($\eta_{\text{M,ISM}} \propto V_{\text{vir}}^{-2}$) or “momentum-driven” ($\eta_{\text{M,ISM}} \propto V_{\text{vir}}^{-1}$; e.g., Murray et al., 2005; Hopkins et al., 2012; Muratov et al., 2015; Christensen et al., 2016). We find that $\eta_{\text{M,ISM}} \propto V_{\text{vir}}^{-2}$ at high redshift, with a significant steepening at low redshift. There appears to be no need to appeal to a “broken” power law as found for the FIRE-1 halos by Muratov et al. (2015). Our power law (particularly at high redshift) is consistent with simple theoretical expectations for energy-conserving winds as laid out in Murray et al. (2005). At lower redshifts, our relation becomes even steeper, consistent with a picture in which there are significant losses in the ISM prior to the wind being launched.

In addition to correlating the loading factors against the global halo virial velocity and stellar mass, it is important to consider instantaneous correlations with properties that explicitly characterize the state of the ISM and inner halo (e.g., as suggested by Fielding et al., 2017a; Li & Bryan, 2020). After all, the virial velocity and stellar mass correlations alone do not unambiguously explain what sets the properties of winds upon initial breakout from the ISM. For example, why do winds in high-redshift dwarfs appear to be energy-conserving (i.e., consistent with a V_{vir}^{-2} scaling)? While painting a fully fleshed out physical picture is beyond the scope of this work, we

found three important trends that can help guide future work using controlled numerical experiments. First and foremost, instantaneous mass loading factors are preferentially higher during more powerful starbursts (i.e., when the peak SFR is more prominent compared to the 1 Gyr-averaged SFR). During such locally bursty SF events, we may expect more strongly clustered SNe (Faucher-Giguère, 2018). The resulting powerful stellar feedback may clear out the denser phase of the ISM while percolating through the less dense phase, ultimately breaking out of the galaxy prior to losing significant energy via radiative cooling (e.g., Fielding et al., 2018).

Correlating instantaneous mass loading factors with dense ISM gas fractions reveals a lot of scatter and effectively no trend, especially if we ignore the m13 halos at high f_{dense} with low η_{M} (SN-driven winds cannot easily escape from these massive halos). The lack of a strong correlation with f_{dense} may reflect the fact that more powerful starbursts are also expected to occur when dense ISM gas fractions are higher, and these in turn may drive more powerful winds despite high f_{dense} . On the other hand, our overly simplistic definition of f_{dense} using only particles with $n > 1000 \text{ cm}^{-3}$ may not be the best diagnostic of ISM breakout conditions: if the warmer volume-filling ISM phase fraction can be reliably measured, that may lead to a more robust correlation (Li & Bryan, 2020). On a related note, the ISM may be more turbulent when the overall gas fraction $M_{\text{gas}}/(M_{\text{gas}} + M_{*})$ is higher, which may make it easier to drive strong outflows (this may help explain why winds become weaker in more massive halos at later times, when their overall gas fractions have decreased; Hayward & Hopkins, 2017). Finally, we find that instantaneous mass loadings are suppressed when the inner halo is virialized

(as in the more massive halos) (Stern et al., 2020). The lack of a virialized inner CGM in dwarfs may allow outflows to propagate relatively unimpeded with minimal energy and momentum losses. Shock heating and entrainment of CGM/IGM gas by these energy-conserving outflows may cause preventative feedback that can suppress future gas accretion and ultimately help reduce the global star formation efficiency of dwarfs (Pandya et al., 2020). Despite this interesting heuristic exercise, we stress that our instantaneous burst analysis groups together outflow episodes occurring in halos of widely different masses and across ~ 10 Gyr of cosmic time. It is an enormously challenging task to simultaneously control for all of the possible interplay between global and local conditions using a fully cosmological simulation, but it is encouraging that we at least see some emergent systematic trends with our simple summary statistics.

3.6.1.2 Comparison to high-resolution idealized simulations

It is difficult to say definitively how our FIRE-2 wind scalings compare to those from resolved ISM idealized simulations. A future analysis of the FIRE-2 simulations closer to the ISM while accounting for the complicated geometries of our galaxies can help place these kinds of comparisons on a firmer footing (e.g., following Gurvich et al., 2020). One seemingly major difference worth commenting on is that Kim et al. (2020a) very clearly predict that cool outflows (with $T < 2 \times 10^4$ K) carry most of the mass whereas hot outflows (with $T > 5 \times 10^5$ K) carry most of the energy in their TIGRESS kpc-scale sub-galactic simulations. In contrast, for our low-redshift MW halos, the hot phase carries both most of the mass and energy (Figure 3.7 and Figure 3.9). The

simplest explanation is that our measurements are made much farther from the galaxy ($0.1R_{\text{vir}}$) than in resolved ISM simulations ($\sim 1 - 2$ disk scale heights), and much of the cold and warm outflows may not be expected to make it to $\sim 0.5R_{\text{vir}}$ anyway. Instead they may recycle as fountain flows much closer to the disk (e.g., Anglés-Alcázar et al., 2017a; Hafen et al., 2020; Gurvich et al., 2020) or get mixed into the hot phase (e.g., Fielding et al., 2020a; Schneider et al., 2020). Note also that FIRE includes additional prescriptions for radiative pressure feedback and photoionization that are not modeled in TIGRESS but which may be crucial for heating the ISM and enabling breakout of hot winds. Nevertheless, it is encouraging that the hot mass loadings of our low-redshift MW halos are still only of order ~ 0.1 , which is similar to the TIGRESS predictions (Kim et al., 2020a).

A related question is why the hot mass loading in our dwarfs is far larger than ~ 0.1 , in fact closer to ~ 10 (the overall mass loading is ~ 100 with a $\sim 10\%$ hot phase fraction for low M_* galaxies; see Figure 3.7). Partially, this may be due to shock heating and entrainment of inner CGM gas by energy-conserving outflows in the dwarfs. The warm phase is even more prominent than the hot phase for winds in the FIRE-2 dwarfs. This may be due to the fact that warm outflows may be able to travel farther into the CGM of dwarfs because of their shallower potential well depths. There can also be a disproportionately larger contribution of swept-up warm ISM and inner CGM gas for outflows in dwarfs compared to more massive halos. Note that since the virial temperatures of dwarfs can be lower than our hot phase threshold temperature of 10^5 K, much of the warm outflows in dwarfs can still be considered “dynamically hot”

(and vice versa for more massive halos). The idealized, high-resolution global dwarf simulations by Hu (2019) show that the warm phase (what they call the ionized phase) is indeed very important: it is the dominant phase beyond a few kpc owing to cooling of the hot phase and shock heating of cooler gas.

As for trends between instantaneous wind loading factors and ISM and SFR surface densities, the clearest correlation we have found is between η_M and the \dot{M} -weighted average Σ_{gas} over individual burst windows. The instantaneous mass loading tends to drop systematically as Σ_{gas} increases, becoming of order ~ 0.1 in the low-redshift MW-mass halos. This is qualitatively consistent with predictions from TIGRESS where the authors find $\eta_M \propto \Sigma_{\text{gas}}^{-1.12}$, albeit much closer to the galaxy (one scale height above/below the disk; Figure 12 in Kim et al., 2020a). The correlation with Σ_{SFR} shows more scatter. This may partially be driven by the fact that we are combining bursts from widely different halos and at many different redshifts (up to $z = 4$). We also measure our loading factors farther from the ISM than sub-galactic simulations do; in fact, our chosen distance of $0.1R_{\text{vir}}$ corresponds to $\sim 25 - 30$ physical kpc for a low-redshift MW-mass halo, which is far beyond the simulation domain of kpc-scale resolved ISM models. Since the properties of a wind may change as it propagates through the inner halo, it is perhaps natural to expect different correlations with more scatter farther from the galaxy (just like we expect halo-scale loadings to be more complicated to interpret). A fruitful avenue for future work will be to combine measurements of loading factors very close to the galaxy (ideally by defining subpatches of the ISM and properly accounting for the more complicated gravitational potential)

with our spherically-averaged loading factors farther out (see also Gurvich et al., 2020).

3.6.1.3 Comparison to other cosmological simulations

Lastly, it is useful to qualitatively discuss our results in the context of other cosmological simulations (both zooms and large-volume). The FIRE-2 simulations are particularly unique for predicting wind properties in a cosmological context due to their explicit stellar feedback model. In contrast, the modeling of SN-driven winds across different cosmological simulations varies dramatically and often involves ad hoc approaches (e.g., decoupling winds from hydrodynamics, artificially delayed cooling, etc.).

Compared to the FIRE-1 results of Muratov et al. (2015, 2017) that we build on (see also Anglés-Alcázar et al., 2017a), our overall FIRE-2 mass and metal loadings scale with stellar mass in qualitatively similar ways (despite our more stringent wind selection criteria; Figure 3.5). However, we have gone further and provided several new insights by explicitly measuring outflow energy and momentum loadings, temperature and velocity distributions, and scalings with quasi-local ISM properties. This allowed us to explicitly demonstrate that winds in dwarfs seem to be energy-conserving whereas the more massive halos show significant outflow energy losses, especially at low redshift. Interestingly, Christensen et al. (2016) also find that outflows in their GASOLINE zoom-in simulations are consistent with the simple energy-driven scaling ($\eta_M \propto V_{\text{vir}}^{-2}$; their Figure 11). Tollet et al. (2019) find a much steeper relation in the NIHAO zoom-in simulations: $\eta_M \propto V_{\text{vir}}^{-4.6}$, which they attribute to the reduced efficiency of SN feedback

in more massive halos. They also cut off their steep scaling for dwarfs with $V_{\text{vir}} < 45$ km/s since there is a lot of scatter which they claim is due to stochastic SF. Instead, they argue that for these dwarfs, the mass loadings must revert to following at most a predicted V_{vir}^{-2} scaling because there is a “maximum” efficiency of SN feedback and because most of their dwarf outflows are cold with radial velocities below the escape velocity. With our more stringent wind selection criterion (which captures the slow component of the hot wind while neglecting cold, turbulent outflows), we find relatively little scatter for the lowest mass dwarfs, which suggests that the steepening of the overall relation with time is indeed driven by higher mass halos. Note that the prominent redshift dependence in Figure 3.6 but weaker redshift dependence when plotting against stellar mass (Figure 3.7) can at least partially be explained by the stellar-to-halo-mass ratio, at fixed halo mass, getting larger at later times (at fixed M_{vir}) whereas V_{vir} does not evolve as dramatically. The redshift evolution in the stellar-to-halo-mass ratio is particularly prominent for more massive halos since SF is so inefficient in dwarfs, and so the steepening against virial velocity with time may be driven by the m12 halos.

As for large-volume simulations, energy and momentum loadings are generally not measured. One notable exception is Mitchell et al. (2020) who measured mass loading factors as well as radial profiles of energy and momentum outflow rates in the EAGLE simulations. They find that $\eta_{\text{M}} \propto V_{\text{vir}}^{-1.5}$ for lower mass halos where stellar feedback dominates, with the normalization increasing towards higher redshift; this scaling steepens for halo-scale mass loadings to roughly match $\propto V_{\text{vir}}^{-2}$ as expected for energy-driven winds. They attribute this steepening to entrainment of CGM gas by

winds, as inferred from average radial profiles of mass, momentum and energy outflow rates. Another interesting result of modern large-volume simulations is that they predict multi-phase winds, as recently shown for the IllustrisTNG model by Nelson et al. (2019). Despite the very different SN feedback subgrid models of FIRE-2 and IllustrisTNG (decoupled kinetic wind model; Springel & Hernquist, 2003), it is encouraging that the outflow temperature distributions vary in similar systematic fashions with stellar mass, in particular that the cold phase is noticeably absent in the lowest mass dwarfs and more prominent in higher mass halos. This relative agreement likely reflects the fact that the physics of radiative cooling, which is similar amongst all these simulations, is predominantly responsible for setting the general properties of the outflow temperature distributions. However, the inclusion of AGN feedback in many large-volume simulations combined with a more phenomenological treatment of winds (plus different wind analysis methods) makes it difficult to draw direct comparisons with FIRE-2. We do note that the EAGLE (Schaye et al., 2015, their equation 6) and IllustrisTNG (Pillepich et al., 2018, their Figure B2) large-volume simulations assume that winds carry less thermal energy when the ISM metallicity is larger, which requires that the emergent energy loading factors should be lower in more massive halos and at later times. While we do not explicitly consider scalings with ISM metallicity, their assumption is consistent with our overall finding. A more granular analysis of winds closer to the ISM in FIRE-2 would shed further light on this assumption, including the role of ISM gas metallicity and density in setting the redshift dependence of some of our wind scalings.

In the future, it will be insightful to compare to large-volume simulations that

“plug in” scalings for mass loadings and wind velocities taken from zoom-in simulations (Davé et al., 2016, 2019; Huang et al., 2020). Our comprehensive multi-dimensional characterization of FIRE-2 winds in terms of their temperature and velocity distributions, and their energy and momentum loadings, will serve as useful inputs and benchmarks for large-volume models with insufficient resolution to capture stellar feedback. Our FIRE-2 scalings can be implemented in SAMs as is classically done (e.g., ISM mass loading factor versus global halo virial velocity), and it may also be possible to use our relations in radially-resolved SAMs where outflow properties are varied according to local ISM properties (such as the gas mass surface density; Forbes et al., 2019). Perhaps most importantly, SAMs will benefit from implementing our multi-phase energy and momentum loading factors, which can be used to drive CGM heating and hence suppress ISM accretion, push ambient CGM gas out of the halo via entrainment, and prevent IGM gas from accreting into the halo in the first place (see more discussion of preventative stellar feedback models in Pandya et al., 2020).

3.6.2 Systematic uncertainties

Although we have taken the first step to characterize the full thermodynamic properties of outflows in fully cosmological simulations (adapting analysis methods commonly used for high-resolution idealized simulations; e.g., Kim et al., 2020a), there are several sources of systematic uncertainty that may impact our results and interpretation. Many of these, which we list here, may be fruitful avenues for future work.

First and foremost, our ISM loading factors are measured with a shell at 0.1 –

$0.2R_{\text{vir}}$, which corresponds to nearly ~ 25 kpc in a MW halo at $z = 0$. This was done partly for simplicity (we can use spherical shells, avoid contamination from dense ISM mass flows, ignore the highly non-trivial geometry of galaxies, especially high-redshift dwarfs, etc.), but this is very far from the midplane of the disk and far beyond the domain of highly resolved ISM simulations (e.g., the ~ 1 kpc scale boxes simulated by Kim et al., 2020a). While it is fundamental to know the properties of outflows this far out near the inner CGM, it is somewhat ambiguous what fraction of our ISM outflows are fresh from the MW versus swept up inner CGM material. A future analysis that considers the properties of outflows directly above and below the galaxy can provide many useful physical insights and sanity checks (this can be done easily at least for MW halos at intermediate and lower redshifts when they have a well-defined disk, but dwarf geometries are more complicated). In parallel, additional metrics for characterizing the inner CGM beyond the $t_{\text{cool}}/t_{\text{ff}}$ proxy of Stern et al. (2020) may help us better understand the role of the CGM in modulating outflows and its own susceptibility to being heated/entrained. Combined with a particle tracking approach (e.g., Anglés-Alcázar et al., 2017a; Hafen et al., 2020), we would also be able to more confidently constrain the distribution of particle travel times, maximum distances, recycling times, and whether particles expected to conserve energy/momentum actually do so.

Furthermore, the fixed mass resolution of the Lagrangian FIRE-2 simulations may lead to unresolved cooling of hot outflows and propagation of cold outflows. We have found that in MW-mass FIRE-2 halos at low redshift, the hot phase carries not only most of the energy but also most of the mass. In highly resolved ISM simulations

closer to the disk, the hot phase carries most of the energy whereas the cool phase carries most of the mass. If we are not resolving the cooling of hot gas, then colder clumps that should form and become entrained in the CGM (leading to smaller outflow rates measured at R_{vir}) may not be properly captured. We have also seen that cold outflows are heavily suppressed except in high-redshift dwarfs; this may be physical in the dwarfs, but it could also partially be due to poor resolution in the CGM, especially in the more massive halos at low-redshift. On the other hand, we emphasize that our wind selection criteria exclude outflows that do not have enough starting energy to reach at least $0.5R_{\text{vir}}$; most of this excluded material is almost certainly slow, cold outflows. In addition, the FIRE-2 subgrid approach of depositing momentum in lieu of thermal energy when the Sedov-Taylor SN phase is unresolved may play a role in setting $\eta_E \ll 1$ for the more massive halos (but see Hopkins et al., 2018b). Recall that the mass resolution in the low-mass (m10) dwarfs is $\sim 250M_{\odot}$ so the Sedov-Taylor phase is likely well-resolved, whereas in the more massive halos it is not ($\sim 2100 - 33000M_{\odot}$). Note, though, that when SNe are clustered, most of the hot gas may be contained in “superbubbles” which are better resolved.

Finally, there are other sources of mass, momentum, energy and metal input beyond type II SNe that are not included in our analytic reference injection rates (but which are included in the FIRE-2 simulations), so our energy and momentum loading factors may be overestimated.¹⁴ It may not be possible to estimate, in a clean way,

¹⁴On a related note, the core FIRE-2 simulations that we use do not include cosmic ray physics, which can significantly affect outflow and CGM properties (Hopkins et al., 2020). Our simulations also do not include a subgrid model for turbulent diffusion which would otherwise allow metals to no longer strictly follow mass. While this can affect the distribution of outflow metallicities for a given episode, our bulk

total injection rates in cosmological simulations due to the number of processes, many of which are approximated using complicated subgrid models. For example, we do not account for type Ia SNe, radiative heating from the stars in the central galaxy, outflows and turbulence stirred by satellites, or gravitational shock heating of infalling gas in more massive halos. We also do not account for the possibly complicated effect of the FIRE-2 stellar feedback subgrid model: a correction to the loading factors may be warranted to account for the lack of thermal energy input and reliance on momentum deposition alone during the unresolved Sedov-Taylor phase of SNe. All of these effects make it more complicated to interpret the absolute values of η_E and η_p , especially when they are high. In addition, the m13 halos, which tend to be outliers in some of our relations, are only run down to $z = 1$ so we are missing roughly half of their evolution in our intermediate redshift bin (down to $z = 0.5$). It is possible that the peculiarly high outflow rates of the m13 halos may decrease significantly at $z < 1$ (as happens for the m12 halos at later times).

3.7 Summary

We have characterized the mass, momentum, energy and metal loading factors of multi-phase galactic winds in the FIRE-2 cosmological “zoom-in” simulations. To accomplish this, we implemented a physically motivated Bernoulli velocity wind selection criterion to account for the bulk kinetic, thermal and potential energy of gas

shell-averaged measurements may be robust (see arguments in Muratov et al., 2017, end of their section 5.4).

particles and exclude slower, turbulent moving outflows from genuine winds. We report instantaneous outflow measurements at two characteristic radii: close to the ISM boundary ($0.1 - 0.2R_{\text{vir}}$) and the halo boundary ($1.0 - 1.1R_{\text{vir}}$). Given the inherently multi-phase nature of galactic winds, we computed loading factors separately for the cold ($T < 10^3\text{K}$), warm ($10^3\text{K} < T < 10^5\text{K}$) and hot ($T > 10^5\text{K}$) phases. In order to minimize systematics due to travel time delays, entrainment, etc., our fiducial loading factors were measured as averages over three relatively large redshift bins: low redshift ($z = 0.0 - 0.5$), intermediate redshift ($z = 0.5 - 2.0$) and high-redshift ($z = 2.0 - 4.0$). We also implemented a robust algorithm to derive instantaneous loading factors for the ISM shell to complement our redshift-averaged measurements and explore correlations with physical properties on short timescales. With the large sample size of the core FIRE-2 suite, we analyzed halos in four mass bins: low-mass dwarfs ($M_{\text{vir}} \sim 10^{10}M_{\odot}$ at $z = 0$), intermediate-mass dwarfs ($\sim 10^{11}M_{\odot}$ at $z = 0$), MW-mass galaxies ($\sim 10^{12}M_{\odot}$ at $z = 0$), and more massive halos at high redshift ($\sim 10^{12.5} - 10^{13}M_{\odot}$ by $z = 1$).

Our main takeaways are as follows:

1. The ISM mass loading factor is preferentially higher for dwarfs (of order ~ 100) compared to more massive halos (below unity). Cold mass loading fractions are negligible in all halos except high redshift dwarfs where it approaches order unity. Warm mass loading fractions dominate over cold and hot mass loading fractions in dwarfs, whereas hot outflows carry most of the mass in the more massive halos. Similarly, the ISM momentum, energy and metal loadings are of order unity in

the dwarfs (especially at high redshift) and significantly lower in the more massive halos. The warm phase tends to carry most of the momentum, energy and metals in the dwarfs whereas the hot phase dominates for the more massive halos. Correlating total ISM mass loadings with the global halo virial velocity results in a V_{vir}^{-2} dependence at high-redshift (consistent with energy-driven winds), but a steeper scaling at later times.

2. The average Bernoulli velocity of hot outflows is $2 - 5\times$ the difference in gravitational potential between $0.1R_{\text{vir}}$ and R_{vir} , especially in high-redshift dwarfs, meaning that we should expect to see substantial outflows at R_{vir} . Indeed, mass outflow rates at R_{vir} are several times larger than mass outflow rates at $0.1R_{\text{vir}}$ in the dwarfs, indicative of swept up CGM gas. In the low-redshift MW halos, this $\eta_{\text{M,halo}}/\eta_{\text{M,ISM}}$ ratio is also of order unity when we consider only “escaping” ISM outflows. Energy outflow rates at R_{vir} are comparable to those at $0.1R_{\text{vir}}$ in dwarfs whereas this $\eta_{\text{E,halo}}/\eta_{\text{E,ISM}}$ ratio is much lower (~ 0.25) in low-redshift MW halos. Halo-scale momentum loading factors exceed ISM momentum loading factors in dwarfs (as expected for energy-conserving outflows) but are comparable in MW-mass halos at later times. Most of the metals that leave the ISM tend to escape from dwarf halos but are retained within low-redshift MW-mass halos.
3. Correlating instantaneous wind loading factors with \dot{M} -weighted physical properties over individual burst windows reveals a few interesting trends. Instantaneous η_{M} shows a clear negative correlation with Σ_{gas} but there is substantially more

scatter versus Σ_{SFR} . In contrast, we see a clear positive correlation between the instantaneous η_{M} and a measure of how locally bursty a SF episode is (defined as the peak SFR within a burst interval divided by the 1 Gyr-averaged SFR). We see a lot of scatter and effectively no correlation between η_{M} and f_{dense} , which may reflect competing trends between how the dense ISM gas fraction affects starburst and wind strengths, and/or that our simple f_{dense} statistic is not an ideal measure of ISM wind breakout conditions. Finally, we see a strong negative correlation between η_{M} and $t_{\text{cool}}/t_{\text{ff}}$ (which is larger than two when the inner halo is virialized): mass loading is preferentially suppressed when the inner halo is virialized (as is the case in massive halos at later times but not in dwarfs or at high redshift).

Our results suggest that the reduced global star formation efficiency of dwarfs may at least partially be driven by their more powerful winds. At the same time, our comprehensive analysis has revealed the multi-phase nature of weaker SN-driven winds in massive halos. Our findings can be used to guide future controlled numerical experiments that aim to clarify the key parameters that determine the properties of galactic winds. In future work, we will use this rich dataset to implement preventative stellar feedback in next-generation SAMs. The traditional approach of relying on mass and metal ejection alone can be improved upon by also considering the energy and momentum injected into the CGM/IGM by SN-driven winds. This may have important physical implications for CGM/IGM heating rates and observable consequences for the redshift evolution of the mass-metallicity relation, the stellar-to-halo-mass relation for

dwarfs, and the chemical enrichment of the CGM/IGM.

Chapter 4

A new semi-analytic model with explicit preventative stellar feedback

4.1 Introduction

This chapter combines insights from the previous chapters and presents our new prototype SAM in which nearly all of the free parameters and functional forms are constrained from the core FIRE-2 simulations (Hopkins et al., 2018a). In Pandya et al. (2020), we speculated that the energy of SN-driven winds may be responsible for shock-heating the CGM and IGM of FIRE-2 dwarfs, leading to their very low measured accretion rates. With our measurements of wind energy from Chapter 3, we can ask whether the SN-driven winds generated in the FIRE-2 simulations are indeed a viable heating source to offset CGM and IGM cooling in SAMs.

This chapter is organized as follows. Section 4.2 briefly describes the FIRE

data and our updated particle tracking measurements of inflow and outflow fluxes. Section 4.3 introduces the prototype SAM. Section 4.4 presents a case study of CGM/IGM heating in a low-mass dwarf. Section 4.5 presents results for the entire FIRE-2 ensemble, including our SAM-derived heating efficiency functions. We discuss our results in Section 4.6 and conclude in Section 4.7. We adopt a standard (Planck Collaboration et al., 2016) cosmology throughout.

4.2 FIRE-2 Simulations: Sample and analysis

4.2.1 Simulation sample

We use the same sample of “core” FIRE-2 halos as in Pandya et al. (2020, 2021). We exclude the anomalously late-forming low-mass dwarf m10v (it has effectively zero star formation bursts until $z \sim 0.5$). This means our sample of core FIRE-2 halos includes m10q, m10y, m10z, m11a, m11b, m11c, m11q, m11v, m11f, m12i, m12f and m12m. Our halo catalogs and merger trees are based on Rockstar and consistent-trees (Behroozi et al., 2013b,c).

4.2.2 Particle tracking fluxes

We have re-measured inflow and outflow rates (mass, momentum, energy and metal mass) using single-adjacent-snapshot particle tracking, which complements our instantaneous flux method used in Pandya et al. (2020, 2021). The motivation is that moving forward, there are several gas recycling parameters in SAMs that cannot be

constrained with instantaneous fluxes but can be pinned down with multi-snapshot particle tracking. At this early stage, we adopt a simple single-adjacent-snapshot tracking method with definitions appropriate for calibrating our SAM as follows (note that this is much simpler than the multi-snapshot FIRE particle tracking by Anglés-Alcázar et al., 2017a; Hafen et al., 2019, 2020).

For any two adjacent snapshots, we identify all particles belonging to three “zones”: (1) the ISM defined as $< 0.1R_{\text{vir}}$, (2) the CGM defined as $0.1 - 1.0R_{\text{vir}}$, and (3) outside the halo defined as $> R_{\text{vir}}$. By computing the intersection of arrays of particle IDs for different zones in the two snapshots, we can identify which particles crossed zones. For example, inflowing particles are those that either crossed from outside the halo to $< R_{\text{vir}}$ by the second snapshot (halo inflows), or from $> 0.1R_{\text{vir}}$ to $< 0.1R_{\text{vir}}$ (ISM inflows). Similarly, outflowing particles are those that crossed from the ISM zone to $> 0.1R_{\text{vir}}$ or, in the case of halo outflows, from $< R_{\text{vir}}$ to $> R_{\text{vir}}$ by the second snapshot. Note that we neglect the possibility that ISM inflows may have become star particles by the very next snapshot; while this may happen, it is likely a negligible fraction of the full inflowing flux.

With inflowing and outflowing particle IDs classified, we can simply compute the mass flux by adding up the mass of all particles that crossed a zone boundary and dividing by the time elapsed between the two snapshots. We can also compute the momentum, energy and metal flux following the method described in Pandya et al. (2021). In particular, for the outflowing energy flux, we can compute the Bernoulli

velocity, v_B , for each particle:

$$v_B = \frac{1}{2}v^2 + \frac{3}{2}c_s^2 \quad (4.1)$$

where v is the norm of the particle’s halo-centric velocity vector and c_s is the particle sound speed computed as $\sqrt{\gamma \frac{kT}{\mu m_p}}$. We assume a monatomic ideal gas so the adiabatic index $\gamma = 5/3$ and we compute the mean molecular weight μ self-consistently for each particle using its metallicity. Equation 4.1 neglects the gravitational term ($-\frac{1}{2}v_{\text{esc}}^2$) since we are interested in the kinetic energy plus enthalpy transported by outflows (this is the energy available for heating the CGM/IGM; see Pandya et al., 2021). v_B is the specific energy and thus the total energy flux crossing a boundary would be the sum of Mv_B^2 for all outflowing particles divided by the snapshot time spacing. Analogously, the metal mass flux transported by either inflows or outflows would be the sum of MZ divided by the time elapsed between the two snapshots; here, Z is the metallicity of each particle. We do not make use of momentum fluxes in this chapter but those can also be computed easily.

Unlike Pandya et al. (2021), we do not impose any minimum Bernoulli velocity cut for outflowing particles so our particle tracking fluxes correspond to the simplest zero-radial-velocity threshold between inflows and outflows. This is done to ensure mass conservation when attempting to reproduce the bulk zone masses. Although some fraction of particles that crossed from, say, the ISM zone to CGM zone may not have a high enough v_B to travel far, they should still be considered outflows and deposited into the CGM for the purposes of a SAM. The SAM would then need an explicit “fountain flow” model for these weak outflows that recycle back into the ISM zone on a much

faster timescale than the CGM cooling time or re-infall rate of gas ejected from the halo (both of which can approach of order a Hubble time). Another way to say this is that if we had assumed a Bernoulli velocity cut to only select the subset of ISM outflows that have enough energy to reach at least $0.5R_{\text{vir}}$ Pandya et al. (as reported by 2021), we would be missing roughly half of one CGM source term (the mass deposited into it by ISM winds), and this would limit our ability to reproduce the time series of CGM mass in FIRE.

We have verified that our outflow fluxes match those reported in Pandya et al. (2020, 2021) using the instantaneous shell flux method with a simple $v_{\text{rad}} > 0$ threshold remarkably well.¹ We have also verified we can reproduce the bulk CGM, ISM and stellar masses of the FIRE halos by simply integrating our particle tracking fluxes (we know the source and sink terms for the different zones, e.g., $\dot{M}_{\text{CGM}} = \dot{M}_{\text{in,halo}} + \dot{M}_{\text{out,ISM}} - \dot{M}_{\text{out,halo}} - \dot{M}_{\text{in,ISM}}$). This increases our confidence in the particle tracking fluxes representing the “ground truth” in some sense – if a SAM can accurately reproduce the time series of our particle tracking fluxes, then it will consequently also reproduce the time evolution of the CGM, ISM and stellar mass (and hence emulate the simulation to an extent, which is our goal).

¹The halo gas inflow fluxes are systematically higher with our particle tracking method (the ISM inflow fluxes agree very well). This is partially because in Pandya et al. (2020, 2021), we were not accounting for the outward motion of the virial shell with the Hubble flow even though we were adding the Hubble velocity term to the particle peculiar velocities (this is not an issue with particle tracking which simply involves intersecting particle ID arrays). In any case, our conclusions from Pandya et al. (2020) regarding $\dot{M}_{\text{in,halo}} \ll f_{\text{b}}\dot{M}_{\text{vir}}$ in dwarfs remain unaffected (the effects of preventative stellar feedback are still apparent as large dips in the halo gas inflow rate coinciding with strong energy-conserving winds).

4.3 A Prototype Simulation-Calibrated Semi-Analytic Model

4.3.1 Guiding philosophy

Many of the free parameters in SAMs control mass and metal flow rates that are observationally unconstrained but which can be measured in simulations. By requiring mass and metal flow rates to broadly agree between SAMs and reference simulations, we may be able to fix the values of many free SAM parameters. This would reduce the otherwise enormous parameter space of SAMs and increase our confidence when making predictions to help interpret observations. Hence, our goals are to: (1) develop a SAM where the parameterizations of standard physical processes are general enough that we may reasonably expect them to apply to any cosmological simulation, and (2) extract relevant summary statistics from the simulations to better inform our choice of SAM parameters.

We may have additional free parameters that encode:

1. missing physics in our model (e.g., not modeling satellites, the multi-phase CGM, or stochasticity of SF/outflows)
2. missing physics in the reference simulations (e.g., AGN feedback)
3. uncertain or incomplete measurements from the simulations (e.g., crude parameterizations, unconstrained outflow recycling rates, and numerical errors in halo accretion rates)
4. uncertain subgrid physics and cosmic variance in the reference simulations (e.g.,

IMF and small sample size of halos)

Some of these free parameters may be further pinned down in the future by analyzing the simulations in more detail whereas other free parameters may be effectively unavoidable sources of systematic error. The SAM should be flexible enough to marginalize over uncertainties in any of these free parameters and robust enough to handle the possibility that different simulations may require different values for both the fixed and free parameters.

We do not want to change everything in classical SAMs right from the start, but we do want to simplify recipes as much as possible. Then, we can go into the reference simulation and derive scaling relations for as many of the relevant physical processes as possible. Our SAM will be a hierarchical Bayesian model so that we can allow the exact values of the free parameters to vary from halo to halo, but the overall functional forms are fixed. Coupled with an efficient sampler (e.g., MCMC or likelihood-free inference), we will be able to map out degeneracies between our free parameters and identify areas where future analysis of the simulations can provide even more constraints for the prototype SAM.

Since we are running the SAM on merger trees from hydrodynamical simulations, we have the option of hard-wiring in the full time series of any baryonic property that will be very hard to accurately reproduce (e.g., bursty SFHs). This exercise would let us ask how well we can reproduce the simulation conditioned on knowing one or more baryonic properties exactly. However, this should be avoided if at all possible because it will cause issues down the road: when we later run the SAM on large-volume

DM-only simulations, the only input we have is the halo merger tree and there will be no baryonic properties to hard-wire in. In the meantime, we can use this hard-wiring approach to guide our development of new parameterizations and provide initial guesses before doing parameter space exploration.

We now describe the equations implemented in the prototype SAM.² Much of this is inspired by the Santa Cruz SAM (Somerville & Primack, 1999; Somerville et al., 2008a, 2015) but we have made changes where needed to better reproduce FIRE. We realize that much of this has already been described in classic papers but we find it useful to review the details here and point out areas requiring future work. The most novel aspect is new preventative feedback terms for IGM and CGM heating owing to the energy loading of SN-driven winds.

4.3.2 Halo inflow

We compute the halo gas inflow rate as the sum of pristine IGM accretion and recycling of previously ejected outflows:

$$\dot{M}_{\text{in,halo}} = \dot{M}_{\text{in,pristine}} + \dot{M}_{\text{in,recycled}} \quad (4.2)$$

Pristine inflow is defined as the universal baryon fraction times the halo DM accretion rate, further reduced by the fraction of gas that is expected to be heated by the UV background (more suppression in dwarfs). The halo DM accretion rate is computed

²The code is currently available in a private github repository. We plan to publicly release an open-source alpha version as part of a future work where we run the zoom-calibrated prototype SAM on large-volume DM-only simulations.

simply by finite differencing the M_{vir} time series from the merger tree.³ This simple finite differencing approach always leads to artificial numerical noise in \dot{M}_{vir} . To mitigate this stochasticity (in many cases unphysical), we boxcar-smooth \dot{M}_{vir} over a ~ 1 Gyr window (roughly of order the halo dynamical time). This smoothed DM accretion rate matches our explicit instantaneous DM inflow rates from particle fluxes remarkably well on average. We find that the smoothing allows us to more clearly see the effect of CGM/IGM heating without muddying the interpretation (since otherwise the artificial stochasticity in \dot{M}_{vir} propagates to the rest of the SAM predictions: ISM inflow, SFH, outflow rates, heating rates, etc.). Note that in this chapter, we only model the most massive progenitor (MMP) branch rather than the full merger tree. The smoothing approach may break down when we eventually model the full merger tree, i.e., account for accretion of non-MMP branch subhalos, but that should not be a concern here.

For the UV background heating, we assume the Kravtsov et al. (2004) filtering mass analytic model with parameters set to reproduce the Okamoto et al. (2008) UV-heating gas suppression results. Then, following the Santa Cruz SAM, we compute

$$\dot{M}_{\text{in,pristine}} = f_{\text{UV}} f_{\text{b}} \dot{M}_{\text{vir}} \quad (4.3)$$

where f_{UV} is the suppression of accretion of UV background-heated gas, f_{b} is the cosmic baryon fraction and \dot{M}_{vir} is the dark matter halo mass growth rate.

The recycled halo gas inflow rate is simply the ejected gas mass divided by the Hubble time at that redshift (this is effectively what the Santa Cruz SAM does). In

³Rockstar provides both the default Mvir estimate and another “sam_Mvir” estimate that has historically been used by some SAMs that do their own subhalo modeling (ignoring subhalos from the finder). We use the default Mvir since sam_Mvir leads to lower \dot{M}_{vir} over large swaths of time.

general one might expect this recycling rate to be a function of halo mass where lower mass halos have longer effective recycling times (e.g., Henriques et al., 2013) – there may be some degeneracies with preventative feedback that we plan to investigate in the future.

$$\dot{M}_{\text{in,recycled}} = \frac{M_{\text{ejected}}}{t_{\text{Hubble}}} \quad (4.4)$$

We reduce the overall $\dot{M}_{\text{in,halo}}$ by a new preventative feedback parameter that accounts for the energy loading of SN-driven winds at R_{vir} , which we describe in subsection 4.3.6 below.

4.3.3 ISM inflow (CGM cooling)

We implement the traditional White & Frenk (1991) CGM cooling flow model. We assume that all of the halo gas is at the same temperature, namely the virial temperature of the DM halo:

$$\frac{T_{\text{vir}}}{K} = 35.9 \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^2 \quad (4.5)$$

where $V_{\text{vir}} = \sqrt{GM(< R_{\text{vir}})/R_{\text{vir}}}$ is the circular velocity at R_{vir} .⁴ We further assume this CGM gas follows a singular isothermal sphere (SIS) density profile:

$$\rho_g(r) = \frac{M_{\text{CGM}}}{4\pi R_{\text{vir}} r^2} \quad (4.6)$$

⁴In FIRE, we can explicitly compute the total enclosed mass within R_{vir} using star, gas and high-resolution DM particles. In a SAM, we have to approximate the enclosed mass with the M_{vir} provided in the merger tree. The two V_{vir} are generally similar but the merger tree one tends to systematically be a couple km/s lower for unclear reasons (perhaps because the Rockstar halo finder does not take baryonic particles into account).

and that all of the gas has the same metallicity, Z_{CGM} . This lets us define a cooling time profile for the hot halo gas:

$$t_{\text{cool}}(r) = \frac{(3/2)\mu m_p k T_{\text{vir}}}{\rho_g(r)\Lambda(T_{\text{vir}}, Z_{\text{CGM}})} \quad (4.7)$$

where Λ is the Sutherland & Dopita (1993) cooling function and we assume the mean molecular weight $\mu = 0.59$ for simplicity. This t_{cool} is the time it would take gas to radiate away its thermal energy purely through radiative cooling (assuming no heating).

We can then derive the classical “cooling radius” as the radius where the cooling time equals the “age” of the system. Historically, the age has been approximated as the time since the last major merger since those events are assumed to re-heat and re-distribute the remnant’s halo gas back into a SIS profile with temperature T_{vir} . In between such events, the CGM gas is assumed to steadily accrete into the ISM via a cooling flow (SAMs generally do not self-consistently adjust the density profile as the inner parts experience a cooling flow). Many other choices for the comparison age have been explored in the past (e.g., Somerville & Primack, 1999), but more recent SAMs instead simply derive the radius where the cooling time equals the halo dynamical time $t_{\text{dyn}} = R_{\text{vir}}/V_{\text{vir}}$. The idea is that gas within this radius will be even denser and have even shorter cooling times, whereas gas beyond this radius will take longer than a halo dynamical time to cool. We choose this approach for simplicity. Hence plugging Equation 4.6 into Equation 4.7 and setting $t_{\text{cool}} = t_{\text{dyn}}$, we can solve for r and denote it the cooling radius R_{cool} :

$$R_{\text{cool}} = \sqrt{\frac{t_{\text{dyn}} M_{\text{CGM}} \Lambda(T_{\text{vir}}, Z_{\text{CGM}})}{6\pi\mu m_p k T_{\text{vir}} R_{\text{vir}}}} \quad (4.8)$$

As intuitively expected, the cooling radius will be less than the virial radius in more massive halos at later times. This is because only a fraction of the CGM gas (i.e., only the gas in the dense inner halo) will be able to cool within a dynamical time. In this case where $R_{\text{cool}} < R_{\text{vir}}$, White & Frenk (1991) and later SAMs generally adopt the mass continuity formula advocated by Bertschinger (1989):

$$\dot{M}_{\text{cool}} = 4\pi R_{\text{cool}}^2 \rho_g(R_{\text{cool}}) \frac{dR_{\text{cool}}}{dt} \quad (4.9)$$

In general, R_{cool} grows with time since more and more of the gas in the outer parts of the halo becomes eligible for cooling (this phenomenon was termed a “cooling wave” by Bertschinger, 1989). In principle, one can get $\frac{dR_{\text{cool}}}{dt}$ by differentiating Equation 4.8 and then plug into Equation 4.9 to get

$$\dot{M}_{\text{cool}} = \frac{M_{\text{CGM}}}{t_{\text{dyn}}} \frac{R_{\text{cool}}}{R_{\text{vir}}} \quad (4.10)$$

This is the standard ISM inflow rate formula used in most SAMs. Sometimes a mysterious factor of 1/2 is applied, and it is claimed that it comes from differentiating the $R_{\text{cool}} \propto \sqrt{t}$ dependence (see derivation in Mo et al., 2010, for an arbitrary density power law). We omit it here to ensure continuity between the “hot mode” and “cold mode” cooling regime as described next (see also Guo et al., 2011).

In this traditional model, a puzzling situation can happen where the cooling radius is larger than the virial radius. This happens for dwarf halos and implies that all of the gas within the halo (plus some more gas outside the halo) will cool within a dynamical time. Hence, the actual radiative cooling time becomes irrelevant and the only limiting factor is the freefall time of the infalling CGM/IGM gas. Over the

years, this $R_{\text{cool}} > R_{\text{vir}}$ scenario has been interpreted as a signpost for “cold mode” accretion (in contrast to the previous case termed as “hot mode” accretion). Many SAMs implicitly assume that this cold mode corresponds to rapid accretion of cold filaments that directly go from the IGM to the ISM without being shock heated to T_{vir} (as has been seen in simulations; Kereš et al., 2005) – even though there is no explicit filament model. The way that SAMs calculate \dot{M}_{cool} when $R_{\text{cool}} > R_{\text{vir}}$ differs from model to model. Some SAMs assume that the newly accreted ISM gas must come entirely from cold IGM filaments, which have nothing to do with the hot halo gas, and so they set $\dot{M}_{\text{cool}} = \dot{M}_{\text{in,halo}}$ (Croton et al., 2006; De Lucia & Blaizot, 2007; Somerville et al., 2008a). This approach means that the CGM mass will effectively never grow in dwarfs (Pandya et al., 2020) and it also violates dynamical constraints (there is no way all the inflowing gas at R_{vir} will reach the ISM in one timestep if the snapshot spacing is only ~ 20 Myr, as in FIRE, when the halo dynamical time is ~ 1 Gyr). We instead adopt the approach advocated by Guo et al. (2011): when $R_{\text{cool}} > R_{\text{vir}}$, assume instead that ISM accretion is limited only by the freefall time of the hot halo gas, so that

$$\dot{M}_{\text{in,ISM}} = \begin{cases} \frac{M_{\text{CGM}}}{t_{\text{dyn}}} \frac{R_{\text{cool}}}{R_{\text{vir}}} & \text{when } R_{\text{cool}} < R_{\text{vir}} \\ \frac{M_{\text{CGM}}}{t_{\text{dyn}}} & \text{when } R_{\text{cool}} \geq R_{\text{vir}} . \end{cases} \quad (4.11)$$

We caution that this is not really a satisfying solution either. The fact that $R_{\text{cool}} > R_{\text{vir}}$ implies that additional gas outside the halo must also be available for infall (perhaps this may take the form of cold IGM filaments). So in some sense, $M_{\text{CGM}}/t_{\text{dyn}}$ is a lower limit for cold mode accretion (some extra fraction of $\dot{M}_{\text{in,halo}}$ should also be added).

However, this approach at least has the advantage that it does not predict extremely low CGM masses in dwarfs – gas cannot just instantaneously accrete into the ISM in a single timestep, instead it must be limited by the freefall (dynamical) time. In the future, we will add an explicit cold filament model with parameters tuned to reproduce simulations (e.g., following Mandelker et al., 2020) and allow both cold and hot accretion modes to occur simultaneously (Lu et al., 2011a).

We additionally suppress $\dot{M}_{\text{in,ISM}}$ by a new preventative feedback parameter that accounts for the energy-loading of SN-driven winds out of the ISM, which we describe in subsection 4.3.6 below.

4.3.4 Star formation

We model the star formation rate as simply

$$\text{SFR} = \frac{M_{\text{ISM}}}{t_{\text{dep}}} \quad (4.12)$$

We can keep track of M_{ISM} but we cannot predict the gas depletion timescale t_{dep} (put another way: we do not have an explicit model for the star formation efficiency). Instead, as shown in Figure 4.1, we have parameterized t_{dep} simultaneously as a function of V_{vir} and redshift for the FIRE-2 halos:

$$\frac{t_{\text{dep}}}{\text{Gyr}} = 10^{6.7} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{-3.0+2.5 \log(1+z)} (1+z)^{-5.9} \quad (4.13)$$

We make the instantaneous recycling approximation for mass return to the ISM by stellar winds and SNe. Fixing the recycling mass fraction to be $f_{\text{recycle}} = 0.4$ (appropriate for the initial mass function of Chabrier, 2003), only the fraction $(1 -$

f_{recycle}) of the SFR gets added as long-lived stellar mass whereas the rest is left as ISM gas in each timestep.

This simple model has the advantage that it can plausibly match the SFH on average, particularly when SF is continuous as in massive halos at late times. However it can fail when SF is bursty, especially for late-forming dwarf halos which can go many Gyr without a single starburst.

We do not attempt multi-phase ISM partitioning or galaxy morphological modeling (distinct mass profiles for the gas disk, stellar disk, stellar bulge, etc.) at this time. We also do not currently model merger-induced starbursts and stellar/gas accretion from satellites.

4.3.5 Stellar feedback (ejective)

In Figure 4.2, we parameterize mass loading factors out of the ISM and the halo as a function of both V_{vir} and redshift:

$$\eta_{\text{M,ISM}} = 10^{7.6} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{-3.7+4.2 \log(1+z)} (1+z)^{-6.3} \quad (4.14)$$

and

$$\eta_{\text{M,halo}} = 10^{9.0} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{-4.5+4.4 \log(1+z)} (1+z)^{-7.1} \quad (4.15)$$

At any given timestep, the prototype SAM can therefore predict outflow rates from the ISM and from the CGM given the current SFR, V_{vir} and redshift: $\dot{M}_{\text{out}} = \eta_{\text{M}} \text{SFR}$.

Our approach of adopting the halo loading factor scaling from FIRE-2 is different from previous SAMs, which instead have free parameters (often a threshold halo




Figure 4.1: ISM gas depletion time as a function of V_{vir} and redshift for the FIRE-2 halos. We can use our continuous fit to estimate SFRs in the prototype SAM.

virial velocity and some power law exponent) that control what fraction of the ISM wind gets deposited into the CGM vs. into the ejected reservoir. This classical approach does not account for CGM entrainment where more wind mass can be leaving the halo than has previously left the ISM – this is especially important in the FIRE dwarfs given their energy-conserving winds (Pandya et al., 2020, 2021). Since we do not yet have a fully predictive model for CGM entrainment by SN-driven winds, our approach for $\eta_{\text{M,halo}}$ lets us account for this phenomenologically.

Our outflow model is currently instantaneous: mass is removed from the ISM and instantly deposited into the CGM in the same timestep as the SF. Then, in that same time step, we also move the estimated halo outflow mass from the CGM to the ejected reservoir immediately. Whenever the halo mass loading is larger than the ISM mass loading, our approach naturally ejects some extra CGM mass (as required beyond the ISM wind mass alone). We are careful to only eject the amount of mass available in the ISM/CGM to prevent negative mass errors.

This approach builds up the ejected mass reservoir, only a fraction of which can re-accrete back into the halo at every timestep (as described earlier). In principle, a large fraction f_{escape} of winds from dwarfs may become unbound forever and should not be added to the ejected reservoir. We do not have constraints yet from FIRE on this but will in the near-future with multi-snapshot particle tracking. In the meantime, allowing for a V_{vir} -dependent recycling parameterization may lead to very long recycling times in dwarfs, which may effectively achieve the same thing (though note that currently we only assume a single universal timescale for all halos).

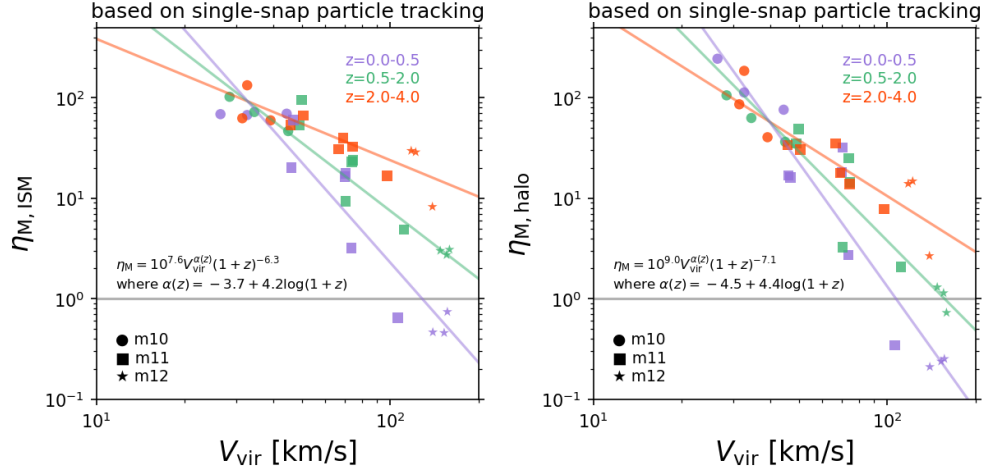


Figure 4.2: ISM mass loading factor (left) and halo mass loading factor (right) as a function of V_{vir} and redshift. Note that these are based on particle tracking fluxes without any cut on minimum Bernoulli velocity (i.e., most closely corresponding to a zero-radial-velocity threshold between outflows and inflows). We can use our continuous fits in the SAM, although for some individual dwarf halos the scatter can be substantial and should be taken into account.

4.3.6 Stellar feedback (preventative)

In Figure 4.3, we parameterize the Bernoulli velocity of ISM and halo outflows as a function of V_{vir} and redshift simultaneously:

$$v_{\text{B,ISM}} = 10^{0.3} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{1.0-0.3 \log(1+z)} (1+z)^{0.4} \quad (4.16)$$

and

$$v_{\text{B,halo}} = 10^{0.3} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{0.8-0.1 \log(1+z)} (1+z)^{0.6} \quad (4.17)$$

By definition, $v_B \equiv \sqrt{\dot{E}_{\text{out}}/\dot{M}_{\text{out}}}$ where E_{out} is the total kinetic energy plus enthalpy transported by the wind (Pandya et al., 2021). Thus, in the SAM, knowing \dot{M}_{out} via our η_M parameterization, we can also assign energy outflow rates for the ISM and halo

winds using our v_B parameterization.

Next, we need to derive “heating efficiency functions” that quantify how much of the available SN wind energy couples to the CGM and IGM to suppress accretion. Pandya et al. (2020) emphasized that the Santa Cruz SAM (representative of other classical SAMs) has much larger accretion rates than FIRE, by 1-2 orders of magnitude for dwarfs. Since it has no preventative feedback, it relies on correspondingly very high mass loading factors (much larger than in FIRE) to roughly match the stellar and ISM mass assembly histories. Pandya et al. (2020) suggested that the energy of SN-driven winds can shock heat gas in the CGM and even outside the halo (in the case of dwarfs). However, they did not have explicit measurements of the available SN wind energy for such heating. Here we implement such heating in the prototype SAM using our FIRE measurements and ask whether there is indeed enough SN wind energy to sufficiently suppress halo/IGM gas cooling, and what that energy coupling efficiency needs to be.

For each halo, at each timestep, we can compute the difference between the halo inflow rate predicted by the prototype SAM (as described above) and the actual FIRE halo inflow rate. If this difference is positive, that means the SAM has some excess accretion rate $\Delta\dot{M}_{\text{in}}$. Preventing this amount of excess accretion would require a heating (unbinding) rate of

$$\dot{E}_{\text{needed}} = \frac{1}{2}\Delta\dot{M}_{\text{in}}V_{\text{vir}}^2 \quad (4.18)$$

For now, we do a very simple heating calculation:

$$\dot{M}_{\text{heat}} = \begin{cases} \Delta\dot{M}_{\text{in}} & \text{if } \dot{E}_{\text{out,halo}} \geq \dot{E}_{\text{needed}} \\ 2\dot{E}_{\text{out,halo}}/V_{\text{vir}}^2 & \text{if } \dot{E}_{\text{out,halo}} < \dot{E}_{\text{needed}} . \end{cases} \quad (4.19)$$

In words, if there is enough halo wind energy available, then we offset the full $\Delta\dot{M}$ excess accretion. Otherwise, we can only offset a smaller amount given the available halo wind energy. In the massive halos at late times, only a fraction of the SN wind energy reaches R_{vir} whereas in the dwarfs, winds tend to be energy-conserving as they propagate from $\sim 0.1R_{\text{vir}}$ to R_{vir} (Pandya et al., 2021). So it is almost a question of energy budgets: how does the available SN wind energy compare to the predicted excess accretion when we run the prototype SAM? Note that this “hard-wired” $\Delta\dot{M}_{\text{in}}$ approach is for illustrative purposes, but it can also be used to parameterize and provide initial guesses for heating efficiencies.

We repeat the exact same procedure to check whether CGM heating is needed in each timestep (does the SAM have an excess ISM inflow rate compared to FIRE or not?) and how much CGM heating is possible given the available $\dot{E}_{\text{out,ISM}}$.

We define a new inflow suppression factor as

$$f_{\text{SN}} = 1 - \frac{\dot{M}_{\text{heat}}}{\dot{M}_{\text{in}}} \quad (4.20)$$

This is calculated separately for the halo inflow and the ISM inflow. In the case of halo inflow, we thus end up with

$$\dot{M}_{\text{in,halo}} = f_{\text{SN,halo}} f_{\text{UV}} f_{\text{b}} \dot{M}_{\text{vir}} + f_{\text{SN,halo}} \dot{M}_{\text{in,recycled}} \quad (4.21)$$

This makes explicit that we have a new preventative feedback parameter f_{SN} for pristine gas owing to SN wind heating in addition to the usual f_{UV} suppression from the UV background. $f_{\text{SN,halo}}$ also multiplies $\dot{M}_{\text{in,recycled}}$ but we expect a degeneracy with the recycling parameters that we will map in a future work. Similarly, in the case of ISM inflow suppression, we multiply $\dot{M}_{\text{in,ISM}}$ by a new $f_{\text{SN,CGM}}$ parameter – it is assumed that the fraction $(1 - f_{\text{SN,CGM}})$ is kept hot at $\sim T_{\text{vir}}$ and remains in the CGM instead of accreting into the ISM.

For energy budgeting purposes, we also keep track of the fraction of available wind energy that is actually used for heating the CGM and IGM at each timestep. One can imagine that in the MW halos, which have much higher accretion rates, a larger fraction of the SN wind energy would be needed to offset even a fraction of the total \dot{M}_{in} (of order $\sim 10M_{\odot}/\text{yr}$ for the m12’s). Whereas in the dwarfs, the lower overall accretion rates (e.g., $\sim 0.01M_{\odot}/\text{yr}$ for the m10’s) would require disproportionately lower SN wind energy to heat. In principle, we could enforce some kind of “energy conservation” such that only the fraction of ISM wind energy that did not couple to the CGM would then still be available for further heating outside the halo. However, it is not clear how to do this since, e.g., energy-conserving winds in dwarfs may entrain a significant fraction of the CGM yet still maintain their overall energy to then heat the IGM (see supplementary movies of the m10 and m11 halos in Pandya et al., 2021).

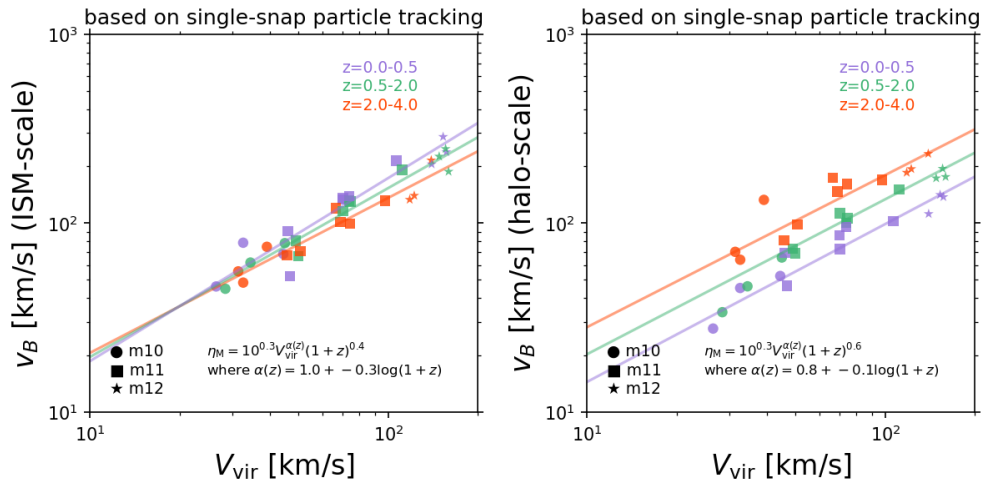


Figure 4.3: Bernoulli velocity of ISM outflows (left) and halo outflows (right). We can assign energy outflow rates to the SAM using the continuous parameterizations as a function of both V_{vir} and redshift (see text). This can then serve as the basis for preventative heating of the CGM and IGM by SN-driven winds (more important in dwarfs).

4.3.7 Chemical evolution

We assume a single “effective yield” parameter y_Z to model chemical evolution in a simple way, combining this with our instantaneous recycling approximation described in subsection 4.3.4. We assume that new stars form with the same metallicity as the ISM gas at that time:

$$\dot{M}_{Z,\text{stars}} = Z_{\text{ISM}}(1 - f_{\text{recycle}})\text{SFR} \quad (4.22)$$

The ISM metal mass decreases by that amount; note that the $(1 - f_{\text{recycle}})$ factor accounts for the fact that SNe and stellar winds instantaneously return metals back into the ISM. The ISM metal mass will also increase because nucleosynthesis within stars ejects newly produced metals (as encapsulated by our y_Z parameter). For this nucleosynthetic yield,

we also use the instantaneous approximation.

The metal masses of the ISM, CGM, and ejected gas will also be affected by metal mass inflow and outflow rates. For this, we can estimate average flow metallicities from FIRE as \dot{Z}/\dot{M} (analogous to what we did for v_B ; Pandya et al., 2021). In Figure 4.4, we parameterize inflow and outflow metallicities at $0.1 - 0.2R_{\text{vir}}$ and $1.0 - 1.1R_{\text{vir}}$ from FIRE as a function of V_{vir} and redshift:

$$Z_{\text{in,halo}}/Z_{\odot} = 10^{-5.2} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{2.2+0.1 \log(1+z)} (1+z)^{-1.6} \quad (4.23)$$

$$Z_{\text{in,ISM}}/Z_{\odot} = 10^{-3.9} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{1.8+0.5 \log(1+z)} (1+z)^{-2.4} \quad (4.24)$$

$$Z_{\text{out,ISM}}/Z_{\odot} = 10^{-4.5} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{2.2-0.5 \log(1+z)} (1+z)^{-0.5} \quad (4.25)$$

$$Z_{\text{out,halo}}/Z_{\odot} = 10^{-6.4} \left(\frac{V_{\text{vir}}}{\text{km/s}} \right)^{-3.0-1.9 \log(1+z)} (1+z)^{2.9} \quad (4.26)$$

With these, the net metal mass change rates for the ISM, CGM and ejected components are respectively given by:

$$\dot{M}_{Z,\text{ISM}} = Z_{\text{in,ISM}}\dot{M}_{\text{in,ISM}} + y_Z(1 - f_{\text{recycle}})\text{SFR} - Z_{\text{ISM}}(1 - f_{\text{recycle}})\text{SFR} - Z_{\text{out,ISM}}\dot{M}_{\text{out,ISM}} \quad (4.27)$$

$$\dot{M}_{Z,\text{CGM}} = Z_{\text{in,halo}}\dot{M}_{\text{in,recycled}} + Z_{\text{out,ISM}}\dot{M}_{\text{out,ISM}} - Z_{\text{out,halo}}\dot{M}_{\text{out,halo}} \quad (4.28)$$

$$\dot{M}_{Z,\text{ejected}} = Z_{\text{out,halo}}\dot{M}_{\text{out,halo}} - Z_{\text{in,halo}}\dot{M}_{\text{in,recycled}} \quad (4.29)$$

We are careful to only accrete/eject the metal mass that is actually available from the different prototype SAM reservoirs to prevent negative metal mass errors.

4.4 Case study of modeling a low-mass dwarf

We begin by illustrating the effects of our new rapid CGM cooling model and SN-driven heating of the IGM and CGM in a low-mass dwarf (m10q).

4.4.1 Impact of new rapid CGM cooling model

Figure 4.5 shows the time evolution of CGM mass and corresponding mass inflow/outflow rates crossing R_{vir} for the low-mass dwarf m10q. As shown by Pandya et al. (2020), the SC-SAM predicts very low CGM masses for dwarfs since in this $R_{\text{cool}} > R_{\text{vir}}$ regime, it assumes that the ISM inflow rate equals the halo gas inflow rate, effectively removing a source term for the CGM. Our new rapid CGM cooling model where the ISM inflow rate is limited by the halo free-fall time is capable of reproducing the CGM mass in FIRE. However, this agreement is still achieved with much higher halo inflow/outflow rates than FIRE (which roughly balance so that the net halo inflow rate should be similar). Next we assess the impact of reducing halo accretion.

4.4.2 Impact of SN heating of the IGM

We find that SN heating of the IGM is capable of reducing halo accretion rates down to the level measured in FIRE (yellow line in middle panel of Figure 4.5). In other words, there is generally enough energy in the SN-driven wind leaving R_{vir} to offset the

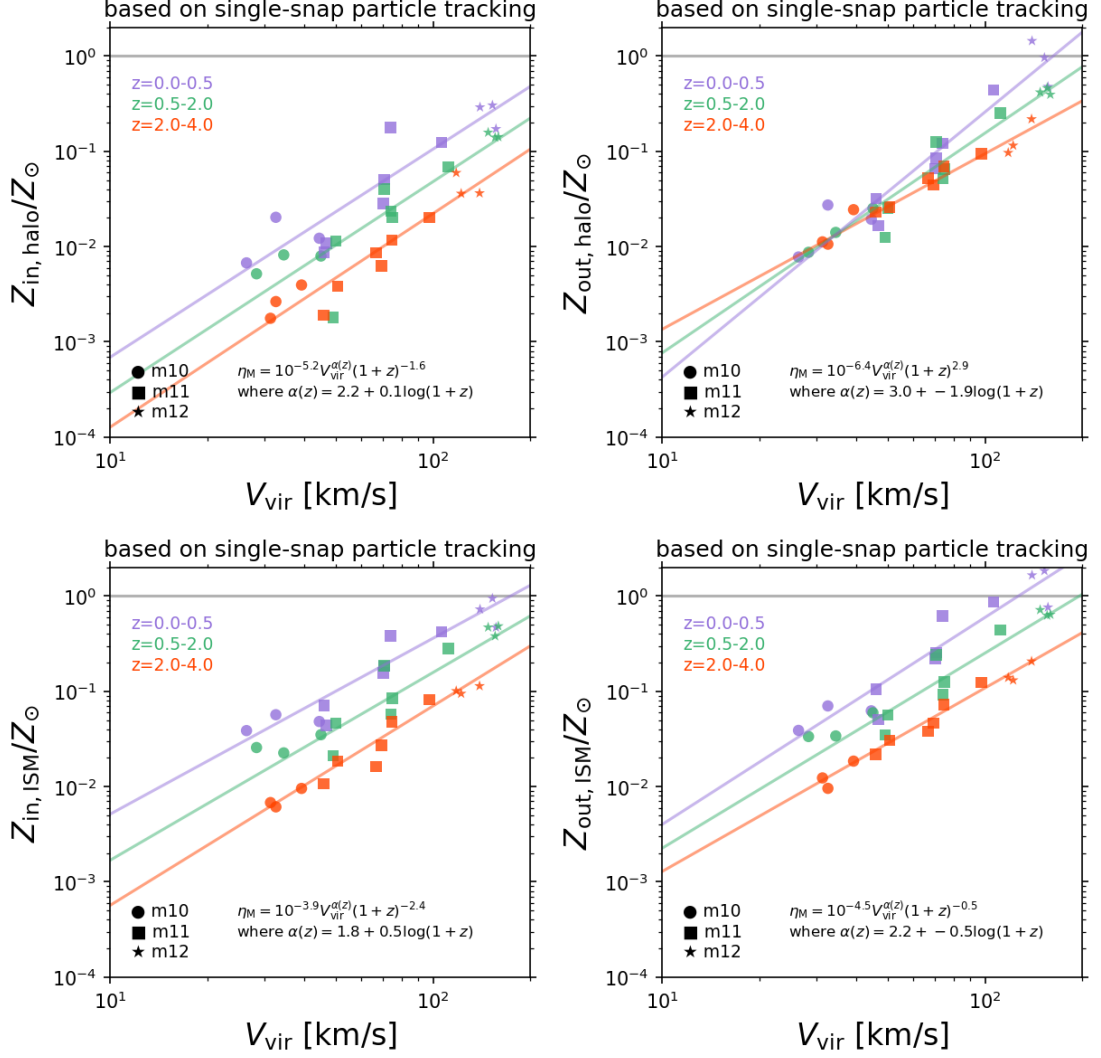


Figure 4.4: Inflow and outflow metallicities ($\dot{Z}/\dot{M}/Z_{\odot}$) for the ISM shell and halo shell. Our parameterizations that depend on V_{vir} and redshift enable easy implementation into the prototype SAM.

$\Delta\dot{M}_{\text{in,halo}}$ between the prototype SAM and FIRE. There are some times where the halo inflow deficit in FIRE is too large and cannot be fully offset in the prototype SAM. In some cases, such as at ~ 6 Gyr in Figure 4.5, this may be due to the prototype SAM underestimating a large spike in the halo outflow and hence available wind energy (this is a consequence of our continuous/smooth SF model). There may also be other sources of heating such as satellites and IGM turbulence that we do not yet have an explicit model for.

Naturally, reducing the halo accretion rate in the prototype SAM leads to a factor of a few lower CGM mass than in FIRE. As will be shown next, the lower CGM mass leads to a lower ISM inflow rate, lower SFR, and hence lower overall outflow rates from both the ISM and at R_{vir} (yellow line in bottom panel of Figure 4.5). Thus, it seems that simply by reducing the amount of gas accreting into the halo to match FIRE, the prototype SAM would also get the halo outflow rate roughly right (versus being almost an order of magnitude higher without IGM heating).

4.4.3 Impact of SN heating of the CGM

While SN heating of the IGM alone is capable of reducing the halo inflow and outflow rates in the prototype SAM to better match FIRE, the CGM mass becomes systematically less than FIRE by a factor of a few. If we enable SN-driven heating of the CGM in addition to the existing IGM heating, some CGM mass is prevented from cooling into the ISM and remains in the CGM. This has the effect of increasing the CGM mass to again qualitatively agree with FIRE (red line in the top panel of Figure 4.5).

The agreement in the halo inflow and outflow rates remains qualitatively the same as the IGM-heating-only case.

It is informative to also look at the ISM mass and associated flow rates with and without heating. Figure 4.6 shows that by only enabling IGM heating, the ISM mass and corresponding fluxes decrease in the prototype SAM to qualitatively match FIRE. The ISM fluxes still remain higher than FIRE by a factor of order unity. At first glance, it would appear that CGM heating has no additional effect. However, recall that when enabling CGM heating, the CGM mass increases – by roughly an order of magnitude relative to the IGM-heating-only case (top panel of Figure 4.5). Since this dwarf halo is in the rapid CGM cooling regime ($R_{\text{cool}} > R_{\text{vir}}$), this larger CGM mass would directly lead to a larger ISM inflow rate via $\dot{M}_{\text{in,ISM}} = M_{\text{CGM}}/t_{\text{dyn}}$. The fact that the ISM inflow rates agree qualitatively between the two heating scenarios (IGM-heating-only versus both CGM+IGM heating; middle panel of Figure 4.6) requires substantial CGM heating. Paradoxically, CGM heating offsets ISM accretion, leading to larger CGM mass, and that increased CGM density leads to an even faster cooling rate requiring more heating to offset.

4.4.4 Stellar mass assembly history

Figure 4.7 shows the stellar mass assembly history for m10q in FIRE, the Santa Cruz SAM, and in the prototype SAM under the various heating scenarios. We see that without any IGM/CGM heating, the prototype SAM (and SC-SAM) systematically over-predict the FIRE stellar mass assembly history by a factor of a few. Enabling

IGM and/or CGM+IGM heating brings the stellar mass assembly history into better agreement with FIRE, primarily because the ISM mass decreases sufficiently (Figure 4.6) for our $SFR = M_{\text{ISM}}/t_{\text{dep}}$ to reproduce FIRE. However, at early times, we still overshoot the FIRE stellar mass which likely happens because we have a continuous SF model whereas there are far fewer discrete SF bursts in very early FIRE dwarfs. Alternatively, extending our t_{dep} parameterization for FIRE to go beyond the $z = 2 - 4$ bin, perhaps by adding one additional $z = 4 - 8$ bin, might be useful.

4.4.5 Wind energy coupling efficiency

Figure 4.8 shows the accretion suppression factor f_{SN} assumed in the prototype SAM for the IGM-heating-only and both CGM+IGM heating scenarios. We see that for this low-mass dwarf, $f_{\text{SN}} \ll 1$ at most times, implying that substantial heating of both the IGM and CGM are needed. We also see that nearly all of the available wind energy is needed to couple to the CGM (assuming SN-driven winds from the central galaxy are the only source of the heating). This supports our argument above that early CGM heating leads to an increase in CGM mass which then translates to a higher cooling rate that must be offset with even more heating. In the case of IGM heating, a substantial fraction of the halo wind energy is also required to couple to gas outside the halo, potentially heating it to above the halo virial temperature and making it harder for the gas to accrete.




Figure 4.5: Time evolution of CGM mass (top), halo gas inflow rate (middle) and halo gas outflow rate (bottom) for the low-mass dwarf m10q. FIRE measurements are in black, Santa Cruz SAM in cyan. There are 3 lines for the prototype SAM: without any IGM/CGM heating (magenta), with only IGM heating (yellow), and with both IGM and CGM heating (red). We see that our new rapid CGM cooling model is capable of roughly reproducing the FIRE CGM mass whereas the SC-SAM is low by orders of magnitude. However, this agreement is achieved still with very high halo inflow/outflow rates, whereas enabling SN heating of the IGM and CGM leads to good agreement in both the mass and flux time series.




Figure 4.6: Analogous to Figure 4.5 but now the time evolution of ISM mass and associated ISM inflow/outflow rates for the low-mass dwarf m10q.




Figure 4.7: Analogous to Figure 4.5 but now the time evolution of stellar mass.

4.5 Results for the entire FIRE-2 ensemble

Following our case study of an individual low-mass dwarf (for which the SC-SAM disagreed most dramatically with FIRE-2; Pandya et al., 2020), here we summarize results for all other FIRE-2 halos (including intermediate-mass dwarfs and MW-mass halos). Condensed figures showing the time evolution of CGM, ISM and stellar properties for individual halos can be found Appendix 4.5.1.

4.5.1 Evolution of individual halos

For completeness, here we show the time evolution of CGM, ISM and stellar mass, and associated mass flow rates, in FIRE and the SAMs for individual halos.




Figure 4.8: Time evolution of the accretion suppression factor (top) and wind energy coupling efficiency (bottom) for the low-mass dwarf m10q. The IGM-heating-only scenario for the prototype SAM is shown with the yellow lines, and the CGM plus IGM heating scenario in red. When $f_{\text{SN}} < 1$, that means there was otherwise excess accretion $\Delta\dot{M}$ that needed to be offset. We see that this low-mass dwarf requires both CGM and IGM heating. Furthermore, we see that nearly all of the available wind energy needs to couple to the CGM to offset excess ISM accretion, and there needs to be substantial wind energy coupling with the IGM as well.

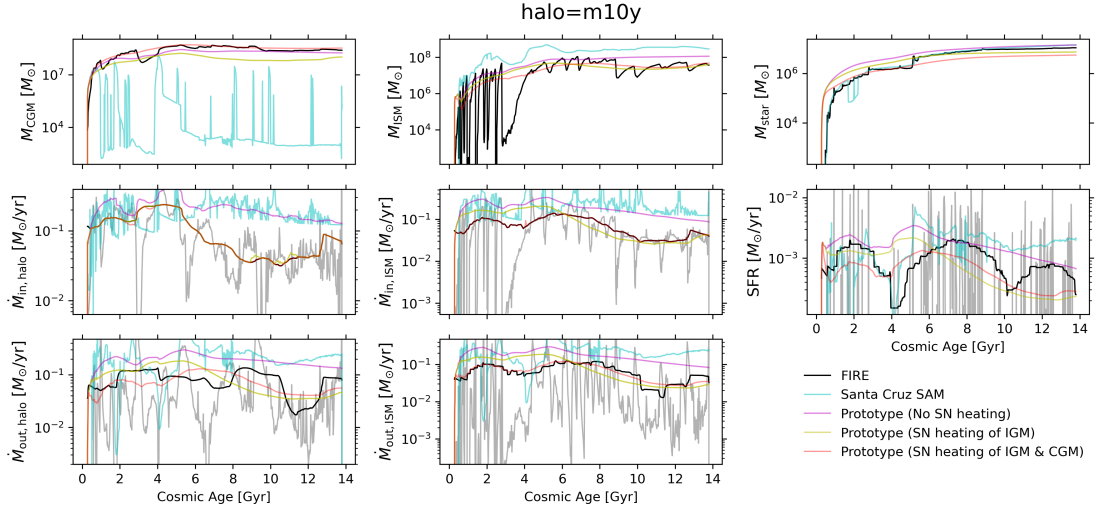


Figure 4.9: Time evolution of CGM mass and associated inflow/outflow rates (left column), ISM mass and corresponding inflow/outflow rates (middle column), and stellar mass and underlying SFR (right column) for the low-mass dwarf m10y. Line styles are the same as in previous figures.

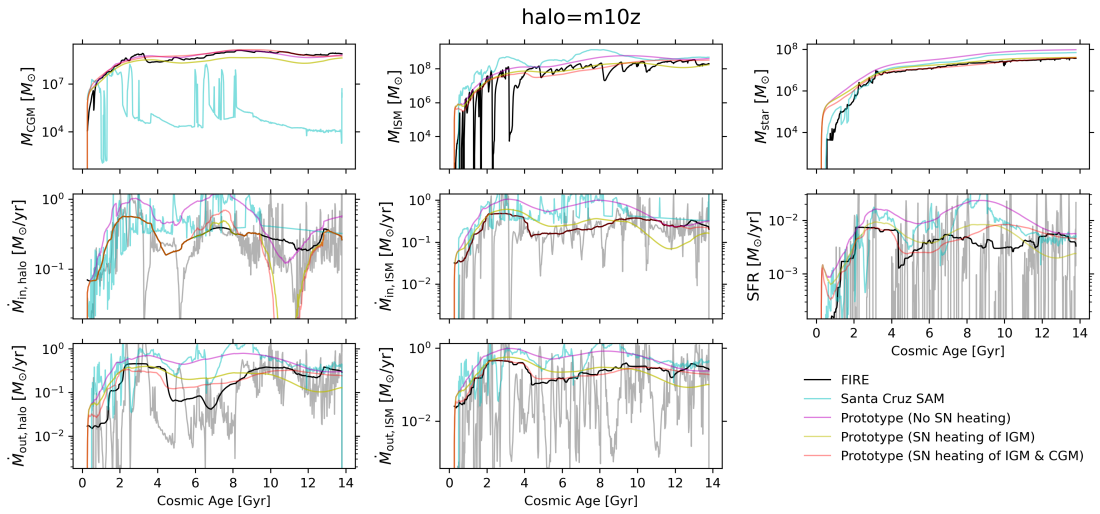


Figure 4.10: Similar to Figure 4.9 but for the low-mass dwarf m10z.




Figure 4.11: Similar to Figure 4.9 but for the intermediate-mass dwarf m11a.




Figure 4.12: Similar to Figure 4.9 but for the intermediate-mass dwarf m11b.




Figure 4.13: Similar to Figure 4.9 but for the intermediate-mass dwarf m11c.

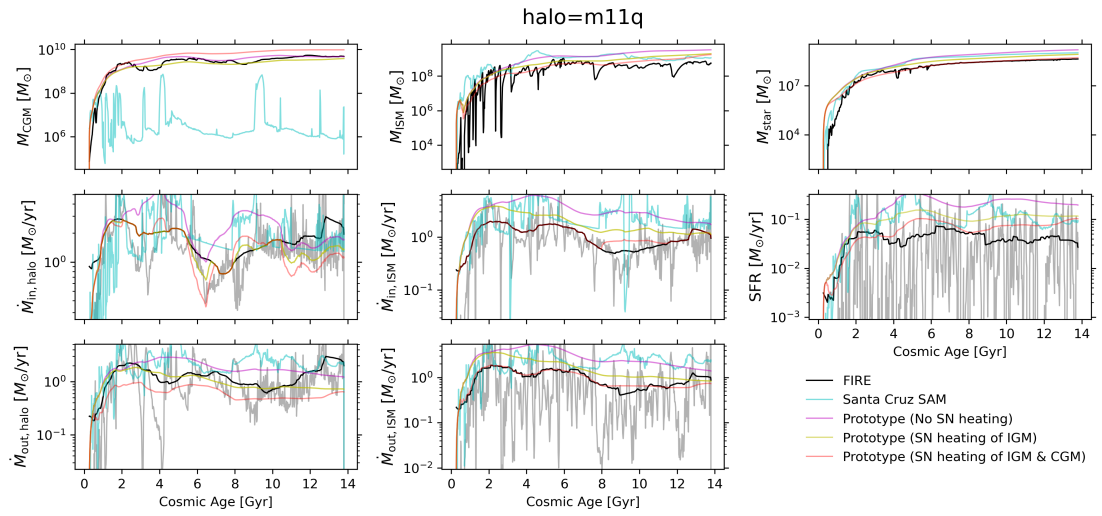


Figure 4.14: Similar to Figure 4.9 but for the intermediate-mass dwarf m11q.




Figure 4.19: Similar to Figure 4.9 but for the MW-mass halo m12m.

4.5.2 IGM heating efficiency function

Figure 4.20 shows the halo gas accretion suppression factor, $f_{\text{SN,IGM}}$, as a function of halo virial velocity in three broad redshift bins. We see that there is an overall trend such that dwarfs require more IGM heating in the prototype SAM, so that f_{SN} is lower in order to bring $f_{\text{SN}} f_{\text{UV}} f_{\text{b}} \dot{M}_{\text{vir}}$ into better agreement with the gas inflow rate at R_{vir} in FIRE. However, there is substantial scatter in the dwarfs that does not appear to correlate simply with redshift. In contrast, higher mass halos uniformly asymptote to $f_{\text{SN}} = 1$ implying IGM heating is unnecessary. This makes sense since the halo virial temperatures are larger than the ambient IGM temperature, meaning more of the gas beyond R_{vir} is bound to the halo and able to accrete. And in the case of high-redshift MW progenitors, those are fed by cold filaments which may not be as

susceptible to heating.

4.5.3 CGM heating efficiency function

Figure 4.21 similarly shows the CGM cooling suppression factor $f_{\text{SN,CGM}}$ assumed in the prototype SAM. The overall CGM heating efficiencies are much larger than for the IGM. This reflects the need to offset the otherwise large CGM cooling rates predicted in the “cold/rapid cooling mode” in dwarfs as $M_{\text{CGM}}/t_{\text{dyn}}$ – in some cases, scaling this cooling rate down by as much as $f_{\text{SN,CGM}} \sim 0.1$. The MW-mass halos at low redshift apparently do not require CGM heating, but this may be due to the prototype SAM lacking a fountain flow term and underproducing the CGM metallicity. Both of those issues lead to a lower estimate of $\dot{M}_{\text{in,ISM}}$ than FIRE by a factor of a few, and it is plausible that some CGM heating would be required once the prototype SAM is suitably modified.

4.5.4 CGM depletion time as a heating diagnostic

Figure 4.22 plots an “effective” CGM depletion time computed as $M_{\text{CGM}}/\dot{M}_{\text{in,ISM}}$ for the FIRE-2 halos and the prototype SAM with and without our SN heating of the IGM/CGM. At high redshift, dwarfs (and MW progenitors) would naively be expected to undergo “cold/rapid mode” accretion. In our model, this means that the scaling of t_{dep} with V_{vir} should be flat since the ISM inflow rate is limited by the free-fall time of the CGM (so $t_{\text{dep}} = t_{\text{dyn}}$ and note that t_{dyn} is independent of V_{vir}). The prototype SAM indeed would predict a flat scaling in the dwarf regime (middle panel). However, dwarfs




Figure 4.20: The halo accretion suppression factor, $f_{\text{SN,IGM}}$ as a function of halo virial velocity in three broad redshift bins: $z = 0.0 - 0.5$ (purple), $z = 0.5 - 2.0$ (green) and $z = 2.0 - 4.0$ (red). A low value of $f_{\text{SN,IGM}}$ indicates greater suppression whereas $f_{\text{SN,IGM}} = 1$ means no heating was assumed by the prototype SAM. There is an overall trend for dwarfs requiring greater suppression albeit with significant scatter. Higher-mass halos do not require significant halo accretion suppression (as expected, since at high-redshift they are fed by cold filaments, and at later times, their virial temperature is higher than the ambient IGM gas enabling efficient accretion).




Figure 4.21: Analogous to Figure 4.20 but now for CGM heating. There is much more heating assumed in the prototype SAM overall for the CGM as compared to the IGM, with some dwarfs requiring the “cold/rapid mode” CGM cooling rate $M_{\text{CGM}}/t_{\text{dyn}}$ to be scaled down by as much as $f_{\text{SN,CGM}} \sim 0.1$. The MW-mass halos at late times do not require as much CGM heating, though that may be due to the lack of a fountain flow term and underproduction of CGM metallicity in the prototype SAM (leading to lower cooling rates versus FIRE to begin with).

in FIRE-2 have systematically higher CGM depletion times which is suggestive of additional energy input that offsets the otherwise expected rapid cooling. Indeed, when we enable SN heating of the CGM and IGM in the prototype SAM, we can qualitatively reproduce the negative-sloped trend at high-redshift.

At lower redshifts, there is substantial scatter in the FIRE-2 data. This may be due to the contribution of fountain flows to the measured $\dot{M}_{\text{in,ISM}}$ thus leading to an overestimate of the \dot{M}_{cool} that we are really interested in. Mergers and turbulence stirred by satellites may also complicate the measurement of an effective CGM depletion time in these cosmological simulations. Nevertheless, on average the CGM depletion time is larger at lower redshifts (e.g., note the smooth redshift evolution of the MW points) which is suggestive of CGM cooling becoming increasingly more inefficient at later times in FIRE.

4.6 Discussion

4.6.1 Significance and implications of preventative SN feedback

We have demonstrated that SAMs without preventative feedback predict high accretion rates into both the halo and the ISM, and that this issue is particularly pronounced for lower mass halos. While it is possible to reproduce the bulk CGM, ISM and stellar masses of dwarfs without heating, the corresponding inflow and outflow rates would be dramatically higher than measured in FIRE. Historically, it has been challenging to preferentially limit early star formation in dwarfs with SAMs that in-




Figure 4.22: The effective CGM depletion time ($M_{\text{CGM}}/\dot{M}_{\text{in,ISM}}$) as a function of halo virial velocity in three broad redshift bins. The FIRE-2 data are shown on the left, the prototype SAM without any SN heating of the IGM/CGM in the middle, and the prototype SAM with full SN heating of both the IGM/CGM on the right. The horizontal dashed lines illustrate the halo dynamical time (at R_{vir}) at the midpoint of the three redshift bins. While there is a lot of scatter at lower redshifts, the high-redshift halos would naively be expected to undergo “cold/rapid mode” accretion where $t_{\text{dep}} = t_{\text{dyn}}$ since the ISM inflow rate is limited by the free-fall time of the CGM (middle panel). However, in FIRE, the high-redshift dwarfs have preferentially longer CGM depletion times suggestive of heating. Indeed, the prototype SAM is able to qualitatively reproduce the negative-sloped high-redshift trend when SN heating of the CGM/IGM is enabled.

clude only ejective feedback. One solution was to preferentially increase the recycling timescale of gas ejected from dwarf halos (Henriques et al., 2013; White et al., 2015). We have presented an alternative solution involving preventative heating, although both mechanisms may be in play.

There are many possible sources of preventative feedback, and previous studies have not distinguished between them. Lu et al. (2015) invoke general “pre-heating” of the IGM and draw a similar conclusion that by reducing the amount of gas accreting into the halo, it is possible to reproduce many observed scaling relations for local disk galaxies whereas an otherwise purely ejective feedback model fails. Here we have shown for the first time that there is enough kinetic energy plus enthalpy in SN-driven winds, at least with the FIRE-2 stellar feedback scheme, to offset the excess accretion in dwarfs predicted by simple SAMs. Furthermore, we have shown that not only is wind heating needed at the halo boundary, but it is in particular needed for halos that would naively be expected to have rapid “cold mode” accretion. Our measurements of the ISM inflow rate in FIRE are far below the naive expectation where the whole CGM cools extremely rapidly and its infall rate is limited by its free-fall time.

We anticipate that if SN-driven heating of the CGM and IGM around dwarfs is indeed significant, then there may be observational implications concerning their metal content. Lu et al. (2017) have already shown that a SAM with both preventative and ejective feedback can simultaneously match the stellar mass function and stellar mass-metallicity relation of nearby galaxies. However, it remains to be seen what the effect of SN heating is on the redshift evolution of these two metrics, the gas-phase mass-

metallicity relation, and the relative abundances of different ions in the CGM and IGM. Given that most of the energy and metals of SN ejecta escapes from dwarfs (Pandya et al., 2021), we may expect the CGM and IGM around dwarfs to be preferentially enhanced in metals compared to MW-mass halos. We may also expect a greater abundance of highly ionized species if the diffuse gas is continually heated or shocked by winds. It will be useful to examine scaling relations from, e.g., the COS-Dwarfs survey to understand differences in the CGM metal content of dwarfs versus MW-mass galaxies.

4.6.2 Remaining uncertainties in the prototype SAM

Here we discuss remaining uncertainties of the prototype SAM – these are avenues for ongoing and future work.

4.6.2.1 Fountain flow recycling

Our parameterized ISM mass loading factors from FIRE (Figure 4.2) do not currently distinguish between high-velocity winds that will propagate far beyond the ISM versus cooler, slower outflows that will rapidly recycle in the inner halo (so-called fountain flows). In the prototype SAM, if we deposit this cool/slow material into the CGM, it is assumed to be re-heated to the halo virial temperature and will take far longer to re-accrete into the ISM compared to the typically much more rapid fountain flow timescale (only ~ 250 Myr; Anglés-Alcázar et al., 2017a).

Our visualization of the FIRE simulations reveals that inner halo recycling is

ubiquitous and must be explicitly accounted for in SAMs. Figure 4.23 shows a dramatic example of the ISM of an intermediate-mass dwarf extending to beyond $0.2R_{\text{vir}}$. This galaxy is experiencing cycles of ISM expansion and re-contraction due to strong SN feedback, which leads to higher measured ISM inflow rates than would be predicted from CGM cooling alone ($M_{\text{CGM}}/t_{\text{dyn}}$).

In the future, we will extend our single-adjacent-snapshot tracking to quantify the distribution of timescales for outflowing particles to return to $< 0.1R_{\text{vir}}$. Then we can define a new CGM fountain flow reservoir with free parameters for its mass budget and recycling timescale (calibrated to reproduce FIRE). The amount of wind mass in our fountain reservoir and its recycling timescale may help interpret observations of gas clouds at the CGM-galaxy interface (e.g., high-velocity HI clouds, covering fractions of different ions, etc.).

4.6.2.2 Underproduction of CGM metallicity

The prototype SAM currently systematically under-predicts the CGM metallicity in FIRE. The disagreement can be as bad as $\sim 2-3$ orders of magnitude or by only a factors of a few in some cases. An order of magnitude deficit in the CGM metallicity can translate to an order of magnitude deficit in the cooling function $\Lambda(T_{\text{vir}}, Z_{\text{CGM}})$, and since $R_{\text{cool}} \propto \sqrt{\Lambda}$ and $\dot{M}_{\text{in,ISM}} \propto R_{\text{cool}}$, the (hot mode) CGM cooling rate would be under-predicted by a factor of $\sqrt{10}$.

Figure 4.24 illustrates that simply integrating the various \dot{Z} terms measured in FIRE along with an instantaneous single-yield ($y = 0.03$) stellar recycling assumption

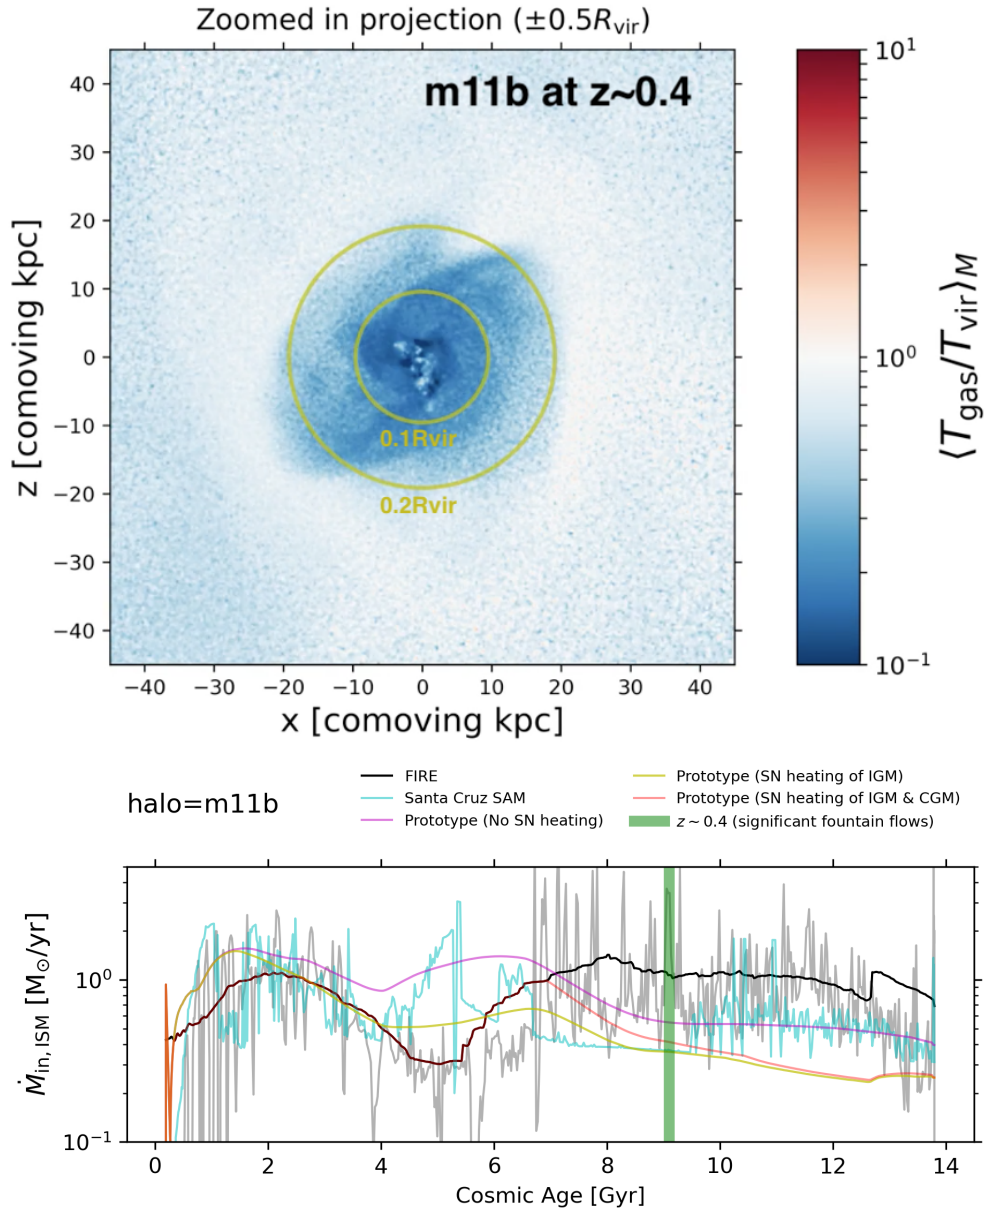


Figure 4.23: Illustration of the ISM extending to $\sim 0.2R_{\text{vir}}$ for an intermediate-mass dwarf due to cycles of expansion and re-contraction driven by strong SN feedback (top). This leads to systematically higher measurements of the ISM inflow rate in FIRE than would be predicted from CGM cooling alone by the prototype SAM (bottom panel). With multi-snapshot particle tracking in the future, we will distinguish first-time ISM accretion from recycled ISM accretion to guide the implementation of a new rapid fountain flow reservoir in the prototype SAM.

mostly reproduces the ISM and stellar metallicity. However, the CGM metallicity is significantly underpredicted. We find this to be the case for all halos. This implies that we are missing a source term for CGM metals. The missing metals must presumably be produced via nucleosynthesis in satellite star particles. We have not yet measured satellite properties, but by simply estimating the total SFR of all subhalos within R_{vir} of the central halo, we can then estimate the contribution of new metals produced in the CGM via satellites (irrespective of whether those satellites eject their metals via outflows).

4.6.2.3 Inadequacies of the standard CGM cooling model

Our CGM cooling model is effectively that of White & Frenk (1991), except in the “cold mode” regime ($R_{\text{cool}} > R_{\text{vir}}$) we set the ISM inflow rate to the CGM mass divided by its freefall time (approximated as the halo dynamical time). We showed that this “cold/rapid” accretion model in dwarfs is capable of reproducing their CGM masses whereas the approach in some other SAMs leads to irreconcilably low CGM masses (even though the predicted ISM inflow rates are similar). One can imagine that our approach gives an upper limit to the CGM cooling rate since gas may spiral inwards with some angular momentum support, and thus there may need to be an order unity correction factor to the free-fall time approximation. On the other hand, since $R_{\text{cool}} > R_{\text{vir}}$, that implies that extra gas outside the halo is also available for direct accretion onto the ISM. That may take the form of cold filaments that deposit already-cold IGM gas into the ISM without shock heating. In other words, while our cold mode model may make

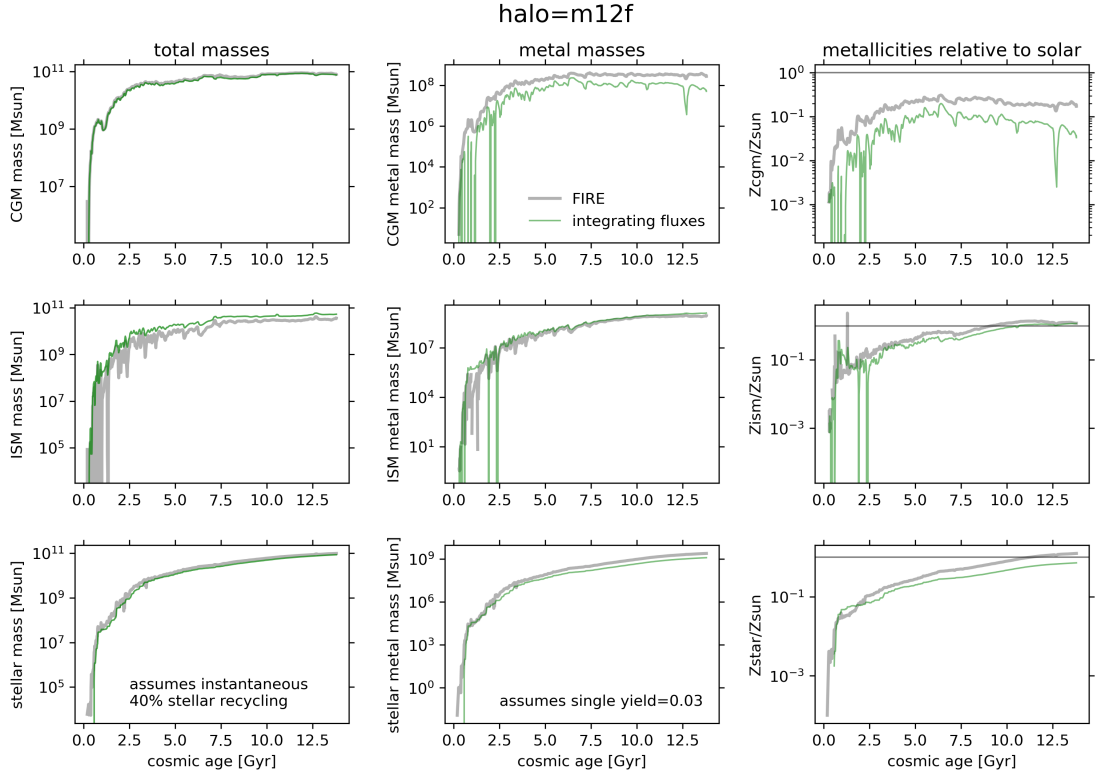


Figure 4.24: Representative example showing that we can roughly reproduce the overall masses of the CGM, ISM and stars as well as the metallicity of the ISM and stars. However, the CGM metallicity is significantly underestimated. The gray curves are bulk measurements from FIRE. The green curves are from simply integrating the \dot{M} and \dot{Z} terms as measured via our particle tracking.

sense for the CGM, we need to consider adding an explicit IGM→ISM cold filament model, which we expect to be substantial in the high-redshift MW progenitors and some intermediate-mass dwarfs.

As for hot mode accretion, we started with the usual assumptions of a uniformly hot CGM at T_{vir} with a singular isothermal density profile and a single metallicity Z_{CGM} . Our results already suggest that the CGM temperature of dwarfs in FIRE may be considerably hotter than T_{vir} (hence our solution was CGM heating via SN winds; Figure 4.22). There is also the issue of a multi-phase CGM partitioning into roughly hydrostatic hot gas that is still cooling versus already-cooled gas that is eligible for accretion into the CGM; this may be important to model explicitly in MW-mass halos (along with cool inner halo fountain flows, cold clumps, precipitation, etc.; e.g., Maller & Bullock, 2004; Voit et al., 2015). It is reasonable to expect that the CGM metallicity will have some radial gradient – it may be more enriched in the inner parts, which may have consequences for computing the cooling rate and interpreting observations. Finally, the CGM density profile, at least in MW-mass halos, may be shallower than r^{-2} which will also impact the cooling rate calculations.

We plan to investigate the implementation and implications of more detailed CGM models for SAMs in the future.

4.6.2.4 Other sources of preventative feedback

In this work we have asked whether SN-driven winds alone can provide the heating necessary to reproduce the low CGM cooling and halo accretion rates in FIRE.

We showed that there is more than enough kinetic energy plus enthalpy in the SN winds in FIRE to offset excess accretion predicted by SAMs. However, there are almost certainly other significant heating sources in the simulations as well. Perhaps the most obvious is outflows and turbulence stirred by satellites in the CGM, which we currently do not model. Neighboring halos may also pre-heat the IGM to temperatures that are hotter than the virial temperatures of our lowest-mass dwarfs, making it harder for them to accrete “pristine” gas. And, at the very highest redshifts, before there is significant star formation, if significant heating is required (e.g., in dwarfs), then that energy must be provided by sources other than Type II SNe.

This begs the question: what do the Bernoulli velocities of winds in other simulations look like? If other simulations predict weaker wind energies than FIRE, then either we would need to model the other heating mechanisms or those simulations may have lower accretion rates in the first place. Finally, of course, the FIRE simulations do not include AGN feedback which may play a role in reducing CGM cooling and SF rates in MW-mass halos (in fact, most SAMs today assume significant AGN heating even in MW-mass halos; Croton et al., 2006; Somerville et al., 2008a). What do high-resolution simulations suggest for the CGM/IGM heating (and pushing) rates from AGN-driven winds and jets, and crucially, how does this depend on the assumptions for SMBH seeding and feeding (e.g., Brennan et al., 2018)?

4.6.2.5 Modeling non-main branches of the full merger tree

Currently in the prototype SAM we only model the most massive progenitor branch. However, the other branches can contain progenitors of comparable mass that can contribute substantial ejected gas metals and perhaps CGM heating via their own outflows and stirring turbulence. In the future we will model these other branches up until they join the main branch (become accreted as subhalos), at which point we will need an explicit model for satellite stripping, orbital dynamics and merger times (perhaps the usual dynamical friction formula used in SAMs can be calibrated to reproduce satellite properties in the simulation; e.g., Jiang et al., 2021).

4.6.2.6 Continuous versus stochastic SF and feedback

Our simple SF model in the prototype SAM is only designed to reproduce the FIRE SFHs on average. This has consequences: we cannot predict stronger than average winds which lead to stronger than average suppression of CGM cooling and halo accretion. Our effectively continuous SF model also overpredicts the stellar mass at early times whereas in FIRE, high-redshift dwarfs only have a few discrete SF episodes. These early discrete SF episodes drive particularly powerful energy-conserving winds that can clear out the entire ISM, which our model also does not capture (since only the average mass loadings are used, and our parameterizations only go up to $z = 4$). It would be useful to implement an alternative stochastic SF model in the prototype SAM that also includes an explicit time delay for outflows to reach the CGM, possibly entrain gas and then leave the halo. While we may not be able to accurately predict the time of

each individual SF burst in FIRE, this model may be a more faithful representation of SF physics in zoom simulations and have implications for selecting galaxies in mock observational surveys.

4.6.2.7 Lack of multi-phase and multi-element partitioning

We currently do not explicitly model different phases of the ISM, CGM, outflows and inflows. While we have characterized the loading factors of cold, warm and hot outflows for FIRE (Pandya et al., 2021), it is not yet clear how to implement wind phases (and mixing) into a SAM. The simplest thing would be to estimate a cool outflow fraction that is expected to recycle back into the ISM on a short timescale – this may help interpret observations. For multi-phase inflows, we can separately track the hot hydrostatic CGM gas and already-cooled halo gas that is eligible to fall into the ISM (or be re-heated). On top of that, we may have cold filaments and we can allow the hot and cold accretion modes to occur simultaneously (as in Lu et al., 2011a). Tracking different phases in the future will facilitate observational comparisons.

Our chemical evolution model is also too simple: we do not follow the yields of individual elements and we ignore the contribution of Type Ia SNe. This means we cannot predict varying $[\alpha/\text{Fe}]$ abundance ratios, which is crucial to be able to simultaneously match observed gas-phase and stellar mass-metallicity relations. In the future, we will fix the yields of multiple elements using the tables assumed in FIRE. We will also do time-dependent stellar recycling and account for both the prompt and delayed metal contribution of type Ia SNe (following, e.g., Arrighi et al., 2010; Yates et al.,

2013).

4.7 Summary

We have presented the prototype of a new semi-analytic model of galaxy formation with explicit implementation of preventative stellar feedback. We measured scaling relations from FIRE-2 for the ISM depletion time, mass loading factors for winds leaving the ISM and halo, outflow Bernoulli velocities (specific energy), and metallicities of inflows and outflows at both the ISM and halo scale. These serve as inputs to our prototype SAM and dramatically reduce the number of free parameters. We assume that preventative feedback comes mainly from SN-driven winds that can shock-heat the CGM and IGM (heating from the UV background is not as important for our halo masses). Our main findings are as follows.

1. Instead of setting the ISM inflow rate equal to the halo accretion rate in the $R_{\text{cool}} > R_{\text{vir}}$ regime (rapid/cold mode accretion), we assume that the ISM inflow rate is limited only by the free-fall time of the CGM: $\dot{M}_{\text{in,ISM}} = M_{\text{CGM}}/t_{\text{dyn}}$. This model is capable of reproducing the dwarf CGM masses measured in FIRE whereas in the previous model the CGM mass effectively does not grow with time. However, it still predicts a very large ISM inflow rate compared to measurements in FIRE, by an order of magnitude for some dwarfs.
2. Standard SAM assumptions for pristine halo accretion (the cosmic baryon fraction times the dark matter accretion rate) plus recycling of SN wind ejecta over a

Hubble time lead to much larger halo inflow rates than measured for FIRE dwarfs. The kinetic energy plus enthalpy transported by SN winds in FIRE is large enough to offset this excess accretion, which in some cases can be larger than an order of magnitude. The winds can shock heat the gas outside the halo to above the halo virial temperature so that it effectively cannot accrete.

3. Similarly, we ask how much wind energy leaving the ISM must be coupled to the CGM to force agreement between the SAM's CGM cooling rate and the inflow rate measured in FIRE. MW-mass halos require little to no CGM heating, although this is highly uncertain due to the prototype SAM's lack of a fountain flow recycling term and its underprediction of CGM metallicity. Dwarf halos require substantially more CGM heating since otherwise $M_{\text{CGM}}/t_{\text{dyn}}$ would lead to overly high cooling rates and ISM masses. Amazingly, the SN wind energy is sufficiently high in FIRE to potentially offset a lot of this excess CGM cooling.
4. We derive heating efficiency functions for the IGM and CGM using the SN wind energy coupling assumed in our SAM. These describe a new f_{SN} parameter that would multiply and suppress the halo accretion rate and CGM cooling rate. Both functions asymptote to $f_{\text{SN}} = 1$ for MW-mass halos implying that no additional heating is necessary. Dwarfs require substantial heating albeit with more scatter: f_{SN} can be as low as 0.5 for the IGM (meaning half of the expected accretion is prevented) and as low as 0.1 for the CGM (meaning only a tenth of the expected CGM cooling rate is actually allowed).

5. We compute an effective CGM depletion time as $t_{\text{dep,CGM}} = M_{\text{CGM}}/\dot{M}_{\text{in,ISM}}$. At high-redshift and in dwarf halos, the FIRE halos would naively be expected to have cold mode accretion and thus have $t_{\text{dep}} \approx t_{\text{dyn}}$ (independent of halo mass). Instead, the CGM depletion time is preferentially longer in lower mass halos. We show that the prototype SAM can qualitatively recover this negative correlation only if we turn on CGM+IGM heating. At lower redshifts, there is considerably more scatter which may be due to the contribution of inner halo fountain flows.

The most important next step will be to separately measure pristine (first-time) ISM accretion versus fountain flows, and model a new inner halo recycling term for SAMs. Since $\dot{M}_{\text{in,ISM}} = \dot{M}_{\text{CGM,cool}} + \dot{M}_{\text{fountain}}$, once we have pinned down $\dot{M}_{\text{fountain}}$, we can tweak $\dot{M}_{\text{CGM,cool}}$ via heating. We also need to identify and correct the source of the prototype SAM’s underprediction of CGM metallicity (at the order of magnitude level) since that may translate to an underprediction of the CGM cooling rate for “hot mode” accretion (this may be due to the lack of modeling satellite branches). Then, we will be able to provide robust CGM and IGM heating functions that can be implemented in SAMs run on DM-only simulations (for which there will be no simulated baryonic properties available for calibration purposes). We will also then be in a position to try alternative hot mode CGM cooling prescriptions, explicitly model gas brought in via cold IGM filaments, constrain outer halo recycling parameters, and ultimately generate mock CGM absorption and emission maps throughout large volumes to complement existing and future surveys. For all of this, it will be crucial to convert the prototype SAM into

a hierarchical Bayesian model where each halo follows the same set of equations, but the exact free parameters may be drawn from some distribution to account for scatter in the galaxy population.

Chapter 5

Conclusions and Future Work

5.1 Summary

In this thesis, we have compared two different approaches to modeling galaxy formation and explored how they can be made to better agree with each other in terms of the overall baryon cycle. Chapter 2 compared the Santa Cruz SAM (Somerville et al., 2015, and references therein) to the FIRE-2 suite of cosmological zoom-in simulations spanning a range of masses (Hopkins et al., 2018a). We demonstrated that the two models agree on the stellar assembly histories of individual galaxies remarkably well but with very different underlying mass flow rates for the CGM and ISM. The SAM predicts much higher accretion rates into the halo and ISM and hence requires correspondingly higher mass outflow rates than measured in FIRE. The CGM masses of dwarfs in the SAM are also orders of magnitude lower than in FIRE, which emphasizes that the CGM (along with mass fluxes) may be an effective discriminator between models of

galaxy formation that otherwise predict roughly similar bulk ISM and stellar masses. We presented a simple model where the energy of SN-driven winds may shock heat gas outside dwarf halos to above their virial temperature, thereby suppressing accretion to levels measured in FIRE.

Chapter 3 presented the most comprehensive analysis of SN-driven winds in fully cosmological simulations to date. In addition to measuring the often-quoted wind mass and metal loading factors in FIRE (now split by phase) , we also measured energy and momentum loading factors to estimate how much of the input SN energy and momentum escapes from the ISM. We showed that SN-driven winds in dwarfs are energy-conserving even as they propagate out of the halo (in fact, the energy, momentum and metal loading factors are all of order unity) whereas winds in MW-mass halos show substantial losses upon first breaking out from the ISM.

Chapter 4 combined insights from the previous two chapters to explicitly model preventative feedback in a SAM where for the first time this process was constrained by the results of high-resolution cosmological simulations and attributed explicitly to shock heating by SN-driven winds. We presented the prototype of a new SAM in which nearly all of the free parameters and functional forms are derived from the FIRE-2 simulations. We showed that it is possible to simultaneously roughly reproduce the CGM, ISM and stellar masses of dwarfs if we allow some fraction of the SN wind energy to couple to the CGM and IGM. We find that there is generally more than enough wind energy available to offset the excess accretion that would otherwise be predicted with simple SAM approximations for halo gas accretion and CGM cooling. However,

we also find that significant uncertainties remain in modeling rapid inner halo recycling from fountain flows, the chemical enrichment of the CGM, and several other physical processes not yet captured in the prototype SAM. Nevertheless, this work demonstrates that it may indeed be possible to develop a simple model that can emulate the key results of high-resolution simulations and sets the foundation for eventually scaling up predictions to statistically large, cosmologically representative volumes.

5.2 Future Work

The planned future work described in this section can roughly be split into two categories: (1) further development and improvements for the prototype SAM, and (2) generating mock observations from both the SAM and hydrodynamical simulations to help interpret a variety of observations, with an emphasis on the CGM.

5.2.1 Observable consequences of preventative stellar feedback

One of the main conclusions of this thesis is that preventative stellar feedback should be ubiquitous in dwarfs, and it is therefore important to ask how observations can test this claim. First, it would be very useful to translate maps of gas density, temperature and metallicity to emission line maps during particularly powerful outflow episodes of FIRE dwarfs. This would require post-processing the simulation particle data with a photoionization model, spectral synthesis and radiative transfer code. It would be especially interesting to do this for high-redshift dwarfs and assess the observability of rest-frame optical lines like $H\alpha$ with the James Webb Space Telescope

(JWST). In addition, it would also be helpful to report column densities for different ions that are within the wavelength range covered by various facilities as a function of redshift. What are the expected column densities and emission fluxes from energy-conserving, rapidly-expanding superbubbles around dwarfs, and how do those compare to the detection limits of current and planned facilities?

At a population level, we can use the new SAM to explore how various scaling relations may change with and without preventative feedback. Lu et al. (2017) have shown it is not possible to simultaneously reproduce the stellar mass function and stellar mass-metallicity relation of nearby dwarfs with only ejective or only preventative feedback, but that a combination of the two is required. Their preventative feedback model postulates general pre-heating of the IGM without invoking specific heating sources. It would be useful to re-visit this claim in our model where it is the SN wind energy that drives the heating, and to make predictions at higher redshifts that can be tested with JWST. It would also be good to simultaneously consider the gas-phase mass-metallicity relation which should also be affected by preventative feedback (and historically, it has been challenging to simultaneously reproduce both the stellar and gas-phase mass-metallicity relations with SAMs; Somerville et al., 2015). Relatedly, it would be useful to constrain the metallicity of observed winds around dwarfs since we predict that the metal loading factor should be of order unity (most of the SN ejecta metals escape) and substantially lower in MW-mass halos. More generally, if star-forming galaxies roughly maintain equilibrium in a bathtub-type model (Davé et al., 2012), then we would predict that in cases of preventative feedback, the mass loading factor should be lower

since some accretion is prevented and less gas is available for star formation. With the advent of large IFU surveys and the ability to estimate mass loading factors for statistical samples of galaxies (e.g., Roberts-Borsani et al., 2020), it would therefore be interesting to re-visit galaxies that have peculiarly low inferred mass loading compared to other galaxies of similar mass, radius, SFR, etc. Note that investigating metallicities will require us to implement a more general chemical evolution model in the prototype SAM in which yields of multiple elements are taken directly from stellar evolution tables and time evolution is taken into account (including the delayed contribution of heavy elements from Type Ia SNe).

Finally, if high-redshift dwarfs do indeed have powerful energy-conserving winds that heat gas outside the halo to above the virial temperature, then this may have implications for cosmic reionization. There is great uncertainty about whether the photons produced by star formation alone can escape in sufficient numbers to reionize the Universe (e.g., see the FIRE-1 perspective provided by Ma et al., 2015). However, if nearly all of the available SN wind energy from dwarfs is able to escape intact to $> 2R_{\text{vir}}$ as our FIRE-2 results suggest, then galactic winds may provide a non-negligible source of IGM heating. It would be interesting to translate the large-scale wind properties in FIRE to potential heating rates. It would also be useful to measure the sizes of the collisionally-ionized and metal-enriched bubbles around FIRE halos. In parallel, it would be insightful to run idealized simulations where highly energy- and mass-loaded winds are propagated into a homogeneous or turbulent IGM so that the emergent wind UV/X-ray spectra and energy coupling efficiencies can be explored over a range of con-

ditions (e.g., following Meiksin et al., 2017).

5.2.2 Wind recycling

It will be crucial to split ISM accretion into fresh CGM cooling versus rapid recycling of fountain flows in the inner halo. Currently, winds in the prototype SAM leaving the ISM are either re-heated to the virial temperature in the CGM or take of order the Hubble time to re-accrete into the halo from the ejected reservoir. However, a substantial fraction of the ISM inflow rate measured in FIRE may come from more rapidly recycled cooler winds, or represent gas flows associated with the expansion and re-contraction of the ISM disk due to feedback in dwarfs. With multi-snapshot particle tracking of the FIRE data, we will be able to isolate the recycling term for the ISM inflow rate and calibrate a new inner halo recycling model for the SAM. More generally, we will also have to re-think the ejected reservoir for winds leaving the halo. Some fraction of winds in dwarfs will become unbound from the halo forever whereas another part of the “ejected” reservoir may actually reside within the CGM but just cool very inefficiently. The ejected gas recycling rate parameterization may also vary with halo mass in FIRE whereas currently we are assuming a fixed recycling timescale (Henriques et al., 2013; White et al., 2015).

Modeling fountain flows will let us immediately compare to observations of extraplanar gas at the disk-CGM interface of nearby galaxies, which is thought to be dominated by fountain flows rather than freshly cooled CGM (e.g., Fraternali & Binney, 2008; Marasco et al., 2019). We will also be able to re-visit the origin of damped

Lyman- α absorbers (DLAs) as a function of redshift: is the large observed number density of absorbers consistent with arising from the extended ISM (e.g., Berry et al., 2014) plus cool fountain flows around galaxies? As for halo-scale wind recycling, it will be interesting to explore gas and DM flow patterns in FIRE up to $\sim 2 - 3R_{\text{vir}}$, determine the radius at which the accretion transitions from recycling-dominated to pristine-dominated, and compare that to various physically-motivated definitions for the halo boundary (including the splashback radius). For dwarfs in particular, it would be useful to constrain how far beyond the virial radius does the gas inflow rate rise back to the expected $f_b \dot{M}_{\text{in,DM}}$. That large-scale “effective feedback radius” within which baryonic effects dominate may have implications for cosmological probe experiments that rely on tracing galaxies.

5.2.3 More faithful modeling of CGM physics

While our CGM cooling model is able to get reasonably close to FIRE in terms of the CGM mass and ISM inflow rate, it still represents an incomplete picture of the simulations. Our model assumes that the CGM follows a singular isothermal density profile, is uniformly at the halo virial temperature, and has a single overall metallicity. Instead, we should allow these three quantities to follow arbitrary power law profiles whose parameters are calibrated from FIRE. It has already been suggested that an r^{-2} power law is too steep for MW-mass halos in FIRE, which instead have a shallower density profile of $r^{-3/2}$ (Stern et al., 2019). It is also reasonable to expect that the CGM is multi-phase and that the inner part is preferentially metal-enriched.

One of the biggest limitations of our CGM cooling model is that it does not allow for the simultaneous occurrence of “cold” and “hot” mode accretion. In other words, we currently only model the cooling rate of the existing CGM gas but not the accretion of cold IGM filaments. For this, we need to implement an explicit filament model with parameters that describe their cold ISM mass deposition rate and frequency as a function of halo mass and redshift (e.g., Mandelker et al., 2020). We can adapt our analysis machinery from Chapter 3 to characterize multi-phase inflows in FIRE instead of outflows.

By incorporating a more physical model of the CGM and accretion of cold IGM filaments, we will unlock opportunities to use the SAM to forward model observations of the CGM in both absorption and emission for large samples of halos. As a baseline, we may assume that the CGM gas is in collisional ionization equilibrium, that it is subject to photoionization from the UV background, and that the mapping from total metallicity to individual elements follows solar abundances. Then we can draw random sightlines as a function of halocentric radius and predict ion column densities versus impact parameter that can be compared to existing and future CGM surveys (see Tumlinson et al., 2017). The emission spectrum from X-ray to optical wavelengths of the radiatively cooling warm/hot CGM can be predicted by assuming a cooling function (e.g., Faerman et al., 2019). We can also make simplified estimates of the Lyman- α emission from cold IGM filaments which together with the cool fountain flows can help interpret the origin of extended neutral gas in high-redshift halos as well as large-scale clustering of Lyman- α emitters (e.g., Dijkstra & Loeb, 2009). Lastly, running on large-volume DM-only

simulations, we can predict gas density profiles around halos and help interpret high-redshift CGM/IGM tomography experiments with existing facilities and future thirty-meter class telescopes (e.g., Rudie et al., 2012; Lee et al., 2018).

5.2.4 Extension to satellites, higher mass halos and SMBHs

Eventually, we will want to use the SAM to make predictions for observables involving the whole galaxy population such as correlation functions, stellar mass functions, and scaling relations extending up to galaxy cluster scales. In order to do this, we will need to model satellite galaxies and higher-mass halos, with SMBH physics being a co-requisite for the latter.

Currently we are only modeling the most massive progenitor branch of the full merger tree, but the other branches contain substantial information. Non-main progenitors can bring in pre-enriched gas (and stars) and help resolve the current underproduction of CGM metallicity in the prototype SAM. After a distinct halo becomes a subhalo such that its branch merges with the main branch, we will need to implement recipes for mergers and tidal stripping processes. To help calibrate these satellite modeling parameters, we can repeat our Chapter 2 analysis for central halos but now tracking the evolution of satellite properties in the simulations. With an explicit satellite model, we can confront observations of satellite colors, clustering, and other statistics describing the evolution of MW-like Local Group-like and cluster-sized halos.

On the other hand, we also need to extend our model to higher mass halos where AGN feedback is crucial. One idea is to add a flexible empirical model that mimics

the effects of AGN feedback and other physical processes exclusive to high-mass halos, thereby artificially suppressing SFRs, CGM cooling rates, halo inflow rates, etc. This flexible model could be calibrated to reproduce various observations. Alternatively, we can run the FIRE-calibrated SAM on merger trees for higher mass halos (e.g., extracted from large DM-only simulations) and compare to predictions from abundance matching or semi-empirical models (e.g., Moster et al., 2018; Behroozi et al., 2019). Then we could estimate how much additional feedback is needed beyond stellar feedback alone to reproduce the observed properties of massive galaxies.

Extending the prototype SAM to satellites and higher mass halos will almost certainly require additional simulations for calibration purposes. For example, AGN feedback is not included in the core FIRE-2 simulations and the simulated halos are also all chosen to be roughly isolated whereas to model satellites we will require simulations of dense clusters. As the available computational power increases, it may be possible to simulate larger volumes with sufficiently high resolution to keep modeling stellar feedback locally and explicitly as in FIRE. This would increase our sample size of halos for calibrating SAMs. However, a fully self-consistent treatment of SMBH seeding, feedback and feedback remains out of reach for now and so any cosmological simulation would require phenomenological, effectively SAM-like recipes for AGN. In this regard, it may be worthwhile to also try calibrating the prototype SAM using large-volume phenomenological simulations since these at least attempt to have a “complete” range of physics. A SAM that can emulate large-volume simulations can be used to interpret their complex emergent results, explore variations in their subgrid recipes, and scale up

the predictions from those simulations to even larger volumes for cosmological probe applications.

5.2.5 Towards an open-source hierarchical Bayesian SAM

Lastly, we plan to invest a significant amount of effort over the next few years on making the prototype SAM code publicly available as open-source code with comprehensive documentation. This will open up new collaborative opportunities for further developing our model and comparing to observations. A major goal will be to translate the prototype SAM into a hierarchical Bayesian model. This is important because we are currently calibrating SAM parameters on an individual halo-by-halo basis. While every galaxy is subject to following the same set of coupled ODEs, the exact free parameter values may vary from object to object. Thus, using a hierarchical Bayesian model will allow us to capture scatter in the galaxy population.

In addition, while many of our free parameters are currently fixed by measuring the relevant fluxes in FIRE, we may require flexible models for other physics that either is not included in FIRE or is difficult to directly measure and parameterize (e.g., recycling, turbulence, effects of cosmic rays and SMBHs, etc.). With a hierarchical Bayesian model coupled to an efficient parameter space exploration technique, we can jointly fit any free parameters and map out degeneracies (and then look for ways to break those degeneracies either with additional simulation or observational data). Ultimately, we would like to re-calibrate all SAM parameters several times and attempt to emulate many different sets of reference simulations, perhaps with the aid of deep learning. This

will allow us to quantify uncertainties in modeling galaxy formation physics and how that translates to uncertainties in various cosmological probes. In this thesis we have taken just the first few promising steps towards this grand ambition and are hopeful that SAMs will eventually be developed into a complete, realistic and standard physical model of galaxy formation.

Bibliography

Anglés-Alcázar, D., Davé, R., Özel, F., & Oppenheimer, B. D. 2014, *ApJ*, 782, 84,
doi: 10.1088/0004-637X/782/2/84

Anglés-Alcázar, D., Faucher-Giguère, C.-A., Kereš, D., et al. 2017a, *MNRAS*, 470, 4698,
doi: 10.1093/mnras/stx1517

Anglés-Alcázar, D., Faucher-Giguère, C.-A., Quataert, E., et al. 2017b, *MNRAS*, 472,
L109, doi: 10.1093/mnrasl/slx161

Angles-Alcazar, D., Quataert, E., Hopkins, P., et al. 2020, arXiv e-prints,
arXiv:2008.12303. <https://arxiv.org/abs/2008.12303>

Arrigoni, M., Trager, S. C., Somerville, R. S., & Gibson, B. K. 2010, *MNRAS*, 402, 173,
doi: 10.1111/j.1365-2966.2009.15924.x

Ayromlou, M., Nelson, D., Yates, R. M., et al. 2020, arXiv e-prints, arXiv:2004.14390.
<https://arxiv.org/abs/2004.14390>

Behroozi, P., Wechsler, R. H., Hearin, A. P., & Conroy, C. 2019, *MNRAS*, 488, 3143,
doi: 10.1093/mnras/stz1182

- Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013a, *ApJ*, 770, 57, doi: 10.1088/0004-637X/770/1/57
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013b, *ApJ*, 762, 109, doi: 10.1088/0004-637X/762/2/109
- Behroozi, P. S., Wechsler, R. H., Wu, H.-Y., et al. 2013c, *ApJ*, 763, 18, doi: 10.1088/0004-637X/763/1/18
- Benson, A. J. 2010, *Phys. Rep.*, 495, 33, doi: 10.1016/j.physrep.2010.06.001
- Benson, A. J., Bower, R. G., Frenk, C. S., et al. 2003, *ApJ*, 599, 38, doi: 10.1086/379160
- Benson, A. J., Pearce, F. R., Frenk, C. S., Baugh, C. M., & Jenkins, A. 2001, *MNRAS*, 320, 261, doi: 10.1046/j.1365-8711.2001.03966.x
- Berry, M., Somerville, R. S., Haas, M. R., et al. 2014, *MNRAS*, 441, 939, doi: 10.1093/mnras/stu613
- Bertschinger, E. 1989, *ApJ*, 340, 666, doi: 10.1086/167428
- Bigiel, F., Leroy, A., Walter, F., et al. 2008, *AJ*, 136, 2846, doi: 10.1088/0004-6256/136/6/2846
- Bigiel, F., Leroy, A. K., Walter, F., et al. 2011, *ApJ*, 730, L13, doi: 10.1088/2041-8205/730/2/L13
- Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, *Nature*, 311, 517, doi: 10.1038/311517a0

- Boselli, A., Cortese, L., Boquien, M., et al. 2014, *A&A*, 564, A66, doi: 10.1051/0004-6361/201322312
- Brennan, R., Choi, E., Somerville, R. S., et al. 2018, *ApJ*, 860, 14, doi: 10.3847/1538-4357/aac2c4
- Bryan, G. L., & Norman, M. L. 1998, *ApJ*, 495, 80, doi: 10.1086/305262
- Calette, A. R., Avila-Reese, V., Rodríguez-Puebla, A., Hernández-Toledo, H., & Pappastergis, E. 2018, *Rev. Mexicana Astron. Astrofis.*, 54, 443. <https://arxiv.org/abs/1803.07692>
- Cattaneo, A., Blaizot, J., Weinberg, D. H., et al. 2007, *MNRAS*, 377, 63, doi: 10.1111/j.1365-2966.2007.11597.x
- Chabrier, G. 2003, *PASP*, 115, 763, doi: 10.1086/376392
- Chan, T. K., Kereš, D., Wetzell, A., et al. 2018, *MNRAS*, 478, 906, doi: 10.1093/mnras/sty1153
- Christensen, C. R., Davé, R., Governato, F., et al. 2016, *ApJ*, 824, 57, doi: 10.3847/0004-637X/824/1/57
- Cicone, C., Maiolino, R., Sturm, E., et al. 2014, *A&A*, 562, A21, doi: 10.1051/0004-6361/201322464
- Cochrane, R. K., Hayward, C. C., Anglés-Alcázar, D., et al. 2019, *MNRAS*, 488, 1779, doi: 10.1093/mnras/stz1736

- Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F., & Zepf, S. E. 1994, MNRAS, 271, 781, doi: 10.1093/mnras/271.4.781
- Crain, R. A., Schaye, J., Bower, R. G., et al. 2015, MNRAS, 450, 1937, doi: 10.1093/mnras/stv725
- Creasey, P., Theuns, T., & Bower, R. G. 2013, MNRAS, 429, 1922, doi: 10.1093/mnras/sts439
- Croton, D. J., Springel, V., White, S. D. M., et al. 2006, MNRAS, 365, 11, doi: 10.1111/j.1365-2966.2005.09675.x
- Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, MNRAS, 486, 2827, doi: 10.1093/mnras/stz937
- Davé, R., Finlator, K., & Oppenheimer, B. D. 2011, MNRAS, 416, 1354, doi: 10.1111/j.1365-2966.2011.19132.x
- . 2012, MNRAS, 421, 98, doi: 10.1111/j.1365-2966.2011.20148.x
- Davé, R., Thompson, R., & Hopkins, P. F. 2016, MNRAS, 462, 3265, doi: 10.1093/mnras/stw1862
- De Lucia, G., & Blaizot, J. 2007, MNRAS, 375, 2, doi: 10.1111/j.1365-2966.2006.11287.x
- Dekel, A., & Silk, J. 1986, ApJ, 303, 39, doi: 10.1086/164050

- Dijkstra, M., & Loeb, A. 2009, MNRAS, 400, 1109, doi: 10.1111/j.1365-2966.2009.15533.x
- Dutton, A. A., Macciò, A. V., Dekel, A., et al. 2016, MNRAS, 461, 2658, doi: 10.1093/mnras/stw1537
- Efstathiou, G. 2000, MNRAS, 317, 697, doi: 10.1046/j.1365-8711.2000.03665.x
- El-Badry, K., Quataert, E., Wetzel, A., et al. 2018, MNRAS, 473, 1930, doi: 10.1093/mnras/stx2482
- Escala, I., Wetzel, A., Kirby, E. N., et al. 2018, MNRAS, 474, 2194, doi: 10.1093/mnras/stx2858
- Faerman, Y., Sternberg, A., & McKee, C. F. 2019, arXiv e-prints, arXiv:1909.09169.
<https://arxiv.org/abs/1909.09169>
- Faucher-Giguère, C.-A. 2018, MNRAS, 473, 3717, doi: 10.1093/mnras/stx2595
- Faucher-Giguère, C.-A., Feldmann, R., Quataert, E., et al. 2016, MNRAS, 461, L32, doi: 10.1093/mnrasl/slw091
- Faucher-Giguère, C.-A., Kereš, D., & Ma, C.-P. 2011, MNRAS, 417, 2982, doi: 10.1111/j.1365-2966.2011.19457.x
- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, ApJ, 703, 1416, doi: 10.1088/0004-637X/703/2/1416
- Feldmann, R. 2015, MNRAS, 449, 3274, doi: 10.1093/mnras/stv552

- Feldmann, R., Hopkins, P. F., Quataert, E., Faucher-Giguère, C.-A., & Kereš, D. 2016, MNRAS, 458, L14, doi: 10.1093/mnrasl/slw014
- Feldmann, R., Quataert, E., Hopkins, P. F., Faucher-Giguère, C.-A., & Kereš, D. 2017, MNRAS, 470, 1050, doi: 10.1093/mnras/stx1120
- Fielding, D., Quataert, E., & Martizzi, D. 2018, MNRAS, 481, 3325, doi: 10.1093/mnras/sty2466
- Fielding, D., Quataert, E., Martizzi, D., & Faucher-Giguère, C.-A. 2017a, MNRAS, 470, L39, doi: 10.1093/mnrasl/slx072
- Fielding, D., Quataert, E., McCourt, M., & Thompson, T. A. 2017b, MNRAS, 466, 3810, doi: 10.1093/mnras/stw3326
- Fielding, D. B., Ostriker, E. C., Bryan, G. L., & Jermyn, A. S. 2020a, ApJ, 894, L24, doi: 10.3847/2041-8213/ab8d2c
- Fielding, D. B., Tonnesen, S., DeFelippis, D., et al. 2020b, ApJ, 903, 32, doi: 10.3847/1538-4357/abbc6d
- Fiore, F., Feruglio, C., Shankar, F., et al. 2017, A&A, 601, A143, doi: 10.1051/0004-6361/201629478
- Fitts, A., Boylan-Kolchin, M., Elbert, O. D., et al. 2017, MNRAS, 471, 3547, doi: 10.1093/mnras/stx1757

- Fluetsch, A., Maiolino, R., Carniani, S., et al. 2019, MNRAS, 483, 4586, doi: 10.1093/mnras/sty3449
- Forbes, J. C., Krumholz, M. R., Burkert, A., & Dekel, A. 2014, MNRAS, 443, 168, doi: 10.1093/mnras/stu1142
- Forbes, J. C., Krumholz, M. R., & Speagle, J. S. 2019, MNRAS, 487, 3581, doi: 10.1093/mnras/stz1473
- Fraternali, F., & Binney, J. J. 2008, MNRAS, 386, 935, doi: 10.1111/j.1365-2966.2008.13071.x
- Garrison-Kimmel, S., Wetzel, A., Bullock, J. S., et al. 2017, MNRAS, 471, 1709, doi: 10.1093/mnras/stx1710
- Genel, S., Vogelsberger, M., Springel, V., et al. 2014, MNRAS, 445, 175, doi: 10.1093/mnras/stu1654
- Gentry, E. S., Krumholz, M. R., Dekel, A., & Madau, P. 2017, MNRAS, 465, 2471, doi: 10.1093/mnras/stw2746
- Gnedin, N. Y., & Kravtsov, A. V. 2011, ApJ, 728, 88, doi: 10.1088/0004-637X/728/2/88
- Guo, Q., White, S., Boylan-Kolchin, M., et al. 2011, MNRAS, 413, 101, doi: 10.1111/j.1365-2966.2010.18114.x

- Guo, Q., Gonzalez-Perez, V., Guo, Q., et al. 2016, MNRAS, 461, 3457, doi: 10.1093/mnras/stw1525
- Gurvich, A. B., Faucher-Giguère, C.-A., Richings, A. J., et al. 2020, MNRAS, 498, 3664, doi: 10.1093/mnras/staa2578
- Haardt, F., & Madau, P. 2001, in Clusters of Galaxies and the High Redshift Universe Observed in X-rays, ed. D. M. Neumann & J. T. V. Tran, 64. <https://arxiv.org/abs/astro-ph/0106018>
- Hafen, Z., Faucher-Giguère, C.-A., Anglés-Alcázar, D., et al. 2019, MNRAS, 488, 1248, doi: 10.1093/mnras/stz1773
- . 2020, MNRAS, 494, 3581, doi: 10.1093/mnras/staa902
- Hayward, C. C., & Hopkins, P. F. 2017, MNRAS, 465, 1682, doi: 10.1093/mnras/stw2888
- Helly, J. C., Cole, S., Frenk, C. S., et al. 2003, MNRAS, 338, 913, doi: 10.1046/j.1365-8711.2003.06152.x
- Henriques, B. M. B., Thomas, P. A., Oliver, S., & Roseboom, I. 2009, MNRAS, 396, 535, doi: 10.1111/j.1365-2966.2009.14730.x
- Henriques, B. M. B., White, S. D. M., Thomas, P. A., et al. 2013, MNRAS, 431, 3373, doi: 10.1093/mnras/stt415

- Hirschmann, M., De Lucia, G., & Fontanot, F. 2016, MNRAS, 461, 1760, doi: 10.1093/mnras/stw1318
- Hirschmann, M., Naab, T., Somerville, R. S., Burkert, A., & Oser, L. 2012, MNRAS, 419, 3200, doi: 10.1111/j.1365-2966.2011.19961.x
- Hopkins, P. F. 2015, MNRAS, 450, 53, doi: 10.1093/mnras/stv195
- Hopkins, P. F., Chan, T. K., Ji, S., et al. 2020, MNRAS, doi: 10.1093/mnras/staa3690
- Hopkins, P. F., Kereš, D., Oñorbe, J., et al. 2014, MNRAS, 445, 581, doi: 10.1093/mnras/stu1738
- Hopkins, P. F., Quataert, E., & Murray, N. 2012, MNRAS, 421, 3522, doi: 10.1111/j.1365-2966.2012.20593.x
- Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018a, MNRAS, 480, 800, doi: 10.1093/mnras/sty1690
- . 2018b, MNRAS, 477, 1578, doi: 10.1093/mnras/sty674
- Hu, C.-Y. 2019, MNRAS, 483, 3363, doi: 10.1093/mnras/sty3252
- Huang, S., Katz, N., Davé, R., et al. 2020, MNRAS, 493, 1, doi: 10.1093/mnras/staa135
- Hummels, C. B., Smith, B. D., Hopkins, P. F., et al. 2019, ApJ, 882, 156, doi: 10.3847/1538-4357/ab378f

- Hung, C.-L., Hayward, C. C., Yuan, T., et al. 2019, MNRAS, 482, 5125, doi: 10.1093/mnras/sty2970
- Jiang, F., Dekel, A., Freundlich, J., et al. 2021, MNRAS, 502, 621, doi: 10.1093/mnras/staa4034
- Jiang, F., & van den Bosch, F. C. 2016, MNRAS, 458, 2848, doi: 10.1093/mnras/stw439
- Kauffmann, G., White, S. D. M., & Guiderdoni, B. 1993, MNRAS, 264, 201, doi: 10.1093/mnras/264.1.201
- Kereš, D., Katz, N., Fardal, M., Davé, R., & Weinberg, D. H. 2009, MNRAS, 395, 160, doi: 10.1111/j.1365-2966.2009.14541.x
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, MNRAS, 363, 2, doi: 10.1111/j.1365-2966.2005.09451.x
- Khullar, S., Krumholz, M. R., Federrath, C., & Cunningham, A. J. 2019, MNRAS, 488, 1407, doi: 10.1093/mnras/stz1800
- Kim, C.-G., & Ostriker, E. C. 2015, ApJ, 802, 99, doi: 10.1088/0004-637X/802/2/99
- . 2018, ApJ, 853, 173, doi: 10.3847/1538-4357/aaa5ff
- Kim, C.-G., Ostriker, E. C., Somerville, R. S., et al. 2020a, ApJ, 900, 61, doi: 10.3847/1538-4357/aba962
- Kim, C.-G., Ostriker, E. C., Fielding, D. B., et al. 2020b, ApJ, 903, L34, doi: 10.3847/2041-8213/abc252

- Kravtsov, A. V., Gnedin, O. Y., & Klypin, A. A. 2004, *ApJ*, 609, 482, doi: 10.1086/421322
- Kroupa, P. 2001, *MNRAS*, 322, 231, doi: 10.1046/j.1365-8711.2001.04022.x
- Lagos, C. d. P., Tobar, R. J., Robotham, A. S. G., et al. 2018, *MNRAS*, 481, 3573, doi: 10.1093/mnras/sty2440
- Lan, T.-W., & Mo, H. 2019, *MNRAS*, 486, 608, doi: 10.1093/mnras/stz839
- Lee, K.-G., Krolewski, A., White, M., et al. 2018, *ApJS*, 237, 31, doi: 10.3847/1538-4365/aace58
- Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, *ApJS*, 123, 3, doi: 10.1086/313233
- Li, M., & Bryan, G. L. 2019, arXiv e-prints, arXiv:1910.09554. <https://arxiv.org/abs/1910.09554>
- . 2020, *ApJ*, 890, L30, doi: 10.3847/2041-8213/ab7304
- Li, M., Bryan, G. L., & Ostriker, J. P. 2017, *ApJ*, 841, 101, doi: 10.3847/1538-4357/aa7263
- Li, M., & Tonnesen, S. 2019, arXiv e-prints, arXiv:1910.14235. <https://arxiv.org/abs/1910.14235>
- . 2020, *ApJ*, 898, 148, doi: 10.3847/1538-4357/ab9f9f

- Lu, Y., Benson, A., Wetzel, A., et al. 2017, *ApJ*, 846, 66, doi: 10.3847/1538-4357/aa845e
- Lu, Y., Kereš, D., Katz, N., et al. 2011a, *MNRAS*, 416, 660, doi: 10.1111/j.1365-2966.2011.19072.x
- Lu, Y., Mo, H. J., & Wechsler, R. H. 2015, *MNRAS*, 446, 1907, doi: 10.1093/mnras/stu2215
- Lu, Y., Mo, H. J., Weinberg, M. D., & Katz, N. 2011b, *MNRAS*, 416, 1949, doi: 10.1111/j.1365-2966.2011.19170.x
- Ma, X., Hopkins, P. F., Faucher-Giguère, C.-A., et al. 2016, *MNRAS*, 456, 2140, doi: 10.1093/mnras/stv2659
- Ma, X., Hopkins, P. F., Wetzel, A. R., et al. 2017, *MNRAS*, 467, 2430, doi: 10.1093/mnras/stx273
- Ma, X., Kasen, D., Hopkins, P. F., et al. 2015, *MNRAS*, 453, 960, doi: 10.1093/mnras/stv1679
- Mac Low, M.-M., & Ferrara, A. 1999, *ApJ*, 513, 142, doi: 10.1086/306832
- Maller, A. H., & Bullock, J. S. 2004, *MNRAS*, 355, 694, doi: 10.1111/j.1365-2966.2004.08349.x
- Mandelker, N., Nagai, D., Aung, H., et al. 2020, *MNRAS*, 494, 2641, doi: 10.1093/mnras/staa812

- Marasco, A., Fraternali, F., Heald, G., et al. 2019, *A&A*, 631, A50, doi: 10.1051/0004-6361/201936338
- Martizzi, D., Fielding, D., Faucher-Giguère, C.-A., & Quataert, E. 2016, *MNRAS*, 459, 2311, doi: 10.1093/mnras/stw745
- Meiksin, A., Khochfar, S., Paardekooper, J.-P., Dalla Vecchia, C., & Kohn, S. 2017, *MNRAS*, 471, 3632, doi: 10.1093/mnras/stx1857
- Mitchell, P. D., Schaye, J., Bower, R. G., & Crain, R. A. 2019, arXiv e-prints, arXiv:1910.09566. <https://arxiv.org/abs/1910.09566>
- . 2020, *MNRAS*, 494, 3971, doi: 10.1093/mnras/staa938
- Mitchell, P. D., Lacey, C. G., Lagos, C. D. P., et al. 2018, *MNRAS*, 474, 492, doi: 10.1093/mnras/stx2770
- Mitra, S., Davé, R., & Finlator, K. 2015, *MNRAS*, 452, 1184, doi: 10.1093/mnras/stv1387
- Mo, H., van den Bosch, F. C., & White, S. 2010, *Galaxy Formation and Evolution* (Cambridge University Press)
- Mo, H. J., Mao, S., & White, S. D. M. 1998, *MNRAS*, 295, 319, doi: 10.1046/j.1365-8711.1998.01227.x
- Moster, B. P., Naab, T., & White, S. D. M. 2013, *MNRAS*, 428, 3121, doi: 10.1093/mnras/sts261

- . 2018, MNRAS, 477, 1822, doi: 10.1093/mnras/sty655
- Motwani, B., Genel, S., Bryan, G. L., et al. 2020, arXiv e-prints, arXiv:2006.16314.
<https://arxiv.org/abs/2006.16314>
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., et al. 2015, MNRAS, 454, 2691,
doi: 10.1093/mnras/stv2126
- . 2017, MNRAS, 468, 4170, doi: 10.1093/mnras/stx667
- Murray, N., Quataert, E., & Thompson, T. A. 2005, ApJ, 618, 569, doi: 10.1086/
426067
- Mutch, S. J., Geil, P. M., Poole, G. B., et al. 2016, MNRAS, 462, 250, doi: 10.1093/
mnras/stw1506
- Naab, T., & Ostriker, J. P. 2017, ARA&A, 55, 59, doi: 10.1146/
annurev-astro-081913-040019
- Narayanan, D., Krumholz, M. R., Ostriker, E. C., & Hernquist, L. 2012, MNRAS, 421,
3127, doi: 10.1111/j.1365-2966.2012.20536.x
- Neistein, E., Khochfar, S., Dalla Vecchia, C., & Schaye, J. 2012, MNRAS, 421, 3579,
doi: 10.1111/j.1365-2966.2012.20584.x
- Nelson, D., Pillepich, A., Springel, V., et al. 2019, MNRAS, 490, 3234, doi: 10.1093/
mnras/stz2306
- Noh, Y., & McQuinn, M. 2014, MNRAS, 444, 503, doi: 10.1093/mnras/stu1412

- Okamoto, T., Gao, L., & Theuns, T. 2008, MNRAS, 390, 920, doi: 10.1111/j.1365-2966.2008.13830.x
- Oppenheimer, B. D., & Davé, R. 2006, MNRAS, 373, 1265, doi: 10.1111/j.1365-2966.2006.10989.x
- . 2009, MNRAS, 395, 1875, doi: 10.1111/j.1365-2966.2009.14676.x
- Oppenheimer, B. D., Davé, R., Kereš, D., et al. 2010, MNRAS, 406, 2325, doi: 10.1111/j.1365-2966.2010.16872.x
- Pandya, V., Somerville, R. S., Anglés-Alcázar, D., et al. 2020, ApJ, 905, 4, doi: 10.3847/1538-4357/abc3c1
- Pandya, V., Fielding, D., Anglés-Alcázar, D., et al. 2021, arXiv e-prints, arXiv:2103.06891. <https://arxiv.org/abs/2103.06891>
- Parsotan, T., Cochrane, R. K., Hayward, C. C., et al. 2021, MNRAS, 501, 1591, doi: 10.1093/mnras/staa3765
- Pawlik, A. H., & Schaye, J. 2009, MNRAS, 396, L46, doi: 10.1111/j.1745-3933.2009.00659.x
- Peeples, M. S., Werk, J. K., Tumlinson, J., et al. 2014, ApJ, 786, 54, doi: 10.1088/0004-637X/786/1/54
- Peeples, M. S., Corlies, L., Tumlinson, J., et al. 2019, ApJ, 873, 129, doi: 10.3847/1538-4357/ab0654

- Pillepich, A., Springel, V., Nelson, D., et al. 2018, MNRAS, 473, 4077, doi: 10.1093/mnras/stx2656
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, A&A, 571, A16, doi: 10.1051/0004-6361/201321591
- . 2016, A&A, 594, A13, doi: 10.1051/0004-6361/201525830
- Popping, G., Somerville, R. S., & Trager, S. C. 2014, MNRAS, 442, 2398, doi: 10.1093/mnras/stu991
- Popping, G., Pillepich, A., Somerville, R. S., et al. 2019, ApJ, 882, 137, doi: 10.3847/1538-4357/ab30f2
- Porter, L. A., Somerville, R. S., Primack, J. R., & Johansson, P. H. 2014, MNRAS, 444, 942, doi: 10.1093/mnras/stu1434
- Qin, Y., Duffy, A. R., Mutch, S. J., et al. 2018, MNRAS, 477, 1318, doi: 10.1093/mnras/sty767
- Roberts-Borsani, G. W., Saintonge, A., Masters, K. L., & Stark, D. V. 2020, MNRAS, 493, 3081, doi: 10.1093/mnras/staa464
- Rodríguez-Puebla, A., Primack, J. R., Avila-Reese, V., & Faber, S. M. 2017, MNRAS, 470, 651, doi: 10.1093/mnras/stx1172
- Rodríguez-Puebla, A., Primack, J. R., Behroozi, P., & Faber, S. M. 2016, MNRAS, 455, 2592, doi: 10.1093/mnras/stv2513

- Rudie, G. C., Steidel, C. C., Trainor, R. F., et al. 2012, *ApJ*, 750, 67, doi: 10.1088/0004-637X/750/1/67
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, 446, 521, doi: 10.1093/mnras/stu2058
- Schneider, E. E., Ostriker, E. C., Robertson, B. E., & Thompson, T. A. 2020, *ApJ*, 895, 43, doi: 10.3847/1538-4357/ab8ae8
- Smith, M. C., Sijacki, D., & Shen, S. 2018, *MNRAS*, 478, 302, doi: 10.1093/mnras/sty994
- Somerville, R. S., & Davé, R. 2015, *ARA&A*, 53, 51, doi: 10.1146/annurev-astro-082812-140951
- Somerville, R. S., Gilmore, R. C., Primack, J. R., & Domínguez, A. 2012, *MNRAS*, 423, 1992, doi: 10.1111/j.1365-2966.2012.20490.x
- Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., & Hernquist, L. 2008a, *MNRAS*, 391, 481, doi: 10.1111/j.1365-2966.2008.13805.x
- Somerville, R. S., Popping, G., & Trager, S. C. 2015, *MNRAS*, 453, 4337, doi: 10.1093/mnras/stv1877
- Somerville, R. S., & Primack, J. R. 1999, *MNRAS*, 310, 1087, doi: 10.1046/j.1365-8711.1999.03032.x

- Somerville, R. S., Primack, J. R., & Faber, S. M. 2001, MNRAS, 320, 504, doi: 10.1046/j.1365-8711.2001.03975.x
- Somerville, R. S., Barden, M., Rix, H.-W., et al. 2008b, ApJ, 672, 776, doi: 10.1086/523661
- Somerville, R. S., Olsen, C., Yung, L. Y. A., et al. 2021, MNRAS, 502, 4858, doi: 10.1093/mnras/stab231
- Sparre, M., Hayward, C. C., Feldmann, R., et al. 2017, MNRAS, 466, 88, doi: 10.1093/mnras/stw3011
- Springel, V., & Hernquist, L. 2003, MNRAS, 339, 289, doi: 10.1046/j.1365-8711.2003.06206.x
- Steinwandel, U. P., Moster, B. P., Naab, T., Hu, C.-Y., & Walch, S. 2020, MNRAS, 495, 1035, doi: 10.1093/mnras/staa821
- Stern, J., Fielding, D., Faucher-Giguère, C.-A., & Quataert, E. 2019, MNRAS, 488, 2549, doi: 10.1093/mnras/stz1859
- Stern, J., Faucher-Giguère, C.-A., Fielding, D., et al. 2020, arXiv e-prints, arXiv:2006.13976. <https://arxiv.org/abs/2006.13976>
- Stevens, A. R. H., Lagos, C. d. P., Contreras, S., et al. 2017, MNRAS, 467, 2066, doi: 10.1093/mnras/stx243

- Stinson, G., Seth, A., Katz, N., et al. 2006, MNRAS, 373, 1074, doi: 10.1111/j.1365-2966.2006.11097.x
- Stringer, M. J., Brooks, A. M., Benson, A. J., & Governato, F. 2010, MNRAS, 407, 632, doi: 10.1111/j.1365-2966.2010.16944.x
- Sutherland, R. S., & Dopita, M. A. 1993, ApJS, 88, 253, doi: 10.1086/191823
- Tacchella, S., Bose, S., Conroy, C., Eisenstein, D. J., & Johnson, B. D. 2018, ApJ, 868, 92, doi: 10.3847/1538-4357/aae8e0
- Tollet, É., Cattaneo, A., Macciò, A. V., Dutton, A. A., & Kang, X. 2019, MNRAS, 485, 2511, doi: 10.1093/mnras/stz545
- Tumlinson, J., Peebles, M. S., & Werk, J. K. 2017, ARA&A, 55, 389, doi: 10.1146/annurev-astro-091916-055240
- van de Voort, F., Schaye, J., Booth, C. M., Haas, M. R., & Dalla Vecchia, C. 2011, MNRAS, 414, 2458, doi: 10.1111/j.1365-2966.2011.18565.x
- van de Voort, F., Springel, V., Mandelker, N., van den Bosch, F. C., & Pakmor, R. 2019, MNRAS, 482, L85, doi: 10.1093/mnrasl/sly190
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014, MNRAS, 444, 1518, doi: 10.1093/mnras/stu1536
- Voit, G. M., Bryan, G. L., O'Shea, B. W., & Donahue, M. 2015, ApJ, 808, L30, doi: 10.1088/2041-8205/808/1/L30

- Walch, S., Girichidis, P., Naab, T., et al. 2015, MNRAS, 454, 238, doi: 10.1093/mnras/stv1975
- Wang, L., Dutton, A. A., Stinson, G. S., et al. 2015, MNRAS, 454, 83, doi: 10.1093/mnras/stv1937
- Wechsler, R. H., & Tinker, J. L. 2018, ARA&A, 56, 435, doi: 10.1146/annurev-astro-081817-051756
- Weinmann, S. M., Pasquali, A., Oppenheimer, B. D., et al. 2012, MNRAS, 426, 2797, doi: 10.1111/j.1365-2966.2012.21931.x
- Wellons, S., Faucher-Giguère, C.-A., Anglés-Alcázar, D., et al. 2020, MNRAS, 497, 4051, doi: 10.1093/mnras/staa2229
- Wetzell, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, ApJ, 827, L23, doi: 10.3847/2041-8205/827/2/L23
- Wheeler, C., Hopkins, P. F., Pace, A. B., et al. 2019, MNRAS, 490, 4447, doi: 10.1093/mnras/stz2887
- White, C. E., Somerville, R. S., & Ferguson, H. C. 2015, ApJ, 799, 201, doi: 10.1088/0004-637X/799/2/201
- White, S. D. M., & Frenk, C. S. 1991, ApJ, 379, 52, doi: 10.1086/170483
- White, S. D. M., & Rees, M. J. 1978, MNRAS, 183, 341, doi: 10.1093/mnras/183.3.341

- Wright, R. J., Lagos, C. d. P., Power, C., & Mitchell, P. D. 2020, MNRAS, doi: 10.1093/mnras/staa2359
- Yates, R. M., Henriques, B., Thomas, P. A., et al. 2013, MNRAS, 435, 3500, doi: 10.1093/mnras/stt1542
- Yoshida, N., Stoehr, F., Springel, V., & White, S. D. M. 2002, MNRAS, 335, 762, doi: 10.1046/j.1365-8711.2002.05661.x
- Yung, L. Y. A., Somerville, R. S., Finkelstein, S. L., Popping, G., & Davé, R. 2019a, MNRAS, 483, 2983, doi: 10.1093/mnras/sty3241
- Yung, L. Y. A., Somerville, R. S., Popping, G., et al. 2019b, MNRAS, 490, 2855, doi: 10.1093/mnras/stz2755