ALMA MATER STUDIORUM · UNIVERSITÀ DI BOLOGNA

Department of Physics and Astronomy - DIFA

SECOND CYCLE DEGREE

REGULATING THE COOLING IN THE GALACTIC CORONA WITH DIFFERENT HEATING MODELS IN SIMULATIONS OF MILKY WAY-LIKE GALAXIES

Defended by:

Francesco Loparco

Supervisor:

Prof. Federico Marinacci

Co-supervisor:

Dott. Filippo Barbani

Graduation Session 07/2026

Academic Year 2025/2026

Abstract

Star-forming galaxies like the Milky Way would exhaust their cold, star-forming gas within a few billion years if it were not continuously replenished. The hot circumgalactic medium (CGM) – the so-called corona – surrounding these galaxies is the natural reservoir for this resupply, but in idealised simulations it cools too efficiently, accreting onto the disc and driving star formation above observed levels. Some form of heating is thought to regulate accretion from corona, yet how it should be modelled and how sensitive the properties of the hot CGM are to the way energy is injected, remains an open question.

This thesis investigates how different prescriptions for heating the corona of an isolated, Milky-Way-like galaxy affect its hot gas, its thermodynamic state, and the gas exchange between corona and disc. Within a single controlled setup, we compare prescriptions that deliver energy in different ways and assess whether the response of the corona influences the gas accretion onto the star-forming disc.

Using the moving-mesh code AREPO with the SMUGGLE model for the ISM and stellar feedback, we evolve isolated galaxy simulations sharing identical initial conditions and differing only in the coronal heating: a baseline no heating simulation (No-Heat), two prescriptions compensating the local radiative losses on different timescales (Heat 1 and Heat 2), and a radio-mode Active Galactic Nuclei (AGN) model injecting energy through discrete bubbles, with the budget set from the halo X-ray luminosity via a luminosity–temperature relation (NF Bubbles). The runs are compared at matched times through the phase-resolved gas mass, the coronal density and temperature profiles, the density–temperature plane, the star formation history, and gaseous inflows (outflows) onto (from) the disc.

In all models the corona eventually cools and feeds the disc, with the main

differences lying in how effectively the various heating prescriptions slow this process. The distributed heating models provide measurable thermal support, with the accumulated-injection variant (Heat 2) affecting the innermost regions of the corona, while the continuous model (Heat 1) remains close to the baseline scenario. In contrast, the adopted bubble-heating model injects power at larger radii than the previous two heating models, having an effect on the distribution of the warm+hot gas in the outer corona. Despite the different evolution of the coronal gas, the star formation rate and the disc inflow and outflow rates change only marginally across the different models over the duration of the simulations.

These results highlight the importance of the AGN feedback implementation in regulating the thermal evolution of the CGM. Future work will explore larger energy budgets for the bubble-heating model and replace the current scaling-relation prescription with self-consistent accretion onto a live black-hole particle, allowing the feedback energy to emerge naturally from the simulated accretion history rather than being imposed externally.

Contents

Chapter 1

Introduction

Galaxies like the Milky Way have been forming stars at a nearly steady pace for most of cosmic history (Haywood et al., 2013; Bland-Hawthorn & Gerhard, 2016). Reproducing this environment in a (self-consistent) numerical model of galactic evolution involves some considerations: the cold gas that fuels star formation is consumed within a few billion years, far shorter than the age of the disc. The galaxy must therefore be continuously resupplied with gas from an external reservoir, which is the hot circumgalactic medium (CGM) that surrounds the disc. The difficulty is to keep a balance between the velocity of resupply high enough to sustain star formation over several Gyrs, but not so fast as to overproduce stars. In idealised simulations an isolated hot CGM evolved on its own, rather than within a full cosmological context, tends to cool faster than the stellar feedback alone can counteract , particularly in massive haloes, supplying the disc with more fuel than is observed (Su et al., 2019).

This thesis studies that difficult balance in an isolated Milky Way-like galaxy and examines how an explicit heating of the hot CGM, added on top of stellar feedback, affects the gas content of this region but also the gas supply to the disc.

1.1 The gas depletion problem in Milky Way-like galaxies

Observations of the Milky Way and of similar local spiral galaxies show a sustained star formation rate (SFR), currently estimated at about $1 - 3\,\mathrm{M}_{\odot}\,\mathrm{yr}^{-1}$ (Kennicutt, 1998; Bland-Hawthorn & Gerhard, 2016). Reconstructions of the Galactic star formation history, based on stellar surveys and chemical evolution models, indicate that the disc has kept forming stars without strong quenching over the last ~ 10 Gyr, at a rate that has remained of order a few $\mathrm{M_\odot\,yr^{-1}}$ (Haywood et al., 2013; Snaith et al., 2015).

The total stellar mass built up in the disc is simply the sum of all the stars ever formed, i.e. the SFR integrated over the age of the galaxy

$$M_* = \int_0^{t_{\rm age}} \mathrm{SFR}(t)\ \mathrm{d}t. \tag{1.1}$$

A SFR of a few $\mathrm{M_\odot\,yr^{-1}}$ sustained over $t_{\rm age} \sim 10$ Gyr therefore accumulates a stellar mass of order

$$M_* \approx 3 \times 10^{10}\,\mathrm{M_\odot}, \tag{1.2}$$

of the same order as the observed stellar mass of the Milky Way disc, $M_{*,\rm disc} \approx 5 \times 10^{10}\,\mathrm{M_\odot}$ (Licquia & Newman, 2015; Bland-Hawthorn & Gerhard, 2016).

A discrepancy appears as soon as this assembly is compared with the amount of cold, star-forming gas available in the disc today. The current mass of the cold interstellar medium (ISM), including both molecular (H_2) and neutral atomic (H) gas, is of order $M_{\rm gas} \sim 3 \times 10^9\,\mathrm{M_\odot}$. If the galaxy were a closed box that acquires no new material, the gas depletion timescale would be the time needed to consume this reservoir at the present rate

$$t_{\rm dep} = \frac{M_{\rm gas}}{\mathrm{SFR}}. \tag{1.3}$$

For representative Milky Way values ($M_{\rm gas} \sim 3 \times 10^9\,\mathrm{M_\odot}$, $\mathrm{SFR} \sim 1.5 - 3\,\mathrm{M_\odot\,yr^{-1}}$), Equation 1.3 gives $t_{\rm dep} \sim 1 - 2$ Gyr. This is roughly an order of magnitude shorter than the ~ 10 Gyr over which the disc has been forming stars. The galaxy therefore cannot be a closed system: to keep forming stars at an almost constant rate over this lifetime, it must be continuously resupplied with gas from an external reservoir.

1.2 The multiphase circumgalactic medium and the cooling problem

The natural source of this fuel is the CGM, an extended, multiphase gas reservoir that surrounds the disc out to the virial radius of the dark matter halo hosting the galaxy (Tumlinson et al., 2017; Faucher-Giguère & Oh, 2023). The CGM contains cold and warm clumps, but in a Milky Way-mass halo ($M_{\rm vir} \sim 10^{12}\,{\rm M}_\odot$) most of its mass is in an extended hot corona(Tumlinson et al., 2017; Faucher-Giguère & Oh, 2023). This gas was shock-heated during the gravitational collapse of the halo, to roughly the virial temperature

$$T_{\rm vir} \approx \frac{\mu m_p v_{\rm vir}^2}{2k_B} \sim 10^5 - 10^6\,{\rm K}, \tag{1.4}$$

where μ is the mean molecular weight, m_p the proton mass, k_B the Boltzmann constant, and $v_{\rm vir}$ the virial velocity, the characteristic speed of gas used to derive the halo potential well. For a halo of $M_{\rm vir} \sim 10^{12}\,{\rm M}_\odot$ the virial velocity is $v_{\rm vir} \approx 170\,{\rm km\,s^{-1}}$; taking $\mu \approx 0.6$ for a fully ionised primordial plasma, Equation 1.4 gives $T_{\rm vir} \approx 10^6\,{\rm K}$, placing the corona in the $10^5 - 10^6\,{\rm K}$ range. This hot, low-metallicity gas gradually radiates away its thermal energy; part of it can cool, condense, and accrete onto the disc, providing a slow but steady supply of fuel to the disc (Tumlinson et al., 2017). The cooling of the coronal gas occurs on a time scale defined by its radiative cooling time, defined as the ratio between the thermal energy density of the gas and its volumetric cooling rate:

$$t_{\rm cool} = \frac{\frac{3}{2}nk_BT}{n_H n_e \Lambda(T,Z)}, \tag{1.5}$$

where n, n_H, and n_e are the total, hydrogen, and electron number densities, and $\Lambda(T,Z)$ is the temperature- and metallicity-dependent cooling function (see, e.g., Sutherland & Dopita, 1993; Wiersma et al., 2009).

Turning this picture into a numerical model is where the difficulty arises. In the idealized configuration previously described, the dense inner corona has a cooling time much shorter than the Hubble time ($t_{\rm cool} \ll t_{\rm Hubble} \sim 14$ Gyr), which means that the gas has enough time to radiate its energy within the life of the galaxy. It is also shorter than the local dynamical time $t_{\rm dyn}$, the time the gas takes to respond dynamically to a change in its pressure, so the gas cools faster than it can re-adjust and cannot settle into a stable hot configuration. Stellar feedback returns some energy to the corona, but especially in the early stages of the evolution this energy is not enough to balance its cooling: the hot gas loses pressure support and descends onto the disc, which receives an excess of gas (see Barbani et al., 2023). Eventually, the galaxy still reaches a steady state, but with a SFR above the observed values (White & Rees, 1978; Somerville & Davé, 2015).

This outcome is not simply a numerical issue to be corrected, but reflects a missing physical ingredient. In real galaxies the corona does not cool undisturbed: energy is returned to it by stellar feedback, by accretion shocks and, in haloes hosting an active nucleus, by AGN feedback. An idealised model that includes none of this additional heatings is therefore physically incomplete. A realistic treatment of the corona requires a mechanism that offsets the cooling of Equation (1.5) and stabilises the hot gas thermodynamically; the form and efficiency of this heating are what this thesis sets out to test.

1.3 The central black hole and the Fermi Bubbles

A natural place to look for this extra heating is the galactic centre. The cold gas that the corona delivers to the disc does not all turn into stars: part of it can flow further inward, towards the central milli-parsecs, where it can feed the supermassive black hole (SMBH) present there (Genzel et al., 2010). For example, at the centre of the Milky Way, there is Sagittarius A* (Sgr A*) with a mass of $M_{\rm BH} \sim 4 \times 10^6\,{\rm M}_\odot$ (Ghez et al., 2008; Gillessen et al., 2009; GRAVITY Collaboration et al., 2022). The amount of energy such a SMBH can release is bounded by its Eddington luminosity, the radiative output at which the outward radiation pressure balances gravity, under the assumptions of steady, spherically symmetric accretion of fully ionised hydrogen (so that radiation couples to the gas through Thomson scattering off free electrons):

$$L_{\rm Edd} = \frac{4\pi G M_{\rm BH} m_p c}{\sigma_T} \approx 1.3 \times 10^{38} \left(\frac{M_{\rm BH}}{{\rm M}_\odot}\right) {\rm erg\,s^{-1}}, \tag{1.6}$$

where G is the gravitational constant, c the speed of light, and σ_T the Thomson cross-section. For the mass of Sgr A* this gives $L_{\rm Edd} \sim 5 \times 10^{44}\,{\rm erg\,s^{-1}}$.

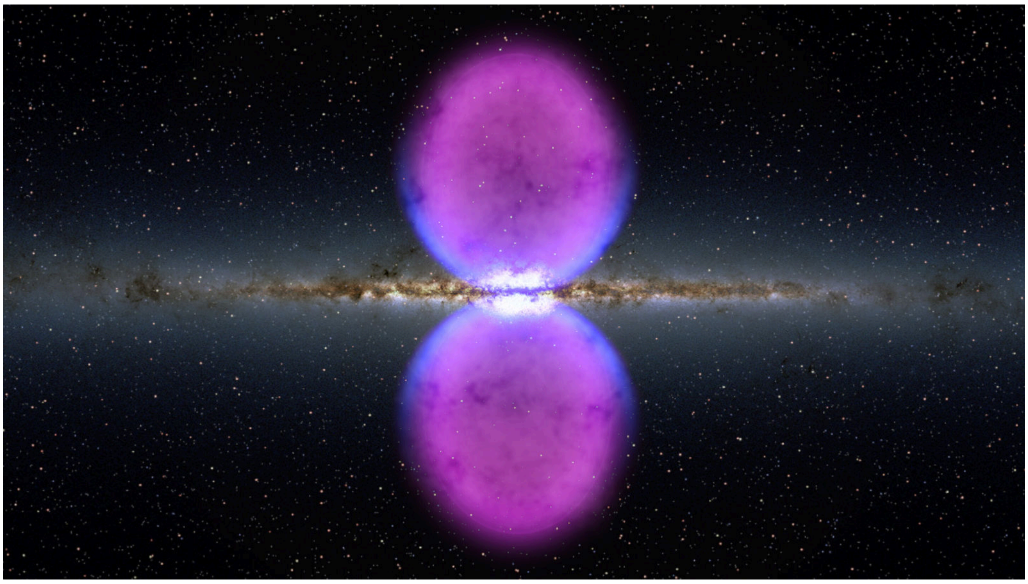

Figure 1.1: The Fermi bubbles extend for about ~ 10 kpc above and below the Galactic disc, as shown in this NASA illustration. X-rays can be seen at the bubbles' edges (blue); gamma rays (magenta) show the bubbles. The origin of these structures is still debated. Image credit: NASA/Goddard Space Flight Centre.

This black hole is currently radiating below this value: its bolometric luminosity (the power radiated by the source integrated over the entire electromagnetic spectrum) is $L_{\rm bol} \sim 10^{36}\,{\rm erg\,s^{-1}}$, corresponding to an Eddington ratio $L_{\rm bol}/L_{\rm Edd} \sim 10^{-9}$ (Genzel et al., 2010). Today it is quiescent, and is not heating the corona in any significant way. Although Sgr A* is quiescent today, there is clear evidence that the Galactic centre was far more active in the recent past. The Fermi Bubbles, two large γ-ray emitting lobes extending about 10 kpc above and below the Galactic plane (Su et al., 2010a; The Fermi-LAT Collaboration, 2014), are the signature of such an episode (see Figure 1.1). Reproducing their size and energetics requires an energy injection of order

$$E_{\rm inject} \sim 10^{55}\text{-}10^{56}\,{\rm erg}, \tag{1.7}$$

released from the central region over a few million years.

The origin of this energy is still debated, between a past accretion episode of Sgr A* (Zubovas et al., 2012) and an intense burst of nuclear star formation (Crocker et al., 2015). For the present work the distinction is not important; what matters is the implication common to both scenarios: the central region of the Galaxy is able to deposit energy of order $10^{55} - 10^{56}$ erg into the surrounding halo.

In this work, this is the only use we make of the Fermi Bubbles: they fix the *scale* of the energy available to heat the corona. We do not model the bubbles themselves, their morphology, emission or kinematics, nor do we model the accretion physics of the central black hole (for instance Bondi–Hoyle–Lyttleton accretion onto a sink particle, for which see Bondi & Hoyle, 1944; Bondi, 1952 and for a review Edgar, 2004); nor do we claim a causal chain linking coronal cooling, black-hole feeding and these outbursts. They serve only as observational evidence that an AGN has actually released an amount of energy of this magnitude. The energy budget of our more realistic model (see Chapter 6) is not taken from this figure: it is set independently from the halo X-ray luminosity through a luminosity-temperature relation (see Section 6.2 for further details). The Fermi Bubble energetics are used only as an a posteriori check on the scale of this budget, discussed quantitatively in Section 6.2.

1.4 The heating models

To test the idea that such an energy could have some effects on the corona cooling flow, we add three separate heating channels to the simulations. Two of these models, that are described in the following and more in detail in Chapter 5, are built on the principle of *cooling compensation*: rather than imposing an arbitrary energy flux, we estimate the thermal energy lost by the corona to radiation and return a comparable amount to the gas,

$$E_{\rm heat} \approx \left| \dot{E}_{\rm cool} \right| \Delta t. \tag{1.8}$$

These two variants (Heat 1 and Heat 2) share the same energy budget and differ only in the timescale over which it is delivered. A third model (in detail in Chapter 6) instead injects energy through discrete hot bubbles, in the spirit of radio-mode AGN feedback (see Sections 1.2 and 1.3), with the energy budget set from the halo X-ray luminosity through a scaling relation rather than from self-consistent black-hole accretion. In all cases the heating is restricted to the diffuse, hot coronal gas (through a density threshold and a temperature cut), so that the dense, star-forming disc is left untouched. We compare these three heating prescriptions against a control run without heating, whose details given in Chapter 4.

1.5 Aims of this thesis

The aim of this thesis is to investigate, through numerical simulations of an isolated Milky Way-like galaxy, how heating the corona affects its gas content and the gas it supplies to the disc. Starting from a control run in which the corona cools without any compensating heating, we introduce three heating prescriptions and ask how each of them changes the thermodynamic state and the spatial distribution of the coronal gas over timescales comparable to the duration of the simulation.

A central question is whether reducing the cooling of the corona translates into a reduced gas supply to the disc, and therefore a lower SFR. As we will show, this connection turns out to be weaker than one might expect. The net inflow of hot gas onto the disc region is essentially the same with or without heating, so the corona keeps feeding the disc at the same rate regardless of the energy injected. The SFR, set largely by stellar feedback within the disc, remains around $3-5\,\mathrm{M_\odot\,yr^{-1}}$ in all models and is only marginally sensitive to the state of the corona. Where the prescriptions do leave a measurable imprint is on the corona itself, and the way the energy is delivered matters more than its total amount. The distributed models (Heat 1 and Heat 2) retain a larger hot-gas reservoir within the inner corona (with respect to the No-Heat), whereas the localised bubble model, drives the strongest spatial redistribution of the coronal gas: it depletes the hot phase within $\sim 80-100$ kpc while pushing hot gas outward, where it accumulates in the outer corona over the timescales investigated. Characterising this relative independence of the disc from the corona, and identifying where each prescription does leave a measurable imprint, is the main goal of the analysis.

1.6 Structure of the thesis

The subsequent chapters of this thesis are organised as follows: Chapter 2 describes the simulation code, the AREPO moving-mesh scheme and the SMUGGLE model for the ISM and stellar feedback. The initial conditions of the isolated galaxy are instead explained into Chapter 3. Chapter 4 presents the control run without heating, called the No-Heat model, establishing the behaviour of the galaxy when the corona cools without compensating heating. Chapter 5 examines the two cooling-compensation models (Heat 1 and Heat 2), while Chapter 6 the bubble model inspired by the AGN feedback. Chapter 7, finally, summarises the results and discusses some possible paths to continue and improve the work of this thesis.

Chapter 2

Numerical Methods

Generally speaking, astrophysical systems are governed by highly non-linear phenomena that cannot be solved entirely through analytical approaches, thus requiring the use of advanced computational methods. In the context of galaxy evolution, the physical processes can be broadly divided into two distinct regimes: the collision-less dynamics of dark matter and stars, and the collisional, thermodynamic evolution of the multiphase gas. In a collision-less system the constituent bodies interact only through the gravitational field they generate, and individual close encounters are negligible over the timescales of interest; this is an excellent approximation for the dark matter and stars of a galaxy. A collisional component such as the interstellar gas, by contrast, evolves through local interactions like pressure gradients, shocks, and radiative cooling and heating, which couple neighbouring fluid elements directly.

To accurately model a galaxy, a numerical simulation must seamlessly combine these two regimes. For this reason, the simulations analysed in this thesis are run using AREPO (Springel, 2010b), a highly efficient moving-mesh code designed to compute both the gravitational interactions and the fluid motions (see Section 2.3). Furthermore, a galaxy simulation faces a problem of scale: any code, including AREPO, can solve the dynamics of gas and matter only down to the size of its resolution elements. Many of the processes that actually govern a galaxy occur below the finer scale that can be resolved in a simulation: the formation of individual stars, single supernova explosions, and the structure of cold molecular clouds all take place on scales of parsecs or less. These sub-grid processes are handled by the SMUGGLE (Stars and Multiphase Gas in Galaxies) model (Marinacci et al., 2019). This comprehensive sub-grid model is specifically designed to account for the processes that occur below the resolution elements and feeds their effects back into the gas that AREPO evolves. AREPO thus provides the dynamics on resolved scales, while SMUGGLE represents the unresolved ISM physics that drives the baryonic evolution.

In this chapter we first describe how the gravitational forces are computed, then the hydrodynamical solver, and finally the SMUGGLE model.

2.1 Numerical methods for gravity

Gravity is the primary force shaping the formation of cosmic structures, being directly responsible for the processes that lead to the formation of cosmic structures. To model the evolution of complex systems such as galaxies, numerical simulations must solve the N-body problem, calculating the mutual gravitational attraction between N discrete mass elements.

This calculation requires an approach including each particle to interact with all the others, scaling as $\mathcal{O}(N^2)$, resulting in a prohibitive computational cost for modern simulations containing millions or several billions of resolution elements. Consequently, the implementation of approximate, yet highly optimized, gravitational solvers becomes a necessity (Vogelsberger et al., 2020). The primary challenge of this implementation of the multiple-particle interactions consists in evaluating the acceleration experienced by each i-th particle to integrate its equation of motion

$$\ddot{\mathbf{x}}_i = -\nabla\Phi(\mathbf{x}_i)\,, \tag{2.1}$$

where $\mathbf{x}_i$ is the coordinate vector of the particle, $\ddot{\mathbf{x}}_i$ is its acceleration, and Φ is the global gravitational potential. The exact potential generated by all the particles in the system can be computed as

$$\Phi(\mathbf{x}_i) = -G\sum_{j\neq i}^{N} \frac{m_j}{\sqrt{\|\mathbf{x}_i - \mathbf{x}_j\|^2 + \epsilon^2}}\,, \tag{2.2}$$

where G is the gravitational constant and m_j is the mass of the j-th particle. Crucially, Equation (2.2) introduces the parameter ϵ, absent in the classical Newtonian formulation. This numerical prescription, known as *gravitational softening*, is necessary because without it, the gravitational force diverges as two particles approach each other. This has two unwanted consequences: it can halt the integration and it produces spurious encounters, that mimic the stellar ones, absent in a truly collisionless system where each resolution element represents an ensemble of mass rather than an individual star. Evaluating the potential directly through Equation (2.2) corresponds to the *direct summation* technique. Once the potential is obtained, the evolution of the system is derived by integrating the following set of coupled ordinary differential equations for each particle

$$\begin{cases} \dot{\mathbf{v}}_i = -\nabla\Phi(\mathbf{x}_i) \\ \dot{\mathbf{x}}_i = \mathbf{v}_i \end{cases} \tag{2.3}$$

These equations are typically integrated using specific solvers, such as the *Leapfrog* scheme (see, e.g., Springel, 2005), which is implemented in AREPO. However, direct summation is highly inefficient. Computing for N particles, the force between all $N(N-1)/2$ pairs, requires $\sim \mathcal{O}(N^2)$ operations per time-step, which becomes prohibitive for the millions of resolution elements considered. To overcome this limitation, several approximated but highly efficient techniques have been developed. Tree-based methods reduce the cost of the gravitational interactions to $\mathcal{O}(N \log N)$, while mesh-based methods scale as $\mathcal{O}(N)$ in the particle operations (mass assignment and force interpolation) at the price of an additional Poisson solve on the grid, of cost $\mathcal{O}(N_g \log N_g)$ for N_g grid cells. These techniques are described in the subsequent Sections.

2.1.1 Particle-Mesh (PM) method

The Particle-Mesh (PM) method (Hockney & Eastwood, 1981) evaluates gravity by solving Poisson's equation

$$\nabla^2\Phi(\mathbf{x}) = 4\pi G\rho(\mathbf{x}) \tag{2.4}$$

on a Cartesian grid. Instead of computing point-to-point interactions, the method relies on Fourier space, where the differential Poisson equation reduces to a simple algebraic product:

$$-k^2\hat{\Phi}(\mathbf{k}) = 4\pi G\hat{\rho}(\mathbf{k})\,, \tag{2.5}$$

with $\mathbf{k}$ being the wave vector. The procedure involves assigning the particle masses to the mesh cells to derive a continuous density field $\rho(\mathbf{x})$. After applying a Fourier Transform (FT) to obtain $\hat{\rho}(\mathbf{k})$, the potential $\hat{\Phi}(\mathbf{k})$ is computed and transformed back to real space via inverse FT. Finally, the accelerations are evaluated using finite differences and interpolated back to the particle positions. The PM method is extremely fast, but its spatial resolution is intrinsically limited by the size of the mesh cells (the force resolution cannot go beyond the size of a single mesh cell). Therefore, it is ideally suited for evaluating long-range forces in almost homogeneous, large-scale matter distributions.

2.1.2 Tree method

The Tree method (Barnes & Hut, 1986) takes a fundamentally different approach, based on the hierarchical subdivision of the simulation domain. The space is recursively divided into eight cubic subdomains (nodes) until each node contains at most one particle (the "leaves" of the tree, forming an *octree* structure).

The core principle is to group distant particles together: when evaluating the force acting on a target particle, the gravitational contribution of a distant node is approximated using a multipole expansion of its mass distribution, treating it as a single massive entity. This drastically reduces the number of force evaluations, replacing the $\mathcal{O}(N^2)$ direct summation with an $\mathcal{O}(N \log N)$ scaling.

The Tree method automatically provides high spatial resolution in dense, highly clustered regions. However, its efficiency drops significantly when dealing with large-scale, homogeneous mass distributions. Whether a distant node can be approximated by its multipole expansion or must instead be opened into its sub-nodes is decided by the *Multipole Acceptance Criterion* (MAC), which compares the angular size of the node to its distance from the target particle: a node that subtends too large an angle is too close to be treated as a single entity and must be opened. In a nearly uniform medium the space around any given particle is filled with mass down to the smallest scales, so the MAC is violated throughout much of the local volume, forcing the algorithm to open nodes all the way down to individual leaves. The local force evaluation then reverts to a direct particle-particle summation, pushing the cost back towards $\mathcal{O}(N^2)$ and degrading the overall performance.

2.1.3 Tree-PM method

The two methods are complementary: each is efficient precisely where the other is not, so modern N-body codes commonly implement a hybrid strategy known as the Tree-PM method (Bode & Ostriker, 2003; Springel, 2005). This algorithm splits the gravitational potential into two distinct components: a short-range force, accurately computed using the high-resolution Tree method, and a long-range force, efficiently evaluated using the PM method on a coarse grid. This combination is particularly powerful and represents the standard choice for large-scale cosmological simulations.

2.2 Numerical methods for hydrodynamics

While dark matter dominates the global gravitational potential of a galaxy, it is the baryonic gas that dictates its observable structure and its long-term evolution. In particular, the ISM and circumgalactic medium (CGM) behave as a highly dynamic, collisional fluid. This gas is intrinsically multiphase, spanning several orders of magnitude in temperature and density, and coexisting in distinct thermal states, from the cold and dense star forming clouds to hot and rarefied corona. Later in this work (see Sections 5.5.1 and 6.8), temperature-density phase diagrams will serve as an important tool to visually track these states, providing direct evidence of how the gas moves across different regimes.

Physical processes at the core of galaxy formation, such as radiative cooling, shock heating, turbulence, and supernova feedback act continuously on this gas, driving mass and energy exchanges between the different phases. Because the fluid dynamics governing these mechanisms are non-linear, analytical solutions are seldom available. Therefore, numerical hydrodynamical methods are strictly required to solve the fluid equations and track the thermodynamic evolution of the gas.

Assuming an inviscid fluid and neglecting magnetic fields, gas dynamics are governed by the Euler equations, a set of hyperbolic, non-linear partial differential equations representing the conservation of mass, momentum, and energy

$$\begin{cases} \dfrac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \\ \\ \dfrac{\partial (\rho \mathbf{v})}{\partial t} + \nabla \cdot (\rho \mathbf{v} \otimes \mathbf{v} + P\mathcal{I}) = 0 \\ \\ \dfrac{\partial E}{\partial t} + \nabla \cdot [(E + P)\mathbf{v}] = 0, \end{cases} \tag{2.6}$$

where ρ is the mass density, $\mathbf{v}$ is the velocity vector, P is the thermal pressure, $\mathcal{I}$ is the identity matrix, and E is the total energy density of the fluid. The total energy density is defined as the sum of kinetic and internal energy densities as follows

$$E = \frac{1}{2}\rho|\mathbf{v}|^2 + \frac{P}{\gamma - 1} \,. \tag{2.7}$$

To close this system of equations, an equation of state relating pressure, density, and specific internal energy e is required. For an ideal fluid, this relation reads

$$P = (\gamma - 1)\rho e \,, \tag{2.8}$$

where γ is the adiabatic index (set to 5/3 for a monoatomic ideal gas). To numerically solve the Euler equations, two historically distinct families of computational methods have been widely adopted: Eulerian and Lagrangian schemes.

2.2.1 Eulerian and Lagrangian codes

In an **Eulerian** approach, the computational domain is partitioned into a number of small, non-overlapping volumes called *cells*; together they constitute what is known as a *mesh*, or grid. Each cell is a stationary geometric region anchored in space, and the fluid is described not element by element but through the average values of its properties (density, momentum and energy) within each cell. The mesh can be *fixed*, with cells of uniform size set once at the start, or *adaptively refined*, meaning that cells are automatically divided into smaller ones wherever finer resolution is required (for instance dense regions) and merged again where it is not. Rather than tracking individual fluid elements, the Eulerian approach computes the fluxes of mass, momentum and energy across the cell interfaces as the fluid flows through them.

The fluid variables in each cell are updated at each timestep, often through finite-volume Godunov schemes (Godunov, 1959), in which the conserved quantities are evolved by solving, at every cell interface, a *Riemann problem* which is the exact evolution of a fluid initially separated into two uniform states by a discontinuity (Toro et al., 1994; Toro, 2013). Solving these local problems is what allows the method to capture sharp features such as shocks. While Eulerian methods are excellent at resolving shocks and fluid instabilities (e.g., Kelvin-Helmholtz, see, e.g., Agertz et al. 2007), their main drawbacks are the lack of *Galilean invariance* which means that the solution depends on the bulk velocity of the fluid relative to the grid, so the same physical flow is resolved with different accuracy depending on how fast it moves across the mesh and the need to refine this explicitly to resolve dense regions (Springel, 2010b).

Conversely, **Lagrangian** methods (such as Smoothed Particle Hydrodynamics, SPH, Lucy 1977; Gingold & Monaghan 1977; Springel 2010a) follow the motion of individual fluid elements using discrete particles. In this perspective, each particle represents a moving parcel of fluid with a fixed mass but a variable volume, shifting the focus from a fixed spatial frame directly to the fluid element itself. The continuous fluid properties are recovered through interpolation over neighbouring particles using a *smoothing kernel*: a weighting function that decreases with distance, so that nearby particles contribute more than distant ones to the local estimate. This approach is intrinsically Galilean-invariant and naturally adapts the resolution to the local mass density, making it highly effective for galaxy evolution simulations with enormous dynamic ranges. However, standard SPH struggles to accurately resolve sharp discontinuities, such as strong shock waves and surface instabilities, due to artificial viscosity and the suppression of fluid mixing (Agertz et al., 2007).

2.2.2 Arbitrary Lagrangian-Eulerian (ALE) codes

To mitigate the respective shortcomings of fixed-grid and particle-based methods, a hybrid approach has been developed: the **Arbitrary Lagrangian-Eulerian (ALE)** methods. In ALE codes, the simulated volume is discretized into cells, but the mesh itself is allowed to move and deform following the bulk flow of the fluid. This dynamic mesh topology naturally adapts the spatial resolution to the fluid density (like SPH) while retaining the highly accurate shock-capturing capabilities of mesh-based solvers. The moving-mesh code **AREPO** (Springel, 2010b; Weinberger et al., 2020), used to carry out the simulations analysed in this work, is a perfect example of this technique.

2.3 The AREPO Code

To accurately solve the coupled equations of gravity and hydrodynamics, the simulations analysed in this work are performed using **AREPO** (Springel, 2010b; Weinberger et al., 2020). AREPO is a massively parallel, highly optimized code designed for gravitational N-body systems and magnetohydrodynamics, capable of operating on both Newtonian and cosmological backgrounds.

As introduced previously, AREPO was specifically developed to bridge the gap between traditional approaches, combining the shock-capturing advantages and high accuracy of finite-volume Eulerian hydrodynamics with the Galilean invariance and spatial adaptability of Lagrangian SPH methods. To achieve this ALE behaviour, the code does not rely on a fixed Cartesian grid, but rather employs an unstructured moving mesh built from a *Voronoi tessellation*. In its standard operational mode, this mesh moves continuously along with the fluid flow in a quasi-Lagrangian fashion, dynamically adjusting the geometry and the volume of the cells according to the local density field.

2.3.1 The Voronoi moving mesh

To partition the simulated volume, AREPO employs a highly adaptable topology built from a set of coordinates called *mesh-generating points*. The mesh is constructed in two steps. First, AREPO builds a *Delaunay triangulation* of these points, composed of tetrahedra (or triangles, in two dimensions) whose vertices are the mesh-generating points themselves with the fundamental property that the circumsphere of any given tetrahedron contains no other generating points. This is the operation that can be carried out most quickly, which is why the code starts

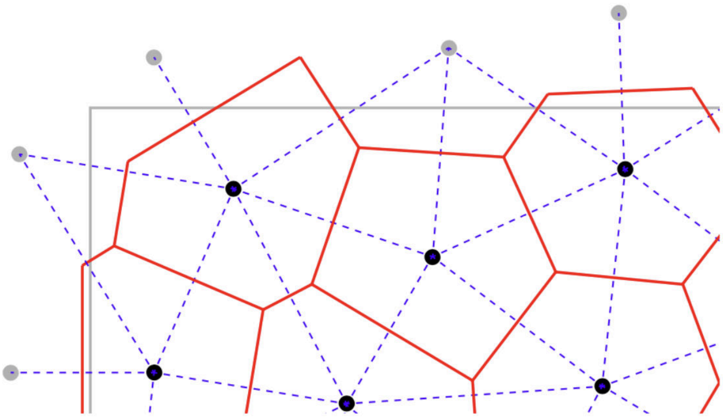

Figure 2.1: Illustration of the Voronoi tessellation in two dimensions. The red lines show the Voronoi cells, the blue dashed lines show the corresponding Delaunay triangulation while the black dots are the mesh generating points. Figure taken from Weinberger et al. (2020).

from it (Springel, 2010b). From the Delaunay triangulation, AREPO then derives its geometric dual, the *Voronoi tessellation*: the partition of space in which the cell associated with each generating point encompasses the region that is strictly closer to that point than to any other (see Figure 2.1).

The defining feature of AREPO is that these mesh-generating points are not static, but are allowed to move dynamically with the fluid flow. At each computational time-step, the position of every generating point is updated according to the local bulk velocity of the gas, and the Voronoi mesh is reconstructed consistently. Specifically, the velocity vector $\mathbf{w}_i$ assigned to the i-th mesh-generating point is defined as

$$\mathbf{w}_i = \mathbf{v}_i + \mathbf{v}_{\mathrm{corr},i} \, , \tag{2.9}$$

where $\mathbf{v}_i$ is the actual fluid velocity evaluated within the cell, and $\mathbf{v}_{\mathrm{corr},i}$ is a corrective velocity.

This correction is necessary to maintain the numerical accuracy of the scheme. During the fluid evolution, the mesh-generating points can drift away from the true centre of mass (geometric centre) of their respective cells, leading to highly distorted, elongated cell shapes. Such topological distortions severely degrade the accuracy of the spatial gradients and the linear reconstruction step at the cell interfaces (Pakmor et al., 2016). The corrective velocity $\mathbf{v}_{\mathrm{corr},i}$ is therefore introduced to dynamically realign the generating points with the centre of mass of their respective cells (Springel, 2010b), ensuring that the mesh remains sufficiently regular and minimizing the rotation rate of the cell faces during the motion.

2.3.2 Hydrodynamic solver

To solve the equations of hydrodynamics, AREPO employs a finite-volume Godunov method specifically adapted for an unstructured, moving Voronoi mesh (Springel, 2010b). Instead of tracking complex continuous fields, the macroscopic state of the fluid is represented by discrete conserved variables—mass, momentum, and total energy—which are volume-averaged within each individual cell.

The temporal evolution of the fluid is governed by the exchange of these conserved quantities across the boundaries of adjacent cells. To accurately evaluate these conservative fluxes, the code reconstructs the conserved variables of the fluid at each cell interface and solves a local *Riemann problem*: the evolution of two initially constant fluid states separated by a discontinuity, which determines how the physical waves and the associated fluxes propagate across the surface. While AREPO also supports approximate solvers (e.g., the HLLC scheme) to reduce computational load, the simulations presented in this work rely on the exact Riemann solver to maximize accuracy, particularly in resolving strong shock waves and contact discontinuities in the ISM.

Finally, the fluid state is advanced in time using a second-order Runge-Kutta scheme (Heun's method). The original version of AREPO (Springel, 2010b) employed a MUSCL-Hancock integrator, which uses the mesh geometry only at the beginning of the time-step and is therefore only first-order accurate for general mesh motions; it was subsequently replaced by the present second-order scheme (Pakmor et al., 2016), which recovers full second-order convergence and substantially improves the conservation of angular momentum. This time integration is carefully formulated to remain strictly conservative, accounting not only for the hydrodynamical fluxes but also for the continuous translation, rotation, and topological deformation of the mesh faces that occur during each time-step.

2.3.3 Time-step constraints

The time-step Δt chosen to integrate the equations of hydrodynamics and gravity is a fundamental parameter that dictates the numerical stability of the simulation. If the integration step is too large, information travel farther than the numerical scheme can follow within one step. Specifically, if a sound wave or a shock front crosses more than one cell within a single timestep, the numerical scheme (which exchanges fluxes only between neighbouring cells) can no longer track the interaction so the solution becomes numerically unstable.

For the hydrodynamical evolution, AREPO enforces the Courant-Friedrichs-Lewy (CFL) stability criterion, which sets the condition to avoid the problem previously presented. The maximum allowed time-step for a given cell is restricted by

$$\Delta t_{\mathrm{hydro}} \leq C_{\mathrm{CFL}} \frac{R_{\mathrm{cell}}}{c_s} , \tag{2.10}$$

where C_{CFL} is the Courant factor (set to 0.3 in our runs), $R_{\mathrm{cell}} = (3V/4\pi)^{1/3}$ is the effective radius of the cell based on its volume V, and $c_s = \sqrt{\gamma P/\rho}$ is the local sound speed. This formulation highlights the main advantage of a moving-mesh approach over traditional Eulerian methods. In a static grid, the maximum timestep must account for how fast the fluid crosses the stationary cells. Consequently, the denominator of the CFL condition must include both the sound speed and the fluid bulk velocity ($c_s + |\mathbf{v}|$). In supersonic flows, the bulk velocity $|\mathbf{v}|$ dominates between the two, forcing the timestep to become small and prohibitive in terms of computational costs. In AREPO, however, the Voronoi mesh moves continuously alongside the fluid. Because the gas is nearly stationary relative to its own moving cell boundaries, the bulk term is largely reduced from the constraint. The time step remains dictated solely by the local sound speed c_s, preserving high computational efficiency regardless of the fluid velocity.

Separately, the gravitational dynamics impose a distinct stability constraint. To ensure an accurate numerical integration of the particle trajectories, the gravitational time-step is strictly limited by the local acceleration $\mathbf{a}$ and the gravitational softening length ϵ

$$\Delta t_{\mathrm{grav}} \leq \sqrt{\frac{2 C_{\mathrm{grav}} \epsilon}{|\mathbf{a}|}} , \tag{2.11}$$

with the empirical tolerance parameter set to $C_{\rm grav} = 0.012$.

Furthermore, to prevent integration errors during sudden fluid accelerations, AREPO implements an additional timestep limiter based on the maximum signal velocity of the Riemann solver (see Springel, 2010b, for details). This actively reduces Δt if a strong shock wave from neighbouring regions is rapidly approaching the cell. Eventually, the time-step actually assigned to a cell is the smallest of these constraints, in order to fulfil the CFL condition presented.

Individual time-stepping

Astrophysical simulations modelling galaxies span a wide dynamic range in gas density and temperature. The dense, cold ISM in the galactic centre requires to set extremely short time-steps into the simulation to avoid numerical instability, whereas the diffuse, hot coronal gas in the outskirts can be safely integrated with much larger steps. Forcing every cell onto the shortest required time-step would make the computation prohibitively expensive.

To overcome this, AREPO employs a local time-stepping scheme, assigning each particle its own timestep based on the local CFL and gravitational constraints. However, to keep the entire simulation synchronized, these values cannot be arbitrary. Instead, every individual timestep is constrained to be the global maximum timestep divided by a power of two (e.g., 1/2, 1/4, 1/8). This allows particles taking multiple short steps to periodically align in time with particles taking a single large step. This temporal alignment ensures that their mutual interactions are computed correctly, improving the performance without losing local precision.

2.4 The SMUGGLE sub-grid model

Simulating realistic galaxy evolution presents a fundamental multi-scale challenge: it requires following the large-scale dynamics of an entire galaxy while, at the same time, resolving the sub-parsec physics of processes such as star formation and supernova explosions. Furthermore, the ISM that mediates this coupling is intrinsically multiphase, with cold, warm and hot gas coexisting and continuously exchanging mass and energy. As a result, formulating a complete theoretical framework for galaxy evolution is a complex procedure (Somerville & Davé, 2015; Naab & Ostriker, 2017; Vogelsberger et al., 2020).

Historically, to bypass the prohibitive computational cost of resolving these processes, galaxy formation simulations relied on an effective Equation of State (eEOS) to treat the ISM as a single fluid with an artificially stiffened pressure, derived from a two-phase description in which cold clouds are assumed to be in pressure equilibrium with a hot ambient medium (Springel & Hernquist, 2003). While numerically stable, this approach is based on an artificial pressure imposed to the gas by construction that smooths this medium. This leads to a problem: the physics that would regulate the collapse of cold, dense gas clouds in a more clumpy structure, would not be captured. This happens because a region of gas collapses under its own gravity once it becomes denser and colder than a characteristic scale – the Jeans scale – below which self-gravity overcomes the supporting pressure. When this scale falls below the size of a resolution element, the code cannot follow the internal structure that would govern this collapse, therefore not accurately modelling the environment where stars actually form.

To overcome these limitations and to capture the true multiphase nature of the ISM, the simulations analysed in this work employ the Stars and MUltiphase Gas in GaLaxiEs model (SMUGGLE, Marinacci et al. 2019), implemented within the AREPO moving-mesh code. Rather than prescribing an effective pressure like the eEOS, SMUGGLE evolves the actual thermal state of each gas cell from explicit cooling and heating processes, so that gas at different temperatures can coexist rather than being forced onto a single temperature and pressure smooth medium. In this SMUGGLE aligns with the modern paradigm of explicit stellar feedback and a resolved multiphase ISM, shared by simulation frameworks, such as FIRE-2 (Hopkins et al., 2018), VINTERGATAN (Agertz et al., 2021) and COLIBRE (Schaye et al., 2025).

SMUGGLE models the multiphase ISM through three coupled components. The first is the tracked radiative cooling and heating of the gas: primordial (hydrogen and helium) and metal-line cooling, but also extended to low-temperature atomic, fine-structure, and molecular processes (with a prescription for the self-shielding of dense gas from the UV background) so that the gas can cool down to $T \sim 10$ K, appropriately balanced by photoelectric and cosmic-ray heating. The second is a stochastic prescription for star formation in gas that exceeds a local density threshold. The third is a detailed treatment of stellar evolution which return mass and metals to the surrounding ISM. This last component is where stellar feedback enters, injecting energy and momentum through three primary channels: core-collapse and Type Ia supernovae, radiative feedback (encompassing photoionization and radiation pressure from massive stars), and stellar winds ejected by young OB and Asymptotic Giant Branch (AGB) stars. Acting together, these channels inject energy and momentum back into the gas, reproducing the observed Kennicutt-Schmidt relation (Marinacci et al., 2019).

To show how these three components govern the ISM, the following subsections will sequentially detail the numerical implementation of the gas cooling, the star formation and evolution criteria, and finally the stellar feedback channels.

2.4.1 Cooling and heating

The thermal state of the multiphase ISM is governed by a delicate balance between radiative energy losses and internal heating processes. This balance is formalized by the net cooling function

$$\Lambda_{\rm net} = \Lambda_{\rm cool}(T, X, Z) - \Lambda_{\rm heat}\,, \tag{2.12}$$

where $\Lambda_{\rm cool}$ and $\Lambda_{\rm heat}$ are, respectively, the cooling and heating rates per unit volume, while $L = n^2\Lambda_{cool}(T, X, Z)\ \mathrm{erg\,cm^3\,s^{-1}}$ is defined as the luminosity per unit volume, and n is the gas number density. The cooling function strongly depends on the gas temperature T, its ionization state X, and its metallicity Z. Accurately capturing this interplay is crucial in galaxy evolution simulations, as the localized thermodynamic state dictates the ability of the gas to collapse and trigger star formation.

Within the SMUGGLE framework, the total cooling rate is modelled as a superposition of multiple physical channels. The model explicitly accounts for primordial cooling (two-body processes from H and He networks) and Compton cooling off the Cosmic Microwave Background (CMB) photons, following the implementation by Vogelsberger et al. (2013). However, in the dense, metal-enriched ISM, the dominant energy loss mechanism is the metal-line cooling. This is computed using pre-compiled tables from the photo-ionization code CLOUDY (Ferland et al., 1998), scaled linearly with the local gas metallicity. Furthermore, to accurately resolve the cold neutral medium, SMUGGLE extends the classical cooling models by adopting low-temperature metal and fine-structure formulation developed for the FIRE-2 model (Hopkins et al., 2018). Importantly, the cooling network is subjected to a spatially uniform, redshift-dependent ionizing UV background (Faucher-Giguère et al., 2009). Since this background radiation would over-heat the dense star-forming clouds, SMUGGLE dynamically applies a local gas self-shielding correction, $f_{\rm ssh}$, based on the phenomenological fits provided by Rahmati et al. (2013, see also Vogelsberger et al. 2013). This shielding factor progressively attenuates the UV photoionization and photo-heating rates as the local hydrogen density increases. Scaled by this factor, the molecular cooling channels allow the dense, neutral gas to dissipate energy and reach very low temperatures ($T \sim 10$ K), effectively enabling gravitational collapse, which is needed for star formation.

To counterbalance these extreme energy losses and maintain the structural stability of the multiphase ISM, the model incorporates two primary environmental heating mechanisms: cosmic ray (CR) heating and photoelectric heating. Cosmic ray heating (modelled following Guo & Oh, 2008) operates primarily in dense gas clouds via Coulomb interactions between cosmic rays and free electrons, as well as through hadronic collisions. Simultaneously, photoelectric heating is driven by the interaction between the interstellar UV radiation field and dust grains, predominantly Polycyclic Aromatic Hydrocarbons (PAHs). Following the framework of Wolfire et al. (2003), UV photons eject electrons from the dust grains via the photoelectric effect; these electrons subsequently thermalize the surrounding gas. Since dust abundance is tightly coupled to the metal content of the ISM, this heating rate scales directly with the local gas metallicity ($Z/Z_\odot$). Ultimately, these combined heating rates ($\Lambda_{\rm heat} = \Lambda_{\rm CR} + \Lambda_{\rm phot}$) are subtracted from the cooling mechanisms to yield the final net cooling function, governing the global thermodynamic evolution of the simulated galaxy.

2.4.2 Star formation

Within the simulation, the process of star formation is implemented by probabilistically converting gas cells into collision-less stellar particles. A gas cell is eligible to undergo star formation only when it satisfies two physical criteria, linking the process to dense and cold gas environments:

- its local density must exceed a strict threshold, $\rho > \rho_{\rm th} = 100\,{\rm cm}^{-3}$. This high-density cut-off effectively restricts star formation to the cold, dense gas, representative of the physical conditions encountered in the cores of Giant Molecular Clouds (GMCs);

- the gas must be gravitationally bound. This condition is evaluated using the local virial parameter, α, which compares the kinetic and thermal support of the gas against its self-gravity
$$\alpha_i = \frac{||\nabla \otimes v_i||^2 + (c_{s,i}/\Delta x_i)^2}{8\pi G \rho_i}, \tag{2.13}$$
where $||\nabla \otimes v_i||$ is the norm of the velocity gradient tensor (capturing macroscopic turbulence), $c_{s,i}$ is the local sound speed (capturing microscopic thermal pressure), Δx_i is the effective size of the cell, and G is the gravitational constant. A cell is considered bound and allowed to form stars only if $\alpha < 1$. Because the sound speed in an ideal gas is proportional to the square root of the temperature ($c_s \propto \sqrt{T}$), this virial condition naturally enforces the low-temperature constraint: if the gas is too hot, the high thermal pressure drives $\alpha \gg 1$, effectively preventing gravitational collapse.

If a gas cell fulfils both criteria, its instantaneous SFR is computed assuming a local volumetric Schmidt law
$$\dot{M}_\star = \varepsilon \frac{M_{\rm gas}}{t_{\rm dyn}}\,, \tag{2.14}$$
where $M_{\rm gas}$ is the mass of the gas cell, and ε is the star formation efficiency per free-fall time. Following empirical observations (e.g., Krumholz & Tan, 2007), this efficiency is set to $\varepsilon = 0.01$. The dynamical time, $t_{\rm dyn}$, corresponds to the local gravitational free-fall time
$$t_{\rm dyn} = \sqrt{\frac{3\pi}{32 G\rho}}\,. \tag{2.15}$$

In terms of numerical implementation, allowing a cell to continuously spawn stellar particles with infinitesimal masses $\dot{M}_\star \Delta t$ would lead to a significant increase in the total number of particles, rapidly degrading the performance of the N-body solver. To maintain computational efficiency, AREPO employs a stochastic conversion scheme introduced by Springel & Hernquist (2003). Rather than creating many small mass stellar particles, the code aims to keep the mass of new stellar particles roughly equivalent to the target mass resolution of the gas (i.e., the typical mass of a single gas cell in the simulation). At each time-step Δt, an eligible gas cell has a probability p to be entirely converted into a single, collision-less star particle, given by

$$p = 1 - \exp\left(-\frac{\dot{M}_\star \Delta t}{M_{\rm gas}}\right) . \tag{2.16}$$

By drawing a uniformly distributed random number $p^* \in [0, 1)$ and executing the full conversion only if $p^* < p$, the code preserves the globally averaged SFR of the galaxy. Because an entire gas cell is completely replaced by exactly one stellar particle of the same mass, this method conserves the total discrete particle count of the simulation.

Once successfully formed, the new stellar particle does not represent a single star, but rather a Simple Stellar Population (SSP) of mass $M_\star$, assumed to be coeval and of uniform initial metallicity. The total mass $M_\star$ of this population is given by the integral of the Initial Mass Function (IMF), $\Phi(m)$, which is assumed to be fully sampled between a minimum stellar mass of $0.1\,{\rm M}_\odot$ and a maximum mass of $100\,{\rm M}_\odot$

$$M_\star(t = t_0) = \int_{M_{min}}^{M_{max}} m\Phi(m)\, dm\, . \tag{2.17}$$

The specific choice of the IMF is a critical parameter for the entire sub-grid model. It directly determines the relative abundance of massive stars, thereby governing the expected rate of core-collapse supernovae (type II-SNe) and the subsequent injection of energy, momentum, and metals back into the ISM. The SMUGGLE model adopts a Chabrier IMF (Chabrier, 2003), which is characterized by a log-normal distribution for low-mass stars and a power-law tail for the massive components and has the form

$$\Phi(m) = \begin{cases} \frac{0.158}{2.3m} \exp\left[-\frac{(\log m - \log 0.08)^2}{2\times 0.69^2}\right] & \text{for } m \leq 1\,{\rm M}_\odot\, , \\ 0.238\, m^{-2.3} & \text{for } m > 1\,{\rm M}_\odot\, . \end{cases} \tag{2.18}$$

Compared to a classical Salpeter IMF (Salpeter, 1955), the Chabrier IMF is relatively top-heavy, so it yields a higher fraction of high-mass stars per unit of total formed stellar mass, consequently driving a higher specific energy and momentum injection rate into the surrounding gas.

2.4.3 Stellar evolution and chemical enrichment

Following their formation, the collision-less stellar particles (representing a coeval SSP) act as dynamic sources of mass and metals for the surrounding ISM. The SMUGGLE model tracks the internal ageing of these SSPs by employing the comprehensive stellar evolution framework, originally developed by Vogelsberger et al. (2013), increased with explicit mass loss prescriptions for massive OB stars. The return of material to the ISM is driven by three primary evolutionary channels: stellar winds from young OB stars, winds from Asymptotic Giant Branch (AGB) stars, and explosive ejections from both core-collapse (Type II) and thermonuclear (Type Ia) supernovae. To compute the mass and metals released at any given time-step Δt, the code relies on the metallicity-dependent stellar lifetime function, $\tau(m, Z)$, tabulated by Portinari et al. (1998). This function dictates the precise mass range of stars within the SSP that are leaving the main sequence at the current simulation time.

Rather than computationally resolving the complex, short-lived post-main-sequence phases, the model assumes that stars inject their mass and chemical yields instantaneously at the end of their main-sequence lifetime. The total ejected mass, in a single time step Δt, is computed simply by integrating the stellar recycling fraction over the Chabrier IMF for the specific mass bin of dying stars as

$$\Delta M_{rec}(t, \Delta t, Z) = \int_{M(t+\Delta t)}^{M(t)} m f_{rec}(m, Z) \Phi(m)\, dm. \tag{2.19}$$

Here, $\Delta M_{\rm rec}$ represents the total mass returned to the ISM by an SSP of age t and initial metallicity Z during the time interval Δt. The integration bounds, $M(t)$ and $M(t + \Delta t)$, denote the main-sequence turn-off masses at the beginning and the end of the time step, respectively. Together, they define the exact mass range of the stars that reach the end of their lives within that specific interval. Inside the integral, m is the initial stellar mass, $\Phi(m)$ is the IMF, and $f_{\rm rec}(m, Z)$ is the recycling fraction, which specifies the exact portion of the initial stellar mass that is not permanently locked into a compact remnant but is instead successfully expelled.

Crucially, the ejected material is not merely a recycling of the initial SSP composition. The model explicitly tracks the nucleosynthetic production of nine individual elements (H, He, C, N, O, Ne, Mg, Si, and Fe). The net chemical yields—representing the newly synthesized mass for each element—are calculated using detailed theoretical tables: Karakas (2010) for AGB stars, Portinari et al. (1998) for core-collapse supernovae (incorporating the nucleosynthesis networks of Woosley & Weaver 1995) and Thielemann et al. (2003) and Travaglio et al. (2004) for Type Ia supernovae. This accounting elemental yields ensures a self-consistent chemical enrichment of the galactic disc, which in turn directly regulates the metal-line cooling rates discussed in Section 2.4.1.

2.4.4 Supernova Feedback

Supernova (SN) feedback represents the primary regulatory mechanism in star-forming galaxies. By injecting energy on the order of $E_{\rm SN} \sim 10^{51}$ erg per event into the ISM (Agertz et al., 2011, 2013), SNe drive turbulence (Martizzi et al., 2016), heat the gas, and eventually launch galactic-scale outflows and fountains (Li et al., 2017), establishing a self-regulating negative feedback loop that prevents runaway star formation.

The SMUGGLE model explicitly tracks two distinct supernova channels, recognizing their different timescales and nucleosynthetic yields:

- **Core-Collapse (Type II) SNe:** Originating from the collapse of massive stars ($> 8\,\mathrm{M_\odot}$), these events occur on short timescales ($\sim$ few Myr) after the birth of a stellar population. They dominate the prompt injection of mass, thermal energy, and α-elements into the ISM.

- **Thermonuclear (Type Ia) SNe:** Arising from accreting or merging white dwarfs in binary systems, these explosions occur over a much broader delay-time distribution. While they inject significantly less total mass compared to Core-Collapse SNe, they are the primary source of iron-peak elements, fundamentally shaping the chemical evolution of the galaxy over cosmic time.

Capturing these effects in a simulation, however, is not straightforward: the way SNe energy couples to the gas depends critically on how the expansion of an individual remnant is resolved. We therefore first review the evolution of a supernova remnant and then describe how SMUGGLE injects feedback when that evolution cannot be resolved directly.

The numerical challenge of SN remnant evolution

To understand the specific numerical implementation of SN feedback in SMUGGLE, it is crucial to consider the physical evolution of a supernova remnant. Following the initial explosion, the ejecta sweeps up the surrounding ISM, accumulating it into a dense, shocked shell of gas, the *swept-up shell*, that expands outward behind the leading shock front. Before radiative losses become dominant, the remnant undergoes the energy-conserving *Sedov-Taylor* phase: the high internal pressure of the hot bubble performs $P\,dV$ mechanical work on the sweeping shell, accelerating it outwards. This mechanical work significantly amplifies the radial momentum of the shell, typically boosting it by a factor of ~ 10 compared to the initial ejecta momentum (as demonstrated by high-resolution hydrodynamical simulations, e.g., Martizzi et al., 2015). Eventually, the shell keeps expanding until it reaches the *cooling radius*, where the temperature of the shocked gas drops into a regime of highly efficient metal-line emission ($T \sim 10^6$ K). At this point, the remnant enters the *radiative phase*: its thermal energy is rapidly radiated away, but the shell conserves the boosted momentum built up in the previous phase, the *terminal momentum*, and it continues to go through the ISM sweeping up further gas.

The fundamental challenge in global galaxy simulations is resolving the Sedov-Taylor phase. The cooling radius of a remnant is small, and if the local resolution cannot capture it, the $\sim 10^{51}$ erg of a SN end up deposited as thermal energy in a single unresolved gas cell. Because the cooling time of such dense gas is extremely short, this energy is radiated away almost immediately, before the hot shell can perform any mechanical work on the surrounding gas. The result is the *over-cooling problem*: the SN energy is lost to radiation, the momentum boost of the Sedov-Taylor phase is never generated, and the feedback becomes ineffective. SMUGGLE resolves this by assessing the local resolution at each event. When the cell is small enough for the Sedov-Taylor phase to be resolved, the energy is injected thermally, as in the physical case. When it is not, the model switches from thermal-energy deposition to a momentum-based coupling: it directly injects the terminal radial momentum that the swept-up shell would have acquired by the end of the energy-conserving phase. In this way the mechanical impact of the SN, the boosted momentum that survives into the radiative phase, is transferred to the gas correctly, independently of the resolution of the simulation being performed.

Chapter 3

Initial conditions

The previous chapter described the numerical methods, the moving-mesh code AREPO and the SMUGGLE feedback and ISM model, with which the galaxy is evolved. These methods, however, do not specify by themselves the state of the system: they evolve whatever configuration is given as a starting point, which is the set of initial conditions subject of the present chapter. We study an isolated galaxy rather than one assembled self-consistently in a cosmological context where the initial state is an input that we must construct, so how well this state is built sets the right evolution of the galaxy.

The main condition such a starting state must satisfy is dynamical equilibrium. For the hot corona, this means hydrostatic equilibrium: at every radius the inward pull of gravity must be balanced by the outward gradient of the gas pressure. If this balance is not satisfied by the initial conditions (ICs), the net force on the gas is non-zero, and the corona responds by contracting or expanding until it settles into a new equilibrium of its own. During this initial relaxation the density and temperature of the coronal gas change on their own, driven by this imperfection of the starting setup.

The effect we set out to measure is the modification of the corona produced by the heating prescriptions, so if a relaxation transient of the ICs can be as large as the effect of the heating model, a change in the coronal gas could no longer be attributed unambiguously. Setting the corona in accurate hydrostatic equilibrium is therefore the condition for a controlled comparison between models. This chapter describes how such an equilibrium configuration is generated, following the approach of Barbani et al. (2023), which in turn builds on the galaxy-modelling framework of Springel et al. (2005) and the equilibrium models of Hernquist (1990) and Springel (2000).

From a physical and numerical standpoint, the components of our model galaxy cannot be treated uniformly. Within the AREPO and SMUGGLE frameworks, the particles and cells representing the system are broadly categorized into two fundamental classes based on their dynamic behaviour:

- **Collision-less components:** This class includes dark matter particles and stars (both in the bulge and the stellar disc). These components are governed by the collision-less Boltzmann equation, meaning that they interact exclusively through the collective, large-scale gravitational potential of the galaxy. Because direct, point-like particle collisions are physically negligible on galactic scales, their dynamical equilibrium cannot rely on thermal pressure. Instead, it must be sustained either by macroscopic rotational support (centrifugal balance, dominant in the stellar disc) or by random orbital motions (velocity dispersion, particularly important in the bulge and in the dark matter halo).

- **Collisional components:** This class comprises the gas fluid elements, which are discretized on the moving Voronoi mesh. Unlike stars and dark matter, the gaseous phase is subject to gas-dynamical forces, hydrodynamic pressure gradients, shocks, and radiative cooling processes governed by the Euler equations. Consequently, the equilibrium of the gaseous components is fundamentally hydrostatic, where the inward pull of gravity is counterbalanced by the outward isotropic thermal pressure of the gas.

All these components must be initialized in an approximate state of dynamical equilibrium. The setup includes a dark matter halo and a stellar bulge, both described by a Hernquist profile, while the stellar disc and the gaseous disc exhibit exponential surface density profiles, coherent with observations. The spatial distribution and kinematics of these components have been generated following the theoretical framework established by Springel et al. (2005), which builds upon the equilibrium models of Hernquist (1990) and Springel (2000).

3.1 Dark matter halo and bulge

The dark matter halo and the central stellar bulge constitute the massive, spherically symmetric structure of our galaxy model. In modern cosmology, high-resolution N-body simulations have firmly established that the density profiles of dark matter halos virialised in a ΛCDM universe are well-described by the Navarro-Frenk-White (NFW) profile (Navarro et al., 1997), defined as

$$\rho_{\rm NFW}(r) = \frac{\rho_s}{\frac{r}{r_s}\left(1+\frac{r}{r_s}\right)^2}\,, \tag{3.1}$$

where r_s represents the scale radius and ρ_s is the characteristic scale density. This characteristic density is typically parametrized as a function of the halo concentration $c = r_{200}/r_s$

$$\rho_s = \frac{200}{3}\rho_{\rm cr}\frac{c^3}{\ln(1+c) - c/(1+c)}\,, \tag{3.2}$$

where r_{200} is the virial radius—defined as the radius enclosing an average density equal to 200 times the critical density of the Universe,

$$\rho_{\rm cr} = \frac{3H_0^2}{8\pi G}\,, \tag{3.3}$$

with G the gravitational constant and H_0 the present time Hubble parameter.

Despite its cosmological accuracy, the NFW profile presents two severe limitations when used to construct the ICs of isolated galaxy models. First, the total mass obtained by integrating the NFW profile diverges logarithmically as $r \to \infty$. This requires the introduction of arbitrary truncation functions that can compromise the smooth gravitational equilibrium of the outer regions. Second, the NFW profile does not possess an analytical phase-space distribution function, making the statistical sampling of particle velocities mathematically complex and computationally expensive.

Using a Hernquist profile (Hernquist, 1990) instead of a NFW is convenient because it has an analytic distribution function, which in some cases can simplify the sampling of the ICs, while the NFW profile does not; moreover it has a converging mass profile since $\rho(r) \propto r^{-4}$ in the outer parts. The Hernquist density profile is formulated as

$$\rho_{\rm DM}(r) = \frac{M_{\rm DM}}{2\pi} \frac{a}{r(a+r)^3}\,, \tag{3.4}$$

where $M_{\rm DM}$ is the total mass of the halo and a is the characteristic scale length.

We assume that the total mass $M_{\rm DM}$ of the Hernquist profile is exactly equal to the virial mass M_{200} enclosed within the NFW virial radius. Furthermore, we require the two density distributions to be virtually identical in the innermost regions ($r \ll r_{200}$), where the visible galaxy resides. By equating the inner behaviour of both profiles, we derive a direct analytical relation linking the Hernquist scale length a to the NFW scale radius r_s and concentration c (Springel & Hernquist, 2003)

$$a = r_s \sqrt{2\left[\ln(1+c) - \frac{c}{1+c}\right]}\,. \tag{3.5}$$

The bulge, being structurally similar to a small elliptical galaxy or a dense central spherical distribution, is modelled as well via Hernquist profile of the form

$$\rho_b(r) = \frac{M_b}{2\pi} \frac{b}{r(b+r)^3}, \tag{3.6}$$

with b the bulge scale length and M_b is the total bulge mass. The bulge mass is scaled to the halo mass (contained within r_{200}) via the dimensionless bulge fraction $m_{\rm b}$, such that $M_{\rm b} = m_{\rm b} M_{200}$.

To optimize computational efficiency and avoid calculating the orbits of millions of background dark matter particles, the dark matter halo is implemented as a static, analytical external potential in our simulations. This choice significantly accelerates the gravitational solver without affecting the structural evolution of the visible baryonic component. Conversely, the central stellar bulge is implemented as a fully "live" component composed of active collision-less particles. To guarantee its stability, the initial velocities of the bulge particles are derived by solving the Jeans equations under the total gravitational potential of the system and approximating the local velocity distribution with a Gaussian function, as detailed in Section 3.3.

3.2 Stellar and gaseous discs

The baryonic disc of the galaxy is modelled as a superposition of a collision-less stellar component and a collisional gaseous fluid. Following observational constraints for spiral galaxies (Freeman, 1970; Portinari & Salucci, 2010), both components are initialized with an exponential radial surface density profile

$$\Sigma_i(R) = \frac{M_i}{2\pi h_i^2} \exp\left(-\frac{R}{h_i}\right) , \tag{3.7}$$

where $R = \sqrt{x^2 + y^2}$ is the cylindrical radius, M_i is the total mass of the component ($i \in \{*, \mathrm{g}\}$ for stars and gas, respectively), and h_i is its characteristic radial scale length. While the radial distributions are mathematically similar, the vertical structures of the two discs are treated fundamentally differently due to their distinct physical natures.

For the collision-less stellar disc, the vertical density distribution is analytically prescribed using an isothermal sheet model, characterized by a sech^2 profile

$$\rho_*(R, z) = \frac{\Sigma_*(R)}{2z_0} \mathrm{sech}^2\left(\frac{z}{z_0}\right) , \tag{3.8}$$

where the stellar scale height, z_0, is assumed to be radially constant and is set as a free parameter of the model (Binney & Tremaine, 2008). Here, an important distinction has to be made from stellar to gas disc, because the scale height in the latter case cannot be set as a free parameter. Since the gaseous disc is a collisional fluid, the vertical distribution is dynamically determined by the vertical hydrostatic equilibrium against the total gravitational potential of the galaxy

$$-\frac{1}{\rho_g}\frac{\partial P_g}{\partial z} = \frac{\partial \Phi}{\partial z}, \tag{3.9}$$

where ρ_g and P_g are the density and pressure of the gas, respectively, while Φ is the total gravitational potential.

Because the total potential Φ depends not only on the dark matter and stellar distributions but also on the gas itself (self-gravity), this equation does not yield a simple analytical solution. Therefore, the initialization code employs an iterative numerical method: it assumes an initial guess for Φ, computes the corresponding vertical gas density $\rho_{\mathrm{g}}(R, z)$, and then recalculates the updated potential using a hierarchical multipole expansion based on a tree-code. This loop is repeated until convergence is reached, strictly ensuring that the vertical integral of the converged density matches the target exponential surface density

$$\Sigma_g(R) = \int_{-\infty}^{+\infty} \rho_g(R,z)dz \tag{3.10}$$

at any given radius. In equation (3.10), $\Sigma_g(R)$ and $\rho_g(R,z)$ are defined as gas surface density and gaseous disc density, respectively (Springel et al., 2005).

3.3 Velocity assignment

In a multi-component galaxy model, the total gravitational potential Φ is determined by the complex superposition of the halo, the bulge, and the discs. Because of this non-trivial geometry, to calculate a phase-space distribution function for the collision-less components, the ICs are constructed using a local Gaussian approximation of the velocity field, based on the solutions of the Jeans equations (Springel et al., 2005). For each point $r = \sqrt{x^2+y^2+z^2}$, the velocities (v_R, v_ϕ, v_z) in cylindrical coordinates follow a triaxial Gaussian distribution, with each component v_i characterised by its own mean streaming velocity $\langle v_\phi \rangle$ and velocity dispersion σ_i, from which they are randomly sampled. With this method we are not computing the exact equilibrium distribution function, but this approximation has proven to be sufficient for generating ICs in approximate equilibrium (see Springel et al., 2005) and validated numerically in Section 3.6.4.

3.3.1 Velocity assignment for halo and bulge

For spherical non-collisional structures, we assume that the distribution function depends only on the energy and on the z-component of the angular momentum L_z. If the distribution function of the system were known, it would be simple to derive the velocities. However, that is not the case for the model considered, for this reason the Gaussian approximation has been used. The first and second moments of this distribution are determined by solving the Jeans equations in cylindrical coordinates (R, ϕ, z). Assuming an isotropic velocity dispersion tensor in the meridional plane

($\sigma_R^2 = \sigma_z^2 = \langle v_R^2 \rangle = \langle v_z^2 \rangle$), the vertical and radial second-order velocity moments are obtained as

$$\langle v_z^2 \rangle = \langle v_R^2 \rangle = \frac{1}{\rho} \int_z^\infty \rho(R, z) \frac{\partial \Phi}{\partial z} dz, \tag{3.11}$$

where ρ is the local density of the considered component and Φ is the total gravitational potential, which includes the contribution of all galaxy components. The azimuthal second-order moment is then derived as

$$\langle v_\phi^2 \rangle = \langle v_R^2 \rangle + \frac{R}{\rho} \frac{\partial (\rho \langle v_R^2 \rangle)}{\partial R} + v_c^2, \tag{3.12}$$

where $v_c^2 = R \frac{\partial \Phi}{\partial R}$ is the square of the circular velocity given by the total potential.

A mean streaming component in the azimuthal direction $\langle v_\phi \rangle$ can also be present. For the bulge this is assumed to be zero, reflecting its pressure-supported spherical structure, while for the disc it is derived from the epicyclic approximation as detailed in Section 3.3.2. The corresponding velocity dispersion in the azimuthal direction is then given by

$$\sigma_\phi^2 = \langle v_\phi^2 \rangle - \langle v_\phi \rangle^2 . \tag{3.13}$$

3.3.2 Velocity assignment for disc components

To decompose the total azimuthal moment into a mean streaming velocity (rotation) $\langle v_\phi \rangle$ and a random velocity dispersion σ_ϕ^2, the model employs the epicyclic approximation, which is valid for quasi-circular orbits in an axisymmetric potential. This approximation links the radial and azimuthal dispersions through the epicyclic frequency

$$\sigma_\phi^2 = \frac{\sigma_R^2}{\eta^2}, \tag{3.14}$$

where the dimensionless ratio η is defined by the radial derivatives of the total potential

$$\eta^2 = \frac{4}{R} \frac{\partial \Phi}{\partial R} \left(\frac{3}{R} \frac{\partial \Phi}{\partial R} + \frac{\partial^2 \Phi}{\partial R^2} \right)^{-1} . \tag{3.15}$$

The mean azimuthal streaming velocity (accounting for the asymmetric drift) is finally set to

$$\langle v_\phi \rangle = \sqrt{\langle v_\phi^2 \rangle - \frac{\sigma_R^2}{\eta^2}} . \tag{3.16}$$

While the bulge is typically assumed to have zero net rotation ($\langle v_\phi \rangle = 0$), the stellar disc essentially relies on this streaming component for its centrifugal support (Hernquist, 1993; Garma-Oehmichen et al., 2021).

Conversely, the kinematics of the collisional gaseous disc are not governed by the collision-less Jeans equations, but rather by the stationary Euler equations. The gas rotates smoothly, supported centrifugally by rotation and radially by its internal thermal pressure gradient and therefore is assigned a rotation velocity following

$$v_{\phi,\mathrm{gas}}^2 = R\left(\frac{\partial \Phi}{\partial R} + \frac{1}{\rho_\mathrm{g}}\frac{\partial P_g}{\partial R}\right), \tag{3.17}$$

where ρ_g is the gas density of the considered component and Φ is the total gravitational potential, while P_g is the gas pressure.

3.4 Disc temperature and metallicity

Finally, to complete the thermodynamic initialization, the gaseous disc is assigned with a uniform initial temperature of 10^4 K and a homogeneous solar metallicity ($Z = Z_\odot$[1]), for simplicity. They start as constant values but as soon as the dynamical simulation begins, the SMUGGLE sub-grid model explicitly tracks the multiphase evolution of the gas and its temperature within a range of $[10^2 - 10^6]$ K. The establishment of a multiphase ISM is a consequence of several of the mechanisms mentioned in Section 2.4, such as stellar feedback, cooling and heating, but also the mixing and accretion typical of a coronal gas. The hot corona is better explained in the following section, in particular how it is modelled within the setup.

3.5 Hot Galactic Corona

The galaxy model is embedded within a hot CGM, commonly referred to as the galactic corona. This extended gas reservoir plays a crucial role in the long-term evolution of the system, acting both as a continuous fuel supply for star formation over cosmological timescales and as the background environment into which galactic outflows and supernova-driven winds propagate. For the initialization of this specific

[1] In our simulations $Z_\odot = 0.0127$ Asplund et al. (2009)

component, we rely on the analytical framework recently developed by Barbani et al. (2023, 2025) taking into account Milky Way observational constraints.

The spatial distribution of the coronal gas is explicitly modelled using a β-profile (Cavaliere & Fusco-Femiano, 1978). The radial density profile is analytically defined as

$$\rho_{\rm cor}(r) = \rho_0 \left[1 + \left(\frac{r}{r_c}\right)^2\right]^{-\frac{3}{2}\beta} , \tag{3.18}$$

where ρ_0 is the central density, r_c is the core radius, and r is the spherical radius. In this setup, we adopt $\beta = 2/3$ and a core radius $r_c = 8\,\mathrm{kpc}$ (see Barbani et al., 2023). These structural parameters yield a total coronal mass of $3 \times 10^{10}\,\mathrm{M}_\odot$ within the virial radius, in excellent agreement with the empirical estimates for the Milky Way (Grcevich & Putman, 2009; Anderson & Bregman, 2010; Gatto et al., 2013; Salem et al., 2015; Das et al., 2020; Putman et al., 2021).

Unlike the static approximations often used in idealized setups, our model incorporates a rotating corona. Observations suggest that the lower CGM partially co-rotates with the galactic disc, although with a significant velocity lag (Marinacci et al., 2011; Miller & Bregman, 2013; Afruni et al., 2022). To capture this kinematic structure, the coronal gas is initialized with an azimuthal rotational velocity that is a fixed fraction of the local circular velocity

$$v_\phi(R, z) = \alpha \sqrt{R \frac{\partial \Phi}{\partial R}} , \tag{3.19}$$

where Φ is the total gravitational potential and R is the cylindrical radius. Setting the proportionality constant to $\alpha = 0.4$ ensures a rotational velocity of approximately $90\,\mathrm{km\ s^{-1}}$ near the disc plane ($z \approx 2\,\mathrm{kpc}$), which is expected for the rotation of the corona with respect to the disc with a lag of $\sim 80 - 120$ km s^{-1} (Marinacci et al., 2011; Afruni et al., 2022).

For what concerns the thermodynamic structure of the corona the starting point is assuming hydrostatic equilibrium. In the case of a purely static corona, the temperature profile is obtained by integrating the spherical hydrostatic equation over the total cumulative mass profile $M(r)$

$$T(r) = \frac{\mu m_p}{k_B} \frac{1}{\rho_{\rm cor}(r)} \int_r^\infty \rho_{\rm cor}(r') \frac{G M(r')}{r'^2} dr' , \tag{3.20}$$

where k_B is the Boltzmann constant, μ is the mean molecular weight, m_p is the proton mass, while ρ_{cor} radial number density defined at the beginning of this section. However, because our ICs includes rotation, this standard spherical equilibrium is modified to account for centrifugal support. The temperature profile is therefore computed assuming hydrostatic equilibrium within an effective gravitational potential, $\Phi_{\rm eff}$, defined as

$$\Phi_{\rm eff}(R,z) = \Phi(R,z) - \int_{\infty}^{R} \frac{v_\phi(R')^2}{R'} dR' \,. \tag{3.21}$$

By substituting $\Phi_{\rm eff}$ into the generalized hydrostatic equilibrium equation for an ideal gas, the initialization algorithm generates a stable temperature gradient that peaks in the dense central regions ($\sim 10^6$ K) and smoothly declines toward the virial radius. Finally, to complete the initialization and track the subsequent chemical enrichment driven by the SMUGGLE model, the coronal gas is assigned with a uniform, sub-solar metallicity of $0.1\,Z_\odot$ (Tosi, 1988; Bogdán et al., 2017).

For further details about the general implementation of the ICs we refer the reader to Barbani et al. (2023).

3.6 Sampling the ICs

After having established the theoretical framework of dark-matter halo, bulge, disc, and the hot corona, composing the initial setup, the subsequent step involves the translation of these continuous analytical models into discrete ICs suitable for numerical simulations. Generating an accurate particle realization of the theoretical profiles is crucial, as an unstable realization would compromise the entire evolution, as discussed at the beginning of this chapter. To achieve this, the system is sampled within the phase space, requiring the explicit assignment of a six-dimensional coordinate vector (position and velocity) to each collision-less particle (dark matter and stars). Furthermore, for the hydrodynamic component (gas), the initialization must also include thermodynamic and chemical scalar fields, specifically the internal energy and the metallicity, which allow to track the cooling rates and the multiphase gas evolution of the surrounding ISM.

3.6.1 Position sampling

The particle positions for each galactic component are initialized using the inverse transform sampling method. This is a standard technique used to generate random numbers following a specific probability density function $f_X(X)$. The idea is that a random number is interpreted as an enclosed-mass fraction, and the radius at which the profile encloses that fraction is the sampled position, so denser regions receive proportionally more particles. The procedure relies on the cumulative distribution function (CDF), defined as

$$F_X(X) = \int_{-\infty}^{X} f_X(\tau)\, d\tau \,. \tag{3.22}$$

By definition, $F_X(X)$ is a monotonically non-decreasing function bounded between 0 and 1. The method consists of drawing a random number u from a uniform distribution $\mathcal{U}(0,1)$, setting $u = F_X(X)$, and analytically or numerically computing the inverse function $X = F_X^{-1}(u)$, we obtain a value of X distributed as $f_X(X)$. In the context of galactic initial conditions, $f_X(X)$ corresponds to the physical density profile of the component, and $F_X(X)$ represents its normalized cumulative mass profile.

Spherical systems

The dark matter halo, the bulge, and the hot corona are all spherically symmetric, and their positions are sampled with the same procedure: the radial coordinate is drawn by inverting the normalized cumulative mass profile of each component, while the angular coordinates are assigned isotropically. The components differ only in the profile that enters the cumulative mass: a Hernquist profile for the halo and the bulge 3.1, and a β-model for the corona 3.5.

For the Hernquist profile the inversion is analytic. The normalized cumulative mass profile is

$$F_M(r) = \left(\frac{r}{r+a}\right)^2 . \tag{3.23}$$

By equating this to a uniform random variable u, the radial coordinate can be explicitly isolated

$$u = \left(\frac{r}{r+a}\right)^2 \quad \Longrightarrow \quad r = \frac{a\sqrt{u}}{1-\sqrt{u}} \,. \tag{3.24}$$

This fixes the radial distance r of the dark-matter halo and of the bulge. The same inversion is applied to the cumulative mass profile of the corona's β-model.

In all three components the angular coordinates, the azimuthal angle ϕ and the polar angle θ, are then sampled isotropically as

$$\phi = 2\pi v\,, \tag{3.25}$$

$$\theta = \arccos(2w - 1)\,, \tag{3.26}$$

where v and w are two additional independent random numbers uniformly distributed between 0 and 1. Finally, the corresponding Cartesian coordinates (x, y, z) are obtained via the standard spherical transformation.

Axisymmetric systems

For the disc components, the geometry is cylindrical (R, ϕ, z). The vertical position z for the stellar disc follows a sech^2 profile, which can also be inverted analytically

$$z = \frac{z_0}{2} \ln\left(\frac{u}{1-u}\right)\,, \tag{3.27}$$

with u drawn from $\mathcal{U}(0, 1)$.

For the gaseous disc, the vertical density profile is determined numerically by integrating the hydrostatic equilibrium equations on the grid (see Section 3.2). The resulting CDF is then constructed by cumulative summation over the grid points and inverted via linear interpolation to sample the vertical particle positions.

The radial mass distribution of the exponential discs cannot be inverted analytically either. The radius R is therefore determined using the Newton-Raphson method. Given the normalized cumulative radial mass profile, we seek the root x_n of the equation $f(x) = 0$, defined iteratively as

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}\,. \tag{3.28}$$

For the stellar disc, defining the dimensionless variable $x = R/h_\star$, the functions to evaluate are

$$f(x_n) = -1 + (1 + x_n)\exp(-x_n) + u\,, \tag{3.29}$$

$$f'(x_n) = -x_n \exp(-x_n)\,. \tag{3.30}$$

Starting from an initial guess (e.g., $x_1 = 1$), the iteration proceeds until convergence is reached ($|x_{n+1} - x_n|/x_n < 10^{-7}$).
For the gaseous disc, the profile is truncated at a maximum radius of 10 h_g. Consequently, the probability distribution must be renormalized to account for this truncation. Once the cylindrical radius $R = x_n h_g$ and the height z are determined, the azimuthal angle ϕ is uniformly sampled, allowing the derivation of the Cartesian coordinates for the disc particles.

3.6.2 Velocity sampling

For each particle, positioned as detailed in Section 3.6.1, the subsequent step is the initialization of their kinematic properties. To ensure dynamical equilibrium, the velocity components of the collision-less particles (dark matter, bulge, and stellar disc) must be sampled from Gaussian distributions, characterized by the mean streaming velocities and the velocity dispersions $(\sigma_R, \sigma_z, \sigma_\phi)$ derived in Sections 3.3.

First, the velocity dispersions in the radial, vertical, and azimuthal directions are computed by solving the Jeans equations on a 2D logarithmic grid in the (R, z) plane. The grid consists of 512×512 points, covering a total box size of approximately 857 kpc. Logarithmic spacing is used so that the inner regions of the galaxy, where the density varies rapidly, are sampled more finely than the outskirts. The precise dispersion values at each individual particle position are then assigned via bilinear interpolation from the grid.

Then, to sample the velocities from a normal distribution, we employ the polar form of the Box-Muller transform (Box & Muller, 1958). This algorithm generates pairs of independent, standard, normally distributed random numbers from a source of uniformly distributed variables. Initially, two pseudo-random numbers y_1 and y_2 are extracted from a uniform distribution $\mathcal{U}(-1, 1)$. We evaluate the squared radius $r = y_1^2 + y_2^2$; if $r = 0$ or $r \geq 1$, the pair is discarded to ensure the values strictly lie within the unit circle, and a new pair is drawn. If accepted, the standard Gaussian variable ξ_0 (with mean $\mu = 0$ and variance $\sigma^2 = 1$) is computed as

$$\xi_0 = y_{1,2}\sqrt{\frac{-2\ln r}{r}}\,. \tag{3.31}$$

The Box-Muller transform naturally yields pairs of independent standard Gaussian variables at each call; by repeating the procedure we obtain one independent standard variable per spatial direction, denoted ξ_R, ξ_z and ξ_ϕ, each drawn from a Gaussian of zero mean and unit variance. To assign the physical velocity components, the standard variable ξ_0 is scaled by the local velocity dispersion and shifted by the mean streaming velocity. For instance, the cylindrical velocity components for the stellar disc particles are assigned as

$$v_R = \xi_R \sigma_R, \tag{3.32}$$

$$v_z = \xi_z \sigma_z, \tag{3.33}$$

$$v_\phi = \xi_\phi \sigma_\phi + \langle v_\phi \rangle. \tag{3.34}$$

The mean radial and vertical velocities are uniformly set to zero ($\langle v_R \rangle = \langle v_z \rangle = 0$), while the azimuthal direction carries the free-streaming rotation $\langle v_\phi \rangle$ as defined in Eq. (3.16). This general prescription also applies to the dark matter halo, whereas for the classical bulge, the net rotation is assumed to be zero ($\langle v_\phi \rangle = 0$).

Different conditions apply to the gaseous components. For the rotationally supported gaseous disc, the random dispersion component is negligible compared to the bulk motion; therefore, the velocities are not sampled from a Gaussian distribution, but the azimuthal velocity is directly imposed according to the rotational equilibrium condition given by Eq. (3.17); the radial and vertical components are set to zero ($v_R = v_z = 0$), consistently with a disc in rotational support. For the hot corona, the radial and vertical velocities are set to zero, the system being supported against gravity by thermal pressure, while the azimuthal component follows the rotational prescription of Eq. (3.21).

As a final step, the assigned cylindrical velocities (v_R, v_ϕ, v_z) are transformed into the Cartesian frame (v_x, v_y, v_z) of the simulation box using the corresponding spatial coordinates (x, y) and the cylindrical radius R as

$$v_x = v_R \frac{x}{R} - v_\phi \frac{y}{R}, \tag{3.35}$$

$$v_y = v_R \frac{y}{R} + v_\phi \frac{x}{R}, \tag{3.36}$$

$$v_z = v_z. \tag{3.37}$$

3.6.3 Metallicity sampling

The last property to assign to the gas particles is the metallicity. This is an important quantity for galaxy evolution: its primary effect is to make radiative cooling more efficient, influencing the gas to cool, collapse and form stars. In the ICs the gas has homogeneous metallicity, set to a different value in the two gaseous components: $Z = Z_\odot$ in the disc and $Z = 0.1\, Z_\odot$ in the hot corona.

The metallicity is treated as a *passive scalar*: each gas cell carries its own metallicity Z_i, advected by the fluid flow without feeding back on the dynamics, according to the continuity equation

$$\frac{\partial(\rho Z)}{\partial t} + \nabla \cdot (\rho Z \mathbf{v}) = 0\,. \tag{3.38}$$

Beyond the global metallicity, the code also tracks the abundances of nine chemical elements (H, He, C, N, O, Ne, Mg, Si, and Fe). Different metallicities correspond to different abundances of the individual elements, obtained in each cell i by interpolating linearly between the primordial and solar mixtures:

$$f_{X,i} = f_{X,\mathrm{prim},i} + \frac{Z_i}{Z_\odot}\left(f_{X,\odot,i} - f_{X,\mathrm{prim},i}\right)\,, \tag{3.39}$$

where $f_{X,\mathrm{prim}}$ and $f_{X,\odot}$ are the primordial and solar abundances of element X, listed in Table 3.1, and $Z_\odot = 0.0127$ is the solar metallicity. In each cell i the abundances sum to $\sum_X f_{X,i} = 1$, ensuring mass conservation.

Table 3.1: Primordial abundances f_{prim} and solar abundances $f_\odot$ of the nine tracked elements (adapted from Barbani et al. (2023)). These are mass fractions, i.e. the ratio between the mass of a given element X and the total mass of all elements. Solar abundances are taken from Asplund et al. (2009).

	H	He	C (10^{-3})	N (10^{-3})	O (10^{-3})	Ne (10^{-3})	Mg (10^{-3})	Si (10^{-3})	Fe (10^{-3})
f_{prim}	0.76	0.24	0	0	0	0	0	0	0
$f_\odot$	0.7388	0.2485	2.4	0.7	5.7	1.2	0.7	0.7	1.3

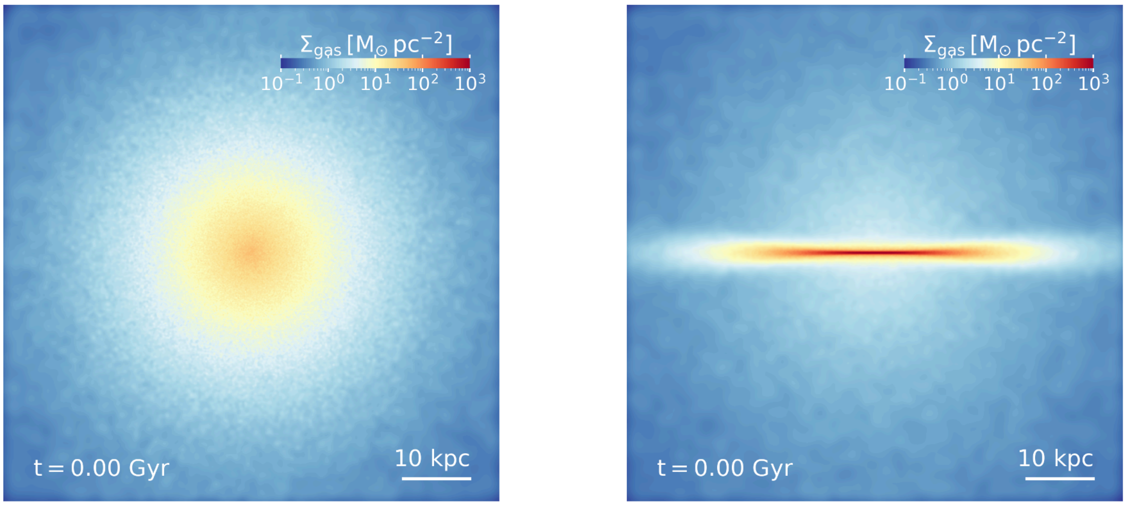

Figure 3.1: Projected gas surface density $\Sigma_{\rm gas}$ of the ICs, shown face-on (left) and edge-on (right). Each panel spans 75×75 kpc, shows a white bar that marks 10 kpc reference scale and the time at which the snapshot was taken, $t = 0$ Gyr. The colour scale is logarithmic and covers $\Sigma_{\rm gas} = [10^{-1}, 10^{3}]\, M_\odot\,{\rm pc}^{-2}$ range.

3.6.4 Numerical validation of the ICs

Before tracking the evolution of the galaxy and exploring the impact of different heating models in an isolated realization of a Milky Way-like galaxy, it is necessary to verify that the initialization algorithms have correctly translated the analytical profiles of the galaxy model into a discrete and stable particle distribution. We perform this check on the initial snapshot (Figure 3.1), which captures the system at $t = 0$, prior to any gravitational or hydrodynamical integration. The face-on gas projection shows a smooth and axisymmetric density distribution declining with radius within a $\Sigma_{\rm gas} \sim 1 - 10^2\, M_\odot\,{\rm pc}^{-2}$ range . The edge-on gas projection instead, shows a radially extended disc, symmetric about the mid-plane and reaching central surface densities of $\Sigma_{\rm gas} \sim 10^2 - 10^3\, M_\odot\,{\rm pc}^{-2}$ embedded in the diffuse hot corona that surrounds it ($\Sigma_{\rm gas} \sim 1 - 10\, M_\odot\,{\rm pc}^{-2}$). We conclude that the ICs reproduce the intended disc-corona structure, with no numerical artifacts, providing a stable starting point for the simulations that are analysed in the remaining parts of this work.

Chapter 4

The baseline scenario: No-Heat model

The four models analysed in this work are designed to investigate the role of different heating prescriptions, with respect to a baseline scenario where no heating is present (besides the process already included in SMUGGLE), in regulating the thermodynamic evolution of the hot CGM in a Milky Way-like galaxy. Each model builds upon the same set of ICs described in Chapter 3, following the framework established by Barbani et al. (2023), and differs exclusively in the heating prescription applied to the coronal gas. This controlled approach allows us to isolate the physical effects of each prescription and to directly compare their impact on the coronal depletion, the SFR, the inflow/outflow balance and the thermodynamic features of the gas.

The **No-Heat** model serves as the baseline scenario: no heating is applied to the corona, which is therefore left to cool radiatively without any compensating energy input. It defines the cooling of the hot CGM against which the heating models are measured. The **Heat 1** and **Heat 2** models share the same cooling-compensation mechanism, in which thermal energy is injected to offset the local radiative losses of the corona; they differ only in the timescale over which the same energy budget is delivered. Heat 1 distributes the energy continuously, in small amounts at each local timestep, so that the corona is heated smoothly. Heat 2 redistributes the same budget over each global timestep, delivering it in fewer, larger injections whose local energy peaks can potentially drive mechanical effects on the surrounding gas. Comparing the two isolates the role of *how* the energy is delivered at fixed total budget. Finally, the **NF Bubbles** model adopts a different approach, in which the heating is tied to an AGN feedback model through a bubble-injection prescription following Nulsen & Fabian (2000) and Grand et al. (2017). Among the models considered, it is the most directly motivated by a physical feedback mechanism (radio-mode AGN), although its energy budget is set from the halo X-ray luminosity through a scaling relation rather than from self-consistent black-hole accretion.

These four models are analysed in the following chapters. In particular, the No-Heat baseline is the subject of the present chapter. The two cooling-compensation models, Heat 1 and Heat 2, are discussed in Chapter 5 while the NF Bubbles model is presented in Chapter 6.

4.1 The No-Heat simulation

The No-Heat model serves as the reference scenario. Its primary objective is to establish a physically motivated starting point to evaluate the regulation capabilities of stellar feedback when acting in isolation, against which the effects of the heating models previously described are then compared. Therefore, we first simulate the galaxy without any additional external heating mechanism for the coronal gas (other than the ones already included in the SMUGGLE implementation, see Section 2.4).

Without an external heating source to compensate for the continuous energy dissipation, the hot CGM is subject to radiative cooling. As the coronal gas radiates away its thermal energy, its pressure drops, breaking the initial hydrostatic equilibrium and driving a continuous infall of gas from the corona onto the galactic disc. This process, is known as *cooling flow.* The resulting gas supply sustains a star formation rate of $\sim$ 4-5 $\mathrm{M_\odot\,yr^{-1}}$ in the No-Heat run (Chapter 4.1), a factor of ~ 2 above the $\sim$ 1-3 $\mathrm{M_\odot\,yr^{-1}}$ observed in quiescent spirals like the Milky Way (Putman et al., 2021; Kennicutt, 1998; Bland-Hawthorn & Gerhard, 2016).

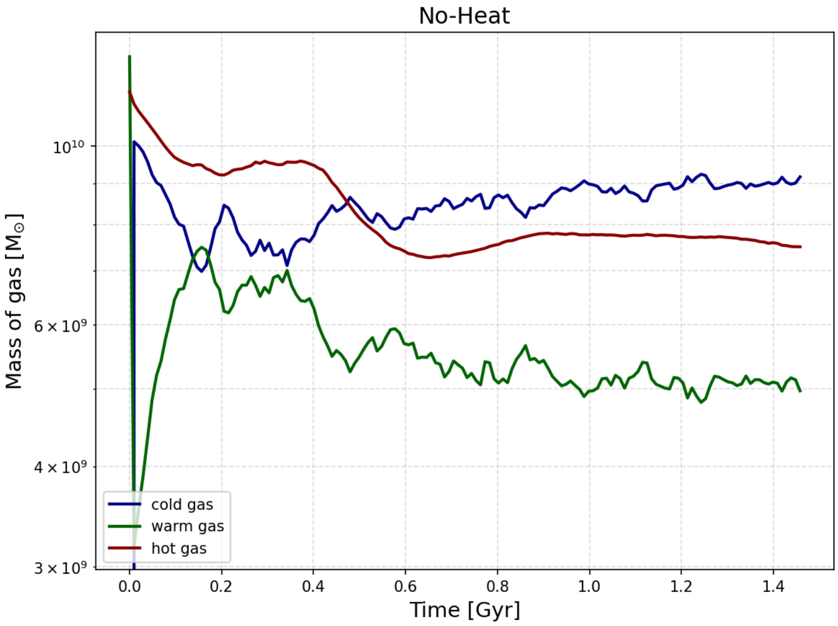

Figure 4.1: Time evolution of the gas mass partitioned by temperature phase for the No-Heat model, computed within a spherical region of $r < 100$ kpc . The cold ($T < 10^4$ K, blue), warm ($10^4 \leq T \leq 5 \times 10^5$ K, green) and hot ($T > 5 \times 10^5$K, red) phases are tracked separately over $t \simeq 1.5$ Gyr of evolution.

4.2 Mass evolution of the gas

The mass evolution of the three gas phases, computed within the inner hot CGM ($R < 100$ kpc), reveals the macroscopic consequences of the radiative cooling of the corona in the absence of any external heating (Figure 4.1). At $t = 0$ the hot ($T > 5 \times 10^5$ K) phase dominates the baryonic budget, with initial masses of $\sim 10^{10}\, M_\odot$. This phase shows the clearest signature of the cooling flow. Its mass drops from $\sim 10^{10}\, M_\odot$ to $\sim$ 7-8 $\times\, 10^9\, M_\odot$ already by $t \approx 0.6$ Gyr, and then settles onto a plateau, after $\sim$ 0.8 Gyr. This depletion reflects the continuous radiative losses of the corona, which are not compensated by any heating source in this run, besides the ones integrated into SMUGGLE.

The cold and warm phases, after a brief initial transient ($\sim$ 0.2 Gyr), evolve in opposite directions: the cold gas grows from $\sim 7 \times 10^9$ to $\sim 9 \times 10^9\, M_\odot$, while the warm gas decreases from $\sim 7 \times 10^9$ to $\sim 5 \times 10^9\, M_\odot$. The comparable magnitude of these two changes is consistent with gas cooling out of the warm phase and accumulating in the cold disc.

Crucially, the cold gas does not grow indefinitely, despite the continuous supply from the cooling corona: it settles to a roughly constant level, suggesting that the disc acts as an efficient processing engine rather than a passive sink for the infalling coronal material. This balance is maintained by two competing mechanisms of the SMUGGLE model. First, as the infalling gas reaches the higher densities of the mid-plane (visually confirmed by the cold spiral structures in the temperature maps, Figure 4.4), it forms molecular hydrogen and is consumed by star formation, transferring mass from the gas to the collision-less stellar component. Second, the resulting supernovae reheat and expel a fraction of the cold gas out of the disc, as shown by the outflows in Figure 4.9 and the density maps (Figures 4.2, 4.3).

4.3 Morphological evolution of the gas disc

Besides the evolution of the gas mass, we also investigate how the evolution of the corona leaves a progressive imprint on the galaxy structure and shape. The galaxy already hosts a rotationally-supported gas disc from the ICs; the cooling flow does not form this disc but feeds it. As the cooling coronal gas reaches the galactic plane, it settles into the already existing disc. This build-up, and the structural evolution that the accretion drives, are captured by the projected density maps at different times (Figures 4.2 and 4.3).

4.3.1 Surface Density maps

The face-on projections (Figure 4.2) reveal the progressive structural evolution of the disc. The spiral and clumpy structure of the disc emerges from the combined effects of gas cooling, star formation and stellar feedback; as we explained in Section 2.4, this ensemble of processes is already implemented in SMUGGLE prescription for the evolution of the ISM. The continuous accretion of cooling coronal gas supplies additional material to the disc, sustaining its growing surface density over time. At $t \sim 0.30$ Gyr the disc already shows forming spiral arms, with surface densities in the range $10 - 10^2\, M_\odot\, \mathrm{pc}^{-2}$, while the inter-arm regions remain at $\sim 1 - 10\, M_\odot\, \mathrm{pc}^{-2}$. By $t \sim 1.32$ Gyr the disc has settled into a more organized configuration, with completely formed spiral arms and a dense central core.

The edge-on projections (Figure 4.3) show two distinct components. A geometrically thin, dense mid-plane which persists throughout the entire simulation, representing the star-forming disc, where the gas reaches surface densities of $10^2 - 10^3\,M_\odot\,\mathrm{pc}^{-2}$; and a gas layer surrounding the disc and extending to $|z| \sim 5$ kpc above and below the mid-plane. This extra-planar component is sustained by gas lifted from the disc plane by supernova feedback, a key ingredient of the galactic fountain cycle: gas is ejected from the mid-plane, rises into the CGM, and subsequently falls back, producing the simultaneous presence of outflowing and inflowing gas visible in the inflow/outflow profiles (Figure 4.9).

4.3.2 Mass-weighted Temperature maps

The mass-weighted temperature maps (Figures 4.4 and 4.5) show the different structure of the disc and the hot corona from a different point of view with respect to the surface density maps. In these figures, we can distinguish that the extended CGM at $|z| > 5$ kpc is dominated by hot gas at $T \sim 10^6$ K throughout the entire simulation, while the disc shows the filaments and the spirals at a much lower temperature.

In the face-on projections, the disc exhibits: cold filaments at $T \sim 10^3 - 10^4$ K that trace the spiral arms and constitute the actual star-forming sites, embedded in a warmer inter-arm medium at $T \sim 10^4 - 10^5$ K.

The edge-on projections confirm the vertical thermal stratification of the galaxy. A thin, cold mid-plane ($T \sim 10^3 - 10^4$K) persists throughout the simulation, surrounded by a thickening layer of warm gas at $T \sim 10^5$ K extending to $|z| \sim 5$ kpc, representing the gas sustained by the galactic fountain cycle (Fraternali, 2017; Barbani et al., 2023, 2025). Moreover, the circumgalactic region of hot gas at $T \sim 10^6$K is more clear and visible from the edge-on perspective.

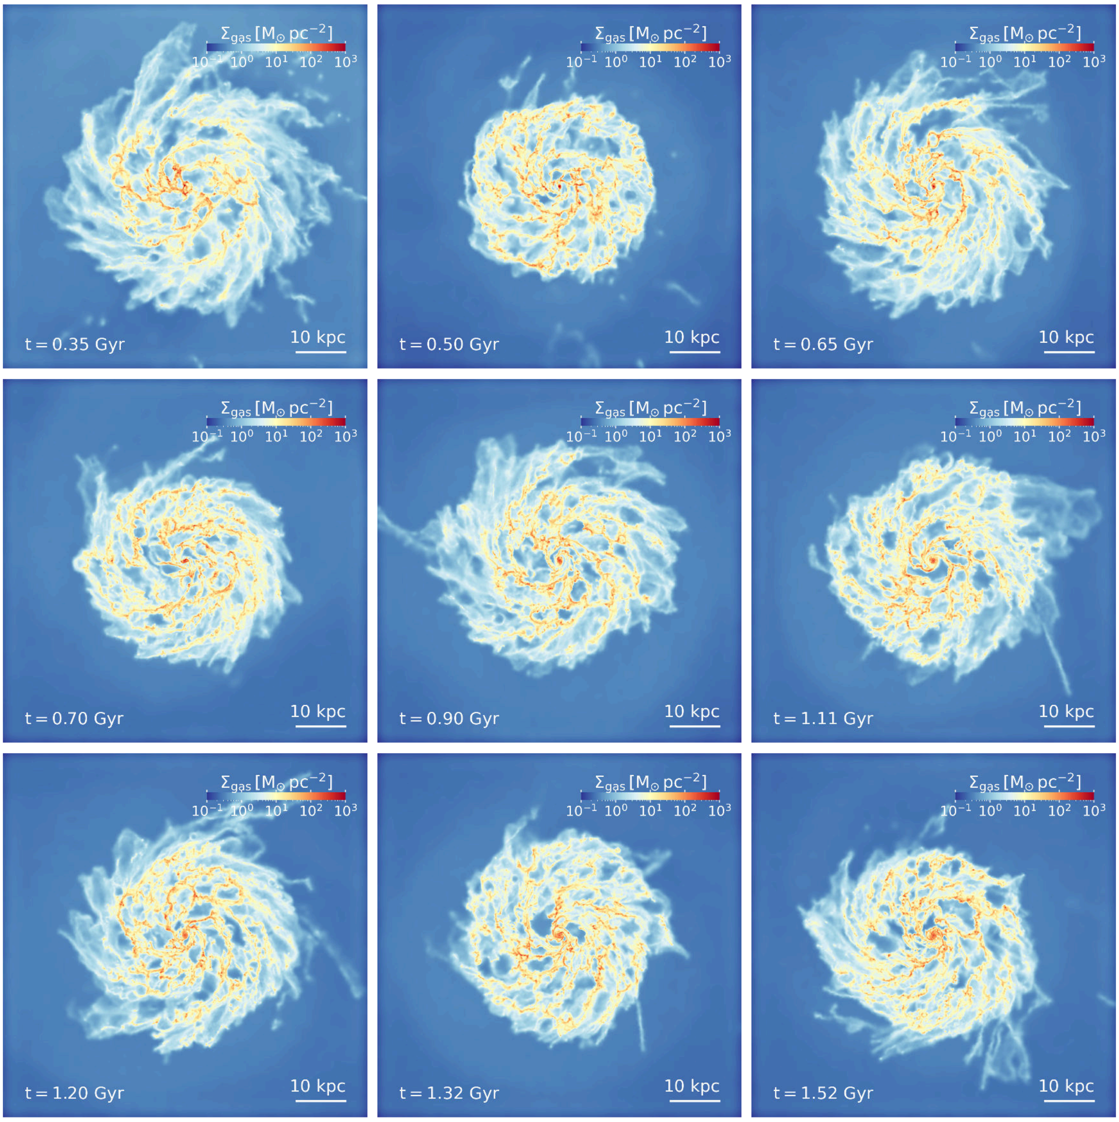

Figure 4.2: Face-on projected gas surface density $\Sigma_{\rm gas}$ of the No-Heat model run from $t = 0.30$ to $t = 1.52$ Gyr to show the evolution of the simulated galaxy. The colour scale is logarithmic and is identical across all panels to allow a direct comparison between different times. The white bar indicates a physical scale of 10 kpc. It is noticeable the evolution of the disc structure and filaments through time.

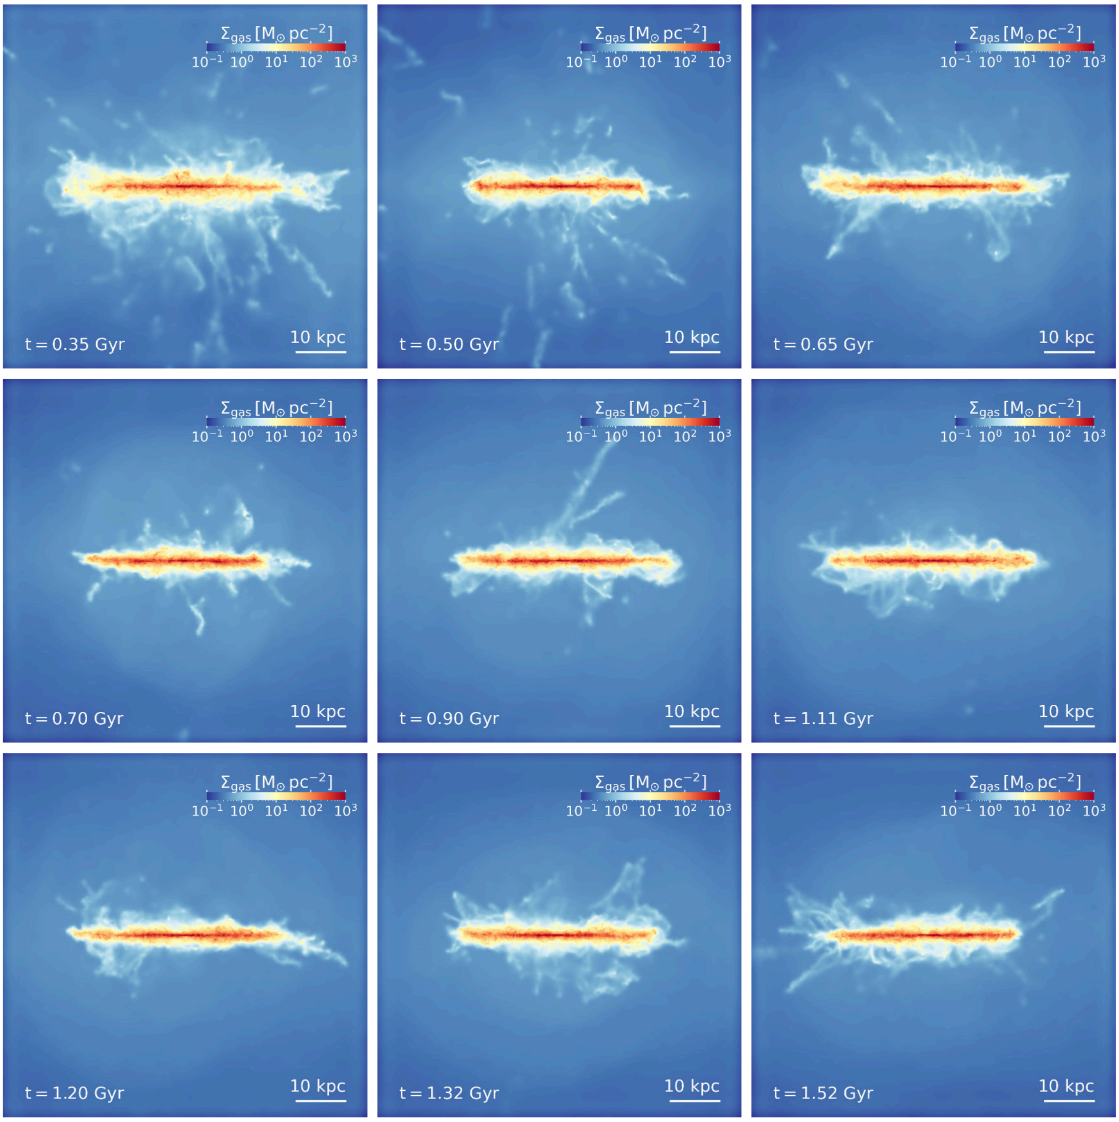

Figure 4.3: Edge-on projected gas surface density $\Sigma_{\rm gas}$ of the No-Heat model run from $t = 0.30$ to $t = 1.52$ Gyr, same evolution time to be consistent with Figure 4.2. The colour scale is logarithmic and is identical across all panels to allow a direct comparison between different times. The white bar indicates a physical scale of 10 kpc. These maps show a thin dense disc evolving together with filaments.

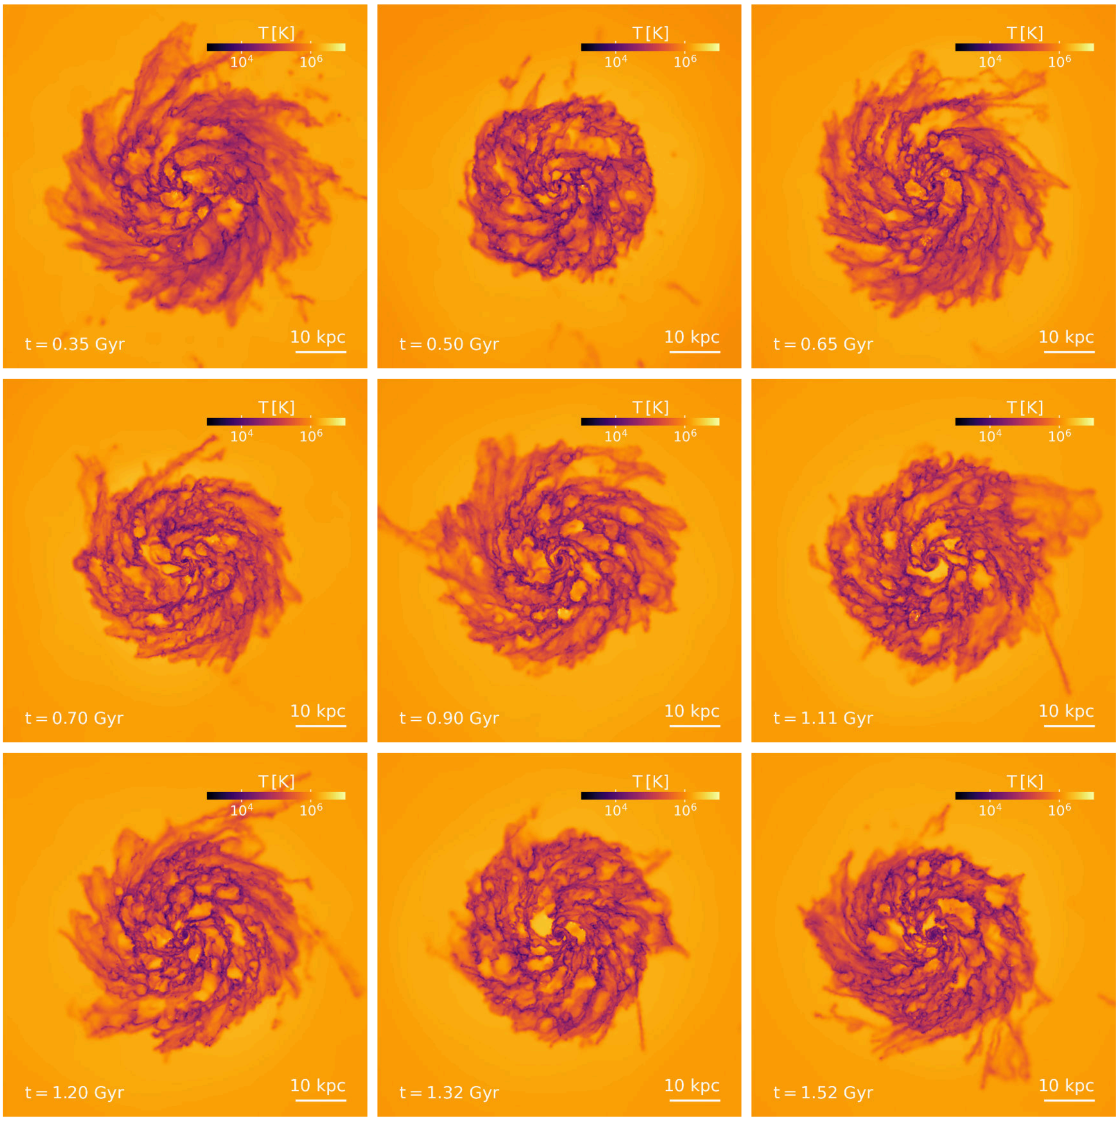

Figure 4.4: Face-on mass-weighted temperature maps of the No-Heat model run from $t = 0.30$ to $t = 1.52$ Gyr, same evolution time to be consistent with Figure 4.2. The colour scale is logarithmic and is identical across all panels to allow a direct comparison between different times. The white bar indicates a physical scale of 10 kpc. It is shown the cold structure structure of the disc embedded in the hotter corona.

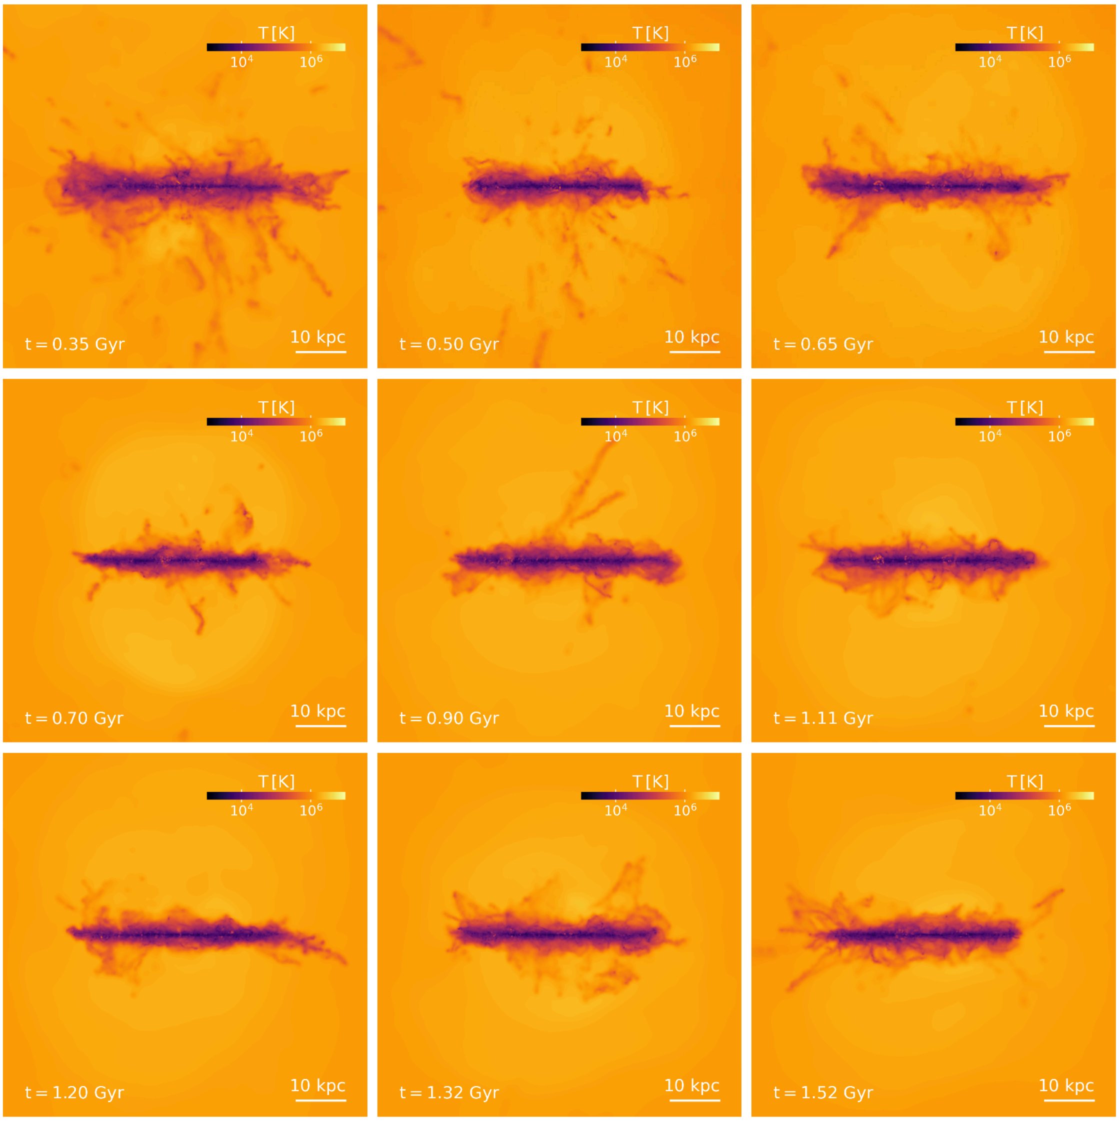

Figure 4.5: Edge-on mass-weighted projected gas temperature T of the No-Heat model run from $t = 0.30$ to $t = 1.52$ Gyr. The colour scale is logarithmic and is identical across all panels to allow a direct comparison between models. The white bar indicates a physical scale of 10 kpc. The dense and cold disc evolves together with filaments both embedded into the hot coronal gas.

4.4 Thermodynamic and Multiphase ISM

4.4.1 Phase diagrams

The phase diagrams, weighted by gas mass, complement the maps displayed previously by showing how the gas is distributed across the thermodynamic regimes and how it is redistributed by cooling and feedback (Figure 4.6). The dashed lines partition the density–temperature plane into the regions separating the coronal gas, the dense star-forming gas, and the intermediate phases (see also Marinacci et al. 2019 and Barbani et al. 2023). The plane is divided into six regions: (a) the hot, rarefied coronal gas ($\log n < -2$, $\log T > 4.5$); (b) the denser coronal gas ($\log n > -2$, $\log T > 4.5$); (c) and (d) the outer disc gas at intermediate temperatures ($\log n > -2$, $3.5 < \log T < 4.5$); (e) the cold, dense star-forming gas ($\log n > 2$, $\log T < 3.5$); and (f) the cold but diffuse gas ($\log n < 2$, $\log T < 3.5$), which lies below the density threshold required for star formation.

At $t \sim 0.35$ Gyr after the beginning of the simulation, the coronal gas occupies region (a), at low density and high temperature. The most prominent feature is a horizontal accumulation at $\log T \simeq 4.1$ extending towards high densities. This arises from the shape of the radiative cooling function: near $T \sim 10^4$ K hydrogen recombines and the cooling function drops steeply, so that below this temperature there are far fewer efficient cooling channels and the gas loses its internal energy much more slowly. Gas cooling from the hot corona therefore reaches $T \sim 10^4$ K quickly but then cools only very slowly beyond it, and accumulates at this temperature. This stalling point acts as a bottleneck in the cooling process, a temperature at which the flow of gas towards the cold, dense regime is throttled by the collapse of the cooling efficiency, visible as the pile-up at $\log T \simeq 4.1$ in the phase diagram.

A diagonal branch then connects this stall to the star-forming region (e) at $\log n > 2$ and $\log T < 3.5$. Comparing the two snapshots, the coronal section in region (a) remains substantially populated even after ~ 1.2 Gyr. This is consistent with the fact that only the inner corona ($r \lesssim 16$ kpc) has a cooling time shorter than the simulated time span, while the outer corona retains its thermal energy (Barbani et al., 2023). The diffuse plume extending from the dense/cold phase, regions (c)

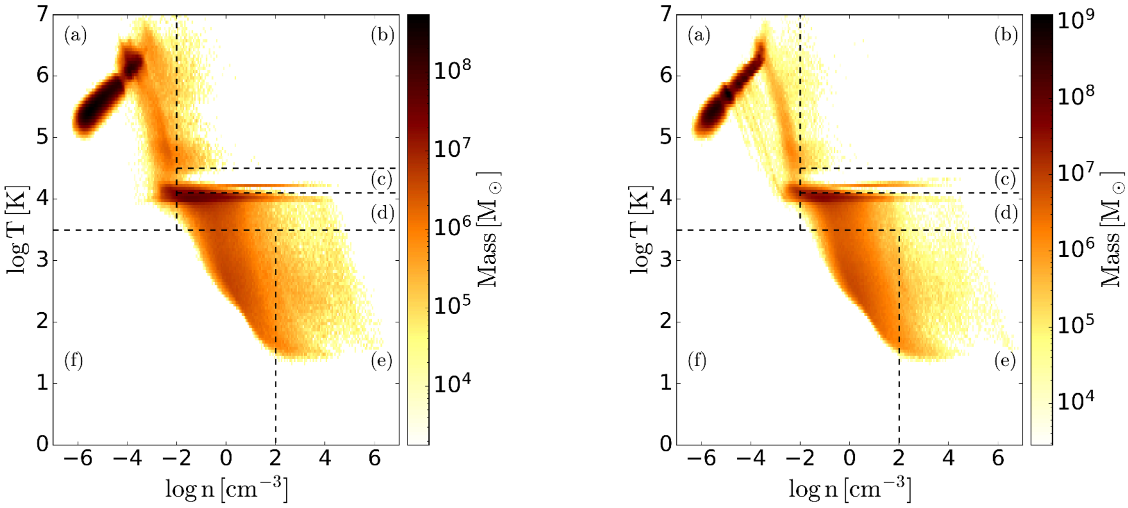

Figure 4.6: Mass-weighted phase diagrams ($\log T$ vs $\log n$) of the total gas, at the $t = 0.35$ Gyr of the simulation (left panel) and after 1.2 Gyr (right panel). The dashed lines partition the plane into six regions: (a) the hot, rarefied coronal gas ($\log n < -2$, $\log T > 4.5$); (b) the denser coronal gas ($\log n > -2$, $\log T > 4.5$), nearly absent; (c) and (d) the outer disc gas at intermediate temperatures ($\log n > -2$, $3.5 < \log T < 4.5$); (e) the cold, dense star-forming gas ($\log n > 2$, $\log T < 3.5$); and (f) the cold but diffuse gas ($\log n < 2$, $\log T < 3.5$), which lies below the density threshold required for star formation. Darker colours indicate higher gas mass.

and (d), back towards lower densities and higher temperatures is the thermodynamic signature of the gas heated and displaced by stellar feedback.

4.4.2 Mass and Volume Probability densities

To quantify the multiphase structure displayed in the phase diagrams, we analyse the distribution of gas mass and volume as a function of the number density n. Figure 4.7 shows mass (left panel) and volume (right panel) distributions, computed over the entire simulation box and partitioned into cold, warm and hot phases. The contrast between the mass-weighted and volume-weighted distributions illustrates the physical state of the No-Heat scenario. At the lowest densities ($\log n \lesssim -4$), both distributions are dominated by the hot gas, filling the outer regions of the corona. Towards higher densities the hot phase declines sharply and by $\log n > -1$ its contribution to both budgets becomes negligible. The warm phase acts as a transition regime between the hot corona and the cold disc, contributing significantly to both the mass and volume budgets across low / intermediate densities ($-4 \lesssim \log n \lesssim 0$). Conversely, the cold phase ($T < 10^4$ K) occupies a small fraction of the total volume but dominates the mass budget at intermediate and high densities ($\log n > -1.5$). The cold-gas mass distribution forms a broad peak around $\log n \simeq 0.5 - 1.0$ and a high-density tail tracing the dense, collapsing molecular gas within the galactic disc where star formation actively occurs. The volume distribution, by contrast, confirms that this cold dense phase fills only a negligible fraction of the simulated volume despite dominating the mass at high densities.

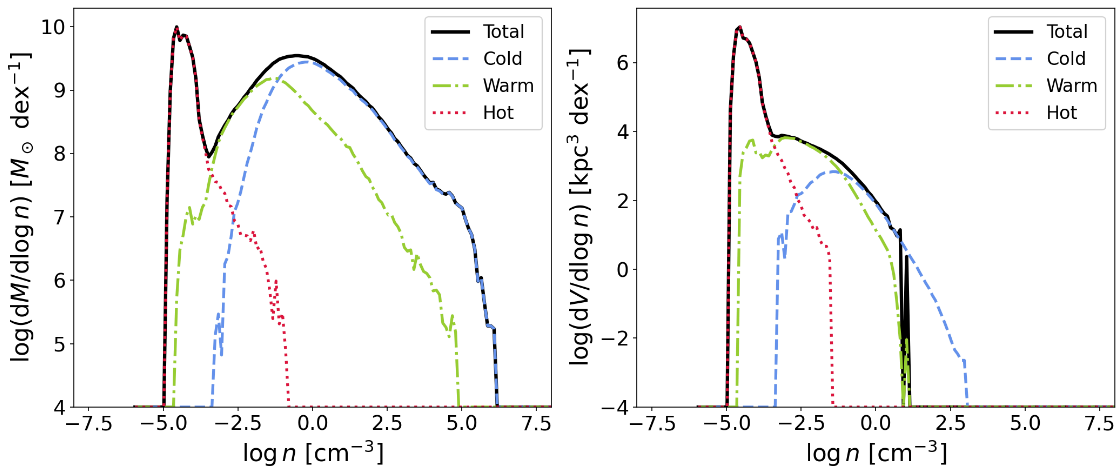

Figure 4.7: Distribution of gas mass (left panel) and volume (right panel) as a function of number density $\log n$ for the No-Heat model at $t \simeq 1.5$ Gyr , computed within a sphere of radius $r = 100$ kpc. The gas is partitioned into total (black solid), cold (blue dashed), warm (green dash-dotted) and hot (red dotted) phases, defined respectively as $T < 10^4$ K, $10^4\,\mathrm{K} < T < 5 \times 10^5\,\mathrm{K}$, and $T > 5 \times 10^5$ K.

4.5 Star formation evolution and self-regulation

The continuous influx of cold gas onto the disc, described in the previous sections, provides abundant fuel for star formation. The SFR evolution (Figure 4.8) shows how this process is regulated by the stellar feedback implemented in the SMUGGLE model. After two initial peaks of $\sim 9\,M_{\odot}\,\mathrm{yr}^{-1}$ at $t \sim 0.1$ and $t \sim 0.2$ Gyr, the SFR declines and stabilizes, oscillating around $4-5\,M_{\odot}\,\mathrm{yr}^{-1}$ for the rest of the simulation. This stable plateau is the signature of the self-regulation between star formation and stellar feedback, whose underlying mechanism is analysed in detail through the gas flows in Section 4.5.1.

This plateau value of $4-5\,M_{\odot}\,\mathrm{yr}^{-1}$ is consistent with the SFRs obtained by Barbani et al. (2023), with the same ICs, which are themselves in reasonable agreement with the value observed in the Milky Way ($\sim 1-3, M_{\odot}\,\mathrm{yr}^{-1}$, Bland-Hawthorn & Gerhard 2016) and in analogous galaxies. The simulation does not aim to reproduce the Milky Way exactly, and a moderate offset from its present-day rate is therefore expected. The initial peaks at $t \sim 0.1$ Gyr and $t \sim 0.2$ should be interpreted with caution, as already noted by Barbani et al. (2023): the corona is initialised in hydrostatic equilibrium, but the activation of the cooling processes causes it to lose pressure support in its central regions, which then accrete rapidly onto the disc. This early, transient burst of accretion drives the initial SFR peak, before star formation settles into its feedback-regulated regime.

4.5.1 Inflow/outflow evolution

To quantify the gas circulation that regulates star formation, we compute the mass inflow and outflow rates at three different locations: across a planar slab close to the disc, to probe the disc-halo exchange; and across two spherical shells in the corona, to probe the gas circulation on larger scales.

We first consider a planar slab placed at a height $z = 2$ kpc above the disc plane. Following the approach of Marinacci et al. (2019) and Barbani et al. (2023), we select all gas cells within a cylindrical radius $R < 30$ kpc and inside a slab of thickness $\Delta z = 0.3$ kpc. For each cell, the contribution to the flow rate is $\dot{M} = m_i\,|v_{z,i}|/\Delta z$, where m_i and $v_{z,i}$ are the mass and vertical velocity of the cell. Cells moving away from the disc ($z\,v_z > 0$) contribute to the outflow, those moving towards it ($z\,v_z < 0$) to the inflow. Figure 4.9 shows the resulting rates as a function of time. After the initial transient, both settle into a fluctuating regime, with a mean inflow rate of $\sim 10\,M_{\odot}\,\mathrm{yr}^{-1}$ and a mean outflow rate of $\sim 7\,M_{\odot}\,\mathrm{yr}^{-1}$. The two curves are anti-correlated: peaks in the outflow correspond to troughs in the inflow and vice versa. This is the kinematic signature of the galactic fountain, i.e. the cycle in which gas is ejected from the disc by stellar feedback, travels through the lower corona, and subsequently falls back onto the disc.

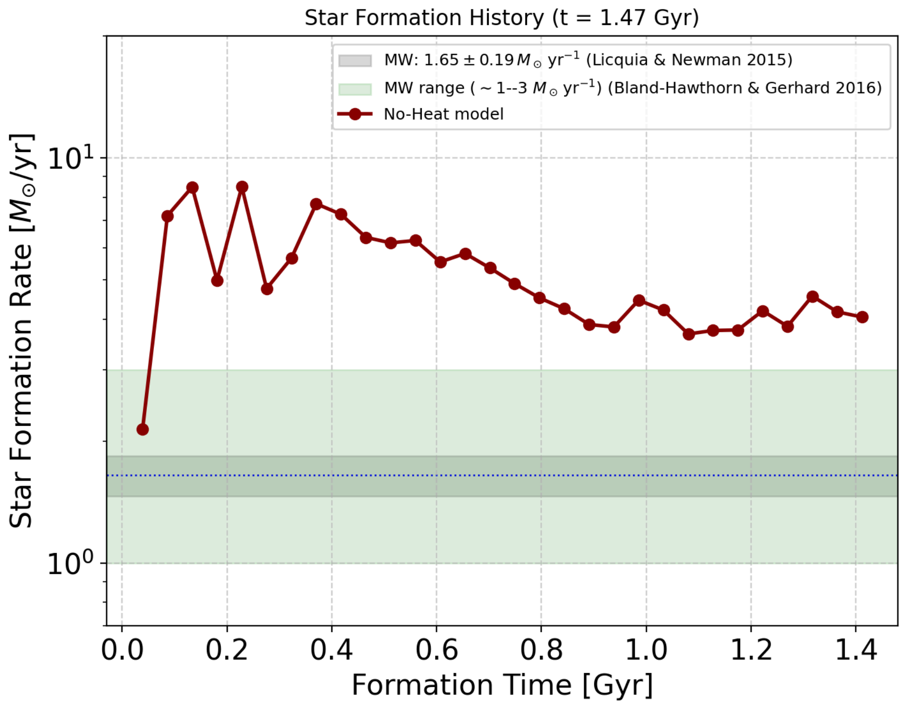

Figure 4.8: SFR as a function of time for the No-Heat model. After an initial peak of $\sim 9\,M_\odot\,\mathrm{yr}^{-1}$ at $t \sim 0.1$ and $t \sim 0.2$ Gyr, driven by the early accretion of coronal gas as the corona loses its initial hydrostatic support, the SFR declines and settles onto a stable plateau at $\sim 4-5\,M_\odot\,\mathrm{yr}^{-1}$, maintained by the self-regulation between star formation and stellar feedback for the rest of the simulation. In light green it is displayed the range of observed values we refer to (Bland-Hawthorn & Gerhard, 2016).

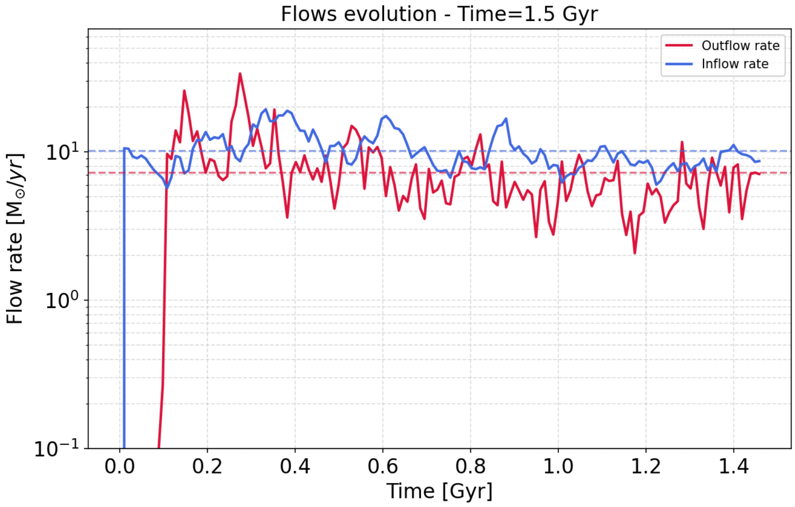

Figure 4.9: No-Heat inflows and outflows computed within a $|z| = 2$ kpc distance from the disc. The gas is selected within a cylindrical radius $R < 30$ kpc and inside a slab of thickness $\Delta z = 0.3$ kpc. The red curve shows the outflows ($\sim 7\, M_\odot\, \mathrm{yr}^{-1}$), while the blue one shows the inflows ($\sim 10\, M_\odot\, \mathrm{yr}^{-1}$). So there is a net increase of gas inward.

To probe the gas circulation on coronal scales, we repeat the measurement across spherical shells at $r = 10$ and $r = 50$ kpc, using the radial velocity $v_r < 0$ for inflow, $v_r > 0$ for outflow. At $r = 10$ kpc the circulation shows a mean inflow of $\sim 14\, M_\odot\, \mathrm{yr}^{-1}$ slightly exceeding the outflow of ~ 12, i.e. a modest net accretion. At $r = 50$ kpc the absolute rates drop by almost an order of magnitude (inflow ~ 3, outflow $\sim 1\, M_\odot\, \mathrm{yr}^{-1}$), while the net inflow remains comparable ($\sim 1.6\, M_\odot\, \mathrm{yr}^{-1}$): the outer corona accretes gently, and the circulation intensifies inward.

These rates reflect the recycling of gas through the galactic fountains : at both radii the inflow and outflow are large and close to each other, so most of the circulating gas is recycled and only a small net amount is accreted inward. This net inflow feeds the inner regions and the disc, where, as the cold and warm reservoirs compensate each other throughout the run (Figure 4.1), it is balanced by conversion into stars, so the disc gas mass settles into equilibrium. This balance is maintained by a self-regulating loop, visible in the temporal relation between the SFR (Figure 4.8) and the gas flows (especially Figures 4.9 and 4.10) where peaks in the SFR are followed, after ~ 0.1 Gyr, by peaks in the outflow. What happens is that SNe from newly formed stars drive gas out of the disc, limiting star formation until the gas

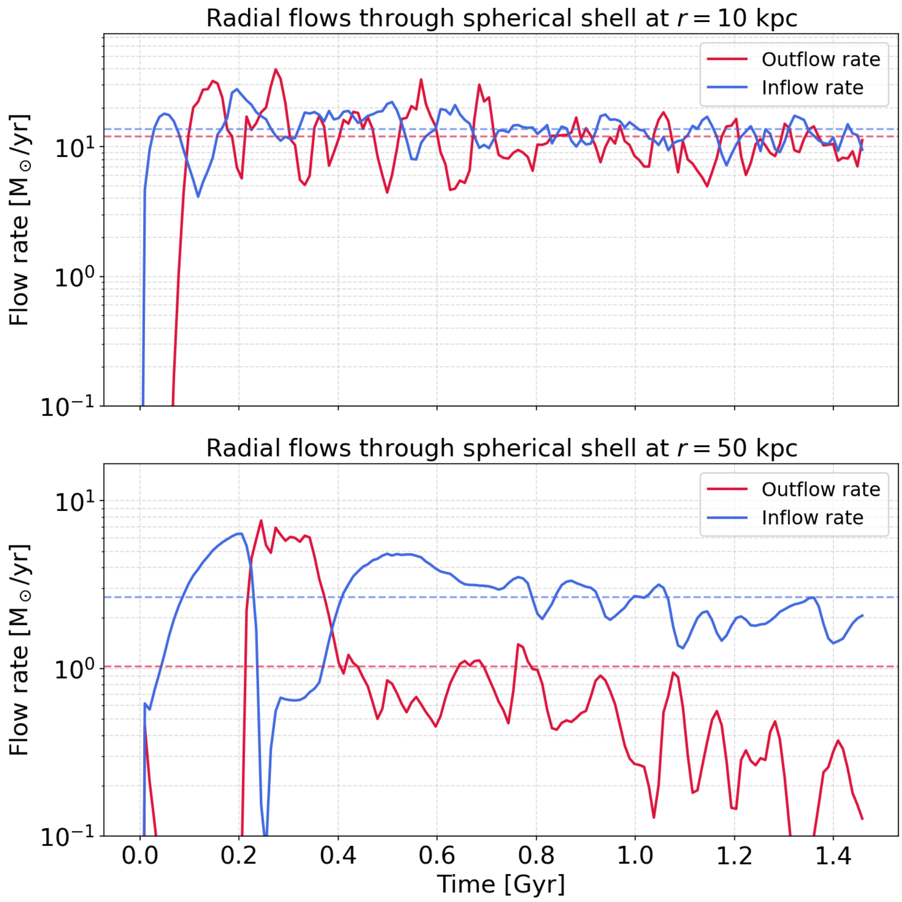

Figure 4.10: No-Heat inflows and outflows computed within a sphere of radius $r = 10$ kpc in the upper panel, and within $r = 50$ kpc in the lower panel. The red curve shows the outflows , while the blue one shows the inflows, which are at $r = 10$ kpc ~ 12 and $\sim 14\,M_\odot\,\mathrm{yr}^{-1}$ respectively , whereas at $r = 50$ kpc they are ~ 1 and $\sim 3\,M_\odot\,\mathrm{yr}^{-1}$. So there is a net increase of gas inward also at these distances.

can cool and accrete again. The SFR correlates tightly with the outflow, with the time lag mentioned before, but only weakly with the inflow.

4.6 Summary

In conclusion, the No-Heat model shows that the SMUGGLE stellar feedback is remarkably efficient at locally self-regulating the disc, maintaining the SFR at a steady, moderate level despite the gas accretion from the cooling corona. By construction, however, this scenario includes no mechanism able to replenish the thermal energy lost by the hot CGM on scales larger than ~ 10 kpc: the stellar feedback acts locally, within and around the star-forming disc, and cannot offset the radiative cooling of the extended hot corona, which therefore continues to cool and accrete onto the galaxy mid-plane.

These results make explicit the central motivation of this thesis. They demonstrate that reproducing a realistic, long-lived Milky Way-like corona requires an additional, spatially diffuse heating channel acting on the coronal gas itself, a channel that is absent from the standard SMUGGLE framework and whose implementation and characterisation constitute the original contribution of the present work. The following chapters introduce and test such heating models, building directly on the baseline established here.

Chapter 5

Distributed-Heating Models of the corona

5.1 Physical motivation: Active Galactic Nuclei (AGN) feedback and Fermi bubbles

The No-Heat model showed that, in the absence of an external energy source, the corona is subject to radiative cooling: stellar feedback alone is able to regulate the multiphase structure of the ISM and star formation in the disc, but does not prevent the inner regions of the corona from cooling and accreting onto it. An additional, large-scale heating channel acting on the coronal gas therefore appears necessary. It is widely accepted that massive galaxies host a central SMBH, which represents the most physically candidate for the heating channel previously mentioned to deposit energy into the surrounding ISM and CGM. We develop a model of this kind, based on AGN-driven bubbles, in Chapter 6, where the mechanism and the relevant literature are discussed in detail. Here, we mention the idea and the reason behind the implementation of such a mechanism. The idea that the energy released by a central SMBH can counteract radiative cooling was developed primarily in the context of the *cooling flow problem* in galaxy groups and clusters (see, e.g., Fabian, 1994; Peterson & Fabian, 2006). Although alternative heating sources have been proposed—including thermal conduction, cosmic rays and turbulent dissipation—there is now broad consensus that AGN feedback is the dominant mechanism balancing radiative losses in these systems (McNamara & Nulsen, 2007; Fabian, 2012). This type of feedback is not merely a theoretical device invoked to suppress excessive cooling in simulations: it is directly observed. X-ray cavities and radio lobes seen in the cores of galaxy clusters are interpreted as bubbles inflated by the central AGN in the surrounding hot gas, providing direct evidence of this energy injection (Fabian, 2012; McNamara & Nulsen, 2007). The Milky Way provides direct evidence that even a galaxy with a relatively low-mass SMBH (Sgr A*, $M_\bullet \sim 4 \times 10^6\, M_\odot$) has undergone past episodes of energetic nuclear activity. The most striking example are the *Fermi bubbles*, two giant γ-ray structures discovered with the *Fermi*-LAT (Su et al., 2010b), extending symmetrically to $|z| \sim 10$ kpc above and below the Galactic plane. Their morphology, sharp edges and hard spectrum are interpreted as the relic of a large-scale outflow originating from the Galactic centre, plausibly driven by a past accretion event onto the SMBH or by a central starburst (Su et al., 2010b). Such structures indicate that the energy injection from the Galactic nucleus is episodic rather than continuous. Before trying to implement such a mechanism, we first take a deliberately simpler route: rather than tying the heating to a specific physical source and reproducing the localised geometry of an AGN outflow, we introduce a distributed heating term that compensates the radiative losses of the coronal gas directly, cell by cell. The purpose of this chapter is therefore twofold. First, to test whether compensating the coronal cooling is by itself sufficient to offset the depletion of the hot corona seen in the No-Heat run. Second, to isolate the role of the *temporal regime* of the energy injection: the two models presented from now on share the same energy budget and differ only in the timescale over which it is delivered. They thus provide a controlled reference for the more realistic bubble model discussed in Chapter 6), with the No-Heat run of Chapter 4 as the common baseline against which all heating models are compared.

5.2 Numerical implementation

The two heating models presented in this chapter, **Heat 1** and **Heat 2**, share the same physical prescription for the rate of energy injected and differ only in the timescale over which it is delivered. The heating is designed to compensate, on a cell-by-cell basis, the radiative cooling of the coronal gas. For each gas cell, we estimate the rate at which it radiates its thermal energy as

$$L_{\mathrm{cool},i} = \frac{m_i\, u_i}{t_{\mathrm{cool},i}}, \tag{5.1}$$

where m_i and u_i are the mass and specific internal energy of the cell and $t_{\mathrm{cool},i}$ is its cooling time. Both Heat 1 and Heat 2 models share this same heating rate, but deliver the corresponding energy $\Delta E_i = L_{\mathrm{cool},i}\, dt$ over different injection intervals dt, set by the hierarchy of nested timesteps used by AREPO to evolve the gas. In either case the injection is restricted to the diffuse coronal gas through a density threshold ($\rho < \rho_{\mathrm{th}} = 0.001\, cm^{-3}$) and a temperature cut ($T > 10^4$ K), so that only the hot, rarefied atmosphere is heated and the dense, star-forming disc is left unaffected.

5.2.1 Heat 1: Continuous injection

In the **Heat 1** model, the energy is injected on the local hydrodynamical timestep of each cell, i.e., the same time interval on which the gas hydrodynamics are evolved. The thermal energy is thus released almost continuously and in small amounts. Because each injection adds only a small amount of internal energy a time, a fraction of it is radiated away before it can accumulate, this is strongest in the inner dense corona where radiative losses are largest. Overall, the continuous compensation of the radiative losses still slows down the depletion of the hot gas mass of the corona with respect to the No-Heat case, but this is a subtle and the only imprint of this model, as better shown in the following sections.

5.2.2 Heat 2: Accumulated injection

In the **Heat 2** model, the same heating rate is instead accumulated over a longer interval, corresponding to the longest active timestep in the simulation hierarchy. The time-averaged power delivered to the gas is the same as in Heat 1, but the energy is released in fewer, larger packets. Because each injection raises the local internal energy by a larger amount at once, a larger fraction of the deposited energy can be retained by the gas before being radiated away, rather than being lost almost immediately as in the continuous case. In practice, however, the evolution of the corona in Heat 2 is very close to that of Heat 1: the major differences are found in the flows computed within a sphere of $r = 10$ kpc radius from the disc, in the phase diagram panels displayed in Section 5.5.1 and a slight but more attenuated decline of the hot gas mass that can be seen in Figure 5.1. Towards the end of the simulated time ($t \approx 1.5$ Gyr) Heat 2 retains marginally more hot gas mass, which may hint at a more sustained corona at later times; however, this trend is small and our simulations do not extend far enough to confirm it.

5.3 Mass evolution

To assess the impact of the two heating prescriptions on the thermodynamic state of the gas, we compare the evolution of the gas mass, split into temperature phases and computed within the inner corona ($r < 100$ kpc), for the Heat 1 and Heat 2 models, using the No-Heat run as a reference (Figure 5.1).

The cold and warm phases evolve almost identically in the three models. After a brief transient (~ 0.2 Gyr), the cold gas mass grows from $\sim 7 \times 10^9$ to $\sim 9 \times 10^9\, M_\odot$, while the warm gas decreases from $\sim 7 \times 10^9$ to $\sim 5 \times 10^9\, M_\odot$. The fact that these two phases follow the same evolution in all three runs confirms that the heating, restricted to the diffuse coronal gas, does not significantly affect the denser gas of the disc.

The differences between the models appear in the hot phase. In the No-Heat run the hot gas mass starts slightly above $10^{10}\, M_\odot$ and drops to $\sim$ 6-7 $\times\, 10^9\, M_\odot$ already by $t \approx 0.6$ Gyr, settling onto a plateau around $7 \times 10^9\, M_\odot$ after ~ 0.8 Gyr. The heating models deplete hot gas more slowly: in Heat 1 the start is same as No-Heat but the decline is more gradual and does not reach a stable plateau within the simulated time range, precisely with the hot gas mass falling below $7 \times 10^9\, M_\odot$ only after $t \approx 1.2$ Gyr. Heat 2 follows very closely the behaviour of Heat 1: both heating prescriptions slow down the depletion of the hot corona with respect to the unheated baseline, confirming that the cooling compensation does provide some thermal support. However, the difference between the two delivery timescales (Heat 1 and Heat 2) remains marginal throughout the simulated time, consistent with the coronal density profiles (Section 5.6), which show closely overlapping distributions

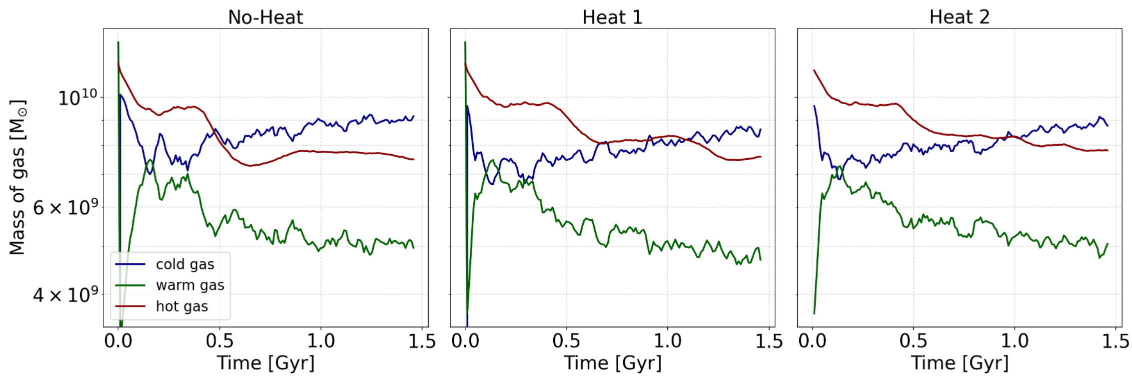

Figure 5.1: Evolution of the gas mass split into cold ($T < 10^4$ K, blue), warm ($10^4 < T < 5 \times 10^5$ K, green) and hot ($T > 5 \times 10^5$ K, red) phases, computed within a sphere of $r = 100$ kpc, for the No-Heat, Heat 1 and Heat 2 models. It is noticeable the effect of the heating prescriptions on the hot gas (red) while the other two phases seem to be really similar between the three models.

for the three models.

5.4 Morphological evolution of the gas disc

Having quantified how the heating affects the global hot gas mass budget, we now turn from the integrated masses to spatial distribution of the simulated galaxy, asking where in the galaxy these differences originate. We begin with the analysis of the disc, and the region extending $\sim$ 5 kpc above and below it, through the projected density and temperature maps, before moving back out to the corona.

5.4.1 Surface Density maps

The face-on density maps show that the structure of the disc (mid-plane density, spiral arms and central core) remains essentially identical in the three models. This is the expected behaviour: the heating is applied only to the diffuse coronal gas (Section 5.2) and by construction does not act on the dense, cold gas of the disc, so the sites of star formation are left unaffected.

The edge-on projections, by contrast, reveal systematic differences in the extra-planar gas, consistently present across the snapshots. In the No-Heat model the region $\sim$ 5 kpc above the disc is filled by a thick, diffuse envelope of gas surrounding

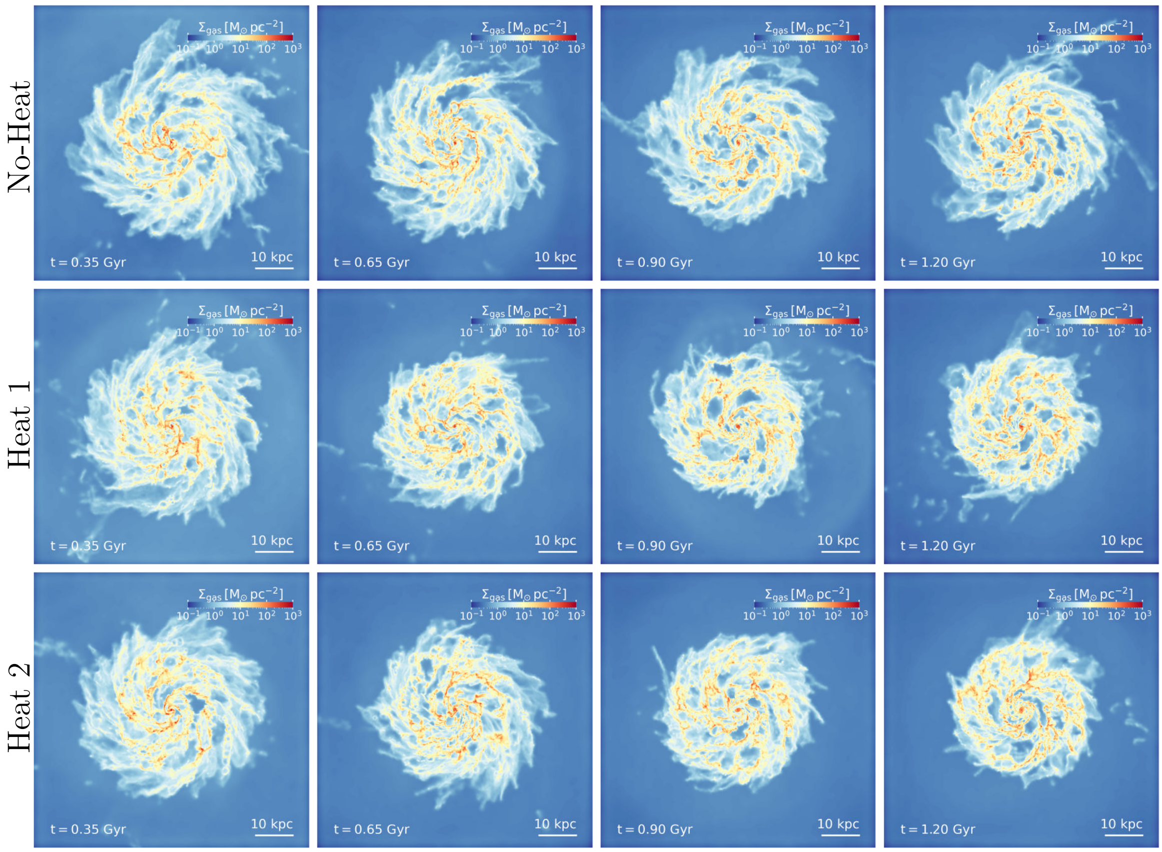

Figure 5.2: Face-on projected gas surface density $\Sigma_{\rm gas}$ for the No-Heat, Heat 1 and Heat 2 models (rows) at four different times (columns). The four times are chosen to match the epochs at which the mass evolution (Figure 5.1) shows the most significant changes in the gas content of the corona, so that the morphology can be examined at the moments when the models begin to differ. The disc structure remains closely similar across the three models at every time, reflecting no significant effects on the dense, star-forming gas from the coronal heating prescription. The only appreciable difference is the more compact structure, from t = 0.65 Gyr onwards, of the two heating models with respect to the unheated one.

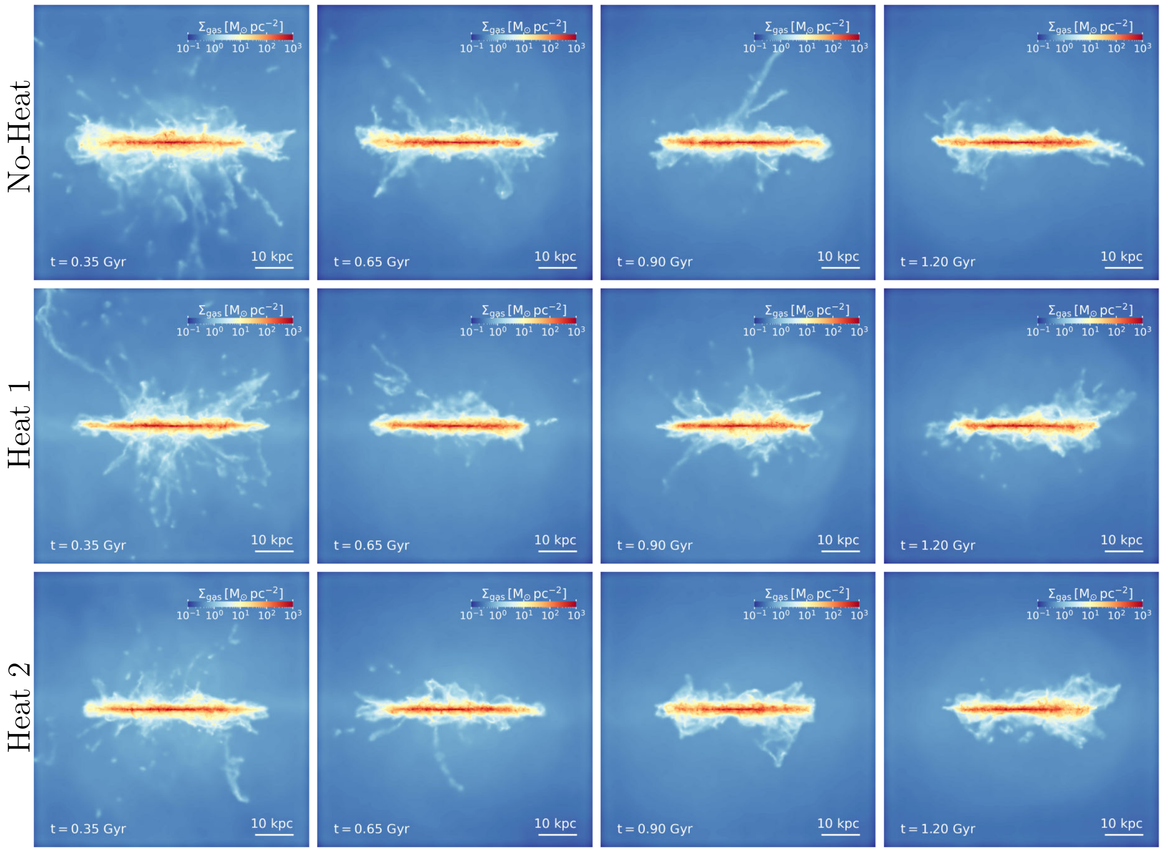

Figure 5.3: Edge-on projected gas surface density $\Sigma_{\rm gas}$ for the No-Heat, Heat 1 and Heat 2 models (rows) at four different times (columns). The four times are chosen to match the epochs at which the mass evolution (Figure 4.1) shows the most significant changes in the gas content of the corona, and also to be consistent with the Face-on density maps. The disc structure remains closely similar across the three models at every time, reflecting no significant effects on the dense, star-forming gas from the coronal heating prescription.

the mid-plane, whereas in Heat 1 the disc appears thinner, and the gas above it seems organised into more defined, filamentary plumes, while in Heat 2 the same region gas is more compact and less filamentary.

It is worth noting that these morphological differences are not reflected in the mass flow rates measured at $|z| = 2$ kpc (Figure 5.10), where the inflow and outflow rates of the models are comparable. This suggests that the differences could develop mainly at larger heights, above the slab used to measure the flows: while a similar amount of gas crosses the 2 kpc plane in all models, the way this gas is structured and distributed at higher altitudes may differ. This investigation is presented in Section 5.7.2.

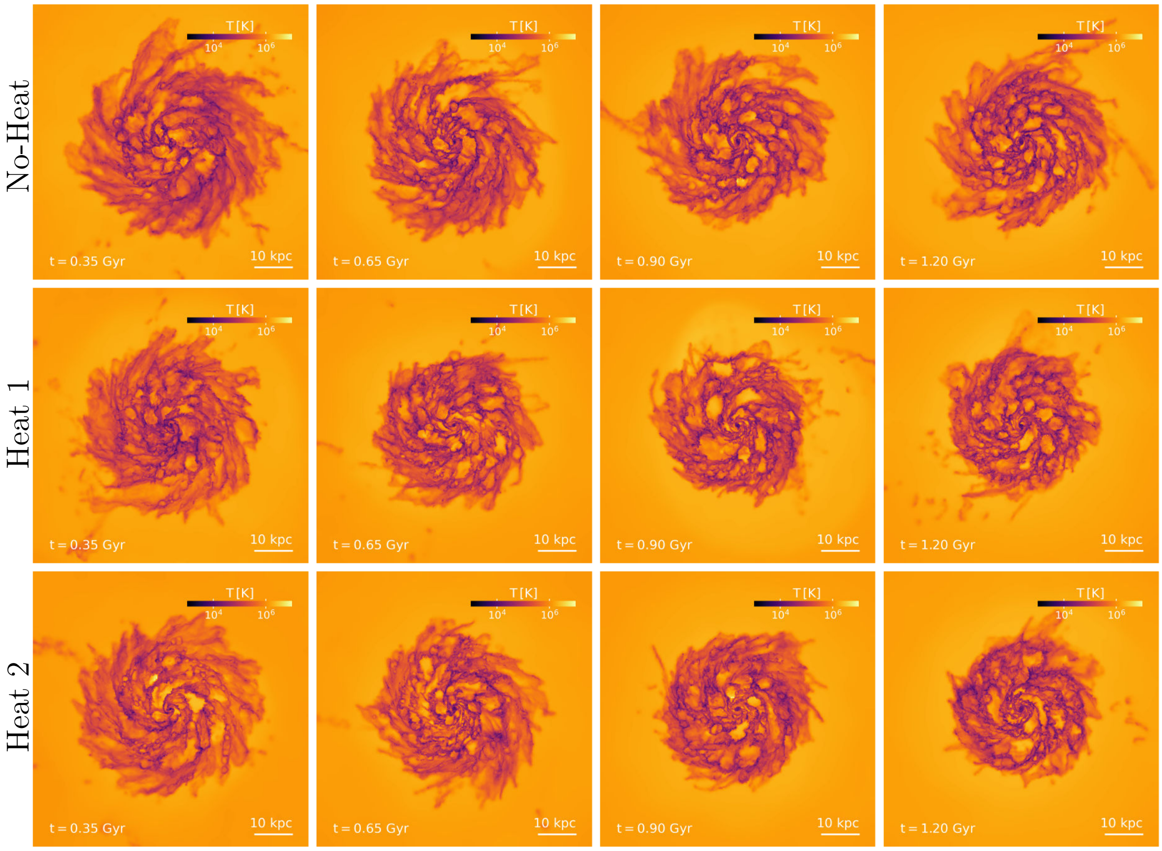

Figure 5.4: Face-on projected, mass-weighted temperature maps for the No Heat, Heat 1 and Heat 2 models (rows) at four different times chosen to match the ones in Figure 5.2 (columns). As for the density maps, the temperature structure of the disc is closely similar across the three models at every time; the heating prescriptions act on the diffuse coronal gas and leave the temperature of the dense, star-forming disc unaffected. The only appreciable difference is the more compact structure, from $t = 0.65$ Gyr onwards, of the two heating models with respect to the unheated one, as for the density face-on maps.

5.4.2 Mass-weighted Temperature snapshots

The face-on temperature maps (Figure 5.4) show that the disc has the same multi-phase structure in the two heating models: cold filaments ($T \sim 10^3 - 10^4$ K) trace the spiral arms and mark the star-forming sites, embedded in a warmer inter-arm medium ($T \sim 10^4 - 10^5$ K), with localised hot patches associated with recent stellar feedback. As for the density, the disc temperature structure appears similar to the No-heat prescription, consistent with the fact that this model acts only on the diffuse coronal gas.

The edge-on temperature maps (Figure 5.5) confirm the morphological picture

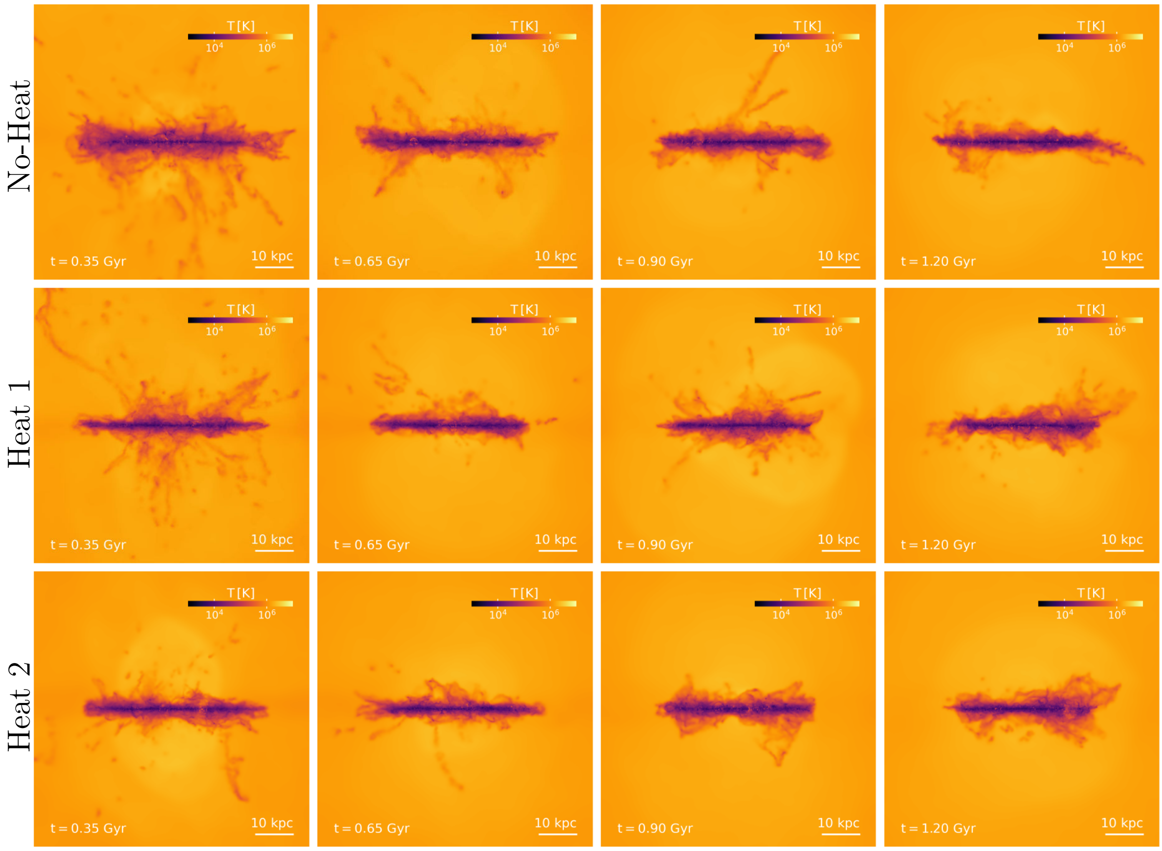

Figure 5.5: Edge-on projected, mass-weighted gas temperature for the No-Heat, Heat 1 and Heat 2 models (rows) at four different times chosen to match the ones in Figure 5.2 (columns). As for the density, the temperature structure of the disc is closely similar across the three models at every time; the heating prescriptions act on the diffuse coronal gas and leave the temperature of the dense, star-forming disc unaffected.

given by the density. In all models a thin, cold disc ($T \lesssim 10^4$ K) is embedded in the hot coronal gas ($T \sim 5 \times 10^5$ K), whose temperature is comparable in the three cases. The difference lies again in the gas above the plane: in the Heat 1 the cooler gas above and below the plane is organised into more extended filaments and plumes than in Heat 2, where it forms a more compact layer surrounding the mid-plane, with fewer filamentary structure. The temperature of the surrounding corona itself is not significantly different between the models; what changes more is the spatial structure of the cooler gas embedded in it.

5.5 Thermodynamic and Multiphase ISM

The surface-density maps show where the gas sits, but not how its mass is distributed among the thermodynamic phases. Two models could share the same projected appearance while differing in how much of their gas is cold, warm or hot, and it is precisely this partition that the heating is expected to modify. So, to test this, we look at the gas configuration into the density–temperature plane, where each phase occupies a distinct region and the effect of the heating on the coronal gas can be read directly. We use this representation to ask whether the cooling-compensation heating leaves a measurable imprint on the thermodynamic state of the gas, relative to the No-Heat baseline.

5.5.1 Phase diagrams

The mass-weighted phase diagrams ($\log T$ versus $\log n$) of the three models at three representative times are shown in Figure 5.6. The configuration of the phases and regions is the same assumed in Section 4.1. In all models the gas segregates into the same regimes: the hot, rarefied coronal gas at low density and high temperature ($\log n \lesssim -3$, $T \gtrsim 10^6$ K), the cold, dense star-forming gas at high density and low temperature ($\log n \gtrsim 0$, $T < 10^4$ K), and the intermediate phases connecting them. The diffuse component extending from the dense, cold gas towards lower densities and higher temperatures traces the gas heated and displaced by stellar feedback and expelled from the disc. Comparing the columns, the diagrams evolve in the same way in all models: after the initial transient the distribution settles into a stable multiphase structure that is maintained throughout the simulation. Comparing the rows, the three models are remarkably similar at each time, indicating that the heating does not introduce any qualitative change in the way the gas populates the density vs temperature plane.

The only appreciable difference appears in the coronal region (a). In the No-Heat and Heat 1 models a low-mass tail of gas extends from the hot corona towards the intermediate temperatures ($10^4 \lesssim T \lesssim 10^6$ K) at low density, tracing coronal gas that is cooling out of the hot phase. In Heat 2 this feature is essentially absent: by compensating the radiative losses, the Heat 2 injection keeps the coronal gas close to its characteristic temperature and prevents it from cooling through these intermediate temperatures. This is the clearest thermodynamic signature of the heating. The fact that the feature is present in Heat 1, remarkably similar from No-Heat, confirms the limits of the continuous injection scheme in offsetting the radiative losses of the inner corona.

5.5.2 Mass and Volume Probability densities

To complete the thermodynamic analysis, we examine the distribution of gas mass and volume as a function of number density for the Heat 1 and Heat 2 models (Figure 5.7). As for the phase diagrams, the two distributions are very similar, confirming that the temporal regime of the energy injection does not change remarkably the way the gas is distributed in density. The mass distribution is dominated by the cold phase ($T < 10^4$ K), which forms a broad peak at densities $\log n \sim 0$ in both models. This reflects again the fact that the heating, applied only to the diffuse coronal gas ($\rho < \rho_{\rm th}$, $T > 10^4$ K; Section 5.2), leaves the dense gas of the disc almost unaffected, so the bulk of the baryonic mass remains in the mid-plane in both cases.

The volume distribution is instead dominated at low densities ($\log n \lesssim -3$) by the hot gas, which occupies most of the volume while contributing little to the mass. This is the expected behaviour of a hot, rarefied corona and, as such, it is common to all the models, including the No-Heat baseline: it is a property of the coronal gas itself rather than an effect of the heating. The mass, and volume, weighted distributions of Heat 1 and Heat 2 are therefore consistent with the picture established in the previous sections, and add no further distinguishing feature between the two heating schemes, exception made for the slight difference in the tail of the cold gas in the mass PDF, in particular the Heat 2 model has a cut at slightly lower densities ($\log n \sim 6$) with respect to the Heat 1 that exceeds this value. This behaviour seems compensated by an increase of the warm gas mass at densities $\log n \sim -4$, but, as already pointed out, it is not a significant difference in the two models.

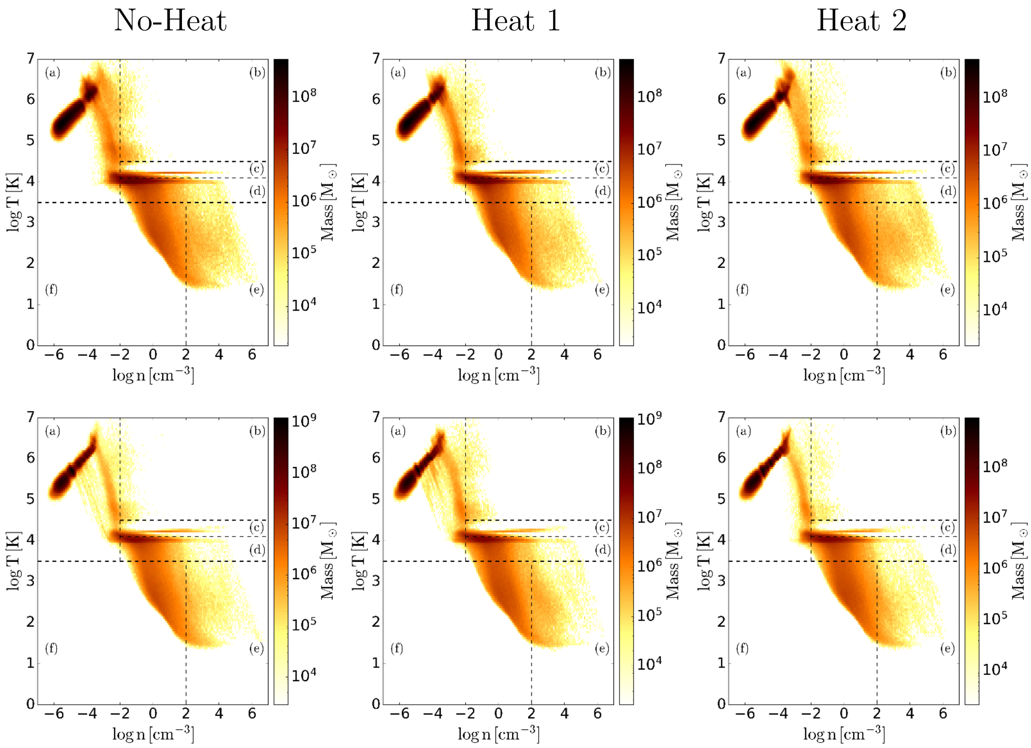

Figure 5.6: Mass-weighted phase diagrams in the density–temperature plane ($\log T$ versus $\log n$) for the No-Heat, Heat 1 and Heat 2 models (columns) evaluated at $t = 0.35$ (top row) and $t = 1.20$ Gyr (bottom row). The dashed lines partition the diagram into the same regions defined in Section 4.1. Darker colours indicate higher gas mass. At each time the three models occupy the same regions with closely similar mass partitions, indicating that the cooling-compensation heating leaves only a faint imprint on the thermodynamic state of the gas. The most significant difference is displayed for the Heat 2 model in the region a) by the absence of the tail, hint of a modest support to the corona by the heating model.

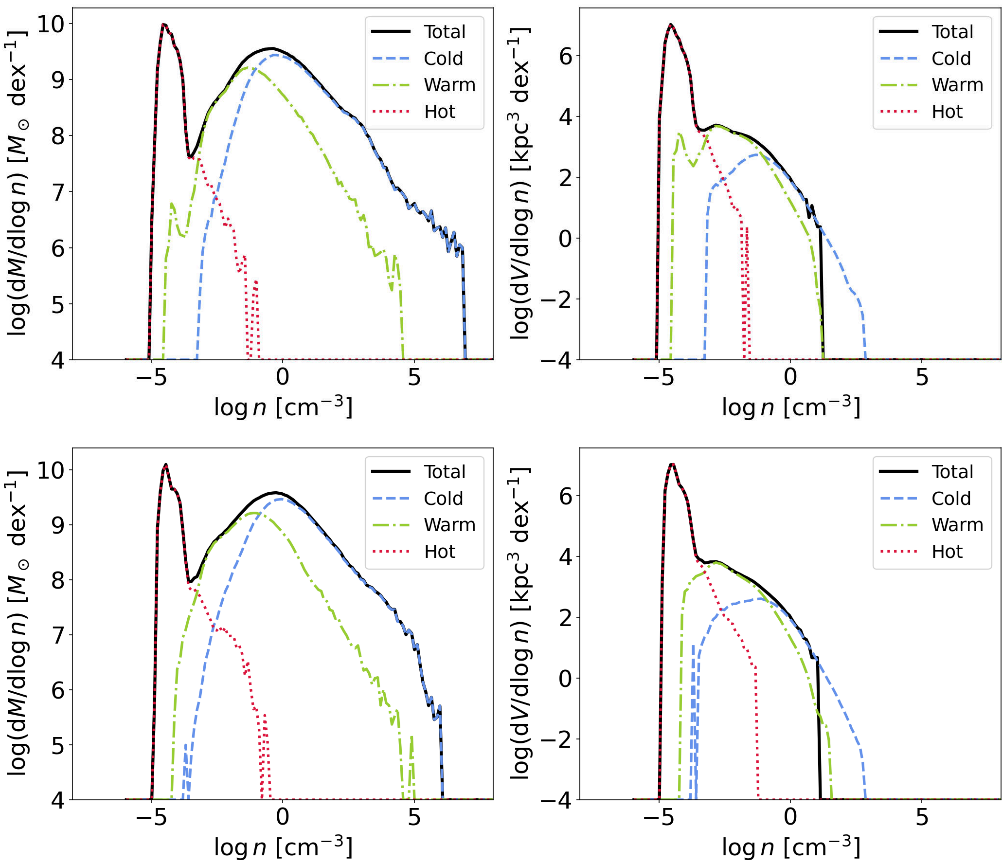

Figure 5.7: Mass (left column) and volume (right column) probability distribution functions of the gas as a function of number density, $\mathrm{d}M/\mathrm{d}\log n$ and $\mathrm{d}V/\mathrm{d}\log n$, for the Heat 1 (top row) and Heat 2 (bottom row) models at $t = 1.5$ Gyr, decomposed into the cold, warm and hot phases (see legend) with the total in black. The mass distribution is bimodal, with a low-density peak from the hot corona and a high-density peak from the cold disc gas, while in the volume distribution the hot gas dominates, filling most of the volume while contributing little mass. The two heating models give similar distributions, confirming that the cooling compensation does not alter how the gas mass and volume are partitioned across the density regimes.

5.6 Coronal gas density profiles: the limits of central heating

The phase diagrams show that the heating leaves a thermodynamic imprint on the coronal gas, while the mass evolution (Section 5.3) shows that the heating slows the depletion of the hot phase. Both diagnostics, however, are integrated over the corona and say nothing about *where* this gas sits. We therefore turn to the radial density profiles of the hot corona, to test whether the slower depletion corresponds to a measurable redistribution of the coronal gas, or whether the difference is confined to the total amount rather than its spatial arrangement.

To quantify the spatial distribution of the hot corona, we analyse the radial density profiles of the coronal gas ($T > 5 \times 10^5$ K) in a spherical shell of radius $r < 240$ kpc. Figure 5.8 compares the three models (No-Heat, Heat 1 and Heat 2) at four representative times, together with the initial β-model profile of Barbani et al. (2023) as a reference for the equilibrium configuration the corona starts from. The four times are chosen to be consistent with the comparison made in previous figures like Figure 5.2. The most striking feature is the similarity of the three profiles. At all the times probed (except for $t = 0.35$ Gyr), the density profiles of No-Heat, Heat 1 and Heat 2 remain closely overlapping across the whole radial range. At the earliest time ($t = 0.35$ Gyr) the inner profiles show larger fluctuations, consistent with the initial transient as the corona settles from its ICs; these smooth out at later times. Only at $t = 0.90$ Gyr does Heat 2 lie marginally above the other two within the innermost radii ($r < 10$ kpc), but this difference is local and does not persist at the other times. Within the inner corona ($r \lesssim 40$ kpc) all three profiles lie below the initial β-model, reflecting the depletion of the inner hot gas as the corona cools and feeds the disc. At larger radii ($r \gtrsim 40$ kpc) the profiles closely follow the initial β-model, indicating that the outer corona remains essentially in its initial equilibrium configuration throughout the simulated time. This behaviour is common to the three models: the heating modifies neither the depletion of the inner corona nor the equilibrium of the outer one in a measurable way.

The near-coincidence of the profiles refines the picture given by the hot-gas mass evolution (Section 5.3). There, the heating models showed a slower depletion of the hot phase than the no heat baseline; here, the radial profiles reveal that this difference does not translate into a measurable redistribution of the coronal gas: the slightly larger hot-gas reservoir retained by the heating models is spread over essentially the same radial structure as in the No-Heat case. At fixed mean injected power, and on the timescales explored, this form of distributed, cooling-compensating heating affects the amount and the timing of the coronal depletion, but not the large-scale spatial structure of the corona.

5.7 Star formation and gas flows

5.7.1 SFR

Having probed the effects of the heating models presented in this chapter on the density profiles of the corona, we now move on the possible effects on the SFR. The star formation histories of the three models are compared in Figure 5.9. They follow the same overall behaviour: after an initial peak, the SFR settles onto a plateau oscillating around $\sim 4-5\,M_{\odot}\,\mathrm{yr}^{-1}$, with comparable mean values in all cases. These rates are higher than the observational estimate for the Milky Way ($\sim 1-3, M_{\odot}\,\mathrm{yr}^{-1}$, Bland-Hawthorn & Gerhard 2016), but they stay on a stable plateau rather than growing without bound, indicating that star formation remains well regulated by stellar feedback.

As pointed out above, none of the heating prescriptions produces a measurable decrease in the amount of coronal gas that cools and accretes onto the disc, so the SFR is essentially unchanged with respect to the No-Heat model. The reason follows directly from the preceding sections: the heating slows the depletion of the hot corona but is too weak to significantly suppress its accretion onto the disc, the hot-gas content and the coronal density profiles remain close to the baseline scenario (Sections 5.3 and 5.6), so the gas supply reaching the disc is essentially unchanged. The SFR is therefore not remarkably sensitive to the heating.

The three histories differ only in their short-term variability. During the plateau phase ($t \gtrsim 0.4$ Gyr) the No-Heat and Heat 1 histories show larger and more frequent fluctuations around the mean, with occasional bursts reaching $\sim 10\,M_{\odot}\,\mathrm{yr}^{-1}$,

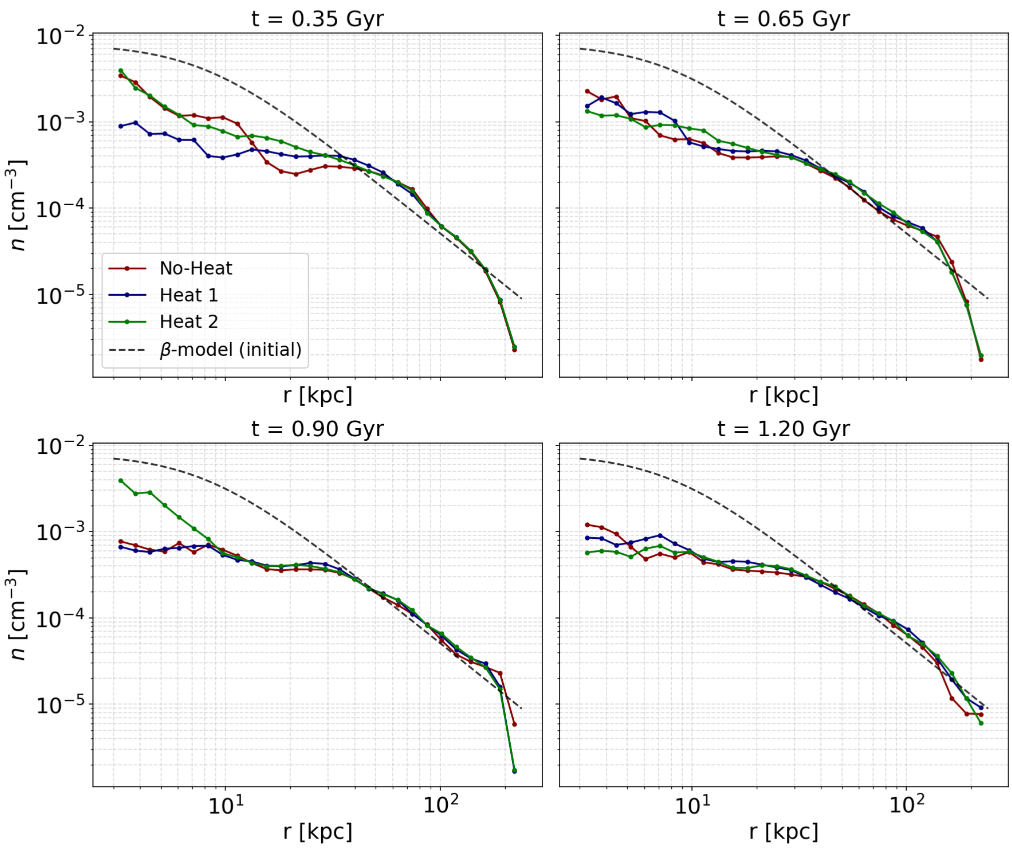

Figure 5.8: Radial density profiles of the hot coronal gas ($T > 5 \times 10^5$ K) for the No-Heat, Heat 1 and Heat 2 models at four representative times, computed within a sphere of radius of $r < 240$ kpc. The dashed black line shows the initial β-model profile of Barbani et al. (2023). The three models remain closely overlapping at all times: within $r \lesssim 40$ kpc the profiles fall below the initial β-model, tracing the depletion of the inner corona, while at larger radii they follow it, indicating that the outer corona stays close to its initial equilibrium. The larger fluctuations of the inner profiles at $t = 0.35$ Gyr reflect the initial transient as the corona settles from its initial conditions. Overall, the heating produces no measurable redistribution of the coronal gas with respect to the baseline.

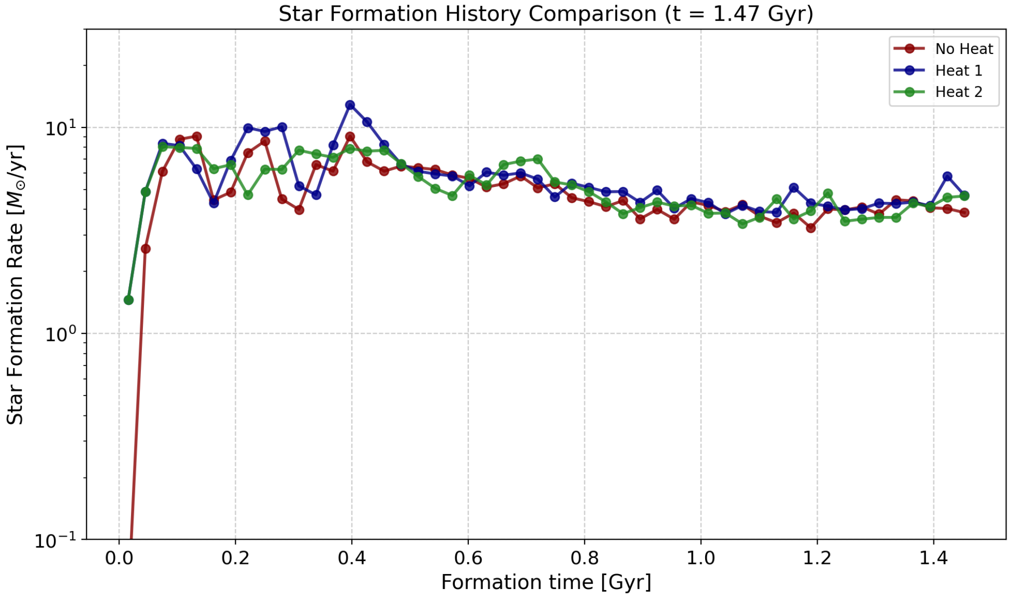

Figure 5.9: SFR for the No-Heat, Heat 1 and Heat 2 models indicated by red, blue and green curves, respectively. The SFR trends are aligned in most of the evolution indicating that this quantity is affected modestly by the heating models. After the initial peaks at $t = 0.1$, $t = 0.3$, $t = 0.4$, they all reach a stable plateau around the value $\sim 4 - 5\, M_\odot\, yr^{-1}$ towards the end of the run.

whereas the Heat 2 history is comparatively smoother, remaining closer to its mean. This is a difference in the amplitude of the oscillations about the plateau, not in the plateau level itself, which is the same in all three models.

5.7.2 Inflow/outflow evolution

To characterise the gas exchange between disc and corona, we compute the inflow and outflow rates at the same three locations as in the No-Heat model (Section 4.5.1): the planar slab at $z = 2$ kpc and the spherical shells at $r = 10$ and $r = 50$ kpc (Figure 5.10 and 5.11). At $z = 2$ kpc, Heat 1 and Heat 2 have mean rate values of $\sim 15\, M_\odot\, \mathrm{yr}^{-1}$ inward and $\sim 10\, M_\odot\, \mathrm{yr}^{-1}$ outward, against the $\sim 17\, M_\odot\, \mathrm{yr}^{-1}$ inflow and $\sim 11\, M_\odot\, \mathrm{yr}^{-1}$ outflow of No-Heat model. From this comparison, we see that at the disc-corona interface there is a modest change of the gas recycling, but not so effective as to change the SFR.

At $r = 10$ kpc the mean rates of Heat 1 are aligned with the No-Heat values, indicating that there are no remarkable differences in the inflows/outflows between

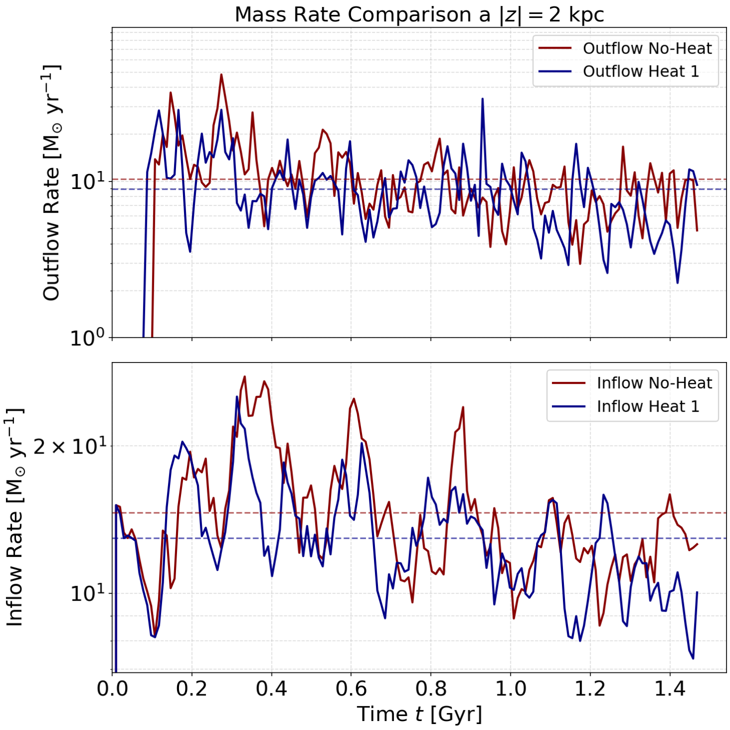

(a) No-Heat vs Heat 1

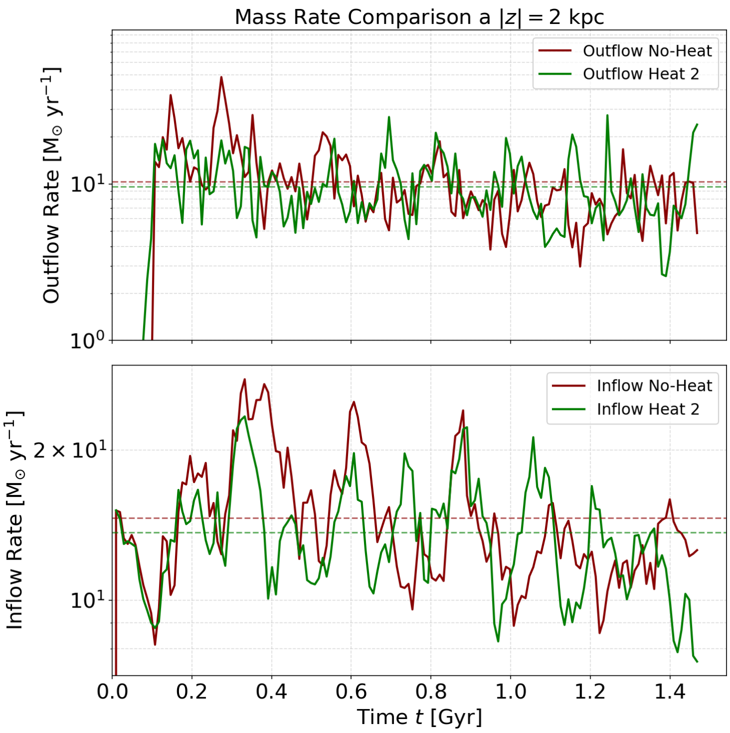

(b) No-Heat vs Heat 2

Figure 5.10: Mass inflow and outflow rates measured across a planar slab at $|z| = 2$ kpc above and below the disc, comparing the No-Heat baseline (dark red) with Heat 1 (blue) and Heat 2 (green). Dashed horizontal lines mark the time-averaged rate of each model. In both comparisons the heated runs follow the baseline closely, in mean level and amplitude of the fluctuations, indicating that the distributed heating does not significantly alter the disc-corona gas exchange in this region.

the two models (inflow $\sim 14\, M_{\odot}\, \mathrm{yr}^{-1}$, outflow $\sim 12\, M_{\odot}\, \mathrm{yr}^{-1}$). Whereas the comparison between No-Heat and Heat 2 shows an increase of the flows both inward and outward for the heating model by $\sim 2\, M_{\odot}\, \mathrm{yr}^{-1}$. At $r = 50$ kpc, by contrast, all three models compared are very similar, with values of the net inflow $\sim 1.7 - 1.8\, \mathrm{M}_{\odot}\, \mathrm{yr}^{-1}$, so still with the same level of net accretion towards the disc.

This shows that the heating impacts the gas circulation with some small changes at the disc-corona interface and at $r = 10$ kpc distance (just for Heat 2), while out to 50 kpc the differences are not remarkable: the corona keeps feeding the disc with a slight different rate depending on the heating prescriptions applied, yet not so effective to change the overall SFR.

5.8 Summary

In this Chapter we tested two distributed heating prescriptions, designed to inject into the coronal gas a power comparable to its cooling luminosity L_{cool}, and differing only in the timescale of the injection.

Both schemes slow modestly the depletion of the inner corona (hot gas) with respect to the No-Heat baseline, as shown by the mass evolution (Section 5.3), confirming that compensating the radiative losses does provide some thermal support. The two delivery timescales, however, prove almost equivalent: the continuous scheme (Heat 1) and the accumulated scheme (Heat 2) leave similar imprints across all the diagnostics considered—mass evolution, morphology, density profiles and gas flows. The only appreciable differences appear in the phase diagrams, where the accumulated injection of Heat 2 keeps the coronal gas closer to its characteristic temperature and suppresses the tail of cooling gas still present in Heat 1 and in the baseline; in the slower depletion of the hot gas previously mentioned; and also in the modest increase of the net flows (for Heat 2 model) found at $r = 10$ kpc distance from the disc. These are the most evident effects of the two distributed-heating models described in this chapter.

These results show that sustaining the corona requires not just enough energy, but energy delivered in a way that escapes immediate radiative losses. The distributed-heating prescriptions adopted here are moreover a deliberate idealisa-

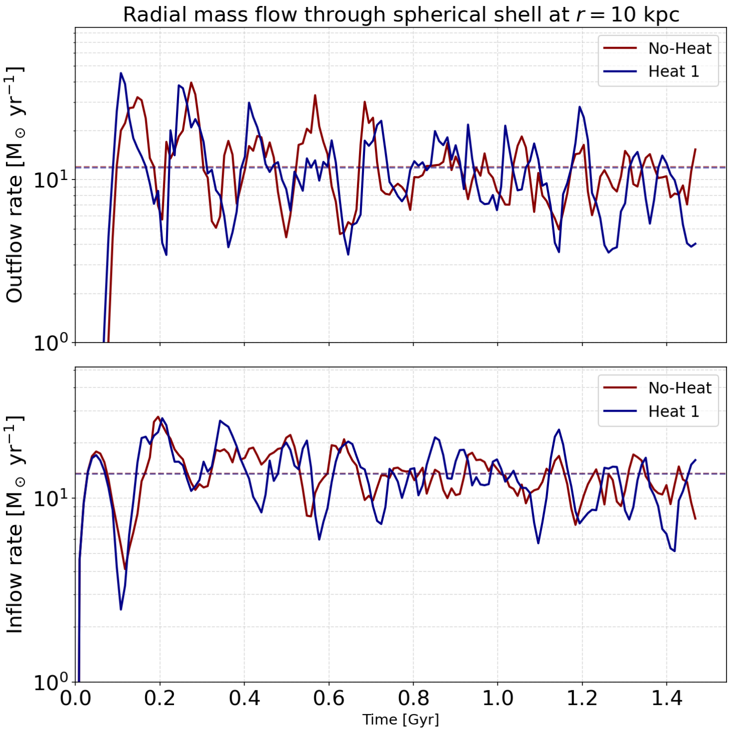

(a) No-Heat vs Heat 1

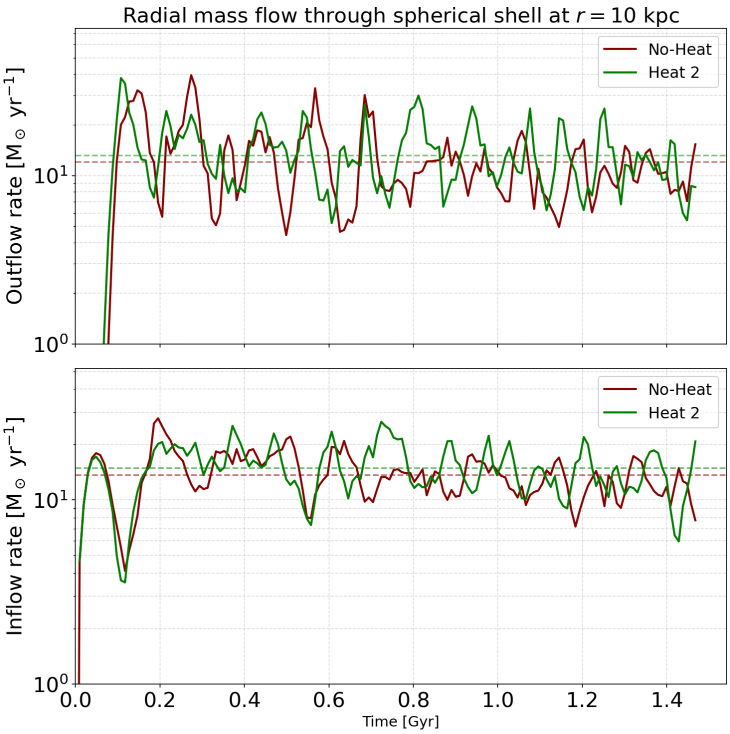

(b) No-Heat vs Heat 2

Figure 5.11: Mass inflow and outflow rates measured within a sphere of radius $r = 10$ kpc, comparing the No-Heat baseline (dark red) with Heat 1 (blue) and Heat 2 (green). Dashed horizontal lines mark the time-averaged rate of each model. The Heat 1 model follows the baseline closely, in mean level and amplitude of the fluctuations, on the contrary, Heat 2 enhance both outflows and inflows, showing a $\sim 2\, M_\odot\, \mathrm{yr}^{-1}$ higher rate with respect to No-Heat .

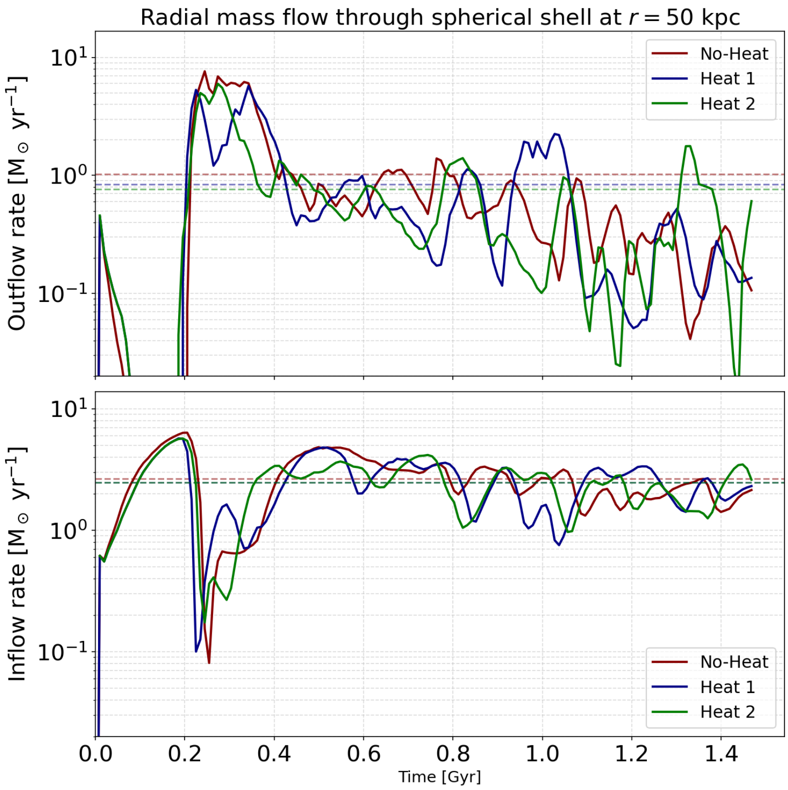

Figure 5.12: Mass inflow and outflow rates measured within a sphere of radius $r = 50$ kpc, comparing the No-Heat baseline (dark red) with Heat 1 (blue) and Heat 2 (green). Dashed horizontal lines mark the time-averaged rate of each model. All three models have similar trends indicating that at this distance the average rates are comparable and the heating considered are not remarkably affecting the corona.

tion: a real AGN does not heat the halo in a distributed fashion, but injects energy in a localised, bipolar way, as the Fermi bubbles of the Milky Way suggest (Section 5.1). To address this aspect, the next chapter is built on the same ICs as Heat 1 and Heat 2, but tests a physically motivated, localised bubble-heating model.

Chapter 6

The Nulsen and Fabian Bubbles Model

The last and most physically motivated of the three heating prescriptions, which we refer to as **NF Model**, implements a simplified version of the AGN feedback of the Auriga galaxy formation model (see Grand et al., 2017), which is in turn based on the bubble-injection scheme of Sijacki et al. (2007) and on the hot-gas flow theory of Nulsen & Fabian (2000). Unlike Heat 1 and Heat 2, where the energy is distributed across the whole coronal gas, here the heating is released into a smaller number of discrete, spatially localised bubbles that are stochastically placed in the hot CGM. This provides a more structured approximation of how a central AGN deposits mechanical energy into the surrounding corona. The total energy budget is set from the halo X-ray luminosity through a luminosity-temperature scaling relation (see Section 6.2).

The physical picture that motivates this model follows from Chapter 1: if the corona cools and feeds the central region, the resulting feedback can return energy to the CGM itself. The same qualitative behaviour is seen in the cosmological Auriga simulations, where switching off the radio-mode feedback leads to more coronal gas condensing onto the central galaxy (Grand et al., 2017). For simplicity, our NF Model brings only the energy-injection side of this mechanism (the deposition of energy into discrete bubbles in the corona) into an isolated setup, while leaving out the black-hole accretion that sets the available energy budget in Auriga. We will now go through some preliminary considerations useful to understand the theoretical basis of the NF model.

6.1 Preliminary considerations

The goal of this section is to set a preliminary path which explains how the cooling of the corona is responsible for the energy budget available for the AGN feedback. We do not do this to compute the budget used in our runs (which is fixed in a different and less sophisticated way, see Section 6.2), but to show that our model is anchored to the same physics as Nulsen & Fabian (2000) and Grand et al. (2017) works, that is the hot-gas flow theory and the radio-mode AGN feedback, respectively.

Before implementing our version of the NF model, we reproduce analytically the scaling relations on which the Auriga radio-mode feedback is built in order to fix the relevant physical scales of the problem: the halo virial velocity, temperature and radius; and the cooling function of the corona. These quantities are needed to obtain a rough estimate of the radio-mode efficiency factor R, defined as the ratio between the injected radio power of the AGN and the coronal X-ray luminosity (Eq. 6.7). This estimate serves only to place our budget within the Auriga framework: as better detailed in Section 6.2, the runs presented here do not use the accretion-based value of R, but adopt the limiting case $R = 1$, in which the reservoir is filled directly at the X-ray luminosity L_X. The accretion-rate calculation below therefore motivates the order of magnitude of our energy budget; it does not set it.

The starting point is the radiative cooling of the hot corona, and in particular the cooling function $\Lambda(T, Z)$, which measures how efficiently the hot gas radiates away its thermal energy: the energy lost per unit volume and time is $n_H n_e \Lambda$, so a larger Λ means faster cooling. The cooling function depends on both temperature and metallicity (see Eq. 1.5). For our estimate, we evaluate Λ at the representative temperature $T_{\rm vir} \sim 10^6$ K , computed in Eq. (1.4), and then treat it as a constant.

At this temperature the gas is collisionally ionised, so its atoms are stripped of electrons by collisions, and it can radiate in two ways: bremsstrahlung, the radiation emitted by free electrons deflected by ions; and line emission from heavy elements, which at $T \sim 10^6$ K still retain bound electrons that can be excited and radiate. Each species therefore adds a cooling channel, so Λ is strongly dependent on metallicity: at this temperatures it increases by roughly an order of magnitude from a metal-poor gas ($Z \lesssim 0.1\, Z_\odot$) to a solar composition ($Z \simeq Z_\odot$). To remain consistent with the a metal-poor corona (as ours), a value of $\Lambda \approx 2 \times 10^{-24}\,\mathrm{erg\,cm^3\,s^{-1}}$ is adopted in these computations.

We stress again that this value enters only in the illustrative estimate of R: the budget used in these simulations is taken directly from L_X (see Section 6.2), and is independent of it.

Having characterised how the gas cools, we now connect this cooling to the rate at which the corona feeds the central black hole, starting from the cooling time which is the quantity that indicates how fast the gas radiates away its thermal energy. From the ideal-gas equation of state, $P = \rho k_B T/(\mu m_p)$, together with the internal-energy relation $\rho u = P/(\gamma - 1)$, it is possible to derive also this timescale as

$$t_{\mathrm{cool}} = \frac{\rho u}{\Lambda\, n_H^2} = \frac{P}{(\gamma-1)\,\Lambda\, n_H^2} = \frac{\rho\, k_B T}{\mu m_p (\gamma-1)\,\Lambda\, n_H^2}, \tag{6.1}$$

where, for a fully ionised gas, the mass density and the hydrogen number density are related by $n_H = \chi\rho/m_p$.

This cooling time is the central ingredient of the Nulsen & Fabian (2000) model in which the gas feeds the central black hole through a cooling flow: as the hot corona radiates, it loses pressure support and settles inward. Two timescales describe this process: the cooling time t_{cool}, just described, and the flow time t_{flow}, that is the time the gas takes to drift to the centre. Nulsen & Fabian (2000) assume that in a steady cooling flow the two are comparable, $t_{\mathrm{cool}} \approx t_{\mathrm{flow}}$, the gas falls inward at roughly the rate at which it cools down. This condition sets the state of the gas at the edge of the accretion region of the central black hole and the resulting accretion rate $\dot{M}$, the amount of mass acquired per unit time, can be thus derived (see Nulsen & Fabian 2000, their Eq. 7, for further details).

Because the balance is anchored to t_{cool}, and $t_{\mathrm{cool}} \propto 1/\Lambda$, the accretion rate inherits the same dependence, $\dot{M} \propto 1/\Lambda(T_{\mathrm{vir}})$, at fixed halo temperature and black-hole mass. This accretion rate and the efficiency factor R are better discussed in Section 6.2.

6.2 Energy budget: from X-ray losses to accretion rate

In this section we explain the transition from the general treatment from which the NF model is based on (Nulsen & Fabian, 2000; Grand et al., 2017; Pratt et al., 2009), to the simplified version we adopted to set the energy budget used in this model.

The energy budget of the feedback is set by the dark-matter halo mass because this determines the X-ray luminosity L_X of the corona, which is therefore used to fix the rate at which energy is made available to the heating. In detail, a dark-matter halo is a gravitational potential well: the more massive it is, the deeper the well. The depth of this well is characterised by the *virial velocity* $v_{\rm vir}$.

A more massive halo features a higher $v_{\rm vir}$, quantitatively,

$$v_{\rm vir} = (10\, G\, H_0\, M_{200})^{1/3}\,, \tag{6.2}$$

where $M_{200} \simeq 1.6 \times 10^{12}\, M_\odot$ is the halo mass (defined as the mass enclosed within the radius where the mean density is 200 times the cosmic mean), while $H_0 = 100\, h\ km\ s^{-1}\ Mpc^{-1}$, with $h = 0.7$ in our simulations.

The temperature the gas must have, to be in equilibrium within the well, is the *virial temperature* $T_{\rm vir}$ and it scales as the square of the virial velocity

$$T_{\rm vir} = 35.9 \left(\frac{v_{\rm vir}}{\rm km\, s^{-1}}\right)^2\ {\rm K}, \tag{6.3}$$

where 35.9 is the pre-factor result of the constants ratio present in Eq.(1.4), computed with an atomic weight $\mu \simeq 0.59$. A gas in equilibrium in a more massive halo therefore settles into a hotter corona. For the halo considered in this work, $T_{\rm vir} \simeq 10^6$ K as assumed in Section 6.1. Since the halo is static in our simulations, $v_{\rm vir}$, $T_{\rm vir}$, the virial radius $r_{\rm vir} = v_{\rm vir}/(10 H_0)$, and the quantities subsequently derived, are constant in time.

A hot, diffuse plasma radiates in X-rays, and a hotter, denser corona shines more brightly. We do not compute this emission from first principles here; instead we use an *observed* relation between the X-ray luminosity and the temperature of the hot gas in galaxy groups and clusters, measured by Pratt et al. (2009),

$$L_X = C \left(\frac{T_{\rm vir}}{T_0} \right)^{2.62}, \tag{6.4}$$

with $T_0 = 5.8 \times 10^7\,{\rm K}$ (5 keV) and $C = 7.26 \times 10^{44}\,{\rm erg\,s^{-1}}$. So once the halo mass is chosen, the X-ray luminosity of its corona is completely determined. In our model galaxy $L_X = 1.86 \times 10^{40}\,{\rm erg\,s^{-1}}$.

We derive in parallel the energy that an AGN would return to the gas through its radio-mode feedback: the energy released by the central black hole is not radiated away efficiently but is deposited into the surrounding gas. This is the channel through which radio-mode AGN are thought to heat the corona (Grand et al., 2017).

The rate at which gas flows onto the black hole is taken from the hot-gas cooling-flow theory of Nulsen & Fabian (2000)

$$\dot{M} = \underbrace{\left[\frac{2\pi Q(\gamma - 1)\, k_B T_{\rm vir}}{\mu m_p\, \Lambda(T_{\rm vir})} \right]}_{\text{thermal state of the corona}} \underbrace{\left[G\, M_{\rm BH}\, y^2 \right]}_{\text{gravity of the black hole}} \mathcal{M}^{3/2}, \tag{6.5}$$

where $\gamma = 5/3$ is the adiabatic index (the ratio of specific heats of the gas), k_B the Boltzmann constant, μm_p the mean mass per gas particle, G the gravitational constant, and $y = m_p/\chi$ is the mass per hydrogen atom (with χ the hydrogen mass fraction introduced in Section 6.1). The constant $Q = 2.5$ is a dimensionless factor that depends on the details of the cooling-flow structure (the shape of the cooling function and how mass is deposited along the flow), for which we adopt the representative value given by Nulsen & Fabian (2000). We assign to the central black hole a mass $M_{BH} = 4 \times 10^6\, M_\odot$, equal to the Sgr A^* mass. Moreover, here $\mathcal{M}$ is the Mach number of the cooling flow at its outer edge, i.e. the ratio of the inflow velocity to the local sound speed (the speed at which pressure disturbances travel through the gas). Rather than determine it dynamically from the flow, we follow the Auriga simulations and fix it to the constant calibration value $\mathcal{M} = 0.0075$ (Grand et al., 2017). In general, the first bracket sets the thermal state of the corona, and in particular the inverse dependence on the cooling function, $\dot{M} \propto 1/\Lambda(T_{\rm vir})$, discussed in the previous section. The second sets the gravitational pull of the black hole. Following this framework, we infer that only a fraction of the rest-mass energy of the accreted gas, $\dot{M}c^2$ (with c the speed of light), is actually returned to the surrounding gas as heat. Two efficiencies set this fraction: the *radiative efficiency* $\epsilon_r = 0.2$, the fraction of the accreted rest-mass energy radiated by the black hole, and the *thermal-coupling efficiency* $\epsilon_f = 0.07$, the fraction of that radiated energy that couples to the gas as heat. Therefore, the radio power injected by the feedback can be estimated as

$$L_{\rm radio} = \epsilon_f \, \epsilon_r \, c^2 \, \dot{M}, \tag{6.6}$$

and the efficiency factor R, computed as the ratio between this radio power and the X-ray luminosity of the corona previously estimated, is given by

$$R = \frac{L_{\rm radio}}{L_X}. \tag{6.7}$$

The purpose of the comparison between these two luminosities is not to set the budget directly, but to *check the consistency* of our simplified treatment against the full Auriga prescription. The efficiency factor we obtain, $R \simeq 0.87$ (with $L_{\rm radio} = 1.61 \times 10^{40}$ and $L_X = 1.86 \times 10^{40}\,{\rm erg\,s^{-1}}$, from Eq. 6.6 and 6.4 respectively), is close to unity: this confirms that our simplified implementation reproduces the radio-mode AGN scaling correctly. Having established this consistency, in the runs we adopt the limiting case $R = 1$, filling the reservoir directly with the full L_X, and without overestimating the budget the radio-mode feedback would return to the corona. This is the simplest choice made and also in this way we use the largest energy budget, equal to the corona full X-ray losses.

We note that L_X is an analytic estimate from the observed $L - T$ relation based only on the virial temperature, whereas L_{cool} is measured in the simulation by summing the radiated power over all diffuse coronal cells $L_i = m_i \, u_i / t_{{\rm cool},i}$. The dense inner corona radiates more than the $L - T$ scaling predicts, yielding a cooling luminosity $\sim 4 - 5$ times larger than L_X . So, even with $R = 1$, the budget is not equal to the cooling luminosity but still is the maximum power the radio-mode feedback can return and is representative of our simplified model.

Linking altogether: the X-ray luminosity tells us how much energy the corona loses to radiation per unit time, the radio-mode AGN feedback of Grand et al. (2017) is built on the assumption that the energy returned to the gas balances these X-ray losses, up to a factor R. We collect this energy in what we call the *energy reservoir*. Since we set $R = 1$, at each global timestep (the highest active timestep), the reservoir is increased by the energy the corona would radiate in X-rays over that interval, by a value of

$$\Delta E = L_X \, \Delta t, \tag{6.8}$$

and the heating is later drawn from this reservoir as detailed in Section 6.3. As an a posteriori check on the scale of this budget, we can integrate the injected energy over the simulation and compare it with the energetics of the Fermi Bubbles (Section 1.3). Over $\Delta t \sim 1.5$ Gyr the total injected energy is $E = L_X \Delta t \approx 10^{57}$ erg, of the same order as the energy inferred for the Fermi Bubbles. This confirms that our budget lies on the energy scale the Galactic centre is observationally known to deliver.

6.3 Bubble injection scheme

The energy stored in the reservoir is not released continuously but accumulated and then injected in discrete events. The idea is that a steady, low-level supply of energy is stored and then released all at once as a set of hot bubbles, mimicking how an AGN deposits energy in intermittent outbursts. As mentioned, at each global timestep the energy $L_X \, \Delta t$ is added to the reservoir (Section 6.2). A heating event is triggered only once the reservoir has grown past a threshold $E_{\rm thr}$. This threshold is the energy required to raise a bubble-sized volume of gas, at a fixed reference density $\rho_{\rm core}$, to a fraction f_E of that energy threshold,

$$E_{\rm thr} = f_E \, u \, \rho_{\rm core} \, \frac{4\pi}{3} \, (f_R \, r_{\rm vir})^3 \, . \tag{6.9}$$

Here $\rho_{\rm core}$ is not the actual coronal density but a high reference density used only to set the scale of the threshold

$$\rho_{\rm core} = 10^4 \, \Omega_b \, \rho_{\rm cr} = 10^4 \, \Omega_b \, \frac{3H_0^2}{8\pi G}, \tag{6.10}$$

i.e. 10^4 times the mean cosmic baryon density, with Ω_b the baryon density parameter and $\rho_{\rm cr} = 3H_0^2/8\pi G$ the critical density, corresponding to $\rho_{\rm core} \approx 4.4 \times 10^{-27}\,{\rm g\,cm^{-3}}$ with $\Omega_b = 0.048$, $H_0 = 70\,{\rm km\,s^{-1}\,Mpc^{-1}}$. The factor $f_E = 0.05$ fixes the nominal energy given to each bubble. Requiring the reservoir to exceed $E_{\rm thr}$ ensures that energy is released only when enough has accumulated to heat at least one bubble. When the threshold is reached, up to $N_b = 15$ bubbles can be placed in the halo.

A single bubble is defined by a radius $R_b = f_R\, r_{\rm vir}$ (a fixed fraction $f_R = 0.1$ of the virial radius $r_{\rm vir} = v_{\rm vir}/(10H_0)$) and a specific energy $u = \frac{3}{4}v_{\rm vir}^2$ (the energy per unit mass needed to heat the gas to the virial temperature). Each is a sphere of radius R_b, positioned at a random, isotropically distributed location within a maximum distance $f_{\rm dist}\, r_{\rm vir}$ from the centre, with $f_{\rm dist} = 0.8$. The bubbles are therefore scattered through the corona rather than confined to the centre. Within each bubble, energy is added as heat (by raising the specific internal energy of the gas) only to cells that satisfy two conditions: the gas must be diffuse, $\rho < \rho_{\rm thr} = 0.01\, cm^{-3}$, which excludes the dense, star-forming disc, and it must already belong to the warm/hot phase, $u_{\rm therm} > 0.1\, u$. Since $u = \frac{3}{4}v_{\rm vir}^2$ corresponds to the virial temperature $T_{\rm vir} \simeq 10^6$ K, this condition is a temperature threshold at $\sim T_{\rm vir}/10 \simeq 10^5$ K. This value sits at the lower edge of the coronal temperature range: gas above it belongs to the warm ($10^4 - 5 \times 10^5$ K) and hot ($> 5 \times 10^5$ K) phases that make up the corona, while the cold disc gas ($< 10^4$ K), lying below this threshold, is excluded. The factor 0.1 is therefore not arbitrary, it places the cut just below the hot corona. In this way the heating targets the coronal gas and leaves the cold disc untouched.

The reservoir builds up over time and nothing happens until it exceeds the threshold $E_{\rm thr}$: as long as the accumulated energy is below $E_{\rm thr}$, it is simply carried over from one timestep to the next. Once the threshold is crossed, the energy is released into one or more bubbles, heating the eligible gas inside them. Whatever amount of energy remains in the reservoir after this release, including the case in which no eligible gas is found inside the bubbles, is not discarded but becomes the starting level from which the reservoir resumes accumulating for the following injection episode. In this way no injected energy is lost: the reservoir behaves as a buffer that stores energy until there is enough to heat the gas, and retains any surplus for later. So overall the mechanism of the NF model is very different from Heat 1 and Heat 2 which have instead a more distributed heating prescription.

6.4 Comparison with the distributed heating models

Heat 1 and Heat 2 models share the same energy budget and the same spatial extent, and differ from each other only in the timing of injection; while NF bubbles differs from both along all features, as summarised in Table 6.1.

	Heat 1	**Heat 2**	**NF Bubbles**
Energy budget (source)	local cooling (cell by cell)	local cooling (cell by cell)	halo L_X (Pratt/NF)
Spatial extent	whole diffuse corona	whole diffuse corona	$N_b = 15$ discrete bubbles, $r < f_{\rm dist} r_{\rm vir}$
Timing	local timestep (continuous)	global timestep (accumulated)	reservoir threshold (impulsive)

Table 6.1: The three heating prescriptions compared along the features that distinguish them: the source of the energy budget, its spatial extent, and the timing of injection. Heat 1 and Heat 2 share budget and spatial extent and differ only in timing; NF Bubbles differs from both along all three features.

The first and most important difference concerns the *source* of the energy budget: in the Heat 1 and Heat 2 the heating compensates the radiative cooling locally, cell by cell, so that the injected power tracks the instantaneous thermodynamic state of each gas element. In NF model the budget is instead a single, global quantity, the halo X-ray luminosity estimated from a scaling relation, which is fixed by the halo mass and is insensitive to the local state of the gas (see Section 6.2). The two distributed models therefore test whether replacing the radiated energy *where it is lost* is sufficient to stabilise the corona, while NF model tests whether a discrete injection from a global budget achieves the same outcome.

The second difference is the *spatial distribution* of the heating: Heat 1 and Heat 2 act smoothly throughout the diffuse corona, whereas NF concentrates its injection into specific localised regions. The expected signatures are correspondingly different: a smooth, distributed prescription should support the corona uniformly, whereas the bubble heating produces a patchy temperature structure whose subsequent evolution is one of the questions the simulations are meant to address.

The third difference is the *timing of injection*: Heat 1 continuously on the local hydrodynamic timestep; Heat 2 accumulates the same budget and releases it on the global timestep; NF model injects whenever the accumulated energy reservoir reaches the set threshold. The injection of NF is therefore not tied to a regular time interval at all, but to the rate at which the budget fills the reservoir.

6.5 Limitations of the model

Several assumptions of the NF Bubbles model should be kept in mind when interpreting its results.

First, the luminosity–temperature relation of Equation 6.4 is calibrated on galaxy clusters, with temperatures of order a few keV, whereas it is applied here to a Milky Way-mass halo with $T_{\rm vir} \sim 0.09$ keV, roughly two orders of magnitude below the 5 keV pivot of the relation.

Second, because the halo is static, the heating budget is constant in time and does not respond to the instantaneous state of the gas or to the actual cooling losses of the corona. In this sense the budget is a fixed input set by the halo mass rather than a self-consistently regulated quantity.

Third, we adopt the limiting case of the efficiency factor $R = 1$ rather than computing it from the fully self-consistent Auriga prescription. Our computed value of $R \simeq 0.9$ is close to unity and represents a valid approximation but it does not follow from a complete treatment of the black hole accretion process.

Therefore, as already pointed out, our implementation of the NF model simplifies the original model of Nulsen & Fabian (2000) adopted in Auriga (Grand et al., 2017), and should be interpreted with approximations considered. In the following Sections, we discuss the outcomes obtained with this model.

6.6 Mass evolution

To follow how the NF heating prescription affects the gas mass distribution, we track this quantity divided in three temperature phases as a function of time: cold ($T < 10^4$ K), warm ($10^4 < T < 5 \times 10^5$ K) and hot ($T > 5 \times 10^5$ K) gas. Figures 6.1 and 6.2 show the gas mass evolution through time of the four models considered, one panel per phase, within two distances respectively, $r < 100$ kpc and $r < 240$ kpc. The two figures are reported together to show that, for the NF model, the comparison between them is a key diagnostic to better picture whether the gas mass is transferred from the inner to the outer corona.

What is noticeable within $r < 100$ kpc is that, after an initial transient ($t \lesssim 0.1$ Gyr) associated with the relaxation of the ICs, all models share the same qualitative behaviour: the warm+hot gas is depleted while the cold phase grows, which is a signature of the cooling flow (Figure 6.1). Within this distance, the two distributed-heating models separate modestly in how effectively they slow this depletion: after $t \sim 0.6$ Gyr Heat 1 and Heat 2 show $\sim 0.5 \times 10^{10}$ $M_{\odot}$ more hot-gas mass than the No-Heat model, but going towards the end of the run they all three reach the same plateau of $\sim 7.7 \times 10^{10}$ $M_{\odot}$. Within $r < 240$ kpc (Figure 6.2), after $t \sim 1.1$ Gyr , Heat 1 and Heat 2 retain a larger hot reservoir than No-Heat ($\sim 0.2 \times 10^{10}$ $M_{\odot}$), consistent with the cell-by-cell compensation of the radiative losses displayed in the previous chapter: the distributed heating provides a measurable, albeit modest, thermal support to the corona (Heat 2 $\gtrsim$ Heat 1 >No-Heat) retaining more hot gas as the run goes to the end.

The NF model behaviour also depend on the distance, like the other models, but it show different features. Within the $r < 100$ kpc distance (Figure 6.1), it shows the lowest hot-gas mass fraction of the four models, depleting this region faster even than the No-Heat run, consistent with a possible outward motion of the hot gas due to the action of the heating bubbles, but also with the cooling flow which drives hot gas inward. Within $r < 240$ kpc (Figure 6.2), conversely, NF model develops the largest hot reservoir of the four models, with a peak around $t \simeq 1.1$ Gyr, before declining again towards the end of the run. This excess is though transient: the bubbles do not establish a new hotter equilibrium of the outer corona, but inflate its hot content temporarily. The inner deficit together with the outer excess indicates that the bubbles are probably transporting hot gas outward rather than enhancing its conversion into colder phases.

At the same time, if we look at the warm phase evolution within $r < 240$ kpc, it appears consistent with this picture: the NF warm gas drops to a pronounced minimum around the time the hot mass peaks. The decline of the warm phase and the rise of the hot phase over this interval are of comparable magnitude, as expected if a substantial fraction of the warm gas is heated across the $T = 5 \times 10^5$ K threshold into the hot phase by the bubbles. It is not established from the mass evolution alone whether the hot excess in the outer corona is mostly due to a warm$\rightarrow$hot conversion rather than a motion outwards of the hot gas; both are compatible with the data. We return this point in Section 6.10.2.

The result of this section is twofold. The distributed prescriptions (Heat 1 and Heat 2) provide a modest thermal support that slows the depletion of the hot corona, most visibly at large radii ($R \gtrsim 100$ kpc) and late times ($t \gtrsim 1.0$ Gyr). Whereas the bubble model, at the adopted budget, shows a depletion of hot gas mass within $R < 100$ kpc , but also the largest reservoir within $R < 240$ kpc, with a possible hint of a motion of gas outwards better investigated in Sections 6.9 and 6.10.2.

6.7 Morphological evolution of the gas disc

Having quantified how the bubble model affects the global gas budget, we now turn from the integrated masses to the spatial distribution of the simulated galaxy, i.e. shape and structure, to investigate whether the NF model has some effects on these features, as previously done in Chapter 5 for the other two heating models.

6.7.1 Surface Density maps

First of all, we describe the structures of the disc and the region ~ 5 kpc above and below it, through the projected density and mass-weighted temperature maps, before moving back out to the corona in Section 6.9.

In the face-on view the two models are remarkably similar: the same spiral arms,

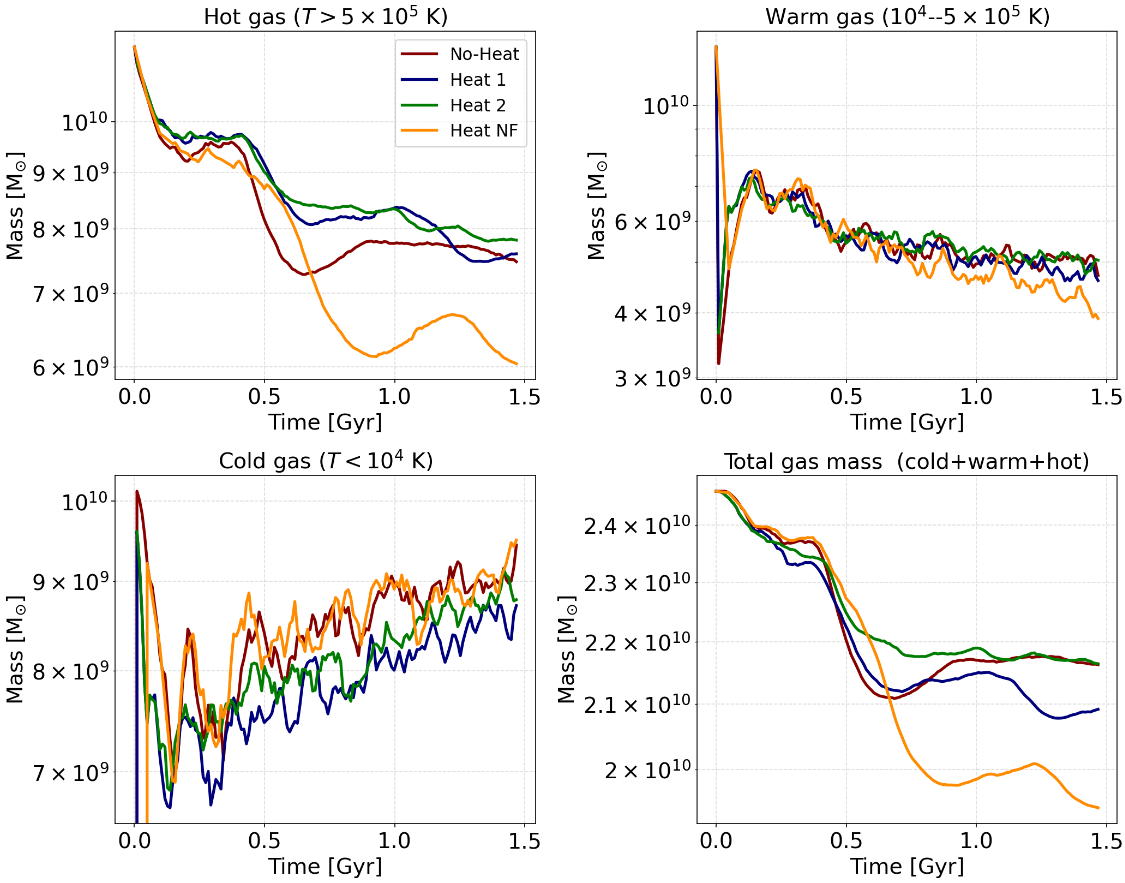

Figure 6.1: Evolution of the gas mass split into cold ($T < 10^4$ K), warm ($10^4 < T < 5 \times 10^5$ K) and hot ($T > 5 \times 10^5$ K) phases, computed within a sphere of radius $r = 100$ kpc, for the No-Heat, Heat 1, Heat 2 and NF models. Each panel shows a phase of the gas, while each curve is representative of each model. It is noticeable how the cold phase is increasing through time, while after $t \sim 0.6$ Gyr the hot phase shows a rapid decrease (in all the models), sign of depletion of the inner corona.

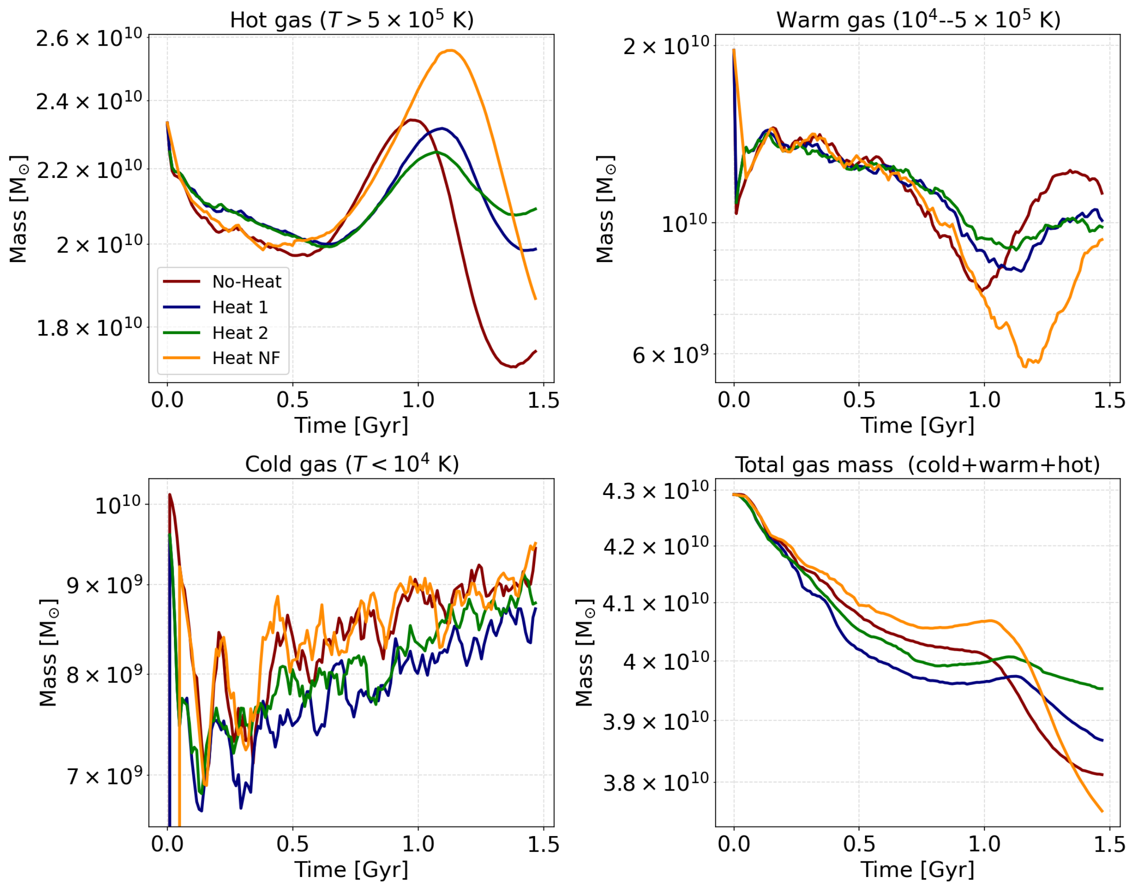

Figure 6.2: Evolution of the gas mass split into cold ($T < 10^4$ K), warm ($10^4 < T < 5 \times 10^5$ K) and hot ($T > 5 \times 10^5$ K) phases, computed within a sphere of radius $r = 240$ kpc, for the No-Heat, Heat 1, Heat 2 and NF models. Each panel shows a phase of the gas, while each curve is representative of each model. It is noticeable how the cold phase is increasing through time, while after $t \sim 1.0$ Gyr the hot phase of all the models shows a pronounced bump, in particular the NF model showing the highest peak.

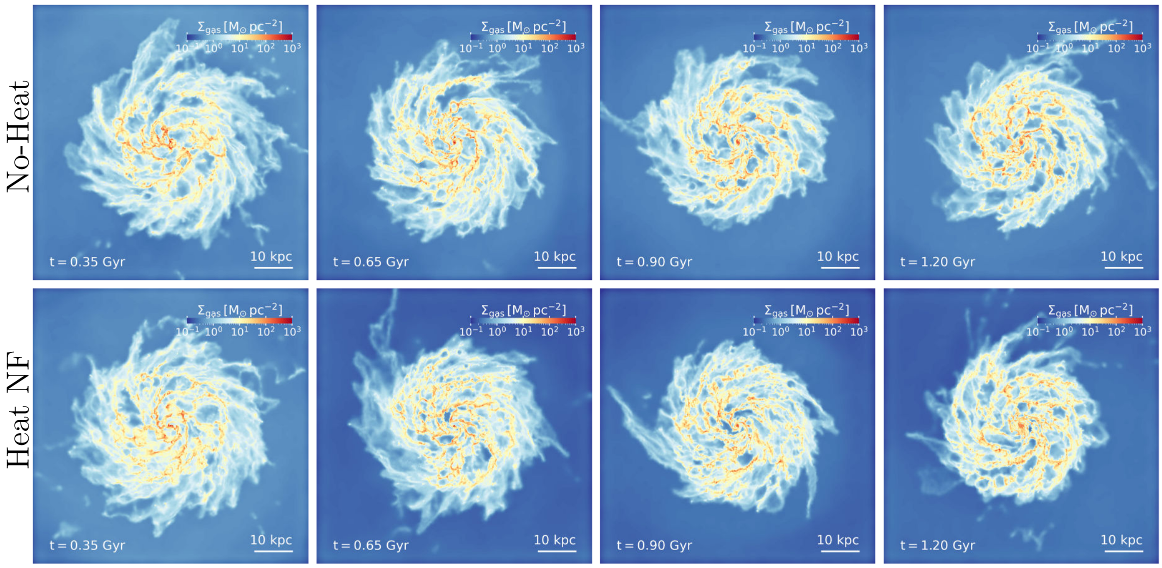

Figure 6.3: Face-on projected gas surface density $\Sigma_{\rm gas}$ for the No-Heat and NF models (rows) at four different times (columns). The four times are chosen to match the epochs displayed also in Figure 5.2. The disc structure remains closely similar across the two models at every time, although the NF model seems to be slightly more compact than the No-Heat model, showing a smaller area of the disc.

clumpy structure and central concentration are present in all of them. The star-forming disc is similar between the runs, confirming that the dense disc is regulated by local processes rather than by the coronal heating prescription, consistent with the evolution of the star formation histories (see Section 6.10.1).

The edge-on view shows that the two models develop a thin dense mid-plane surrounded by an extra-planar layer (the region $\sim$ 5 kpc above and below it) of diffuse gas and filaments, the imprint of the galactic fountain. The detailed structure of this extra-planar gas varies with time and model: No-Heat and NF models share very similar structure evolving through time towards a more compact and less filamentary one. In general, there are not significant differences between the models for the disc and the extra-planar region.

6.7.2 Mass-weighted temperature maps

We now investigate the mass-weighted temperature maps, to have a visual comparison between the No-Heat and NF model disc structure, but this time to check the

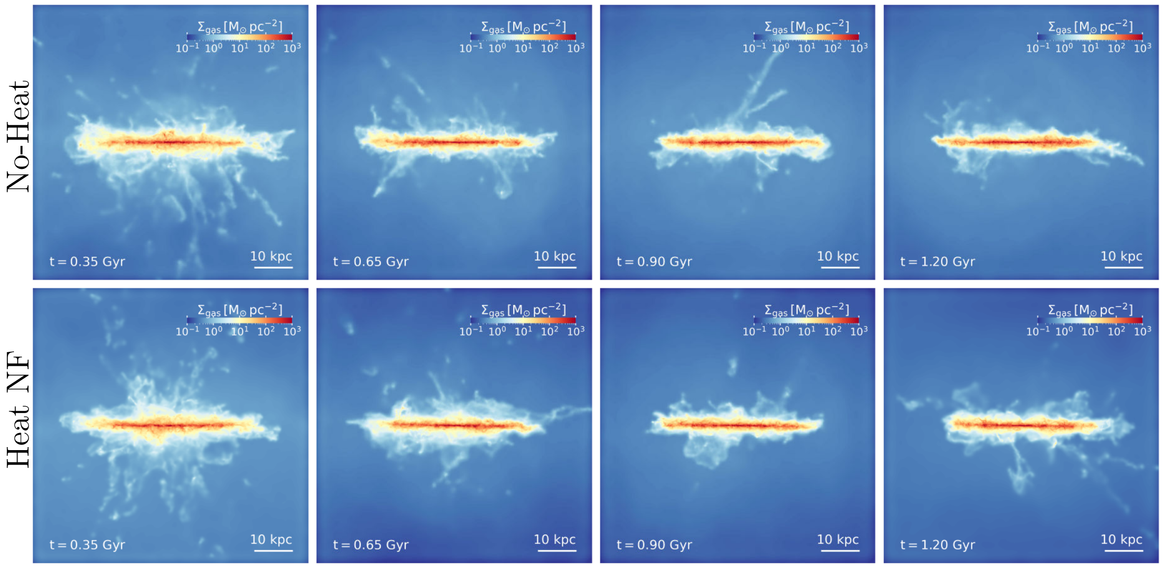

Figure 6.4: Edge-on projected gas surface density $\Sigma_{\rm gas}$ for the No-Heat and NF models (rows) at four different times (columns). The four times are chosen to match the epochs displayed also in Figure 5.2. The disc structure remains closely similar across the two models at every time, reflecting no remarkable differences with the features found through the face-on maps.

different phases of the gas through the different temperatures.

The face-on temperature maps (Figure 6.5) show that the disc has the same multiphase structure between the two models displayed: cold filaments ($T \sim 10^3 - 10^4$ K) trace the spiral arms and mark the star-forming sites, embedded in a warmer inter-arm medium ($T \sim 10^4 - 10^5$ K), with localised hot patches associated with recent stellar feedback. As for the density, the disc temperature structure appears similar to the No-heat prescription, consistent with the fact that this model acts more on coronal gas at larger radii ($r > 100$ kpc).

The edge-on temperature maps (Figure 6.6) confirm the morphological picture given by the density. In both models a thin, cold disc ($T \lesssim 10^4$ K) evolves through the time (from $t = 0.35$ Gyr to $t = 1.20$ Gyr) together with filaments really similar in the two prescriptions. Moreover the disc is embedded in the hot coronal gas ($T \sim 5 \times 10^5$K), whose temperature is also comparable in the two cases considered, showing no distinct differences in these temperature maps.

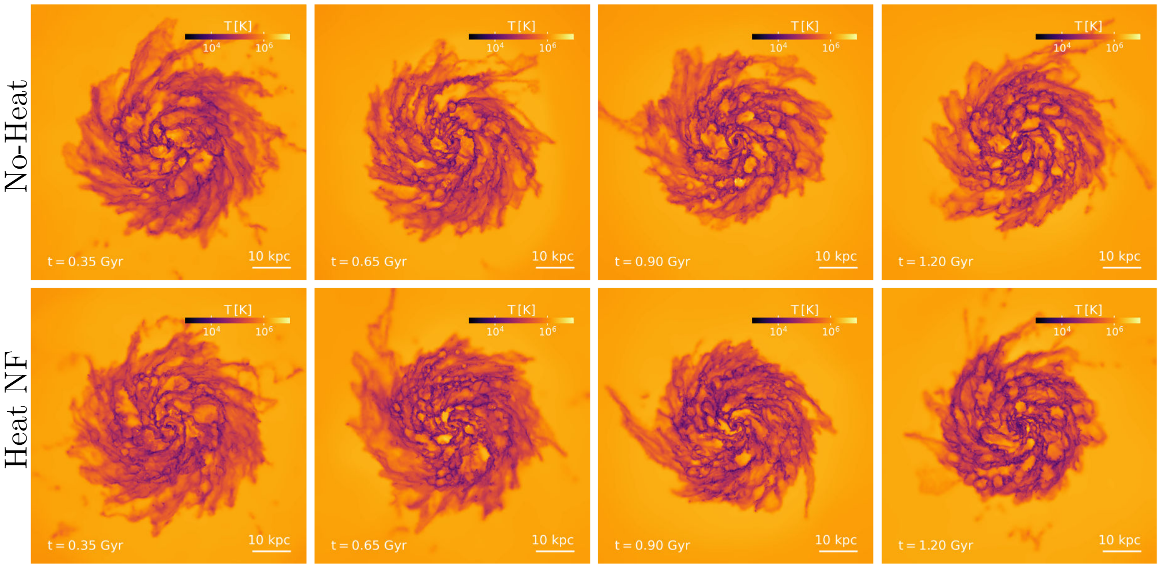

Figure 6.5: Face-on projected, mass-weighted gas temperature for the No Heat and NF models (rows) at the same four times as Figure 5.2 (columns). As for the density, the temperature structure of the disc is closely similar across the two models at every time; the heating prescriptions act on the diffuse coronal gas and leave the temperature of the dense, star-forming disc unaffected.

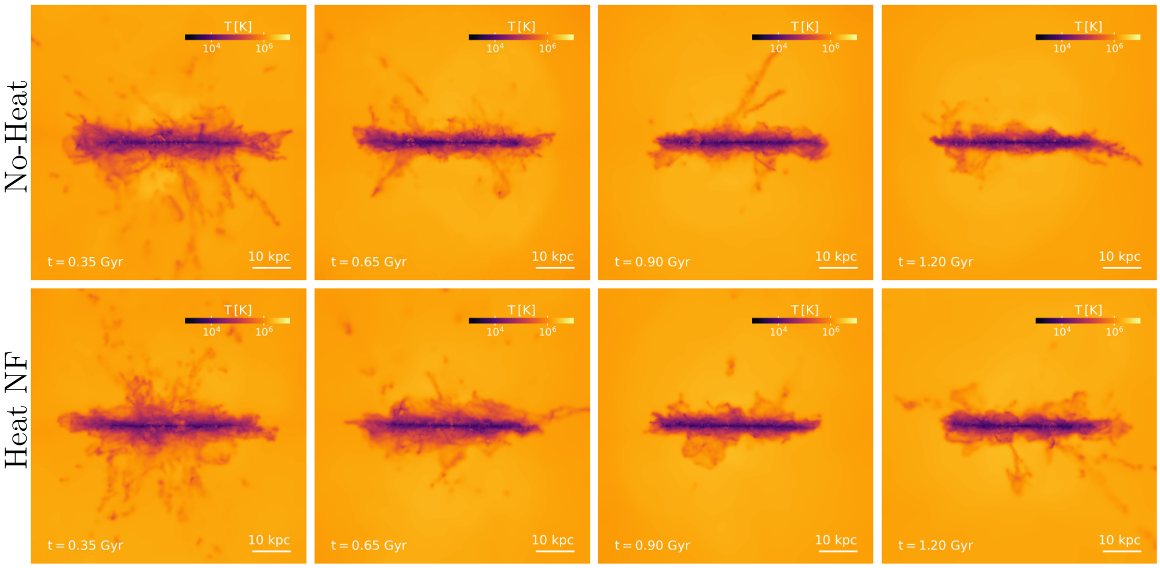

Figure 6.6: Edge-on projected, mass-weighted gas temperature for the No Heat and NF models (rows) at the same four times as Figure 5.2 (columns). Seen edge-on, the structure of the disc and the surrounding hot gas is closely similar across the two models at every time, denoting that the NF model is not affecting remarkably the disc temperature distribution.

6.8 Thermodynamic and Multiphase ISM

The structure of the disc and the extra-planar region has been shown to be remarkably similar among the two models compared. We now move to the investigation of a potential thermodynamic imprint of the NF model on the cooling flow from the coronal region to the dense, cold disc.

6.8.1 Phase Diagrams

Figure 6.7 shows the mass-weighted phase diagrams of the four models compared, at $t = 0.35$ and 1.20 Gyr to better capture possible differences, with the density–temperature plane divided into gas phases, using the same definitions as in Section 5.5.1. In all four models the gas occupies the same regions: the hot coronal gas in region (a), the horizontal accumulation at $\log T \simeq 4.1$ where hydrogen recombination slows the cooling, the diagonal branch connecting it to the cold, dense star forming gas, and the diffuse plume of feedback-heated gas.

The bubble heating model introduces no remarkable changes in how the gas mass populates the plane. Except for a subtler difference that distinguishes NF (as well as Heat 2) from the No-Heat and Heat 1 models within the coronal region (a): at $t = 1.20$ Gyr in No-Heat and Heat 1 a faint tail extends below the hot coronal feature towards lower temperatures ($\sim 10^4$ K), tracing gas cooling out of the hot phase, while for both Heat 2 and NF models, this tail is not present. As we explained in Section 5.5.1, in the former model, the distributed heating compensates enough the local cooling to keep the coronal gas near its characteristic temperature, so the gas does not show a cooling flow as prominent as the unheated and Heat 1 models. For what concerns the NF model, in the previous sections we already showed, at $t = 1.20$ Gyr, a hot gas reservoir larger than the other three models, probably by bringing gas outwards from the inner corona. So, in the lack of a tail of gas transit from region (a) towards regions (c) and (d), we find a visual confirmation of where the heating model is acting (see also Sections 6.6 and 6.10.2). The same visual feature of Heat 2 and NF models, thus reflects two different behaviours: thermal support in one case and transport of hot gas outwards in the other.

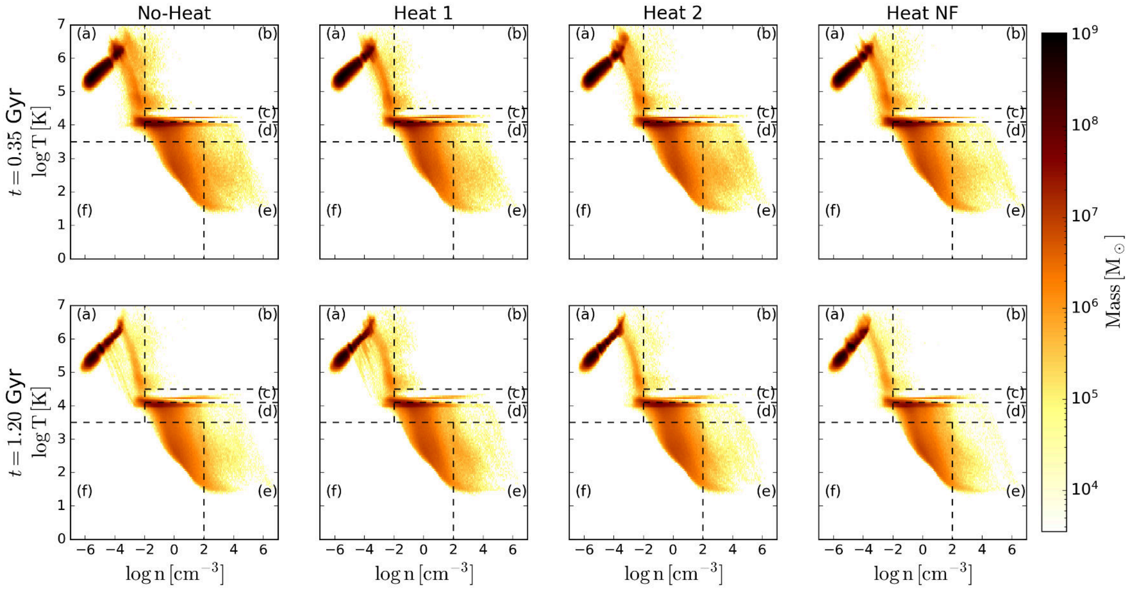

Figure 6.7: Phase diagrams evaluated at two times, $t = 0.35$ (top row) and $t = 1.20$ Gyr (bottom row), for the No-Heat, Heat 1, Heat 2 and NF models (from left to right). They share the same features for most of the regions depicted, exception made for a slight difference in region (a) where the Heat 2 and NF models do not show the tail propagating from the high-mass spot towards regions of lower temperatures, (c) and (d).

6.8.2 Mass and Volume Probability densities

Figure 6.8 shows the distribution of gas mass (left panel) and volume (right panel) as a function of number density, partitioned into cold, warm and hot phases, for the NF model. The structure is the familiar multiphase picture. The mass distribution is bimodal: a peak at low density ($\log n \simeq -4.5$) made up of the hot coronal gas, and a broad peak at high density ($\log n \simeq 0$) dominated by the cold disc gas, with the warm phase bridging the two at intermediate densities. The volume distribution is instead dominated at low density by the hot gas, which fills most of the volume while contributing little mass, whereas the cold dense gas dominates the mass at high density but occupies a negligible volume. Compared with the No-Heat baseline PDF (Figure 4.7), the distributions are essentially unchanged: the same phases occupy the same density ranges with the same relative weights. The bubble heating does not alter how the baryonic mass and volume are partitioned across the density regimes, so the probability distributions add no distinguishing feature between the models. The heating bubbles have a more important effect on the redistribution of

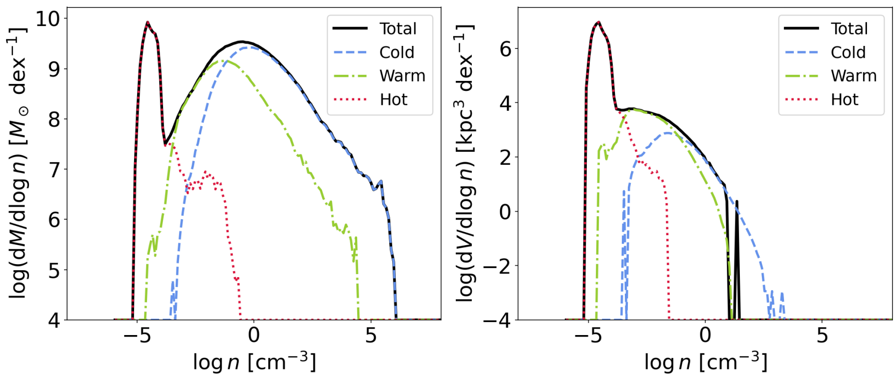

Figure 6.8: Distribution of gas mass (left panel) and volume (right panel) as a function of number density $\log n$ for the NF model at $t \simeq 1.5$ Gyr , computed within a sphere of radius $r = 100$ kpc. The gas is partitioned into total (black solid), cold (blue dashed), warm (green dash-dotted) and hot (red dotted) phases, defined as gas in the temperature ranges $T < 10^4$ K, $10^4\,\mathrm{K} < T < 5 \times 10^5\,\mathrm{K}$, and $T > 5 \times 10^5$ K, respectively.

the corona which is better captured by the mass evolution profiles, cumulative mass profiles and the inflows/outflows (see Sections 6.6, 6.9, and 6.10.2).

6.9 Coronal gas density profiles

The mass evolution shows the amount of hot gas retained by each model within a given distance, but not where this gas is located. To investigate this, we compute the spherically-averaged density profile of the coronal gas ($T > 5 \times 10^5$ K) in logarithmically-spaced shells out to the virial radius. If the heating models add hot gas, this will show up as a vertical shift of the profiles at fixed radius; if instead they move gas from one region to another, it will show up as a depletion at some radii compensated by an excess at others. The profiles, at $5 < r < 100$ kpc and $100 < r < 240$ kpc, for the four models at four reference times are shown in Figures 6.9 and 6.10.

Figure 6.9 shows that, at the earliest time ($t = 0.35$ Gyr), the four profiles are highly affected by the initial transient, on the contrary, they are broadly overlapped in Figure 6.10. From $t \simeq 0.65$ Gyr onward, the coronal profiles follow different trends as a function of radius: for the NF model in the inner part ($r \sim 10 - 100$ kpc), the hot-gas density is positioned below the other three models; beyond $r \simeq 100$ kpc all four models, instead, overlap at the two earliest times, while at the latest time ($t = 1.20$ Gyr) the NF profile sits slightly ($\sim 10^{-4}$ cm^{-3}) above the others across the outer shells. That excess is more appreciable in the integrated hot mass presented in Section 6.6, in this analysis is a less remarkable feature. Taken together, the density profiles of the corona and the integrated masses, are consistent with the picture that some fraction of the hot gas mass has moved outwards.

Instead, the distributed-heating models behave differently. Heat 1 and Heat 2 maintain a hot-gas density comparable to, or slightly above, the No-Heat baseline within $5 < r < 100$ kpc distance, consistent with the modest local thermal support identified from the mass evolution (Figure6.1). It has to be said that, within $5 < r < 10$ kpc, Heat 1, Heat 2 and No-Heat (but also NF) coronal density profiles are very noisy due to the interaction with the disc, the transient enhancement of Heat 2 at the smallest radii ($r \lesssim 10$ kpc) seen at $t = 0.90$ Gyr, is a signature of this behaviour and indeed does not persist to $t = 1.20$ Gyr. After $r \sim 10$ kpc, instead, the distributed-heating models overlap with No-Heat into more stable and accurate trends, exception made for the earliest time where the initial transient noise is still strong even at $r \sim 30$ kpc. The profiles therefore deliver one half of the picture: at $r \lesssim 100$ kpc the NF model is depleted relative to the other three. The complementary half, where the displaced gas ends up, it cannot be deduced from these kind of plots.

A clearer view comes from two changes to the analysis followed so far. First, instead of tracking the hot gas alone, we sum the warm and hot phases: if the NF bubbles push hot gas outward and let part of it cool into the warm phase, then following the hot gas alone would lose track of the gas that has moved but is no longer hot, so summing the warm and hot phases follows all the displaced gas. Second, instead of the density profile we use the *cumulative mass* $M_{\rm warm+hot}(< r)$, which is the total warm+hot gas enclosed within a radius r that is the gas summed from the centre out to r. By construction this quantity grows with r, since larger radii enclose more gas.

The cumulative curves for each model are shown in Figure 6.11 and divided into four panels showing the same four times of analysis consistent with the evaluation of the coronal profiles previously displayed. At the two earliest times (0.35 and 0.65

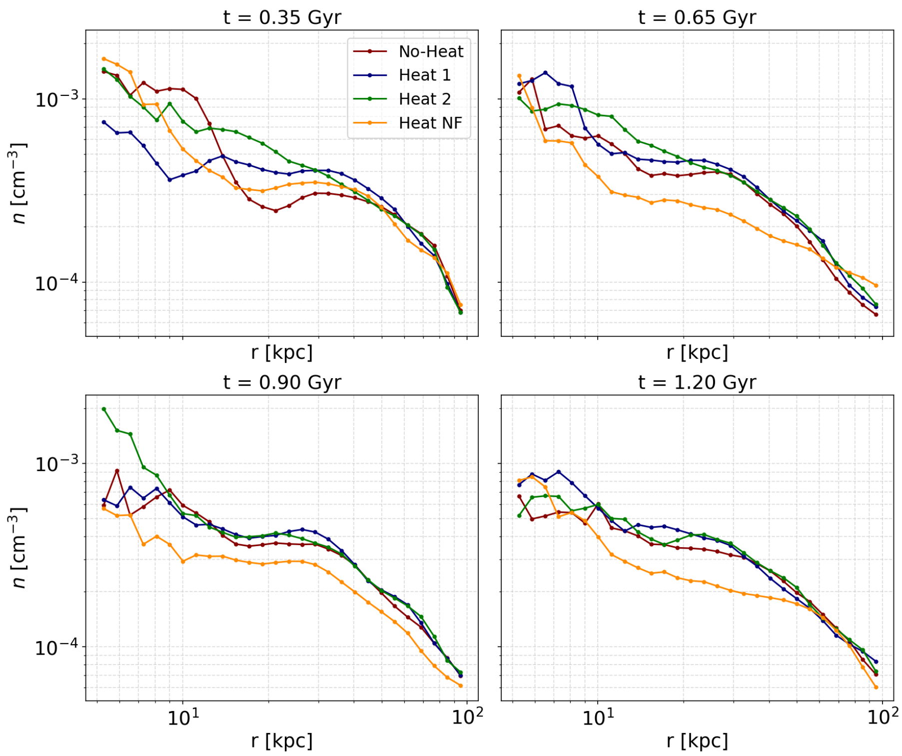

Figure 6.9: Evolution of the hot-gas ($T > 5 \times 10^5$ K) density profile, computed within a shell of width $5 < r < 100$ kpc, for the No-Heat, Heat 1, Heat 2 and NF models. Each panel shows four curves representing these models, while the panels are representative of four specific times (indicated at the top of each panel) consistent with Figure 5.8. The times are therefore chosen to coincide with the most significant changes in the gas content of the corona, and to remain consistent with the description given in the previous Chapter.

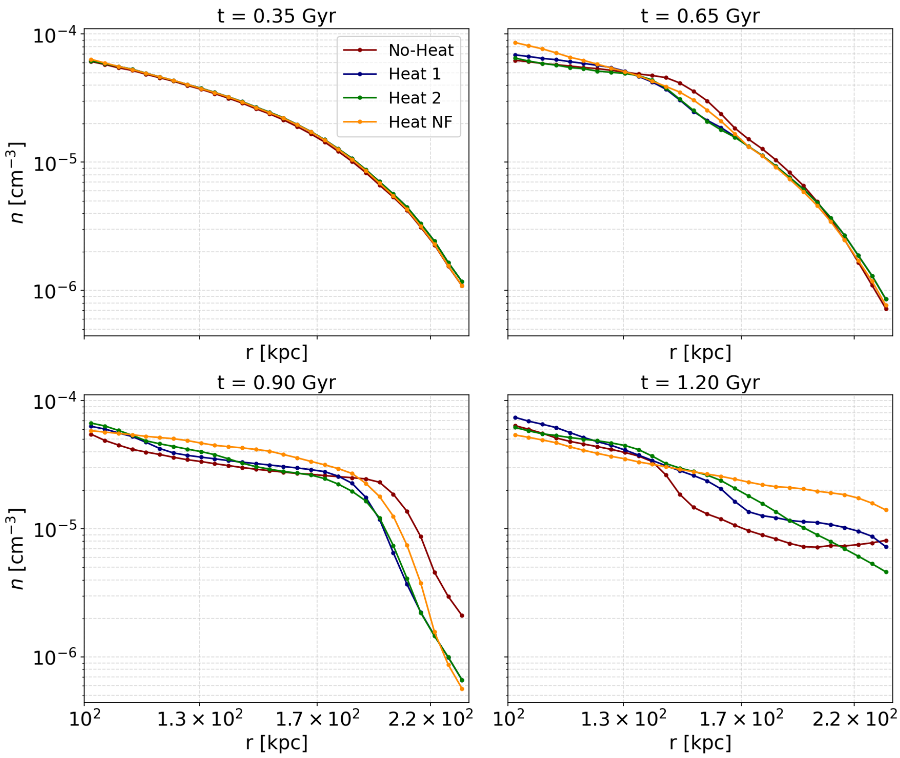

Figure 6.10: Evolution of the hot-gas ($T > 5 \times 10^5$ K) density profile, computed within a shell of width $100 < r < 240$ kpc, for the No-Heat, Heat 1, Heat 2 and NF models. The figure shows four curves representing these models, while the panels are representative of four specific times consistent with Figure 5.8. The times are therefore chosen to coincide with the most significant changes in the gas content of the corona, and to remain consistent with the description given in the previous Chapter.

Gyr) the four cumulative curves nearly coincide at all radii, so no displacement is yet measurable. At $t = 0.90$ Gyr and $t = 1.2$ Gyr, the NF curve lies below the others within $10 < r < 180$ kpc and $10 < r < 230$ kpc radius respectively. A lower cumulative curve means less gas is enclosed up to that point. The NF model then holds less warm+hot gas than the other models, in this range. Further out, instead, the curve rises again and crosses the others, and because the total mass (the value reached at the outermost radius) is similar in all models, this indicates that it ends up enclosing more mass overall at large radii. The NF profile therefore shows an inner depletion of warm+hot gas compensated by an outer excess, which we intepret as a signature of gas displaced outward.

While the inner-deficit / outer-excess shape establishes that gas is displaced outward, to quantify how far the displacement reaches, we define the *crossover radius* as the radius at which the NF cumulative profile meets the No-Heat one: it marks the boundary between the depleted inner region and the outer region, and thus the extent reached by the displaced gas at each time. At the earliest time this radius is not well defined, since the curves overlap throughout, but for the other times the crossover migrates outward with time, from $\simeq 90$ kpc at 0.65 Gyr, to $\simeq 150$ kpc at 0.90 Gyr, and $\simeq 210$ kpc at 1.20 Gyr. This outward migration is driven by the NF heating, which which keeps pushing hot gas out of the inner corona so that the front of displaced warm+hot gas reaches progressively larger radii, the growing crossover radius is therefore a measure of the outward transport of gas over time.

6.10 Star formation and gas flows

6.10.1 Star Formation

Having characterised how the bubble heating reshapes the distribution of hot-gas mass throughout the corona, we now ask whether the effects of this model reach the star-forming disc, altering the SFR.

As in the distributed heating models (Chapter 5), the bubble model leaves the SFR mostly unchanged, as shown in Figure 6.12: after the initial peak driven by the early accretion of coronal gas (see Barbani et al. 2023), all four histories settle

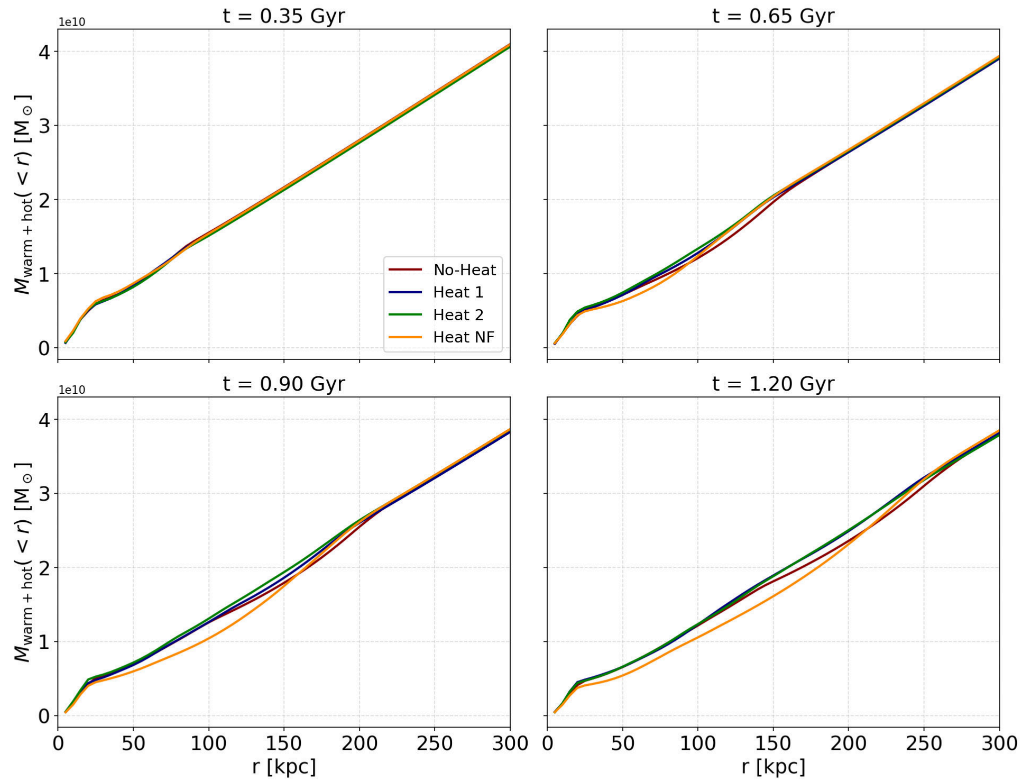

Figure 6.11: Cumulative mass profile for the No-Heat, Heat 1, Heat 2 and NF models computed within a sphere of radius $r = 300$ kpc, larger than the virial radius to check for excess of gas beyond this distance. Each panel corresponds to a specific time (consistent with Figures 6.9 and 6.10) in the evolution of the cumulative mass, while each curve represents a specific model. The profile is a direct measure of how the warm+hot gas mass is distributed within a region of radius r: the larger the radius, the larger the mass enclosed.

onto a comparable plateau, oscillating around $\sim 3-5\,M_{\odot}\,\mathrm{yr}^{-1}$. This shows that the NF model settles onto a marginally lower plateau than the distributed-heating models. This offset is small and the trend is not sharp, but it is systematic and consistent with the slightly reduced cold-gas content by depleting the inner corona, so the NF model leaves marginally less gas available to cool and reach the disc. Accordingly, the NF SFR is the closest to the observational estimate for the Milky Way ($\sim 1-3\,M_{\odot}\,\mathrm{yr}^{-1}$, Bland-Hawthorn & Gerhard 2016); like the other models it stays bounded to the plateau level rather than growing, signature of the feedback-regulated star formation, but it still exceeds he observed values in the first ~ 0.6 Gyr. The overall picture, consistent with the rest of the chapter, is that the bubble heating acts on the hot corona on large scales, depleting its inner part and driving hot gas outward (as explained in Sections 6.6 and 6.9 and further investigated in Section 6.10.2), while its impact on the disc is much weaker. It does not significantly alter the supply of gas reaching the disc, the effect on the SFR is real but minor. The large-scale rearrangement of the coronal gas couples only weakly to the star formation in the disc: the NF heating reshapes the hot corona, yet this leaves only a minor imprint on the gas supply and hence on the SFR.

6.10.2 Inflow / outflow evolution

We investigate now if the gas supply to the disc remains unchanged, as the SFR, and we verify at the same time if the hot-gas flows can establish the origin of the outer hot-mass excess found in the gas mass evolution of the NF model (see Figure 6.2). We compute the inflow and outflow rates at the same locations used for the distributed models (Chapter 5): the planar slab at $|z| = 2$ kpc, probing the disc-inner corona exchange, and two spherical shells of the corona within $r = 50$ kpc and $r = 100$ kpc.

At the disc-inner corona interface the picture is the same as for every other model. The inflow and outflow across the $|z| = 2$ kpc slab above and below the disc, are similar to the No-Heat case, both in their mean values and in their amplitude: the bubble heating does not alter the disc-halo cycle in the inner regions, consistent with the very similar levels of SFR in the two simulations.

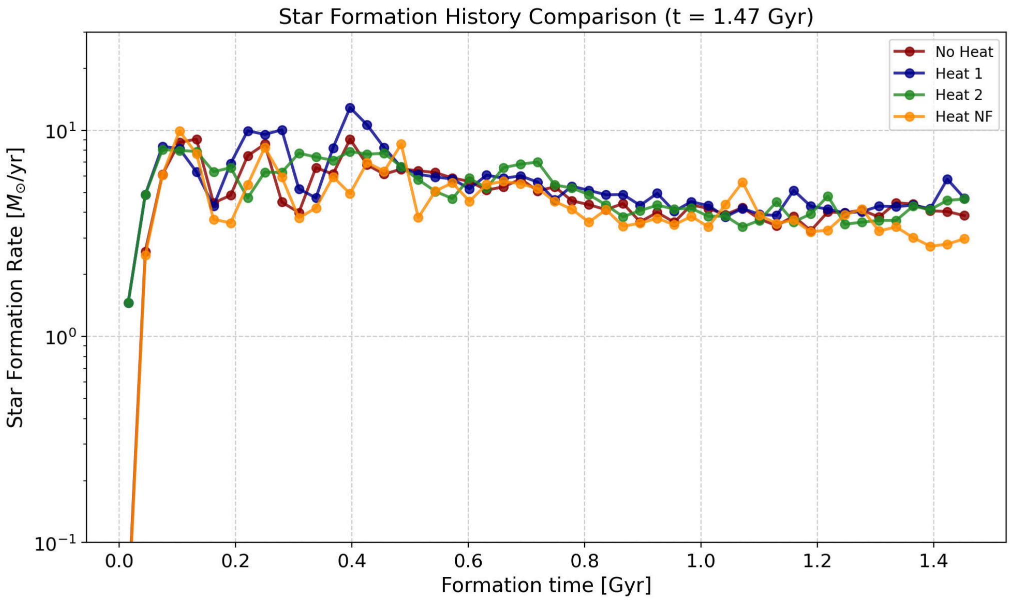

Figure 6.12: SFR for the No-Heat, Heat 1, Heat 2 and NF models shown by red, blue, green and yellow curves, respectively. All four histories follow a comparable trend, indicating that the SFR is not strongly affected by the heating models; yet the NF run settles onto a marginally lower plateau, consistent with its slightly reduced gas content in the inner regions of the corona.

For the coronal shells we restrict the measurement to the *hot gas* ($T > 5 \times 10^5$ K): the mass evolution and the density profiles showed that, in NF, the hot phase is the one that behaves differently with respect to the other two phases, so the hot-gas flow will be investigated further to understand the inner depletion and outer excess mentioned in Sections 6.6 and 6.9. Tracking the total flow would dilute this signal with the warm and cold phases, which behave as in the No-Heat model. We choose to compare NF only with the No-Heat run because the distributed-heating models have flows really similar to the unheated one (Chapter 5) and adding them would make figures 6.14 and 6.15 more difficult to read.

In these two figures, the hot gas flows in the corona, by contrast with the gas in the region close to the disc, differ between the two models compared and this difference accounts for the distance-dependent hot mass of Section 6.6. In the net rate the two are similar: both are dominated by the inward cooling flow and remain net-inflowing for most of the run, with comparable averages ($\langle \dot{M}_{\rm net} \rangle \simeq 1.7$ and $0.5\, M_\odot\,{\rm yr}^{-1}$ for No-Heat and NF at 100 kpc) except for two opposite behaviours at

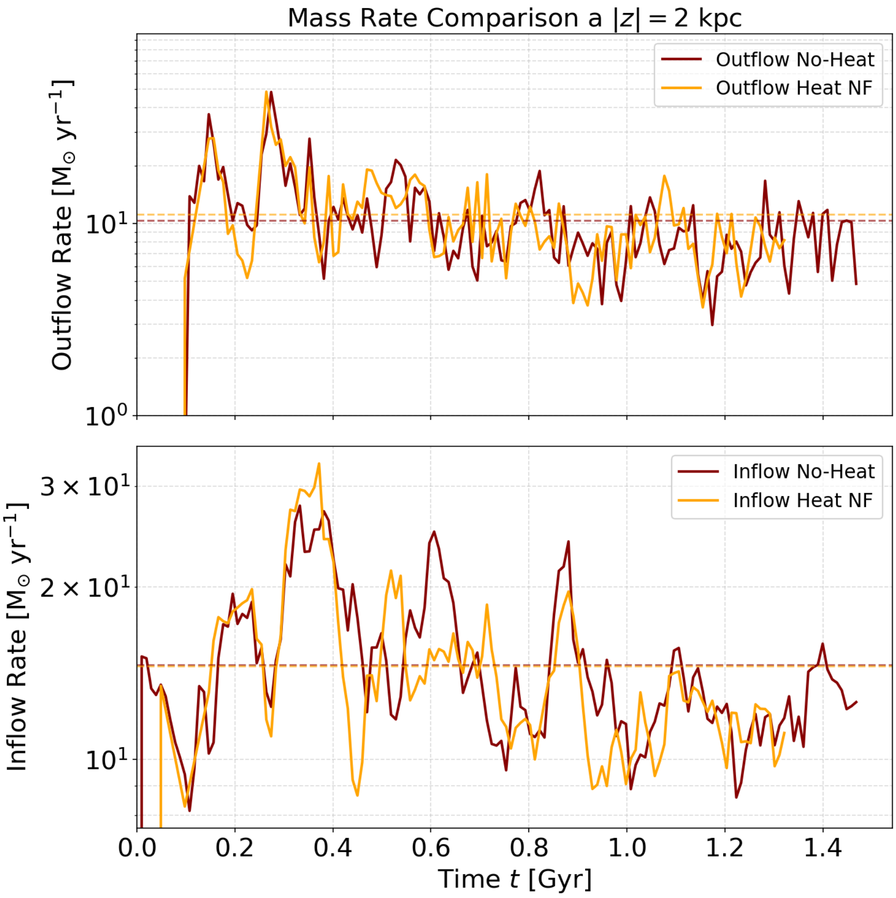

Figure 6.13: Comparison between NF and No Heat inflows (lower panel) / outflows (upper panel). The similar trends of the two models considered probe that at $|z| = 2$ kpc distance from the disc the gas recycling is not changed remarkably by adding the NF heating channel.

$t \sim 0.6$ and $t \sim 1.2$ Gyr, indicating that one model, with respect to the other, was dominated by inflows rather than outflows at that specific time. The most noticeable difference is actually in the *outflowing* component. At $r = 50$ kpc (Figure 6.14), the NF hot outflow is only marginally above the baseline, but at $r = 100$ kpc (Figure 6.15), it lies systematically above it from $t \simeq 0.6$ Gyr to the end of the run, by up to a factor of ~ 4.5 (median 1.6 vs $0.35\, M_\odot\, \mathrm{yr}^{-1}$). This outward flux is the dynamical counterpart of the outer hot-mass excess presented in Section 6.6, and it precedes the peak of that larger reservoir ($t \simeq 1.1$ Gyr). Both models show an early outward episode at 100 kpc (No-Heat at $t \simeq 0.42$, NF at $t \simeq 0.65$ Gyr) tied to the initial rearrangement of the corona; what distinguishes NF is that its outflow stays elevated afterwards, whereas No-Heat stays at a lower value.

These measurements answer the two questions posed at the start. First, at $|z| = 2$ kpc from the disc the inflows and outflows are close to the No-Heat case, matching the small SFR difference: the gas supply to the disc is only marginally affected. Second, the strong hot outflow at $r = 100$ kpc explains the outer hot mass excess of Section 6.6. This hot mass is moved to the outer corona because the action of the heating has pushed hot gas outward from the inner regions.

6.11 Summary

In this chapter we examined the NF bubble model, in which the corona is heated by discrete and localised energy injections whose budget is set from the halo X-ray luminosity. Our aim was to establish whether or not this heating had some effects on the corona and its cooling flow, on the disc structure and on the SFR.

We approached the question through several complementary diagnostics: the mass evolution and the coronal density profiles were used to track where the hot gas resides as a function of time; the cumulative mass profiles tested whether the warm+hot gas was redistributed outwards from the inner to the outer corona; and the gas flows quantified the rate of this displacement. In parallel, we investigated the disc morphology and the SFR to check whether the heating leaves any imprint on the disc structure, and we used the gas phase diagrams to follow the changes in both temperature and density of the coronal gas, finding the only difference at

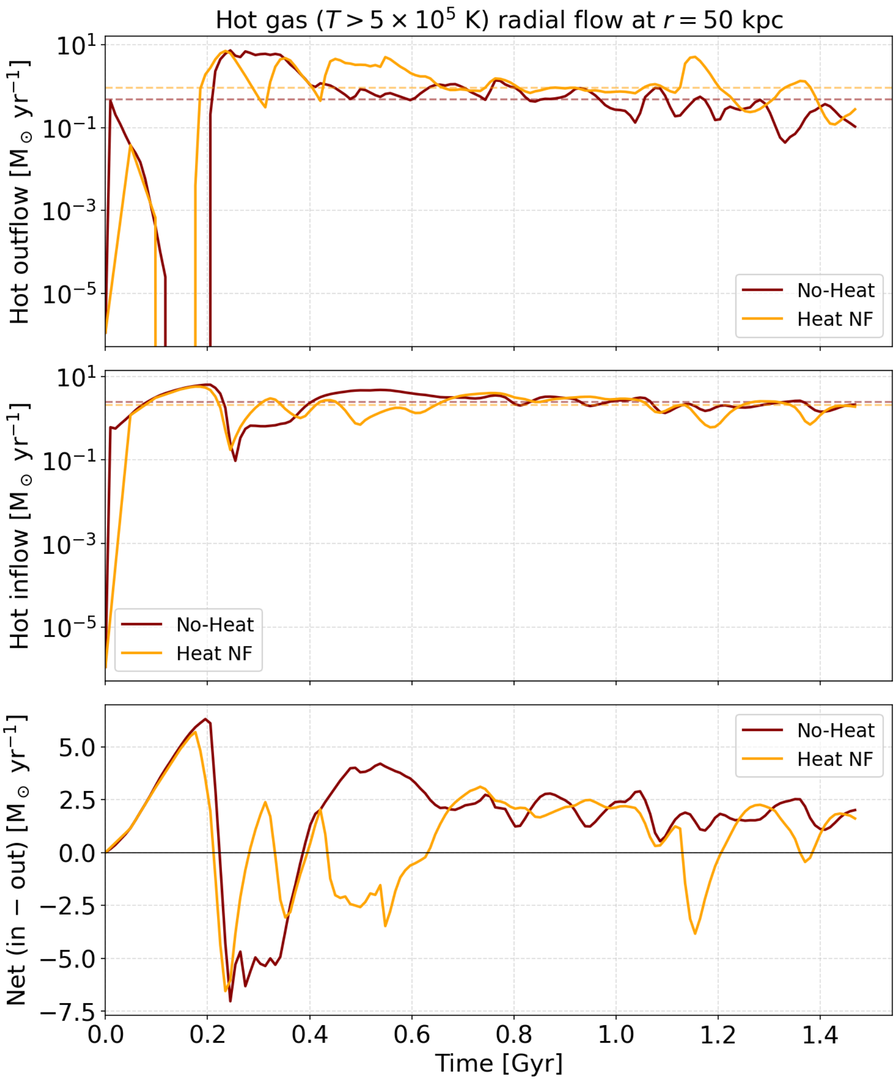

Figure 6.14: Comparison between NF and No-Heat models inflows/outflows at 50 kpc just of the hot gas ($T > 5 \times 10^5$K). In the upper panel, is displayed the outflows comparison where the dark orange curve is the NF model that stays over the dark red curve (No-Heat) for the whole time after $t = 0.4$ Gyr probing it shows a larger flow of gas outwards. The inflow rate, in the middle panel, shows instead very similar trends. The combination of the first two panels results into a net rate flow exploiting better some differences in the trends of the two models, in particular in correspondence of the NF model-outflow peaks. This is showed in the bottom panel.

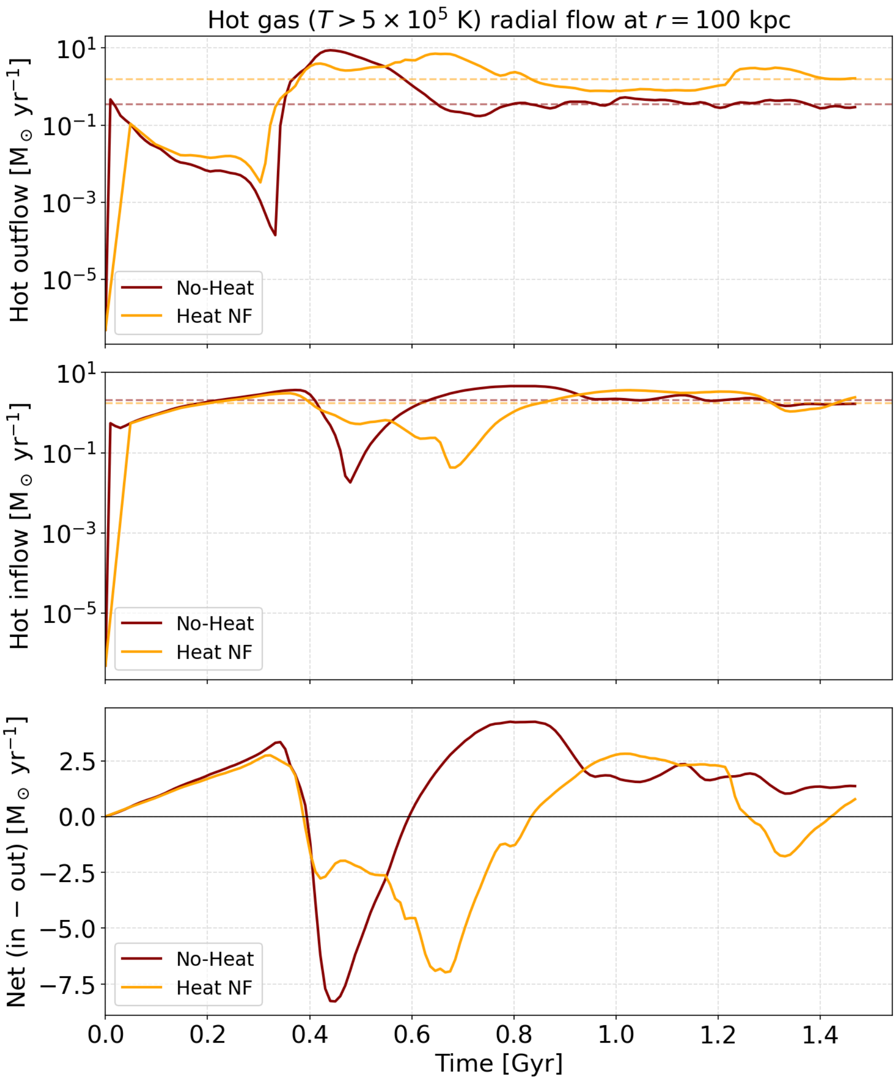

Figure 6.15: Comparison between NF and No-Heat models inflows/outflows at 100 kpc for the hot gas ($T > 5 \times 10^5$K). In the upper panel, is displayed the outflows comparison where the dark orange curve is the NF model that stays over the dark red curve (No-Heat) for the whole time after $t = 0.6$ Gyr probing it has a larger flow of gas outwards. The inflow rate, in the middle panel, shows instead very similar trends, but also an out of phase pronounced dip. The combination of the first two panels results into a net flow rate (bottom panel) showing some differences in the trends of the two models, in particular in correspondence of the NF model peaks or wells.

t = 1.2 Gyr in the lack of a tail from the high temperature - low density region towards the more dense regions of the disc.

The two main results are: the SFR and the gas supply to the disc mostly unaffected by the heating of the corona, and the displacement of an amount of the hot gas from the inner to the outer corona. The SFR remains around $3 - 5\ M_{\odot}\ yr^{-1}$, similar but slightly lower than No-Heat model, as well as the net inflow of hot gas at 50 and 100 kpc. The heating does not essentially change how much gas the corona delivers to the disc, leaving the structure of the disc without significant changes with respect to the unheated model. By contrast, the cumulative-mass profiles show the important imprint of depletion of the NF model on the hot gas of the inner corona (within $\sim$ 100 kpc), and the excess towards the outer region as well. The radial flows confirm and quantify this, with the outward flux of hot gas at 100 kpc a factor $\sim$ 4.5 higher than in the No-Heat run.

Taken together, these results show that the NF model acts without altering the disc or its star formation in a significant way, rather the heating has an effect on the amount of hot gas present in the inner / outer corona.

Chapter 7

Conclusions

Star-forming galaxies like the Milky Way are surrounded by a hot gaseous corona at the virial temperature $T_{vir} \sim 10^6$ K, whose radiative cooling and accretion onto the disc provide a fundamental supply of fuel for star formation (Barbani et al., 2023, 2025). In this scenario, the corona does not represent a passive reservoir but an active regulator of the gas recycling: it cools, it accretes, and it replenishes the cold gas that the disc consumes. A natural question follows, and it is the one this thesis sets out to address: if the corona is subject to a heating source that counteracts its cooling (other than the ones integrated into SMUGGLE), does this alter the evolution of the galaxy, and the amount of gas the corona provides to the disc?

To answer this question, we compared the **No-Heat** model (Chapter 4) with three models that differ in how the corona is heated. The **Heat 1** and **Heat 2** models (Chapter 5) implement a cooling-compensation prescription, in which energy is injected to offset the radiative losses of the coronal gas, differing in the timescale over which this budget is distributed. The **NF Bubbles** model (Chapter 6) adopts a simplified version of the works developed by Nulsen & Fabian (2000) and Grand et al. (2017). Our version is a discrete energy injection prescription based on the placement of bubbles in the corona whose budget is derived from the X-ray luminosity of the coronal gas via the $L - T$ relation expressed in Eq. (6.4). Since this thesis aims to answer whether the action of these three heating models on the corona alters the evolution of the galaxy, we analysed several of its features through some diagnostics, such as the gas mass evolution per phase, the phase diagrams and the inflow/outflow rates. Our main results are summarized in the following.

7.1 Summary of the results

The central result of this work concerns *where*, within the hot CGM, the heating prescriptions leave their mark and where they do not. The three heating models have different effects on the hot gas of the corona, but they have in common that none of them alters significantly the cooling flow that ultimately reaches the star-forming disc. As a consequence the SFR is only marginally different across the four models: after an initial peak, all settle onto a shared plateau of $\sim 3-5\,M_{\odot}\,\mathrm{yr}^{-1}$, (closer to $\sim 3\,M_{\odot}\,\mathrm{yr}^{-1}$ for NF model). Star formation is set mainly by the self-regulation of the disc through stellar feedback, as built into the SMUGGLE model, and is not mainly driven by how the hot CGM is heated, still the way each prescription acts on the corona differs markedly, so tracing those differences is what we aim to display next with a summary of the results for each heating model.

Heat 1: local-continuous injection. Heat 1 is closely similar to the No-Heat baseline at every radius examined. We can infer that the injected amount of heat is not capable of replenishing the energy lost by cooling and is therefore radiated away before it can do any work (see also Chapter 5). The only measurable trace is that the hot gas within 100 kpc declines slightly more slowly than in No-Heat; beyond this, Heat 1 offers no further indication of any influence on the coronal gas, and leaves no imprint on the gas flows at either 10 or 50 kpc.

Heat 2: focus on the inner corona. Sharing the same energy budget as Heat 1 but accumulating it over the global timestep before the release, Heat 2 does affect the corona, but its action is concentrated in the *innermost* regions of the hot CGM. As in Heat 1, the hot gas within 100 kpc does decline more slowly than in No-Heat, but Heat 2 leaves also an additional signature at $\sim$ 10 kpc: both the inflow and the outflow rates exceed those of No-Heat by $\sim 2-3\,M_{\odot}\,\mathrm{yr}^{-1}$, while at $\sim$ 50 kpc and at 2 kpc from the disc, the flows revert to the unheated-like behaviour. The heat thus reaches the inner corona and perturbs its kinematics. The injected energy enhances the local circulation, lifting more gas outward while a comparable amount of it falls back, keeping the *net* flow and the supply to the disc preserved. This is consistent with the fact that, within a ~ 100 kpc distance from the disc, the cold gas keeps increasing while the warm gas decreases (see Figure 5.1): the two trends mirror each other, as warm gas cools and feeds the cold phase, so the cooling flow remains only marginally different from the No-Heat and the SFR similar.

The behaviour of Heat 2 at ~ 10 kpc can be understood in terms of the energy budget. Heat 2 injects more energy than Heat 1, and this is what allows it to leave a measurable imprint on the inner corona. The energy it deposits, however, is set to compensate the local radiative cooling rather than to exceed it: by construction it offsets the losses, but does not provide the surplus that would be needed to expel the gas outwards. As a result, the heated gas is lifted outward but remains bound, and falls back rather than being expelled.

As a conclusion, we deduce that the level of star formation is more dependent on the stellar feedback already present in SMUGGLE than on the Heat 2 energy injection reaching the inner corona regions.

NF Bubbles: action on the outer corona. The discrete bubble prescription acts in a different manner, and on a different region, with respect to the distributed-heating models. Its most evident effect is the outward transport of hot gas, depleting the intermediate region between ~ 10 and ~ 100 kpc, and carrying $\sim 0.2 \times 10^{10}\, M_\odot$ more than the other models beyond 100 kpc (see Figures 6.2 and 6.15). This is seen clearly also in the cumulative warm+hot mass profile, where the NF curve is placed below the others in the intermediate region of the corona, while it reaches again the other models at distances $r \gtrsim 240$ kpc, indicating more mass in the outer corona. The gas mass distribution of the corona is thus rearranged: depleted at intermediate radii and enhanced outward, yet it is worth mentioning that the cooling flow and the amount of cold gas exchanged with disc is instead only marginally different from the No-Heat, resulting in relatively small differences in SFR and inflows/outflows between the NF model and No-Heat.

A unifying picture. Taken together, the four models trace a consistent picture. Heat 1 leaves no remarkable imprint on the corona, Heat 2 acts on its innermost regions, and NF Bubbles acts on the coronal gas at large radii beyond ~ 100 kpc: three distinct effects on the hot CGM, none of which change remarkably the rate at which the disc forms stars. This conclusion remains the same even when we increase the NF injection threshold by a factor of ten (the NF-new run, Section 7.3), releasing the same budget in fewer but more energetic events, the SFR remains consistent with the other models within the fluctuations. The lack of a significant effect on star formation is therefore not a consequence of an under-powered budget, since also an increase of an order of magnitude seems to be not enough to cause a remarkable change in the SFR (see Section 7.3).

This result may seem at odds with Barbani et al. (2023), who established that the corona is the primary reservoir sustaining star formation in Milky Way-like galaxies. Yet this contradiction is only apparent. In all four models the corona remains present and continues to accrete: none of the heating prescriptions offsets the cooling onto the disc, and the cold gas that fuels star formation keeps growing over time in every model. The supply of cold gas to the disc is never interrupted, which is why the SFR is similar to the No-Heat model. Our result can be seen as a complementary to the ones presented in Barbani et al. (2023): they showed that the presence of the corona is what sustains star formation over several Gyr, whereas we find that, once the corona is in place, heating its hot gas has only a marginal influence on the level of star formation of the disc.

7.2 Physical interpretation and outlook

We can now draw a general picture of what this thesis contributes to the study of the interaction between star-forming galaxies and their hot CGM. The central result is that, at fixed budget, *where* and *how* the energy is injected into the corona matters more than *how much*. Heat 1 and Heat 2 share the same budget but differ in injection timescale, yet only Heat 2 measurably affects the inner corona; the NF model, injecting through discrete bubbles at large radii, transports hot gas outward beyond ~ 100 kpc and drives a warm-to-hot conversion, thereby setting where the hot gas is redistributed. In none of these cases, however, does the effect of the heating reach the cold gas near the disc: heating the outer corona simply moves hot gas further out and converts warm gas into hot, without significantly influencing the cold reservoir that feeds star formation, so the cooling flow onto the disc is never quenched. Whichever way the corona is heated, the SFR stays close to the unheated case, governed by the self-regulation of the disc through stellar feedback.

This may seem to contradict the observed role of AGN feedback in quenching star formation in real galaxies, but the tension possibly comes from some of the simplifications of our setup. Our heating budget is fixed and does not react to the state of the gas, there is no black hole that accretes self-consistently, but also the galaxy evolves in isolation, outside a cosmological context. The quenching seen in real Milky Way-like galaxies may also rely on these missing mechanisms. Since heating the outer corona leaves the cold reservoir without remarkable changes, an effective feedback channel would instead have to act on the cold gas itself, or within the inner regions ($\lesssim$ 10 kpc), rather than on the diffuse outer corona. Testing whether feedback deposited in these regions would actually regulate star formation is left to future work.

7.3 Future perspectives

The simplifications discussed above, such as a fixed heating budget, the absence of a self-consistently accreting black hole, and the isolated setup, naturally point to how the present work could be extended. We outline two directions, both coming from the NF model we implemented, that would bring the current setup closer to a more advanced and physically realistic one.

The first line of development concerns the energy *injection threshold* displayed in Section 6.3: in our implementation the energy accumulates in a reservoir and is released only once a fixed threshold is reached. The value of this threshold therefore does not control the total budget but the frequency of release: a low threshold producing frequent, low-energy bubbles, a high threshold producing rare but more energetic ones. Varying this threshold would test, for example, whether the outward redistribution of warm+hot gas depends on how the same energy is injected in time.

There is the possibility that many "gentle" injections and a few more energetic ones leave the same imprint on the corona, or this difference can cause a large difference in the hot gas distribution. In Figure 7.1 is displayed a first result of the gas mass distribution obtained with an injection threshold ten times larger than in the NF model, presented in Chapter 6. Also in this case the mass of cold gas keeps increasing in time, indicating that the cooling flow onto the disc still evolves at a similar pace to the No-Heat model; consistently, the SFR remains within the scatter of the other models (see Figure 7.2). A full comparison with the remaining diagnostics used in this work would be valuable. However, this analysis is still ongoing and is therefore left for future work.

The second direction is to let the energy budget of the feedback emerge from the simulation rather than being imposed. We fixed the budget to the full X-ray luminosity ($R = 1$) as the simplest choice, treating it as an external input. A natural step forward is to couple the bubble injection to a black hole that grows and accretes self-consistently, so that the radio-mode energy is set by the accretion rate onto the black hole rather than fixed by an *ad hoc* prescription. The heating would also vary in time, as in the implementation used in this work, but would follow the accretion history of the black hole instead of a fixed budget. This is the step that would most directly narrow the distance between the idealized, isolated setup used here and the fully self-consistent AGN feedback of cosmological simulations such as Auriga (Grand et al., 2017), where the black hole, its accretion, and the surrounding corona evolve together.

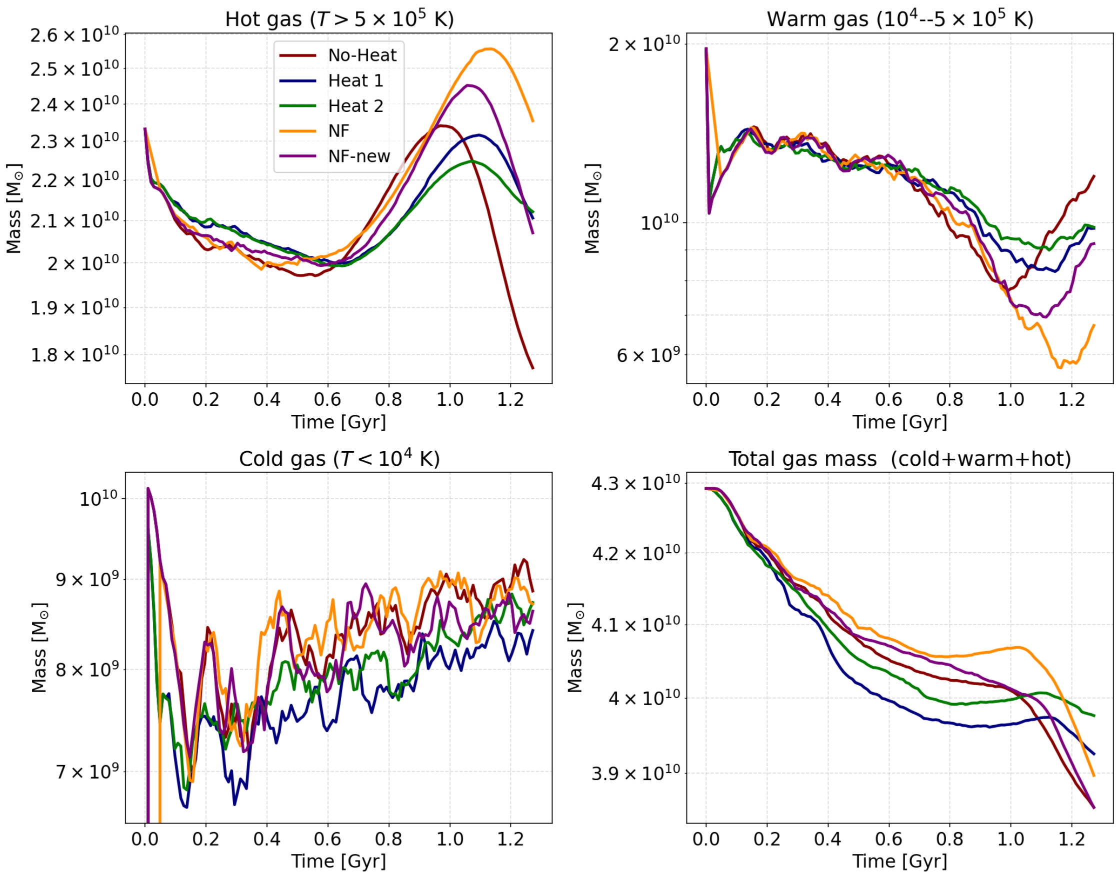

Figure 7.1: Evolution of the gas mass split into cold ($T < 10^4$ K), warm ($10^4 < T < 5 \times 10^5$ K) and hot ($T > 5 \times 10^5$ K) phases, computed within a sphere of radius $r = 240$ kpc, for the No-Heat, Heat 1, Heat 2, NF and NF-new models. Each panel shows a phase of the gas, while each curve is representative of each model. It is noticeable how the cold phase is increasing through time, but after $t \sim 1.0$ Gyr the hot phase shows a pronounced bump too.

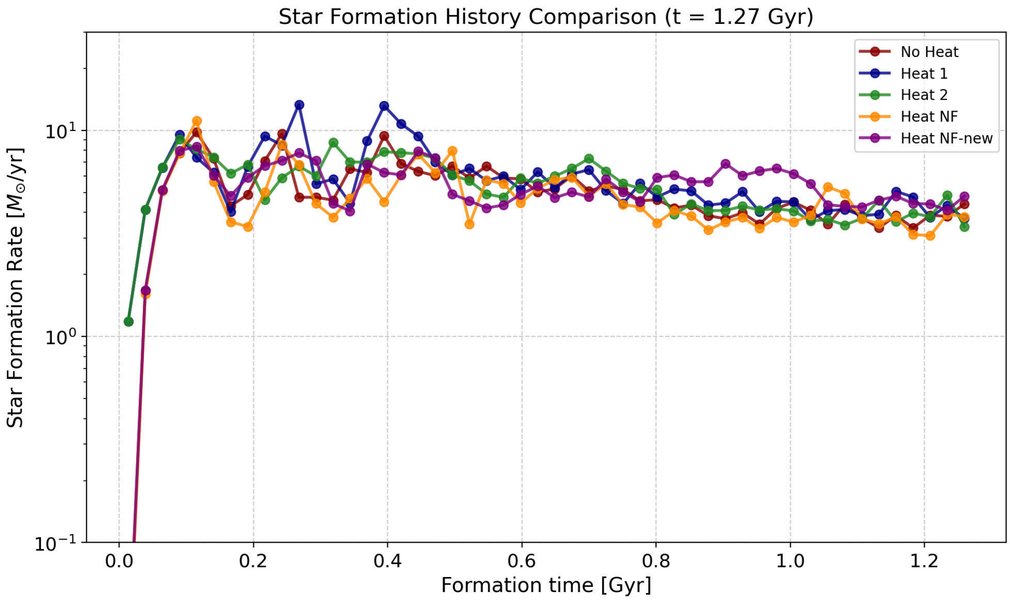

Figure 7.2: SFR of the five models compared: No-Heat, Heat 1, Heat 2, NF and NF-new, within a $t \sim 1.3$ Gyr time. Every curve represents one model, as labelled in the legend. The NF-new model is comparable to the other models studied in this thesis.

Bibliography

Afruni A., Fraternali F., Pezzulli G., 2022, Astronomy and Astrophysics, 662, A8

Agertz O., et al., 2007, Monthly Notices of the Royal Astronomical Society, 380, 963

Agertz O., Teyssier R., Moore B., 2011, Monthly Notices of the Royal Astronomical Society, 410, 1391

Agertz O., Kravtsov A. V., Leitner S. N., Gnedin N. Y., 2013, The Astrophysical Journal, 770, 25

Agertz O., et al., 2021, Monthly Notices of the Royal Astronomical Society, 503, 5826

Anderson M. E., Bregman J. N., 2010, The Astrophysical Journal, 714, 320

Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, Annual Review of Astronomy and Astrophysics, 47, 481

Barbani F., Pascale R., Marinacci F., Sales L. V., Vogelsberger M., Torrey P., Li H., 2023, Monthly Notices of the Royal Astronomical Society, 524, 4091

Barbani F., Pascale R., Marinacci F., Torrey P., Sales L. V., Li H., Vogelsberger M., 2025, Astronomy & Astrophysics, 697, A121

Barnes J., Hut P., 1986, Nature, 324, 446

Binney J., Tremaine S., 2008, Galactic Dynamics: Second Edition. Princeton University Press

Bland-Hawthorn J., Gerhard O., 2016, Annual Review of Astronomy and Astrophysics, 54, 529

Bode P., Ostriker J. P., 2003, The Astrophysical Journal Supplement Series, 145, 1

Bogdán Á., Bourdin H., Forman W. R., Kraft R. P., Vogelsberger M., Hernquist L., Springel V., 2017, The Astrophysical Journal, 850, 98

Bondi H., 1952, Monthly Notices of the Royal Astronomical Society, 112, 195

Bondi H., Hoyle F., 1944, Monthly Notices of the Royal Astronomical Society, 104, 273

Box G. E. P., Muller M. E., 1958, The Annals of Mathematical Statistics, 29, 610

Cavaliere A., Fusco-Femiano R., 1978, Astronomy and Astrophysics, 70, 677

Chabrier G., 2003, Publications of the Astronomical Society of the Pacific, 115, 763

Crocker R. M., Bicknell G. V., Taylor A. M., Carretti E., 2015, The Astrophysical Journal, 808, 107

Das S., Mathur S., Gupta A., Nicastro F., Krongold Y., 2020, The Astrophysical Journal, 891, 133

Edgar R., 2004, New Astronomy Reviews, 48, 843

Fabian A. C., 1994, Annual Review of Astronomy and Astrophysics, 32, 277

Fabian A. C., 2012, Annual Review of Astronomy and Astrophysics, 50, 455

Faucher-Giguère C.-A., Oh S. P., 2023, Annual Review of Astronomy and Astrophysics, 61, 131

Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, The Astrophysical Journal, 703, 1416

Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, Publications of the Astronomical Society of the Pacific, 110, 761

Fraternali F., 2017, Gas Accretion via Condensation and Fountains. Springer, pp 323–353, doi:10.1007/978-3-319-52512-9˙14

Freeman K. C., 1970, The Astrophysical Journal, 160, 811

GRAVITY Collaboration Abuter R., Aimar N., et al., 2022, Astronomy & Astrophysics, 657, L12

Garma-Oehmichen L., Martinez-Medina L., Hernández-Toledo H., Puerari I., 2021, Monthly Notices of the Royal Astronomical Society, 502, 4708

Gatto A., Fraternali F., Read J. I., Marinacci F., Lux H., Walch S., 2013, Monthly Notices of the Royal Astronomical Society, 433, 2749

Genzel R., Eisenhauer F., Gillessen S., 2010, Reviews of Modern Physics, 82, 3121

Ghez A. M., et al., 2008, The Astrophysical Journal, 689, 1044

Gillessen S., Eisenhauer F., Trippe S., Alexander T., Genzel R., Martins F., Ott T., 2009, The Astrophysical Journal, 692, 1075

Gingold R. A., Monaghan J. J., 1977, Monthly Notices of the Royal Astronomical Society, 181, 375

Godunov S. K., 1959, Math. Sbornik, 47, 271

Grand R. J. J., et al., 2017, Monthly Notices of the Royal Astronomical Society, p. stx071

Grcevich J., Putman M. E., 2009, The Astrophysical Journal, 696, 385

Guo F., Oh S. P., 2008, Monthly Notices of the Royal Astronomical Society, 384, 251

Haywood M., Di Matteo P., Lehnert M. D., Katz D., Gómez A., 2013, Astronomy and Astrophysics, 560, A109

Hernquist L., 1990, The Astrophysical Journal, 356, 359

Hernquist L., 1993, The Astrophysical Journal Supplement Series, 86, 389

Hockney R. W., Eastwood J. W., 1981, Computer Simulation Using Particles. https://ui.adsabs.harvard.edu/abs/1981csup.book.....H

Hopkins P. F., et al., 2018, Monthly Notices of the Royal Astronomical Society, 480, 800

Karakas A. I., 2010, Monthly Notices of the Royal Astronomical Society, 403, 1413

Kennicutt Jr. R. C., 1998, Annual Review of Astronomy and Astrophysics, 36, 189

Krumholz M. R., Tan J. C., 2007, The Astrophysical Journal, 654, 304

Li M., Bryan G. L., Ostriker J. P., 2017, The Astrophysical Journal, 841, 101

Licquia T. C., Newman J. A., 2015, The Astrophysical Journal, 806, 96

Lucy L. B., 1977, The Astronomical Journal, 82, 1013

Marinacci F., Fraternali F., Nipoti C., Binney J., Ciotti L., Londrillo P., 2011, Monthly Notices of the Royal Astronomical Society, 415, 1534

Marinacci F., Sales L. V., Vogelsberger M., Torrey P., Springel V., 2019, Monthly Notices of the Royal Astronomical Society, 489, 4233

Martizzi D., Faucher-Giguère C.-A., Quataert E., 2015, Monthly Notices of the Royal Astronomical Society, 450, 504

Martizzi D., Fielding D., Faucher-Giguère C.-A., Quataert E., 2016, Monthly Notices of the Royal Astronomical Society, 459, 2311

McNamara B. R., Nulsen P. E. J., 2007, Annual Review of Astronomy and Astrophysics, 45, 117

Miller M. J., Bregman J. N., 2013, The Astrophysical Journal, 770, 118

Naab T., Ostriker J. P., 2017, Annual Review of Astronomy and Astrophysics, 55, 59

Navarro J. F., Frenk C. S., White S. D. M., 1997, The Astrophysical Journal, 490, 493

Nulsen P. E. J., Fabian A. C., 2000, Monthly Notices of the Royal Astronomical Society, 311, 346

Pakmor R., Springel V., Bauer A., Mocz P., Muñoz D. J., Ohlmann S. T., Schaal K., Zhu C., 2016, Monthly Notices of the Royal Astronomical Society, 455, 1134

Peterson J. R., Fabian A. C., 2006, Physics Reports, 427, 1

Portinari L., Salucci P., 2010, Astronomy & Astrophysics, 521, A82

Portinari L., Chiosi C., Bressan A., 1998, Astronomy and Astrophysics, 334, 505

Pratt G. W., Croston J. H., Arnaud M., Böhringer H., 2009, Astronomy and Astrophysics, 498, 361

Putman M. E., Zheng Y., Price-Whelan A. M., Grcevich J., Johnson A., 2021, The Astrophysical Journal, 913, 53

Rahmati A., Pawlik A. H., Raicevic M., Schaye J., 2013, Monthly Notices of the Royal Astronomical Society, 430, 2427

Salem M., Besla G., Bryan G. L., Putman M. E., van der Marel R. P., Tonnesen S., 2015, The Astrophysical Journal, 815, 77

Salpeter E. E., 1955, The Astrophysical Journal, 121, 161

Schaye J., Chaikin E., Schaller M., Ploeckinger S., et al., 2025, Monthly Notices of the Royal Astronomical Society

Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, Monthly Notices of the Royal Astronomical Society, 380, 877

Snaith O., Haywood M., Di Matteo P., Lehnert M. D., Combes F., Katz D., Gómez A., 2015, Astronomy and Astrophysics, 578, A87

Somerville R. S., Davé R., 2015, Annual Review of Astronomy and Astrophysics, 53, 51

Springel V., 2000, Monthly Notices of the Royal Astronomical Society, 312, 859

Springel V., 2005, Monthly Notices of the Royal Astronomical Society, 364, 1105

Springel V., 2010a, Annual Review of Astronomy and Astrophysics, 48, 391

Springel V., 2010b, Monthly Notices of the Royal Astronomical Society, 401, 791

Springel V., Hernquist L., 2003, Monthly Notices of the Royal Astronomical Society, 339, 289

Springel V., Di Matteo T., Hernquist L., 2005, Monthly Notices of the Royal Astronomical Society, 361, 776

Su M., Slatyer T. R., Finkbeiner D. P., 2010a, The Astrophysical Journal, 724, 1044

Su M., Slatyer T. R., Finkbeiner D. P., 2010b, The Astrophysical Journal, 724, 1044

Su K.-Y., Hopkins P. F., Hayward C. C., et al., 2019, Monthly Notices of the Royal Astronomical Society, 487, 4393

Sutherland R. S., Dopita M. A., 1993, The Astrophysical Journal Supplement Series, 88, 253

The Fermi-LAT Collaboration 2014, The Astrophysical Journal, 793, 64

Thielemann F.-K., et al., 2003, Nuclear Physics A, 718, 139

Toro E. F., 2013, Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction. Springer Science & Business Media

Toro E. F., Spruce M., Speares W., 1994, Shock Waves, 4, 25

Tosi M., 1988, Astronomy & Astrophysics, 197, 47

Travaglio C., Hillebrandt W., Reinecke M., Thielemann F.-K., 2004, Astronomy and Astrophysics, 425, 1029

Tumlinson J., Peeples M. S., Werk J. K., 2017, Annual Review of Astronomy and Astrophysics, 55, 389

Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, Monthly Notices of the Royal Astronomical Society, 436, 3031

Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, Nature Reviews Physics, 2, 42

Weinberger R., Springel V., Pakmor R., 2020, The Astrophysical Journal Supplement Series, 248, 32

White S. D. M., Rees M. J., 1978, Monthly Notices of the Royal Astronomical Society, 183, 341

Wiersma R. P. C., Schaye J., Smith B. D., 2009, Monthly Notices of the Royal Astronomical Society, 393, 99

Wolfire M. G., McKee C. F., Hollenbach D., Tielens A. G. G. M., 2003, The Astrophysical Journal, 587, 278

Woosley S. E., Weaver T. A., 1995, The Astrophysical Journal Supplement Series, 101, 181

Zubovas K., King A. R., Nayakshin S., 2012, Monthly Notices of the Royal Astronomical Society, 421, 3464