

**TURBULENCE, INJECTION, AND BUOYANCY: WHEN AND WHERE
COSMIC RAYS MATTER**

by

Roark S. Habegger

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

Doctor of Philosophy

(Astronomy)

at the

UNIVERSITY OF WISCONSIN–MADISON

2026

Date of final oral examination: 04/17/2026

The dissertation is approved by the following members of the Final Oral Committee:

Dr. Ellen G. Zweibel, Professor, Astronomy & Physics

Dr. Vladimir Zhdankin, Assistant Professor, Physics

Dr. Snezana Stanimirovic, Professor, Astronomy

Dr. Sebastian Heinz, Professor, Astronomy

Dr. Mateusz Ruszkowski, Professor, Astronomy (University of Michigan)

© Copyright by Roark S. Habegger 2026
All Rights Reserved

*Dedicated to my Mom and Dad,
For teaching me that I can do anything I set my mind to.*

ACKNOWLEDGMENTS

While this work bears my name alone, it could never have been completed without the support, guidance, and advice others provided. I have a debt to many people who made this work possible. I would also like to acknowledge funding from NSF and NASA, as well as from UW-Madison. I am extremely thankful for the opportunity to pursue my research, and it would not have been possible without the monetary and computational resources from those organizations.

First and foremost, I would like to thank my advisor, Ellen, for the years of guidance, attention, and engaging discussion. I balk at the thought of trying to complete this Ph.D. without such a supportive and kind advisor. The countless hours I have spent in your office, in group meetings, and just talking about science with you have encouraged my curiosity and given me an endless supply of research ideas. Beyond that, you gave me the freedom and resources to develop my professional network, to build collaborations outside of Madison. I know I took us off the rails on tangential ideas many times, but I am so thankful that you always encouraged me to think deeper and keep asking how this nonlinear, complex universe works. You give me hope and you have inspired me to be the best scientist, mentor, and teacher I can be.

I also want to thank my other primary mentors in my scientific career so far: Mateusz Ruzkowski and Fabian Heitsch. Fabian, I am thankful for the care and time you put into teaching me computational methods and introducing me to theoretical research. I use the tools and lenses of thought that you taught me every day to separate the physics from the numerics. Mateusz, I am thankful for the many hours you have given me on zoom to discuss cosmic rays and their interaction with turbulence and clouds, picking out the physics piece-by-piece. Also, for reminding me that just getting our thoughts straight is a productive activity. Knowing the story is just as important as writing it. I hope these past two years are just the beginning of a long and fruitful collaboration.

In addition to Ellen and Mateusz, I would like to thank my other thesis committee members: Snez, Sebastian, and Vladimir. Thank you for taking the time to read this thesis and offer me feedback on it.

I would also like to thank all the other professors who have taught me (even outside of formal courses) for being amazing teachers and giving me the background necessary to be an astrophysicist. It was a lucky opportunity to learn from you, and I am so grateful for the breadth of knowledge I gained as a result.

To my friends and colleagues, I am so thankful to have had your insight during

my graduate career. The number of people who have had an impact on me is large, and I cannot thank you all enough, here. From office and hallway conversations with Mohan, Hanjue, Isaac, Zach, Francisco, Bindesh, and Evan, to running with Aaron, climbing with Max, discussing AI with Nick, and statistics with Michael, you all have helped me make breakthroughs in my research and keep my head up while I struggled through paper writing and job applications. Similarly, I would like to thank the undergrads I have had the opportunity to mentor, Sherry and Natalie. Your curiosity has kept me on my toes and reminded me how much fun it is to learn (and teach) something new.

Ka Ho and Ka Wai: Thank you both for the many exciting late-night discussions about turbulence and cosmic rays, which led to our paper together. You two showed me that doing good science can be downright fun and exciting, especially when done with friends.

Sophie: Thank you for your friendship and for getting me through a challenging time in my Ph.D. — a moment when I felt like I had no good ideas and no path forward. You reminded me that it was fun to invent and work through new ideas, even when it turns out someone already thought of them fifty years ago and put them in the appendix of a random textbook.

Patrick: Thank you for being my best friend. I would have lost my sanity without our canoe trips, sharing good beer, and all the nonsense we got up to trying to fix our cars and espresso machines. Thanks for encouraging me to get out, be social, and have fun throughout my time in graduate school

My brother, Luke: All the rounds of school I have done without you around have been difficult. Calling every week has definitely made this time easier. Some days, I wish we could go back to running around the woods, building Legos, and romping around the house. I am just thankful that we still get to make each other laugh while we find our way through work and life.

My Mom & Dad: Thank you for always being there for me. I will always remember the letter you wrote me when I was a kid, saying that I should follow my curiosity and do whatever work I find joy in. I am so lucky to have your support, pushing me to trust that my hard work would pay off, while also making sure I enjoy life. Everything I create during my life, I owe to your belief in me.

My love, Caitlin: I would not have been able to do this without your love, your intelligence, and your laughter over these last two years of my Ph.D. You bring joy to my life every day, keeping me sane during the difficult points and multiplying my excitement during the good parts.

CONTENTS

List of Figures	vi
List of Tables	viii
1 Introduction	1
1.1 <i>A Brief History of Cosmic-Ray Astronomy</i>	4
1.2 <i>Cosmic-Ray Feedback</i>	9
2 Theoretical Background	15
2.1 <i>Cosmic Rays as Particles</i>	16
2.2 <i>Dynamics of the Cosmic-Ray Distribution Function</i>	24
2.3 <i>Cosmic-Ray Magnetohydrodynamics</i>	27
2.4 <i>Instabilities and the Energy Principle</i>	27
2.5 <i>Thermal Instability in Astrophysical Media</i>	27
2.6 <i>Magnetized Turbulence</i>	29
3 Cosmic-Ray Feedback on ISM Turbulence	31
3.1 <i>Abstract</i>	31
3.2 <i>Introduction</i>	32
3.3 <i>Methods</i>	37
3.4 <i>Results</i>	49
3.5 <i>Discussion</i>	69
3.6 <i>Conclusions</i>	77
3.7 <i>Appendix</i>	80
4 Cosmic-Ray Feedback on CGM/ICM Turbulence	82
4.1 <i>Abstract</i>	82
4.2 <i>Introduction</i>	83
4.3 <i>Physical Background</i>	86
4.4 <i>Simulation Details</i>	95
4.5 <i>Results</i>	104
4.6 <i>Discussion</i>	127
4.7 <i>Conclusion</i>	131
5 Single Cosmic-Ray Injections	135
5.1 <i>Abstract</i>	135

5.2	<i>Motivation</i>	136
5.3	<i>Background</i>	141
5.4	<i>Setup</i>	150
5.5	<i>Results</i>	159
5.6	<i>Discussion</i>	164
5.7	<i>Conclusion</i>	172
5.8	<i>Appendix</i>	175
6	Many Cosmic-Ray Injections	183
6.1	<i>Abstract</i>	183
6.2	<i>Introduction</i>	184
6.3	<i>Methods</i>	188
6.4	<i>Results</i>	200
6.5	<i>Discussion</i>	232
6.6	<i>Conclusion</i>	239
7	Cosmic-Ray Buoyancy	242
7.1	<i>Abstract</i>	242
7.2	<i>Introduction</i>	243
7.3	<i>Buoyant Stability</i>	246
7.4	<i>Compressibility of the ISM</i>	251
7.5	<i>Results</i>	255
7.6	<i>Discussion</i>	269
7.7	<i>Summary & Conclusions</i>	271
7.8	<i>Appendix</i>	272
8	Conclusion	282
	Bibliography	288

LIST OF FIGURES

1.1	Cosmic Rays in Cloud Chambers	5
1.2	The Cosmic Ray Energy Spectrum	8
2.1	Cosmic Ray Motion and Scattering	20
2.2	Magnetic Mirrors	22
2.3	Cloudy Cooling Curves	30
3.1	Thermal phase space with radiative cooling	42
3.2	Energy evolution in each simulation	50
3.3	Clumping factor evolution in each simulation	51
3.4	Simulation slices with magnetic field structure	52
3.5	Density distribution	58
3.6	Column density distribution	59
3.7	Phase diagrams in saturated state	60
3.8	Turbulent energy spectrum for each simulation	63
3.9	Density spectrum for each simulation	64
3.10	Diffusion coefficient distribution	66
3.11	Average diffusion coefficient and column density distribution	67
3.12	Lower resolution energy evolution	79
3.13	Lower resolution turbulent spectra	80
4.1	Critical Length Scales for Turbulent Reacceleration	90
4.2	Gamma-ray emissivity in phase space	94
4.3	Gamma-ray luminosity evolution	104
4.4	Density and Gamma-ray luminosity renders	105
4.5	Energy evolution	109
4.6	Thermodynamic, kinematic, and magnetic phase space plots	110
4.7	Growth rate of Gamma-ray luminosity	113
4.8	Cosmic-ray energy growth rate slices	115
4.9	Cosmic-ray energy density growth time histograms	117
4.10	Cells gaining CR energy over cells losing CR energy	118
4.11	Correlation factor	119
4.12	Phase space distribution for select simulations	122
4.13	Movement of gamma-ray emissivity weighted centroids	126
5.1	Single Injection Schematic	144
5.2	Parker Instability Phase Space Schematic	146

5.3	Single Injection Flux Tube Evolution	158
5.4	Slices from Single Injection Diffusion Simulation	162
5.5	Cosmic-Ray Pressure Evolution for Single Injections	163
5.6	Flux Tube Dynamics for Q1	166
5.7	Flux Tube Dynamics for Q2	168
5.8	Flux Tube Dynamics for Q3	170
5.9	Flux Tube Mass Loading and Height	170
5.10	Flux Tube Momentum	171
5.11	Simulation Cuts of Mass Loading	171
5.12	2D Parker Instability	176
5.13	The effect of Boundary Conditions on Single Injections	178
5.14	The Effect of Dimensionality on Single Injections	180
5.15	Single Injection Numerical Convergence	181
6.1	Multiple Injection Simulation Renders	185
6.2	Cooling and Thermodynamics in Tallbox	195
6.3	Multiple Injection Time Evolution	201
6.4	Multiple Injection Outflow Rates	202
6.5	Turbulent Energy Spectrum	205
6.6	Parker Instability Structure	208
6.7	Dense-gas formation	210
6.8	Vertical Energy Growth and Decorrelation of Cosmic Rays	211
6.9	Late-Stage Vertical Density Profile	214
6.10	Late-Stage Vertical Velocity Profile	214
6.11	Vertical Acceleration Components Spacetime Diagram	215
6.12	Late-Stage Average Vertical Acceleration Components	216
6.13	Late-Stage Net Vertical Force	217
6.14	Late-Stage Vertical Magnetic Field	220
6.15	Late-Stage Phase Space Diagrams	221
6.16	Fitted Vertical Profiles	225
6.17	Gas Scale Heights	226
6.18	F_{cal} Evolution	230
6.19	Comparison of Resolution and F_{cal}	231
7.1	Galactic Atmosphere	246
7.2	Compressibility	258
7.3	Instability Parameter Contours	260
7.4	Basic Galactic Atmosphere	263
7.5	Marginally Stable Atmosphere	265

7.6	Net Buoyancy Force	266
-----	------------------------------	-----

LIST OF TABLES

3.1	Simulation parameters	47
3.2	Radiative energy and cosmic-ray energy loss/gain rates	53
3.3	Filling factors in saturated state	61
3.4	Turbulent energy spectral index results	63
4.1	Simulation Parameters	97
5.1	Single Injection Simulation Parameters	154
6.1	Multiple Injection Simulation Parameters	192

TURBULENCE, INJECTION, AND BUOYANCY: WHEN AND WHERE COSMIC RAYS MATTER

Roark S. Habegger

Under the supervision of Professor Ellen G. Zweibel
At the University of Wisconsin-Madison

Cosmic rays are massive charged particles traveling at relativistic speeds, throughout our universe. From molecular clouds in the interstellar medium to the diffuse gas of the intracluster medium, cosmic rays have an impact on the dynamics and evolution of astrophysical media. Like many other natural phenomena, cosmic rays feed back on the processes and objects that create them. Utilizing numerical simulations alongside theoretical analysis, I show that cosmic rays eschew energy from multiphase turbulence, accelerate galactic outflows, and drive buoyant mixing between the galactic disk and halo. I also show how these effects depend severely on cosmic-ray transport and interaction with the background plasma. Ranging in length scale from parsecs to kiloparsecs and in time scale from kiloyears to gigayears, this thesis covers several avenues of cosmic-ray feedback found in our universe and identifies when cosmic rays can have a significant, observable impact on astrophysical systems and their dynamics.

1 INTRODUCTION

Ellen: So, what do you know about cosmic rays?

Roark: They are charged particles, traveling really fast. That's about it.

— (2020), over Zoom

Cosmic rays are particles traveling near the speed of light, flying throughout interstellar space. They are a factor of a billion less numerous than the non-relativistic particles that make up the interstellar medium (ISM). However, because they are a billion times more energetic, the total energy in cosmic rays is close to the total energy in other components of the ISM. That idea is the basic one from which I will build: the rarity of cosmic rays is outweighed by their large energy, making them a significant component of the ISM.

As an additional energy source and sink, cosmic rays give and take energy to and from the background plasma. That interaction facilitates *feedback*. Specifically, the properties of the background plasma determine how cosmic rays give and take energy. As cosmic rays give and take energy, they change the properties of the background plasma. This process leads to a circular effect of the cosmic rays “feeding back” on the plasma.

There are other avenues of cosmic-ray feedback due to their unique interaction with the magnetic field. First, cosmic rays exert an additional pressure on the plasma. This extra forcing due to cosmic rays changes the stability of plasma and can directly accelerate or decelerate bulk flows. Second, the cosmic-ray streaming instability can provide additional heating to the background plasma. Third, if a cosmic ray

physically contacts another particle, it can deposit significant energy, which triggers chemical reactions or the emission of high-energy photons, neutrinos, leptons, and other particles.

The details of cosmic-ray feedback have been studied and debated for decades, but a renewed focus has resulted because of their involvement in stellar feedback. Stars, as a component of galaxies, also affect the same processes that form them. They form from cold, dense, gravitationally unstable plasma in the ISM. At the end of a star's life, it can create a supernova explosion, which heats and accelerates the surrounding plasma. That heating and forcing from a supernova impacts the rate of star formation in a galaxy.

The generally accepted theory of the origin of galactic cosmic rays is that they are accelerated to relativistic energies in supernova blast waves. With that source in mind, cosmic-ray feedback in a galaxy can be viewed as a subset of stellar feedback because the cosmic rays fit into the same overarching process that affects star formation in a galaxy. By studying cosmic-ray feedback, we are further improving our understanding of how galaxies evolve, as well as how the ISM within galaxies is structured.

If the story were as simple as I have laid out so far, then there would not be much work to do, and I would not have spent my thesis studying cosmic-ray feedback. But the interplay of cosmic-ray transport, magnetic field structure, and multiphase turbulence produces complex effects on the evolution of astrophysical media like the ISM, intracluster medium (ICM), and circumgalactic medium (CGM). By impacting these media, cosmic rays also affect the evolutionary history of a given galaxy. While

this impact adds to the significance of the research in this thesis, the research is focused on the fundamental physical processes. Tying all of these processes back into a single model that explains changes in a galaxy’s star formation rate, metallicity, etc., is out of the scope of the current work. Instead, I focus on distilling the effects of cosmic-ray feedback on astrophysical media.

The thesis is divided into two key parts: first, an introduction and theoretical background setting the stage for the second part, my research on cosmic-ray feedback. The introduction (Chapter 1) covers observations of cosmic rays and their transport, along with the question that has interested me throughout the course of this research “*How does cosmic-ray transport affect cosmic-ray feedback?*”. In the theoretical background section (Chapter 2), I cover the details of cosmic-ray transport, discussing particle transport and moving up to the cosmic-ray “fluid” approximation of cosmic-ray magnetohydrodynamics (CRMHD). From there, I present an overview of CRMHD instabilities, along with providing more background detail on thermal instabilities and turbulence in magnetized astrophysical plasmas.

In Chapters 3 through 8, I present the research I have performed under the guidance of Dr. Ellen Zweibel. I begin with the results from simulations of turbulent, periodic boxes of a cosmic-ray fluid embedded in a magnetized, multiphase plasma (Chapters 3 and 4). Then, I present my studies of cosmic-ray injections in stratified boxes (Chapters 5 and 6). Those simulations inform the final section of the research, which discusses cosmic-ray buoyancy and its applications to the structure and stability of the galactic halo (Chapter 7). The conclusion section (Chapter 8) provides a summary of all the research, identifying the key results and takeaways for other

researchers.

1.1 A Brief History of Cosmic-Ray Astronomy

Cosmic rays are flying through my body as I write this and through your body as you read it. We were unaware of this ambient “radiation” until the discovery of radioactivity created ways of detecting it. With the ability to detect high-energy radiation, Victor Hess used hot air balloons and showed that this ambient background radiation increased with altitude, implying a “cosmic” source. This discovery in 1911-1912 started the field of cosmic-ray astronomy.

On the Earth’s surface, the cosmic rays appeared as γ -rays. They were uncharged and could produce air showers, as shown in the cloud chambers in Figure 1.1 from (Anderson et al., 1934). Figure 1.1 shows the tracks of cosmic rays through a cloud chamber. On the left of Figure 2.1, an air shower (the cascade of particles and their interactions) from an uncharged 2.5 GeV cosmic ray hitting a lead plate is shown. The cosmic ray was initially uncharged, since its track is absent in the upper panel. For some perspective, if it had been a proton cosmic ray, the proton would have a Lorentz factor of $\gamma = 2.66$, implying a speed of $v = 0.926c = 2.77 \times 10^8 \text{ m s}^{-1}$ ($\gamma = 4982$ and $v = 0.99999998c$ if the particle was an electron with the same total energy). On the right of Figure 1.1, there is an electron which, upon hitting the lead plate, changes into a positron (same mass as an electron with a positive charge) and gyrates counterclockwise instead of its initial clockwise rotation. Not only could this style of observation measure the flux of cosmic rays at different altitudes and




Figure 1.1: In the 1930s, cloud chambers were used to visualize the tracks of high-energy particles. These images come from early experiments that attempted to observe cosmic rays (Anderson et al., 1934). Left: A particle shower resulting from a 2.5 GeV Cosmic Ray hitting a lead plate (bar in the middle). Right: a high-energy electron becomes a positron when passing through the plate, leading to a change in its rotation direction.

latitudes, but it also provided a way to learn about the nature of the cosmic radiation.

Astrophysicists knew there had to be a source for this radiation. Millikan went so far as to claim they were the signature of matter creation by God¹. Ironically, Millikan and his collaborators would help disprove this theory and other ideas. Eventually, researchers showed the flux of cosmic rays varied with latitude, with fewer detected at the equator and more detected at higher latitudes. This result proved the flux depended on the Earth's magnetic field, which would only shield the equato-

¹If he had not died in 1953, then we would have been in for a real treat in 1964 when Penzias and Wilson discovered the cosmic microwave background, the radiation from the beginning of the universe.

rial regions from charged particles. The uncharged cosmic rays had to be the result of the charged particles hitting the Earth's atmosphere and producing high-energy γ -rays, which were then detected. Combining this with the anisotropy of cosmic rays incident on the Earth, observations in the 1930's already suggested that cosmic rays must primarily be protons traveling at relativistic speed.

One theory has stood the test of time and remains our best explanation for the production of cosmic rays: acceleration in supernova explosions. Baade & Zwicky (1934a) proposed this idea building from the fact that the observations — emission uncorrelated with any possible source objects — suggested a “sporadic” production mechanism. With a simple energy accounting, they showed that putting just 10% of each supernova explosion's energy into cosmic rays was enough to explain the observed flux. This theory seemed probable, but there was still no acceleration mechanism. The energy was there, ready to be taken.

Eventually, Fermi (1949) presented a viable mechanism: if cosmic rays scattered off of moving magnetic mirrors, then they could gain energy during those interactions. Since head-on collisions (mirrors moving towards the cosmic ray) are more common than receding collisions (mirrors moving in the same direction as the cosmic ray), the cosmic rays would gain energy through stochastic scatterings. Now, we refer to that process as second-order Fermi acceleration, because the energy gain is proportional to the squared velocity of the magnetic mirrors. Later, this theory would be improved by Bell (1978): if the cosmic rays were to scatter off Alfvén waves on the upstream and downstream sides of a shock front, then they would continuously gain energy from every jump over the front, taking it away from the shock. This process, known

as diffusive shock acceleration or first-order Fermi acceleration, is generally viewed as the key process for the production of most cosmic rays, as it will be efficient in supernova explosions (some of the most energetic shock waves in the universe). The process is first order because the energy gain is proportional to the velocity of the shock.

Alongside a viable acceleration mechanism for the most common cosmic rays, observations of cosmic rays continued to improve. On the Voyager 1 and 2 probes, the cosmic-ray subsystem (Stone et al., 1977) provided some early results for the cosmic ray flux in the solar system. The most interesting results from those probes came as Voyager 2 left the heliopause (the edge of the Sun's wind). Upon crossing this boundary, into the ISM, the flux of galactic cosmic rays increased (Stone et al., 2019). So, the solar wind protects us from cosmic rays, just as the Earth's magnetic field does.

Now, cosmic ray observatories have developed detailed models and precise measurements of the cosmic-ray spectrum at Earth. In Figure 1.2, we show the most accurate modern cosmic-ray spectrum, which shows the flux of cosmic rays at a given energy. Each particle species has a different flux, with cosmic-ray protons of an energy ~ 10 GeV being the most common. Other than the peak of the curve, there are two other interesting parts: the knee and the ankle. Currently, we still do not know what causes the knee, a slight bump in the spectrum at PeV energies (I worked on a possible source, interacting shock waves, with Dr. Aerdker and Dr. Merten (Aerdker et al., 2025)). The ankle above EeV energies is caused by extragalactic cosmic rays.

With these detailed observations, astrophysicists have been able to discover more




Figure 1.2: Combined measurements of the cosmic-ray energy spectrum incident on Earth, from Evoli (2020). Protons with an energy between 1 GeV and 100 GeV dominate the distribution. The distribution has several interesting components, including the labeled “knee” and “ankle.”

about cosmic rays. In particular, by comparing ratios of the amount of higher atomic number cosmic rays (e.g., boron and carbon (Evoli et al., 2019)), we can estimate the confinement time of cosmic rays — how long they have been bouncing around the ISM. The best estimates of the diffusion coefficient of the ~ 10 GeV proton cosmic rays in the ISM are now at $\sim 10^{28} \text{ cm}^2\text{s}^{-1}$ (Evoli et al., 2019). This value implies those cosmic rays are taking a diffusive random walk with a mean free path between scatterings of approximately $10^{18} \text{ cm} \sim 0.3\text{pc}$, with each scattering occurring approximately every 3yr. On the scales of the ISM (from 1 pc to ~ 1 kpc, with dynamical times ~ 1 Myr), an individual cosmic ray will scatter millions of times on smaller scales than the dynamics we want to observe. This reality makes the study of the dynamical impact of proton cosmic rays multi-scale and difficult to solve.

1.2 Cosmic-Ray Feedback

Others have cataloged and synthesized the multitude of avenues for cosmic-ray feedback in extremely useful review papers (Zweibel, 2017; Ruszkowski & Pfrommer, 2023; Owen et al., 2023), so I do not attempt to give a complete review here. Instead, I will provide an overview of the process and explain where the work in this thesis fits into our modern understanding of cosmic-ray feedback.

The peak of the cosmic-ray spectrum alone provides an interesting result. The spectrum in Figure 1.2 is the flux of cosmic-ray energy. To get an estimate of the total cosmic-ray energy, we need to divide this by the speed at which cosmic rays are traveling and integrate over solid angle. Since that speed is effectively the

speed of light, we get an energy density in the proton cosmic rays on the order of $4\pi \times 10^{-5} \text{GeV m}^{-3} = 10^{-1} \text{eV cm}^{-3}$. This energy density is equivalent to a thermal gas with an average density of the ISM (1cm^{-3}) and a temperature of $T = 10^4 \text{K}$, which matches the warm neutral medium (WNM). So, the energy density in cosmic rays in the ISM is about the same as the thermal energy density in the ISM. That similarity is the first clue to cosmic-ray feedback. Just as a ball thrown up in the air will have equal gravitational potential energy and kinetic energy in the middle of its rise and in the middle of its fall, the ISM is actively moving energy between cosmic rays and the thermal gas (and magnetic energy, turbulent energy, radiative energy, etc.). This “energy equipartition” which has been observed not only in the Milky Way but also in other galaxies, is the first step: if there is an additional source or sink of energy, an additional place to store energy, then it could affect the dynamics of the overall system.

In the study of stellar feedback, there are two ways in which ionizing radiation, stellar winds, and supernovae affect the future production of stars: direct and indirect. Direct refers to the physical removal of gas from a galaxy — if it is not in the galaxy, but instead sent out into the galactic halo and CGM, then it is less likely that it will cool and eventually form stars. Indirect refers to heating — if there is ample heating due to feedback, then the gas in the galaxy will not be able to cool and eventually become gravitationally unstable, forming stars. Cosmic rays contribute to both the indirect and direct parts of stellar feedback. Increasing the cosmic-ray pressure increases the total pressure in the ISM — this pressurization will adjust the dynamics of the background plasma. Any gradients in cosmic-ray pressure can

accelerate the gas in the ISM, helping to drive galactic winds. In addition to those direct methods, cosmic rays can generate additional heat: if cosmic rays create the same magnetic fluctuations they scatter off of (the self-confinement model), those fluctuations will damp and heat the background plasma.

Stellar feedback is expected to be the primary regulator of star formation in galaxies less massive than our own. For more massive galaxies, black hole feedback is likely the primary regulator. However, since the supermassive black holes in the center of the galaxy create feedback by launching magnetized, jetted outflows, they can also accelerate cosmic rays. Therefore, cosmic-ray feedback will operate across the full range of galaxy masses, not just in small galaxies regulated by stellar feedback.

The forcing and heating by cosmic rays depend on how they are transported on microscales. In particular, the evolution of the cosmic-ray distribution function (Chapter 2.2) determines how the cosmic-ray pressure will evolve and change. Therefore, we first need to model these kinetic scales before we can move to a simplified fluid dynamical picture of cosmic-ray pressure for the effect of cosmic rays on large scales.

This thesis focuses on this fundamental question: *“How does cosmic-ray transport affect cosmic-ray feedback?”* Through the use of cosmic-ray magnetohydrodynamics (Chapter 2.3), we have examined several cases where changes in assumptions about kinetic physics (evolution of the cosmic-ray distribution function) change the dynamical impact of cosmic rays on the background thermal plasma.

First, we examine cosmic-ray feedback on turbulence in the ISM (Chapter 3) and on larger scales, in the CGM and ICM (Chapter 4). In both cases, we show

that cosmic rays can extract energy from turbulent motion, as well as from the condensation mode of thermal instability. This extraction of energy reduces the total amount of energy radiated away by condensing gas. Instead, the cosmic-ray pressure grows. In the case of extrinsic confinement by magnetized turbulence, the cosmic ray energy will grow until it becomes large enough that the plasma will light up, producing significant γ -ray luminosity, which disagrees with observations. Therefore, the cosmic rays need to be transported faster, or confined by the waves they generate by themselves. In that case, the cosmic rays deposit the additional energy in the rest of the plasma, not the parts condensing due to thermal instability.

Second, we examine cosmic-ray feedback due to injections of cosmic rays from supernova explosions (Chapters 5 and 6). We find that these injections can drive instability and produce large-scale outflows similar to the Parker instability, except on a much faster timescale. The faster timescale primarily results from the loading of magnetic flux tubes in the galactic disk. When a supernova injects cosmic rays, they will be transported away from the injection site, along the mean magnetic field. This propagation increases the cosmic-ray pressure along that magnetic flux tube, causing it to buoyantly float upward. That buoyancy still depends on the state of the initial background in the single injection case (Chapter 5). Once we put in a realistic number of injections (matching the solar neighborhood), the entire mid-plane of the galaxy becomes over-pressured due to the injections occurring at random locations. Eventually, the flux tubes erupt out of the disk, driving outflows and inflows. The outflows are driven by cosmic-ray pressure gradients, and the inflows form massive clouds of cold gas, which could conceivably drive significant star formation at later

times as they sink into the galactic disk.

Third, we present a theory of how cosmic-ray transport predisposes the galactic disk and halo to mixing. Cosmic-ray transport is more diffusive in the disk, so the plasma there is more compressible and more likely to be unstable to buoyancy. Moving outward from the disk into the halo, the cosmic-ray pressure becomes the dominant pressure, and the cosmic-ray transport switches to become more advective. This plasma is then harder to compress because it takes more energy to compress the cosmic rays. When there is more compressible material underneath less compressible material, the more compressible material has a tendency to rise. This process drives mixing in a variety of cases (including marginal stability) and could be an important factor in the study of the motion of gas in the ISM.

All the work in this thesis identifies and focuses on the effects of cosmic-ray transport, testing the three primary avenues: advection, streaming, and diffusion. These are the avenues of the transport of cosmic-ray energy across large scales — each has its own assumptions about the actual particle transport on small scales. But in general, they all assume the cosmic rays take a diffusive random walk. Recent work has highlighted that the picture might be more complicated, with sub-diffusion, super-diffusion, and scattering off of intermittent structures playing an important role (Kempski et al., 2023a; Lemoine, 2023). Even in the case that we must revise the equations of cosmic-ray magnetohydrodynamics, the results in this thesis will be useful because we have shown how some of the dominant transport avenues affect feedback. Also, for the ~ 10 GeV cosmic rays, these additional transport effects are less important. Since those particles dominate feedback, our results will remain

important as the study of cosmic-ray feedback continues and we begin to account for all the complexity hidden in the lives of these energetic particles.

2 THEORETICAL BACKGROUND

A journey of a thousand miles begins with a single step.

— *Tao Te Ching* (4th Century BCE) by Lao Tzu

Fundamentally, cosmic rays are charged particles. Their motion can be solved utilizing the Lorentz Force equation alongside Maxwell’s equations. However, to study larger scales and many cosmic rays, we need a more tenable treatment than integrating the motions of the $N_{\text{CR}} \sim 10^{-9} \text{ cm}^{-3} \times (1 \text{ pc})^3 = 3 \cdot 10^{46}$ particles in one cubic parsec. That need brings us first to the evolution of the *distribution function* of cosmic rays. Next, we simplify to the evolution of *moments* of the cosmic-ray distribution function, under the assumption that we are studying length scales much larger than the mean free path of cosmic rays. I refer to the final set of equations as cosmic-ray magnetohydrodynamics (CRMHD), as it appends some equations related to cosmic rays onto the ideal magnetohydrodynamic (MHD) equations. This chapter walks through the derivation of CRMHD, drawing from the work of others (Zweibel, 2017; Thomas & Pfrommer, 2019; Ruszkowski & Pfrommer, 2023).

Once we have the equations of CRMHD, we apply them to some interesting astrophysical setups and analyze their stability. Then, I cover thermal instability and magnetized turbulence at an introductory level. Those two topics are particularly useful in understanding the research on cosmic-ray feedback in turbulence.

The derivations and presentation of this material build from past reviews on cosmic-ray propagation and feedback (Zweibel, 2013, 2017; Ruszkowski & Pfrommer, 2023; Owen et al., 2023), as well as from astrophysics textbooks (Kulsrud, 2004;

Draine, 2011). Additionally, many of these ideas trace back to meetings with Dr. Zweibel and Dr. Ruszkowski, and the many professors and mentors with whom I have had the privilege of working and studying. When possible (based on my memory), I make a note of those who suggested a given idea or line of reasoning.

2.1 Cosmic Rays as Particles

2.1.1 Particle Motion

Each cosmic ray is a particle with position \mathbf{x} and momentum \mathbf{p} . Since the cosmic rays are traveling near the speed of light, we need to use the relativistic treatment of particle motion. Therefore, the cosmic ray's momentum, in the lab frame, is

$$\mathbf{p} = \gamma m \mathbf{v} = \gamma m \frac{d\mathbf{x}}{dt} \quad (2.1)$$

where $\gamma = 1/\sqrt{1 - v^2/c^2}$ is the Lorentz factor. While the particle's velocity will never exactly reach the speed of light c , its momentum will be infinitely large as its velocity approaches c . Similarly, the cosmic ray's energy

$$E = \sqrt{p^2 c^2 + m^2 c^4} = \gamma m c^2 \quad (2.2)$$

can reach arbitrarily high values. This difference, between using velocities and momenta (or energy or Lorentz factor), becomes more useful when we study the distribution function of cosmic rays.

As a cosmic ray moves through space, it experiences forces exerted on it by the

electromagnetic field. Specifically, the Lorentz force:

$$\frac{d\mathbf{p}}{dt} = q \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right). \quad (2.3)$$

Solving this for an arbitrary electric and magnetic field is difficult, but has been done extensively for nearly a century. Instead, we are going to use simple examples to identify key conserved quantities, focusing on the physical mechanisms underlying cosmic-ray transport and acceleration.

For example, if we dot Equation 2.3 with the velocity vector \mathbf{v} , we will find

$$\mathbf{v} \cdot \frac{d\mathbf{p}}{dt} = q\mathbf{v} \cdot \mathbf{E} \quad \rightarrow \quad \frac{1}{2} \frac{dp^2}{dt} = q\mathbf{p} \cdot \mathbf{E} \quad (2.4)$$

because the momentum vector is always parallel to the velocity vector, making it perpendicular to the vector $\mathbf{v} \times \mathbf{B}$. Equation 2.4 provides a key conclusion: the only way a particle's energy can change is as the result of an electric field acting on it. The magnetic field cannot change its energy.

If we have the electric field vanish, then the particle's motion is only a function of the magnetic field at its position:

$$\frac{d\mathbf{p}}{dt} = \frac{q}{\gamma mc} \mathbf{p} \times \mathbf{B}(\mathbf{x}). \quad (2.5)$$

Conveniently, we already know the total energy E is conserved, so the Lorentz factor γ is constant too. Splitting the momentum into components parallel and perpendicular to the magnetic field $\mathbf{p} = p_{\parallel} \hat{\mathbf{b}} + \mathbf{p}_{\perp}$ (note we use $\hat{\mathbf{b}} = \mathbf{B}/B$ to denote the unit

vector pointing in the direction of the magnetic field), we get two equations:

$$\frac{dp_{\parallel}}{dt} = 0 \quad \frac{d\mathbf{p}_{\perp}}{dt} = \frac{q}{\gamma mc} \mathbf{p}_{\perp} \times \mathbf{B}(\mathbf{x}). \quad (2.6)$$

Since \mathbf{p}_{\perp} exists only on the two-dimensional plane perpendicular to the local magnetic field direction, we can split it into components $p_1 \hat{x}_1$ and $p_2 \hat{x}_2$ such that the unit vectors $(\hat{x}_1, \hat{x}_2, \hat{b})$ form an orthonormal basis of \mathbb{R}^3 . This decomposition results in a pair of coupled equations

$$\frac{dp_1}{dt} = \frac{qB}{\gamma mc} p_2 \quad \frac{dp_2}{dt} = -\frac{qB}{\gamma mc} p_1. \quad (2.7)$$

We see that for a constant magnetic field, the solutions take the form $p_i = A \sin(\omega_c t) + B \cos(\omega_c t)$ with a frequency called the *cyclotron frequency*:

$$\omega_c := \frac{qB}{\gamma mc}. \quad (2.8)$$

Altogether, this gives a position vector of the form

$$\mathbf{x}(t) = \mathbf{x}_0 + \frac{p_{\perp}}{\gamma m \omega_c} \sin(\omega_c t) \hat{x}_1 - \frac{p_{\perp}}{\gamma m \omega_c} \cos(\omega_c t) \hat{x}_2 + \frac{p_{\parallel}}{\gamma m} t \cdot \hat{b} \quad (2.9)$$

where I have assumed a starting position \mathbf{x}_0 and the perpendicular momentum vector initially parallel to the \hat{x}_1 direction. The motion is therefore a helix, wrapping around the magnetic field line at a constant radius. Instead of splitting into perpendicular and parallel momentum, it is also common to define the cosine of the pitch angle θ

of the particle:

$$\mu := \cos \theta = \frac{p_{\parallel}}{p}. \quad (2.10)$$

From this definition, $p_{\parallel} = \mu p = p \cos \theta$ and $p_{\perp} = p \sin \theta = p \sqrt{1 - \mu^2}$. We often decompose the momentum space of the particle into total momentum p and cosine μ , as this captures the key aspects of the particle's motion.

The radius of gyration of the orbit (the Larmor radius or gyro-radius) is set by the perpendicular momentum, specifically:

$$r_L = \frac{p_{\perp}}{\gamma m \omega_c} = 3.13 \times 10^{12} \text{ cm} \frac{1 \mu\text{G}}{B} \frac{m}{m_p} \frac{e}{q} \sqrt{\gamma^2 - 1} \sqrt{1 - \mu^2}. \quad (2.11)$$

For a \sim GeV proton cosmic ray with $\gamma = 2$, this radius is 0.3AU. Figure 2.1 illustrates the gyration of the cosmic ray at this Larmor radius in illustrations 1, 2, and 3, from different viewing angles. In illustration 4, we show the breakdown of the particle's velocity into parallel and perpendicular components.

But the universe is not threaded by a uniform magnetic field — all the interesting parts wiggle and wave through space. Illustration 5 in Figure 2.1 shows an apt analogy for how cosmic rays interact with these wiggles. Cosmic rays barely interact with large-scale waves with wavelength $\lambda \gg r_L$, because the magnetic field appears uniform over every gyro-orbit. The cosmic rays fly right over them, like a cyclist over a mountain or hill. Similarly, the cosmic rays will average over small-scale fluctuations with wavelength $\lambda \ll r_L$ during their gyro-orbit — just like a cyclist riding over a gravel road. Once a cyclist comes across a hill that is equivalent in size to the length of their bike or the size of their tire, then they will wipe out or

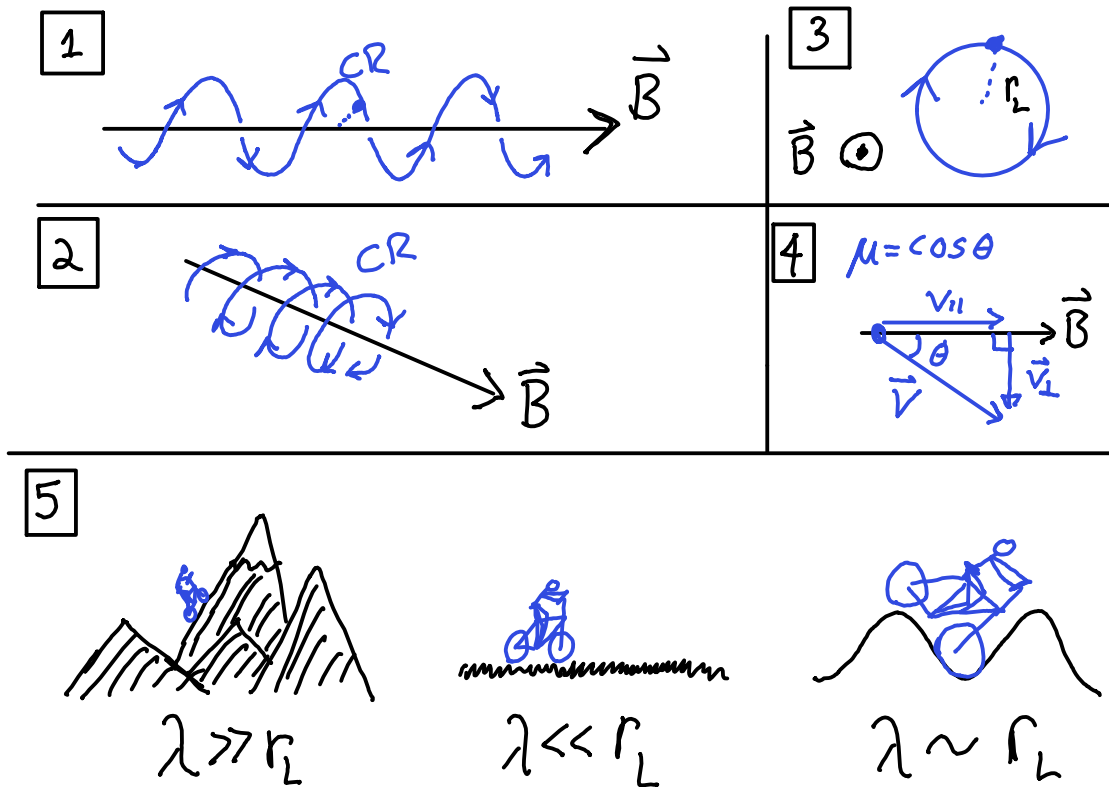


Figure 2.1: Diagrams illustrating the motion of a cosmic ray in a uniform magnetic field. First, an edge-on view of a cosmic ray gyrating around a magnetic field. Second, an off-angle viewing of the gyration. Third, a face-on view of a cosmic ray's gyration around the magnetic field. Fourth, a diagram breaking the particle's motion into components parallel and perpendicular to the magnetic field. Fifth, an analogy of gyroresonant scattering of cosmic rays, illustrating that a cyclist can ride up a mountain (long wavelength fluctuations $\lambda \gg r_L$), or over gravel (short wavelength fluctuations $\lambda \ll r_L$), but a cyclist will have difficulty going over hills with a wavelength equal to the size of their bike ($\lambda \sim r_L$). This process matches the experience of cosmic rays coming across fluctuations in the magnetic field (analogy from E. G. Zweibel).

wreck their bike. Similarly, cosmic rays coming across variations in the magnetic field with wavelengths $\lambda \sim r_L$ will scatter in a random direction off the fluctuation (this analogy comes from E. G. Zweibel).

We will examine the resonant interaction of an entire population of cosmic rays with electromagnetic fluctuations in Chapter 2.2, but we will complete our examination of particle motion by studying how a single cosmic ray would interact with a magnetic bottle. This example will provide some useful insight into the process of cosmic-ray scattering, as well as how cosmic rays can be accelerated without a large-scale electric field.

2.1.2 Magnetic Bottle Reflection

A magnetic bottle is a structure often constructed in plasma physics using strong magnets, but they also appear naturally in turbulent astrophysical environments where compressions in the magnetic field form due to the cooling of plasma (see Chapter 2.5) or through gravitational instability. Because the magnetic field is divergence-less ($\nabla \cdot \mathbf{B} = 0$), magnetic field lines never converge to an end point. Combine this with flux freezing in ideal magnetohydrodynamics, where the flux of magnetic field through the plasma must be conserved, and any compressions of gas will pinch magnetic field lines.

The ideal magnetic bottle structure is illustrated in Panel 1 of Figure 2.2: two regions of strong magnetic field, connected by a region of weak magnetic field. The Larmor radius (Equation 2.11) will decrease in the stronger field regions. Each side of the magnetic bottle is its own *magnetic mirror*. Note that the system is symmetric

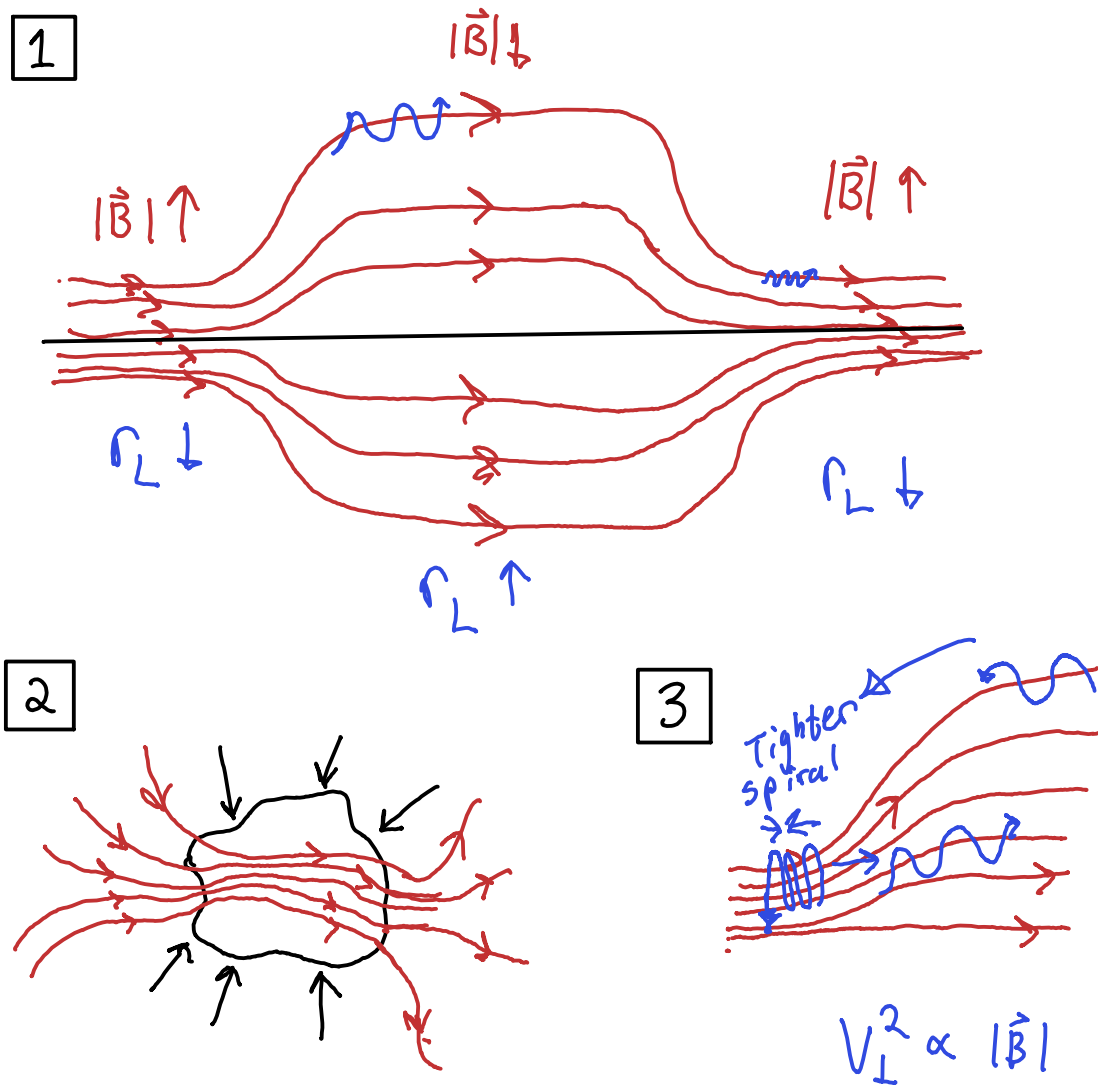


Figure 2.2: Charged particles in magnetic bottles. Panel 1: A magnetic bottle is a region of weak magnetic field between two regions of strong magnetic field, which are threaded with the same field lines. The Larmor radius of the particle gets smaller in the strong fields. Panel 2: Compressions of the magnetic field, which are ubiquitous in astrophysical plasmas, can produce these mirror structures. Panel 3: Due to the conservation of the particle's energy, the perpendicular velocity of the particle must increase. Eventually, the parallel velocity will go to zero, and then negative, reflecting off the magnetic mirror.

about the \hat{z} axis, creating a cylindrical structure. Panel 1 only shows a single slice of the cylinder. In a moment, we will discuss how these stronger field regions actually reflect particles. But the structure in Panel 1 will keep some particles stuck, bouncing between the two strong-field regions.

Panel 2 of Figure 2.2 shows a formation mechanism for a magnetic mirror in a turbulent plasma. Any bulk compression perpendicular to the magnetic field will cause a region of increased field strength, as all the lines are pinched together. Similarly, an expansion or rarefaction perpendicular to the magnetic field could create the bottle structure in Panel 1. Given the ubiquity of turbulence in astrophysics, along with the effects of gravitation and condensation (Chapter 2.5), these magnetic mirrors should form easily and be common.

Panel 3 of Figure 2.2 shows a better description of the process of mirroring. As the particle moves towards a region of a stronger magnetic field, its energy will remain constant, since a steady magnetic field cannot change the energy of the particle. Despite the energy being constant, there is an additional force due to the additional component of the magnetic field in the \hat{x} and \hat{y} directions.

2.1.3 Fermi Acceleration

2.2 Dynamics of the Cosmic-Ray Distribution Function

The number density of cosmic rays in the interstellar medium is approximately 10^{-9} cm^{-3} . In the volume of the solar system, approximately $(100 \text{ AU})^3 = 3 \times 10^{45} \text{ cm}^3$, there are over 10^{36} particles. This number is similar to the approximately 10^{37} molecules of water in the 0.5 km^3 volume of Lake Mendota. Just like simulating the motion of every water molecule in Lake Mendota is a fool's errand, the same is true for cosmic rays in the ISM. The problem becomes worse as we move to larger scales, like $1 \text{ pc} = 2.06264 \times 10^5 \text{ AU}$, kpc, and so on.

With a large number of particles, we can move from the individual-particle picture to a statistical picture, which tracks the *distribution function* of particles. The distribution function gives the number of particles per (physical) volume, per momentum space volume. By momentum space, I mean the 3D space of all possible momentum vectors a given particle could have. I.e., all combinations of p_{\parallel} and the 2D vector \mathbf{p}_{\perp} . The distribution function is given by

$$f(\mathbf{x}, \mathbf{p}, t) = \frac{d\mathcal{N}}{d^3x d^3p}. \quad (2.12)$$

By construction, this function can tell us information, like the number density of

cosmic rays $n(\mathbf{x}, t)$ if we integrate over momentum space:

$$n(\mathbf{x}, t) = \int d^3p [f(\mathbf{x}, \mathbf{p}, t)]. \quad (2.13)$$

Or we can calculate the energy density of cosmic rays by taking a moment of the distribution function (multiplying it by some quantity related to momentum):

$$E_c(\mathbf{x}, t) = \int d^3p [Ef(\mathbf{x}, \mathbf{p}, t)] \quad (2.14)$$

where E is the particle energy (Equation 2.2). We will look at these moments of the distribution function more in Chapter 2.3.

The differential equation governing the evolution of the distribution function is the Boltzmann equation:

$$\frac{df}{dt} \equiv \frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{\partial \mathbf{p}}{\partial t} \cdot \frac{\partial f}{\partial \mathbf{p}} = \left(\frac{\delta f}{\delta t} \right)_c. \quad (2.15)$$

The equation results from the conservation of phase space volume. The only deviation from phase space volume conservation comes from collisions with other particles, given by the collision operator $(\delta f / \delta t)_c$. While cosmic rays can collide with other particles, those collisions are rare over the course of their propagation. Instead, cosmic rays are an effectively collisionless plasma — we will study the phase space dynamics (kinetics) while neglecting collisions.

The cosmic rays will experience the Lorentz force (Equation 2.3), along with feeling the effects of any external gravitational field. Since the cosmic rays contribute

minimally to the overall mass distribution, we do not need to consider self-gravity. Therefore, the only gravitational force they experience will be an external field created by other masses. That additional acceleration is straightforward to add at any point, so we neglect it for the rest of this derivation. Instead, we just focus on their interaction with the electromagnetic fields, leaving us with the following differential equation:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + q \left(\mathbf{E} + \frac{\mathbf{v} \times \mathbf{B}}{c} \right) \cdot \frac{\partial f}{\partial \mathbf{p}} = 0. \quad (2.16)$$

2.2.1 Quasi-linear Theory

2.2.2 Self Confinement or Extrinsic Confinement?

2.3 Cosmic-Ray Magnetohydrodynamics

2.3.1 Full Generality

2.3.2 Two Moment: A Numerically Driven Solution

2.4 Instabilities and the Energy Principle

2.4.1 Energy Principle in CRMHD

2.4.2 Convective Instability (Parker)

2.5 Thermal Instability in Astrophysical Media

Thermal instability plays a key role in the dynamics of the ISM, acting as a bridge to the higher densities necessary for gravitational collapse and star formation. Specifically, non-adiabatic thermal instability. Using the first law of thermodynamics:

$$dU = dQ - PdV \tag{2.17}$$

where U is the specific thermal energy, Q is the heating, P is the gas pressure, and V is the specific volume. Translating this to the fluid equation form with thermal

energy density $E = U\rho$ and $V = \rho^{-1}$, we find

$$dE = \rho \cdot dQ + (P + E)\frac{d\rho}{\rho}. \quad (2.18)$$

Generally, in fluid dynamics, we will assume a closed system with no gain in heat — an *adiabatic* system with $dQ = 0$. Combining this with an equation of state like $P = E(\gamma - 1)$ will lead to the traditional fluid energy equation in Lagrangian form.

However, astrophysical media are more complex. There are additional heating and cooling processes involved. In particular, there is radiative cooling and heating, as the plasma can release energy in the form of light which escapes from the system. Cosmic rays can provide ionization, the photo-electric effect on dust grains can heat the thermal gas, and other processes can heat the ISM non-adiabatically. The original idea for how these processes could lead to thermal instability comes from Field (1965). Field defined a heat-loss function \mathcal{L} such that $dQ/dt = -\mathcal{L}$. That prescription leads to the fluid dynamical equation

$$\frac{\partial E}{\partial t} + \mathbf{u} \cdot \nabla E = -\rho\mathcal{L} - \gamma E \nabla \cdot \mathbf{u} \quad (2.19)$$

where we have assumed $P = E(\gamma - 1)$ and a bulk fluid velocity \mathbf{u} .

Across astrophysics, we have developed useful tables and models of the heating and cooling processes in the ISM, CGM, and ICM. The common choice is to define a heating rate per particle Γ , and a volumetric cooling rate Λ . Those definitions imply a heat-loss function

$$\rho\mathcal{L} = n^2\Lambda - n\Gamma. \quad (2.20)$$

In the simplest scenarios, Γ is a constant set by the cosmic-ray ionization rate or the photo-electric heating rate. The cooling Λ is more difficult to model, since it is dominated by the optically-thin cooling of plasma. The cooling rate depends on the abundance of each atomic element in the ISM, the density and temperature of the plasma, and the ionization fraction. The plasma will emit photons, which will escape because the low density leaves it optically thin to radiation. That limit primarily applies to temperatures above $T \sim 100\text{K}$ and densities below $n \sim 10^2\text{cm}^{-3}$. Those low temperatures and high densities correspond to molecular clouds, and those can become optically thick to radiation. In particular, self-shielding of molecular hydrogen is an important process that only occurs once the gas becomes optically thick.

In Figure

2.5.1 Condensation Mode

2.6 Magnetized Turbulence

2.6.1 Kolmogorov's Cascade

2.6.2 Critical Balance

2.6.3 Compressive Modes and Cosmic Rays




Figure 2.3: Cooling function and its properties, calculated using the Cloudy photoionization code (Chatzikos et al., 2023). The top plot shows the cooling function Λ as a function of temperature, with each line corresponding to a metallicity and hydrogen number density given by the colorbar and legend on the top of the plot. The middle panel shows the slope of the cooling function, along with a dashed black line denoting the condensation mode stability criterion from Field (1965) — plasma above the dashed black line is thermally stable, whereas gas underneath it is unstable, with a growing condensation mode. The bottom plot shows the cooling time for some densities. These cooling times can be compared to the sound-crossing time of the plasma, shown with black dashed lines as a function of length scale and Mach number.

3 COSMIC-RAY FEEDBACK ON ISM TURBULENCE

Of the deluge and how to represent it in a picture: Let the dark and gloomy air be seen buffeted by the rush of contrary winds and dense from the continued rain mingled with hail and bearing hither and thither an infinite number of branches torn from the trees and mixed with numberless leaves.

— *The Notebooks of Leonardo Da Vinci* (1880) by Jean Paul Richter.

This chapter is a modified copy of the paper Habegger et al. (2024), titled *Cosmic-Ray Feedback on Bi-stable Interstellar Medium Turbulence*, published in The Astrophysical Journal in 2024. No text has been changed; the figures and equations have just been reformatted to accommodate the submission of this thesis document.

I deeply appreciate the contributions of my co-authors Dr. Ka Wai Ho, Dr. Ka Ho Yuen, and Dr. Ellen Zweibel, to this work.

3.1 Abstract

While cosmic rays ($E \gtrsim 1$ GeV) are well coupled to a galaxy’s interstellar medium (ISM) at scales of $L > 100$ pc, adjusting stratification and driving outflows, their impact on small scales is less clear. Based on calculations of the cosmic ray diffusion coefficient from observations of the grammage in the Milky Way, cosmic rays have little time to dynamically impact the ISM on those small scales. Using numerical simulations, we explore how more complex cosmic ray transport could allow cosmic

rays to couple to the ISM on small scales. We create a two-zone model of cosmic ray transport, with the cosmic ray diffusion coefficient set at the estimated Milky Way value in cold gas but smaller in warm gas. We compare this model to simulations with a constant diffusion coefficient. Quicker diffusion through cold gas allows more cold gas to form compared to a simulation with a constant, small diffusion coefficient. However, slower diffusion in warm gas allows cosmic rays to take energy from the turbulent cascade anisotropically. This cosmic ray energization comes at the expense of turbulent energy which would otherwise be lost during radiative cooling. Finally, we show our two-zone model is capable of matching observational estimates of the grammage for some transport paths through the simulation.

3.2 Introduction

Cosmic rays, along with magnetic fields, turbulent energy, thermal energy, and radiation pressure, are a major constituent of the interstellar medium (ISM, see McKee & Ostriker (2007)). Cosmic ray pressure, provided primarily by protons of a few GeV, contributes to the vertical stratification of galactic disks, drives galactic winds, impedes star formation, and alters the evolution of supernova remnants (see Zweibel (2017), Ruszkowski & Pfrommer (2023), and Owen et al. (2023) for recent reviews and references).

One of the major unresolved issues in cosmic ray astrophysics is how the cosmic rays derive their energy. The favored and dominant mechanism is acceleration in shocks driven by supernovae, with a complementary role for shocks in stellar winds

(Axford et al., 1977; Bell, 1978; Blandford & Ostriker, 1978). But the extent to which cosmic rays tap other sources of energy (e.g., energy in large-scale turbulence (Ptuskin, 1988; Bustard & Oh, 2022)) throughout their lifetime is still unclear.

A mechanism for converting turbulent energy to cosmic ray energy was suggested by Ptuskin (1988) and has recently been revived by Commerçon et al. (2019), Dubois et al. (2019), Bustard & Oh (2022), Bustard & Oh (2023). This mechanism works as follows: cosmic rays are scattered by kinetic scale fluctuations in the ambient magnetic field (a standard feature of cosmic ray transport theory since the foundational work of Parker (1965) & Jokipii (1966)). On scales larger than the mean free path between scatterings the cosmic rays behave as a diffusive fluid. This diffusion allows cosmic rays to extract energy from compressive fluctuations much as a thermally conducting gas absorbs heat from sound waves. The energy gain is maximized when the diffusion time across a turbulent eddy is comparable to the eddy turnover time.

The value of the cosmic ray diffusion coefficient in the Milky Way disk is derived mainly from the observed abundances of cosmic rays produced by spallation reactions between cosmic rays and interstellar gas. Those abundances are essentially a column density measurement over the convoluted paths cosmic rays follow. Recent measurements estimate the diffusion coefficient to be $\kappa_{\parallel} \sim 10^{28} \text{ cm}^2 \text{ s}^{-1}$ (Evoli et al., 2019, 2020). Taking a reasonable driving scale for ISM turbulence of 100 pc and a realistic phase velocity of 10 km s^{-1} gives a characteristic eddy turnover time of $\tau \sim 10 \text{ Myr}$. On the same length scale, the cosmic rays will diffuse in just $\tau \sim 0.3 \text{ Myr}$. This fundamental mismatch suggests that turbulent energy extraction and feedback on ISM scales by cosmic rays is weak. This conclusion is corroborated by Commerçon et al.

(2019); Bustard & Oh (2022, 2023). Achieving maximum coupling would require a diffusivity of order $\sim 10^{26} \text{ cm}^2 \text{ s}^{-1}$.

The effects of cosmic rays on turbulence are not limited to extraction of energy. In a bi-stable medium, decreasing the cosmic ray diffusion coefficient to a small value will limit the production of cold gas (Commerçon et al., 2019). This result comes from cosmic rays being trapped in condensing clouds instead of diffusing out. While stuck in the clouds, they provide a significant pressure against condensation and collapse. When cosmic ray energy gain is maximized (when the eddy turnover time and cosmic ray diffusion time scale are similar), there are some circumstances under which the turbulent energy cascade is truncated at $L \lesssim 10 \text{ pc}$ scales (Bustard & Oh, 2023). Both of these effects would be in conflict with observations, which show plenty of cold gas formation and no cutoff in the turbulent energy cascade down to a scale of $L \lesssim 1 \text{ pc}$ (Pingel et al., 2018, 2022).

These past works on cosmic ray coupling to turbulence generally assume a constant diffusion coefficient. To test if the conclusions remain true for more complex cosmic ray transport, we examine the case where κ_{\parallel} is not a spatial constant. This idea is similar to two-zone and multi-zone models used in studies of cosmic ray propagation (Guo et al., 2016; Jóhannesson et al., 2019; De La Torre Luque et al., 2023), and has a sound physical basis: in weakly ionized gas, the small scale magnetic field fluctuations which cosmic rays scatter off are strongly damped by friction between ions and neutrals (ion-neutral damping, see Kulsrud & Cesarsky (1971)). Therefore, it is plausible that κ_{\parallel} is larger in the denser, more neutral phases of the ISM.

Instead of tracking or evolving ionization in our simulations, we use temperature

dependence as a proxy. We have the cosmic rays diffuse slowly through warm gas and quickly through cold gas as a result of ion-neutral damping. This model is derived from a study of cosmic ray driven galactic outflows and star formation feedback in global galaxy simulations by Farber et al. (2018). In that paper, the cosmic ray diffusion coefficient was assumed to change between hot and warm gas, resulting in lower density winds and more star formation.

Our goal with implementing a variable diffusion coefficient and applying it to a bi-stable ISM is to reconsider some of the results from Commerçon et al. (2019); Bustard & Oh (2022, 2023). **More specifically, can we form cold gas while still having cosmic rays significantly impact the properties of the turbulent, multiphase ISM?**

In this paper, we demonstrate that the two zone model allows cosmic rays to extract significant energy from turbulence in warm, low density gas while reducing their dynamical role in cold gas. Additionally, the cosmic rays accumulate the column density required to reproduce observed values of the grammage. This result suggests that observations of light element abundances might not imply cosmic rays are dynamically insignificant on small scales ($L \lesssim 100$ pc).

Up to now, we have not addressed the source of the fluctuations that scatter cosmic rays, leading to a diffusive random walk through the ISM. They could be generated by the cosmic rays themselves, through the streaming instability (Kulsrud & Pearce, 1969) or simply be part of an extrinsic turbulent cascade. There are significant differences: streaming transport introduces an extra heating term which adjusts how cosmic rays feed back on the thermal gas (Zweibel, 2013, 2017). Additionally,

opposing cosmic ray pressure and plasma density gradients can result in cosmic ray pressure plateaus, or bottlenecks (Wiener et al., 2017a; Bustard & Zweibel, 2021). However, to keep the problem simpler and better connect with past work, we neglect the streaming instability in this paper. For the impact of streaming on cosmic ray feedback in the turbulent ISM, we direct readers to Dubois et al. (2019) and Bustard & Oh (2023).

We compare and contrast four 3D turbulent box ($L \sim 100$ pc) simulations: one with no cosmic rays, one with constant diffusivity set to the canonical Milky Way value $\kappa_{\text{MW}} = 3 \times 10^{28} \text{cm}^2 \text{s}^{-1}$ (note this is slightly larger than recent observations suggest), another with constant diffusivity set to the critical value for maximum energy gain found in previous studies ($\sim 10^{26} \text{cm}^2 \text{s}^{-1}$), and one which uses a temperature dependent cosmic ray energy diffusion coefficient which switches between the two values. Other than cosmic ray energy and transport, the simulations have identical initial conditions. They even use the same random seed for turbulent driving. Using identical setups implies the differences we observe are the result of including cosmic ray energy and varying its transport. We run our simulations for more eddy turnover times than previous works, adding to the statistical robustness of our analysis. While the exact diffusion coefficients in some of our simulations may not apply to the Milky Way galaxy’s ISM, our work improves the physical understanding of cosmic ray feedback on a turbulent cascade.

This paper is structured as follows: in Section 3.3, we detail our computational methods and choice of simulation parameters; in Section 3.4 we show results from the simulations; in Section 3.5 we discuss our results, along with their significance; and

in Section 3.6 we provide conclusions and a short summary of the work. In Appendix 3.7.1 we provide some additional results, from lower resolution simulations, which add to the robustness of our conclusions.

3.3 Methods

In this section, we detail the methods used for cosmic ray hydrodynamics in MHD, discuss the limits of this implementation of cosmic ray hydrodynamics, our use of turbulent driving, the implementation of radiative cooling, the addition of a scalar dye for tracking cold gas formation, and our method for cosmic ray decoupling in cold gas. Finally, we derive and discuss our simulation parameters.

We use a recent release of the `Athena++` magnetohydrodynamic (MHD) code which evolves total cosmic ray energy and flux alongside the other hydrodynamic variables (Jiang & Oh, 2018; Stone et al., 2020). We start the simulation with a 3D homogeneous box of equal sides in each direction ($L = 100$ pc) with periodic boundary conditions. The free parameters of the initial setup are the temperature T and density ρ of the thermal gas, the plasma beta $\beta \equiv 8\pi P_g/B^2$ (where P_g is the thermal pressure and B is the magnetic field strength), and the cosmic ray beta $\beta_{\text{cr}} \equiv P_g/P_c$ (where P_c is the cosmic ray pressure). The magnetic field is initially uniform and directed in the along \hat{x} . The simulations presented here have a resolution $\Delta x \sim 0.39$ pc, but we include results from lower resolution simulations ($\Delta x \sim 0.78$ pc) in Appendix 3.7.1.

3.3.1 MHD & CR Hydrodynamics

The Athena++ implementation from Jiang & Oh (2018) solves the following equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad (3.1)$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) = 0 \quad (3.2)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot \left(\rho \mathbf{u} \mathbf{u} + \mathbb{1} \left(P_g + \frac{B^2}{2} \right) - \mathbf{B} \mathbf{B} \right) = \sigma_c \cdot \left(\mathbf{F}_c - \frac{4}{3} E_c \mathbf{u} \right) \quad (3.3)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left(\left(E + P_g + \frac{B^2}{2} \right) \mathbf{u} - (\mathbf{B} \cdot \mathbf{u}) \mathbf{B} \right) = \mathcal{L}(n, T) - \frac{1}{3} \mathbf{u} \cdot \nabla E_c \quad (3.4)$$

$$\frac{\partial E_c}{\partial t} + \nabla \cdot \mathbf{F}_c = \frac{1}{3} \mathbf{u} \cdot \nabla E_c \quad (3.5)$$

$$\frac{1}{V_m^2} \frac{\partial \mathbf{F}_c}{\partial t} + \frac{1}{3} \nabla E_c = -\sigma_c \cdot \left(\mathbf{F}_c - \frac{4}{3} E_c \mathbf{u} \right). \quad (3.6)$$

In the above equations, the variables evolved in time are the thermal gas density ρ , bulk velocity \mathbf{u} , gas pressure P_g , magnetic field \mathbf{B} , cosmic ray energy E_c , and cosmic ray energy flux \mathbf{F}_c . The total magnetohydrodynamic energy density (excluding E_c) is a combination of those variables: $E \equiv \rho u^2/2 + P_g/(\gamma_g - 1) + B^2/2$. Throughout this work the thermal gas is treated as ideal with ratio of specific heats $\gamma_g = 5/3$ and the cosmic rays are treated as an ultrarelativistic fluid with $\gamma_c = 4/3$. The radiative heating and cooling function $\mathcal{L}(n, T)$ depends on number density $n = \rho/\bar{m}$ (where \bar{m} is the mean particle mass) and gas temperature $T = P_g/(nk_B)$ (where k_B is the Boltzmann constant). This function is specified in Section 3.3.4.

For the evolution of cosmic ray energy, there are two other key variables. First, the modified speed of light V_m appears where the speed of light should be. For numerical accuracy, the only requirement on V_m is it needs to be the fastest speed in the simulation, which is essentially requiring:

$$V_m^2 > u^2 + \frac{P_g + P_c + \frac{B^2}{2}}{\rho}. \quad (3.7)$$

Second, we define the cosmic ray transport matrix σ_c to allow a split of parallel and perpendicular diffusion of cosmic rays:

$$\sigma_c = \frac{1}{\kappa_{\perp}} \mathbb{1} + \left(\frac{1}{\kappa_{\parallel}} - \frac{1}{\kappa_{\perp}} \right) \hat{b}\hat{b}. \quad (3.8)$$

The magnetic field direction \hat{b} is calculated locally, in each cell at each time, and is not in general the mean magnetic field direction. We set κ_{\perp} to be a small value, $10^{18} \text{ cm}^2 \text{ s}^{-1}$, so that perpendicular diffusion across a cell width will take $10\times$ the simulation run time. This choice of perpendicular diffusion coefficient is justified by modern cosmic ray transport models. Quasilinear diffusion theory predicts $\kappa_{\perp}/\kappa_{\parallel} \sim (r_g/l)^2 \sim 10^{-11}$ for gyroradius $r_g \sim 1 \text{ AU}$ and parallel mean free path $l \sim 1 \text{ pc}$ (Desiati & Zweibel, 2014). From another perspective, we could account for perpendicular transport due to unresolved fluctuations (below $\sim 1 \text{ pc}$) in the magnetic field. This method still finds a decreased perpendicular diffusion coefficient ($\kappa_{\perp}/\kappa_{\parallel} \propto M_A^n \lesssim (0.5)^n$ for $n = 2 - 4$ depending on models, see Matthaeus et al. (2003); Shalchi et al. (2006); Yan & Lazarian (2008)). Because we use an extremely low value of κ_{\perp} , cosmic ray energy barely diffuses perpendicular to the field lines over the course of

our simulation. We allow a more complex definition for the parallel transport κ_{\parallel} , which is detailed in Section 3.3.6.

3.3.2 Limits of the Fluid Prescription for Cosmic Rays

Since we use a fluid description of cosmic rays rather than a full kinetic description or evolution of particle motion, it would be counterproductive to resolve physical scales shorter than the mean free path of scatterings ($l \sim \kappa_{\parallel}/c$). For the largest κ_{\parallel} we consider, the length scale is $l \sim 0.32$ pc. The resolution of our simulation is $\Delta x = 0.39$ pc L_{100}/N_{256} where $N_{256} = N/256$ is the number of cells in each direction in terms of our base simulations and $L_{100} = L/(100\text{pc})$ is the simulation box length. The resolution we use is 256^3 ; higher resolution would not be meaningful when using the fluid treatment of cosmic rays. We also include 128^3 resolution results in Appendix 3.7.1.

Going beyond a fluid description, e.g. by running MHD-PIC (Particle In Cell, e.g. Sun & Bai (2023)) simulations, which use PIC for cosmic ray transport, could eliminate this restriction on resolution and would self-consistently evolve the cosmic ray distribution function. But that would further limit the time step beyond what we manage and require an exorbitant amount of computational resources to match our simulations' size.

3.3.3 Turbulent Driving

We drive our initially static system at large scales by injecting energy at wavenumbers $m = 1/L = k/(2\pi) = \{1, 2\}$ (equivalent to $L = 100, 50$ pc in physical units), with a

k^{-2} dependence. We use a constant total energy injection rate

$$\dot{E} = \frac{E_{\text{th}}}{\tau} = 5.5 \cdot 10^{48} \frac{\text{erg}}{\text{Myr}} \left(\frac{11 \text{ Myr}}{\tau} \right) T_4 n_0 L_{100}^3 \quad (3.9)$$

where $T_4 = T/(10^4 \text{ K})$, $n_0 = n/(1 \text{ cm}^{-3})$, and τ is the characteristic time to inject energy equal to the initial thermal energy in the simulation volume.

The **Athena++** turbulent driving module uses an Ornstein-Uhlenbeck process (Uhlenbeck & Ornstein, 1930) with a compressive-solenoidal splitting (Eswaran & Pope, 1988). In our case, the energy is injected as velocity perturbations using purely compressive driving, with a forcing function \mathbf{F} which is derived from the velocity perturbations under the assumption that the energy input rate, $\mathbf{u} \cdot \mathbf{F}$, is constant.

The forcing can be decaying (only injected at the start of the simulation), continuous (driven at every time step with a user defined time decorrelation), or impulsive (driven at a time interval set by the user). We use impulsive driving because of the small time step imposed by the modified speed of light V_m in Equation 3.6. The impulses are driven at a timestep of $\delta t = 0.005$ in computational units. This timestep is near the value of the timestep necessary to integrate fluid motion without the cosmic ray energy equations. Therefore we are doing the equivalent of continuous driving in a simulation without cosmic rays.

3.3.4 Radiative Cooling

We adopt the simple model of optically thin radiative cooling from Inoue et al. (2006) which accounts for the emission of $\text{Ly}\alpha$ photons (important in warm gas) and the




Figure 3.1: Comparison of energy sources and sinks in thermal gas phase space. Dashed black line shows where the gas is in thermal equilibrium ($\mathcal{L} = 0$ in Equation 3.10). Dashed white lines show constant temperature contours for $T = 10^{3.74}\text{K} = 5.5 \cdot 10^3\text{K}$, $T = 10^3\text{K}$, and $T = 10^2\text{K}$. Color map compares net radiative energy change (either loss via cooling or gain via heating) to the total turbulent energy we inject (Equation 3.9). The green dot shows where the simulations start, slightly below the equilibrium curve. Three phases (a warm phase, an unstable phase, and a cold phase) with different temperatures can exist in pressure equilibrium for pressures between $P/k_B \sim 10^3 \text{ K cm}^{-3}$ and $P/k_B \sim 10^4 \text{ K cm}^{-3}$.

emission of [CII] photons (important in cold gas). We assume heating as a result of the photoelectric effect on dust grains. These assumptions give the following radiative heating and cooling function:

$$\mathcal{L}(n, T) = n\Gamma - n^2\Lambda(T) \quad (3.10)$$

where the cooling function is

$$\Lambda(T) = \left(7.3 \cdot 10^{-21} \exp\left[\frac{-118400}{T + 1500}\right] + 7.9 \cdot 10^{-27} \exp\left[\frac{-92}{T}\right] \right) \frac{\text{erg cm}^3}{\text{s}} \quad (3.11)$$

and the heating rate is

$$\Gamma = 2 \cdot 10^{-26} \frac{\text{erg}}{\text{s}}. \quad (3.12)$$

The equilibrium curve in pressure-density phase space created by these heating and cooling terms is shown as a black dashed line in Figure 3.1. Using this radiative heating and cooling setup creates a gas which has two thermally stable phases (with

different temperatures) which can exist in pressure equilibrium with each other. There is also a thermally unstable phase in the same pressure range. This bi-stable model of the ISM captures the important physical process of thermal instability (Field, 1965) in its simplest form: the transition between a warm phase and a cold phase.

Figure 3.1 also shows a colormap comparing the net heating and cooling rate with the turbulent energy density injection rate. The radiative energy rate is calculated by taking the absolute value of $\mathcal{L}(n, T)$ in Equation 3.11, whereas the turbulent energy density injection rate is a constant in (P, n) space (take Equation 3.9, and divide by simulation size $L^3 = (100 \text{ pc})^3$). While near thermal equilibrium (the black dashed line) the turbulent energy injection dominates, its contribution is subdominant in most of the domain (except for the blue wedge at low densities and low pressures). Note that the turbulent energy input is given for comparison purposes only; it is not a heating mechanism in the same sense as the photoelectric effect on dust grains. When the only energy gains and losses are those represented by Equation 3.10, heating dominates below the thermal equilibrium curve and cooling dominates above it.

We start the simulation in thermally stable equilibrium with the gas in the warm phase at a temperature of 6388 K and a density of 1 cm^{-3} . The turbulent stirring drives compression and rarefaction, with some gas getting pushed into thermally unstable states. A portion of the unstable gas then collapses into small scale, high density clouds via the thermal instability.

3.3.5 Scalar Dye for Dense Gas Formation

Although we do not include self gravity or mass sinks representing star particles in our simulations, we do attempt to track the formation of cold gas which could eventually form stars. We do this by injecting a passive scalar dye at a constant rate of $1M_{\odot} \text{ Myr}^{-1}$ into cells which hit the temperature floor ($T_{\text{floor}} = 20 \text{ K}$). This rate comes from assuming a star formation rate of $1M_{\odot} \text{ yr}^{-1}$ in a galaxy volume of about 10^{12} pc^3 , but then adjusting for our simulation volume of 10^6 pc^3 .

The gas density in thermal equilibrium at the temperature floor $T = 20 \text{ K}$ is 252 cm^{-3} . For our $N = 256^3$ simulations, the gas mass in a single cell at that density and temperature is about $3.5M_{\odot}$. The Jeans mass at that density and temperature is about $50M_{\odot}$. Therefore, we would need to produce a contiguous agglomeration of ~ 15 cells in order to meet the conditions for gravitational instability.

Another consideration is self-shielding by molecular hydrogen in dense clouds. This effect changes the cooling function and allows more molecular gas to form, which could make it easier for a cloud to gravitationally collapse. Self-shielding becomes important at an atomic hydrogen column density of $N_H \sim 5 \cdot 10^{20} \text{ cm}^{-2}$ (Draine, 2011), which corresponds to a number density

$$n_H \sim \frac{N_H}{R} = 162 \text{ cm}^{-3} \frac{N_H}{5 \cdot 10^{20} \text{ cm}^{-2}} \frac{R}{1 \text{ pc}} \quad (3.13)$$

assuming a cloud of radius $R = 1 \text{ pc}$. This effect may therefore become important in the most central regions of any cold gas regions in our simulations. However, we certainly do not resolve these molecular regions, given our resolution of $\Delta x = 0.39 \text{ pc}$

in the 256^3 simulations. We do not include their formation or the effect hydrogen self shielding because of this resolution limit.

While we do not simulate star formation, the effects of self gravity, or the formation of molecular gas, the scalar dye traces dense gas formation over the simulation's evolution. Overall, it allows us to track how much total dense gas is formed in each simulation.

3.3.6 Two Zone Cosmic Ray Diffusion Coefficient

Our two-zone cosmic ray transport model is motivated by ion-neutral damping (Xu et al., 2016). We assume the cosmic rays are scattered by Alfvén waves from an extrinsic turbulent cascade. When the ionization of the ISM decreases in cold clouds, those Alfvén waves are damped by collisions between ions and neutrals (Kulsrud & Pearce, 1969; Kulsrud & Cesarsky, 1971), increasing the effective cosmic ray diffusion coefficient in that gas. As we mentioned in Section 3.2, the effect of an increased diffusion coefficient due to ion-neutral damping in cold gas has been already examined in the context of global simulations of cosmic ray feedback (Farber et al., 2018).

To implement the temperature dependent cosmic ray energy diffusion coefficient, we adjust the method used in Farber et al. (2018), which used a discontinuous, piecewise constant function for κ_{\parallel} . This prescription produces numerical instability if the discontinuity is too large. To increase numerical stability, we introduce a smooth function for the transition between two values of κ_{\parallel} in the cold and warm gas; κ_c and κ_w , respectively. To allow larger variations in κ we build the transition

in logarithmic space

$$\log_{10} \kappa_{\parallel}(T) = (\log_{10}(\kappa_c) - \log_{10}(\kappa_w)) f_c(T) + \log_{10}(\kappa_w)$$

using a switching function

$$f_c(T) = \frac{1}{2} \left(1 - \tanh \left(\frac{T - T_0}{\Delta T} \right) \right). \quad (3.14)$$

This function requires a transition temperature T_0 between the diffusion coefficient values, along with a transition width ΔT in temperature. Our implementation allows the diffusion coefficient to change by several orders of magnitude with more numerical stability.

We use 10^3K for transition temperature T_0 because it cuts through the middle of the unstable neutral medium in the pressure-density phase space (see the $T = 10^3\text{K}$ contour in Figure 3.1). We use a width of $\Delta T = 0.1T_0$ which is wide enough to allow for smooth transitions and steep enough to limit the number of cells with a diffusion coefficient in between the two constant values.

3.3.7 Parameter Choices

After considering all these different physical processes, we still have several parameters to set for the simulations. First, we set the gas density, temperature, plasma beta, and cosmic ray beta which describe the initially homogeneous fluid. Then we discuss possible values of the cosmic ray diffusion coefficient. Finally, we set the diffusion coefficient value in each simulation, including its value in the warm and

Simulation	T_i (K)	β	β_{cr}	$\kappa_w(\text{cm}^2 \text{s}^{-1})$	$\kappa_c(\text{cm}^2 \text{s}^{-1})$
No CR	6388	1	∞	-	-
Milky Way	6388	1	5	$3 \cdot 10^{28}$	$3 \cdot 10^{28}$
Critical	6388	1	5	$5.82 \cdot 10^{25}$	$5.82 \cdot 10^{25}$
Two Zone	6388	1	5	$5.82 \cdot 10^{25}$	$3 \cdot 10^{28}$

Table 3.1: Start-up parameters and cosmic ray energy diffusion coefficients for each simulation. Each of these setups were run at two resolutions (128, 256), with a modified speed of light $V_m = (1/60, 1/30)c$ respectively. See Appendix 3.7.1 for analysis of the $N = 128$ simulations.

cold gas (κ_w and κ_c in Equation 3.14) for the simulation where we use our two-zone model.

The parameters for each simulation are shown in Table 3.1. We set the initial plasma parameters as $T_i = 6388 \text{ K}$, $n = 1 \text{ cm}^{-3}$, $\beta = 1$, and $\beta_{\text{cr}} = 5$ (in the simulation without cosmic rays, $\beta_{\text{cr}} = \infty$). The choice of $\beta_{\text{cr}} = 5$ comes from the assumption that the cosmic ray energy increases with time and the radiative cooling causes thermal energy to decrease with time. Therefore, the β_{cr} should decrease with time. If we want there to be a significant amount of time where the thermal, magnetic, and cosmic ray energy are all a similar order of magnitude (i.e. $\beta_{\text{cr}} \sim 1$), then we have to start with a larger $\beta_{\text{cr}} > 1$. An approximate energy equipartition mimics the state of the ISM in the Milky Way.

For the diffusion coefficients, we start by considering the average value in the Milky Way. The canonical Milky Way cosmic ray diffusion coefficient value of $\kappa_{\text{MW}} \approx 3 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ comes from older measurements of the grammage X , which has dimensions of mass per area (Berezinskii et al., 1990; Jones et al., 2001). Grammage is a measure of the amount of material a cosmic ray passes through on its path through the ISM and is calculated by the observed ratio of primary and secondary cosmic

rays. The average diffusion coefficient $\bar{\kappa}_{\parallel}$ can be calculated from the grammage using

$$X \approx \bar{\rho} c \tau = \bar{\rho} c \frac{H^2}{\bar{\kappa}_{\parallel}} \rightarrow \bar{\kappa}_{\parallel} \approx \bar{\rho} c \frac{H^2}{X}. \quad (3.15)$$

where H is the root-mean-square displacement of cosmic rays from their source during their confinement time in the Milky Way, $\bar{\rho}$ is the average density of the ISM, and τ is the average time it takes the cosmic rays to diffuse. Taking displacement to be the half-thickness of the Milky Way $H \sim 250$ pc, mean density to be $\bar{\rho} = m_p \text{ cm}^{-3}$, and $X \sim 1 \text{ g cm}^{-2}$ for the low energy (~ 1 GeV) cosmic rays, gives an order of magnitude estimate of $\bar{\kappa}_{\parallel} \sim 3 \cdot 10^{28} \text{ cm}^2 \text{ s}^{-1} = \kappa_{\text{MW}}$. We stick with this older, canonical, estimate (Jones et al., 2001) for our simulations. However, more modern estimates suggest a slightly larger grammage $X \sim 3 \text{ g cm}^{-2}$ for the lowest energy cosmic rays, reducing the average diffusion coefficient by the same factor.

Another interesting value is the critical diffusion identified in Commerçon et al. (2019); Bustard & Oh (2022, 2023). Near the critical diffusion coefficient the cosmic rays will significantly disrupt the turbulent cascade. Bustard & Oh (2022) use analytical arguments to determine the dependence of the critical rate on various parameters, finding

$$\begin{aligned} \kappa_{\text{crit}} &= 0.15 v_{\text{ph}} L_0 = 0.15 L_0 \sqrt{\frac{P_g + P_B + P_c}{\rho}} = 0.15 L_0 \sqrt{\frac{k_B T}{m}} \sqrt{1 + \beta^{-1} + \beta_{\text{cr}}^{-1}} \\ \kappa_{\text{crit}} &= 5.82 \cdot 10^{25} \frac{\text{cm}^2}{\text{s}} \frac{L_0}{100 \text{pc}} T_4^{1/2} \sqrt{\frac{1 + \beta^{-1} + \beta_{\text{cr}}^{-1}}{3}} \end{aligned} \quad (3.16)$$

where L_0 is the outer scale of the turbulent cascade, v_{ph} is the phase velocity of

waves, T_4 is gas temperature in units of 10^4K , and 0.15 is a fitted coefficient found by Bustard & Oh (2022) using simulation results.

This diffusion coefficient is too low compared to κ_{MW} . If κ_{crit} were the actual diffusion coefficient in the Milky Way, then the grammage would be increased to $X \sim 500 \text{ g cm}^{-2}$. Even though κ_{crit} is physically unrealistic for an average value in the Milky Way, it is the other diffusion coefficient of interest because it leads to rapid cosmic ray energization (Bustard & Oh, 2022).

In addition to a base simulation without cosmic rays (`No CR`) we run simulations with constant diffusion coefficients at $\kappa_{\text{MW}} = 3 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ and $\kappa_{\text{crit}} = 5.82 \times 10^{25} \text{ cm}^2 \text{ s}^{-1}$, labelled `Milky Way` and `Critical` respectively (see Table 3.1). Finally, we run a simulation, labelled `Two Zone`, which uses our two zone model detailed in Section 3.3.6. So we can more easily compare with the other simulations, we set the warm gas cosmic ray diffusion coefficient (κ_w in Equation 3.14) to the critical value κ_{crit} in Equation 3.16, and the cold gas cosmic ray diffusion coefficient (κ_c in Equation 3.14) to the Milky Way value κ_{MW} .

3.4 Results

In this section we cover our simulations in depth. For a short summary of key takeaways, see Section 3.6. Each simulation is run for 100 Myr. That time frame allows the turbulent state to saturate and covers ~ 8 sound crossing times $\tau_s = L_0/\sqrt{k_B T_i/m} \approx 13 \text{ Myr}$. When calculating median quantities in Sections 3.4.2, 3.4.5, & 3.4.4 we set 25 Myr to be the start of the saturated state and take the median over




Figure 3.2: Energy evolution in each simulation. Solid black lines show the total thermal energy, dash-dotted green lines show the total cosmic ray energy, solid red lines show total kinetic energy, and dashed blue lines show total magnetic energy. All energies are normalized by the initial thermal energy ($E_{th,0} = 3.9 \cdot 10^{49}$ erg) in the simulation. While qualitatively similar, each run has unique characteristics. For example, the *Two Zone* and *Critical* simulations both have large increases in cosmic ray energy, but the cosmic ray energy in the *Two Zone* simulations grows at a slower rate (see Table 3.2).

a time range of $[25, 100]$ Myr with data at each 1 Myr. This saturated state after $t = 25$ Myr (shown as a dotted vertical line in the top plot of Figure 3.3) is apparent in Figures 3.2 and 3.3. However, the saturated state is also highly intermittent. When studying that state we default to taking the median over time, instead of averaging, because it is more robust to large variations (e.g. the set of numbers




Figure 3.3: *Top two plots:* Gas density clumping factor (as defined in Commerçon et al. (2019)) for each simulation. Solid black line shows the No CR simulation, dash-dotted green line shows the Milky Way simulation, dash-dotted red line shows the Critical simulation, and dashed blue line shows the Two Zone simulation. To make the plot clearer, we separate the Critical and Two Zone from the Milky Way and No CR runs. Initially, the runs track with each other almost exactly. The Milky Way run ends up slightly out of phase quickly, likely because the cosmic rays can rapidly react to any compression and escape along the magnetic field. The Two Zone and Critical runs track with each other almost exactly for 20 Myr. Then, the Two Zone run develops enough cold gas such that the cosmic ray transport differs significantly and more clumps of cold gas form. The Two Zone simulation then reaches clumping factors similar to those in the No CR and Milky Way simulations. *Bottom plot:* Average value of scalar dye in each simulation (with units of $M_{\odot} \text{pc}^{-3}$, see Section 3.3.5). The Critical run produces the least amount of scalar dye at all times. The Two Zone simulation tracks with the Milky Way simulation once it has developed cold gas, while the Milky Way and No CR simulations fluctuate near one another for the entire simulation run.

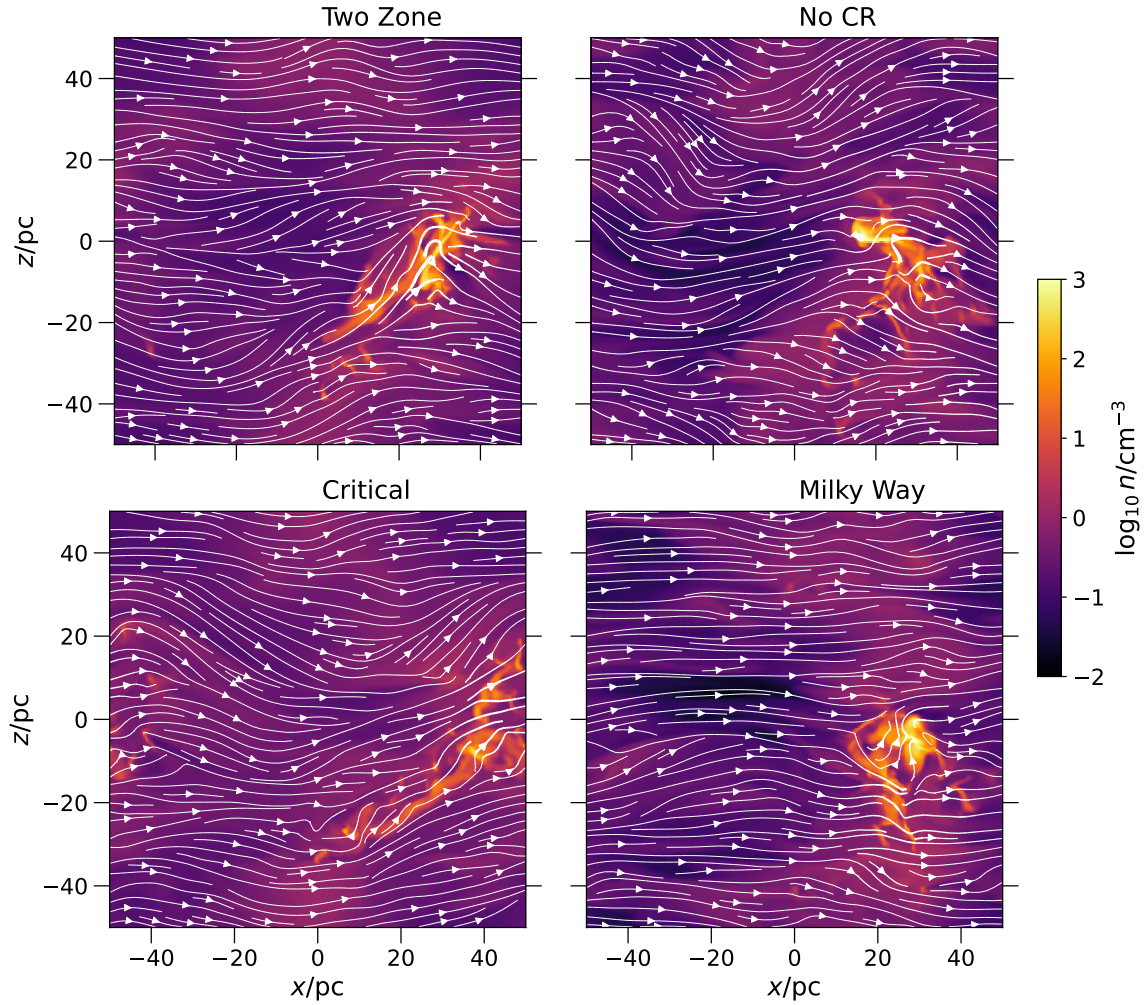


Figure 3.4: Slices from each simulation showing the gas number density. White lines show the magnetic field, and have a thickness proportional to the strength of the magnetic field (thicker line means stronger field). Each slice is at a time of $t = 85 \text{ Myr}$ and at $y = 0$. The density colormap is the same for each plot and is shown in the colorbar to the right.

Simulation	$\dot{E}_{\text{rad}}/\dot{E}_{\text{turb}}$	$\dot{E}_{\text{cr}}/\dot{E}_{\text{turb}}$
No CR	-1.3 ± 0.3	NA
Milky Way	-1.3 ± 0.5	0.010 ± 0.005
Critical	-1.1 ± 0.3	0.21 ± 0.07
Two Zone	-1.1 ± 0.5	0.17 ± 0.06

Table 3.2: Median time derivatives of radiative energy and cosmic ray energy over the saturated regime from $t = 25$ Myr to $t = 100$ Myr. The standard deviation during that time frame is shown after the \pm . By comparing the net cooling of the four simulations, it is apparent that having cosmic rays transported at the critical rate in warm gas siphons some energy, decreasing the energy lost to optically thin cooling. The **Two Zone** and **Milky Way** simulations have the largest variation in radiative energy, likely because of their formation of dense structures (see Section 3.4.2).

(1, 1, 11, 1, 1) has a median of 1 but an average of 3). Here, we show results from the $N = 256$ simulations, but these results are also apparent in the $N = 128$ simulations (see Appendix 3.7.1).

3.4.1 Energy Evolution

The evolution of thermal, cosmic ray, kinetic, and magnetic energy for each simulation is shown in Figure 3.2. The compressive driving initially heats the gas (increasing thermal energy) before some of the gas is pushed into a thermally unstable regime and collapses into cold clouds. This collapse leads to a loss of thermal energy as the gas cools via optically thin radiative cooling (see Section 3.3.4).

The median radiative energy loss rate during the saturated regime is shown in Table 3.2, normalized by the turbulent energy injection rate in Equation 3.9. The **No CR** and **Milky Way** simulations emit more radiation than the **Critical** and **Two Zone** simulations.

Given the nearly constant loss of radiative energy (see Table 3.2) along with the steady thermal, kinetic, and magnetic energy (see Figure 3.2), it is clear that most of

the energy injected is radiated away in the `No CR` and `Milky Way` simulations. The only difference in the `Critical` and `Two Zone` simulation is the steady increase of cosmic ray energy. Since the other energies reach a similar steady state as in the `No CR` simulation, there must be less radiation in the simulations with cosmic rays transported at the critical diffusion coefficient.

In Figure 3.2, the cosmic ray energy increases almost linearly in the three simulations which include cosmic rays, but the rate is dependent on the diffusion coefficient. In the `Critical` and `Two Zone` simulations, the increase in cosmic ray energy is faster than in the `Milky Way` simulation. In Table 3.2, we show the median rate of change of cosmic ray energy over the saturated regime for each simulation. The rate of increase in the `Critical` and `Two Zone` simulations is approximately $20\times$ the rate of increase in the `Milky Way` simulation. This difference is the result of using the critical diffusion coefficient in the warm gas $\kappa_w = \kappa_{\text{crit}}$. Because the warm gas is volume filling, κ_w is the diffusion rate for a majority of the simulation's volume. With a significant volume near the critical value identified in Bustard & Oh (2022) and Bustard & Oh (2023), there is cosmic ray energization (or reacceleration) which creates a larger energy sink.

One striking feature of Figure 3.2 is the nearly constant values of magnetic, kinetic, and thermal energy from $t = 25$ Myr to $t = 100$ Myr. For each simulation, we do statistical analysis over that time frame to better understand the effect of including cosmic rays. A caveat to this is the increasing cosmic ray energy in the `Critical` and `Two Zone` simulations. However, the energy growth rate is consistent over the same time frame, and is a result of siphoning of energy from the rest of

the ISM. The cosmic ray pressure gradients will be the important dynamical factor, and those can stay at the same order of magnitude while the total cosmic ray energy increases.

3.4.2 Structure Formation

Next, we examine the formation of cold, dense gas in each simulation. Commerçon et al. (2019) use a unitless clumping factor $\langle \rho^2 \rangle / \langle \rho \rangle^2$ to average over the volume of each of their simulations and quantify the appearance of dense structures. We calculate this clumping factor for each of our simulations, and show its evolution over time in Figure 3.3. The burstiness in time exhibited by the clumping factor highlights the necessity of frequent data dumps from the simulations. The lines in Figure 3.3 are made up of data points from every 10^{-2} Myr, as we set the clumping factor to be calculated as part of **Athena++** history output.

The simulations track with each other early on, but eventually depart from each other in Figure 3.3. This initial consistency is because the simulations all use the same random seed to generate turbulence. Once they depart from each other, the **Two Zone** simulation consistently produces larger clumping factors than the **Critical** simulation, even reaching a higher peak value than the **No CR** and **Milky Way** simulations. This increase compared to the **Critical** simulation reflects the ability of the **Two Zone** simulation to create larger dense structures than the **Critical** simulation. If we label any of the increases in clumping factor as “compressive episodes”, then it is clear those episodes occur at the same time in the **Critical** and **Two Zone** simulation. However, the adjustment of κ_{\parallel} in cold gas allows higher density structures

to form in the `Two Zone` simulation when compared to the `Critical` simulation.

The `Milky Way` and `No CR` runs also depart from exact agreement, even ending up completely out of phase in their compressive episodes after $t > 70$ Myr. However, the variation during these runs is similar. The early deviation of compressive episodes is likely caused by the rapid diffusion of cosmic rays out of compressed regions. This diffusion along the magnetic field allows cosmic ray pressure gradients to form perpendicular to the magnetic field and affect the evolution of the gas density. However, in further analysis of the simulations (e.g. Figure 3.5), we will see the `Milky Way` and `No CR` runs are very similar when averaged over time.

Additionally, we show the evolution of the total amount of scalar dye (χ) (see Section 3.3.5) in Figure 3.3. There is no decay or loss mechanism for χ , so it tracks the total amount of cold dense gas produced over the course of each simulation. The `Critical` simulation produces the least amount of cold dense gas at all times, whereas the `Two Zone`, `Milky Way`, and `No CR` simulations cross each other at different points. This crossing is caused by the intermittency of dense cloud formation, which is shown in the time evolution of the clumping factor in top plot of Figure 3.3. At the end of the simulation, the `Two Zone` simulation has produced over 20% more cold dense gas than the `Critical` simulation, and the only difference between those simulations was the increased diffusion coefficient in cold gas.

In Figure 3.4, we show slices of each simulation at the same positions ($y = 0$ pc) and time ($t = 85$ Myr). These slices show the gas density with overlaid magnetic field as white streamlines. The thickness of the lines corresponds to the magnetic field strength at that point. The resolution is large enough to illustrate some of

the smaller scale structure within the high density regions. The magnetic field is significantly distorted in the dense structures, where it is also the strongest. Overall, the magnetic field is predominantly in the \hat{x} direction in each simulation.

The slices illustrate some of the key differences between the simulations. The `No CR` and `Milky Way` simulations reach lower densities in the volume filling warm gas (those slices have more black regions corresponding to $n \sim 10^{-2}\text{cm}^{-3}$). The dense clump in the `Critical` simulation is not as distinct from the rest of the gas. The other simulations form more compact clumps along the magnetic field, with higher peak densities. This difference in structure likely results from the dense gas being less (or not) affected by cosmic ray pressure in the `Two Zone`, `Milky Way`, and `No CR` simulations.

3.4.3 Gas Density & Pressure Phase Space

Next we examine the phase space (gas pressure and number density) of each simulation. In Figure 3.5 we show the histogram of gas number density for each simulation. Each shaded region corresponds to the full range of histograms for a given simulation, between $t = 25$ Myr and $t = 100$ Myr with 1 Myr time steps. Each line shows the median histogram over that time frame. Overall, the `No CR` and `Milky Way` simulations match, and the `Critical` and `Two Zone` simulations match. The `Two Zone` and `Critical` simulations have very little gas below $n = 0.1\text{cm}^{-3}$, compared to the other simulations.

While the difference is not statistically significant, the `Critical` simulation has less high density gas, and more thermally unstable gas (near $n = 10\text{cm}^{-3}$). This

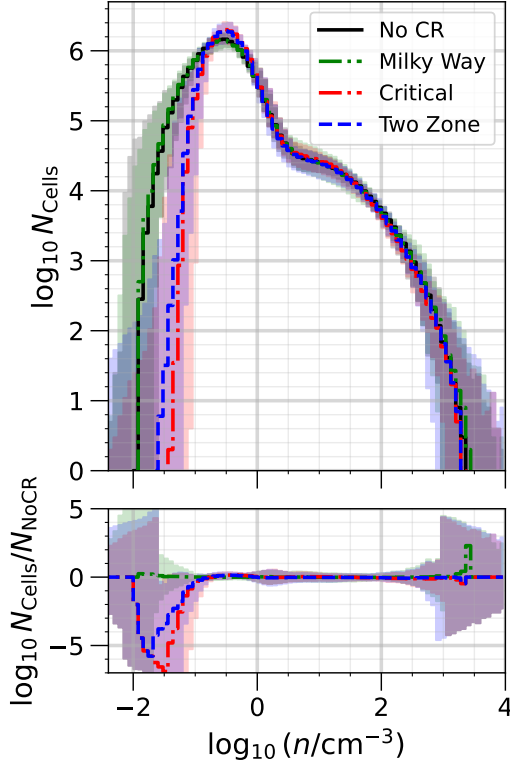


Figure 3.5: Number density histogram for each simulation. The lines show the median histogram over the entire period of saturation ($t = [25, 100]$ Myr). The shaded regions show the full range of histogram values for each simulation. The No CR and Milky Way simulation are nearly identical. The Two Zone and Critical simulations are also nearly identical, except the Two Zone simulation has more spread in the low density regime. The Two Zone and Critical simulations both produce less low density warm gas than the Milky Way and No CR simulations.

change likely results from the gas being more difficult to compress. With the smaller diffusion coefficient in the *Critical* simulation, the cosmic ray energy will not diffuse away fast enough during compressions. The cosmic ray pressure gradient then builds up and resists the compression. The compressibility of cosmic rays also limits the production of low density gas in the *Two Zone* simulation. However, because the diffusion coefficient changes and the cosmic ray pressure can rapidly diffuse along the magnetic field, the unstable gas in that simulation can collapse to the cold phase. This process leads to the *Two Zone* simulation's density histogram agreeing more with the *Milky Way* and *No CR* simulations in the higher density regime.

In Figure 3.6 we show the distribution of column densities, integrated over one

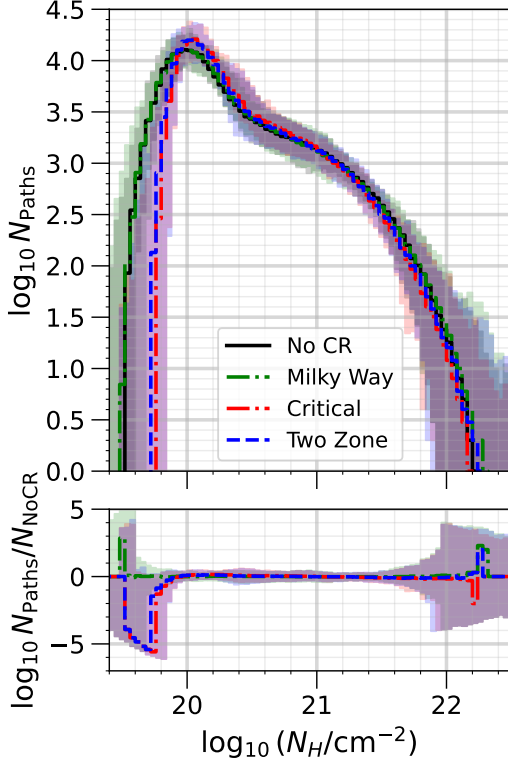


Figure 3.6: Column density histogram for each simulation. The line shows the median histogram over the entire period of saturation ($t = [25, 100]$ Myr). The shaded regions show the full range of histogram values for each simulation. Each column density value is calculated by integrating along either the \hat{x} , \hat{y} or \hat{z} axis. This integration gives 256^2 paths for each direction. The histogram is made from considering all 3×256^2 paths we get from integrating along each axis individually.

direction in each simulation. The distributions are, overall, similar to the number density histograms in Figure 3.5. However, N_H is an observable and we can now compare our distribution to realistic Milky Way column densities $N_H \sim 10^{21} \text{ cm}^{-2}$. Most integrated paths through our simulation do have a lower column density by an order of magnitude. These paths have a decreased column density because they either do not intersect a cold, dense cloud or cover enough distance. To match the Milky Way's column densities ($N_H = 10^{21} \text{ cm}^{-2}$) for a majority of paths through our simulation, we would need to integrate across several realizations of our simulation volume. Then, more paths would intersect a cold cloud or pass through enough warm gas to reach observed column densities. This point should be kept in mind

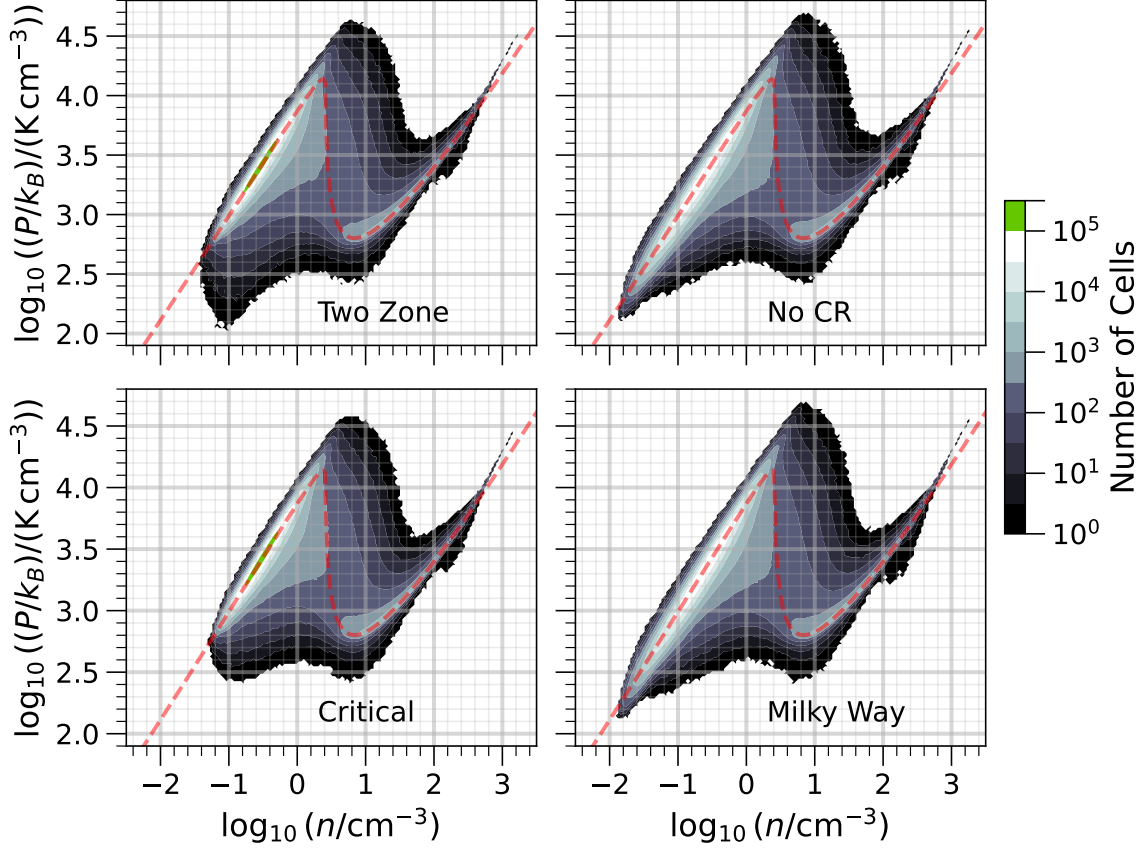


Figure 3.7: Pressure-Density phase diagrams for each simulation. The histograms illustrate the distribution of gas in phase space and along the thermal equilibrium curve ($\mathcal{L} = 0$), which is shown as a dashed red line. Each diagram uses the same coloring for the contours, which is illustrated in the colorbar to the right of the plots. We take the median of the histograms over the saturated time frame, between 25 Myr and 100 Myr. The `Two Zone` and `Critical` simulations have more warm gas concentrated above a number density of 10^{-1} cm^{-3} . Overall, the phase diagrams are similar across all the simulations.

when evaluating the verisimilitude of the cosmic ray transport model in the `Two Zone` simulation.

Next, we look at pressure-density phase space. We produce phase space histograms for each 1 Myr of evolution in the saturated regime, and then we take the median of those histograms to produce a single plot for each simulation. Those me-

Simulation	VFF (%)			MFF (%)		
	Cold	Warm	Unst.	Cold	Warm	Unst.
No CR	1.6 ± 0.2	92 ± 1	6 ± 1	61 ± 2	30 ± 1	9 ± 3
Milky Way	1.6 ± 0.3	93 ± 1	6 ± 1	60 ± 2	31 ± 1	10 ± 2
Critical	1.6 ± 0.2	93 ± 1	5 ± 1	55 ± 2	36 ± 1	9 ± 2
Two Zone	1.5 ± 0.2	93 ± 1	5 ± 1	56 ± 2	35 ± 1	9 ± 2

Table 3.3: Volume Filling Fraction (VFF) and Mass Filling Fraction (MFF) in each simulation, written as percentages. The quantities are calculated by averaging the MFF and VFF over the saturated time frame. We separate the phases by temperature: the cold phase is any gas below 200 K, the unstable (Unst.) gas is anything with temperatures between the warm and cold gas, and the warm gas is any gas with temperature greater than 5500 K. We also provide the standard deviation (1σ) of the filling fractions over that 75 Myr of evolution. The only significant change due to cosmic rays and their transport is in the MFF of cold and warm gas. This change is related to the lower production of cold dense gas shown in Figure 3.3.

dian phase space distributions are shown in Figure 3.7. For each simulation, we also plot the thermal equilibrium curve as a dashed red line. The curve is determined by solving Equation 3.10 with $\mathcal{L} = 0$.

Each simulation’s histogram is visually similar, with large populations of cold and warm gas in the thermally stable regimes of the equilibrium curve. There is more warm gas than cold gas in each simulation. We see the same cutoff of $n \sim 0.1 \text{ cm}^{-3}$ for the **Two Zone** and **Critical** simulations which was apparent in the density histograms (Figure 3.5). We also see the warm gas distribution in the **Two Zone** simulation extends to lower pressures at that density. This extension was also apparent at low densities in the density histograms, where the **Two Zone** simulations had slightly more low density gas than the **Critical** simulation. This extension is unique to the **Two Zone** simulation.

The turbulent driving is compressive, creating regions with increased density and pressure. This driving is more apparent in the **No CR** and **Milky Way** phase

diagrams. The extension toward the top of the phase diagram in these simulations is larger because the cosmic ray energy is taking away less of the injected energy during a compression. In the `Two Zone` and `Critical` simulations, the warm gas has the critical diffusion coefficient, allowing the cosmic ray energy to serve as a sink for some of the injected turbulent energy, reducing the amount of compression which takes place.

Finally, in Table 3.3, we show the volume filling fractions (VFFs) and mass filling fractions (MFFs) of the different gas phases in each simulation. The listed quantities are percentages, and they are calculated by taking the median over the saturated time frame. The listed errors are the standard deviation during the 75 Myr saturated time frame. Overall, the simulations all have the same VFF for each phase. The only impact cosmic rays have is on the MFF of cold and warm gas. This decrease in cold gas MFF (and corresponding increase in warm gas MFF) is likely tied to the lower total amount of cold gas formation over the course of the simulation (see scalar dye evolution in Figure 3.3). The change in MFF with no equivalent adjustment of VFF suggests the median density in cold gas is lower in the `Critical` and `Two Zone` simulations, and similarly the median density in the warm gas is higher. This difference also shows up in the phase diagrams (see Figure 3.7) and the density histograms (see Figures 3.5 & 3.6).

3.4.4 Power Spectra

To study the effect of the cosmic ray energy on the saturated turbulence, we examine how the power spectrum of kinetic energy varies between each simulation. In

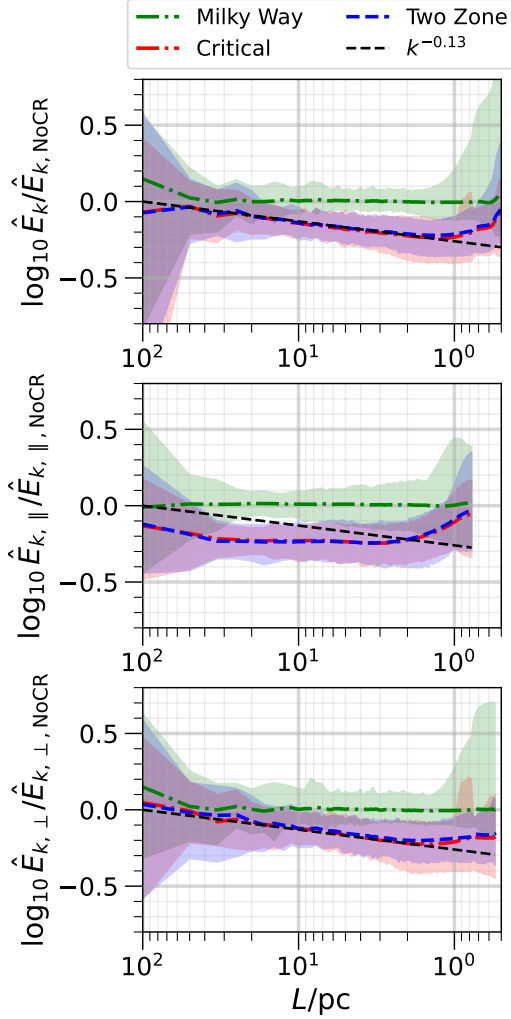


Figure 3.8: Kinetic energy per mass spectra from the Milky Way, Critical, and Two Zone simulations compared to the No CR simulation’s spectrum. Lines show the median spectrum, whereas the shaded regions show the full variation of the spectrum over the saturated timeframe ($t = [25, 100]$ Myr). *Top plot:* The total kinetic energy per unit mass. *Middle plot:* The kinetic energy per unit mass parallel to the mean magnetic field. *Bottom plot:* The kinetic energy per unit mass perpendicular to the mean magnetic field.

Simulation	Kinetic Energy Slope (α in $k^{-\alpha}$)
No CR	1.79 ± 0.07
Milky Way	1.78 ± 0.06
Critical	1.91 ± 0.08
Two Zone	1.91 ± 0.07

Table 3.4: Median inertial range spectral index for each simulation with 1σ variations. The dependence on κ_w is clear, as simulations **Critical** and **Two Zone** have a steeper slope and are the simulations with cosmic ray energy diffusion at the critical rate κ_{crit} through the warm gas. Steepening of the spectrum as a result of cosmic ray transport matches results from Bustard & Oh (2023) even though we have multiple gas phases. Additionally, the slope of our No CR run is similar to the slope of the equivalent run in Bustard & Oh (2023).




Figure 3.9: *Top plot:* Fourier transform (spectrum) of gas density for each simulation, scaled by a power law of $k^{0.95}$. The lines are the median spectrum, whereas the shaded regions show the full variation over the saturated time frame. *Bottom plot:* The ratio of other simulations' median spectrum to the No CR simulation spectrum.

Figure 3.8, we plot the total kinetic energy per mass ($v^2/2$) spectrum, the spectrum parallel to the mean magnetic field (meaning the \hat{x} direction, so $v_x^2/2$), and the spectrum perpendicular to the mean magnetic field. The solid lines correspond to the median spectrum over the saturated time frame and the shaded regions show the full variation. We divide the spectra from the Milky Way, Critical, and Two Zone simulations by the median spectrum of the No CR simulation.

Figure 3.8 illustrates that the Two Zone and Critical simulations have a significantly different turbulent cascade than the No CR simulation. The cosmic rays have little to no effect in the Milky Way simulation, where the cosmic rays are not well coupled to the outer eddy scale. Additionally, the middle and bottom panels of Figure 3.8 illustrate that the well coupled cosmic rays in the Two Zone and Critical simulations have an anisotropic impact on the turbulent cascade. While the parallel

kinetic energy at all scales is reduced, the perpendicular kinetic energy experiences an actual change in slope. The steepening of the slope is similar to some results in Bustard & Oh (2023), even though our simulations have multiple gas phases. The change in slope is approximately $k^{-0.13}$ in the `Critical` and `Two Zone` simulations. For completeness, we calculate the original slope of each spectrum’s inertial range, and these are listed in Table 3.4. For each 1 Myr time dump, we fit a power law $k^{-\alpha}$ over the length scales $L \in [5, 20]$ pc. We then take the median of the power law index α over the saturated time frame to get the median slope values and their 1σ variation.

In the top plot of Figure 3.9 we show the Fourier transform of the gas density in each simulation. In the bottom plot of Figure 3.9 we show the ratio of the median Fourier transform for each simulation, relative to the median of the `No CR` simulation. Overall, the simulations produce a similar characteristic density spectrum, with a power law $\sim k^{-0.95}$ in the inertial range ($L \in [5, 20]$ pc, see Figure 3.8). The `Critical` simulation has less dense gas (lower Fourier amplitudes) at all scales. However, as the bottom of Figure 3.9 shows, the variations in the inertial range are on the scale of 10% when comparing the simulations. From the top plot, it is clear that the variation over time for each simulation (the shaded regions) is even larger than the variation between simulations.

3.4.5 Diffusion Coefficient Statistics

In Figure 3.10, we show histograms of the ratio of the cosmic ray energy diffusion coefficient κ_{\parallel} (calculated from Equation 3.14) to the local critical diffusion coefficient




Figure 3.10: Histograms of $\kappa_{\parallel} / \kappa_{\text{crit}}$ for the Two Zone (top plot, blue lines), Critical (bottom plot, red lines), and Milky Way (bottom plot, green lines) simulations. At each cell, we calculate κ_{crit} from Equation 3.16 using the cell’s pressure and density values. We then take the actual diffusion coefficient in that cell and divide by the local critical coefficient. The histograms are shown at several times, and stars mark where $\kappa_{\parallel} = \kappa_{\text{crit}}$.

κ_{crit} (calculated at each cell using Equation 3.16) for the Two Zone, Critical, and Milky Way simulations. The blue lines in the upper plot are from the Two Zone simulation. The red lines in the lower plot are from the Critical simulation. The green lines in the lower plot are from the Milky Way simulation. Stars mark the value of the histogram at $\kappa_{\parallel} = \kappa_{\text{crit}}$.

The stars move down with time in both the Two Zone and Critical simulations.




Figure 3.11: Maximum distribution of average cosmic ray diffusion coefficient and gas column density along the mean magnetic field (in the \hat{x} direction) for the **Two Zone** simulation. For each 1 Myr time dump in the saturated state ($t = [25, 100]$ Myr), we calculate the histogram in $(N_H, \bar{\kappa}_{\parallel})$ space. To produce this figure, we take the maximum of each bin across time. The resulting distribution, shown here, presents the full range of diffusion coefficients and column densities produced in the simulation. In solid black lines, we plot contours of grammage (see Equation 3.15), illustrating that at certain times in the simulation we get realistic values (order unity in units of $\text{g} \cdot \text{cm}^{-2}$) of the grammage for some paths along the mean magnetic field. We also show the asymptotic values of the diffusion coefficient, $\kappa_{\text{crit}} = 5.82 \cdot 10^{25} \text{cm}^2 \text{s}^{-1}$ and $\kappa_{\text{MW}} = 3 \cdot 10^{28} \text{cm}^2 \text{s}^{-1}$, as horizontal dashed lines.

This decrease in the amount of gas with cosmic ray transport at the critical rate is the result of increasing cosmic ray energy in the simulations (see Figure 3.2). As the cosmic ray energy increases, the critical diffusion coefficient κ_{crit} increases. This increase leaves a majority of the gas with a diffusion coefficient below the critical value.

The distributions of the `Milky Way` and `Critical` simulations primarily reflect the gas temperature distribution, because they have constant diffusion coefficient κ_{\parallel} but a critical diffusion coefficient κ_{crit} dependent on gas temperature, cosmic ray pressure, and magnetic pressure (see Equation 3.16). The `Two Zone` simulation looks similar, except the cold gas component has been shifted to the right because of its larger diffusion coefficient value.

In Figure 3.11, we show the relationship between average cosmic ray diffusion coefficient $\bar{\kappa}_{\parallel}$ and column density N_H for paths along the mean magnetic field in the `Two Zone` simulation. For each 1Myr time dump in the saturated time frame ($t = [25, 100]$ Myr), we integrate along the \hat{x} direction, parallel to the mean magnetic field, to find the column density and average diffusion coefficient. For each time dump, this gives $256^2 = 65536$ paths. Then, we take the maximum across all times in each bin of $(N_H, \bar{\kappa}_{\parallel})$ space and plot those in Figure 3.11. Note that the `Milky Way` and `Critical` simulations would appear as single horizontal lines in this figure, located at their particular value of the diffusion coefficient (either κ_{MW} or κ_{crit} , shown as dashed lines in Figure 3.11).

A majority of the paths travel only through warm gas and keep the original critical diffusion coefficient value, showing up as a thin region of yellow bins in the bottom of Figure 3.11. However, for the paths which pass through some cold gas, the average diffusion coefficient changes significantly. This change also affects the grammage, which can be calculated using Equation 3.15. Contours of the grammage are shown as black lines in Figure 3.11, and some of the paths end up with realistic values of the grammage X (order unity in units of $\text{g} \cdot \text{cm}^{-2}$ according to observable

estimates, Jones et al. (2001); Evoli et al. (2019)).

Figure 3.11 is a lower limit on the distribution of average diffusion coefficient $\bar{\kappa}_{\parallel}$. Using mass-weighted average or actual path integration along the magnetic field, the paths would pass through more cold gas because there is some bending in the field lines (see Figure 3.4). This change in paths would cause more gas at high diffusivity because more magnetic field lines pass through the cold gas (where the field strength is larger). Additionally, our lines of integration do not pass through as much of the ISM as cosmic rays in the Milky Way (see column density histogram in Figure 3.6). If every path passed through a cold cloud, then there would be many more paths with a realistic mean diffusion coefficient $\bar{\kappa}_{\parallel}$.

Figure 3.11 helps illustrate the key takeaway from our use of a temperature dependent diffusion coefficient in the `Two Zone` simulation. The temperature dependence allows cosmic rays to significantly affect the turbulent energy cascade in warm gas (see Figure 3.8). However, this was previously only possible if the diffusion coefficient was too low compared to Milky Way observations (e.g. the results of Commerçon et al. 2019 and Bustard & Oh 2023). By having the diffusion coefficient increase in cold gas, we keep the effect on the cascade while the average diffusion coefficient, and resulting grammage, is closer to observational estimates.

3.5 Discussion

This section roughly follows the order of results presented in Section 3.4. For our final conclusions and summary, see Section 3.6.

Figure 3.2 together with Table 3.2 show that when cosmic rays are weakly coupled or not present at all, the system reaches a steady state with turbulent energy input balanced by compression and radiative cooling. When strongly coupled cosmic rays are added to the mix, as in the `Critical` and `Two Zone` simulations, some turbulent energy input goes into compressive heating of cosmic rays. In that case, less of the turbulent energy input is lost to radiative cooling.

However, there is rarefaction as well as compression, and one might ask why the cosmic rays gain energy instead of returning it back to the flow during rarefaction. A full answer to this question is beyond the scope of the present paper and we plan to address it in future work. However, we provide a partial answer by considering how cosmic ray diffusion affects small amplitude acoustic waves.

Consider the linearized cosmic ray pressure equation for waves propagating along the ambient magnetic field:

$$\frac{\partial \delta P_c}{\partial t} = -\gamma_c P_c \frac{\partial \delta v}{\partial x} + \kappa_{\parallel} \frac{\partial^2 \delta P_c}{\partial x^2}. \quad (3.17)$$

Assume the solution is a plane wave with frequency ω and wavenumber k . Cosmic rays do work on the gas at the rate $\dot{E} = -\delta v \partial \delta P_c / \partial x$ (note E refers to energy density, not energy). Using the linearized equation to determine δP_c in terms of δv allows us to calculate the energy change

$$\dot{E} = -\gamma_c P_c |k \delta v|^2 \frac{\kappa_{\parallel} k^2}{\omega^2 + \kappa_{\parallel}^2 k^4}. \quad (3.18)$$

Equation (3.18) shows that cosmic rays steadily extract energy from the wave, much

as a gas which conducts heat extracts energy from sound waves. The effect is maximized at $\omega/k = \kappa_{\parallel}k$. If we identify ω/k with the phase velocity and k with L^{-1} we recover the argument used to estimate the critical diffusivity κ_{crit} (see Equation 3.16) in Bustard & Oh (2022).

This effect has a counterpart in the kinetic theory description of cosmic ray propagation in the limit of short mean free path (Skillin, 1975). It is akin to first-order Fermi acceleration because the cosmic rays are scattering off converging fluctuations. Note that our simulations do not include actual shock acceleration. We would need strong shocks as well as an explicit term (e.g., shock acceleration term in Dubois et al. (2019)) which more accurately models the acceleration process on subgrid scales.

We reiterate that this linear perturbation argument only considers interaction with a simple longitudinal sound wave. This derivation does not account for the interaction of the cosmic rays with a turbulent cascade. The apparent dependence of the total energy of the thermal gas on the diffusion coefficient (κ_{\parallel}^{-1}) would change when considering the entire cascade. Specifically, the compression in multiphase turbulence is much more frequent than in adiabatic turbulence (Ho et al., 2023). We expect that the $\nabla \cdot v$ term (first term of right-hand side of Equation 3.17) will have a stronger effect in terms of energization of cosmic rays.

Returning to the effect of cosmic rays on an entire turbulent cascade, we see that during the formation of dense gas, the cosmic rays can diffuse independently of the gas, quickly redistributing energy from compressed regions. Those dense regions can collapse further because any cosmic ray pressure gradients have been decreased

or removed. This process is clearest in the evolution of the gas density clumping factor and scalar dye χ in Figure 3.3. We see the `Two Zone`, `Milky Way`, and `No CR` simulations have large spikes and variation in clumping as opposed to the evolution in the `Critical` simulation where cosmic rays provide significant pressure support against the formation of cool clouds. While cool clouds still form in the `Critical` simulation, there are fewer of them and they do not reach as high densities, which is clear from the `Critical` simulation’s smaller production of scalar dye χ . This process agrees with the results presented in Commerçon et al. (2019), which showed that decreasing the diffusion coefficient leads to smoother gas distributions and fewer cold clouds.

Moving on to cosmic ray transport through this multiphase ISM, the `Two Zone` simulation allowed us to examine how a temperature dependent diffusion coefficient would affect the grammage cosmic rays pass through on their way through the ISM. Observations of the primary-to-secondary ratio give stringent requirements on the grammage X that GeV cosmic rays will experience on average (Evoli et al., 2019, 2020). In our `Two Zone` simulation, the cosmic rays mainly experience a grammage set by the volume filling warm gas. However, Figure 3.11 shows that some paths will have physically realistic grammage values when only the cold gas has a realistic diffusion coefficient. Additionally, most paths through the simulation do not produce a realistic column density (see Figure 3.6). Performing the integration over multiple realizations of the simulation box would produce more sightlines with higher column densities. Those higher column density sightlines could also have a higher average diffusion coefficient $\bar{\kappa}_{\parallel}$ if they intersect cold clouds.

Therefore, for Milky Way column densities, the diffusion coefficient in the cold gas could become the dominant contributor to the mean cosmic ray diffusion coefficient. This conclusion adds to the importance of cosmic ray transport modelling which accounts for changes in temperature, ionization, or magnetic field strength. For example, in the realistic case of the diffusion coefficient in the warm gas and cold gas being $\kappa_{\parallel} \sim \kappa_{\text{MW}}$, the diffusion coefficient in hot and warm ionized gas could be a much lower value (possibly near κ_{crit}) without changing the average cosmic ray diffusion coefficient calculated from the observed primary-to-secondary ratio.

While not a physically significant result, we find that the numerical stability of the **Two Zone** simulation depends on how we calculated the diffusion coefficient. Applying the switching function in logarithmic space leads to a smoother transition in regions with a large temperature gradient. This may be useful for other studies which model inhomogeneous diffusion without the added computational expense of calculating κ_{\parallel} *ab initio* and on the fly from gas parameters.

We show the cosmic ray diffusion coefficient has a significant impact on the turbulent cascade in Section 3.4.4. This results from a simple relationship in the original CR+MHD equations. A rapid flattening and smoothing of cosmic ray pressure, caused by a large diffusion coefficient, minimizes the effect of cosmic rays because there is less time for the pressure gradient ∇P_c to change the motion of gas (Equation 3.3). A smaller diffusion coefficient means large peaks in the cosmic ray pressure can exist longer and create large pressure gradients which adjust the gas flow.

The steepening of the kinetic energy spectrum by a factor of $k^{-0.13}$ in Figure 3.8 is clearly the result of the critical diffusion coefficient in the warm gas because

the same results appear for simulations `Critical` and `Two Zone`. The reduction in parallel kinetic energy at all scales is because of the decreased diffusion parallel to the magnetic field, i.e. the $\kappa_{\parallel} \nabla^2 E_c$ term contained in the $\nabla \cdot \mathbf{F}_c$ term of Equation 3.5. The lower diffusion coefficient means there is enough time for the turbulent flows to drive more energization before the cosmic rays diffuse away.

The steepening in spectral slope for the perpendicular (and total) kinetic energy is more difficult to physically explain. If the $k^{-1.79}$ dependence of the total kinetic energy spectrum in simulation `No CR` is the result of our turbulent driving being too weak to develop a Kolmogorov ($k^{-5/3}$) or Iroshnikov-Kraichnan ($k^{-3/2}$) spectrum, then the further adjustment towards Burgers turbulence (k^{-2}) in the `Two Zone` and `Critical` simulations suggests transport at the critical diffusion coefficient may increase the frequency of ‘shocks’ in the direction perpendicular to the magnetic field. Although, we see little to no formation or propagation of shocks. Instead, it is possible the sharp density discontinuities between the cold and warm phases have an effect on turbulent transport similar to that of shocks in Burgers turbulence (instantaneous transfer of energy from large to small scales rather than a true cascade).

A secondary point from the examination of the kinetic energy spectra is that the slope of the `No CR` run is similar to the slope found in Bustard & Oh (2023), despite our inclusion of optically thin radiative cooling. This similarity is likely because motions on inertial scales ($20 \text{ pc} > L > 5 \text{ pc}$) are mostly determined by the warm gas in our simulations. This warm gas follows an equation of state similar to the simple adiabatic law in Kritsuk et al. (2017) & Bustard & Oh (2023). To see the differences created by radiative cooling, we would need to extend the inertial scale below 1 pc.

This extension requires increasing the resolution of our simulations. However, an increase in resolution will make our treatment of diffusive transport of E_c inaccurate (see Section 3.3.1). This limitation highlights the importance of future simulations which evolve the cosmic ray distribution function and/or actual particle trajectories (e.g. using methods like MHD-PIC, Sun & Bai 2023).

For clarity, we would like to quickly highlight the key differences and similarities between our study and previous studies of Commerçon et al. (2019), Dubois et al. (2019), Bustard & Oh (2022), and Bustard & Oh (2023). First, no other study has quantified and shown the anisotropic nature of cosmic ray feedback on a multiphase ISM turbulent energy cascade (see Figure 3.8). Second, our `Two Zone` simulation, which incorporates a temperature dependent diffusion coefficient (motivated by the global simulations of Farber et al. 2018), is unique; it has no counterpart in previous studies. Thirdly, the idea that cosmic rays can take energy which would otherwise be radiated during thermal instability (see Table 3.2) has not been shown previously.

There are clear similarities between our work and previous works. Our `No CR`, `Milky Way`, and `Critical` simulations are directly comparable to some simulations in Commerçon et al. (2019). Our turbulent cascade results are similar to those in Bustard & Oh (2023). However, our simulations are run for a longer simulation time and on larger scales than Commerçon et al. (2019), and Bustard & Oh (2023) did not incorporate a multiphase ISM or thermal instability. Therefore, even where our results are similar, our work is different such that it adds to the robustness of some key ideas and conclusions in Commerçon et al. (2019) and Bustard & Oh (2023). Overall, we view our work as an extension of these previous works, while

also illustrating new results and possibilities.

Prior to concluding, it is important to note that a significant caveat to our results is our simulations are longer than the cosmic ray confinement time in the galaxy. Therefore, most of the cosmic ray energy should escape the simulation box. Over that time frame, cosmic ray losses to hadronic and Coulomb interactions could also be significant on the time scales we simulate. Using a hadronic loss term integrated over cosmic ray energy from Enßlin et al. (2007); Bustard & Zweibel (2021), the timescale for losing cosmic ray energy to hadronic interactions with the thermal gas is

$$\tau_{\text{Hadron}} = \frac{E_{\text{CR}}}{\dot{E}_{\text{CR}}} = cn\bar{\sigma}_{pp}K_p = 66 \text{ Myr}. \quad (3.19)$$

We decided to not include the effect of hadronic losses, or a realistic confinement time, to focus on the coupling between the thermal gas and cosmic ray fluid already inherent in Equations 3.1 to 3.6. Since we did not include these effects, it is important to note the significant cosmic ray energy in the **Critical** and **Two Zone** simulations would likely escape the galaxy, rather than being kept in the ISM, because of the small confinement time of cosmic rays. If the confinement time were longer than the hadronic loss time, then the energy would end up in photons emitted during hadronic interactions.

An additional caveat is our focus on diffusive and advective transport of cosmic rays. We do not include streaming transport, nor its heating effect. Some simulations in Bustard & Oh (2023) examine the impact of streaming on the turbulent cascade. However, for the high densities inside the galactic disk ($n > 1\text{cm}^{-3}$), it is expected that cosmic rays are transported diffusively (Armillotta et al., 2021, 2022; Thomas

et al., 2023; Armillotta et al., 2024). Therefore, focusing on diffusive transport in this study is a reasonable reduction in scope.

3.6 Conclusions

We presented simulations of a bi-stable ISM including the effects of cosmic ray feedback for three different cosmic ray diffusion coefficient models, along with a simulation with no cosmic ray energy. We used a long baseline for statistical analysis of the saturated turbulent state. We showed that cosmic rays can decrease the amount of energy radiated away via optically thin cooling. This energy siphoning is mediated by a process similar to first order Fermi acceleration in regions with a non-zero divergence of the velocity field. We illustrated the impacts of a temperature dependent (as a proxy for ionization dependence) cosmic ray diffusion coefficient on the formation of cold dense gas and cosmic ray grammage. We re-examined the adjustment of the turbulent energy cascade by cosmic rays detailed in Bustard & Oh (2022, 2023) and find the cosmic rays have an anisotropic effect on the cascade.

From these simulations, our analysis, and our discussion, the key conclusions are:

1. Cosmic ray transport at the critical rate allows cosmic rays to remove energy from collapsing gas via first order Fermi acceleration, decreasing the total radiated energy (see Table 3.2).
2. Increasing the diffusion coefficient only within cold gas allows $\sim 20\%$ more cold, dense gas to form over time than in a simulation with a constantly low diffusion coefficient (see Figure 3.3).

3. Even when the volume filling warm gas has a small diffusion coefficient some paths through our simulation box have an average grammage near Milky Way values ($\gtrsim 1 \text{ g cm}^{-2}$, see Figure 3.11). Given that our column densities are low (see Figure 3.6), we actually underestimate the number of paths through the ISM with a larger diffusion coefficient. This underestimation also applies to the number of paths with realistic grammage.
4. Cosmic ray energy transported via diffusion at the critical rate identified in Bustard & Oh (2023) decreases kinetic energy parallel to the mean magnetic field at all scales by nearly a factor of 2, while also changing the spectral slope of kinetic energy perpendicular to the mean magnetic field by a factor of $k^{-0.13}$ (see Figure 3.8).

Acknowledgments

We would like to thank Ryan Farber and Chad Bustard for sharing their expertise and advice during the course of this work. We also acknowledge Peng Oh for extensive discussions during the final stage of the paper. We thank Michael Halfmoon for granting additional NERSC time needed for this project. Finally, we would like to thank the referee for many useful comments and suggestions which significantly improved this work.

RH acknowledges funding from NASA FINESST grant 80NSSC22K1749 and NSF grant AST-2007323 during the course of this work. Research presented in this article was supported by the LDRD program of LANL with project # 20220107DR (KWH)




Figure 3.12: Same as Figure 3.2, but for the lower resolution $N = 128$ simulations.

& 20220700PRD1 (KHY), and a U.S. DOE Fusion Energy Science project. This research used resources provided by the LANL Institutional Computing Program (y23_filaments), which is supported by the DOE NNSA Contract No. 89233218CNA000001. This research also used resources of NERSC with award numbers FES-ERCAP-m4239 (PI: KHY) and m4364 (PI: KWH).

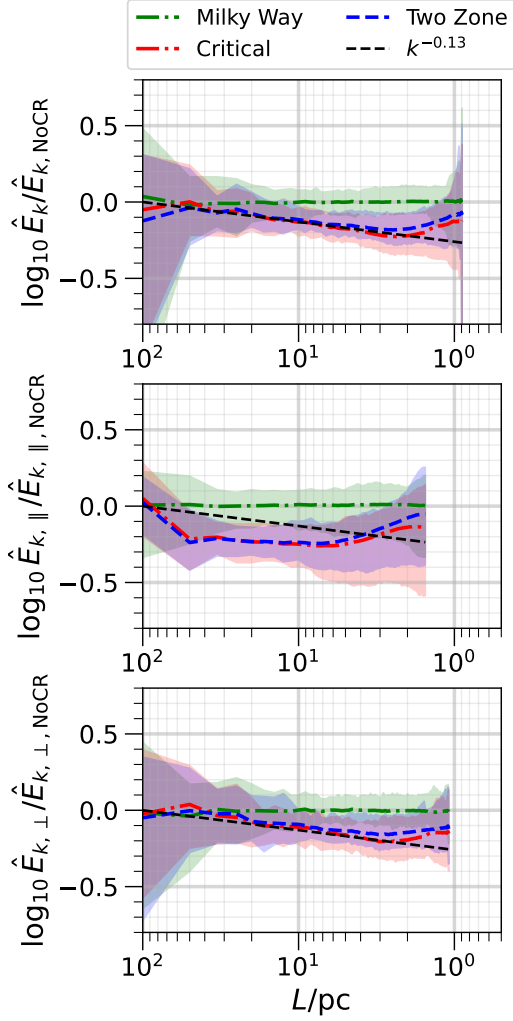


Figure 3.13: Same as Figure 3.8, except for the lower resolution $N = 128$ simulations.

3.7 Appendix

3.7.1 Lower Resolution Results

We performed lower resolution simulations ($N = 128$, $\Delta x \sim 0.78$ pc) with the same setup as the simulations listed in Table 3.1. The results exhibit little to no difference from the higher resolution simulations ($N = 256$, $\Delta x \sim 0.39$ pc), illustrating the

robustness of our conclusions.

Figure 3.12 shows the evolution of energy in the $N = 128$ simulations. The time evolution of each energy is similar to the $N = 256$ simulations, with the **Two Zone** and **Critical** simulations having significant cosmic ray energization.

Figure 3.13 shows the kinetic energy cascade, and its parallel and perpendicular components. Even in the $N = 128$ simulations with a smaller inertial range, the change in slope is the same for the overall cascade and for the parallel and perpendicular components.

4 COSMIC-RAY FEEDBACK ON CGM/ICM TURBULENCE

Wherever there exists a difference of temperature... it is possible to have also the production of impelling power.

— *Reflections on the Motive Power of Heat* (1824) by Sadi Carnot.

This chapter is a modified copy of the paper Habegger et al. (2025), titled *The Impact of Cosmic-Ray Transport on the Gamma-Ray Luminosity of Diffuse Gas*, currently under review in the *Astrophysical Journal*. No text has been changed; the figures and equations have just been reformatted to accommodate the submission of this thesis document.

I deeply appreciate the contributions of my co-authors Dr. Mateusz Ruszkowski and Dr. Ellen Zweibel, to this work.

4.1 Abstract

In astrophysical media, relativistic particles (cosmic rays) collide with thermal particles and produce γ -rays through pion decay. Therefore, observations of γ -rays from diffuse gas provide an opportunity to study the distribution of cosmic rays in different astrophysical environments. The γ -ray luminosity of the plasma depends on where cosmic rays are: if they are in denser gas, they produce more γ -rays. In this work, we study how different cosmic-ray transport mechanisms impact the γ -ray luminosity of a turbulent, multiphase medium formed from an initially diffuse medium ($n \sim 10^{-2}\text{cm}^{-3}$). Overall, we find that cosmic rays must escape cold dense

regions in order to produce less γ -ray emission and be consistent with observations. Our simulations with fast transport mechanisms (either diffusion or streaming) are degenerate: they each produce a lower γ -ray luminosity compared to slow transport simulations, by two orders of magnitude. This result means that fast transport (particularly in dense clumps) is necessary for simulations to agree with weak detections and non-detections of γ -ray emission from diffuse media, like the CGM or ICM. We show that turbulent reacceleration is a key driver of this γ -ray emission, because it increases the cosmic-ray energy density. Additionally, the correlation of cosmic-ray energy density and gas density causes observational diagnostics to give over-estimates (of the order $\sim 1 - 5$ times) of the true volume-averaged cosmic-ray energy density. Finally, we have also found a new reacceleration mechanism: the condensation mode of thermal instability re-accelerates cosmic rays and is more efficient than turbulent reacceleration.

4.2 Introduction

The most common cosmic rays are protons with energies $\sim 1 - 10$ GeV. They make up a majority of the total cosmic-ray energy density in our local interstellar medium (ISM) (Corti et al., 2016; Evoli et al., 2019). While detectors like Voyager and AMS directly detect them in the solar neighborhood and provide this estimate, it is challenging to analyze cosmic rays outside the nearby ISM.

Instead, studies of cosmic rays on large scales, in the diffuse circumgalactic medium (CGM) and intracluster medium (ICM), rely on observing the radiation

emitted as a result of cosmic rays interacting with thermal plasma (see recent reviews by (Owen et al., 2023) and (Ruszkowski & Pfrommer, 2023)). As cosmic-ray protons pass through thermal plasma, they can hit other nuclei and trigger a hadronic interaction which produces pions (Enßlin et al., 2007; Guo & Oh, 2008). Some of these pions decay into γ -rays, which can propagate through the gas and eventually reach us. Only cosmic rays with an energy above a threshold energy $E_{\text{th}} = 1.22 \text{ GeV}$ in the rest frame of the observer can trigger this interaction, but the common $\sim \text{GeV}$ protons meet this criterion. Since they are the most common, they dominate this γ -ray emission. As a result, observations of γ -ray emission can be used to constrain the total cosmic-ray energy density in our own galaxy, as well as in galaxies and galaxy clusters far away from us (Fabian et al., 2006; Ackermann et al., 2014). In this work, we focus on the denser, cool-core cluster regime matching the inside of the Perseus cluster. Because of the relatively higher densities, it is more likely that diffuse (not point source emission from an AGN Abdo et al. (2009)) γ -ray emission will eventually be detected in this regime.

However, there is a problem with directly calculating the cosmic-ray energy density from the γ -ray emission. The γ -ray luminosity L_γ is determined by the product of the gas density and the cosmic-ray energy density. The average of a product ($\langle fg \rangle$) can deviate significantly from the product of averages ($\langle f \rangle \langle g \rangle$) for a multiphase gas (e.g., Boettcher et al. (2013)). When averaging over a volume of multiphase gas, we separate the total luminosity into three quantities: the average cosmic-ray energy density, the average gas density, and the correlation of cosmic-ray energy density and gas density. The correlation of the cosmic-ray energy density and gas density

(quantified in Section 4.3.2) depends on the thermodynamic, magnetic, and turbulent properties of the gas, and the cosmic-ray transport through the gas. We separate the contribution of this correlation, showing that it causes an overestimate of the true volume-averaged cosmic-ray energy density. However, the increase in cosmic-ray energy density due to turbulent reacceleration has an even stronger effect on γ -ray emission.

The interaction of cosmic rays with a turbulent medium is a timely problem, as Ptuskin’s acceleration of cosmic rays by long-wavelength turbulence (Ptuskin (1988), hereafter P88) is being reexamined with modern computational techniques (Bustard & Oh (2022, 2023), hereafter BO22, BO23, Habegger et al. (2024), hereafter H24, Sampson et al. (2025)). This acceleration is capable of increasing the cosmic-ray energy, and therefore the γ -ray luminosity. We will review the expected growth rates, but the important takeaway from previous works is that the mechanism should be most optimal and have its largest impact in the CGM, as opposed to the ISM and ICM (BO23).

In this work, we examine how cosmic-ray transport influences the γ -ray luminosity of diffuse, magnetized, and turbulent plasma. We show that transport influences the γ -ray luminosity through two mechanisms: (1) increasing the average cosmic-ray energy density through turbulent reacceleration, and (2) changing the correlation of cosmic-ray energy density and gas density. If the cosmic rays do not escape dense gas, then the γ -ray luminosity rapidly increases as a result of both cosmic-ray acceleration and correlation with dense gas. In that case, our simulations would disagree with null detections of diffuse γ -ray emission from diffuse gas (e.g., in galaxy clusters).

Therefore, we can place limits on the cosmic-ray transport in diffuse media such as the CGM and ICM. We find that the diffusion coefficient must be $\gtrsim 10^{30} \text{ cm}^2\text{s}^{-1}$ or that streaming (self-confinement model of cosmic-ray transport) must dominate transport in diffuse gas. We also find that reacceleration by condensing, cold clouds is more efficient at reacceleration of cosmic rays than turbulent reacceleration in a single-phase medium.

In Section 4.3, we provide background on the calculation of the γ -ray luminosity, turbulent acceleration, and correlation. In Section 4.4, we provide details about the simulations we run with Athena++, including our choice of cosmic-ray transport and heating/cooling physics. In Section 4.5, we present our results before discussing them in the context of previous work in Section 4.6. Finally, in Section 4.7 we summarize the outcomes of this work.

4.3 Physical Background

The γ -ray luminosity of a plasma with nucleon density n and cosmic-ray energy density E_c is (Enßlin et al., 2007; Guo & Oh, 2008):

$$L_\gamma \propto \int_V d^3x \eta_\pi \propto \int_V d^3x n E_c \quad (4.1)$$

where η_π is the energy lost per unit volume, per unit time by cosmic rays due to hadronic interactions, and V is the volume of plasma. We stick to a proportionality because the constant which would be necessary is set by many assumptions about the cosmic-ray spectrum (Enßlin et al., 2007; Guo & Oh, 2008). However, regardless

of the microphysics, the dependence of L_γ on n and E_c is the same.

Assuming a constant (integrated over the cosmic-ray energy spectrum) L_0 , which captures those assumptions, the total γ -ray luminosity from that volume V of gas is

$$L_\gamma = \frac{L_0}{V} \int_V d^3x \left(\frac{E_c}{\text{eV cm}^{-3}} \right) \left(\frac{n}{\text{cm}^{-3}} \right) = L_0 \left\langle \left(\frac{E_c}{\text{eV cm}^{-3}} \right) \left(\frac{n}{\text{cm}^{-3}} \right) \right\rangle \quad (4.2)$$

where we have used angle brackets to denote volume averaging. From Equation 4.2, it becomes easier to analyze the γ -ray emission through a statistics lens: the emission is determined by the correlation of the two variables. After Equation 4.2, we express E_c and n in units of eV cm^{-3} and cm^{-3} .

If the density and cosmic-ray energy density are independent variables, then the luminosity is trivial to calculate:

$$L_\gamma \rightarrow L_0 \langle E_c \rangle \langle n \rangle. \quad (4.3)$$

However, we know that cosmic-ray energy and electron density are not independent variables. They are coupled together by a variety of physical processes.

Depending on the cosmic-ray transport, magnetic field structure, and dynamics of the diffuse gas, the correlation varies significantly. For the sake of simplicity, we define a normalized expectation of the product of cosmic-ray energy density and electron density:

$$\mathcal{C} = \frac{\langle E_c n \rangle}{\langle E_c \rangle \langle n \rangle}. \quad (4.4)$$

Then the total luminosity is just

$$\frac{L_\gamma}{L_0} = \langle E_c \rangle \langle n \rangle \mathcal{C}. \quad (4.5)$$

The reason we formulate the luminosity in the way of Equation 4.5 is because it isolates the three factors in determining the γ -ray luminosity. First, the average cosmic-ray energy density $\langle E_c \rangle$. Second, the average density $\langle n \rangle$. Third, the correlation of thermal plasma and cosmic-ray energy density, which is captured in the term \mathcal{C} . Each of these factors could be adjusted independently from one another, providing a useful separation of the possible physical mechanisms behind an increase (or decrease) in γ -ray luminosity.

It is important to note that all previous estimates assume $\mathcal{C} = 1$. For global cluster observations, profiles relating $\langle E_c \rangle$ to $\langle n \rangle$ are prescribed or modeled. We are suggesting that there is an additional factor beyond that global density dependence, which is affected by the interaction of cosmic rays and turbulence. The correlation factor illustrates how far off the true volume-averaged cosmic-ray energy density is from the observed value. Observations of γ -ray luminosity probe the combined quantity $\langle E_c \rangle \mathcal{C}$, not just $\langle E_c \rangle$.

For a set volume of diffuse gas, the number density $\langle n \rangle$ will change very little unless a runaway cooling process drives the system to form dense, neutral gas. Even then, the integrated number density of particles in a closed system is constant because it is proportional to the total mass. Instead, for a periodic turbulent box as we study here, $\langle n \rangle$ does not change, while the other two factors, $\langle E_c \rangle$ and \mathcal{C} , vary significantly as a function of the cosmic-ray transport and turbulent properties of the diffuse gas.

4.3.1 Evolution of Average Cosmic-Ray Energy Density

First, we discuss how the average cosmic-ray energy density $\langle E_c \rangle$ can change. There could be actual cosmic-ray sources (e.g., shocks) that accelerate thermal ions. Cosmic-ray losses, such as those due to hadronic or Coulomb interactions, can reduce the average cosmic-ray energy density. Finally, turbulent reacceleration could cause cosmic-ray energy to increase (P88).

In this work, we focus on the evolution of a periodic box of diffuse gas which is not actively being shocked or injected with cosmic rays. Our simulations start with a set amount of cosmic-ray energy, and do not include the sources themselves. This treatment allows us to more easily track the interaction of the background cosmic rays with the entire turbulent cascade, agnostic to the sources and rate of cosmic-ray production from the many possible accelerators.

Hadronic and Coulomb interactions are not dominant enough in diffuse gas to completely erase the cosmic-ray background. The loss timescale due to hadronic interactions is ~ 5 Gyr at the initial density and cosmic-ray energy density we consider. Since this timescale is only about five times the simulation run times, we include the loss of energy to hadronic interactions. But this loss plays a subdominant role in our simulations because the loss time is so long.

Finally, turbulent reacceleration is an avenue by which cosmic-ray energy density can increase in any chunk of diffuse gas. If the cosmic rays probe turbulent eddies on a large scale, but decouple from smaller-scale eddies, then the cosmic rays can sap energy away from the turbulent cascade at the length scale where they begin to decouple (P88). The process derives from the different length scale dependence of the

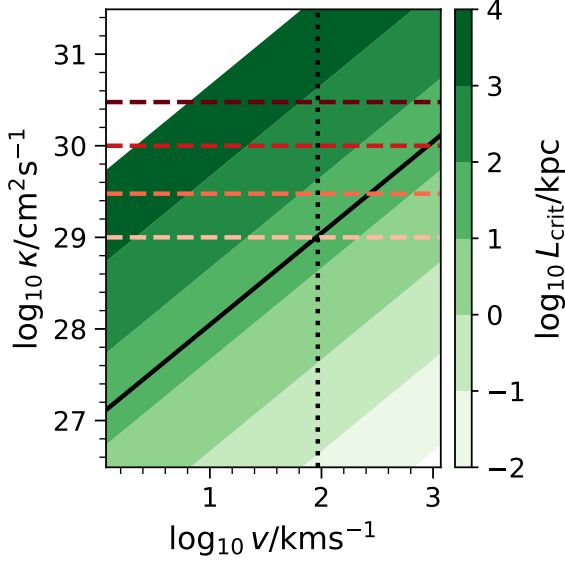


Figure 4.1: Contours of the critical outer length scale L_{crit} , where cosmic-ray reacceleration by long-wavelength turbulence is optimal. Green contours show the dependence of this length scale on the cosmic-ray diffusion coefficient κ and the phase velocity of the turbulence v . The solid black line shows the outer scale of turbulent driving used in this work, 24 kpc, and the dotted black line shows the sound speed of the initial homogeneous medium. The dashed lines correspond to each of the diffusion cases we examine (a different hue of red for each diffusion coefficient we examine). The smallest diffusion coefficient we use, $10^{29} \text{ cm}^2 \text{ s}^{-1}$, is at the optimal value for our simulation setup.

eddy turnover time ($\tau \sim L/v_{\text{ph}}$) and the cosmic-ray diffusion timescale ($\tau_c \sim L^2/\kappa$). The acceleration is optimized when the outer eddy scale is at a critical length scale L_{crit} . This critical length is related to the diffusion coefficient of cosmic rays κ and the turbulent phase velocity ($v_{\text{ph}} = \sqrt{(P_g + P_c + B^2/8\pi)/\rho}$) by

$$\kappa \approx 0.15 L_{\text{crit}} v_{\text{ph}}, \quad (4.6)$$

which comes from the numerical experiments in BO22.

In Figure 4.1, we show a contour map of critical length scale L_{crit} as a function of diffusion coefficient and turbulent phase velocity. We also plot the outer turbulent driving scale (24 kpc) we use in our simulations as a solid black line. Each diffusion coefficient we consider is shown as a dashed red line (each with a unique shade of red). We chose these coefficients to test a range of cases from the strongest reacceleration case to the more physically realistic regime with larger diffusion coefficients, where

reacceleration occurs at a slower rate. Finally, we also plot a dotted black line at the turbulent phase velocity in our simulations. We include Figure 4.1 to make the criteria necessary for turbulent reacceleration clearer. In some of our simulations, we include streaming transport and spatially variable cosmic-ray transport — both of these processes will lead to a deviation from Equation 4.6, in agreement with previous works (BO23; H24).

Figure 4.1 highlights a key aspect of turbulent acceleration: it primarily happens on large scales $> \text{kpc}$. This requirement is set by the cosmic-ray diffusion coefficient being a relatively large quantity. For example, observations of the local interstellar medium suggest a value of $10^{28} \text{ cm}^2 \text{ s}^{-1}$ (Evoli et al., 2019). With a reasonable ISM turbulent phase velocity of $\sim 10 \text{ km s}^{-1}$, the critical length scale will be $> 10 \text{ kpc}$, significantly larger than typical ISM scales. Instead, a much lower diffusion coefficient is necessary for turbulent acceleration to operate efficiently (H24).

Only on large scales ($> \text{kpc}$) and in gas with a larger v_{ph} could turbulent acceleration operate efficiently. In fact, we can estimate the growth rate of cosmic-ray energy as a result of this cosmic-ray acceleration via turbulence. Following the work of BO22; BO23, the growth time for anisotropic, magnetic-field-aligned transport is

$$\frac{E_c}{\dot{E}_c} \sim \frac{p^2}{D_{pp}} \equiv t_{\text{grow}} \sim \frac{9}{2} \frac{\kappa_{\parallel}}{v^2} \left(\frac{vL_0}{\kappa_{\parallel}} \right)^{1/2}. \quad (4.7)$$

We use the notation from BO22: p is cosmic-ray momentum, D_{pp} is the momentum diffusion coefficient, v is the root-mean-square velocity, κ_{\parallel} is the spatial diffusion coefficient parallel to the mean magnetic field, and L_0 is the outer scale of the turbulent cascade. BO22 derive this equation by augmenting the usual acceleration time for

a large diffusion coefficient (P88) with the number of eddies a cosmic ray interacts with, causing the growth time to decrease $t_{\text{grow}} \rightarrow t_{\text{grow}} N^{-1}$. They estimate this number is $N_{\text{eddy}} \sim \sqrt{\kappa_{\parallel}/(vL_0)}$, which is then included in the acceleration timescale (Equation 4.7).

We will show turbulent acceleration occurs in the hot diffuse phase of plasma in our simulations, in agreement with Equation 4.7. We also identify a separate effect due to the transition between the hot and cold phases, which has a faster growth rate. Just as in a single phase of gas, the cosmic rays can sap energy from a critical length scale, where they diffuse from an eddy before it breaks up into smaller scales; they can also sap energy from thermally unstable compressions. Converging flows due to the condensation mode lead to the growth of cosmic-ray energy. As long as the cosmic rays diffuse before a heating source causes the condensed cloud to expand, then they sap energy from the condensation and cooling of gas. This result is similar to what was found in the ISM case (H24).

It is also important to note that the turbulent reacceleration by long-wavelength turbulence we study here differs from the second-order Fermi acceleration implemented in other studies (e.g., (Böss et al., 2024)). We do not include any second-order Fermi term deriving from the individual scatterings of cosmic rays off magnetic fluctuations (for the derivation, see (Zweibel, 2017)). All reacceleration here comes via the interaction of the cosmic-ray energy density E_c and the bulk velocity of the thermal plasma \mathbf{u} . The primary terms which handle the transfer of energy between turbulent energy and cosmic-ray energy are $-\gamma_c E_c \nabla \cdot \mathbf{u}$ in the cosmic-ray energy equation and ∇P_c in the momentum equation. However, the diffusion term in the

cosmic-ray energy equation allows the cosmic rays to dissipate energy on small scales, making the transfer process irreversible at scales where diffusion is the dominant cosmic-ray transport mechanism (over advection).

Therefore, both turbulent reacceleration by long wavelength turbulence and the acceleration through condensing thermal instability we find in our simulations operate through the same term. In the case of condensation, instead of sapping energy from the turbulent cascade, the cosmic rays are taking kinetic energy from condensing gas, which would be lost via radiative cooling (H24).

4.3.2 Correlation

The correlation factor \mathcal{C} defined in Equation 4.4 has not yet been significantly examined for a turbulent medium. The only previous discussion is in Pinzke & Pfrommer (2010), where they account for the radial variation of the cosmic-ray pressure to gas pressure ratio $X_c = P_c/P_g$, which combined with a density profile could be translated into the correlation factor \mathcal{C} as a function of radius in a galaxy or galaxy cluster. However, previous works are specifically prescribing the dependence of $\langle E_c \rangle$ on $\langle n \rangle$, ignoring any local correlation of gas on scales unresolved by the observations.

In general, \mathcal{C} will be determined by the system's distribution in E_c - n phase space, illustrated in Figure 4.2. The γ -ray emissivity is proportional to the hadronic loss rate η_π in Equation 4.1. The full formula for η_γ is (Guo & Oh, 2008):

$$\eta_\gamma = \frac{\eta_\pi}{3} = 1.95 \times 10^{-16} \frac{\text{cm}^3}{\text{s}} \times E_c n = 3.12 \times 10^{-28} \frac{\text{erg}}{\text{cm}^3 \text{s}} \left(\frac{E_c}{\text{eV cm}^{-3}} \right) \left(\frac{n}{\text{cm}^{-3}} \right). \quad (4.8)$$

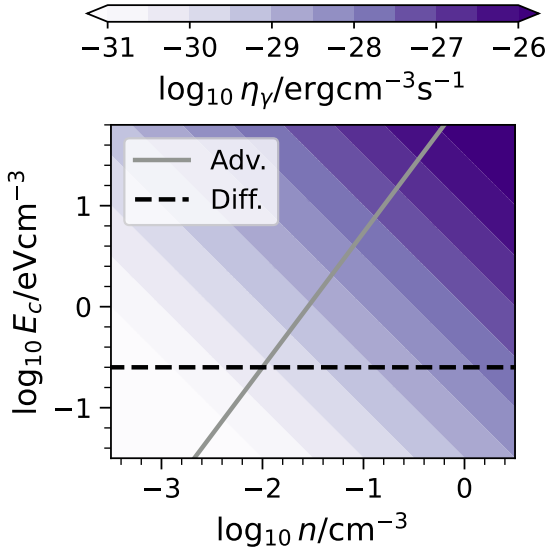


Figure 4.2: γ -ray emissivity as a function of cosmic-ray energy density E_c and plasma density n . The emissivity is proportional to the product of these quantities. The transport of cosmic rays determines the polytropic relationship between E_c and n . Advective transport, when cosmic rays are locked to the gas, leads to a $n^{4/3}$ dependence, illustrated with a solid gray line. More diffusive cosmic-ray transport leads to an n^0 dependence, which is shown as a dashed black line. Streaming along magnetic flux tubes would produce a $n^{4/3}$ dependence. As gas compresses and density increases, different cosmic-ray transport mechanisms produce significantly different γ -ray emissivity. Specifically, more advective (slower) transport leads to more γ -ray production.

Note our η_π is equivalent to Γ_h in Guo & Oh (2008). Following Guo & Oh (2008), we assume only one third of hadronic interactions produce neutral pions, which will then decay into γ -rays. So only the neutral pion channel contributes to the γ -ray emissivity η_γ . Additionally, Equation 4.8 can be used to calculate L_0 from Equation 4.2.

Depending on cosmic-ray transport, the distribution of the system will move and adjust in different ways in this space. Specifically, for slow transport, the cosmic rays are tied to the thermal gas and E_c follows an $n^{4/3}$ dependence. But for fast transport (diffusive or streaming), the cosmic rays are able to move through gas faster, leading to an n^0 dependence where the cosmic-ray pressure becomes uniform despite density variations. In reality, cosmic-ray transport will depend on the local plasma properties, magnetic field strength, etc. As a result, the distribution in the E_c - n phase space can be complex. In our simulation, we find correlation factors ~ 5 . This difference is significant because it means the values derived from observational

diagnostics can be consistently off from the actual average cosmic-ray energy density.

A final important facet of E_c - n phase space is how γ -ray luminosity heavily favors dense gas with high cosmic-ray energy density. These regions will produce significantly more γ -rays. So, to avoid any small regions of denser gas dominating the γ -ray luminosity, the cosmic rays must decorrelate from the dense gas. We test the importance of this decorrelation by probing different rates of cosmic-ray transport. We will increase the speed of cosmic-ray transport from the estimate of the Milky Way diffusion coefficient of $\kappa_{\parallel} \sim 10^{28} \text{ cm}^2 \text{ s}^{-1}$ in two ways. First, we will simply increase the diffusion coefficient. Second, we will examine streaming transport (the self-confinement model), when the Alfvén speed $B/\sqrt{4\pi\rho}$ dominates transport. Finally, in another simulation, we will increase Alfvén speed in cold gas to account for the effect of ion-neutral damping.

4.4 Simulation Details

We examine the evolution of γ -ray luminosity in a turbulent box of multiphase gas using Athena++ (Jiang & Oh, 2018; Stone et al., 2020). Athena++ solves the following equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \rho \mathbf{u} = 0 \quad (4.9)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot \left(\rho \mathbf{u} \mathbf{u} + \mathbb{1} \left(P_g + \frac{B^2}{2} \right) - \mathbf{B} \mathbf{B} \right) = \sigma_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{u} E_c) \quad (4.10)$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{u} \times \mathbf{B}) \quad (4.11)$$

$$\begin{aligned} \frac{\partial E_{\text{MHD}}}{\partial t} + \nabla \cdot \left(\mathbf{u} \left(E_{\text{MHD}} + P_g + \frac{B^2}{2} \right) - \mathbf{B} (\mathbf{B} \cdot \mathbf{u}) \right) \\ = (\mathbf{u} + \mathbf{v}_s) \cdot \sigma_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{u} E_c) - \rho \mathcal{L} \end{aligned} \quad (4.12)$$

$$\frac{\partial E_c}{\partial t} + \nabla \cdot \mathbf{F}_c = - (\mathbf{u} + \mathbf{v}_s) \cdot \sigma_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{u} E_c) - \frac{E_c}{\tau_h} \quad (4.13)$$

$$\frac{1}{V_m^2} \frac{\partial \mathbf{F}_c}{\partial t} + (\gamma_c - 1) \nabla E_c = - \sigma_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{u} E_c). \quad (4.14)$$

In the above, ρ is the plasma density, \mathbf{u} is the bulk velocity of the plasma, P_g is the gas pressure, \mathbf{B} is the magnetic field (with strength B), σ_c is the cosmic-ray transport matrix (see Equation 4.17 in Section 4.4.2), \mathbf{F}_c is the cosmic-ray energy flux density, $\gamma_c = 4/3$ is the adiabatic index of the cosmic-ray fluid, τ_h is the loss-rate of cosmic-ray energy density via hadronic interactions (see Equation 4.20 in Section 4.4.2), $E_{\text{MHD}} = \rho u^2/2 + B^2/2 + P_g/(\gamma_g - 1)$ is the total magnetohydrodynamic energy, $\rho \mathcal{L}$ is the heat-loss function (see Equation 4.21 in Section 4.4.3), \mathbf{v}_s is the cosmic-ray streaming velocity (see Equation 4.16 in Section 4.4.2), and V_m is the modified speed of light.

The computational methods we use within Athena++ are standard. However, we still specify the non-thermal physics related to cosmic rays and optically thin radiative cooling and heating for completeness.

4.4.1 Setup

Our simulation's initial setup is a homogeneous box with side lengths of 48 kpc and $N = 128$ cells on each side, filled with diffuse ionized plasma, a uniform magnetic

Simulation	$\kappa_{\parallel}/\text{cm}^2\text{s}^{-1}$	Streaming?	IND?
1e29	$1 \cdot 10^{29}$	No	No
3e29	$3 \cdot 10^{29}$	No	No
1e30	$1 \cdot 10^{30}$	No	No
3e30	$3 \cdot 10^{30}$	No	No
Str	$3 \cdot 10^{28}$	Yes	No
StrIND	$3 \cdot 10^{28}$	Yes	Yes

Table 4.1: Simulation Parameters. IND is short for ion-neutral damping. We include this effect through a parameterization of the ionization fraction f_i in Equation 4.19.

field, and a background of cosmic rays. We stir this box, using an Ornstein-Uhlenbeck process, and force the velocity perturbations to all together have a combined kinetic energy injection rate of $\dot{E}_{\text{inj}} = 1.5 \times 10^{41} \text{ erg s}^{-1}$. This low level of energy injection is enough to drive the sub-sonic turbulence in the diffuse plasma. We attempt to get Mach numbers $\mathcal{M} < 1$ because that matches observations of the ICM in the Perseus cluster (Hitomi Collaboration et al., 2018).

The initial parameters for each simulation are $T_0 = 10^7 \text{ K}$, $\rho_0 = 10^{-2} m_H \text{ cm}^{-3}$, $\beta_0 = 8\pi P_g/B^2 = 50$, and $X_{\text{cr}} = P_c/P_g = 0.01$. These correspond to a magnetic field strength of $B = 2.6 \mu\text{G}$ and a cosmic-ray pressure of $P_{\text{cr}} = 8.6 \times 10^{-2} \text{ eV cm}^{-3}$. In Table 4.1, we show the cosmic-ray transport parameters for each simulation.

These conditions differ from commonly used values for the ICM, because we are aiming to study cool-core cluster conditions, like the multiphase turbulence in the Perseus cluster. So, we assume a higher density ($10^{-2} m_H \text{ cm}^{-3}$, not $10^{-4} m_H \text{ cm}^{-3}$) and a lower temperature (10^7 K , not 10^8 K) than commonly used values for the general ICM. Similarly, the magnetic field strength and cosmic-ray pressure are larger because we fix the β and X_{cr} values at reasonable values for the Perseus cluster. The direct application of the parameters we use would appear in the inside of cool-

core clusters, where temperatures can be 1 keV and densities are higher $\gtrsim 10^{-2} \text{ cm}^{-3}$ (Fabian et al., 2006; Hitomi Collaboration et al., 2016, 2018).

Also, given we now have a set box size and initial conditions, we can define L_0 from Equation 4.2 by multiplying Equation 4.8 by the simulation volume:

$$L_0 = 2.63 \times 10^{39} \frac{\text{erg}}{\text{s}} X_{\text{cr}} \left(\frac{n}{10^{-2} \text{ cm}^{-3}} \right)^2 \left(\frac{T}{10^7 \text{ K}} \right). \quad (4.15)$$

The luminosity in Equation 4.15 corresponds to a flux of $\sim 1 \text{ GeV}$ photons of $\sim 1.7 \times 10^{-12} \text{ cm}^{-2} \text{ s}^{-1}$ assuming the simulation is 100 Mpc away from the detector. That estimate is more useful for comparisons with galaxy cluster observations (Ackermann et al., 2014), even though our simulation size is significantly smaller than a cluster. In fact, for studying an entire cluster, we would need larger scale simulations with a large scale density profile ranging from our initial density to lower densities in the outskirts of the cluster 10^{-4} cm^{-3} . Others have studied cosmic-ray feedback on large scales with realistic environments and structures, but we simplify the setup in this work to study how cosmic-ray transport affects the γ -ray luminosity of an arbitrary volume of gas in the regions (like a cluster core) where more γ -rays would be found.

4.4.2 Cosmic-Ray Physics

We use the two-moment method implemented in Jiang & Oh (2018). This implementation allows us to include the impacts of multiple cosmic-ray transport processes.

First, it includes the advection of cosmic rays with the background plasma, at the bulk flow velocity \mathbf{u} in Equations 4.9-4.14. Advection derives from the fact that

cosmic rays are scattering off magnetic field fluctuations. As the cosmic rays scatter in pitch angle, the cosmic-ray distribution function evolves in the frame co-moving with the waves, which includes any bulk flow speed (Zweibel, 2017).

Second, the wave frame also includes the streaming speed \mathbf{v}_s . The exact value of this speed depends on the source of the magnetic fluctuations which scatter the cosmic rays (Zweibel, 2017). In the self-confinement model, the fluctuations are Alfvén waves which are generated by the cosmic rays themselves as they stream down their pressure gradient, in which case $\mathbf{v}_s = v_A$ and is directed down ∇P_c . In the extrinsic turbulence model, the fluctuations are generated by a turbulent cascade which is generally taken to be balanced; i.e., the waves propagate up and down ∇P_c with equal intensity, and $\mathbf{v}_s \equiv 0$. More general constructions allow for a continuous range between the self-confinement and extrinsic turbulence models, but in this setup, we switch between the two models discontinuously. In Table 4.1, the simulations with “no” streaming have $\mathbf{v}_s = 0$, which corresponds to the extrinsic turbulence model. In simulations with streaming (“Yes” in the third column of Table 4.1), we follow the self-confinement model and include the streaming speed as

$$\mathbf{v}_s = -\frac{\mathbf{B}}{\sqrt{4\pi\rho_i}} \frac{\mathbf{B} \cdot \nabla P_c}{|\mathbf{B} \cdot \nabla P_c|}, \quad (4.16)$$

where ρ_i is the ion density. Except in the simulation with ion-neutral damping (**StrIND**), we assume the plasma is fully ionized regardless of temperature ($\rho_i = \rho$).

Finally, we also include spatial diffusion, which derives from the pitch angle scattering of cosmic rays. This diffusion appears regardless of the source of the scattering magnetic field fluctuations (Zweibel, 2017). Assuming the magnetic field is resolved,

the transport will be anisotropic, with cosmic rays diffusing faster in the direction of the magnetic field than in the directions perpendicular to it. To include this effect, we adjust the combined transport matrix σ_c in Equations 4.9-4.14 to be

$$\sigma_c = \frac{1}{\kappa_{\perp}} \mathbb{1} + \left(\frac{1}{\kappa_{\parallel} - \frac{4E_c \mathbf{v}_s \cdot \hat{\mathbf{b}}}{\hat{\mathbf{b}} \cdot \nabla E_c}} - \frac{1}{\kappa_{\perp}} \right) \hat{\mathbf{b}} \hat{\mathbf{b}} \quad (4.17)$$

where $\hat{\mathbf{b}} = \mathbf{B}/|\mathbf{B}|$ is the direction of the magnetic field, κ_{\perp} is the cosmic-ray diffusion coefficient perpendicular to the magnetic field, and κ_{\parallel} is the cosmic-ray diffusion coefficient parallel to the magnetic field.

An approximate way of including the effect of ion-neutral damping in the self-confinement model (IND in Table 4.1) is to adjust the streaming velocity to account for the appearance of neutrals in thermal gas with a temperature below a certain value. Similar approaches have been used previously (Ruszkowski et al., 2017). Only in the **StrIND** simulation, we model the ion density as

$$\rho_i = \rho f_i(T) \quad (4.18)$$

where ρ is the local gas density and $f_i(T)$ is the ion fraction at a given temperature. We model the ion fraction with a simple switching function $S(x)$:

$$f_i(T) = 1 + (f_{\min} - 1) S \left(\frac{T - T_i}{\Delta T} \right). \quad (4.19)$$

We choose $S(x) = 0.5 [1 + \tanh(x)]$, $T_i = 10^5$ K, and $\Delta T = 10^4$ K. This formulation is similar to the setup others have used to study the effect of changes in ionization

(Bustard & Zweibel, 2021).

For a fully realistic implementation, the temperature of the switch from fully ionized to partially ionized should be lowered to $T_i \sim 10^4$ K (Sutherland & Dopita, 1993; Bustard & Zweibel, 2021). We instead chose the parameters for the switching function based on our cooling function. Our cooling curve peaks at 10^5 K (see Section 4.4.3), and that temperature sets the dividing line between hot and cold plasma in our simulation. Choosing the switching temperature to match that phase transition will more realistically model the transition between a hot ionized phase and a colder neutral phase.

We also include the losses due to hadronic interactions, with a simple energy loss time τ_h (see Equation 4.13). Using Equation 12 from Guo & Oh (2008) for the energy lost to hadronic interactions per unit volume Γ_h , the loss time is:

$$\frac{1}{\tau_h} = \frac{\Gamma_h}{E_c} = 5.86 \times 10^{-16} \text{ s}^{-1} \left(\frac{n}{\text{cm}^{-3}} \right) = 0.185 \text{ Gyr}^{-1} \left(\frac{n}{10^{-2} \text{ cm}^{-3}} \right). \quad (4.20)$$

The rate $5.86 \times 10^{-16} \text{ s}^{-1}$ comes directly from Guo & Oh (2008). Since the total loss rate is proportional to nE_c , the loss time is inversely proportional to just the density n . It is also important to note that this rate leads to a loss of energy because of the sign of E_c/τ_π in Equation 4.13.

4.4.3 Heating & Radiative Cooling

In this section, we detail the heating and cooling in the heat loss function $\rho\mathcal{L}$ (Field, 1965) in Equation 4.12. We use the common formulation, which assumes a constant

per-particle heating rate Γ and a temperature-dependent volumetric cooling function:

$$\rho\mathcal{L} = n^2\Lambda(T) - n\Gamma. \quad (4.21)$$

We calculate the heating and cooling terms separately in the simulations, using an operator splitting approach.

It is possible to connect the evaluation of each term into a single calculation when Γ is a constant or a function of temperature (Tan & Fielding, 2024). However, it is not so simple when attempting to keep the entire simulation in constant energy balance. Enforcing conservation of energy when including the effects of optically thin radiative cooling requires adjusting Γ to counter the average cooling rate throughout the simulation. After evaluating the loss of energy due to the $n^2\Lambda(T)$ term in Equation 4.21 everywhere, we sum the total energy loss over the entire simulation. Then we subtract off the turbulent energy injection rate. The remaining energy is added back through the Γ in Equation 4.21, a constant rate for every cell in the simulation, conserving total thermal energy. If the turbulent energy injection is larger than the cooling, then we do not add any heating. This setting causes net energy gain, but for a short time at the beginning of the simulation, before some compressions form with larger cooling rates. Once large enough temperature variations form as the result of turbulent driving, the regions with more cooling than the average will cool, and regions with less cooling than average will become hotter. We also decouple this heating and cooling accounting from the cosmic-ray energy. While it is reasonable to assume the energy from thermal radiative cooling could be reabsorbed within the $(48 \text{ kpc})^3$ volume, the hadronic losses by cosmic rays (Equation 4.20) would not be

returned to the background thermal gas. All our simulations have total energy loss because of this additional energy outlet.

For the calculation of cooling, we implement the exact integration method from Townsend (2009). This process is straightforward for power-law cooling curves, so we use a common parameterization of the collisional-ionization equilibrium (CIE) curve from Rosen & Bregman (1995). Specifically, a piecewise power law given by:

$$\Lambda(T) = \begin{cases} 0 & T \leq 300 \\ 2.2380 \cdot 10^{-32} T^2 & 300 \leq T < 2000 \\ 1.0012 \cdot 10^{-30} T^{1.5} & 2000 \leq T < 8000 \\ 4.624 \cdot 10^{-36} T^{2.867} & 8000 \leq T < 10^5 \\ 1.78 \cdot 10^{-18} T^{-0.65} & 10^5 \leq T < 4 \cdot 10^7 \\ 3.2217 \cdot 10^{-27} T^{0.5} & 4 \cdot 10^7 \leq T \end{cases} \quad (4.22)$$

where T is in units of K and Λ is in units of $\text{erg cm}^3 \text{s}^{-1}$. We do adjust the function in Equation 4.22, by moving the floor temperature up to $T = 2000$ K instead of 300 K.

With the heating set to balance the radiative cooling of the thermal, there is only one true loss process in our simulations once turbulence is saturated: the hadronic losses of cosmic-ray energy. Since some simulations end up with significantly larger cosmic-ray energy densities on average, they end up with a larger total energy lost to hadronic interactions. This loss causes the total energy in the simulation to decrease over time, but we stop the simulations before lost energy becomes greater than 10% of the initial total energy.

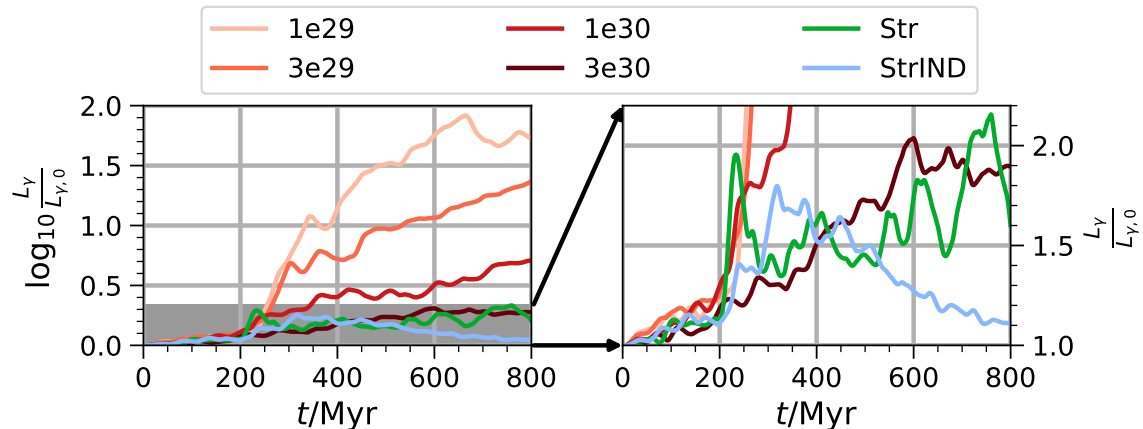


Figure 4.3: Evolution of γ -ray luminosity over time for each simulation. Left side shows the evolution of the logarithm of L_γ relative to the initial state. Right side shows a zoomed-in, linear scale version (the gray region in the left figure) to highlight the differences between the various fast transport cases. While slow transport (lighter red lines) leads to a nearly 100 times larger L_γ (left plot), there are order unity discrepancies between the different fast transport results (right plot).

4.5 Results

Our primary result is illustrated in Figure 4.3. Each simulation’s volume integrated γ -ray luminosity is shown with respect to the value at $t = 0$, which is $L_0 = 2.63 \times 10^{39} \text{erg s}^{-1}$ (Equation 4.15) for all simulations. The left-side plot in Figure 4.3 uses a log-scaled y-axis, whereas the right-side plot zooms in to a linear scale y-axis. The right-side plot clarifies the variation in γ -ray luminosity across the three fast transport cases (simulations 3e30, Str, and StrIND in Table 4.1). The lines with different hues of red are the simulations with only diffusive transport, whereas the green and blue lines correspond to the two simulations with streaming transport and subdominant diffusion.

Figure 4.3 shows that a simple change in cosmic-ray transport leads to variations in the total γ -ray luminosity of 100 times the initial γ -ray luminosity. Fast transport




Figure 4.4: Line-of-sight averaged gas density and cosmic-ray energy density adjacent to integrated γ -ray luminosity for each simulation. As the diffusion coefficient increases, cosmic rays decouple from the dense gas, and the diffuse γ -ray luminosity decreases. Note that the cosmic-ray energy density contours are normalized by the total cosmic-ray energy in the simulation to account for the effects of reacceleration, which cause the total luminosity in the $1e29$ run to be 100 times larger than the fast transport runs (see Figure 4.3). For the streaming runs, the γ -ray luminosity is low, but the cosmic rays are even more decorrelated from dense gas when the effects of ion-neutral damping are included.

is necessary for our simulations to be in reasonable agreement with observational upper limits on diffuse emission. However, among the fast transport processes, there is variation. Therefore, our simulations additionally illustrate how the microphysics of cosmic-ray transport significantly affect observables.

Compared to our estimate of flux from an object 100 Mpc away from us ($1.7 \times 10^{-12} \text{ cm}^{-2} \text{ s}^{-1}$, see section 4.4.1), the two orders of magnitude increase in luminosity for the **1e29** simulation is alarming. Propagating through, this simulation (significantly smaller than a galaxy cluster) would have produced a detectable flux $10^{-10} \text{ ph cm}^{-2} \text{ s}^{-1}$ (see Section 4.4.1).

The only simulation with a long-term decrease in γ -ray luminosity is the **StrIND** simulation with increased streaming speed in denser gas. Since this simulation assumes the self-confinement model, cosmic rays are generating their own waves, which then dissipate into the background plasma. This process leads to a decrease in cosmic-ray energy, especially since the heating rate (proportional to the Alfvén speed) increases in dense gas. So, once the simulations form enough dense gas, the cosmic rays begin to lose a significant amount of their energy to the background plasma. But as we will see in Section 4.5.1, that extra heating is marginal from the perspective of the background plasma.

The only other consistent decrease occurs at the end of the **1e29** simulation. Later, when examining the energy evolution, it is apparent that cosmic-ray energy has grown to become a significant energy component by $t \gtrsim 650 \text{ Myr}$. This new partition of energy pushes the **1e29** simulation into a different regime, which limits the growth of cosmic-ray energy through turbulent reacceleration. Additionally, the

γ -ray luminosity increases to be much larger, causing the total energy lost to become significant. This regime pushes beyond what is realistic for cool clusters or the CGM, which is already illustrated in Figure 4.3 as the γ -ray luminosity becomes very large. Therefore, we limit our analysis to the time periods prior to this late regime $t \gtrsim 650$ Myr in the **1e29** simulation.

Figure 4.4 shows actual snapshots of the simulations at $t = 600$ Myr, illustrating how the cosmic-ray energy density is able to de-correlate from the dense gas. This process helps keep the γ -ray luminosity low in the fast transport simulations. The left plots for each simulation show the average gas density along the (not pictured) \hat{z} axis as a colormap, with contours overlaid highlighting where the average cosmic-ray energy density (again along the \hat{z} axis) is above or below the mean value for the entire simulation. The right plot of each simulation shows the γ -ray luminosity L_γ with respect to the luminosity of the homogeneous initial conditions.

While the contours of large cosmic-ray pressure track well with dense gas in the **1e29** simulation, as we move between simulations, we see the cosmic rays and dense gas become more and more decorrelated. The largest decorrelation appears in the **StrIND** and **3e30** simulations. The plots of L_γ also change significantly as a function of cosmic-ray transport.

Figure 4.4 also illustrates different morphologies for the different cosmic-ray transport cases. In the lower diffusion coefficient cases, cosmic rays can help support the condensing against collapse, but the larger effect is through their reacceleration. As we will show in Section 4.5.2, turbulent reacceleration of cosmic rays and cold cloud reacceleration of cosmic rays are more efficient in the lower diffusion coefficient cases.

Therefore, more energy ends up in cosmic rays, where it escapes our accounting for heating. Since the cosmic rays can take energy from the condensing gas, the condensing gas loses less energy to radiative cooling.

This decrease in radiative cooling also decreases the amount of heating we add back into the simulation, allowing the thermal gas to continue to cool even more. As a result, we end up with larger regions of cold gas in the lower diffusion coefficient simulations. Note that this numerical effect is robust, as the avenues for cosmic-ray energy loss (e.g., hadronic and coulomb losses) differ from radiative cooling and would not directly add back into the thermal energy.

In the rest of this section, we provide more detailed results. We check the thermodynamic state of each simulation to confirm that the only driver for the variation in γ -ray luminosity is cosmic-ray transport. Then, we examine the causes of the increase in γ -ray luminosity. We find the primary contributor is turbulent acceleration by condensing cold clouds. However, the correlation of cosmic rays and cold gas (defined in Equation 4.4) also has an impact. Finally, we illustrate how the pockets of cold gas, instead of the volume-filling diffuse gas, set the γ -ray luminosity in our simulation.

4.5.1 Different Gamma-ray Luminosity, Similar Thermodynamics

As illustrated in Figure 4.3, the simulations have significantly different γ -ray luminosities, by as much as an order of magnitude. However, they are similar in most other aspects. In Figure 4.5, we show the evolution of each energy component in the

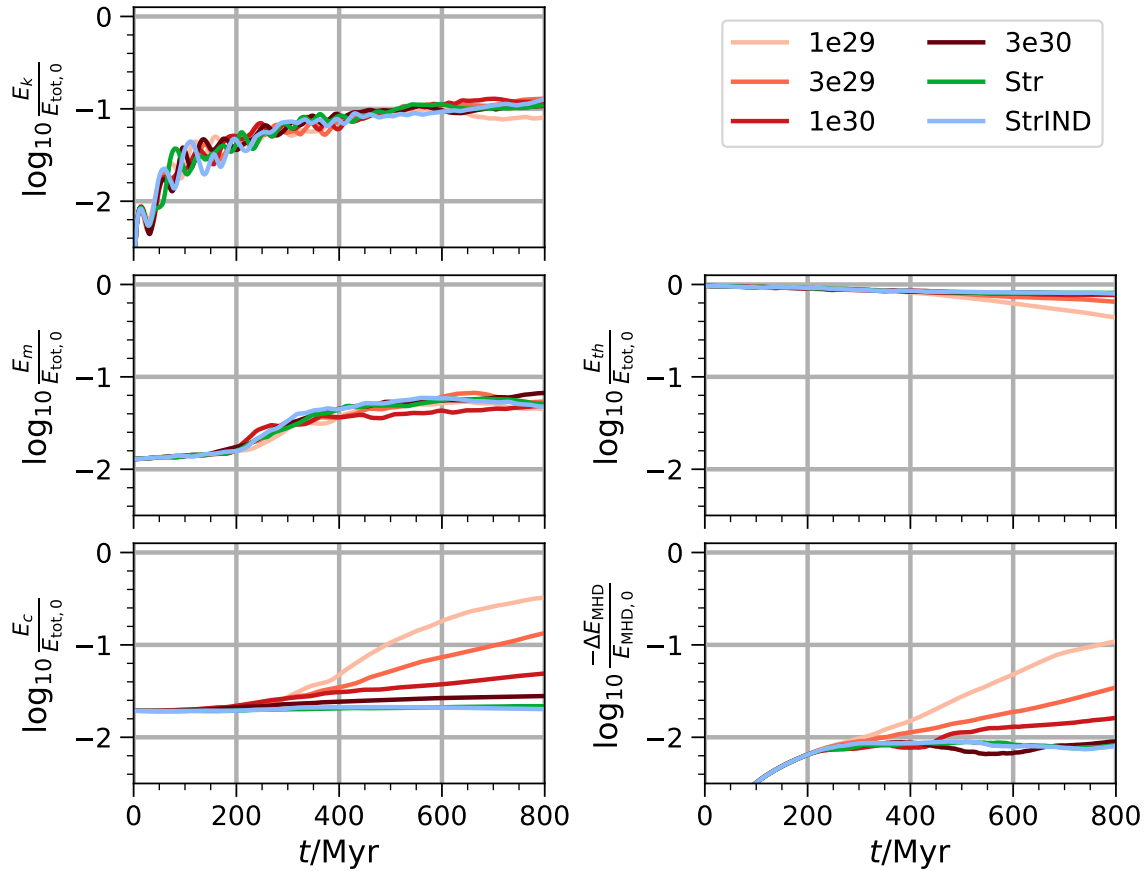


Figure 4.5: Evolution of total kinetic energy (upper left), total magnetic energy (middle left), total cosmic-ray energy (lower left), total thermal energy (middle right), and total energy change (lower right) for each simulation. The simulations have similar kinetic and magnetic energy evolution. The cosmic-ray energy evolution differs across the simulations because of the different cosmic-ray transport in each simulation. Similarly, the thermal energy barely changes, except for in the slow cosmic-ray transport cases. This change is clearer in the total energy loss plot (bottom right). The total energy decreases because of γ -ray emission, which increases proportionally with the cosmic-ray energy density. So, when the cosmic-ray energy density increases (bottom left plot), the total energy loss also increases.




Figure 4.6: Thermodynamic, kinematic, and magnetic properties of the simulations at $t = 600 \text{ Myr}$. Left column shows T - ρ phase space, middle column shows T - \mathcal{M} phase space, and the right column shows T - β phase space. In the top row and right of the third column, we show the marginal distributions of gas density ρ , Mach number \mathcal{M} , plasma beta β , and temperature T . Except for some small deviations in amplitude, the marginal distributions and overall phase space distributions are similar for all simulations.

simulations. Total kinetic energy and total magnetic energy are consistent across the simulations (left column, top and middle graphs). However, due to turbulent acceleration of cosmic rays, the total cosmic-ray energy grows in some simulations (left column, bottom plot). This energy gain comes at the expense of thermal energy (right column, middle plot).

Finally, in those simulations where cosmic-ray acceleration occurs, we have a larger change in the energy accounting (right column, bottom plot). Our heating only accounts for the cooling terms and turbulent driving, making sure that the net

energy lost via cooling is injected back into the simulation in either kinetic or thermal energy. However, we also included the loss of cosmic-ray energy due to hadronic interactions. This energy loss is proportional to the total cosmic-ray energy, so the hadronic losses increase with time in those simulations where cosmic-ray acceleration occurs. Since we do not add the hadronic losses back in as heating, the total energy of the simulations decreases (see Sections 4.4.2 and 4.4.3). The initial change in energy is the result of turbulence not being saturated (see the evolution of kinetic energy in the left column, top plot). During this time, the heating is weakened because the accounting assumes the turbulent cascade is in a steady state. This saturation time delay leads to a loss of energy of order 1%.

Next, we examine the thermodynamic and turbulent phase spaces for the simulations. In Figure 4.6 we show the simulations mostly overlap in T - ρ phase space (bottom row, left plot), T - \mathcal{M} (Mach number) phase space (bottom row, middle-left plot), and T - β phase space (bottom row, middle-right plot). We also show the marginal distributions of temperature T (right-most plot), density ρ (top row, left plot), Mach number \mathcal{M} (top row, middle plot), and beta β (top row, right plot).

In the 2D-phase spaces, each simulation is plotted with a black colormap (shown in the colorbar on the right of Figure 4.6). The histogram of each simulation has the same colormap and shading. The simulations almost completely overlap in each phase space. This similarity in thermodynamic, turbulent, and magnetic states is also apparent in the marginal distributions. The biggest differences are in the Mach numbers of the cold gas in the `1e29` simulation, and some additional heating in diffuse gas in the `StrIND` simulation. Similarly, the most common temperature values vary

slightly in the marginal temperature distributions. This extra cooling is present primarily because of the cosmic-ray acceleration and subsequent loss of cosmic-ray energy to hadronic interactions in the diffusion simulations (see total energy error plot in Figure 4.5).

Otherwise, the phase space structures in Figure 4.6 are reasonable, and consistent with our expectations and previous results (H24). Rapid cooling leads to a clump of cells between $T \sim 10^4$ K and $T \sim 10^5$ K. Those cells have a larger Mach number, since their sound speed is lower, but the turbulent velocities are similar to the surrounding medium. Those colder, denser cells also have a lower β , as the magnetic field becomes significantly compressed and strengthened. Additionally, the magnetic field in the dense clouds remains connected to the diffuse medium, so cosmic rays can move between the two phases along field lines as described in Section 4.4.2, suggesting numerical diffusion does not contribute significantly to the evolution of the magnetic field.

Overall, the conclusion is that adjusting the cosmic-ray transport will primarily impact the γ -ray luminosity of the diffuse plasma, not the thermodynamic properties. Any changes to the thermodynamic, turbulent, and magnetic properties of the plasma are minimal. As we discussed in Section 4.3.1, the primary way for cosmic rays to influence the thermodynamics and turbulence is via turbulent acceleration, but even this has a small impact because the total loss to hadronic interactions due to the increasing cosmic-ray energy density is small compared to the total plasma energy. Again, we are in the limit of $X_c \lesssim 1$, which is necessary for the diffuse γ -ray emission to agree with observed upper limits (Ackermann et al., 2014).




Figure 4.7: The evolution of L_γ is partly a function of the cosmic-ray acceleration timescale. Solid lines show the growth rate (logarithmic time derivative) of the total cosmic-ray energy for each simulation, with line colors matching the legend in Figure 4.3. In dashed lines, we also include the analytic theory expectation for the growth of cosmic-ray energy via turbulent reacceleration (BO22).

4.5.2 Turbulent Reacceleration, Cold Cloud

Reacceleration, or Correlation?

Now, we quantify the separate impacts of turbulent reacceleration and correlation on the γ -ray luminosity. We find that when turbulent reacceleration occurs in the slow diffusion simulations, it dominates the γ -ray luminosity. For fast transport, correlation is the primary factor influencing the γ -ray luminosity.

In Figure 4.7, we plot the inverse of the growth time of cosmic-ray energy during each simulation. We also plot the inverse growth time from analytic theory for the diffusive simulations (see Equation 4.7), following the work of BO22. There are significant differences in the shape of the simulated curves and the theoretical ones, which derive from the theory not accounting for a multiphase medium. However, the ordering of the curves (i.e., the dependence on κ_{\parallel}) at late times is consistent with analytic theory. At some points, our simulations even get faster acceleration rates than analytic theory predicts. For the streaming simulations, this turbulent acceleration is negligible, and the simulations spend more time with actual negative

growth rates (lines dive off the bottom of the logarithmic plot). For reference, we smoothed the growth rates from the simulations with a simple moving average (0th order Savitzky-Golay filter) with a window of 50 Myr.

Until now, we have been referring directly to Ptuskin’s turbulent reacceleration as the source of the increase in cosmic-ray energy. However, the deviation from analytic theory suggests there could be a different mechanism involved in reaccelerating cosmic rays. To discern the cause of this deviation from theory, we identify where reacceleration occurs in the simulation.

In Figure 4.8, we examine the spatial distribution of cosmic-ray energy growth time at $t = 300$ Myr for every simulation. The colormap shows where the cosmic-ray energy is either increasing (red, positive growth rate) or decreasing (blue, negative growth rate). The growth rate is defined as

$$\frac{1}{t_{\text{grow}}} = \frac{1}{E_c} \frac{\partial E_c}{\partial t} = -\frac{1}{E_c} (\nabla \cdot \mathbf{F}_c + (\mathbf{u} + \mathbf{v}_s) \cdot \sigma_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{u} E_c)). \quad (4.23)$$

Note for this calculation, we have used Equation 4.13 and removed the term related to hadronic losses. This definition of t_{grow} allows us to identify how energy is transferred between the cosmic rays and the gas. This calculation differs from that of Figure 4.7, which uses the total change in average cosmic-ray energy density, including the hadronic losses.

In Figure 4.8, we also plot contours at $T = 10^5$ K, highlighting any condensing clouds. Around these clouds, the growth time is mostly positive in all simulations, except **StrIND**. There are also plenty of red regions far away from the condensing clouds, as compressive waves wash over the diffuse gas. Those regions are where

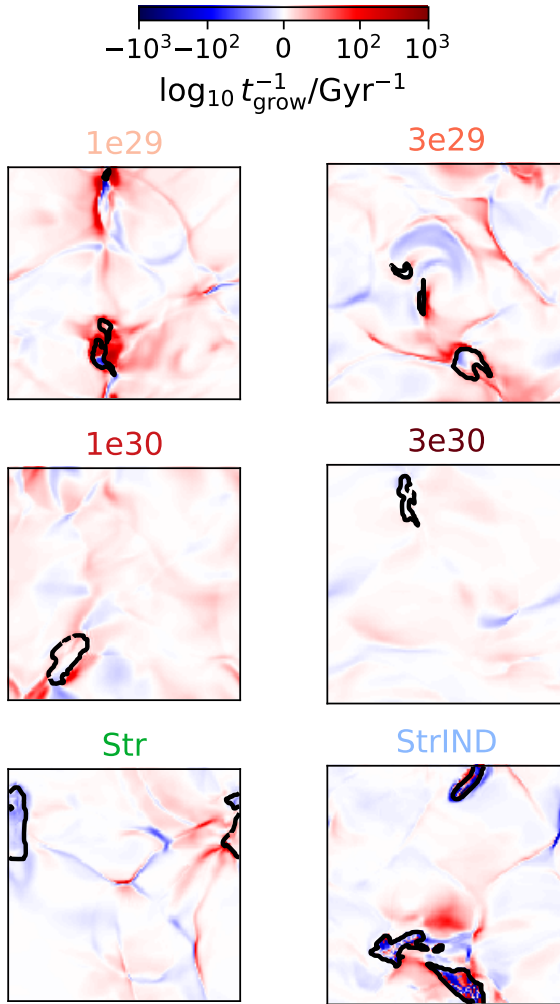


Figure 4.8: Cosmic-ray energy growth time (negative values are energy loss) in a slice of the simulations at $t = 300$ Myr. Black contours are where $T = 10^5$ K. The growth time calculation ignores hadronic losses, focusing on the interaction with turbulence and the cooling-induced collapse of clouds. The regions on and around the contours are predominantly red, meaning that cooling, condensing clouds preferably increase cosmic-ray energy density.

we expect turbulent acceleration to occur, with cosmic rays getting some net energy gain from large-scale turbulence when the bulk movement of the cosmic-ray energy density decouples from small-scale turbulent fluctuations. But the fastest growth rates in Figure 4.8 are generally near the edges of the condensing clouds. The cosmic rays initially get stuck in the clouds and experience compression because the cooling time is short compared to the diffusion timescale. This process reaccelerates them

just as a large-scale turbulent eddy might. Then, on the diffusion timescale, the cosmic rays can escape the cold cloud before it is destroyed by heating or fragments into smaller clouds.

However, Figure 4.8 is just one snapshot. In Figure 4.9, we show how all cells over the entire simulation are distributed in temperature and t_{grow} space. The vertical symmetry of every simulation is readily apparent. In general, in the turbulent box, cosmic-ray energy is increasing just as often as it is decreasing — the cosmic rays and thermal gas are, on average, coupled and evenly exchanging energy through adiabatic (reversible) expansion and compression.

There are small deviations which explain the cosmic-ray energy growth apparent in the bottom plot of the left column of Figure 4.5. For example, all the diffusion simulations show a slightly bright vertical streak near $T \sim 10^7$ K. Most of the simulation volume is at this high temperature, so most cells occupy this region, and most of them have minimal cosmic-ray energy growth or decay. The bright streak extending into the positive growth time regime illustrates the effect of turbulent reacceleration. This bright streak does not have a corresponding streak in the negative growth time regime, illustrating how cosmic rays are gaining energy on average.

The differences in the cold gas regime ($T \lesssim 10^5$ K) are harder to see because fewer cells end up in this temperature regime. One apparent difference is in the **StrIND** simulation (bottom row, right column of Figure 4.9). There are more cells in the negative growth rate regime for cold temperatures, showing that the increase in streaming velocity as a result of the change in ionization fraction leads to cosmic rays giving energy back to the thermal gas on average. The purple area extends




Figure 4.9: Average histograms of temperature and cosmic-ray energy-density-growth-time space. All cells over all time dumps (every 10 Myr) are included here. Overall, there is symmetry between cosmic-ray energy density growth and loss. Most cells are at a temperature near 10^7 K, with a smaller population of cells in colder gas at 3×10^4 K. But taking the ratio of the upper and lower sections of this plot highlights the differences and net long term effects (see Figure 4.10).




Figure 4.10: Ratio of the number of cells with cosmic-ray energy increase and the number of cells with cosmic-ray energy decrease, averaged over the simulation's runtime. Red regions show temperature regimes where cosmic rays are accelerated, and blue regions show where cosmic rays lose energy on average. The slight red region near $T \sim 10^7$ K is due to traditional turbulent reacceleration, but the dark red regions at 10^5 K is due to condensing clouds of gas. This condensing cloud acceleration only disappears when the cosmic rays decouple from the gas via ion-neutral damping (**StrIND**) at the same temperature as the peak of the cooling curve.

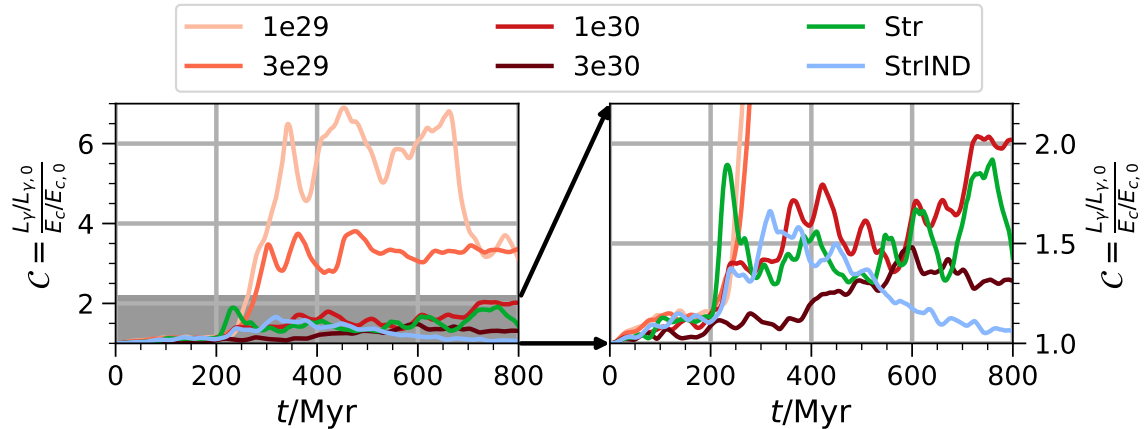


Figure 4.11: The left plot shows the evolution of the correlation coefficient C we defined in Equation 4.4 for each simulation. This coefficient divides out the effect of cosmic-ray acceleration, isolating the role of the correlation of cosmic-ray energy density and gas density in setting the γ -ray luminosity in the multiphase medium which forms after $t \sim 300$ Myr. The maximum effect the correlation has is a factor of ~ 6 , in the slowest diffusion coefficient simulation. However, this does not account for the nearly two orders of magnitude increase illustrated in Figure 4.3. Instead, turbulent reacceleration is the source of that large increase in γ -ray luminosity. For the fast transport when turbulent reacceleration is inefficient (right plot), the correlation becomes the dominant factor in setting the γ -ray luminosity (compare to the right panel of Figure 4.3).

further back into the $T \sim 10^5$ K regime. In the diffusion simulations, the opposite is true: the purple areas in the cold gas regime extend towards $T \sim 10^5$ K *only in the positive growth time regime*. Therefore, there is a reacceleration of cosmic rays occurring in this regime.

In Figure 4.10, we show the ratio of the positive and negative growth time regimes for each simulation to highlight the temperature regimes where there is a net exchange of energy between the thermal gas and cosmic rays. These plots highlight regimes where cosmic rays are preferentially gaining energy (red) or preferentially losing energy (blue). It is important to note that regions higher in this plot are gaining energy at a faster rate — the growth time gets as low as ~ 1 Myr in some cases.

All the simulations show some amount of turbulent reacceleration in the hot gas regime, with a red region around 10^7 K but extending across at least two orders of magnitude in growth time. In general, there are about 10 times as many cells with cosmic rays gaining energy in this temperature regime at any time dump. However, every simulation except **StrIND** shows a darker red region right at 10^5 K, illustrating that on average there are 100 times as many cells in this region gaining energy as opposed to losing it. In contrast, the **StrIND** simulation has a large blue regime where $T \lesssim 10^5$ K, because the streaming heating term is amplified in that temperature regime. Extra cosmic-ray energy is rapidly put back into heating the thermal gas through ion-neutral damping wherever there is a cosmic-ray pressure gradient.

The reacceleration at 10^5 K coincides with the peak of the cooling curve in our simulations. That temperature is the contour level used in Figure 4.8, so we know these cells are always around or near cooling clouds. However, given this is the peak of the cooling curve, we know gas at $T \sim 10^5$ K does not stay at that temperature. It is a transitory state as the gas is cooling. However, that state is when compression and condensation is strongest, leading to a large and negative $\nabla \cdot \mathbf{v}$ which reaccelerates cosmic rays. **Not only does turbulent acceleration increase the cosmic-ray energy of our simulations, but there is a similar process driven by the formation of cold clouds in a multiphase medium.**

Now that we understand where the increase in cosmic-ray energy comes from, we can focus on the correlation of cosmic rays and dense gas. To isolate the effect of dense gas-cosmic ray correlation, we divide out the evolution of the total cosmic-ray energy density from the γ -ray luminosity. Those curves are shown in Figure 4.11 in

the same style as Figure 4.3. This quantity is equivalent to the correlation \mathcal{C} defined in Equation 4.4. The correlation increases and peaks during the initial formation of cold gas and with the saturation of turbulence. However, the curves then level off at different values determined by the cosmic-ray transport rate. The lowest diffusion coefficient run **1e29** has a correlation of ~ 6 , whereas the fastest diffusion run **3e30** only reaches a maximum correlation of ~ 1.5 .

From Figure 4.11, it is clear that while correlation plays a role in setting the luminosity, the two order of magnitude increase in γ -ray luminosity illustrated in Figure 4.3 must primarily result from the growth of cosmic-ray energy due to reacceleration (see Figures 4.7 and 4.10). The correlation only accounts for a small amount of this increase.

The story is different when comparing the fast transport simulations (**3e30**, **Str**, and **StrIND**). On the right side of Figure 4.11, the lines of correlation \mathcal{C} are much closer to the amplitude in the change of the γ -ray luminosity. This result makes sense, since the total cosmic-ray energy barely changes for those simulations (see Figures 4.5 and 4.7). Therefore, while reacceleration determines the growth of γ -ray luminosity when cosmic rays are transported slowly (left side of Figure 4.3), it is the correlation between dense gas and cosmic rays that determines the γ -ray luminosity when cosmic rays are transported rapidly.

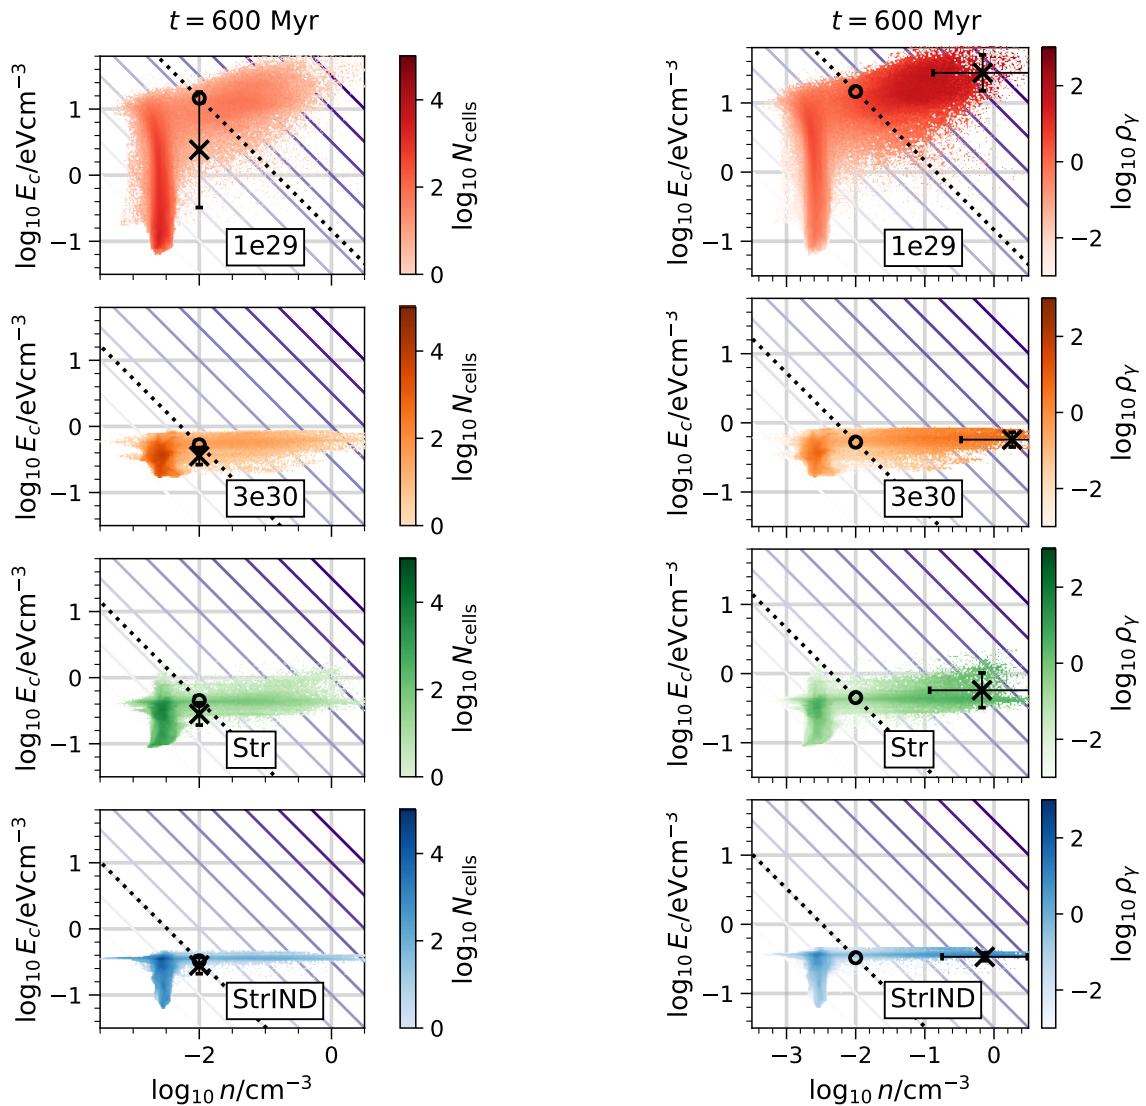


Figure 4.12: Cosmic-ray energy-gas density phase space distributions for the 1e29 (top row), 3e30 (second row), Str (third row), and StrIND (bottom row) simulations at $t = 600$ Myr. For convenience, we plot the γ -ray emissivity contours and polytropic relationships, shown in Figure 4.2. We plot a dotted purple line along the average emissivity (proportional to the γ -ray luminosity) of each simulation at $t = 600$ Myr, along with a black circle where assuming $\mathcal{C} = 1$ puts the average cosmic-ray energy density, given the average gas density of 10^{-2} cm^{-3} and average emissivity. The left column shows the raw distribution of computational cells, a volume-weighted histogram. The right column re-weights the distribution using the γ -ray emissivity. Because the denser gas emits more γ -rays, that part of the distribution appears darker. For each weighting (volume and γ -ray emissivity), we calculate the average of n and E_c , and mark this with an X. The error bars show the standard deviation of the average values.

4.5.3 Cosmic-Ray Energy Density in Dense Gas sets Gamma-Ray Luminosity

Now that we have an estimate for the effect of correlation on the γ -ray luminosity, we will take a closer look at the distribution of the simulations in E_c - n phase space. As discussed in Section 4.3.2, the distribution of gas in E_c - n phase space determines the γ -ray luminosity. Therefore, all the changes due to cosmic-ray transport we have examined so far in Section 4.5 will be apparent in E_c - n phase space.

In Figure 4.12, we illustrate how each simulation's volume is distributed in E_c - n phase space at $t = 600$ Myr. For simulations **1e29**, **3e30**, **Str**, and **StrIND**, we show how many cells are in each region of cosmic-ray energy density vs. gas density phase space in the left column of Figure 4.12. We also plot the contour lines of γ -ray emissivity and the advective vs. diffusive transport polytropic indices from Figure 4.2. The distributions change significantly across the different diffusion rates and transport mechanisms.

All the simulations are predominantly diffuse gas ($n \lesssim 10^{-2} \text{ cm}^{-3}$) and not dense gas, as we saw in Figure 4.6. However, in the case where reacceleration is important (the **1e29** simulation in the top plot), the cosmic-ray energy density is much higher in the dense gas. Similarly, the cosmic-ray energy density stays even lower in the dense gas of the **StrIND** simulation than in the **3e30** and **Str** simulations, as a result of the increase in the Alfvén speed due to the change in ionization.

Each plot has an X marking the mean values of cosmic-ray energy density and gas density. The error bar shows one standard deviation in $\langle E_c \rangle$, but we do not show the error bar in $\langle n \rangle$. Looking back at the marginal density distribution in Figure 4.6

illustrates the density is bimodal, so using a single standard deviation is inaccurate. Calculating it leads to an error bar stretching the entire span of density values, from 10^{-4} cm^{-3} to 1 cm^{-3} . However, since mass is conserved in the simulations, this average density of 10^{-2} stays constant during the entire simulation. Therefore, the final distribution of densities appears “unconstrained” because of the large range in density values, but the average value is extremely constrained.

Each plot also has a dotted purple line. This line traces along the average emissivity (proportional to the total luminosity) for each simulation at the given time. The separation of this line from the location of the average density and average cosmic-ray energy is captured by the correlation factor \mathcal{C} (see Figure 4.11).

We also plot a black circle (or O), marking where a γ -ray detection of our simulation, combined with the average density 10^{-2} cm^{-3} , would place the average cosmic-ray energy density. The disagreement of the X and O points in the left column is the correlation \mathcal{C} .

The left column of Figure 4.12 shows that a single measurement of γ -ray luminosity is not a reliable measure of the distribution of gas in E_c - n phase space. This luminosity will fall in between the dense gas and the diffuse gas, and attempting to calculate $\langle E_c \rangle$ from that luminosity and $\langle n \rangle$ without considering this distribution will cause a severe overestimate of the cosmic-ray energy density. The overestimate, the difference between the X and O in these plots, is equivalent to the correlation factor (Section 4.3.2).

Across all the plots in the left column of Figure 4.12, the true average (the X) also does not appear where there is significant data. Instead, the filling factor of

the dense gas causes the diffuse component to be at a lower density, such that when averaged with the higher density gas, the expected value lands in between the two phases. Very little gas exists in this state because it is thermally unstable.

In the right column of Figure 4.12, we show the same histograms, reweighted by the γ -ray emissivity of each cell. We then calculate the expectation value or mean value of the cosmic-ray energy density and gas density with these weights to find which part of the distribution contributes the most to the final γ -ray luminosity. Those points are marked with an X and have error bars corresponding to one standard deviation of each variable (cosmic-ray energy density and gas density). Across all transport mechanisms, the dense gas dominates the γ -ray luminosity, despite it not being volume filling (illustrated by there being lower values of N_{cells} in the dense gas regime in the histograms of the left column of Figure 4.12).

We again plot a dark purple, dotted line corresponding to the average emissivity. The weighted average of cosmic-ray energy density and gas density is much higher than the average emissivity. Effectively, the cosmic-ray energy density in this dense gas pulls the average luminosity upward.

We also plot the black circle marking a mock observational result. Across all transport mechanisms, the observed value of E_c is at a similar value to the emissivity weighted average cosmic-ray energy density — not the volume weighted average cosmic-ray energy density. In other words, the X marks in the right column of Figure 4.12 are always near the observed value of E_c , shown by the O marks. In the left-hand column, the volume weighted E_c marked with an X is always below the observed estimate. **Therefore, if a medium is multiphase and has some dense**




Figure 4.13: Following up on the calculation of the η_γ -weighted centroids in Figure 4.12, we plot the movement of the emissivity weighted centroids (X markers on the right side of Figure 4.12) over time for each simulation. The diffusion and advection polytropic laws intersect at the initial state of each simulation. Each curve follows a spectrum of slopes, with slower cosmic-ray transport leading to movement close to the advection polytropic law. Overall, transport dictates what the cosmic-ray energy density is in dense gas, and this significantly adjusts the total γ -ray output.

gas, the γ -ray luminosity probes the average cosmic-ray energy density in the dense phase — not the cosmic-ray energy density in the diffuse phase.

It is important to note that the γ -ray-emissivity weighted mean values of gas density and cosmic-ray energy density are different from the mean values used in calculating the γ -ray luminosity. Returning to Equation 4.1, we can see the weighting is purely based on volume. Therefore, taking mean values of gas density and cosmic-ray energy on the left side column of Figure 4.12 is still a reasonable way to estimate the γ -ray luminosity. But the right side of Figure 4.12 highlights how it is important to consider any (possibly unresolved) cold and dense gas when observing the γ -ray emission from diffuse gas. If the medium is multiphase with small volumes of dense gas, then those cold, dense clumps could dominate the γ -ray emission.

We also can examine how the γ -ray-emissivity weighted phase space centroids in the right column of Figure 4.12 evolve through time. Figure 4.13 shows how the weighted average density and cosmic-ray energy density evolve with time in E_c -

n phase space. Overall, the centroids follow the simple polytropic laws but have key deviations. Initially, they all follow the advection curve. Once diffusion takes over, the lines move to the right, with lower diffusion coefficients staying closer to the advection case. The polytropic law is also apparent in the simple streaming case without ion-neutral damping (simulation **Str**, green line in Figure 4.13), which follows a 2/3 polytropic index before cold gas forms. That index is expected for streaming at the Alfvén speed and results from the $\sqrt{\rho}$ in the Alfvén speed when examining a steady state system.

Note, we do not plot the standard deviations in Figure 4.13. The dense gas ends up with a variety of densities in the **3e30**, **Str**, and **StrIND** simulations, but a significantly lower emissivity weighted cosmic-energy density since the cosmic rays can easily escape the dense regions in these simulations.

4.6 Discussion

In this section, we discuss the implications of our results. In particular, we discuss how our examination of the correlation factor \mathcal{C} can be used in estimating the average cosmic-ray energy density in diffuse gas and how turbulent reacceleration could have a significant impact on the dynamics of turbulence in diffuse media like the CGM or ICM.

As mentioned in Section 4.2, observations of γ -rays from environments like the ICM and CGM provide one of the few probes of the cosmic-ray energy density E_c . While we have direct measurements of the cosmic rays ($\sim 1 - 10$ GeV) in

the solar neighborhood, we lack a direct probe in these diffuse, large-scale environments. Therefore, we must understand all the assumptions we make about the diffuse medium and how those assumptions introduce error into derived quantities like an average cosmic-ray energy density.

Previous works have focused on the full structure of galaxy clusters, allowing for a direct comparison with γ -ray observations (Beckmann et al., 2022). The end conclusion from most of these simulations is to continue decreasing the average cosmic-ray energy density, so that the estimated γ -ray emission is within a reasonable value. That is still necessary, but our work addresses another possible takeaway from observations: that fast cosmic-ray transport is necessary to keep γ -ray emission low. In particular, our simulations illustrate that the cosmic-ray energy density needs to decouple from turbulent fluctuations in dense gas regions in order to have a low γ -ray luminosity. These results primarily apply to observations of the cores of galaxy clusters, where the initial conditions of our simulations are found, but there are also other objects (outskirts of CGM, massive elliptical galaxies, etc.) where similar initial conditions and length scales could be found.

The simplest estimate of cosmic-ray energy density is the ratio of the observed γ -ray luminosity to the average gas density (ideally measured with another observable quantity besides γ -ray luminosity). That estimation method is shown in Equation 4.3. But as we have illustrated, the turbulent, multiphase nature of astrophysical plasmas limits the accuracy of this method. A good first correction is to consider the correlation of cosmic rays with dense gas, and include this as a factor \mathcal{C} . This correction can be estimated from an assumed distribution of gas in E_c - n phase space.

We have already shown that the value of \mathcal{C} can range from 6-1.5, depending on the cosmic-ray transport in the medium. For more complex simulations with more realistic cosmic-ray transport, the value could be different. Overall, the correlation factor allows for a conversion from the γ -ray luminosity observations, which probe the cosmic-ray energy density in the denser gas components of a multiphase plasma, to the true average value of cosmic-ray energy density.

We also illustrated how the γ -ray luminosity could be increased as a result of turbulent reacceleration (P88). This mechanism caused the largest change in γ -ray luminosity (see Figures 4.3 and 4.7), and was an even stronger effect than the correlation factor. Although the quantities are not completely independent. Turbulent reacceleration produced more complex and extended E_c - n phase space distributions (Figure 4.12), which increased the correlation factors.

There is a nuance to turbulent reacceleration which we have not discussed. The reacceleration only matters to γ -ray luminosity if the cosmic rays gaining energy exceed the energy threshold for hadronic interaction ($E_{\text{thr}} = 1.282 \text{ GeV}$). Additionally, if reacceleration pushes cosmic rays to higher energies such that the emission is outside observable bands of γ -ray emission, then the reacceleration will not impact the γ -ray luminosity. Overall, the evolution of the cosmic-ray distribution function under the influence of turbulent reacceleration would have an impact on the evolution of the γ -ray luminosity. However, we are unable to address that process, because our simulations do not evolve separate bands of the cosmic-ray distribution function. Instead, they evolve the total cosmic-ray energy density. Therefore, the details of the turbulent reacceleration we observe require further study. But fundamentally, our

results show that reacceleration exacerbates the γ -ray luminosity problem identified in other works. If cosmic-ray transport does not become fast in colder, denser gas, then the γ -ray luminosity becomes too large to agree with observations. Turbulent acceleration further increases the cosmic-ray energy density in dense gas, further increasing the γ -ray luminosity.

We expect these results regarding turbulent reacceleration to extend to more complex cosmic-ray transport, like two-zone transport models (H24). When diffusive transport increases in dense gas, turbulent reacceleration still occurs in the diffuse gas (H24). However, since the cosmic-ray energy is able to escape the dense gas, it should lead to lower total γ -ray luminosities, like those seen in the **StrIND** simulation. Except, instead of the energy going back into thermal gas, the energy will still be in cosmic rays, but they will emit fewer γ -rays because their residence time in the dense gas will be shorter in a two-zone model of diffusive transport. Future work could determine if a two-zone model could allow the increase in cosmic-ray energy we observe in the slower transport cases without a corresponding increase in γ -ray luminosity.

The final point of discussion regards how the cosmic-ray energy growth time in our simulations (Figure 4.7) slightly disagreed with the analytic estimates in P88; BO22; BO23. We found the cause of disagreement by looking at where cosmic-ray energy was increasing the most. We found that $T \sim 10^5$ K gas was the dominant contributor to cosmic-ray reacceleration in Figure 4.10. That temperature corresponds to cooling, condensing gas clouds (see Figure 4.8) as that temperature is where our simple cooling curve peaks (see Equation 4.22). However, if the cosmic rays effectively decouple

from the gas at the same temperature (due to ion-neutral damping **StrIND**), then the reacceleration via condensing cold clouds does not occur.

In reality, the cooling function and cosmic-ray transport in multiphase plasma are more complex than what we implement here. The reacceleration via condensing clouds could still occur, depending on how changes in ionization correspond to changes in the cooling rate. The peak in the cooling curve for solar metallicity gas is near $T \sim 10^5$ K, but ionization is more likely to shift at $T \sim 10^4$ K. Therefore, one could end up with models where cosmic rays are coupled to regions of condensing warm ionized medium (WIM) surrounding regions of warm neutral medium (WNM). The cosmic-ray reacceleration would occur in this shell of WIM, but the cosmic rays could decouple (due to ion-neutral damping) in the WNM and cold neutral medium (CNM) of dense clouds of gas. In the future, we will examine the collapse of an individual cloud from diffuse gas in ICM, CGM, and ISM conditions, including more realistic and complex physical descriptions of cooling and cosmic-ray transport. Without those detailed simulations, it is difficult to conclude that cosmic-ray reacceleration by condensing clouds actually occurs in any astrophysical medium. However, that mechanism does describe the reacceleration in our simulations, which deviated from the analytic estimates and simulations of single-phase turbulent media.

4.7 Conclusion

In this work, we presented an in-depth look at how cosmic-ray transport influences the γ -ray luminosity of diffuse gas. Overall, we have illustrated how the interpre-

tation of γ -ray luminosity observations of diffuse gas requires several assumptions about cosmic-ray transport, the interaction of cosmic rays with turbulence, and the correlation of cosmic rays with dense gas. Below, we provide a list of the individual key results from this paper. We also point to the figures that illustrate each conclusion.

- Cosmic-ray transport needs to be rapid in diffuse gas on large length scales, with either a large diffusion coefficient $\gtrsim 10^{30} \text{ cm}^2\text{s}^{-1}$, or via streaming at the Alfvén speed. Otherwise, the γ -ray luminosity will be extremely large. Specifically, fast transport reduces γ -ray luminosity by a factor of at least 100 compared to slower transport. (See Figure 4.3)
- The correlation of cosmic rays with dense gas causes a factor of 1.5-6 difference in estimates of the average cosmic-ray energy density. This correlation itself is a subdominant factor in determining the boost in γ -ray luminosity. Instead, it is the turbulent and cooling-induced reacceleration that drives the level of γ -ray production. (See Figures 4.4, 4.11, and 4.12)
- Turbulent reacceleration by long wavelengths (P88) in diffuse gas will increase the average cosmic-ray energy by factors ranging from 2 to 20, thereby driving the γ -ray luminosity up. This effect shows it is important to check the efficiency of turbulent reacceleration (See Figure 4.7).
- We show that cosmic-ray energization by condensing, cooling clouds could be more efficient than single-phase turbulent reacceleration suggested by P88. Additionally, it operates at length scales associated with the cooling processes,

instead of only being determined by cosmic-ray transport. The cosmic-ray transport can stifle this cold cloud reacceleration if ion-neutral damping transfers cosmic-ray energy to the thermal plasma at the same temperature as the phase transition. (See Figures 4.8,4.9, and 4.10)

- At the initial cosmic-ray energy densities studied in this paper, the gas dynamics (temperatures and velocities) and energy balance are not greatly affected by the cosmic rays (Figures 4.5 and reffig:thermodynamics). Except in one case: when the hadronic losses experienced by cosmic rays reach the same order of magnitude as the thermal energy at late times in the `1e29` simulation.
- Even when cold, denser gas is not volume filling, it is the dominant factor in setting the γ -ray luminosity. This conclusion results directly from the dependence of γ -ray luminosity on the local gas density and cosmic-ray energy density. However, it implies that estimates of average cosmic-ray energy density from observations of γ -ray luminosity are estimates of the cosmic-ray energy density in any dense gas clumps in the medium, as opposed to the cosmic-ray energy density in the diffuse component. (See Figures 4.12 and 4.13)

Acknowledgements

The authors would like to thank Karol Fulat, Ka Wai Ho, Mohan Richter-Addo, Aaron Tran, Bindesh Tripathi, and Ka Ho Yuen for their insight and helpful conversations which improved this work. We also appreciate the thoughtful referee whose

advice greatly improved this manuscript, particularly by identifying areas where we could improve the clarity of our results and conclusions.

We appreciate the computational time on NCSU Delta provided by NSF ACCESS allocation PHY240310.

RH and EZ greatly appreciate funding from NASA FINESST grant 80NSSC22-K1749, NSF grant AST-2007323, and NSF grant PHY 2409224 which supported this work. MR acknowledges support from the National Science Foundation Collaborative Research Grant NSF AST-2009227. This work was performed in part at the Aspen Center for Physics (ACP) during the “*Cosmic Ray Feedback in Galaxies and Galaxy Clusters*” summer program. The ACP is supported by the National Science Foundation grant PHY-2210452, and by grants from the Simons Foundation (1161654, Troyer) and Alfred P. Sloan Foundation (G-2024-22395). This research was supported in part by grant NSF PHY-2309135 to the Kavli Institute for Theoretical Physics (KITP).

Software: Athena++ (Stone et al., 2020; Jiang & Oh, 2018), Matplotlib (Hunter, 2007), NumPy (Van Der Walt et al., 2011; Harris et al., 2020), AstroPy (Astropy Collaboration et al., 2013, 2018)

5 SINGLE COSMIC-RAY INJECTIONS

They [must] develop a sense for energy relationships. They [must] observe that energy soaks up the patterns of things and builds with those patterns.

— *Children of Dune* (1976) by Frank Herbert

This chapter is a modified copy of the paper Habegger et al. (2023), titled *The Impact of Cosmic Ray Injection on Magnetic Flux Tubes in a Galactic Disk*, published in *The Astrophysical Journal* in 2023. No text has been changed; the figures and equations have just been reformatted to accommodate the submission of this thesis document.

I deeply appreciate the contributions of my co-authors Dr. Ellen Zweibel and Sherry Wong, to this work.

5.1 Abstract

In a seminal paper, Parker (1966) showed the vertical stratification of the interstellar medium (ISM) is unstable if magnetic fields and cosmic rays provide too large a fraction of pressure support. Cosmic ray acceleration is linked to star formation, so Parker’s Instability and its nonlinear outcomes are a type of star formation feedback. Numerical simulations have shown the instability can significantly restructure the ISM, thinning the thermal gas layer and thickening the magnetic field and cosmic ray layer. However, the timescale on which this occurs is rather long (~ 0.4 Gyr). Furthermore, the conditions for instability depend on the model adopted for cosmic

ray transport. In this work, we connect the instability and feedback problems by examining the effect of a single, spatially and temporally localized cosmic ray injection on the ISM over ($\sim 1 \text{ kpc}^3$) scales. We perform cosmic ray magnetohydrodynamic simulations using the `Athena++` code, varying the background properties, dominant cosmic ray transport mechanism, and injection characteristics between our simulation runs. We find robust effects of buoyancy for all transport models, with disruption of the ISM on timescales as short as 100 Myr when the background equilibrium is dominated by cosmic ray pressure.

5.2 Motivation

In 1966, Eugene Parker showed the interstellar medium's (ISM's) stratification was unstable when non-thermal pressure support by magnetic field and cosmic rays was introduced (Parker, 1966). The free energy source for the instability is the gravitational potential energy of gas held above its natural scale height. Parker argued that undular perturbations of the field lines, which on average are tangent to the Galactic plane, would allow gas to slide into the magnetic troughs, while the crests were free to buoyantly rise. The end result would be the production of cosmic ray inflated magnetic lobes rising perpendicular to the galactic disk and pockets of dense gas in the Galactic plane.

In order for the instability to occur, the equation of state (EoS) of the medium must be soft enough to allow compression into pockets. Parker assumed the cosmic ray pressure would remain constant along field lines and that the gas would follow an

isothermal EoS. Under these conditions, *any* pressure support from magnetic fields and/or cosmic rays renders the system unstable.

The Milky Way's thermal energy density, magnetic energy density, and cosmic ray energy density are all on the order of $\sim 10^{-12}$ erg cm⁻³; e.g. (Ferrière, 2001; Draine, 2011). In general, the far infrared - radio correlation (Bell, 2003) suggests that star formation, cosmic ray acceleration, and magnetic field generation are linked. This correlation means the Parker instability connects star formation to the disruption of the ISM and the production of outflows. So, the Parker instability is a form of star formation feedback and could play a key role in galactic evolution. It could mediate the formation of dense clouds, but also heat and disrupt them (Bustard & Zweibel, 2021). It could launch mass from the disk as a wind or fountain (Everett et al., 2008; Uhlig et al., 2012; Hanasz et al., 2013; Girichidis et al., 2016; Ruszkowski et al., 2017; Wiener et al., 2017b; Farber et al., 2018; Zhang, 2018; Hopkins et al., 2018; Chan et al., 2019; Hopkins et al., 2021). Perhaps less directly related to feedback, it could contribute to cosmic ray escape from the disk, and could influence the galactic dynamo by growing the component of the field perpendicular to the disk (Hanasz & Lesch, 2000; Hanasz et al., 2004, 2009).

While the non-thermal pressure in cosmic rays and the aforementioned buoyancy mechanism provide a way for cosmic rays to affect galactic evolution and the structure of the ISM, there is evidence against the Parker instability being a dominant effect in galaxies. The original Parker instability model has been revised to include more physical processes and more realistic setups (Mouschovias, 1974; Zweibel & Kulsrud, 1975; Asseo et al., 1978; Giz & Shu, 1993; Heintz & Zweibel, 2018), and analyzed in

the nonlinear regime with numerical simulations (Heintz et al., 2020; Tharakkal et al., 2023). Overall, these results lead to more stringent conditions for instability and show that the instability requires a longer time to vertically restructure the ISM; typically over 400 Myr (Heintz et al., 2020; Tharakkal et al., 2023). This long timescale means the effects of the instability could be overshadowed by other processes with shorter timescales.

One key difference between the different analyses lies in the treatment of cosmic ray transport and the choice of cosmic ray EoS. A stiffer adiabatic index (γ_g (gas) or γ_c (cosmic rays)) makes the system more stable; a softer EoS is destabilizing. Parker’s original assumption, $\gamma_c = 0$, makes the system very unstable (Heintz & Zweibel, 2018). Of course, not all cosmic ray effects can be expressed through an EoS, diffusion being an obvious exception. We defer a detailed discussion of transport to Section 5.3.3, but the main choices are between the self confinement model (also known as the streaming model), confinement by extrinsic turbulence, which is diffusive, and the limit of zero diffusion (also known as advection).

As a result of the EoS dependence on transport, the chosen transport model determines whether the cosmic rays stabilize or destabilize the system (Zweibel, 2017; Heintz & Zweibel, 2018). In particular, cosmic rays are *stabilizing* under the advection model. This result is puzzling, given the simple physical basis for cosmic ray buoyancy, and partly motivates the work described in this paper.

The fundamental question connecting research on the Parker instability to research on star formation is, “Under what conditions do cosmic rays change the ISM of a galaxy?” To probe this question, we examine how a single localized injection of

cosmic rays (e.g., particles accelerated in the shock waves created by supernova explosions) disrupts the ISM. We take this route instead of a traditional Parker instability seeded by small velocity perturbations because it directly connects our setup to star formation feedback. Localized cosmic ray injection has been studied previously as a driver for the galactic dynamo (Hanasz & Lesch, 2000; Hanasz et al., 2004, 2009). It produces structure similar to the Parker instability, with an extended magnetic lobe and gas condensed in parts of the disk. Focusing on the sources of cosmic rays instead of the background led us to ask the following: “*What is the timescale on which a single cosmic ray injection changes the vertical structure of the surrounding ISM?*” We break this question into four sub-questions concerning a single injection’s impact on the ISM:

- (Q1) How does the background medium’s stability change the injection’s impact?
- (Q2) How does the choice of cosmic ray transport model change the injection’s impact?
- (Q3) How do the injection’s strength and vertical location change its impact?
- (Q4) How does cosmic ray injection differ from heat injection by a thermal explosion?

By answering these questions for a range of models within an extended parameter space, our simulations provide a detailed look at how the injection of cosmic rays in a single event changes the ISM on ~ 1 kpc length scales and over times shorter than

the ~ 400 Myr required for restructuring by the Parker instability (Heintz et al., 2020; Tharakkal et al., 2023).

Overall, we show that for a stiff cosmic ray EoS, and transport by advection or diffusion, cosmic ray injection can loft magnetic fields and cosmic rays to heights exceeding a kiloparsec on timescales ≈ 100 Myr. Cosmic ray transport dominated by streaming (i.e., a strong magnetic field) results in buoyant rising, producing structure similar to the Parker instability on timescales $\gtrsim 170$ Myr. While injecting cosmic rays in a burst replicates some features of the Parker Instability, the medium responds on a timescale that is several times faster, and the dependence on transport is entirely different. Taken together our simulations show that cosmic ray injection can cause significant dynamical change on a ~ 100 Myr timescale, with the change being largest if diffusion is the dominant cosmic ray transport process, the magnetic field is weak, and the background cosmic ray pressure is large.

The paper is split into six sections including this motivation section. In Section 5.3, we provide background on the Parker instability and cosmic ray transport. In Section 5.4 we cover our numerical methods, our initial conditions for magnetohydrodynamic simulations, and our parameter choices for each simulation run. In Section 5.5 we present the overall results of our simulations. In Section 5.6 we compare the runs and discuss the implications of our simulations, while answering our four questions concerning various astrophysical parameters. In Section 5.7 we summarize the paper and provide key takeaways from the work. Readers primarily interested in our results should focus on Figures 5.5, 5.9, 5.10, & 5.11, Table 5.1, and Sections 5.6 & 5.7.

In the appendices, we provide further detail on our simulations. In Appendix 5.8.1, we compare a Parker instability simulation run with our code to one obtained previously with a significantly different code and show they give consistent results. In Appendix 5.8.2, we discuss the boundary conditions and dimensionality of our simulations. In Appendix 5.8.3, we discuss the dependence of our simulation results on various numerical parameters and our parameter choices for the primary simulations.

5.3 Background

5.3.1 Magnetohydrostatic Equilibrium

The Parker instability disrupts a magnetohydrostatic equilibrium in the galactic disk Parker (1966). With a few exceptions (Asseo et al., 1978; Boulares & Cox, 1990), the equilibrium quantities are assumed to be functions of z alone and the magnetic field lines are horizontal and straight. The condition for equilibrium is then

$$\frac{d}{dz} \left(P_g + P_c + \frac{B^2}{8\pi} \right) = \rho g. \quad (5.1)$$

Parker solved the stability problem for a particularly simple class of equilibria: fixed, constant gravity g , constant mean squared random gas velocity P_g/ρ , and uniform ratios of nonthermal to thermal gas pressures parameterized by two constants α and

β^3

$$P_B = \frac{B^2}{8\pi} = \alpha P_g \quad P_c = \beta P_g. \quad (5.2)$$

While other features of the equilibrium setup were modified in later work, the (α, β) parameterization has generally endured, and we use it to describe the pre-injection state in our simulations. However, we use the equilibrium setup from Giz & Shu (1993) as our initial ISM background. In this model the gravitational profile is smooth

$$\mathbf{g}(z) = -g_* \hat{z} \tanh\left(\frac{z}{H_*}\right). \quad (5.3)$$

instead of the discontinuous $-g_* \text{Sign}(z)$ profile Parker (1966) used. The smooth function in Equation 5.3 poses fewer numerical difficulties than a discontinuous profile in numerical simulations which extend above and below the galactic plane (Heintz et al., 2020). In Equation 5.3, the asymptotic vertical gravitational acceleration $g_* > 0$ and the gravitational scale height H_* depend on the structure of the galactic disk's stellar population.

Assuming a plane-parallel hydrostatic equilibrium, Equation 5.1 becomes

$$\frac{dP_{\text{tot}}}{dz} \equiv (1 + \alpha + \beta) \frac{dP_g}{dz} = \rho(z)g(z) \quad (5.4)$$

where P_{tot} is the total pressure; the sum of gas, magnetic, and cosmic ray pressure. We solve for the equilibrium under the influence of the gravitational profile Equation 5.3, using an isothermal equation of state $P_g = c_s^2 \rho$ with constant sound speed c_s ,

³Note Parker's β is **not** the usual plasma beta. Instead, the parameter α is the inverse of the plasma beta $\beta_{\text{pl}} = P_g P_B^{-1} = \alpha^{-1}$.

and midplane values $P_g(0) = P_{g,0}$, $\rho(0) = \rho_0$. The solution to Equation 5.4 is

$$\frac{P_g(z)}{P_{g,0}} = \frac{\rho(z)}{\rho_0} = f(z) = \operatorname{sech}^\eta \left[\frac{z}{\eta H} \right]. \quad (5.5)$$

The solution depends on how the scale height of the gas H relates to the gravitational scale height. The scale height of the gas is

$$H = \frac{c_s^2}{g_*} (1 + \alpha + \beta). \quad (5.6)$$

The ratio of the two scale heights $\eta = H_*/H$ is a constant. Taking a limit of Equation 5.5 as $\eta \rightarrow 0$ (or equivalently, $H_* \rightarrow 0$) recovers the solutions in Parker (1966).

The geometry and coordinate system for this equilibrium are illustrated in Figure 5.1. The pressures, including magnetic pressure, are homogeneous in surfaces parallel to the xy plane, which is shown as a green plane in Figure 5.1. For our simulations, we orient the initial magnetic field in the \hat{x} direction. The magnetic field is shown as a blue arrow in Figure 5.1. Using Equations 5.2 & 5.5 the magnetic field is

$$\mathbf{B}(z) = \hat{x} \sqrt{8\pi\alpha P_{g,0} f(z)}. \quad (5.7)$$

Similarly, the cosmic ray pressure is

$$P_c(z) = \beta P_{g,0} f(z). \quad (5.8)$$

In Figure 5.1, the coordinates $(\hat{x}, \hat{y}, \hat{z})$ approximately map to a galactic disk's cylindrical coordinates. The \hat{x} direction is parallel or anti-parallel to the azimuthal




Figure 5.1: This schematic shows our initial setup for hydrostatic equilibrium. We have a gravitational field (red arrows) in the \hat{z} direction, which flips sign at the midplane $z = 0$ (see Equation 5.3). The magnetic field (blue arrow) is oriented along the \hat{x} direction and creates a magnetic pressure supporting vertical hydrostatic equilibrium (Equations 5.4,5.7). We then inject cosmic rays, shown by an orange circle (see Section 5.4.2 for injection profile), above the midplane, which is shown as a green rectangle.

direction $\hat{\phi}$ (depending on whether the magnetic field is oriented clockwise or counterclockwise around the galactic center), the \hat{y} direction is parallel or anti-parallel to the radial direction \hat{r} (again, depending on magnetic field orientation), and the \hat{z} direction is parallel to the cylindrical \hat{z} direction.

5.3.2 Parker Instability

We can assess the stability of Parker's equilibrium to ideal (energy-conserving) small amplitude perturbations with the generalized Schwarzschild convection criterion (Newcomb, 1961; Boulares & Cox, 1990; Zweibel, 2017). This criterion for instability in a vertically stratified atmosphere is

$$-\frac{d \ln \rho}{dz} < \frac{\rho g_*}{\gamma_g P_g}. \quad (5.9)$$

This criterion is applicable above the midplane with $z > 0$. Assuming we are well above the midplane ($z > H_*$) where the gravitational acceleration is constant ($-g_*$), then the logarithmic derivative is the scale height H of the gas. This assumption gives the same criterion as Parker’s original gravity profile. In this limit, the criterion for instability to occur is

$$\alpha + \beta > \gamma_g - 1. \quad (5.10)$$

If we include the compressibility of cosmic rays, then we need to replace $\gamma_g P_g$ with $\gamma_g P_g + \gamma_c P_c$ in Equation 5.9. Using this change and evaluating the criterion again, we get

$$\alpha + \beta(1 - \gamma_c) > \gamma_g - 1. \quad (5.11)$$

The above criterion illustrates a confusing result noted in Zweibel (2017), and again in Heintz & Zweibel (2018). Generically, one would guess increasing non-thermal pressures makes the system more unstable. However, Equation 5.11 shows there is a way for increasing cosmic ray pressure (β) to make the system more stable. Since the cosmic ray fluid is relativistic gas, it has $\gamma_c = \frac{4}{3}$. In that case, the left hand side would be smaller with increasing β . The physical explanation for this effect is that compressing the cosmic rays requires work.

Self confinement by streaming leads to the polytropic relation $P_c \propto \rho^{2/3}$ along magnetic flux tubes. Using this value in Equation 5.11 would predict a stability threshold intermediate between “classic” Parker ($\gamma_c = 0$) and $\gamma_c = 4/3$. Instead, it is shown through a modal analysis in Heintz & Zweibel (2018) that self confinement leads to a even larger domain of instability and faster growth rates than “classic”




Figure 5.2: This schematic shows the parameter space of magnetic pressure $\alpha = P_B/P_g$ and cosmic ray pressure $\beta = P_c/P_g$. For advective transport, we show the regions of parameter space where the initial hydrostatic equilibrium (see Section 5.3.2) is unstable to the Parker instability. The gray region is stable if cosmic rays have an adiabatic index $\gamma_c = 4/3$. The black region is stable if the cosmic rays are an infinitely compressible fluid with index $\gamma_c = 0$. The S, U, B, and C points are where the simulations, listed in Table 5.1, are located in this parameter space.

Parker with $\gamma_c = 0$. In fact, the analysis used to derive Equation 5.9 does not apply to the streaming model, due to the relative drift between cosmic rays and thermal gas, and the heating that accompanies streaming. Similarly, Equation 5.9 does not apply to diffusion dominated cosmic ray transport.

We show a schematic for advection dominated transport in Figure 5.2. This figure is not quantitatively predictive of our simulation setup, which has a smooth g_* given by Equation 5.3, a non-ideal perturbation, and in some cases, transport dominated by streaming or diffusion. The schematic shows the (α, β) space of equilibrium solutions with an adiabatic gas exponent $\gamma_g = 5/3$. Using Equations 5.10 and 5.11, we show different regions of stability determined by the cosmic ray adiabatic index γ_c . There is a strictly Parker unstable region (unshaded), a strictly Parker stable region (black), and a gray region which is stable if the cosmic rays take work to compress (assuming

$\gamma_c = 4/3$). For smaller values of γ_g , the fan of stability boundaries shifts downward so the boundaries intersect on the vertical axis where $\alpha = \gamma_g - 1$.

To reiterate: We choose simulation parameters throughout the (α, β) parameter space to probe different background medium conditions. This short review of stability of the background medium illustrates how our different background medium models would be classified in a Parker instability analysis, and provides a short review of the Parker instability and its dependence on cosmic ray transport.

5.3.3 Cosmic Ray Transport

The stability of the system described in Section 5.3.2 and illustrated in Figure 5.2 depends on the effective compressibility of the cosmic rays, and therefore on cosmic ray transport. Cosmic rays are transported throughout the interstellar medium by three mechanisms: advection, diffusion, and streaming. The effects of these different transport mechanisms have already been examined in global galactic models (Girichidis et al., 2016; Wiener et al., 2017b).

All three mechanisms are based on the assumption that the cosmic rays exchange momentum and energy by scattering from fluctuations in the magnetic field with a wavelength of order the cosmic ray gyroradius (so-called gyroresonant scattering). Advection, which applies in the limit of infinitely short mean free path and scattering by fluctuations with no preferred propagation direction along the background magnetic field lines, is the most straightforward: if the thermal gas has a bulk flow in a particular direction, the cosmic rays should flow with the gas. If advection is dominant, then the cosmic rays behave like a relativistic fluid, with negligible inertia,

following the flow of the thermal (non-relativistic) gas.

Diffusion of cosmic rays applies when the fluctuations again have no preferred propagation direction, but the mean free path to scattering is large enough to allow the cosmic rays to leak through the thermal gas. If the measured residence time and vertical scale height of cosmic rays in the Milky Way are explained by diffusive propagation, the implied diffusion coefficient is $\kappa_c \sim 3 \cdot 10^{28} \text{ cm}^2 \text{ s}^{-1}$.

Streaming, or self confinement, applies when the magnetic fluctuations are generated by the cosmic rays themselves as a result of directional anisotropy in the frame of the fluctuations. It can be shown that in the short mean free path regime, the streaming direction is down the cosmic ray pressure gradient. This instability appears in the kinetic theory for cosmic ray transport, and it drives the cosmic rays into a bulk flow at the Alfvén speed $\mathbf{v}_A \equiv B/\sqrt{4\pi\rho}$ (Kulsrud & Pearce, 1969; Wentzel, 1969; Skilling, 1975). The streaming instability also heats the gas: the cosmic rays transfer energy to hydromagnetic waves through gyro-resonance, which then dissipate energy into the thermal gas (Kulsrud & Pearce, 1969; Wentzel, 1969; Skilling, 1975). This heating appears as a source term in the equation for the thermal energy density and the cosmic ray energy density. The heating rate from the streaming instability is $\mathbf{v}_A \cdot \nabla P_c$ (see Zweibel (2017) for discussion and references).

Each of these transport methods drives a different characteristic response in the ISM when cosmic rays are injected. Advection and streaming both drive steep fronts of gas and cosmic ray energy. Whereas advection will result in the cosmic ray fluid moving at the flow speed, bringing the gas with it, streaming allows the cosmic rays to move ahead of the gas if the Alfvén speed is faster than the sound speed. In the

streaming case, a front of cosmic rays will move through the gas while heating the gas according to the Alfvénic heating term $\mathbf{v}_A \cdot \nabla P_c$. Diffusion, in contrast to both streaming and advection, drives a smoother flow of cosmic rays while also exerting a force on the gas along the cosmic ray pressure gradient.

We provide the following useful timescales for each transport mechanism to move cosmic rays through a distance L of the ISM:

$$\tau_{\text{adv}} = L \cdot \left(\frac{\gamma_g P_0}{\rho_0} \right)^{-1/2} \approx 65 L_{\text{kpc}} T_4^{-1/2} \text{Myr}, \quad (5.12)$$

$$\tau_{\text{str}} = L \cdot \left(\frac{B_0}{\sqrt{4\pi\rho_0}} \right)^{-1} \approx 35 L_{\text{kpc}} \rho_{-24}^{1/2} B_\mu \text{Myr}, \quad (5.13)$$

$$\tau_{\text{diff}} = \frac{L^2}{\kappa_c} \approx 30 L_{\text{kpc}}^2 \kappa_{28}^{-1} \text{Myr}, \quad (5.14)$$

where L_{kpc} , T_4 , B_μ , ρ_{-24} , and κ_{28} are the length scale, gas temperature, gas density (assuming full ionization), and cosmic ray diffusivity in units of 1 kpc, 10^4 K, $1\mu\text{G}$, $10^{-24} \text{g cm}^{-3}$, and $10^{28} \text{cm}^2 \text{g}^{-1}$. In our simulations, the gas flows are generally subsonic, so the advection timescale is an imperfect measure: the correct timescale will be larger. For advection dominated simulations, we set the simulation parameters such that the streaming and diffusion timescales are much larger than an order unity multiple of the advection timescale in Equation 5.12. However, when the advection and streaming timescales are similar according to Equations 5.12 & 5.13, the streaming transport dominates the dynamics in our simulations because of this overestimated flow speed. We only use these timescales to determine the dominant transport mechanism, defaulting to streaming when the streaming and advection

timescales are similar (i.e. Simulation U - see Table 5.1).

5.4 Setup

5.4.1 Basic Equations and Numerical Methods

Our simulations are based on fluid equations for the thermal gas - cosmic ray - magnetic field system Breitschwerdt et al. (1991); Zweibel (2017). These equations accommodate cosmic ray transport by both diffusion and streaming; for the latter, it is assumed that the scattering waves propagate down the cosmic ray pressure gradient. Enforcing this feature creates numerical difficulties near extrema in cosmic ray pressure. Although this can be handled through a smoothing length (Sharma et al., 2010), we adopt the two moment formulation of streaming transport developed by Jiang & Oh (2018) (see also Thomas et al. (2021)) as implemented in the Athena++ code (Stone et al., 2020), which modifies the equations presented in Breitschwerdt et al. (1991) through a time dependent evolution equation for the cosmic ray flux.

The implementation of Athena++ from Jiang & Oh (2018) solves the following equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (5.15)$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{v} \times \mathbf{B}) = 0, \quad (5.16)$$

$$\frac{\partial (\rho \mathbf{v})}{\partial t} + \nabla \cdot \left(\rho \mathbf{v} \mathbf{v} + \left(P_g + \frac{1}{2} B^2 \right) \mathbb{1} - \mathbf{B} \mathbf{B} \right) = \rho \mathbf{g} + \hat{\boldsymbol{\sigma}}_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{v} E_c), \quad (5.17)$$

$$\begin{aligned} \frac{\partial E}{\partial t} + \nabla \cdot \left(\mathbf{v} \left(E + P_g + \frac{1}{2} B^2 \right) - \mathbf{B} (\mathbf{B} \cdot \mathbf{v}) \right) \\ = \rho \mathbf{g} \cdot \mathbf{v} - (\gamma_c - 1) (\mathbf{v} + \mathbf{v}_s) \cdot \nabla E_c, \end{aligned} \quad (5.18)$$

$$\frac{\partial E_c}{\partial t} + \nabla \cdot \mathbf{F}_c = (\gamma_c - 1) (\mathbf{v} + \mathbf{v}_s) \cdot \nabla E_c, \quad (5.19)$$

$$\frac{1}{V_m^2} \frac{\partial \mathbf{F}_c}{\partial t} + (\gamma_c - 1) \nabla E_c = -\hat{\boldsymbol{\sigma}}_c \cdot (\mathbf{F}_c - \gamma_c \mathbf{v} E_c). \quad (5.20)$$

The dynamical variables are gas density ρ , bulk flow velocity \mathbf{v} , gas pressure P_g , magnetic field \mathbf{B} , the combined internal and kinetic energy density $E = P/(\gamma_g - 1) + \rho v^2/2$, the cosmic ray energy density $E_c = P_c/(\gamma_c - 1)$, and the cosmic ray energy flux \mathbf{F}_c . We include the effects of a gravitational acceleration \mathbf{g} , given by Equation 5.3. We use $\gamma_c = 4/3$ and set the modified speed of light parameter to $V_m = 0.3c$. For more discussion on convergence according to this parameter, see Appendix 5.8.3.

5.4.2 Initial Conditions

The initial profiles of density, gas pressure, magnetic field, and cosmic ray pressure are described in Section 5.3. We use midplane density and pressure values of $\rho_0 = 10^{-24} \text{ g cm}^{-3}$ and $P_{g,0} = 10^{-12} \text{ g cm}^{-3}$. We use an asymptotic gravitational acceleration $g_* = 4 \cdot 10^{-9} \text{ cm s}^{-2}$ and a gravitational scale height $H_* = 100 \text{ pc}$ (Heintz et al., 2020). This height is smaller than what is used in Parker instability simulations (see App. 5.8.1 and Rodrigues et al. (2016); Heintz et al. (2020)). Additionally, this height decreases our initial profile's height ratio η when compared to Milky Way and Parker instability simulations (Rodrigues et al., 2016; Heintz et al., 2020). This adjustment means our injection can occur in a region where the gravity has nearly

reached the asymptotic value g_* . As the flux tube of interest rises, it will move in a constant gravitational acceleration.

The simulation grid is $200 \times 50 \times 300$ cells (ordered with respect to $(\hat{x}, \hat{y}, \hat{z})$). The cell size is $10 \times 20 \times 10 \text{ pc}^3$, giving a total simulation volume $2 \times 1 \times 3 \text{ kpc}^3$. The third dimension (\hat{z}) is off-center of the midplane ($z = 0$), extending from $z = -1 \text{ kpc}$ to $z = +2 \text{ kpc}$. The cosmic rays are injected above the midplane ($z_0 > 0$), so this off-centering focuses the computational resources on the injection and the resulting flows. We do not use any adaptive or static mesh refinement. We adopt outflow boundary conditions in the \hat{x} and \hat{y} directions and vacuum (diode) boundary conditions in the \hat{z} direction to minimize the effect of domain size. See Appendix 5.8.2 for additional discussion.

The equilibrium is setup to within single-precision floating point numerical error, which is less than the density and pressure floors of our simulations. There are some waves created by the interaction of the equilibrium with the vertical boundaries where there is an abrupt transition to vacuum. These waves are small and have no long term effect on the dynamics of our simulation.

The cosmic ray energy density injection profile is a 3D Gaussian function:

$$\varepsilon_c(x, y, z) = \frac{E_{\text{SN}}}{(2\pi r_{\text{SN}}^2)^{3/2}} \cdot \exp\left[\frac{-1}{2r_{\text{SN}}^2} |\mathbf{r} - \mathbf{r}_0|^2\right]. \quad (5.21)$$

The parameters of the injection are its position \mathbf{r}_0 , radius r_{SN} , and total energy injected E_{SN} . We add this perturbation onto the background cosmic ray pressure profile in Equation 5.5. Integrating the profile, Equation 5.21, over volume gives E_{SN} as the total energy injected. The injection occurs at different heights above the

midplane depending on the other parameters under consideration (see Table 5.1).

In each simulation, we use the same injection radius $r_{\text{SN}} = 50 \text{ pc}$. This radius is large for a single supernova shock, but reasonable for our default energy injection from ~ 10 supernovae (see Section 5.4.3). This larger radius allows our simulation (with 10–20 pc resolution) to reliably sample the injection profile. To avoid sampling errors (aliasing) which will change the total injected energy, we center the injection in the xy plane on a cell center. The grid has cell faces along $x = 0$ and $y = 0$ planes, so we choose to center the injection at $x_0 = 5 \text{ pc}$ and $y_0 = 10 \text{ pc}$. As a result of this placement, the peak of the profile occurs in the center of a computational cell. The height z_0 and total energy E_{SN} are varied across our simulations (see 5.1)

For reference, if the pressure of 10^{51} erg injected cosmic rays were uniformly distributed throughout the volume of the injection tube (radius 50 pc, length 2 kpc) is $\sim 0.73 \times 10^{-12} \text{ dynes cm}^{-2}$, corresponding roughly to a doubling of background cosmic ray pressure. In this strictly 1D situation, with no horizontal pressure gradient, the tube would expand radially. This expansion would reduce its density and cause it to float upward buoyantly, with a characteristic rise time $\sim 100 \text{ Myr}$ (for our adopted gravitational field and scale height).

5.4.3 Explanation of Parameters

We present the results of ten simulations of cosmic ray injection. The initial parameters for each simulation are shown in Table 5.1. We choose the parameters for the simulations to help us answer the four questions, presented at the end of Section 5.2

Run Name	α	β	γ_g	H (pc)	η	z_0 (pc)	Weight (M_\odot)	E_{SN} (erg)	CR	$\kappa_{c,\parallel}$ ($\text{cm}^2 \text{s}^{-1}$)	Question
S	0.1	0.1	5/3	97	1.03	45	$1.28 \cdot 10^4$	10^{51}	Adv.	$3 \cdot 10^{24}$	Q1
U	1.0	1.0	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{51}	Str.	$3 \cdot 10^{24}$	Q1, Q3, Q4
Uiso	1.0	1.0	1.1	243	0.41	105	$2.24 \cdot 10^4$	10^{51}	Str.	$3 \cdot 10^{24}$	Q1
Cadv	0.1	1.9	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{51}	Adv.	$3 \cdot 10^{24}$	Q1, Q2
Bstr	1.9	0.1	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{51}	Str.	$3 \cdot 10^{24}$	Q1, Q2
Cdiff	0.1	1.9	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{51}	Diff.	$3 \cdot 10^{28}$	Q1, Q2
Bdiff	1.9	0.1	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{51}	Diff.	$3 \cdot 10^{28}$	Q2
Uheavy	1.0	1.0	5/3	243	0.41	25	$3.27 \cdot 10^4$	10^{51}	Str.	$3 \cdot 10^{24}$	Q3
Ublast	1.0	1.0	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{52}	Str.	$3 \cdot 10^{24}$	Q3
Utherm	1.0	1.0	5/3	243	0.41	105	$2.24 \cdot 10^4$	10^{51*}	Str.	$3 \cdot 10^{24}$	Q4

Table 5.1: This table shows the parameters we vary for each simulation run we perform. The first column shows each simulation run name. The next columns show background medium parameters: magnetic pressure to gas pressure ratio α , cosmic ray pressure to gas pressure ratio β , thermal adiabatic index γ_g , gas scale height H , and profile exponent η (see Equation 5.5). The second set of columns shows injection parameters: initial height of the injection z_0 , the weight in the column above the injection, and the volume integrated injection energy E_{SN} (see Equation 5.21). The third set of columns shows dominant cosmic ray transport mechanism and the parallel cosmic ray diffusion coefficient $\kappa_{c,\parallel}$. The final column shows which questions the simulation applies to (see Sections 5.2 & 5.4.3). The parameters held constant for all simulations are: midplane density $\rho_0 = 10^{-24} \text{ erg cm}^{-3}$, midplane pressure $P_0 = 10^{-12} \text{ erg cm}^{-3}$, sound speed $c_s = 10^6 \text{ cm s}^{-1}$, temperature $T = 1.21 \cdot 10^4 \text{ K}$, stellar scale height $H_* = 100 \text{ pc}$, asymptotic gravitational acceleration $g_* = 4 \cdot 10^{-9} \text{ cm s}^{-2}$, injection radius $r_{\text{SN}} = 50 \text{ pc}$, perpendicular cosmic ray diffusion coefficient $\kappa_{c,\perp} = 3 \cdot 10^{24} \frac{\text{cm}^2}{\text{s}}$, and modified speed of light $V_m = 0.3c$. The midplane magnetic field for each simulation is given by $B_0 = 5.06 \mu\text{G}\sqrt{\alpha}$. The advective, streaming, and diffusive timescales are $\tau_{\text{adv}} = 76 \text{ Myr } (\gamma_g/(5/3))^{-1/2}$, $\tau_{\text{str}} = 69 \text{ Myr } (\alpha)^{-1/2}$, and $\tau_{\text{diff}} = 10 \text{ Myr } (\kappa_{c,\parallel}/10^{28} \text{ cm}^2 \text{ s}^{-1})^{-1}$, respectively.

and labeled Q1, Q2, Q3, and Q4.

We first considered what values of α and β are reasonable in the solar neighborhood. Measurements suggest the pressures (thermal, cosmic ray, and magnetic) are nearly equal (Ferrière, 2001). Therefore, we choose a base run with $\alpha = \beta = 1$. Simulations with these values are named with a prefix U. In Figure 5.2, this is on a stability boundary when cosmic ray streaming and diffusion are not included. However, each simulation in Table 5.1 includes transport by diffusion (even when it is not dominant), and most include streaming. Therefore, the stability implied by Figure 5.2 does not strictly apply, suggesting this equal pressure background medium is Parker unstable. To compare this background with a strictly Parker stable background, we run a simulation with $\alpha = \beta = 0.1$. We use a prefix ‘S’ to refer to this combination of values. Our first two simulations, labelled U and S, allow us to address Q1: how the stability of the background medium changes the effect of cosmic ray injection. For both these simulations, we use a gas adiabatic constant consistent with a monatomic ideal gas $\gamma_g = 5/3$, a negligible cosmic ray diffusion constant $\kappa_c = 3 \cdot 10^{24} \text{ cm}^2 \text{ s}^{-1}$, and an injection energy equivalent to the estimated cosmic ray injection energy of ~ 10 supernovae (Caprioli & Spitkovsky, 2014), such as might be expected from an association of coeval massive stars.

Our third simulation, Uiso, has the same parameters as U except for the adiabatic gas constant $\gamma_g = 1.1$, which brings the system closer to the isothermal case originally considered by Parker (1966) and is expected to be more unstable. This simulation also addresses Q1.

To answer Q2, how cosmic ray transport affects the behavior of the system post

injection, we need a way to differentiate the dominant cosmic ray transport mechanism. We do this by considering each transport timescale, given by Equations 5.12, 5.13, 5.14. It turns out that the flows in our simulations are subsonic and our advection timescale assumes propagation at the sound speed. So, when streaming and advection timescales are close, we assume streaming dominates. Simulations **U** and **S**, are instances of streaming and advection dominance, respectively.

To better isolate the impact of each transport mechanism, we probe two other points in the (α, β) plane of Figure 5.2. We run a simulation with a large magnetic field $\alpha = 1.9$ and low amount of background cosmic rays $\beta = 0.1$, so it is in the top left corner of Figure 5.2. This simulation is labelled **Bstr**, and streaming dominates the transport of the injected cosmic rays because the Alfvén speed is much higher than the flow speed and the diffusion rate. We also run a simulation **Cadv** with $\alpha = 0.1$ and $\beta = 1.9$, placing it in the bottom right hand corner of Figure 5.2. With this simulation, we probe a medium which is Parker stable due to the stiff equation of state with large non-thermal pressure where streaming and diffusion are subdominant transport mechanisms relative to advection. The simulation **Cdiff** has the same values of α and β , but uses a diffusion constant $\kappa = 3 \cdot 10^{28} \text{ cm}^2 \text{ s}^{-1}$ which is close to estimated Milky Way values.

While these three simulations allow us to examine how streaming, advection, and diffusion affect the response of the medium to cosmic ray injection, they do not completely eliminate the influence of the background state; while the nonthermal pressures are equal in all three, the role of magnetic tension in **Bstr** is larger than in **Cadv** or **Cdiff** due to the stronger background magnetic field. Therefore, to determine

whether the streaming cosmic rays or the strong magnetic field was more critical, we ran `Bdiff`. That simulation had the same initial parameters as `Bstr`, but with a higher diffusion coefficient and streaming terms removed. Although diffusion and advection do not have an ‘off-switch’ in the implementation from Jiang & Oh (2018), streaming transport does, and we make use of it when running `Bdiff`.

`Bstr` and `Cadv` also address Q1 because they are at different points in the (α, β) plane. They have different background cosmic ray and magnetic pressures from simulations `U` or `S`, allowing us to learn how that aspect of the background medium changes the impact of cosmic ray injection.

Question 3 probes the effects of injection properties, so we adjust parameters related to the injection. Using simulation `U` as a control case, we compare both of the following simulations to it, and only change single parameters. First, we use simulation `Uheavy` with the cosmic ray injection at a lower height in the disk, meaning it has more weight above it. Second, we use a simulation, `Ublast`, with the strength of 10% of 100 supernovae to understand how different amounts of energy injection change the results. Other than the weight above the injection and the injection energy, we keep the variables the same as simulation `U`, which we compare each of these cases against.

For Q4, which compares the effects of cosmic rays to direct energization of thermal gas, we run a simulation, `Utherm`, with the same initial parameters as the Parker unstable simulation `U`, except the injection is in thermal pressure instead of in cosmic ray pressure.




Figure 5.3: After injecting cosmic rays on a magnetic flux tube, we track that tube’s movement. For each simulation shown, the solid black line is maximum height of the flux tube at a given time. The green dotted line is the average mass in the region $x < |250 \text{ pc}|$, as a fraction of the initial average mass. The purple solid, dashed, and dash-dotted lines are the gas, magnetic, and cosmic ray pressure, respectively, averaged in the same region. The second y axis on the right side applies to the lines of mass and pressures, with each quantity $Q \in \{\langle M \rangle, \langle P_g \rangle, \langle P_B \rangle, \langle P_c \rangle\}$ normalized by either $P_0 = 10^{-12} \text{ erg cm}^{-3}$ or the initial value (for mass $\langle M \rangle$).

5.5 Results

We focus on a flux tube enclosing the initial injection. Since we use the ideal MHD equations and because there is effectively no perpendicular cosmic ray diffusion, the gas and cosmic rays stay on the flux tube. All the simulations share some common features. The cosmic rays in the injection sphere are overpressured by more than a factor of 20. This creates a low density cavity within a few Myr and launches a pressure driven flow away from the injection sphere. The cavity itself is buoyant because of reduced density and the outflow reduces the density along a progressively longer segment of the tube. The result is a rising magnetized arch which becomes ever more buoyant as gas is drained from the tube by a gravity driven downflow. The only forces that counter the rise are magnetic tension due to field line curvature, the inertia of the overlying gas, and adiabatic cooling of both the thermal and cosmic ray gases as they expand. The configuration has strong up-down asymmetry, because the effects which slow the flux tube's rise increase near the midplane.

For each simulation, we examine physical quantities near the center of the tube ($x = \pm 250$ pc). This restriction allows us to focus on the dynamics in the center of the flux tube, where the initial cosmic ray injection occurred. Additionally, it minimizes the impact of the boundary conditions on our results for early times in the simulation (See Appendix B for a discussion of boundary conditions). In Figure 5.3 we show the maximum height along the flux tube as a solid black line, the average gas pressure as a solid purple line, the average magnetic pressure as a dashed purple line, the average cosmic ray pressure as a dash-dotted purple line, and the average mass as a dotted green line. The average flux tube height in the ± 250 pc region is approximately the

same as the maximum height. These quantities help us understand how the injection evolves through time and how it causes the flux tube to change.

The most striking part of the results in Figure 5.3 are the plots for simulations **Cadv**, **Cdiff**, **Ublast**, and **Utherm**. In these four simulations, the flux tubes rise over 400 pc in less than 100 Myr. The rise is accompanied by decreases in mass and pressure in the central region of the flux tube.

The other simulations do not exhibit such a quick, drastic rise in flux tube height, nor decreases in pressure and mass in the central region of the flux tube. Simulations **S** and **Uheavy** end up being stable to the perturbation (meaning the flux tube rise is limited). **S** even finds a new equilibrium after an initial rise. Simulations **U**, **Uiso**, and **Bstr** all eventually begin to buoyantly rise, but only rising to 200 pc, at a time of 150 Myr after the cosmic ray injection. These simulations each have streaming as the dominant cosmic ray transport - instead of creating a violent disruption in the center of the magnetic flux tube, streaming puts energy into heating the entire length of the magnetic flux tube. However, the stronger injection in **Ublast** still launches material rapidly when compared to **U**, even though streaming is dominant.

To more clearly illustrate the dynamics in our simulations, we show 2D cuts from simulation **Cdiff** in Figure 5.4. The first column shows gas density, the second column shows vertical momentum, and the third column shows horizontal momentum. The green lines are magnetic field lines along the flux tube analyzed in Figure 5.3. The first row is a snapshot at 4 Myr, the second row at 60 Myr, and the third row at 120 Myr. Once the cosmic rays diffuse, they create an over-pressured flux tube which begins to move upward (middle column, top row). After the flux tube bends,

gas begins to fall down the curved lines (right column, middle row). After some time, this process builds until a large up-flow in the center and down-flow along the magnetic field lines. The final density plot (left column, bottom row) also shows how the flux tube is able to lift some mass out of the disk.

We show the cosmic ray pressure distributions for simulations **Cadv**, **Bstr**, and **Cdiff** in Figure 5.5. The first column shows simulation **Cadv**, the second column shows **Bstr**, and the third column shows **Cdiff**. The first and second rows show the same time dumps for each simulation (top row is 3 Myr, middle row is 50 Myr). The final row shows the later development of the injection for each simulation. The **Bstr** simulation does not grow as much vertically because magnetic tension holds the flux tube down for most of the simulation. Only once the field lines are able to bend does the flux tube begin to rise, because gas is able to leave the flux tube at a faster rate. This effect is also shown in the second row of plots in Figure 5.3. Regardless of cosmic ray transport by streaming or diffusion, the strong magnetic field simulations (**Bstr** and **Bdiff**) exhibit slowed or negligible rise of the flux tube. Streaming appears to be the most effective in the large magnetic field case because the gas gets heated in addition to being overpressured.

5.5.1 A Caveat on Simulation *Utherm*

In simulation **Utherm**, instead of injecting cosmic rays, we injected energy as thermal heating. This thermal energy injection had the same magnitude, 10^{51} erg, as the cosmic ray injection in simulation **U**.

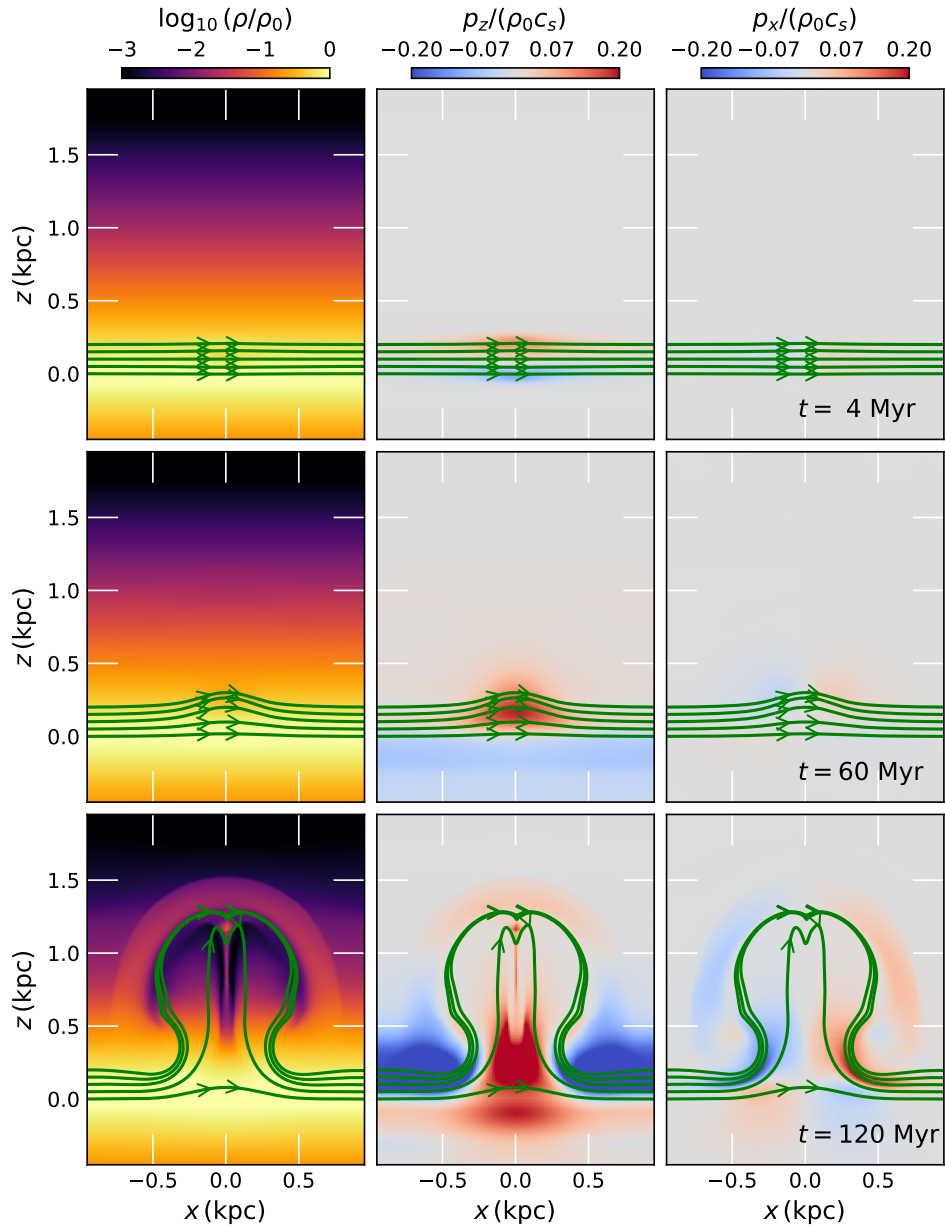


Figure 5.4: Plots of gas density (first column), vertical momentum (second column), and horizontal (\hat{x}) momentum (third column) at a $y = 0$ cut of simulation *Cdiff*. The cosmic rays diffuse along the flux tube, leading to an over-pressured region which rises buoyantly. The normalization constants are $P_0 = 10^{-12} \text{ erg cm}^{-3}$ and $\rho_0 c_s = 10^{-18} \text{ g cm}^{-2} \text{ s}^{-1}$. The green lines are magnetic field lines defining the flux tube where the injection initially occurred.

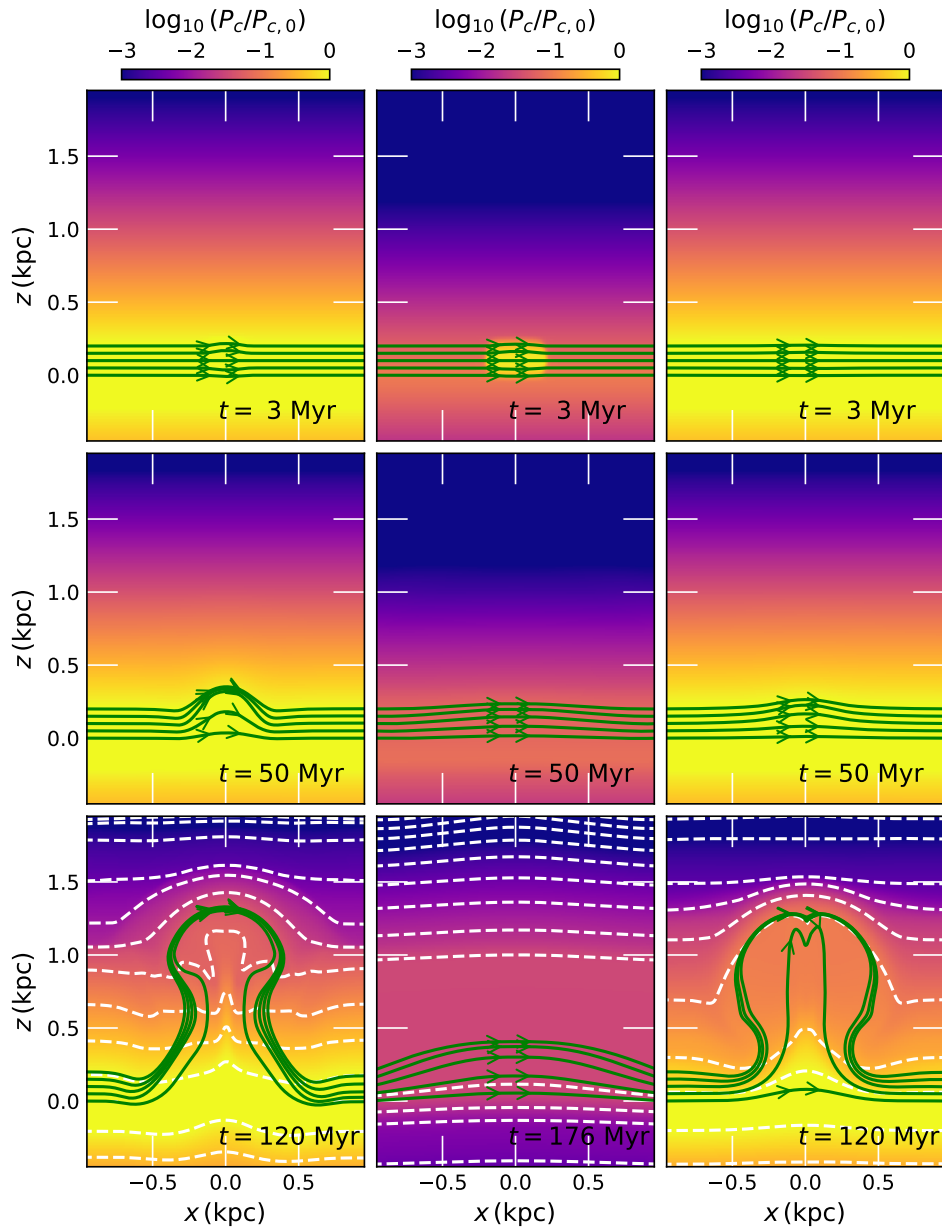


Figure 5.5: Plots of cosmic ray pressure at a $y = 0$ cut of simulations **Cadv** (first column), **Bstr** (second column), and **Cdiff** (third column). The normalization constant is $P_0 = 10^{-12} \text{ erg cm}^{-3}$. The green lines are magnetic field lines. The dashed white lines are contours of equal cosmic ray pressure. In the first two rows, cuts are all from the same time snapshots, but the final row shows simulation **Bstr** at a later time snapshot than the other two simulations.

While this simulation highlights the effectiveness of thermal energy injection, our simulations lack radiative cooling. This injected gas should cool by emitting radiation while it expands. We estimate a radiative cooling time scale to determine the effect of radiative cooling on the injection, using Equation 34.4 from Draine (2011):

$$t_{\text{cool}} = 1.1 \cdot 10^5 \text{ yr} \left(\frac{T}{10^6 \text{ K}} \right)^{1.7} \left(\frac{n_{\text{H}}}{\text{cm}^3} \right)^{-1} \quad (5.22)$$

For simulation **Utherm**, the peak temperature in the center of the injection is $10^{5.04}$ K, compared to the background $T_0 \sim 10^4$ K. Therefore, the radiative cooling time for this injection is $t_{\text{cool}} \approx 2.2$ kyr. This time scale is so short that this thermal injection should dissipate before causing a disruption in the ISM like we see in simulation **Utherm**. While temperatures are even higher in actual supernova remnants ($\sim 10^7$ K for the Sedov-Taylor phase (Draine, 2011)), our injection represents a long term average impact of those remnants. Radiative cooling will limit the impact of the thermal injection, as compared to a cosmic ray injection. A similar effect is seen in models of supernova heated gas launched from galactic disks, which tend to be fountains rather than winds (Shapiro & Field, 1976; Rosen & Bregman, 1995; Bustard et al., 2018, 2020).

5.6 Discussion

To the questions posed in Section 5.2, we now have the following answers:

- (Q1) A purely Parker stable medium limits the effect of cosmic ray injection. An isothermal medium is more prone to disruption by a cosmic ray injection.

See Figure 5.6.

- (Q2) Cosmic ray advection and diffusion drive changes in the ISM on ~ 100 Myr time scales, through explosive launching and buoyancy, respectively. Large magnetic field strength (which also implies streaming dominated transport) delays the flux tube's rise, but eventually causes buoyant rising. In the large magnetic field case, streaming is more disruptive than diffusion. See Figure 5.7.
- (Q3) Stronger injections drive more explosive flows, and injections closer to the midplane take longer to launch the flux tube. See Figure 5.8.
- (Q4) Thermal injection drives buoyant rising of the flux tube on a shorter time scale than cosmic ray injection. Cosmic rays decrease the average mass along the flux tube at a slower rate. Eventually, cosmic ray injections overtake thermal injections in height. However, simulation `Utherm` overestimates the efficacy of thermal injection because it lacks cooling (see Section 5.5.1). See Figures 5.8, 5.9, 5.10, 5.11.

In the following subsections, we provide more complete explanations for these answers.

5.6.1 Dependence on Background Medium (Q1)

We explored how changes in the background medium affected the evolution of a cosmic ray injection. The results are shown in Figure 5.6, which plots the average mass and height along the flux tube of each simulation.




Figure 5.6: Plot of maximum tube height (black) and average mass (green) within $x = \pm 250$ pc of the injection location for simulations related to Q1. Note the vertical axis has been extended compared to Figure 5.3. Solid lines show simulation **S**, dashed lines show simulation **U**, dash-dotted lines show simulation **Uiso**, dotted lines show simulation **Bstr** and dots show simulation **Cadv**. The streaming dominated simulations, **U** and **Bstr**, rise late. The advection dominated simulations, **Cadv** and **S**, have different results: the medium with more cosmic ray pressure is significantly disrupted. Simulation **Uiso** with $\gamma_g = 1.1$ allows more mass loss later in the simulation, when compared with simulation **U**, which had $\gamma_g = 5/3$.

The Parker stable simulation (Simulation **S**) weathers the cosmic ray injection, at least under transport by advection and streaming, reaching a new stable equilibrium. When the background is Parker unstable according to the original criterion (Equation 5.10), the injection drives significant change. For the simulations **U**, **Bstr**, and **Cadv**, we can use Equation 5.10 to show the combined nonthermal pressures dominate the gas compressibility ($\alpha + \beta = 2$) $>$ ($\gamma_g - 1 = 2/3$). The simulations with a background biased towards magnetic pressure (**Bstr**) or cosmic ray pressure (**Cadv**) launch the flux tube faster than the equipartition case (**U**). Therefore, the initial nonthermal pressures, given by (α, β) , determine how prone the system is to disruption by cosmic ray injection. This dependence also suggests an ISM with non-uniform α, β would provide a complex environment for cosmic ray injection, since different directions and positions could be more (or less) prone to being disrupted.

The case of an isothermal-like background with $\gamma_g = 1.1$ responds on a slightly shorter timescale than a background medium with $\gamma_g = 5/3$. This conclusion comes from considering the average height achieved by the flux tubes in simulations `Uiso` and `U`. In simulation `Uiso`, buoyant rising begins after 80 Myr and stays significant through the end of the simulation. For simulation `U`, buoyant rising becomes dominant only after 140 Myr. The average mass also decreases at a faster rate at late times in simulation `Uiso`. Because γ_g is lower, it takes less work to compress the gas and push it off the flux tube once buoyancy kicks in. This change allows mass to flow at a faster rate, driving the buoyancy force to become larger.

5.6.2 Dependence on Cosmic Ray Transport (Q2)

We ran four simulations to determine the impact of cosmic ray transport: simulation `Bstr` is streaming dominated, simulation `Cadv` is advection dominated, `Cdiff` is diffusion dominated, and `Bdiff` is diffusion dominated with a strong magnetic field. With these simulations, it is clear that streaming does a poor job of launching material when compared to diffusion and advection, which both disrupt the ISM on a short time scale ~ 100 Myr. However, `Bdiff` shows this delay is not a result of the cosmic ray transport by streaming. Instead, the large magnetic field (necessary for streaming to be dominant) resists any bending created by the injection. Streaming is more effective than diffusion at disrupting the vertical structure when magnetic tension is a dominant force. Figure 5.7 shows the flux tube averaged mass and height for these three simulations.

Simulation `Bstr` does not start to rise buoyantly until a significant amount of




Figure 5.7: Plot of maximum tube height (black) and average mass (green) within $x = \pm 250$ pc of the injection location for simulations related to Q2. Note the vertical axis has been extended compared to Figure 5.3. Solid lines show simulation Cdiff, dashed lines show simulation Cadv, dash-dotted lines show simulation Bstr, and dotted lines show simulation Bdiff. The streaming case has the flux tube rise buoyantly, but only after a long time $\gtrsim 100$ Myr. In contrast, the advection and diffusion dominated cases rise early $\lesssim 100$ Myr and cause a significant amount of mass loss from the flux tube. The streaming case is primarily inhibited by the magnetic field strength, not the streaming transport, because in simulation Bdiff the flux tube takes an even longer time to rise.

mass has been lost from the flux tube. In comparison, the buoyant rise in simulation Cdiff happens quickly (see panels of Figure 5.3). The advective simulation Cadv is mainly driven by an explosive launching, instead of buoyancy. However, it and the diffusion case produce similar results in terms of flux tube movement. The difference between those two simulations is more apparent in the beginning, when the launching is different.

Of these simulations, the diffusion case is the most surprising. Initially, one might think cosmic ray diffusion may have a smaller effect on feedback processes than streaming and advection, because there is less time for cosmic rays to impact the ISM (See Equation 5.14) when using the Milky Way value of the diffusion coefficient. However, our simulation Cdiff shows a large cosmic ray injection can generate

enough force through the cosmic ray pressure gradient to move mass out of the disk and bend the magnetic field.

5.6.3 Dependence on Injection Characteristics (Q3 & Q4)

We ran three simulations focused on the injection characteristics. The first, `Uheavy`, placed the injection lower in the galactic disk. While the flux tube rose later than in other simulations, this simulation gave similar results when compared to `U`. Simulation `Ublast` examined how increasing the injection energy to 10^{52} erg would change the dynamics. This large injection caused significant change in the ISM, causing the flux tube to rise rapidly. This simulation suggests that future simulations, with multiple injections, could lead to rapid disruption. In Simulation `Utherm`, we replaced the cosmic ray injection with a thermal injection of similar magnitude. This injection contained no mass, but was an over-pressure ($T \sim 10^5$ K $>$ $T_0 \sim 10^4$ K) region. This thermal energy injection caused rapid mass flow off the flux tube, driving a buoyant rise. However, the simulation is physically unrealistic because we do not include radiative cooling (see Section 5.5.1). The tube averaged mass and height of these simulations are shown as part of Figure 5.8.

5.6.4 Fast Disruption: Cosmic rays vs. Thermal Injection

The simulations which produced the largest changes in the shortest time for an injection energy 10^{51} erg were `Cadv`, `Cdiff`, and `Utherm`. These injections had flux tubes rise faster than a traditional Parker Instability, and caused the flux tube to




Figure 5.8: Plot of maximum tube height (black) and average mass (green) within $x = \pm 250$ pc of the injection for simulations related to Q3 and Q4. Note the vertical axis has been extended compared to Figure 5.3. Solid lines show simulation U, dashed lines show simulation Uheavy, dash-dotted lines show simulation Ublast, and dotted lines show simulation Utherm. Simulation Uheavy does eventually rise from the disk after 200 Myr. Simulation Ublast drives an explosive launching faster than the thermal injection in Utherm.

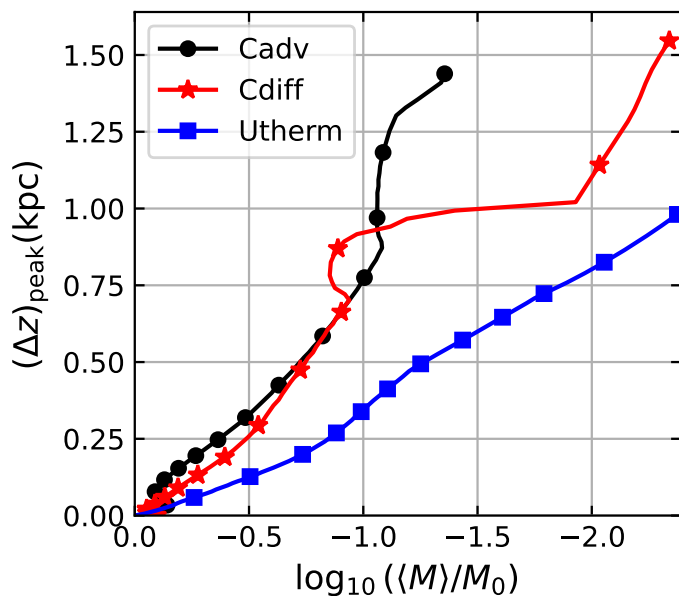


Figure 5.9: The evolution over time of each flux tube's height and mass. Solid lines and filled markers correspond to the tube's peak height. Each simulation is shown until 130 Myr, with markers placed at 10 Myr increments. Black circles show simulation Cadv, red stars show simulation Cdiff, and blue squares show simulation Utherm. While the thermal injection in simulation Utherm initially moves material up faster, at late times the cosmic ray injections in simulations Cadv and Cdiff overtake simulation Utherm. After 130 Myr, the diffusion transported cosmic ray injection has created the most vertically extended flux tube.




Figure 5.10: For the same simulations, we show the average momentum in the vertical direction in the region of $z > 1.5 \text{ kpc}$. Cdiff creates the largest mass flux. The decrease around 140 Myr is when the flux tube leaves the simulation through the upper boundary.

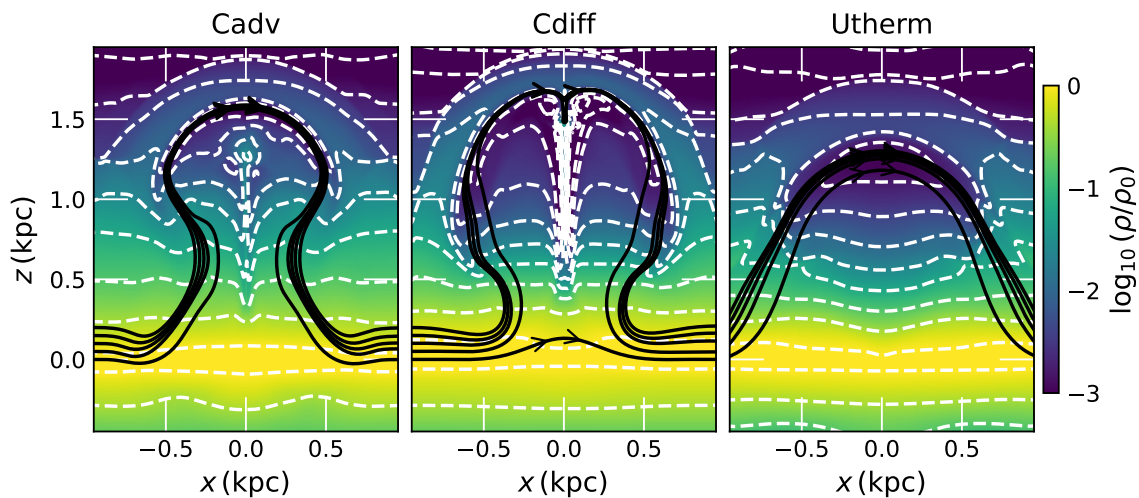


Figure 5.11: Cuts of simulations Cadv, Cdiff, and Utherm. These simulations were the fastest at disrupting the vertical structure, with each image shown at a time of 130 Myr. The black lines are magnetic field lines and the white dashed lines are density contours. These plots are cuts of our 3D simulations at $y = 0$. Simulation Cdiff reaches the highest point, with significantly adjusted magnetic field structure.

lose approximately half its mass in $\lesssim 100$ Myr. In Figure 5.9 we show the evolution of peak tube height against mass. The markers on each line denote 10 Myr steps, whereas each line is made with points at a resolution of 1 Myr. The final marker is at 130 Myr for each simulation. Even though the cosmic ray injection simulations take a longer time to start rising, they rise at a faster rate than the thermal injection. They overtake the thermal injection after approximately 90 Myr, and they have more mass at that time. The density at $y = 0$ of the three simulations at this final time, 130 Myr, is shown in Figure 5.11. Each simulation produces a different magnetic field structure. Considering simulation `Utherm` is an overestimate of thermal injection’s impact (see cooling time argument in Section 5.5.1), we can focus on the difference between `Cadv` and `Cdiff`. The left and middle plots in Figure 5.11 show that cosmic ray transport changes the both the flow of gas around the rising flux tube and the shape of the magnetic field lines. Figure 5.10 shows the average vertical momentum of the cells with $z > 1.5$ kpc for each of the three simulations. Following the red dashed line for simulation `Cdiff`, it is clear the diffusion dominated simulation creates the largest mass flux from the disk. While this gas may fall down given enough time and vertical expansion distance, we are unable to continue to follow that gas because it leaves out of the top of our simulation ($z = 2$ kpc).

5.7 Conclusion

We ran ten simulations of cosmic ray injection into a vertically stratified medium, using the `Athena++` code. These simulations illustrated the effect of cosmic ray

injection in a galactic disk, on intermediate scales (~ 1 kpc) larger than the ISM's fine structure and smaller than the entire galaxy. By exploring an extended parameter space, we produced the highest resolution picture yet of localized cosmic ray injection on these mesoscales. We also showed that cosmic ray transport dictates the impact of cosmic rays on the ISM. Below, we provide the key points and results of this work:

- Cosmic ray diffusion of a cosmic ray injection can change the ISM's vertical structure and a galaxy's magnetic field on timescales < 100 Myr.
- The large magnetic field strength necessary for cosmic ray streaming dominance over diffusion limits the rapid disruption of the flux tube because of increased magnetic tension. Streaming is effective at disrupting the vertical structure, but takes a longer time than in weak magnetic field cases.
- A flux tube disrupted by cosmic ray injection will rise faster at later times than one disrupted by thermal injection, producing a larger mass flow of material out of the galactic disk.

Our simulations provided useful heuristic results by considering cosmic ray injection in a stratified medium. These simulations are clearly limited in their realism because galactic disks are not uniformly stratified. The ISM in galactic disks is multiphase and clumpy. Additionally, by neglecting the dynamics of the stellar gravitational potential, we lack forcing terms which could change the effects of these energy injections. Future work may need to focus on diffusion as the primary cosmic ray transport mechanism, along with implementation of more realistic ISM conditions (multiple gas phases, galactic rotation, etc.). The non-constant α , β of the

injection have a significant impact. This non-uniformity should be extended to the background medium, instead of using the constant α, β assumption originally introduced by Parker (1966). Variable magnetic field strength is particularly important, as it could amplify the importance of cosmic ray streaming in low β_{pl} regions.

In these simulations, we neglected radiative heating and cooling, but those processes are an important consideration in the ISM. The radiative cooling would be particularly important for the gas heated by cosmic ray streaming. Cooling will minimize, and possibly remove, the impact of the thermal injection in simulation `Utherm`. Additionally, having only a single cosmic ray injection by supernovae in our simulation volume over a time of 100 Myr is unrealistic. In future work, we aim to include multiple injections at different times and locations. Separating the injections in space opens up a variety of other situations which make it difficult to isolate the buoyancy the injection creates, hence why we only considered a single injection location in this work.

For large scale (cosmological or galactic) simulations, there needs to be some consideration of cosmic ray injection. Even simulations which include cosmic rays generally do not resolve their injection into the ISM. Our work shows a cosmic ray injection can drive an upward flow in < 100 Myr after their injection, and most of that upward movement actually happens in < 20 Myr. A resolution of ~ 100 pc would be enough to illustrate the impact of an upward flow along the magnetic field (consider Figure 5.11). A sub-grid physics module related to cosmic ray injection could initiate this flow before letting it evolve independently. In the y direction, across the magnetic field and in the plane of the disk, higher resolution would be

necessary. The width of the flux tube barely reaches 100 pc in that direction, and the overall dynamics in that direction are minimal. The main impact the y direction had in our simulations is to allow mass above the flux tube to move out of the path of the rising flux tube. The actual resolution of that dimension is less significant (See Appendix 5.8.2).

Our work also shows cosmic ray injection is an important part of dynamics in a galaxy. Moving, heating, and compressing gas all have an impact on where stars form and on galactic structure. Since the eventual rise of the flux tube happens in a short time (< 20 Myr), the dynamics created by cosmic ray injection are less likely to be washed out by galactic rotation and other dynamical processes. Our results also shed light on the respective roles of cosmic ray compressibility and buoyancy alluded to in Section 5.3.2. Finally, our results suggest that including effects of localized cosmic ray injection in global simulations could be a necessary step in accurately modeling galactic outflows and evolution.

5.8 Appendix

5.8.1 Comparison with Parker Instability Simulations

To test our numerical methods and initial conditions, we ran a 2D simulation matching one in Heintz et al. (2020). Those simulations used the FLASH code (Fryxell et al., 2000), along with the streaming transport method in Sharma et al. (2010). Matching those parameters (originally based on Rodrigues et al. (2016)), we use $\rho_0 = 6.76 \cdot 10^{-25} \text{ g cm}^{-3}$, $P_{g,0} = 8.19 \cdot 10^{-13} \text{ erg cm}^{-3}$, $g_* = 2 \cdot 10^{-9} \text{ cm s}^{-2}$, $\eta = 2$,




Figure 5.12: *Top row:* Vertical profiles of horizontally averaged gas density (solid red line), cosmic ray pressure (dash-dotted blue line), and magnetic pressure (dotted green line). The initial vertical profile for the quantities is shown as well (dashed black line). The initial profile is the same for each quantity because we normalize by the initial midplane value of each quantity, and each physical variable is proportional to the other (see Section 5.3). *Bottom row:* Gas density at corresponding time dump to the vertical profile above each image. The total mass in the simulation is conserved, with a loss of $\sim 0.1\%$ through the diode boundary conditions in the vertical (\hat{z}) directions.

$$H = 250 \text{ pc}, \alpha = 10/3, \beta = 1.25, \text{ and } V_m = 0.01c.$$

We find that that our simulation method is consistent with that of Heintz et al. (2020), and our results are illustrated in Figure 5.12. This figure shows horizontally (\hat{x}) averaged profiles of the gas density, cosmic ray pressure, and magnetic pressure at several time dumps, compared with the original profile for the quantities. Below the averaged profiles, we show the gas density at each selected time dump. The magnetic and cosmic ray pressures decrease more slowly away from the plane as the magnetic

field becomes bent, while gas is compressed towards the midplane. Gravity pulls the gas down along the magnetic flux tubes, which have turned perpendicular to the disk in several locations. Our simulation matches expected behavior (see Figure 11 of Heintz et al. (2020), a similar plot) and evolves on a similar time scale to Tharakkal et al. (2023). These profiles illustrate the transition from a linear growth regime to a nonlinear regime examined in depth by Tharakkal et al. (2023). Additionally, these 2D simulations illustrate the cosmic ray method implemented in `Athena++` by Jiang & Oh (2018) is useful in studying the Parker instability, since it is in reasonable agreement with other numerical simulations.

5.8.2 Boundary Conditions and Dimensionality

In the direction of stratification (\hat{z}), we use vacuum (also known as diode) boundary conditions. These conditions make it impossible for inflow to occur, since the boundary cells are set to the density and pressure floor of the numerical simulation. We also avoid setting up a steep discontinuity at this boundary by extending the simulation several scale heights in the vertical direction. Because of this extension, the cells near the boundary are already almost at the density and pressure floors when the simulation starts. Then, any dynamical activity beyond that initial equilibrium profile will propagate out of the simulation.

In directions in the plane of the disk (\hat{x} and \hat{y}) we use outflow boundary conditions. While these allow inflow from boundary cells after the crest of a wave passes the boundary, there is very little error if the gas is quickly moving out of the simulation at that boundary. Our current problem satisfies this assumption because the flow is




Figure 5.13: Plots matching Figure 5.3, but comparing simulations Bstr and Cdiff to identical simulations with periodic boundary conditions, BstrPer and CdiffPer. For each simulation, the solid black line is maximum height of the flux tube at a given time. The green dotted line is the average mass in the region $x < |250 \text{ pc}|$, as a fraction of the initial average mass. The purple solid, dashed, and dash-dotted lines are the gas, magnetic, and cosmic ray pressure, respectively, averaged in the same region. The second y axis on the right side applies to the lines of mass and pressures, with each quantity $Q \in \{\langle M \rangle, \langle P_g \rangle, \langle P_B \rangle, \langle P_c \rangle\}$ normalized by either $P_0 = 10^{-12} \text{ erg cm}^{-3}$ or the initial value (for mass $\langle M \rangle$).

either static, or large (when the perturbation pushes gas down a magnetic flux tube and toward the boundary). Ideally, we would again use vacuum boundary conditions to stop inflow. However, we cannot use a vacuum in this direction since this would create a significant discontinuity near the midplane of the simulation, where there are gas densities and pressures above the floor values. Therefore, outflow boundary conditions are the best ones for treating a single, spatially and temporally isolated,

injection.

With the goal of minimizing numerical errors, simulations often exploit periodic boundary conditions. In the current work, we wanted to avoid tying down the flux tubes at the boundaries of the simulation in this study, because that can have a stabilizing effect (Zweibel & Bruhwiler, 1992). However, to make sure our simulation results were not significantly different when using periodic boundaries, we ran two simulations, `BstrPer` and `CdiffPer`, with periodic boundary conditions but the same initial conditions as `Bstr` and `Cdiff`, respectively. The flux tube dynamics of those simulations are shown in Figure 5.13. Simulations `Cdiff` and `CdiffPer` show very similar results. For simulations `Bstr` and `BstrPer`, the difference is larger: the periodic boundary conditions tie down the magnetic flux tube at the boundaries, making it more difficult for the tube to rise after 100 Myr. Overall, the trends are the same, but the timing is delayed when using periodic boundaries.

We also examined the effect of dimensionality on our simulations. Initial 2D simulations differed from our results in 3D simulations significantly, with tubes taking much longer to rise in 2D simulations. This delay is caused by flux tubes above our injection being unable to move out of the way of the rising flux tube where the injection took place. This effect is also apparent in Figure 5.14, where we examine different resolutions and numerical sizes in the third dimension, \hat{y} . For 8 cells in the \hat{y} direction and resolution $\Delta y = 62.5$ pc, the tube slowly rises before bursting upward. This trend is similar to our 2D simulations, whose initial rise takes longer. In these simulations, the rise only happens quickly because we used a large injection energy $E_{\text{inj}} \approx 10^{51}$ erg similar to simulation `Ublast`. For 16 cells and a resolution




Figure 5.14: Height of flux tube in simulations with large injection energy $E_{\text{inj}} \approx 10^{51}$ erg, similar to simulation `Ublast`. These simulations were run to examine the impact of the third dimension, \hat{y} , on our results. The red solid line shows a simulation with only 8 numerical cells in the \hat{y} direction, the blue dashed line used 16 numerical cells, the green X markers are from a simulation with 32 numerical cells, and the purple plus markers are from a simulation with 64 numerical cells. The simulations with more room in the \hat{y} direction are well converged at a length of ± 0.5 kpc. We chose to use 50 numerical cells in our production runs, with a resolution of $\Delta y = 20$ pc.

$\Delta y = 31.25$ pc, we see a smooth well-behaved flux tube rise. By increasing the width of the box to 32 and 64 numerical cells with the same resolution as the 16 cell run, we get clear convergence of results. From these simulations, we chose a width of ± 0.5 kpc and resolution of $\Delta y = 20$ pc, resulting in 50 numerical cells.

5.8.3 Simulation Convergence

To mitigate the significance of purely numerical parameters we performed several convergence tests. We used 2D simulations to narrow our choice of time integrator, Courant-Friedrich-Lewy (CFL) number, resolution, and the reduced speed of light constant V_m (see Equation 5.20). Using 2D simulations allowed us to save computational resources while exploring the convergence of these simulations. We illustrated

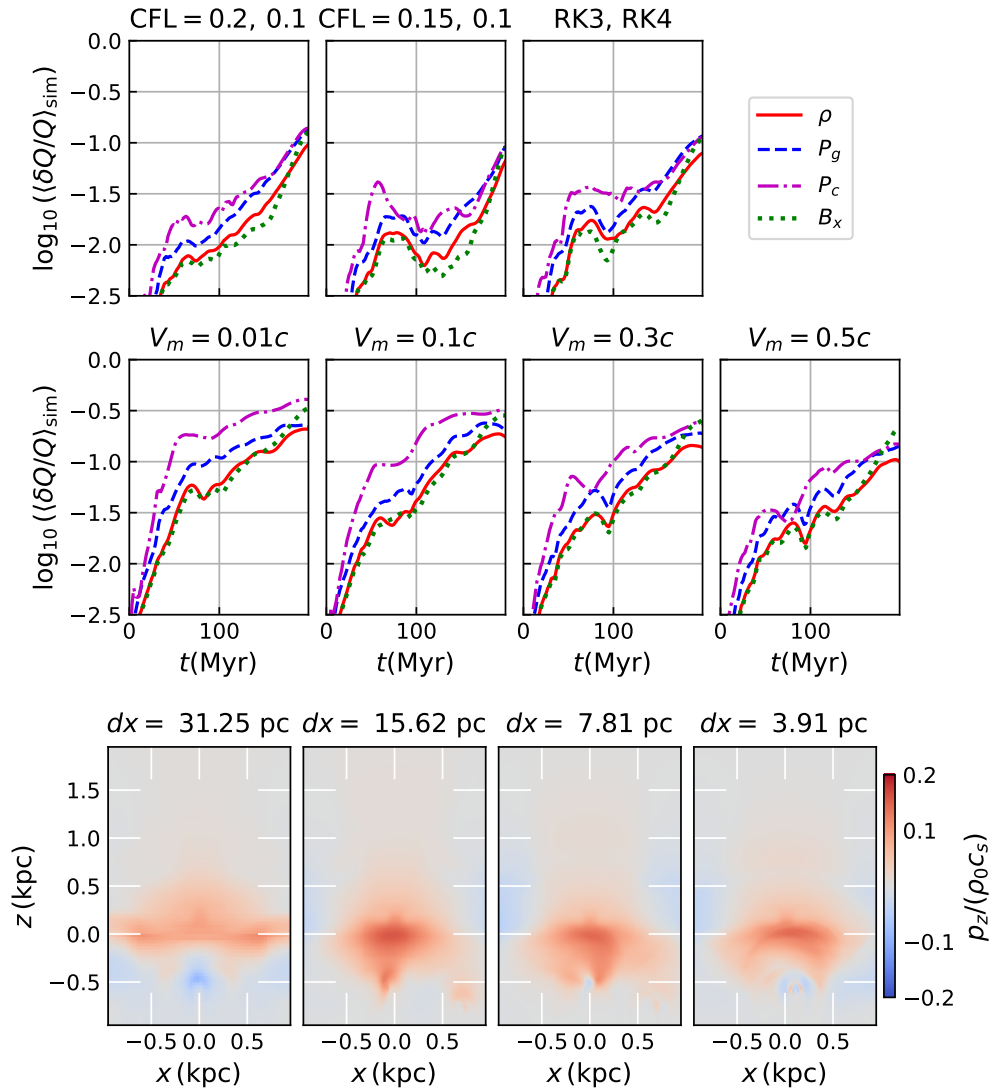


Figure 5.15: Convergence of key numerical parameters. *Top row, first plot:* Average cell-by-cell numerical error between CFL = 0.2 and CFL = 0.1. *Top row, second plot:* Average cell-by-cell numerical error between CFL = 0.15 and CFL = 0.1. *Top row, third plot:* Average cell-by-cell numerical error between a run with a third order Runge-Kutta time integrator and a run with a fourth order Runge-Kutta time integrator. *Middle row:* Average cell-by-cell numerical error between runs with varying V_m . Each has error measured with respect to a $V_m = c$ simulation. *Bottom row:* Plot of vertical momentum p_z in simulations with different resolutions $dx = dz$.

in Section 5.8.2 that 2D simulations would be less accurate because they limit the movement of magnetic flux tubes, but most of the fast, dynamic flow is still in the xz -plane. For numerical parameters, the motion in that plane is where we need be concerned. All the convergence tests here used a weak injection of 10^{50} erg with other parameters equal to those of simulation U. With the weaker injection and 2D restricted motion, the flux tube begins to rise buoyantly by 200 Myr.

The time integrator choice of 3rd order Runge-Kutta and CFL number (0.2) were well converged, as the average cell-by-cell error is under 10% in the top row of Figure 5.15. The error associated with the modified speed of light is larger: in the middle row, the cell-by-cell error for various values of V_m are shown, each compared to a simulation with $V_m = c$. From these, we pick $V_m = 0.3c$ because there is not a huge increase in accuracy by going to higher values of V_m . Increasing V_m beyond $0.3c$ would decrease the timestep below 0.1 kyr, increasing computational resource requirements without an equivalent increase in accuracy. Finally, in the bottom row of Figure 5.15, we show the vertical momentum p_z for simulations with varying resolution. Instead of comparing these simulations via interpolation of the high-resolution simulations, we focus on how similar the dynamics are between each simulation. Each simulation had an aspect ratio of 1, with resolution $\Delta x = \Delta z$. Any higher resolution than 15 pc appears to reproduce a similar vertical momentum on the flux tube. Also, the structures are similar between resolutions 15.62 pc, 7.81 pc, and 3.91 pc. From these, we decided to run our performance simulations with a resolution $\Delta x = \Delta z = 10$ pc.

6 MANY COSMIC-RAY INJECTIONS

In science, moreover, the work of the individual is so bound up with that of his scientific predecessors and contemporaries that it appears almost as an impersonal product of his generation.

— *The Scientific Monthly* (1921), paraphrasing Albert Einstein

This chapter is a modified copy of the paper Habegger & Zweibel (2025), titled *Cosmic-Ray Feedback from Supernovae in a Parker-Unstable Medium*, published in *The Astrophysical Journal* in 2025. No text has been changed; the figures and equations have just been reformatted to accommodate the submission of this thesis document.

I deeply appreciate the contributions of my co-author and advisor, Dr. Ellen Zweibel, to this work.

6.1 Abstract

Supernova energy drives interstellar medium (ISM) turbulence and can help launch galactic winds. What difference does it make if 10% of the energy is initially deposited into cosmic rays? To help answer this question and study cosmic-ray feedback, we perform galactic patch simulations of a stratified ISM in a low star formation rate, high magnetic field regime. We compare two magnetohydrodynamic and cosmic ray (MHD+CR) simulations, which are identical except for how each supernova's energy is injected. In one, 10% of the energy is injected as cosmic-ray energy. In the

other case, energy injection is strictly thermal and kinetic. We find that cosmic-ray injections drive a faster, hotter, and more massive outflow long after the injections occur. Both simulations show the formation of cold clouds (with a total mass fraction $> 50\%$) through the Parker instability and thermal instability. The Parker instability simultaneously produces high mass loading factors $\eta > 10^3$ as it does not directly require star formation and supernovae. We also show how the Parker instability naturally leads to a decorrelation of cosmic-ray pressure and gas density. This decorrelation leads to a significant decrease in the calorimetric fraction for injected cosmic rays, but it depends on having a highly resolved magnetic field.

6.2 Introduction

A tenth, 10%, or 0.1 of anything is a small fraction. Neglecting such a small component is a reasonable assumption in many situations. An example of this small component is cosmic rays in the interstellar medium (ISM): 10% of each supernova's energy is all that is needed to explain the amount of these high-energy particles in our galaxy's ISM (Baade & Zwicky, 1934b; Blandford & Ostriker, 1978). More recent work has illustrated that the fraction produced in a shock could be as high as 20% (Caprioli & Spitkovsky, 2014). That fraction could still be considered negligible, especially on the long time scales ($t \gtrsim 100$ Myr) and long length scales ($L \gtrsim 1$ kpc) of galaxy evolution, since supernovae drive turbulence and heat gas at small scales ($L \sim 10$ pc).

However, other work has shown that cosmic rays are not negligible on large

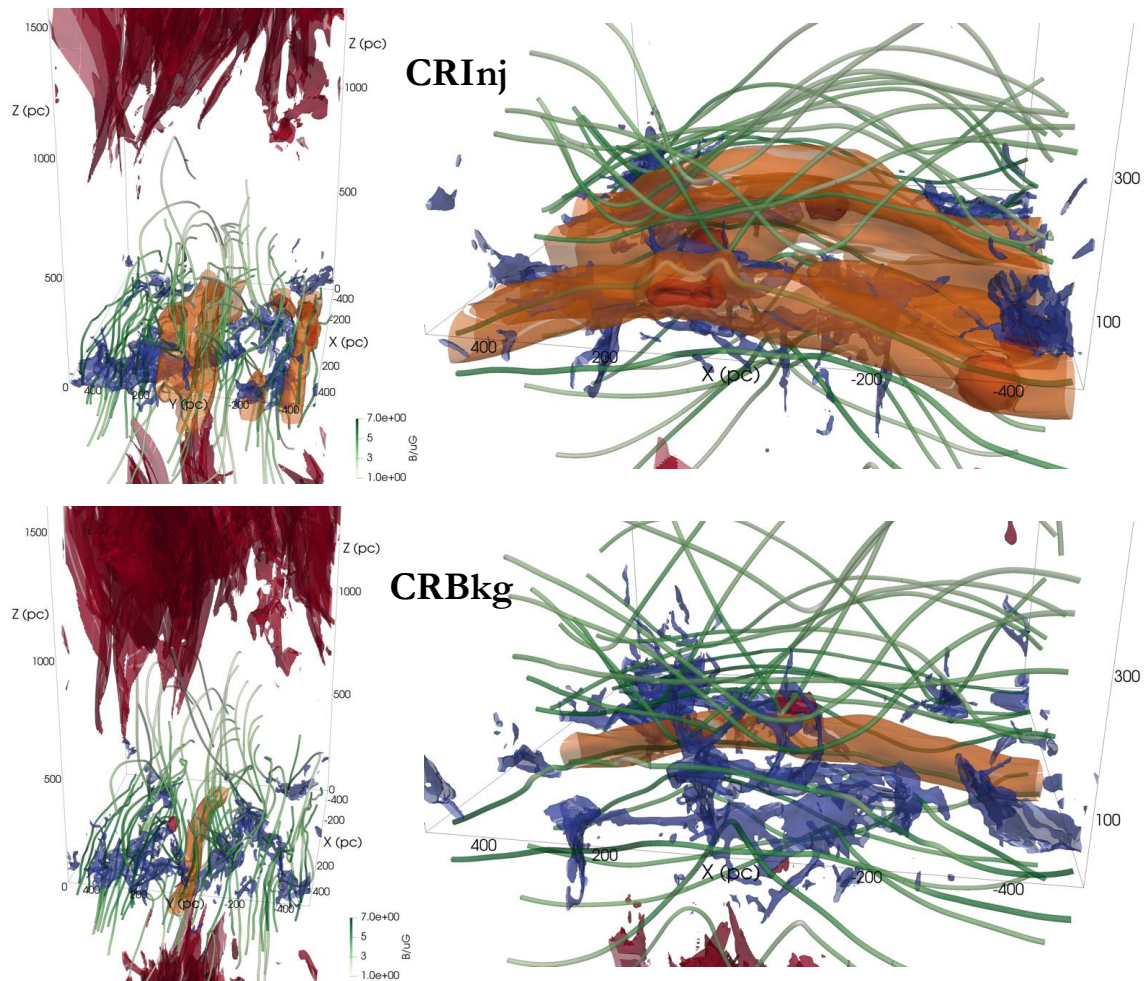


Figure 6.1: 3D snapshots of each simulation at $t = 200$ Myr (see Table 6.1). The top row shows volume renders from our CRInj simulation, and the bottom row shows renders from our CRBkg simulation. The left side figures show a larger scale picture illustrating the hot wind and extended magnetic field, whereas the right side figures show a zoom-in of the midplane region. The blue structures are contours around cold gas ($T \leq 500$ K), and the red contours are around hot gas ($T \geq 5 \times 10^4$ K). Green lines trace the magnetic field, and darker green corresponds to stronger magnetic field (see colorbar in left column figures). The orange contours in the $z = 0$ surround regions of high cosmic-ray pressure. The red blobs in the midplane show expanding supernova. In the CRInj simulation, these supernovae are encased by orange contours showing where they dumped cosmic ray energy along the local magnetic field.

scales (see recent reviews Zweibel (2017); Ruszkowski & Pfrommer (2023); Owen et al. (2023)). They are in approximate pressure/energy equipartition with the thermal gas and magnetic field (Ferrière, 2001). Gradients in cosmic-ray pressure can drive outflows or fountain flows from a galactic disk, removing gas that could otherwise form stars (Chan et al., 2022; Tsung et al., 2023; Armillotta et al., 2024; Thomas et al., 2025). Bottlenecks created by cosmic-ray streaming can destroy and disrupt cold gas clouds (Wiener et al., 2017a, 2019; Bustard & Zweibel, 2021; Tsung et al., 2022). Cosmic-ray transport is faster parallel to the magnetic field, and this anisotropic transport can impact the turbulent energy cascade of the thermal gas (Habegger et al., 2024).

These feedback mechanisms are stifled if cosmic rays are transported out of the galaxy too quickly. In order to keep global galaxy simulations (and cosmological zoom-in simulations) in agreement with observed γ -ray luminosities, a cosmic-ray diffusion coefficient of about one order of magnitude higher than the standard value of $3 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ derived from the B/C ratio (Jones et al., 2001; Evoli et al., 2020) is often invoked, following the simulations of Chan et al. (2019). This diffusion coefficient amplification means that the cosmic rays escape the galaxy before they ever affect the dynamics of the thermal gas. This result brings us back to the idea that cosmic rays produce only a marginal effect; including them just makes stellar feedback slightly more efficient.

In this work, we examine how individual cosmic-ray sources eventually adjust the dynamics and vertical structure of the multiphase ISM, becoming a significant component of stellar feedback. We previously examined how a single localized cosmic-ray

energy injection on a magnetic flux tube could disrupt a stratified ISM (Habegger et al., 2023). This work builds on that study to consider multiple injections throughout the galactic midplane. Our simulations evolve a slab of stratified ISM, with midplane parameters that match the solar neighborhood. After a significant fountain flow develops in the first 200 Myr, with magnetic field geometry similar to the Parker instability (Parker, 1966), the slab reaches a steady state (lasting $\gtrsim 100$ Myr). During the transition to the steady state, the cosmic rays are able to escape quickly without interacting with the thermal gas. This escape significantly decreases the γ -ray luminosity.

Figure 6.1 is a set of volume renders from our simulations at a time of $t = 200$ Myr, which illustrate the complex temperature and magnetic field structures that form over time in the stratified, multiphase, and magnetized ISM. The temperature varies over several orders of magnitude (blue contours surround gas with $T < 500$ K and red contours surround gas with $T > 5 \cdot 10^4$ K). Regions of high cosmic-ray pressure (orange contours) follow the magnetic field (in green), inflating magnetic flux tubes. Some magnetic field lines stretch perpendicular to the disk, over 1 kpc in length, despite initially only being directed horizontally. Considering the full scale of a galaxy, those magnetic field lines would connect the galactic plane and the galactic halo. We find that these field lines act as escape highways for injected cosmic rays, taking the simulation below the calorimetric limit. These extended field lines result from expanding plumes that push the horizontal magnetic field out of the galactic plane. Those plumes drive the creation of cold gas in valleys, the same structure predicted for the Parker instability (Parker, 1966). These valleys result in

a decorrelation of cosmic-ray pressure and gas density. The decorrelation is visible in Figure 6.1, where the cold gas (blue contours) and high cosmic-ray pressure regions (orange contours) alternate in the direction perpendicular to the mean magnetic field. While the Parker instability drives this decorrelation in our simulations, the general process of accreting cold gas from the galactic halo should have the same effect as long as the magnetic field is well-resolved.

In Section 6.3, we detail our simulation setup, numerical methods, and implementation of physical processes. In Section 6.4, we cover the primary results from our simulations. In Section 6.5, we discuss several implications derived from our results and analysis. Finally, in Section 6.6, we provide a short summary and list of key conclusions.

6.3 Methods

We use the astrophysical magnetohydrodynamic (MHD) simulation code `Athena++` to evolve a thermal fluid alongside a cosmic-ray fluid (Stone et al., 2020; Jiang & Oh, 2018). The code solves the following system of equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad (6.1)$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) = 0 \quad (6.2)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot \left(\rho \mathbf{u} \mathbf{u} + \mathbb{1} \left(P_g + \frac{B^2}{2} \right) - \mathbf{B} \mathbf{B} \right) = \rho \mathbf{g} + \sigma_c \cdot \left(\mathbf{F}_c - \frac{4}{3} E_c \mathbf{u} \right) \quad (6.3)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left(\left(E + P_g + \frac{B^2}{2} \right) \mathbf{u} - (\mathbf{B} \cdot \mathbf{u}) \mathbf{B} \right) = \rho \mathbf{g} \cdot \mathbf{u} + \mathcal{L}(n, T) - \frac{1}{3} (\mathbf{u} + \mathbf{v}_s) \cdot \nabla E_c \quad (6.4)$$

$$\frac{\partial E_c}{\partial t} + \nabla \cdot \mathbf{F}_c = \frac{1}{3} (\mathbf{u} + \mathbf{v}_s) \cdot \nabla E_c \quad (6.5)$$

$$\frac{1}{V_m^2} \frac{\partial \mathbf{F}_c}{\partial t} + \frac{1}{3} \nabla E_c = -\sigma_c \cdot \left(\mathbf{F}_c - \frac{4}{3} E_c \mathbf{u} \right). \quad (6.6)$$

These equations describe the evolution of a magnetized two-fluid system. The first fluid, thermal gas, has a mass density ρ , velocity \mathbf{u} , pressure P_g . The second fluid, made up of cosmic rays (i.e. a relativistic, non-thermal fluid), is described by a total energy density E_c and flux \mathbf{F}_c . The cosmic-ray transport coefficient σ_c is a diagonal matrix which governs how cosmic-ray energy moves in the simulation, \mathbf{v}_s is the streaming velocity, and the modified speed of light V_m sets the maximum transport speed. We evolve the magnetic field \mathbf{B} from Faraday's Law in the ideal MHD approximation (Eqn. 6.2). The remaining variables are the total magnetohydrodynamic energy density of the first fluid $E = P_g/(\gamma - 1) + \rho u^2/2 + B^2/2$, the gravitational field \mathbf{g} , and the radiative heating and cooling function $\mathcal{L}(n, T)$, which depends on the number density ($n = \rho/m$, where m is particle mass) and temperature ($T = P_g m / (\rho k_B)$) of the thermal gas. We assume hydrogen gas with a constant molecular weight of $\mu = 1$, so the particle mass is $m = \mu m_H = m_H$.

For the cosmic-ray transport, we allow the cosmic rays to diffuse along the local magnetic field at the canonical Milky Way value of $\kappa_{\parallel} = 3 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ (Jones et al. (2001), for more recent determinations see Evoli et al. (2019, 2020)). This anisotropic transport is created by splitting the transport coefficient into components perpendicular and parallel to the local magnetic field direction \hat{b} in each cell (see Jiang

& Oh (2018)):

$$\sigma_{c,\text{diff}} = \frac{1}{\kappa_{\perp}} \mathbb{1} + \left(\frac{1}{\kappa_{\parallel}} - \frac{1}{\kappa_{\perp}} \right) \hat{b}\hat{b}. \quad (6.7)$$

For the perpendicular direction, we set the transport rate at $\kappa_{\perp} = 3 \times 10^{18} \text{ cm}^2\text{s}^{-1}$ so there is no diffusion perpendicular to the field over the course of the simulation (Desiati & Zweibel, 2014). See Habegger et al. (2024) for an extended discussion on choosing the cosmic-ray diffusion coefficient in this anisotropic, resolved magnetic field treatment.

We also include the effects of cosmic-ray streaming transport. Streaming at the local Alfvén speed is dominant when cosmic rays generate the same Alfvén waves that they gyro-resonantly scatter off (Zweibel, 2017). This process leads to additional heating because the Alfvén waves dissipate after scattering the cosmic rays, taking energy from the cosmic rays into the thermal gas. The heating of gas (and decrease in cosmic ray energy) are proportional to the streaming speed \mathbf{v}_s and appear in the energy equations (MHD energy in Equation 6.4, cosmic-ray energy in Equation 6.5). The additional transport appears in the transport matrix by including the following term:

$$\sigma_{c,\text{str}} = -\hat{b}\hat{b} \frac{(\gamma_c - 1)\hat{b} \cdot \nabla E_c}{\gamma_c E_c \mathbf{v}_s \cdot \hat{b}} = -\hat{b}\hat{b} \frac{\hat{b} \cdot \nabla E_c}{4E_c \mathbf{v}_s \cdot \hat{b}}. \quad (6.8)$$

Additionally, it is important to note that the streaming velocity is the Alfvén velocity directed down the cosmic ray pressure gradient (McKenzie & Voelk, 1982; Breitschwerdt et al., 1991; Zweibel, 2017; Jiang & Oh, 2018):

$$\mathbf{v}_s = -\frac{\mathbf{B}}{\sqrt{4\pi\rho}} \frac{\mathbf{B} \cdot \nabla P_c}{|\mathbf{B} \cdot \nabla P_c|}. \quad (6.9)$$

Note that we use the full gas density ρ in Equation 6.9, not the ion density. In reality, the ion density should appear in the definition of the Alfvén speed for plasma with an ionization fraction less than one. This simplification only impacts the cold gas which forms in our simulations.

In cold gas, the ionization fraction decreases so the gas density and ion density disagree. Additionally, Alfvén waves are damped making all gyroresonant scattering of cosmic rays weaker. Other works have successfully implemented and studied the effect of temperature dependent cosmic ray transport (Farber et al., 2018; Bustard & Zweibel, 2021; Armillotta et al., 2024; Habegger et al., 2024).

However, we neglect this effect in these simulations because our chosen diffusion rate is high enough for the cold gas. Diffusion across an 100pc cloud occurs at a speed $\kappa_{\parallel}/L = 10^3 \text{kms}^{-1}$. For smaller clouds, the rate of diffusion is even faster. Increasing the rate further by including more complete physics for cosmic ray transport would have a minimal impact on our results because diffusion already rapidly smooths cosmic ray pressure gradients in cold gas.

The total transport parallel to the magnetic field is the addition of the reciprocals of the transfer terms:

$$\frac{1}{\hat{b} \cdot \sigma_c \cdot \hat{b}} = \frac{1}{\hat{b} \cdot \sigma_{c,\text{diff}} \cdot \hat{b}} + \frac{1}{\hat{b} \cdot \sigma_{c,\text{str}} \cdot \hat{b}} = \kappa_{\parallel} - \frac{4E_c \mathbf{v}_s \cdot \hat{b}}{\hat{b} \cdot \nabla E_c}. \quad (6.10)$$

This combination gives a final, total transport matrix of

$$\sigma_c = \frac{1}{\kappa_{\perp}} \mathbb{1} + \left(\frac{1}{\kappa_{\parallel} - \frac{4E_c \mathbf{v}_s \cdot \hat{b}}{\hat{b} \cdot \nabla E_c}} - \frac{1}{\kappa_{\perp}} \right) \hat{b} \hat{b}. \quad (6.11)$$

Parameter	CRInj	CRBkg
T_0/K		10^4
$\rho_0/(m_p \text{ cm}^{-3})$		1
β		1
β_{cr}		1
$B_0/(\mu\text{G})$		5.89
$P_c/(\text{eV cm}^{-3})$		0.862
$\Sigma_*/(M_\odot \text{ pc}^{-2})$		50
H_*/pc		100
$\kappa_{\parallel}/(\text{cm}^2 \text{ s}^{-1})$		$3 \cdot 10^{28}$
$N_{\text{SN}}/\text{Myr}^{-1}$		1
$E_{\text{inj,th}}/\text{erg}$	$0.9 \cdot 10^{51}$	$0.1 \cdot 10^{51}$
$E_{\text{inj,cr}}/\text{erg}$	$1.0 \cdot 10^{51}$	0

Table 6.1: Simulation parameters: initial values of temperature of gas T_0 , midplane gas density ρ_0 , plasma beta β , cosmic-ray beta β_{cr} , stellar surface density Σ_* (which sets gravitational acceleration), stellar scale height H_* , parallel cosmic-ray diffusion coefficient κ_{\parallel} , supernova injection rate N_{SN} , thermal energy per injection $E_{\text{inj,th}}$, and cosmic-ray energy per injection $E_{\text{inj,cr}}$

This streaming and diffusion implementation is standard, and follows directly from Jiang & Oh (2018). We include it here in an effort to make Equations 6.5 & 6.6 clear to readers unfamiliar with Jiang & Oh (2018).

In the simulations used for this work, we set $V_m = 0.1c$. See Appendix C of Habegger et al. (2023) for an extended V_m convergence study. Our largest flow velocities are on the order of 100 km s^{-1} , much lower than our chosen V_m .

6.3.1 Setup & Initial Conditions

We present two simulations, with identical initial conditions. The only difference is how supernova energy is injected. In one simulation, named **CRInj**, each supernova has 10% of its energy injected as cosmic-ray energy instead of thermal energy. In the second simulation, named **CRBkg**, no cosmic-ray energy is injected with each

supernova. Instead, the only cosmic rays are from the initial background we put in to match the observed estimates that cosmic ray pressure is similar in value to the thermal gas pressure (i.e. $P_c \sim P_g$.) The important parameters for each simulation are shown in Table 6.1.

The simulations are on a Cartesian grid, with the \hat{x} and \hat{y} directions spanning the midplane of the slab and the \hat{z} direction extending perpendicular to the midplane. The midplane has an area $1 \text{ kpc} \times 1 \text{ kpc}$, but the simulation extends $\pm 2.4 \text{ kpc}$, above and below the midplane, for a total simulation volume of $1 \text{ kpc}^2 \times 4.8 \text{ kpc}$. In the x and y directions, the boundary conditions are periodic. In the z direction, we use a vacuum or ‘diode’ boundary condition which only allows outflow and no inflow.

Each direction has a resolution of $\Delta x = \Delta y = \Delta z = 5 \text{ pc}$, resulting in 200 resolution elements in the \hat{x} and \hat{y} directions, but 960 in the \hat{z} direction. This resolution, combined with momentum injection (see section 6.3.3), is necessary for an accurate model of supernova feedback in a multiphase ISM (Kim & Ostriker, 2015). In particular, both are necessary for resolving the Sedov-Taylor phase and the transition to the snowplow phase (Draine, 2011).

The simulations are initially in hydrostatic equilibrium, and the setup is identical to the one in Habegger et al. (2023). We use a gravitational acceleration profile

$$\mathbf{g}(z) = -2\pi G \Sigma_* \hat{z} \tanh\left(\frac{z}{H_*}\right), \quad (6.12)$$

which assumes a stellar mass density Σ_* with a $\text{sech}^2(z/H_*)$ vertical distribution. The hyperbolic tangent profile is particularly useful in numerical studies because it is smooth at $z = 0$ (Giz & Shu, 1993).

Taking the cosmic-ray pressure P_c and magnetic pressure $P_B = B^2/2$ to be proportional to the gas pressure P_g , and then assuming the equilibrium is isothermal with $P_g \propto \rho$, the initial condition is

$$\frac{\rho(z)}{\rho_0} = \frac{P_g(z)}{P_{g,0}} = f(z) = \operatorname{sech}^\eta \left(\frac{z}{\eta H} \right), \quad (6.13)$$

where $\eta = H_*/H$ is the ratio of stellar scale height H_* to gas scale height

$$H = \frac{k_B T_0}{2\pi G \Sigma_* m_p} \left(1 + \frac{1}{\beta} + \frac{1}{\beta_{\text{cr}}} \right). \quad (6.14)$$

Using the function $f(z)$ in Equation 6.13, we can define the cosmic-ray pressure with a cosmic-ray beta $\beta_{\text{cr}} = P_{g,0}/P_{c,0}$ and the magnetic field with a plasma beta $\beta = 2P_{g,0}/B_0^2$:

$$P_c = \frac{P_{g,0}}{\beta_{\text{cr}}} f(z) \quad (6.15)$$

$$\mathbf{B} = \hat{x} \sqrt{\frac{2P_{g,0}}{\beta} f(z)}. \quad (6.16)$$

All velocities are initially set to zero, and the \hat{z} and \hat{y} components of the magnetic field are set to zero everywhere.

6.3.2 Heating and Radiative Cooling

We include heating along with optically thin radiative cooling of the ISM in our simulations. We use a source term defined by the net heat-gain function (negative




Figure 6.2: *Top:* Our combined radiative heating and cooling rate \mathcal{L} in Equation 6.17 gives equilibrium states shown as a black line in this gas pressure-gas density phase diagram. We plot our temperature floor as a dashed line in the bottom right of the plot. The contours show the net heating rate in red contours and net cooling rates in blue contours, with darker colors being faster rates. *Bottom:* We use the radiative cooling rate Λ , as a function of gas temperature T , given by the black line. This rate stitches together two functions, one from Inoue et al. (2006) (light green curve) as well as a polynomial fit to the Collisional Ionization Equilibrium (CIE) cooling rate from Rosen & Bregman (1995). We transition between these curves at their intersection of $T = 1.40413 \cdot 10^4\text{ K}$.

of the heat-loss function):

$$\mathcal{L}(n, T) = n\Gamma - n^2\Lambda(T). \quad (6.17)$$

We adopt the cooling function (in units of $\text{erg cm}^3 \text{s}^{-1}$, with temperature T in units of K)

$$\Lambda(T) = \begin{cases} \Lambda_{\text{IK06}}(T) & T < 14041.3 \\ 4.624 \cdot 10^{-36} T^{2.867} & 14041.3 \leq T < 10^5 \\ 1.78 \cdot 10^{-18} T^{-0.65} & 10^5 \leq T < 4 \cdot 10^7 \\ 3.2217 \cdot 10^{-27} T^{0.5} & 4 \cdot 10^7 \leq T \end{cases} \quad (6.18)$$

which combines the cooling function from Rosen & Bregman (1995) along with the cooling function from Inoue et al. (2006), which has the functional form

$$\Lambda_{\text{IK06}}(T) = 7.3 \times 10^{21} \exp\left[\frac{-118400}{T + 1500}\right] + 7.9 \times 10^{27} \exp\left[\frac{-92}{T}\right]. \quad (6.19)$$

These two exponential dependencies are designed to match the expected cooling rates due to $\text{Ly}\alpha$ and CII emission (Inoue et al., 2006).

The cooling functions from Rosen & Bregman (1995) and Inoue et al. (2006) are plotted in the bottom graph of Figure 6.2. We stitch them together at $T = 1.40413 \cdot 10^4 \text{ K}$, where they intersect. The resulting curve, our cooling function, is shown as a black line in the bottom figure of Figure 6.2.

We use the heating rate from Inoue et al. (2006), a constant $\Gamma = 2 \cdot 10^{-26} \text{ erg s}^{-1}$

from photoelectric heating by dust grains (Inoue et al., 2006; Draine, 2011). With all the parameters of Equation 6.17 set, we have an equilibrium curve defined by $\Gamma = n\Lambda(T)$. We show this curve in pressure-density phase space as a solid black line in the top plot of Figure 6.2. As mentioned in Section 6.3.1, we start with an isothermal atmosphere. The number density ranges from $n = 1$ to approximately $n \approx 10^{-4}$. This isothermal atmosphere quickly ($\lesssim 1$ Myr) relaxes to the equilibrium curve in that density range before the energy injections significantly adjust the structure.

In Figure 6.2, we also plot contours for the net isochoric cooling and heating time:

$$\tau = \frac{k_B T}{\gamma - 1} \frac{1}{\Gamma - n\Lambda(T)}. \quad (6.20)$$

In Figure 6.2, the red areas are where $\tau > 0$ in Equation 6.20, giving net heating. The blue areas are where $\tau < 0$, resulting in net cooling. We plot the isochoric time because we use an operator split approach, including the heating and radiative cooling as an isochoric process. Adiabatic compression and expansion are already included in the energy evolution (see Equation 6.4). This treatment is common, e.g. see Townsend (2009). Unlike the exact method in Townsend (2009), we use an explicit integration for the heating and radiative cooling, with order set by the MHD time integrator. Explicit integration is reasonable for our simulations because our cooling time is longer than the simulation's time step ($\Delta t = \text{CFL} \times \Delta x / V_m \sim 32.6$ yr), which is set by the modified speed of light V_m .

Examining Figure 6.2, we see at higher densities $n \gtrsim 10^{-2} \text{ cm}^{-3}$ the gas cools or heats on short timescales, until it again reaches the equilibrium curve where heating and cooling balance. We also plot our floor temperature of $T = 30$ K as a dashed

line in the top plot of Figure 6.2. This floor is necessary to avoid numerical errors, as well as to stop the formation of molecular gas which we do not accurately model. Whenever the gas reaches the temperature floor, we turn off the radiative heating and cooling terms. If the gas is compressed or expanded by gas motions and brought above the temperature floor, then the terms are again included.

6.3.3 Energy Injections

We inject energy at random locations in the midplane. We restrict the midplane to be cells with vertical coordinate between $z = \pm H_* = \pm 100$ pc. We select a random x, y, z position following that restriction. This random selection neglects the fact that star formation occurs in dense gas (Schmidt, 1959). However, it is a reasonable choice; others have found that injecting supernovae primarily in dense gas can lead to peak-driving which reduces the impact they have when they occur in low-density gas, and thus fails to create a realistic ISM (Walch et al., 2015).

The number of injections in each timestep Δt is taken from a Poisson distribution, with parameter $\lambda = \dot{N}_{\text{SN}}\Delta t$ (see Table 6.1). After calculating the number of injections, we generate their locations using the aforementioned method. This randomization of number of injections makes our simulations stochastic in time, despite having a fixed injection rate \dot{N}_{SN} . We keep \dot{N}_{SN} constant so we can isolate the impact of including cosmic-ray injections. For a more realistic and complete stellar feedback loop, we would need to tie the injection rate to the amount of cold gas in the simulation (assuming the amount of cold gas is a tracer of star formation rate).

It is important to note that we are considering a low rate of supernovae, $\dot{N}_{\text{SN}} =$

1Myr^{-1} in comparison to other tallbox simulations. This rate approximately corresponds to a star formation rate of $100M_{\odot}\text{Myr}^{-1}\text{kpc}^{-2} = 10^{-4}M_{\odot}\text{yr}^{-1}\text{kpc}^{-2}$. This rate is reasonable in the $\sim \text{kpc}^2$ area near the solar neighborhood, between spiral arms (Elia et al., 2022; McCallum et al., 2025).

Once the number of injections and their locations are determined, we inject thermal energy equally in a sphere of radius 20 pc. We also inject kinetic energy (with velocity directed radially outward from the injection location) following the prescription given by Equation 34 in Kim & Ostriker (2015). Finally, we inject cosmic ray energy in a sphere of radius 50 pc centered on the same location. Both simulations we present inject the same kinetic energy with each supernova, but we vary the thermal energy and cosmic ray energy injected. In `CRBkg`, we only inject thermal energy, and in `CRInj` we inject some cosmic ray energy (see Table 6.1).

6.3.4 A Note on Parker Instability

Our initial conditions are overall similar to setups used to examine the Parker instability (Parker, 1966). We tested multiple setups in the single injection simulations and found disruption by the injection only occurs for setups unstable to the Parker instability (Habegger et al., 2023). Our initial setup in this work (see Table 6.1) is linearly unstable to the Parker instability. Since we begin with an isothermal atmosphere, the linear instability criterion (only considering a plane-parallel atmosphere) is (Newcomb, 1961; Zweibel, 2017; Heintz & Zweibel, 2018):

$$1 + \frac{1}{\beta} + \frac{1}{\beta_{\text{cr}}} > \gamma_g + \gamma_c \frac{1}{\beta_{\text{cr}}} \quad (6.21)$$

and plugging in the values from Table 6.1, the instability criterion becomes $\gamma_c < \frac{4}{3}$. Since we include non-advective transport like streaming and diffusion, we expect to satisfy this criteria and the setup is then unstable to the Parker instability.

We can also estimate the growth rate of the most unstable mode. Following the dispersion relations derived in Heintz & Zweibel (2018), we use the 3D Modified Parker (non-zero γ_c) with streaming transport case. For our setup, we find the most unstable Fourier mode has a wavelength along the magnetic field of $L_x \sim 0.81$ kpc, and a growth rate of $\tau \sim 32$ Myr. Additionally, all the the wavelengths above $L_x \geq 0.32$ kpc are unstable. Those large scale wavelengths are all above the characteristic size of a fully evolved supernova supernova remnant, which generally expand to be 100 pc.

6.4 Results

We now describe the results of our two simulations, named CRInj and CRBkg. A key similarity between the simulations is the development of two stages in time. We describe this time evolution before a detailed examination of each stage. The first stage is related to the exponential growth of the Parker instability, while the second stage is a steady state which develops after the Parker instability has disrupted the initial hydrostatic equilibrium. Finally, we compare our simulations to observational diagnostics.

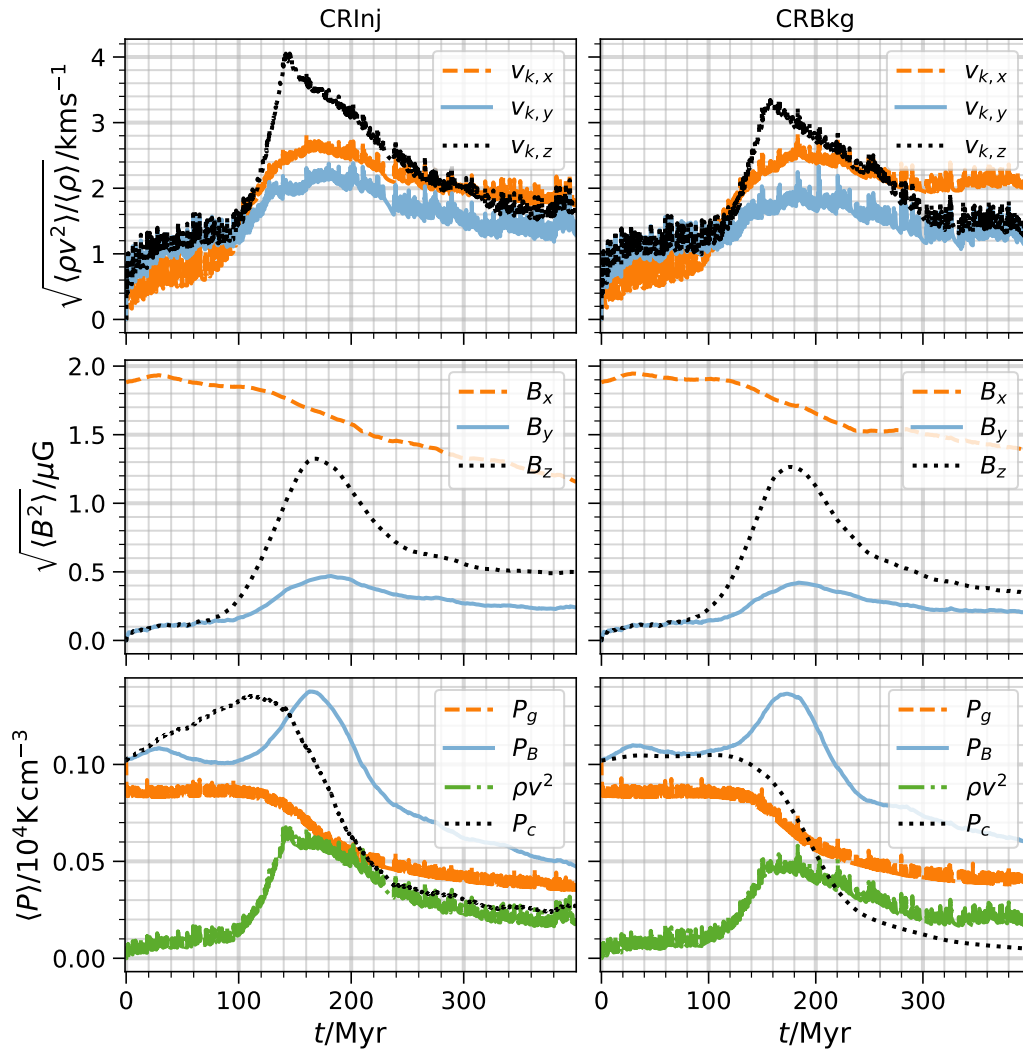


Figure 6.3: The time evolution of simulations CRInj (left column) and CRBkg (right column), with each quantity averaged over the entire simulation volume. The top row shows the average velocity divided into its directional components (\hat{x} , \hat{y} , \hat{z}), and the middle row shows the same for magnetic field strength. Compared to the simulation with only an initial background of cosmic rays, additional cosmic ray injections amplify the vertical velocity. The exponential growth of vertical magnetic field is related to the Parker instability (see Section 6.4.2). The bottom row compares thermal, magnetic, turbulent, and cosmic-ray pressures. The average cosmic ray pressure reaches a steady state near the kinetic pressure and below the average gas and magnetic pressures.




Figure 6.4: *Top:* Mass outflow and inflow rates for regions above and below the midplane ($|z| > 1$ kpc). Solid lines show the CRInj simulation and dashed lines show the CRBkg simulation. The outflows and inflows are out of phase, suggesting the simulations produce a fountain flow. Since the effective star formation rate in our simulation volume is $\sim 10^{-4} M_{\odot} \text{yr}^{-1} \text{kpc}^{-2}$, the outflows have mass loading factors $\eta = \dot{M}_{\text{out}}/\dot{M}_{\text{SFR}} > 10^3$. *Bottom:* Mass fraction of different phases in terms of the initial mass in the simulation. After $t = 0$, the sum of these phases can be less than one because mass can flow out of the top and bottom boundaries. Again, the solid lines show the CRInj simulation and the dashed lines show the CRBkg simulation.

6.4.1 Time Evolution

We run the simulations for 400 Myr. Initially, the entire setup is destabilized by many energy injections. These perturbations drive the Parker instability, immediately producing a fountain flow which becomes nonlinear at approximately 100 Myr into the simulation. At 200 Myr the simulation begins to settle into a steady state structure. This steady state lasts from approximately 250 Myr until the end of the

simulation at 400 Myr.

These evolutionary stages are apparent in Figure 6.3, which illustrates how the average velocity, magnetic field strength, and pressure components in the simulation vary over time. The left column of Figure 6.3 shows the results from the `CRInj` simulation, and the right column shows the results from the `CRBkg` simulation.

In the top row of Figure 6.3, we show the directional components of velocity averaged over the simulation volume. During the initial destabilization, most of the motion ends up being vertical, perpendicular to the midplane. Before $t \sim 250$ Myr, the `CRInj` simulation produces more motion in every component direction than the `CRBkg` simulation, but most significantly in the vertical direction. Therefore, the additional cosmic-ray pressure from the injections has a significant impact on the dynamics of the gas. In the eventual steady state after $t \sim 250$ Myr the average velocity in each direction levels off to a constant value. The velocity parallel to the mean magnetic field (\hat{x} direction) ends up being the largest. The spikes are due to individual energy injections. A key difference between the simulations in the steady state is the ordering. In the `CRInj` simulation, the vertical motion is distinctly larger than the planar motion perpendicular to the magnetic field (\hat{y} direction). However, the `CRBkg` simulation has these velocities at about the same value.

The middle row of Figure 6.3 shows the time evolution of the directional components of magnetic field averaged over the entire simulation volume. Both simulations have a similar evolution - the magnetic field starts only in the \hat{x} direction, and this energy decays because magnetic flux in the \hat{x} and \hat{y} direction can escape out of the top and bottom of the simulation box. The magnetic field strength in the vertical

direction increases, exponentially, until approximately $t \sim 180$ Myr. In Section 6.4.2, we show this increase is tied to the Parker instability's characteristic structure. After $t \sim 180$ Myr, the magnetic field decays. This decay might be offset by amplification due to shearing if we had included galactic rotation in our models. Additionally, we observe numerical magnetic reconnection, which contributes to the decrease in magnetic field strength. Overall, the loss in horizontal magnetic field illustrates how the Parker instability can efficiently decrease, redistribute, and reorient the magnetic field

The bottom row of Figure 6.3 compares the evolution of thermal, magnetic, turbulent, and cosmic-ray pressures averaged over the entire simulation volume. Initially, the midplane gas, magnetic, and cosmic ray pressures are all $10^4 k_B \text{Kcm}^{-3}$ (see Table 6.1). The total turbulent pressure is initially zero, but it rises to the order of magnitude of the other pressure components. The thermal and turbulent pressures have continual spikes as a result of the injections. The thermal pressure also has an underlying decrease as cold, dense gas forms from optically thin radiative cooling. The total magnetic pressure reflects the decrease in the \hat{x} directed component and the spike in vertically directed magnetic field. Finally, the cosmic-ray pressure shows the biggest difference between the two simulations. The `CRInj` simulation shows a linear increase early on as the cosmic-ray energy injections add to the background cosmic-ray pressure. Both simulations show a decrease in cosmic-ray pressure which is tied to the increase in vertical magnetic field. Since the cosmic rays move along the direction of the local magnetic field, there needs to be magnetic flux directed out of the top and bottom of the simulation to allow the cosmic-ray energy to escape.




Figure 6.5: *Left:* Kinetic energy spectrum $\tilde{v}(k)$ (in units of km s^{-1}) for the CRInj simulation at each 1 Myr time dump. The color map shows the amplitude of velocity fluctuations at each wavelength $L = 2\pi/k$ and the green line traces the peak wavelength of the spectrum. The spectrum is primarily driven by supernovae, at a scale of $L \sim 20$ pc. Eventually, the larger scale $L \gtrsim 0.33$ kpc mode dominates. *Right:* Amplitude over time of four key Fourier modes, along with line showing the growth rate of the $L \gtrsim 0.5$ kpc. The $L = 0.5$ kpc mode grows as a result of the Parker instability, independent of the large scale modes until it reaches their level. While that mode never becomes dominant, it feeds the smaller scale modes once it grows to the same strength at $t \sim 80$ Myr, causing the $L \gtrsim 0.33$ kpc mode to dominate. While the $L \gtrsim 0.5$ kpc mode grows consistently until joining the other modes, the small scale modes saturate after $t \sim 20$ Myr.

As the vertical magnetic field strength decreases, the cosmic-ray pressure levels off in the CRBkg simulation. The CRInj simulation also levels off, but to a non-zero value due to continuing cosmic-ray energy injections.

Moving on from energy evolution, we focus on mass evolution in the simulations. The top graph of Figure 6.4 shows the mass inflow and outflow rates outside of the midplane ($|z| > 1$ kpc). The solid lines show the CRInj simulation and the dashed lines show the CRBkg simulation. The cosmic-ray injections lead to a larger peak mass outflow (and inflow) rate, and they drive a larger steady outflow at late times

($t > 300$ Myr). In the CRBkg simulation, the inflow and outflow nearly balance each other at late times. There also appears to be a fountain flow component, where outflows are followed by inflows at a delay of about ~ 30 Myr. Finally, the two simulation curves are almost exactly the same until $t \sim 100$ Myr, when the Parker instability begins to change the magnetic field orientation.

In the bottom of Figure 6.4, we show the mass fractions of different gas phases. Warm gas has a temperature $4 \cdot 10^3 \text{K} \lesssim T \lesssim 10^5 \text{K}$, whereas unstable gas has a temperature $100 \text{K} \lesssim T \lesssim 4 \cdot 10^3$. Cold gas has a temperature $T \lesssim 100 \text{K}$. In pressure-density phase space (see Figure 6.2) the warm gas stretches between the minimum of the equilibrium curve up to the peak above a gas density of 1cm^{-3} . The unstable gas is between that peak and the minimum just below a density of 10cm^{-3} , and the cold gas ranges down to the temperature floor of $T = 30 \text{K}$. It should be noted that the mass fractions in Figure 6.4 do not add to unity at all times. There is a small loss of mass out of the outflow boundaries which causes the total mass to decrease by a factor of $\lesssim 3\%$ over the 400 Myr of evolution.

Overall, we see that initially ($t \lesssim 120$ Myr) some thermally unstable gas is created, which eventually becomes cold gas. The lack of a significant mass fraction of unstable gas is characteristic of weak multiphase turbulence (Ho et al., 2024). That conclusion is consistent with our simulations - most of the kinetic energy in our simulations are vertical, and associated with the bulk outflow in the top panel of Figure 6.4. Driving the gas into the strong turbulence regime would require more supernova injections.

After the unstable gas forms, the cold gas quickly becomes the dominant component of the overall mass fraction after $t \sim 250$ Myr for both simulations. There is

a slight delay in the CRBkg simulation, which is shown with dashed lines. The peak in unstable gas formation tracks with the inflow rate in the top plot of Figure 6.4. This correlation suggests a connection between the fountain flow and the formation of cold gas via thermal instability. In Section 6.4.2, we show this connection comes from a combination of the Parker instability and Field’s thermal instability.

6.4.2 Parker Instability

Now we examine the initial development of the Parker instability during the beginning of the simulation, $t \lesssim 200$ Myr. First, we illustrate that the growth of the Parker instability is due to large scale velocity fluctuations which grow linearly out of the small scale energy injections we input. The primary difference between our simulations and previous simulations is that our perturbations come from discrete supernovae instead of small velocity perturbations (e.g. Heintz et al. (2020)).

Figure 6.5 shows the evolution of the kinetic energy spectrum. We calculate the spectrum in a 1 kpc^3 volume centered on the midplane - specifically, we only include cells with $|z| \leq 0.5 \text{ kpc}$, matching the length of the x and y directions. On the left side of Figure 6.5, we show the time evolution of the kinetic energy spectrum. Most of the energy is at small scales $L \leq 1/3 \text{ kpc}$. The wavelength with the most energy at each time is plotted as a green line. For the first 100 Myr, the dominant mode oscillates between 0.1 kpc and 1/3 kpc. However, the 1/3 kpc then becomes dominant.

Looking at the right side of Figure 6.5, it is apparent that the small scale modes quickly saturate after just 20 Myr. Then, the larger scale modes (specifically $L =$




Figure 6.6: *Left:* Vertical velocity colormap with gas density contours and magnetic field lines overlaid. The large plume centered at $x \sim 0.1 \text{ kpc}$ is moving upward at a fast rate ($v_z \sim 50 - 100 \text{ km s}^{-1}$). *Right:* Colormap of cosmic-ray pressure with gas density contours and magnetic field lines overlaid. The density contours highlight the thermally unstable gas phase which forms below where the plume expands into itself (because of periodic boundary conditions). The plume is driven upward partly by the cosmic-ray pressure gradient, which is apparent in the right side plot. Both the left and right simulation slices are taken from the CRInj simulation at $y = 352.5 \text{ pc}$ and at a time of $t = 135 \text{ Myr}$. The crimson box in the right hand figure shows the zoom-in region pictured in Figure 6.7.

1/2 kpc) eventually rise to join them at $t \sim 100$ Myr. results from the growth of the larger scale modes. This large scale mode drives a cascade which amplifies the small scale modes, causing $L = 1/3$ kpc to become dominant. While the supernova injections are the clear driver of the kinetic energy at early times on small scales $L \leq 1/3$ kpc, it is not obvious what drives the larger scale.

The large scale mode $L = 1/2$ kpc grows exponentially at a rate of $\tau = 15$ Myr initially, before saturating at a value lower than the smaller scale modes. While this initial growth could be caused by the fadeaway of supernova explosions, the supernova explosions do not explain the further driving of $L = 1/2$ kpc. At $t \sim 60$ Myr, the $L = 1/2$ kpc grows exponentially again, at a slower rate of $\tau = 25$ Myr. This rate is closer to the $\tau = 33$ Myr rate predicted from the linear theory of the Parker instability (see Section 6.3.4) than it is to the lifetime of a supernova explosion or to our injection rate.

Eventually, even the driving by the Parker instability saturates. This saturation is apparent in the right side of Figure 6.5 for $t \gtrsim 160$ Myr. The saturation is significantly faster than ~ 0.5 Gyr saturation timescale which previous Parker instability simulations have found (Heintz et al., 2020; Tharakkal et al., 2023). After the saturation point, all the modes decay in amplitude.

Although the saturation time is impacted by our limited box size along the magnetic field direction, the rapid outflow begins a little before the saturation (see top plot of Figure 6.4). This outflow suggests the dynamical impact of the Parker instability still happens sooner than previously expected. In fact, it happens at a rate similar to what we saw from single cosmic ray energy injections in a Parker unstable




Figure 6.7: A zoom-in on dense gas from Figure 6.6. The condensing gas is in regions of light purple and orange, highlighting that it forms in low cosmic-ray pressure regions. This process decorrelates gas density and cosmic ray pressure. The sinking dense gas also drives an increase in magnetic field strength as it compresses magnetic field lines together. Note that the magnetic field line colors now follow a linear scale as opposed to the logarithmic scale used in Figure 6.6.

medium (Habegger et al., 2023).

Now, we focus on the growth and saturation of the large scale modes of the Parker instability. Figure 6.6 shows slices of the CRInj simulation at time $t = 135$ Myr and $y = 352.5$ pc. These slices illustrate saturation of the Parker instability, with plumes which erupt from the simulation’s midplane, also producing valleys of cold gas (Parker, 1966). The left plot of Figure 6.6 shows the velocity of the gas in the simulation, whereas the right plot shows the cosmic-ray pressure. Both plots have magnetic field lines and gas density contours overlaid. The gas density contours are selected to only be around gas which is thermally unstable.




Figure 6.8: *Top:* Growth of vertical kinetic and magnetic energy density in the CRInj simulation is exponential with a characteristic time of $\tau \sim 18\text{ Myr}$. *Middle:* Column density colormap and average cosmic-ray pressure contours at $t = 90\text{ Myr}$. *Bottom:* Same as middle plot, but at $t = 140\text{ Myr}$. The middle panel shows a relatively uniform medium at the time before the magnetic energy density reaches the kinetic energy density. The bottom panel shows the decorrelation of cosmic ray pressure and gas column density after the Parker instability saturates. The cosmic ray pressure is large in the expanding plumes.

The plume is moving upward at $v_z \lesssim 100 \text{ km s}^{-1}$. The cosmic-ray pressure is higher inside of the rising plume of gas, showing that the cosmic-ray pressure gradient is helping to drive the gas flow. The rising plume pushes gas out of the way. Some gas rolls off the top of the plume and falls back to the midplane along magnetic field lines. This gas compresses at $z \lesssim 0.25 \text{ kpc}$, forming thermally unstable gas. The Parker instability picture is generally focused on the formation of these cold gas valleys. However, the roll off of magnetic flux tubes (and gas) *perpendicular* to the mean magnetic field greatly facilitates the rise of the buoyant flux tubes, highlighting the importance of the third dimension (Habegger et al., 2023). I.e. a two-dimensional simulation would have resulted in lower values of v_z since the expanding flux tube would need to lift all gas directly above it (Habegger et al., 2023).

Looking at the right plot in Figure 6.6, it is also apparent that the unstable gas is sinking into gas with a higher cosmic ray pressure. We present a zoom in on this region in Figure 6.7. Figure 6.7 shows the individual clouds have lower cosmic ray pressures, with colors of light purple, white and orange, whereas the midplane is a dark purple. This results in a decorrelation of cosmic ray pressure and gas density. Additionally, between the gas clouds, there is a magnetic field aligned cosmic ray pressure gradient, allowing cosmic ray pressure to leave the midplane. As the gas sinks, it also compresses the magnetic field, leading to an amplification at the edge of the largest cloud in Figure 6.7.

Figure 6.8 further illustrates this decorrelation by the Parker instability. During the exponential growth of vertical kinetic and magnetic energy density (top plot), the variations in column density and average cosmic-ray pressure in the \hat{z} direction

change significantly (bottom two plots). Additionally, we find a growth rate of approximately $\tau = \dot{E}/E = 18 \text{ Myr}$ which is faster than the $\tau \sim 33 \text{ Myr}$ growth rate predicted using a linear Parker growth rate. At $t = 90 \text{ Myr}$ (middle plot of Figure 6.8), the gas is fairly uniform in column density and average cosmic-ray pressure. These quantities are shown with a colormap and contour lines, respectively.

By $t = 140 \text{ Myr}$ (bottom plot of Figure 6.8), the Parker instability has decreased the column density of the outflowing plume regions and decreased the average cosmic-ray pressure along columns where dense gas is forming. The dense gas regions also correspond to the footpoint regions where the magnetic field has become mostly vertical (see right side of Figure 6.6). This vertical field allows cosmic-ray energy to rapidly escape toward higher values of z as they stream down the magnetic field aligned pressure gradient. Eventually that cosmic-ray energy leaves the simulation through the vacuum boundaries at the top and bottom of our simulation.

The Parker instability and this decorrelation may not remain in a fully turbulent galactic disk, with spiral arms driving additional gas compressions and stirring the magnetic field. But in general, the decorrelation could appear wherever the magnetic field becomes predominantly vertical outside of the midplane. When vertical field appears, it gives the cosmic rays a pathway to escape and the gas a pathway to sink into the galactic plane and condense.

6.4.3 Saturated State

The simulations develop a saturated and steady state after the Parker instability disrupts the original hydrostatic equilibrium. We focus on the time period between




Figure 6.9: The average vertical density profile over the period $t = [260, 400]$ Myr. The solid green line is from the CRInj simulation and dashed purple line is from the CRBkg simulation. We calculated the profile at every 1 Myr time dump by averaging over each xy -plane and then averaging those vertical profiles over time. The shaded regions show the full variation (min and max average density value at a given z position) for each simulation. Both simulations form a thin cold gas disk ($n > 1 \text{ cm}^{-3}$), a thin warm disk ($n < 1 \text{ cm}^{-3}$), and a thick warm gas disk ($n < 10^{-1} \text{ cm}^{-3}$).




Figure 6.10: xy -plane averaged outflow/inflow velocities from both simulations. The averaged velocity profile over the steady state time frame ($t \in [260, 400]$ Myr) is shown for each simulation. The solid green line is from the CRInj simulation and dashed purple line is from the CRBkg simulation. The shaded regions show the full variation of the profile over the steady state time frame. The CRInj simulation has a faster outflow which also starts from closer to the midplane.

$t = 260$ Myr and $t = 400$ Myr, even though the steady state appears to start as early as $t = 200$ Myr in Figure 6.3. We start the analysis at $t = 260$ Myr to avoid the transition period from the nonlinear evolution of the Parker instability. This transition is clearer in Figure 6.4. There, the outflow and inflow rates, as well as mass fractions of each phase, only level off to consistent values by $t = 260$ Myr.

In the steady state, we focus on the vertical structure which appears in our simulation. Figure 6.9 shows average vertical density profiles over time. The density profiles are calculated by averaging over the xy -plane at every value of z . In Figure

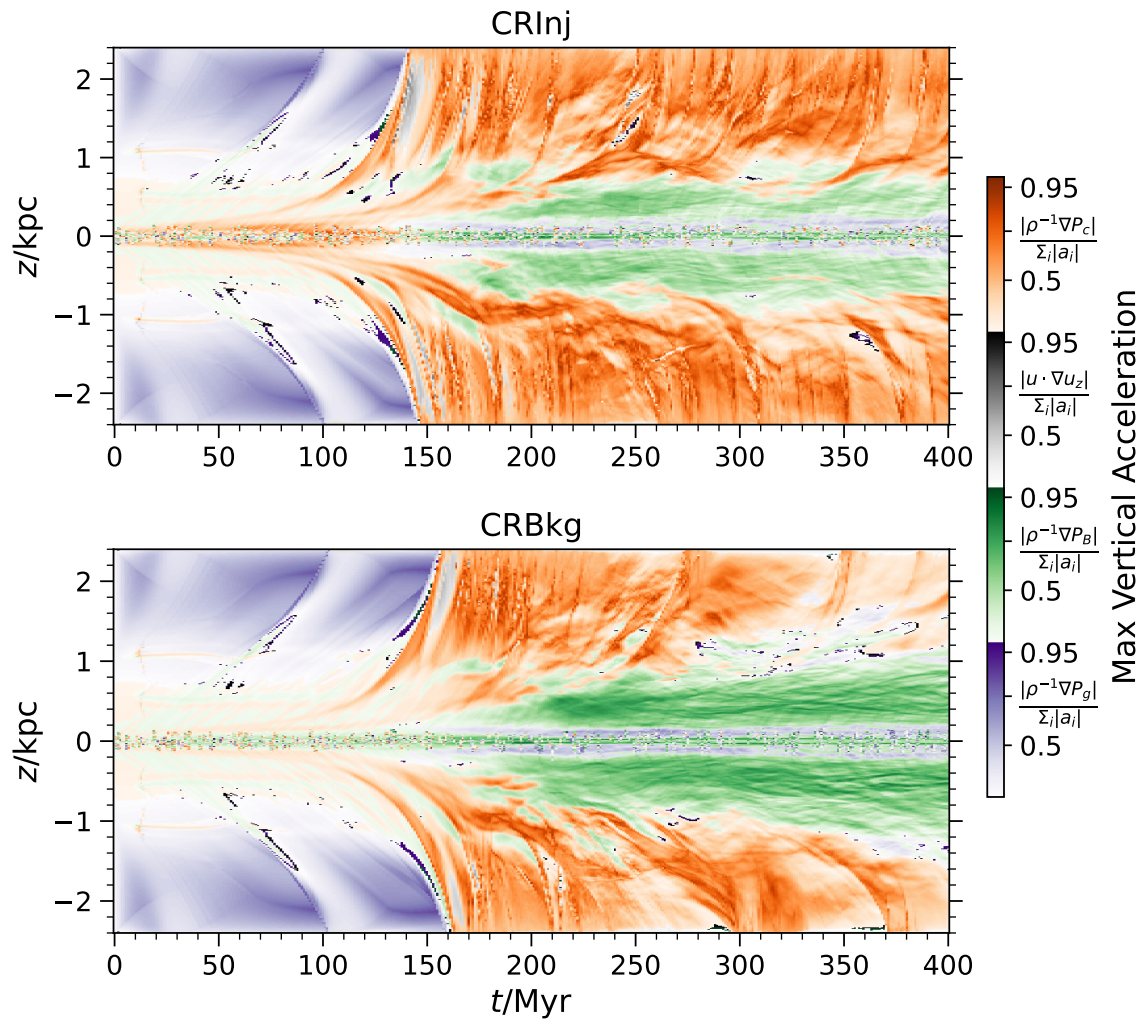


Figure 6.11: Spacetime diagrams of the dominant acceleration components in the CRInj simulation (top figure) and the CRBkg simulation (bottom figure). The lower quarter of the colorbar (purple) is for regions where the xy -plane averaged vertical acceleration is predominantly driven by the gas pressure gradient, the next quarter (green) is when the magnetic pressure gradient is dominant, the following quarter (black) is where the advective motion dominates the acceleration, and the top quarter (orange) is where the cosmic-ray pressure is dominant. The CRInj simulation has much stronger cosmic-ray pressure gradients at later times, once the magnetic field is mainly vertical.




Figure 6.12: Vertical acceleration by various forces in the CRInj simulation (left plot) and in the CRBkg simulation (right plot). The solid black line shows the gravitational acceleration from the background profile, the solid orange line shows the acceleration from the vertical cosmic-ray pressure gradient, the dashed purple line shows the acceleration from the gas pressure gradient, and the dash-dotted green line shows the vertical magnetic pressure gradient. Each line (except the constant gravitational acceleration) has a shaded region showing the full variation of the acceleration from the given force during the steady state time frame ($t \in [260, 400]$ Myr).

6.9 the lines show the average profile over the time frame $t = [260, 400]$ Myr for each simulation. The shaded regions show the full variation of the profile over that time frame. The dotted line shows the initial hydrostatic equilibrium.

Figure 6.9 shows there is both more gas in the midplane and more gas above $|z| \sim 1$ kpc than at the start of the simulations. The spike in gas density near $z \sim 0$ is cold gas that has settled after forming via thermal instability. There is a switch at large vertical positions, where gas has been pushed out by the Parker instability. This restructuring means a smaller scale height for cold, high density gas and a larger scale height for warm, low density gas. This new quasi-equilibrium appears in both




Figure 6.13: *Bottom:* Net acceleration in each simulation, calculated by adding up the forces in Figure 6.12. *Top:* Ratio of net acceleration between the CRInj simulation and the CRBkg simulation. Except for near the midplane, the CRInj simulation always has a larger net acceleration directed away from the midplane (i.e. negative acceleration for $z < 0$ and positive acceleration for $z > 0$).

simulations. So, the cosmic-ray injections have a minimal impact on the resultant density stratification, because the averaged profiles are similar in both simulations. The CRInj does end up with slightly denser gas in the region $z > 1\text{kpc}$, adjusting the scale height slightly (see Section 6.4.4).

Next, we can look at the flow speed in the saturated state. Figure 6.10 shows the vertical profile of outflow speed. We calculate the average outflow speed at every value of z , integrating over the xy -plane. We multiply this value by the sign of the z coordinate to visualize outflow (if the quantity is positive) and inflow (if the quantity is negative). Figure 6.10 shows the outflow velocity profiles for each simulation averaged over the steady state time frame $260\text{ Myr} < t < 400\text{ Myr}$, and the shaded regions show the full variation in the profile over that time frame. The

CRInj simulation shows a steady outflow beginning at $|z| \sim 0.5$ kpc which could contribute to a galactic wind. This steady outflow contrasts the state in the **CRBkg** simulation which shows more fluctuation between outflow and inflow in the time-averaged profile, only having a dominant outflow for $|z| \geq 1.2$ kpc. The fluctuation between outflow and inflow is also apparent in the volume integrated mass outflow and inflow rates plotted in Figure 6.4.

Presumably, the cosmic-ray injections are driving the outflow in the **CRInj** simulation, since that is the only difference between it and the **CRBkg** simulation. To prove this, we calculate the average vertical acceleration for gas pressure, cosmic-ray pressure, magnetic pressure (only the x and y components of the magnetic field contribute), and nonlinear acceleration $\mathbf{u} \cdot \nabla u_z$ in every xy plane. We show the maximum vertical acceleration at each vertical position z at each time dump t in Figure 6.11. The strength of each color shows the fraction of the total acceleration for the dominant component. The notation $\sum_i |a_i|$ just refers to summing over all the accelerations included in Figure 6.11. Note that we use the absolute value of each acceleration component so that there is no cancellation. The primary difference between the two simulations is the large cosmic-ray pressure gradient supported regions in the steady state period of the **CRInj** simulation. The cosmic-ray pressure gradient is the dominant acceleration term outside of the midplane ($|z| \gtrsim 1$ kpc) by a large amount, $\geq 95\%$. This cosmic-ray pressure gradient dominance contrasts with the steady state of the **CRBkg** simulation, where the cosmic-ray pressure gradient is less dominant, at only $\lesssim 50\%$ of the acceleration for late times.

Both simulations show a pinched region of magnetic support in the midplane,

where the gas density is highest, surrounded by regions of gas pressure support. This switch in pressure support within the midplane derives from the switch from cold to warm gas - the cold gas has pressure support from its magnetic field but the warm gas is supported by its own gas pressure. Outside of $|z| \sim 250$ pc, the magnetic pressure again takes over in both simulations before falling off.

Taking a closer look at the accelerations, we plot each acceleration component averaged over the saturated time frame in Figure 6.12. The lines show the average acceleration due to the background gravitational field (solid black line) and due to the gradients of cosmic-ray pressure (solid orange line), gas pressure (purple dashed line), magnetic pressure (green dash-dotted line), and nonlinear acceleration (black dotted line). The shaded regions show the full variation of the acceleration over the saturated time frame. The `CRInj` simulation has much more acceleration from the cosmic-ray pressure gradient term, which vanishes in the midplane. That is the only difference - the magnetic pressure and gas pressure acceleration profiles in the simulations are similar. They both peak at small z values and then fall off with height.

The net acceleration in each simulation is shown in the bottom plot of Figure 6.13. Both simulations have mostly outward (from the midplane) acceleration, which matches with the volume integrated inflow/outflow integration in Figure 6.4. These accelerations are only estimates because we do not account for the magnetic tension in the calculation. However, the tension force was predominantly zero at most times because the quantities are averaged over the xy plane at each value of z . The average tension shows up as flat lines when we include them in Figure 6.12, so we remove it to




Figure 6.14: xy -plane averaged fraction of the magnetic field directed vertically. The lines show the average vertical magnetic field fraction for the steady state time period ($t \in [260, 400]$ Myr). The solid green line is from the CRInj simulation and dashed purple line is from the CRBkg simulation. The shaded regions show the full variation of the profile over the steady state time frame. The CRInj simulation produces more vertical magnetic field in the midplane, resulting in a slightly lower height where the vertical field strength becomes dominant.

make the figure easier to understand. However, the tension does have some variation ($\sim \pm 0.2 \text{ km s}^{-1} \text{ Myr}^{-1}$) which could increase or decrease the net acceleration.

The top plot of Figure 6.13 shows the ratio of the accelerations, illustrating that the CRInj always has more acceleration. In some positions it has as much as $10\times$ the acceleration of the CRBkg simulation at large values of $|z|$. In the regions above 0.5 kpc , the cosmic ray injections provide a factor of $3\times$ acceleration compared to the CRBkg simulation. This increased acceleration further illustrates the ability of cosmic rays to drive outflows from a galactic disk, even if they do not take over as a forcing mechanism until gas reaches the galactic halo.

The dominance of cosmic-ray pressure support in the extraplanar gas deserves more scrutiny. Why does the cosmic-ray pressure support only become significant outside the midplane? To answer this question, we have to examine cosmic-ray transport throughout our simulations.

First, we examine the magnetic field structure, which can impact the cosmic-ray transport. The rapid falling off of magnetic pressure with height in Figures 6.11 &




Figure 6.15: Average phase diagrams for the CRInj simulation (left column) and the CRBkg simulation over the steady state time frame ($t > 260$ Myr). *Top Row:* cosmic-ray pressure-gas number density phase diagram, which show predominately diffusive transport in the high density midplane gas and predominantly streaming transport in the diffuse extraplanar gas. *Middle Row:* Gas pressure-gas number density phase diagram. The histogram follows the heating cooling equilibrium curve, shown as a dashed green line (see Figure 6.2), except the CRInj deviates at high gas density where Alfvénic heating from the cosmic rays pushes the gas off the equilibrium curve. *Bottom Row:* Temperature-gas number density phase diagram. The CRInj ends up with higher temperature gas due to the Alfvénic heating from cosmic-ray streaming. We also plot the cut-off temperature we use to remove cold and thermally unstable gas in Section 6.4.4.

6.12 is partly because of a change in magnetic field orientation. The beginnings of that orientation change were apparent in the Parker instability growth pictures in Figure 6.6. Focusing on the averaged vertical structure, Figure 6.14 shows the xy -plane averaged fraction of the magnetic field in the vertical direction $b_z = |B_z|/B$. We use the absolute value because there is an equal amount of positive and negative vertical magnetic flux when integrating over each xy -plane. Above $|z| \sim 1$ kpc, the magnetic field becomes predominantly vertical in both simulations during the saturated time frame. There is more vertical magnetic field within the midplane of the CRInj simulation compared to the CRBkg simulation (although it is still subdominant to the horizontal field). As a result, there is less magnetic pressure available to support gas in the CRInj simulation. In the CRBkg simulation, there is more horizontal field in the midplane, but the vertical field still takes over in the extraplanar gas at $|z| \gtrsim 1$ kpc.

Second, we can examine the cosmic-ray transport in the midplane compared to the transport in the extraplanar gas. In the top row of Figure 6.15, we show the cosmic-ray pressure-gas number density phase space. The histograms are averaged over the steady state time frame of $t \in [260, 400]$ Myr. These figures show a clear differentiation in cosmic-ray transport between the dense gas and diffuse gas. In the dense gas, cosmic rays rapidly diffuse, ending up almost uniformly distributed across orders of magnitude in gas density. Once the cosmic rays escape the dense gas, they enter the extraplanar gas which has a lower density, and high speeds, allowing advective and streaming transport to dominate. This transition in transport matches what has been seen in previous tallbox simulations with post-processing of cosmic-ray

transport, highlighting a transition to advective transport (Armillotta et al., 2021, 2024).

In both our simulations, the streaming transport is dominant over advection. This streaming transport is apparent in the histograms, which do not follow the $\gamma_c = 4/3$ polytropic law ($P_c \propto n^{\gamma_c}$) for cosmic rays advected with the thermal gas. Instead, it follows a $\gamma_c = 2/3$ index, which is expected for a steady state system with streaming transport (Wiener et al., 2019). This result is unexpected given our system is not in a perfect steady state and other transport mechanisms (advection and diffusion) are also included in the simulation. Likely, since the magnetic field is predominantly vertical and aligned with the cosmic-ray pressure gradient, the streaming becomes dominant. Additionally, the simulations have a slightly stronger magnetic field than other works, and the turbulence is weak. These changes amplify streaming as a transport mechanism over advection.

Figure 6.15 also shows the gas pressure-gas number density phase space (middle row) and the gas temperature-gas number density phase space (bottom row). The histograms are averaged over the steady state time period. For the most part, the profiles in both simulations follow the thermal equilibrium curve of $\Gamma = n\Lambda$ shown in Figure 6.2 and plotted as a dashed green line in middle row plots. However, the CRInj simulation's histogram departs from this curve at low gas densities, and this low density gas is also at a higher temperature in the temperature-density phase space. This heating comes from the streaming transport of cosmic rays, which takes over in the diffuse gas. At the lower density, the net cooling time increases and is unable to offset cosmic-ray heating, leading to increased temperature. While we have not done

any simulations in which streaming is present but cosmic-ray heating is removed from the energy equation, the relative paucity of low density, high temperature in the CRBkg model demonstrates that this is an important effect.

After this examination of the steady state, a clear picture of the interplay between magnetic field orientation and cosmic-ray feedback emerges. The shift to a predominantly vertical magnetic field outside the midplane opens the door for a cosmic-ray driven outflow. While injected cosmic rays diffuse in the midplane, they stream and provide a large pressure gradient outside the midplane, producing a faster wind compared to the simulation without cosmic-ray injections. The streaming of cosmic rays also causes them to heat the thermal gas, producing a hotter outflow.

Before we conclude, we test our simulations against observational results. To do this, we calculate and estimate some observable properties of our simulated galactic patch.

6.4.4 Observational Comparison

Using our xy -plane averaging process from Section 6.4.3, we can fit vertical profiles and measure scale heights for the gas density, gas pressure, cosmic-ray pressure, and magnetic pressure.

Figure 6.16 shows the average vertical profiles and our median fit to those profiles during the saturated, steady state. The averaged data are plotted within a shaded region, which shows the total variation over the saturated time frame. We removed the cold and thermally unstable gas disk in the midplane to isolate the impact on warm and hot gas, by using only gas cells with a temperature $T > 4000$ K for the




Figure 6.16: Average vertical profiles from the CRInj simulation, fitted with two independent scale heights. The red shaded regions show the standard deviation of the vertical profiles during the saturated time frame, and the solid red lines show the median profile. The dashed green lines show the initial hydrostatic equilibrium, and the dashed black lines show the best fit to each profile. The profiles are produced by first removing all cold and unstable gas with $T < 3000$ K

xy -plane averaging. We use a linear combination of the base hydrostatic equilibrium profile (see Section 6.3.1):

$$f(z|A, H) = A \left(\operatorname{sech} \left[\frac{z}{H_*} \right] \right)^{H_*/H}. \quad (6.22)$$

At each 1 Myr time dump, this gives us a pair of scale heights H_{inner} & H_{outer} and amplitudes A_{inner} & A_{outer} . We chose the above profile because it is the equilibrium




Figure 6.17: Gas scale heights in CRInj simulation and CRBkg simulation. Black lines show the inner warm gas density scale heights, and red lines show the outer warm gas density scale heights. The green lines show the scale height for the magnetic energy density, and blue lines show the scale height of the thin, cold disk. Solid lines show CRInj simulation, and dashed lines show the CRBkg simulation.

profile under the influence of our chosen gravitational acceleration (Equation 6.12). The fitted profiles are pictured for gas density and gas pressure in Figure 6.16, with a central peak near the midplane and large wings outside the midplane.

While the fitted scale heights can be related to an exponential scale height for $z \geq H_*$, we do not expect them to always match observed scale heights in the Milky Way. This expectation is because the gravitational acceleration profile we use is a simplification of the true acceleration profile. Instead, we encourage a focus on the structure which our simulations produce: a thinner, dense disk and a thicker disk inflated by non-thermal pressures. A two scale height model matches with the modern understanding of the structure of the Milky Way’s thermal gas (McClure-Griffiths et al., 2023; Rybarczyk et al., 2024).

For the non-thermal pressure, there is a different story. The cosmic-ray pressure profile in the third row of Figure 6.16 looks significantly different and cannot be fit

by the hydrostatic equilibrium profile in Equation 6.22. The magnetic energy density profile in the fourth plot still has an inner profile and an outer profile, each following Equation 6.22. However, the outer profile only takes over at $z \approx \pm 1.5$ kpc instead of at $z \approx \pm 0.4$ kpc like the gas density and pressure. A similar structure also appears for the turbulent pressure in the bottom plot. Figure 6.16 also illustrates that the gas density, magnetic field, and cosmic rays do not trace each other well.

In Figure 6.17, we show the time evolution of scale heights of gas density and magnetic energy density over the saturated state for both simulations. We also include the scale height for the cold gas, which we find by fitting a single component with the form of Equation 6.22. We plot these quantities because they are measurable (e.g. gas density via HI emission and absorption, and magnetic energy density through synchrotron radiation). We only plot the inner profile for the magnetic energy density because the values we measure for the outer scale height ($\gtrsim 2$ kpc) are untrustworthy. The outer scale height of the magnetic energy density we find is close to the simulation's vertical box size (4.8kpc), and the outer scale height only takes over above 1.5kpc (see Figure 6.16). The only point to be taken away from the outer scale height of the magnetic energy is that it flattens out due predominantly vertical magnetic field at large z (see Figure 6.14).

In Figure 6.17, it is apparent that the inner scale height of the warm gas, the inner scale height of the magnetic energy density, and the scale height of cold gas, are relatively constant during the steady state time frame ($t > 260$ Myr). The outer scale height of the warm gas varies significantly in both simulations. The inner and outer warm gas scale heights differ by a large factor ($\times 5$), and the magnetic

energy density's inner scale height falls in between the two warm gas components. On average, the CRInj has a slightly larger outer warm gas scale height component and larger magnetic scale height. But the simulations are in overall agreement with each other, and both illustrate the need for two components to fit the warm gas' structure.

While the cold gas component initially has a scale height of $\lesssim 50$ pc, it eventually decays to a value equal to the resolution of our simulation ~ 5 pc. Without a more complete feedback prescription (i.e. increase the supernovae rate as more cold gas forms), there is nothing stopping the cold gas from falling to $z = 0$ (see Section 6.5.3).

Another observable quantity to consider is γ -ray luminosity due to hadronic collisions of cosmic rays with thermal particles. These collisions produce a cascade of reactions which includes the emission of γ -rays. This luminosity is proportional to the cosmic-ray energy density, providing a useful tool for probing the cosmic-ray content of galaxies other than the Milky Way (Guo & Oh, 2008). In Figure 6.18, we show the evolution of the total energy lost in all hadronic interactions for our simulations, not just the energy emitted in γ -rays. We normalize this by the cosmic-ray energy injection rate in the CRInj simulation, following the definition of the $F_{\text{cal}} = L_{\pi}/\dot{E}_{\text{cr}}$ parameter in Lacki et al. (2011). We calculate the hadronic loss rate L_{π} using the formula for the hadronic loss rate per volume Γ_h in Equation 12 of Guo & Oh (2008):

$$L_{\pi} = \int_V d^3x \Gamma_h = 8.70 \cdot 10^{44} \frac{\text{erg}}{\text{Myr}} \sum_{\text{cells}} \left(\frac{E_c}{\text{eV cm}^{-3}} \right) \left(\frac{n}{\text{cm}^{-3}} \right) \quad (6.23)$$

For any subsection of the simulations, we can apply the formula in Equation 6.23.

When calculating L_π , we partition the each simulation into two regions: the midplane (within a stellar scale height, $|z| \leq H_*$) and everything except the midplane (i.e. cells with $|z| > H_*$).

In the bottom plot of Figure 6.18, we show F_{cal} for both simulations. The value can be larger than one in our simulations because we do not actually include the loss rate in our simulations. Additionally, these values are not the true measure in the case of the **CRBkg** simulation which actually has $\dot{E}_{\text{cr,inj}} = 0$.

Since most of the gas density is in the midplane and L_π is proportional to gas density (see Equation 6.23), most of the hadronic losses occur in the midplane. Initially, the losses are large because of the initial hydrostatic equilibrium we forced on the system (see Section 6.3.1). After $t \approx 250$ Myr, the hadronic losses in the **CRBkg** simulation vanish because much of the initial cosmic-ray background has escaped out of the top and bottom of the simulation. The emission outside the midplane in the **CRInj** simulation also vanishes. However, in the saturated, steady state we examined in Section 6.4.3, the hadronic losses level off to a steady value in the midplane of the **CRInj** simulation.

To get to a physically realistic and comparable F_{cal} parameter, we need to get rid of the initial background cosmic-ray pressure. For an approximation of this, we subtract the **CRBkg** hadronic loss rate off of the **CRInj** hadronic loss rate. This net calorimetric fraction is plotted in the top of Figure 6.18 for the midplane, outside the midplane, and their combination. Initially, as cosmic-ray injections are dumped into the simulation volume, the net F_{cal} increases at a constant rate. Then, at 100 Myr, when the Parker instability begins producing a strong vertical magnetic field, (see




Figure 6.18: *Bottom:* Fraction of cosmic-ray energy injected which would be lost to hadronic interactions, following the definition of F_{cal} in Lacki et al. (2011). We plot the results for both simulations, with solid lines showing the CRInj simulation and dashed lines showing the CRBkg simulations. We separate the emission from the midplane $|z| \leq H_*$, shown in purple, from the emission outside of the midplane, shown in green. *Top:* We calculate the Net F_{cal} by subtracting the CRBkg simulation's F_{cal} from the CRInj simulation's F_{cal} . This subtraction leaves only the F_{cal} associated with our cosmic-ray injections. For comparison, the blue shaded region shows observational estimates of F_{cal} in starbursting galaxies NGC 253 and M82 (Lacki et al., 2011).

Figures 6.6 and 6.8), the net F_{cal} decreases nearly as rapidly as it rose. With the vertical magnetic field, the cosmic-ray injections are able to escape the simulation faster, producing less hadronic interactions.

As the simulation saturates, we see a steady growth of the net F_{cal} parameter. However, its growth is at a lower rate than early in the simulation, despite our cosmic-ray injection rate being constant (the injections are stochastic in time, but the average rate is constant - see Section 6.3.3). Since there is still some vertical magnetic fields at these times, there is more escape than at the beginning of the simulation. But that escape is still not enough to agree with observations from Lacki et al. (2011). The only time our calorimetric fraction gets to the observational values is at the nonlinear saturation of the Parker instability, near $t \sim 180$ Myr when the vertical magnetic field strength is largest (see Figure 6.3).




Figure 6.19: *Top plot:* Net F_{cal} as calculated in Figure 6.18, in both our low resolution and high resolution simulations. *Bottom plot:* Slices of gas density and magnetic fields in the low resolution and high resolution simulations at $t = 100\text{ Myr}$ and $t = 180\text{ Myr}$. The vertical magnetic field facilitates the escape of cosmic rays from the midplane, which decreases the γ -ray luminosity and the net F_{cal} .

This minimum point in evolution of F_{cal} also depends on resolution. In Figure 6.19, we show the evolution of the net calorimetric fraction over time for a 10 pc resolution simulation and for our fiducial 5 pc resolution `CRInj` simulation. Both simulations follow a similar evolution over time, although the higher resolution simulation shows a steeper decrease in net F_{cal} during the saturation of the Parker instability.

The bottom part of Figure 6.19 shows how F_{call} decreases: the magnetic field becomes vertical between $t = 100$ Myr and $t = 180$ Myr. These field lines allow cosmic rays to escape the simulation, decreasing the γ -ray luminosity. After the pictured slices, the vertical field lines begin to numerically reconnect. This numerical reconnection is not physical, but it does disconnect the cosmic rays in the midplane from the outflow boundary conditions in the \hat{z} direction. This disconnect leads to an overestimation of γ -ray luminosity.

6.5 Discussion

Returning to the original goal of understanding how cosmic ray injections impact stellar feedback, we have seen their primary effect is in accelerating gas at $|z| \gtrsim 0.5$ kpc, building the base behind a galactic wind. This steady outflow lasted for $\gtrsim 100$ Myr and was primarily driven by the cosmic-ray pressure gradient. That gradient was sourced by cosmic rays streaming along vertical magnetic field lines and it was amplified when cosmic ray injections were included. This result is similar to what has been seen in galaxy scale simulations like Farber et al. (2018) and

Thomas et al. (2025), where cosmic rays drive a low density outflow and dense gas is left behind. In fact, our result could form the base of these large scale outflows, or source an outflow along an open magnetic flux tube (Breitschwerdt et al., 1991; Shimoda & Inutsuka, 2022).

However, there were also similarities between the two simulations. The cosmic-ray injections had a minimal impact on the vertical density structure (see Figure 6.9). The vertical density profiles of each simulation in the steady state were similar, with the only difference being in the evolution of the scale height of warm gas outside the midplane.

Additionally, the Parker instability developed in both simulations and drove a fast outflow while reorienting the magnetic field. The instability launched a short lived and massive outflow before the simulations settle into a steady, but weaker, outflow. The outflow slowed because the free energy (contained in the gravitational stratification and non-thermal pressures) of the initial conditions was released during the nonlinear saturation of the instability. The end effect of the Parker instability in the simulations was the creation of a vertical magnetic field, which facilitated cosmic ray escape and the formation of cold dense gas.

Our results also suggest some solutions to other questions and problems. The CRInj simulation illustrated that the orientation of the magnetic field has a significant impact on cosmic-ray transport. The change in transport has implications for the cosmic-ray calorimeter problem and the vertical stability of the ISM. Before we address these in Sections 6.5.1 & 6.5.2, note that these implications are not isolated to the Parker instability. The change in cosmic-ray transport we see is due to the

vertical magnetic field outside of the midplane. This vertical orientation could result from the Parker instability, as it did in our simulations, or from other mechanisms like superbubble chimneys (O’Neill et al., 2024).

6.5.1 γ -Ray Bright Galaxy Problem

The steady state structure and change in cosmic-ray transport we see in the `CRInj` simulation could provide a new solution to the problem of large γ -ray luminosities in global galaxy simulations. This problem was examined in depth with the FIRE simulations in Chan et al. (2019). They found that the only way to match the observed γ -ray luminosities of lower stellar surface density galaxies was to increase the cosmic-ray diffusion coefficient beyond the observationally estimated value.

The brightness problem can also be couched as a cosmic-ray calorimeter problem. Observational estimates have shown that smaller (lower Σ_*) galaxies are not cosmic-ray calorimeters, requiring $F_{\text{cal}} < 1$ (Lacki et al., 2011). But all of the streaming and Milky-Way diffusion coefficient simulations in Chan et al. (2019) were at or above the calorimetric limit. Therefore, using the observationally expected transport processes, simulated galaxies end up being cosmic-ray calorimeters, producing too many γ -rays.

We calculated the F_{cal} parameter from Lacki et al. (2011) and found that our simulation reaches their observational estimates for starbursting galaxies (see top panel of Figure 6.18), once we account for our initial cosmic ray background.

Our net F_{cal} parameter gets close to the observed values without cosmic ray diffusion amplified above the observed Milky Way value. Cosmic rays are transported diffusively, and via streaming, along the local magnetic field direction in

our simulation, like some of the simulations in Chan et al. (2019). Combining this anisotropic transport with a higher resolution in the diffuse gas allowed the cosmic rays to escape along vertically oriented magnetic field lines. We suspect resolution is the reason this has not been noted before. Since we resolve the entire simulation at $\Delta x = \Delta y = \Delta z = 5$ pc, we resolve more of the magnetic field structure outside the midplane. For comparison with finite mass codes, a gas mass resolution of $m_g \approx 0.129 M_\odot$ would be necessary to resolve $n = 10^{-2} \text{cm}^{-3}$ gas at a resolution of 5 pc. This mass resolution is much smaller than the $m_g \approx 10^3 M_\odot$ used in galaxy simulations like Chan et al. (2019).

However, direct calculation of the F_{cal} parameter (not the net value) suggests our CRInj simulation still has all cosmic ray energy lost to hadronic collisions (see bottom plot of Figure 6.18). This coincides with a rise in the net F_{cal} once numerical reconnection occurs, reducing the amount of vertical field. The numerical reconnection occurs where the Parker instability plumes run into each other. While there is still some vertical flux which survives (see Figure 6.14), the reconnection decreases the total magnetic energy density (see Figure 6.3). It is necessary to study this process in higher resolution to determine what happens when the effects of numerical reconnection are minimized.

We suggest spatially resolving the diffuse gas outside the midplane of the galaxy is an important ingredient in solving the problem of simulated galaxies being too bright in γ -ray luminosity. While the dynamics of diffuse gas outside the midplane does not require a high resolution, high resolution might be necessary to produce a realistic magnetic field structure.

Naively, we would expect magnetic field lines should stretch from the midplane, out and up into the galactic halo, and eventually the circumgalactic medium. Without higher resolution in the diffuse gas, these vertical field lines numerically reconnect and the regions become disconnected from the midplane. But if the vertical field lines were still there, connecting the midplane to the diffuse gas, then they would provide an escape highway for cosmic rays. These hidden highways could be the true solution to the calorimeter problem, instead of an increased cosmic-ray diffusion coefficient.

6.5.2 Cosmic-Ray Transport and Stability of the Disk-Halo Interface

In our simulations, the cosmic rays are predominantly diffusive at high densities but are transported by streaming at low densities (see top row of Figure 6.15). At lower densities (in completely ionized gas), the Alfvén speed gets larger, allowing the streaming transport to be larger than diffusive transport. Because of the steady stratification, this density dependence is also a vertical position dependence. This change in transport matches the recent results from other groups who identify a similar switch in cosmic-ray transport outside the midplane (Armillotta et al., 2024; Thomas et al., 2025).

A change in cosmic-ray transport also means a change in feedback. When cosmic rays are streaming at the Alfvén speed, they heat the gas (see middle and bottom rows of Figure 6.15) by scattering off of Alfvén waves they generate themselves. Additionally, the streaming process is an advective process instead of a diffusive process. Instead of smoothing out cosmic-ray pressure gradients like diffusion, any

steep cosmic-ray pressure gradients are advected at the Alfvén speed. So the switch to streaming means larger cosmic-ray pressure gradients, creating a larger force on the thermal gas, which can help drive outflows (see Figure 6.11). These impacts of streaming have been discussed extensively in other works (e.g. Tsung et al. 2023).

The P_c - n diagrams in the top row of Figure 6.15 point to a new idea regarding cosmic-ray feedback at the disk-halo interface. When cosmic-ray transport changes, it also changes the effective polytropic index of cosmic rays. A change in polytropic index changes the stability of the system (note the appearance of γ_c in Equation 6.21). Our simulations show an increase in polytropic index with height. So, the change in cosmic-ray transport with z could naturally shift the stability of gas parcels moving between the galactic disk and halo.

We developed this interpretation after reading Hosking et al. (2025), which illustrated how a plasma’s compressibility can impact the structure of metastable, stratified MHD atmospheres. A detailed analysis on the stability of a CR+MHD atmosphere is out of the scope of this paper. Instead, we just share the idea here and examine it in detail in a forthcoming paper (Habegger, Hosking & Zweibel, in prep.).

6.5.3 Caveats

Before concluding, we would like to detail some caveats to this work and the assumptions we made in designing the simulations.

Firstly, the gravitational profile we chose for this equilibrium (Equation 6.12) is unrealistic at large $|z|$ outside of the midplane, because (1) it does not include a

contribution from the dark matter halo and (2) it assumes an infinite plane of stars. At large $|z|$, the dark matter halo will have more of an impact and the finite size of the galaxy will become important. Overall, these effects will decrease the vertical component of gravitational acceleration at large $|z|$, increasing the outflow rate we find in the simulations, helping to launch a large-scale cosmic-ray driven galactic wind.

Next, our radiative heating mechanism of a constant Γ for all gas densities and vertical positions is unrealistic. The heating rate should fall off with vertical position z , allowing more radiative cooling to occur outside the midplane. This choice makes it harder for gas to cool at large z where gas densities are also lower. However, choosing a constant Γ makes the thermodynamic equilibrium curve (top panel of Figure 6.2) the same for the entire simulation volume, simplifying the analysis.

Also, we do not include cosmic-ray momentum boost of supernovae explosions (Diesing & Caprioli, 2018), or a local decrease in the cosmic-ray diffusion coefficient in star formation regions (Semenov et al., 2021). These additional complexities in cosmic-ray transport would primarily change cosmic-ray feedback in the midplane, but would have a smaller impact on larger length scales ($L \gtrsim 1$ kpc) and longer timescales ($t > 1$ Myr) we consider.

Finally, our work does not complete the full stellar feedback loop to show how cosmic-ray injections impact the star formation rate. To complete that loop, the injection rate would at least need to be determined by star formation rates (e.g. like in the galaxy scale simulations of Chan et al. 2019 or other tallbox simulations like Rathjen et al. 2021; Kim et al. 2023). Because we do not complete the loop, the

physical process we highlight of vertical magnetic field lines being cosmic ray escape highways is just a physical explanation of decreased γ -ray luminosity. The overall effect on star formation rate of magnetic field lines connecting the disk and halo is still unclear.

However, we suspect the true picture of cosmic ray feedback on star formation rates will require the increased magnetic field resolution. For example, reconsider Figure 6.6. In cases where the Parker instability is a dominant dynamical process, the decorrelation of cosmic rays and gas density could significantly impact star formation rates. However, that decorrelation was caused by field lines being vertical and not undergoing numerical reconnection. The vertical field lines also allowed cosmic rays to escape the midplane before they could interact with the cold gas, decreasing the γ -ray luminosity. So, the fundamental mechanism is the production of a vertical field, connecting our results to other setups which may not undergo the Parker instability but have a different way to produce vertical magnetic field.

6.6 Conclusion

The simulations we have run, `CRInj` and `CRBkg`, have illustrated the importance of cosmic-ray injections in driving outflows and structuring the ISM. Here, we list some key results from this study of cosmic-ray feedback:

1. The Parker instability drives overturning of gas which turns our initial hydrostatic equilibrium into a two-scale-height vertical gas density structure, in agreement with observations of the Milky Way. This overturning and satu-

ration happens on a shorter timescale than previously predicted: $\sim 150\text{Myr}$ instead of $\sim 500\text{Myr}$ (Heintz et al., 2020)). This faster rate is the result of supernovae injection.

2. This overturning of gas combined with thermal instability drives the production of thermally unstable gas above the midplane, which then sinks to $z \sim 0$ while cooling. This process occurs between outflowing plumes driven, in part, by cosmic-ray pressure gradients. Our simulations show this process decorrelates cosmic-ray pressure and gas density.
3. During the nonlinear saturation of the Parker instability, a more vertically oriented magnetic field forms just outside the midplane. This magnetic field structure allows for more efficient cosmic-ray escape alongside significant cosmic-ray feedback and a decrease in the calorimetric fraction F_{cal} .
4. The same vertical magnetic field structure also allows for a cosmic ray pressure gradient driven outflow to survive for a long time period ($> 100\text{Myr}$). This outflow could form the base of a large scale galactic wind, stretching down to $|z| \sim 500\text{pc}$.
5. The outflow is amplified in both the Parker instability regime and in the steady state that forms afterwards by the individual cosmic-ray injections included in the CRInj simulation.
6. We found that resolution was a significant factor in decreasing the calorimetric fraction F_{cal} . Numerical reconnection of magnetic field lines disconnects the

midplane from the high values of $|z|$, trapping cosmic rays in the midplane. Without high resolution or accurate subgrid modeling of vertical magnetic field, simulations will overpredict γ -ray luminosity of galaxies.

7 COSMIC-RAY BUOYANCY

Then, I felt a sense of mystery. I had shed light on a small corner of nature. Other scientists had illuminated larger corners. But there were almost certainly vast chambers and ballrooms that remained in the dark.

— *A Sense of the Mysterious* (2005), Alan Lightman

This chapter is a modified copy of a paper to be submitted to the *Astrophysical Journal*, titled *A Shift in Cosmic-Ray Transport Produces Regions of Instability in Galactic Atmospheres, Driving Mixing Between the Disk and Halo*. No text has been changed; the figures and equations have just been reformatted to accommodate the submission of this thesis document.

I deeply appreciate the contributions of my co-authors, Dr. Ellen Zweibel and Dr. David Hosking, to this work.

7.1 Abstract

Cosmic-ray transport changes as a function of thermal plasma properties. In particular, recent work has identified a shift in transport between the galactic disk and the galactic halo. Previously, the effect of this change in cosmic-ray transport has not been included in studies of the buoyant instability of gas in galaxies. In this work, we examine the stability of galactic atmospheres (the disk and halo) that exhibit a shift in cosmic-ray transport. We show that the combination of (1) a cosmic-ray pressure scale height larger than the gas pressure scale height and (2) a shift from

diffusive to streaming/advective transport leads to a region of buoyant instability, facilitating mixing between the disk and halo. We then examine the cases where this region does not appear, and study the nonlinear stability of those atmospheres. We find these are metastable, prone to restructuring of the atmosphere if perturbations are large enough (e.g., supernovae). In either the linearly unstable or nonlinearly unstable cases, a shift in cosmic-ray transport facilitates the mixing of gas between the galactic disk and the galactic halo.

7.2 Introduction

The stability of a fluid to convective motion is a cornerstone of astrophysical dynamics. The most common applications are in stellar structure and planetary atmospheres (Schwarzschild, 1958; Goody & Yung, 1989). However, galaxies also exhibit stratification of gas. Understanding the inflows to and outflows from galaxies is necessary in modeling their evolution and growth. Whether a galaxy's structure is buoyantly stable or unstable will play a part in the production of inflows and outflows.

In recent years, observations have measured the scale height of the different pressure components in galaxies, including the non-thermal components like magnetic and cosmic-ray pressure. Specifically, the non-thermal components of the interstellar medium (ISM) have larger scale heights than the gas (Ferrière, 2001; Kalberla & Kerp, 2009; Krause et al., 2018; McClure-Griffiths et al., 2023; Heesen et al., 2025). Cosmic rays contribute to the stratification, likely becoming more dynami-

cally important in the galactic halo and circumgalactic medium (CGM) (Ruszkowski & Pfrommer, 2023; Owen et al., 2023). This ordering of scale heights remains true when comparing to galaxies other than the Milky Way (Heesen et al., 2025).

The effect of cosmic rays on thermal gas is dependent on how they are transported. In the disk, it is expected that low-energy (\sim GeV) cosmic rays take a diffusive random walk, with a diffusion coefficient $\kappa \sim 10^{28} \text{cm}^2 \text{s}^{-1}$ (Evoli et al., 2019). Although our understanding of this process is changing (Kempski et al., 2023a; Lemoine, 2023), the transport must still be rapid compared to the gas dynamics in the disk (Armiglotta et al., 2024; Habegger et al., 2024; Habegger & Zweibel, 2025). Outside of the galactic midplane, the cosmic rays are transported via advection with the thermal plasma, or they stream at the local Alfvén speed (Armiglotta et al., 2024; Thomas et al., 2025; Habegger & Zweibel, 2025). This shift is likely correlated with a change in the orientation of the magnetic field (Habegger & Zweibel, 2025), which becomes more vertical outside the galactic midplane (Stein et al., 2025). The change in transport and magnetic field orientation is illustrated in Figure 7.1. The plasma in the galactic halo and the plasma in the galactic disk, while connected to each other, have fundamentally different properties.

The different cosmic-ray transport mechanisms affect the compressibility of the cosmic rays as a fluid (Zweibel, 2017). If they rapidly diffuse, then compressive motions do little to adjust the cosmic ray pressure. If they are locked to the gas, cosmic-ray pressure gradients develop during compression, requiring more total energy to compress the thermal and cosmic-ray fluids. We parameterize this change in cosmic-ray transport with an effective polytropic index $\gamma_{c,\text{eff}}$. Through the changing

polytropic index, cosmic-ray transport shifts the buoyant stability of gas between the galactic disk and the galactic halo.

In this paper, we examine this shift in stability. In general, gas with diffusive cosmic rays is more compressible and therefore more unstable to buoyant motion. Gas with advective or streaming cosmic rays is harder to compress and therefore stable to buoyancy. Since the shift in transport occurs when moving between low-density and high-density gas, the transport becomes particularly important in the transition from the galactic disk (high density) to the galactic halo (low density). In general, the cosmic ray and gas system in the disk will be more buoyant than gas in the galactic halo because of this change.

In Section 7.3, we apply the work of Hosking et al. (2025) to the galactic disk & halo system. We present the derivation in a way that we hope encourages readers to apply the general buoyancy force analysis in Hosking et al. (2025) to other astrophysical systems. In Section 7.4, we derive a simplified model for the compressibility of the ISM and the galactic halo, but leave details in the appendix. In Section 7.5, we use the compressibility model to study the buoyant stability of the galactic disk and halo, examining both local stability of a gas parcel and the overall stability of an atmosphere. In Section 7.6, we discuss how our results affect the inflow and outflow of gas from galaxies, as well as how the stability analysis could explain the existence of cold gas at high latitudes. Finally, in Section 7.7 we summarize the work and highlight our key conclusions.




Figure 7.1: The galactic disk and the halo have different plasma properties. In the disk, the magnetic field is perpendicular to the gravitational acceleration, increasing the importance of magnetic pressure. Additionally, cosmic rays diffuse rapidly through and out of the disk. In the halo, the magnetic pressure is less important for vertical stratification since the magnetic field is parallel to the acceleration vector. The cosmic rays are also more stuck to the gas, and bulk transport is at a linear combination of the local flow speed and the local Alfvén speed.

7.3 Buoyant Stability

There are many ways to examine the linear stability of a fluid dynamical system. One could perform a linear perturbation analysis, or use the energy principle (Bernstein et al., 1958) if dissipative effects are unimportant, but here we simply calculate the force due to buoyancy (Hosking et al., 2025). Hosking et al. (2025) examined 2D MHD stratification, but their general buoyancy force calculation allows an extension beyond the linear regime. Specifically, they study nonlinear evolution and metastable states of their chosen system. While all the methods agree in the linear perturbation limit for the restricted set of perturbations considered here, the calculation of the net buoyant force on a displaced fluid element extends to the nonlinear regime and identifies the physical importance of the compressibility of the fluid. The downside of net buoyancy force, however, is that it only produces a necessary, not a sufficient, condition for stability.

Hosking et al. (2025) define the compressibility as

$$\kappa := \left. \frac{\partial \ln \rho}{\partial \ln \mathcal{P}} \right|_Q, \quad (7.1)$$

where the total pressure $\mathcal{P} = \sum_i P_i$ includes all pressure components, ρ is just the fluid (or plasma) density, and \mathbf{Q} represents a state vector of quantities which are conserved in an ideal fluid, such as specific entropy or magnetic flux. For a non-magnetized, ideal gas with adiabatic index γ_g , the compressibility is just $\kappa = 1/\gamma_g$. For a larger γ_g , it takes more energy to compress the gas, thereby decreasing the compressibility.

The importance of the compressibility becomes most apparent when calculating the force density on a fluid parcel moving from $z_1 \rightarrow z_2$. Using the notation from Hosking et al. (2025), the net force density is

$$\mathbf{f} = g\hat{z} (\rho(\mathcal{P}_2, \mathbf{Q}_2) - \rho(\mathcal{P}_2, \mathbf{Q}_1)) \quad (7.2)$$

where $g > 0$ is the strength of the downward gravitational acceleration (i.e. $\mathbf{g} = -g\hat{z}$), and the subscripts denote whether the given quantity is evaluated at z_1 or z_2 , e.g. $\mathbf{Q}_1 = \mathbf{Q}(z_1)$. In writing Equation 7.2, it is implicitly assumed that the displaced parcel is in pressure equilibrium with its surroundings. Additionally, so that buoyancy is the only net force acting on the parcel, we assume hydrostatic equilibrium:

$$\rho g = -\frac{\partial \mathcal{P}}{\partial z}. \quad (7.3)$$

Because of our assumed direction of gravitational acceleration (downward), we only look at atmospheres with a total pressure decreasing with height.

Following a similar procedure to Hosking et al. (2025), we can differentiate Equation 7.2 with respect to z_2 and use the fundamental theorem of calculus to get the

force density as

$$\mathbf{f} = g\hat{z}\rho(P_2, \mathbf{Q}_1) \int_{z_1}^{z_2} dz \frac{\partial}{\partial z} \left(\frac{\rho(P, \mathbf{Q})}{\rho(P, \mathbf{Q}_1)} \right). \quad (7.4)$$

After some algebra and application of the chain rule (with ρ treated as a function of P and Q), Equation 7.4 becomes

$$\mathbf{f} = g\hat{z}\rho(P_2, \mathbf{Q}_1) \int_{z_1}^{z_2} dz \left[\frac{\rho(P, \mathbf{Q})}{\rho(P, \mathbf{Q}_1)} \left(-\frac{\partial \ln \mathcal{P}}{\partial z} \right) \times (\epsilon + \kappa(P, \mathbf{Q}_1) - \kappa(P, \mathbf{Q})) \right]. \quad (7.5)$$

In Equation 7.5, we introduce a quantity we call the instability parameter. It is equivalent to the *negative* of the same variable used in Hosking et al. (2025) for states at marginal stability ($\epsilon \sim 0$). It is defined as

$$\epsilon := - \left(\frac{\partial \ln \mathcal{P}}{\partial z} \right)^{-1} \frac{\partial \ln \rho}{\partial \mathbf{Q}} \Big|_P \cdot \frac{\partial \mathbf{Q}}{\partial z} = \kappa - \frac{\partial \ln \rho}{\partial z} \left(\frac{\partial \ln \mathcal{P}}{\partial z} \right)^{-1} \quad (7.6)$$

In the first expression, it is clear that the instability parameter isolates the effect of vertical changes in the conserved quantities \mathbf{Q} . The benefit of the second expression is that it isolates the fluid properties related to compression (in κ) from the dependence on atmospheric structure (contained in the ratio of total pressure and gas density scale heights).

Until this point, the force calculation has made no assumption about z_1 and z_2 except that $z_2 > z_1$. To get the linear stability criterion, we just take the limit $z_2 \rightarrow z_1$ of Equation 7.5. In this limit, we find

$$\mathbf{f} \rightarrow \hat{z}(\delta z) \rho g \left(-\frac{\partial \ln \mathcal{P}}{\partial z} \right) \epsilon. \quad (7.7)$$

Equation 7.7 highlights why we introduced the quantity ϵ and named it the instability parameter. As mentioned previously, the total pressure with height is negative, so the quantity in parentheses is positive ($-\partial \ln \mathcal{P} / \partial z > 0$). If $\epsilon > 0$, a fluid parcel will be linearly unstable to perturbations because the force will be in the same direction as the perturbation δz . If $\epsilon < 0$, then the force will be in the opposite direction of the perturbation δz , returning the parcel to its starting location. If $\epsilon = 0$, the parcel will tend to stay at its perturbed position (marginal stability).

Although ϵ is introduced through a mathematical derivation, it has a physical interpretation. As we remarked below Equation 7.1, $\kappa = \gamma_g^{-1}$ for an adiabatic gas. If we define a “structural” $\gamma_s \equiv (\partial \ln P / \partial z) / (\partial \ln \rho / \partial z)$, then $\epsilon = 1/\gamma - 1/\gamma_s$, and the system is unstable if $\gamma_s > \gamma$, i.e. if the stratification is super-adiabatic. This result is known in stellar structure theory as the Schwarzschild criterion for convection (Schwarzschild, 1958). So, fundamentally, ϵ directly encodes the traditional linear instability criterion into this buoyant force analysis.

7.3.1 Marginal Stability

If the atmosphere is marginally stable in some region (e.g. $\epsilon(z) = 0$ for $z_1 < z < z_2$), then the factor determining stability in that region is the change in compressibility.

Equation 7.5 becomes:

$$\mathbf{f} = g \hat{z} \rho(P_2, \mathbf{Q}_1) \int_{z_1}^{z_2} dz \left[\frac{\rho(P, \mathbf{Q})}{\rho(P, \mathbf{Q}_1)} \left(-\frac{\partial \ln \mathcal{P}}{\partial z} \right) \times (\kappa(P, \mathbf{Q}_1) - \kappa(P, \mathbf{Q})) \right]. \quad (7.8)$$

This case leads to an important conclusion from Hosking et al. (2025): if a fluid

parcel with compressibility κ_1 is moved through a marginally stable atmosphere, and the surrounding fluid has a lower compressibility $\kappa < \kappa_1$, then the parcel will experience a destabilizing buoyancy force. This treatment allows for the consideration of nonlinear stability in the case of marginal stability.

For studying the nonlinear stability of marginally stable atmospheres, it is more convenient to make ϵ an adjustable value. Hosking et al. (2025) do this by specifying only the profiles of the ratio of the conserved quantities which make up the vector \mathbf{Q} in the case of 2D-MHD. The profiles of each conserved quantity (and pressure) can then be calculated based on a given value of ϵ .

We use a slightly different formulation, assuming the profiles of the pressure ratios $P_i/P_0 := X_i = \beta_i^{-1}$ to a reference component P_0 (e.g., the gas pressure). This setup is more convenient and applicable to astrophysical systems where various observables can be tied to the pressure components. We still transform to the mass coordinate m defined as (Hosking et al., 2025)

$$m(z) = \int_z^\infty dz' \rho(z'). \quad (7.9)$$

This definition implies $\partial m / \partial z = -\rho(z)$ and turns the hydrostatic equilibrium condition into $\partial \mathcal{P} / \partial m = g$. Therefore $\mathcal{P} = mg$. If all the pressure ratio profiles $X_i(m)$ are known, the reference pressure P_0 is automatically given:

$$P_0(m) = mg \left(1 + \sum_{i \neq 0} X_i(m) \right)^{-1} \quad (7.10)$$

Returning to the instability parameter in Equation 7.6, we can shuffle the equa-

tion around to calculate the density gradient as a function of m as well:

$$\frac{\partial \rho}{\partial m} = \frac{\rho}{m} (\kappa(\rho, \mathcal{P}, X_i) - \epsilon). \quad (7.11)$$

It is straightforward to integrate this function numerically for various values of ϵ .

Then, using Equation 7.9, we can calculate the value of z for each value of m :

$$z = - \int_{m(z=0)}^{m(z)} \frac{dm'}{\rho(m')}. \quad (7.12)$$

The benefit of this entire treatment and the density and pressure profiles as a function of mass coordinate is the ability to examine the metastability of a marginally stable setup with small ϵ . As done in Hosking et al. (2025), this requires calculating the enthalpy as a function of \mathcal{P} and the quantities conserved with fluid motion \mathbf{Q} . By permuting equal-sized mass bins, we can find the minimum energy state, which is likely to be reached when the system experiences a large enough perturbation.

7.4 Compressibility of the ISM

All that remains to apply the formalism in Hosking et al. (2025) to the galactic disk and halo is to calculate the compressibility of the ISM. From there, we can test different stratification models and identify interesting setups and results.

There are several pressure components in the total pressure \mathcal{P} : the thermal gas pressure, radiation pressure, magnetic pressure, turbulent pressure, and cosmic-ray pressure. For each pressure component, we define a generalized conserved quantity

Q_i such that

$$Q_i = \frac{P_i}{f_i(\rho)}. \quad (7.13)$$

In most cases, one can simply take $f \propto \rho^\gamma$. However, for a polytropic index which depends on density,

$$\gamma_{i,\text{eff}} := \left. \frac{\partial \ln P_i}{\partial \ln \rho} \right|_{Q_i}, \quad (7.14)$$

the relationship is more complex. Specifically, f_i is

$$f_i(\rho) \propto \exp \left[\int d \ln \rho (\gamma_{i,\text{eff}}(\rho)) \right]. \quad (7.15)$$

Equation 7.15 only simplifies to ρ^γ if $\gamma_{i,\text{eff}}$ is a constant with respect to density ρ . For any non-constant dependence, f_i is a more complex function.

Before specifying the f_i and the $\gamma_{i,\text{eff}}$ for each component, we can calculate the total compressibility (Equation 7.1 for an arbitrary number of components:

$$\kappa = \frac{\sum_i P_i}{\sum_i \gamma_{i,\text{eff}} P_i}. \quad (7.16)$$

This equation illustrates that the compressibility is the inverse of a pressure-weighted average of the polytropic indices of the components. So, components with a higher pressure will contribute more to the compressibility. While pressure gradients will determine the forcing, the value of the compressibility does depend on the actual value of the pressure. Therefore, to estimate the compressibility, one must estimate the total pressure and its distribution into the different components in the galactic disk and halo.

Radiation — While radiation from stars in a galaxy can create dust-driven and line-driven winds, the actual local pressure in the radiation field in the galactic disk and halo is low compared to the other pressures. Therefore, it provides a negligible contribution to the compressibility.

Magnetic Field — We set $\gamma_{B,\text{eff}} = 0$, which is consistent with the most unstable mode identified using the MHD energy principle (Newcomb, 1961) and corresponds to sliding along the background field with little vertical compression or bending. It is also consistent with observations of the ISM for gas with densities $n \lesssim 1\text{cm}^{-3}$, where the $B - \rho$ correlation is weak or nonexistent. (Crutcher et al., 2010). For more details, see Appendix 7.8.1.

Cosmic Rays — We consider three different values of $\gamma_{c,\text{eff}}$. In the limiting case of rapid diffusive cosmic-ray transport (diffusion time much shorter than the characteristic displacement time), the effective polytropic index is zero. For the other limiting case of advection with the gas, the index is $4/3$ (assuming displacement times much shorter than diffusion times). Finally, for the case of steady-state streaming transport, $\gamma_{c,\text{eff}}$ takes the value $2/3$ (Wiener et al., 2017a). In the galactic disk and in neutral gas, it is expected that diffusion is dominant, so $\gamma_{c,\text{eff}} = 0$ there. In the galactic halo and in ionized gas, the cosmic rays become more locked to the gas and are advected, or they stream at the Alfvén speed (Armillotta et al., 2024; Thomas et al., 2025; Habegger & Zweibel, 2025). Therefore, the galactic halo’s effective polytropic index is between $\gamma_{c,\text{eff}} = 2/3$ and $\gamma_{c,\text{eff}} = 4/3$. For more details, see Appendix 7.8.2.

Thermal Pressure — The effective polytropic index for the gas pressure can

take a variety of values. In the adiabatic case, it is $\gamma_{g,\text{eff}} = 5/3$. For isothermal gas, it is $\gamma_{g,\text{eff}} = 1$. In the ISM, where non-adiabatic heating and radiative cooling timescales are shorter than dynamical times, the value depends on the slope of the cooling curve $\alpha = \partial \ln \Lambda / \partial \ln T$. Specifically, one finds

$$\gamma_{g,\text{eff}} = 1 - \frac{1}{\alpha}. \quad (7.17)$$

However, for the gas to be thermally stable, we need $\alpha > 1$ (Field, 1965). Therefore, gas in thermal equilibrium between heating and radiative cooling, which is also stable against condensation, will have an effective polytropic gas index $0 < \gamma_{g,\text{eff}} \lesssim 1$. Appendix 7.8.3 derives the results for $\gamma_{g,\text{eff}}$ in more detail.

Turbulence — The turbulent pressure will have a compressibility of 5/3 and behave similarly to the gas pressure from a 1D plane parallel atmosphere perspective. In this work, we consider a variety of gas pressure values and vertical profiles, even varying the compressibility. These encompass the effects of turbulent pressure on the vertical stratification. A turbulent pressure greater than gas pressure effectively increases the gas pressure and keeps the adiabatic index at 5/3.

In Appendix 7.8.3 we also discuss the effective gas polytropic index in the case where cosmic-ray streaming transport is dominant, and the collisionless heating due to streaming and the acoustic instability (Zweibel, 2017) is the dominant heating process. That regime could be important in the galactic halo, as the photoelectric effect on dust decreases further from the disk. It could also be important at the edges of cold clouds (Bustard & Zweibel, 2021). However, we do not apply it in this current work because its primary impact is to further decrease $\gamma_{g,\text{eff}}$ for thermally

stable gas (see Appendix 7.8.3). Additionally, there are other buoyancy instabilities which could occur in this regime (Kempski et al., 2023b), separate from what we pursue in this work.

Total Compressibility — Altogether, the compressibility we calculate for the general galactic disk and galactic halo is

$$\kappa = \frac{P_g + P_c + P_B}{\gamma_{g,\text{eff}}P_g + \gamma_{c,\text{eff}}P_c} = \frac{1 + \beta_{\text{cr}}^{-1} + \beta^{-1}}{\gamma_{g,\text{eff}} + \gamma_{c,\text{eff}}\beta_{\text{cr}}^{-1}} \quad (7.18)$$

where we have introduced the plasma beta $\beta = P_g/P_B$ and the cosmic-ray beta $\beta_{\text{cr}} = P_g/P_c$. Note that the inverse of the cosmic-ray beta β_{cr}^{-1} is the cosmic-ray pressure to gas pressure fraction $X_{\text{cr}} = P_c/P_g$ commonly used in studies of the intracluster medium.

7.5 Results

We present the analysis of our proposed disk and halo model in this section, beginning by studying linear instability. In section 7.5.1, we begin by showing that the instability parameter reduces to the case of Parker (1966). In Section 7.5.2, we study the dependence of the compressibility (Equation 7.18) on β , β_c , $\gamma_{g,\text{eff}}$, and $\gamma_{c,\text{eff}}$, highlighting that the halo gas has a significantly lower compressibility than the disk gas. In Section 7.5.3, we identify parameter regimes (of cosmic-ray scale height ratio H_c/H_g and cosmic-ray pressure ratio $\beta_c^{-1} = X_c$) where the change in cosmic-ray transport and magnetic orientation between the galactic disk and halo will lead to a change in sign of the instability parameter. In Section 7.5.4, we study

the stability of a realistic galactic stratification, showing that a shift in cosmic-ray transport from diffusion to advection produces a region of instability between the disk and halo regimes. Finally, in Section 7.5.5, we examine the case of marginal stability, where the cosmic-ray transport occurs too low in the galactic atmosphere to allow the instability region identified in Section 7.5.4 to occur. We find that even the marginal stable system is nonlinearly unstable to mixing between the disk and halo.

7.5.1 Linear (Parker) Instability

Using the compressibility from Section 7.4, we return to the instability parameter (Equation 7.6):

$$\epsilon = \frac{\mathcal{P}}{\gamma_{g,\text{eff}}P_g + \gamma_{c,\text{eff}}P_c} - \left(\frac{\partial \ln \mathcal{P}}{\partial z} \right)^{-2} \frac{1}{\mathcal{P}} \frac{\partial^2 \mathcal{P}}{\partial z^2}. \quad (7.19)$$

Evaluating Equation (7.19) when the pressure components all have the same scale height and follow an exponential profile $P_i \propto \exp(-|z|/H)$ gives a generalization of the Parker instability criterion (Parker, 1966; Zweibel, 2017). In that case, the instability criterion ($\epsilon > 0$) is

$$(1 - \gamma_{g,\text{eff}})P_g + (1 - \gamma_{c,\text{eff}})P_c + P_B > 0. \quad (7.20)$$

Parker (1966) assumed an isothermal gas $\gamma_{g,\text{eff}} = 1$ and diffusive cosmic rays $\gamma_{c,\text{eff}} = 0$, leading to the instability criterion

$$P_c + P_B > 0. \quad (7.21)$$

This implies the instability occurs for any non-zero combination of cosmic-ray pressure and magnetic pressure. Later work addressed the assumptions on $\gamma_{g,\text{eff}}$ and $\gamma_{c,\text{eff}}$ (Zweibel & Kulsrud, 1975; Heintz & Zweibel, 2018), illustrating that the instability would not be as common as suggested in Parker (1966).

However, we have found that things become more interesting when examining a change in the effective polytropic index of cosmic-ray pressure with height. As discussed in the introduction and depicted in Figure 7.1, the galactic disk and galactic halo have different cosmic-ray transport properties (likely due in part to the change in magnetic field orientation (Habegger & Zweibel, 2025)). When we incorporate this, the change in cosmic-ray transport to advection or streaming reduces the buoyancy of gas outside of the galactic disk.

7.5.2 Changes in cosmic-ray Transport Adjust the Compressibility

First, let's consider the compressibility in Equation 7.18. Figure 7.2 shows the dependence of the compressibility on $\beta_{\text{cr}}^{-1} = P_c/P_g$. The left plot of Figure 7.2 shows the compressibility with no magnetic pressure, and the right plot includes magnetic pressure with $\beta = 1$. Of course, the magnetic field cannot be exactly zero; if it were,




Figure 7.2: The compressibility κ of the stratified cosmic ray & thermal plasma system for various effective polytropic indices $\gamma_{g,\text{eff}}$ and $\gamma_{c,\text{eff}}$. The horizontal axis is $\log_{10}(\beta_{\text{cr}}^{-1}) = \log_{10}(P_c/P_g)$, so moving right moves the system to be cosmic-ray pressure dominated. The left plot shows the compressibility neglecting the contribution of magnetic fields. The right plot shows the case of $\beta = 1$. Solid lines show the adiabatic limit $\gamma_{g,\text{eff}} = 5/3$ and dashed lines show the isothermal/non-adiabatic limit $\gamma_{g,\text{eff}} = 1$. Black lines show diffusive cosmic-ray transport, blue lines show streaming transport, and red lines show advective transport. The lowest compressibility in both plots is achieved in the adiabatic ($\gamma_{g,\text{eff}} = 5/3$), advective ($\gamma_{c,\text{eff}} = 4/3$) case with no magnetic pressure $\beta^{-1} = 0$. This case matches our picture of the galactic halo gas. The compressibility of the galactic disk (where cosmic rays are diffusive and magnetic pressure supports the gas against gravity) is much higher, matching the black solid line in the panel on the right.

cosmic rays would be decoupled from the thermal gas and not confined to the Galaxy, so $\beta^{-1} = 0$ should be understood as a limiting case of a field purely parallel to the gravitational acceleration direction. The primary impact of the magnetic field is to increase compressibility, because it adds pressure with no corresponding increase in density (see Appendix 7.8.1).

The solid lines on each plot assume the adiabatic limit with $\gamma_{g,\text{eff}} = 5/3$. The dashed lines show an isothermal limit $\gamma_{g,\text{eff}} = 1$, which is similar to the non-adiabatic limit where heating and radiative cooling are dominant (see Appendix 7.8.3). Holding

other parameters constant, the non-adiabatic regime is more compressible.

Finally, the color of the lines shows the different cosmic-ray transport assumptions. Black lines show diffusive transport, blue lines show streaming transport, and red lines show advective transport. The primary takeaway from this plot is that increasing $\gamma_{c,\text{eff}}$, making the cosmic rays more locked to the gas, drives the system to have a lower compressibility. This result is expected and obvious: a lower compressibility means it takes more energy to compress a chunk of gas and cosmic rays.

The lowest compressibility in both plots is achieved in the adiabatic ($\gamma_{g,\text{eff}} = 5/3$), advective ($\gamma_{c,\text{eff}} = 4/3$) case with no magnetic pressure $\beta^{-1} = 0$. This case matches our picture of the galactic halo gas in Figure 7.1, assuming the cooling time is large, and the non-adiabatic effects are negligible.

The case of advective transport is in direct contrast to diffusive transport, shown with black lines. Looking at the black lines in the right plot of Figure 7.2, we see that the compressibility of the galactic disk (where cosmic rays are diffusive and magnetic pressure supports the gas against gravity) is much higher.

Because the instability parameter ϵ is proportional to the compressibility κ , the disk gas is more likely to be unstable for a given atmosphere. Additionally, in the metastable case, a higher compressibility gas moving through a lower compressibility gas will be buoyant (Hosking et al., 2025). That scenario appears when gas from the galactic disk is perturbed upward into the galactic halo without changing its intrinsic properties other than maintaining pressure balance with its surroundings.

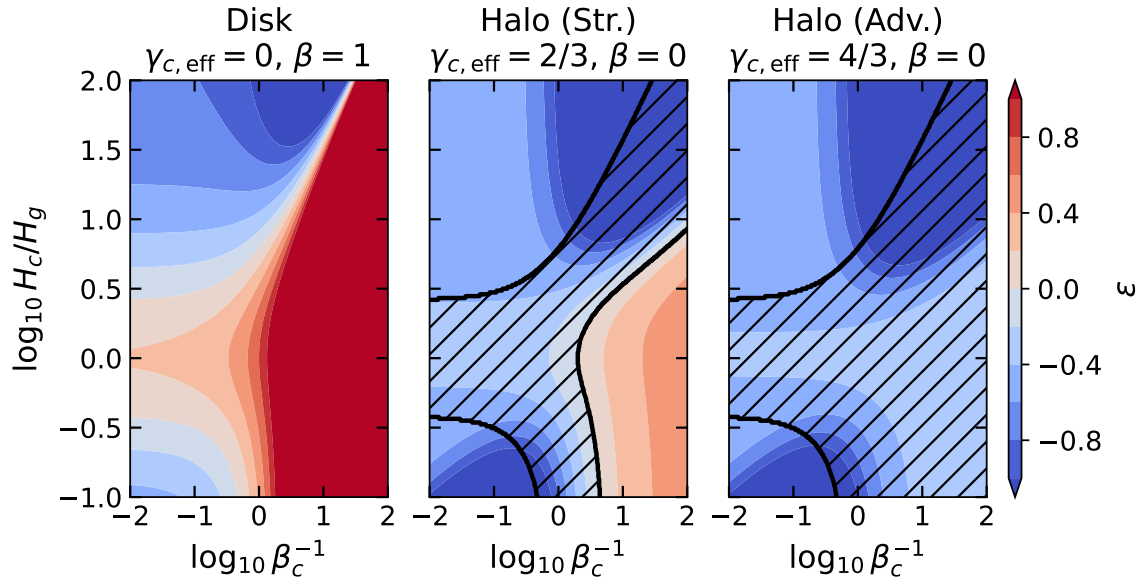


Figure 7.3: Contours of instability parameter in the $(H_c/H_g, \beta_c^{-1})$ space for exponential scale heights (Equation 7.22). Red contours correspond to instability ($\epsilon > 0$) and blue contours correspond to stability ($\epsilon < 0$). From left to right, we present models for (1) the galactic disk, (2) the galactic halo with streaming transport, and (3) the galactic halo with advective transport. In the middle and right plots, we also identify parameter regimes where the sign of ϵ changes from positive to negative in comparison to the galactic disk case (left plot) with black contours and hatching.

7.5.3 Structures which Change Stability due to a Shift in cosmic-ray Transport

Most observations of our own galaxy and others point towards another problem in the derivation by Parker (1966): the assumption that the pressure components all have the same scale height. If we allow the cosmic-ray pressure and magnetic pressure to have different scale heights than the gas, then the instability parameter is

$$\epsilon = \frac{1 + \beta_c^{-1} + \beta^{-1}}{\gamma_g + \gamma_c \beta_c^{-1}} - \frac{\left(1 + \frac{H_g^2}{H_c^2} \beta_c^{-1} + \frac{H_g^2}{H_B^2} \beta^{-1}\right) (1 + \beta_c^{-1} + \beta^{-1})}{\left(1 + \frac{H_g}{H_c} \beta_c^{-1} + \frac{H_g}{H_B} \beta^{-1}\right)^2} \quad (7.22)$$

The above leaves six undetermined parameters: the pressure ratios β^{-1} and β_c^{-1} , the scale height ratios H_g/H_B and H_g/H_c , and the polytropic indices γ_g and γ_c .

Instead of studying the full 6D parameter space, we can make simplifying assumptions. First, we assume the cosmic-ray pressure scale height and the magnetic scale height are equal to one another ($H_c = H_B$). This assumption agrees with observations of the far-infrared-radio (FIR) correlation, which suggest cosmic-ray pressure and the magnetic pressure decay at a similar rate when moving away from a galaxy. Next, we set $\gamma_g = 5/3$. Finally, we chose γ_c and β based on our simplified models of the disk and halo in Figure 7.1. These assumptions leave two undetermined variables: the cosmic-ray pressure (β_c^{-1}) and cosmic-ray pressure scale height (H_c/H_g) normalized by the gas pressure and scale height, respectively.

In Figure 7.3, we show this parameter space for the case of the disk (horizontal magnetic field with $\beta = 1$, diffusive cosmic-ray transport $\gamma_{c,\text{eff}} = 0$), and two models of the halo (vertical magnetic field with $\beta^{-1} = 0$): one with streaming transport ($\gamma_{c,\text{eff}} = 2/3$) and one with advective transport ($\gamma_{c,\text{eff}} = 4/3$). Consistent with Parker's original analysis (Parker, 1966), diffusive transport (left panel) is unstable ($\epsilon > 0$) for many values of cosmic-ray pressure and cosmic-ray pressure scale heights. In the halo, where the magnetic field is vertical, and cosmic rays are harder to compress, the region of instability shrinks in streaming transport (middle panel) and completely disappears in the case of advective transport (right panel).

In the middle and right panels of Figure 7.3, we also plot black hatched contours showing parameter values for which the sign of the instability parameter ϵ changes from positive to negative. For those values of β_c and H_c/H_g , the plasma would

switch between being buoyantly unstable and stable depending on whether it was in the galactic halo or the galactic disk. The region covers a large region of parameter space, in particular the region where cosmic-ray pressure is higher than the gas pressure — it is relatively independent of the cosmic-ray pressure scale height.

Therefore, if there is a vertical position, moving up from the galactic midplane, where cosmic-ray pressure becomes dominant over the gas pressure, then the system will be unstable. *Unless* a shift in cosmic-ray transport causes it to become buoyantly stable. In the long-term, one might expect any region of instability to drive the growth of the vertical component of the magnetic field (Habegger & Zweibel, 2025), possibly changing the cosmic-ray transport from diffusive to advective or streaming, making that region buoyant against stability. A steady-state would then result when cosmic-ray transport switches at the height where the system would be unstable to buoyancy, except for the switch in transport.

7.5.4 Stability of Vertical Structure

Staying with the simple exponential scale height pressure profiles, we now look at how the instability parameter changes with height for a galactic atmosphere. In Figure 7.4, we show a simple galactic atmosphere for the solar neighborhood. The top plot of Figure 7.4 shows the assumed pressure profiles, which have a midplane ($z = 0$) gas pressure $P_g/k_B = 2 \times 10^4 \text{K cm}^{-3}$, cosmic-ray pressure $P_c = 1 \text{eV cm}^{-3}$ (equivalent to $1.1 \times 10^4 \text{K cm}^{-3}$), and magnetic field strength $B = 2 \mu\text{G}$ (equivalent to $2.6 \times 10^3 \text{K cm}^{-3}$). We assume the cosmic-ray pressure and magnetic pressure have scale heights of $H_c = H_B = 1 \text{kpc}$, but the gas pressure has a scale height of $H_g = 0.3 \text{kpc}$.




Figure 7.4: Top plot: Vertical profiles of gas pressure, cosmic-ray, pressure, magnetic pressure, and total pressure for a simple exponential scale height model. Middle plot: Resulting density profile assuming hydrostatic equilibrium and a gravitational acceleration of $g = 2\pi G \cdot 50M_{\odot}\text{pc}^{-2} = 4.4 \cdot 10^{-9}\text{cm}^2\text{s}^{-1}$. Bottom plot: Instability parameter ϵ for different values of $\gamma_{c,\text{eff}}$. We also plot the instability parameter for a model which switches from diffusion to advection for cosmic-ray transport at a density of $n = 0.1\text{cm}^{-3}$.

We take a gravitational acceleration of $g = 2\pi G \cdot 50M_{\odot}\text{pc}^{-2} = 4.4 \cdot 10^{-9}\text{cm}^2\text{s}^{-1}$ to calculate the density profile in the middle plot of Figure 7.4. These estimates come from estimates of the neutral hydrogen (HI) scale height (McClure-Griffiths et al., 2023), common radio scale heights found in Milky-Way type galaxies (Krause et al., 2018), and commonly assumed parameters for the solar-neighborhood (McKee et al., 2015).

In the bottom plot of Figure 7.4, we show how the instability parameter ϵ varies as a function of height for different values of $\gamma_{c,\text{eff}}$. The diffusive case of $\gamma_{c,\text{eff}} = 0$ becomes unstable as soon as the cosmic-ray pressure becomes the dominant pressure component. The advective case of $\gamma_{c,\text{eff}} = 4/3$ remains stable for all shown values of z . The streaming case of $\gamma_{c,\text{eff}} = 2/3$ becomes unstable at a larger height of ~ 850 pc.

For the green dash-dotted line, we show how the instability parameter would change as a result of the vertical change in cosmic-ray transport. If cosmic-ray transport automatically switched at $n = 0.1 \text{ cm}^{-3}$, then gas which was buoyantly unstable would suddenly become buoyantly stable at a height of around 600 pc. Overall, for our simple model of a switch in cosmic-ray transport, we would expect a buoyant region (akin to convective regions in stellar structure) between $z = 200$ pc and $z = 600$ pc for these fiducial parameters.

While the size and extent of this region may change, the primary factors are (1) where the cosmic-ray pressure becomes dominant and (2) where the cosmic-ray transport shift occurs. If the transport becomes less diffusive at a higher vertical position than where the cosmic-ray pressure becomes dominant, then there will be a region of buoyant instability. If the switch is from diffusive to streaming transport, then there could be further instability at higher vertical positions where the case of $\gamma_{c,\text{eff}} = 2/3$ is unstable.

7.5.5 Metastability Analysis

Looking at Figure 7.4, it is clear that if the switch in cosmic-ray transport is below the height where cosmic-ray pressure becomes the dominant component, then the system




Figure 7.5: Top plot: Vertical profiles of gas pressure and cosmic-ray pressure in the marginally stable setup we study in Section 7.5.5. Middle plot: Density profile under different prescriptions for $\gamma_{c,\text{eff}}$ (see text for more details, and Figure 7.6 for the $\gamma_{c,\text{eff}}$ profiles). Bottom plot: The prescribed profiles for β_c^{-1} and ϵ .

would remain stable despite the switch in transport. The buoyant mixing region would disappear. Now, we examine that regime in the case of marginal stability.

Utilizing the framework detailed in Section 7.3.1, we set the cosmic-ray pressure to the gas pressure profile with respect to mass coordinate m , and the profile of ϵ . We then solve for $\rho(m)$. Note, we do not include magnetic pressure in this setup so that we can eventually solve for the minimum energy state in the same manner as

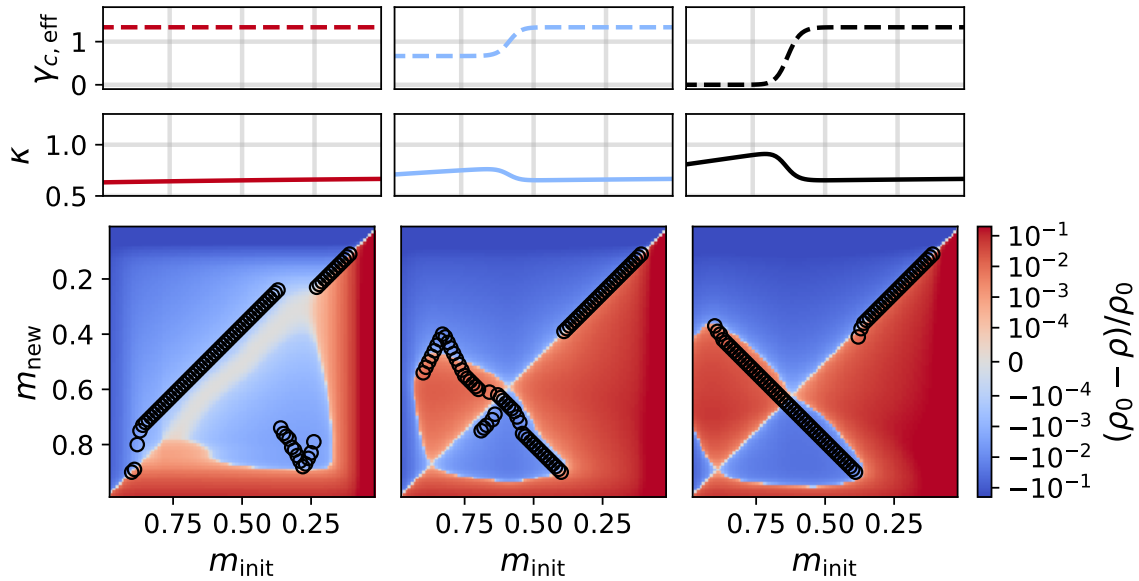


Figure 7.6: Net buoyancy force for a mass parcel moved in pressure equilibrium from m_{init} to m_{new} . These calculations are done for the pressure and density profiles shown in Figure 7.5. From left to right, we vary the value of $\gamma_{c,\text{eff}}$ in dense gas from $4/3$ (left column), to $2/3$ (middle column), to 0 (right column). In all cases, $\gamma_{c,\text{eff}} = 4/3$ in the lower density gas. In the second row, we also plot the compressibility κ to illustrate how the setup has more compressible fluid below less compressible fluid in the middle and right columns. Black circles in the bottom row of buoyancy force diagrams identify the minimum energy state across all possible orderings of the mass bins.

Hosking et al. (2025). With only two pressure components, solving for the buoyancy force assuming pressure equilibrium and the resulting specific enthalpy of the parcel at a given height is a well-defined problem. With many pressure components, solving for pressure balance is not possible in general. That step is necessary for solving the linear sum assignment problem, which shuffles equal mass bins to find the minimum energy state (Hosking et al., 2025).

We use the following profiles (note m is the integrated mass above a given height

z , so the $m = 1$ is the midplane and $m = 0$ is $z \rightarrow \infty$):

$$\beta_c^{-1}(m) = \frac{2}{3} \left(\frac{3}{2} - m \right) \quad (7.23)$$

$$\epsilon(m) = \tanh \left(\frac{m - 0.05}{0.05} \right) - \tanh \left(\frac{m - 0.95}{0.05} \right) - 2. \quad (7.24)$$

Therefore, $P_c \leq P_g$ for the entire atmosphere. So the region discussed in Section 7.5.4 would disappear.

Figure 7.5 shows the resulting atmosphere, in a similar format to Figure 7.4. In the middle plot, we show three curves for different profiles of $\gamma_{c,\text{eff}}(\rho)$. Specifically, we define the function f_c (Equation 7.15) as

$$f_c(\rho) = \exp \left\{ \frac{\gamma_1 + \gamma_2}{2} \log \left(\frac{\rho}{\rho_s} \right) + (\gamma_2 - \gamma_1) \Delta \log \left(\cosh \left[\frac{1}{\Delta} \log \left(\frac{\rho}{\rho_s} \right) \right] \right) \right\}. \quad (7.25)$$

This formula corresponds to a broken power law which switches from γ_1 to γ_2 at a density ρ_s . The switch is in logarithmic space and follows a hyperbolic tanh function with width Δ . The resulting effective polytropic index is

$$\begin{aligned} \gamma_{c,\text{eff}} = \gamma_1 + (\gamma_2 - \gamma_1) \frac{1}{2} \left(1 + \tanh \left[\frac{1}{\Delta} \log \left(\frac{\rho}{\rho_s} \right) \right] \right) \\ + \log(\rho) \frac{\gamma_2 - \gamma_1}{2} \Delta \operatorname{sech}^2 \left[\frac{1}{\Delta} \log \left(\frac{\rho}{\rho_s} \right) \right]. \end{aligned} \quad (7.26)$$

This formulation matches commonly used broken-power law setups, except it uses a hyperbolic tangent switching function and not a sigmoid function.

In the middle plot of Figure 7.5, we show the density profiles for (1) a constant

$\gamma_{c,\text{eff}}$, (2) a switch from $2/3$ to $4/3$, and (3) a switch from 0 to $4/3$. These profiles are shown with a red line, blue line, and black line, respectively, similar to Figure 7.2. All the profiles switch at $\rho_s = 0.5$, which is shown as a gray dashed line. The second and third profiles switch cosmic-ray transport at around a mass coordinate of 0.6 .

With the profiles set, we can solve for the change in density when parcels moved from position m_{init} to a new position m_{new} in pressure equilibrium with their surroundings, conserving $Q_c = P_c/f_c(\rho)$ and $Q_g = P_g/\rho^{5/3}$ during the movement. These changes in density are shown in the bottom three panels of Figure 7.6, with blue corresponding to a net downward buoyancy force (higher displaced parcel density ρ than new surroundings ρ_0) and red corresponding to a net upward buoyancy force (lower displaced parcel density ρ than new surroundings ρ_0). Above these plots, we also show the prescribed $\gamma_{c,\text{eff}}$ profile and the resulting compressibility κ , both as a function m (using the density and pressure profiles in Figure 7.5).

In each net buoyancy force plot in Figure 7.6, we also show the minimum energy state (the smallest total energy for any shuffling of the mass bins to different heights) with black circles. In the case of a constant $\gamma_{c,\text{eff}} = 4/3$, the minimum energy state involves the gas with a higher cosmic-ray pressure to gas ratio sinking, because that gas has a higher compressibility than the gas below it. However, for the case of switching cosmic-ray transport from either streaming (middle column of Figure 7.6) or diffusion (right column of Figure 7.6) to advective with $\gamma_{c,\text{eff}} = 4/3$, the system is unstable instead unstable in both directions. The plasma below the transport switch has a higher compressibility and will move upward if perturbed enough. The gas right above the switch will move downward because of its lower compressibility.

Therefore, give enough time and energy injection to drive the plasma up and down from its initial equilibrium, and the end state will be a movement of gas from the disk to the halo, and vice versa. Just as the buoyant region in Figure 7.4 facilitates mixing between the disk and halo, a marginally stable system will also tend to flip the two regimes across the region where cosmic-ray transport changes.

7.6 Discussion

The analysis in Section 7.5 illustrates that the proposed setup in Figure 7.1 is inherently predisposed to linear or nonlinear instability. The plasma in the galactic disk is more compressible because of its magnetic field perpendicular to the gravitational acceleration and because of diffusive cosmic rays. In contrast, the galactic halo is much more compressible because the magnetic field is more aligned with the gravitational acceleration and the cosmic rays are more locked to the fluid, with a larger effective polytropic index $\gamma_{c,\text{eff}}$. Therefore, our modern understanding and estimates of the non-thermal components of the ISM suggest galaxies are predisposed to buoyant instability.

Our work focused primarily on the switch in cosmic-ray transport, because the switch in magnetic field orientation has a minimal effect unless the magnetic field is strong ($\beta^{-1} \gtrsim 1$). The cosmic-ray pressure, on the other hand, is similar to the thermal and turbulent pressures in the galactic disk, but has a larger scale height. Therefore, there will always be a vertical position such that cosmic-ray pressure is dominant and the cosmic-ray transport determines the compressibility of the back-

ground plasma. Where the cosmic-ray pressure becomes dominant ($\beta_c^{-1} \geq 1$), and the cosmic-ray transport at that vertical position will determine the stability of the stratification.

Our results also have implications for the multiphase structure of the ISM. Cosmic-ray transport does not depend directly on vertical position — it depends on the properties of the background plasma, in particular the plasma density ρ (Armillotta et al., 2024; Thomas et al., 2025; Habegger & Zweibel, 2025). Therefore, denser gas (i.e., the cold neutral medium (CNM) will be easy to compress since cosmic rays can move through it rapidly. As shown by Hosking et al. (2025), a higher compressibility gas will undergo buoyant motion when embedded in a lower compressibility gas. Therefore, this switch in cosmic-ray transport formalism could be an explanation for the production of neutral atomic hydrogen at high latitudes (McClure-Griffiths et al., 2023; Wenger et al., 2024; Rybarczyk et al., 2024). Since cold gas is more prone to buoyant motion, any small perturbation will grow over time (Equation 7.7). We plan to test this hypothesis using more detailed analytic theory and simulations in the future, but it is outside the scope of this current paper, which is focused on the structural impact of a shift in cosmic-ray transport.

This work simplifies much of the dynamics of the ISM plasma down to polytropic laws for gas, cosmic rays, and magnetic fields. Our choices of these laws are consistent with observations of the ISM and with simulations that include complex cosmic-ray transport. Allowing departures from our selected polytropic laws will certainly have an impact, but it is unlikely to completely erase the impact on buoyancy when examining multiple phases of gas with disparate gas densities.

Another important caveat is our assumption of hydrostatic equilibrium. While the galactic halo and the disk are turbulent, dynamically evolving structures, studying the hydrostatic equilibrium is still useful. In particular, studying the compressibility, which is a fundamental property of the plasma, is a useful endeavor because it still quantifies how much energy is required to do work on the gas. The effective polytropic indices that appear in the compressibility determine how much kinetic energy becomes thermal energy. Therefore, the highly compressible neutral gas with diffusive cosmic rays will be easier to push and pull on. This gas will therefore be more affected by any forces generated by gravity, supernovae, or relativistic jets.

7.7 Summary & Conclusions

Overall, our examination here identifies how non-thermal components like cosmic rays and magnetic fields create significant changes in the buoyant stability of gas in the ISM and galactic halo. In particular, the polytropic index of cosmic-ray pressure depends on cosmic-ray transport. As cosmic-ray transport changes between ionized and neutral gas, it increases the compressibility of the bulk medium. This change influences the buoyancy of the system, leading to a counterintuitive effect: cold, dense gas is buoyantly unstable, but hot gas is buoyantly stable. When examining vertical stratification, the galactic disk gas is buoyantly unstable, and the galactic halo gas is buoyantly stable. This setup predisposes the halo and disk to mixing, even when the setup is forced to be marginally stable. For realistic vertical profiles of gas pressure, cosmic-ray pressure, and magnetic pressure in the solar neighborhood,

we identify there could be a region of buoyant instability and mixing between the galactic disk and galactic halo, ranging from $z \sim 200$ pc up to $z \sim 600$ pc.

Acknowledgements

The authors would like to thank Sophie Aerdker, Karol Fulat, Ka Wai Ho, Mohan Richter-Addo, Mateusz Ruszkowski, Aaron Tran, Bindesh Tripathi, and Ka Ho Yuen for their insight and helpful conversations, which improved this work. The authors greatly appreciate funding from NASA FINESST grant No. 80NSSC22K1749, and NSF grants AST-2007323 and PHY-2409224, which supported this work.

Software: Matplotlib (Hunter, 2007), NumPy (Van Der Walt et al., 2011; Harris et al., 2020), AstroPy (Astropy Collaboration et al., 2013, 2018), SciPy (Virtanen et al., 2020)

7.8 Appendix

7.8.1 Effective Polytropic Index for Magnetic Forces

Polytropic relationships for the magnetic field strength are often used to study the compression of a plasma. In studies of the ISM and molecular clouds, there is even the famous $B - \rho$ relationship (Crutcher et al., 2010; Seta & McClure-Griffiths, 2025) which is further justified by basic theories of compression of magnetic fields (Tritsis et al., 2015).

Observations in the ISM show a steepening relationship between B and gas num-

ber density n with increasing n . Assuming a power law $B \propto n^\alpha$, observational estimates are $\alpha \sim 0 - 0.25$ for gas with density $n \lesssim 1\text{cm}^{-3}$ and $\alpha \sim 0.5$ for higher densities (Crutcher et al., 2010; Whitworth et al., 2024; Seta & McClure-Griffiths, 2025).

To match our definition of the effective polytropic index across each pressure component (Equation 7.13), we instead calculate $\gamma_{B,\text{eff}}$ defined by

$$Q_B = \frac{P_B}{\rho^{\gamma_{B,\text{eff}}}} = \frac{B^2}{8\pi\rho^{\gamma_{B,\text{eff}}}}. \quad (7.27)$$

As a result, the effective polytropic index we use is double the value of the commonly used power law slope:

$$\gamma_{B,\text{eff}} = 2\alpha \quad (\text{for } B \propto n^\alpha). \quad (7.28)$$

The values are self-consistent and result simply from the magnetic pressure being the square of the magnetic field strength.

Since we are not considering the stratification of a molecular cloud and are instead focused on the disk-halo interface, it would be reasonable to focus on the low-density limit and take $\gamma_{B,\text{eff}}$ to be between 0 and 0.5.

It is important to note that this choice of $\gamma_{B,\text{eff}}$ differs from the 2D MHD system studied in Hosking et al. (2025). In the case of 2D motion in the plane perpendicular to the mean magnetic field, one expects $B \propto \rho$, or $\gamma_{B,\text{eff}} = 2$. Naively, that setup is similar to the picture we illustrated for the galactic disk in Figure 7.1. However, we do not expect the motions in the galactic disk to be strictly perpendicular to the magnetic field.

Instead, applying the MHD energy principle (Bernstein et al., 1958) to the same setup we illustrate in the galactic disk, Newcomb (1961) found a different result. Taking the proper limit for long wavelengths parallel to the magnetic field, the instability criterion is consistent with $\gamma_{B,\text{eff}} = 0$. This result appears when accounting for quasi-interchanges of magnetic flux tubes (Newcomb, 1961). In a sense, there is always an unstable mode corresponding to long wavelength flux tube interchanges, and the system will find that mode.

After considering all of the above, we take $\gamma_{B,\text{eff}} = 0$ in our models of the disk and halo. This assumption means that any increase in magnetic field strength will put a system into a more unstable regime, because the magnetic field is infinitely compressible.

While we set $\gamma_{B,\text{eff}}$ to be the same in the galactic disk and halo, the orientation of the field does still change. In the galactic halo, since the magnetic field is more parallel to the gravitational acceleration, the magnetic pressure plays less of a role. We choose to remove it altogether in the galactic halo. As a result, increased magnetic field strength in the halo does not impact the buoyant stability of the fluid. However, it can still impact cosmic-ray transport by increasing the Alfvén speed in the galactic halo.

7.8.2 Effective Polytropic Index for Cosmic-Ray Pressure

To find an effective polytropic index for cosmic-ray pressure, we examine the evolution of cosmic-ray pressure from (Zweibel, 2017):

$$\begin{aligned} \frac{d \ln P_c}{dt} = & -\gamma_c \nabla \cdot \mathbf{u} - (\mathbf{w} - \mathbf{u}) \cdot \nabla \ln P_c - \gamma_c \nabla \cdot (\mathbf{w} - \mathbf{u}) \\ & + \nabla \cdot (\boldsymbol{\kappa} \cdot \nabla \ln P_c) + (\nabla \ln P_c) \cdot \boldsymbol{\kappa} \cdot \nabla \ln P_c, \end{aligned} \quad (7.29)$$

where $\gamma_c = 4/3$ is the adiabatic index of the cosmic rays, \mathbf{u} is the local fluid velocity, \mathbf{w} is the wave frame of the cosmic rays, and $\boldsymbol{\kappa}$ is the diffusion tensor for the cosmic rays (not to confused with the unbolded κ we use for compressibility).

Next, we use the assumption that the wave frame of the cosmic rays depends on the local Alfvén speed:

$$\mathbf{w} - \mathbf{u} = f_s \frac{\mathbf{B}}{\sqrt{4\pi\rho}} = f_s \mathbf{v}_A. \quad (7.30)$$

The streaming terms then become

$$\begin{aligned} \frac{d \ln P_c}{dt} = & -\gamma_c \nabla \cdot \mathbf{u} - f_s \mathbf{v}_A \cdot \nabla \ln P_c + \frac{\gamma_c}{2} f_s \mathbf{v}_A \cdot \nabla \ln \rho - \gamma_c \mathbf{v}_A \cdot \nabla f_s \\ & + \nabla \cdot (\boldsymbol{\kappa} \cdot \nabla \ln P_c) + (\nabla \ln P_c) \cdot \boldsymbol{\kappa} \cdot \nabla \ln P_c. \end{aligned} \quad (7.31)$$

In the case of weak diffusion and $f_s = 0$, the cosmic rays become locked to the gas:

$$\frac{d \ln P_c}{dt} = -\gamma_c \nabla \cdot \mathbf{u} = \gamma_c \frac{d \ln \rho}{dt} \quad (7.32)$$

This equation can be rearranged to find the quantity Q_c conserved with the fluid

motion:

$$Q_c = \frac{P_c}{\rho^{\gamma_{c,\text{eff}}}} = s^{\gamma_{c,\text{eff}}} \quad (7.33)$$

where $\gamma_{c,\text{eff}} = \gamma_c = 4/3$.

If diffusion is weak and the Alfvén speed is faster than the local flow, then the cosmic-ray pressure equation is

$$\frac{\partial \ln P_c}{\partial t} = -f_s \mathbf{v}_A \cdot \nabla \ln P_c + \frac{\gamma_c}{2} f_s \mathbf{v}_A \cdot \nabla \ln \rho - \gamma_c \mathbf{v}_A \cdot \nabla f_s. \quad (7.34)$$

On time scales longer than $\tau_A = (\mathbf{v}_A \cdot \nabla \ln P_c)^{-1}$, the cosmic-ray pressure will reach a steady state given by

$$f_s \mathbf{v}_A \cdot \nabla \ln P_c = \frac{\gamma_c}{2} f_s \mathbf{v}_A \cdot \nabla \ln \rho - \gamma_c \mathbf{v}_A \cdot \nabla f_s. \quad (7.35)$$

If f_s is a constant dependent on the direction of the local cosmic-ray pressure gradient, then the equation implies an effective adiabatic index for the cosmic-ray pressure given by $\gamma_{c,\text{eff}} = \gamma_c/2 = 2/3$. This case is applicable in general when the flow speed is much lower than the Alfvén speed. In that scenario, the system's longer dynamical evolution will still hold true to $\gamma_{c,\text{eff}} = 2/3$ because the cosmic-ray pressure will rapidly adjust to balance the density gradient, satisfying Equation 7.35. This effective polytropic index even appears in non-steady state simulations with a strong magnetic field because of the timescale ordering (Habegger & Zweibel, 2025).

The final regime is that of dominant cosmic-ray diffusion. This regime obeys the

equation

$$\frac{\partial \ln P_c}{\partial t} = \nabla \cdot (\boldsymbol{\kappa} \cdot \nabla \ln P_c) + (\nabla \ln P_c) \cdot \boldsymbol{\kappa} \cdot \nabla \ln P_c. \quad (7.36)$$

Given that this requires diffusion be faster than the local fluid velocity and the Alfvén speed, the cosmic-ray pressure must evolve on a faster timescale than the thermal plasma. Since diffusion tends to flatten fluctuations, any local peaks in cosmic-ray pressure will diffuse to be a constant on a faster timescale than any dynamics in the background plasma. Therefore, the cosmic-ray pressure itself can be treated as a constant. That is equivalent to taking $\gamma_{c,\text{eff}} = 0$.

In our models, we take $\gamma_{c,\text{eff}} = 0$ in the galactic disk and neutral gas because it is expected that diffusion is the dominant transport process there. This index changes in the galactic halo and in ionized gas. Simulations have illustrated a shift to advective and streaming transport in the galactic halo and well into the circumgalactic medium (CGM) as density decreases (Armillotta et al., 2024; Thomas et al., 2025; Habegger & Zweibel, 2025). In our models, we set $\gamma_{c,\text{eff}}$ to be either 2/3 or 4/3, to explore both the advective and streaming regimes. As discussed above, the dominant transport will be set by determining whether the flow speed or the Alfvén speed is higher. In a super-Alfvénic flow, $\gamma_{c,\text{eff}} = 4/3$, but in a sub-Alfvénic flow, $\gamma_{c,\text{eff}} = 2/3$.

7.8.3 Effective Polytrropic Index for the Gas Pressure

The relationship between pressure and density can be derived from the First Law of Thermodynamics, combined with the equation of mass conservation

$$\frac{dP_g}{dt} = -\gamma_g P_g \nabla \cdot \mathbf{u} + \frac{\rho \mathcal{L}}{\gamma_g - 1}, \quad (7.37)$$

where \mathcal{L} is a heat-loss function (Field, 1965) denoting all non-adiabatic heating and cooling terms.

In what follows, we treat two limits: the adiabatic limit and the non-adiabatic limit. In the adiabatic limit, the fluid divergence term $\gamma_g P_g \nabla \cdot \mathbf{u}$ dominates the gas's thermal evolution. In the non-adiabatic limit, the heating and cooling effects encoded in \mathcal{L} dominate the thermal evolution. We ignore thermal conduction, as it is generally unimportant at the temperatures we consider here, except at interfaces between different phases of the ISM.

In the adiabatic limit, with no heating or radiative cooling effects (an ideal gas), Equation 7.37 implies a conserved quantity related to the specific entropy s_g (Hosking et al., 2025):

$$Q_g = \frac{P_g}{\rho^{\gamma_g}} = s_g^{\gamma_g}. \quad (7.38)$$

Therefore, in the adiabatic limit, we have $\gamma_{g,\text{eff}} = \gamma_g = 5/3$ for a (mostly) hydrogen gas or plasma.

For the ISM, where non-ideal effects are important, we can define an effective polytrropic index $\gamma_{g,\text{eff}}$ and solve for it assuming heating and radiative cooling are the dominant processes. In this non-adiabatic limit, we initially consider cases where

the gas is in thermal equilibrium. That case corresponds $\mathcal{L} = 0$. Therefore, all non-adiabatic heating and cooling must balance. In notation commonly found in studies of the ISM, we get

$$\Gamma + \Gamma_{cr} = n\Lambda(T) \quad (7.39)$$

where Γ is the total heating rate per particle (e.g., from heating due to the photoelectric effect on dust grains), Γ_{cr} is the heating due to cosmic-ray streaming, and $\Lambda(T)$ is the volumetric cooling rate. Equation 7.39 along with the ideal gas law $P = nk_B T$ (assuming just a hydrogen gas), creates a closed system for us to calculate $\gamma_{g,\text{eff}}$.

Perturbing the heating rate, gas density, and temperature, magnetic field strength, etc., we can differentiate Equation 7.39 with and get an equation for $\gamma_{g,\text{eff}}$:

$$= \frac{\delta n}{n} \left(1 + (\gamma_{g,\text{eff}} - 1) \frac{\partial \ln \Lambda}{\partial \ln T} \right). \quad (7.40)$$

In the above, we have already assumed $dT/T = (\gamma_{g,\text{eff}} - 1)dn/n$ for an ideal gas equation of state. Rearranging this to solve for $\gamma_{g,\text{eff}}$, we find:

$$\gamma_{g,\text{eff}} = 1 + \left(\frac{\partial \ln \Lambda}{\partial \ln T} \right)^{-1} \times \left(\frac{\delta \Gamma + \delta \Gamma_{cr}}{\Lambda(T)\delta n} - 1 \right). \quad (7.41)$$

If there is no heating from cosmic-ray streaming, and using a power law cooling function $\Lambda \propto T^\alpha$ and constant heating rate ($\delta \Gamma \rightarrow 0$), then the effective polytropic index is

$$\gamma_{g,\text{eff}} = 1 - \frac{1}{\alpha}. \quad (7.42)$$

However, if cosmic-ray heating is dominant in keeping thermal equilibrium, then

we need to calculate the perturbation of the cosmic-ray heating term. The cosmic-ray heating term is

$$\Gamma_{cr} = \frac{BP_c}{n\sqrt{4\pi\rho}} \left| \hat{b} \cdot \nabla \ln P_c \right| = \frac{\gamma_{c,\text{eff}} B_0 P_{c,0}}{n_0 \sqrt{4\pi\rho_0}} \times \left(\frac{\rho}{\rho_0} \right)^{\gamma_{B,\text{eff}}/2 + \gamma_{c,\text{eff}} - 3/2} \left| \hat{b} \cdot \nabla \ln \rho \right|, \quad (7.43)$$

where we have used the polytropic indices from Appendix 7.8.1 and Appendix 7.8.2 for the magnetic pressure and the cosmic-ray pressure. Perturbing Equation 7.43, we find

$$\frac{\delta\Gamma_{cr}}{\Gamma_{cr}} = (\gamma_{B,\text{eff}}/2 + \gamma_{c,\text{eff}} - 3/2) \frac{\delta\rho}{\rho} + \frac{\delta \left| \hat{b} \cdot \nabla \ln \rho \right|}{\left| \hat{b} \cdot \nabla \ln \rho \right|}. \quad (7.44)$$

We will neglect the term related to the density gradient, assuming a constant gradient in the region where we are considering this perturbation from equilibrium (Equation 7.39). Assuming $\gamma_B = 0$ as discussed in Appendix 7.8.1, and assuming $\gamma_{c,\text{eff}} = 2/3$ since we are assuming cosmic-ray streaming is significant, we get

$$\frac{\delta\Gamma_{cr}}{\Gamma_{cr}} = -\frac{5}{6} \frac{\delta\rho}{\rho} \quad (7.45)$$

The negative power law relationship between ρ and Γ_{cr} agrees with results from Wiener et al. (2013), which follows a similar line of reasoning but with different choices for effective polytropic indices.

Plugging Equation 7.45 into Equation 7.41, ignoring the non-cosmic-ray heating, the effective gas index is

$$\gamma_{g,\text{eff}} = 1 - \frac{11}{6\alpha}. \quad (7.46)$$

Overall, the Alfvénic heating amplifies the effect of the cooling function $\Lambda \propto T^\alpha$

compared to Equation 7.42.

This Alfvénic heating rate could be dominant outside the disk, where there is less starlight to create photo-ionization or to illuminate dust grains to heat the gas via the photoelectric effect. Additionally, it assumes the Alfvén speed is high so that cosmic rays are transported at the Alfvén speed. While these conditions are rare, we note this calculation in order to identify the maximal impact of that heating term on the effective adiabatic index of the gas. A similar treatment of dominant cosmic-ray heating appears in Kempster & Quataert (2020); Kempster et al. (2023b), works that studied galaxy cluster halos. There are additional buoyancy instabilities in that regime, but we do not consider them here.

Overall, we can examine the case of an adiabatic gas with $\gamma_{g,\text{eff}} = 5/3$, or a gas in thermal equilibrium between heating and radiative cooling. In that non-adiabatic regime, we can either treat the gas as isothermal, with $\gamma_{g,\text{eff}} = 1$, or decrease the value as a result of the cooling function slope (see Equations 7.42 and 7.46). However, the cooling function slope plays an important role in determining the thermal stability of the gas — if the gas is thermally unstable, then it is unlikely that the buoyancy analysis in Section 7.3 is applicable. So, we need to isolate a thermally stable gas, which is why we focused on $\gamma_{g,\text{eff}} = 5/3$ for a majority of the analysis in this paper.

8 CONCLUSION

I leave Sisyphus at the foot of the mountain!

One always finds one's burden again.

— *The Myth of Sisyphus* (1955) by Albert Camus

In conclusion, we have examined several interesting problems in the study of cosmic-ray feedback and identified instances where different avenues of cosmic-ray transport have a dynamical impact and leave observable imprints on the thermal plasma. We focused on two setups: (1) periodic, turbulent boxes and (2) stratified boxes. In this chapter, I synthesize the key takeaways from the novel work presented in this thesis (Chapters 3 through 7).

There are three key avenues for the transport of cosmic-ray energy across large scales. The cosmic-ray energy can be advected with the background plasma, it can stream along magnetic field lines, or it can diffuse through the background plasma. Advection with the gas leads to a direct forcing through the cosmic-ray pressure gradient. Within the two non-advective transport mechanisms, streaming and diffusion, each process will cause different dynamical effects. In general, diffusion allows the cosmic-ray energy to move independently from the thermal plasma, effectively decoupling the cosmic-ray fluid from the thermal plasma and decreasing the cosmic-ray pressure gradient, reducing the forcing on the background plasma. Streaming also allows this decoupling, except with an additional heating term because it rests on the self-confinement assumption: if cosmic rays generate the waves that scatter them, then the waves will eventually damp and leave energy behind in the thermal

plasma. Additionally, streaming leads to sharp cosmic-ray pressure gradients, which could still provide significant forcing.

The work presented in this thesis has actually illustrated that the above picture of the relationship between transport and feedback is not so simple. Diffusion not only works to decrease feedback — it can facilitate the transfer of cosmic-ray energy to larger scales, eschewing energy from a turbulent cascade, moving cosmic-ray energy to other places where it will provide forcing and heating. In a separate vein, streaming transport provides an additional avenue for cosmic-ray energy to go back into the thermal plasma, keeping the overall energy accounting coupled even if the cosmic-ray energy is spatially decoupled from the plasma.

This energy accounting is clearest in the turbulence work of Chapters 3 and 4. In those works, the diffusion of cosmic-ray energy density can allow the cosmic-ray energy to decouple from the thermal and kinetic energy of the background plasma. This decoupling will occur on a critical scale set by the diffusion coefficient and the phase speed of the turbulence. From there, cosmic-ray energy will grow unabated until it creates significant pressurization, allowing the energy to be returned to the bulk velocity of the thermal plasma.

Additionally, when non-adiabatic effects like optically-thin radiative cooling and photoelectric heating are introduced, the compressions of thermal plasma become non-adiabatic, with higher density gas being at a lower temperature. These multi-phase plasmas may be in pressure equilibrium, but the cosmic-ray energy can experience a completely different effect. While the gas pressure decreases because of the cooling, the cosmic-ray pressure will increase due to the compression if the cosmic-ray

energy is advected with the gas. After the compression, the thermal gas can remain in the condensed, cool state for a long period of time. Given enough time, cosmic-ray diffusion will take over and allow the cosmic-ray energy to diffuse outward. Since the initial compression of the cosmic rays and gas was held in pressure equilibrium, the radiative cooling was limited by the pressurization by cosmic rays. The cosmic rays are thereby able to take away energy that would have been radiated away, eventually returning it to larger scales.

In the case of a stratified atmosphere, the transport of cosmic rays determines the compressibility of the cosmic-ray fluid. The compressibility then determines the dynamical effect of cosmic-ray injections, either for an individual injection or many injections. Even though the simulations in Chapters 5 and 6 provide a more detailed picture, the analytic work in Chapter 7 identifies how the results can be distilled to the simple physics of buoyancy. If diffusion is dominant, then cosmic-ray energy is easier to compress and tends to be more buoyantly unstable (in either the linear or nonlinear case). If advection is dominant, then cosmic-ray energy is difficult to compress, but not as difficult as the gas, leading to systems that are more stable against buoyant motion. Finally, for streaming, the cosmic-ray energy does not compress as easily as for diffusion, but it is still more compressible than the advection case, leading to more buoyant instability than in the case of advection.

Those buoyancy results apply to the stratified box simulations. For example, in Chapter 5 we show that if the background is buoyantly stable, then a single cosmic-ray injection will not drive a large buoyant flux tube. It will cause the flux tube to fluctuate, but it will sink down to a similar location to where it started. The

only way to break that behaviour would be through many injections, inflating the flux tube and atmosphere until it is completely unstable. This result is illustrated in Chapter 6, which shows the disruption of an entire slab of gas once the Parker instability has had ~ 100 Myr to grow.

More work remains to be completed in comparing the analytic results of Chapter 7 to simulations. In particular, the shift in cosmic-ray transport occurs as a function of plasma density, and additional shifts could occur based on magnetic field strength. Multiphase, 3D simulations with adjustable cosmic-ray transport and physics are necessary to identify the full extent of the effects of the compressibility argument. If dense gas is actually more compressible because the cosmic rays rapidly diffuse through it, then that would have dynamical effects on the motions of cold clouds in the galaxy.

Finally, future improvements of our understanding of cosmic-ray feedback and transport will need to utilize the available observational results. We have many probes of the cosmic-ray spectrum here on Earth (see Chapter 1.1). It is more difficult to observe cosmic rays far away from our home, in the ISM outside the solar neighborhood, in other galaxies, or on the larger scales covered by the CGM or ICM. For these other locales, observations of γ -ray emission and synchrotron radiation due to the interaction of cosmic rays with the thermal plasma and the magnetic field provide useful information to constrain the total cosmic-ray energy in far-off places. These indirect methods of observing cosmic rays provide several clues to the interaction of cosmic rays with the thermal plasma.

In Chapter 4, we focus on how cosmic-ray transport sets the γ -ray luminosity of a

large-scale, diffuse plasma. We illustrate that fast cosmic-ray transport is necessary to keep γ -ray luminosity low. Otherwise, the correlation of cosmic rays with dense gas and the turbulent reacceleration of cosmic rays will drive the γ -ray luminosity up. Theoretical analyses like those are necessary to interpret observations of γ -ray luminosity, which generally provide only upper limits on the diffuse emission of γ -rays produced by cosmic rays.

This result on large scales, requiring fast cosmic-ray transport, actually moves in the opposite direction from results on small scales. Observations of nearby γ -ray and synchrotron emission suggest cosmic rays need to be confined near their sources for a certain amount of time to explain observations. These differing results highlight the necessity of examining cosmic-ray transport in a variety of environments, across many scales — the transport and feedback of cosmic rays will depend on their environment and the scales of interest. This concept has appeared again and again over the chapters of this thesis; from turbulent reacceleration to the buoyant rise of cosmic-ray-loaded flux tubes, cosmic-ray transport changes the resulting feedback.

This work focused in general on those large-scale problems, where there are fewer opportunities to constrain the physics of cosmic-ray feedback against observations. In the future, I plan to study the transport and feedback of cosmic rays on their most likely sources: supernova remnants. With over ~ 300 of these scattered throughout the Milky Way, emitting synchrotron, X-ray, and γ -rays, along with particle messengers like neutrinos, there is an opportunity to constrain cosmic-ray feedback and transport on the key scales where they are created. The challenge is that these sources sit on the edge, between the kinetic scale treatment and the fluid treatment

of cosmic-ray transport and evolution. To probe the physics of their acceleration and feedback on these mesoscales, we need to evolve the distribution function of cosmic rays along with the overall cosmic-ray pressure. Those simulations will also need to be informed by fundamental kinetic simulations, which evolve the motions of the particles on even smaller scales (e.g., Particle-In-Cell simulations).

I am hopeful that there is enough complexity in that problem to keep me busy for a lifetime. If not, then I hope I have the strength that Ellen has: to push the boulder up again, always excited to find new tidbits of knowledge lurking around the next ledge.

BIBLIOGRAPHY

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, *ApJ*, 699, 31, doi: 10.1088/0004-637X/699/1/31
- Ackermann, M., Ajello, M., Albert, A., et al. 2014, *ApJ*, 787, 18, doi: 10.1088/0004-637X/787/1/18
- Aerdker, S., Habegger, R., Merten, L., Zweibel, E., & Becker Tjus, J. 2025, *A&A*, 703, A51, doi: 10.1051/0004-6361/202453593
- Anderson, C. D., Millikan, R. A., Neddermeyer, S., & Pickering, W. 1934, *Physical Review*, 45, 352, doi: 10.1103/PhysRev.45.352
- Armiglotta, L., Ostriker, E. C., & Jiang, Y.-F. 2021, *ApJ*, 922, 11, doi: 10.3847/1538-4357/ac1db2
- . 2022, *ApJ*, 929, 170, doi: 10.3847/1538-4357/ac5fa9
- Armiglotta, L., Ostriker, E. C., Kim, C.-G., & Jiang, Y.-F. 2024, *ApJ*, 964, 99, doi: 10.3847/1538-4357/ad1e5c
- Asseo, E., Cesarsky, C. J., Lachize-Rey, M., & Pellat, R. 1978, *ApJL*, 225, L21, doi: 10.1086/182784
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: 10.1051/0004-6361/201322068
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: 10.3847/1538-3881/aabc4f
- Axford, W. I., Leer, E., & Skadron, G. 1977, in *International Cosmic Ray Conference*, Vol. 11, International Cosmic Ray Conference, 132
- Baade, W., & Zwicky, F. 1934a, *Proceedings of the National Academy of Science*, 20, 259, doi: 10.1073/pnas.20.5.259
- . 1934b, *Physical Review*, 46, 76, doi: 10.1103/PhysRev.46.76.2

- Beckmann, R. S., Dubois, Y., Pellissier, A., et al. 2022, *A&A*, 665, A129, doi: 10.1051/0004-6361/202142527
- Bell, A. R. 1978, *MNRAS*, 182, 147, doi: 10.1093/mnras/182.2.147
- Bell, E. F. 2003, *ApJ*, 586, 794, doi: 10.1086/367829
- Berezinskii, V. S., Bulanov, S. V., Dogiel, V. A., & Ptuskin, V. S. 1990, *Astrophysics of cosmic rays*
- Bernstein, I. B., Frieman, E. A., Kruskal, M. D., & Kulsrud, R. M. 1958, *Proceedings of the Royal Society of London Series A*, 244, 17, doi: 10.1098/rspa.1958.0023
- Blandford, R. D., & Ostriker, J. P. 1978, *ApJL*, 221, L29, doi: 10.1086/182658
- Boettcher, E., Zweibel, E. G., Yoast-Hull, T. M., & Gallagher, III, J. S. 2013, *ApJ*, 779, 12, doi: 10.1088/0004-637X/779/1/12
- Böss, L. M., Dolag, K., Steinwandel, U. P., et al. 2024, *A&A*, 692, A232, doi: 10.1051/0004-6361/202348339
- Boulares, A., & Cox, D. P. 1990, *ApJ*, 365, 544, doi: 10.1086/169509
- Breitschwerdt, D., McKenzie, J. F., & Voelk, H. J. 1991, *A&A*, 245, 79
- Bustard, C., & Oh, S. P. 2022, *ApJ*, 941, 65, doi: 10.3847/1538-4357/aca021
- . 2023, *ApJ*, 955, 64, doi: 10.3847/1538-4357/aceef9
- Bustard, C., Pardy, S. A., D’Onghia, E., Zweibel, E. G., & Gallagher, III, J. S. 2018, *ApJ*, 863, 49, doi: 10.3847/1538-4357/aad08f
- Bustard, C., & Zweibel, E. G. 2021, *ApJ*, 913, 106, doi: 10.3847/1538-4357/abf64c
- Bustard, C., Zweibel, E. G., D’Onghia, E., Gallagher, III, J. S., & Farber, R. 2020, *ApJ*, 893, 29, doi: 10.3847/1538-4357/ab7fa3
- Caprioli, D., & Spitkovsky, A. 2014, *ApJ*, 783, 91, doi: 10.1088/0004-637X/783/2/91
- Chan, T. K., Kereš, D., Gurvich, A. B., et al. 2022, *MNRAS*, 517, 597, doi: 10.1093/mnras/stac2236

- Chan, T. K., Kereš, D., Hopkins, P. F., et al. 2019, *MNRAS*, 488, 3716, doi: 10.1093/mnras/stz1895
- Chatzikos, M., Bianchi, S., Camilloni, F., et al. 2023, *RMxAA*, 59, 327, doi: 10.22201/ia.01851101p.2023.59.02.12
- Commerçon, B., Marcowith, A., & Dubois, Y. 2019, *A&A*, 622, A143, doi: 10.1051/0004-6361/201833809
- Corti, C., Bindi, V., Consolandi, C., & Whitman, K. 2016, *ApJ*, 829, 8, doi: 10.3847/0004-637X/829/1/8
- Crutcher, R. M., Wandelt, B., Heiles, C., Falgarone, E., & Troland, T. H. 2010, *ApJ*, 725, 466, doi: 10.1088/0004-637X/725/1/466
- De La Torre Luque, P., Gaggero, D., Grasso, D., et al. 2023, *A&A*, 672, A58, doi: 10.1051/0004-6361/202243714
- Desiati, P., & Zweibel, E. G. 2014, *ApJ*, 791, 51, doi: 10.1088/0004-637X/791/1/51
- Diesing, R., & Caprioli, D. 2018, *PhRvL*, 121, 091101, doi: 10.1103/PhysRevLett.121.091101
- Draine, B. T. 2011, *Physics of the Interstellar and Intergalactic Medium* (Princeton University Press)
- Dubois, Y., Commerçon, B., Marcowith, A., & Brahim, L. 2019, *A&A*, 631, A121, doi: 10.1051/0004-6361/201936275
- Elia, D., Molinari, S., Schisano, E., et al. 2022, *ApJ*, 941, 162, doi: 10.3847/1538-4357/aca27d
- Enßlin, T. A., Pfrommer, C., Springel, V., & Jubelgas, M. 2007, *A&A*, 473, 41, doi: 10.1051/0004-6361:20065294
- Eswaran, V., & Pope, S. B. 1988, *Computers and Fluids*, 16, 257, doi: 10.1016/0045-7930(88)90013-8
- Everett, J. E., Zweibel, E. G., Benjamin, R. A., et al. 2008, *ApJ*, 674, 258, doi: 10.1086/524766
- Evoli, C. 2020, *The Cosmic-Ray Energy Spectrum*, Zenodo, doi: 10.5281/zenodo.4396125

- Evoli, C., Aloisio, R., & Blasi, P. 2019, *PhRvD*, 99, 103023, doi: 10.1103/PhysRevD.99.103023
- Evoli, C., Morlino, G., Blasi, P., & Aloisio, R. 2020, *PhRvD*, 101, 023013, doi: 10.1103/PhysRevD.101.023013
- Fabian, A. C., Sanders, J. S., Taylor, G. B., et al. 2006, *MNRAS*, 366, 417, doi: 10.1111/j.1365-2966.2005.09896.x
- Farber, R., Ruszkowski, M., Yang, H.-Y. K., & Zweibel, E. G. 2018, *ApJ*, 856, 112, doi: 10.3847/1538-4357/aab26d
- Fermi, E. 1949, *Physical Review*, 75, 1169, doi: 10.1103/PhysRev.75.1169
- Ferrière, K. M. 2001, *Reviews of Modern Physics*, 73, 1031, doi: 10.1103/RevModPhys.73.1031
- Field, G. B. 1965, *ApJ*, 142, 531, doi: 10.1086/148317
- Fryxell, B., Olson, K., Ricker, P., et al. 2000, *ApJS*, 131, 273, doi: 10.1086/317361
- Girichidis, P., Naab, T., Walch, S., et al. 2016, *ApJL*, 816, L19, doi: 10.3847/2041-8205/816/2/L19
- Giz, A. T., & Shu, F. H. 1993, *ApJ*, 404, 185, doi: 10.1086/172267
- Goody, R. M., & Yung, Y. L. 1989, *Atmospheric radiation : theoretical basis*
- Guo, F., & Oh, S. P. 2008, *MNRAS*, 384, 251, doi: 10.1111/j.1365-2966.2007.12692.x
- Guo, Y.-Q., Tian, Z., & Jin, C. 2016, *ApJ*, 819, 54, doi: 10.3847/0004-637X/819/1/54
- Habegger, R., Ho, K. W., Yuen, K. H., & Zweibel, E. G. 2024, *ApJ*, 974, 17, doi: 10.3847/1538-4357/ad67da
- Habegger, R., Ruszkowski, M., & Zweibel, E. G. 2025, arXiv e-prints, arXiv:2510.24622, doi: 10.48550/arXiv.2510.24622
- Habegger, R., & Zweibel, E. G. 2025, *ApJ*, 990, 75, doi: 10.3847/1538-4357/adf4d7

- Habegger, R., Zweibel, E. G., & Wong, S. 2023, *ApJ*, 951, 99, doi: 10.3847/1538-4357/accf8e
- Hanasz, M., Kowal, G., Otmianowska-Mazur, K., & Lesch, H. 2004, *ApJL*, 605, L33, doi: 10.1086/420697
- Hanasz, M., & Lesch, H. 2000, *ApJ*, 543, 235, doi: 10.1086/317077
- Hanasz, M., Lesch, H., Naab, T., et al. 2013, *ApJL*, 777, L38, doi: 10.1088/2041-8205/777/2/L38
- Hanasz, M., Wóltański, D., & Kowalik, K. 2009, *ApJL*, 706, L155, doi: 10.1088/0004-637X/706/1/L155
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: 10.1038/s41586-020-2649-2
- Heesen, V., Stein, M., Pourjafari, N., et al. 2025, *A&A*, 699, A243, doi: 10.1051/0004-6361/202554046
- Heintz, E., Bustard, C., & Zweibel, E. G. 2020, *ApJ*, 891, 157, doi: 10.3847/1538-4357/ab7453
- Heintz, E., & Zweibel, E. G. 2018, *ApJ*, 860, 97, doi: 10.3847/1538-4357/aac208
- Hitomi Collaboration, Aharonian, F., Akamatsu, H., et al. 2016, *Nature*, 535, 117, doi: 10.1038/nature18627
- . 2018, *PASJ*, 70, 9, doi: 10.1093/pasj/psx138
- Ho, K. W., Yuen, K. H., & Lazarian, A. 2023, *MNRAS*, 521, 230, doi: 10.1093/mnras/stad481
- . 2024, arXiv e-prints, arXiv:2407.14199, doi: 10.48550/arXiv.2407.14199
- Hopkins, P. F., Chan, T. K., Squire, J., et al. 2021, *MNRAS*, 501, 3663, doi: 10.1093/mnras/staa3692
- Hopkins, P. F., Wetzell, A., Kereš, D., et al. 2018, *MNRAS*, 480, 800, doi: 10.1093/mnras/sty1690
- Hosking, D. N., Wasserman, D., & Cowley, S. C. 2025, *Journal of Plasma Physics*, 91, E35, doi: 10.1017/S0022377824001521

- Hunter, J. D. 2007, *Computing in Science and Engineering*, 9, 90, doi: 10.1109/MCSE.2007.55
- Inoue, T., Inutsuka, S.-i., & Koyama, H. 2006, *ApJ*, 652, 1331, doi: 10.1086/508334
- Iroshnikov, P. S. 1964, *Soviet Ast.*, 7, 566
- Jiang, Y.-F., & Oh, S. P. 2018, *ApJ*, 854, 5, doi: 10.3847/1538-4357/aaa6ce
- Jóhannesson, G., Porter, T. A., & Moskalenko, I. V. 2019, *ApJ*, 879, 91, doi: 10.3847/1538-4357/ab258e
- Jokipii, J. R. 1966, *ApJ*, 146, 480, doi: 10.1086/148912
- Jones, F. C., Lukasiak, A., Ptuskin, V., & Webber, W. 2001, *ApJ*, 547, 264, doi: 10.1086/318358
- Kalberla, P. M. W., & Kerp, J. 2009, *ARA&A*, 47, 27, doi: 10.1146/annurev-astro-082708-101823
- Kempski, P., Fielding, D. B., Quataert, E., et al. 2023a, *MNRAS*, 525, 4985, doi: 10.1093/mnras/stad2609
- Kempski, P., & Quataert, E. 2020, *MNRAS*, 493, 1801, doi: 10.1093/mnras/staa385
- Kempski, P., Quataert, E., & Squire, J. 2023b, *MNRAS*, 524, 1893, doi: 10.1093/mnras/stad1744
- Kim, C.-G., Kim, J.-G., Gong, M., & Ostriker, E. C. 2023, *ApJ*, 946, 3, doi: 10.3847/1538-4357/acbd3a
- Kim, C.-G., & Ostriker, E. C. 2015, *ApJ*, 802, 99, doi: 10.1088/0004-637X/802/2/99
- Kolmogorov, A. 1941, *Akademiia Nauk SSSR Doklady*, 30, 301
- Kraichnan, R. H. 1965, *Physics of Fluids*, 8, 1385, doi: 10.1063/1.1761412
- Krause, M., Irwin, J., Wiegert, T., et al. 2018, *A&A*, 611, A72, doi: 10.1051/0004-6361/201731991

- Kritsuk, A. G., Ustyugov, S. D., & Norman, M. L. 2017, *New Journal of Physics*, 19, 065003, doi: 10.1088/1367-2630/aa7156
- Kulsrud, R., & Pearce, W. P. 1969, *ApJ*, 156, 445, doi: 10.1086/149981
- Kulsrud, R. M. 2004, *Plasma Physics for Astrophysics* (Princeton University Press)
- Kulsrud, R. M., & Cesarsky, C. J. 1971, *Astrophys. Lett.*, 8, 189
- Lacki, B. C., Thompson, T. A., Quataert, E., Loeb, A., & Waxman, E. 2011, *ApJ*, 734, 107, doi: 10.1088/0004-637X/734/2/107
- Lemoine, M. 2023, *Journal of Plasma Physics*, 89, 175890501, doi: 10.1017/S0022377823000946
- Matthaeus, W. H., Qin, G., Bieber, J. W., & Zank, G. P. 2003, *ApJL*, 590, L53, doi: 10.1086/376613
- McCallum, L., Wood, K., Benjamin, R., Krishnarao, D., & McLeod, A. F. 2025, *MNRAS*, 541, 2324, doi: 10.1093/mnras/staf1022
- McClure-Griffiths, N. M., Stanimirović, S., & Rybarczyk, D. R. 2023, *ARA&A*, 61, 19, doi: 10.1146/annurev-astro-052920-104851
- McKee, C. F., & Ostriker, E. C. 2007, *ARA&A*, 45, 565, doi: 10.1146/annurev.astro.45.051806.110602
- McKee, C. F., Parravano, A., & Hollenbach, D. J. 2015, *ApJ*, 814, 13, doi: 10.1088/0004-637X/814/1/13
- McKenzie, J. F., & Voelk, H. J. 1982, *A&A*, 116, 191
- Mouschovias, T. C. 1974, *ApJ*, 192, 37, doi: 10.1086/153032
- Newcomb, W. A. 1961, *Physics of Fluids*, 4, 391, doi: 10.1063/1.1706342
- O'Neill, T. J., Zucker, C., Goodman, A. A., & Edenhofer, G. 2024, *ApJ*, 973, 136, doi: 10.3847/1538-4357/ad61de
- Owen, E. R., Wu, K., Inoue, Y., Yang, H.-Y. K., & Mitchell, A. M. W. 2023, *Galaxies*, 11, 86, doi: 10.3390/galaxies11040086
- Parker, E. N. 1965, *Planet. Space Sci.*, 13, 9, doi: 10.1016/0032-0633(65)90131-5

- . 1966, *ApJ*, 145, 811, doi: 10.1086/148828
- Pingel, N. M., Lee, M.-Y., Burkhart, B., & Stanimirović, S. 2018, *ApJ*, 856, 136, doi: 10.3847/1538-4357/aab34b
- Pingel, N. M., Dempsey, J., McClure-Griffiths, N. M., et al. 2022, *PASA*, 39, e005, doi: 10.1017/pasa.2021.59
- Pinzke, A., & Pfrommer, C. 2010, *MNRAS*, 409, 449, doi: 10.1111/j.1365-2966.2010.17328.x
- Ptuskin, V. S. 1988, *Soviet Astronomy Letters*, 14, 255
- Rathjen, T.-E., Naab, T., Girichidis, P., et al. 2021, *MNRAS*, 504, 1039, doi: 10.1093/mnras/stab900
- Rodrigues, L. F. S., Sarson, G. R., Shukurov, A., Bushby, P. J., & Fletcher, A. 2016, *ApJ*, 816, 2, doi: 10.3847/0004-637X/816/1/2
- Rosen, A., & Bregman, J. N. 1995, *ApJ*, 440, 634, doi: 10.1086/175303
- Ruszkowski, M., & Pfrommer, C. 2023, *A&A Rv*, 31, 4, doi: 10.1007/s00159-023-00149-2
- Ruszkowski, M., Yang, H.-Y. K., & Zweibel, E. 2017, *ApJ*, 834, 208, doi: 10.3847/1538-4357/834/2/208
- Rybarczyk, D. R., Wenger, T. V., & Stanimirović, S. 2024, *ApJ*, 975, 167, doi: 10.3847/1538-4357/ad79f7
- Sampson, M. L., Beattie, J. R., Teyssier, R., et al. 2025, arXiv e-prints, arXiv:2506.03768, doi: 10.48550/arXiv.2506.03768
- Schmidt, M. 1959, *ApJ*, 129, 243, doi: 10.1086/146614
- Schwarzschild, M. 1958, *Structure and evolution of the stars*.
- Semenov, V. A., Kravtsov, A. V., & Caprioli, D. 2021, *ApJ*, 910, 126, doi: 10.3847/1538-4357/abe2a6
- Seta, A., & McClure-Griffiths, N. M. 2025, *MNRAS*, 539, 1024, doi: 10.1093/mnras/staf520

- Shalchi, A., Bieber, J. W., Matthaeus, W. H., & Schlickeiser, R. 2006, *ApJ*, 642, 230, doi: 10.1086/500728
- Shapiro, P. R., & Field, G. B. 1976, *ApJ*, 205, 762, doi: 10.1086/154332
- Sharma, P., Colella, P., & Martin, D. F. 2010, *SIAM Journal on Scientific Computing*, 32, 3564, doi: 10.1137/100792135
- Shimoda, J., & Inutsuka, S.-i. 2022, *ApJ*, 926, 8, doi: 10.3847/1538-4357/ac4110
- Skilling, J. 1975, *MNRAS*, 172, 557, doi: 10.1093/mnras/172.3.557
- Stein, M., Kleimann, J., Adebahr, B., et al. 2025, *A&A*, 696, A112, doi: 10.1051/0004-6361/202452322
- Stone, E. C., Cummings, A. C., Heikkila, B. C., & Lal, N. 2019, *Nature Astronomy*, 3, 1013, doi: 10.1038/s41550-019-0928-3
- Stone, E. C., Vogt, R. E., McDonald, F. B., et al. 1977, *SSRv*, 21, 355, doi: 10.1007/BF00211546
- Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. 2020, *ApJS*, 249, 4, doi: 10.3847/1538-4365/ab929b
- Sun, X., & Bai, X.-N. 2023, *MNRAS*, 523, 3328, doi: 10.1093/mnras/stad1548
- Sutherland, R. S., & Dopita, M. A. 1993, *ApJS*, 88, 253, doi: 10.1086/191823
- Tan, B., & Fielding, D. B. 2024, *MNRAS*, 527, 9683, doi: 10.1093/mnras/stad3793
- Tharakkal, D., Shukurov, A., Gent, F. A., et al. 2023, *MNRAS*, 525, 5597, doi: 10.1093/mnras/stad2610
- Thomas, T., & Pfrommer, C. 2019, *MNRAS*, 485, 2977, doi: 10.1093/mnras/stz263
- Thomas, T., Pfrommer, C., & Pakmor, R. 2021, *MNRAS*, 503, 2242, doi: 10.1093/mnras/stab397
- . 2023, *MNRAS*, 521, 3023, doi: 10.1093/mnras/stad472
- . 2025, *A&A*, 698, A104, doi: 10.1051/0004-6361/202450817

- Townsend, R. H. D. 2009, *ApJS*, 181, 391, doi: 10.1088/0067-0049/181/2/391
- Tritsis, A., Panopoulou, G. V., Mouschovias, T. C., Tassis, K., & Pavlidou, V. 2015, *MNRAS*, 451, 4384, doi: 10.1093/mnras/stv1133
- Tsung, T. H. N., Oh, S. P., & Bustard, C. 2023, *MNRAS*, 526, 3301, doi: 10.1093/mnras/stad2720
- Tsung, T. H. N., Oh, S. P., & Jiang, Y.-F. 2022, *MNRAS*, 513, 4464, doi: 10.1093/mnras/stac1123
- Uhlenbeck, G. E., & Ornstein, L. S. 1930, *Physical Review*, 36, 823, doi: 10.1103/PhysRev.36.823
- Uhlig, M., Pfrommer, C., Sharma, M., et al. 2012, *MNRAS*, 423, 2374, doi: 10.1111/j.1365-2966.2012.21045.x
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Computing in Science and Engineering*, 13, 22, doi: 10.1109/MCSE.2011.37
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: 10.1038/s41592-019-0686-2
- Walch, S., Girichidis, P., Naab, T., et al. 2015, *MNRAS*, 454, 238, doi: 10.1093/mnras/stv1975
- Wenger, T. V., Rybarczyk, D. R., & Stanimirović, S. 2024, *ApJ*, 966, 206, doi: 10.3847/1538-4357/ad3923
- Wentzel, D. G. 1969, *ApJ*, 156, 303, doi: 10.1086/149965
- Whitworth, D. J., Srinivasan, S., Pudritz, R. E., et al. 2024, arXiv e-prints, arXiv:2407.18293, doi: 10.48550/arXiv.2407.18293
- Wiener, J., Oh, S. P., & Zweibel, E. G. 2017a, *MNRAS*, 467, 646, doi: 10.1093/mnras/stx109
- Wiener, J., Pfrommer, C., & Oh, S. P. 2017b, *MNRAS*, 467, 906, doi: 10.1093/mnras/stx127
- Wiener, J., Zweibel, E. G., & Oh, S. P. 2013, *ApJ*, 767, 87, doi: 10.1088/0004-637X/767/1/87

- Wiener, J., Zweibel, E. G., & Ruszkowski, M. 2019, *MNRAS*, 489, 205, doi: 10.1093/mnras/stz2007
- Xu, S., Yan, H., & Lazarian, A. 2016, *ApJ*, 826, 166, doi: 10.3847/0004-637X/826/2/166
- Yan, H., & Lazarian, A. 2008, *ApJ*, 673, 942, doi: 10.1086/524771
- Zhang, D. 2018, *Galaxies*, 6, 114, doi: 10.3390/galaxies6040114
- Zweibel, E. G. 2013, *Physics of Plasmas*, 20, 055501, doi: 10.1063/1.4807033
- . 2017, *Physics of Plasmas*, 24, 055402, doi: 10.1063/1.4984017
- Zweibel, E. G., & Bruhwiler, D. L. 1992, *ApJ*, 399, 318, doi: 10.1086/171927
- Zweibel, E. G., & Kulsrud, R. M. 1975, *ApJ*, 201, 63, doi: 10.1086/153858