Draft version December 6, 2016
Preprint typeset using LaTeX style emulateapj v. 12/16/11

REPRODUCING SUPERNOVA STATISTICS IN SIMULATIONS OF GALAXY EVOLUTION

Alex Gurvich[1]
[1] Department of Physics and Astronomy and CIERA, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

(Dated: December 5, 2016)
Draft version December 6, 2016

ABSTRACT

We consider a parameterized pseudo-random model for seeding supernovae in a galactic disk in order to reproduce the clusering statistics relevant to the non-linear production of momentum in the Sedov-Taylor phase of supernova remnant evolution. We believe that a non-linear amplification of momentum related to the overlapping of multiple supernova remnants, so called "superbubbles" can explain the observed sensitivity of the mass loading factor, η, on the fiducial random seeding scheme in previous studies of galaxy evolution. We find that it is indeed possible to reproduce the emergent clustering statistics of the self-consistent supernova seeding scheme in the FIRE simulations with a simpler pseudo-random model enabling future studies of ultra-high resolution to study the effects of stellar feedback without explicitly modeling the large scale stellar evolution processes. However, there are indications that matching the morphology of the self-consistently modeled superbubbles using this first order scheme for synthetically generated superbubbles is not automatically satisfied, leaving room for improvement in order to fully capture the details of the clustering statistics in the FIRE simulation.

1. INTRODUCTION

One of the most important questions in studies of galaxy evolution is the link between the energy and momentum injected by supernovae and the suppression of new star formation by expelling potentially star forming gas. The exact mechanism by which this process is regulated is still hotly debated (Faucher-Giguère *et al.* 2013; Krumholz and Burkhart 2016). However, previous studies on the subject have shown that the rate of expelled gas sensitively depends on the way that supernovae are placed, or seeded, in the galactic disk (Martizzi *et al.* 2016)– varying by an order of magnitude when comparing the simplest assumptions. However, both seeding models compared in Martizzi *et al.* (2016) result in an under-estimate for the value of the mass loading factor, η, when compared to the more physically motivated, and self-consistent, FIRE simulations (Muratov *et al.* 2015). The FIRE simulations represent a suite of of high-resolution "zoom in" simulations with initial conditions for simulated galaxy "halos" that are generated from a low resolution simulation of a cosmological sample. However, because of computational limitations feedback processes are often estimated using "sub-grid models" that encapsulate the physics of the process without explicitly resolving it. These sub-grid models are typically derived either from analytic arguments or from higher resolution, and smaller scale, simulations.

In order to address the question, we consider one of these higher resolution simulations of an individual galaxy and explicitly track the location and time of individual supernova events. By comparing the spatial and temporal correlations of this "self-consistent" sample with those of random models parameterized we can characterize how "un-random" the self-consistent case is. Specifically, we develop a recursive friends-of-friends algorithm that that uses a dynamic linking length. That length is determined by the radius to which the supernova grows before mixing with the surrounding interstellar medium, which is a function of the ambient gas density at the location of the supernova following the prescription for the "cooling radius," R_{cool}, in an inhomogenous interstellar medium from Martizzi *et al.* (2015). We develop a fully automatic and self-consistent framework with which to consider an arbitrary PDF for the "clustering strength," a measure of the number of superbubbles produced of a given size. From this PDF one can generate clusters of a prescribed internal structure according to a second arbitrary PDF for the fiducial linking radii of the synthetic cluster– not to be confused with the cooling radius, which is a physical quantity that is a function of the ambient density wherein the fiducial linking radius is an arbitrary statistic used here only to generate the synthetic clusters.

The structure of this paper is as follows: in Section 2 we outline our methods for parameterization, in Section 3 we consider the effectiveness and flexibility of this parameterization, in Section 4 we present our discussion of these results, and in Section 5 we give our conclusions and discuss the next steps in this work.

2. METHODOLOGY

2.1. *Measuring the Number and Size of Superbubbles*

We define the "clustering strength" as the functional form of the of the relationship between the size of a "cluster" of supernovae and the number of clusters of that size– where a cluster is defined as a group of supernovae who are members of a friends of friends group. Supernovae are determined to be members of a friends of friends group if and only if the epicenter of another member of the group is within their cooling radius (or if a supernova's epicenter is contained within a member of the group's cooling radius). An important caveat to mention as well, is that a supernova is only considered to be "within" the cooling radius of another remnant if

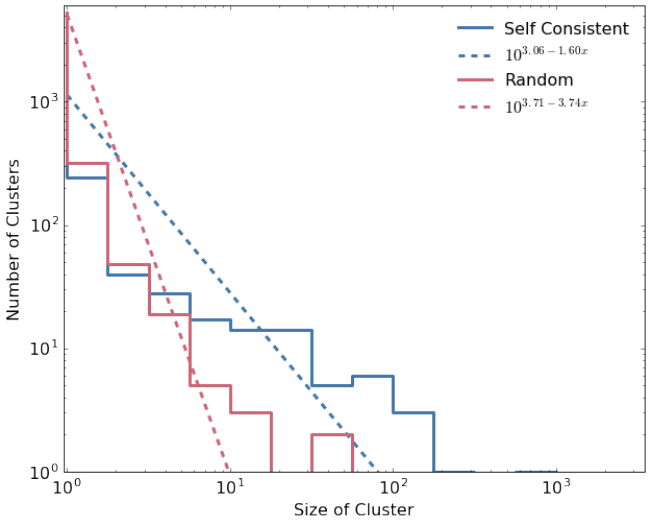

FIG. 1.— The clustering strength as realized when using the self consistent model that is also used in the FIRE simulations compared to that when supernovae are seeded purely randomly. Here, a best fit power law is overplotted for each as determined by a linear regression with uncertainties estimated as $\sqrt{N}$ (here suppressed from the plot). This crude estimate is a first order approximation to the clustering strength.

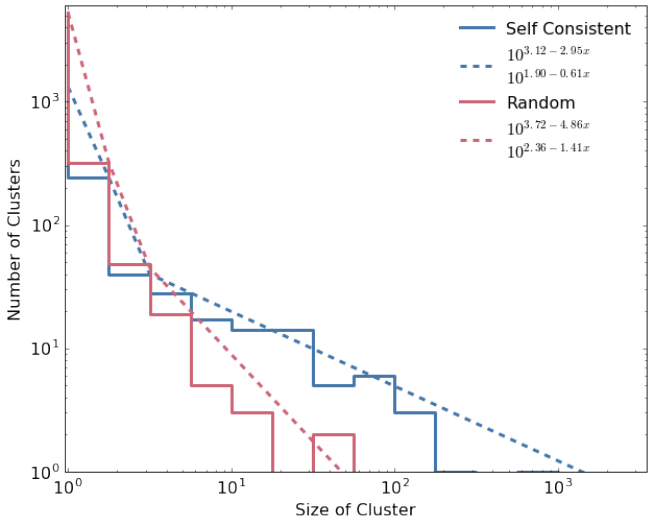

FIG. 2.— The clustering strength as realized when using the self consistent model that is also used in the FIRE simulations compared to that when supernovae are seeded purely randomly. Here, a best fit broken power law is overplotted for each as determined by a linear regression with uncertainties estimated as $\sqrt{N}$ (here suppressed from the plot). The position of the joint is determined by considering the set of broken power laws that correspond to a linear regression using points on either side of the joint for any joint position. The one with the smallest χ^2 is chosen as the "best fit." Compared to Figure 1 this gives a much better description of the clustering strength.

and only if the two also overlap in time– only supernovae and remnants that coexist at the same time are also considered to be overlapping if one's remnant contains the other.

Cooling radii and times (the lifetime of the remnant) are determined by measuring the local gas density at the epicenter of the supernova following the prescription for an inhomogenous interstellar medium from Martizzi *et al.* (2015). We believe that by correctly capturing the number of superbubbles of a given size it is likely that the amount of momentum generated, and thereby the amount of mass ejected from the galaxy, will be captured as well. Therefore, it is imperative to the science to, at the very least, correctly reproduce the clustering strength realized in the simulation.

Figures 1 and 2 compare the clustering strength realized in the simulation in the "Self Consistent" case and the case where supernovae are seeded totally randomly. Figure 1 fits a single power law over the entire range of cluster size which is an alright approximation for the totally random case in red but a pretty obviously bad one for the self consistent case plotted in blue. A better description is given when you allow for a second power law for "large" cluster sizes, here this is determined to be above N=4. We determine this by minimizing the total χ^2 as a function of the position of the joint, using independent linear regression fits for the points on either side. Broken power laws are thus computed for every possible joint position, excluding the endpoints for obvious reasons, and the one that results in the smallest value of χ^2 is chosen as the best fit. In both cases, uncertainties are estimated as $\sqrt{N}$ where N is the number of clusters in each bin. This is a gross approximation that serves only to apply some level of weight to the early as their contribution is otherwise washed out when weighted equally with the more numerous later points.

2.2. *Seeding and Populating Synthetic Clusters*

We seed supernovae by following a relatively simple prescription for generating ad hoc synthetic clusters that are placed into random locations in the simulation. Cluster "heads" are placed randomly in space within some slab representing the area where the bulk of the galactic disk is. All supernovae are assigned a launch time uniformly, though in principle this too is a PDF that is a function of the shape of the star formation rate of that galaxy.

After placing the cluster head, subsequent members are added iteratively by randomly selecting a member of the cluster and placing the next member at a random location within a fiducial radius that is specified with an arbitrary PDF, here we use a single Gaussian distribution but in principle one could let this PDF become a (many) degree of freedom(s), until the cluster reaches its predetermined size. The number of synthetic clusters that are produced of a given size is determined by a chosen PDF for the clustering strength, which is taken as an input. This process, hereafter referred to as "running an input PDF" is laid out as:

1. Choose an input PDF for the clustering strength.
2. Generate a uniformly distributed random number between [0,1].
3. use the CDF of your input PDF to determine what cluster size that random number corresponds to.
4. Place the head of the cluster at a random location within a 3D slab encompassing the bulk of the galactic disk.
5. Choose a random member of the cluster (the first time you reach this step the only option is the head).

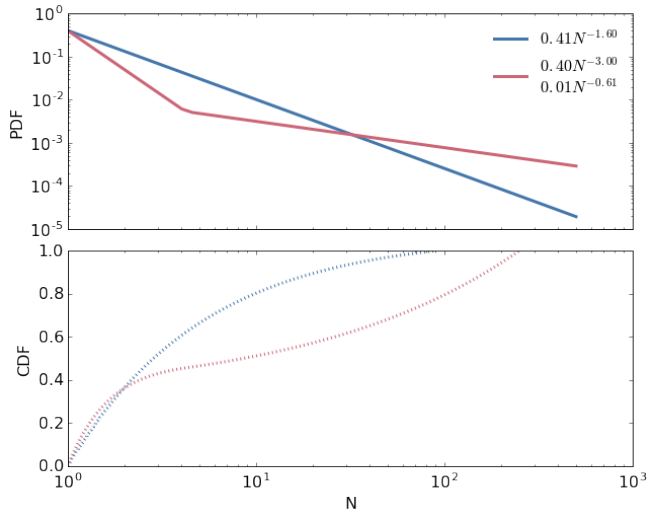

FIG. 3.— Here we show two sample input PDFs for the clustering strength used to generate synthetic supernova clusters. In the top panel, a single power law is plotted in blue while a broken power law is plotted in red. The same color scheme is used in the bottom panel, where the CDF, approximated from numerical integration of the curves in the top panel, is plotted as well. The points in the bottom panel are those used to draw points from the PDF when numerically inverted for uniformly distributed values between [0,1].

6. Generate another uniformly distributed random number between [0,1].

7. Use the CDF of your chosen fiducial radial PDF to determine how far away you should place the new member.

8. Randomly place the new member within that radius.

9. Return to Step 5 until the total size of the cluster reaches the size determined from Step 3.

10. Return to Step 2 until the total number of supernovae, summed over all the synthetic clusters, is above some predetermined value. Here we use the total number of supernovae launched within an average 1Myr time interval in the self consistent case.

In Figure 3 we show two sample input PDFs for the clustering strength. The bottom panel shows the CDF, which we numerically integrate. Points on the CDF are successively approximated as points are drawn from the distribution, increasing the resolution of the numerical inversion as time goes on. This way, the algorithm and code structure we use can be applied to any arbitrary PDF, regardless of its invertibility.

From here, the synthetic supernovae are treated identically to the randomly and self consistently generated supernovae and passed into the same friends of friends grouping algorithm. Any information about fiducial radii or synthetic cluster identity is lost as new linking radii, corresponding to the ambient density around the synthetically generated supernovae, are assigned and cluster identity is determined as described in the beginning of Section 2.

3. RESULTS

After running a number of input PDFs (see steps in Section 2) a best fit model can be computed. In principle, this can be automated and input arbitrary PDFs can be generated iteratively after comparison utilizing a Monte Carlo evolution algorithm though (unfortunately) this is outside the scope of this work.

Instead, we run a comprehensive set of power law and broken power law PDF tuning their various parameters until the measured clustering strength independently match that of their self consistent counterparts. Here we report the results of the best fit power law and broken power law input PDF, the specific coefficients of which are found in Figure 3, and the resulting clustering strength in Figures 5 and 6, which was used to determine the fit, and the emergent distribution of cooling radii, shown in Figure 7, as a function of cluster size, a rough descriptor of the morphology of the supperbubbles as a function of their population.

3.1. *Measuring the Clustering Strength*

After running a comprehensive series of input PDFs for both power law and broken power law forms we find the best fit to the self consistent clustering strength for both cases respectively. We note however that while these clustering strength relationships agree fairly well that does not necessitate that the underlying distribution is recovered.

This is apparent in Figure 5 where the two best fits agree fairly well but the synthetic population has signifcantly more smaller clusters than the self consistent case and consequently fewer larger clusters to make up for it. This has the combined effect of washing out any cohesive correlation that might be visible when the populations are visualized, as in Figure 4. In the end, this results in a distribution that is qualitatively more similar to the completely random case, even if there is some coherent structure hidden in the quantitative details.

On the other hand, the broken power law formulation for the input PDF results in a starkly better description of the clustering power, shown in Figure 6. Though there is a vertical slight offset in the second half of the best fit broken power law with respect to the self consistent case this is consistent with the fluctuations at that low number level of cluster size. A single large cluster could push the fit up marginally without changing its overall shape, which is what we see here. Comparing the distributions themselves, they are remarkably similar over the entire range of cluster size N. Additionally, the high level of correlation in the visualization of the broken power law input PDF in Figure 4 is qualitatively similar to the self consistent case and markedly different from the random case, obviously lacking the same inherent spiral structure.

3.2. *Measuring the Distribution of Cooling Radii*

In addition to a simple census census of superbubbles as described in the previous sections in detail, the internal structure and morphology of the cluster can play an important role in the amount of momentum generated by the cluster and consequently the mass ejected from the disk. The consequences for a difference in morphology are discussed in Section 4 in more detail but it suffices to say that a cluster whose supernovae primarily have small R_{cool}, representing a long chain of small remnants, is qualitatively different than a cluster made up of

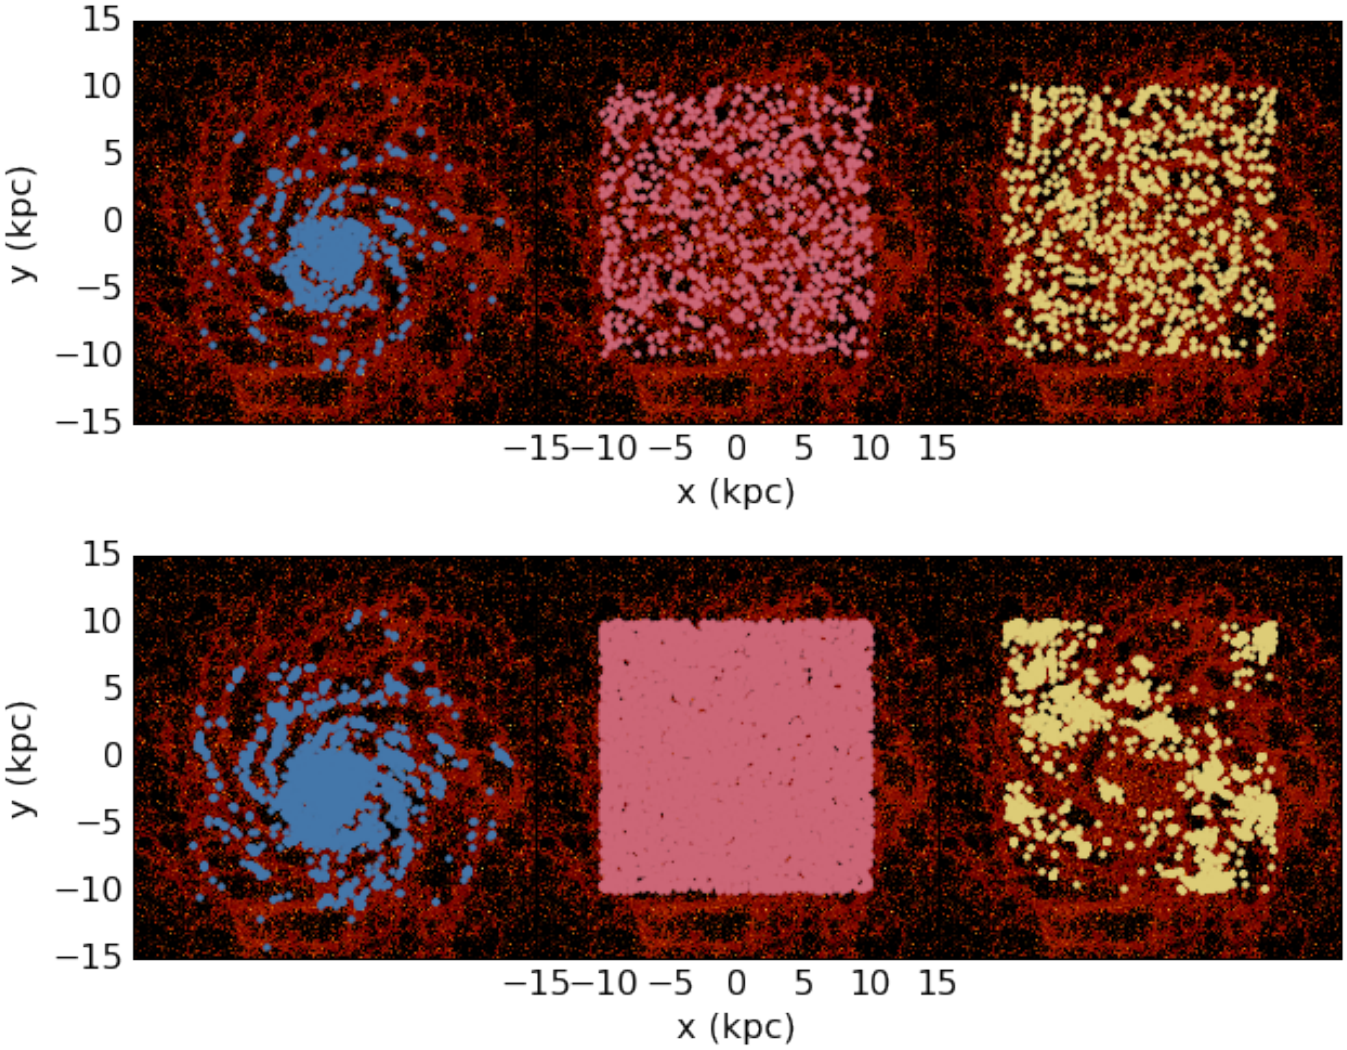

FIG. 4.— Thumbnails of the volume weighted gas density distribution with decimated supernova populations overlaid in different colors. The leftmost panel shows the self-consistent population in blue which clearly follows the spiral structure and morphology of the galaxy. In the middle panel the totally random population is plotted in red, showing no structure or correspondence to the underlying gas density distribution. The rightmost panel shows the synthetic population in yellow. The top row is decimated, plotting only $\frac{1}{8}$th of the points while the bottom is not decimated at all. The synthetic panel of the top row is drawn from the best fit power law input PDF while the bottom row is drawn from the best fit broken power law PDF. Allowing the joint in the input PDF results in significantly more correlation and structure. Decimation is performed on the top row because otherwise the middle and right panels would be indistinguishable (re: both are a solid block). This is not necessary for the bottom row as the structure is emergent from the input PDF.

a few supernovae with large R_{cool} encompassing a large number of smaller remnants. These differences would be apparent on a plot of the distribution of R_{cool} versus cluster size. This is shown in Figure 7, where each panel represents a bin in cluster size, whose edges are denoted in that panel's subtitle. All supernovae that belong to a cluster of size N that fall within that panel's bin are plotted, regardless of if they belong to the same cluster or not. This has the effect of washing out a lot of the effects morphology might have on a plot like this by average over clusters of different morphology. Regardless, if a specific morphology dominates clusters of a given size this should still be apparent through the apparent dilution with other sorts of morphologies.

Qualitatively speaking, the synthetic broken power law PDF input, plotted in green, matches the self consistent case (blue) much better than the synthetic power law PDF input (yellow). This is most apparent in the [3,6),[6,10),[100,178), and [178,316) panels. In the latter two, there arren't actually any synthetic power law input PDF clusters of that size, which is a result of the analysis presented in the previous subsection. However, a stark contrast between the green and blue curves presents itself in the final two panels, where even the broken power law input PDF doesn't produce a cluster of $N > 1000$. Most interestingly however is the difference in the R_{cool} PDFs in the [316,562) panel where the morphologies of the two are clearly different. Where the blue curve extends to cooling radii of ≈ 1 kpc and then down to much smaller cooling radii of tens of pc the green curve is comprised of entirely cooling radii of ≈ 100 pc. The former suggests a single (or few) large supernovae encompassing an area cleared away and disturbed by many more much smaller supernovae where the latter describes a scenario where a chain of intermediate supernovae all go off in a similarly dense medium. It becomes necessary at this

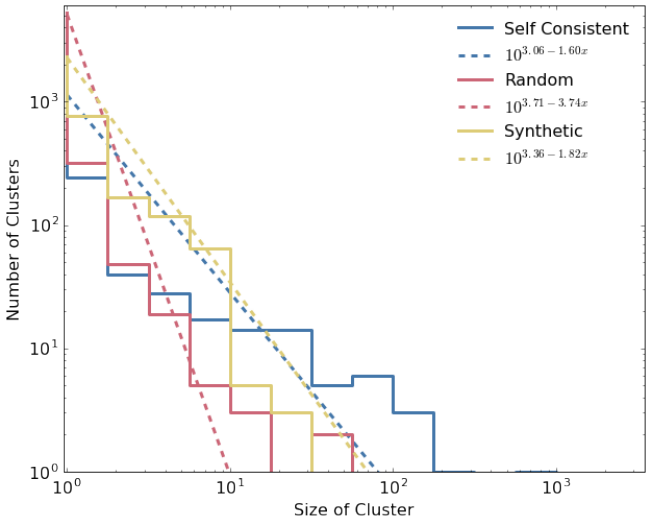

FIG. 5.— Analogous to Figure 1 this shows the measured clustering strength with the best fit power law (dashed) derived from a linear regression for the self consistent (blue), random (red), and synthetic using a power law input PDF (yellow) supernova populations. While the self consistent and synthetic best fit powerlaws may be similar, the distributions themselves look qualitatively different indicating that the rough approximation of a single power law for an input PDF is not enough.

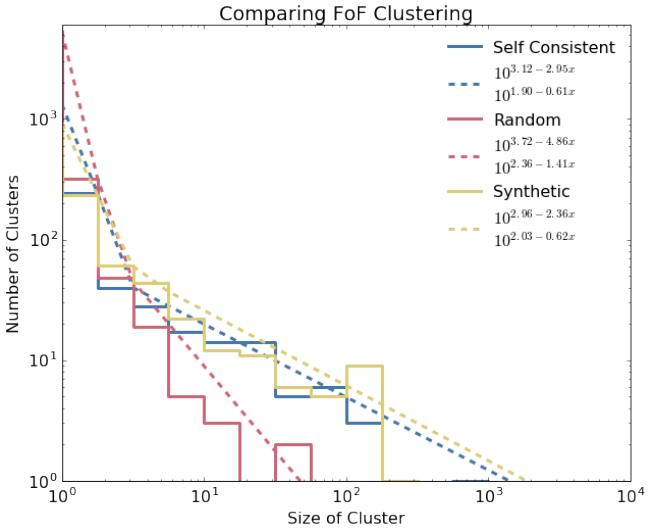

FIG. 6.— Analogous to Figure 2 this shows the measured clustering strength with the best fit broken power law (dashed) derived from a linear regression for the self consistent (blue), random (red), and synthetic using a broken power law input PDF (yellow) supernova populations. Here, a small offset in the second part of the best fit broken power law is well within the uncertainty (here suppressed, but crudely estimated as $\sqrt{N}$, in that regime. The distributions themselves look very similar throughout the entire range of cluster size N. This shows marked improvement over Figure 5, giving good support for a broken power law input PDF.

point to remind that plots like these necessarily average over time, as well as in this case space when considering different clusters of identical length. The implications of this important qualitative difference are discussed in the next section.

4. DISCUSSION

Looking at the details of the differences between the power law and broken power law input PDFs and why they result in categorically different populations of supernovae is relatively straight forward. Figure 3 shows the best fit power law and broken power law input PDFs and their corresponding CDFs. Of note, though the input PDFs are very different, in the regime before the joint in the broken PDF the integral is very similar showing that for small clusters a power law is not a bad approximation. However, the power law severely underestimates the number of large clusters as its CDF quickly rises to 1. We know that the large clusters represent the very massive superbubbles we believe are important to generating the large amounts of momentum required to expel mass at the observed rate of galactic outflows. Therefore a power law PDF by nature will suppress the number of superbubbles created, and the ensuing mass loading factor.

Figure 4 shows that the macroscopic orientation of the supernova clusters is very different in the self consistent and broken power law input PDF cases. However, it is not obviously clear that the spiral structure of the supernova clusters, mimicking the gas density distribution, is an important factor in the amount of mass ejected from the galaxy disk. Supernovae that are launched in a relatively evacuated medium quickly sweep up the low density mass without much change in energy, generating a lot of momentum in the process, see the Sedov-Taylor phase of supernova remnant expansion. Therefore, though the self consistent case results in supernovae that follow the overall spiral structure of the galactic disk it may be true that supernovae launched in regions outside the spiral arms of the disk expand quickly without sweeping up much gas or losing much energy, hitting the disk in the same energy-momentum phase space that it would've been launched in anyway if it were launched within the disk itself. The effect of this difference could be measured by running additional simulations using this new pseudo-random seeding scheme and looking at the rate at which gas is expelled in both the self consistent and pseudo-random cases.

To consider the difference that the relatively "small-scale" cluster morphology might have on the momentum generated take for example, a superbubble made up of a long chain of small supernovae is qualitatively different than having a similar superbubble constituted of the same number of supernovae but instead with one single massive remnant that encompasses many smaller remnants in a more spherical geometry. Naively one could imagine that the chain of small remnants wouldn't form a cohesive single bubble the pushes outward from the disk while the latter case presents as many single bubbles clearing out the area for a single massive bubble to sweep up the disturbed remains without hitting any particularly dense patches which would slow it down. This is relevant because we see in Figure 7 that the largest clusters produced in the self consistent case are made up one relatively few very large supernovae with many more smaller supernovae filling the bulk of the rest of the population of the cluster. We know this is true because at this size of cluster there are few if even more than a single cluster of this size so averaging effects across different clusters are minimal. Of course one could in principle make a similar plot where each panel was its own cluster, but 12 panels proved large enough for our purposes, 1000 would necessitate its own paper– or textbook.

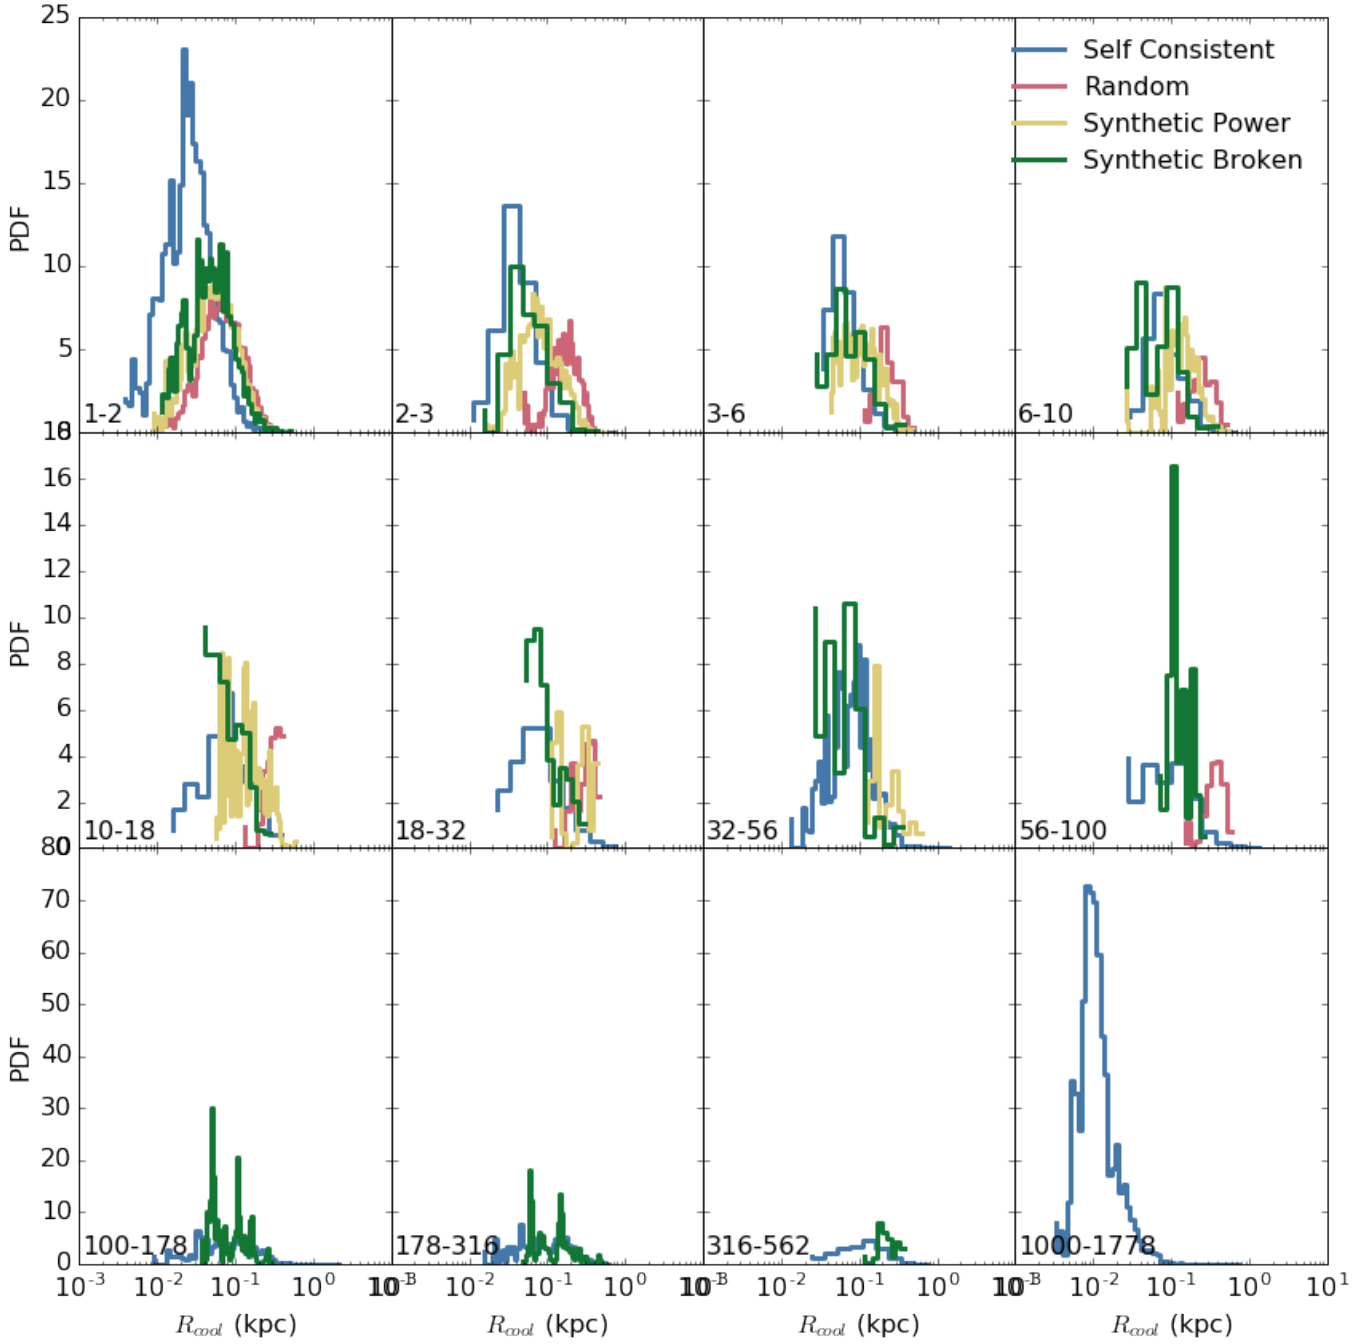

FIG. 7.— Here we show the normalized PDF of cooling radius for each of the self consistent, random, and synthetic cases (both the power law and broken power law input PDF cases) dividing the supernova population into bins based on their cluster identity. Each panel is plotted using only supernovae who belong to a cluster of a given size, noted in the subtitle of the each panel with the right edge excluded. For example, the top left panel is made up of supernovae that are entirely isolated and not a member of any cluster while the bottom right panel are the supernovae of the largest cluster formed in the self consistent case (it turns out there is only one of that size). Note that while the x axes are shared for every panel the y axes are only shared by row in order to show the small amplitude of the intermediate sized clusters in the middle row properly.Of note is the markedly better agreement of the blue curve with the green curve when compared to the yellow curve across all panels. The synthetic power law input PDF doesn't actually produce a cluster larger than 100 members, on that basis alone arguing for their agreement. Of course, there is room for improvement and this is discussed in section 4.

5. CONCLUSION

We conclude that overall, based on the agreement in Figure 6 of the synthetic and self consistent distributions alone, a more correct description of the clustering strength comes from considering a broken power law input PDF. This statement is evidenced as well by the notion that the self consistent case is better described by a best fit broken power law, shown in Figure 2. Additionally, the comparison between the top and bottom right-most panels in the visualization in Figure 4 makes a clear qualitative demarcation between the two in favor of the broken power law as well. On top of that, the morphology of the clusters, as roughly described in Figure 7, is more accurately captured by the broken power law input PDF.

We note that the specific values of the broken power law input PDF are a function of the galaxies macroscopic properties and could potentially be prescribed by a star formation history of the galaxy though future work in that direction would need to be done to say for certain. Further, though the qualitative agreement of the cluster morphologies is *better* for the broken power law case there is still a significant difference in a way that could potentially dictate whether the "correct" amount of momentum would actually be generated for a self consistently modeled cluster when approximated by a pseudo-random cluster of equivalent size. There is also the point that the pseudo-random case is blind to the underlying gas density distribution, it is not obvious that this is necessarily a problem, but future work might consider launching the head of the cluster given a PDF consistent with the underlying gas density distribution in order to capture the actual positions of the cluster as well– however this could quickly become very complex as larger clusters may categorically appear in different areas of the galaxy when compared to smaller clusters, so a PDF for the cluster head might itself be a function of the size of the cluster.

There are also more complicated mathematical frameworks that might serve to more consistently model the self consistent launching of supernovae pseudo- randomly. Approximating the self consistent population as a realization of a 4-D Gaussian field in space and time could allow you to calculate the correlation statistics in both time and space, e.g. by using a two point correlation function, and create a new model that exhibits the same clustering statistics. You could imagine further tuning this model to more heavily weight the underlying gas density distribution in order to doubly handle the questions of the correlation of clusters with the galaxy structure and their morphologies simultaneously in a more rigorous mathematical framework.

All of this and more can be quantified and addressed with further simulations of isolated galaxies utilizing different supernova launching schemes in order to isolate the most fundamental piece of the physics that is causing the enhanced momentum to be generated. Whether that is the distribution of cluster sizes, their morphologies, or their locations within the context of the galaxy– or some combination therein, remains to be seen.

REFERENCES

C.-A. Faucher-Giguère, E. Quataert, and P. F. Hopkins, MNRAS **433**, 1970 (2013), arXiv:1301.3905.

M. R. Krumholz and B. Burkhart, MNRAS **458**, 1671 (2016), arXiv:1512.03439.

D. Martizzi, D. Fielding, C.-A. Faucher-Giguère, and E. Quataert, MNRAS **459**, 2311 (2016), arXiv:1601.03399.

A. L. Muratov, D. Kereš, C.-A. Faucher-Giguère, P. F. Hopkins, E. Quataert, and N. Murray, MNRAS **454**, 2691 (2015), arXiv:1501.03155.

D. Martizzi, C.-A. Faucher-Giguère, and E. Quataert, MNRAS **450**, 504 (2015), arXiv:1409.4425