

Environmental variation of the low-mass IMF

Tabassum S. Tanvir,¹★ Mark R. Krumholz^{1,2} and Christoph Federrath^{1,2}

¹*Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia*

²*Australian Research Council Centre of Excellence in All Sky Astrophysics (ASTRO3D), Canberra, ACT 2611, Australia*

Accepted 2022 September 10. Received 2022 September 7; in original form 2022 June 10

ABSTRACT

We use a series of magnetohydrodynamic simulations including both radiative and protostellar outflow feedback to study environmental variation of the initial mass function (IMF). The simulations represent a carefully-controlled experiment whereby we keep all dimensionless parameters of the flow constant except for those related to feedback. We show that radiation feedback suppresses the formation of lower mass objects more effectively as the surface density increases, but this only partially compensates for the decreasing Jeans mass in denser environments. Similarly, we find that protostellar outflows are more effective at suppressing the formation of massive stars in higher surface density environments. The combined effect of these two trends is towards an IMF with a lower characteristic mass and a narrower overall mass range in high surface density environments. We discuss the implications for these findings for the interpretation of observational evidence of IMF variation in early type galaxies.

Key words: magnetic fields – radiative transfer – stars: formation – stars: luminosity function, mass function – stars: protostars.

1 INTRODUCTION

The initial mass function (IMF) describes the mass distribution of stars at formation. It plays a critical role in a wide range of fields, from galaxy formation and evolution to nuclear astrophysics to planetary science. For this reason, measuring the IMF and understanding its origin has been one of the central issues in the study of the star formation process. Salpeter (1955) first attempted to measure the IMF by using Solar neighbourhood stars. He adopted a single slope power-law distribution of the form $dN \propto M^{-\alpha} dM$, and found a best fitting slope $\alpha = 2.35$. Six decades later, this distribution is still considered the standard for stars $> 1 M_{\odot}$. However, as this function approaches zero, it diverges, indicating that there must be a turnover in the IMF at lower masses. Miller & Scalo (1979) introduced the idea of a lognormal distribution between 0.1 and $\simeq 30 M_{\odot}$, motivated by a clear flattening in the observed Solar neighbourhood mass distribution below $1 M_{\odot}$. Today, the two most widely used forms of the IMF are the broken power-law distribution introduced by Kroupa (2001), and the combination of lognormal (at low mass) and power law (at high mass) distribution by Chabrier (2005) – see Offner et al. (2014) for a recent review. Both of these forms feature a broad plateau at masses from $\approx 0.1 - 1 M_{\odot}$, and then a power-law decline at higher masses.

From a theoretical point of view the question is what physical processes are responsible for setting the shape of the IMF. We know from observations that the IMF is close to universal in the Milky Way and nearby local galaxies (e.g. Lee et al. 2020), which suggests a simple and universal physical mechanism. Turbulence is a natural choice, and there have been many theoretical models of the IMF based on the properties of turbulence. For example, Padoan,

Nordlund & Jones (1997) and Padoan & Nordlund (2002) propose a model in which interacting shocks produce dense filaments and sheets where star formation occurs; the mass function is determined by the distribution of post-shock properties, which is a function of turbulent power spectrum. Hennebelle & Chabrier (2008, 2009) propose a model based on the Press & Schechter (1974) formalism, in which the IMF is treated as a barrier-crossing problem: turbulence generates a spectrum of density fluctuations at different scales, and when one of these fluctuations crosses the barrier to becoming self-gravitating, it collapses to form a star. (Hopkins 2012, 2013) develops a model based on excursion set theory, for which the basic physical setup – fluctuations induced by turbulence leading to barrier crossing – is quite similar to that of Hennebelle & Chabrier, but where some of the detailed assumptions, for example about the scale-dependence of the density statistics, are different.

A common feature of these models is that they are based on isothermal gas under the influence of turbulence. While such a theory can very naturally explain the power-law tail of the IMF, and there is good numerical evidence that turbulence does indeed determine this feature (Nam, Federrath & Krumholz 2021), isothermal turbulence is scale-free over the inertial range (Kolmogorov 1941; Burgers 1948), with characteristic lengths only at the driving, dissipation, and sonic scales (Federrath et al. 2021). It therefore cannot explain why the IMF turns over at a particular mass that is independent of the total or Jeans mass of the parent star-forming cloud (Krumholz 2014). Again, numerical evidence supports this conclusion. Guszejnov, Krumholz & Hopkins (2016), Guszejnov et al. (2018) show that fragmentation of non-magnetized isothermal gas produces a pure power-law mass distribution, while Guszejnov et al. (2020) show that magnetized isothermal turbulence leads to a mass function with a peak that is simply proportional to the initial sonic mass, and therefore not plausibly related to the nearly-universal peak observed in the local IMF. These findings strongly suggest that any viable

★ E-mail: tabassum.tanvir@anu.edu.au

theory for the non-power-law low-mass IMF requires a deviation from the assumption of isothermal gas (however, see Haugbølle, Padoan & Nordlund 2018 for a contrary view).

The question then becomes what physical process induces the deviation from isothermality. Deviations can occur due to gas becoming optically-thick to its own cooling radiation at a density of $n \sim 10^{10} \text{ cm}^{-3}$ (Low & Lynden-Bell 1976; Rees 1976; Masunaga, Miyama & Inutsuka 1998) or the onset of dust-gas coupling at a density $n \sim 10^5 \text{ cm}^{-3}$ (Spaans & Silk 2000; Jappsen et al. 2005; Larson 2005; Elmegreen, Klessen & Wilson 2008). The former leads to a mass scale that is plausibly associated with the smallest possible brown dwarfs, $\sim 10^{-3} M_{\odot}$, but not plausibly associated with the observed IMF peak. The latter is more promising in terms of the mass scale, but Guszejnov et al. (2016) show that if one modifies hydrodynamic simulations simply by introducing an EOS stiffening at some density, the resulting IMF is extremely sensitive to its initial condition, which renders these models unable to provide the observed universal mass scale.

For this reason many authors have considered stellar radiation feedback as a potential mechanism to set a characteristic mass scale. In this picture, low mass young stars radiate due to accretion, raising the surrounding gas temperature and therefore the Jeans mass, strongly suppressing fragmentation (Krumholz 2006, 2011; Bate 2009, 2012; Offner et al. 2009; Krumholz, Klein & McKee 2012; Krumholz et al. 2016; Federrath, Krumholz & Hopkins 2017; Cunningham et al. 2018; Mathew & Federrath 2020, 2021). This effect suppresses the formation of low-mass objects by ensuring that the mass around them is unable to fragment further, and is instead available to accrete. Additional important physical processes that control the peak of the IMF include a combination of magnetic fields and protostellar jet feedback (e.g. Shu, Li & Allen 2004; Federrath et al. 2014; Mathew & Federrath 2021), and possibly tidal effects of the first Larson cores (e.g. Lee & Hennebelle 2018; Hennebelle, Lee & Chabrier 2019).

However, there have been limited efforts thus far to determine the implications of these models for variation in the location of the IMF peak with star-forming environment. Krumholz & McKee (2008) argue from analytic models, and Krumholz et al. (2010) and Myers et al. (2011) confirm with simulations, that higher surface densities make it easier to form massive stars on the tail of the IMF, because higher optical depths trap radiation more effectively, suppressing fragmentation and allowing massive cores to undergo a processes closer to monolithic collapse. However, neither of these studies address the peak of the IMF, as opposed to the high-mass tail. Sharda & Krumholz (2022) do study the IMF peak, and predict a lower characteristic mass in metal-rich regions at high gas pressure / surface density, but these analytic model have yet to be checked by simulations, and in any event they predict only the approximate location of the IMF peak, not the full functional form of the IMF.

The problem of environmental variation is becoming urgent, however, because tentative evidence has started to emerge that, in the most extreme star-forming environments, the location of the IMF peak does change slightly. The strongest evidence for this shift has emerged from massive, early-type galaxies – see Smith (2020) for a recent review. In these systems, spectroscopic (e.g. van Dokkum & Conroy 2010; Spiniello et al. 2012; La Barbera et al. 2013; Conroy, van Dokkum & Villaume 2017), dynamical (e.g. Cappellari et al. 2012; Newman et al. 2017; Oldham & Auger 2018), and gravitational lensing (Treu et al. 2010; Spiniello et al. 2015) all point to an IMF with a lower peak than is found in the Milky Way, though there remain some inconsistencies in the measurements as to exactly where and in which galaxies this shift in the IMF occurs

(Smith 2014). There are also claims of a shift of the peak to higher masses in low-mass galaxies (Geha et al. 2013; Gennaro et al. 2018), though in these systems the small numbers of observed stars has led to considerable debate regarding the statistical significance of the result (El-Badry, Weisz & Quataert 2017). Regardless, the fact that evidence is emerging for variations in the location of the IMF peak represents both a challenge for theory and an opportunity, since reproducing and explaining these observations offers a strong test of models.

This paper aims to examine the effect of environmental variation on the IMF, focusing on the interaction of the environment with the two feedback effects that are thought to be most important near the IMF peak: radiation feedback and protostellar outflows. In order to isolate environmental effects, we identify the key dimensionless parameters that govern a star-forming system, and vary them systematically in order to perform a clean experiment that isolates the interaction of feedback and environmental effects from all other physical processes. We do so in the context of simulations that form a full star cluster with a measurable IMF, rather than just a single, massive core as in the earlier experiments of Krumholz et al. (2010). In Section 2, we describe the numerical method and initial conditions of our simulations. In Section 3, we examine the results of the simulations. We discuss the implications of our findings for the IMF in general in Section 4 and we summarize in Section 5.

2 NUMERICAL METHODS AND INITIAL CONDITIONS

2.1 Numerical Methods

The numerical methods we use are identical to those of Myers et al. (2014) and Cunningham et al. (2018), and we refer readers there for a more detailed description; here we merely summarize the key points. We use the ORION2 adaptive mesh refinement (AMR) code (Li et al. 2021) to carry out our simulations. The code uses the scheme of Li et al. (2012) to solve the equations of ideal magnetohydrodynamics coupled to self-gravity (Truelove et al. 1998; Klein et al. 1999) and radiation transfer (Krumholz, Klein & McKee 2007) in the two-temperature, mixed-frame, grey, flux-limited diffusion approximation. We use the tabulated density- and temperature-dependent Rosseland- and Planck-mean opacities provided by Semenov et al. (2003).

The code uses sink particles (Krumholz, McKee & Klein 2004) to replace regions where protostars are forming and that are collapsing beyond our ability to resolve. Each of these particles runs an instance of the one-zone protostellar evolution model described by Offner et al. (2009), which prescribes the instantaneous properties (radius, luminosity, polytropic index, etc.) of the star that particle represents. The luminosity becomes a source term in our radiative transfer equations, and we also include feedback due to protostellar outflows via momentum sources around each sink particle. We use the outflow model described in Cunningham et al. (2011): when mass is accreted on to a sink particle, a fraction f_w of it is ejected back into the simulation in the form of an outflow, presumably from an unresolved inner disc. The outflow material is launched with a speed v_w . In our simulation, we use the same wind model parameters described in Hansen et al. (2012) and Cunningham et al. (2018): $f_w = 0.3$ and $v_w = \min(v_{\text{kep}}, 60 \text{ km s}^{-1})$, where v_{kep} is the Kepler speed at the surface of the protostar. These parameters are motivated by observations of the momentum budget of protostellar outflows (Richer et al. 2000).

Table 1. Simulation parameters, from left to right: run name, mass in the computational box, size of the computational box, mean density in the computational box, 3D velocity dispersion, initial magnetic field strength, mean free-fall time, number of cells per linear dimension on the base grid, maximum level of refinement, smallest cell size, turbulent crossing time, and surface density.

Name	$M_{\text{box}} (M_{\odot})$	L (pc)	ρ_0 (g cm^{-3})	σ_v (km s^{-1})	B_0 (mG)	t_{ff} (kyr)	N	\mathcal{L}_{max}	Δx (AU)	t_{cross} (Myr)	Σ (g cm^{-2})
L	2000	0.92	1.74×10^{-19}	2.4	0.36	160	512	2	46	0.4	0.5
M	1000	0.46	6.96×10^{-19}	2.4	0.73	80	512	2	23	0.2	1
H	500	0.23	2.784×10^{-18}	2.4	1.45	40	512	2	11	0.1	2

2.2 Initial Conditions

All our simulations take place in a periodic domain filled with Solar metallicity molecular gas with a mean molecular weight of $2.33m_p$ (corresponding to a gas of 90 per cent H and 10 per cent He by number) and an initial temperature of 10 K. This combination of molecular weight and temperature corresponds to a sound speed of 0.19 km/s. Based on the findings of Krumholz & McKee (2008), Krumholz et al. (2010), and Myers et al. (2011), gas surface density is a key environmental parameter affecting the IMF, and we therefore wish to explore the effect of varying it. This in turn requires that we vary the mass and size of the box. However, since our goal is to conduct a clean experiment where the *only* possible reason for differences in the outcome is feedback, we must vary these dimensional parameters in such a way as to ensure that the simulations remain re-scaled versions of one another, i.e. so that all the dimensionless numbers describing the system *in the absence of feedback* (Mach number, virial parameter, etc.) remain unchanged. For a box of material with mean density ρ , length L , and mean magnetic field strength B , the re-scaling that accomplishes this is (Krumholz 2014)

$$\rho' = f\rho, \quad (1)$$

$$L' = f^{-1/2}L, \quad (2)$$

$$B' = f^{1/2}B. \quad (3)$$

This transformation leaves the Mach number, Alfvén Mach number, virial parameter, number of Jeans masses, and all other dimensionless parameters of the box unchanged for an arbitrary positive number f . However, this rescaling does change the surface density by a factor $(\rho'/\rho)(L'/L) = f^{1/2}$; it therefore changes the optical depth through the domain, and thus the response of the gas to feedback. Rescaling also changes the mass contained in the box, by a factor $(\rho'/\rho)(L'/L)^3 = f^{-1/2}$, and the free-fall time $t_{\text{ff}} = \sqrt{3\pi/32G\rho}$, by a factor $f^{-1/2}$.

For our study, we consider three different boxes: a ‘medium’ density box (run M hereafter) containing a mass $M_{\text{box}} = 1000 M_{\odot}$ with surface density $\Sigma = 1 \text{ g cm}^{-2}$ (corresponding to a mean density $\rho_0 = 7.0 \times 10^{-19} \text{ g cm}^{-3}$, box length $L = 0.46 \text{ pc}$), and ‘low’ and ‘high’ density boxes (L and H) hereafter, which are rescaled versions of run M with $f = 1/4$ and $f = 4$, respectively. Table 1 summarizes the three simulation types. The medium density case is roughly intended to correspond to regions of vigorous massive star cluster formation found in the Galaxy; the size and surface density are comparable to those observed for cold ATLASGAL clumps (Urquhart et al. 2018). The low-density case is intended to more closely represent a nearby low-mass star-forming region such as Perseus of Ophiuchus, while the high-density case is intended to mimic the formation environments characteristic of extra-galactic super star clusters such as R136 or the protoclusters of NGC 253 (Leroy et al. 2018).

All our simulations use gravity and magneto-hydrodynamics with periodic boundary conditions. For the radiation field, we use Marshak boundary conditions, whereby radiation can freely leave the domain,

but the domain is also bathed with an external radiation field that corresponds to a 10 K isotropic blackbody.

The simulation itself is divided into two phases. In the first phase, which begins with uniform density and magnetic field, we turn-off gravity and radiation, set $\gamma = 1.0001$ so that the gas is close to isothermal, and drive turbulence in the box using the method of Mac Low (1999). The driving field is purely solenoidal (Federrath, Klessen & Schmidt 2008), with constant power per wavenumber $P(k)$ over the range $1 \leq kL/2\pi \leq 2$. At every time step we apply a force with a spatial distribution matching the driving spectrum, and with an amplitude chosen to maintain a roughly constant mass-weighted RMS thermal Mach number of $\mathcal{M} = \sigma_v/c_s = 12.6$ where σ_v is the three-dimensional (3D) velocity dispersion. The corresponding virial parameter is (e.g. Federrath & Klessen 2012)

$$\alpha = \frac{5\sigma_v^2 R}{GM_{\text{box}}} = 1, \quad (4)$$

where we have taken the characteristic radius $R = L/2$. The corresponding turbulent crossing time is

$$t_{\text{cross}} = \frac{L}{\sigma_v} = (0.4, 0.2, 0.1) \text{ Myr} \quad (5)$$

for L, M, and H runs, respectively.

In addition to \mathcal{M} and α , the final dimensionless number that characterizes our simulation is the plasma β . We fix this parameter by initialising the uniform magnetic field at the start of the first phase of the simulation to the value B_0 required to have (e.g. Federrath & Klessen 2012)

$$\beta = \frac{8\pi\rho c_s^2}{B_0^2} = 2 \left(\frac{\mathcal{M}_A}{\mathcal{M}} \right)^2 = 0.012. \quad (6)$$

The corresponding Alfvén Mach number $\mathcal{M}_A = 1$. As noted above, by fixing \mathcal{M} , α , \mathcal{M}_A , and β , we also fix all other dimensionless numbers in the absence of radiation and feedback.

We evolve the initially uniform gas with driving for two crossing times, and then treat the cloud state at the end of the two crossing times as the initial conditions for the turbulent cluster simulation. At this point we turn on gravity, sink particles, and radiation and turn-off turbulence driving. We also set $\gamma = 5/3$ instead of 1.0001, which allows the temperature to vary with the outcome of the radiative transfer calculations. We run these simulations until they reach 5 per cent star formation efficiency (SFE), i.e. until 5 per cent of the initial gas has been accreted by star particles.

2.3 Resolution, refinement, and sink particles

We initialise the AMR hierarchy on a 512^3 base grid which we denote as $\mathcal{L} = 0$. During the driving phase, we disable AMR, so no higher levels exist. We use a 512^3 grid during this phase because this resolution is sufficient to capture a well-resolved inertial cascade down to the scale of $L/30$ (Federrath et al. 2010, 2011; Li et al. 2012). This scale is one order of magnitude less than the initial Jeans length

in our simulations, so during the driving phase this choice ensures that we have well-resolved turbulence down to scales well below the scale at which self-gravity is expected to become important.

Once we turn on gravity we allow the grid to adaptively refine to a maximum level $\mathcal{L}_{\max} = 2$. We use two different criteria to decide where to add resolution. First, we increase the resolution wherever fewer than eight zones resolve the local Jeans length, $J > 1/8$ (Truelove et al. 1997). This criterion triggers refinement to the finest level at a density of 1/4 of that required to trigger the insertion of a sink particle at the finest level, thereby ensuring that strongly collapsing regions get refined to the finest level. The second condition is that we refine any cell that contains a poorly resolved gradient of the radiation energy density, $|\nabla E_r| \Delta x / E_r > 0.125$, where Δx is the grid spacing. This ensures that we resolve the radiation field and temperature structure. Our choice $\mathcal{L}_{\max} = 2$ corresponds to a maximum resolution $\Delta x = 23$ AU for the medium run, and factors of two smaller (larger) for the high (low) surface density runs.

In our simulations, sink particles form in any zone on the finest AMR level where the gas becomes dense enough to reach a local Jeans number

$$J = \sqrt{\frac{G\rho\Delta x^2}{\pi c_s^2}} > \frac{1}{4}. \quad (7)$$

At a temperature of 10 K this corresponds to a density $\rho_{\text{sink}} = 8.8 \times 10^{-15} \text{ g cm}^{-3}$ for run M, and a factor of 4 lower (higher) for run L (H). However, in practice the gas near the sink creation threshold is often significantly warmer than 10 K once the simulation is underway, and thus the actual maximum resolvable density is substantially higher. Once formed, these sink particles evolve and interact with the gas via gravity and accretion following the procedure described in Krumholz et al. (2004). This has been updated to include the effects of the magnetic field on the gas accretion rate on sink particles (Lee et al. 2014).

There is a limitation in the method we use for radiation transfer: we assume that the gas and dust temperatures are equal. Since the gas and dust temperatures become identical at densities above $10^4 - 10^5 \text{ cm}^{-3}$ (e.g. Goldsmith 2001), this is a good assumption for most of the gas in the simulation box. However, in the low density, non-self-gravitating regions, the density may be too low to allow efficient gas-grain coupling, allowing gas to be either hotter or cooler than the dust, depending on the environment. However, this is unlikely to be important for the purpose of determining the IMF, since in the regions where gas is collapsing and fragmenting, the density is high enough that grains and gas are well-coupled. Simulations that include an explicit treatment of imperfect grain-gas coupling (Bate & Keto 2015) find that its effects on fragmentation are minimal (Bate 2019).

2.4 Parameter study

We carry out two different realizations for each simulation type (L, M, and H). For a given simulation type, the two realizations are identical in their mean properties and differ only in their randomly-generated turbulent driving fields. We refer to simulations using the first random realization as L1, M1, and H1, and those using the second random realization as L2, M2, and H2. Runs L1, M1, and H1 all use the same driving pattern, as do runs L2, M2, and H2. We also carry out one additional simulation, *ISO1*, which uses the same driving pattern and initial conditions as M1, but where we have disabled both outflow feedback and radiative transfer, and we continue to use an isothermal equation of state ($\gamma = 1.0001$) throughout the collapse phase.

We emphasize that since L1, M1, H1, and *ISO1* all use the same driving field (as do L2, M2, and H2), have identical dimensionless parameters, use the same initial and maximum resolution, and use the same Jeans refinement criterion (in the sense that the factor by which a given cell must be overdense relative to the box mean in order to be flagged for refinement is the same in all cases), the initial conditions and simulation setup are *identical to machine precision*. The only physics elements of L1, M1, H1, and *ISO1* that are not identical are radiation and outflow feedback;¹ these break the symmetry between the runs, and introduce characteristic length and mass scales that would not exist in their absence. Any differences in outcome between the various runs must therefore be due to one of these factors.

3 RESULTS

Here, we present the results of our simulations. We first give an overview of the global evolution of the simulations in Section 3.1. In Section 3.2 we discuss the sink particle mass distribution, and identify differences between the various runs. In Sections 3.3 and 3.4 we discuss the roles of radiation feedback and protostellar outflows, respectively, in shaping the mass distribution.

3.1 Overview of simulations

In Figs 1 and 2, we show the column density and density-weighted mean temperature of runs L1, M1, and H1, one of our two turbulent realizations; while the exact morphology is of course different for realization 2, the qualitative discussion we present here applies equally well to that case. Since star formation occurs at slightly different rates in each of these simulations, we show results at matching star formation efficiencies rather than matching times, though the difference is relatively small (provided time is measured in units of the simulation free-fall time). In all three runs the turbulence has created dense filamentary structures, which over time further fragment into isolated gas cores that collapse to form stars. As can be seen from Fig. 1, star formation activity is mostly confined in the filamentary regions. The overall morphologies of the L1, M1, and H1 runs are very similar, which is not surprising since, absent the effects of feedback, the runs would be identical. However, some small differences are clearly visible: as the column density of the runs increases the filamentary shapes become more disrupted; filaments are narrower and straighter in run L1, and broader and more dispersed in run H1.

To better understand the collapse morphology, fragmentation and star formation we look at the temperature structure of the gas shown in Fig. 2. This shows much larger differences between the runs than Fig. 1. Compared to run L1, run H1 is warmer because it has a higher surface density, which leads to a higher optical depth that traps the radiation produced more effectively. The difference between the runs increases with SFE as the influence of radiation feedback on the temperature grows.

We show the time evolution of the SFE in Fig. 3, and the total number of stars present as a function of the SFE in Fig. 4. We define the SFE as the fraction of initial box mass converted to stars, i.e.

$$\text{SFE} = \frac{m_\star}{M_{\text{box}}} = \frac{m_\star}{M_{\text{gas}} + m_\star}, \quad (8)$$

¹Exact identity to machine precision is also broken slightly by differences in numerical truncation error caused by the fact that the time steps in the different runs are not completely identical. However, this is not a significant physical effect.

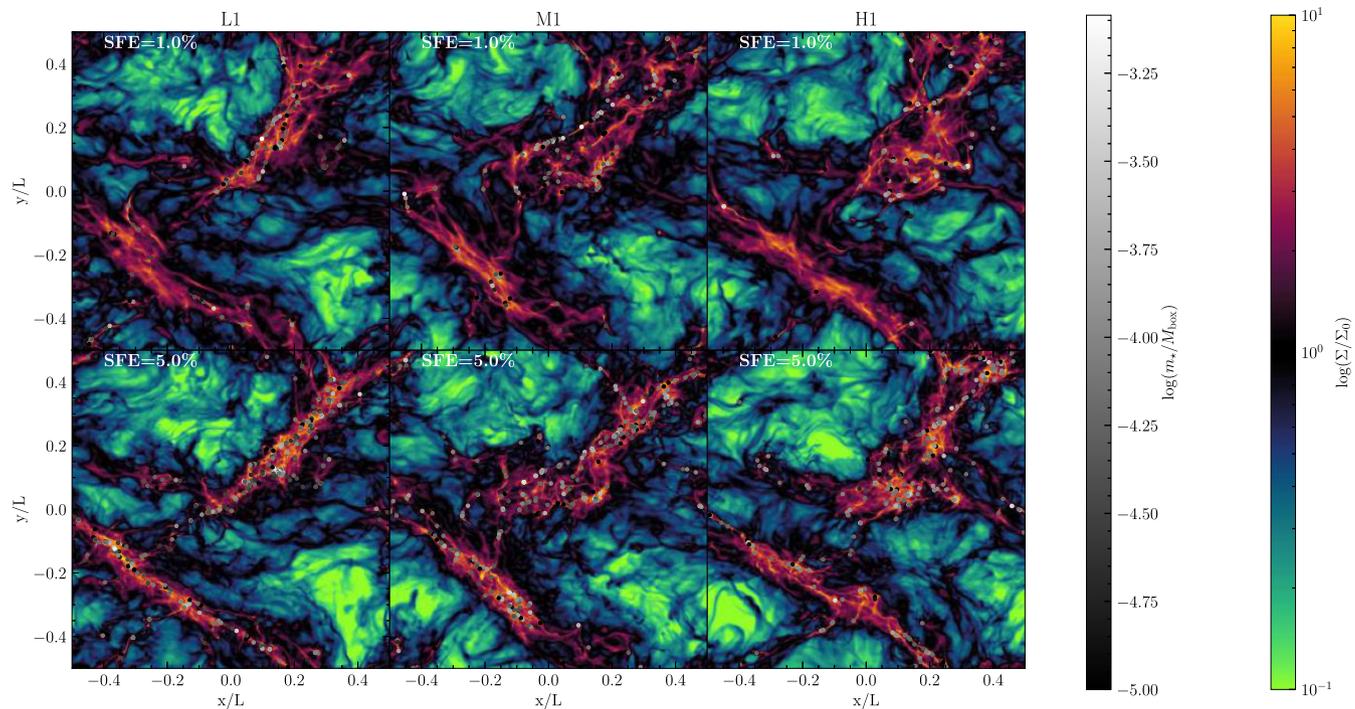


Figure 1. Column densities in simulations L1, M1, and H1 (left to right). The top row shows the simulation state at 1 per cent SFE, the bottom row at 5 per cent SFE. The colour scale goes from $\log(\Sigma/\Sigma_0) = -1$ to 1, where $\Sigma_0 = \rho_0 L$ and ρ_0 is the mean density in the simulation domain. Circles show star particles, and are colour-coded by mass m_* from $\log(m_*/M_{\text{box}}) = -5$ to -3 , where M_{box} is the total mass of the simulation box.

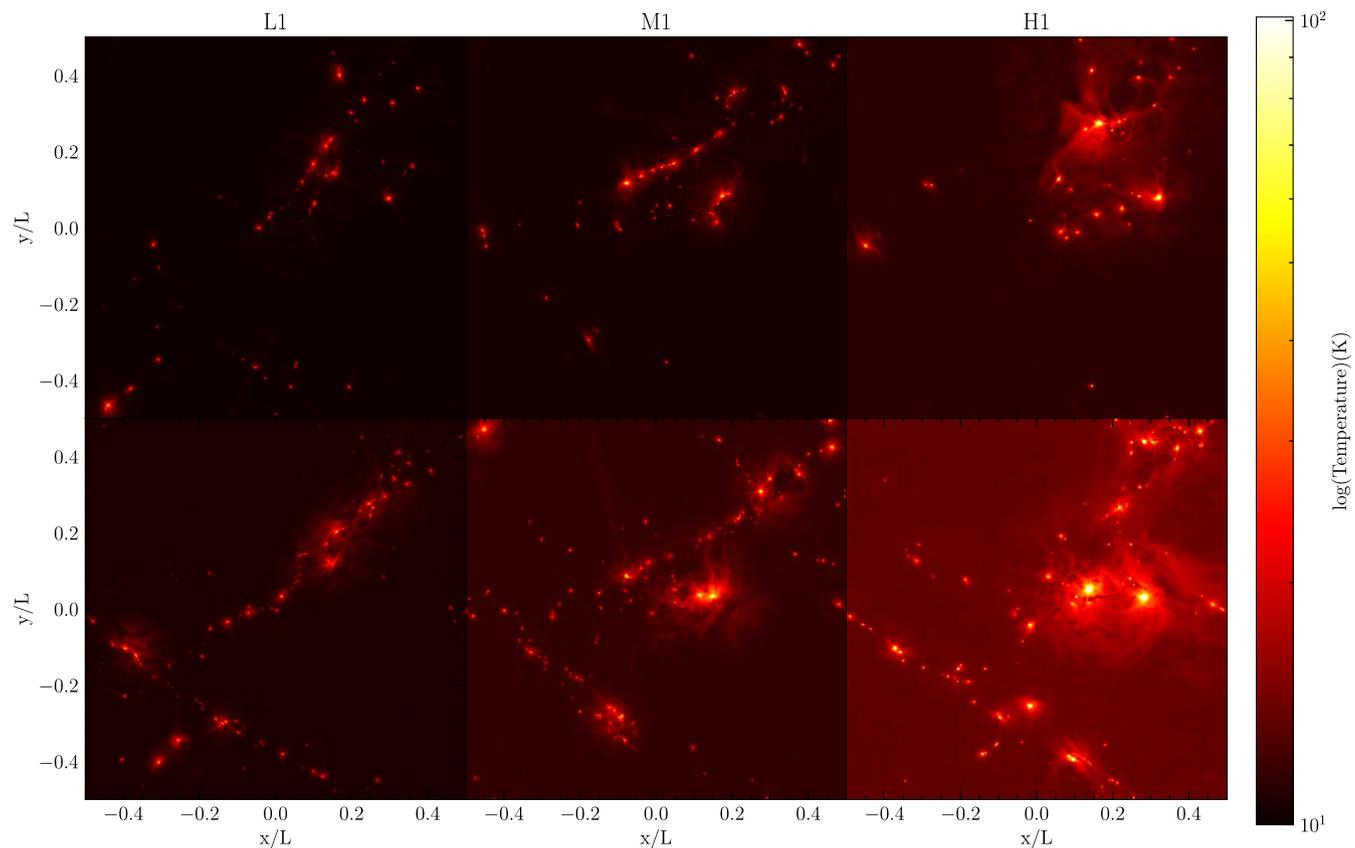


Figure 2. Same as Fig. 1, but showing density-weighted projected temperature rather than column density.

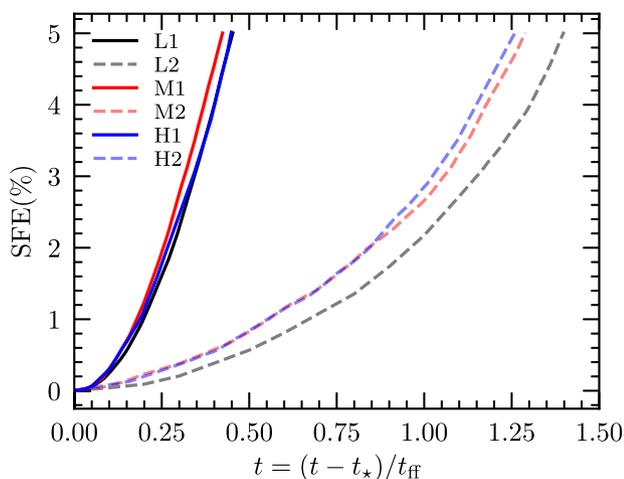


Figure 3. SFE as a function of time since the formation of the first star t_* , measured in units of the free-fall time t_{ff} . The solid lines are realization 1 and the dashed lines are for realization 2.

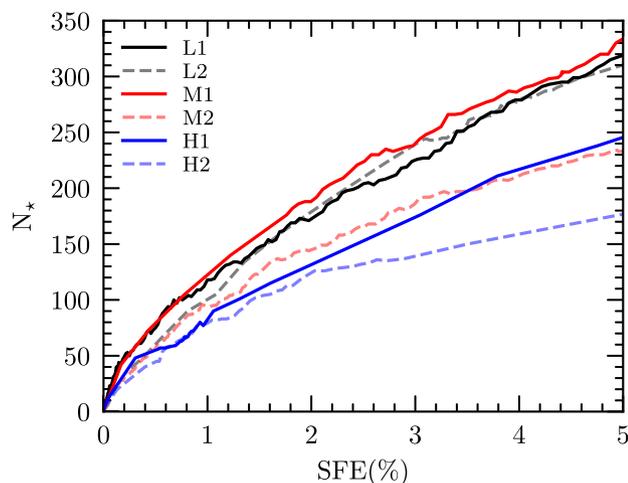


Figure 4. Number of stars formed as a function of SFE. The solid lines are realization 1 and the dashed lines are realization 2.

where m_* is the total stellar mass and M_{box} is the total of the gas and stellar mass in the simulation (which is always equal to the initial mass in the box). From the onset of star formation (which is itself $\sim 0.5 t_{\text{ff}}$ after we turn on self-gravity), it takes approximately $0.5 t_{\text{ff}}$ for run L1 to reach 5 per cent SFE. For run L2 it takes almost 3 times as long to reach the same SFE. The difference is simply due to the different random realizations in the driving pattern, and is within the normal range of variation with driving pattern (e.g. see fig. 4 of Nam et al. 2021). By contrast, there is fairly little difference between runs with the same driving pattern but different surface densities. This suggests that, whatever effects feedback is having in reducing the SFE, those effects are almost independent of surface density.

If we examine Fig. 4, which shows the number of stars versus SFE, we see a somewhat different pattern. The differences between the two realizations (solid versus dashed lines) are relatively modest, but there is a clear trend with surface density: at fixed SFE, the high surface density run has produced the fewest stars, and the low surface density run the most. The one exception to this is runs L1 and M1, where the numbers of stars are quite similar. Thus, we see that the SFE does not depend strongly on the absolute surface density (in our

scaled experiments), but the number of stars at fixed SFE, and by implication the stellar mass distribution, does.

3.2 Stellar mass distribution

We now examine the stellar mass distributions in the simulations more closely. In all the figures we present in this section, we show results using two different mass scales: absolute, and relative to box mass (i.e. stellar masses are expressed as $m = m_*/M_{\text{box}}$). In the absence of both radiation feedback and outflow feedback the relative mass distribution would be identical for all the simulations, and the absolute mass distributions would simply scale as $f^{-1/2}$, i.e. we should find that the absolute mass distributions in the H runs have the same shape as in the corresponding M runs, but shifted to lower masses by a factor of 2; similarly, L runs should be shifted a factor of 2 higher in absolute mass than M runs.

We begin by looking at the two sets of plots in Fig. 5, which shows the evolution of the median and 25th to 75th percentile ranges of the sink particle mass distributions for all sets of simulations. Here we measure the median and percentiles with respect to stellar mass rather than number, i.e. if m_{25} is the 25th percentile mass, this means that the stars with masses $m < m_{25}$ together constitute 25 per cent of the stellar mass in the simulation at that time, not that they are 25 per cent of the stars by number.² From the plots we can see that the median masses are relatively converged by the time we reach 5 per cent SFE, and are increasing only very slowly. Though most of the sink particles at 5 per cent SFE are still accreting, this is counterbalanced by the fact that new sink particles are forming at the same time. Therefore, the population as a whole has reached nearly a steady state distribution. When measured on an absolute scale (top panels), we see that the median mass, along with the 25th to 75th percentile range, decreases as the surface density increases, as we would expect. However, the decrease is not by exactly the amount that would be expected for scale-free behaviour; when measured relative to the box mass, the stellar mass distributions are not identical but instead show clear variations.

To explore these differences further, in Fig. 6 we show the cumulative mass functions of the runs at 5 per cent SFE, again on both an absolute and a relative mass scale; as with Fig. 5, we measure the CDF with respect to mass, not number. From Fig. 6 it is clearly evident that the radiation feedback and outflow feedback have broken the symmetry between these runs. However, the effect is not simply a shift in the median, but is rather a change in the shape as well. The mass range in the higher surface density cases is narrower than for lower surface density runs. Measured relative to the box mass, we see a characteristic pattern that the higher surface density runs (blue) have fewer low-mass stars (i.e. the CDF begins to rise above zero at higher mass), but the same runs *also* have fewer very high mass stars (i.e. the CDF approaches unity at lower mass). Viewed on an absolute scale, the higher surface density runs have IMFs that are both lower in mean mass and narrower in range of mass. The no-feedback run appears as a continuation of this trend, i.e. the absolute mass scale is even larger than for run L1, and the mass distribution is even broader.

²We prefer to work with percentiles computed by mass rather than number, because the latter are quite sensitive to the numerical details of the sink particle algorithm, such that different choices for parameters such as the Jeans number used when calculating the sink particle creation density threshold can lead to different numbers of very low-mass objects (e.g. see Appendix C of Haugbølle et al. 2018). Percentiles computed by mass are far less sensitive to such numerical details.

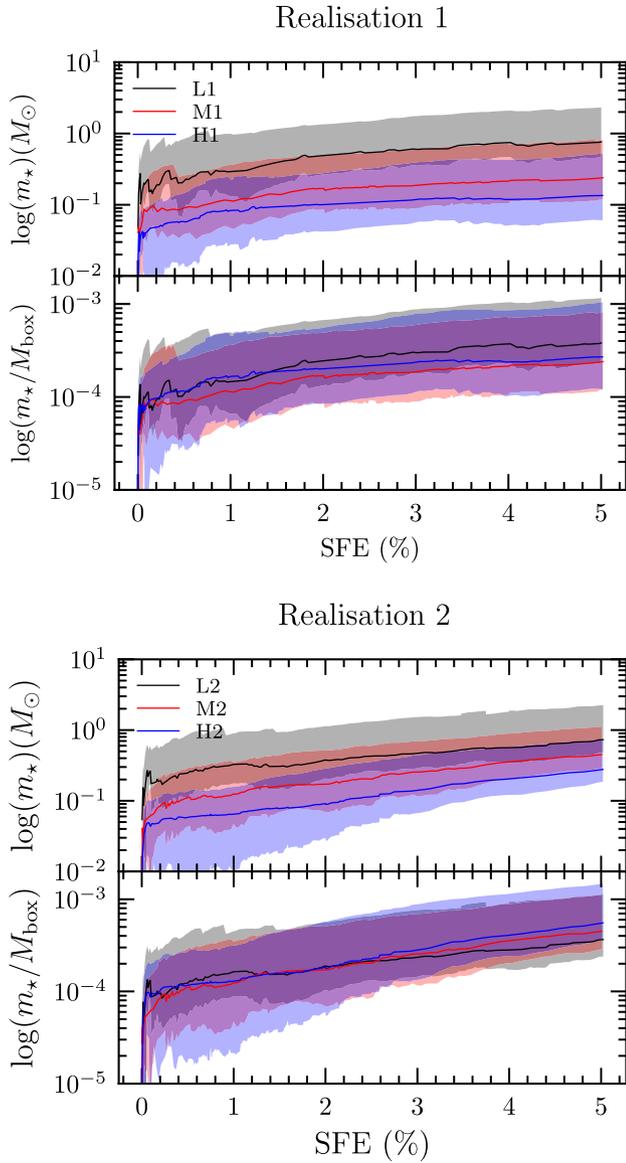


Figure 5. Solid lines show the evolution of the median of the sink particle mass distribution for realization 1 (top two panels) and realization 2 (bottom two panels), while the shaded regions around them show the 25th to 75th percentile range; here medians and percentiles are measured with respect to mass rather than number, i.e. the mass m_* plotted is the one for which half the stellar mass is in stars with masses $< m_*$. In each pair of panels, the upper one shows the absolute mass expressed in M_\odot , while the lower one shows the mass expressed relative to the box mass M_{box} . The black, red, blue lines and shaded regions correspond to the L, M, and H runs, respectively.

3.3 The effect of radiation feedback on the simulations

We now seek to understand the mechanisms by which feedback breaks the symmetry between our different cases and produces the differences in the IMF that we have observed. We begin in this section by studying the effects of radiation feedback, which alters the gas temperature distribution. Fig. 7 shows the distribution of the gas mass with both density and temperature in runs L1, M1, and H1 at SFEs of 1 per cent and 5 per cent. In this and all similar figures in this section, realization 2 is qualitatively the same as realization 1, so we only show plots of realization 1 for reasons of space. From the plot it is clear that as the surface density increases the simulations

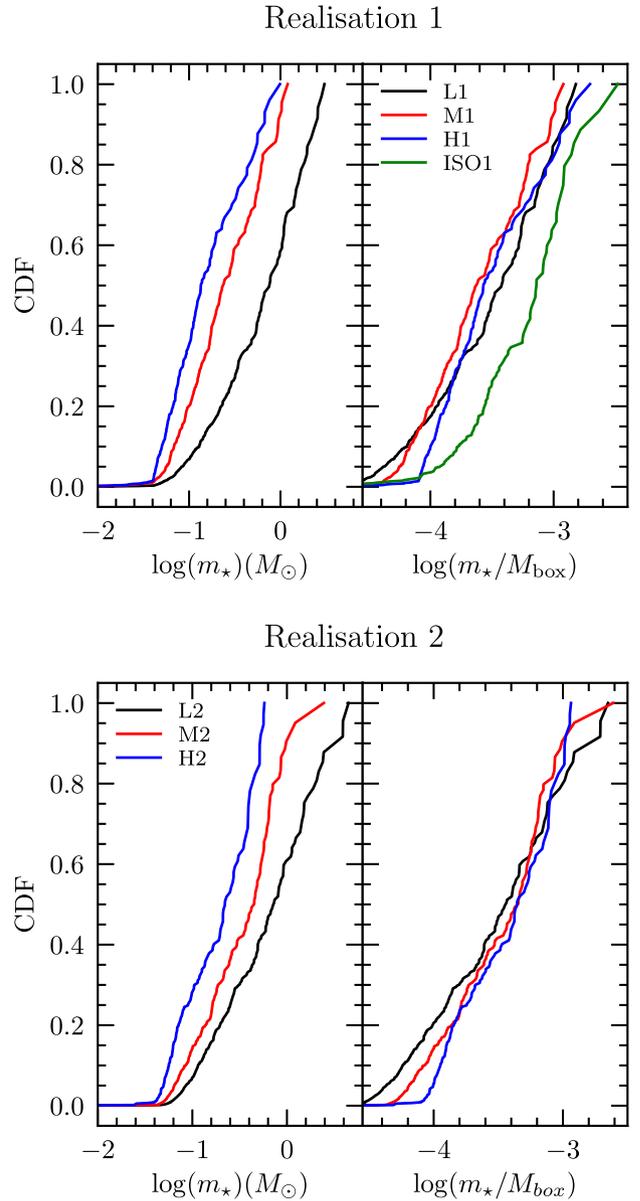


Figure 6. Cumulative distribution function (CDF) of the sink particle masses of the simulations; we show realization 1 in the upper pair of panels, and realization 2 in the lower pair. The black, red and blue lines refer to the low, medium and high surface density runs. The green line denotes the run with no feedback. The left sides of both pairs of panels show the CDF with respect to absolute stellar mass, while the right sides show the CDF for mass measured relative to the box mass M_{box} . The low surface density runs have a broader mass range and higher mean mass than the higher surface density runs. The no-feedback run is a continuation of this trend with a broader mass range and higher mean mass than the lower surface density runs.

become progressively warmer, therefore suppressing fragmentation. This is consistent with earlier studies of radiation feedback, both numerical (e.g. Krumholz et al. 2010; Myers et al. 2011; Mathew & Federrath 2020, 2021) and analytic (e.g. Chakrabarti & McKee 2005; Krumholz 2006, 2011; Krumholz & McKee 2008; Guszejnov et al. 2016) which find that higher surface densities trap radiation more effectively, leading to systematically higher gas temperatures at fixed ratio of stellar mass to gas mass. To help visualize this effect, the dashed lines in Fig. 7 illustrate loci of constant box-normalized Jeans

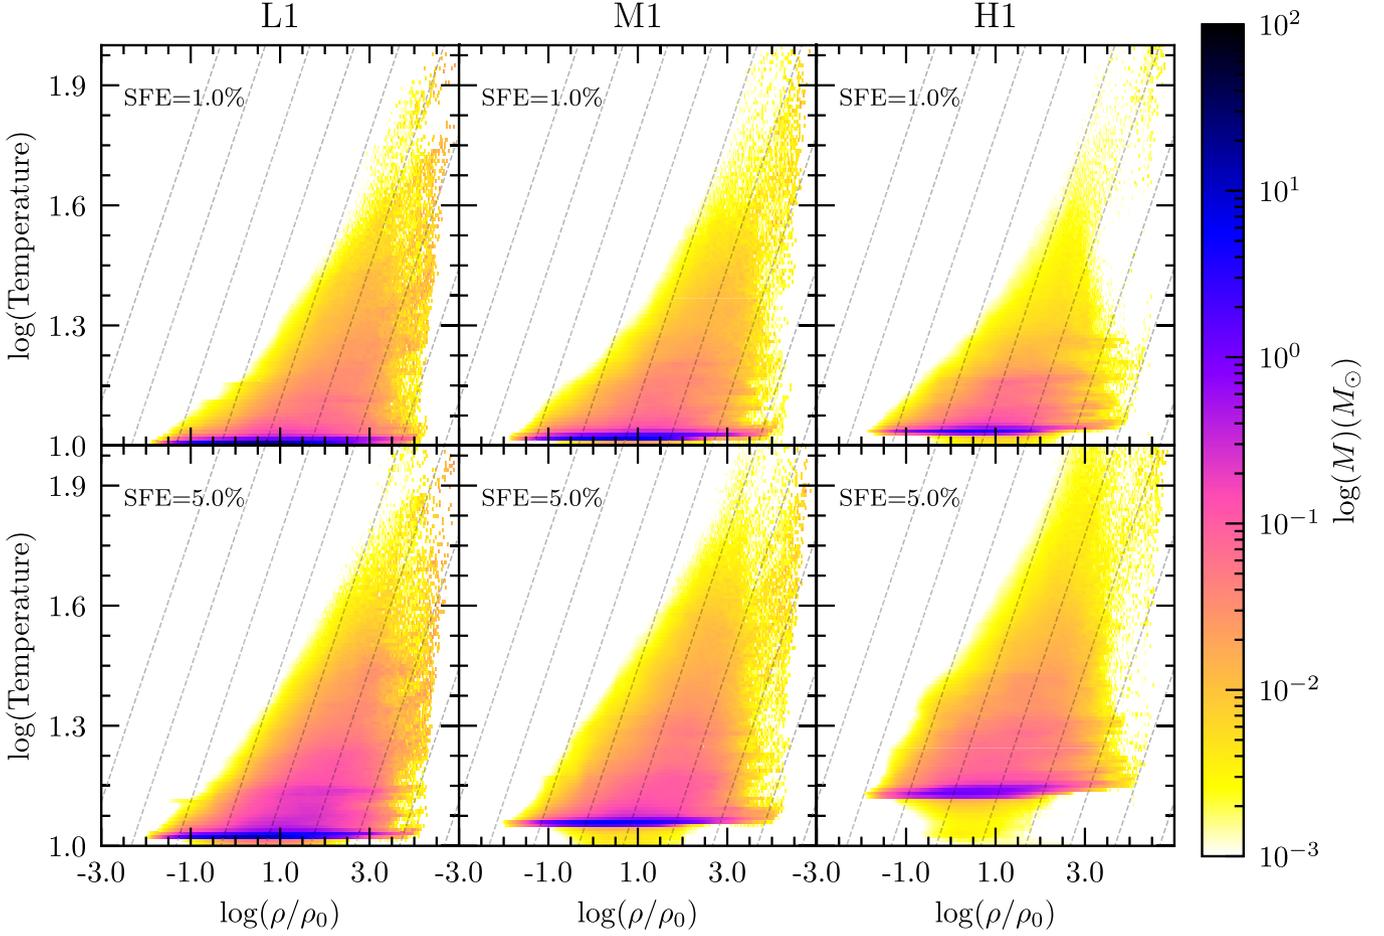


Figure 7. The joint distribution of normalized density ρ/ρ_0 and temperature T for runs L1, M1, H1 (left to right). The top row shows the state of the simulations at 1 per cent SFE and the bottom row at 5 per cent SFE. The colour bar shows how much mass is in each density-temperature bin. Dashed lines indicate loci of constant box-normalized Jeans mass M_j/M_{box} ; lines are spaced logarithmically at intervals of 0.5 dex, with the left-most line corresponding to $\log M_j/M_{\text{box}} = -2$.

mass M_j/M_{box} , where

$$M_j = \frac{\pi}{6} \left(\frac{k_B T}{G \mu} \right)^{3/2} \frac{1}{\rho^{1/2}} \quad (9)$$

is the Jeans mass. It is clear that, as we move from L1 to M1 to H1, the mass distribution shifts systematically to the left relative to these lines, indicating less mass at low M_j/M_{box} .

To explore this effect further, we show the 1D mass-weighted CDF of M_j/M_{box} in Fig. 8. The way to interpret this plot is that each point on the horizontal axis corresponds to one of the dashed lines in Fig. 7, with higher values of M_j/M_{box} (moving right in Fig. 8) corresponding to moving left in Fig. 7. The value on the vertical axis in Fig. 8 indicates what fraction of the mass in the simulation lies to the right of the corresponding dashed line in Fig. 7. We see a clear trend with surface density: in run M1, the distribution extends to lower M_j/M_{box} than in L1, which in turn extends to lower values than H1.

To visualize how this trend might impact fragmentation, in the plot we also show the 1-1 relation (dashed lines). The significance of the 1-1 lines becomes clear if we recall that the CDF is also a quantity measured in units of a fraction of the total simulation mass M_{box} . Consider for example $M_j/M_{\text{box}} = 10^{-5}$ at a SFE of 1 per cent (left-hand panel). Examining the plot, we see that the CDF of run H1 is below the 1-1 line at this value of M_j/M_{box} ; the CDF value

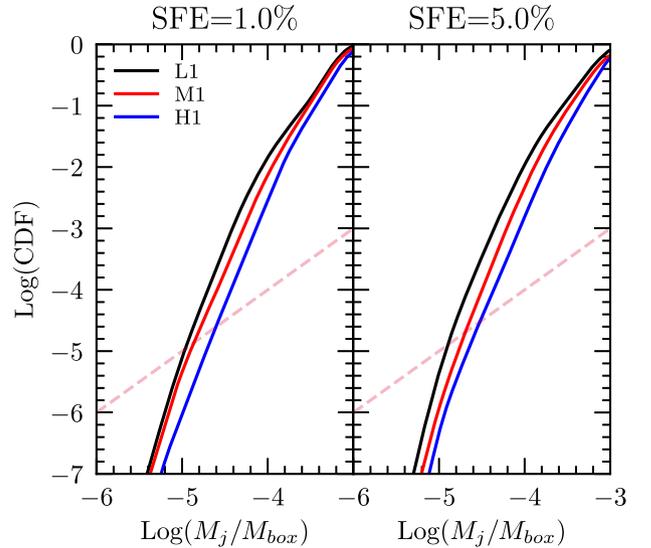


Figure 8. CDFs of M_j/M_{box} in runs L1, M1, and H1 (black, red and blue lines) at SFEs of 1 per cent (left) and 5 per cent (right). The dashed one-to-one lines indicate a minimum condition for fragmentation: for any mass M_j/M_{box} for which the CDF falls below the dashed line, there is less than a single mass of Jeans material in the box, and thus it is impossible to create an object of that mass via gravitational collapse.

here is $\approx 10^{-6.2}$, i.e. the amount of mass in the box for which $M_J < 10^{-5}M_{\text{box}}$ is $\approx 10^{-6.2}M_{\text{box}}$. Thus, in the entire simulation box, there is only $\approx 10^{-6.2}/10^{-5} = 6.3$ per cent of a Jeans mass of material dense and cold enough that its Jeans mass is $< 10^{-5}M_{\text{box}}$ – there is too little mass in the box to make even a single object this small via gravitational collapse. By contrast, at the same value of M_J/M_{box} , run L1 is above the 1-1 line, and has a CDF value of $\approx 10^{-4.8}$, so there is $1 \approx 10^{-4.8}/10^{-5} = 160$ per cent of a Jeans mass of material. Of course this is only a necessary condition for the creation of such a small collapsed object, not a sufficient one; the CDF is a point statistic, and tells us nothing about how this mass is arranged. It may be split up into a dozen locations spread throughout the simulation volume, such that no single contiguous region has enough mass to be gravitationally unstable and collapse. Nonetheless, it is clear that runs L1 and H1 are fundamentally different at this mass: run L1 can in principle make collapsed objects with masses as small as $10^{-5}M_{\text{box}}$, while run H1 cannot. The intersection of the dashed 1-1 line with the CDF therefore represents at least a rough estimate of the minimum object mass that it is possible for gravity to produce.

Armed with this understanding, we can now make sense of the differences we observe at the low-mass ends of the IMFs produced in the different simulations. Fig. 8 shows that fragmentation to the smallest objects is suppressed more at the higher surface density runs than the lower surface density runs. Thus radiative feedback suppresses the IMFs at the low mass end, which is exactly what we observe: when measured in box-normalized units, run H has fewer very low-mass stars, and run L has more. However, we remind readers that this statement applies to masses measured in normalized units – while radiation feedback does push the low end of the IMF upward in run H1 compared to L1, it does not do so by enough to fully compensate for the overall lower value of the mean Jeans mass in run H1 compared to L1. Thus the lowest-mass stars in run H1 are still (slightly) smaller than those in run L1 – just by much less than the factor of four difference that would be expected given the differences in mean density between the two cases.

However, our analysis thus far does not explain what is happening at the high-mass end of the IMF. If radiation were the only actor, we would expect that at higher surface density, more mass should be accreted on to the high mass stars as well, leading the overall IMF to shift upwards (in normalized units). That is not what we observe in Fig. 6. This suggests that outflow rather than radiative feedback is dominant at the high mass end of the IMF. We turn to this question next.

3.4 The effect of the outflow feedback on the simulations

If radiation feedback were the only feedback mechanism operating, our analysis from the previous section suggest that we should see a uniform shift of the IMF towards higher (normalized) masses in the high surface density simulations. However, this is not what we observe: the higher surface density simulations produce fewer low-mass stars (i.e. the IMF is shifted upward at the low-mass end), but *also* fewer – or at least an equal number of – high-mass stars. There is no obvious mechanism by which radiation could produce such an effect, since in the high surface density runs are uniformly warmer, and higher temperature favours more massive objects. Instead, the most likely explanation is that the decrease, or lack of increase, in the abundance of high-mass stars in the high surface density runs is that protostellar outflow feedback is breaking

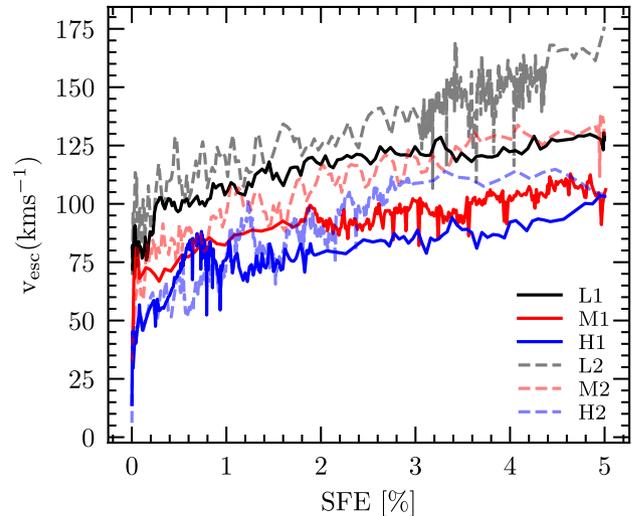


Figure 9. Accretion weighted mean surface escape speed of the simulations against SFE. The solid lines are realization 1 and the dashed lines are realization 2.

up the gas more effectively in these runs, making it harder for more massive objects to grow.

To understand how and why outflow feedback depends on the surface density, in Fig. 9 we show the evolution of accretion rate-weighted mean surface escape speed of the runs as a function SFE. This quantity matters for outflows, because our outflow prescription links the outflow velocity to the stellar surface escape speed – stars with higher escape speeds have faster outflows – and thus the accretion rate-weighted mean surface escape speed is directly proportional to the outflow momentum flux. The correlation between surface escape speed and outflow speed implemented in our model is both seen observationally (and references therein Richer et al. 2000) and expected theoretically (and references therein Konigl & Pudritz 2000), since outflows are launched from close to the stellar surface; see Cunningham et al. (2011) and Federrath et al. (2014) for details.

From the plot it is clear that the lower surface density runs have systematically higher escape speed than the high surface density runs, and thus inject more outflow momentum (in normalized units). This difference occurs because in physical time units, the lower surface density runs take the longest to reach any given SFE, which gives the stars more time to contract toward the main sequence. Just as the optical depth of the domain is a quantity that is not conserved as we rescale from the L to the M to the H runs, and this breaks the symmetry with respect to radiation feedback, the number of Kelvin-Helmholtz times per free-fall time is also not conserved under rescaling, and this breaks the symmetry with respect to outflow feedback.

To see how this symmetry breaking manifests in the simulations, in Fig. 10 we show the outflow kinetic energy, scalar momentum and volume occupied normalized to the box kinetic energy, momentum and volume; formally, we normalize the kinetic energy by $(1/2)M_{\text{box}}\sigma_v^2$ (where σ_v is the initial 3D velocity dispersion), the momentum to $M_{\text{box}}\sigma_v$, and the volume to L^3 . For the purposes of this plot, we define outflow material using ORION2’s passive scalar capability: material that is accreted by stars and then re-inserted into the computational domain by our stellar outflow prescription is tagged with a passive scalar that is conservatively transported thereafter, and thus in every cell at every time we know the partial density ρ_{out} of outflow material. The outflow kinetic energy, scalar momentum and

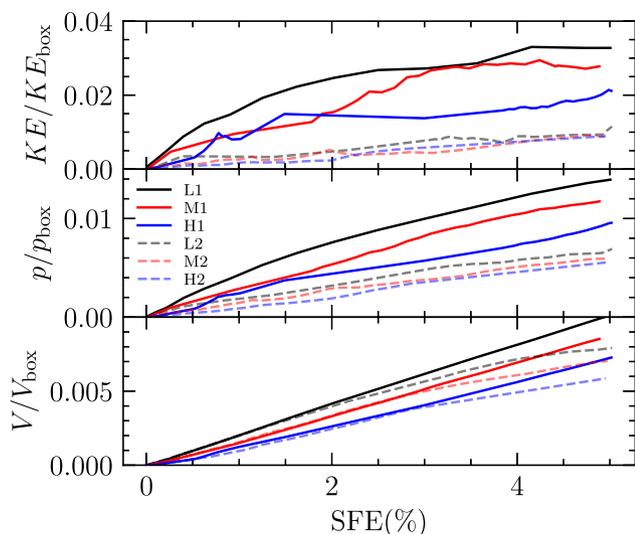


Figure 10. Outflow kinetic energy, outflow momentum and outflow volume relative to the box kinetic energy, momentum, and volume, all as a function of SFE. The solid lines represent realization 1 and the dashed lines represent realization 2.

volume in a given cell of total density ρ , velocity v , and volume V are then $(1/2)\rho_{\text{out}}v^2$, $\rho_{\text{out}}v$, and $(\rho_{\text{out}}/\rho)V$.

Examining Fig. 10, we see that outflows in the low surface density runs carry almost three times as much kinetic energy and twice as much scalar momentum, and occupy twice as much volume, as in the high surface density runs. Again, we remind the reader that, due to the symmetry between the L, M, and H runs, by construction the only possible sources for these differences are the differing rates of outflow momentum injection (c.f. Fig. 9), or a secondary effect of the differing temperature distributions that results from the broken symmetry of stellar radiation feedback. Regardless of the source of the difference, one might expect based on Fig. 10 that outflow feedback would be more effective in the L runs, not the H ones as we seem to require. However, the higher outflow occupation volume, coupled with the more fragmented and less filamentary density distribution we see in Fig. 1, suggests a solution. The greater outflow injection momentum in the low surface density runs makes it easier for the outflows to punch holes in the collapsing cores in these cases. The outflows rapidly break out of the dense gas, transferring less of their momentum in the process, and more easily reaching low-density regions where they are able to expand and occupy more volume. This leads to less fragmentation of the dense gas in the low surface density runs, and therefore makes outflow feedback less effective, leading to a heavier IMF at the massive end.

4 DISCUSSION

4.1 Implications for IMF variations

In this section, we compare our IMF with the observable IMFs. For the purpose of quantitative comparison between simulation and observation we compare the CDFs from our simulations with the Chabrier (2005) IMF. Since our simulations do not produce enough stars to fully sample the high mass end of the mass distribution, it is important to perform this comparison accounting for stochastic effects. For this purpose we generate 50 000 clusters with masses of 100, 50, and 25 M_{\odot} , corresponding to the total masses in our L, M, and H runs. To create these clusters we draw stars from a Chabrier

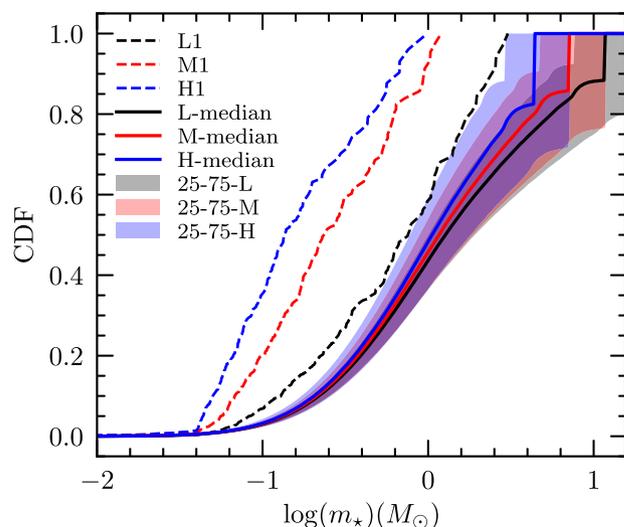


Figure 11. CDF of the sink particle masses of the simulations (dashed lines) compared to the CDF for a population of stars with the same total stellar mass as the simulations drawn from a Chabrier (2005) IMF; for the Chabrier IMFs, the central solid line is the 50th percentile value, and the shaded region shows the 25th–75th percentile range. The black, red and blue lines refer to the low, medium and high surface density runs. Note that the step-like discontinuity at high mass in the Chabrier IMF is a real effect, caused by the fact that there is a single most massive star in any real sample, which for the small clusters we are producing contains a not insignificant fraction of the total mass.

IMF truncated at 0.01 M_{\odot} on the lower end and 120 M_{\odot} on the higher end. In this manner we generate 50 000 sample clusters for each of our target masses; we then compute the CDF for each of the samples, from which we can compute the p th percentile value of the CDF at any given stellar mass m_* , which we can compare to our simulation CDFs as shown in Fig. 6. We show this comparison in Fig. 11, plotting the 25th, 50th and 75th percentiles (again computed by mass) for the sample clusters.

Fig. 11 shows the CDFs of runs L1, M1, and H1, along with the CDFs produced by drawing stellar populations of equal mass from a Chabrier IMF. It is obvious from the plot that run L1 is reasonably close to a Chabrier IMF, while both runs M1 and H1 are shifted to much lower masses. Much of this shift is at the high mass end, where our simulations do not form as many massive stars as would be expected from an unbiased draw from the IMF. In the L1 runs 40 per cent of the total stellar mass is stored in stars $> 1 M_{\odot}$ while the H run produces no stars $> 1 M_{\odot}$. This is at least in part a time effect, in the sense that, while we have run long enough for the median mass in our simulations to stabilise, we have not run long enough for the top of the mass distribution to do so; indeed, this is obvious simply from the fact that, since we stop the simulations at 5 per cent SFE, we could not have formed a star more massive than 50 M_{\odot} in run M (100 M_{\odot} in L, 25 M_{\odot} in H) even if all the collapsed mass had gone into a single object. Accounting for this effect, it is clear that the typical stellar mass in run L1 is reasonably close to what might be expected from a Chabrier IMF, while that in runs M1 and H1 is substantially lower.

4.2 Implications for the mass to light ratio in early type galaxies

As discussed in Section 1, one of the primary motivations for this work is to attempt to understand the variations in apparent IMF that have been observed in early type galaxies (ETGs). While there are multiple lines of evidence for this variation, the most straightforward

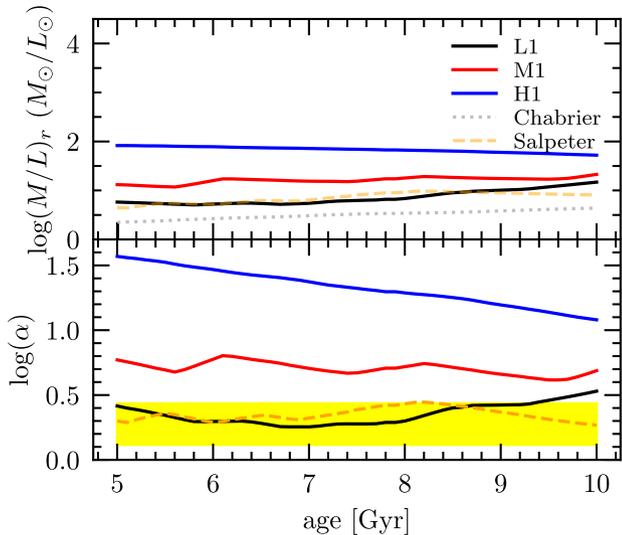


Figure 12. The top panel shows the mass to light ratio of the stellar population formed in the simulations in Solar units, and for comparison the mass to light ratio of stars drawn from a Chabrier (2005) IMF and a Salpeter (1955) (dashed lines), all as a function of age. The black, red, blue solid lines are of the simulations L1, M1, H1 respectively. The bottom panel shows the IMF mismatch parameter $\alpha = (M/L)/(M/L)_{\text{Chabrier}}$, again as a function of stellar population age. The yellow band shows the range of α values of 41 ETG from Gu et al. (2022).

to extract from our simulations is the mass to light ratio of the stellar populations we produce. In this section we therefore look at the mass to light ratio in our simulations for the purpose of comparing to that in observed galaxies. We calculate this by using the SLUG stellar population synthesis code (da Silva, Fumagalli & Krumholz 2012; Krumholz et al. 2015) to generate isochrones at stellar population ages from 5 Gyr to 10 Gyr, using the MIST stellar evolution tracks (Choi et al. 2016) and Starburst99-style stellar atmosphere models (Leitherer et al. 1999). Each isochrone provides a prediction of present-day mass, bolometric luminosity, and luminosity in a range of photometric filters as a function of initial mass for stars with initial mass $\geq 0.1 M_{\odot}$; we assume that the luminosities of stars with initial masses less than $0.1 M_{\odot}$ are negligible, and that these stars also experience negligible mass loss. We further assume that all stars with initial mass $< 8 M_{\odot}$ (which are all the stars formed in our simulations) that reach the end of their lives leave behind $0.7 M_{\odot}$ white dwarf remnants. We use the isochrone to calculate the luminosity and present-day mass of all the stars formed in each of our simulations at ages from 5 – 10 Gyr, and from these we calculate the mass to light ratio of the stellar population as a function of age for each of our simulations. For comparison, we use the same isochrones to calculate the mass to light ratio of Chabrier (2005) and Salpeter (1955, truncated at a lower mass limit of $0.1 M_{\odot}$) IMFs at the same ages. We use the SDSS r band for this calculation, but results are qualitatively similar in other filters.

The upper panel of Fig. 12 shows the mass to light ratio for the stellar populations produced in our simulations, and for comparison the dashed and dotted lines show the mass to light ratios of populations drawn from Chabrier and Salpeter IMFs at the same age. We see that L1 has a mass to light ratio intermediate similar to or slightly lower than Salpeter, M1 is a few tenths of a dex heavier than Salpeter, while H1 is much heavier. Thus the IMF variations we measure in our simulations would manifest as mass to light ratio variations in observations at levels comparable to those

observed between ETGs and local spirals. The differences between the simulations, and between the simulations and the Chabrier and Salpeter IMF, are much larger than can plausibly be explained by sampling effects associated with the difference in total stellar mass between the runs, as we demonstrate in Appendix A.

To emphasize the similarity between the M/L variations we find in our simulations and those seen in observations, the lower panel of Fig. 12 shows the IMF mismatch parameter, $\alpha = (M/L)/(M/L)_{\text{Chabrier}}$; i.e. $\alpha = 1$ corresponds to a stellar population that has the same M/L as a Chabrier IMF. For comparison, the yellow band shows the range of α measured for 41 ETGs by Gu et al. (2022); the band shown shows the range that contains 90 per cent of their sample values. The dashed line shows the mismatch parameter of Salpeter IMF. Gu et al. find that most of their sample galaxies have IMFs that deviate from their reference Chabrier IMF. In Fig. 12, we can see that run L1’s mismatch parameter is similar to that measured by Gu et al., while M1 has a slightly heavier IMF and H1 a much heavier one. The absolute value of α is probably not of tremendous significance here, due to the systematic uncertainties in both the simulations and the observations. The more significant point is that our L and M runs are producing differences in mass to light ratio comparable to those observed between ETGs and local spirals. They therefore provide a plausible explanation for how such M/L variations could arise.

4.3 Effects of systematic variations of SFE

In this paper, we have run our simulations up to 5 per cent SFE. At this fixed SFE, the median and the percentile ranges of the stellar mass do approach a steady state. However, the IMF is not fully converged at this point, because the simulations do not produce enough massive stars to fully sample the IMF at the massive end. This has potentially important implications for variations of the IMF, because Guszejnov et al. (2022) find that the final SFE of a cloud varies systematically with the cloud surface density. They argue that this variation suppresses variations in the IMF, because in their simulations the lower stellar masses found in higher surface density clouds at fixed SFE are largely cancelled by a systematic increase in stellar mass with SFE, coupled with higher SFEs in higher surface density clouds.

While our simulations do not contain massive stars capable of disrupting clouds and thus limiting the final SFE, even if there were systematic SFE variations, it seems unlikely that this effect could fully remove the IMF variations we have found. To demonstrate this, in Fig. 13 we show the mass functions of run L1 and M1 at 1 per cent, 3 per cent, and 5 per cent SFE; we omit H1 and realization 2 to avoid clutter, but the results are qualitatively similar. Significantly, we see that (1) the CDFs at 3 per cent and 5 per cent are very similar, so replacing one with the other would make relatively little difference, and (2) even if the SFE were to vary much more extremely, for example reaching 1 per cent in the low surface density case and 5 per cent in the medium surface density case, the IMF at medium surface density would *still* be more bottom-heavy than at low surface density. This plot clearly shows that while a systematic increase of SFE with surface density would decrease but not completely remove systematic variations of the IMF.

We also caution against accepting uncritically the conclusion that the final SFE in a real galactic environment will match that predicted by simulations of isolated molecular clouds. These simulations assume that clouds can be disrupted only by internal feedback, and that the external environment plays no role. However, it is far from clear that this is true in realistic galactic environment. For example, Jeffreson & Kruijssen (2018) and Jeffreson et al. (2020) find that

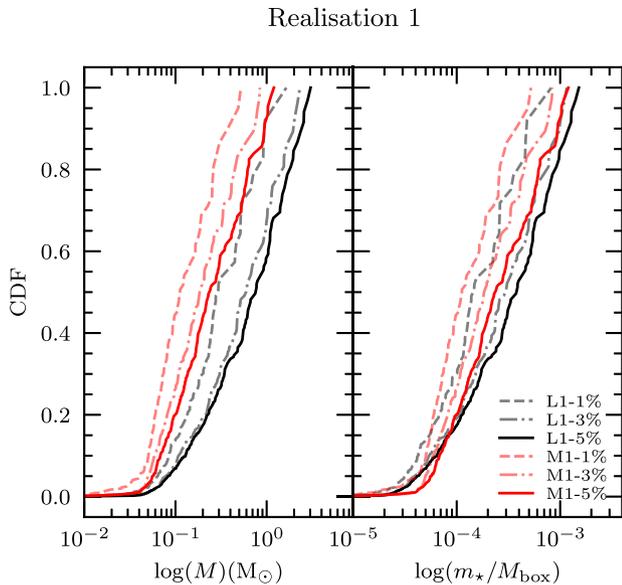


Figure 13. Same as Fig. 6, but showing CDFs of the sink particle masses of runs L1 and M1 at three different SFEs. The red and black solid lines shows the CDF at 5 per cent SFE while the dashed and dash-dotted line represents the CDF at 1 per cent and 3 per cent SFE.

giant molecular cloud lifetimes are mostly set by internal feedback in low-pressure Milky Way outer disc-like environments, but that at higher pressures, as would be relevant for our M and H runs, environmental effects such as galactic shear, external pressure, and ongoing accretion become increasingly important. It is therefore not clear that the final SFE in a realistic environment will follow the isolated-cloud scaling found in Guszejnov et al. (2022).

4.4 Alternative scaling

In our simulations we have kept the dimensionless parameters constant by holding the gas temperature T (or equivalently the sound speed c_s) constant while the box mass and box size vary. In the absence of any kind of feedback mechanism we expect the stellar mass to scale just exactly like the Jeans mass, $M_J \propto \rho^{-1/2} T^{3/2}$, which, given the relationship between density and surface density in our scalings, implies $M_J \propto \Sigma^{-1}$. Radiation feedback flattens this relationship considerably in our simulations, because radiation is trapped more effectively and thus provides stronger feedback at higher Σ . This leads to a characteristic stellar mass that still decreases with Σ , but significantly less strongly than Σ^{-1} .

However, holding T constant while varying the box size and mass is not the only possible way to hold the dimensionless numbers (Mach number, Alfvén Mach number, virial parameter) constant. It is therefore of interest to consider how the stellar mass might vary if instead of having a constant background temperature and sound speed with a variable box mass, we instead varied the sound speed and temperature while keeping a constant mass in the box. Constant box mass implies that the box length L and density ρ vary with surface density as $L \propto \Sigma^{-1/2}$ and $\rho \propto L^{-3} \propto \Sigma^{3/2}$, respectively. Since the Mach number is constant, the velocity dispersion must scale with sound speed as $\sigma_v \propto c_s$, and thus with temperature as $\sigma_v \propto T^{1/2}$. Similarly, since the box mass and virial ratio are constant, equation (4) implies $\sigma_v \propto L^{-1/2} \propto \Sigma$, and therefore we can also deduce that $T \propto \Sigma^{1/2}$. Combining these scalings, since the Jeans mass $M_J \propto T^{3/2} \rho^{-1/2}$, in this scenario in the absence of any feedback

mechanism, M_J will be independent of Σ . We can think of this as being due to a cancellation: a smaller cloud of fixed mass has higher density, which favours fragmentation to smaller masses, but also requires higher velocity dispersion to remain at constant virial ratio, which in turn implies higher temperature if the Mach number is also constant. This higher temperature suppresses fragmentation, exactly cancelling the density effect.

As in our fiducial scaling where we hold T constant, this provides a zeroth-order answer to how we might expect the characteristic stellar mass to vary with the surface density of the environment. The first-order answer must then incorporate the effects of feedback, and in particular how radiative suppression of fragmentation will operate as a function of Σ . The main effect we have identified in this paper will still operate, in that higher Σ clouds will trap radiation more effectively, and thus we expect radiation feedback to be more effective at higher Σ . However, in the alternative scaling scenario where we vary T , there is a countervailing effect that we must consider. In a cloud where the background temperature is lower, we expect that radiation feedback will be more effective at pushing the characteristic mass upwards, since injecting the same amount of thermal energy from starlight will lead to a larger perturbation to the temperature. Conversely, radiation feedback will be less effective when the background temperature is higher, resulting in a lower characteristic mass relative to the box mass despite the jeans mass being higher. Since $T \propto \sqrt{\Sigma}$ in this alternative scaling, we might expect radiation feedback to be *less* effective rather than more effective at higher Σ . It is unclear, and we probably cannot determine absent full simulations, whether this effect is stronger or weaker than the optical depth effect, which points in the opposite direction. Thus it is unclear whether the resulting characteristic stellar mass with feedback would be slightly increasing or slightly decreasing with Σ ; in either case, though, it seems likely to scale with Σ less strongly than for our fiducial case.

We emphasize, however, that while this alternative scaling is useful as a thought exercise, it is unlikely to be astrophysically realized. While there are good physical reasons for the temperatures in star-forming clouds to be only weakly dependent on galactic environment (mostly that CO and dust cooling provide very effective thermostats – e.g. see Sharda & Krumholz 2022), there is no obvious reason why clouds should have fixed Mach numbers. Indeed, observations suggest that they clearly do not, since the observed molecular gas velocity dispersions in some starburst galaxies easily reach ten times those found in the Milky Way (e.g. Scoville et al. 2017), but the gas temperatures are not 100 times larger, as would be required for constant \mathcal{M} . Nonetheless, to the extent that the background temperature does increase with Σ , even if not enough to maintain constant \mathcal{M} , our analysis above suggests that this would somewhat flatten the dependence of stellar mass on Σ compared to our numerical results.

5 CONCLUSIONS

In this paper we present a new set of magnetohydrodynamic simulations including both radiation and protostellar outflow feedback, which we use to explore environmental variations of the IMF. The goal of these simulations is to isolate the interaction of feedback with the star-forming environment, something we achieve by setting up experiments in which we keep the dimensionless parameters constant so that, in the absence of feedback, our simulations representing different star-forming environments would all simply be rescaled versions of one another, and thus would yield statistically identical mass distributions when normalized to the total simulation mass. We

then systematically vary the surface density of the star-forming environment, carrying out two different realizations for each environment of interest to obtain more robust statistics.

We find that as the surface density increases the median stellar mass decreases and the mass range of the IMF becomes narrower. Both the radiation feedback and protostellar outflow feedback contribute to these differences, but in different ways. Radiation feedback suppresses fragmentation to smaller objects, and does so with increasing effectiveness as the surface density rises and radiation is trapped more effectively. This pushes the low end of the IMF upwards in higher surface density environments, though not by enough to fully compensate for the decrease in Jeans mass with increasing density.

However, if this were the only feedback mechanism, we would expect the entire IMF to shift upward due to radiation, which is not what we find. Instead, we find that outflow feedback is dominant at the high mass end of the IMF. When the surface density is low, stars contract further towards the main sequence during star formation, yielding more powerful outflows that can easily escape to low density regions where they are able to expand and occupy more volume. Conversely, this makes them less effective as a feedback mechanism, since it leads to less fragmentation of the dense gas. By contrast, weaker outflows are trapped in dense gas, where they breakup massive clumps more effectively, suppressing massive star formation. The combination of these two effects – more effective outflow suppression of high-mass stars at high surface density, and more effective radiative suppression of low mass stars – together lead to a narrower IMF at higher surface density.

We also explore the implications of our simulations for understanding the mass to light ratio variations observed in ETGs. The IMF variations we find produce shifts in the mass to light ratio of old stellar populations that are qualitatively consistent with those inferred from observations. The shifts are also directionally correct, in the sense that denser star-forming environments, as might be expected to characterize the environments in which the stars of ETGs formed, give rise to stellar populations with higher mass to light ratios once they reach ~ 10 Gyr ages. Thus our results provide a potential explanation for the origin of these variations.

ACKNOWLEDGEMENTS

MRK acknowledges support from the Australian Research Council through its Future Fellowship and Discovery Projects funding schemes, awards FT180100375 and DP190101258. CF acknowledges funding provided by the Australian Research Council (Future Fellowship FT180100495), and the Australia-Germany Joint Research Cooperation Scheme (UA-DAAD). This research was undertaken with the assistance of resources and services from the National Computational Infrastructure (NCI, grants jh2 and ek9), which is supported by the Australian Government, awarded through the National Merit, ANU Merit, and Astronomy Supercomputing Time grant schemes.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Bate M. R., 2009, *MNRAS*, 392, 1363
 Bate M. R., 2012, *MNRAS*, 419, 3115

- Bate M. R., 2019, *MNRAS*, 484, 2341
 Bate M. R., Keto E. R., 2015, *MNRAS*, 449, 2643
 Burgers J. M., 1948, *Adv. Appl. Mech.*, 1, 171
 Cappellari M. et al., 2012, *Nature*, 484, 485
 Chabrier G., 2005, *The Initial Mass Function: From Salpeter 1955 to 2005*, p. 41
 Chakrabarti S., McKee C. F., 2005, *ApJ*, 631, 792
 Choi J., Dotter A., Conroy C., Cantiello M., Paxton B., Johnson B. D., 2016, *ApJ*, 823, 102
 Conroy C., van Dokkum P. G., Villaume A., 2017, *ApJ*, 837, 166
 Cunningham A. J., Klein R. I., Krumholz M. R., McKee C. F., 2011, *ApJ*, 740, 107
 Cunningham A. J., Krumholz M. R., McKee C. F., Klein R. I., 2018, *MNRAS*, 476, 771
 da Silva R. L., Fumagalli M., Krumholz M., 2012, *ApJ*, 745, 145
 El-Badry K., Weisz D. R., Quataert E., 2017, *MNRAS*, 468, 319
 Elmegreen B. G., Klessen R. S., Wilson C. D., 2008, *ApJ*, 681, 365
 Federrath C., Klessen R. S., 2012, *ApJ*, 761, 156
 Federrath C., Klessen R. S., Schmidt W., 2008, *ApJ*, 688, L79
 Federrath C., Roman-Duval J., Klessen R. S., Schmidt W., Mac Low M. M., 2010, *A&A*, 512, A81
 Federrath C., Sur S., Schleicher D. R. G., Banerjee R., Klessen R. S., 2011, *ApJ*, 731, 62
 Federrath C., Schrön M., Banerjee R., Klessen R. S., 2014, *ApJ*, 790, 128
 Federrath C., Krumholz M., Hopkins P. F., 2017, *J. Phys.: Conf. Ser.*, 837, 012007
 Federrath C., Klessen R. S., Iapichino L., Beattie J. R., 2021, *Nature Astron.*, 5, 365
 Geha M. et al., 2013, *ApJ*, 771, 29
 Gennaro M. et al., 2018, *ApJ*, 855, 20
 Goldsmith P. F., 2001, *ApJ*, 557, 736
 Gu M., Greene J. E., Newman A. B., Kreisch C., Quenneville M. E., Ma C.-P., Blakeslee J. P., 2022, *ApJ*, 932, 103
 Guszejnov D., Krumholz M. R., Hopkins P. F., 2016, *MNRAS*, 458, 673
 Guszejnov D., Hopkins P. F., Grudić M. Y., Krumholz M. R., Federrath C., 2018, *MNRAS*, 480, 182
 Guszejnov D., Grudić M. Y., Hopkins P. F., Offner S. S. R., Faucher-Giguère C.-A., 2020, *MNRAS*, 496, 5072
 Guszejnov D., Grudić M. Y., Offner S. S. R., Faucher-Giguère C.-A., Hopkins P. F., Rosen A. L., 2022, *MNRAS*, 515, 4929
 Hansen C. E., Klein R. I., McKee C. F., Fisher R. T., 2012, *ApJ*, 747, 22
 Haugbølle T., Padoan P., Nordlund Å., 2018, *ApJ*, 854, 35
 Hennebelle P., Chabrier G., 2008, *ApJ*, 684, 395
 Hennebelle P., Chabrier G., 2009, *ApJ*, 702, 1428
 Hennebelle P., Lee Y.-N., Chabrier G., 2019, *ApJ*, 883, 140
 Hopkins P. F., 2012, *MNRAS*, 423, 2037
 Hopkins P. F., 2013, *MNRAS*, 430, 1653
 Jappsen A.-K., Klessen R. S., Larson R. B., Li Y., Mac Low M.-M., 2005, *A&A*, 435, 611
 Jeffreson S. M. R., Kruijssen J. M. D., 2018, *MNRAS*, 476, 3688
 Jeffreson S. M. R., Kruijssen J. M. D., Keller B. W., Chevance M., Glover S. C. O., 2020, *MNRAS*, 498, 385
 Klein R. I., Fisher R. T., McKee C. F., Truelove J. K., 1999, in Miyama S. M., Tomisaka K., Hanawa T., eds, *Numerical Astrophysics*. Kluwer Academic, Boston, Mass., p. 131
 Kolmogorov A. N., 1941, *Dokl. Akad. Nauk SSSR*, 32, 16
 Konigl A., Pudritz R. E., 2000, in Mannings V., Boss A. P., Russell S. S., eds, *Protostars and Planets IV*. Univ. Arizona Press, Tucson, p. 759
 Kroupa P., 2001, *MNRAS*, 322, 231
 Krumholz M. R., 2006, *ApJ*, 641, L45
 Krumholz M. R., 2011, *ApJ*, 743, 110
 Krumholz M. R., 2014, *Phys. Rep.*, 539, 49
 Krumholz M. R., McKee C. F., 2008, *Nature*, 451, 1082
 Krumholz M. R., McKee C. F., Klein R. I., 2004, *ApJ*, 611, 399
 Krumholz M. R., Klein R. I., McKee C. F., 2007, *ApJ*, 656, 959
 Krumholz M. R., Cunningham A. J., Klein R. I., McKee C. F., 2010, *ApJ*, 713, 1120
 Krumholz M. R., Klein R. I., McKee C. F., 2012, *ApJ*, 754, 71

Krumholz M. R., Fumagalli M., da Silva R. L., Rendahl T., Parra J., 2015, *MNRAS*, 452, 1447

Krumholz M. R., Myers A. T., Klein R. I., McKee C. F., 2016, *MNRAS*, 460, 3272

La Barbera F., Ferreras I., Vazdekis A., de la Rosa I. G., de Carvalho R. R., Trevisan M., Falcón-Barroso J., Ricciardelli E., 2013, *MNRAS*, 433, 3017

Larson R. B., 2005, *MNRAS*, 359, 211

Lee Y.-N., Hennebelle P., 2018, *A&A*, 611, A89

Lee Y.-N., Offner S. S. R., Hennebelle P., André P., Zinnecker H., Ballesteros-Paredes J., Inutsuka S.-I., Kruijssen J. M. D., 2020, *Space Sci. Rev.*, 216, 70

Lee A. T., Cunningham A. J., McKee C. F., Klein R. I., 2014, *ApJ*, 783, 50

Leitherer C. et al., 1999, *ApJS*, 123, 3

Leroy A. K. et al., 2018, *ApJ*, 869, 126

Li P. et al., 2021, *J. Open Source Softw.*, 6, 3771

Li P. S., Martin D. F., Klein R. I., McKee C. F., 2012, *ApJ*, 745, 139

Low C., Lynden-Bell D., 1976, *MNRAS*, 176, 367

Mac Low M.-M., 1999, *ApJ*, 524, 169

Masunaga H., Miyama S. M., Inutsuka S.-I., 1998, *ApJ*, 495, 346

Mathew S. S., Federrath C., 2020, *MNRAS*, 496, 5201

Mathew S. S., Federrath C., 2021, *MNRAS*, 507, 2448

Miller G. E., Scalo J. M., 1979, *ApJS*, 41, 513

Myers A. T., Krumholz M. R., Klein R. I., McKee C. F., 2011, *ApJ*, 735, 49

Myers A. T., Klein R. I., Krumholz M. R., McKee C. F., 2014, *MNRAS*, 439, 3420

Nam D. G., Federrath C., Krumholz M. R., 2021, *MNRAS*, 503, 1138

Newman A. B., Smith R. J., Conroy C., Villaume A., van Dokkum P., 2017, *ApJ*, 845, 157

Offner S. S. R., Klein R. I., McKee C. F., Krumholz M. R., 2009, *ApJ*, 703, 131

Offner S. S. R., Clark P. C., Hennebelle P., Bastian N., Bate M. R., Hopkins P. F., Moraux E., Whitworth A. P., 2014, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, *Protostars and Planets VI*. Univ. Arizona Press, Tucson, p. 53

Oldham L., Auger M., 2018, *MNRAS*, 474, 4169

Padoan P., Nordlund A., Jones B. J. T., 1997, *MNRAS*, 288, 145

Padoan P., Nordlund Å., 2002, *ApJ*, 576, 870

Press W. H., Schechter P., 1974, *ApJ*, 187, 425

Rees M. J., 1976, *MNRAS*, 176, 483

Richer J. S., Shepherd D. S., Cabrit S., Bachiller R., Churchwell E., 2000, in Mannings V., Boss A. P., Russell S. S., eds, *Protostars and Planets IV*. Univ. Arizona Press, Tucson, p. 867

Salpeter E. E., 1955, *ApJ*, 121, 161

Scoville N. et al., 2017, *ApJ*, 836, 66

Semenov D., Henning T., Helling C., Ilgner M., Sedlmayr E., 2003, *A&A*, 410, 611

Sharda P., Krumholz M. R., 2022, *MNRAS*, 509, 1959

Shu F. H., Li Z.-Y., Allen A., 2004, *ApJ*, 601, 930

Smith R. J., 2014, *MNRAS*, 443, L69

Smith R. J., 2020, *ARA&A*, 58, 577

Spaans M., Silk J., 2000, *ApJ*, 538, 115

Spiniello C., Trager S. C., Koopmans L. V. E., Chen Y. P., 2012, *ApJ*, 753, L32

Spiniello C., Koopmans L. V. E., Trager S. C., Barnabè M., Treu T., Czoske O., Vegetti S., Bolton A., 2015, *MNRAS*, 452, 2434

Treu T., Auger M. W., Koopmans L. V. E., Gavazzi R., Marshall P. J., Bolton A. S., 2010, *ApJ*, 709, 1195

Truelove J. K., Klein R. I., McKee C. F., Holliman J. H. I., Howell L. H., Greenough J. A., 1997, *ApJ*, 489, L179

Truelove J. K., Klein R. I., McKee C. F., Holliman J. H. I., Howell L. H., Greenough J. A., Woods D. T., 1998, *ApJ*, 495, 821

Urquhart J. S. et al., 2018, *MNRAS*, 473, 1059

van Dokkum P. G., Conroy C., 2010, *Nature*, 468, 940

APPENDIX A: MASS TO LIGHT RATIOS WITH FINITE MASSES

Our simulations differ from each other in the total mass of stars they produce, and the finite masses found in them also differ from the infinitely-sampled limit we use to calculate our fiducial mass to light ratios for the Salpeter (1955) and Chabrier (2005) IMFs. It is therefore important to verify that the differences in mass to light ratio that we find are not due to the effects of sampling different total stellar masses in each case. To carry out this check, we repeat the SLUG calculations described in Section 4.2 to create synthetic clusters with total masses 100, 50, and 25 M_{\odot} , corresponding to the total stellar masses in L, M, and H runs. For this experiment we draw stars from a Chabrier (2005) IMF truncated at 0.01 M_{\odot} at the low mass end and 120 M_{\odot} at the high mass end. We generate 50 000 clusters for each of our target masses, and for each run we compute M/L as a function of age from 5–10 Gyr, exactly as in Section 4.2.

Fig. A1 shows the mean mass to light ratios of these finite-mass experiments, in comparison to the infinitely-sampled limit. We see that, not surprisingly, the effects are largest for run H, which has the smallest total mass, but that the effects overall are never large. The maximum difference between the finite-mass and infinitely-sampled results never exceeds ≈ 0.05 dex. Comparison to Fig. 12 shows that such differences are much smaller than the differences we find between the different runs, and between them and the Chabrier (2005) IMF.

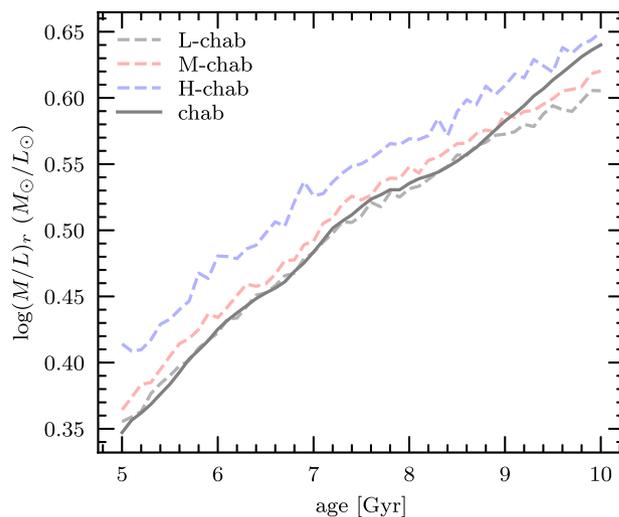


Figure A1. The solid black line shows the mass to light ratio of an infinitely-massive simple stellar population with a Chabrier (2005) IMF as a function of age. The black, red, and blue dashed lines show the mean mass to light ratios of stellar populations drawn from the same IMF with total masses of 100, 50, and 25 M_{\odot} , corresponding to the total stellar masses in runs L, M, and H.

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