

How to model supernovae in simulations of star and galaxy formation

Philip F. Hopkins,^{1★} Andrew Wetzel,^{1,2,3†} Dušan Kereš,⁴
Claude-André Faucher-Giguère,⁵ Eliot Quataert,⁶ Michael Boylan-Kolchin,⁷
Norman Murray,⁸ Christopher C. Hayward^{9,10} and Kareem El-Badry⁶

¹TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

²The Observatories of the Carnegie Institution for Science, Pasadena, CA 91101, USA

³Department of Physics, University of California, Davis, CA 95616, USA

⁴Department of Physics, Center for Astrophysics and Space Science, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA

⁵Department of Physics and Astronomy and CIERA, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

⁶Department of Astronomy and Theoretical Astrophysics Center, University of California Berkeley, Berkeley, CA 94720, USA

⁷Department of Astronomy, The University of Texas at Austin, 2515 Speedway, Stop C1400, Austin, TX 78712, USA

⁸Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, ON M5S 3H8, Canada

⁹Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

¹⁰Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

Accepted 2018 March 12. Received 2018 March 12; in original form 2017 July 14

ABSTRACT

We study the implementation of mechanical feedback from supernovae (SNe) and stellar mass loss in galaxy simulations, within the Feedback In Realistic Environments (FIRE) project. We present the FIRE-2 algorithm for coupling mechanical feedback, which can be applied to any hydrodynamics method (e.g. fixed-grid, moving-mesh, and mesh-less methods), and black hole as well as stellar feedback. This algorithm ensures manifest conservation of mass, energy, and momentum, and avoids imprinting ‘preferred directions’ on the ejecta. We show that it is critical to incorporate both momentum and thermal energy of mechanical ejecta in a self-consistent manner, accounting for SNe cooling radii when they are not resolved. Using idealized simulations of single SN explosions, we show that the FIRE-2 algorithm, independent of resolution, reproduces converged solutions in both energy and momentum. In contrast, common ‘fully thermal’ (energy-dump) or ‘fully kinetic’ (particle-kicking) schemes in the literature depend strongly on resolution: when applied at mass resolution $\gtrsim 100 M_{\odot}$, they diverge by orders of magnitude from the converged solution. In galaxy-formation simulations, this divergence leads to orders-of-magnitude differences in galaxy properties, unless those models are adjusted in a resolution-dependent way. We show that *all models that individually time-resolve SNe converge to the FIRE-2 solution at sufficiently high resolution ($< 100 M_{\odot}$)*. However, in both idealized single-SN simulations and cosmological galaxy-formation simulations, the FIRE-2 algorithm converges much faster than other sub-grid models *without* re-tuning parameters.

Key words: stars: formation – galaxies: active – galaxies: evolution – galaxies: formation – cosmology: theory.

1 INTRODUCTION

Stellar feedback is critical in understanding galaxy formation. Without it, gas accretes into dark matter haloes and galaxies, cools rapidly on a time-scale much faster than the dynamical time, collapses, frag-

ments, and forms stars on a free fall-time (Bournaud et al. 2010; Hopkins, Quataert & Murray 2011; Dobbs, Burkert & Pringle 2011; Harper-Clark & Murray 2011; Tasker 2011), inevitably turning most of the baryons into stars on cosmological time-scales (Katz, Weinberg & Hernquist 1996; Somerville & Primack 1999; Cole et al. 2000; Springel & Hernquist 2003b; Kereš et al. 2009). But observations imply that, on galactic scales, only a few percent of gas turns into stars per free-fall time (Kennicutt 1998), while individual giant molecular clouds (GMCs) disrupt after forming just a few

* E-mail: phopkins@caltech.edu

† Caltech-Carnegie Fellow

percent of their mass in stars (Zuckerman & Evans 1974; Williams & McKee 1997; Evans 1999; Evans et al. 2009). Similarly, galaxies retain and turn into stars just a few percent of the universal baryon fraction (Conroy, Wechsler & Kravtsov 2006; Behroozi, Conroy & Wechsler 2010; Moster et al. 2010), and both direct observations of galactic winds (Martin 1999; Heckman et al. 2000; Sato et al. 2009; Steidel et al. 2010; Coil et al. 2011) and indirect constraints on the inter-galactic and circum-galactic medium (IGM/CGM; Aguirre et al. 2001; Pettini et al. 2003; Songaila 2005; Oppenheimer & Davé 2006; Martin et al. 2010) require that a large fraction of the baryons have been ‘processed’ in galaxies via their accretion, enrichment, and expulsion in supergalactic outflows.

Many different feedback processes contribute to these galactic winds and ultimately the self-regulation of galactic star formation, including protostellar jets, photo-heating, stellar mass loss (O/B and AGB-star winds), radiation pressure, and supernovae (SNe) Types Ia and II (see Evans et al. 2009; Lopez et al. 2011, and references therein). Older galaxy-formation simulations could not resolve the effects of these different processes (even on relatively large scales within the galactic disc), so they used simplified prescriptions to model galactic winds. However, a new generation of high-resolution simulations has emerged with the ability to resolve multi-phase structure in the ISM and so begin to directly incorporate these distinct feedback processes (Hopkins, Quataert & Murray 2011, 2012; Tasker 2011; Agertz et al. 2013; Kannan et al. 2014). One example is the Feedback In Realistic Environments (FIRE)¹ project (Hopkins et al. 2014). These and similar simulations have demonstrated predictions in reasonable agreement with observations for a wide variety of galaxy properties (e.g. Feldmann et al. 2016; Ma et al. 2016; Wetzel et al. 2016; Sparre et al. 2017).

In a companion paper, Hopkins et al. (2017, hereafter Paper I), we presented an updated version of the FIRE code. We refer to this updated FIRE version as ‘FIRE-2’ and the older FIRE implementation as ‘FIRE-1’. We explored how a wide range of numerical effects (resolution, hydrodynamic solver, details of the cooling, and star formation algorithm) influence the results of galaxy-formation simulations. We compared these to the effects of feedback and concluded that mechanical feedback, particularly from Type-II SNe, has much larger effects on galaxy formation (specifically properties such as galaxy masses, star formation histories, metallicities, rotation curves, sizes, and morphologies) compared to the various numerical details studied. This is consistent with a number of previous studies (Abadi et al. 2003; Governato et al. 2004; Robertson et al. 2004; Stinson et al. 2006; Zavala, Okamoto & Frenk 2008; Scannapieco et al. 2012). However, in galaxy-formation simulations, the actual implementation of SNe feedback, and the physical assumptions associated with it, often differ significantly between different codes. This can have significant effects on the predictions for galaxy formation (see Scannapieco et al. 2012; Kim et al. 2016; Rosdahl et al. 2017).

In this paper, we present a detailed study of the algorithmic implementation of SNe feedback and its effects, in the context of the FIRE-2 simulations. We emphasize that there are two *separate* aspects of mechanical feedback that must be explored.

First, the *numerical* aspects of the algorithmic coupling. Given some feedback ‘products’ (mass, metals, energy, momentum) from a star, these must be deposited in the surrounding gas. Any good

algorithm should respect certain basic considerations: conservation (of mass, energy, and momentum), statistical isotropy² (avoiding imprinting preferred directions that depend on either the numerical grid axes or the arbitrary gas configuration around the feedback source), and convergence. We will show that accomplishing these is non-trivial, and that many algorithms in common use (including the older algorithm that we used in FIRE-1³) do not respect all of them.

Second, the *physics* of the coupling must be explored. At any finite resolution, there is a ‘sub-grid scale’ – the space or mass between a star particle and the centre of the nearest gas resolution element, for example. An ideal implementation of the feedback coupling should exactly reproduce the converged solution, if we were to populate that space with infinite resolution – in other words, our coupling should be equivalent to ‘down-grading’ the resolution of a high-resolution case, given the *same* physical assumptions used in the larger-scale simulation. We use a suite of simulations of isolated SNe (with otherwise identical physics to our galaxy-scale simulations) to show that a well-posed algorithm of this nature must account for *both* thermal *and* kinetic energy of the ejecta as they couple in a specific manner. This forms the basis for the default treatment of SNe in the FIRE simulations (introduced in Hopkins et al. 2014), and similar to subsequent implementations in simulations by e.g. Kimm & Cen 2014; Rosdahl et al. 2017). In contrast, we show that coupling only thermal or kinetic energy leads to strongly resolution-dependent errors, which in turn can produce order-of-magnitude too-large or too-small galaxy masses. To predict reasonable masses, such models must be modified (a.k.a. ‘re-tuned’) at each resolution level. This is even more severe in ‘delayed cooling’ or ‘target temperature’ models which are explicitly intended for low-resolution applications, and are not designed to converge to the exact solution at high resolution. This explains many seemingly contradictory conclusions in the literature regarding the implementation of feedback. In contrast, we will show that the mechanical feedback models proposed here reproduce the high-resolution solution in idealized problems at *all* resolution levels that we explore, converge much more rapidly in cosmological galaxy-formation simulations, and (perhaps most importantly) represent the solution towards which other less-accurate ‘sub-grid’ SNe treatments (at least those which do not artificially modify the cooling physics) converge at very high resolution.

Our study here is relevant for simulations of the ISM and galaxy formation with mass resolution in the range $\sim 10 - 10^6 M_{\odot}$; we will show that at resolution higher than this, the numerical details have weak effects because early SN blastwave evolution is explicitly well-resolved. Conversely, at lower resolution than this, treating

² Throughout the text, we use the term ‘statistical isotropy’ to refer to a specific, desirable property of the numerical feedback-coupling algorithm. Namely, that the algorithm does not un-physically systematically bias the ejecta into certain directions (or otherwise ‘imprint’ preferred directions) for numerical reasons. Of course, ejecta may be intrinsically anisotropic in the SN frame, and there can be global anisotropies sourced by e.g. pressure gradients and galaxy morphology, but these can only be captured properly if the ejecta-coupling *algorithm* is statistically isotropic.

³ To be specific (this will be discussed below): the FIRE-1 algorithm used the ‘non-conservative method’ defined in Section 2.2.4, with a less-accurate SPH approximation of the solid angle subtended by neighbours (ω_b defined in Section 2.2.3 set $\propto m_b/\rho_b$), and only coupled to the nearest neighbours for each SN instead of using the bi-directional search defined in Section 2.2.2 and needed to ensure statistical isotropy.

¹ See the FIRE project website: <http://fire.northwestern.edu>. For additional movies and images of FIRE simulations, see: <http://www.tapir.caltech.edu/~phopkins/Site/animations>.

Table 1. FIRE-2 simulations run to $z = 0$ used for our case studies.

Simulation Name	$M_{\text{halo}}^{\text{vir}}$ (M_{\odot})	R_{vir} (kpc)	M_{*} (M_{\odot})	$R_{1/2}$ (kpc)	$m_i, 1000$ ($1000 M_{\odot}$)	$\epsilon_{\text{gas}}^{\text{MIN}}$ (pc)	$r_{\text{DM}}^{\text{conv}}$ (pc)	Notes
m10q	8.0e9	52.4	1.8e6	0.63	0.25	0.52	73	Isolated dwarf in an early-forming halo. Forms a dSph with a bursty SFH.
m12i	1.2e12	275	6.5e10	2.9	7.0	0.38	150	‘Latte’ primary halo from Wetzel et al. (2016). Thin disc with a flat SFH.

Parameters describing the FIRE-2 simulations from Hopkins et al. (2017) that we use for our case studies. Halo and stellar properties listed refer only to the original ‘target’ halo around which the high-resolution region is centred. All properties listed refer to our highest-resolution simulation using the standard, default FIRE-2 physics and numerical methods. All units are physical. (1) Simulation Name: Designation used throughout this paper. (2) $M_{\text{halo}}^{\text{vir}}$: Virial mass (following Bryan & Norman 1998) of the ‘target’ halo at $z = 0$. (3) R_{vir} : Virial radius at $z = 0$. (4) M_{*} : Stellar mass of the central galaxy at $z = 0$. (5) $R_{1/2}$: Half-mass radius of the stars in the central M_{*} at $z = 0$. (6) $m_i, 1000$: Mass resolution: the baryonic (gas or star) particle/element mass, in units of $1000 M_{\odot}$. The DM particle mass is always larger a factor ≈ 5 (the universal ratio). (7) $\epsilon_{\text{gas}}^{\text{MIN}}$: Minimum gravitational force softening reached by the gas in the simulation (gas softenings are adaptive so always exactly match the hydrodynamic resolution or inter-particle spacing); the Plummer-equivalent softening is $\approx 0.7 \epsilon_{\text{gas}}$. (8) $r_{\text{DM}}^{\text{conv}}$: Radius of convergence in the dark matter (DM) properties, in DM-only simulations. This is based on the Power et al. (2003) criterion using the best estimate from Hopkins et al. (2017) as to where the DM density profile is converged to within < 10 per cent. The DM force softening is much less important and has no appreciable effects on any results shown here, so is simply fixed to 40 pc for all runs here.

individual SN events becomes meaningless (necessitating a different sort of ‘sub-grid’ approach).

In Section 2 we provide a summary of the FIRE-2 simulations (Section 2.1), a detailed description of the numerical algorithm for mechanical feedback coupling (Section 2.2), and a detailed motivation and description of the physical breakdown between kinetic and thermal energy (Section 2.3). We note that Paper I includes complete details of all aspects of the simulations here, necessary to fully reproduce our results. In Section 3 we validate the numerical coupling algorithm (conservation, statistical isotropy, and convergence) and explore the effects of alternative coupling schemes on full galaxy formation simulations. In Section 4 we validate the physical breakdown of coupled kinetic/thermal energy, compare this to simulations of individual SN explosions at extremely high resolution, and explore how different choices which neglect these physics alter the predictions of full galaxy formation simulations. We briefly discuss non-convergent alternative models (e.g. ‘delayed cooling’ and ‘target temperature’ models) but provide more detailed tests of these in the Appendices. In Section 5 we summarize our conclusions. Additional tests are discussed in the Appendices.

2 METHODS AND PHYSICAL MOTIVATION

2.1 Overview and methods other than mechanical feedback

The simulations in this paper were run as part of the Feedback in Realistic Environments (FIRE) project, using the FIRE-2 version of the code detailed in Paper I. Our default simulations are exactly those in Paper I; we will vary the SNe algorithm to explore how this alters galaxy formation, but all other simulation properties, physics, and numerical choices are held fixed. For detailed exploration of how those numerical details alter galaxy formation, we refer to Paper I. The simulations were run using GIZMO⁴ (Hopkins 2015), in its meshless finite-mass MFM mode. This is a mesh-free, finite-volume Lagrangian Godunov method which provides adaptive spatial resolution together with conservation of mass, energy, momentum, and angular momentum, and the ability to accurately capture shocks and fluid mixing instabilities (combining advantages of both grid-based and smoothed-particle hydrodynamics methods). For extensive test problems, see Hopkins (2015); Hopkins & Raives (2016); Hopkins (2016, 2017); for tests of the methods specific to these simulations, see Paper I.

⁴ A public version of GIZMO is available at <http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>.

These simulations are cosmological ‘zoom-in’ runs that follow the Lagrangian region that surrounds a galaxy at $z = 0$ (out to several virial radii) from seed perturbations at $z = 100$. Gravity is solved for collisional (gas) and collisionless (stars and dark matter) species with adaptive gravitational softening so hydrodynamic and force softening are always matched. Gas cooling is followed self-consistently from $T = 10 - 10^{10}$ K including free–free, Compton, metal-line, molecular, fine-structure, dust collisional, and cosmic ray processes, photoelectric and photo-ionization heating by both local sources and a uniform but redshift-dependent meta-galactic background, and self-shielding. Gas is turned into stars using a sink-particle prescription (gas which is locally self-gravitating at the resolution scale following Hopkins, Narayanan & Murray 2013, self-shielding/molecular following Krumholz & Gnedin 2011, Jeans unstable, and denser than $n_{\text{crit}} > 1000 \text{ cm}^{-3}$ is converted into star particles on a free-fall time). Star particles are then treated as single-age stellar populations with all IMF-averaged feedback properties calculated from STARBURST99 (Leitherer et al. 1999) assuming a Kroupa (2001) IMF. We then explicitly treat feedback from SNe (both Types Ia and II), stellar mass loss (O/B and AGB winds), and radiation (photo-ionization and photoelectric heating and UV/optical/IR radiation pressure), with implementations at the resolution-scale described in Paper I and here.

Paper I provides a complete description of all aspects of the numerical methods. In this paper, we study the mechanical feedback algorithm, used for SNe and stellar mass loss. In a companion paper (henceforth Paper III), we study the radiation feedback algorithm.

For simplicity, we focus our study here on two example galaxies: **m10q** is a dwarf galaxy and **m12i** is a Milky Way (MW)-mass galaxy. Table 1 lists their properties. Both were studied extensively in Paper I. The star formation history, stellar mass, and mean stellar-mass weighted metallicity of each galaxy as a function of cosmic time, as well as the $z = 0$ baryonic and dark matter mass profiles and rotation curves, will be discussed below. We have explicitly verified that the conclusions drawn here regarding mechanical feedback from our **m10q** and **m12i** simulations are robust across simulations of several different galaxies/haloes at dwarf and MW mass scales, respectively.

2.2 Mechanical feedback coupling algorithm

2.2.1 Determining when events occur

Once a star particle forms, the SNe rate is taken from stellar evolution models, assuming the particle represents an IMF-averaged

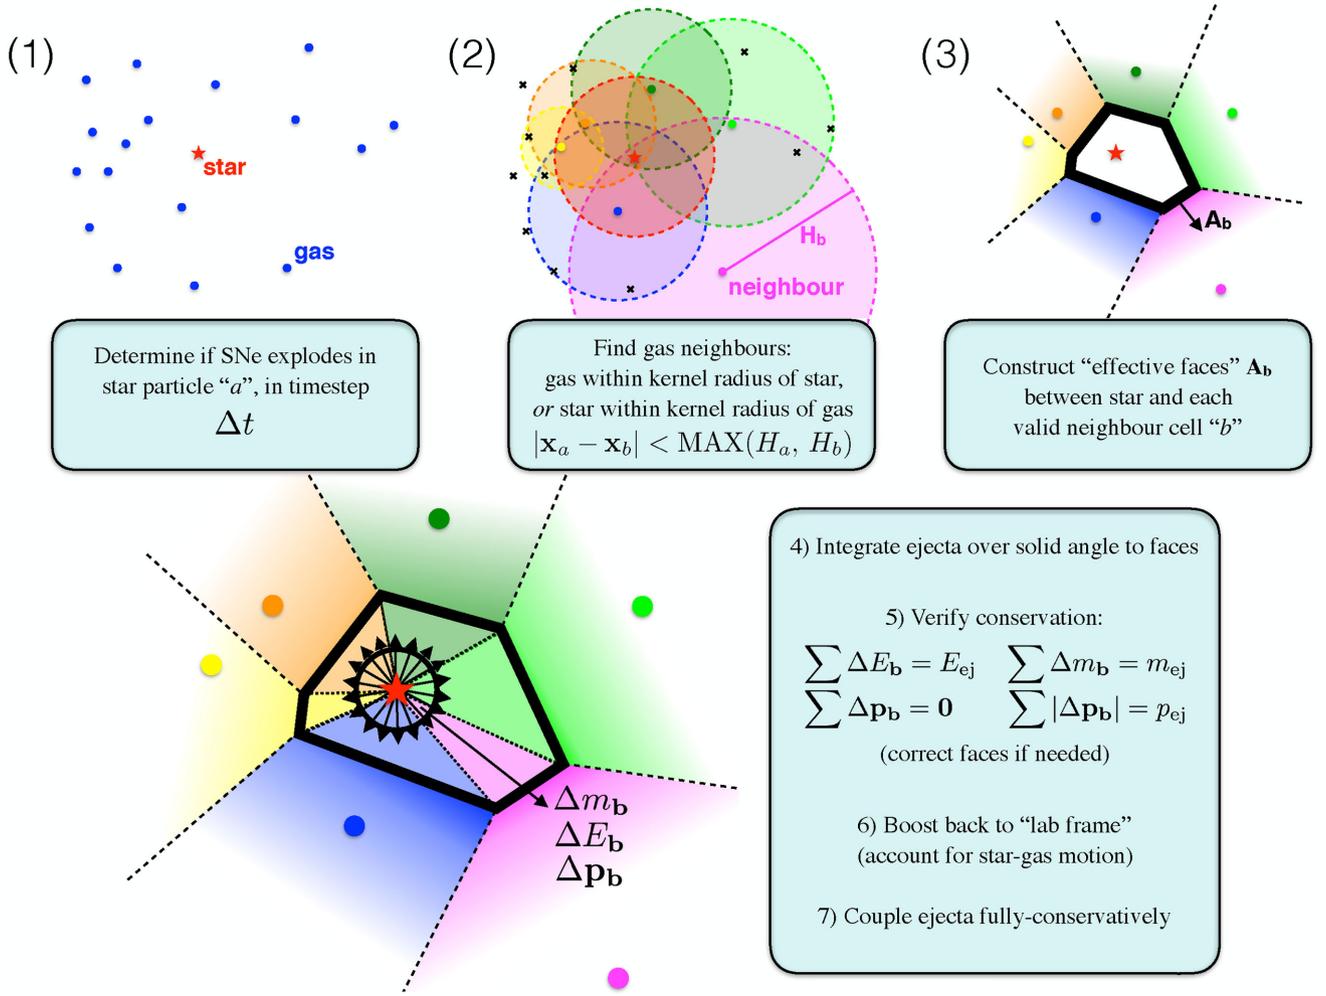


Figure 1. Cartoon illustrating the numerical algorithm for coupling mechanical feedback, described in detail in Section 2.2. (1) We determine from stellar evolution tracks (Paper I, Appendix A) whether a star particle is a source of mechanical feedback (SNe, stellar mass loss, short-range radiation pressure) at a given time-step. (2) We identify valid interacting neighbours for the star (Section 2.2.2). The search includes not only gas within the nearest neighbour search radius, H_a , of the star, but *also* gas for which the star falls within the gas element’s search radius, H_b . This typically corresponds to lower-density gas, and in this example, the two neighbours (magenta+lime) to the right of the star would not be included if only gas within H_a were used. This would artificially prevent us from coupling momentum in that direction, violating statistical isotropy. (3) Construct the ‘effective faces’ of the interacting gas elements, as seen by the star (Section 2.2.3). This example uses a Voronoi tessellation, which is similar to the result from our default MFM method. (4) Integrate the feedback quantities (mass, metals, energy, momentum), assumed to be isotropically emitted from the star (in its rest frame), over the solid angle subtended by the effective face of each gas element, to determine the fluxes into each cell. Integrate through to the face to account for the PdV work done between star and face (Section 2.3). (5) Verify that the fluxes maintain machine-accurate mass/energy/momentum conservation: if not, re-normalize the faces to correct them such that this is satisfied (Section 2.2.4). (6) Boost to the lab frame, to account for any relative star-gas motion (Section 2.2.5). (7) Finally, couple the fluxes, maintaining exact conservation, giving the updated mass, metallicity, thermal and kinetic energy, and momentum to the gas element.

population of a given age (since it formed) and abundances (inherited from its progenitor gas element). Given the particle masses and time-steps ($\Delta t \sim 100\text{--}1000$ yr) for young star particles, the expected number of SNe per particle per time-step is always $\ll 1$. To determine if an event occurs, we therefore draw from a binomial distribution at each time-step given the expected rate $\langle N \rangle = (dN/dM_* dt) m_i \Delta t$, where $(dN/dM_* dt)$ is the IMF-averaged SNe rate per unit mass for a single stellar population of the age and metallicity of the star particle and m_i is the star particle mass. For continuous mass-loss processes such as O/B or AGB winds, an ‘event’ occurs every time-step, with mass loss $\Delta M_* = \Delta t \dot{M}_*$ and the associated kinetic luminosity. See Paper I for details and tabulations of the relevant rates.

Consider a time t_a (time-step Δt), during which a mechanical feedback ‘event’ occurs sourced at some location \mathbf{x}_a (e.g. the location of a star particle ‘a’ in which an SN explodes). Our focus in this paper is how to treat this event. Fig. 1 provides an illustration of our algorithm. We first define a set of conserved quantities: mass m_{ej} , metals $m_{Z, ej}$, momentum $p_{ej} = m_{ej} v_{ej}$, and energy E_{ej} , which must be ‘injected’ into the neighbouring gas via some numerical fluxes.

2.2.2 Finding neighbours to couple

We define an effective neighbour number N_* the same as for the hydrodynamics, $N_* = (4\pi/3) H_a^3 \bar{n}_a(H_a)$, where,

$\bar{n}_a = \sum W(\mathbf{x}_{ba} \equiv \mathbf{x}_b - \mathbf{x}_a, H_a)$, W is the kernel function, and H_a is the search radius around the star (set by N_* , which is the ‘fixed’ parameter).⁵ Thus we obtain all gas elements b within a radius $|\mathbf{x}_{ba}| < H_a$.

However, severe pathologies can occur if feedback is coupled *only* to the nearest neighbouring gas to the star. For example, in an infinitely thin, dense disc of gas surrounding the star particle, with a tenuous atmosphere in the vertical direction above/below the disc, the closest N_* elements to \mathbf{x}_a likely will be in the disc – so searching only within H_a will fail to ‘see’ the vertical directions, thus coupling all feedback within the disc, despite the fact that the disc subtends a vanishingly small portion of the sky as seen from the star. Our solution to this is to use the same approach used in the hydrodynamic solver (in all mesh-free methods; Smoothed-Particle Hydrodynamics (SPH) and MFM/MFV): we include *both* elements with $|\mathbf{x}_{ba}| < H_a$ and $|\mathbf{x}_{ba}| < H_b$. That is, we additionally include any gas elements whose kernel encompasses the star. In the disc example, the closest ‘atmosphere elements’ above/below the disc necessarily have their own kernel radii, H_b , that overlap the disc, so this guarantees ‘covering’ by elements in the vertical direction. This is demonstrated in Fig. 1. The importance of including these elements is validated in our tests below, where we show that failure to include these neighbours artificially biases the feedback deposition.

We impose a maximum cut-off radius, r_{\max} , on the search, to prevent pathological situations for which there is no nearby gas so feedback would be deposited at unphysically large distances. Specifically, we impose $r_{\max} = 2 \text{ kpc}$. This corresponds to where the ram pressure of free-expanding ejecta falls below the thermal pressure in even low-density circum-galactic conditions ($T \sim 10^4 \text{ K}$ at $n \gtrsim 0.001 \text{ cm}^{-3}$). However, our results are not sensitive to this choice, because it affects a vanishingly small number of events.

2.2.3 Weighting the deposition: the correct ‘effective area’

Having identified interacting neighbours, b , we must deposit the injected quantities according to some weighting scheme. Each neighbour resolution element gets a weight $\tilde{\omega}_b$ that determines the fraction of the injected quantity it receives. Of course, this must be normalized to properly conserve quantities, so we first calculate an un-corrected weight, ω_b , and then assign

$$\tilde{\omega}_b \equiv \frac{\omega_b}{\sum_c \omega_c} \quad (1)$$

so that $\sum_b \tilde{\omega}_b = 1$, exactly.

Naively, a simple weight scheme might use $\omega_b = 1$, or $\omega_b = W(\mathbf{x}_{ba}, H_a)$. However, for quasi-Lagrangian schemes for which the different gas elements have approximately equal masses ($m_b \sim \text{constant}$), this is effectively mass-weighting the feedback deposition, which is not physical. In the example of the infinitely thin disc, because most of the neighbour elements lie within the disc, the disc-centred elements would again receive most of the feedback, despite the fact that they cover a vanishingly small portion of the sky from the source.

If the feedback is emitted statistically isotropically from the source \mathbf{x}_a , the correct solution is to integrate the injection into

⁵ In this paper we will use a cubic spline for W , but other choices have weak effects on our conclusions (because W will be re-normalized anyways in the assignment of ‘weights’ for feedback). We adopt $N_* = 64$ for reasons discussed below. The equation for $N_*(H_a)$ is non-linear, so it is solved iteratively in the neighbour search; see Springel & Hernquist (2002).

each solid angle and determine the total solid angle $\Delta\Omega_b$ subtended by a given gas resolution element, i.e. adopt $\omega_b = \Delta\Omega_b/4\pi$. This is shown in Fig. 1. Given a source at \mathbf{x}_a and neighbours at \mathbf{x}_b , we can construct a set of faces that enclose \mathbf{x}_a with some convex hull. Each face has a vector oriented area \mathbf{A}_b ; if the face is symmetric it subtends a solid angle on the sky as seen by \mathbf{x}_a of

$$\omega_b \equiv \frac{1}{2} \left(1 - \frac{1}{\sqrt{1 + (\mathbf{A}_b \cdot \hat{\mathbf{x}}_{ba}) / (\pi |\mathbf{x}_{ba}|^2)}} \right) \approx \frac{\Delta\Omega_b}{4\pi} \quad (2)$$

(This simply interpolates between $\sim A_b/4\pi r_b^2$ for $r_b^2 = |\mathbf{x}_{ba}|^2 \gg A_b \equiv |\mathbf{A}_b \cdot \hat{\mathbf{x}}_{ba}|$, and $1/2$ for $r_b^2 \ll A_b$).⁶

No unique convex hull exists. One solution, for example, would be to construct a Voronoi tessellation around \mathbf{x}_a , with both the star particle \mathbf{x}_a and the locations of all neighbours \mathbf{x}_b as mesh-generating points. However, we already have an internally consistent value of \mathbf{A}_b , namely, the definition $\mathbf{A}_b^{\text{hydro}}$ of the ‘effective faces’ used in the hydrodynamic equations (the faces that appear in the discretized Euler equations: e.g. $d\mathbf{U}_a/dt = -\sum_b \mathbf{F}_{ab}(\mathbf{U}) \cdot \mathbf{A}_b$, where \mathbf{U} is a conserved quantity and \mathbf{F} is its flux). For a Voronoi moving-mesh code (e.g. AREPO), this is the Voronoi tessellation. For SPH as implemented in GIZMO, this is $\mathbf{A}_b^{\text{hydro}} = [\bar{n}_a^{-2} \partial W(r_b, H_a)/\partial r_b + \bar{n}_b^{-2} \partial W(r_b, H_b)/\partial r_b] \hat{\mathbf{x}}_{ba}$. For MFM/MFV the expression is more complicated but is given in Eq. 18 in Hopkins (2015).⁷ We therefore adopt $\mathbf{A}_b = \mathbf{A}_b^{\text{hydro}}$ – the ‘effective face area’ that the neighbour gas elements would share with \mathbf{x}_a in the hydrodynamic equations if the source (star particle) were a gas element. Fig. 2 demonstrates that this is sufficient to ensure the coupling into each solid angle is statistically isotropic in the frame of the SN.

While we find that weighting by solid angle is important, at the level of accuracy here, the exact values of $\mathbf{A}_b^{\text{hydro}}$ given by SPH, MFM, or Voronoi formalisms differ negligibly, and we can use them interchangeably with no detectable effects on our results. This is not surprising: Hopkins (2015) showed that the Voronoi tessellation is simply the limit for a sharply peaked kernel of the MFM faces.

2.2.4 Dealing with vector fluxes (momentum deposition)

If we were only considering sources of scalar conserved quantities (e.g. mass m_{ej} or metals $m_{z, \text{ej}}$), we would be done. We simply define a numerical flux $\Delta m_b = \tilde{\omega}_b m_{\text{ej}}$ into each neighbour element (subtracting the same from our ‘source’ star particle), and we are

⁶ Equation (2) is exact for a face \mathbf{A}_b , which is rotationally symmetric about the axis $\hat{\mathbf{x}}_{ba}$; for asymmetric \mathbf{A}_b , evaluating $\Delta\Omega_b$ exactly requires an expensive numerical quadrature. If this is done exactly, equation (1) is unnecessary: $\sum_b \omega_b = 1$ is guaranteed. We have experimented with an exact numerical quadrature, but it is extremely expensive and has no measurable effect on our results compared to simply using equations (1) and (2) for all \mathbf{A}_b (equation (2) is usually accurate to < 1 per cent, and the most severe discrepancies do not exceed ~ 10 per cent, and these are normalized out by equation (1)).

⁷ In MFM/MFV, the effective face \mathbf{A}_b is given by

$$\mathbf{A}_b \equiv \bar{n}_a^{-1} \hat{\mathbf{q}}_b(\mathbf{x}_a) + \bar{n}_b^{-1} \hat{\mathbf{q}}_a(\mathbf{x}_b) \quad (3)$$

$$\hat{\mathbf{q}}_b(\mathbf{x}_a) \equiv \mathbf{E}_a^{-1} \cdot \mathbf{x}_{ba} W(\mathbf{x}_{ba}, H_a) \quad (4)$$

$$\mathbf{E}_a \equiv \sum_c (\mathbf{x}_{ca} \otimes \mathbf{x}_{ca}) W(\mathbf{x}_{ca}, H_a) \quad (5)$$

where ‘ \cdot ’ and ‘ \otimes ’ denote the inner and outer products, respectively.

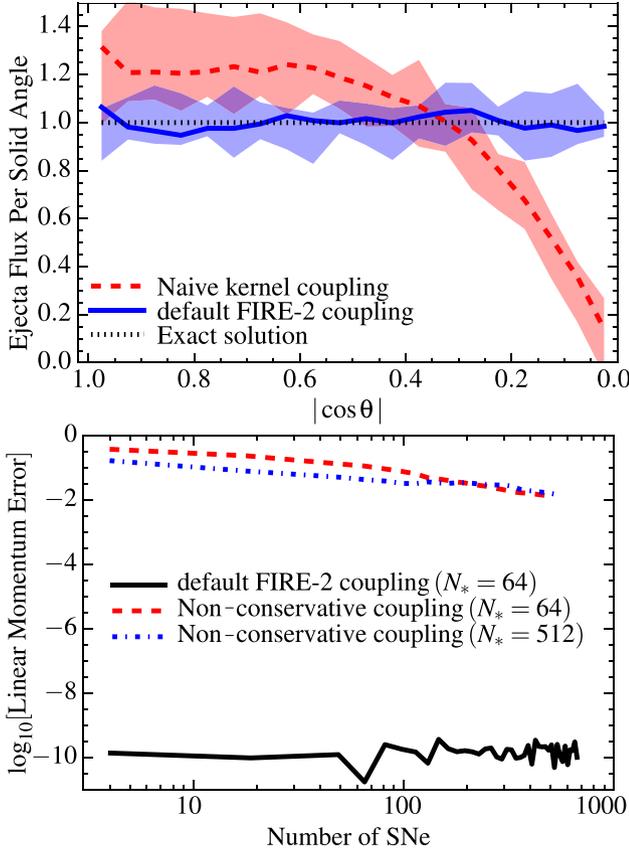


Figure 2. Numerical tests of the mechanical feedback coupling algorithm in Section 4. *Top:* Statistical Isotropy Test: We detonate a single SN in the centre of a thin disc, generated by randomly sampling a vertical Gaussian profile with gas particles, and we measure the resulting momentum/mass/metal flux deposited to neighbour gas elements as a function of polar angle θ ($\cos \theta = \pm 1$ is midplane, $\cos \theta = 0$ polar). We repeat 100 times and show the median (lines) and 95 per cent interval (shaded). The result *should* be statistically isotropic (‘Exact Solution’). Our ‘Default FIRE-2 Coupling’ method (Fig. 1) recovers this with noise owing to the finite number of particles coupled. The ‘Naive Kernel Coupling’ model only includes neighbours within the search radius H_* of the SN (not those where $H_* < |\mathbf{x}_* - \mathbf{x}_{\text{gas}}| < H_{\text{gas}}$; Section 2.2.2) and weights deposition by a simple kernel function (effectively mass-weighting) instead of the solid angle subtended by the element (Section 2.2.3). This naive coupling biases ejecta to couple into the mid-plane and suppresses coupling in the polar direction. *Bottom:* Momentum Conservation Test: We detonate SNe at random locations in the same system and measure the total fractional error in the linear momentum of the box (error $L_1(t) \equiv |\sum_a m_a \mathbf{v}_a(t)| / \sum p_{\text{coupled}}$, where $p_{\text{coupled}} = N(t) p_{\text{ej}}$ is the total magnitude of the momentum injected by all events). This is the net deviation from exact momentum conservation, relative to the total coupled. Our ‘Default FIRE-2 Coupling’ uses a tensor re-normalization scheme to keep these errors at machine accuracy (see Section 2.2.4). The ‘Non-Conservative Coupling’ scheme removes this re-normalization (but is otherwise identical); fractional conservation errors for a single event can then be order-unity! The fractional error declines with SNe number as $\sim N_{\text{SNe}}^{-1/2}$ because of cancellations; increasing the coupled neighbour number N_* reduces the errors but inefficiently.

guaranteed both machine-accurate conservation ($\sum_b \Delta m_b = m_{\text{ej}}$) and the correct spatial distribution of ejecta.

However, the situation is more complex for a vector flux, specifically here, momentum deposition. If the ejecta have some uniform radial velocity, $\mathbf{v}_{\text{ej}} = v_{\text{ej}} \hat{\mathbf{r}}$, away from the source, \mathbf{x}_a , then one might naively define

the corresponding momentum flux $\Delta \mathbf{p}_b = \tilde{\omega}_b m_{\text{ej}} v_{\text{ej}} \hat{\mathbf{x}}_{ba} = p_{\text{ej}} \tilde{\omega}_b \hat{\mathbf{x}}_{ba}$. However, then $\sum_b \Delta \mathbf{p}_b = p_{\text{ej}} \sum_b \tilde{\omega}_b \hat{\mathbf{x}}_{ba}$. But this is *not* guaranteed to vanish: the deposition can violate linear momentum conservation, if $\boldsymbol{\psi}_a \equiv \sum_b \tilde{\omega}_b \hat{\mathbf{x}}_{ba} \neq \mathbf{0}$. The correct $\boldsymbol{\psi}_a = \mathbf{0}$ is only guaranteed if (1) the coupled momentum $\Delta \mathbf{p}_b$ is the *exact* solution of the integral of $p_{\text{ej}} (4\pi |\mathbf{r}|^2)^{-1} \hat{\mathbf{r}} \cdot d\mathbf{A}_b(\theta, \phi)$ (where \mathbf{r} is the vector from \mathbf{x}_a to a location \mathbf{x} on the surface \mathbf{A}_b), and (2) the faces of the convex hull close exactly ($\sum_b \mathbf{A}_b = \mathbf{0}$). Even in a Cartesian grid (which trivially satisfies (2)), condition (1) can only be easily evaluated if we assume (incorrectly) that the feedback event occurs exactly at the centre or corner of a cell; in Voronoi meshes and mesh-free methods (1) is only possible to satisfy with an expensive numerical quadrature, and (2) is only satisfied up to some integration accuracy.

In practice, $\Delta \mathbf{p}_b = p_{\text{ej}} \tilde{\omega}_b \hat{\mathbf{x}}_{ba}$ is a good approximation to the integral in condition (1), and is again exact for faces symmetric about $\hat{\mathbf{x}}_{ba}$, and (2) is satisfied up to second-order integration errors in our MFM/MFV methods, so the dimensionless $|\boldsymbol{\psi}_a| \ll 1$ is small. However, we wish to ensure machine-accurate conservation, so we must impose a *tensor* re-normalization condition, not simply the scalar re-normalization in equation (1): we therefore define the six-dimensional vector weights $\hat{\mathbf{x}}_{ba}^\pm$:

$$\hat{\mathbf{x}}_{ba} \equiv \frac{\mathbf{x}_{ba}}{|\mathbf{x}_{ba}|} = \sum_{+, -} \hat{\mathbf{x}}_{ba}^\pm \quad (6)$$

$$(\hat{\mathbf{x}}_{ba}^+)^{\alpha} \equiv |\mathbf{x}_{ba}|^{-1} \text{MAX}(\mathbf{x}_{ba}^{\alpha}, 0) \Big|_{\alpha=x, y, z} \quad (7)$$

$$(\hat{\mathbf{x}}_{ba}^-)^{\alpha} \equiv |\mathbf{x}_{ba}|^{-1} \text{MIN}(\mathbf{x}_{ba}^{\alpha}, 0) \Big|_{\alpha=x, y, z} \quad (8)$$

i.e. the unit vector component in the plus (or minus) x, y, z directions (α refers to these components), for each neighbour. We can then define a vector weight $\bar{\mathbf{w}}_b$:

$$\bar{\mathbf{w}}_b \equiv \frac{\mathbf{w}_b}{\sum_c |\mathbf{w}_c|} \quad (9)$$

$$\mathbf{w}_b \equiv \omega_b \sum_{+, -} \sum_{\alpha} (\hat{\mathbf{x}}_{ba}^{\pm})^{\alpha} (f_{\pm}^{\alpha})_a \quad (10)$$

$$(f_{\pm}^{\alpha})_a \equiv \left\{ \frac{1}{2} \left[1 + \left(\frac{\sum_c \omega_c |\hat{\mathbf{x}}_{ca}^{\mp}|^{\alpha}}{\sum_c \omega_c |\hat{\mathbf{x}}_{ca}^{\pm}|^{\alpha}} \right)^2 \right] \right\}^{1/2} \quad (11)$$

This is evaluated in two passes over the neighbour list.⁸

It is straightforward to verify (and we show explicitly in tests below) that the approach above guarantees momentum conservation to machine accuracy. Ignoring these correction terms can (if the

⁸ The function (f_{\pm}) in equation (11) is derived by requiring $\mathbf{0} = \sum \Delta \mathbf{p}_b$. Component-wise, this becomes $0 = \sum (\Delta \mathbf{p}_b)^{\alpha} = p_{\text{ej}} / (\sum_c |\mathbf{w}_c|) [(f_+^{\alpha} \sum_b \omega_b (\hat{\mathbf{x}}_{ba}^+)^{\alpha} + f_-^{\alpha} \sum_b \omega_b (\hat{\mathbf{x}}_{ba}^-)^{\alpha})]$. Since p_{ej} and $\sum_c |\mathbf{w}_c|$ are positive-definite, the term in brackets must vanish ($f_+^{\alpha} \psi_+^{\alpha} = f_-^{\alpha} \psi_-^{\alpha}$, if we define $\psi_{\pm}^{\alpha} \equiv \sum_b \omega_b |\hat{\mathbf{x}}_{ba}^{\pm}|^{\alpha}$). But we also wish to minimize the effect of the correction factor f_{\pm} on the total momentum coupled (ensuring $f_{\pm} \approx 1$), so we minimize the least-squares penalty function $\Delta_{\psi}^2 = \|[(f_+^{\alpha} \psi_+^{\alpha})^2 + (f_-^{\alpha} \psi_-^{\alpha})^2] - [(\psi_+^{\alpha})^2 + (\psi_-^{\alpha})^2]\|$. The f_{\pm} in equation (11) is the unique function which simultaneously guarantees $\mathbf{0} = \sum \Delta \mathbf{p}_b$ (i.e. $f_+^{\alpha} \psi_+^{\alpha} = f_-^{\alpha} \psi_-^{\alpha}$) and $\Delta_{\psi}^2 = 0$. It is easy to see that $f_{\pm} \rightarrow 1$, as it should, if $\psi_+ = \psi_-$, i.e. when $\sum \Delta \mathbf{p}_b = \mathbf{0}$ without the need for an additional correction.

neighbours are ‘badly ordered,’ e.g. all lie the same direction), lead to order-unity errors in momentum conservation, and the fractional error $|\sum_b \Delta \mathbf{p}_b|/p_{ej} = |\boldsymbol{\psi}_a|$ depends only on the spatial distribution of neighbours in the kernel, not on the resolution.

Physically, we should think of the vector weights $\bar{\mathbf{w}}$ as accounting for asymmetries about the vector $\hat{\mathbf{x}}_{ab}$ in the faces \mathbf{A}_b . If the faces were all exactly symmetric (e.g. the neighbour elements were perfectly isotropically distributed), then the net momentum integrated into each face would indeed point exactly along $\hat{\mathbf{x}}_{ab}$. But, typically, they are not, so we must account for this in order to properly retain momentum conservation.

2.2.5 Assigning fluxes and including gas-star motion

Finally, we can assign fluxes:

$$\Delta m_b = |\bar{\mathbf{w}}_b| m_{ej} \quad (12)$$

$$\Delta m_{Z,b} = |\bar{\mathbf{w}}_b| m_{Z,ej} \quad (13)$$

$$\Delta E_b = |\bar{\mathbf{w}}_b| E_{ej} \quad (14)$$

$$\Delta \mathbf{p}_b = \bar{\mathbf{w}}_b p_{ej} \quad (15)$$

which the definitions above guarantee will *exactly* satisfy:

$$\sum \Delta m_b = m_{ej} \quad (16)$$

$$\sum \Delta m_{Z,b} = m_{Z,ej} \quad (17)$$

$$\sum \Delta E_b = E_{ej} \quad (18)$$

$$\sum |\Delta \mathbf{p}_b| = p_{ej} \quad (19)$$

$$\sum \Delta \mathbf{p}_b = \mathbf{0} \quad (20)$$

Our definitions also ensure that the fraction of ejecta entering a gas element is as close as possible (as much as allowed by the strict conservation conditions above) to the fraction of solid angle subtended by the element, as would be calculated self-consistently by the hydrodynamic method in the code, i.e.

$$|\bar{\mathbf{w}}_b| \approx \frac{\Delta \Omega_b^{\text{hydro}}}{4\pi} \quad (21)$$

Moreover, in the limit where equation (2) is exact (the faces \mathbf{A}_b are symmetric about $\hat{\mathbf{x}}_{ba}$), and they close exactly ($\sum_b \mathbf{A}_b = \mathbf{0}$; i.e. good element order), then $(f_{\pm}) = 1$ and $\sum_c |\mathbf{w}_c| = 1$, i.e. $\bar{\mathbf{w}}_b \rightarrow \omega_b \hat{\mathbf{x}}_{ba}$ and our naive estimate is both exact and conservative, and no normalization of the weights is necessary. In practice, as noted above, we find that the deviations (in the sum) from this perfectly ordered case are usually small (percents-level), but there are always pathological element configurations where they can be large, and maintaining good conservation requires the corrected terms above.

Implicitly, we have been working in the frame moving with the feedback ‘source’ ($\mathbf{x}_a = \mathbf{0}$, $\mathbf{v}_a \equiv d\mathbf{x}_a/dt = \mathbf{0}$), in which the source is statistically isotropic. However, in coupling the fluxes to surrounding gas elements, we also must account for the frame motion. Boosting back to the lab/simulation frame, the total ejecta velocity entering an element is of course $\Delta m_b^{-1} \Delta \mathbf{p}_b + \mathbf{v}_a$. This change of frame has no effect on the mass fluxes, but it does modify the momentum and energy fluxes: to be properly conservative, we must take:

$$\Delta m'_b \equiv \Delta m_b, \quad \Delta m'_{Z,b} \equiv \Delta m_{Z,b} \quad (22)$$

$$\Delta \mathbf{p}'_b \equiv \Delta \mathbf{p}_b + \Delta m_b \mathbf{v}_a \quad (23)$$

$$\Delta E'_b \equiv \Delta E_b + \frac{1}{2 \Delta m_b} (|\Delta \mathbf{p}'_b|^2 - |\Delta \mathbf{p}_b|^2) \quad (24)$$

where the prime (e.g. ‘ $\Delta m'_b$ ’) notation denotes the lab frame. Note that the extra momentum added to the neighbours ($\sum_b \Delta m_b \mathbf{v}_a = m_{ej} \mathbf{v}_a$) is exactly the momentum lost by the feedback source a , by virtue of its losing m_{ej} in mass.⁹

These fluxes are simply added to each neighbour in a fully conservative manner:

$$m_b^{\text{new}} = m_b + \Delta m'_b \quad (25)$$

$$(Z m_b)^{\text{new}} = Z^{\text{new}} m_b^{\text{new}} = (Z m_b) + \Delta m'_{Z,b} \quad (26)$$

$$\mathbf{p}_b^{\text{new}} = m_b^{\text{new}} \mathbf{v}_b^{\text{new}} = \mathbf{p}_b + \Delta \mathbf{p}'_b \quad (27)$$

$$E_b^{\text{new}} = E_{\text{kinetic}}^{\text{new}} + U_{\text{internal}}^{\text{new}} = E_b + \Delta E'_b \quad (28)$$

So the updated vector velocity \mathbf{v} of the element follows from its updated momentum and mass (and its metallicity follows from its updated metal mass and total mass); the energy E here is a *total* energy, so the updated internal energy U of the element follows from its updated total energy (E), kinetic energy (from \mathbf{v}), and mass (this is the usual procedure in finite-volume updates with conservative hydrodynamic schemes).

The terms accounting for the relative gas-star motion are necessary to ensure exact conservation. For SNe, they have essentially no effect. However, for slow stellar winds (e.g. AGB winds with $v_{\text{wind}} \sim 10 \text{ km s}^{-1}$), the relative star-gas velocity can be much larger than the wind velocity ($|\mathbf{v}_b - \mathbf{v}_a| \gg v_{\text{wind}}$), which means the shock energy and post-shock temperature of the winds colliding with the ISM is much higher than would be calculated ignoring these terms, which may significantly change their role as a feedback agent (Conroy, van Dokkum & Kravtsov 2015).

2.3 Sub-grid physics: unresolved Sedov–Taylor phases

A potential concern if naively applying the above prescription for SNe is that low-resolution simulations are unable to resolve the Sedov–Taylor (S–T) phase, during which the expanding shocked bubble is energy-conserving (the cooling time is long compared to the expansion time) and does $P dV$ work on the gas, converting energy into momentum, until it reaches some terminal radius where the residual thermal energy has been lost and the blastwave becomes a cold, momentum-conserving shell. This would, if properly resolved, modify the input momentum (Δp_b) and energy (ΔE_b) felt by the gas element b .

2.3.1 Motivation: individual SN remnant evolution

Idealized, high-resolution simulations (with element mass $m_i \ll M_{\odot}$) have shown that there is a robust radial terminal momentum, p_t , of the swept-up gas in the momentum-conserving phase, from a single explosion, given by

$$\frac{p_t}{M_{\odot} \text{ km s}^{-1}} \approx 4.8 \times 10^5 \left(\frac{E_{ej}}{10^{51} \text{ erg}} \right)^{\frac{13}{14}} \left(\frac{n}{\text{cm}^{-3}} \right)^{-\frac{1}{4}} f(Z)^{\frac{3}{2}} \quad (29)$$

⁹ The de-boosted energy equation, equation (24), assumes that the gas surrounding the star has initial gas-star relative velocities small compared to the ejecta velocity. A more general expression is presented in Appendix E.

$$f(Z) \approx \begin{cases} 2 & (Z/Z_\odot < 0.01) \\ (Z/Z_\odot)^{-0.14} & (0.01 \leq Z/Z_\odot) \end{cases} \quad (30)$$

where n and Z are the gas number density and metallicity, respectively, surrounding the explosion. The expression above is from Cioffi, McKee & Bertschinger (1988) (where we restrict $f(Z)$ to the minimum metallicity they consider), but similar expressions have been found in a wide range of other studies (for discussion, see Draine & Woods 1991; Slavin & Cox 1992; Thornton et al. 1998; Iffrig & Hennebelle 2015; Kim & Ostriker 2015; Li et al. 2015; Martizzi, Faucher-Giguère & Quataert 2015; Walch & Naab 2015; Haid et al. 2016; Hu et al. 2016; Gentry et al. 2017), with variations up to a factor of ~ 2 , which we explore below.¹⁰

We validate this expression in simulations below. But physically, this follows from simple cooling physics: taking $E_0 \sim E_{51} 10^{51}$ erg and converting an order-unity fraction to thermal energy within a swept-up mass M gives a temperature $T \sim 10^6 \text{ K} (M/3000 E_{51} M_\odot)^{-1}$, so when $M \gtrsim M_{\text{cool}} \sim 3000 E_{51} M_\odot$, T drops to $< 10^6$ K and the gas moves into the peak of the cooling curve where radiative losses are efficient (Rees & Ostriker 1977). While energy-conserving, the shell momentum scales as $p \sim M v \sim \sqrt{2 M E_0}$, so the terminal momentum is $p_t \sim \sqrt{2 M_{\text{cool}} E_0} \sim 5 \times 10^5 E_{51} M_\odot \text{ km s}^{-1}$.

One important caveat: these scalings (and our implementations below) are developed for single ‘events’ (e.g. explosions), as opposed to continuous events (e.g. approximately constant rates of stellar mass-loss over long time periods). ‘Continuous’ feedback can, in principle, produce different scalings (see e.g. the discussion in Weaver et al. 1977; McKee, van Buren & Lazareff 1984; Freyer, Hensler & Yorke 2006; Faucher-Giguère & Quataert 2012; Gentry et al. 2017). It is still the case that winds must either expand in some energy-conserving fashion (doing PdV work) or cool, and so a scaling qualitatively like those here must still apply – however, details of when cooling occurs (which set the exact ‘terminal momentum’), in continuous cases, are much less robust to the environment, density profile, ability of the surrounding medium to confine the wind, and temperature range of the reverse shock (see references above and e.g. Harper-Clark & Murray 2009; Rosen et al. 2014). Moreover, there is growing evidence that stellar mass loss is highly ‘bursty’ or ‘clumpy’ with most of the kinetic luminosity associated with smaller time-or-spatial scale ejection events and/or clumps (Fong, Meixner & Shah 2003; Young et al. 2003; Repolust, Puls & Herrero 2004; Ziurys et al. 2007; Agúndez, Cernicharo & Guélin 2010; Cox et al. 2012). In those cases, treating each ‘event’ with the scalings above is appropriate. Because the kinetic luminosity in stellar mass-loss is an order-of-magnitude lower than that associated with SNe, even relatively large changes in our treatment of stellar mass-loss (e.g. assuming the ejecta are entirely radiative, so the terminal momentum is the initial momentum) have little effect on galaxy scales (if SNe are also present). We therefore, for simplicity, apply the same scalings to all mechanical feedback. But this certainly merits more detailed study in future work.

2.3.2 Numerical treatment

To account for potentially unresolved energy-conserving phases, we first calculate the momentum that *would* be coupled to the

¹⁰ We adopt the specific expression from Cioffi et al. (1988), as opposed to that from more recent work, for consistency with the previous FIRE-1 simulations.

gas element, assuming the blastwave were energy conserving throughout that *single* element, which is simply $\Delta p'_b \rightarrow \Delta p'_b (1 + m_b/\Delta m_b)^{1/2}$. We then compare this to the terminal momentum p_t (assume each neighbour b sees the appropriate ‘share’ of the terminal momentum according to its share of the ejecta mass), and assign the actual coupled momentum to be the smaller of the two.¹¹ In other words:

$$\Delta \mathbf{p}_b^{\text{new}} \equiv \text{MIN} [\Delta \mathbf{p}_b^{\text{energy-conserving}}, \Delta \mathbf{p}_b^{\text{terminal}}] \quad (31)$$

$$= \Delta \mathbf{p}_b \text{ MIN} \left[\sqrt{1 + \frac{m_b}{\Delta m_b}}, \frac{p_t}{p_{\text{ej}}} \right] \quad (32)$$

(where recall $p_{\text{ej}} = \sqrt{2 m_{\text{ej}} E_{\text{ej}}}$). Because the coupled ΔE is the *total* energy and is not changed, this remains manifestly energy-conserving (the energy that implicitly goes into the PdV work increasing \mathbf{p} is automatically moved from thermal to kinetic energy). This is done in the rest frame (before boosting back to the lab frame).

Consider the two limits: (1) when $p_t/p_{\text{ej}} < (1 + m_b/\Delta m_b)^{1/2}$, the physical statement is that the cooling radius is un-resolved. Because $\Delta \mathbf{p}_b = p_{\text{ej}} \bar{\mathbf{w}}_b$, multiplying by p_t/p_{ej} simply replaces the ‘at explosion’ initial p_{ej} with the terminal p_t – in other words, exactly the momentum that the element b *should* see, if we had properly resolved the S–T phase between \mathbf{x}_a and \mathbf{x}_b . On the other hand: (2) when $p_t/p_{\text{ej}} > (1 + m_b/\Delta m_b)^{1/2}$, the cooling radius is resolved; so we simply assume the blastwave is energy-conserving at the location of coupling. Because, by definition, the coupled momentum will be less than p_t , the actual momentum coupling is, in this limit, largely irrelevant – we essentially couple thermal energy and rely on the hydrodynamic code to actually *solve* for the correct PdV work as the blastwave expands.¹²

Strictly speaking, the expressions in equation (31)–(32) are expected if the relative gas-star velocities ($\mathbf{v}_b - \mathbf{v}_a$) surrounding the explosion are either (a) small or (b) uniform. In Appendix E we present the more exact scalings, as well as the appropriate boost/deboost corrections for momentum and energy, accounting for arbitrary gas-star motions.

In Section 4, we show Figs 5 and 6 that this algorithm reproduces the exact result of much higher-resolution, converged simulations of SN blastwaves even when the coupling is applied in lower-resolution simulations – just as intended.

To be fully consistent, we also need to account for the loss of thermal energy (via radiation) in limit (1), when the cooling radius is un-resolved. The effective cooling radius R_{cool} is exactly determined by the expression for p_t , because at the end of the energy-conserving phase ($R_{\text{shock}} = R_{\text{cool}}$), $(1/2)(m_{\text{ej}} + m_{\text{swept}}[R_{\text{cool}}])v_f^2 = (1/2)m_{\text{ej}}v_{\text{ej}}^2$ and $p_t = m_{\text{swept}}[R_{\text{cool}}]v_f$, giving $R_{\text{cool}} \approx 28.4 \text{ pc} (n/\text{cm}^{-3})^{-3/7} (E_{\text{ej}}/10^{51} \text{ erg})^{2/7} f(Z)$ for p_t in equation (29). Following Thornton et al. (1998), the post-shock thermal energy outside R_{cool} decays $\propto (r/R_{\text{cool}})^{-6.5}$, so we first calculate

¹¹ Kimm & Cen (2014) introduce a smooth interpolation function rather than a simple threshold in equation (31); we have experimented with variations of this and find no detectable effects.

¹² Note that we do not need to make any distinction between the free-expansion radius, post-shock (reverse shock) radius, etc., in our formalism, because the fully conservative coupling – which *exactly* solves the elastic two-body gas collision between ejecta and gas resolution element – automatically assigns the correct values in either limit. For example, if $m_b \ll m_{\text{ej}}$, our coupling will automatically determine that element b should simply be ‘swept up’ with velocity $\mathbf{v}_b \approx \mathbf{v}_{\text{ej}}$ (free-expansion); if $m_b \gg m_{\text{ej}}$, the gas is automatically assigned the appropriate post-shock temperature.

the post-shock thermal energy of element b that would be added by the ejecta, $\Delta U_b = E_{\text{ej}} - \Delta \text{KE}$ (where ΔKE is the change in kinetic energy, i.e. based on the coupled energy and momentum) in our usual fully conservative manner, then if $r_b \equiv |\mathbf{x}_{ba}| > R_{\text{cool}}$ we reduce this accordingly: $\Delta U_b \rightarrow \Delta U_b (|\mathbf{x}_{ba}|/R_{\text{cool}})^{-6.5}$. In practice, because *by definition* this correction to ΔU_b only appears outside the cooling radius (where the post-shock cooling time is short compared to the expansion time), we find that the inclusion/exclusion of this correction term has no detectable effects on our simulations (see Section 4.2); if we do not include it, the thermal energy is simply radiated away in the next time-step, as it should be. Still, we include the term for consistency.

We can (and do, for the sake of consistency) apply the full treatment described above to continuous stellar mass loss as well as SNe, using the differential E_{ej} (and enforcing $p_t \geq p_{\text{ej}}$), but the ‘multiplier’ is small because the winds are injected continuously so the E_{ej} in a single time-step is small.

Finally, the calculations of p_t and equations (31)–(32) are done independently for each neighbour b . In effect, we are considering each solid angle face $\Delta \Omega_b$ to be an independent ‘cone’ with its own density and metallicity, in which an independent energy-conserving solution is considered. Haid et al. (2016) have performed a detailed simulation study of SNe in inhomogeneous environments and showed explicitly that almost all of the (already weak) effect of different inhomogeneous initial conditions (in e.g. turbulent, clumpy, multiphase media) in their study and others is properly captured by considering each element surrounding the SN as an independent cone, which is assigned its own density-dependent solution according to the single homogeneous scaling above. In fact, once the density and metallicity dependence are accounted for as we do, residual systematic uncertainties in equation (29) are remarkably small (~ 10 – 50 per cent) – much smaller than uncertainties in the SNe rate itself!

2.3.3 Implied resolution requirements

Equation (32) demonstrates that with sufficiently small element mass (m_b below some critical m_{crit}), the cooling radius is resolved – i.e. we are in limit (2) above: $(1 + m_b/\Delta m_b)^{1/2} < p_t/p_{\text{ej}}$. This corresponds to $m_b \ll m_{\text{crit}} = 2500 M_{\odot} |\bar{\mathbf{w}}| E_{51}^{6/7} (n/\text{cm}^{-3})^{-2/7} f(Z)^3$. Because the kernel function is strongly peaked, most of ejecta energy/momentum/mass is deposited in the nearest few elements, so $|\bar{\mathbf{w}}| \sim 1/\text{few}$; so $m_{\text{crit}} \sim 500 M_{\odot} (n/\text{cm}^{-3})^{-2/7}$. This is a *mass* resolution criterion: as noted above, the cooling *radius* depends on density, $R_{\text{cool}} \sim n^{-1/3}$, such that an almost invariant mass $M_{\text{cool}} \equiv m_{\text{swept}}(R_{\text{cool}}) \sim (\text{a few}) m_{\text{crit}}$ is enclosed inside R_{cool} .

Similar results are found in Hu et al. (2016) (their Appendix B): they show, for example, that with $\sim 100 M_{\odot}$ resolution, the blastwave momentum is almost perfectly recovered (within < 10 per cent of simulations with element/particle masses $\sim 0.01 M_{\odot}$). Even higher-order effects such as the blastwave mass-loading, velocity structure, shell position, etc., are recovered almost perfectly once the shell has propagated into the momentum-conserving phase.

In our cosmological simulations of isolated dwarf galaxies, we begin to satisfy $m_b < m_{\text{crit}}$. However, in MW-mass simulations, this remains unattainable for now. Therefore, ignoring the correction for an unresolved S–T phase in massive galaxies can significantly under-estimate the effects of feedback. We consider explicit resolution tests below which validate these approximate scalings.

3 NUMERICAL TESTS: THE COUPLING ALGORITHM

We now consider detailed numerical tests of the SNe coupling scheme. Specifically, we first consider tests of the pure algorithm used to deposit feedback from Sections 2.2.2–2.2.5, independent of the feedback physics (energy, momentum, rates, etc.).

3.1 Validation: ensuring correct coupling isotropy, weights, and exact conservation

Fig. 2 considers two simple validation tests (for conservation and statistical isotropy) of our algorithm in a pure hydrodynamic test problem. We initialize a periodic box of arbitrarily large size centred on $\mathbf{x} = \mathbf{0}$, filled with particles of equal mass, m , meant to represent a patch of a vertically stratified disc. There is no gravity and the gas is forced to obey an exactly isothermal equation of state with vanishingly small pressure. The particles are laid down randomly with a uniform probability distribution in the x and y dimensions and probability dp along the z dimension such that $dp \propto m^{-1} \langle \rho(z) \rangle dx dy dz$, where $\langle \rho(z) \rangle = \exp(-z^2/2h^2)$. Initial velocities are zero. We define m and code units such that h is equal to the mean inter-particle separation in the midplane. The desired density distribution is therefore obeyed on average but with a noisy particle distribution, as in a real simulation.

The top panel of Fig. 2 shows the results after a single SN detonated at the centre of the box, using the standard FIRE-2 coupling scheme to deposit its ejecta. Because of the enforced equation-of-state, the coupled thermal energy is instantly dissipated – all that is retained is momentum, mass, and metals. We measure the amount deposited in each direction – each unit solid angle ‘as seen by’ the SN. By construction, our algorithm is *supposed* to couple the ejecta statistically isotropically. But because the ejecta must be deposited discretely in a finite number of neighbours, in any single explosion the deposition is noisy: it occurs only along the directions where there are neighbours. We therefore re-generate the box and repeat 100 times, and plot the resulting mean distribution and scatter. We confirm that our default algorithm correctly deposits ejecta statistically isotropically, on average. However, if we instead consider a simpler algorithm where the search for neighbours to couple the SN (Section 2.2.2) is done only using particles within a nearest-neighbour radius H_a of the SN (excluding particles outside H_a but for which the SN is inside *their* nearest-neighbour radius H_b), or if we weight the deposition ‘per neighbour’ by a simple kernel weight (Section 2.2.3), in this case the cubic spline kernel ($\omega_b = W(\mathbf{x}_{ba}, H_a)$); then we obtain a biased ejecta distribution. The bias is as expected: most of the ejecta go into the disc midplane direction, because on average there are more particles in this direction, and they are closer, as opposed to in the vertical direction, where the density decreases. In a real simulation, this is a serious concern: momentum and energy would be preferentially coupled in the plane of the galaxy disc, rather than ‘venting’ in the vertical direction as they should, simply because more particles are in the disc!

In the bottom panel of Fig. 2, we repeat our set-up, but now we repeatedly detonate SNe throughout the box at fixed time intervals, each in a random position. After each SN we measure the total momentum of all gas elements, $|\mathbf{p}| \equiv |\sum m_a \mathbf{v}_a|$, and define the dimensionless, fractional linear momentum error as the ratio of this to the total ejecta momentum that has been injected, $L_1 = |\mathbf{p}|/\sum p_{\text{ej}}$. Linear momentum conservation demands $\mathbf{p} = \mathbf{0}$. In our standard FIRE-2 algorithm, we confirm momentum is conserved to machine accuracy. However, re-running without the tensor renormalization

in Section 2.2.4, we see quite large errors, with $L_1 \sim 0.1 - 1$ for a single SN, decreasing slowly with the number of SNe in the box only because the errors add incoherently (so L_1 gradually decreases $\propto N_{\text{SNe}}^{-1/2}$ as a Poisson process). We can decrease L_1 in the non-conservative algorithm by increasing the number of gas neighbours used for the SN deposition, but this is inefficient and reduces the spatial resolution.

3.2 Tests in fire simulations: effects of algorithmic SNE coupling

In Figs 3 and 4,¹³ we examine how the algorithmic choices discussed above alter the formation history of galaxies in cosmological simulations. We compare:

(i) *Default*: Our default FIRE-2 coupling. This manifestly conserves mass, energy, and momentum; correctly deposits the ejecta in an unbiased (statistically isotropic) manner; and accounts for the Lagrangian distribution of particles in all directions.

(ii) *Non-conservative*: Coupling that neglects the tensor correction from Section 2.2.4, which Fig. 2 showed was necessary to maintain exact momentum conservation. We stress that the scalar mass and energy from SNe are still manifestly conserved here; only vector momentum is imperfectly added.

(iii) *FIRE-1 coupling*: Our older scheme from FIRE-1, which used the non-conservative formulation, conducted the SNe neighbour search only ‘one-directionally’ (ignoring neighbours with at distances $>H_a$), as defined in Section 2.2.2, and scaled the deposition ‘weights’ ω_b defined in Section 2.2.3 with volume ($\omega_b \propto m_b/\rho_b$; the ‘SPH-like’ weighting; see Price 2012), as opposed to solid angle. Fig. 2 shows this leads to unphysically anisotropic momentum deposition.

Fig. 3 (left side) shows that the detailed choice of coupling algorithm has essentially no effect in dwarf galaxies, because of their stochastic, bursty star formation and outflows and irregular/spheroidal morphologies. That is, a ‘galaxy wide explosion’ remains such regardless of exactly how individual SN are deposited. Indeed, we find that this independence from the coupling algorithm persists at any resolution that we test. We do not show visual morphologies of dwarf galaxies in Fig. 4, because they are essentially the same in all cases (see also Paper I). For MW-mass haloes, we find only a weak dependence of galaxy properties in Fig. 3 on the SNe algorithm (see Appendix A for demonstration of this at various resolution levels). The non-conservative implementations generally show a lower central stellar density at <1 kpc, owing to burstier intermediate-redshift star formation, because the momentum conservation errors allow more ‘kicking out’ of material in dense regions, as discussed further below.

At low and intermediate resolutions, the MW-mass simulations all exhibit ‘normal’ discy visual morphologies, without strong dependence on the SNe algorithm. However, at high resolution the ‘non-conservative’ run essentially destroys its disc! This is in striking contrast to the ‘default’ run, where the disc continues to become thinner and more extended at higher resolution (a trend seen in several MW-mass haloes studied in Paper I). Note that the formation

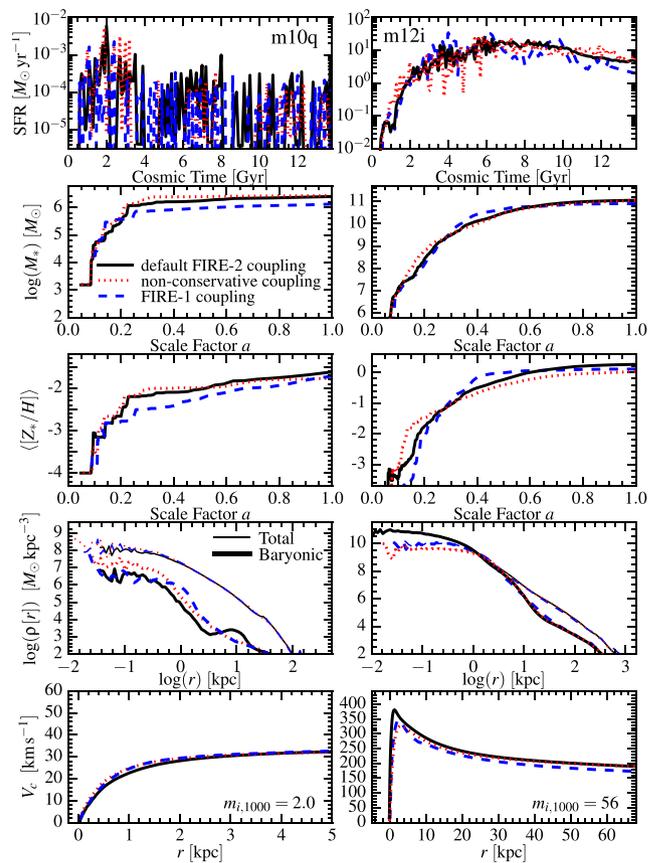


Figure 3. Comparison of different algorithms for coupling the same mechanical feedback to gas, in zoom-in cosmological simulations. *Top*: Star formation history (averaged in 100 Myr intervals) of the primary $z = 0$ galaxy. *Second*: Total stellar mass in box (dominated by primary galaxy) versus scale factor $a = 1/(1+z)$. *Middle*: Stellar mass-weighted average metallicity versus scale factor. *Third*: Baryonic (*thick*) and total (*thin*) mass density profiles (averaged in spherical shells) as a function of radius around the primary galaxy at $z = 0$. *Bottom*: Rotation curves (circular velocity V_c versus radius) in the primary galaxy. In each panel we compare three models: (1) *Default*: Our most accurate, fully conservative algorithm (Section 2.2). Ejecta are deposited in neighbouring gas elements weighted by the solid angle subtended ‘as seen by’ the SN, guaranteeing mass, energy, and statistical isotropy. (2) *Non-conservative coupling*: ‘Naive’ coupling with momentum deposited along the vector connecting the centre-of-mass from star to neighbour gas element, without tensor renormalization (Section 2.2.4). While simple, this algorithm can violate linear momentum conservation (imparting *net* momentum to the gas). (3) *FIRE-1 coupling*: SNe algorithm from FIRE-1: it used the non-conservative method (Section 2.2.4), a less-accurate SPH approximation of the solid angle (essentially a volume-weighting, $\omega_b \propto m_b/\rho_b$; Section 2.2.3), and only used the nearest neighbours for each SN instead of the bi-directional search needed to ensure statistical isotropy (Section 2.2.2). Despite these algorithmic differences, results at this resolution for dwarfs and massive galaxies are similar.

history and mass profile are not dramatically different in the two runs, so what has ‘gone wrong’ in the non-conservative case? The problem is, as noted in Section 2.2.4, the momentum conservation error in the non-conservative algorithm is zeroth-order – it depends only on the spatial distribution of and number of neighbour gas elements within the kernel, not on the absolute mass/spatial scale of that kernel. Because we keep the number of neighbours seen by the SN fixed with changing mass resolution, this means that the fractional errors (i.e. the net linear momentum error deposited per SN) does not

¹³ Mock images in Fig. 4 are computed as *ugr* composites, ray-tracing from each star after using its age and metallicity to determine the intrinsic spectrum from Leitherer et al. (1999) and accounting for line-of-sight dust extinction with an MW-like extinction curve and dust-to-metals ratio following Hopkins et al. (2005).

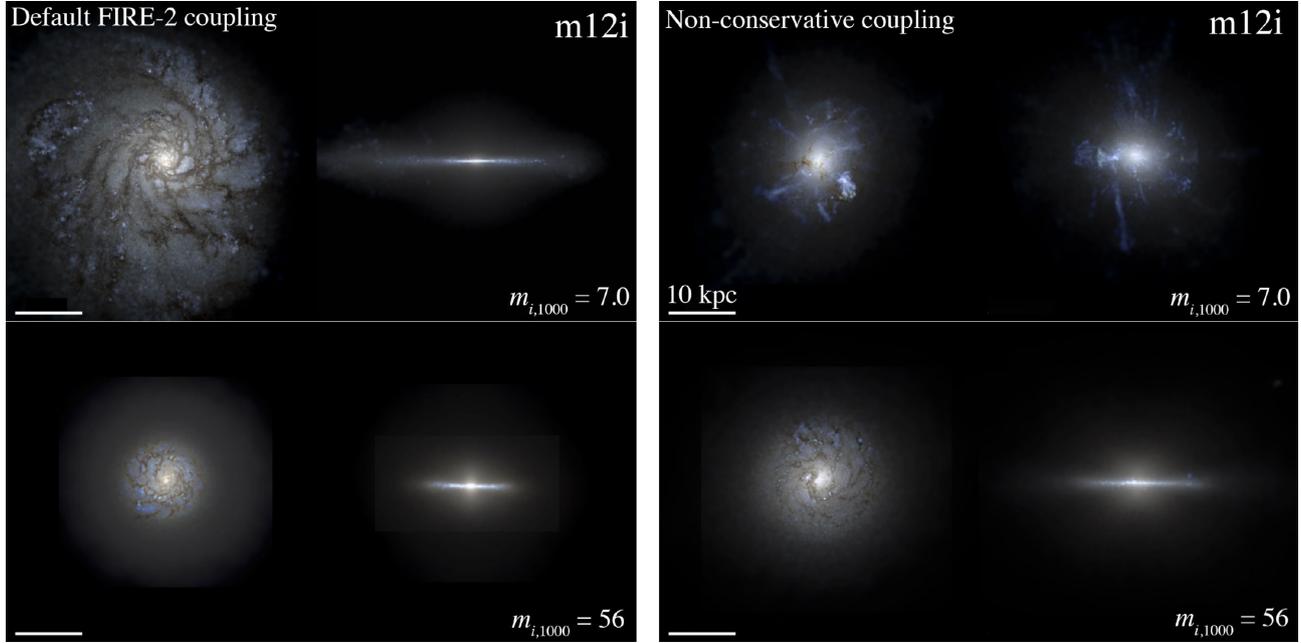


Figure 4. Mock images in HST bands of our **m12i** run at $z = 0$ as a function of mass resolution but for the ‘default’ (left) and ‘non-conservative’ (right) SNe coupling algorithm from Fig. 3. In the non-conservative algorithm, the error in *momentum* is independent of resolution, so the (worst-case) error in *velocity* scales as $|\Delta \mathbf{v}_{\text{err}}| \sim 100 \text{ km s}^{-1} (7000/m_{i,1000})$. At low resolution ($m_{i,1000} = 56$ shown, or $m_{i,1000} = 450$ in Fig. 3) the maximum velocity error is $\sim 10 \text{ km s}^{-1}$ (comparable to random motions in the ISM), so the galaxy is well-behaved. However, at high resolution $|\Delta \mathbf{v}_{\text{err}}| \sim 100 \text{ km s}^{-1}$, so the errors blow up the galaxy! Specifically, some young star-forming regions accumulate a coherent error in momentum conservation and are launched out of the galaxy, generating the visible streams and preventing thin-disc formation. In our default coupling these errors vanish, so the galaxy is well behaved: at higher resolution the disc is even thinner and more extended.

converge away. Meanwhile, the individual gas element masses get smaller at high resolution – so the net linear velocity ‘kick’ becomes larger. The ‘worst-case’ error for a single SNe would be an order-unity fractional violation of momentum conservation, implying a kick $|\Delta \mathbf{v}_{\text{err}}| \sim p_i/m_a \sim 100 \text{ km s}^{-1} (7000/m_{i,1000})$; at low and intermediate resolution even this worst-case gives $|\Delta \mathbf{v}_{\text{err}}| \lesssim 10 \text{ km s}^{-1}$ (comparable to the thin-disc velocity dispersion) so this is not a serious issue. But at our highest resolution, the non-conservative ‘worst-case scenario’ occurs where in some star-forming regions, net momentum is coherently deposited all in one direction owing to a pathological local particle distribution: the cloud then coherently ‘self-ejects’ or ‘bootstraps’ itself out of the disc. The thin disc is destroyed in the process, and the most extreme examples of this are visibly evident as ‘streaks’ of stars from self-ejected clumps flying out of the galaxy centre!

We also re-ran a ‘non-conservative’ simulation of **m12i** at high resolution ($m_{i,1000} = 7.0$) with a crude ‘cap’ or upper limit arbitrarily imposed for the fraction of the momentum allowed to couple to any one particle, and to the maximum velocity change per event (of 50 km s^{-1}). This is presented in Appendix B. In that case, the system does indeed form a thin, extended disc, similar to our default coupling. This confirms that the ‘self-destruction’ of the disc is driven by rare cases with large momentum errors, rather than small errors in ‘typical’ cases.

As noted above, our older FIRE-1 algorithm used the ‘non-conservative’ formulation. The MW-mass simulations published with that algorithm were all lower-resolution, where $|\Delta \mathbf{v}_{\text{err}}| \lesssim 10 \text{ km s}^{-1}$, so these errors were not obvious (at dwarf masses, the lower metallicities and densities meant the cooling radii of blast-waves were explicitly resolved, so as Fig. 3 shows, the effects were even smaller, and their irregular morphologies meant perturbations

to thin discs were not possible). However, running that algorithm in MW-mass haloes at higher resolution led to similar errors as shown in Fig. 4. This, in fact, motivated the development of the new FIRE-2 algorithm.

We have confirmed that all of the conclusions above are not unique to the two haloes above: we have re-ran haloes **m09** ($\sim 10^9 M_\odot$), **m10v** ($\sim 10^{10} M_\odot$), **m11q**, **m11v** ($\sim 10^{11} M_\odot$), **m12f** and **m12m** ($\sim 10^{12} M_\odot$) from Paper I with ‘Default’ and ‘Non-conservative’ implementations. All haloes $\sim 10^9$ – $10^{11} M_\odot$ show the same *lack* of effect from the coupling scheme as our **m10q** run here; the $\sim 10^{12} M_\odot$ haloes all show the same systematic dependencies as our **m12i** run.

In Appendix C we briefly discuss algorithms that ensure manifest momentum conservation by simply coupling a pre-determined momentum in the Cartesian $\pm x$, $\pm y$, $\pm z$ directions (independent of the local mesh or particle geometry). We do not adopt such a method because (a) it ignores the physically correct geometry of the mesh in irregular-mesh or mesh-free methods and (b) it imprints preferred directions on to the simulation, which forces discs to align with the simulation coordinate axes, introducing spurious numerical torques that can significantly reduce disc angular momentum (as often seen in grid-based codes).

4 NUMERICAL TESTS: SUB-GRID PHYSICS AND THE NEED TO ACCOUNT FOR THERMAL AND KINETIC ENERGY

Having tested the algorithmic aspect of SNe coupling above, we now consider tests of the physical scalings in the feedback coupling, specifically how it assigns momentum versus thermal energy as described in Section 2.3.

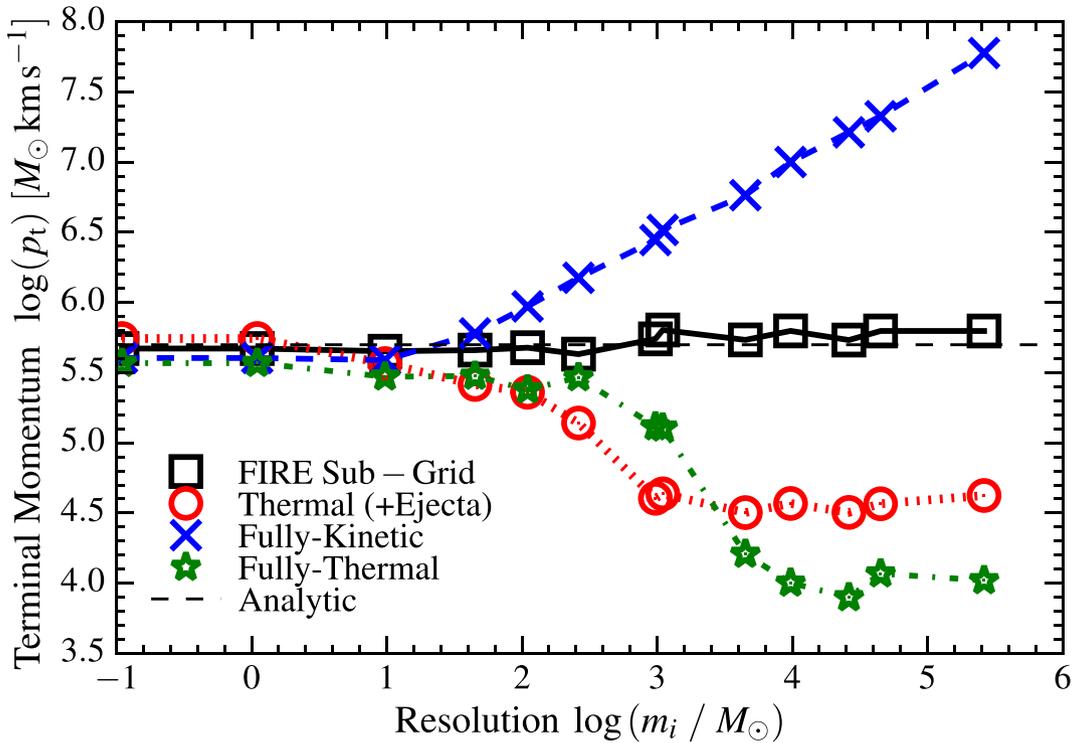


Figure 5. Convergence of ‘sub-grid’ implementations of SNe feedback (Section 2.3). We explode a single SN (ejecta mass = $10.4 M_{\odot}$, kinetic energy = 10^{51} erg, yields in Paper I Appendix A) in a box of uniform density and metallicity ($n = 1 \text{ cm}^{-3}$, $Z = Z_{\odot}$), with the same cooling physics as our standard FIRE-2 simulations, using varying gas particle mass resolution, across $m_i = 0.1 - 10^6 M_{\odot}$, as shown on the x-axis. Once the blastwave is well into the momentum-conserving stage, we measure the terminal gas momentum p_t . We compare: (1) *Analytic*: the analytic result for p_t (equation (29)). (2) *FIRE sub-grid*: Our default implementation. This couples SN mass, metals, energy, and momentum, in a *manifestly* conservative, statistically isotropic manner, with the coupled momentum following equation (32). (3) *Thermal (+ejecta)*: The coupling algorithm is the same, but the coupled momentum is *only* the original ejecta momentum (no PdV work is accounted for) and the energy is always the initial 10^{51} erg at coupling (no un-resolved radiation assumed). At low resolution this means the energy coupled is almost entirely thermal. (4) *Fully kinetic*: We couple 100 per cent of the 10^{51} erg as kinetic energy (all in momentum), regardless of resolution. (5) *Thermal only*: We couple 100 per cent of the 10^{51} erg as thermal energy (none in momentum). At $m_i \ll 10 M_{\odot}$, the ejecta free-expansion phase is resolved and all methods produce an identical, well-resolved S–T phase and terminal momentum in excellent agreement with the analytic value. At $m_i \gtrsim 100 M_{\odot}$, the cooling radius becomes un-resolved. At this low resolution, ‘Thermal (+Ejecta)’ (‘Fully Thermal’) underestimates the terminal momentum by a factor of ~ 15 (~ 60), because the PdV work done in the energy-conserving phase is missed ‘Thermal (+Ejecta)’ simply returns the original ejecta momentum; ‘Fully Thermal’ produces a small residual). At low resolution ‘Fully Kinetic’ overestimates the terminal momentum by a factor of $\sim 100 (m_i/10^5 M_{\odot})^{1/2}$, because it assumes that PdV work continues well after the remnant should cool. Our FIRE sub-grid model, by construction, agrees within ~ 10 per cent with the exact/high-resolution solution, *independent of resolution*.

4.1 Validation: ensuring ‘sub-grid’ scalings reproduce high-resolution simulations in resolution-independent fashion

In Figs 5 and 6, we consider an idealized test problem that validates the sub-grid SNe treatment used in FIRE. We initialize a periodic box of arbitrarily large size with uniform density $n = 1 \text{ cm}^{-3}$ and metallicity $Z = Z_{\odot}$, with constant gas particle mass m_i (so the inter-particle separation is given by $\rho = m_i/h_i^3$, i.e. $h_i \sim 16 \text{ pc} (m_i/100 M_{\odot})^{1/3}$), and with our full FIRE-2 cooling physics (with the $z = 0$ meta-galactic background) and hydrodynamics, but no self-gravity. We then detonate a single SN explosion at the centre of the box, using *exactly* our default FIRE-2 algorithm (same SN energy = 10^{51} erg, ejecta mass = $10.4 M_{\odot}$, metal content, ejecta momentum, and algorithmic coupling scheme from Fig. 1 and Sections 2.2.2–2.2.5). We also test several additional schemes for how to deal with the thermal versus kinetic (energy/momentum) component of the SN.

(i) *FIRE sub-grid*: This is our default FIRE-2 treatment from Section 2.3 (equation (31)), where we account for the PdV work done by the expanding blastwave out to the minimum of either the

coupling radius or cooling radius (where the resulting momentum reaches the terminal momentum p_t in equation (29), and we assume any remaining thermal energy is dissipated outside the cooling radius). The coupled momentum ranges, therefore, between $p_{\text{ejecta}} \leq p_{\text{coupled}} \leq p_{\text{terminal}}$ and *total* (kinetic+thermal) energy coupled ranges from $0 < E_{\text{coupled}} \leq E_{\text{ejecta}} = 10^{51}$ erg, according to the total mass enclosed within the single gas particle (the smallest possible ‘coupling radius’). Recall, at small particle mass, this becomes identical to coupling exactly the SN ejecta energy and momentum. At large particle mass, this reduces to coupling the terminal momentum and radiating (instantly) all residual (post-shock) thermal energy.

(ii) *Thermal (+ejecta)*: This couples only the ejecta momentum ($p_{\text{coupled}} = p_{\text{ejecta}} < p_{\text{terminal}}$): any additional energy is coupled as thermal energy (not radiated away in the coupling step; $E_{\text{coupled}} = E_{\text{ejecta}}$). This ignores any accounting for whether the coupling is inside/outside the cooling radius, or any PdV work done by the un-resolved blastwave expansion. It is equivalent to dropping the terms from Section 2.3 completely. A method like this was used in some previous work with non-cosmological simulations (Hopkins et al. 2012).

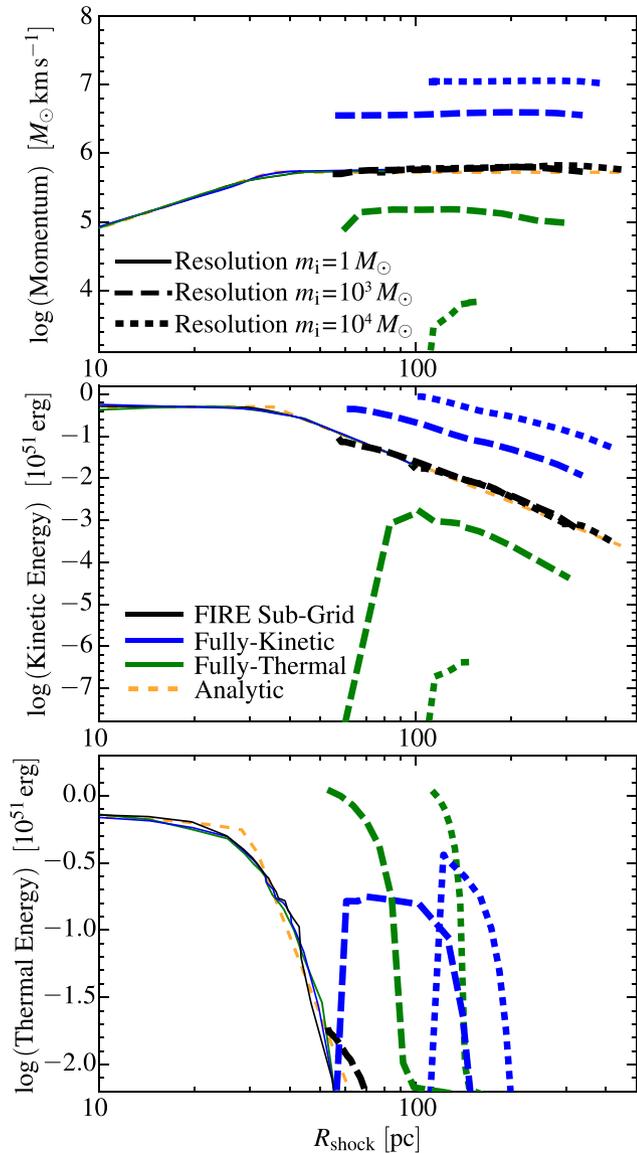


Figure 6. Time evolution of blastwave momentum (*top*), kinetic energy (*middle*), and thermal energy (*bottom*), as a function of shock position $R_{\text{shock}}(t)$, in the test simulations of SNe treatments from Fig. 5. For clarity we show just three resolution levels. ‘Analytic’ assumes the explosion is energy-conserving until the terminal momentum from equation (29) is reached, and subsequently momentum-conserving. At high resolution, all methods resolve the blast and agree with the analytic curve. Because density is fixed, poor mass resolution means the smallest R_{shock} (earlier stages) are not resolved, but the FIRE method reproduces the high-resolution solution at the same radius at all resolution levels. ‘Fully Thermal’ or ‘Fully Kinetic’ models systematically under- or overestimate the momentum and kinetic energy, in a strongly resolution-dependent manner. (The excess thermal energy at low resolution in ‘Fully Kinetic’ models is caused by the shock from the blastwave moving faster than it should, physically.)

(iii) *Fully kinetic*: We assume that 100 percent of the ejecta energy is converted into kinetic energy, i.e. coupled in ‘pure momentum’ form ($p_{\text{coupled}} = \sqrt{2 E_{\text{ejecta}} m_{\text{coupled}}} \geq p_{\text{ejecta}}$, $E_{\text{coupled}} = E_{\text{kinetic}} = E_{\text{ejecta}}$). This ignores any un-resolved cooling. This is similar (algorithmically) to many common implementations in e.g. Aguirre et al. (2001); Springel & Hernquist (2003a); Cen, Nagamine & Ostriker (2005); Dalla Vecchia & Schaye (2008); Vo-

gelsberger et al. (2013) (although most of these authors alter the fraction of energy coupled).

(iv) *Fully thermal*: We assume 100 percent of the ejecta energy is converted into thermal energy, with zero momentum (i.e. $E_{\text{coupled}} = E_{\text{thermal}} = E_{\text{ejecta}}$, $p_{\text{coupled}} = 0$). This also ignores any un-resolved cooling. This is a common implementation used in e.g. Katz, Hernquist & Weinberg (1992); Ceverino & Klypin (2009); Ceverino et al. (2014); Kim et al. (2016).

We evolve the explosion until well after it reaches an asymptotic terminal momentum: when the momentum changes by < 1 per cent over a factor of > 2 increase in the shock radius, or – if this occurs before the shock reaches > 10 inter-particle spacings – when the shock radius moves by < 1 per cent over a factor of 2 increase in time.

In Fig. 5 we plot the terminal momentum in each simulation and compare to our analytic scaling from equation (29). In Fig. 6, we plot the radial profile of the shock properties as the shock radius expands: the total radial momentum, kinetic, and thermal energy (these depend on the time since explosion, so we plot each resolution at different times). We consider particle masses ranging from $m_i < 0.1 M_{\odot}$, sufficient to resolve even the free-expansion phase of explosion (let alone the cooling radius), to $m_i > 10^6 M_{\odot}$.

At sufficiently high resolution, all of the schemes above give identical, well-converged solutions – as they should, since in all cases (at high enough resolution) they generate a shock with the same initial energy, which undergoes an energy-conserving S–T type expansion (in which case the asymptotic solution is fully determined by the ambient density and total blastwave energy). In this limit, the shock formation, S–T phase, conversion of energy into momentum, cooling radius, snowplow phase, and ultimate effective conversion of energy into momentum via PdV work are explicitly resolved, so it does not matter how we initially input the energy. Reassuringly, equation (29) agrees well with the *predicted* terminal momentum in the highest-resolution simulations – in other words, given the cooling physics in FIRE-2, we are using the correct p_t .

At poor resolution, the different treatments diverge, as predicted in Section 2.3.3. For ‘Thermal (+Ejecta)’ and ‘Fully Thermal’ couplings, when the particle mass $m_i \gtrsim 100 M_{\odot}$, the predicted momentum and kinetic energy drop rapidly compared to the converged, exact solutions. Physically, the cooling radius – which is roughly the radius enclosing a fixed mass $m_{\text{cool}} \sim 1000 M_{\odot}$ (see Section 2.3.3) – becomes unresolved. Spreading only thermal energy among this large a gas mass leads to post-shock temperatures below the peak in the cooling curve, so the energy is immediately radiated before much work can be done to accelerate gas (increase the momentum). The terminal momentum and kinetic energy are under-estimated by constant factors of ~ 60 and ~ 3600 , respectively. With the ‘Thermal (+Ejecta)’ case, the same problem occurs, but the initial ejecta momentum remains present, so the terminal momentum and kinetic energy are underestimated by factors of ~ 14 and ~ 200 .

The ‘Fully Kinetic’ coupling errs in the opposite direction at poor resolution: assuming perfect conversion of energy to momentum and ignoring cooling losses gives $p_{\text{coupled}} = \sqrt{2 E_{\text{ejecta}} m_{\text{coupled}}}$, so $p_{\text{coupled}}/p_{\text{terminal}} \propto m_{\text{coupled}}^{1/2}$, and the terminal momentum is over-estimated by a factor of $\sim 20 (m_i/10^4 M_{\odot})^{1/2}$. The kinetic energy is overestimated by a corresponding factor $\sim 400 (m_i/10^4 M_{\odot})$.

In contrast, the FIRE sub-grid model reproduces the high-resolution exact solutions correctly, *independent* of resolution (within < 10 per cent in momentum, kinetic and thermal energy, even at $m_i \sim 10^6 M_{\odot}$). This is the desired behaviour in a ‘good’ sub-grid model. Of course, at poor resolution, the cooling radius

is un-resolved, so the simulation cannot capture the early phases where gas is shock-heated to large temperatures. However the sub-grid treatment captures the correct behaviour of the high-resolution blastwave once it has expanded to a mass or spatial resolution scale which *is* resolved in the low-resolution run. (For similar experiments which reach the same conclusions, see e.g. Fig. 6 in Kim & Ostriker 2017).

4.2 Effects in FIRE simulations: correctly dealing with energy and momentum matters

Having seen in Section 4.1 that correctly accounting for unresolved PdV work in expanding SNe is critical to resolution-independent solutions, we now apply the different treatments therein to cosmological simulations.

Figs 7 and 8 show the results for both dwarf- and MW-mass galaxies in cosmological simulations as well as controlled re-starts of the same MW-mass galaxy at $z \sim 0$ (to ensure identical late-time ICs).

At the mass resolution scales in Figs 7 and 8, our default FIRE-2 coupling scheme reproduces accurately the SN momentum, kinetic and thermal energies from much higher-resolution idealized simulations. In contrast, the ‘Thermal (+Ejecta)’ and ‘Fully Kinetic’ models severely under- and overestimate, respectively, the kinetic energy imparted by SNe (relative to high-resolution simulations and/or analytic solutions). Not surprisingly, then, this is immediately evident in the galaxy evolution. ‘Fully Thermal’ and ‘Thermal (+Ejecta)’ cases resemble a ‘no SNe’ case – because the cooling radii are unresolved, the energy is radiated away immediately, and the terminal momentum that *should have been resolved* is not properly accounted for – so SNe do far less work than they should and far more stars form. The Fully Kinetic case, on the other hand, wildly overestimates the conversion of thermal energy to kinetic (and ignores cooling losses), so star formation is radically suppressed.

Given this strong dependence, one might wonder whether the exact details of our FIRE treatment might change the results. However these are not so important. In Appendix F, we consider a ‘no implicit cooling’ model: here we take our standard FIRE-2 coupling (the coupled momentum, mass, and metals are unchanged), but even if the cooling radius is un-resolved, we still couple the full ejecta energy (i.e. we do not assume, implicitly, that the ejecta thermal energy has radiated away if we do not resolve the cooling radius, so couple a total thermal plus kinetic energy $=E_{\text{ejecta}}$). This produces no detectable difference from our default model, which is completely expected. If the cooling radius is resolved, our default model does not radiate the energy away; if it is unresolved, ‘keeping’ the thermal energy in the SNe coupling step simply leads to its being radiated away explicitly in the simulation cooling step on the subsequent time-step.

Fig. 8 considers the effects of changing the analytic terminal momentum p_t in equation (29), by a factor of ~ 4 . As discussed in Section 2.3, while there are physical uncertainties in this scaling owing to uncertain microphysics of blastwave expansion, they are generally smaller. But in any case, the effect on our galaxy-scale simulations is relatively small, even at low resolution. As expected, smaller p_t leads to higher SFRs, because the momentum coupled per SN is smaller, so more stellar mass is needed to self-regulate. In a simple picture where momentum input self-regulates SF and wind generation (see e.g. Ostriker & Shetty 2011; Faucher-Giguère, Quataert & Hopkins 2013; Hayward & Hopkins 2017), we would expect the SFR to be inversely proportional to p_t at low resolution. However, because of non-linear effects, and the fact that even at low

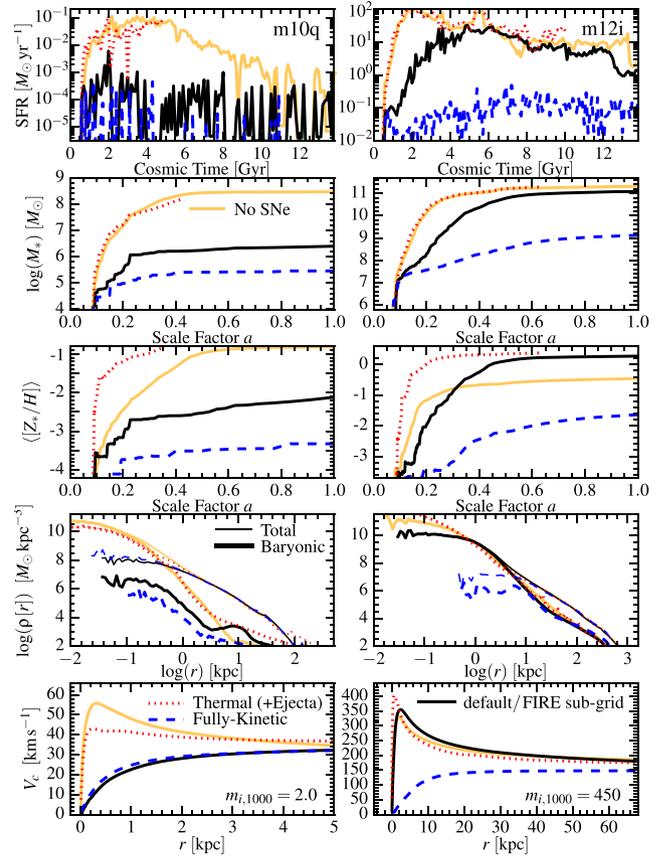


Figure 7. Comparison of the alternative sub-grid assumptions (Section 2.3) from Figs 5 and 6, in zoom-in cosmological simulations. (1) *Default*: the ‘FIRE Sub-Grid’ model, which matches the correct momentum and energy from high-resolution solutions of individual blastwave explosions at all resolution levels. (2) *Thermal (+Ejecta)*: This couples just the initial ejecta momentum (ignoring any un-resolved PdV work done), with the remaining energy in pure thermal form. From our high-resolution tests, this *underestimates* the final momentum (kinetic energy) of SNe by a factor of ~ 16 (~ 250) at the low resolution of the tests here. Therefore our predictions resemble a ‘no SNe’ case, and stars form extremely rapidly in the early Universe. (3) *Fully kinetic*: This couples the 100 per cent of the SNe ejecta energy as kinetic energy: our high-resolution tests show this *overestimates* the momentum (kinetic energy) by a factor of ~ 10 (~ 40) for **m10q** (with $m_{i,1000} = 2$) and ~ 160 (~ 2500) for **m12i** (with $m_{i,1000} = 450$). Not surprisingly, star formation is overwhelmingly suppressed (a MW-mass halo forms a $< 10^9 M_{\odot}$ dwarf). (4) *No SNe*: No supernovae included.

resolution the simulations resolve massive superbubbles (where p_t does not matter because the cooling radius for overlapping explosions is resolved), the actual dependence is sub-linear, $M_* \propto p_t^{-0.3}$. So given the (small) physical uncertainties, this is not a dominant source of error.¹⁴

¹⁴ To be clear, in Fig. 8 we alter *only* the terminal momentum, so e.g. if the cooling radius of superbubbles is resolved the change has no effect whatsoever, and other feedback mechanisms (e.g. radiative feedback) are also un-altered. In contrast, in Orr et al. (2017) (Appendix A) we show the results of multiplying/dividing *all* feedback mechanisms and strengths (total energy and momentum) by a uniform factor $=3$. Not surprisingly this produces a stronger effect closer to the expected inverse-linear dependence; however non-linear effects still reduce the dependence to somewhat sub-linear.

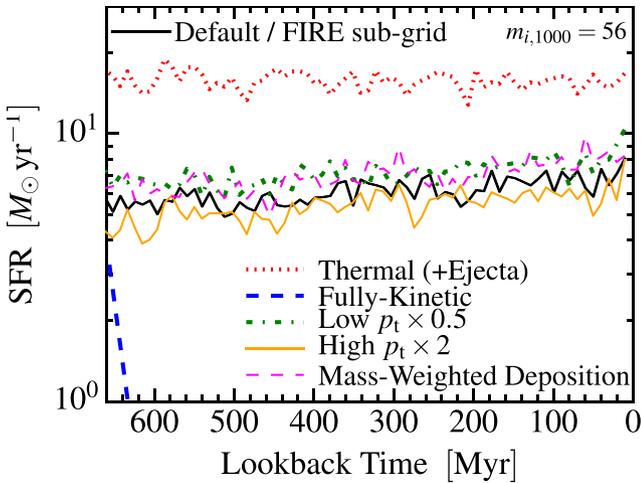


Figure 8. Comparison of the sub-grid SNe treatments from Figs 5–7. In each case, we restart our ‘default’ **m12i** simulation (resolution $m_{i,1000} = 56$) at $z = 0.07$ and run to $z = 0$ (~ 700 Myr) using different feedback implementations. This allows us to examine how different physics change the predicted SFR within the galaxy on smaller-than-cosmological time-scales given an *identical* initial galaxy (removing the non-linear effects over cosmological time present in Fig. 7). As in Fig. 7, we compare our default model, which correctly captures the momentum, thermal, and kinetic energy of SNe at this resolution level ($m_{i,1000} = 56$), to the ‘Thermal (+Ejecta)’ model which underestimates the momentum (kinetic energy) of SNe by a factor of ~ 16 (~ 250) at this resolution, and the ‘Fully Kinetic’ model which overestimates the momentum (kinetic energy) by a factor of ~ 50 (~ 400) at this resolution. As expected, the ‘Thermal (+Ejecta)’ model produces much higher SFRs (the SFR cannot become much higher than $\sim 20 M_{\odot} \text{ yr}^{-1}$, here, because it is limited by the free-fall time), while the ‘Fully Kinetic’ model overwhelmingly suppresses the SFR (the gaseous disc essentially ‘explodes’ in the first dynamical time). We then re-run our ‘default FIRE sub-grid’ model, but systematically increase (decrease) the terminal momentum p_t in equation (29) by a factor of 2 (much larger than physical uncertainties; see Section 2.3). Higher p_t (i.e. larger assumed cooling radii) produce slightly more-suppressed star formation (as expected) but with only an ~ 50 per cent change in SFR for a factor of ~ 4 change in p_t .

Recently, Rosdahl et al. (2017) performed a similar experiment, exploring different SNe implementations in the adaptive mesh refinement (AMR) code RAMSES. They used a different treatment of cooling and star formation, non-cosmological simulations, and no other feedback. However, their conclusions are similar, regarding the relative efficiencies of the ‘Fully Thermal,’ ‘FIRE-sub-grid’ (in their paper, the ‘mechanical’ model), and ‘Fully Kinetic’ treatments of SNe. Our conclusions appear to be robust across a wide range of conditions and detailed numerical treatments.

Again, we have repeated these tests in other haloes to ensure that our conclusions are not unique to a single galaxy. Specifically we have compared a ‘Fully Thermal’ and ‘Fully Kinetic’ run in halo **m10v** and **m12f** from Paper I, and compared re-starts from $z = 0.07$ of an **m12f** run with $m_{i,1000} = 56$ using the same set of parameter variations as Fig. 8. The results are nearly identical to our studies with **m10q** and **m12i**.

4.3 Convergence: incorrect sub-grid treatments converge to the resolution-independent FIRE scaling

In Fig. 9, we consider another convergence test of the SNe coupling scheme, but this time in cosmological simulations. We re-run our

m10q simulation with standard FIRE-2 physics, considering our default SNe treatment as well as the ‘Thermal (+Ejecta)’ and ‘Fully Kinetic’ models, with mass resolution varied from 30×10^5 to $1.3 \times 10^5 M_{\odot}$.

Not only does our default FIRE treatment of SNe produce excellent convergence in the star formation history across this entire resolution range, but *both* the ‘Thermal (+Ejecta)’ model (which suffers from overcooling, hence excessive SF, at low resolution because the SNe energy is almost all coupled thermally) and the ‘Fully Kinetic’ model (which overestimates the kinetic energy of SNe, hence oversuppresses SF, at low resolution) converge to our FIRE solution at higher resolution, especially at $m_i \lesssim 100 M_{\odot}$. Of course, even at our highest resolution, details of SNe shells and venting can differ in the early stages of ejecta expansion, so convergence is not perfect – but the trends clearly approach the ‘default’ model.

4.4 On ‘Delayed-Cooling’ and ‘Target-Temperature’ models

Given the failure of ‘Fully Thermal’ models at low resolution, a popular ‘fix’ in the galaxy formation literature is to artificially suppress gas cooling at large scales, either explicitly or implicitly. This is done via (a) ‘delayed cooling’ prescriptions, for which energy injected by SNe is not allowed to cool for some large time-scale $\Delta t_{\text{delay}} \gtrsim t_{\text{dynamical}} \sim 10^{7-8}$ yr (as in Thacker & Couchman 2000, 2001; Stinson et al. 2006, 2013; Dubois et al. 2015), or (b) ‘target temperature’ prescriptions, where SNe energy is ‘stored’ until sufficient energy is accumulated to heat (in a single ‘event’) a large resolved gas mass to some high temperature $T_{\text{target}} \gg 10^7$ K (as in Gerritsen 1997; Mori et al. 1997; Dalla Vecchia & Schaye 2012; Crain et al. 2015).

Although these approximations may be useful in low-resolution simulations with $m_i \gtrsim 10^6 M_{\odot}$ (typical of large-volume cosmological simulations), where ISM structure and the clustering of star formation cannot be resolved, they are fundamentally ill-posed for simulations with resolved ISM structure, for at least three reasons. (1) Most importantly, they are *non-convergent* (at least as defined here). This is easy to show rigorously, but simply consider a case with arbitrarily good resolution: then either (a) turning off cooling for longer than the actual shock-cooling time, or (b) enforcing a ‘target temperature’ that does not exactly match the initial reverse-shock temperature will produce un-physical results. Strictly speaking, there is no define-able convergence criterion for these models: they do not interpolate to the correct solution as resolution increases, but to some other (non-physical) system. (2) They do not represent the converged solution in Fig. 6 at any low-resolution radius/mass. Once an SN has swept through, say, $\sim 10^6 M_{\odot}$ of gas, it should, correctly, be a cold shell, not a hot bubble. Thus we are not reproducing the higher-resolution solutions correctly, at some finite practical resolution. (3) They introduce an additional set of parameters: Δt_{delay} or T_{target} , and the ‘size’ (or mass) of the region that is influenced. Both of these strongly influence the results. For example, by increasing the region size, one does not simply ‘spread’ the same energy among neighbours differently, but rather, because the models are binary, one either (a) increases the mass that cannot cool or (b) must change the number of SNe ‘stored up’ (hence the implicit cooling-delay-time) to reach T_{target} .

In Appendix D we consider some implementations of these models, at the resolutions studied here. As expected, we show that they do not converge as we approach resolution $\sim 100 M_{\odot}$, and that certain galaxy properties (metallicities, star formation histories) exhibit biases that are clear artefacts of the un-physical nature of these

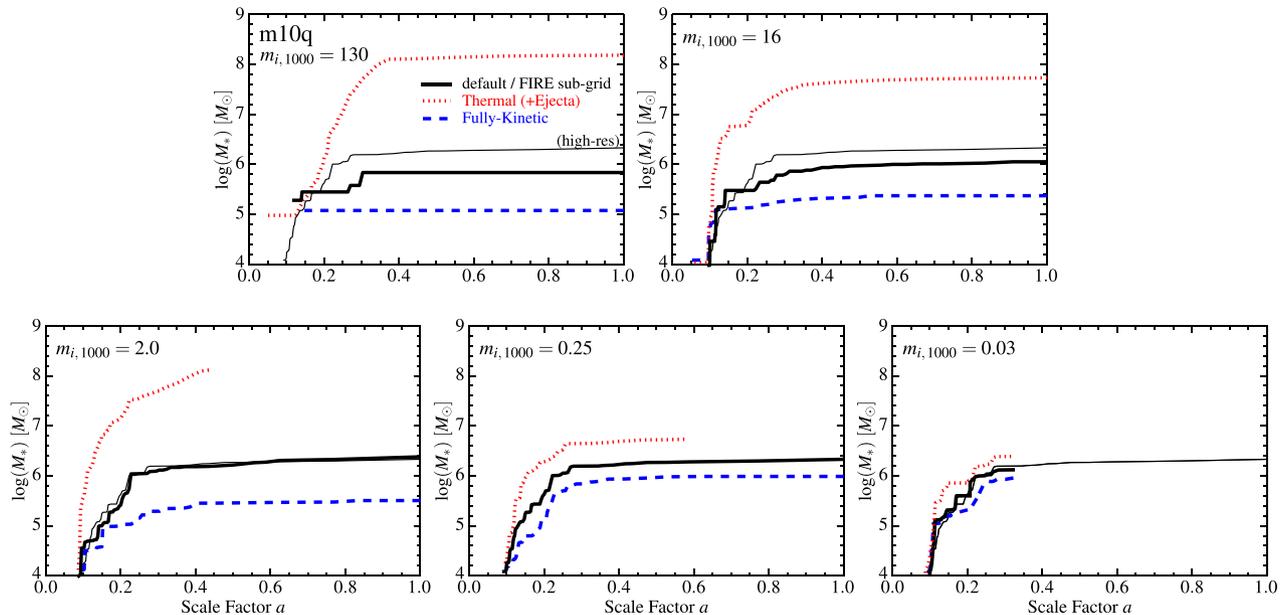


Figure 9. Convergence properties of the sub-grid SNe treatments from Figs 5–7 in a cosmological simulation of a dwarf galaxy. We show stellar mass, M_* , versus cosmic time (as Fig. 7) for each treatment of unresolved phases of SNe explosions. We did not run our highest-resolution simulations (with $\sim 30 M_\odot$ resolution) to $z = 0$ because of computational expense. At low mass resolution, the ‘Thermal (+Ejecta)’ assumption underestimates the strength of feedback (producing *too-high* M_*), while the ‘Fully Kinetic’ assumption overestimates the strength of feedback (producing *too-low* M_*). Our default FIRE sub-grid treatment resembles the high-resolution converged solution even at low mass resolution. Once we reach mass resolution $\ll 1000 M_\odot$ (*bottom*), the cooling radius of individual SN blastwaves starts to become explicitly resolved, and the different treatments converge towards one another – as predicted from our idealized tests of single explosions in Fig. 5. Critically, the ‘Thermal (+Ejecta)’ (similar to ‘Fully Thermal’) and ‘Fully Kinetic’ treatments *both* converge to our ‘default FIRE sub-grid’ treatment at higher resolution.

coupling schemes at high resolution. We therefore do not focus on them further.

5 DISCUSSION AND CONCLUSIONS

We have presented an extensive study of both numerical and physical aspects of the coupling of mechanical feedback in galaxy formation simulations (most importantly, SNe, but the methods are relevant to stellar mass-loss and black hole feedback). We explored this in both idealized calculations of individual SN remnants and the FIRE-2 cosmological simulations at both dwarf and MW mass scales. We conclude that there are two critical components to an optimal algorithm, summarized below.

5.1 Ensuring conservation and statistical isotropy

It is important to design an algorithm that is statistically isotropic (i.e. does not numerically bias the feedback to prefer certain directions), and manifestly conserves mass, metals, momentum, and energy. This is particularly non-trivial in mesh-free numerical methods. In particular, naively distributing ejecta with a simple kernel or area weight to ‘neighbour’ cells or particles – as is common practice in most numerical treatments – can easily produce violations of linear momentum conservation and bias the ejecta so that in, for example, a thin disc, feedback preferentially acts (incorrectly) in the disc plane instead of venting out. This is especially important for *any* numerical method for which the gas resolution elements might be irregularly distributed around a star (e.g. moving-mesh codes, SPH, or AMR if the star is not at the exact cell centre). If these constraints are not met, we show that spurious numerical torques or outflow geometries can artificially remove disc angular momentum

and bias predicted morphologies. Worse yet, the momentum conservation errors may not converge and can become more important at high resolution.

In fact, as discussed in detail in Section 3.2, our older published ‘FIRE-1’ simulations suffered from some of these errors, but (owing to lower resolution) they were relatively small. Higher-resolution tests, however, demonstrated their importance, motivating the development of the new FIRE-2 algorithm.

In Section 2.2 we present a general algorithm (used in FIRE-2) that resolves all of these issues (as well as accounting for relative star-gas motions), and can trivially be applied in any numerical galaxy formation code (regardless of hydrodynamic method), for any mechanical feedback mechanism.

5.2 Accounting for energy and momentum from un-resolved ‘PdV Work’

At the mass (m_i) or spatial resolution (h_i) of current cosmological simulations, it is *physically incorrect* to couple SNe to the gas as entirely thermal energy (heating-only) or as entirely kinetic energy (momentum transfer only) or as the initial ejecta mix of momentum and energy. Because the SN blast wave has implicitly propagated through a region containing mass $\sim m_i$, it *must* have either (a) done some mechanical (‘PdV’) work, increasing the momentum of the blastwave, and/or (b) radiated its energy away. In Hopkins et al. (2014), we proposed a simple way to account for this in simulations, which we provide in detail in Section 2.3. This method is used in all FIRE simulations, and was further tested in idealized simulations by Martizzi et al. (2015), and similar methods have been developed and used in galaxy formation simulations by e.g. Kimm & Cen (2014); Rosdahl et al. (2017). Essentially, we account

for the PdV work by imposing energy conservation up to a terminal momentum (equation (29)), beyond which the energy is radiated, with the transition occurring at the cooling radius of the blastwave.

In this paper, we use high-resolution (reaching $<0.1 M_{\odot}$) simulations of individual SN to show that this implementation, independent of the resolution at which it is applied, reproduces the exact, converged high-resolution simulation of a single SN blastwave, given the *same* physics. In other words, taking a high-resolution simulation of an SN in a homogenous medium and smoothing it at the resolved coupling radius produces the same result as what is directly applied to the large-scale simulations. Perhaps most importantly, we show that this method of partitioning thermal and kinetic energy leads to relatively rapid convergence in predicted stellar masses and star formation histories in galaxy-formation simulations.

In contrast, coupling only thermal or kinetic energy (or the initial ejecta partitioning of the two) will over- or underpredict the coupled momentum by orders of magnitude, in a strongly resolution-dependent fashion (Fig. 5). Briefly, at poor resolution, coupling $\sim 10^{51}$ erg as thermal energy (e.g. including no momentum or only the initial ejecta momentum) spreads the energy over an artificially large mass, so the gas is barely heated and efficiently radiates the energy away without resolving the PdV work. But simply converting all (or any resolution-independent fraction) of this energy into kinetic energy, on the other hand, ignores the cooling that should have occurred and will always, at sufficiently poor resolution, overestimate the correct momentum generated in a resolution-dependent manner (since for fixed kinetic energy input, the *momentum* generated is a function of the mass resolution). This in turn leads to strongly resolution-dependent predictions for galaxy masses (Fig. 9). In principle, one could compensate for this by introducing explicitly resolution-dependent ‘efficiency factors’ that are re-tuned at each resolution level to produce some ‘desired’ result, but this severely limits the predictive power of the simulations and will still fail to produce the correct mix of phases in the ISM and outflows (because the correct thermal-kinetic energy mix is not present). Using cosmological simulations reaching $\sim 30 M_{\odot}$ resolution, we show that *all* of these studied coupling methods do converge to the same solution when applied at sufficiently high resolution. The difference is that the proposed method in Section 2.3 from the FIRE simulations converges much more quickly (at a factor of ~ 1000 lower resolution), while the unphysical ‘Fully Thermal’ or ‘Fully Kinetic’ approaches require mass resolution $\ll 100 M_{\odot}$.

5.3 Caveats and future work

While the SNe coupling algorithm studied here reproduces the converged, high-resolution solution at any practical resolution, it is of course possible that the actual conditions under which the SNe explode (the local resolved density, let alone density sub-structure) continue to change as simulation resolution increases. The small-scale density structure of the ISM might in turn depend on other physics (e.g. H_{II} regions, radiation pressure), which could have different convergence properties from the SNe alone.

We stress that our conclusions are relevant for simulations of the ISM or galaxies with mass resolution in the range $10 M_{\odot} \lesssim m_i \lesssim 10^6 M_{\odot}$. Below $\ll 100 M_{\odot}$, simulations directly resolve early stages of SNe remnant evolution, and it is less important that the coupling is done accurately because the relevant dynamics will be explicitly resolved. Above $\gg 10^6 M_{\odot}$, it quickly becomes impossible to resolve even the largest scales of fragmentation and multi-phase structure in the ISM. Such star formation cannot cluster and SNe are not individually time-resolved (i.e. a resolution element

has many SNe per time-step), so there is no possibility of explicitly resolving overlap of many SNe into superbubbles, regardless of how the SNe are treated. In that limit, it is necessary to implement a *galaxy-wide* sub-grid model for SNe feedback (e.g. a model that directly implements a mass-loading of galactic winds as presented in e.g. Vogelsberger et al. 2013; Davé, Thompson & Hopkins 2016; Davé et al. 2017).

The scalings above for un-resolved ‘ PdV work’ are well-studied for SNe, but much less well-constrained for quasi-continuous processes such as stellar mass-loss (OB and AGB winds) and AGN accretion-disc winds. In both cases, the problem is complicated by the fact that the structure and time-variability (e.g. ‘burstiness’) of the mass-loss processes themselves is poorly understood. Especially for energetic AGN-driven winds, more work is needed to better understand these regimes.

Finally, new physics not included here could alter our conclusions. For example, magnetic fields, or anisotropic thermal conduction, or plasma instabilities altering fluid mixing, or cosmic rays, could all influence the SNe cooling and expansion. Different stellar evolution models could change the predicted SNe rates and/or energetics. It is not our intent to say that the solution here includes all possible physics. However, *independent* of these physics, the two key points (Sections 5.1–5.2) must still hold! And the goal of any ‘sub-grid’ representation of SNe should be to represent the converged solution *given the same physics* as the large-scale simulation – otherwise convergence cannot even be defined in any meaningful sense. So in future work it would be valuable to repeat the exercises in this paper for modified physical assumptions. However, the extensive literature studying the effect of different physical conditions on SNe remnant evolution (see references in Section 2.3.1) has shown that the terminal momentum is weakly sensitive to these additional physics.

ACKNOWLEDGEMENTS

We thank our referee, Joakim Rosdahl, for a number of insightful comments. Support for PFH and co-authors was provided by an Alfred P. Sloan Research Fellowship, NASA ATP Grant NNX14AH35G, and NSF Collaborative Research Grant #1411920 and CAREER grant #1455342. AW was supported by a Caltech-Carnegie Fellowship, in part through the Moore Center for Theoretical Cosmology and Physics at Caltech, and by NASA through grant HST-GO-14734 from STScI. CAFG was supported by NSF through grants AST-1412836 and AST-1517491, and by NASA through grant NNX15AB22G. DK was supported by NSF Grant AST1412153 and a Cottrell Scholar Award from the Research Corporation for Science Advancement. The Flatiron Institute is supported by the Simons Foundation. Numerical calculations were run on the Caltech compute cluster ‘Wheeler,’ allocations TG-AST120025, TG-AST130039 and TG-AST150080 granted by the Extreme Science and Engineering Discovery Environment (XSEDE) supported by the NSF, and the NASA HEC Program through the NAS Division at Ames Research Center and the NCCS at Goddard Space Flight Center.

REFERENCES

- Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003, *ApJ*, 597, 21
 Agertz O., Kravtsov A. V., 2015, *ApJ*, 804, 18
 Agertz O., Kravtsov A. V., Leitner S. N., Gnedin N. Y., 2013, *ApJ*, 770, 25
 Aguirre A., Hernquist L., Schaye J., Weinberg D. H., Katz N., Gardner J., 2001, *ApJ*, 560, 599

- Agúndez M., Cernicharo J., Guélin M., 2010, *ApJL*, 724, L133
- Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Hopkins P. F., Quataert E., Murray N., 2017, *MNRAS*, 470, 4698
- Behroozi P. S., Conroy C., Wechsler R. H., 2010, *ApJ*, 717, 379
- Bournaud F., Elmegreen B. G., Teyssier R., Block D. L., Puerari I., 2010, *MNRAS*, 409, 1088
- Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80
- Byerly Z. D., Adelstein-Lelbach B., Tohline J. E., Marcello D. C., 2014, *ApJS*, 212, 23
- Cen R., Nagamine K., Ostriker J. P., 2005, *ApJ*, 635, 86
- Ceverino D., Klypin A., 2009, *ApJ*, 695, 292
- Ceverino D., Klypin A., Klimek E., Trujillo-Gomez S., Churchill C. W., Primack J., Dekel A., 2014, *MNRAS*, 442, 1545
- Cioffi D. F., McKee C. F., Bertschinger E., 1988, *ApJ*, 334, 252
- Coil A. L., Weiner B. J., Holz D. E., Cooper M. C., Yan R., Aird J., 2011, *ApJ*, 743, 46
- Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, *MNRAS*, 319, 168
- Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201
- Conroy C., van Dokkum P., Kravtsov A., 2015, *ApJ*, 803, 77
- Cox N. L. J. et al., 2012, *A&A*, 537, A35
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Dalla Vecchia C., Schaye J., 2008, *MNRAS*, 387, 1431
- Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
- Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
- Davé R., Rafieferantsoa M. H., Thompson R. J., Hopkins P. F., 2017, *MNRAS*, 467, 115
- Davis S. F., 1984, *J. Comput. Phys.*, 56, 65
- de Val-Borro M. et al., 2006, *MNRAS*, 370, 529
- Dobbs C. L., Burkert A., Pringle J. E., 2011, *MNRAS*, 413, 528
- Draine B. T., Woods D. T., 1991, *ApJ*, 383, 621
- Dubois Y., Volonteri M., Silk J., Devriendt J., Slyz A., Teyssier R., 2015, *MNRAS*, 452, 1502
- Evans II N. J., 1999, *ARA&A*, 37, 311
- Evans N. J. et al., 2009, *ApJS*, 181, 321
- Faucher-Giguère C.-A., Quataert E., 2012, *MNRAS*, 425, 605
- Faucher-Giguère C.-A., Quataert E., Hopkins P. F., 2013, *MNRAS*, 433, 1970
- Faucher-Giguère C.-A., Hopkins P. F., Keres D., Muratov A. L., Quataert E., Murray N., 2015, *MNRAS*, 449, 987
- Faucher-Giguère C.-A., Feldmann R., Quataert E., Kereš D., Hopkins P. F., Murray N., 2016, *MNRAS*, 461, L32
- Feldmann R., Hopkins P. F., Quataert E., Faucher-Giguère C.-A., Kereš D., 2016, *MNRAS*, 458, L14
- Fong D., Meixner M., Shah R. Y., 2003, *ApJL*, 582, L39
- Freyer T., Hensler G., Yorke H. W., 2006, *ApJ*, 638, 262
- Gentry E. S., Krumholz M. R., Dekel A., Madau P., 2017, *MNRAS*, 465, 2471
- Gerritsen J. P. E., 1997, PhD thesis, Groningen Univ., the Netherlands
- Governato F. et al., 2004, *ApJ*, 607, 688
- Hahn O., Teyssier R., Carollo C. M., 2010, *MNRAS*, 405, 274
- Haid S., Walch S., Naab T., Seifried D., Mackey J., Gatto A., 2016, *MNRAS*, 460, 2962
- Harper-Clark E., Murray N., 2009, *ApJ*, 693, 1696
- Harper-Clark E., Murray N., 2011, in Alves J., Elmegreen B. G., Girat J. M., Trimble V., eds., *Computational Star Formation*. Vol. 270. Cambridge Univ. Press, Cambridge, UK, p. 235
- Hayward C. C., Hopkins P. F., 2017, *MNRAS*, 465, 1682
- Heckman T. M., Lehnert M. D., Strickland D. K., Armus L., 2000, *ApJS*, 129, 493
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., 2016, *MNRAS*, 462, 576
- Hopkins P. F., 2017, *MNRAS*, 466, 3387
- Hopkins P. F., Raives M. J., 2016, *MNRAS*, 455, 51
- Hopkins P. F., Hernquist L., Martini P., Cox T. J., Robertson B., Di Matteo T., Springel V., 2005, *ApJL*, 625, L71
- Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950
- Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, 421, 3488
- Hopkins P. F., Narayanan D., Murray N., 2013, *MNRAS*, 432, 2647
- Hopkins P. F., Keres D., Onorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F. et al., 2017, *MNRAS*, submitted ([arXiv:1702.06148](https://arxiv.org/abs/1702.06148))
- Hu C.-Y., Naab T., Walch S., Glover S. C. O., Clark P. C., 2016, *MNRAS*, 458, 3528
- Iffrig O., Hennebelle P., 2015, *A&A*, 576, A95
- Kannan R., Stinson G. S., Macciò A. V., Brook C., Weinmann S. M., Wadsley J., Couchman H. M. P., 2014, *MNRAS*, 437, 3529
- Katz N., Hernquist L., Weinberg D. H., 1992, *ApJL*, 399, L109
- Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
- Kennicutt R. C., Jr., 1998, *ApJ*, 498, 541
- Kereš D., Katz N., Davé R., Fardal M., Weinberg D. H., 2009, *MNRAS*, 396, 2332
- Kim J.-h. et al., 2016, *ApJ*, 833, 202
- Kim C.-G., Ostriker E. C., 2015, *ApJ*, 802, 99
- Kim C.-G., Ostriker E. C., 2017, *ApJ*, 846, 133
- Kim C.-G., Ostriker E. C., Raileanu R., 2017, *ApJ*, 834, 25
- Kimm T., Cen R., 2014, *ApJ*, 788, 121
- Kroupa P., 2001, *MNRAS*, 322, 231
- Krumholz M. R., Gnedin N. Y., 2011, *ApJ*, 729, 36
- Leitherer C. et al., 1999, *ApJS*, 123, 3
- Li M., Ostriker J. P., Cen R., Bryan G. L., Naab T., 2015, *ApJ*, 814, 4
- Lopez L. A., Krumholz M. R., Bolatto A. D., Prochaska J. X., Ramirez-Ruiz E., 2011, *ApJ*, 731, 91
- Ma X., Hopkins P. F., Faucher-Giguère C.-A., Zolman N., Muratov A. L., Kereš D., Quataert E., 2016, *MNRAS*, 456, 2140
- Martin C. L., 1999, *ApJ*, 513, 156
- Martin C. L., Scannapieco E., Ellison S. L., Hennawi J. F., Djorgovski S. G., Fournier A. P., 2010, *ApJ*, 721, 174
- Martizzi D., Faucher-Giguère C.-A., Quataert E., 2015, *MNRAS*, 450, 504
- McKee C. F., van Buren D., Lazareff B., 1984, *ApJL*, 278, L115
- Mori M., Yoshii Y., Tsujimoto T., Nomoto K., 1997, *ApJL*, 478, L21
- Moster B. P., Somerville R. S., Maudtsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010, *ApJ*, 710, 903
- Muratov A. L., Kereš D., Faucher-Giguère C.-A., Hopkins P. F., Quataert E., Murray N., 2015, *MNRAS*, 454, 2691
- Muratov A. L. et al., 2017, *MNRAS*, 468, 4170
- Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
- Orr M. et al., 2017, *MNRAS*, in press ([arXiv:1701.01788](https://arxiv.org/abs/1701.01788))
- Ostriker E. C., Shetty R., 2011, *ApJ*, 731, 41
- Pettini M., Madau P., Bolte M., Prochaska J. X., Ellison S. L., Fan X., 2003, *ApJ*, 594, 695
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, *MNRAS*, 338, 14
- Price D. J., 2012, *J. Comput. Phys.*, 231, 759
- Rees M. J., Ostriker J. P., 1977, *MNRAS*, 179, 541
- Repolust T., Puls J., Herrero A., 2004, *A&A*, 415, 349
- Robertson B., Yoshida N., Springel V., Hernquist L., 2004, *ApJ*, 606, 32
- Rosdahl J., Schaye J., Dubois Y., Kimm T., Teyssier R., 2017, *MNRAS*, 466, 11
- Rosen A. L., Lopez L. A., Krumholz M. R., Ramirez-Ruiz E., 2014, *MNRAS*, 442, 2701
- Sato T., Martin C. L., Noeske K. G., Koo D. C., Lotz J. M., 2009, *ApJ*, 696, 214
- Scannapieco C. et al., 2012, *MNRAS*, 423, 1726
- Shen S., Madau P., Conroy C., Governato F., Mayer L., 2014, *ApJ*, 792, 99
- Slavin J. D., Cox D. P., 1992, *ApJ*, 392, 131
- Somerville R. S., Primack J. R., 1999, *MNRAS*, 310, 1087
- Songaila A., 2005, *AJ*, 130, 1996
- Sparre M., Hayward C. C., Feldmann R., Faucher-Giguère C.-A., Muratov A. L., Kereš D., Hopkins P. F., 2017, *MNRAS*, 466, 88
- Springel V., Hernquist L., 2002, *MNRAS*, 333, 649
- Springel V., Hernquist L., 2003a, *MNRAS*, 339, 289
- Springel V., Hernquist L., 2003b, *MNRAS*, 339, 312
- Steidel C. C., Erb D. K., Shapley A. E., Pettini M., Reddy N., Bogosavljević M., Rudie G. C., Rakic O., 2010, *ApJ*, 717, 289
- Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, *MNRAS*, 373, 1074

Stinson G. S., Brook C., Macciò A. V., Wadsley J., Quinn T. R., Couchman H. M. P., 2013, *MNRAS*, 428, 129
 Tasker E. J., 2011, *ApJ*, 730, 11
 Thacker R. J., Couchman H. M. P., 2000, *ApJ*, 545, 728
 Thacker R. J., Couchman H. M. P., 2001, *ApJL*, 555, L17
 Thornton K., Gaudlitz M., Janka H., Steinmetz M., 1998, *ApJ*, 500, 95
 Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
 Walch S., Naab T., 2015, *MNRAS*, 451, 2757
 Weaver R., McCray R., Castor J., Shapiro P., Moore R., 1977, *ApJ*, 218, 377
 Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, *ApJL*, 827, L23
 Williams J. P., McKee C. F., 1997, *ApJ*, 476, 166
 Young P. A., Hightberger J. L., Arnett D., Ziurys L. M., 2003, *ApJL*, 597, L53

Zavala J., Okamoto T., Frenk C. S., 2008, *MNRAS*, 387, 364
 Ziurys L. M., Milam S. N., Apponi A. J., Woolf N. J., 2007, *Nature*, 447, 1094
 Zuckerman B., Evans N. J., II, 1974, *ApJL*, 192, L149

APPENDIX A: ADDITIONAL RESOLUTION TESTS

Extensive resolution tests of our ‘default’ algorithm, at both dwarf and MW mass scales (and considering both mass and spatial resolution, and additional haloes), are presented in [Paper I](#). The main text here also directly compares the different sub-grid treatments of un-resolved cooling (‘Fully Thermal’, ‘Fully Kinetic’, etc. models) as a function of resolution. Here we simply note that we have re-run

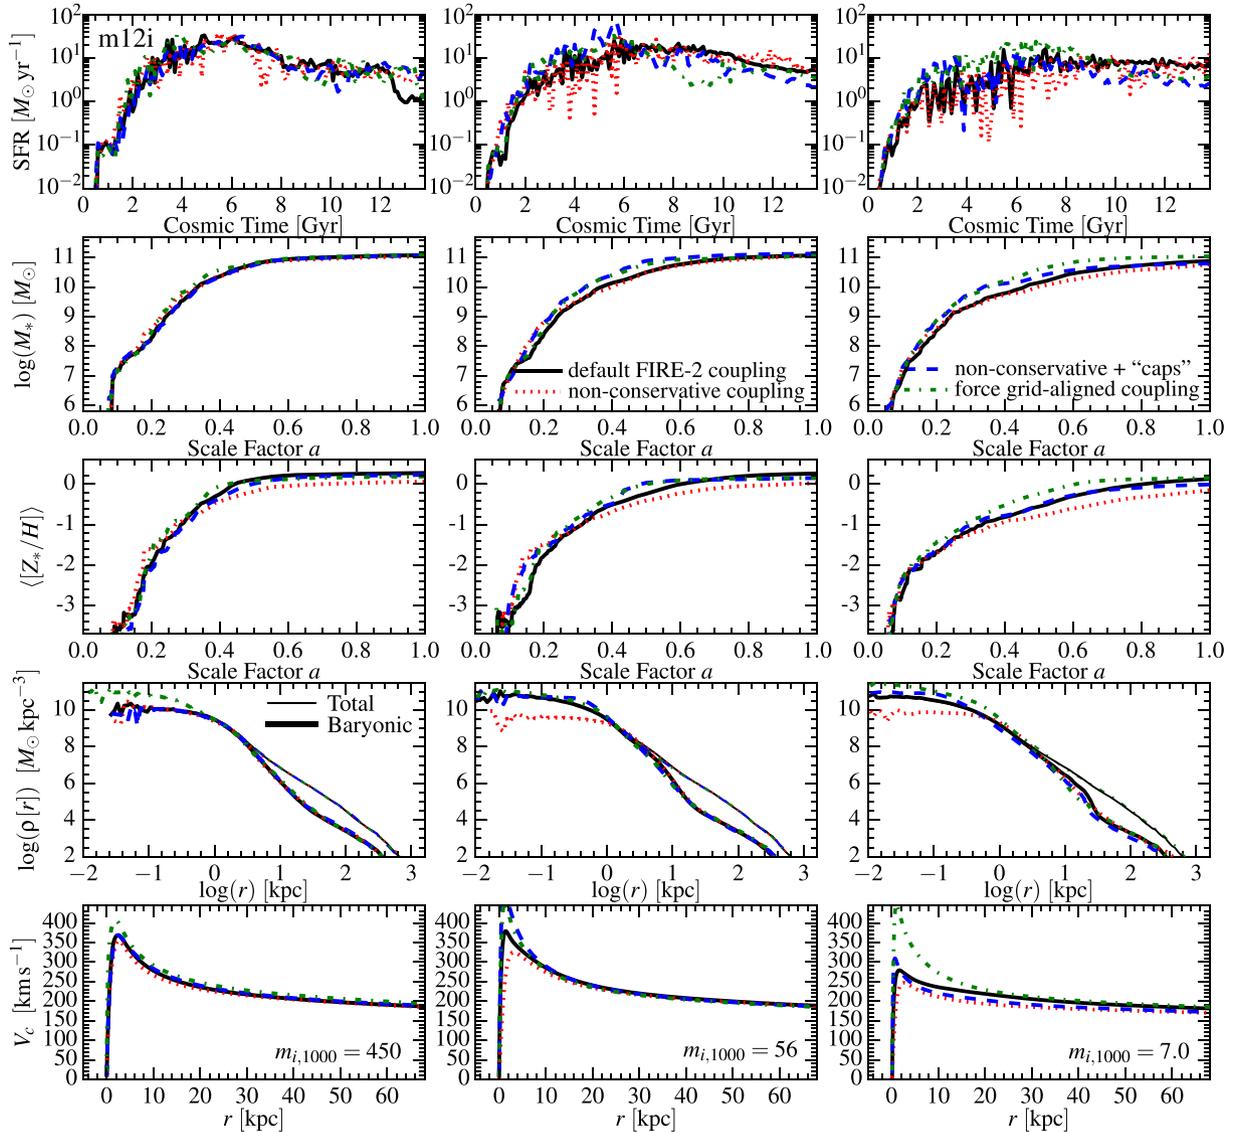


Figure A1. Additional mechanical FB coupling tests, as Fig. 3. We focus on MW-mass galaxies, where differences are maximized. Here we compare the ‘default FIRE-2 coupling’ and the ‘non-conservative coupling’ algorithms from the text, at three resolution levels (labeled). We also compare: *Non-conservative+‘caps’*: a modified version of the non-conservative algorithm which adds simple ‘caps’ to the feedback to prevent conservation errors from getting too large, at the expense of not always coupling the feedback that should be present (Section B). The artificial suppression slightly weakens feedback but resolves the larger differences between ‘default’ and ‘non-conservative’ algorithms. *Force grid-aligned coupling*: a scheme where we artificially inject all feedback in the Cartesian $x/y/z$ grid directions (regardless of the physical gas particle/cell shapes and distribution around the star), mimicking simple coupling in a fixed grid (Section C). This artificially forces alignment of the disc and winds with coordinate axes, generating spurious torques that remove angular momentum and make the disc more compact. Our conclusions from the text do not change with resolution.

tests of the different purely numerical SNe coupling schemes from Figs 3 to 4, at both dwarf and MW mass scales, at several resolution levels. In both cases we find our conclusions from the main text are not sensitive to resolution. In the dwarf case this is unsurprising, since there was no significant effect from the coupling algorithm. For MW-mass haloes, we demonstrate this explicitly in Fig. A1.

APPENDIX B: CONFIRMATION THAT THE ERRORS IN THE NON-CONSERVATIVE ALGORITHM ARE DOMINATED BY EXTREME EVENTS

In Section 3.2, we showed that at sufficiently high resolution, the non-conservative algorithm can produce momentum errors which destroy the thin-disc morphology of a simulated MW-mass galaxy. Here we confirm that the dramatic effects seen there are dominated by the smaller number of ‘extreme’ or ‘worst case’ events, rather than smaller errors that occur more ubiquitously in a non-conservative algorithm.

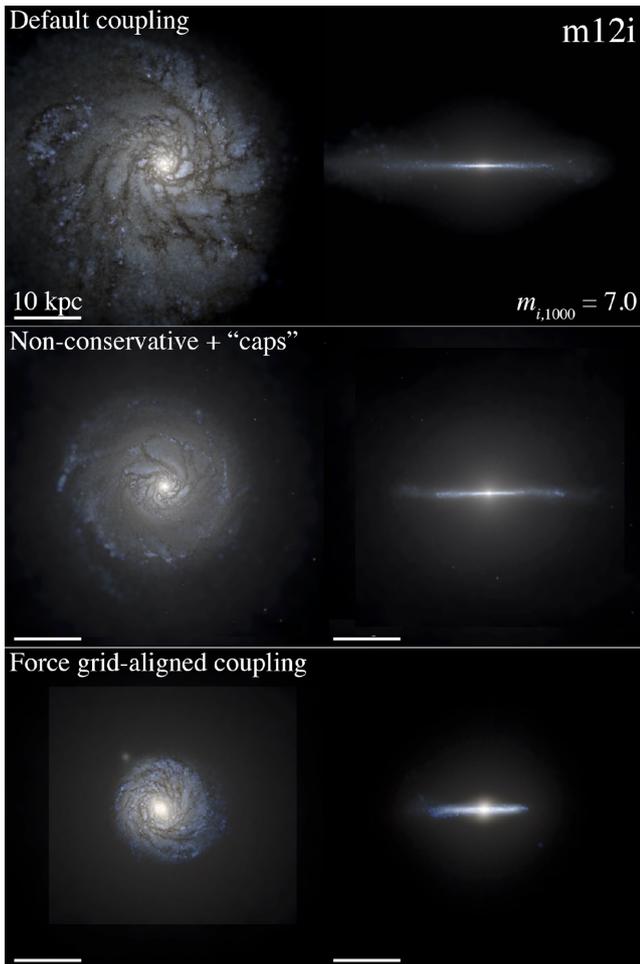


Figure B1. Mock images in HST bands of our **m12i** run at $z = 0$ at our highest resolution ($m_{i,1000} = 7.0$), for the alternative SNe coupling tests in Fig. A1. With the ‘caps’ added to the non-conservative method, the catastrophic errors in Fig. 4 are suppressed and the morphology agrees with our ‘default’ run reasonably well. In the ‘force grid-aligned coupling’ run, the spurious torques from numerically forcing the winds along the coordinate axes (incorrectly) drive the disc into alignment with these axes, removing angular momentum from recycling material and producing a more compact disc.

Specifically, in Figs A1 and B1, we conduct the same tests as Figs 3 and 4, but with a modified non-conservative algorithm (‘non-conservative + caps’). Here we take the non-conservative formulation from Section 3.2 and – for testing purposes only – limit the most serious errors by enforcing an upper limit to the fraction of SNe momentum coupled to any one particle ($= h_b^2/4 [\text{MAX}(h_a, h_b)^2 + |\mathbf{x}_{ba}|^2]$; where $h_a^3 \equiv m_a/\rho_a$) and an upper limit to the maximum velocity change of $\sim 50 \text{ km s}^{-1}$ (per ‘event’).

Figs A1 and B1 clearly demonstrate that it is only the most severe, pathological local coupling cases in the ‘non-conservative’ algorithm which generate the ‘disc destruction’ (as opposed to an integrated sum of small errors). Running this ‘capped’ model at the same resolution, we see a reasonable, clearly thin-disc morphology emerge, in good agreement with our default run. So long as we control (or better yet, eliminate) these errors at a reasonable level, they do not corrupt our solutions. This is why at lower resolution (where the ‘worst-case’ kick magnitude was much smaller, $< 10 \text{ km s}^{-1}$ for a single gas element), as we studied in FIRE-1, we do not see problematic behaviour.

APPENDIX C: PROBLEMS WITH EXPLICITLY-GRID-ALIGNED FEEDBACK COUPLING

In Section 3.2, we discuss the effects of the purely numerical mechanical feedback coupling algorithm. We discussed the importance of algorithms which respect statistical isotropy. Here we compare another algorithm which is not statistically isotropic, for a different reason.

In Fig. A1, we conduct the same tests as Figs 3–4, but we consider a ‘Force grid-aligned coupling’ model. The coupling follows our default algorithm, except we treat the particles around the SN *as if* they were distributed in a perfect Cartesian lattice with the SN at the centre (as if the SN exploded at the exact centre of a cell in a Cartesian grid code), and so enforce the exact same coupling in the $\pm x$, $\pm y$, $\pm z$ coordinate directions. This trivially ensures momentum conservation but is not the correct solution given the actual non-grid distribution of particles. Moreover it imprints the coordinate axes of the simulation directly on to the galaxy – it is a fundamentally non-statistically isotropic algorithm. But this is useful for comparison, because such ‘preferred directions’ are generic to Cartesian grid-based simulations (e.g. AMR) and their SNe coupling schemes.

The ‘grid-aligned’ implementation shows a higher central V_c , especially at our highest resolution ($m_{i,1000} = 7$), owing to a more compact disc. This is evident in Fig. B1, where we compare the $z = 0$ visual morphologies of the MW-mass simulations run at our highest resolution. In the ‘grid-aligned’ implementation (uniquely), the disc is nearly perfectly aligned with the simulation coordinate axes – not surprising given that feedback is forcibly aligned in this case. This artificial alignment generates strong torques on outflowing/recycling material, as well as material within the disc (it must be torqued from its ‘natural’ orientation); as winds recycle and the disc first forms, this in turn produces a significant loss of angular momentum. As a result, the late-time inflowing/recycled material (which forms the disc) has lower angular momentum in this run, and produces a more compact disc, with a much higher central V_c . Note that the error is essentially independent of resolution (whereas the central V_c decreases with resolution, in all other algorithms tested), because the grid alignment is resolution-independent.

We show this to emphasize that this can be a serious worry for fixed-grid or AMR codes, where grid-alignment of discs is a

ubiquitous and well-known problem (Davis 1984), even at extremely high resolution and independent of feedback (because the hydrodynamics themselves are grid-aligned; see e.g. de Val-Borro et al. 2006; Byerly et al. 2014; Hopkins 2015), especially in simulations of cosmological disc formation (see Hahn, Teysseier & Carollo 2010). This may bias these simulations to smaller, more compact galaxies.

Given the highly irregular dSph morphology of **m10q**, there is not an obvious difference in that galaxy with this algorithm (there is no thin disc to torque); we therefore do not show a detailed comparison.

We have re-run haloes **m09** and **m10v** (both dwarfs) and **m12f** and **m12m** (MW-mass) from Paper I with this algorithm to confirm the results are robust across haloes at both dwarf and MW mass scales.

APPENDIX D: DELAYED-COOLING AND TARGET-TEMPERATURE MODELS: TESTS

We briefly discussed ‘delayed-cooling’ and ‘target-temperature’ models in the text in Section 4.4. There we emphasized that such models are fundamentally ill-posed at high resolution. Here we demonstrate this explicitly in zoom-in cosmological simulations at dwarf and MW mass scales.

We compare four simple models, which resemble common implementations in the literature.

(i) *Delayed-cooling (physical)*: Here we take the ‘Fully Thermal’ model from the text (injecting the full 10^{51} erg per SNe as thermal energy), but particles which are heated by the SNe are not allowed to cool for a time Δt_{delay} . Physically, the cooling time of an explosion is (by definition) the time it takes to reach R_{cool} : since it is in an energy-conserving Sedov-phase before this, the time is $t_{\text{cool}}^{\text{shock}} = (2/5) R_{\text{cool}}/v(R_{\text{cool}})$, where $v(R_{\text{cool}}) \equiv p_i/M_{\text{cool}}$ is the velocity at this stage. Given the expression for p_i in the text (equation (29)), this is $t_{\text{cool}}^{\text{shock}} = 5 \times 10^4 \text{ yr } (n/\text{cm}^{-3})^{-4/7} (Z/Z_{\odot})^{-5/14}$ (this is, as it should be, approximately the physical cooling time for metal-enriched gas with the post-shock temperature appropriate for a shock velocity $v(R_{\text{cool}}) \sim 200 \text{ km s}^{-1}$). So we adopt $\Delta t_{\text{delay}} = t_{\text{cool}}^{\text{shock}}$. Note, though, that $t_{\text{cool}}^{\text{shock}}$ is ~ 1000 times shorter than the gas dynamical time ($1/\sqrt{G\rho}$) – so unless the cooling radius (stage of blastwave expansion where the expansion time is shorter than $t_{\text{cool}}^{\text{shock}}$) is resolved, this will do little work.¹⁵

(ii) *Delayed-cooling (300×Physical)*: We take the ‘Delayed-Cooling (Physical)’ model, but arbitrarily multiply the delay time-scale by a factor of 300. This brings it to $\gtrsim 10^7$ yr, comparable to the galaxy dynamical time.

(iii) *Target-temperature (physical)*: We take the ‘Fully Thermal’ model from the text but wish to heat the targeted gas particles as close as possible to some desired target temperature, $T_{\text{target}} \sim 10^{7.5}$ K, without artificially changing the physics. Therefore we adjust the number of neighbour particles on-the-fly as needed to get as close as possible to this goal. However, even putting 10^{51} erg into a single particle can only increase the temperature by $\Delta T \approx 2.4 \times 10^6 (m_i/1000 M_{\odot})^{-1}$ K. So typically this amounts to putting

100 percent of the energy of each SN into the neighbour particle which is closest to (but still below) T_{target} .

(iv) *Target-temperature (store SNe)*: We require that all gas particles heated by an SN receive sufficient energy such that their temperature rises by $T_{\text{target}} = 10^{7.5}$ K. We follow Dalla Vecchia & Schaye (2012)¹⁶ and achieve this by implicitly turning off cooling – we ‘store’ SNe until a sufficient number have accumulated in order to heat a target gas mass by the desired T_{target} . Then all the SNe energy is deposited ‘at once’ in that gas in a thermal-energy dump. To minimize the number of SNe which must be ‘stored,’ we set a target gas mass (for each ‘heating event’) of just 10 gas particles. Given this, the number of SNe which must be ‘stored’ and then injected simultaneously is $N_{\text{SNe}} \sim 10^5 (m_i/10^6 M_{\odot})$; this is physically similar to delaying cooling for ~ 30 Myr (while the SNe accumulate) for a gas particle surrounded by ~ 10 star particles.

Figs D1 and D2 repeat the experiments from Section 4.4 in the main text, for these models. Not surprisingly, at the resolution shown ($m_i \gg 100 M_{\odot}$), the ‘physically motivated’ models (either delayed-cooling or target-temperature) resemble the ‘Fully Thermal’ model from the text, which itself resembled the ‘no SNe’ result. Turning off cooling only for the real cooling time, or heating gas only to the correct physical temperature, ignoring momentum, leads to overcooling at low resolution.

Of course, in this class of models we can simply adjust the model parameters until a reasonable stellar mass is obtained. The ‘Delayed-cooling (300×Physical)’ and ‘Target-temperature (Store SNe)’ models manage to produce order-of-magnitude similar galaxy masses to our converged default model at low resolution. However there are serious issues.

(i) The actual explicit or implicit ‘cooling turnoff times’ are wildly unphysical ($\gtrsim 10$ Myr) – many orders of magnitude larger than physical in both cases (see Agertz & Kravtsov 2015; Martizzi et al. 2015). Thus the solutions we ‘insert’ on large scales do *not* in any way resemble a ‘down-sampled’ high-resolution simulation and nor can the relevant parameters be predicted a priori from higher-resolution simulations. Note that such unphysically long-delayed cooling times are what are actually used in most simulations with these ‘delayed cooling’ models (e.g. Stinson et al. 2006; Shen et al. 2014; Crain et al. 2015).

It has been suggested that these models, while obviously unphysical for a single SN explosion, could represent the result of SNe which are strongly clustered in both space and time. However, all the simulations here, by allowing resolved cooling into GMCs, explicitly resolve stellar clustering (and if anything, we show in Paper I that low resolution tends to overestimate clustering, owing to discrete star-particle sampling). Therefore if such clustering were to occur, one would not need to artificially turn off cooling or store SNe (one could simply allow the explosions to occur rapidly and create a superbubble, as occurs in our default models). In contrast, these models *impose*, rather than predict, a strong and *explicitly resolution-dependent* assumption about clustering: for the target temperature model it is that SNe explode in ‘units’ in both time and space of $\sim 10^5$ SNe ($m_i/10^6 M_{\odot}$).

Walch & Naab (2015); Martizzi et al. (2015); Kim & Ostriker (2017) and most explicitly Kim, Ostriker & Raileanu (2017) have

¹⁵ In our delayed-cooling experiments, we have considered both turning off *all* cooling for a particle, and tracking a separate reservoir of SNe-injected energy, which is not allowed to cool (while other energy can cool). Both give similar results for our comparison here. We also ‘reset’ the delay time Δt_{delay} whenever a new SNe injects energy into a gas element.

¹⁶ The Dalla Vecchia & Schaye (2012) ‘target-temperature’ implementation released the SNe energy stochastically rather than deterministically after a fixed time – we have implemented this as well and the results are identical to the ‘target-temperature (store SNe)’ implementation shown.

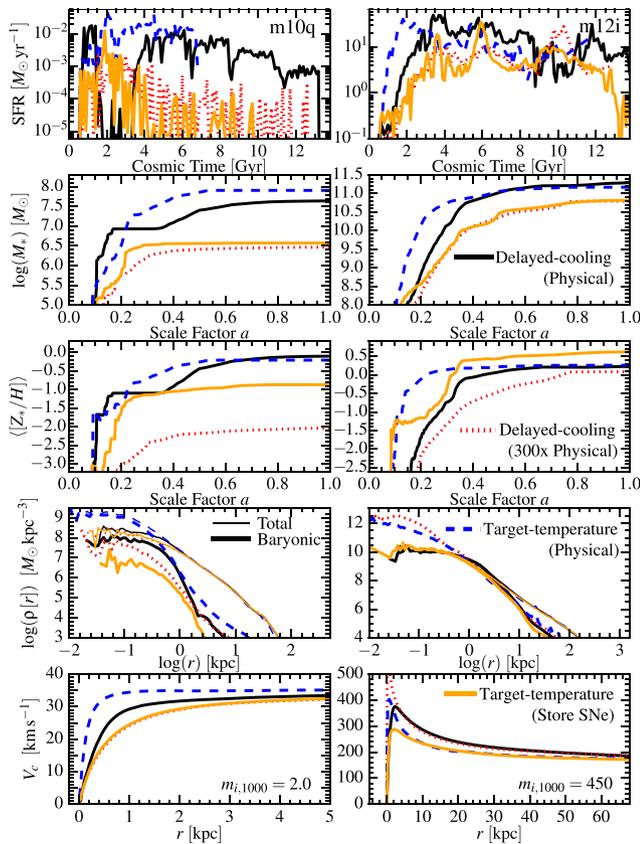


Figure D1. Tests of ‘Delayed-cooling’ or ‘Target-temperature’ models (Section D), as Fig. 3. In these models cooling is artificially ‘turned off’ either explicitly (not allowing SN-heated gas to cool for some Δt_{delay}) or implicitly (by ‘storing’ SNe energy for an arbitrarily long time, until sufficient energy has accumulated to heat gas particles to some target temperature T_{target}). We compare: *Delayed-cooling (Physical)*: Cooling is turned off only for the Δt_{delay} corresponding to the actual blastwave cooling time (time when the blast reaches the cooling radius). If the cooling radius is un-resolved, this time-scale is necessarily much shorter than the dynamical time. *Delayed-cooling (300xPhysical)*: Δt_{delay} is increased by an arbitrary factor of ~ 300 to force it to $\gtrsim 10^7$ yr, comparable to the galaxy dynamical time. *Target-temperature (Physical)*: SNe energy is deposited in such a way (distributed among neighbours) to bring ‘coupled’ neighbours as close as possible to a target temperature $T_{\text{target}} = 10^{7.5}$ K, but without artificially turning off cooling or ‘saving’ SNe after they should explode. *Target-temperature (Store SNe)*: SNe are ‘stored’ until the gas neighbours can all be raised to exactly $T_{\text{target}} = 10^{7.5}$ K (roughly equivalent to a cooling-delay time of $\Delta t_{\text{delay}} \sim 30$ Myr; or forcing 10^5 SNe to explode simultaneously at a single place and time in the **m12i** run). With a physical cooling time or post-SNe temperature, these models are very similar to the ‘Fully Thermal’ model in the text, and produce severe overcooling at low resolution (as expected). By making the effective cooling-delay or ‘storage’ time very large, cooling becomes inefficient and galactic outflows are driven; however this requires un-physical values that do not resemble the solution for resolved explosions. As a result, the star formation histories, rotation curves, metal abundance distributions, and gas phase structure do not resemble the converged solutions in the main text.

demonstrated this in greater detail, in studies of idealized single-SN explosions or clustered SNe in a sub-volume of the ISM. There, these authors demonstrate more explicitly that delayed cooling or target-temperature models are not a good approximation to the ‘down-sampled’ results of high-resolution simulations.

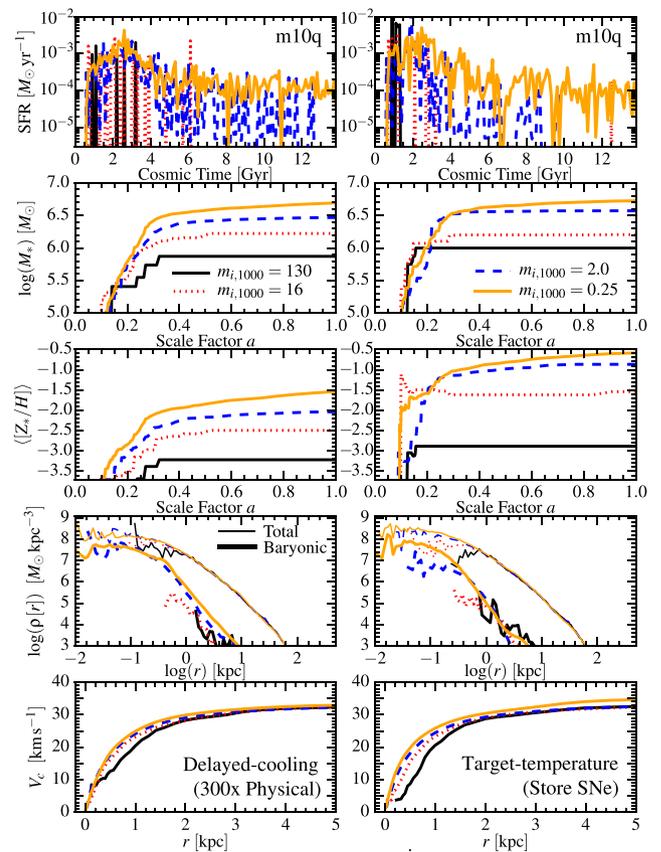


Figure D2. Comparison of the ‘Delayed-cooling (300× Physical)’ (left) and ‘Target-temperature (Store SNe)’ (right) models from Fig. D1, as a function of resolution (in **m10q**). Neither exhibits convergence down to mass resolution $m_i = 250 M_{\odot}$: stellar masses, metallicities, central galaxy (and dark matter) densities, rotation curves in the central \sim kpc, and late-time star formation rates all increase systematically with resolution. Because the models do not physically correspond to any stage of a physical SNe explosion, they in fact *should not* converge to the correct solution even at infinite resolution (for example, even if the cooling radius is resolved, the ‘Delayed-cooling (300×Physical)’ model will necessarily prevent it from cooling for ~ 300 times longer than it should). For this reason neither model is well-suited to the convergence studies in the main text, and they should not be used in simulations which can resolve superbubble overlap and physical cooling radii.

(ii) As also noted by Agertz & Kravtsov (2015), this un-physical feedback coupling produces several artefacts in the galaxy properties. (1) Shapes of the star formation histories are biased: in dwarfs the star formation in both cases is much more concentrated at early times, compared to our converged solutions in Fig. 9. (2) In massive galaxies, the ‘delayed cooling’ model accumulates a massive reservoir of gas (with its cooling turned off by successive generations of SNe) at the galaxy centre, which finally (because of the dependence of t_{delay} on density and metallicity) achieves a short cooling time even with the imposed ‘delay’, then forms a strong starburst (at time ~ 10 Gyr) and leaves an extremely compact bulge (the $\sim 500 \text{ km s}^{-1}$ rotation-curve peak). (3) The metal abundances are highly sensitive to the ‘delayed cooling’ and ‘target temperature’ model implementations, and vary by several orders of magnitude in the variants explored here. The metallicities for dwarfs are extremely high in the ‘target temperature’ models, because the ‘stored’ SNe inject a

huge metal mass simultaneously,¹⁷ which is ejected from the galaxy but is so metal-rich that it re-cools and preferentially forms the next generation of stars. We have verified this feature remains regardless of whether we include or exclude explicit ‘turbulent metal mixing’ (numerical metal diffusion) terms as described in Paper I. (4) The gas phase structure is quite different from our converged solutions in the text. Since these models rely only on hot gas, there is little or no cool ($\sim 10^4$ – 10^5 K) or cold ($\lesssim 10^4$ K) gas in the outflows here, unlike our default simulations (see Muratov et al. 2015, 2017; Faucher-Giguere et al. 2015; Faucher-Giguère et al. 2016; Anglés-Alcázar et al. 2017), although Rosdahl et al. (2017) show that alternative ‘delayed cooling’ implementations err in the opposite manner and produce far too much cold, dense gas in outflows.

(iii) The solutions are non-convergent. Fig. D2 shows this explicitly, re-running **m10q** at resolution from 250 to $10^5 M_\odot$, with the ‘Delayed-cooling (300×Physical)’ and ‘Target-temperature (Store SNe)’ models. In both, the stellar masses, metallicities, central galaxy (and dark matter) densities, rotation curves in the central \sim kpc, and late-time star formation rates all increase systematically as the resolution increases.

Of course, this owes to the explicit resolution dependence of the assumed clustering and blastwave structure of SNe. In target-temperature models, the SNe cluster and are synchronized in time and space in an explicitly particle-mass dependent manner. In delayed-cooling models, the ‘cooling mass’ M_{cool} is essentially defined to be the mass of the kernel over which the SNe are distributed (some multiple of the particle mass): since the terminal momentum for an energy-conserving blast (which this is forced to be, by not allowing cooling) is $p_f \sim \sqrt{E_{\text{SNe}} M_{\text{cool}}}$, the momentum injected increases $\propto M_{\text{cool}}^{1/2} \propto m_i^{1/2}$, so feedback becomes more efficient at lower resolution (analogous to the ‘Fully Kinetic’ models discussed in the text).

Interestingly, while the lack of convergence for delayed-cooling models is ‘smooth’, the ‘target temperature’ models exhibit false convergence in some properties (such as stellar mass) at low resolution, then ‘jump’ in the predicted values once a critical mass resolution (here $\sim 2000 M_\odot$) is reached. That is, of course, the mass resolution where the *physical* cooling radii of SNe begin to be resolved: so the fundamental meaning and behaviour of the sub-grid model changes. At even higher resolution, the ‘target temperature’ of $\sim 10^{7.5}$ K would actually become *lower* than the correct, resolved blastwave temperatures: this would lead one to ‘store’ < 1 SN at a time. Clearly, in this limit the ‘delayed cooling’ and ‘target temperature’ models simply become ill-defined.

APPENDIX E: ENERGY-CONSERVING SOLUTIONS ACCOUNTING FOR ARBITRARY STAR-GAS MOTIONS

In the text, we noted that, for a spherically symmetric blastwave propagating into a medium initially at rest, converting an energy $E_{\text{ej}} = (1/2)m_{\text{ej}}v_{\text{ej}}^2$ into kinetic energy (pure radial momentum), after coupling to a total mass m_b , simply implies a final kinetic energy $p_{\text{final}}^2/(2(m_b + m_{\text{ej}})) = E_{\text{ej}}$, giving $p_{\text{final}} = (1 + m_b/m_{\text{ej}})^{1/2} m_{\text{ej}}(2E_{\text{ej}}/m_{\text{ej}})^{1/2} = (1 + m_b/m_{\text{ej}})^{1/2} p_{\text{ej}}$, where $p_{\text{ej}} = m_{\text{ej}}v_{\text{ej}}$ is the initial ejecta momentum. The situation

¹⁷ If we ‘store up’ SNe each with $\sim 10^{51}$ erg until we can heat a discrete mass Δm to a temperature $\sim 10^{7.5}$ K, then if each SN deposits $\sim 2 M_\odot$ worth of metals, the mass Δm will be immediately enriched to metallicity $Z \approx 2 Z_\odot$!

is more complex if we allow for arbitrary initial gas and stellar velocities.

First recall the mass conservation condition $\sum \Delta m_b = m_{\text{ej}}$ is un-altered by star or gas motion. The momentum condition is, in the rest frame of the star, $\sum \Delta \mathbf{p}_b = \mathbf{0}$, which in the lab/simulation frame becomes $\sum \Delta \mathbf{p}'_b = m_{\text{ej}} \mathbf{v}_a$, trivially satisfied by the boost $\Delta \mathbf{p}'_b = \Delta \mathbf{p}_b + \Delta m_b \mathbf{v}_a$. In these two cases, no net mass or linear momentum is created/destroyed. For energy, we must account for the energy injected. Consider a hypothetical instant ‘just after’ explosion, but ‘before’ coupling. Then the mass of the star particle is $m_a - m_{\text{ej}}$, moving at \mathbf{v}_a . Gas neighbours b have their ‘unperturbed’ mass and velocities, etc. In the rest-frame of the star, the ejecta contain the energy $E_{\text{ej}} = (1/2)m_{\text{ej}}v_{\text{ej}}^2$. Assume the ejecta have negligible initial internal energy, then v_{ej} is the real radial velocity. If the ejecta are isotropic in the rest frame, each parcel in some solid angle $d\Omega$ carries mass $dm = m_{\text{ej}}/(4\pi) d\Omega$, with velocity $\mathbf{v}_{\text{ej}} = v_{\text{ej}} \hat{r}$ (where \hat{r} points from the star radially outward). If the star is moving initially at velocity \mathbf{v}_a , the whole system is boosted, and $\mathbf{v}'_{\text{ej}} = v_{\text{ej}} \hat{r} + \mathbf{v}_a$. To calculate $E'_{\text{ej}} = (1/2) \int |\mathbf{v}'_{\text{ej}}|^2 dm$ in the lab frame, note $|\mathbf{v}'_{\text{ej}}|^2 = v_{\text{ej}}^2 + 2 \mathbf{v}_{\text{ej}} \cdot \mathbf{v}_a + v_a^2 = v_{\text{ej}}^2 + v_a^2 + 2 v_{\text{ej}} v_a \cos \theta_{ea}$ (where we can define standard spherical coordinates such that $\hat{r} \cdot \hat{\mathbf{v}}_a \equiv \cos \theta_{ea}$). Using $dm = m_{\text{ej}}/(4\pi) d\Omega = m_{\text{ej}}/(4\pi) d\phi d\cos \theta_{ea}$, we trivially obtain that the ‘cross-term’ $\mathbf{v}_{\text{ej}} \cdot \mathbf{v}_a$ vanishes (integrating over all ejecta), so we have $E'_{\text{ej}} = (m_{\text{ej}}/2)(v_{\text{ej}}^2 + v_a^2)$.

Now assume we couple some energy and momentum to all the neighbours b – the exact *discrete* energy conservation condition must be satisfied, summed over all elements which receive some ejecta mass/energy/momentum. In the simulation/lab frame, the energy conservation condition can be written:

$$\begin{aligned} E_{\text{initial}} + E'_{\text{ej}} &= \frac{(m_a - m_{\text{ej}})}{2} v_a^2 + \sum_b \frac{\mathbf{p}_b^2}{2m_b} + \frac{m_{\text{ej}}}{2} (v_a^2 + v_{\text{ej}}^2) + U_0 \\ &= \frac{m_a}{2} v_a^2 + \sum_b \frac{\mathbf{p}_b^2}{2m_b} + \frac{m_{\text{ej}}}{2} v_{\text{ej}}^2 + U_0 \\ &= E_{\text{final}} = \frac{(m_a - m_{\text{ej}})}{2} v_a^2 + \sum_b \frac{|\mathbf{p}_b + \Delta \mathbf{p}'_b|^2}{2(m_b + \Delta m_b)} + U_f \end{aligned} \quad (\text{E1})$$

where (as in the main text) we use x_b to denote the pre-coupling value of x . Here U_0 and $U_f \equiv U_0 + \Delta U$ collect the non-kinetic energy terms (e.g. thermal energy, discussed below).

Now using the fact that $\Delta \mathbf{p}'_b \equiv \Delta \mathbf{p}_b + \Delta m_b \mathbf{v}_a$, we can write this in terms of the coupled momenta $\Delta \mathbf{p}_b$. Use the identities $\mathbf{v}_b = \mathbf{p}_b/m_b$, $\sum \Delta m_b = m_{\text{ej}}$, $\sum \Delta \mathbf{p}_b = \mathbf{0}$, $\mu_b \equiv \Delta m_b/m_b$, $1/(1+x) = 1-x/(1+x)$, and following through with some tedious algebra, we can re-arrange terms and write the energy conservation condition as:

$$\epsilon = \sum_b \left[\frac{|\Delta \mathbf{p}_b|^2}{2m_b(1+\mu_b)} + \frac{\mathbf{v}_{ba} \cdot \Delta \mathbf{p}_b}{(1+\mu_b)} \right] \quad (\text{E2})$$

where $\mathbf{v}_{ba} \equiv \mathbf{v}_b - \mathbf{v}_a$ and the energy ϵ is defined by

$$\epsilon \equiv \frac{1}{2} m_{\text{ej}} \left(v_{\text{ej}}^2 + \sum_b w'_b |\mathbf{v}_{ba}|^2 \right) - \Delta U \quad (\text{E3})$$

$$w'_b \equiv \frac{1}{1+\mu_b} \left(\frac{\Delta m_b}{m_{\text{ej}}} \right). \quad (\text{E4})$$

This makes it clear that the dynamics depend only on the *relative* velocity \mathbf{v}_{ba} of gas relative to the star (i.e. a uniform boost will not change the dynamics, as it should not). In ϵ , the term in \mathbf{v}_{ba}^2 reflects the additional energy generated by relative gas-star motion

– since $\sum w'_b \approx 1$, this is negligible for SNe where $\mathbf{v}_{ba}^2 \ll v_{ej}^2$, but potentially important for slow winds.

Now without loss of generality, define the coupled momentum $\Delta \mathbf{p}_b$ as the value we used in the text (for the case where the gas is not moving relative to the star) multiplied by an arbitrary function ψ_b :

$$\Delta \mathbf{p}_b \equiv \psi_b \Delta m_b \left(1 + \frac{m_b}{\Delta m_b}\right)^{1/2} \left(\frac{2\epsilon}{m_{ej}}\right)^{1/2} \Delta \hat{\mathbf{p}}_b \quad (\text{E5})$$

where $\Delta \hat{\mathbf{p}}_b \equiv \Delta \mathbf{p}_b / |\Delta \mathbf{p}_b|$. Inserting this into the energy conservation condition in equation (E2), we obtain the constraint equation in terms of ψ :

$$1 = \sum_b \psi_b^2 \frac{\Delta m_b}{m_{ej}} + 2 \sum_b \psi_b \cos \theta_{ba} \left(\frac{w'_b m_b |\mathbf{v}_{ba}|^2}{2\epsilon}\right)^{1/2} \quad (\text{E6})$$

$$\cos \theta_{ba} \equiv \hat{\mathbf{v}}_{ba} \cdot \Delta \hat{\mathbf{p}}_b = \frac{\mathbf{v}_{ba} \cdot \Delta \mathbf{p}_b}{|\mathbf{v}_{ba}| |\Delta \mathbf{p}_b|} \quad (\text{E7})$$

If $\mathbf{v}_{ba} = 0$ (no initial gas-star motion), then the term in $\cos \theta_{ba}$ vanishes, and this is trivially solved for $\psi_b = 1$, and equation (E5) reduces to our solution for a spherically symmetric explosion in a stationary medium (as it should). More generally, any ψ_b and $\Delta \hat{\mathbf{p}}_b$ must still produce $\sum \Delta \mathbf{p}_b = \mathbf{0}$. If we have defined a set of vector weights, as in the main text, such that this is true for the $\mathbf{v}_{ba} = 0$ (stationary) case, then the simplest choice which guarantees $\sum \Delta \mathbf{p}_b = \mathbf{0}$ is preserved is to take $\psi_b = \psi$, such that $\sum \Delta \mathbf{p}_b \rightarrow \psi \sum \Delta \mathbf{p}_b^{\text{stationary}} = \mathbf{0}$. Of course, in detail the true solution for a blastwave in an inhomogeneous medium with locally varying velocities could feature variable ψ_b or changes in the direction $\Delta \hat{\mathbf{p}}_b$ (i.e. work being done in different directions from the initial ejecta expansion), but by definition this sub-structure is un-resolved at the coupling radius so we should think of ψ as an average over the un-resolved structure. The solution to equation (E6) is then simply:

$$\psi_b = \psi \equiv \sqrt{1 + \beta_\psi^2} - \beta_\psi \quad (\text{E8})$$

$$\beta_\psi \equiv \sum_b \cos \theta_{ba} \left(\frac{w'_b m_b |\mathbf{v}_{ba}|^2}{2\epsilon}\right)^{1/2} \quad (\text{E9})$$

This gives us the desired expression in the strictly energy-conserving limit. But at low resolution (large particle masses) the solution is not energy-conserving (the blastwave has reached the terminal momentum and radiated energy away). Following the text, return to equation (E2) and insert the terminal momentum $\Delta \mathbf{p}_b = \Delta \mathbf{p}_b^{\text{terminal}} = \phi_b \Delta \mathbf{p}_b^{\text{terminal}}(\mathbf{v}_{ba} = 0) = \phi_b (p_t / p_{ej}) \Delta \mathbf{p}_b^{\text{initial}} = \phi_b (p_t / p_{ej}) \Delta m_b (2\epsilon / m_{ej})^{1/2} \Delta \hat{\mathbf{p}}_b$, where ϕ_b is an arbitrary constant analogous to ψ_b . Following the solution above, we will take $\phi_b \rightarrow \phi$. Since some energy has been radiated, the right-hand side of equation (E2) must be $< \epsilon$ (the constraint is an inequality). This is solved by

$$\phi = \text{MIN} \left[1, \alpha_\phi^{-1} \left(\sqrt{\beta_\phi^2 + \alpha_\phi} - \beta_\phi \right) \right] \quad (\text{E10})$$

$$\alpha_\phi \equiv \sum_b w'_b \frac{\Delta m_b}{m_{ej}} \frac{p_t}{m_b v_t} \quad (\text{E11})$$

$$\beta_\phi \equiv \sum_b w'_b \cos \theta_{ba} \frac{|\mathbf{v}_{ba}|}{v_t}, \quad (\text{E12})$$

where $v_t \equiv 2\epsilon / p_t$ (approximately the velocity at which the blastwave becomes radiative). In the limit where the terminal momentum

is reached, $p_t \ll m_b v_t$ by definition, so α_ϕ is vanishingly small and $\phi \approx \text{MIN}[1, 1/2\beta_\phi]$ (for $\beta_\phi > 0$). This has a simple interpretation then: β_ϕ is just the (kernel-averaged) ratio of the net outward gas velocity from the SN to the velocity where the blastwave becomes radiative. If the recession velocity exceeds $v_t \sim 200 \text{ km s}^{-1}$ (the velocity at which the terminal momentum is reached, for a stationary surrounding medium), the SN must reach terminal momentum earlier (at a higher velocity therefore lower terminal momentum) before the ambient medium ‘outruns’ the blastwave: mathematically $\beta_\phi \gtrsim 1$ and $\phi \lesssim 1$, accordingly.

Having computed ψ and ϕ , we can then decide which limit (the energy-conserving or terminal-momentum limit) a gas element should be in, as in equations 31 and (32), by comparing the corrected $\Delta \mathbf{p}_b^{\text{energy-conserving}}$ (equation (E5)) and corrected $\Delta \mathbf{p}_b^{\text{terminal}}$ above. If $|\Delta \mathbf{p}_b^{\text{energy-conserving}}| > |\Delta \mathbf{p}_b^{\text{terminal}}|$ (the momentum implied by the energy-conserving limit exceeds the terminal momentum), or $m_b > m_{\text{cool}}^b \equiv |\Delta \mathbf{p}_b^{\text{terminal}}| / v_t$ (the mass of particle b exceeds the swept-up mass at which the energy-conserving solution would decelerate to below $v_t = 2\epsilon / p_t$), then the terminal solution is applied (otherwise the energy-conserving solution is applied). If the gas-particle motion is negligible ($\psi \approx \phi \approx 1$), these two conditions are *exactly* equivalent; more generally we need to check both.

Physically, note that for $\mathbf{v}_{ba} = 0$ (non-moving cases or uniformly boosting the whole simulation), $\beta_\psi = \beta_\phi = 0$ so $\psi = \phi = 1$ and we obtain the stationary case as expected. For $\mathbf{v}_{ba} = \text{constant}$ (assuming $\Delta m_b \ll m_b$), the condition $\sum \Delta \mathbf{p}_b = \mathbf{0}$ becomes mathematically identical to $\beta_\psi = \beta_\phi = 0$, so there is no change in the coupled momentum or energy relative to what would occur in the stationary case (this is easiest to see by returning to equation (E2) and simply taking the \mathbf{v}_{ba} term outside the sum). In a medium moving with uniform velocity relative to the star, the blastwave produce a stronger, slower-moving shock in the ‘upwind’ direction and weaker, faster-moving shock in the ‘downwind’ direction, but it is easy to verify that the difference in the energies produced in both directions cancel one another (and of course, the momentum imparted in both directions must, by conservation, be equal). In a turbulent medium, different velocities will tend to cancel, so β_ψ, ϕ will both be small.

However, when there is a large net inflow/outflow motion around the star, β_ψ, ϕ can be non-negligible. Consider a spherically symmetric case with $\mathbf{v}_{ba} = v_r \hat{\mathbf{r}}$ so $\cos \theta_{ba} = 1$ if $v_r > 0$, or $\cos \theta_{ba} = -1$ if $v_r < 0$, and assume the total mass to which the ejecta is coupled is distributed in a shell with mass M_{coupled} . Then $|\beta_\psi| \sim (M_{\text{coupled}} v_r^2 / 2\epsilon)^{1/2} \sim (K E_{\text{initial}} / K E_{\text{ejecta}})^{1/2}$, i.e. β_ψ^2 scales with the ratio of the initial (pre-coupling) kinetic energy of the surrounding gas elements (across which the ejecta are distributed) to the ejecta energy. Although $|\mathbf{v}_{ba}| = |v_r| \ll v_{ej}$ for SNe, the kinetic energy is weighted by the particle mass, so it is *not* necessarily negligible: $|\beta_\psi| \gtrsim 1$ if the typical $|v_r| \gtrsim (m_{\text{ejecta}} / M_{\text{coupled}})^{1/2} v_{ej} \sim 350 \text{ km s}^{-1} (m_i / 100 M_\odot)^{-1/2}$ for typical core-collapse SNe. At sufficiently high resolution, then, this term becomes negligible (the gas velocities are never so coherently large), but at low resolution it can be important. Of course, at low resolution, the cooling radius is un-resolved and we should use the terminal momentum expression, where $\beta_\phi \gtrsim 1$ requires $|v_r| \gtrsim v_t \sim 200 \text{ km s}^{-1}$ – the expression becomes resolution-independent in this limit (once m_i exceeds a few hundred solar masses). In either case, if such large v_r is reached, $\psi \approx 1/2\beta_\psi$ (or $\phi \approx 1/2\beta_\phi$) becomes < 1 . This comes from the $\mathbf{p}_b \cdot \Delta \mathbf{p}_b$ term, which dominates over $\Delta \mathbf{p}_b^2$ in this limit – physically, it requires more energy to accelerate a shell which is already moving rapidly away from the origin. Conversely, when $|v_r|$ is large and $\beta_\psi < 0$, $\psi \approx 2|\beta_\psi| \gg 1$, i.e. this implies a larger momentum injection, with the energy for the additional PdV

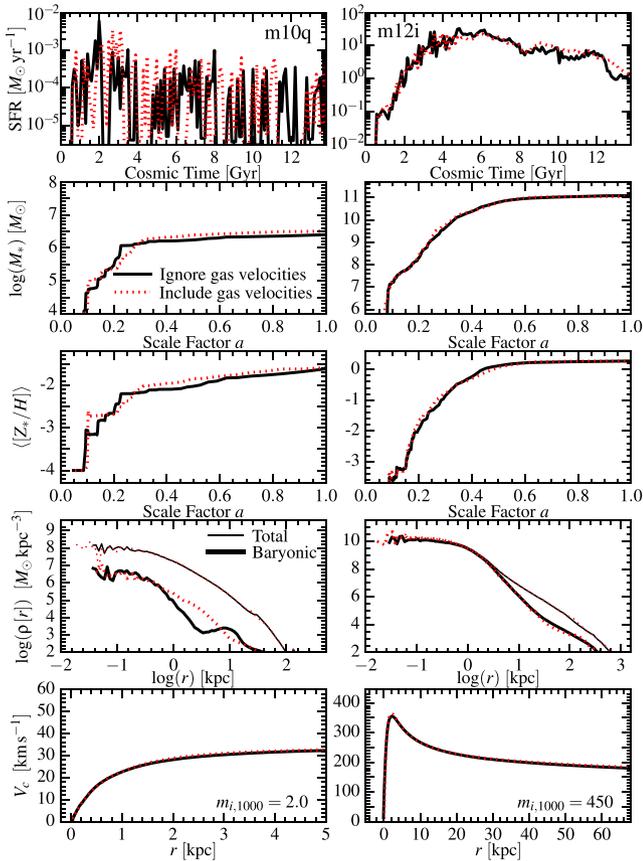


Figure E1. Comparison of our ‘default’ feedback implementation from the text (as Fig. 3), with and without accounting for the relative gas-star motions as described in Section E. The ‘Ignore gas velocities’ implementation treats the momentum injection by SNe as if the ambient medium is at rest (i.e. $\psi = \phi = 1$, in Section E); so the kinetic energy input is variable. The ‘Include gas velocities’ implementation treats the energy injection as fixed and rescales the injected momentum if there are large gas-star velocities. The latter can have significant effects on the coupled momentum at sufficiently low resolution, if the gas is uniformly approaching or receding from SN locations at high velocities ($\gg 200 \text{ km s}^{-1}$). Both implementations become identical at sufficiently high resolution. Even at low resolution (shown here), the differences in galaxy properties are extremely small.

work coming from the shocked external medium falling on to the shock.

If we choose to keep our simple momentum scaling from the main text (setting $\psi = \phi = 1$ always) – i.e. assume that the momentum scaling of SNe is robust across variations in the surrounding gas velocity field – then this necessarily means the kinetic energy coupled by SNe varies, with larger kinetic energy coupled in cases with a net ‘outflow’ motion around the star, and smaller kinetic energy change in cases with net ‘inflow’ motion around the star. This is not necessarily unphysical – it depends, to some extent, on whether the more robust property of SNe blastwaves in a non-

uniform flow is their kinetic energy or their momentum. Similarly, it is quite possible that the general scaling for the terminal momentum p_t from the text could have a complicated dependence on the detailed structure of the velocity field, although simulations in turbulent media discussed in Section 2.3.1 suggest that, on average, broadly similar results are obtained as in simulations where the background is stationary. Clearly, future work is warranted to explore these conditions in more detail.

In practice we find that whether we include this more detailed correction, or set $\psi = \phi = 1$, almost always has a small effect on galaxy properties at all resolution levels: some examples are shown in Fig. E1. Galaxy masses, star formation histories, mass profiles, visual morphologies, metal abundance distribution functions, rotation curves, CGM gas content, and mean outflow rates are essentially unchanged (with at most a systematic ~ 0.1 dex shift in the masses of very low-mass dwarfs, and smaller effects in higher-mass galaxies). We have specifically tested this in the **m10q** and **m12i** galaxies in this paper as well as galaxies **m10v**, **m11q**, **m12f**, and **m12m** from Paper I; we have compared all properties discussed in this paper and in Paper I. The fact that these corrections produce such small effects owes to the fact that coherent, large inflow/outflow velocities around star particles are rare and, even when they occur, tend to average out over time and space. Even in the worst-possible-case (maximal β_ψ) scenario, namely violent post-starburst outflow episodes around dwarf galaxies at low resolution, where most of the ISM of the galaxy is evacuated, the net change in kinetic energy of the gas setting $\psi = \phi = 1$ only differs from the kinetic energy coupled with the exact formulation here by a factor ~ 2 . And, critically, the difference between methods vanishes ($\beta_\psi, \beta_\phi \rightarrow 0$) at sufficiently high resolution.

APPENDIX F: DETAILS OF UNRESOLVED COOLING DO NOT INFLUENCE PREDICTIONS OF OUR DEFAULT MODEL

As noted in Section 4.2, we have verified in a number of tests that, within the context of our default FIRE sub-grid model, the details of how we treat the ‘unresolved cooling phase’ when the simulation does not resolve the local cooling radius are secondary, so long as the correct momentum is coupled to the gas. Fig. F1 shows this explicitly for both **m10q** and **m12i** simulations. In this figure we compare a model where we take our standard sub-grid coupling (the momentum, mass, and metals are unchanged) but always couple the ‘full’ total energy – we do not assume (as in our default model) that the residual thermal energy has been radiated away when the cooling radius is unresolved. As expected, this produces nearly identical results to our default model – in this limit, *by definition*, the cooling time is shorter than the dynamical time at the radius where the energy is deposited. So the code simply radiates away the energy in the next few time-steps, without doing significant work. This is a non-trivial statement, however, in that it clearly shows that in this regime, the *momentum* coupled, not the thermal energy, is the important physical ingredient.

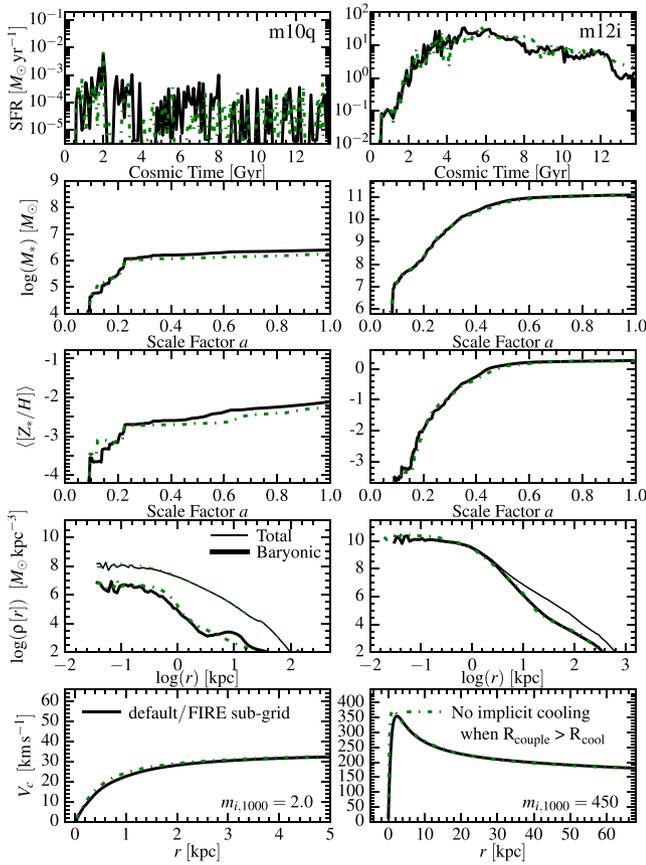


Figure F1. Comparison of the default ‘FIRE sub-grid’ model (as Fig. 7) to a variant implementation in which we do not assume the residual thermal energy from the initial energy-conserving phase has been radiated away if the cooling radius is not resolved (i.e. coupling radius $R_{\text{couple}} > R_{\text{cool}}$). Instead we simply couple it explicitly, using the same kinetic-thermal solution as the default model, and rely on the code cooling physics to radiate it explicitly. This has no effect because, by definition, in this regime the code will radiate the thermal energy away rapidly in either case.

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