Energy Diffusion-Advection Models of Nonthermal Particle Acceleration in Simulations of Relativistic Plasma Turbulence

by

K. W. Wong

B.Eng., Queensland University of Technology, 2011

M.S., University of Colorado at Boulder, 2019

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2024

Committee Members:

Dmitri Uzdensky, Chair

Gregory Werner

Michael Litos

Yuan Shi

Maria Kazachenko

Wong, K. W. (Ph.D., Physics)

Energy Diffusion-Advection Models of Nonthermal Particle Acceleration in Simulations of Relativistic Plasma Turbulence

Thesis directed by Prof. Dmitri Uzdensky

Relativistic nonthermal plasmas are ubiquitous in high-energy astrophysical systems featuring turbulence such as pulsar wind nebulae and active galactic nuclei, as inferred from broadband nonthermal emission spectra. The underlying turbulent nonthermal particle acceleration (NTPA) processes have traditionally been modelled with a Fokker-Planck (FP) diffusion-advection equation for the particle energy distribution. In this dissertation, I test FP-type NTPA theories by analysing three-dimensional (3D) particle-in-cell (PIC) simulations of magnetised turbulence in collisionless relativistic pair plasma. By tracking the energy histories of large numbers of particles in several simulations with different initial magnetisation σ_0 and system size, I first test the energy-diffusion assumption of the FP framework, finding simple diffusion throughout the parameter space. I then measure the FP energy diffusion and advection coefficients (D and A, respectively) as functions of particle energy γmc^2, and compare their dependence on initial and instantaneous system parameters to theoretical predictions. In the high-energy nonthermal tail, I find, robustly with respect to system size and σ_0, that $D \sim \gamma^2$, with a more complicated but generally shallower scaling at thermal and subthermal energies which varies qualitatively depending on σ_0. Hence, I fit $D = D_0\gamma^2$ in the nonthermal region and find that the scaling of D_0 with the instantaneous magnetisation $\sigma(t)$ is consistent with $D_0 \sim \sigma^{3/2}$, although this flattens somewhat at higher $\sigma \sim 1$. I also measure the evolution of the power-law index $\alpha(t)$ of the particle energy distribution and find that it is well-described by an exponential convergence in time. I then build and test an analytic model connecting the FP coefficients and the observed power-law evolution, predicting that $A(\gamma) \sim \gamma \log(\gamma/\gamma_A^*)$. This is consistent with my measurements of $A(\gamma, t)$, and I furthermore find that the measured $A(\gamma, t)$ can acceptably predict $\alpha(t)$ through the model relations. These results suggest that the basic 2nd-

order Fermi acceleration model, which predicts $D_0 \sim \sigma$, may not be a complete description of NTPA in turbulent collisionless relativistic plasmas. My findings encourage further application of tracked particle methods and FP coefficient measurements as a diagnostic in kinetic simulations of various physical situations including collisionless shocks and magnetic reconnection, with relevance to astrophysical plasmas.

Acknowledgements

Thanks to my advisor, Dmitri Uzdensky, and collaborators Vladimir Zhdankin, Greg Werner, and Mitch Begelman.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Overview

Relativistic plasmas with nonthermal power-law energy distributions are ubiquitous in astrophysical systems such as pulsar wind nebulae (PWN) (Meyer et al., 2010; Bühler and Blandford, 2014), jets from active galactic nuclei (AGN) (Begelman et al., 1984; Hartman et al., 1992) and their radio lobes (Hardcastle et al., 2009), and black-hole accretion-disk coronae (Yuan et al., 2003). A specific example is the Crab Nebula, which is a pulsar wind nebula: a supernova remnant with a central pulsar emitting an energetic wind of charged particles. The Crab Nebula features a striking nonthermal emission spectrum, with a power-law scaling segment spanning multiple decades in photon energies from visible to gamma-ray wavelengths (Meyer et al., 2010). This nonthermal emission is widely believed to be synchrotron radiation from relativistic electrons, as originally suggested by Shklovskii (1953), and implies a corresponding nonthermal distribution of particle energies. Meyer et al. (2010) computes an electron spectrum $f(\gamma) \equiv dN/d\gamma$ with spectral index $\alpha = -d\log f/d\log\gamma$ slightly greater than 3, with ultra-relativistic Lorentz factors $\gamma \equiv 1/\sqrt{1 - v^2/c^2}$ from about 10^5–10^9.

One of the most remarkable manifestations of relativistic particle populations in the universe is cosmic rays: particles, mostly ultra-relativistic protons, with extremely high energy, existing throughout the galaxy and with sufficient density as to constitute a substantial portion of energy. While most of the relativistic nonthermal plasmas are associated with so-called relativistic compact objects, namely, neutron stars and black holes, nonrelativistic systems can also produce ultra-relativistic particles. The shockwaves of supernova explosions spend a few percent of the overall

supernova energy budget on accelerating cosmic rays, meaning that they not only produce high energy particles but do so very efficiently. The majority of cosmic rays are produced by such supernova explosions in our own galaxy. The resulting particle energy distribution has a power-law spanning over six orders of magnitude in particle energy, from the proton rest-mass energy (1GeV) to above PeV. From about 1GeV–3×10^6GeV, the spectral index is 2.6 (Zweibel, 2013). The cosmic-ray power-law spectrum continues, albeit with a somewhat steeper slope, to even higher energies, up to (the famous GZK cutoff) 10^{21}eV. These very-high-energy and ultra-high-energy cosmic rays are believed to be produced by relativistic jets from active galactic nuclei powered by supermassive black holes and in gamma-ray bursts.

Substantial effort has been put towards explaining these types of observations. The underlying relativistic nonthermal particle acceleration (NTPA) processes have been studied theoretically for decades; proposed mechanisms include collisionless shocks (Blandford and Eichler, 1987), turbulence (Kulsrud and Ferrari, 1971), and magnetic reconnection (Hoshino and Lyubarsky, 2012). Dissipation of magnetised plasma turbulence is a particularly attractive mechanism for explaining particle energy distributions with nonthermal power-law tails. This is because turbulence causes vigorous multiscale plasma motion characterised by power-law scaling of fluctuation strength with respect to the fluctuation length scale, as I will further discuss in Section 1.2. Particles could be energised by interactions with turbulent fluctuations of similar size to the particle Larmor radius (wave-particle interaction), giving rise to a power-law spectrum of particle energies. Furthermore, many of the above-mentioned nonthermal astrophysical systems are known to be manifestly turbulent through resolved high-resolution images (e.g., PWN, Hester 2008; and AGN jets, e.g., the M87 jets).

The most common turbulent NTPA models, which I will examine in Section 1.3, posit that particles gain energy in a stochastic process that can be modelled using a Fokker-Planck (FP) advection-diffusion equation in momentum space (Fermi, 1949; Kulsrud and Ferrari, 1971; Melrose, 1974; Skilling, 1975; Blandford and Eichler, 1987; Schlickeiser, 1989; Miller et al., 1990; Chandran, 2000; Cho and Lazarian, 2006). Stochastic particle acceleration models in relativistic plasma tur-

bulence have been applied to astrophysical systems such as PWN (Bucciantini et al., 2011; Tanaka and Asano, 2017), AGN accretion flows (Dermer et al., 1996; Kimura et al., 2015), AGN jets (Rieger et al., 2007; Asano et al., 2014), and gamma ray bursts (Dermer and Humi, 2001). This sustained relevance makes it important to test these relativistic NTPA theories, and the most accessible and direct method for doing so is by computational simulation and analysis.

Numerical tests of the FP framework for NTPA were originally performed by injecting test particles into fluid, namely magnetohydrodynamic (MHD), simulations (Dmitruk et al., 2003, 2004; Kowal et al., 2012; Lynn et al., 2014; Kimura et al., 2016; Isliker et al., 2017b) or artificially prescribed complex electromagnetic fields (Arzner et al., 2006; O'Sullivan et al., 2009). These test-particle simulations are relatively inexpensive, but have physical limitations such as ad-hoc particle injection and the absence of particle feedback on the fields, which can only be resolved by considering more physically complete simulations.

Recently, first-principles kinetic (and also hybrid, where ions are treated kinetically but electrons as a fluid) particle-in-cell (PIC) simulations have confirmed that turbulence (Kunz et al., 2016; Makwana et al., 2017; Zhdankin et al., 2017, 2018b, 2019; Comisso and Sironi, 2018; Hankla et al., 2022), shocks (Hoshino et al., 1992; Amato and Arons, 2006; Spitkovsky, 2008; Sironi and Spitkovsky, 2011; Marcowith et al., 2016), and relativistic reconnection (Zenitani and Hoshino, 2001; Jaroschek et al., 2004; Lyubarsky and Liverts, 2008; Sironi and Spitkovsky, 2014; Guo et al., 2014, 2016; Werner et al., 2016; Werner et al., 2018; Werner and Uzdensky, 2017; Werner and Uzdensky, 2021; French et al., 2023) can generate efficient relativistic NTPA in collisionless plasma. PIC simulations contain complete microphysical information including the self-consistent trajectories and energy histories of individual particles. However, this wealth of data has not yet been employed directly to test stochastic acceleration models (e.g., FP) or to measure the energy diffusion and advection coefficients.

This dissertation investigates relativistic NTPA in PIC simulations of relativistic electron-positron pair plasma turbulence. The physical regime and computational scheme are selected for the following reasons:

- Relativistic collisionless pair plasmas exist in natural astrophysical systems likely to have nonthermal turbulent particle acceleration such as pulsar wind nebulae.
- Relativistic pair plasma is a particularly computationally feasible regime for PIC and such simulations have been shown to produce nonthermal particle populations (see Section 1.4).
- Kinetic PIC simulations are required to be confident that all collisionless NTPA physics is captured self consistently.
- PIC simulations produce particle trajectories which can be used to create novel tools for analysing NTPA in PIC simulations and connecting with analytical theory to test FP-type models.

The remainder of this Chapter describes the topics mentioned so far in more detail followed by an outline of the rest of the dissertation.

1.2 Magnetised turbulence theory

This Section briefly examines the plasma turbulence theory relevant to the collisionless relativistic pair plasma regime that was described in Section 1.1. The resulting turbulence properties are necessary inputs to the turbulent NTPA theories, as I will review in Section 1.3. The theory also informs the simulation setup (to be described in Section 1.4) in terms of, e.g., the method of driving and necessary physical scales to resolve.

Turbulence is a fluid phenomenon characterised by irregular, unpredictable flow with fluctuations over a wide range of time and space scales (Matthaeus, 2021). This is true for both classical hydrodynamic turbulence and for plasma turbulence. Turbulence is a broad and rich field, but I will focus mainly on the concept of the turbulent energy cascade. This is the process by which excess energy inserted at large spatial scales causes fluctuations that transfer repeatedly to adjacently smaller scales, until reaching the final dissipative scales: an inherently self-similar and spatially multiscale phenomenon. As particle acceleration is essentially a type of dissipation, the turbulent cascade is one of the most relevant turbulence concepts.

Collisionless plasma turbulence partitions into two scale regimes delineated by the microscopic plasma kinetic scales of the Larmor radius $\rho_e \equiv \gamma mc^2/eB$ and plasma skin depth $d_e = \sqrt{\gamma mc^2/4\pi ne^2}$ (where n is the total charged particle density). Fluctuations much larger than the kinetic lengths belong to the inertial range, which is well-described by MHD, while smaller fluctuations require lower-level models such as the Vlasov-Maxwell equations. For orientation, and also to establish the governing equations, I will present the MHD equations and Vlasov-Maxwell equations, before proceeding with the discussion of turbulence theory.

As I will not be working in detail with MHD in this study, I show the simpler non-relativistic MHD equations instead of the relativistic ones. The non-relativistic incompressible MHD equations including viscosity and resistivity in Gaussian units are

$$\rho_m(\partial_t + \mathbf{v}\cdot\nabla)\mathbf{v} = \mathbf{J}\times\mathbf{B}/c - \nabla P + \rho_m\nu\nabla^2\mathbf{v}, \tag{1.1}$$

$$\partial_t\mathbf{B} = \frac{\eta c^2}{4\pi}\nabla^2\mathbf{B} + \nabla\times(\mathbf{v}\times\mathbf{B}). \tag{1.2}$$

The variables are: flow velocity $\mathbf{v}$, current density $\mathbf{J}$, magnetic field $\mathbf{B}$, speed of light c, density ρ_m, pressure P, viscosity ν, and resistivity η. The equation of state $P(\rho_m)$ is in general determined by thermodynamics (it may be adiabatic, isothermal, etc.). For the incompressible case, P is determined to enforce the condition that $\nabla\cdot\mathbf{v} = 0$. Also, $\mathbf{J} = c\nabla\times\mathbf{B}/4\pi$ in the non-relativistic MHD case.

The Vlasov-Maxwell equations describe the time-evolution of the phase-space distribution functions of charged particle species under self-consistent electromagnetic interactions with no collisions (it can be augmented by adding a collision term and is then referred to as the Boltzmann equation, but this is not necessary for my study). The Vlasov-Maxwell system of equations in

Gaussian units is

$$\partial_t f_s = -\mathbf{v}_s \cdot \nabla f_s - Z_s e(\mathbf{E} + \mathbf{v}_s \times \mathbf{B}/c) \cdot (\partial f_s/\partial \mathbf{p}), \tag{1.3}$$

$$\partial_t \mathbf{E} = c\nabla \times \mathbf{B} - 4\pi \mathbf{J}, \tag{1.4}$$

$$\partial_t \mathbf{B} = -c\nabla \times \mathbf{E}, \tag{1.5}$$

$$\nabla \cdot \mathbf{E} = 4\pi\rho, \tag{1.6}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{1.7}$$

$$\rho = e \int d^3p \sum_s Z_s f_s, \tag{1.8}$$

$$\mathbf{J} = e \int d^3p \sum_s \mathbf{v}_s Z_s f_s, \tag{1.9}$$

$$\mathbf{v}_s = \mathbf{p}c/\sqrt{p^2 + m_s^2 c^2}. \tag{1.10}$$

The variables are: time t, position $\mathbf{x}$, species velocity $\mathbf{v}_s$, momentum $\mathbf{p}$, species distribution function $f_s(\mathbf{x}, \mathbf{p}, t)$, species mass m_s, elementary charge e, signed species charge number Z_s, speed of light c, electric field $\mathbf{E}(\mathbf{x}, t)$, magnetic field $\mathbf{B}(\mathbf{x}, t)$, charge density $\rho(\mathbf{x}, t)$, and current density $\mathbf{J}(\mathbf{x}, t)$. The first (kinetic) equation states that the distribution functions evolve according to spatial advection by $\mathbf{v}_s$ and momentum-space advection by the Lorentz force. The remainder are the Maxwell equations along with the usual charge and current density equations and relativistic velocity momentum relation. For pair plasmas, the species charge numbers are 1 and -1 and both species masses are the electron mass m. External forcing can be effected by adding external charges and currents to the right hand sides of (1.8) and (1.9).

I now return to the turbulence theory of the inertial range, governed by MHD. The phenomenology begins with Goldreich and Sridhar (1995), a theory of strong incompressible non-relativistic MHD turbulence. This model employs a turbulent cascade in wavevectors of fluctuations in both the magnetic field and the fluid velocity, set against a dominant background magnetic field. The energy transfer between adjacent scales occurs through the interaction of counter-propagating Alfvén wave packets. An important aspect of the theory is that the turbulent fluctuations are anisotropic depending on whether they are parallel or perpendicular to the background magnetic

field. Due to a timescale matching argument called critical balance, the level of anisotropy is scale dependent, with smaller scale fluctuations becoming progressively elongated along the parallel direction relative to the perpendicular direction. This foundational theory of MHD turbulence has subsequently been refined with concepts such as dynamic alignment (Boldyrev, 2006) and intermittency, which I will not focus on. An extension to ultra-relativistic MHD was performed by Thompson and Blaes (1998), finding that the interactions are dominated by Alfvén waves. Reviews of plasma turbulence include Schekochihin (2022) and Matthaeus (2021).

The kinetic regime theory is less developed but is less important to the study because the inertial range is most relevant for explaining nonthermal particle distributions at high energies. Studies such as Schekochihin et al. (2009) investigate how the turbulent energy cascade can proceed in the kinetic range, which may involve a kinetic Alfvén wave cascade.

1.3 Nonthermal particle acceleration theories

Historically, the study of NTPA was originally motivated by the observation of cosmic rays. The origin of cosmic rays was addressed by Fermi (1949), by a mechanism now called second-order Fermi acceleration, which I describe here largely referencing the corresponding section in Longair (2011).

Second-order Fermi acceleration envisions particles scattering off randomly-moving magnetic clouds, these clouds being effectively enormously massive rigid colliders. A head-on collision between a particle and a cloud causes the particle to gain energy, while a receding collision causes the particle to lose energy. Due to the greater relative velocity, a head-on collision is more probable than a receding collision, and this results in an increase in energy on average. Importantly, the rate of energy gain is proportional to the particle's energy, resulting in an efficient and fast exponential increase which helps to explain the high observed energies of cosmic rays. This energy increase tendency is accompanied by random fluctuation in energy and hence Fermi acceleration is also associated with the term stochastic acceleration. The second-order acceleration is so termed because the particle's average fractional rate of energy gain is proportional to the magnetic cloud velocity

squared.

Fermi acceleration was later modified to the case of particle acceleration in shock waves, termed first-order Fermi acceleration. It was found here that a mechanism could be formulated such that particles only receive head-on collisions, obtaining an average fractional rate of particle energy gain proportional to the shock speed to the first power. Both of these acceleration mechanisms are highly general and apply to particle acceleration situations apart from cosmic rays and also outside of plasma physics.

The original Fermi (1949) calculations focused on the average estimated rate of energy gain. Subsequent, more detailed studies formally derive an energy diffusion-advection (FP) equation embodying the stochastic acceleration process. The key signature for a second-order-Fermi-type process is an energy diffusion coefficient proportional to energy squared. This characteristic (which is a major subject of my investigation) results in the required exponential rate of energy gain. In general, FP equations can be formulated in full vector momentum-space, or reduced coordinates such as cylindrical (parallel and perpendicular momentum), polar (momentum magnitude and pitch angle), or energy-only, obtained through corresponding partial or full isotropy assumptions. This is important for rigorous analytical studies such as Schlickeiser (1989), which calculates momentum-space FP coefficients using quasilinear theory, also obtaining a momentum-squared type diffusion coefficient upon integrating the directional component. These approaches can be combined with quantities such as estimated turbulent spectra to produce concrete, astrophysically-relevant estimates regarding particle acceleration (e.g., Jokipii, 1977). While the discussion here has centred on cosmic ray literature, similar techniques apply to other particle acceleration settings (Demidem et al., 2020).

1.4 Numerical methods

Plasma simulations follow a hierarchy of computational detail parallel to the hierarchy of plasma analytical descriptions (Birdsall and Langdon, 1991). The highest level of abstraction treats the plasma as a fluid, solving the MHD equations numerically, while more detailed and compu-

tationally expensive simulations operate at the kinetic or particle level. Intermediate complexity levels are also available, such as hybrid approaches which treat different components of the plasma by different fluid or kinetic methods. The choice of simulation type depends on the required level of microphysical detail, which then trades off against the scale and complexity of the simulation setup in order to fit within a computational budget.

In this study, a kinetic treatment is essential on physical grounds, and in particular, a PIC treatment is ideal for analysis capabilities. Firstly, in high-energy astrophysical turbulence, plasma is likely to be collisionless and at relativistic temperatures. This makes it likely to develop small-scale structures and phenomena such as current sheets and kinetic instabilities. These can influence the turbulence properties, energy dissipation, and NTPA, hence calling for a kinetic simulation. Secondly, PIC simulations necessarily generate self-consistent particle trajectories and energy histories, which is vital to a rigorous computational analysis of NTPA. Approaches where the particles and fields are not self-consistent, such as injecting test particles into fluid simulations or artificially prescribed fields, inevitably raise questions regarding the absence of particle feedback on the fields and whether the test particle initialisation is statistically representative.

I now briefly outline the PIC technique, drawing from Birdsall and Langdon (1991). The PIC method is a numerically approximate solution of the Vlasov-Maxwell system of equations (1.3–1.10), obtained by evolving a large population of particles over time through self-consistent electromagnetic fields. In PIC simulations, particles have continuous positions and momenta, while electric and magnetic fields are computed only on a discrete grid. The fields are calculated by solving the Maxwell equations on charge and current densities assigned to the grid points from the particle configuration. The fields are interpolated to the individual particle positions to obtain the Lorentz force on each particle, which is then integrated in time to advance the velocity and position. This is followed by the calculation of new fields based on the updated particle distribution, and the cycle repeated iteratively to progress the simulation. This is only an approximate description. In practice, techniques such as interleaving velocity and position time steps (the "leap-frog method") and interleaving electric and magnetic field grid points (the Yee mesh; Yee, 1966) are used to improve

accuracy.

I now describe some typical aspects of PIC turbulence simulations. As mentioned in Section 1.1, relativistic pair plasma is an astrophysically applicable regime. It is also a computationally advantageous regime, as a mass ratio of one reduces the number of scales to separate, and relativistic motion reduces the separation between wave and particle speeds. This has prompted a large number of PIC simulation studies in this parameter space. Nevertheless, due to the computational cost of PIC simulations, it is only feasible to simulate a small segment of plasma with limited or no externally-imposed structure. Hence, the typical simulation is of a box with periodic boundary conditions. A parameter that often varies between different NTPA studies is the type of large-scale turbulent driving. This dissertation considers driven turbulence, as in (Zhdankin et al., 2017, 2018a,b). In these simulations, the turbulence is driven continuously on large scales, comparable to the box size, much larger than kinetic scales. The driving is volumetric (i.e., not at a boundary but everywhere in the volume), exciting several large-scale Alfvén wave modes electromagnetically. Some more specific details about the simulation setup are given in Section 3.3. Other driving schemes include continuously driven imbalanced turbulence (Hankla et al., 2022), and decaying turbulence (Comisso and Sironi, 2019). The explored parameter space also includes radiative turbulence (Zhdankin et al., 2021) and electron-ion plasmas (Zhdankin et al., 2019).

Figure 1.1 shows 3D colour renderings of physical quantities at a snapshot in time on the surface of the largest simulation cube analysed in this dissertation. The upper panel is the relative particle number density and the lower panel is the electric current density component parallel to the background guide magnetic field. The particle density structures tend to be large scale while the current density ones tend to be small scale and thin, clearly showing current-sheet-like structures.

1.5 Outline

The remainder of this thesis is as follows. Chapter 2 is a study demonstrating, for the first time through first-principles simulations, that an energy diffusion-advection model suitably describes nonthermal particle acceleration. This is established primarily by statistically analysing

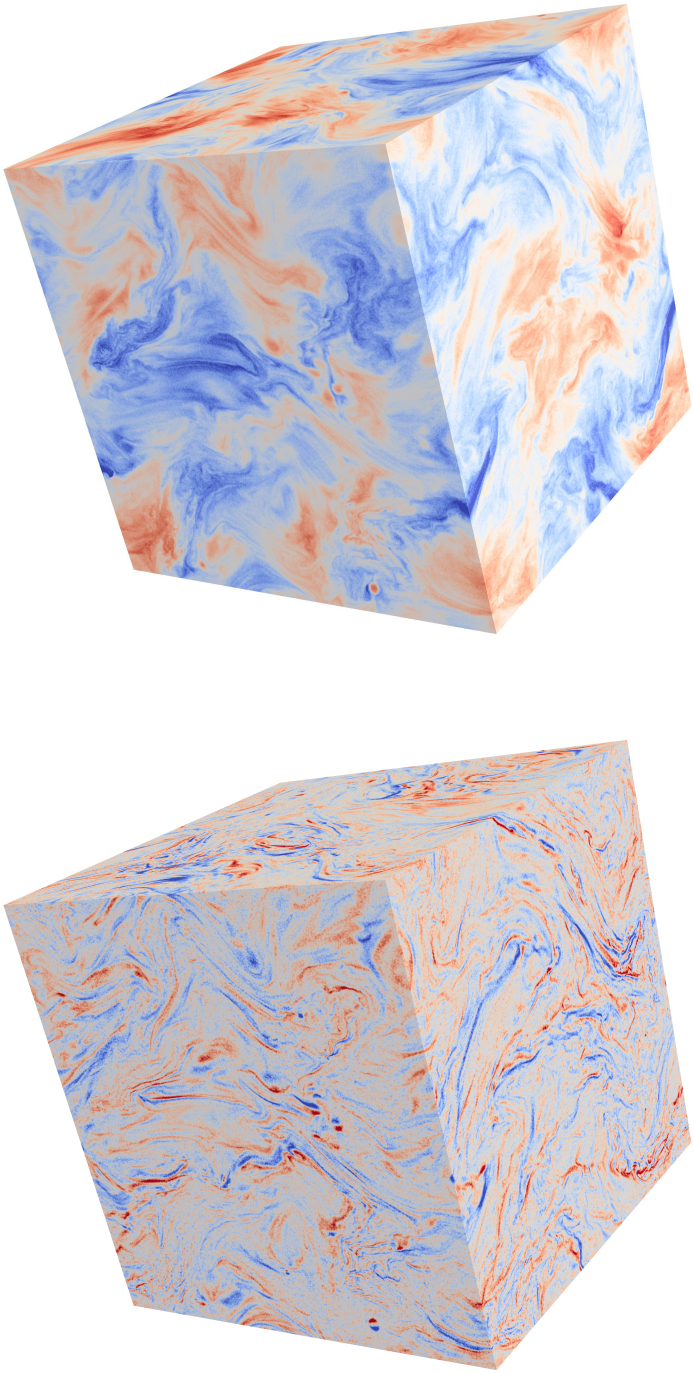

Figure 1.1: 3D colour renderings of quantities on the surface of the simulation cube at a snapshot in time. The upper image is the relative particle number density where red indicates over-density and blue indicates under-density. The lower image is the electric current density component parallel to the background guide magnetic field, with red indicating parallel and blue anti-parallel.

particle trajectories and energy histories from one large simulation. Chapter 2 largely adapts Wong et al. (2020), with the addition of small segments on spatial transport (Subsection 2.3.2) and the effect of the turbulence driving scale (Subsection 2.3.1). Chapter 3 is a more extensive analysis which considers many simulations with different initial system parameters. It shows that the results of Chapter 2 and Wong et al. (2020) hold for a large range of system parameters, and measures the diffusion and advection coefficients in each case, subsequently distilling them into summary parameters that are then compared against analytical theory. This includes a novel first-principles simulation measurement of the diffusion coefficient scaling as a function of instantaneous magnetisation, and also reveals the functional form of the advection coefficient. Furthermore, I measure the nonthermal power-law index, showing that it evolves close to exponentially in time, and relate the power law index to the diffusion and advection coefficients by a simple analytical model. Finally, Chapter 4 summarises the findings and presents a future outlook.

Chapter 2

First-principles demonstration of diffusive particle acceleration

2.1 Introduction

This Chapter is based on Wong et al. (2020). In this Chapter, I use tracked particles to demonstrate stochastic acceleration and directly measure the Fokker-Planck (FP) coefficients in three-dimensional (3D) particle-in-cell (PIC) simulations of driven turbulence in collisionless relativistic plasma. I consider pair plasma both for theoretical and computational simplicity, and for its relevance to high-energy astrophysical systems like pulsar wind nebulae (PWN) and active galactic nuclei (AGN) jets. However, these methods also apply to future investigations of nonthermal particle acceleration (NTPA) in turbulent non-relativistic and electron-ion plasmas, as well as to other processes, e.g., magnetic reconnection.

2.2 Method

I analyze 3D simulations [performed with the PIC code Zeltron (Cerutti et al., 2013)] of externally driven turbulence in relativistic pair plasma (Zhdankin et al., 2018b). I focus on the largest simulation with 1563^3 grid cells and 64 particles per cell (electrons and positrons combined), totalling $\sim 2.4 \times 10^{11}$ particles; smaller simulations give similar results. The simulation domain is a periodic cube of size L, with an initially uniform magnetic guide field $B_0\hat{\mathbf{z}}$. The plasma is initially uniform and isotropic, with total charged particle density n_0, and a Maxwell-Jüttner thermal distribution with a relativistically hot temperature of $T_0 = 100mc^2$, corresponding to the average Lorentz factor $\overline{\gamma}_{\rm init} \approx 3T_0/mc^2 = 300$. The initial magnetization is $\sigma_0 = B_0^2/16\pi n_0 T_0 = 3/8$. In

the fiducial simulation, the normalized system size is $L/2\pi\rho_{e_0} = 163$, where $\rho_{e_0} \equiv \overline{\gamma}_{\rm init} mc^2/eB_0$ is the initial characteristic Larmor radius. Turbulence is electromagnetically driven (TenBarge et al., 2014) and becomes fully developed after several light crossing times (Zhdankin et al., 2018a), with rms turbulent magnetic fluctuations $\delta B_{\rm rms} \sim B_0$. The turbulence is essentially Alfvénic (Zhdankin et al., 2018a), with initial Alfvén velocity $v_{A0}/c \equiv [\sigma_0/(\sigma_0+1)]^{1/2} \simeq 0.52$.

Previous studies (Zhdankin et al., 2017, 2018b) have shown that such turbulence reliably produces nonthermal power-law particle spectra. In this section, I describe my procedure to investigate NTPA in relation to the FP framework. First, I examine gyro-scale oscillations in particle energy and explain their physical origin. Then, I present my methodology for averaging out these oscillations, which is critical for accurately measuring energy diffusion. Finally, I detail the tests of diffusive NTPA in these turbulence simulations, and my procedure for measuring the energy diffusion and advection coefficients as functions of particle energy.

My analysis tracks the positions, momenta, and local electromagnetic field vectors for a statistical ensemble of 8×10^5 randomly-chosen particles. I observe order-unity oscillations in particle energy, γmc^2, at the gyrofrequency, as shown for a representative particle in Fig. 2.1a. The energy oscillates once per gyro-orbit (Fig. 2.1b) because of the large-scale electric field accelerating and decelerating the particle as it gyrates. To describe this analytically, I consider a charged particle moving in constant, uniform electromagnetic fields. I use primed variables for the frame moving with the $\mathbf{E}\times\mathbf{B}$ drift velocity, $\mathbf{v}_D$, given by $c\mathbf{v}_D/(c^2+v_D^2) = \mathbf{E}\times\mathbf{B}/(E^2+B^2)$.

In the primed frame, where $\mathbf{B}'$ and $\mathbf{E}'$ are parallel, the particle gyrates about $\mathbf{B}'$ while being accelerated along $\mathbf{B}'$ by $\mathbf{E}'$. Typically, $E' \ll B'$, and so $\gamma' mc^2$ is slowly-varying on the oscillation timescale. Then, the motion in the primed frame is approximately a simple gyration with $\mathbf{E}' \approx 0$, and, applying the inverse Lorentz transformation, the lab-frame energy can be found:

$$\gamma(t) = \gamma_D \gamma' \left(1 + \beta_D \frac{v'_\perp}{c} \cos\omega' t'\right), \tag{2.1}$$

with $t = t_0 + \gamma_D(t' + \beta_D v'_\perp \sin\omega' t'/\omega' c)$. Here, t is the coordinate time, $v'_\perp$ is the particle's primed-frame velocity perpendicular to $\mathbf{B}'$, $\beta_D = v_D/c$, $\gamma_D = (1 - v_D^2/c^2)^{-1/2}$, $\omega' = eB'/\gamma' mc$ is the

cyclotron frequency with $B' = B/\gamma_D$, and t_0 is a phase. Since I am considering relativistic particles ($v'_\perp \sim c$) and relativistic turbulence ($E_{\rm rms} \sim B_0$ and $\beta_D \sim 1$), the predicted oscillation magnitude is comparable to γ (see Fig. 2.1a), and cannot be ignored.

Without further processing, these energy oscillations are incompatible with a random-walk-based energy-diffusion model. However, I aim to measure the statistical properties of NTPA on the Alfvénic timescale $\sim L/v_A$ relevant to the formation of the nonthermal power law. Since this is generally much longer than the gyroperiod, I analyze just the secular component of the energy histories.

My oscillation removal procedure is informed by Eq. 2.1, which indicates that in the idealized case of uniform constant fields, the secular component of the lab-frame energy is $\gamma_D\gamma'$. However, $\mathbf{v}_D$ fluctuates as the particle traverses small-scale fields, so I average $\mathbf{v}_D$ over the smoothed gyroperiod $2\pi\gamma' mc/eB$, where $\gamma' = \gamma_D\gamma(1-\mathbf{v}_D\cdot\mathbf{v}/c^2)$ is obtained from boosting the lab-frame four-velocity $\gamma\mathbf{v}$ by $\mathbf{v}_D$. I denote by $\langle\mathbf{v}_D\rangle$ the average of $\mathbf{v}_D$ over this smoothed period. I then define the smoothed particle energy to be $\langle\gamma_D\rangle\langle\gamma'\rangle$, where $\langle\gamma'\rangle = \langle\gamma_D\rangle\gamma(1-\langle\mathbf{v}_D\rangle\cdot\mathbf{v}/c^2)$ and $\langle\gamma_D\rangle \equiv (1-\langle\mathbf{v}_D\rangle^2/c^2)^{-1/2}$. This transformed energy (Fig. 2.1a) has greatly reduced oscillations. Thus, this procedure extracts the secular component of particle energy, allowing us to test the FP picture of NTPA. Hereafter, γ and "energy" refer to $\langle\gamma_D\rangle\langle\gamma'\rangle$, except in the case of the overall particle energy distributions (Fig. 2.4a) and magnetic energy spectra (Fig. 2.4b).

I now describe my tests of diffusive acceleration using tracked particles. I use $t_0 = 10.0L/c$ as a fiducial initial time for my measurements, by which point a power-law particle energy spectrum has fully formed. I then bin tracked particles by their energy at t_0 in logarithmic intervals spaced by 10%. For each bin, I measure the standard deviation, $\delta\gamma_{\rm rms}$, and the mean, $\overline{\gamma}$, of the particle energy distribution as a function of subsequent times $\Delta t \equiv t - t_0$. For a classical diffusion process, one expects $\delta\gamma_{\rm rms}(\Delta t) \propto \sqrt{\Delta t}$ if the diffusion coefficient $D(\gamma, t)$ varies slowly compared with Δt and $\delta\gamma_{\rm rms}$.

I then measure the diffusion and advection coefficients, $D(\gamma)$ and $A(\gamma)$, respectively, for the

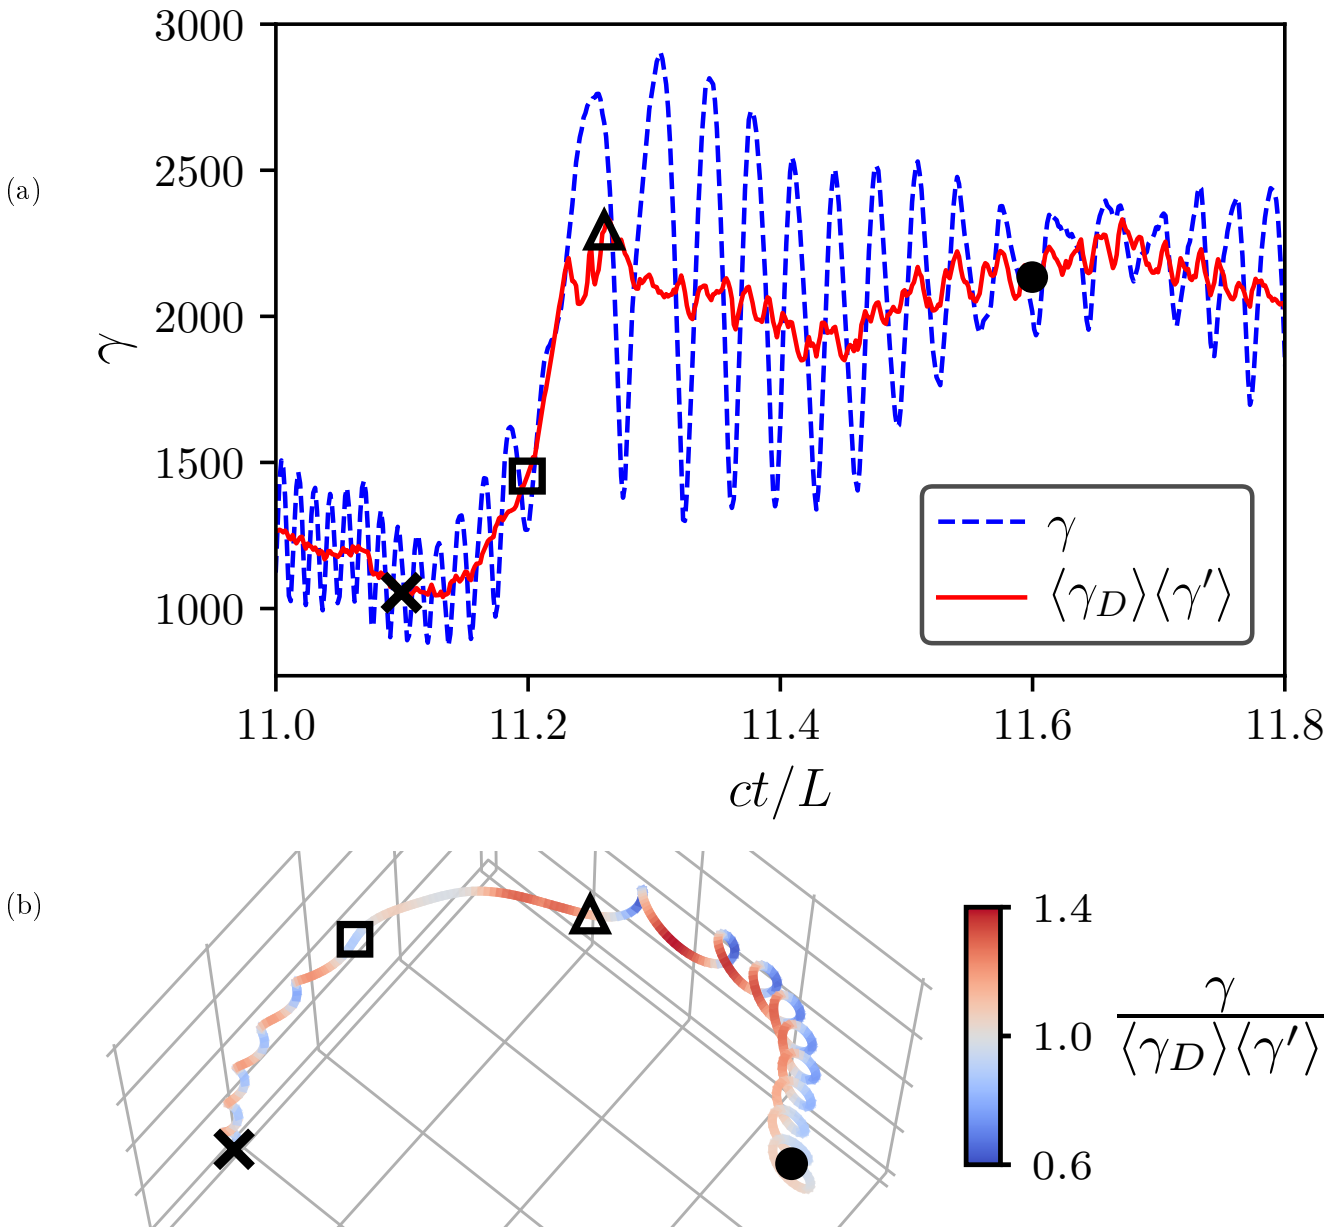

Figure 2.1: (a) Partial energy history of a tracked particle (blue dashed line), showing oscillations which are removed by my transformation (red solid line). (b) Trajectory of the same particle, colored by the instantaneous ratio of the lab-frame energy to the smoothed energy. Markers show time instances in (a) corresponding to particle positions in (b).

simplest FP equation for the energy distribution $f(\gamma, t)$, ignoring pitch angle (with respect to $\mathbf{B}$):

$$\partial_t f = \partial_\gamma (D \partial_\gamma f) - \partial_\gamma (A f). \tag{2.2}$$

Limiting the measurements to times where $\Delta t \lesssim L/v_A$, $\delta\gamma_{\rm rms} \ll \gamma_0$, and $\Delta\overline{\gamma} \equiv \overline{\gamma}(t) - \overline{\gamma}(t_0) \ll \gamma_0$, I approximate the bin distribution as narrow and the coefficients as constant in time. Applying Eq. 2.2, I find:

$$\Delta\overline{\gamma}(\gamma_0, \Delta t) = [\partial_\gamma D|_{\gamma_0} + A(\gamma_0)]\Delta t \equiv M(\gamma_0)\Delta t \tag{2.3}$$

$$\delta\gamma_{\rm rms}(\gamma_0, \Delta t) = \sqrt{2D(\gamma_0)\Delta t}. \tag{2.4}$$

I first measure $D(\gamma)$ and $M(\gamma)$ by applying Eqs. 3.23 and 3.24 to each energy bin and then calculate $A(\gamma) \equiv M - \partial_\gamma D$.

2.3 Results

I first describe the evolution of the overall lab-frame distribution, $f(\gamma)$. Starting from a thermal distribution, $f(\gamma)$ acquires a power-law tail extending to the system-size limit, $\gamma_{\rm max} \equiv LeB_0/2mc^2 \simeq 1.5 \times 10^5$, and gradually hardening over time—its index converges to approximately -3 by $12.3L/c$. At the start of my measurements, $t_0 = 10.0L/c$, the index is approximately -3.2, the peak of $f(\gamma)$ is at $\gamma_{\rm pk} \simeq 520$, and the mean at $\gamma_{\rm avg} \simeq 1170$. Because the system lacks an energy sink, $\gamma_{\rm pk}$ and $\gamma_{\rm avg}$ increase at a rate of about $40c/L$ and $100c/L$, respectively.

I now present tests of energy diffusion. These are extended with more comprehensive diagnostics in Section 3.6. For illustration, Fig. 2.2a shows the evolution of the energy distribution of a single bin of particles with $\gamma_0 = 5 \times 10^3$, deep in the power-law section. I find that $\overline{\gamma} \propto \Delta t$ for small Δt (Fig. 2.2b), while $\delta\gamma_{\rm rms}(\Delta t) \propto \sqrt{\Delta t}$ (Fig. 2.2c), consistent with simple diffusion.

Fig. 2.3 shows $\delta\gamma_{\rm rms}(\Delta t)$ for several bins along with corresponding $\sqrt{\Delta t}$ fits. To avoid artifacts of the smoothing procedure, each fit begins after one gyroperiod $T_{\rm L}(\gamma_0, B_{\rm rms}) \equiv 2\pi\gamma_0 mc/eB_{\rm rms}$. To ensure Eqs. 3.23 and 3.24 are valid, each fit ends when $\delta\gamma_{\rm rms}/\gamma_0$ reaches 0.3, $\Delta\overline{\gamma}/\gamma_0$ reaches 0.1, or $\Delta t = 2L/c$, whichever is earliest. Under these criteria, almost all fits end before $\Delta t = 1L/c$. The

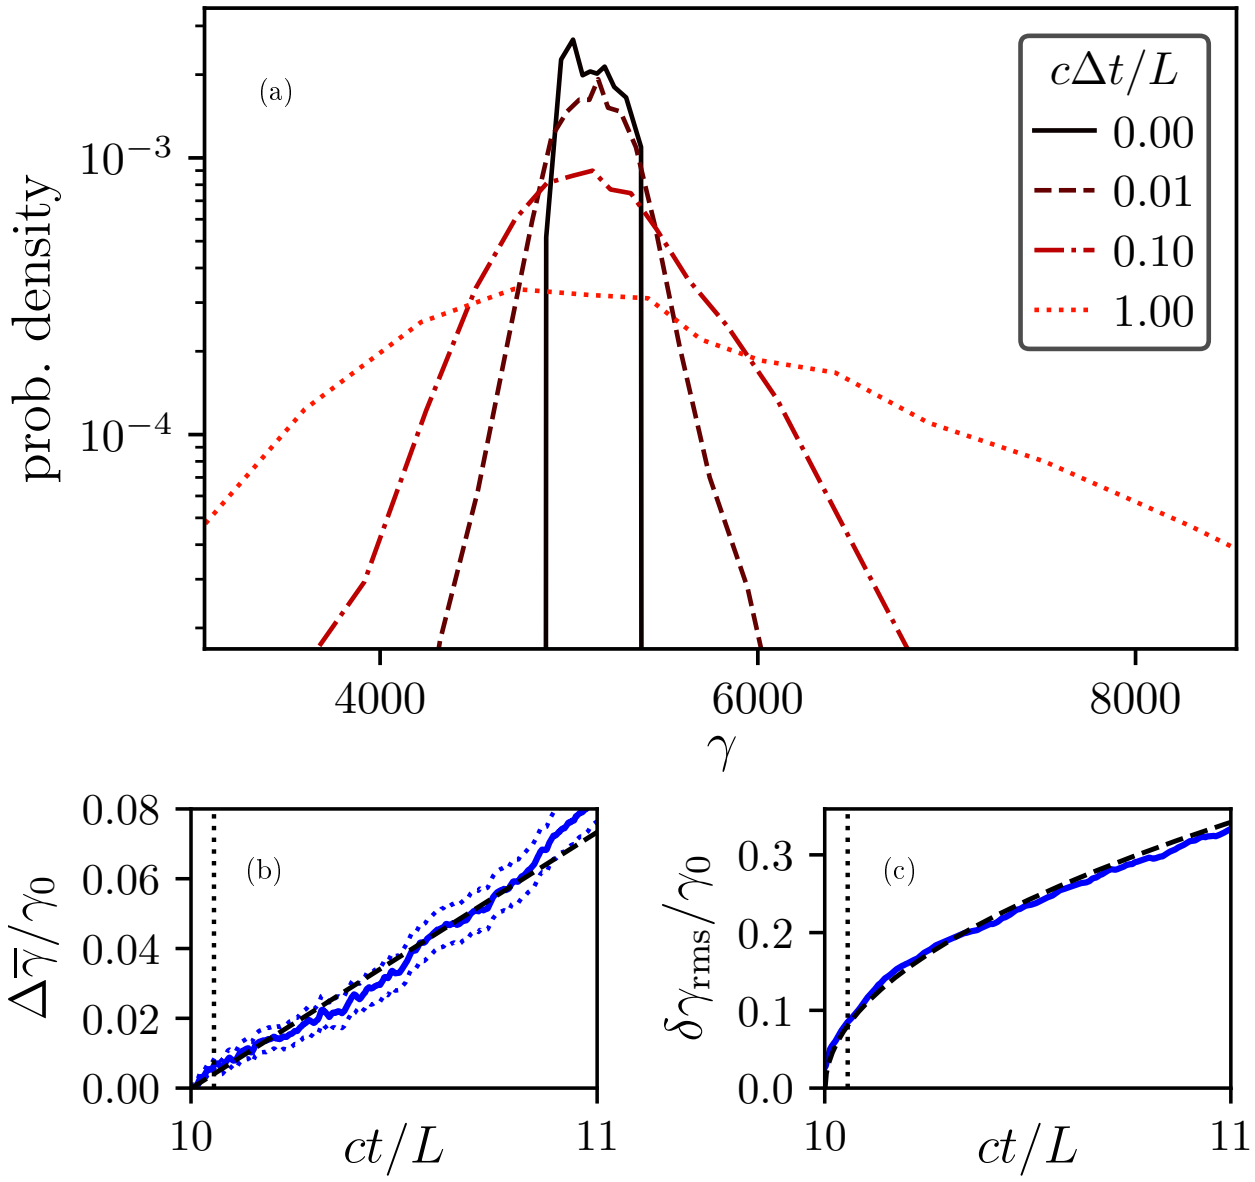

Figure 2.2: (a) Time evolution of the energy distribution for a bin of $N{\sim}2200$ particles with bin-center energy $\gamma_0 = 5 \times 10^3$. For this bin, (b) shows $\overline{\gamma}(t)/\gamma_0$ (solid) with standard error ranges $(\overline{\gamma}(t) \pm \delta\gamma_{\rm rms}(t)/\sqrt{N})/\gamma_0$ (dotted) and a linear Δt fit (dashed), and (c) shows $\delta\gamma_{\rm rms}(t)/\gamma_0$ (solid) with a $\sqrt{\Delta t}$ fit (dashed). In (b) and (c), a vertical dotted line is placed at $\Delta t = T_{\rm L}(\gamma_0, B_{\rm rms})$, the gyroperiod corresponding to γ_0.

fits generally agree well with the data over the fitted intervals. While some of the plotted $\delta\gamma_{\rm rms}(\Delta t)$ have intervals of weakly anomalous energy diffusion these regimes are not my current focus, and I refer to studies of anomalous energy diffusion in plasma turbulence (Isliker et al., 2017b,a). In summary, Figs. 2.2 and 2.3 confirm the expectations of a standard diffusive process, supporting the FP model of turbulent NTPA.

I now report on my measurements of $D(\gamma)$ and $A(\gamma)$. Fig. 2.4c shows $D(\gamma)$, extracted from the fits of $\delta\gamma_{\rm rms}(\Delta t)$ using Eq. 3.24. In the high-energy nonthermal tail ($2\times10^3 \lesssim \gamma \lesssim 3\times10^4$, see Fig. 2.4a), $D = 0.06(c/L)\gamma^2$ is an excellent fit, while for lower energies $\gamma \lesssim \gamma_{\rm pk}$, there is a much shallower scaling roughly consistent with $D \propto \gamma^{2/3}$. I observe that $D \propto \gamma^2$ for particles gyroresonant with fluctuations in the inertial range of the magnetic energy spectrum (Fig. 2.4b), while the lower-energy scaling corresponds to the sub-inertial range of turbulence. In simulations with different magnetization, system size, and number of particles per cell (Section 3.7), the high-energy scaling of $D \propto \gamma^2$ is maintained while the low-energy behavior varies slightly.

The high-energy scaling of $D \propto \gamma^2$ is commonly predicted by NTPA theories (Kulsrud and Ferrari, 1971; Skilling, 1975; Blandford and Eichler, 1987; Schlickeiser, 1989; Chandran, 2000; Cho and Lazarian, 2006). I compare my high-energy fit, $D/\gamma^2 = 0.06c/L$, to the theoretical prediction from 2nd-order Fermi acceleration, $D/\gamma^2 = u_A^2/3c\lambda_{\rm mfp}$, for ultra-relativistic particles interacting with isotropic scatterers moving at the Alfvén velocity, where $u_A = v_A/(1-v_A^2/c^2)^{1/2}$ and $\lambda_{\rm mfp}$ is the mean free path between scattering events (Blandford and Eichler, 1987; Longair, 2011). At $t_0 = 10.0L/c$, $v_A = 0.51c$, giving a theoretical scaling of $D/\gamma^2 = 0.12c/\lambda_{\rm mfp}$, which agrees with my fit if $\lambda_{\rm mfp} \sim 2L$. I also find, from simulations with reduced driving wavelength (Subsection 2.3.1), that the turbulence driving scale, and not the box size, sets D/γ^2 (and hence $\lambda_{\rm mfp}$). It is also possible to directly probe $\lambda_{\rm mfp}$; I show a preliminary investigation in Subsection 2.3.2. While the value of D/γ^2 is consistent with 2nd-order Fermi acceleration from gyroresonant scattering of particles by Alfvén modes, further work is needed to test this view. This includes studying the effect of varying magnetization (Subsection 3.7.1), as well as more comprehensive analysis of the spatial transport characteristics.

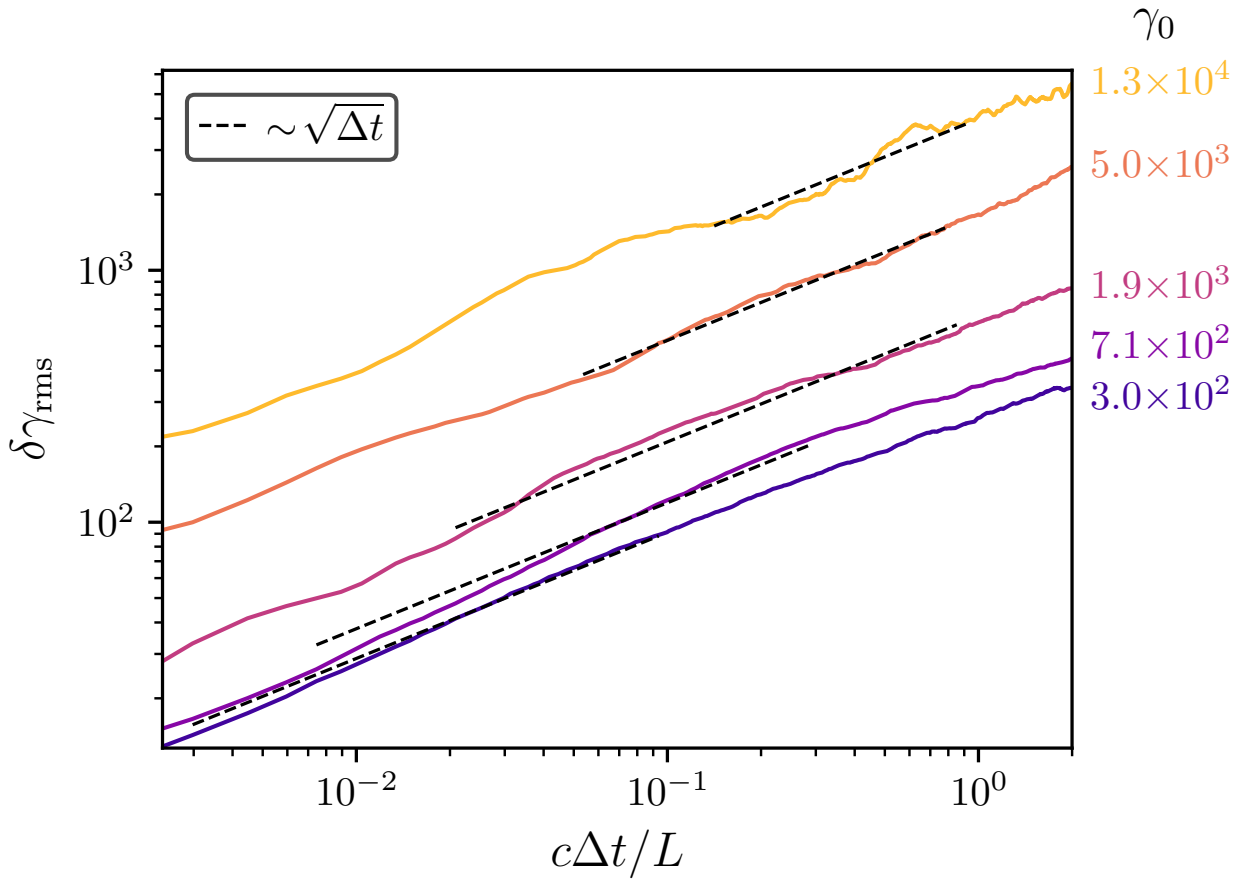

Figure 2.3: The standard deviation, $\delta\gamma_{\rm rms}$, of the particle energies in several bins (solid lines), with corresponding $\sqrt{\Delta t}$ fits (dashed lines), each annotated by the initial bin-center energy γ_0. Each fit is drawn over the corresponding fitting interval, which begins at $\Delta t = T_{\rm L}(\gamma_0, B_{\rm rms})$, the gyroperiod for the bin-center energy of the corresponding bin.

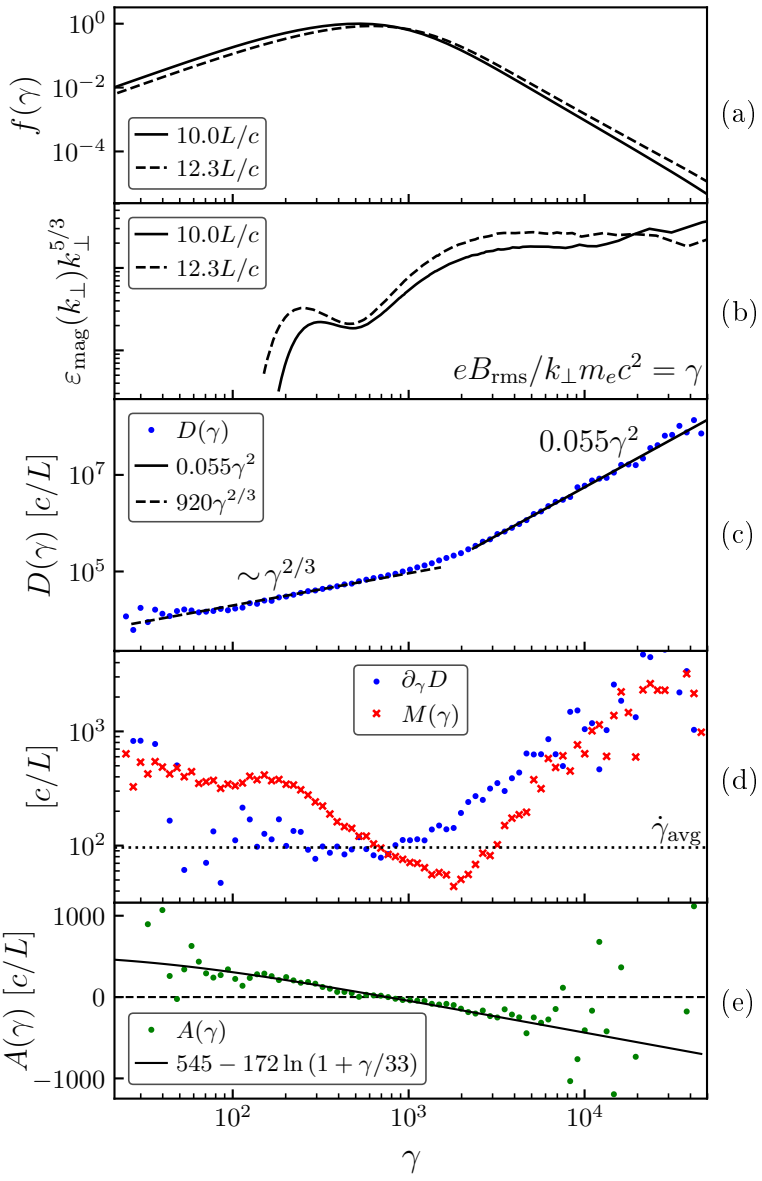

Figure 2.4: (a) The overall particle energy distribution $f(\gamma)$ at the start of the measuring interval ($t = 10.0L/c$, solid line), and a short time later ($t = 12.3L/c$, dashed line). (b) The magnetic energy spectrum $\epsilon_{\rm mag}(k_\perp)$ compensated by $k_\perp^{5/3}$ vs. $eB_{\rm rms}/k_\perp mc^2$, the Lorentz factor corresponding to a perpendicular gyroradius equal to the inverse of the perpendicular (to $\mathbf{B_0}$) wavenumber $k_\perp$. (c) The diffusion coefficient $D(\gamma)$ (dots), with power-law fits of index 2 (solid line) in the nonthermal region, and index 2/3 (dashed line) in the low-energy region. (d) The acceleration rate $M(\gamma)$ (crosses) and the contribution to $M(\gamma)$ by $\partial_\gamma D$ (dots). For reference, the overall average rate of energy gain $\dot\gamma_{\rm avg} \simeq 100c/L$ is shown as a dotted line. (e) The advection coefficient $A(\gamma) \equiv M - \partial_\gamma D$ (dots) with logarithmic fit (solid line).

To compare the effects of the first- and second-order derivative terms in Eq. 2.2, I separate the contributions of A and $\partial_\gamma D$ to the average acceleration rate $M(\gamma) = \partial_t \overline{\gamma}$. I extract $M(\gamma)$ from linear fits of $\overline{\gamma}(\Delta t)$ [see Eq. 3.23], using the same time intervals as those used for fitting $\delta\gamma_{\rm rms}(\Delta t)$ to measure $D(\gamma)$. As shown in Fig. 2.4d, $M(\gamma)$ is positive, as expected with external energy injection, and has a minimum near $\gamma_{\rm avg} \approx 1200$.

I then compute the contribution of energy diffusion to acceleration, $\partial_\gamma D$, (Fig. 2.4d) and subsequently the advection coefficient $A(\gamma) \equiv M - \partial_\gamma D$ (Fig. 2.4e). In the high-energy power-law section ($\gamma > \gamma_{\rm pk}$), A is negative, while for low energies ($\gamma < \gamma_{\rm pk}$), A is positive. Overall, A tends to pull particle energies towards $\gamma_{\rm pk}$, narrowing $f(\gamma)$. I find that A is reasonably approximated by a logarithmic scaling with energy. For a momentum-space FP equation containing only diffusion [see, e.g., Ramaty (1979)], the expected energy-space advection coefficient for Eq. 2.2 is $A = 2D/\gamma$. The measured negative logarithmic scaling thus implies that the momentum-space FP equation contains a significant advection term. In the nonthermal section, the magnitude of A is generally smaller than, but still comparable to that of $\partial_\gamma D$. Hence, the evolution of the nonthermal population in this simulation cannot be interpreted as being due to D alone. This also complicates estimates of the acceleration time based on D/γ^2. The rate of overall energy increase from advection is $\int d\gamma A f \approx -23c/L$ while that from diffusion is $\int d\gamma (\partial_\gamma D) f \approx 140c/L$. The general effects of systematic and stochastic acceleration (due to A and D respectively) are discussed in, e.g., Pisokas et al. (2018) and Vlahos and Isliker (2019). I note that my measurement of A has considerable scatter outside of the central range, $10^2 \lesssim \gamma \lesssim 10^4$, as it depends on the difference between two noisy quantities.

Finally, I test whether the FP equation with coefficients measured by my methodology can reproduce the overall evolution of $f(\gamma, t)$ from the PIC simulations. I first repeat the measurements of D and A at almost 200 different t_0 evenly spaced over the entire simulation, thus obtaining time-dependent coefficients. I then insert the measured coefficients into the FP equation [Eq. 2.2] and solve it numerically using a finite-volume method. I use linear interpolation in γ and nearest-neighbor interpolation in t. The FP coefficients are extrapolated as constant in γ for energies without

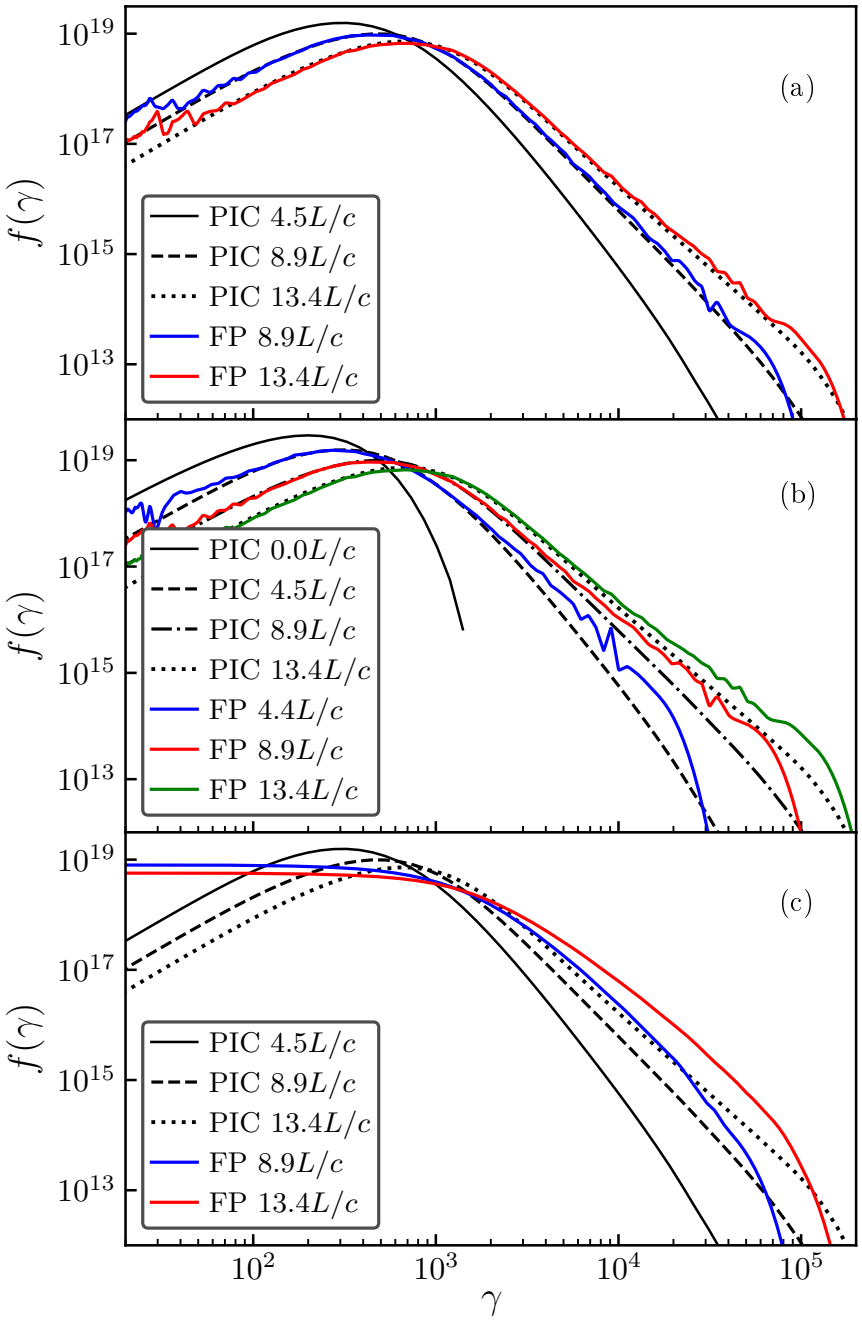

Figure 2.5: (a) Particle distributions $f(\gamma)$ from the PIC simulation at $ct/L \in \{4.5, 8.9, 13.4\}$ (black solid, dashed, and dotted lines, respectively), and the FP solution (initialized at $ct/L = 4.5$) at $ct/L \in \{8.9, 13.4\}$ (red and blue solid lines, respectively). (b) Similar to (a) but for an FP solution initialized at $ct/L = 0$. (c) Similar to (a) but with the advection coefficient artificially set to zero.

enough particles to measure them. Comparing with the particle energy distributions obtained from the PIC simulation, I find that the FP equation gives excellent agreement at all subsequent times when initialized with the corresponding PIC distribution as early as $t = 4.5L/c$ (Fig. 2.5a), with moderate errors if initialized at $t = 0$ (Fig. 2.5b). This may be due to the turbulence not being fully developed at early times. Finally, evolving from $t = 4.5L/c$ with A set to zero shows a clear mismatch with the PIC distributions (Fig. 2.5c), giving us confidence that the measurements of A are reasonably accurate and that A significantly affects the evolution of $f(\gamma, t)$. Thus, in this regime of fully-developed turbulence, I find that the FP equation with my coefficient measurement methodology is appropriate for modeling NTPA. I leave further investigation of the early-time behavior to future work.

2.3.1 Driving scale dependence

This Subsection present results from simulations with the driving wavelength $L_d = 2\pi/k_d$ reduced relative to the box size L. I define the prefactor of the diffusion coefficient γ^2 scaling in the nonthermal range to be $D_0 \equiv D/\gamma^2$. Second-order Fermi acceleration theory predicts D_0 to be inversely proportional to the mean free path between scattering events, λ_{mfp}. The two natural length scales with which λ_{mfp} could possibly be identified are L and L_d. Thus, this simulation series tests the relationship of λ_{mfp} to L and L_d.

Figure 2.6 plots D_0 against instantaneous system magnetisation $\sigma(t)$ for four simulations corresponding to the four combinations of $L/\rho_{e0} \in \{256, 512\}$ with either full or halved driving scale $L_d \in \{L, L/2\}$. When D_0 is plotted in units of c/L, the full and half driving scale points form separate bands, whereas when D_0 is plotted in units of c/L_d, the scatter is reduced and all of the points fall within one band. This indicates that the main determiner of the effective scattering mean free path is the driving scale L_d, and not the box size L. This is consistent with an interpretation of the scattering length being the length scale for which the fluctuations in the magnetic field become comparable to the magnitude of the field itself.

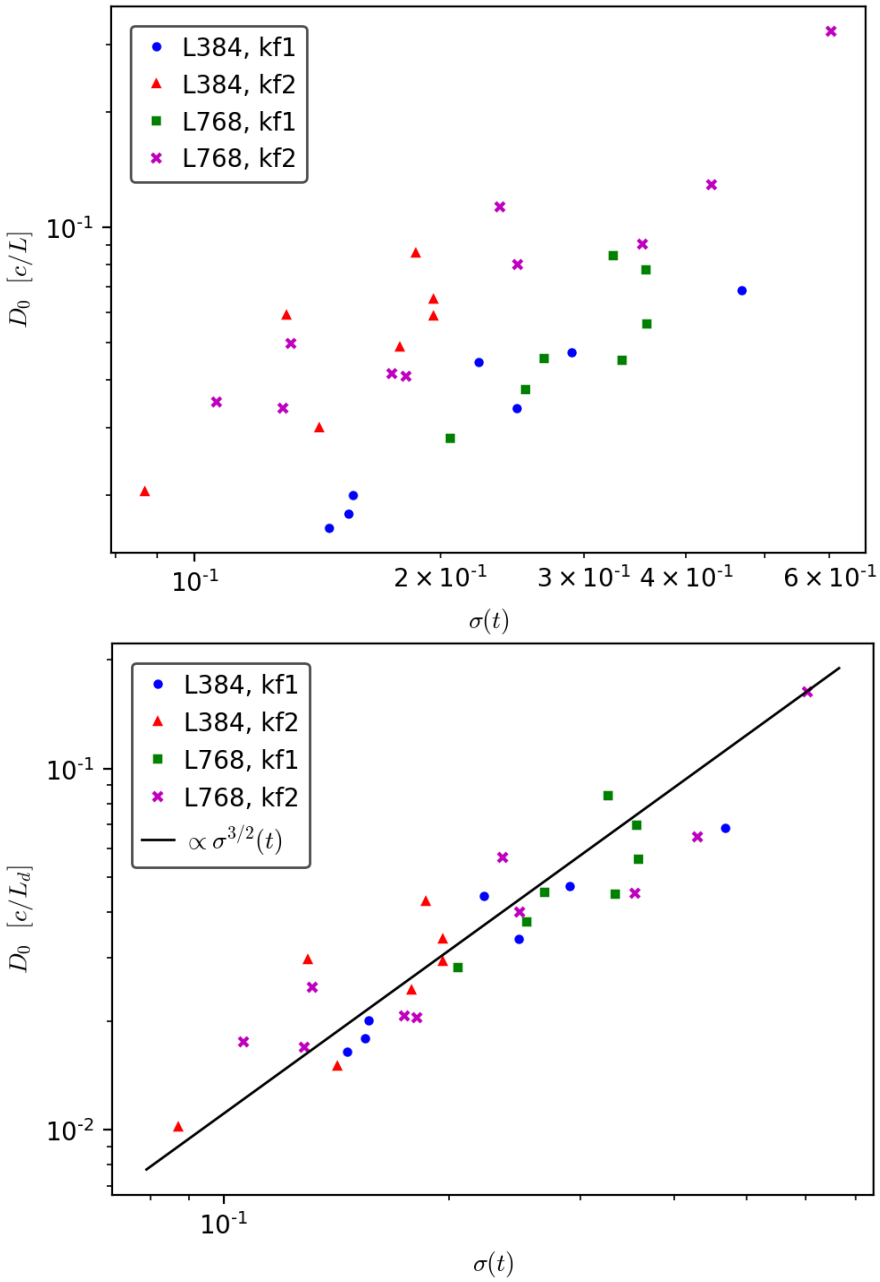

Figure 2.6: Driving scale dependence of $D_0 \equiv D/\gamma^2$ in the high-energy tail. The left panel has D_0 in units of c/L whereas the right panel has D_0 in units of c/L_d. kf1 and kf2 indicate $L_d = L$ and $L_d = L/2$ respectively.

2.3.2 Spatial transport

This Subsection studies $\lambda_{\rm mfp}$ by inspecting the spatial trajectories of tracked particles. The quantity of interest here is $\delta\mathbf{r}_{\rm rms}(\Delta t)$, the vector root mean square displacement of tracked particles between times t_0 and $t_0 + \Delta t$. Figure 2.7 shows $\delta\mathbf{r}_{\rm rms}(\Delta t)$ for all tracked particles in the $N = 1536$, $\sigma_0 = 3/8$ simulation, beginning at $t_0 = 9L/c$ and $t_0 = 10L/c$. At both of these times, turbulence should be fully developed. The spatial transport is anisotropic, with significantly different behaviour in the direction parallel to the guide field. I observe nearly ballistic transport (rms spread $\propto \Delta t$) in all directions up to $\Delta t \sim 0.3L/c$. This elapsed time is significantly larger than the gyro-period for the average particle energy, which is around $0.01L/c$. The overall speed in this phase is roughly $0.32c$ in each direction, corresponding to an overall magnitude of $\sqrt{3 \cdot 0.32^2} \sim 0.55c$, which is close to the Alfvén speed, suggesting an interpretation as the ballistic motion of bulk flows or guiding centres. Motion in the z-direction remains essentially ballistic through to $\Delta t \sim 4L/c$. This is consistent with a very large mean free path in the parallel direction. For the x- and y-directions, the rms spread transitions from ballistic, through a possibly diffusive region, before plateauing at a value somewhat smaller than the system size. The confinement of particles to displacements significantly smaller than the simulation size could be caused by large scale circularisation of bulk flows by the background field.

To obtain further insight, I examine in Figure 2.8 the spatial rms spread for a bin of particles with $3000 \leq \gamma < 3300$, roughly in the middle of the nonthermal range. In comparison to Figure 2.7, I observe an additional ballistic phase. This first phase corresponds to isotropic ballistic motion at the speed of light ($\sqrt{3 \cdot 0.58^2} \sim 1$), i.e., of the particles themselves, which terminates at the average gyroradius of that bin. The second ballistic phase corresponds to Alfvénic ballistic motion of the guiding centres. The third stage is where the z component deviates from the xy components. The xy components pass through a possible diffusive stage before becoming subdiffusive and plateauing. Meanwhile, the z component remains closer to ballistic than diffusive for the remainder of the plotted duration.

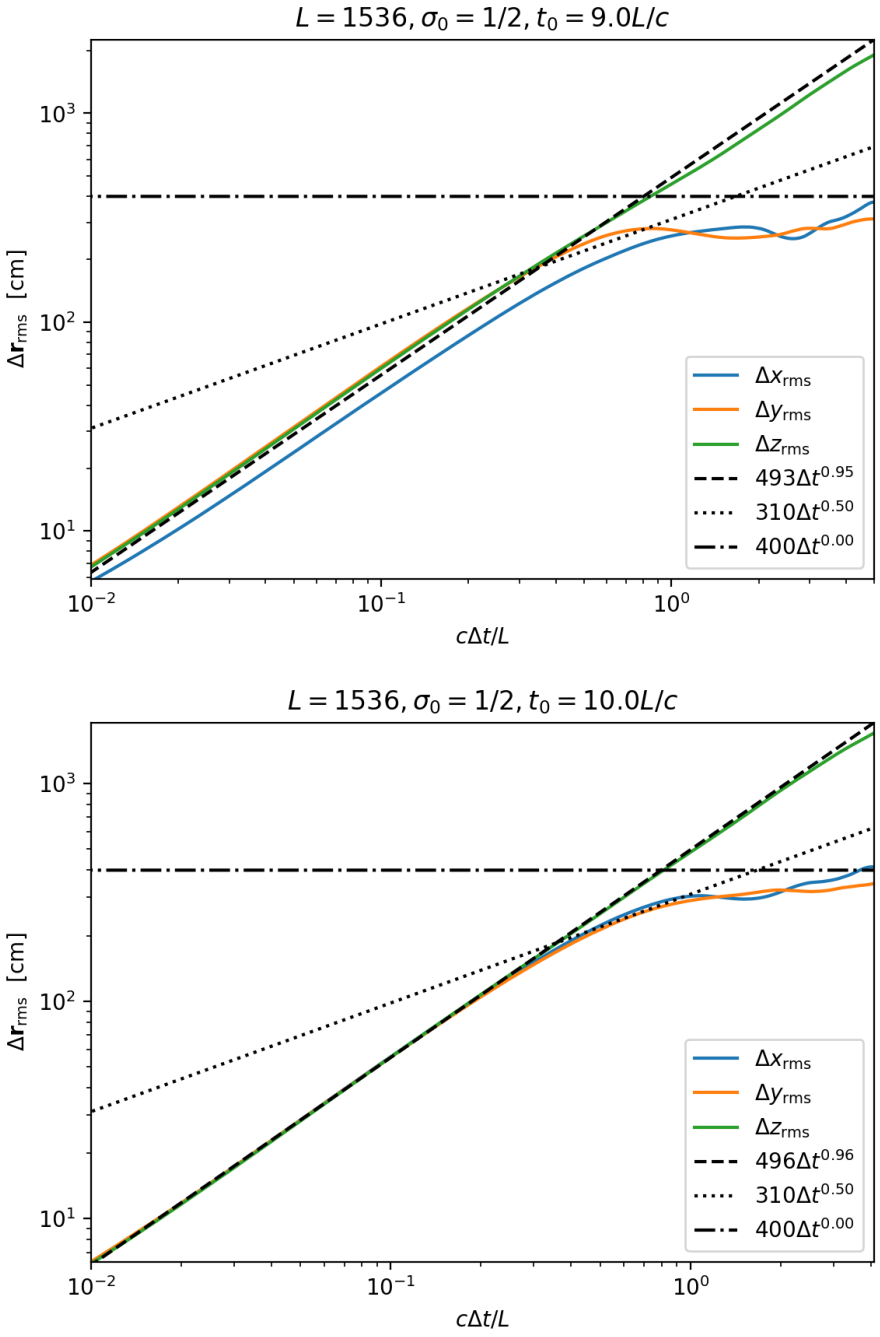

Figure 2.7: Spread of spatial displacements for all tracked particles in the $N = 1536$, $\sigma_0 = 3/8$ simulation. The y-axis units are $1\text{cm} = \Delta x = L/N$.

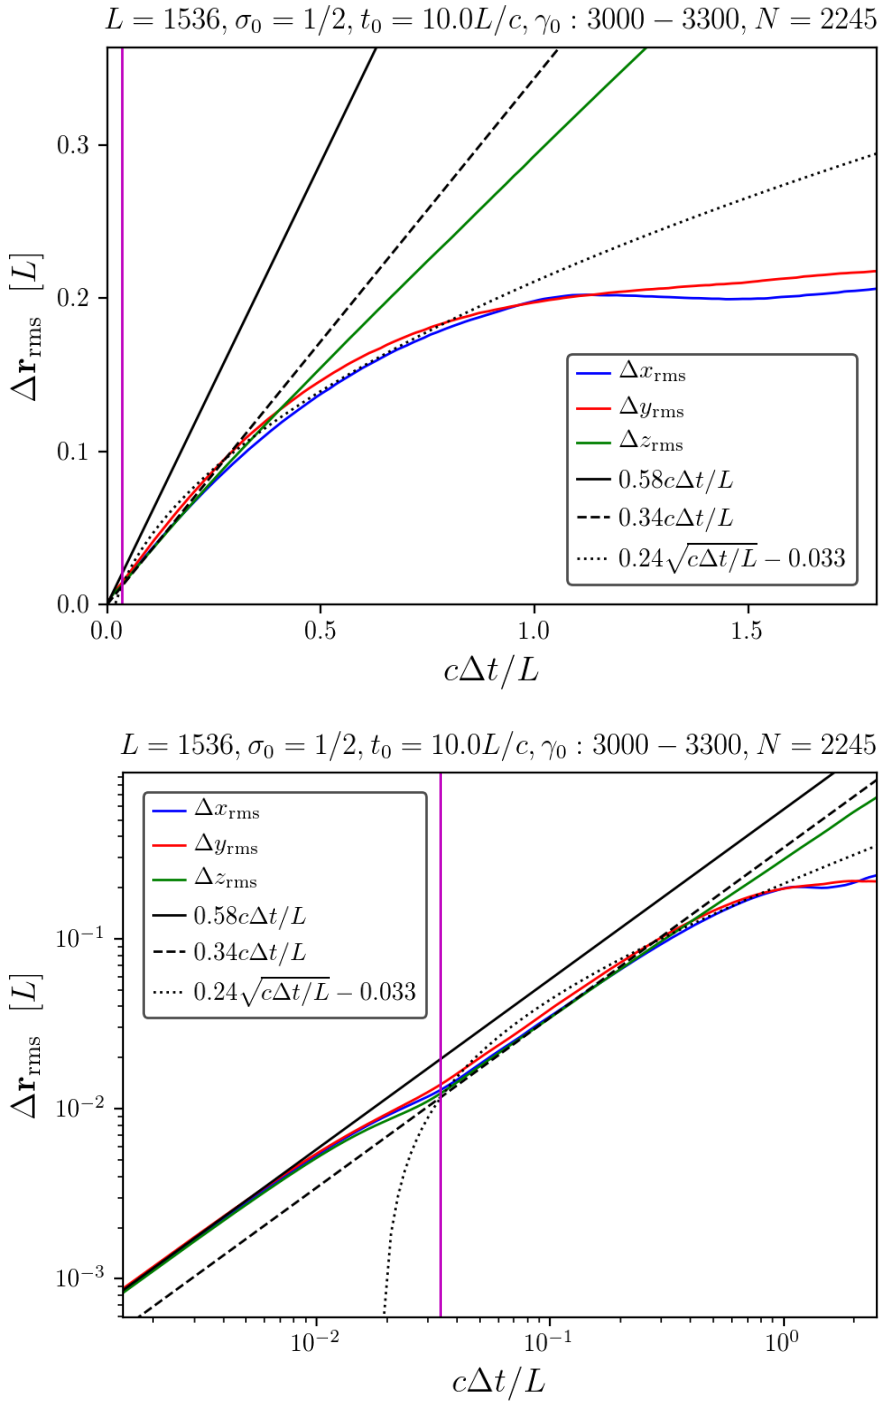

Figure 2.8: Spatial root-mean-displacement for tracked particles with $3000 \leq \gamma < 3300$ in the $N = 1536$, $\sigma_0 = 3/8$ simulation. The pink vertical line is drawn at the average gyroradius of this bin of particles. The top panel has linear scales to show the early behaviour while the bottom panel has logarithmic scales to show the different stages of spatial transport as explained in the main text.

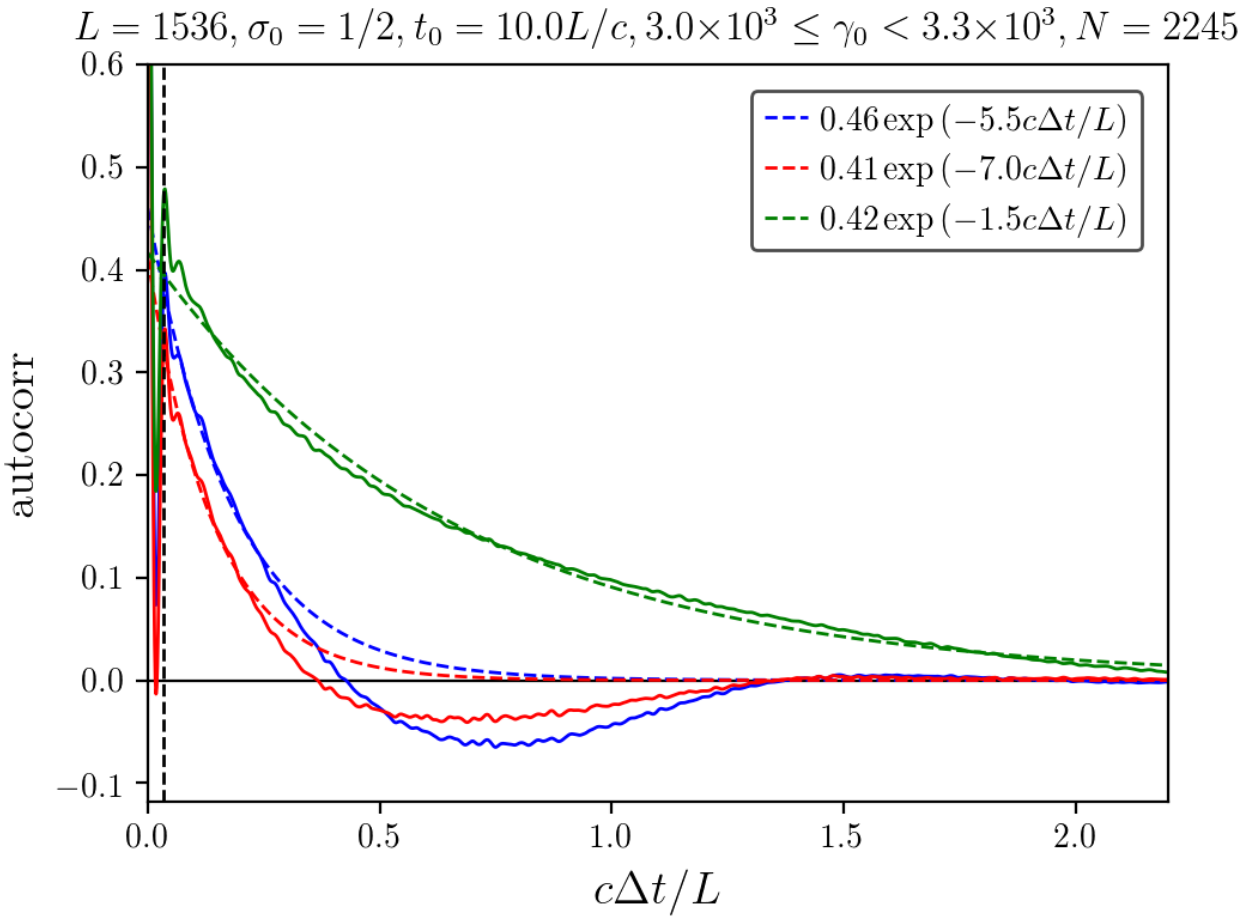

Figure 2.9: Velocity autocorrelations for tracked particles with $3000 \leq \gamma < 3300$ in the $N = 1536$, $\sigma_0 = 3/8$ simulation. The dashed line is placed at $c\Delta t/L$ equal to the bin average gyroperiod.

The perpendicular rms displacement spread transitions from ballistic to essentially constant over a small window such that it is difficult to identify a significant period where it scales with the square root of time (i.e. diffusively). Therefore it is hard to identify a characteristic scattering length or mean free path. In the direction along the background field, the particles are essentially unconfined, and all that can be said is that the scattering length in z is at least as large as the system size. This would be consistent with the implication by the Section 2.3 results that the free path is on the order of the system size or larger.

Supposing that the perpendicular components undergo a diffusive phase, I estimate the mean free path based on the square-root fit in Figure 2.8, using the relation $\Delta x_{\mathrm{rms}} = \sqrt{\lambda_x v_x t}$ (and similar for y). I assume that the secular speed is that of the guiding centres in the second ballistic phase: $v_x = v_y = 0.34c$. This gives $\lambda_x = \lambda_y = 0.16L$. The parallel mean free path λ_z is difficult to estimate using Δz_{rms}, because it shows no signs of a diffusive regime in the viewed timeline. It is possible that λ_z is under-resolved for simulations with $L_d = L/2\pi$, possibly requiring $L_d = L/4\pi$ or better.

Another possible method for obtaining λ_z is through velocity autocorrelations, which are shown in Figure 2.9 for all three velocity components. Each autocorrelation series displays oscillations at the gyro period due to gyromotion, overlaid on a longer-term damping. Using exponential fits starting from after these oscillations, I find that the decorrelation times for the perpendicular and parallel velocity components are respectively $\sim 0.2L/c$ and $\sim 0.7L/c$. Combining with the secular velocity of the guiding centres, this suggests decorrelation lengths $\lambda_x = \lambda_y \sim 0.34c \cdot 0.2L/c = 0.07L$ and $\lambda_z \sim 0.34c \cdot 0.7L/c = 0.24L$. The mean free path suggested by the D/γ^2 measurements in Section 2.3 is roughly $2L$. This is within an order of magnitude to the dominant λ_z.

2.4 Conclusions

In this study I rigorously demonstrate, for the first time, diffusive nonthermal particle acceleration (NTPA) in first-principles particle-in-cell (PIC) simulations of driven relativistic plasma turbulence, through direct statistical measurements using large numbers of tracked particles. I introduce a procedure to suppress large-amplitude gyro-oscillations of particle energy, which is critical

for revealing the diffusive nature of NTPA and measuring the Fokker-Planck (FP) coefficients. I find that the energy diffusion coefficient D scales with particle energy γmc^2 as $D \simeq 0.06(c/L)\gamma^2$ in the high-energy nonthermal power-law region, in line with theoretical expectations (Kulsrud and Ferrari, 1971; Skilling, 1975; Blandford and Eichler, 1987; Schlickeiser, 1989; Chandran, 2000; Cho and Lazarian, 2006), while there is a much shallower scaling at energies below the peak of the energy distribution. I also measure the energy advection coefficient $A(\gamma)$, though with more uncertainty. I find that A is not negligible, and tends to narrow the distribution by accelerating low-energy particles and decelerating high-energy particles. Furthermore, a numerical solution of the FP equation with the measured coefficients reproduces the evolution of the particle energy spectrum from the PIC simulation over a significantly longer time interval than was used for measuring the coefficients. This suggests that this simple FP model can fully account for NTPA in these simulations. These results thus lend strong first-principles numerical support to a broad class of turbulent NTPA theories.

My new methodology can also be applied to future tracked-particle studies of NTPA in other contexts such as shocks or magnetic reconnection, and over broader ranges of physical regimes. Future work may investigate the effects of system parameters such as magnetization, plasma beta, and guide field strength; radiative cooling; relativistic vs non-relativistic regimes; and plasma composition (e.g., pair vs electron-ion plasma). In addition, the analysis can be extended to include pitch-angle dependence and scattering. Characterizing stochastic NTPA in various regimes has important implications in a broad range of contexts such as solar flares, the solar wind, pulsar wind nebulae, active galactic nuclei, gamma ray bursts, and cosmic ray acceleration in supernova remnants. Thus, this study will facilitate more detailed tests of NTPA theories against PIC simulations exploring various physical situations, thereby advancing our understanding of space, solar, and high-energy astrophysical phenomena.

Chapter 3

Energy diffusion and advection coefficients and their relation to the nonthermal power-law tail

3.1 Introduction

This Chapter is adapted from my paper in progress. Proposed NTPA mechanisms include shocks, magnetic reconnection, and turbulence, and in all of these cases the predominant theoretical picture is one in which a particle's interactions with scattering structures such as shock fronts and turbulent eddies cause its energy to accumulate in a biased random walk. The resulting energisation is captured through the Fokker-Planck (FP) diffusion-advection equation in momentum space:

$$\frac{\partial F}{\partial t} = \frac{\partial}{\partial \boldsymbol{p}} \cdot \left(\mathbf{D}_{pp} \cdot \frac{\partial F}{\partial \boldsymbol{p}} \right) - \frac{\partial}{\partial \boldsymbol{p}} \cdot (\boldsymbol{A}_p F), \tag{3.1}$$

where $F(\boldsymbol{p}, t)$ is the (spatially averaged) particle momentum distribution, and $\mathbf{D}_{pp}(\boldsymbol{p}, t)$ and $\boldsymbol{A}_p(\boldsymbol{p}, t)$ are respectively the momentum-space diffusion coefficient tensor and advection coefficient vector. Differences between nonthermal particle acceleration (NTPA) theories are encapsulated in their predictions for the FP coefficients. These FP models of particle acceleration are used widely in space physics and astrophysics (e.g., Miller et al., 1990; Nayakshin and Melia, 1998; Blasi, 2000; Summers and Ma, 2000; Becker et al., 2006; Liu et al., 2009; Mertsch and Sarkar, 2011; Asano et al., 2014; Kimura et al., 2015; Bian et al., 2014; Kundu et al., 2021). Numerical simulations offer the opportunity to first test the diffusive nature of particle energisation, and to then measure the FP coefficients to constitute a test of the underlying theory. However, such computational analysis has not yet been conducted extensively, with just a handful of preliminary pioneering studies published

recently (as described below).

Magnetised turbulence has been studied both with magnetohydrodynamic (MHD) simulations, and more recent kinetic simulations. To study particle acceleration in MHD simulations, test particles are inserted and their trajectories are followed through the simulated fields (e.g., Dmitruk et al., 2003, 2004; Kowal et al., 2012; Lynn et al., 2013, 2014; Kimura et al., 2016, 2019; Medina-Torrejón et al., 2021; Sun and Bai, 2021; Zhang and Xiang, 2021; Bresci et al., 2022). This method is also applied to cases where the fields are prescribed and not simulated (e.g., Vlahos et al., 2004; Isliker et al., 2017a). While less computationally expensive than particle-in-cell (PIC) simulations, this approach suffers from serious drawbacks such as arbitrary particle injection and non-self-consistent particle motion. Kinetic effects could change the acceleration process qualitatively, by influencing the injection of particles into the acceleration process, and backreaction of energetic particles onto the electromagnetic fields. In contrast to MHD, first-principles PIC simulations naturally provide a unique capability for detailed diagnostics through the complete self-consistent history of a particle's trajectory in phase space. However, such simulations are still relatively recent, and initial analysis has been focused on demonstrating the presence of NTPA through largely global diagnostics such as the existence of a nonthermal power-law particle energy distribution (e.g., Zhdankin et al., 2017, 2018b; Comisso and Sironi, 2018; Hankla et al., 2022; Vega et al., 2022; Nättilä and Beloborodov, 2022; Meringolo et al., 2023). Hence there is an opportunity for more detailed tests of NTPA theories through direct inspection of tracked particles.

My previous paper, Wong et al. (2020), described in Chapter 2 of this thesis, is a numerical analysis of an ensemble of tracked particles in a large three-dimensional (3D) PIC simulation which established for the first time that the FP framework was indeed suitable for modelling NTPA in a kinetic simulation of driven relativistic pair-plasma turbulence. I found that the simple diffusion-advection model works even in fully-kinetic simulations of strong turbulence, providing firm first-principles computational evidence of the validity of FP-type particle acceleration models. Recent work by Comisso and Sironi (2019) provided further evidence for diffusive acceleration in PIC simulations (of decaying relativistic turbulence) and found that the diffusion coefficient scales with

magnetization in a way that is consistent with a second-order Fermi acceleration process.

In this Chapter, I further test the FP model and conduct a detailed study of time-dependent FP energy diffusion and advection coefficients in PIC simulations of turbulent relativistic pair plasma, while varying the physical system parameters of the initial plasma magnetisation and the simulation box size relative to the initial average Larmor radius. Since the plasma magnetization and the average Larmor radius (and hence the scale separation between the macroscopic driving scale and the plasma kinetic microscales) evolve in time in these nonradiative simulations due to continuous turbulent heating, I can investigate the dependence of the FP coefficients not only on the initial values of these physical parameters, but also on their time-dependent instantaneous values. In addition, I measure the power-law index of the nonthermal part of the particle energy distribution as a function of time, and relate it to the FP coefficients with an analytical model. I thereby uncover how the power-law index depends on both initial and instantaneous parameters, which results in a more complete understanding of how nonthermal distributions come about in these simulations. These insights, in turn, inform FP-based models of NTPA in space and astrophysical systems, including those used in conjunction with global MHD simulations, and thus have important observational implications. In particular, they provide a solid, first-principles physics basis for formulating concrete usable prescriptions for spectral modelling of relativistic plasma environments around neutron stars and black holes, such as pulsar wind nebulae and black-hole jets and coronae.

The remaining Sections are as follows. Section 3.2 outlines previous analytical research on turbulent NTPA, and presents an analytical model relating the FP coefficients to the evolution of the power-law index of the particle energy distribution. Section 3.3 describes the PIC simulation setup and parameter scans. Section 3.4 discusses general features of the turbulent particle energisation and their time evolution in these simulations. Section 3.5 presents my measurements of the power-law index of the particle energy distribution as a function of time and the scalings of the key parameters describing this evolution with the dimensionless system parameters. Section 3.6 tests whether particle energies can be modelled as diffusive, and whether the measured FP coefficients reproduce the evolution of the particle energy distribution from the PIC simulations. Section 3.7

presents findings on the energy diffusion coefficient as a function of time and system parameters. Section 3.8 shows measurements of the energy advection coefficient as a function of time and system parameters, and examines how it relates to the theoretical model described in Section 3.2. Finally, Section 3.9 summarises the results.

3.2 Theory

3.2.1 Fokker-Planck theories of nonthermal particle acceleration

The full momentum-space FP equation is given in (3.1). There, the diffusion and advection coefficients are, respectively, a tensor ($\mathbf{D}_{pp}$) and a vector ($\boldsymbol{A}_p$), and encapsulate all possibilities of diffusion and advection of the particle distribution between different components of momentum (including, e.g., pitch-angle scattering). In this paper, I will now assume the form $\boldsymbol{A}_p = A_p\hat{\boldsymbol{p}}$ and $\mathbf{D}_{pp} = D_{pp}\mathbf{I}$, where A_p and D_{pp} are the scalar advection and diffusion coefficients, respectively, with $\hat{\boldsymbol{p}}$ being the unit vector along $\boldsymbol{p}$ and $\mathbf{I}$ the identity matrix.

Furthermore, I will limit the analysis to a reduced description of the FP equation for the global particle distribution $f(\gamma, t)$ in energy only, neglecting the pitch angle dependence:

$$\partial_t f = \partial_\gamma (D\partial_\gamma f - Af), \tag{3.2}$$

where the particle density is $dn = f(\gamma)d\gamma$, and the energy diffusion and advection coefficients D and A are scalar functions of particle energy γmc^2. However, as it is common to state the FP coefficients in momentum space, it is useful to convert between the energy-space and momentum-space forms. Beginning with (3.1), where the particle density $dn = F(\mathbf{p})d^3\mathbf{p}$, and assuming for simplicity that F, D_{pp}, and A_p are isotropic (i.e., depend only on $p \equiv |\mathbf{p}|$, thus neglecting any pitch angle dependence) results in

$$\frac{\partial F(p)}{\partial t} = \frac{1}{p^2}\frac{\partial}{\partial p}\left(p^2 D_{pp}\frac{\partial F(p)}{\partial p}\right) - \frac{1}{p^2}\frac{\partial}{\partial p}[p^2 A_p F(p)], \tag{3.3}$$

where dn simplifies to $4\pi p^2 F(p)$. This is the form used in, e.g., Schlickeiser (1985, 1989). Finally, in the ultrarelativistic limit $\gamma \gg 1$ that I will focus on in this paper, the particle energy γmc^2 is just

directly proportional to its momentum p, i.e., $\gamma = p/mc$; I can then substitute $f(\gamma) = 4\pi p^2 F(p)$ and thus write

$$\frac{\partial f(\gamma)}{\partial t} = \frac{\partial}{\partial \gamma}\left[D_{pp}\left(\frac{\partial f(\gamma)}{\partial \gamma} - 2\,\frac{f(\gamma)}{\gamma}\right)\right] - \frac{\partial}{\partial \gamma}(A_p f(\gamma)). \tag{3.4}$$

Comparing this to (3.2), I see that the momentum-space coefficients D_{pp} and A_p are related to the energy-space coefficients (without subscript) by $D = D_{pp}$ and $A = A_p + 2D_{pp}/\gamma$.

3.2.2 Analytical model of the nonthermal power-law tail

Previous numerical simulation (PIC) studies have shown the formation and gradual evolution of a relativistic nonthermal power-law section in the particle energy distribution in driven collisionless plasma turbulence (e.g., Zhdankin et al., 2017, 2018b; Wong et al., 2020). That is, there is a significant time interval where a segment of $f(\gamma)$ is well approximated by a power law. I wish to understand the relationship between the FP coefficients and the time evolution of the power-law index of the nonthermal section of $f(\gamma)$. Hence, I consider the FP equation (3.2), where f, D, and A are functions of γ and t. I will focus here only on the nonthermal region, and only for the times when this power law is well-formed and distinct, so that the high-energy part of $f(\gamma)$ can be said to be an evolving power law

$$f(\gamma, t) = K(t)\gamma^{-\alpha(t)}. \tag{3.5}$$

Here, K should not be viewed as a normalisation factor, since the number of particles in the power law does not have to be conserved. Instead, I imagine reservoirs of particles at low and high energies. At low energies, this is naturally interpreted as the thermal bulk. At high energies, this could be a form of escape; in simulation terms, this may be a high-energy pileup as described in Zhdankin et al. (2017, 2018b). Particles moving in and out of these reservoirs can change the total number of particles in the nonthermal tail.

Theoretical models commonly predict $D(\gamma, t) \propto \gamma^2$ in the nonthermal range (e.g., Kulsrud and Ferrari, 1971; Melrose, 1974; Skilling, 1975; Blandford and Eichler, 1987; Chandran, 2000; Cho

and Lazarian, 2006; Demidem et al., 2020; Lemoine, 2021). Hence, for this model, I assume

$$D(\gamma, t) = D_0(t)\gamma^2. \tag{3.6}$$

In the absence of an advection coefficient, this diffusion coefficient causes the mean energy $\overline{\gamma}$ of a collection of particles to increase at the instantaneous rate of $\dot{\overline{\gamma}} \equiv \partial_t \overline{\gamma} = 2D_0\overline{\gamma}$: an exponential pace, in accordance with the Fermi acceleration theory. Hence, D_0 may be interpreted as the inverse diffusive acceleration time, and the scaling $D \propto \gamma^2$ implies that this characteristic acceleration time is independent of the particle energy ($\dot{\overline{\gamma}}/\overline{\gamma} = \text{const}$).

While I could also assume a functional form for $A(\gamma, t)$, specifying independent forms for all of f, D and A will almost certainly be inconsistent due to overconstraining the problem. Instead, since A is by far the most uncertain theoretically, I use (3.2) with the above adopted functional forms (3.5) and (3.6) for $f(\gamma, t)$ and $D(\gamma, t)$ respectively to solve for $A(\gamma, t)$ analytically, and then later compare it to my numerical measurements of A in Section 3.8. That is, I infer A from consistency conditions on the existence of a power-law range in the energy distribution.

First, I integrate the FP equation (3.2) from γ' to ∞. Keeping in mind (3.5) and (3.6), I assume $\alpha > 1$ and $Af \to 0$ as $\gamma \to \infty$ so that all the upper integral limits can be discarded, and then replace γ' with γ. The non-advection terms evaluate to:

$$\begin{aligned} &D\partial_\gamma f = D_0\gamma^2(-\alpha/\gamma)f = -D_0\alpha\gamma f, \\ &\int \partial_t f \, d\gamma = \partial_t \left(K\frac{\gamma^{1-\alpha}}{1-\alpha} \right) = \frac{\gamma f}{1-\alpha}\left(-\dot{\alpha}\log\gamma + \frac{\dot{K}}{K} + \frac{\dot{\alpha}}{1-\alpha} \right). \end{aligned} \tag{3.7}$$

I collect into (3.2) and eliminate the common factors of f to get

$$A = -\frac{\gamma}{1-\alpha}\left(-\dot{\alpha}\log\gamma + \frac{\dot{K}}{K} + \frac{\dot{\alpha}}{1-\alpha} + D_0(\alpha - \alpha^2) \right). \tag{3.8}$$

This can be recast in a convenient, compact form as

$$A(\gamma, t) = A_0(t)\,\gamma\,\log[\gamma/\gamma_A^*(t)], \tag{3.9}$$

with:

$$A_0(t) = \frac{\dot{\alpha}}{1-\alpha}, \tag{3.10}$$

$$\gamma_A^*(t) = \exp\left[\frac{\dot{K}/K + D_0(\alpha - \alpha^2)}{\dot{\alpha}} + \frac{1}{1-\alpha}\right], \tag{3.11}$$

where all right-hand-side variables in (3.10) and (3.11) are functions of t only. The result (3.9) is the prediction for the functional form of $A(\gamma, t)$ required to maintain a power-law distribution. The γ-dependence is remarkably simple and contained only in (3.9), with (3.10) and (3.11) functions of t alone. As a consequence of the assumed quadratic diffusion coefficient $D \propto \gamma^2$, only one term in the large parentheses of (3.8) has γ-dependence, and this directly leads to a generic $\gamma \log \gamma$ energy dependence of the advection coefficient A, as long as $\dot{\alpha} \neq 0$. If instead $\dot{\alpha} = 0$, then A is simply proportional to γ by (3.8). Since $\dot{\alpha}$ is usually negative in the PIC turbulence simulations, with the power law tending to harden over time, A_0 is positive. Hence, (3.9) predicts that A is negative at low energies and positive at high energies, with the sign change occurring at γ_A^*. Furthermore, A has a minimum at $\gamma = \gamma_A^*/e$, with value $A_{\min} = -A_0\gamma_A^*/e$. I note that at high energies, A is faster than first-order Fermi acceleration (where $A \propto \gamma$) by a factor of $\log \gamma$. This is a surprising property that appears to be generic for maintaining a time-evolving power-law particle energy distribution given $D \propto \gamma^2$.

Up to this point, I have introduced fixed γ-dependencies for the particle energy distribution f and FP coefficients D and A, parameterised by five functions of time K, α, D_0, A_0, and γ_A^*, which are associated to the former three quantities as follows:

$$f : K, \alpha; \quad D : D_0; \quad A : A_0, \gamma_A^*.$$

While all of these ultimately source from the PIC simulation, the particle energy distribution group of variables and the FP coefficient group are measured in two substantially different ways: the former use the overall particle energy distribution and the latter, tracked particles (see Subsection 3.3.2 for some details on tracked particles).

I later find in Section 3.8 that A_0 and γ_A^* both fluctuate substantially in time (or at least, their measurements do), whereas their effect on the power law is cumulative over time, innately

smoothing out the fluctuations, and hence an integral comparison can be less noisy and more illuminating. The tradeoff for this smoothing is that I will have to select the integration constants (somewhat arbitrarily). To enable this, I rearrange the A-variables equations (3.10) and (3.11) into time integrals for the power-law variables $K(t)$ and $\alpha(t)$. From (3.10),

$$\frac{\alpha(t)-1}{\alpha(t_0)-1} = \exp\left[-\int_{t_0}^{t} A_0(t')dt'\right], \tag{3.12}$$

where t_0 is an arbitrary reference time and $\alpha(t_0)$ is an integration constant. Similarly, (3.11) rearranged for $\dot{K}/K$ integrates to

$$\frac{K(t)}{K(t_0)} = \exp\left\{\int_{t_0}^{t}\left[D_0(\alpha^2-\alpha)+\dot{\alpha}\left(\log\gamma_A^* + \frac{1}{\alpha-1}\right)\right]dt'\right\}. \tag{3.13}$$

Here, t_0 is again a reference time, $K(t_0)$ is an integration constant and D_0, α, $\dot{\alpha}$, and γ_A^* are all functions of time t'. One may optionally then trade $\dot{\alpha}$ for A_0 using (3.10). While theoretically equivalent, $\dot{\alpha}$ and A_0 are obtained from different data (f and A, respectively) and one or the other may be preferable for noise or other technical considerations. I also note that the exponential form of (3.12) has a suggestive (although not quite fully compatible) similarity to the exponential fit to $\alpha(t)$ introduced in Section 3.5.

To compute the advection coefficient $A(\gamma,t)$ in these PIC simulations, I will first measure the energy-dependent average rate of particle energy gain $M \equiv A + \partial_\gamma D$ (see Subsection 3.6.1). Later, I will find it convenient to compare directly between theoretical and measured M rather than A because M does not depend on D and its associated measurement noise. Hence I derive analogues of (3.9–3.11) for M using $\partial_\gamma D = 2D_0\gamma$ which comes from (3.6):

$$M(\gamma,t) = M_0\gamma\log(\gamma/\gamma_M^*), \tag{3.14}$$

where $M_0 = A_0$ is the same as in (3.10) (it is convenient to have a different symbol for this prefactor as later fits to A and M will not necessarily result in the same value) and

$$\gamma_M^*(t) = \exp\left[\frac{\dot{K}/K + D_0(-\alpha^2+3\alpha-2)}{\dot{\alpha}} + \frac{1}{1-\alpha}\right], \tag{3.15}$$

differing from γ_A^* only by the factor multiplying D_0. An equivalent relation is $\log\gamma_A^* - \log\gamma_M^* = 2D_0/A_0$. The corresponding integral forms are:

$$\frac{\alpha(t)-1}{\alpha(t_0)-1} = \exp\left[-\int_{t_0}^{t} M_0(t')dt'\right], \tag{3.16}$$

$$\frac{K(t)}{K(t_0)} = \exp\left\{\int_{t_0}^{t}\left[D_0(\alpha^2-3\alpha+2) + \dot{\alpha}\left(\log\gamma_A^* + \frac{1}{\alpha-1}\right)\right]dt'\right\}. \tag{3.17}$$

In summary, from assuming that the particle energy distribution is a time-dependent power-law $f(\gamma,t) = K(t)\gamma^{-\alpha(t)}$ (3.5) and that the energy diffusion coefficient $D(\gamma,t) = D_0(t)\gamma^2$ (3.6), I obtain a prediction for the energy advection coefficient of $A(\gamma,t) = A_0(t)\,\gamma\,\log[\gamma/\gamma_A^*(t)]$ (3.9), where A_0 is a function of α and $\dot{\alpha}$ (3.10); and γ_A^* is a function of $\dot{K}/K$, D_0, α, and $\dot{\alpha}$ (3.11).

My analytical model does not explain why a power-law distribution arises, but this might be expected on general grounds from the underlying scale invariance of turbulence at MHD scales. The assumption of a power-law distribution is also supported by recent maximum-entropy models of particle acceleration that derive such distributions, with the index being connected to the characteristic momentum scale of irreversible dissipation (e.g., Zhdankin, 2022, and references therein).

3.3 Simulations

I analyse 3D PIC simulations of externally driven turbulence in collisionless relativistic pair plasma performed with the code Zeltron (Cerutti et al., 2013). The system is a periodic cube of side length L and a guide magnetic field $B_0\hat{\mathbf{z}}$. The plasma, with total (electrons and positrons combined) charged particle density n_0, is initialised from a uniform isotropic Maxwell-Jüttner relativistic thermal distribution. I fix the initial temperature to $T_0 \equiv \theta_0 mc^2 = 100\,mc^2$ (where m is the electron rest mass), corresponding to an average Lorentz factor of $\gamma_{\mathrm{avg}_0} \approx 3\theta_0 = 300$.

I define the (time-dependent) characteristic Larmor radius $\rho_e(t) \equiv \gamma_{\mathrm{avg}} mc^2/eB_{\mathrm{rms}}$ and skin depth $d_e(t) \equiv (\gamma_{\mathrm{avg}} mc^2/4\pi n_0 e^2)^{1/2}$, where $\gamma_{\mathrm{avg}}(t)$ is the instantaneous average Lorentz factor and $B_{\mathrm{rms}}(t)$ is the rms magnetic field strength. The corresponding initial values are denoted ρ_{e0} and d_{e0}. The system then has three dimensionless physical parameters (with any choice of two independent) which are ratios of the length scales ρ_e, d_e, and L. These are the ("hot") magnetisation

$\sigma \equiv 3B_{\rm rms}^2/16\pi n_0 \gamma_{\rm avg} mc^2 = (3/4)(d_e/\rho_e)^2$, which is the ratio of the magnetic enthalpy $B_{\rm rms}^2/4\pi$ to the plasma enthalpy $(4/3)n_0\gamma_{\rm avg}mc^2$; the system length relative to the Larmor radius L/ρ_e; and the system length relative to the skin depth L/d_e. For the initial ultra-relativistic Maxwell-Jüttner distribution, the initial magnetisation $\sigma_0 = B_0^2/16\pi n_0 T_0$ corresponds to an initial plasma-to-magnetic pressure ratio of $\beta_0 = 8\pi n_0 T_0/B_0^2 = 1/2\sigma_0$.

Turbulence is continuously electromagnetically driven (TenBarge et al., 2014) at wavenumbers $k_d = 2\pi/L$ and becomes fully developed after several light-crossing times; the turbulence-driving implementation and other technical simulation details are described in Zhdankin et al. (2018a). The driving strength is chosen so that the rms level of the turbulent magnetic fluctuations $\delta B_{\rm rms}$ is comparable to the magnitude of the guide field B_0. The turbulence is essentially Alfvénic (Zhdankin et al., 2018a), with characteristic Alfvén velocity $v_A/c \equiv [\sigma/(\sigma+1)]^{1/2}$.

The simulation is performed on a uniform grid of N^3 cells. The grid spacing $\Delta x \equiv L/N$ is chosen to resolve the smaller of ρ_{e0} and d_{e0}, as well as the initial Debye length $\lambda_{D0} = \sqrt{T_0/4\pi n_0 e^2}$ (for an ultra-relativistic Maxwell-Jüttner distribution, $d_e/\lambda_D = \sqrt{3}$). More precisely, I choose $\Delta x = \min(\rho_{e0}, \sqrt{2}d_{e0})/1.5$. Except where otherwise stated, the combined number of electron and positron macroparticles per cell is $N_{\rm ppc} = 64$. The simulation durations t_f range from about 10 to 100 light-crossing times L/c depending primarily on σ_0 (lower σ_0 require longer runs, due to the longer principal dynamical timescale set by L/v_A), with $t_f \sim 20L/c$ being most common.

3.3.1 Parameter scans and simulation list

One of the primary aims of this study is to extend the work of Wong et al. (2020) by analysing and comparing the results of numerous simulations with a range of initial system parameters; these simulations are summarised in Table 3.1. They are grouped as follows:

(1) A system-size scan at fixed $\sigma_0 = 3/8$ (where $\rho_{e0} = \sqrt{2}d_{e0}$) with L/ρ_{e0} ranging from 171–1024 corresponding to N values 256–1536. This group contains the largest simulation where $L/\rho_{e0} = 1024$, with $N^3 = 1536^3$ grid cells.

Table 3.1: List of simulations and their parameters.

	N	σ_0	$\frac{L}{\rho_{e0}}$	$\frac{\rho_{e0}}{\Delta x}$	$\frac{L}{d_{e0}}$	$\frac{d_{e0}}{\Delta x}$	$\frac{ct_f}{L}$	N_{ppc}
σ_0 scan	384	3/2	256	1.5	181	2.1	7.8	64
$N = 384$	384	3/8	256	1.5	362	1.1	22.3	64
	384	3/32	128	3.0	362	1.1	64.9	64
	384	3/128	64	6.0	362	1.1	122.3	64
σ_0 scan	768	3/2	512	1.5	362	2.1	22.3	64
$N = 768$	768	3/4	512	1.5	512	1.5	22.3	64
	768	3/8	512	1.5	724	1.1	22.3	64
	768	3/16	362	2.1	724	1.1	20.9	64
	768	3/32	256	3.0	724	1.1	55.2	64
	768	3/64	181	4.2	724	1.1	59.7	64
	768	3/128	128	6.0	724	1.1	117.8	64
N_{ppc} scan	384	3/8	256	1.5	362	1.1	22.3	8
	384	3/8	256	1.5	362	1.1	22.3	16
	384	3/8	256	1.5	362	1.1	22.3	32
	384	3/8	256	1.5	362	1.1	22.3	128
	384	3/8	256	1.5	362	1.1	22.3	256
system-size	256	3/8	171	1.5	241	1.1	16.7	64
scan	384	3/8	256	1.5	362	1.1	22.3	64
	512	3/8	341	1.5	483	1.1	22.3	64
	768	3/8	512	1.5	724	1.1	16.7	64
	1024	3/8	683	1.5	965	1.1	16.7	64
	1536	3/8	1024	1.5	1448	1.1	14.2	64

(2) Sets of smaller simulations with varying σ_0: one with $N = 768$ and $\sigma_0 \in \{3/2, 3/4, ..., 3/128\}$, and another with $N = 384$ and $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$. As mentioned above, I set $\Delta x = \min(\rho_{e0}, \sqrt{2}d_{e0})/1.5$ to resolve the smaller of the two kinetic scales while maximising the inertial range for given grid size (which is closely related to computational cost), leading to L/ρ_{e0} and L/d_{e0} varying as shown in the table.

(3) A scan to investigate the convergence of the results with the number of particles per cell, N_{ppc}, for a representative case with $\sigma_0 = 3/8$ and $L/\rho_{e0} = 256$ ($N = 384$).

I discuss briefly some complications caused by having three initial dimensionless ratios: L/ρ_{e0}, L/d_{e0}, and $d_{e0}/\rho_{e0} \propto \sqrt{\sigma_0}$. A parameter scan must select at least two of these ratios to change with at most one held fixed (as fixing two ratios would fix all three). When aiming to isolate one variable, the handling of the other two parameters can still be consequential. For the system-size scan, the choice is obvious: as L is the privileged (always largest) scale, I fix d_{e0}/ρ_{e0}. However, the σ_0 scan is more complicated: as ρ_{e0} and d_{e0} have no particular ordering, it is unclear how best to set L/ρ_{e0} and L/d_{e0}. For example, increasing σ_0 with fixed L/ρ_{e0} will decrease L/d_{e0}, and, supposing $d_{e0} > \rho_{e0}$, also decrease the separation between the driving scale and largest microscale (i.e., the inertial range). Then, even though the change in L/d_{e0} is mathematically predetermined due to fixed L/ρ_{e0}, it is still important to consider carefully whether some change in the results would be best attributed to the higher magnetisation or to the shorter inertial range.

A full two-dimensional parameter scan is unfortunately impractical. However, many of the results related to the inertial range are expected to become insensitive to L/ρ_{e0} and L/d_{e0} in the limit of $L/\rho_{e0} \gg 1$ and $L/d_{e0} \gg 1$. Thus, the system-size scan results can provide some assurance that the simulations are in this regime and hence the effect of the varying scale hierarchies in the σ_0 scans is minor. Nevertheless, one should keep in mind this limitation of this study with respect to the σ_0 scans. Accordingly, I refrain from declaring primacy to one or the other of L/ρ_{e0} or L/d_{e0} and refer to the σ_0 scans by their fixed N, which, although not a fully-fledged physical parameter, is (by the Δx choice) almost equal to the initial ratio of L to the smallest microscale.

3.3.2 Tracked particles and energy oscillation removal

My analysis uses large numbers of tracked particles, for which the position, momentum, and local electromagnetic field vectors are recorded at fine enough time intervals to resolve details such as particle orbits. While memory and storage constraints prevent recording these data for all the $\sim 10^{11}$ particles in the simulations, the $\sim 10^6$ randomly selected tracked particles are sufficient to form a high-quality statistical ensemble representative of the overall particle distribution.

There are large (order unity) oscillations in particle energy due to $\mathbf{E} \times \mathbf{B}$ drift, particularly significant in relativistic turbulence with $E_{\rm rms} \sim B_0$. These, along with a removal procedure which is necessary to facilitate FP-type analysis, are described in Wong et al. (2020) and Section 2.2. I perform that procedure on this data, extracting the secular component of particle energy, greatly reducing oscillations and enabling tests of the FP picture of NTPA. Hereafter, γ and "energy" refer to the smoothed energy, except for in global particle energy distributions, where the difference is minor anyway.

3.4 System overview

Before beginning the main analysis, I first provide a qualitative overview of how the driven turbulent system evolves. I briefly present the time evolution of the overall particle energy spectrum and average energy, and then examine a particle acceleration event superimposed on the turbulent fields.

Figure 3.1 shows the evolution of the overall particle energy spectrum $f(\gamma, t)$ for simulations with $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ and $N = 768$, which have the following common behaviour. Beginning at a relativistic thermal distribution, each spectrum develops a power-law nonthermal tail after a few Alfvén crossing times. This nonthermal segment begins near the average particle energy and extends to roughly the system-size limit $\gamma_{\rm max} \equiv LeB_0/2mc^2 = L\gamma_{\rm avg}/2\rho_{e0}$, where the gyroradius reaches the system size. For example, the typical runs with $L/\rho_{e0} = 512$ have $\gamma_{\rm max} \simeq 7.7 \times 10^4$, while the largest simulation has $L/\rho_{e0} = 1024$ and $\gamma_{\rm max} \simeq 1.5 \times 10^5$. As the

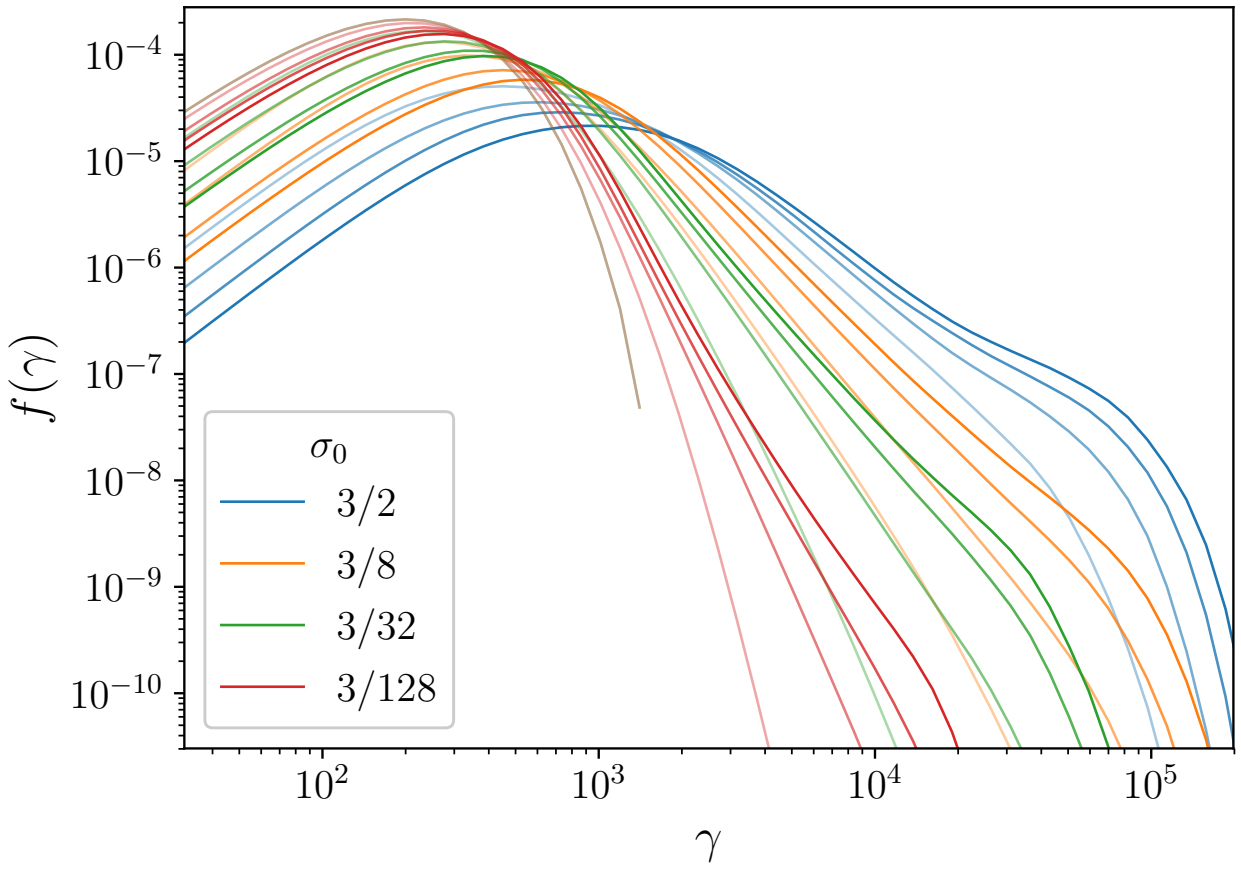

Figure 3.1: Evolution of particle energy distribution for multiple $N = 768$ simulations with different $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$. Lines fade in from the start to the end of each simulation.

system receives continuous energy input from turbulent driving and there is no radiative cooling, the plasma heats over time and so the extent of the nonthermal power-law segment shrinks as the simulation progresses. As particles are accelerated to the system-size limit, they form a pileup population near $\gamma_{\rm max}$, appearing as a bump at the end of the distribution (Zhdankin et al., 2018b).

Figure 3.2 shows the plasma heating history, — i.e., the time evolution of the rate of change of the global average particle energy $\dot{\gamma}_{\rm avg}$, after applying a central moving time-averaging of width L/v_{A0} (based on the initial Alfvén velocity v_{A0}), for the $N = 768$ series of simulations with $\sigma_0 \in \{3/2, 3/4, ..., 3/128\}$. While there are significant fluctuations in time, an overall statistical steady state is reached after a few Alfven-crossing times. This indicates that $\gamma_{\rm avg}$ predominantly increases linearly with time. There is an initial segment of several Alfvén crossing times with a slower rate of increase, corresponding to the period when some of the energy injected by driving goes towards establishing turbulent electromagnetic fields. Once turbulence is fully established, these fields become statistically quasi-steady, and thereafter the supplied energy is predominantly converted into the plasma internal energy (heat and nonthermal particle acceleration), corresponding to a long-term linear increase in $\gamma_{\rm avg}$ with time. The roughly even separation of the $\dot{\gamma}_{\rm avg}$ lines in Figure 3.2 corresponding to simulations spaced by fixed multiples in σ_0 imply that the rate of energy gain is roughly proportional to σ_0. To aid visual comparison, lines with constant $\dot{\gamma}_{\rm avg} \propto \sigma_0$ are overlaid on Figure 3.2.

Figure 3.3 shows the particle trajectory before, during, and after an acceleration event, for a single particle in the simulation with $N = 1536$ ($L/\rho_{e0} = 1024$) and $\sigma_0 = 3/8$ (selected arbitrarily from the tracked particle sample). The local magnetic field is also shown (blue arrows), taken at the time of the primary acceleration event (red dot). Before this energisation event, the particle enters the acceleration region while gyrating around a magnetic field line. It then encounters a shear in the magnetic field, which causes the gyro-centre motion to reflect, in a manner akin to the Fermi process (Fermi, 1949), as considered in recent models such as Lemoine (2021, 2022). The particle then leaves the region with a significantly larger Larmor radius than it began with, indicating an increase in the momentum component perpendicular to the magnetic field. While this is only a

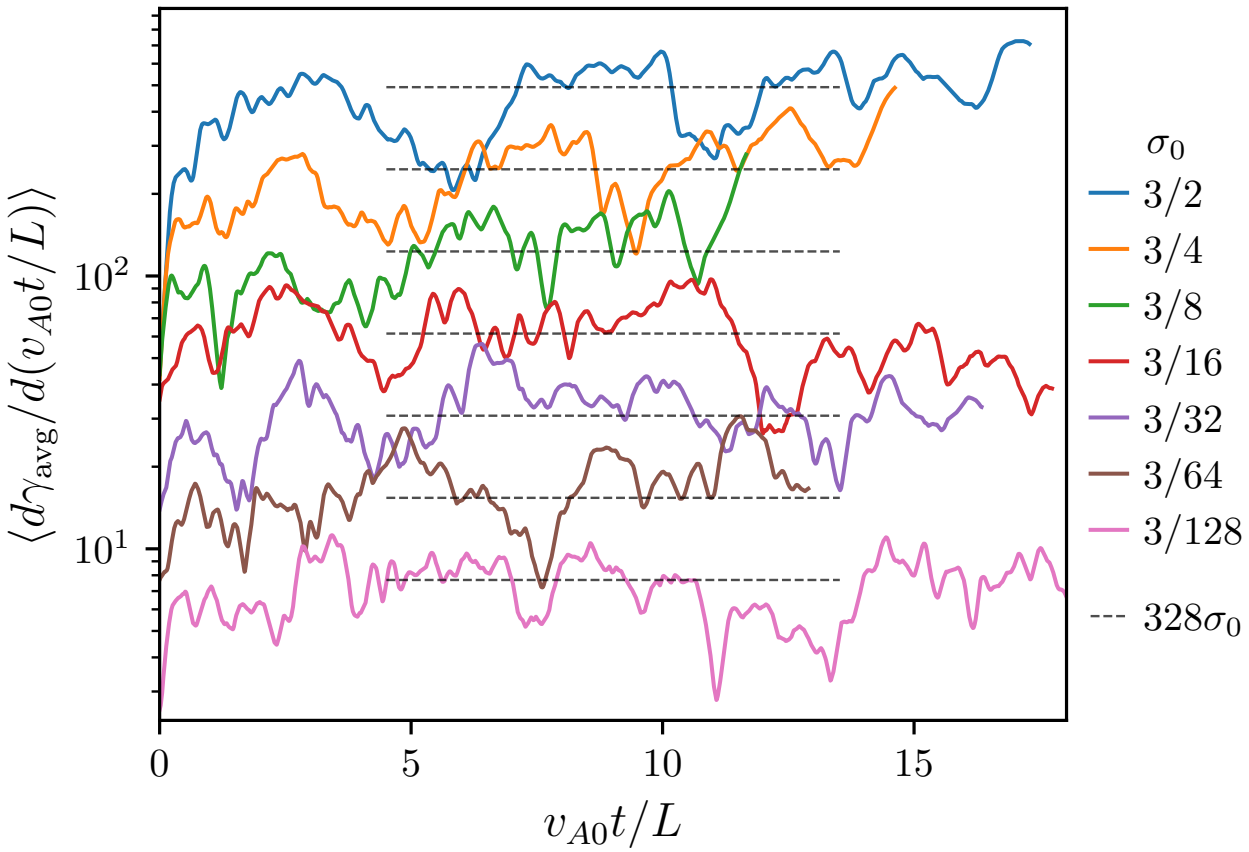

Figure 3.2: Moving average of rate of change of average particle energy for simulations with $N = 768$ and $\sigma_0 \in \{3/2, 3/4, ..., 3/128\}$. The width of the moving average is L/v_{A0}, and this filtering is denoted in the y-axis label by angled brackets. Black dashed lines correspond to $\dot{\gamma}_{\rm avg} \propto \sigma_0$.

single example, I expect acceleration events such as this one to be generic within relativistic plasma turbulence.

3.5 Particle energy spectra

Connecting the evolution of the particle energy distribution with the measured Fokker-Planck coefficients is important to better our understanding of turbulent particle acceleration. If the system were steady-state, this would be simplified by the left-hand-side of (3.2) being zero. However, the time-dependence of $f(\gamma)$ described in Section 3.4 complicates the interpretation of the Fokker-Planck coefficients in the context of a steady-state energy spectrum. Nevertheless, I still wish to understand the slowly evolving, quasi-steady state achieved at late times in these non-radiative turbulence simulations. Understanding, even empirically, the key characteristics of $f(\gamma, t)$, and in particular, the nonthermal tail, will help connect the time-dependent behaviour of the distribution function with that of the FP coefficients measured in later sections.

A fundamental result of kinetic simulations of turbulent NTPA is the generation of particle energy distributions with extended nonthermal ranges (see Figure 3.4a). For a particle energy distribution at some time t, the local power-law index can be defined as $\alpha_{\rm loc}(\gamma, t) \equiv -\partial \log f / \partial \log \gamma$ (and hence a pure power-law segment is characterised by $\partial \alpha_{\rm loc}/\partial \gamma = 0$). At a given instant, the local indices $\alpha_{\rm loc}(\gamma)$ can be summarised into a representative single value, $\alpha(t)$, characterising the predominant value of $\alpha_{\rm loc}(\gamma, t)$ in the high-energy nonthermal tail. Moreover, a special value of $\alpha(t)$, which I denote as $\alpha_{\rm fin}$, is often chosen to represent the late-time asymptotic (or "final") value of $\alpha(t)$. These values $\alpha(t)$ and $\alpha_{\rm fin}$ are main quantities of interest and have observational implications and theoretical significance. For example, the system-size dependence of $\alpha_{\rm fin}$ is critical to understand because it affects our ability to extrapolate the findings of these PIC simulations to astrophysical systems (Zhdankin et al., 2018b). Therefore, in this Section, I investigate the time dependence of α, develop a prescription for measuring $\alpha_{\rm fin}$, and explore their dependence on system parameters such as the system size and initial magnetisation.

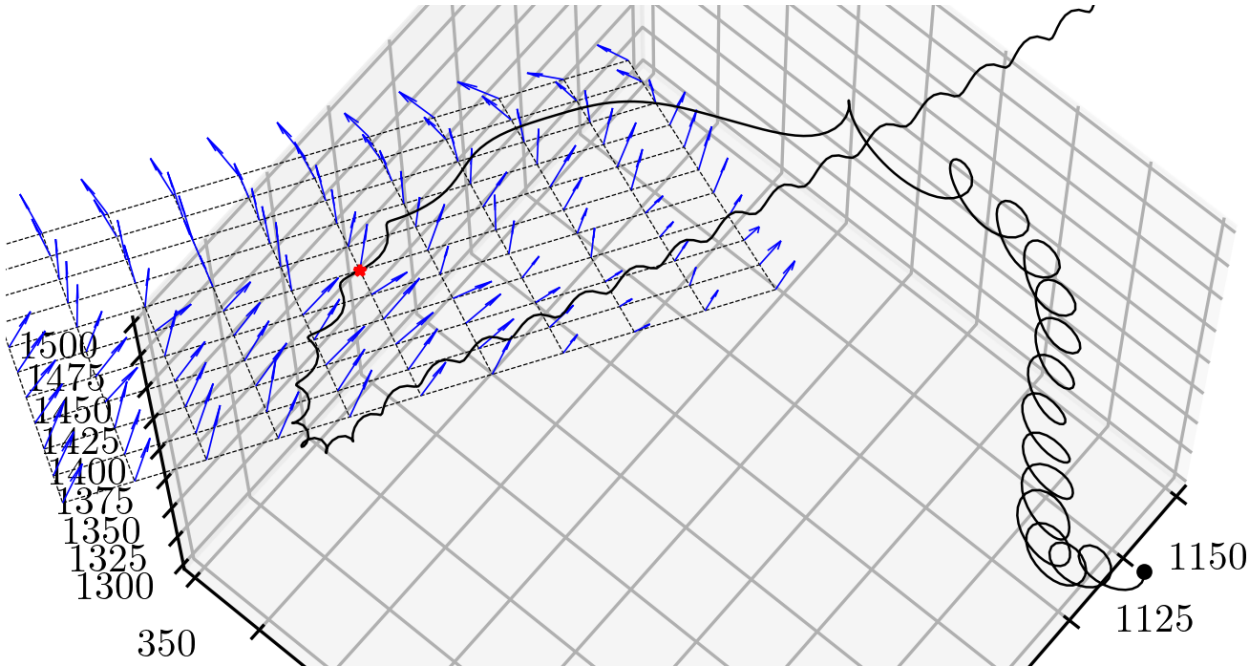

Figure 3.3: Particle trajectory in the vicinity of an acceleration event, overlaid with magnetic field (blue vectors) at the time when the particle was located at the red dot. The latest time in the trajectory is marked by the black dot.

3.5.1 Fitting the power law index

The process of obtaining $\alpha(t)$ from $\alpha_{\rm loc}(\gamma, t)$ at a given time t is complicated by the necessity to select the section of $\alpha_{\rm loc}(\gamma)$ to fit a straight line to. This section must exclude, on the left, the thermal particles, and on the right, the high-energy pileup due to finite simulation size, all the while being as long as possible to maximise the accuracy of the fit. Furthermore, one must then decide how to obtain the late-time asymptotic value $\alpha_{\rm fin}$. For example, past studies have selected $\alpha_{\rm fin}$ based on "the time with the longest fitted power-law segment" (Zhdankin et al., 2017) or a "time that is logarithmic with system size" (Zhdankin et al., 2018b). These decisions limit the precision with which the evolution of different particle energy distributions can be compared, and thus a more systematic scheme is desirable.

Here, I describe a procedure to distill a unique value of $\alpha(t)$ for a given distribution $f(\gamma, t)$, which would be well-suited to the particle energy distributions produced by these simulations and reasonably insensitive to fitting choices. This scheme yields values that are similar to those obtained by the method developed by Werner et al. (2018), but is more robust against the maximum-energy pileup at late times.

At intermediate and late times, when the nonthermal tail has fully developed and a high-energy pileup has appeared, there is a local minimum in the local power-law index $\alpha_{\rm loc}(\gamma)$ (Figure 3.4b). Even as the thermal peak continues to move to the right and the maximum-energy pileup increases in size, the energy γ at which this local minimum is attained remains very stable. Furthermore, precisely because this minimum is a local stationary point, the function $\alpha_{\rm loc}(\gamma)$ varies very little in its vicinity. That is, even if the minimum is not located exactly, the resulting difference in the selected value of the local index α is small. This stability suggests considering the local minimum value of $\alpha_{\rm loc}(\gamma, t)$ as the overall (time-dependent) power-law index $\alpha(t)$ characterising the nonthermal part of the spectrum. I therefore construct an automatic procedure to extract the local minimum of $\alpha_{\rm loc}(\gamma, t)$. I expect this procedure to be broadly applicable to other simulation studies which produce relatively clean and smooth particle distributions with similar structure, exhibiting

a thermal bulk followed by a nonthermal power-law tail and finally a pileup and cutoff at highest energies.

The procedure works as follows. I locate the local minimum by fitting a cubic, in $\log\gamma$ space, (Figure 3.4b, dashed lines) to the γ-range within a factor of five of the geometric mean of the average energy $\gamma_{\rm avg}$ and the system-size limited energy $\gamma_{\rm max}$ (Figure 3.4b, crosses). The Lorentz factor corresponding to the desired local minimum (Figure 3.4b, circles) is then the first turning point of the cubic, or the inflection point if there are no turning points. The cubed term is forced to be negative so that the first turning point is a local minimum (if it exists), and points where the corresponding energy spectrum density $f(\gamma)$ are less than 10^{-9} times the peak value at that time are excluded due to noise. I expect that other methods for finding local minima or inflection points would produce very similar results. When the fit does not have a significant flat or inverted section, α is uncertain, but this simply reflects the fact that the power-law tail is not yet well developed at that time (e.g., Figure 3.4, blue line).

This procedure does not have any physical or theoretical justification, and has so far only been applied to smooth spectra from simulations. However, it has some favourable properties. As can be seen in Figure 3.4a, it very closely fits the centre of the power-law range, and tends to trim off the late-time pileup at $\gamma_{\rm max}$. Also, as it fits a turning point in $\alpha_{\rm loc} \equiv -\partial\log f/\partial\log\gamma$, it can be said to fit the straightest location of the power-law section of $f(\gamma)$. The main idea of this procedure is to adopt the first local minimum value of $\alpha_{\rm loc}$ as the preferred single-value α for a given energy spectrum; the exact method for locating this local minimum is secondary.

3.5.2 Application to the system size scan and statistical samples

I apply my procedure for identifying the single-value power-law index α on $f(\gamma,t)$ snapshots from the system-size scan simulations, which have fixed $\sigma_0 = 3/8$ and varying $L/\rho_{e0} \in \{1024, 683, 512, 341, 256, 171\}$. I ignore data points before $t = 2L/v_{A0}$, when turbulence has not yet fully developed. Figure 3.5 shows the resulting $\alpha(t)$. The vertical line segments on each plotted point are obtained as follows: expand a symmetric $\alpha_{\rm loc}$-range around α until it covers a surrounding

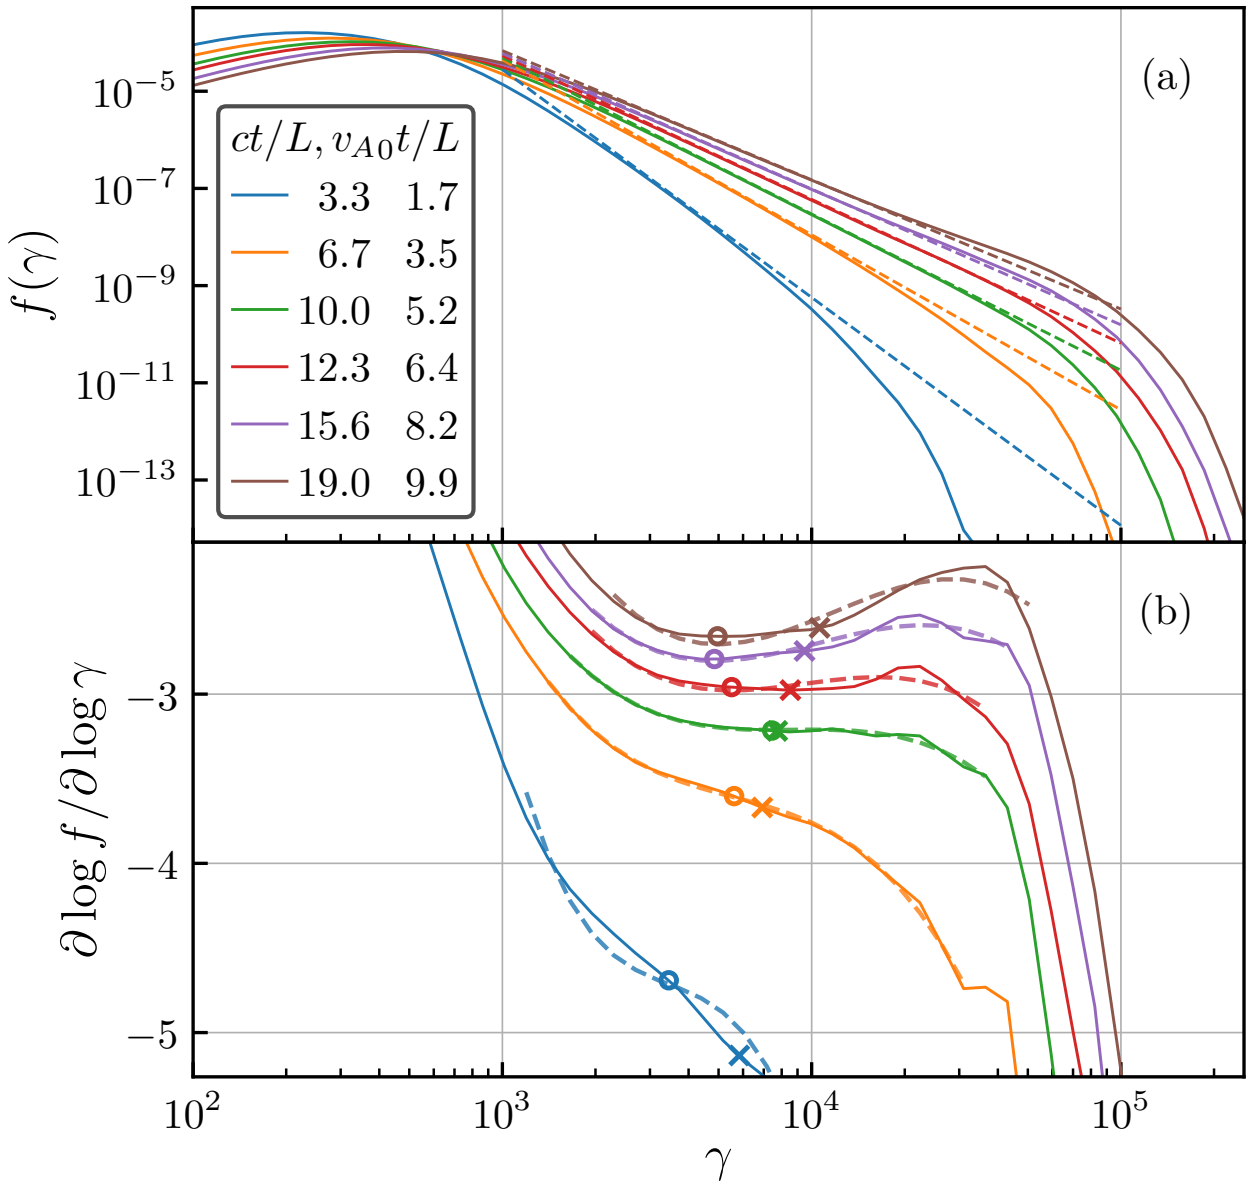

Figure 3.4: Particle energy distribution (a) and the local power-law index (b) at several different times for a representative simulation with $\sigma_0 = 3/8$ and $L/\rho_{e0} = 512$ from the N=768 σ_0-scan. In (b), dashed lines are the fitted cubics, crosses mark the geometric mean of $\gamma_{\rm avg}$ and $\gamma_{\rm max}$, while circles mark the first turning point of the cubic or the inflection point if there are no turning points. See the main text for more details.

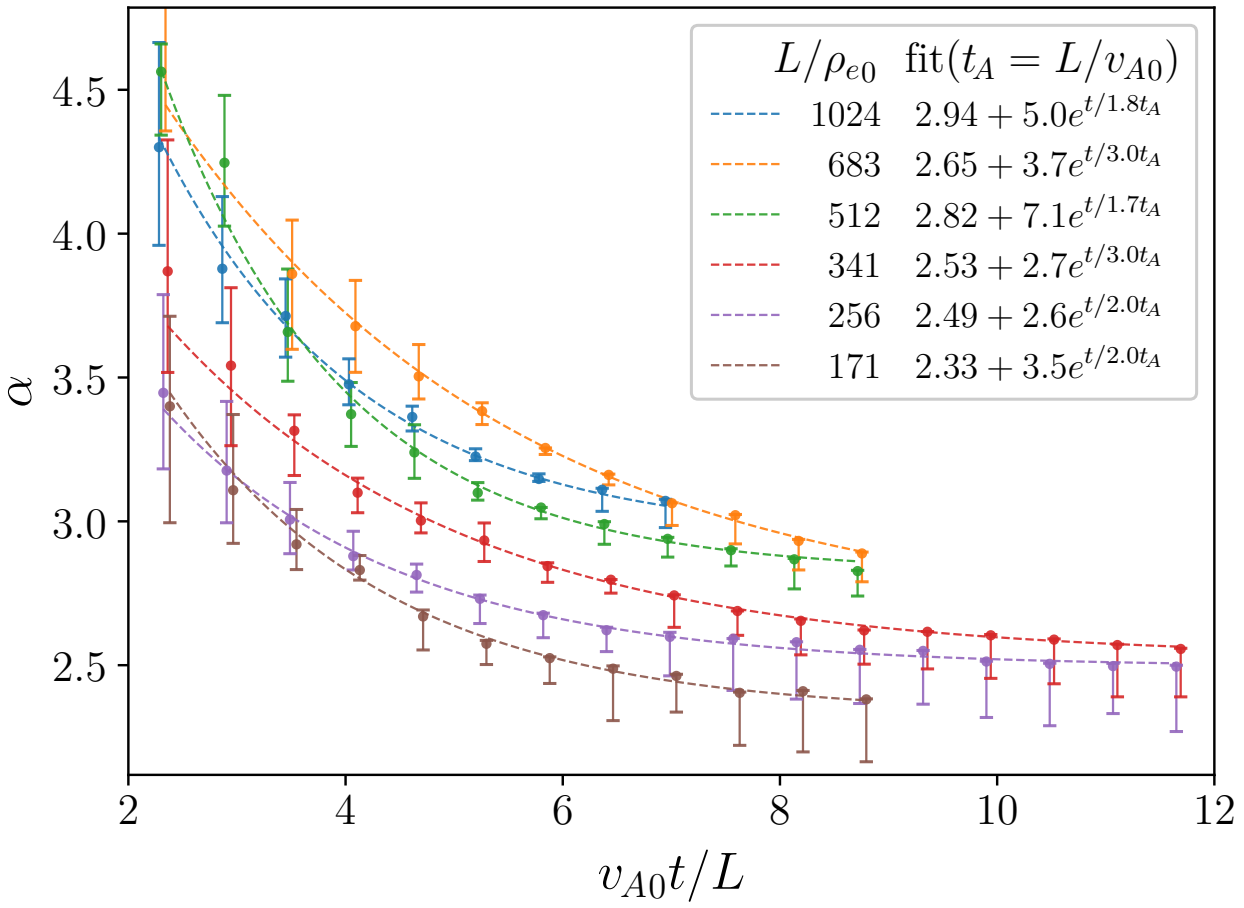

Figure 3.5: Time evolution of power-law indices α for simulations with varying system size L/ρ_{e0} and the same $\sigma_0 = 3/8$, with exponential fits. Each simulation is slightly time-shifted for visibility.

contiguous γ-range of at least half a decade; the vertical lines then indicate the actual $\alpha_{\rm loc}$-range included in this γ-range (neither of which must be symmetric). These "error bars" indicate the range of surrounding $\alpha_{\rm loc}$ values and should not be overly interpreted in a statistical sense. A strongly single-sided "error bar" indicates that α lies in a local minimum of $\alpha_{\rm loc}(\gamma)$. I discard any data point with "error bar" range greater than 40% of α as this indicates an insufficiently converged power-law segment.

Figure 3.5 shows that $\alpha(t)$ decays with time in every simulation. Consequently, I fit each $\alpha(t)$ with a three-parameter exponential decay:

$$\alpha(t) = \alpha_\infty + \Delta_\alpha \exp(-t/\tau), \tag{3.18}$$

and find that the fits are excellent. Although having just warned against statistical interpretation of the "error bars", I nevertheless inversely weigh the data points by their extent, as one would for true statistical standard deviations.

The exponential fit clearly suggests α_∞ as a natural candidate for the late-time "convergent" value of α, i.e., for $\alpha_{\rm fin}$, with few arbitrary parameters. The parameters Δ_α and τ should be considered with respect to the convergence of the power-law tail, separately from its initial formation. Thus, Δ_α is the degree to which α evolves over the convergence timescale τ. By far the most robust parameter is α_∞, which is essentially unchanged for any reasonable variant of the fitting parameters, whereas Δ_α and τ are somewhat sensitive to the fit starting time. The fits are reasonably insensitive to the maximum fitted time as long as the total duration is longer than about $8L/v_{A0}$. Hence, it is reasonable to expect that the trends fitted to the entire simulation as-is would continue if the simulation were somewhat longer. However, one would not expect it to hold indefinitely, as the nonthermal power-law range would eventually be squeezed from both sides until it disappears entirely.

I use the fitted exponential parameters to quantify and compare the per-simulation time-evolution characteristics versus the relative system size L/ρ_{e0}. Figure 3.6 shows that as L/ρ_{e0} increases through a factor of six, α_∞ increases modestly, while Δ_α and τ have no particular trend.

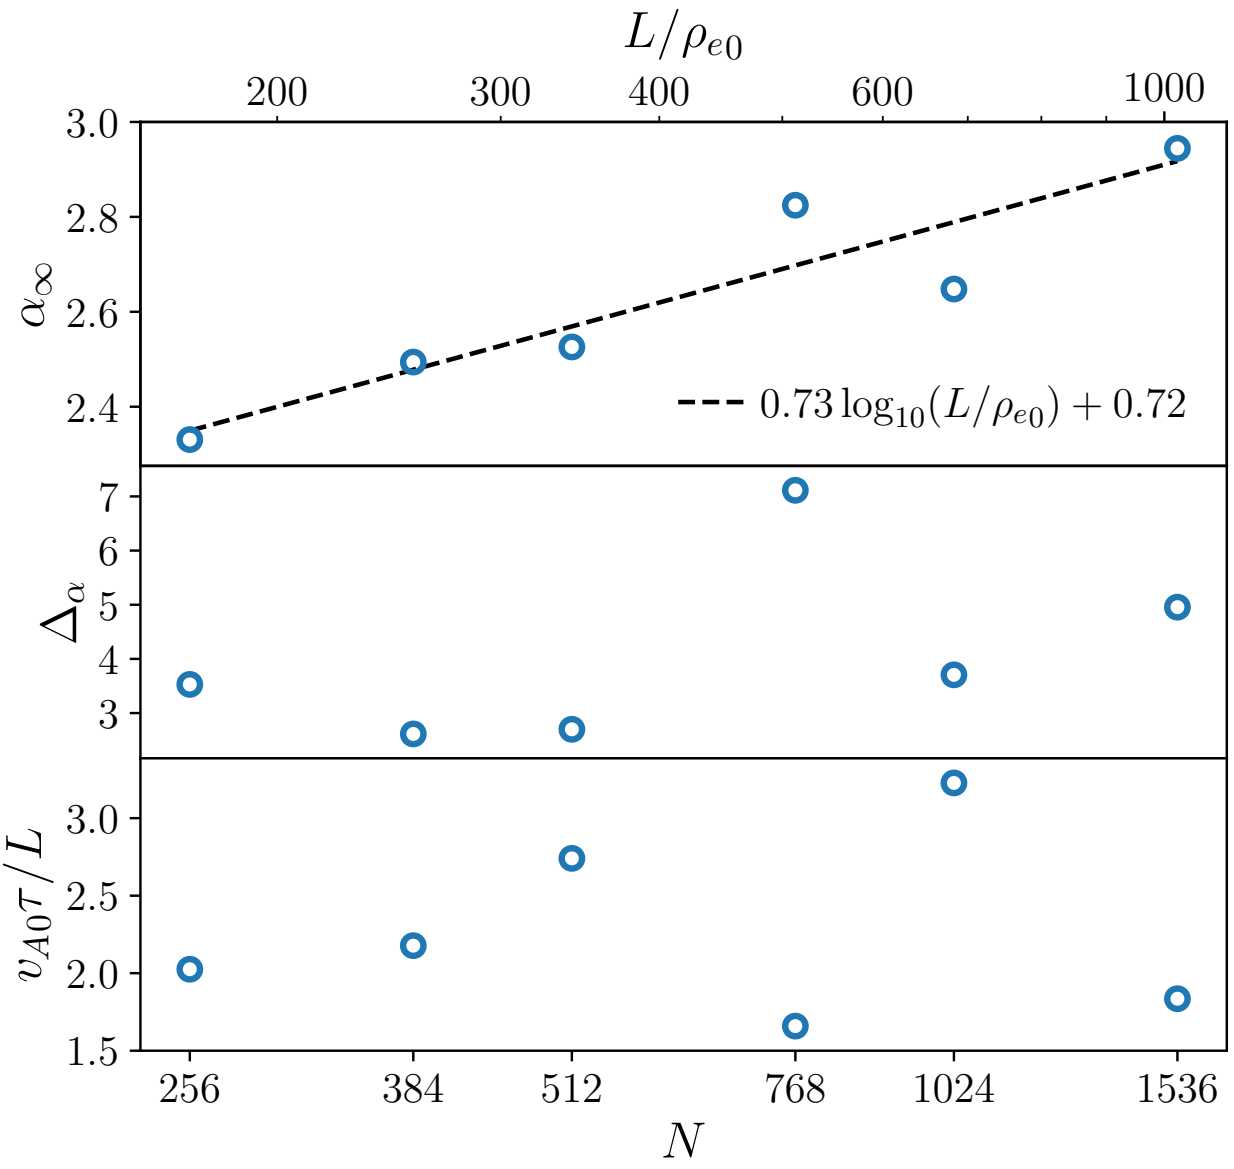

Figure 3.6: Dependence of the three exponential fit parameters of Eq. 3.18 on the relative system size L/ρ_{e0} for simulations with $\sigma_0 = 3/8$.

Table 3.2: Statistical variation in $\alpha(t)$ exponential fit parameters for sets of repeated simulations with the same physical parameters, listed as mean±standard deviation. Both sets had $\sigma_0 = 3/8$ and $N_{\mathrm{ppc}} = 32$.

L/ρ_{e0}	count	α_∞	Δ_α	$v_{A0}\tau/L$
512	7	3.00 ± 0.06	8.9 ± 4	1.3 ± 0.2
256	16	2.68 ± 0.05	7.6 ± 4	1.4 ± 0.4

This matches the basic observation in Figure 3.5 that $\alpha(t)$ tends to shift upwards in magnitude with increasing L/ρ_{e0}, but otherwise behaves similarly between simulations. The convergence timescale τ being a few Alfvén times matches the natural expectation for the system to evolve on Alfvénic timescales. The α_∞ data suggests that the simulations are not converged at the lower L/ρ_{e0} values, but there is not enough information to decide either way for the higher explored L/ρ_{e0}. A reasonable fit for α_∞ as a function of L/ρ_{e0} is

$$\alpha_\infty = 0.73 \log_{10}(L/\rho_{e0}) + 0.72, \tag{3.19}$$

and this is displayed in Figure 3.6. The tendency for larger system size to result in steeper power laws may be attributed to the longer time (in terms of large-scale crossing times) required for particles to reach the system-size-limited energy (Zhdankin et al., 2018b).

To get an idea of the statistical variability of the results, I analyse two sets of computationally cheaper simulations where the same system parameters were used multiple times with different random seeds. One set used 7 simulations with $\sigma_0 = 3/8$, $N_{\mathrm{ppc}} = 32$, and $N = 768$ ($L/\rho_{e0} = 512$); the other, consisting of 16 simulations, differs by using $N = 384$ ($L/\rho_{e0} = 256$). The results are presented in Table 3.2. The scatter in α_∞ is only a few percent, whereas the variation levels of Δ_α and τ are both substantial. Therefore, it is plausible that the α_∞-trend observed in Figure 3.6 is real, but the variations in Δ_α and τ are noise.

3.5.3 Application to the σ_0 scan

I repeat the above analysis on the N=768 σ_0-scan and the N=384 σ_0-scan and combine the results. Figure 3.7 shows the $\alpha(t)$ with exponential fits while Figure 3.8 shows the exponential fit

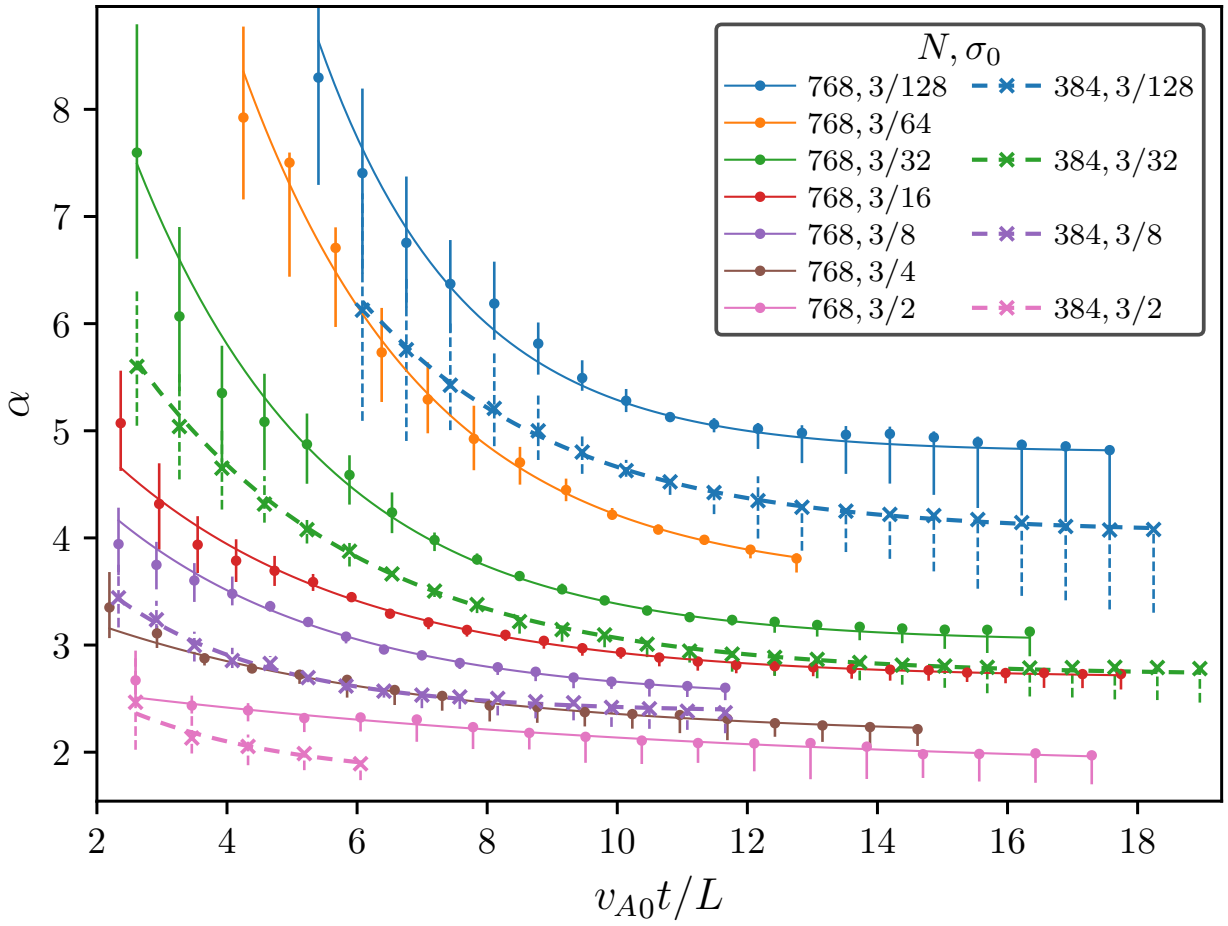

Figure 3.7: Across simulations of varying grid size and initial magnetisation, the particle energy spectrum power-law index time evolution $\alpha(t)$ is well-described by exponential fits of Eq. 3.18. Figure 3.8 presents the best-fit parameters.

parameters versus σ_0, grouped by L/ρ_{e0}. As before, the exponential fits are good in all cases. As σ_0 increases, α_∞ and Δ_α decrease; while τ forms a roughly constant cluster around 3, but gains a large scatter for $\sigma_0 \geq 3/4$. These high-σ_0 values of τ should not be given too much credence. One can observe in Figure 3.7 that ($N = 384$, $\sigma_0 = 3/2$) has few points, while ($N = 768$, $\sigma_0 = 3/2$) has a narrow range of α-values leading to a curve too shallow to reliably characterise the exponential. Furthermore, Table 3.2 shows that τ already varies substantially for $\sigma_0 = 3/8$ and one expects even more relative variation for a shallower curve. Overall, the illustrated trends in the fit parameters reflect the harder power-law distributions and faster system dynamics (with respect to L/c) obtained at higher σ_0, as previously noted by Zhdankin et al. (2017).

Even after combining the two σ_0 scans, there is not enough parameter variation "orthogonal" to σ_0 to properly characterise the effect of L/ρ_{e0} with this data set. However, I note that for the four σ_0 values where there are both $N = 768$ and $N = 384$ simulations, α_∞ is slightly lower in the $N = 384$ case, which is consistent with the Subsection 3.5.2 findings. Figure 3.7 also shows this relationship clearly, with $\alpha(t)$ tending to be lower in magnitude for $N = 384$ simulations compared to their $N = 768$ counterparts.

To quantify the dependencies of α_∞ and Δ_α on the system parameters observed in these simulations, I offer the following simple and practical empirical fits for these functions. Firstly, a test-fit to α_∞ of a log-bilinear function of σ_0 and L/ρ_{e0} indicated that the addition of L/ρ_{e0} was statistically insignificant. This is due to multicollinearity in the parameter space, not lack of effect, as Subsection 3.5.2 shows a significant trend for α_∞ against L/ρ_{e0}. Consequently, I proceed with fits against just σ_0, noting also that partitioning the data by L/ρ_{e0} or L/d_{e0} would not give substantially different results.

The best log-linear fit of α_∞ as a function of σ_0 is:

$$\alpha_\infty \approx 1.9 - 1.2 \log_{10} \sigma_0 \,. \tag{3.20}$$

Alternatively, the offset power-law

$$\alpha_\infty \approx 1 + \sigma_0^{-1/3} \tag{3.21}$$

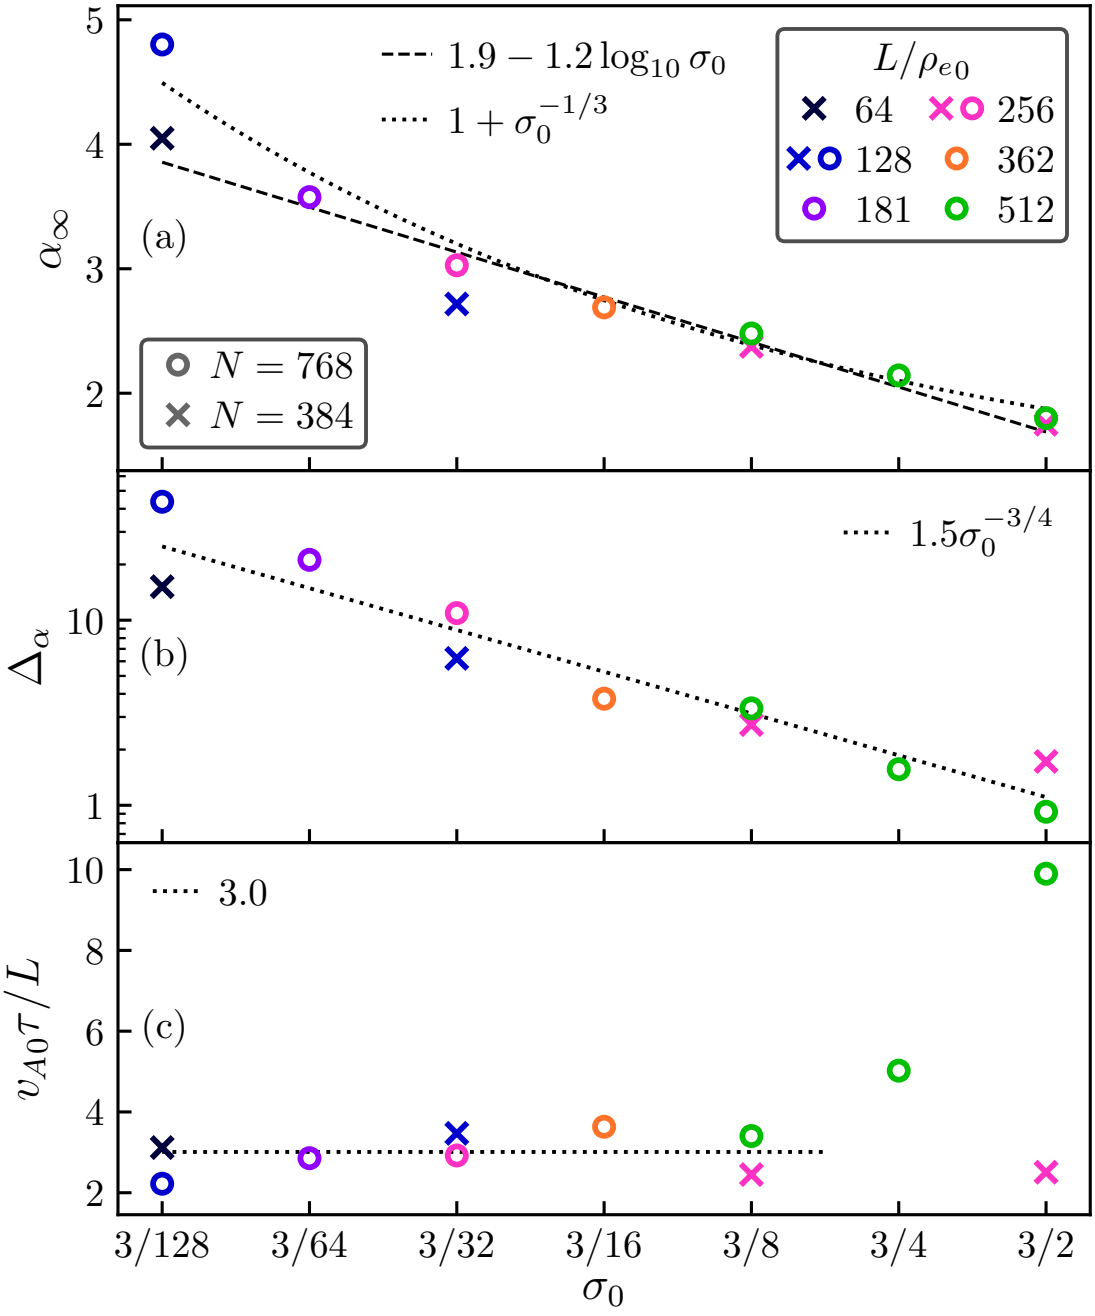

Figure 3.8: Dependence of the three exponential fit parameters of Eq. 3.18 on the initial dimensionless system parameters σ_0 and L/ρ_{e0}.

is also a good fit (the motivation for the constant term being that it is a physical lower limit for α). These fits are displayed in Figure 3.8a.

I also find that

$$\Delta_\alpha = 1.5\sigma_0^{-3/4} \tag{3.22}$$

is a good fit for Δ_α, and this is plotted in Figure 3.8b.

I stress that more simulations would be needed to fully map out the dependence of the fit parameters over the 2D $(\sigma_0, L/\rho_{e0})$ parameter space, and also to further quantify the random variation of repeated simulation runs with the same initial system parameters. I leave these further simulations required for more comprehensive results to future work.

3.6 Tests of diffusive particle acceleration

3.6.1 Test procedure

This Section details my procedure to measure the FP coefficients and their subsequent use in testing if an energy diffusion-advection equation (3.2) is appropriate to model particle acceleration in these turbulence simulations. This was found to be true for a single simulation in my previous study Wong et al. (2020), and I improve on it here with more comprehensive tests on a larger parameter space.

The tracked particles described in Subsection 3.3.2 are integral to this analysis. First, I remove the large $\mathbf{E}\times\mathbf{B}$-drift-induced particle energy oscillations — which are incompatible with an energy diffusion model — by transforming to the particle's $\mathbf{E}\times\mathbf{B}$-drift frame. This is also described in Subsection 3.3.2. Then, for each simulation, I select approximately 100 evenly-distributed time instants, and at each instant, divide particles into energy bins with edges at 10% intervals, at $\gamma = 10(1.1^0, 1.1^1, 1.1^2, ...)$. I refer to each bin by its originating time instant t_0 and its arithmetic bin centre energy γ_0. For the population of particles in each bin, I compute the following quantities as functions of elapsed time $\Delta t \equiv t - t_0$: the energy standard deviation $\delta\gamma_{\rm rms}$, energy variance $\overline{\delta\gamma^2} \equiv \delta\gamma_{\rm rms}^2$, mean energy $\overline{\gamma}$, and change in the mean energy $\Delta\overline{\gamma} \equiv \overline{\gamma}(t) - \overline{\gamma}(t_0)$ (where $\overline{\gamma}(t_0) \approx \gamma_0$).

Examining the evolution of the bin particle population and these moments is a standard technique for analysing systems for FP-type behaviour and subsequently measuring the FP coefficients (Siegert et al. 1998, Friedrich et al. 2000). It is these procedures that require tracked particles and that cannot be done by just using summary statistics such as the overall particle energy distribution.

For the classical Brownian-type diffusion of the FP equation, the bin energy variance increases linearly with time in the limit of infinitesimal elapsed time: $\lim_{\Delta t \to 0} \overline{\delta\gamma^2} \propto \Delta t$. Accordingly, a primary contraindication to an FP model would be if the energy trajectories exhibit anomalous diffusion instead, where the energy histories are dominated by rare but large jumps such that the energy variance intrinsically scales nonlinearly with elapsed time. This has been suggested to be the case in plasma turbulence, for instance, in Isliker et al. (2017b) and Isliker et al. (2017a). Distinguishing definitively between conventional and anomalous diffusion can, however, be difficult in practice, especially if anomalous effects are marginal, or only manifest rarely, in isolated or intermittent situations. There are several reasons for this, but the most relevant one here is that if the FP coefficients vary with energy, this can cause nonlinear variance scaling at finite (but arbitrarily short) elapsed time while retaining a classical Brownian diffusion-advection equation (this can happen with relevant and non-pathological scalings; I will give some concrete examples shortly, see also Friedrich et al. 2002). Therefore, I take the approach of a consistency check: I measure the FP coefficients assuming standard Brownian diffusion, and then see if a numerical solution using the measured coefficients reproduces the PIC simulation particle energy distribution evolution both globally and for individual bins (Subsection 3.6.3). If the agreement is good, one can be confident that if anomalous diffusion is present, it does not substantially affect the particle acceleration process, at least from the point of view of overall distribution characteristics like power law tails (keeping in mind that this is the astrophysical observable).

To measure the FP coefficients from the bin moments, I use the following idealised equations for the latter's time evolution, which are obtained from the FP equation (3.2) with a $\delta(\gamma - \gamma_0)$

initial condition and in the infinitesimal elapsed time approximation:

$$\Delta\overline{\gamma}(\gamma_0, \Delta t) = [\partial_\gamma D|_{\gamma_0} + A(\gamma_0)]\Delta t \equiv M(\gamma_0)\Delta t \tag{3.23}$$

$$\overline{\delta\gamma^2}(\gamma_0, \Delta t) = 2D(\gamma_0)\Delta t. \tag{3.24}$$

However, the smallest elapsed time for which the moments are reliable is on the order of a gyroperiod. This is because, although the $\mathbf{E} \times \mathbf{B}$-based procedure mentioned above removes the bulk of the oscillation, some residual small amplitude oscillation remains, possibly due to other types of drifts. I also use linear fits to the moments to reduce noise, rather than measuring at a single point, and so, for the reason just mentioned, I begin these fits after one gyroperiod $2\pi\gamma_0 mc/eB_{\text{rms}}$ (with B_{rms} sampled at t_0). The fit end times need further consideration, because longer fits better suppress noise, but also amplify the finite time effects which lead to nonlinear time dependence of the bin moments (Gottschall and Peinke, 2008).

To limit the impact of the nonlinear moments evolution from finite-time effects, I terminate the fit once $\delta\gamma_{\text{rms}}/\gamma_0 = 0.3$ (such that $\overline{\delta\gamma^2}/\gamma_0^2 \approx 0.1$) or $|\Delta\overline{\gamma}|/\gamma_0 = 0.1$. This uses γ_0 as an intuitive reference for the bin being relatively thin and close to its original position, or equivalently, as a heuristic for how rapidly the FP coefficients vary in energy space. These limits work for Fermi acceleration, where $D \propto \gamma^2$ (temporarily ignoring A for simplicity). Here, the strong $D(\gamma)$ scaling produces effective superdiffusion that indeed begins when $\delta\gamma_{\text{rms}} \sim \gamma_0$ (one may verify this by referring to the analytical solutions for the equivalent geometric Brownian motion). On the other hand, consider the Ornstein–Uhlenbeck process, where D is constant with respect to γ while A has a negative slope and arbitary intercept (I show in Subsection 3.6.3 that this is a reasonable conceptual model in certain energy ranges). In this process, the variance becomes subdiffusive after an elapsed time inversely proportional to the advection coefficient slope and hence at no particular value of $\delta\gamma_{\text{rms}}/\gamma_0$ or $\Delta\overline{\gamma}/\gamma_0$. However, there is little one can do about this without turning to a complicated iterative procedure, and so, since $D \propto \gamma^2$ is the most important case for this study, I keep the fit limits referenced to γ_0 as described above. The consistency check should also help to indicate whether this fitting range is reasonable.

One must also consider the FP coefficients' time dependence because these simulations are not true steady states. This is due to overall system heating, which is described in Section 3.4. This heating occurs on the Alfvén timescale, so I cap the fit time to $\Delta t < L/v_{A0}$. Finally, I enforce a minimum duration of $0.1L/c$, for the rare cases where the above limits result in a zero or negligible length fit.

I use the above fit range prescription to fit $\Delta\overline{\gamma}$ and $\overline{\delta\gamma^2}$ to linear functions of elapsed time, then read off D and M from the fitted slopes using (3.23) and (3.24). I repeat for all bins to obtain D and M for every point in the grid of (γ_0, t_0) values, and elide the subscripts to consider them plainly functions of energy and time: $D(\gamma, t)$ and $M(\gamma, t)$. Finally, I obtain the advection coefficient using $A = M - \partial_\gamma D$, according to (3.23).

The rest of this Section examines the bin energy moments (Subsection 3.6.2) and the consistency check results (Subsection 3.6.3), leaving detailed analysis of the FP coefficients to Section 3.7 and Section 3.8.

3.6.2 Bin energy moments

Figure 3.9 shows the bin energy variance as a function of elapsed time, $\overline{\delta\gamma^2}(\Delta t)$, compensated by $c\Delta t/L$, for a selection of simulations from the N=768 σ_0-scan. The bins shown are initialised at $v_{A0}t_0/L = 6$ and range from thermal ($\gamma_0/\gamma_{\rm pk} \approx 1$) to nonthermal ($\gamma_0/\gamma_{\rm pk} \gg 1$) energies. The plots show the linear fits that will be used to obtain D via (3.24), which seem reasonable considering the heavy late-time bias inherent in a logarithmic scale. They also show contours of constant $\delta\gamma_{\rm rms}/\gamma_0$ and $\Delta\overline{\gamma}/\gamma_0$ (see caption for details), which give a sense of the relative evolution timescales as a function of initial bin energy.

Generally, while $\delta\gamma_{\rm rms}/\gamma_0 \ll 1$, $\overline{\delta\gamma^2}$ increases roughly linearly with Δt, suggestive of classical diffusion. For low initial magnetisations $\sigma_0 \lesssim 3/32$, there is substantial energy subdiffusion at thermal energies $\gamma_0 \lesssim \gamma_{\rm pk}$, with power law indices ∼0.7–0.9 depending on σ_0, and which tends to become pronounced after intermediate times when $\delta\gamma_{\rm rms}/\gamma_0 > 0.3$. Subsection 3.6.3 shows that FP evolution reproduces this. There is also superdiffusion at late times, once $\delta\gamma_{\rm rms} \sim \gamma_0$, which is

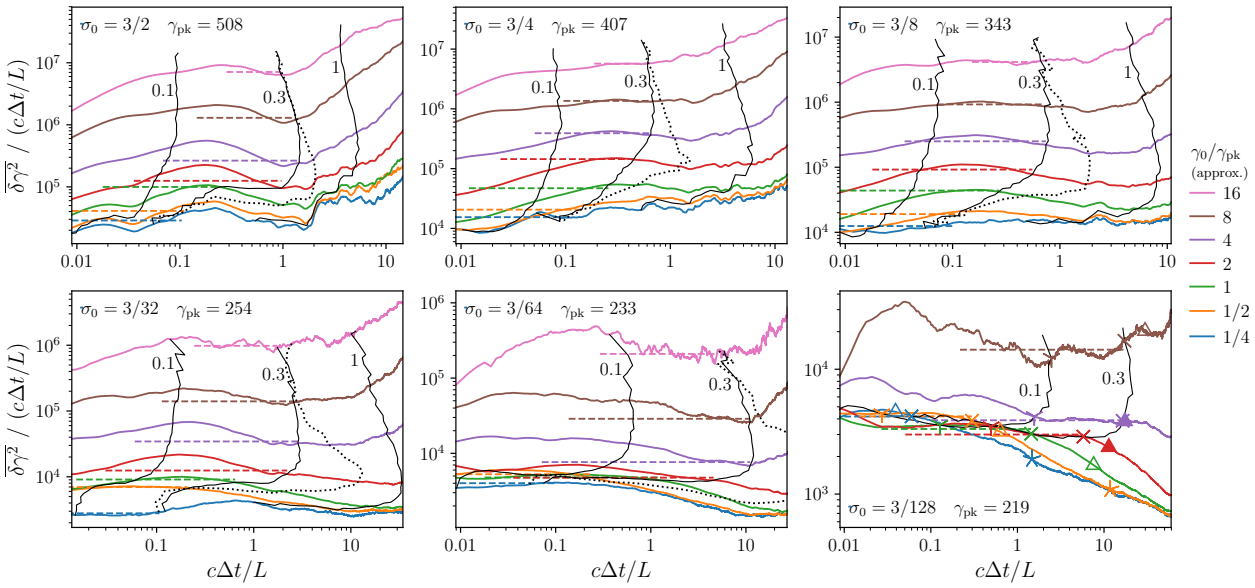

Figure 3.9: Bin energy variance $\overline{\delta\gamma^2}(\Delta t)$, compensated by $c\Delta t/L$, where each panel corresponds to a simulation with different σ_0. All simulations have $N = 768$ and the bins are initialised at $v_{A0}t_0/L = 6$. Coloured dashed lines show the fits I use to obtain D; the line extent shows the fit range. Black solid lines are contours of constant $\delta\gamma_{\rm rms}/\gamma_0$ annotated by the corresponding $\delta\gamma_{\rm rms}/\gamma_0$ value $\in \{0.1, 0.3, 1\}$. Black dotted lines indicate $\Delta\overline{\gamma}/\gamma_0 = +0.1$. These contours also use data from bins in between the ones I selected for plotting variance. The $\sigma_0 = 3/128$ panel does not show the $\Delta\overline{\gamma}/\gamma_0$ contour because $\Delta\overline{\gamma}/\gamma_0$ is substantially negative for some bins in that simulation. For that panel only, there are additional marks: for each variance line; an open triangle for $\Delta\overline{\gamma}/\gamma_0 = +0.1$; a filled triangle for $\Delta\overline{\gamma}/\gamma_0 = -0.1$; and three, four, and five armed spindles for $\delta\gamma_{\rm rms}/\gamma_0 = \{0.1, 0.3, 1\}$ respectively.

particularly apparent at nonthermal energies and at higher σ_0. As mentioned in Subsection 3.6.1, this would be consistent with (but not exclusive to) $D \propto \gamma^2$. It is possible to read off the rough diffusion coefficient energy scaling from the contours of constant $\delta\gamma_{\rm rms}/\gamma_0$. A vertical contour implies, through (3.24), that $D \propto \gamma^2$; there are such near-vertical segments at high energies in all simulations. Similarly, a contour slanting from lower left to upper right, or lower right to upper left, mean respectively a shallower or steeper scaling than γ^2.

Figure 3.10 shows the change in mean bin energy as a function of elapsed time, $\Delta\overline{\gamma}(\Delta t)$. The traces are generally linear. The particle bin with minimum $\Delta\overline{\gamma}(\Delta t)$ slope for each simulation (including negative slope as in the case of $\sigma_0 = 3/128$) tends to occur at central energies $\gamma_0/\gamma_{\rm pk} \approx 2$–$4$ rather than the lowest energies. For the lowest $\sigma_0 = 3/128$ simulation, the $\gamma_0/\gamma_{\rm pk} \leq 1/2$ bins have flattening curves. This occurs at long times and is likely due to coefficient energy dependence; I will check in detail for the $\gamma_0/\gamma_{\rm pk} = 1/2$ bin in the next Subsection.

3.6.3 Consistency checks

In this Subsection, I test if numerical FP evolution with the measured coefficients reproduces the PIC tracked particle distribution evolution, first for a single bin, and then for the entire particle energy distribution.

I choose a bin which shows subdiffusion, corresponding to the $\gamma_0/\gamma_{\rm pk} \approx 1/2$ line in the $\sigma_0 = 3/128$ panel of Figure 3.9, and set the initial condition to a matching narrow rectangular distribution. Figure 3.11 shows the subsequent energy distribution evolution along with the nearby FP coefficient values while Figure 3.12 shows the bin energy moments as a function of time. The excellent match suggests that any effective departure from classical diffusion is due to the FP coefficients' energy dependence effected over finite elapsed time, and not anomalous diffusion. I note that the coefficients are remarkably stable over this time period, as is consistent with the slow dynamics of low magnetisation simulations, and so coefficient time dependence is likely unimportant.

To get a better idea of what aspect of the FP coefficients' energy dependence results in this subdiffusion, I compare to the Ornstein–Uhlenbeck process mentioned in Subsection 3.6.1.

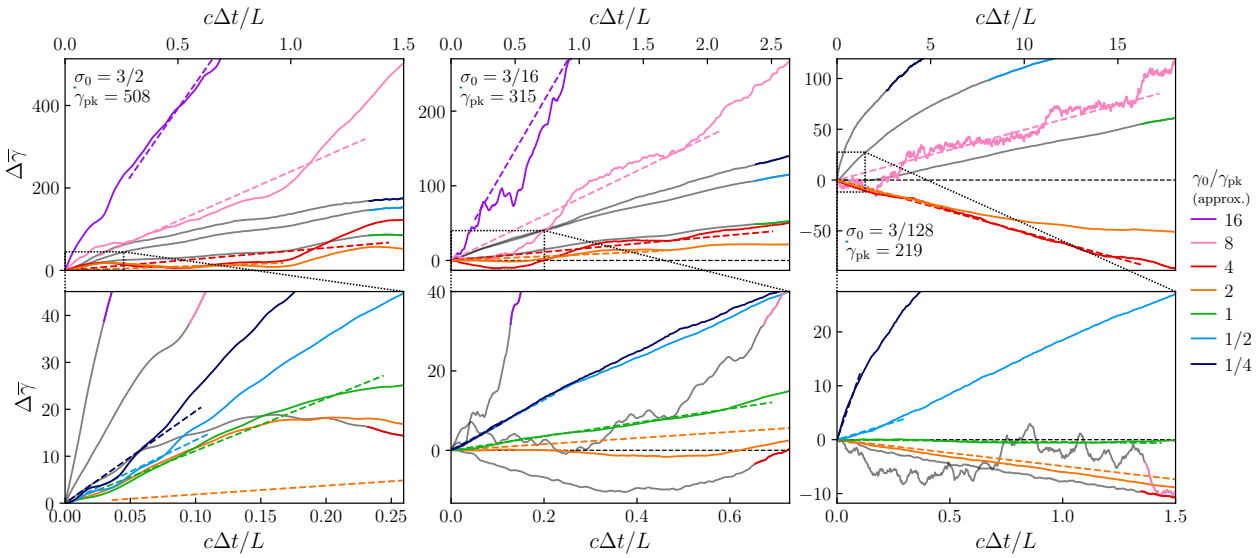

Figure 3.10: Change in bin mean energy $\Delta\overline{\gamma}(\Delta t)$ for $\sigma_0 \in \{3/2, 3/16, 3/128\}$ simulations with $N = 768$, and bins initialised at $v_{A0}t_0/L = 6$. Each column is for one simulation. The top row is for viewing high-energy bins while the bottom row zooms in on the low-energy bins; the zoomed area is shown by a black dotted box. In the top row, high-energy bin lines are fully coloured while only the tips of low-energy bin lines are coloured; this is reversed in the bottom row.

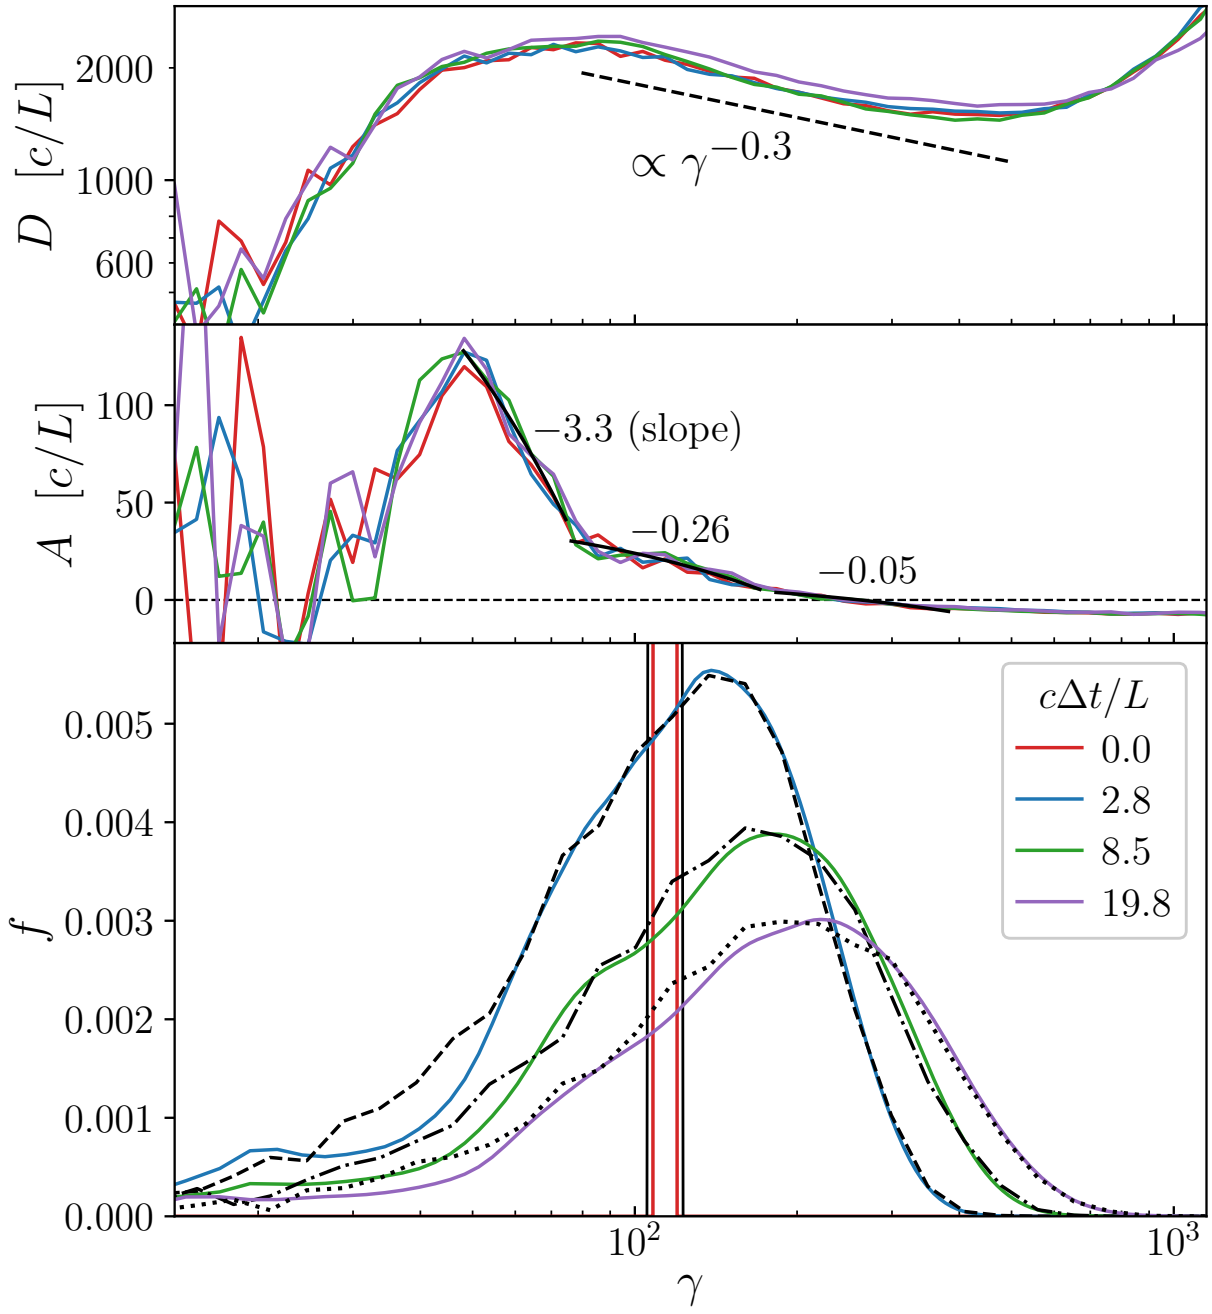

Figure 3.11: Histogram of single-bin tracked particle energy over time compared to the evolution of a matching initial narrow distribution using the FP equation with measured coefficients. There is an excellent match except at lowest energies, where $\partial_\gamma D$ is noisy and hence so is A. The local measured FP coefficients are also shown: the negative slope of A produces effective subdiffusion. The bin is taken at $v_{A0}t_0/L = 6$ and has centre energy $\gamma_0/\gamma_{\rm pk} \approx 1/2$. The simulation has $N = 768$ and $\sigma_0 = 3/128$. The bin distribution function is normalised such that it integrates to 1.

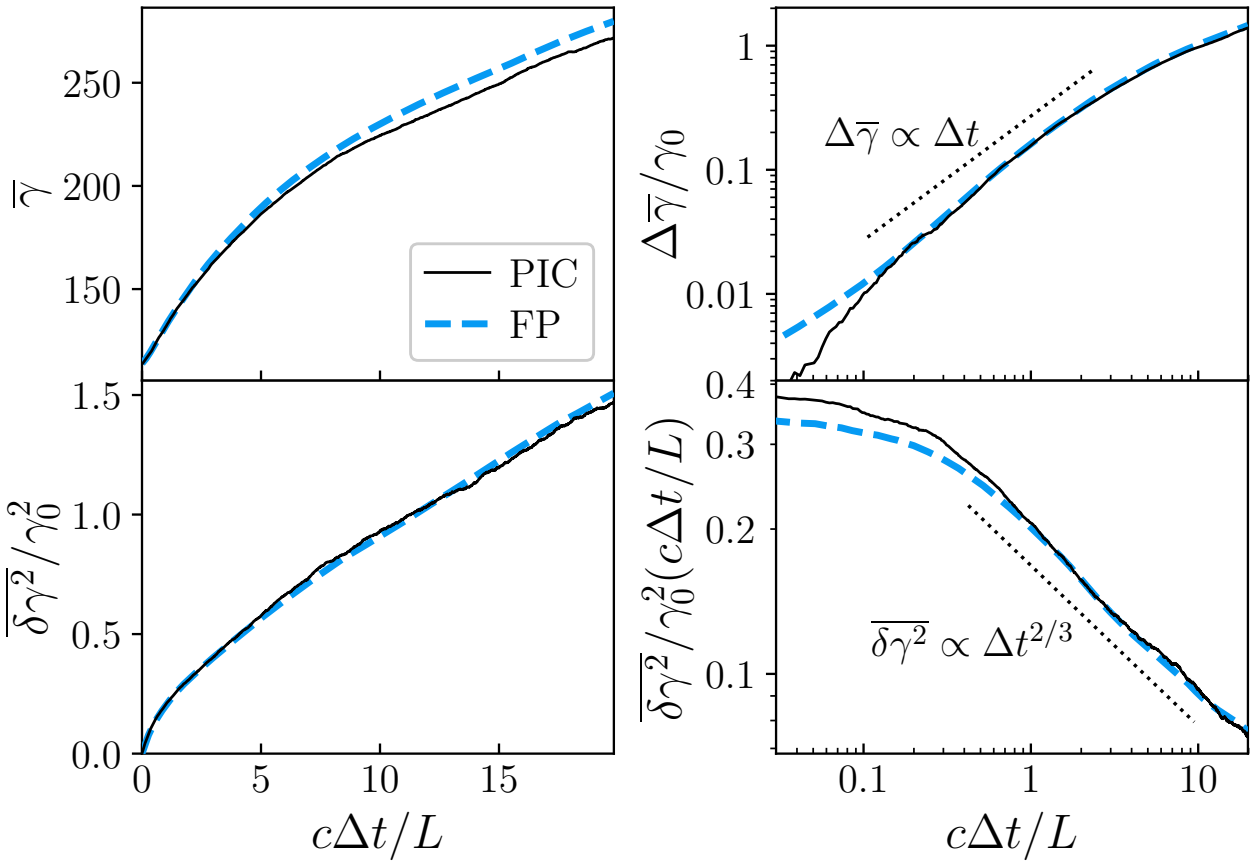

Figure 3.12: Comparison of PIC tracked particle and FP evolution single-bin mean and variance versus time in linear (left) and log (right) scales. The log-scale variance panel is compensated by $c\Delta t/L$. There is excellent agreement between the moments from the binned tracked particles and those obtained by evolving a matching initial narrow distribution using the FP equation with measured coefficients. For large elapsed time, the variance evolves with a subdiffusive scaling $\overline{\delta\gamma^2} \propto \Delta t^{2/3}$. The mean change does not have a definite power-law relationship with time ($\propto \Delta t$ line shown for reference). The bin is taken at $v_{A0}t_0/L = 6$ and has centre energy $\gamma_0/\gamma_{\rm pk} \approx 1/2$. The simulation has $N = 768$ and $\sigma_0 = 3/128$.

This should be a reasonable model because in the vicinity of this bin, the diffusion coefficient has only a weak scaling of roughly $\gamma^{-0.3}$ while the advection coefficient has a negative slope. For an advection coefficient slope of $-\theta$ (i.e., $A = \text{const} - \theta\gamma$), the Ornstein–Uhlenbeck process elicits substantial variance subdiffusion after roughly $\Delta t \sim |2\theta|^{-1}$. At the chosen bin centre energy, $\theta \approx -0.3c/L$, giving a nonlinear timescale of $\sim 2L/c$, which is a reasonable match to when the variance in Figure 3.12 becomes significantly nonlinear, considering that this estimate heavily simplifies the FP coefficients' energy dependence. Furthermore, Figure 3.9 shows that, up to at least $\gamma_0/\gamma_{\rm pk} = 2$, the subdiffusive departure becomes later as initial bin energy increases, which is consistent with Figure 3.11 showing that the magnitude of θ decreases with increasing energy. These observations suggest that the observed subdiffusion is caused mainly by the negative slope in A.

I now test whether the numerical solution of the FP equation with these tracked-particle-based coefficients can reproduce the actual evolution of the full particle energy distribution observed in the PIC simulations. Figure 3.13 shows the numerical FP evolution results for the simulations with $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ from the N=768 σ_0-scan. The initial condition is the initial thermal distribution. I find that the FP equation with the measured time-dependent coefficients accurately reproduces the evolution of $f(\gamma)$ from the PIC simulations over the entire σ_0-range of simulations. This provides solid evidence that the particle acceleration resulting in the nonthermal distribution can be completely modelled by the FP equation (3.2).

It is worth noting that, without the advection coefficient, the FP-based distributions have a substantially altered shape, with excessive heating and generally slightly harder power laws compared to the PIC results. As was stated in Wong et al. (2020), but here shown for a range of σ_0, the advection coefficient plays an important role in maintaining the shape of the particle energy distribution.

3.7 Diffusion coefficient

My previous study, Wong et al. (2020), measured the energy diffusion coefficient D in a single large simulation with moderate initial magnetisation, finding that D scales as γ^2 in the

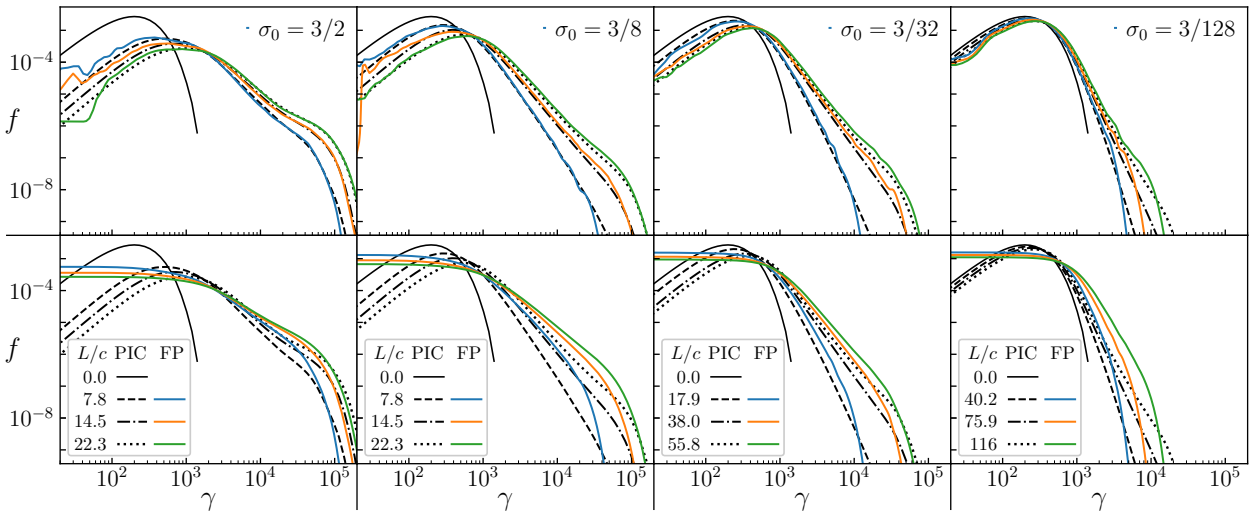

Figure 3.13: Top row: numerical FP evolution with measured coefficients vs PIC particle energy distributions for the global particle energy distribution. The FP evolution was initialised with the initial thermal distribution. The agreement between FP and PIC is excellent across a large range of σ_0. Bottom row: same as top row except with $A = 0$. This results in clear discrepancies between FP and PIC.

nonthermal region and somewhat more shallowly at lower energies. This Section characterises $D(\gamma, t)$ in simulations with different parameters; I examine the time and energy dependence of D, and to what extent the aforementioned scaling characteristics persist. Finding universal γ^2 scaling in the nonthermal range, I extract the γ^2 prefactor D_0 and compare it to NTPA theory predictions.

To investigate the effects of σ_0 and L/ρ_{e0}, I employ the simulations from the N=768 σ_0-scan and the system-size scan. I also use a smaller N=384 $N_{\rm ppc}$-scan to test convergence. Please refer to Table 3.1 and surrounding text for the complete parameter sets.

It is only possible to measure D for energy bins with enough tracked particles, which is dictated by the hardness and extent of the particle energy distribution at the measurement time. To limit statistical noise, results from bins with less than 10 particles are discarded. Nevertheless, the results are primarily established by measurements at energies for which there are ample statistics, with thousands of tracked particles per bin.

3.7.1 Magnetisation dependence

This Subsection uses the N=768 σ_0-scan to analyse the dependence of $D(\gamma, t)$ on both the initial magnetisation σ_0 and the instantaneous magnetisation $\sigma(t)$. Figure 3.14 shows $f(\gamma, t)$ and $D(\gamma, t)$ for the simulations with $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ at times $t_0 \approx \{3, 6, 9\}L/v_{A0}$. These plots have several types of markers indicating different physical length scales. Each length is converted to the particle energy that would result in a gyroradius equal to that length by multiplying by $eB_{\rm rms}/mc^2$. The length scales are:

(1) The grid resolution Δx. At energies near and below the grid resolution energy, data are likely to be affected by numerical noise.

(2) The system-size gyroradius limit $L/2$, corresponding to the system-size energy limit $\gamma_{\rm max}(t) \equiv LeB_{\rm rms}/2mc^2$. This is close to the driving wavelength ($L_d = L$ for the simulations in this Subsection) and so in the vicinity of $\gamma_{\rm max}(t)$ one may see both finite system size and finite driving scale effects.

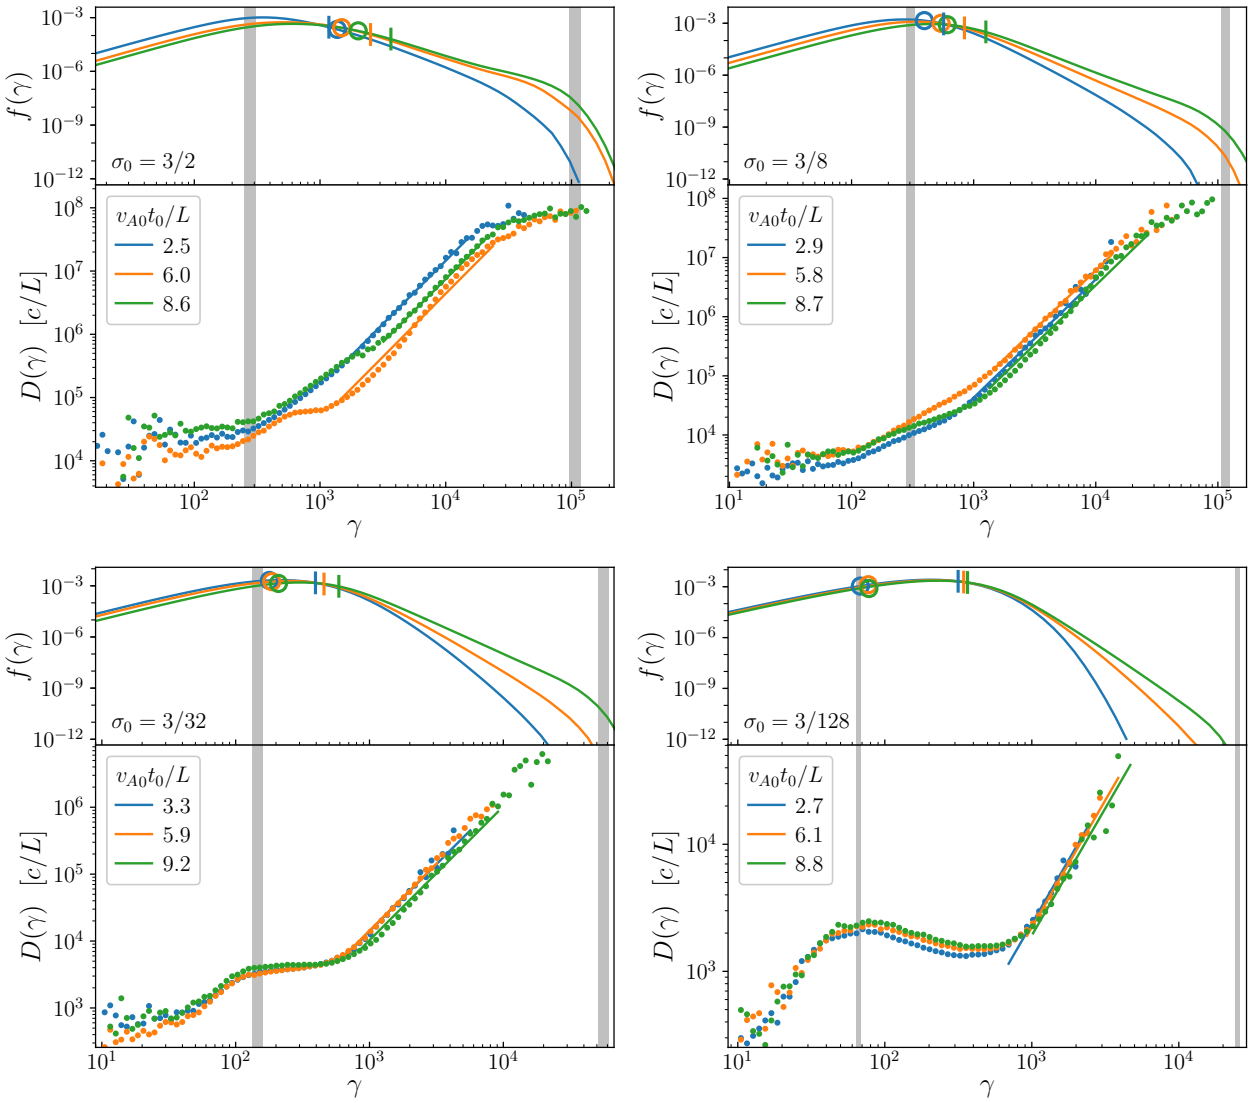

Figure 3.14: The diffusion coefficient D (lower subpanels) and particle energy distribution f (upper subpanels) as a function of energy and time across the N=768 σ_0-scan, displaying the relevant plasma and simulation length scales (converted to particle energies via the Larmor radius) as well as the $D \sim \gamma^2$ fits to the nonthermal region that are used to extract D_0 (solid coloured lines in lower subpanels). The plasma length scales are the average Larmor radius, corresponding to the average particle energy (small vertical bars); and the skin-depth (open circles). The simulation length scales are the grid resolution Δx (vertical grey strip at lower energy); and the system-size limit $L/2$, corresponding to the high-energy cutoff $\gamma_{\rm max}$ (vertical grey strip at higher energy). The strips cover the 10th to 90th percentile values for the respective quantity recorded between the times corresponding to the first and last plotted t_0 instance in each plot.

(3) The skin depth $d_e(t) \equiv (\gamma_{\rm avg} mc^2/4\pi n_0 e^2)^{1/2}$.

(4) The characteristic gyroradius $\rho_e(t) \equiv \gamma_{\rm avg} mc^2/eB_{\rm rms}$ (corresponding to $\gamma_{\rm avg}$ by definition).

I observe that $D(\gamma)$ can be divided into two segments by the beginning of the $f(\gamma)$ nonthermal power law. In the nonthermal power-law, $D \propto \gamma^2$ for all σ_0 (flattening slightly at highest energies, perhaps due to encroaching on $\gamma_{\rm max}$). The low-energy behaviour depends on σ_0. For higher $\sigma_0 \gtrsim 3/8$, the low-energy scaling is consistent with a power law shallower than γ^2. For lower $\sigma_0 \lesssim 3/32$, there is an intermediate energy region with flatter scaling (see $\sigma_0 = 3/32$) or inverted scaling (see $\sigma_0 = 3/128$), and then a steeper section at lowest energies.

A natural explanation for the existence of the intermediate energy region in lower-σ_0 simulations is that certain length-energy scales are substantially separated only in those simulations. However, near each end of the intermediate region there are multiple physical energy scales, making it difficult to unambiguously identify which of them are relevant. For instance, some plausible delineators of the intermediate region are: d_e and ρ_e, the grid scale and the larger of d_e and ρ_e, and the grid scale and the beginning of the nonthermal power law. A more thorough analysis with a targeted series of simulations would be needed to disentangle the effects of the different length-energy scales on $D(\gamma)$. I leave such work to future studies.

Putting aside the low-energy behaviour, I now focus on the prefactor of the nonthermal scaling $D_0 \equiv D(\gamma)/\gamma^2$, its time evolution, and dependence on σ_0 and $\sigma(t)$. The value of D_0 and its relation to other system parameters is important for testing particle acceleration theories, including the ones described previously in Subsection 3.2.2. Handily, D_0 is the inverse of the diffusive acceleration time for a particle in the nonthermal energy range; its energy doubles in a time of order D_0^{-1}, absent other effects (which can be seen by inserting $D = D_0\gamma^2$ into Equation 3.23).

I now study the relationship between $D_0(t)$ and the instantaneous magnetisation $\sigma(t)$. This is of interest because analytical NTPA theories predict the scaling of D_0 with parameters such as the instantaneous magnetisation, or equivalently, the Alfvén velocity. For example, simple standard second-order Fermi acceleration theories (see Subsection 3.2.1) yield $D_0 \propto u_A^2$ where $u_A \equiv v_A(1 -$

$v_A^2/c^2)^{-1/2}$, whereas more sophisticated modern analyses predict $D_0 \propto u_A^3$ (Demidem et al., 2020).

I measure D_0 as a fine-grained function of time by fitting $D = D_0\gamma^2$ to the high-energy region in each sample of $D(\gamma, t)$. Figure 3.15 presents $D_0(t)$ and $\sigma(t)$ for the N=768 σ_0-scan. The resemblance between the two panels suggests a direct relationship between D_0 and σ, in accordance with theoretical predictions. Nevertheless, let us first examine them separately.

The $\sigma(t)$ trajectories show a characteristic pattern of an initial increase to roughly twice σ_0 followed by decay. The initial increase is caused by δB going from zero to roughly B_0 as the turbulence is established, thus doubling the total magnetic energy density (as described in Section 3.3, $\delta B \sim B_0$ is by choice). Thereafter, as the plasma heats over time in these simulations, the relativistic enthalpy density increases, and this causes the magnetisation (being inversely proportional to the enthalpy) to decline. I refer to section 4.2 of Zhdankin et al. (2018a) for further details about $\sigma(t)$ time dynamics. The initial decay in σ is much more rapid for higher-σ_0 simulations; this is because the ratio of dissipated magnetic energy (per dynamical time) to initial thermal energy increases in proportion to σ, as described quantitatively in section 4.2 of Zhdankin et al. (2018a). As a consequence, simulations beginning at $\sigma_0 \geq 3/8$ have similar σ at late $v_{A0}t/L$. Moreover, in absolute time units ct/L, the highest σ_0 simulations (e.g., $\sigma_0 = 3/2$) arrive at lower $\sigma(t)$ values faster than slightly lower σ_0 simulations (e.g., $\sigma_0 = 3/4$).

As alluded to earlier, the time evolution of D_0 is similar to that of σ: initially, D_0 increases briefly, and thereafter it tends to decline. This initial increase lasts for around one to three Alfvén crossing times. The subsequent decline reflects the overall heating trend with both $D(\gamma)$ and $f(\gamma)$ moving to the right as the simulation progresses. However, there are significant and sustained ($\sim L/v_{A0}$ duration) fluctuations in D_0, up to the level of about a factor of two, such that this $D(\gamma)$ movement is not entirely monotonic. Overall, the values of D_0 are higher for higher-σ_0 simulations. However, after about five Alfvén crossing times, the $D_0(t)$ lines for $\sigma_0 \geq 3/8$ simulations significantly overlap. Meanwhile, those for $\sigma_0 \leq 3/32$ are more distinctly separated.

Cross referencing the values of $D_0(t)$ and $\sigma(t)$, I obtain $D_0(\sigma)$, and plot this in Figure 3.16. I exclude points earlier than two Alfvén crossing times, before which turbulence has not fully

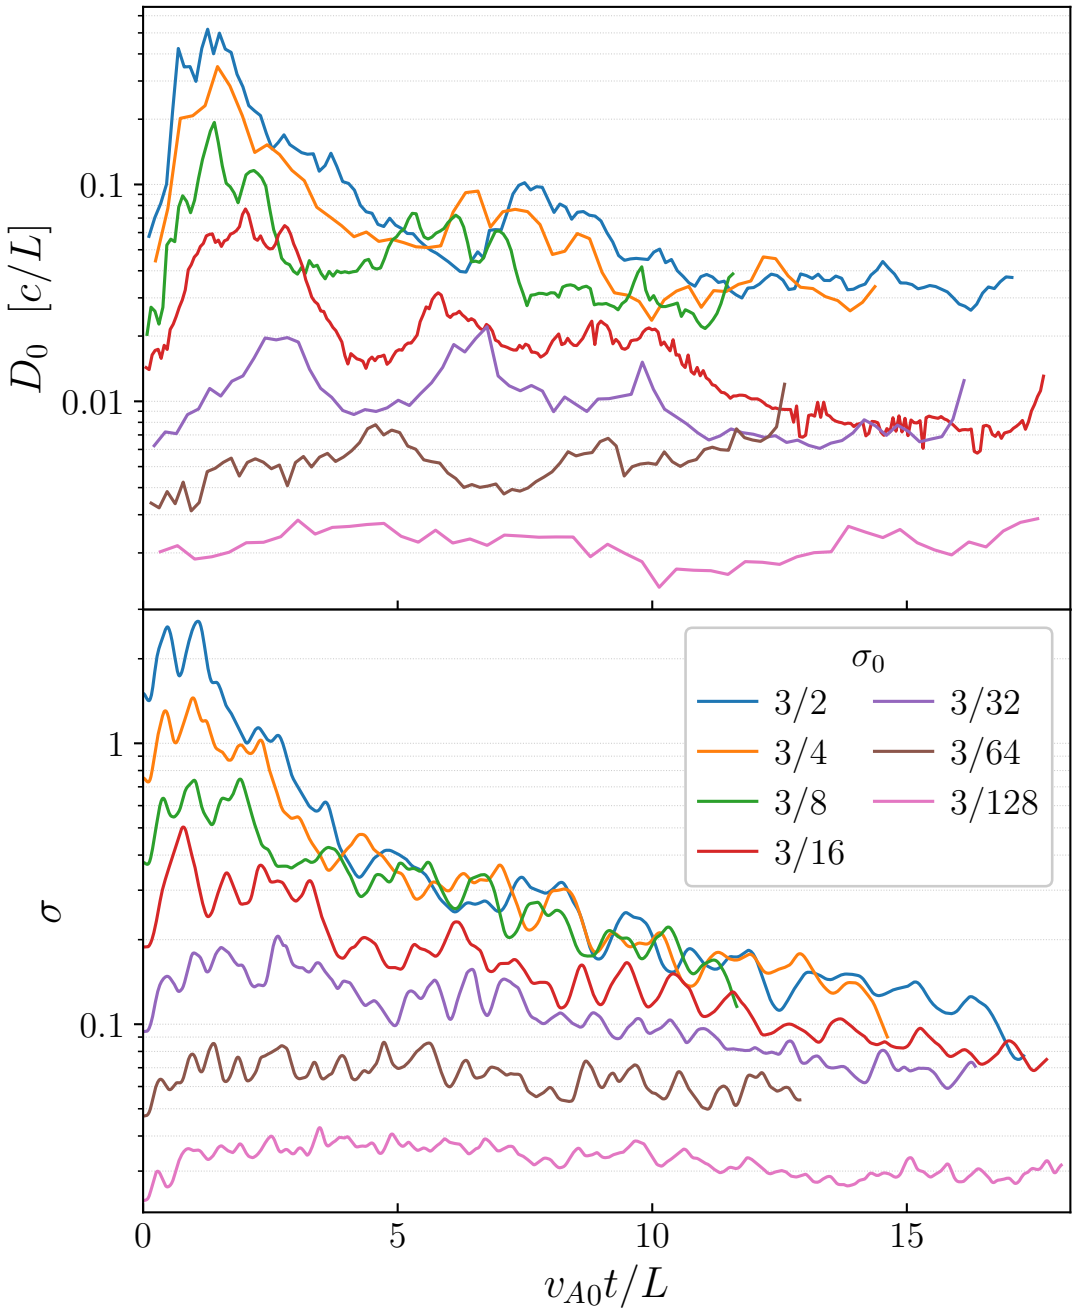

Figure 3.15: (a) The diffusion coefficient γ^2 prefactor $D_0(t)$, and (b) the instantaneous magnetisation $\sigma(t)$, for the N=768 σ_0-scan. The $D_0(t)$ and $\sigma(t)$ curves are qualitatively similar. Both show an initial increase followed by decay, with higher σ also corresponding to higher D_0.

developed. I also exclude points after the normalised system size $L/2\pi\rho_e(t)$ drops below 15, the size needed to start to see the inertial range and obtain converged particle distributions in the previous radiative turbulence study Zhdankin et al. (2020) (while that study states that $L/2\pi\rho_e$ should be greater than 25, slightly lower values are acceptable for practical purposes).

Here, I take advantage of the time dependence resulting from the simulation setup. As the time evolution of $\sigma(t)$ from simulations with different σ_0 overlap, I ask whether D_0 also coincides at these times. Indeed, I find that D_0 mainly depends on instantaneous $\sigma(t)$ rather than σ_0. This is primarily supported by the points at $\sigma \gtrsim 0.1$ from simulations with $\sigma_0 \geq 3/16$, but some overlap is also seen for points with $\sigma_0 \geq 3/64$. For this range of simulations, the points with coinciding $\sigma(t)$ overlap even though they come from simulations with different σ_0. At the lowest σ, however, there is no overlap between different simulations. This is because the $\sigma(t)$ time evolution is much slower for lower σ_0. Future work may benefit from more tightly spaced simulations in the low-σ_0 range.

I find that $D_0 = 0.4(v_A/c)^3\, c/L = 0.4[\sigma/(1+\sigma)]^{3/2}\, c/L$ (solid line in Figure 3.16) is a reasonably good fit over the entire σ range. I also present a power law fit with index $3/2$, which also fits well. The standard Fermi acceleration theory predicts only $D_0 \propto \sigma$. However, these numerical results provide a clear evidence against a linear σ scaling for $\sigma \lesssim 0.5$, although it may be applicable for $\sigma \gtrsim 0.5$. The preferred scaling $D_0 \propto v_A^3$, which becomes $D_0 \propto \sigma^{3/2}$ in the small-σ limit, was predicted by Demidem et al. (2020) for Alfvénic turbulence. There, the extra power of v_A compared to the standard Fermi acceleration theory was attributed to resonance broadening. Thus, these results provide direct numerical support for this theory.

3.7.2 System-size dependence

Figure 3.17 shows the diffusion coefficient at a single point of time $t_0 = 10L/c$ for the system-size scan simulations, which have the same $\sigma_0 = 3/8$ and varying $L/\rho_{e0} \in \{1024, 683, 512, 341, 256, 171\}$. One could also choose time points with a logarithmic dependence on system size, in accordance with the convergence time in Zhdankin et al. (2018b), but the difference is minor. The agreement in the nonthermal range between the simulations of different L/ρ_{e0} is excellent. At this

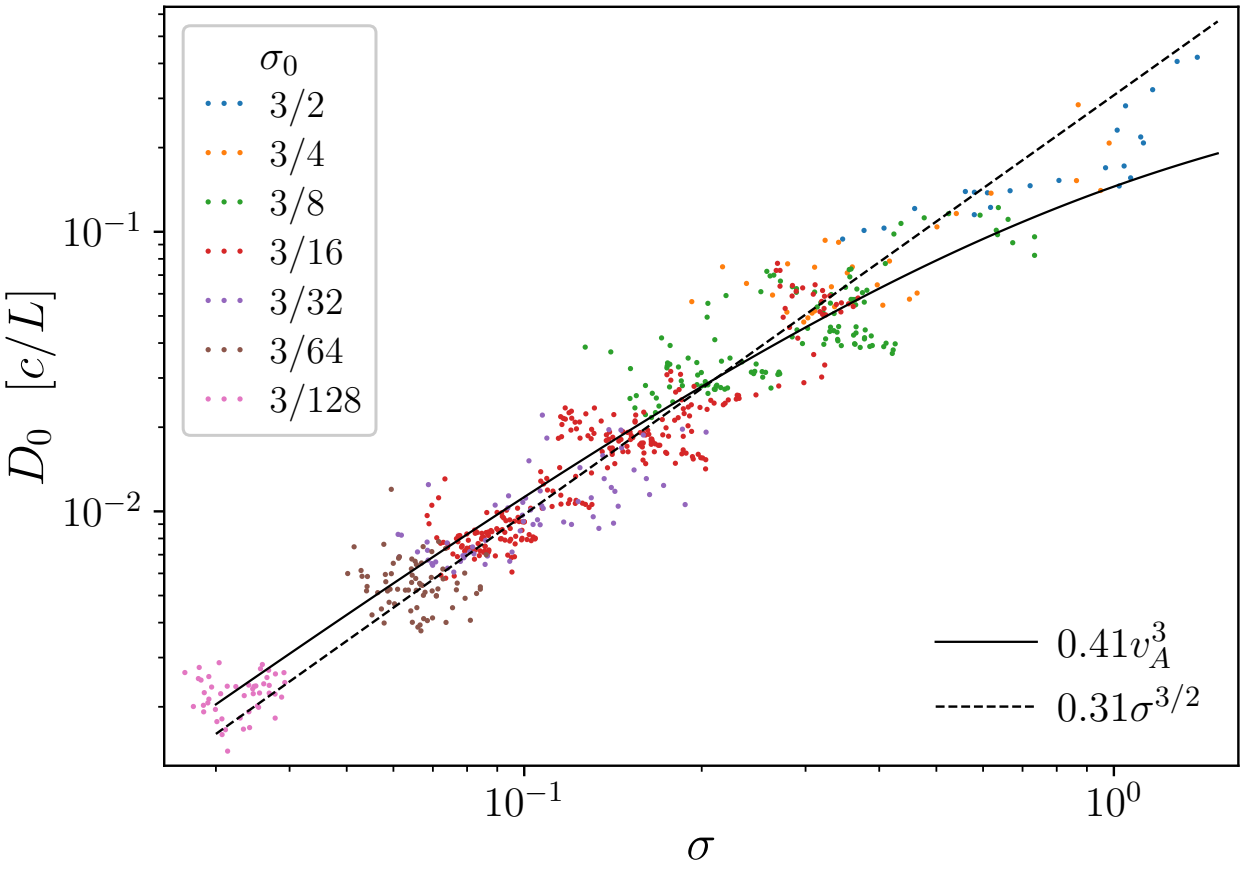

Figure 3.16: D_0 vs instantaneous $\sigma(t)$, combining the results from several simulations with $N = 768$.

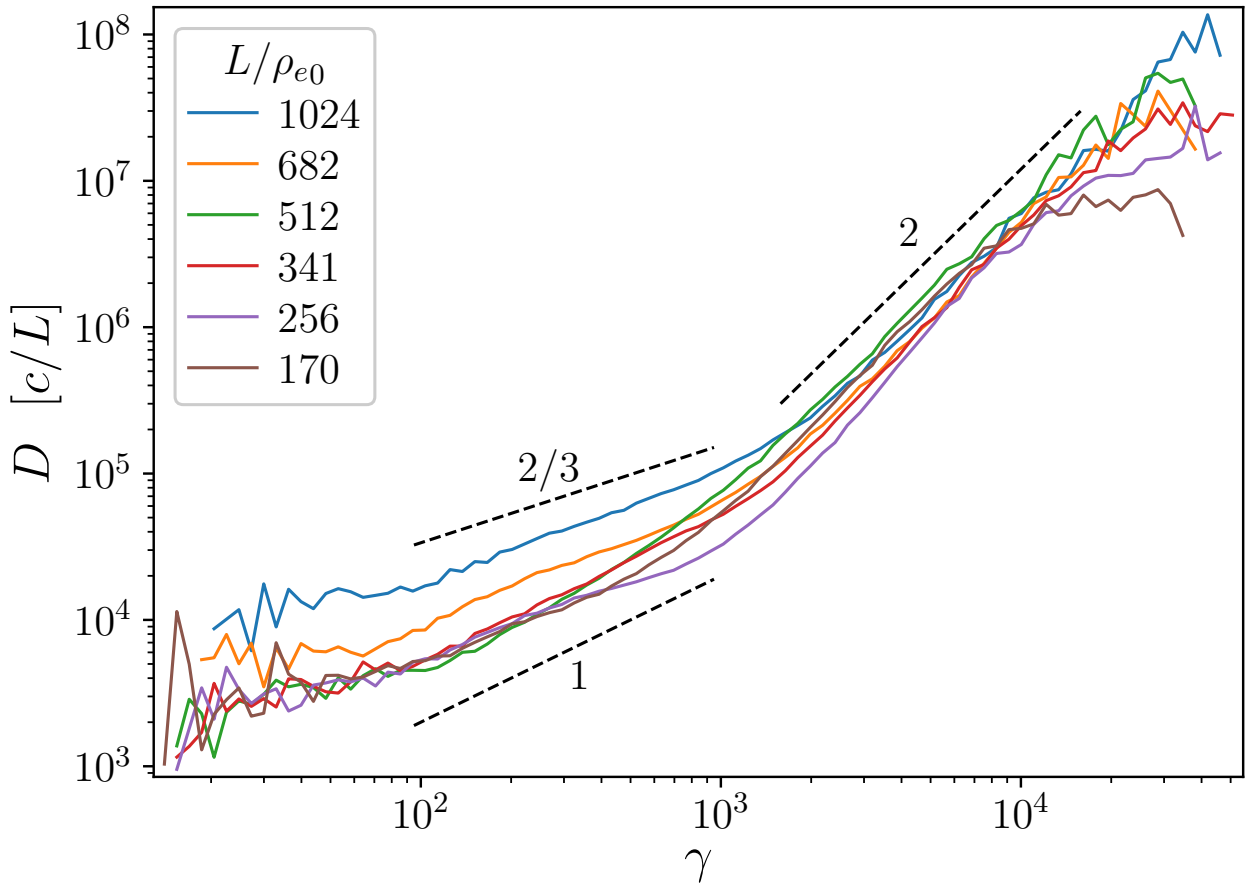

Figure 3.17: $D(\gamma)$ at $t_0 = 10L/c$ for each simulation in the system-size scan, where $\sigma_0 = 3/8$.

$\sigma_0 = 3/8$, $D(\gamma)$ follows a broken power law, where the low-energy part is shallower, and varies somewhat between simulations.

Figure 3.18 shows $D_0(t)$ and $\sigma(t)$ for the system-size scan. The upper panel shows that for simulations with the same σ_0 but different system sizes L/ρ_{e0}, the time-histories of $D_0(t)$ are largely overlapping, consistent with being statistically the same. The lower panel shows that system size has only a weak effect on $\sigma(t)$ when σ_0 is held constant. Repeated runs of simulations with the same initial system parameters but different random seed would be required to test if there is any statistically significant difference in $D_0(t)$ for simulations between which only the system size varies. I leave such an investigation for future work. The convergence of D in the nonthermal range with increasing system size L/ρ_{e0} is good, which gives confidence in the rest of the results.

3.7.3 Particles-per-cell convergence

Figure 3.19 shows the dependence of $D(\gamma)$ at fixed $t_0 = 10L/c$ on the number of particles per cell $N_{\rm ppc}$ for simulations with $\sigma_0 = 3/8$ and $L/\rho_{e0} = 256$ ($N = 384$). The convergence in the high-energy, $D \propto \gamma^2$, range is excellent. The extent of this range shortens slightly as $N_{\rm ppc}$ decreases, with the beginning occurring at slightly higher γ. Lower energies also have a power-law scaling, which, however, becomes shallower with decreasing $N_{\rm ppc}$. Importantly, however, even this lower-energy segment shows very good convergence with increased $N_{\rm ppc}$, with essentially full convergence achieved for $N_{\rm ppc} \gtrsim 64$ for the entire range of particle energies. Note that the effect of decreasing $N_{\rm ppc}$ in Figure 3.19 resembles the trend of increasing system size from Figure 3.17; this resemblance may be explained by PIC noise having a stronger effect on small scales (low energy diffusion) in simulations with a larger inertial range.

The high quality of D measurements in the nonthermal range in simulations with both reduced L/ρ_{e0} and $N_{\rm ppc}$ is somewhat surprising. With reduced system size, the extent of the inertial range should be reduced. However, the dynamic range for which there are good D measurements with $D \sim \gamma^2$ scaling is almost unchanged. This suggests that the highest energy for which D can be measured reliably in these simulations is limited by particle statistics due to the $f \propto \gamma^{-\alpha}$ falloff

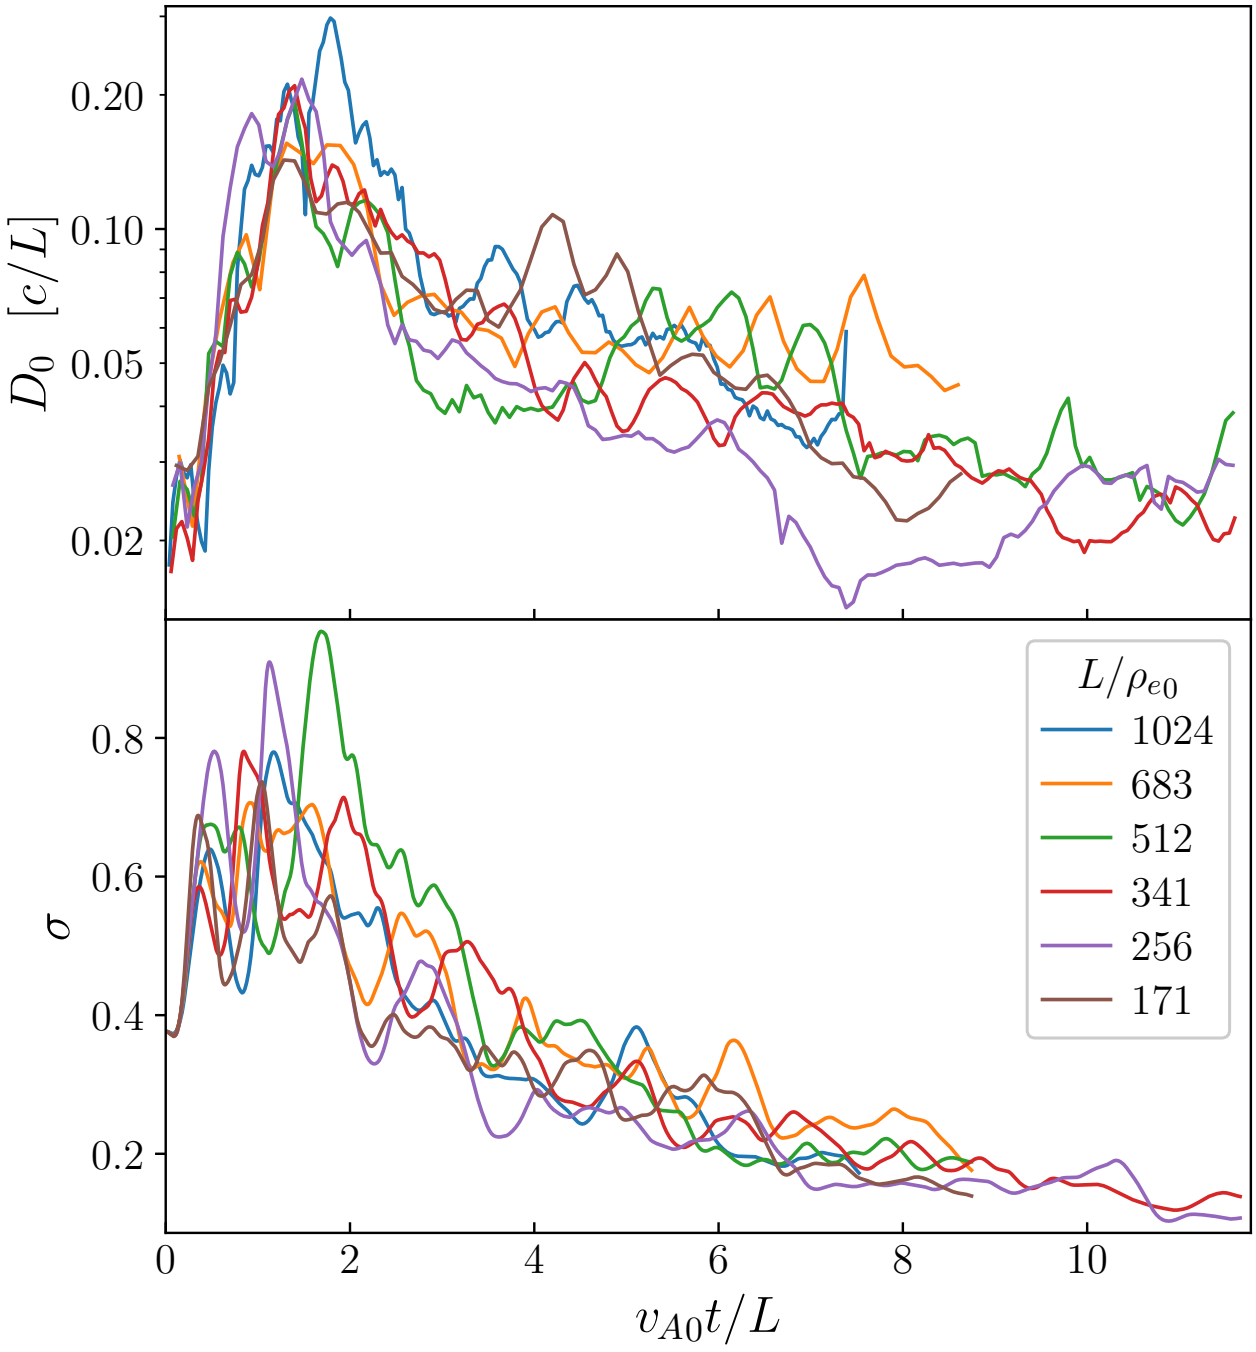

Figure 3.18: (a) $D_0(t)$, and (b) $\sigma(t)$, for the system-size scan, where the simulations have the same $\sigma_0 = 3/8$ but different $L/\rho_{e0} \in \{1024, 683, 512, 341, 256, 171\}$.

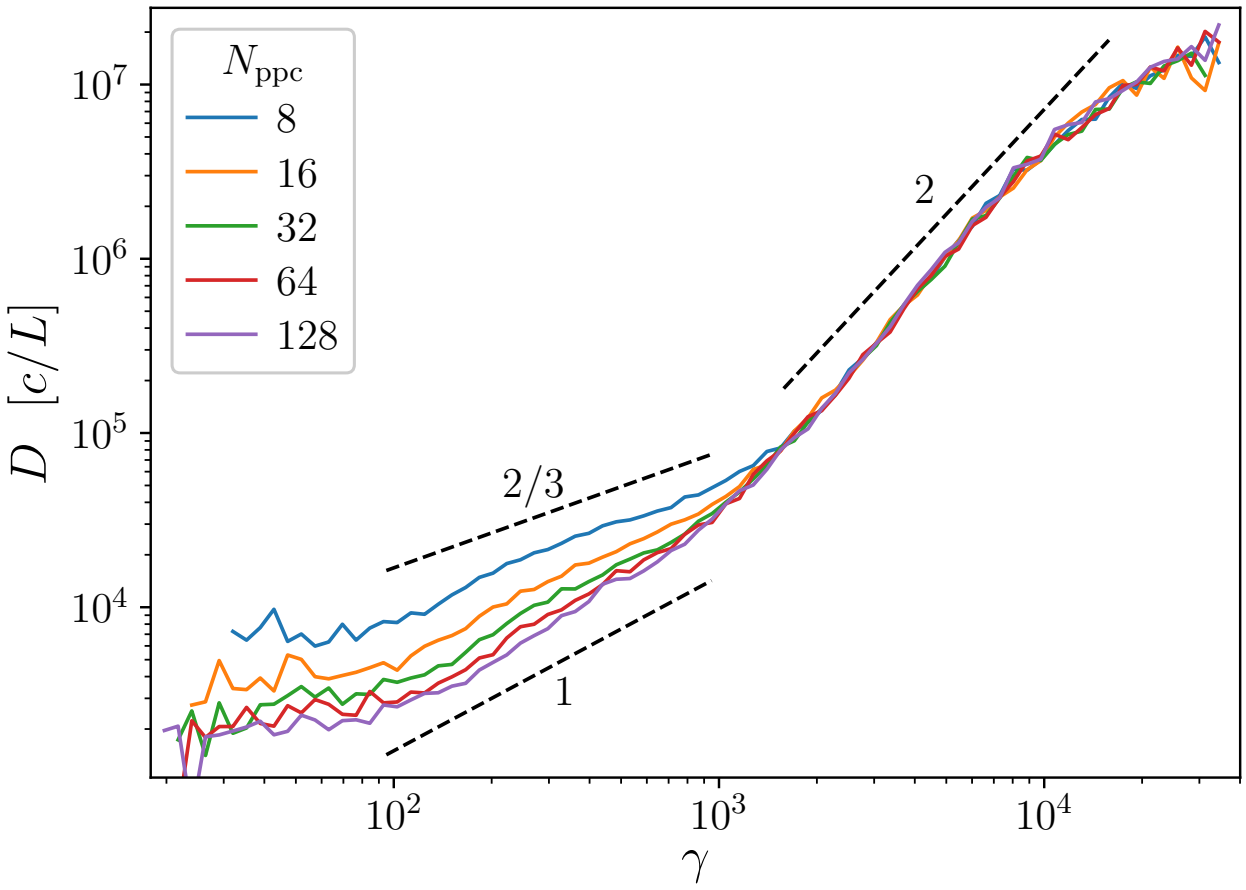

Figure 3.19: $D(\gamma)$ at $t_0 = 10L/c$ for the $N_{\rm ppc}$-scan simulations, which have $L/\rho_{e0} = 256$ and $\sigma_0 = 3/8$. In the nonthermal region, there is excellent convergence. At lower energies, lower-$N_{\rm ppc}$ simulations have higher D.

rather than by the extent of the inertial range or by $\gamma_{\rm max}$.

3.8 Advection coefficient

In Wong et al. (2020), I observed that A was positive at $\gamma < \gamma_{\rm pk}$ and negative at $\gamma > \gamma_{\rm pk}$. However, the data at highest nonthermal energies seemed too noisy to analyse in detail. Now, my model in Subsection 3.2.2 suggests a concrete energy-dependence $A = A_0 \gamma \log\left(\gamma/\gamma_A^*\right)$ (3.9) in the nonthermal range, and also predicts relationships between the FP coefficient model parameters and the power law index. This Section will examine the advection coefficient in the context of that model.

First, I reduce the noise in A by time-averaging it over intervals of $0.5L/v_{A0}$ and also by merging consecutive pairs of bins above $\gamma_0 = 10^4$. Figure 3.20 shows the resulting $A(\gamma, t)$ at times $t_0 \approx \{3, 6, 9\}L/v_{A0}$ for the simulations with $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ from the N=768 σ_0-scan. Although still considerably noisy in the nonthermal range, the qualitative behaviour is clear and is also similar throughout the depicted simulations and time instants: A is positive at highest and lowest energies, and negative in an intermediate interval. The highest-energy behaviour is not visible if the power law is too steep to have enough high-energy tracked particles to measure A there; this occurs for the $\sigma_0 = 3/128$ simulation and early times in the $\sigma_0 = 3/8$ and $\sigma_0 = 3/32$ simulations. However, in these cases the lower-energy data still agrees with the aforementioned pattern. The first zero crossing, from positive to negative near $\gamma_{\rm pk}$, was the initial observation of Wong et al. (2020): that the advection coefficient tends to gather particles together in energy space. This is at thermal energies and so out of the Subsection 3.2.2 model's validity range. The second zero crossing, from negative to positive, is a key feature of the Subsection 3.2.2 model that is borne out in these simulation results. This sign change is unlikely to be an artefact of the particle pileup at the system size limit $\gamma_{\rm max}$ because it occurs roughly an order of magnitude in energy lower than $\gamma_{\rm max}$ (which is depicted as the rightmost grey bar in Figure 3.20). Figure 3.20 also shows fits of A to (3.9), limited to the nonthermal range. As long as there is adequate high-energy data, the fits are consistent with the measurements inside the model validity range of the power-law tail; deviations occur near $\gamma_{\rm pk}$

and $\gamma_{\max}$, which is expected.

While the FP equation is in terms of A, it is more convenient to work with M because $A = M - \partial_\gamma D$ (from Equation 3.23) is contaminated by noise in $\partial_\gamma D$ (with the γ-derivative in particular amplifying the noise), whereas M comes only from $\Delta\overline{\gamma}$ measurements (also by Equation 3.23). Hence, I switch to analysing M from here on. I rewrite (3.14) into

$$M/\gamma = M_0 \log\gamma - M_0 \log\gamma_M^*. \tag{3.25}$$

This now simply states that M/γ is linear in $\log\gamma$. This form is convenient for visual inspection and also favourable for fitting because linear regression is robust and well understood.

Figure 3.21 shows M/γ vs γ on linear-log scales so that a linear equation in $\log\gamma$ such as (3.25) appears as a straight line. I show this for a selection of simulations from the N=768 σ_0-scan at measurement times $t_0 \approx \{3, 6, 9\}L/v_{A0}$. Generally, at the earliest shown instance $t_0 \approx 3L/v_{A0}$, the power-law is in a not quite fully formed intermediate state, whereas at the later measurement times, the power-law is fully established. I also restrict the plot view to the high-energy region above $\gamma_{\rm pk}$. I find, in all cases, a remarkably good linear segment spanning about a decade of energies. This is followed by a decline in M/γ at energies near the system-size limit (or near the end of the particle energy distribution at the time), adulterated by varying amounts of noise: less for high σ_0 and more for low σ_0. Even the lowest $\sigma_0 = 3/128$ simulation plot shows a high-quality straight-line region despite the lack of high-energy particles which caused the high-energy region in the corresponding $A(\gamma, t)$ panel (Figure 3.20d) to be noisy. The linear M/γ region also appears clearly before the power law has fully formed. The high-quality linear segment over a wide range of initial magnetisation values gives strong qualitative support to the Subsection 3.2.2 model.

I repeatedly fit M with functions of the form (3.25), obtaining M_0 and γ_M^*. Figure 3.22 shows $M_0(t)$ and $\gamma_M^*(t)$ for the N=768 σ_0-scan. Since γ_M^* is noisy, I apply a geometric Gaussian smoothing with standard deviation of $0.5L/v_{A0}$. I find that M_0 is higher for higher σ_0 simulations, and it tends to decrease over time. There is no clear trend in γ_M^*, although the values are around $\gamma = 1000$ which is typically within the lower part of the nonthermal section.

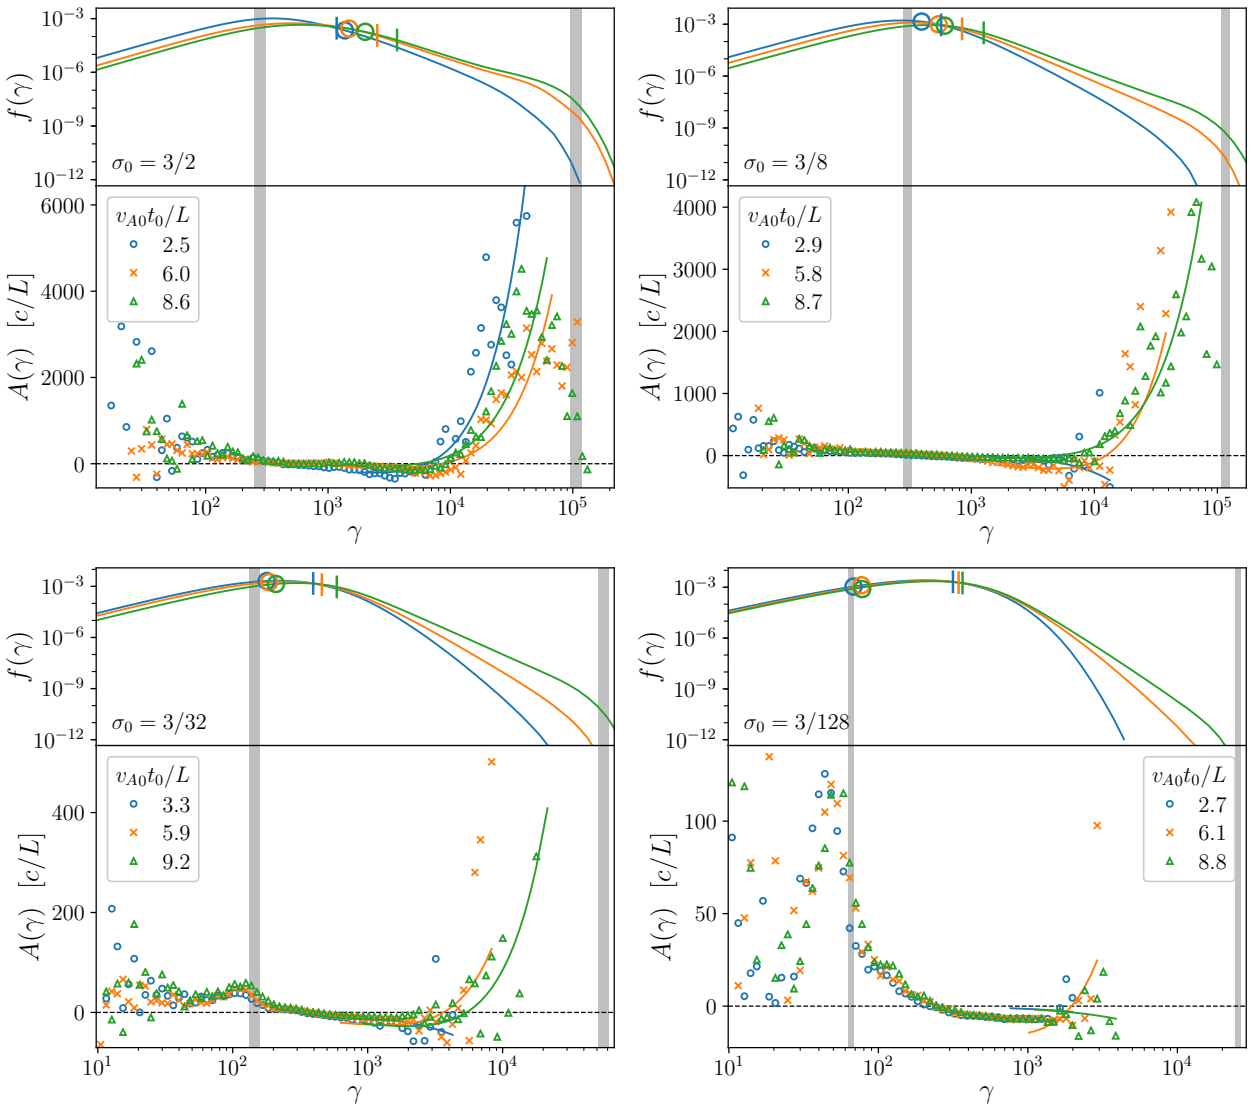

Figure 3.20: A at different points in time for simulations with varying initial magnetisation. Solid lines are fits of $A = A_0\gamma\log\left(\gamma/\gamma_A^*\right)$. Noise at high energies has been reduced by averaging A over $0.5L/v_{A0}$ and combining every two bins after $\gamma = 10^4$. See the Figure 3.14 caption for a description of the length-energy scale markings.

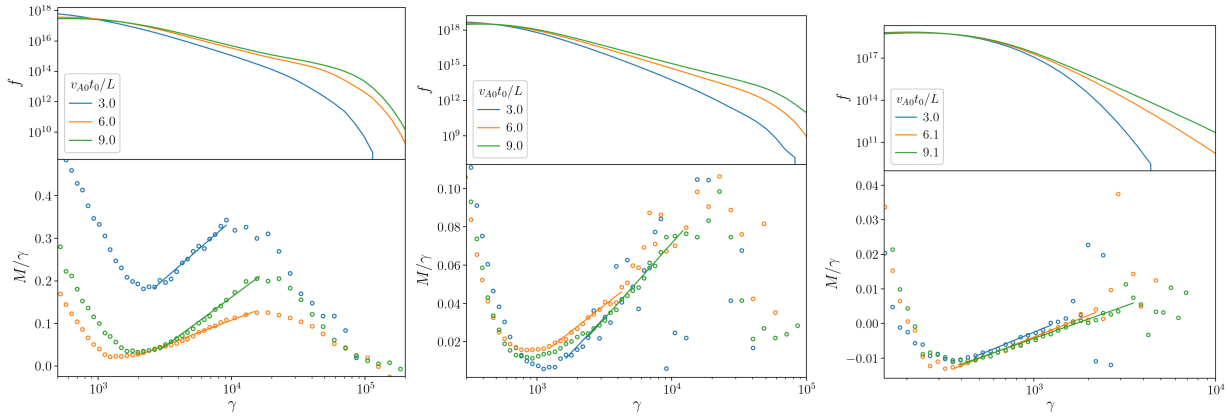

Figure 3.21: M/γ for simulations with $\sigma_0 = 3/2, 3/16, 3/128$, from left to right. M/γ is linear in $\log\gamma$ for a substantial part of the power-law tail. Lower panel solid lines are linear fits in $\log\gamma$.

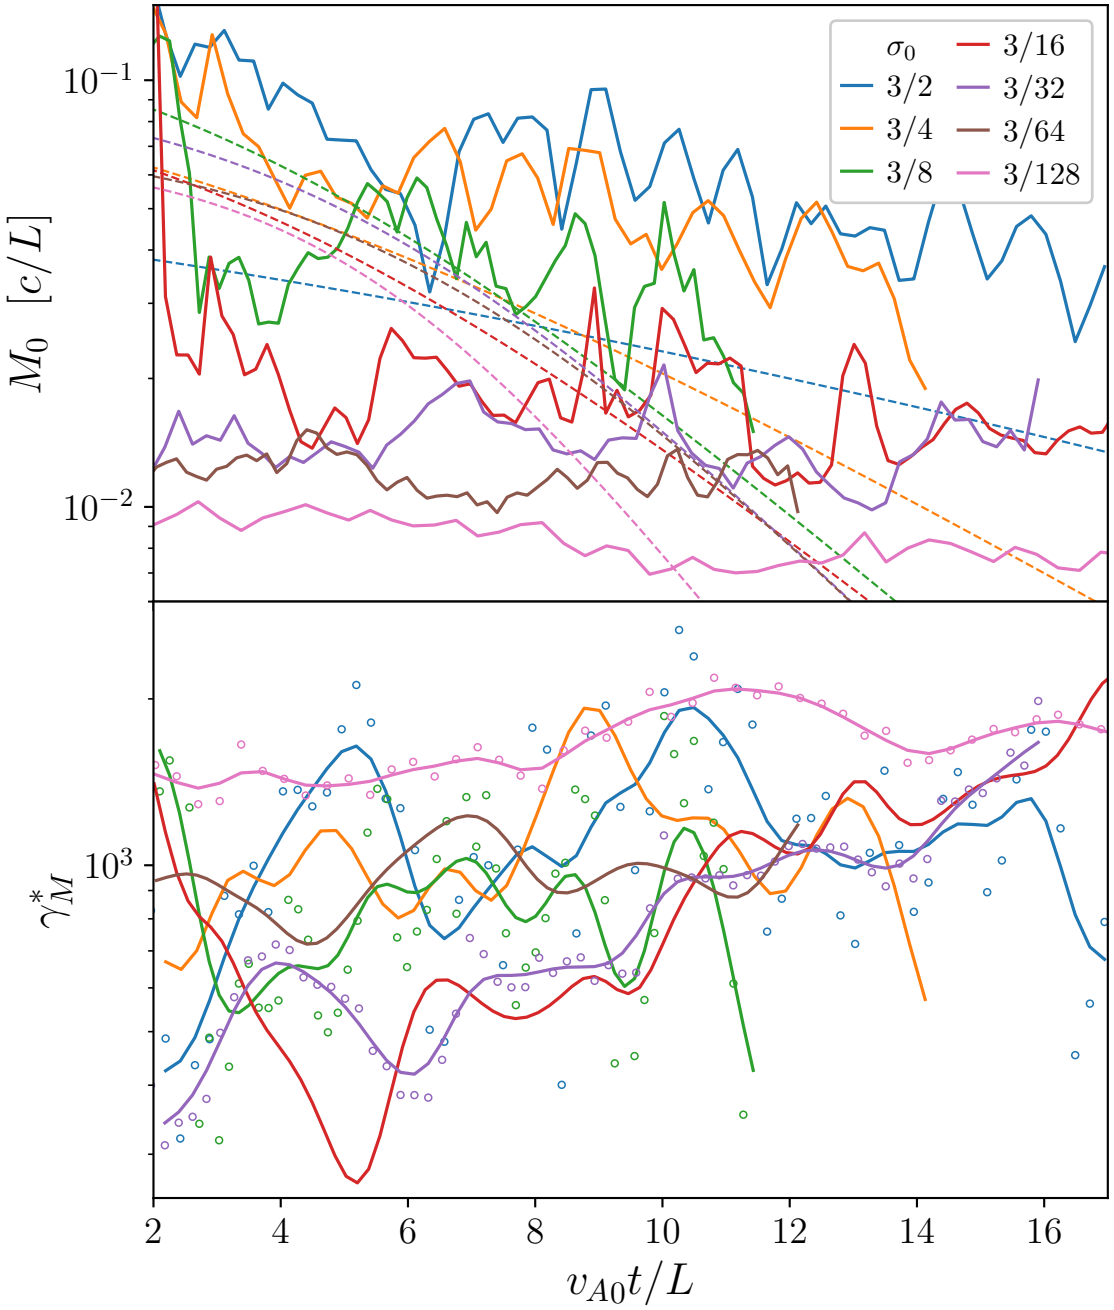

Figure 3.22: (a) Solid lines are M_0 measured from fits to M. Dashed lines are M_0 predictions from α exponential fits through (3.26). M_0 is higher for higher σ_0 simulations, and it tends to decrease over time. (b) Solid lines are smoothed γ^*_M measurements. Open circles are the original γ^*_M measurements. These are only shown for every other σ_0 value, to reduce visual clutter, while giving a sense of the spread.

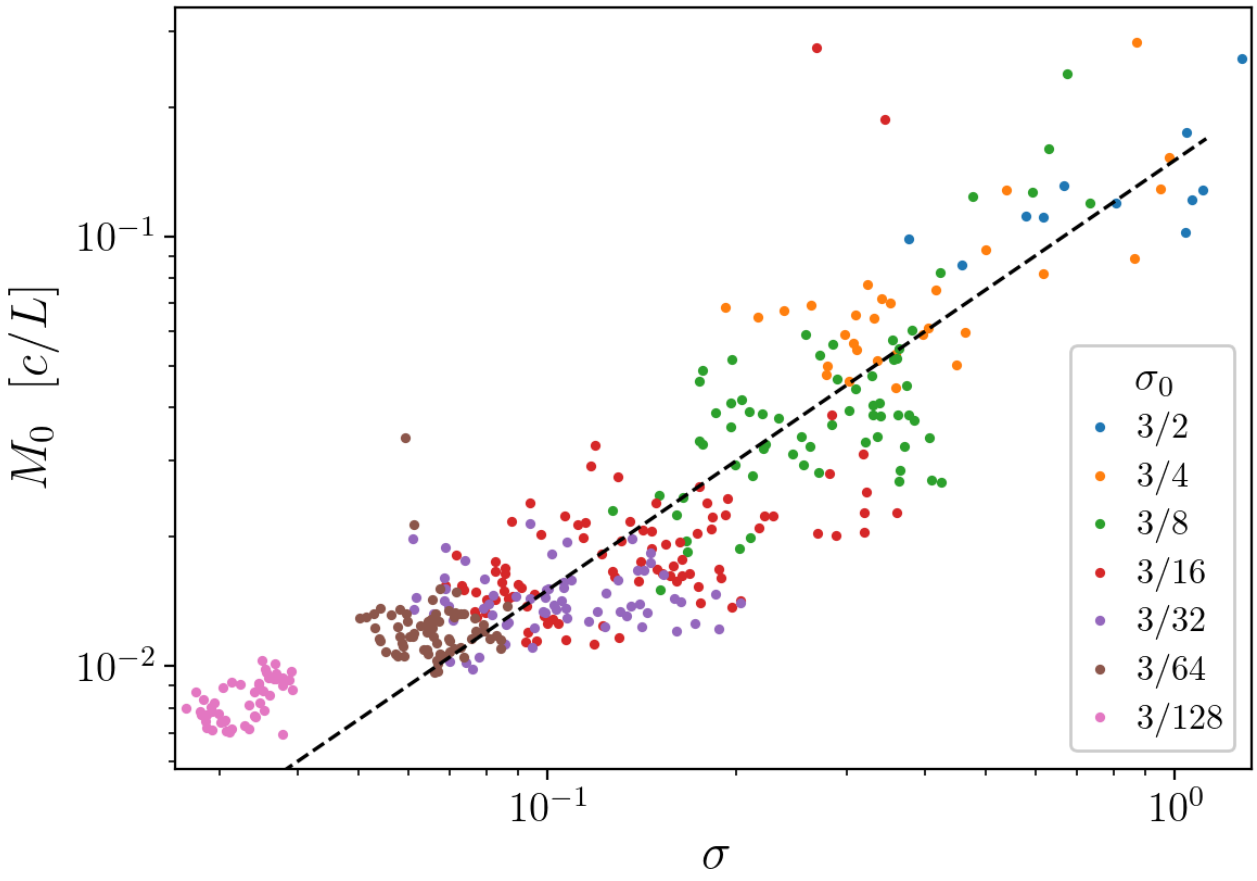

Figure 3.23: M_0 vs instantaneous σ, with a dashed line proportional to σ.

Subsection 3.2.2 does not give any predictions for M_0 and γ_M^* as functions of system parameters in the way that analytical models predict the scaling of (the diffusion coefficient γ^2 prefactor) D_0 with instantaneous σ (see Subsection 3.7.1). Nevertheless, by analogy to the diffusion coefficient, I plot M_0 against σ in Figure 3.23. The scaling is consistent with a power law index of about 1.

3.8.1 Comparison to analytical model of power law evolution

Subsection 3.2.2 derived relationships between the FP coefficient parameters D_0, M_0, and γ_M^* and the power-law parameters K and α. These variables are related by (3.10, 3.15) and their integral forms (3.16, 3.17). In this Subsection, I test the consistency of these formulae with the acquired data.

The FP coefficient parameters come from fits of the FP coefficients to (3.6) and (3.9). Section 3.7 describes the D_0 fit. Section 3.5 described how to measure $\alpha(t)$. The normalisation coefficient K is found by extending the process used for finding α, as follows: The γ-value selected for α obtained from the cubic fit to $-\partial \log f / \partial \log \gamma$ (as described in Section 3.5) also gives a point $[f(\gamma), \gamma]$ which, when substituted into $f(\gamma) = K\gamma^{-\alpha}$, yields $K = f(\gamma)\gamma^{\alpha}$.

I first examine the differential relationships. Figure 3.22a shows the fitted $M_0(t)$ values, along with the prediction from the nonthermal power-law evolution, $-\dot{\alpha}/(\alpha - 1)$, Equation 3.10. For this equation, I use the exponential fits described in Section 3.5 instead of the individual power law index measurements. This is because α evolves slowly at late times, and so small measurement noise in α can cause $\dot{\alpha}$ to fluctuate above and below zero misleadingly, whereas the exponential is an excellent fit and is guaranteed to be smooth and monotonic. The analytical expression for M_0 obtained by substituting the empirical fit $\alpha(t) = \alpha_\infty + \Delta_\alpha \exp(-t/\tau)$ (Equation 3.18) into Equation 3.10 is

$$M_0(t) = \frac{1}{\tau} \frac{1}{1 + \frac{\alpha_0 - 1}{\Delta_\alpha} e^{t/\tau}} . \tag{3.26}$$

The predictions from the α integrals are generally lower than the M_0 values obtained from fitting to M. This is investigated further below.

I note that M_0 and γ_M^* can both fluctuate in time significantly while still producing relatively

smooth $\alpha(t)$ and $K(t)$ because the latter two quantities are obtained through the cumulative effect of the FP coefficients acting on the particle energy distribution. This is, the above differential equation comparisons may be misleading in terms of how much the FP coefficient fluctuations affect the power law index. This idea is embodied by rearranging the differential equations (3.10, 3.11, 3.15) into integral equations for $\alpha(t)$ and $K(t)$: (3.12, 3.13, 3.16, 3.17).

Figure 3.24 compares α and K measured from the particle energy distribution with those obtained through the time integrals (3.12, 3.13) from the FP coefficient model parameters D_0, A_0, and γ_A^*. I choose the integration constants so that the integral lines and α measurement lines intersect where the α "error bar" is smallest. This corresponds roughly to where the power law is flattest, just before the high-energy pileup appears (see Section 3.5), which is essentially the "inflection time" discussed by Zhdankin et al. (2018b). The two agree reasonably well, especially at late times. The integral result for α undershoots the measurement after the constructed intersection point. A possible explanation is if the FP coefficients are partially sensitive to the high-energy pileup. The pileup contributes to the uncertainty in α and would tend to push α to smaller values.

These comparisons between the α and FP measurements stretch the quality of the data with substantial uncertainty on both sides; for α, due to the limited extent of the power law, and for the FP coefficients, due to the limited number of high-energy tracked particles. Nevertheless, this consistency result is significant because the FP coefficients come from tracked particles while the power-law parameters come from the global particle energy distribution, two substantially different data acquisition methods. Hence, their qualitative agreement provides reasonable confidence in the basic validity of the Subsection 3.2.2 model. Furthermore, the remarkable quality of the linear segment in M/γ versus $\log\gamma$ in Figure 3.21, and the fact that the assumptions underlying the Subsection 3.2.2 model are simple, are both encouraging.

3.8.2 Momentum space advection coefficient

As stated in Subsection 3.2.1, the FP model is often formulated in momentum space. For an isotropic model where the differential number of particles is $dn = 4\pi p^2 f(p)$, the diffusion and advec-

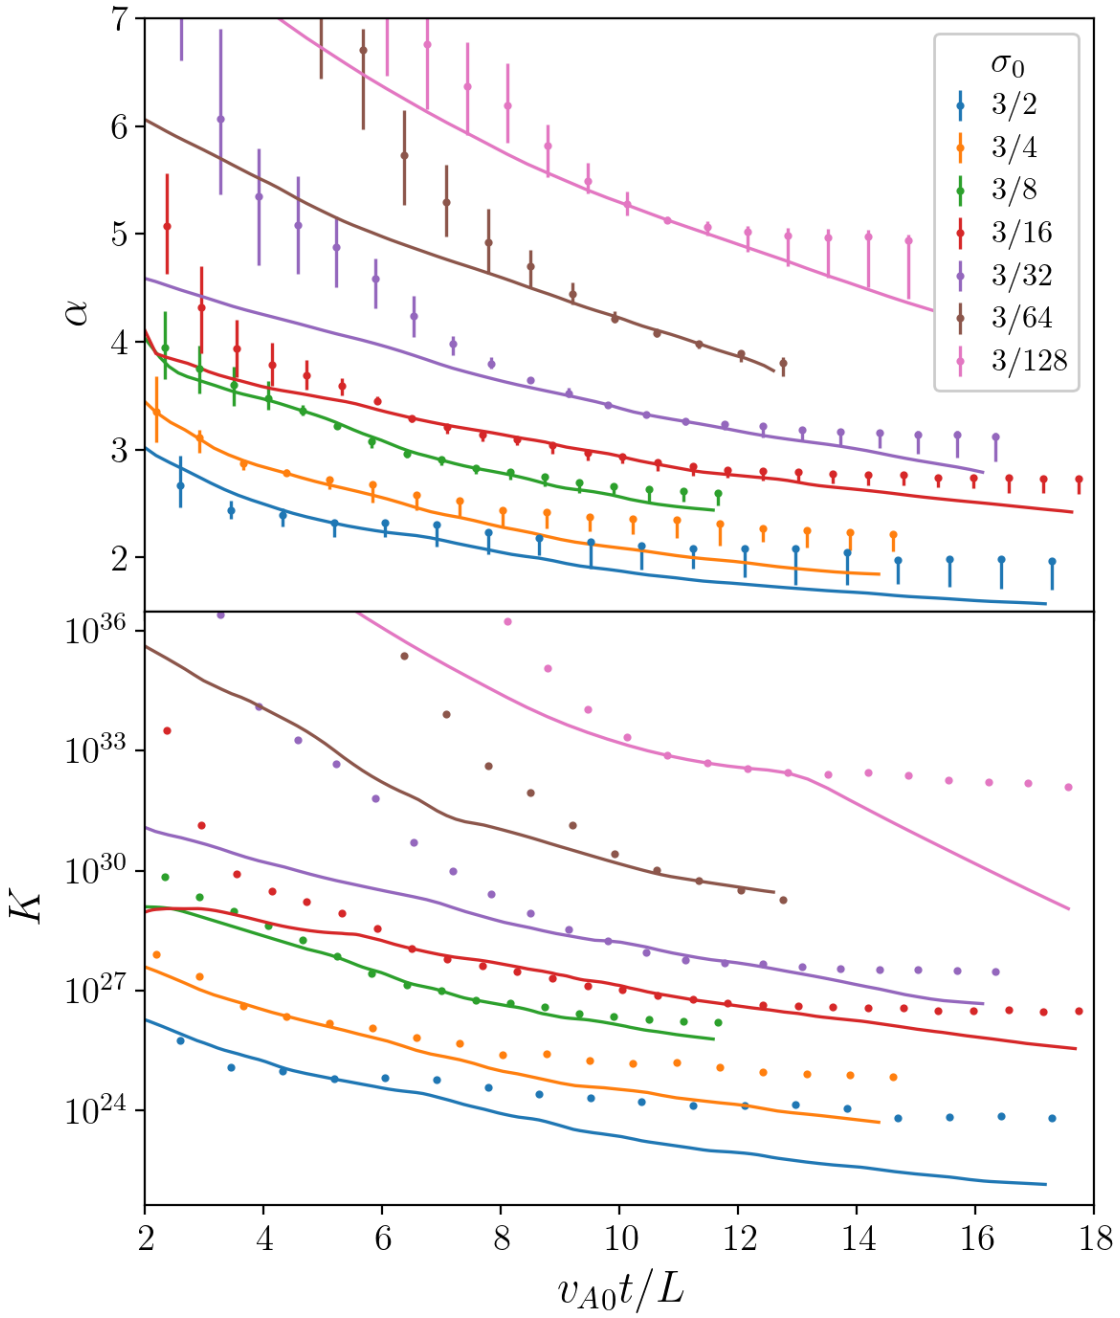

Figure 3.24: α and K measurements vs integral prediction. Points are measurements coming from the particle energy distribution and dashed lines are from the integral formulas as described in the main text. "Error bars" are the same as described in Section 3.5.

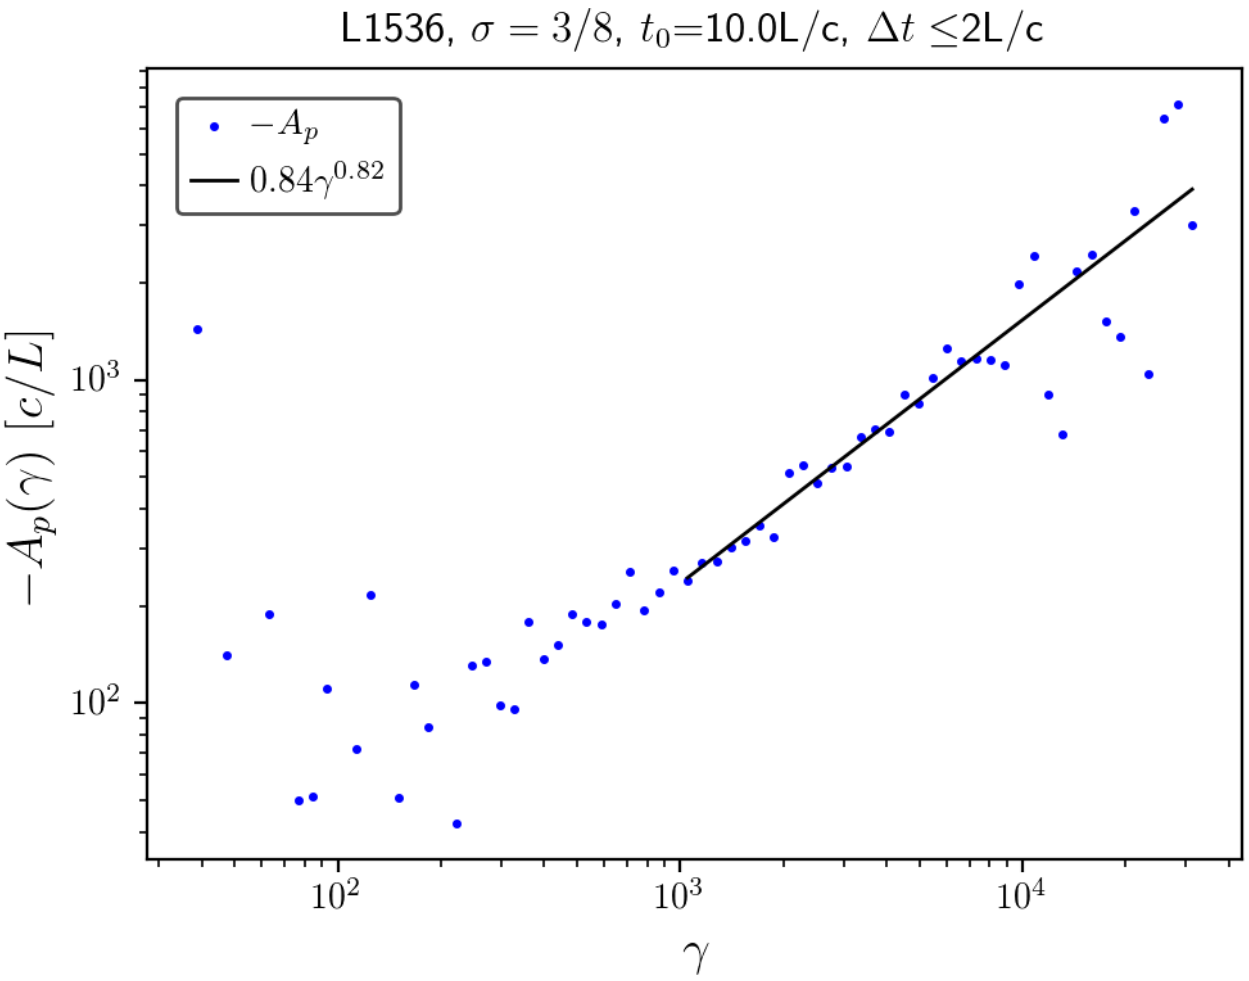

Figure 3.25: Momentum space advection coefficient $A_p = A - 2D/\gamma$.

tion coefficients D_{pp} and A_p correspond to energy-space coefficients (without subscript) $D = D_{pp}$ and $A = A_p + 2D_{pp}/\gamma$. There is a common result or assumption that $A_p = 0$ and consequently $A = 2D/\gamma$ for the case of balanced turbulence (see, e.g. Schlickeiser, 1989, for a quasilinear derivation that retains an advection coefficient).

Figure 3.25 shows the momentum space advection coefficient for the $N = 1536$, $\sigma_0 = 3/8$ simulation at $t_0 = 10L/c$. It is consistent with being everywhere negative and having a power-law scaling with index around 0.8. A single power-law fit is adequate for both the low-energy and nonthermal ranges. While this requires further investigation, the preliminary finding is that the results are statistically inconsistent with the common belief that $A_p = 0$.

3.9 Conclusions

In this Chapter, I investigated the evolution of nonthermal particle populations in first principles particle-in-cell (PIC) simulations of collisionless, magnetised, driven relativistic pair plasma turbulence through the lens of the Fokker-Planck (FP) energy diffusion-advection model. Measurements of the particle energy distribution nonthermal power-law index find behaviour consistent with an exponentially decaying time evolution. I then examine the behaviour of particle energies by tracking large numbers of particles. Particle energies behave diffusively with respect to time across broad system parameter ranges. I measure the energy diffusion and advection coefficients D and A as a function of particle energy and time using the tracked particle data. When fed back in to the FP equation, these coefficients successfully reproduce the evolution of the overall particle energy distribution obtained from the PIC simulations. This suggests that the FP model is an adequate description of turbulent particle acceleration for the entire investigated system parameter space. I find that the diffusion coefficient generally scales with energy squared in the nonthermal range. The coefficient of this scaling, D_0, depends primarily on the instantaneous magnetisation σ rather than the initial magnetisation, and this relationship is consistent with $D_0 \propto v_A^3$. The interpretation of the advection coefficient measurements is guided by an analytical model relating the nonthermal power-law index to the FP coefficients in the nonthermal range. This predicts a scaling

of $A = A_0 \gamma \log\left(\gamma/\gamma_A^*\right)$, which is borne out by the measurements. In addition, quantitative relationships between parameters fitted to the FP coefficient variables and the index and normalisation of the power law tail are fulfilled reasonably well.

This work opens various avenues for further investigation. The diffusion coefficient scaling is somewhat different from the expected scaling of $D_0 \propto \sigma$. This could be further tested through simulation and experiment, as well as investigated analytically. The novel scaling of $A = A_0 \gamma \log\left(\gamma/\gamma_A^*\right)$ has significant effects on the evolution of the nonthermal energy spectrum. This may be of interest as nonthermal particle acceleration models commonly assume that the momentum-space advection coefficient is zero (leading to $A = 2D/\gamma$). While I have focused on the behaviour of the FP coefficients in the nonthermal region, I also measured D and A at thermal and subthermal energies, finding qualitatively different behaviour. These may influence nonthermal particle acceleration (NTPA) by providing intermediate energy particles for stochastic acceleration, but regardless, further investigation may shed light on other physical phenomena relevant at those energies and scales. Apart from turbulence, other mechanisms such as magnetic reconnection and shocks also produce NTPA. My methods may be adapted to simulations of these different types of systems in order to test theories of particle acceleration in those respective environments. Future work may investigate the asymptotic particle distribution function by analysing simulations that include cooling, in which a true statistical steady state is achieved (Zhdankin et al., 2020, 2021), which is an ideal starting point for studying the steady-state particle distribution in relation to the FP model.

This study of magnetised turbulent nonthermal particle acceleration over a wide range of system parameters enhances our understanding of widespread and long-studied plasma physical processes, and has implications for space, solar, and astrophysical systems.

Chapter 4

Summary and outlook

This dissertation investigated energy diffusion-advection models of relativistic nonthermal particle acceleration (NTPA) in particle-in-cell (PIC) simulations of relativistic collisionless magnetised pair plasma turbulence by statistically analysing large numbers of self-consistent particle trajectories and energy histories. I analysed simulations over a wide parameter space, with the main system parameters being the initial magnetisation, which controls how relativistic the turbulent bulk motions are, and the system size, which sets the separation between the driving and kinetic scales and hence the extent of the inertial range. The outcomes of this project divide into two main areas: physical findings regarding the relativistic NTPA process, and the development and demonstration of new tools for analysing NTPA in PIC simulations.

The major physical findings are as follows:

- In these turbulence simulations, relativistic particle energies evolve in a manner consistent with an energy diffusion-advection model where the diffusion is of the classical Brownian type. The measured energy diffusion and advection coefficients, $D(\gamma, t)$ and $A(\gamma, t)$ — functions of particle Lorentz factor γ and time t — reproduce the evolution of the overall particle energy distribution when evolved numerically in the Fokker-Planck (FP) equation.

- In the nonthermal range, the energy diffusion coefficient is proportional to particle energy squared: $D = D_0\gamma^2$ (defining D_0 as the proportionality factor), in agreement with second-order Fermi acceleration. This is true for the whole explored simulation parameter space which covers a wide range of different initial magnetisation and system size values.

- The prefactor D_0 is roughly proportional to $\sigma^{3/2}$ (where σ is the instantaneous magnetisation), independent of the initial magnetisation σ_0. This is slightly but noticeably steeper than most Fermi-type models which predict a linear scaling with magnetisation, but is consistent with the recent theoretical prediction of Demidem et al. (2020), which attributes the steeper scaling to resonance broadening of Alfvénic modes in critically balanced magnetohydrodynamic (MHD) turbulence.

- In the nonthermal range, the energy advection coefficient scaling is $A \propto \gamma \log(\gamma/\gamma_A^*)$, with γ_A^* a constant, or equivalently, A/γ is linear in $\log \gamma$ (with intercept). This scaling is entirely novel and was initially suggested by my analytical model connecting the nonthermal power-law index of the particle energy distribution with the FP coefficients in the inertial range.

- The power-law index of the relativistic particle spectrum evolves close to exponentially in time over a wide range of initial magnetisation and system size values, and the asymptotic value tends to be lower (i.e., harder) for higher initial magnetisation. This implies that relativistic turbulence is a particularly effective particle accelerator, especially in producing hard power-law spectra of high-energy particles.

This dissertation uses the relativistic pair plasma regime as it is the simplest, most basic model that still has the critical, key physics for relativistic NTPA. The long-term impact of this study is not, however, limited to just relativistic pair plasma astrophysical systems as these foundational methods can now be extended to other plasma-physical processes such as magnetic reconnection, shocks, and plasma instabilities, as well as different plasma regimes including:

- Radiative turbulence. In particular, to investigate whether the diffusion coefficient retains its energy-squared scaling, and how the radiative losses (e.g., synchrotron and/or inverse-Compton) affect the FP coefficients, for instance, in the advection coefficient. This is particularly intriguing because the radiative turbulence studies of Zhdankin et al. (2020) indicate that the particle energy distribution may lose a power-law shape and become quasi-thermal in the strong-cooling regime. This is therefore an attractive opportunity to take advantage

of particle tracking as it goes beyond using only the overall particle energy distribution. This would allow it to reveal whether a second-order Fermi acceleration is operating even if the formation of a nonthermal power law of particle energies is being quenched by vigorous, energy-dependent cooling.

- Electron-ion plasma, which has applications in the relativistic or semi-relativistic regime to black-hole radiatively-inefficient accretion flows and jets; and in the nonrelativistic regime, to, e.g., the solar wind. There are also possible basic plasma-physics motivations such as understanding how electron and ion NTPA interact in the absence of binary collisions, via turbulent electromagnetic fields. The electron-ion case brings in additional complexity, both computational — having to deal with an additional separation of electron and ion microscales — and for the analysis: having to investigate two sets (e and i) of particle energy distribution power-law indices and two sets of FP coefficients.

Apart from different physical setups and regimes, the tracked particle analysis methods can be extended to improve the capability to investigate NTPA. Some possible such steps are:

- Inclusion of particle pitch angle in addition to the current 1-dimensional model in energy only. This is important because pitch-angle scattering is a vital part of FP-type NTPA theories, and studies such as Schlickeiser (1989) compute tensor diffusion coefficients including pitch angle coordinates. Thus, measuring the pitch-angle scattering through quantities like the momentum-space (rather than energy-space) diffusion and advection coefficients would give another, complementary way to test NTPA theories.

- Particle trajectories can also be used to study spatial transport, which I only examined preliminarily in Subsection 2.3.2. This can be used to investigate, e.g., spatial diffusion of energetic particles, an important topic in cosmic-ray physics. This is likely to benefit from smoothing or averaging the spatial gyration, just as removing energy oscillations was necessary for the energy-diffusion-advection analysis.

These results and ones generated through subsequent studies such as those described above will help develop more realistic, physics-based FP models of NTPA in relativistic turbulent plasma environments like PWN, AGN jets and radio-lobes, and black-hole accretion flows, including, in particular, the Event Horizon Telescope targets: the supermassive black holes at the centre of our galaxy (Sgr A*) and M87. In particular, the incorporation of first-principles PIC-based FP-model prescriptions into the existing and future 3D global relativistic MHD (including general-relativistic MHD [GRMHD] in the case of black holes) numerical models of these enigmatic astrophysical systems will enable us to more rigorously treat nonthermal populations of both electrons and ions, putting on a firmer footing the connection between those sophisticated dynamical models with the actual observations. Specifically, calculating relativistic electron NTPA using these methods in GRMHD simulations of black holes and their coronae and jets will allow us to predict (and then compare with observations) the observable broad-band (from radio to gamma-rays) electromagnetic radiation signatures from these systems. Calculating ion NTPA is particularly exciting in light of the ongoing multi-messenger revolution in astrophysics, because it will help us confidently assess the potential of various astrophysical systems as possible cosmic-ray accelerators and as sources of very-high-energy neutrinos (e.g., detected by the IceCube neutrino observatory at the South Pole) produced by accelerated ultra-relativistic protons via hadronic processes.

In summary, in this dissertation I showed the presence of FP-type energy diffusion leading to NTPA in PIC simulations of relativistic collisionless turbulence, and measured the FP energy diffusion and advection coefficients using self-consistent tracked particles, both for the first time. In doing so, I also demonstrated the viability of an emerging suite of simulation diagnostics exploiting the wealth of microphysical information generated by PIC simulations that heretofore has been underutilised. Inevitably, these types of tracked particle techniques will join the standard analytical palette, which will help obtain plasma- and astrophysical insights, rigorously test analytical NTPA models, and thus advance our understanding of high-energy astrophysical phenomena.

Bibliography

E. Amato and J. Arons. Heating and Nonthermal Particle Acceleration in Relativistic, Transverse Magnetosonic Shock Waves in Proton-Electron-Positron Plasmas. Astrophys. J., 653(1):325–338, Dec 2006.

Kaspar Arzner, Bernard Knaepen, Daniele Carati, Nicolas Denewet, and Loukas Vlahos. The effect of coherent structures on stochastic acceleration in MHD turbulence. Astrophys. J., 637(1):322–332, jan 2006.

Katsuaki Asano, Fumio Takahara, Masaaki Kusunose, Kenji Toma, and Jun Kakuwa. Time-dependent Models for Blazar Emission with the Second-order Fermi Acceleration. The Astrophysical Journal, 780:64, January 2014. ISSN 0004-637X.

Peter A. Becker, Truong Le, and Charles D. Dermer. Time-dependent Stochastic Particle Acceleration in Astrophysical Plasmas: Exact Solutions Including Momentum-dependent Escape. The Astrophysical Journal, 647:539–551, August 2006. ISSN 0004-637X.

Mitchell C. Begelman, Roger D. Blandford, and Martin J. Rees. Theory of extragalactic radio sources. Rev. Mod. Phys., 56:255–351, Apr 1984.

Nicolas H. Bian, A. Gordon Emslie, Duncan J. Stackhouse, and Eduard P. Kontar. The Formation of Kappa-distribution Accelerated Electron Populations in Solar Flares. The Astrophysical Journal, 796:142, December 2014. ISSN 0004-637X.

C. K. Birdsall and A. B. Langdon. Plasma Physics via Computer Simulation. January 1991.

Roger Blandford and David Eichler. Particle acceleration at astrophysical shocks: A theory of cosmic ray origin. Physics Reports, 154(1):1–75, 1987.

Pasquale Blasi. Stochastic Acceleration and Nonthermal Radiation in Clusters of Galaxies. The Astrophysical Journal Letters, 532:L9–L12, March 2000. ISSN 0004-637X.

Stanislav Boldyrev. Spectrum of Magnetohydrodynamic Turbulence. Physical Review Letters, 96: 115002, March 2006. ISSN 0031-9007.

Virginia Bresci, Martin Lemoine, Laurent Gremillet, Luca Comisso, Lorenzo Sironi, and Camilia Demidem. Nonresonant particle acceleration in strong turbulence: Comparison to kinetic and mhd simulations. Physical Review D, 106(2):023028, 2022.

N. Bucciantini, J. Arons, and E. Amato. Modelling spectral evolution of pulsar wind nebulae inside supernova remnants. Mon. Not. R. Astron Soc., 410(1):381–398, Jan 2011.

R Bühler and R Blandford. The surprising crab pulsar and its nebula: a review. Rep. Prog. Phys., 77(6):066901, 2014.

B. Cerutti, G. R. Werner, D. A. Uzdensky, and M. C. Begelman. Simulations of Particle Acceleration beyond the Classical Synchrotron Burnoff Limit in Magnetic Reconnection: An Explanation of the Crab Flares. The Astrophysical Journal, 770:147, June 2013. ISSN 0004-637X.

Benjamin DG Chandran. Scattering of energetic particles by anisotropic magnetohydrodynamic turbulence with a goldreich-sridhar power spectrum. Physical Review Letters, 85(22):4656, 2000.

Jungyeon Cho and A Lazarian. Particle acceleration by magnetohydrodynamic turbulence. The Astrophysical Journal, 638(2):811, 2006.

Luca Comisso and Lorenzo Sironi. Particle acceleration in relativistic plasma turbulence. Physical review letters, 121(25):255101, 2018.

Luca Comisso and Lorenzo Sironi. The interplay of magnetically dominated turbulence and magnetic reconnection in producing nonthermal particles. The Astrophysical Journal, 886(2):122, 2019.

Camilia Demidem, Martin Lemoine, and Fabien Casse. Particle acceleration in relativistic turbulence: A theoretical appraisal. Physical Review D, 102:023003, July 2020. ISSN 1550-7998.

Charles D. Dermer and Mayer Humi. Adiabatic Losses and Stochastic Particle Acceleration in Gamma-Ray Burst Blast Waves. Astrophys. J., 556(1):479–493, Jul 2001.

Charles D. Dermer, James A. Miller, and Hui Li. Stochastic Particle Acceleration near Accreting Black Holes. Astrophys. J., 456:106, Jan 1996.

P Dmitruk, WH Matthaeus, N Seenu, and Michael R Brown. Test particle acceleration in three-dimensional magnetohydrodynamic turbulence. the astrophysical journal, 597(1):L81, 2003.

Pablo Dmitruk, WH Matthaeus, and N Seenu. Test particle energization by current sheets and nonuniform fields in magnetohydrodynamic turbulence. The Astrophysical Journal, 617(1):667, 2004.

Enrico Fermi. On the origin of the cosmic radiation. Physical review, 75(8):1169, 1949.

Omar French, Fan Guo, Qile Zhang, and Dmitri A. Uzdensky. Particle Injection and Nonthermal Particle Acceleration in Relativistic Magnetic Reconnection*. The Astrophysical Journal, 948(1): 19, April 2023. ISSN 0004-637X.

R. Friedrich, S. Siegert, J. Peinke, St. Lück, M. Siefert, M. Lindemann, J. Raethjen, G. Deuschl, and G. Pfister. Extracting model equations from experimental data. Physics Letters A, 271(3): 217–222, June 2000. ISSN 03759601.

R. Friedrich, Ch. Renner, M. Siefert, and J. Peinke. Comment on "Indispensable Finite Time Corrections for Fokker-Planck Equations from Time Series Data". Physical Review Letters, 89 (14):149401, September 2002. ISSN 0031-9007, 1079-7114.

P. Goldreich and S. Sridhar. Toward a Theory of Interstellar Turbulence. II. Strong Alfvenic Turbulence. The Astrophysical Journal, 438:763, January 1995. ISSN 0004-637X.

Julia Gottschall and Joachim Peinke. On the definition and handling of different drift and diffusion estimates. New Journal of Physics, 10(8):083034, August 2008. ISSN 1367-2630.

Fan Guo, Hui Li, William Daughton, and Yi-Hsin Liu. Formation of Hard Power Laws in the Energetic Particle Spectra Resulting from Relativistic Magnetic Reconnection. Phys. Rev. Lett., 113(15):155005, Oct 2014.

Fan Guo, Xiaocan Li, Hui Li, William Daughton, Bing Zhang, Nicole Lloyd-Ronning, Yi-Hsin Liu, Haocheng Zhang, and Wei Deng. Efficient Production of High-energy Nonthermal Particles during Magnetic Reconnection in a Magnetically Dominated Ion-Electron Plasma. Astrophys. J. Lett., 818(1):L9, Feb 2016.

Amelia M Hankla, Vladimir Zhdankin, Gregory R Werner, Dmitri A Uzdensky, and Mitchell C Begelman. Kinetic simulations of imbalanced turbulence in a relativistic plasma: Net flow and particle acceleration. Monthly Notices of the Royal Astronomical Society, 509(3):3826–3841, 2022.

M. J. Hardcastle, C. C. Cheung, I. J. Feain, and Ł. Stawarz. High-energy particle acceleration and production of ultra-high-energy cosmic rays in the giant lobes of Centaurus A. Mon. Not. R. Astron Soc., 393:1041–1053, March 2009.

R. C. Hartman, D. L. Bertsch, C. E. Fichtel, S. D. Hunter, G. Kanbach, D. A. Kniffen, P. W. Kwok, Y. C. Lin, J. R. Mattox, H. A. Mayer-Hasselwander, P. F. Michelson, C. von Montigny, H. I. Nel, P. L. Nolan, K. Pinkau, H. Rothermel, E. Schneid, M. Sommer, P. Sreekumar, and D. J. Thompson. Detection of high-energy gamma radiation from quasar 3C 279 by the EGRET telescope on the Compton Gamma Ray Observatory. Astrophys. J. Lett., 385:L1–L4, January 1992.

J. Jeff Hester. The Crab Nebula: An Astrophysical Chimera. Annual Review of Astronomy and Astrophysics, 46(1):127–155, September 2008. ISSN 0066-4146, 1545-4282.

M. Hoshino and Y. Lyubarsky. Relativistic Reconnection and Particle Acceleration. Space Sci. Rev., 173:521–533, November 2012.

Masahiro Hoshino, Jonathan Arons, Yves A. Gallant, and A. B. Langdon. Relativistic Magnetosonic Shock Waves in Synchrotron Sources: Shock Structure and Nonthermal Acceleration of Positrons. Astrophys. J., 390:454, May 1992.

Heinz Isliker, Theophilos Pisokas, Loukas Vlahos, and Anastasios Anastasiadis. Particle Acceleration and Fractional Transport in Turbulent Reconnection. The Astrophysical Journal, 849(1):35, October 2017a. ISSN 1538-4357.

Heinz Isliker, Loukas Vlahos, and Dana Constantinescu. Fractional Transport in Strongly Turbulent Plasmas. Physical Review Letters, 119:045101, July 2017b. ISSN 0031-9007.

C. H. Jaroschek, H. Lesch, and R. A. Treumann. Relativistic Kinetic Reconnection as the Possible Source Mechanism for High Variability and Flat Spectra in Extragalactic Radio Sources. Astrophys. J. Lett., 605(1):L9–L12, Apr 2004.

J. R. Jokipii. Fermi Acceleration and the Structure of Interstellar Turbulence. 1:429, January 1977.

Shigeo S. Kimura, Kohta Murase, and Kenji Toma. Neutrino and Cosmic-Ray Emission and Cumulative Background from Radiatively Inefficient Accretion Flows in Low-luminosity Active Galactic Nuclei. The Astrophysical Journal, 806:159, June 2015. ISSN 0004-637X.

Shigeo S Kimura, Kenji Toma, Takeru K Suzuki, and Shu-ichiro Inutsuka. Stochastic particle acceleration in turbulence generated by magnetorotational instability. The Astrophysical Journal, 822(2):88, 2016.

Shigeo S Kimura, Kengo Tomida, and Kohta Murase. Acceleration and escape processes of high-energy particles in turbulence inside hot accretion flows. Monthly Notices of the Royal Astronomical Society, 485(1):163–178, 2019.

Grzegorz Kowal, Elisabete M de Gouveia Dal Pino, and Alex Lazarian. Particle acceleration in turbulence and weakly stochastic reconnection. Physical Review Letters, 108(24):241102, 2012.

Russell M Kulsrud and Attilio Ferrari. The relativistic quasilinear theory of particle acceleration by hydromagnetic turbulence. Astrophysics and Space Science, 12(2):302–318, 1971.

Sayan Kundu, Bhargav Vaidya, and Andrea Mignone. Numerical modeling and physical interplay of stochastic turbulent acceleration for nonthermal emission processes. The Astrophysical Journal, 921(1):74, 2021.

Matthew W. Kunz, James M. Stone, and Eliot Quataert. Magnetorotational turbulence and dynamo in a collisionless plasma. Phys. Rev. Lett., 117:235101, Dec 2016.

Martin Lemoine. Particle acceleration in strong mhd turbulence. Physical Review D, 104(6):063020, 2021.

Martin Lemoine. First-principles fermi acceleration in magnetized turbulence. Physical Review Letters, 129(21):215101, 2022.

Wei Liu, Vahé Petrosian, and John T. Mariska. Combined Modeling of Acceleration, Transport, and Hydrodynamic Response in Solar Flares. I. The Numerical Model. The Astrophysical Journal, 702:1553–1566, September 2009. ISSN 0004-637X.

Malcolm S. Longair. High Energy Astrophysics. Cambridge University Press, 2011. ISBN 0521756189.

Jacob W Lynn, Eliot Quataert, Benjamin DG Chandran, and Ian J Parrish. The efficiency of second-order fermi acceleration by weakly compressible magnetohydrodynamic turbulence. The Astrophysical Journal, 777(2):128, 2013.

Jacob W Lynn, Eliot Quataert, Benjamin DG Chandran, and Ian J Parrish. Acceleration of relativistic electrons by magnetohydrodynamic turbulence: implications for non-thermal emission from black hole accretion disks. The Astrophysical Journal, 791(1):71, 2014.

Yuri Lyubarsky and Michael Liverts. Particle Acceleration in the Driven Relativistic Reconnection. Astrophys. J., 682(2):1436–1442, Aug 2008.

Kirit Makwana, Hui Li, Fan Guo, and Xiaocan Li. Dissipation and particle energization in moderate to low beta turbulent plasma via pic simulations. J. Phys. Conf. Ser., 837(1):012004, 2017.

A. Marcowith, A. Bret, A. Bykov, M. E. Dieckman, L. O'C Drury, B. Lembège, M. Lemoine, G. Morlino, G. Murphy, G. Pelletier, I. Plotnikov, B. Reville, M. Riquelme, L. Sironi, and A. Stockem Novo. The microphysics of collisionless shock waves. Reports on Progress in Physics, 79(4): 046901, Apr 2016.

W. H. Matthaeus. Turbulence in space plasmas: Who needs it? Physics of Plasmas, 28(3):032306, March 2021. ISSN 1070-664X, 1089-7674.

Tania E Medina-Torrejón, Elisabete M de Gouveia Dal Pino, Luis HS Kadowaki, Grzegorz Kowal, Chandra B Singh, and Yosuke Mizuno. Particle acceleration by relativistic magnetic reconnection driven by kink instability turbulence in poynting flux–dominated jets. The Astrophysical Journal, 908(2):193, 2021.

DB Melrose. Resonant scattering of particles and second phase acceleration in the solar corona. Solar Physics, 37(2):353–365, 1974.

Claudio Meringolo, Alejandro Cruz-Osorio, Luciano Rezzolla, and Sergio Servidio. Microphysical plasma relations from special-relativistic turbulence. The Astrophysical Journal, 944(2):122, 2023.

Philipp Mertsch and Subir Sarkar. Fermi Gamma-Ray "Bubbles" from Stochastic Acceleration of Electrons. Physical Review Letters, 107:091101, August 2011. ISSN 0031-9007.

M. Meyer, D. Horns, and H.-S. Zechlin. The Crab Nebula as a standard candle in very high-energy astrophysics. Astronomy & Astrophysics, 523:A2, November 2010. ISSN 0004-6361, 1432-0746.

James A. Miller, Nidhal Guessoum, and Reuven Ramaty. Stochastic Fermi acceleration in solar flares. The Astrophysical Journal, 361:701–708, October 1990. ISSN 0004-637X.

Joonas Nättilä and Andrei M Beloborodov. Heating of magnetically dominated plasma by alfvén-wave turbulence. Physical Review Letters, 128(7):075101, 2022.

Sergei Nayakshin and Fulvio Melia. Self-consistent Fokker-Planck Treatment of Particle Distributions in Astrophysical Plasmas. The Astrophysical Journal Supplement Series, 114:269–288, February 1998. ISSN 0067-0049.

S. O'Sullivan, B. Reville, and A. M. Taylor. Stochastic particle acceleration in the lobes of giant radio galaxies. Mon. Not. R. Astron Soc., 400(1):248–257, 2009.

Theophilos Pisokas, Loukas Vlahos, and Heinz Isliker. Synergy of stochastic and systematic energization of plasmas during turbulent reconnection. Astrophys. J., 852(2):64, jan 2018.

R. Ramaty. Energetic particles in solar flares. In J. Arons, C. McKee, and C. Max, editors, Particle Acceleration Mechanisms in Astrophysics, volume 56 of American Institute of Physics Conference Series, pages 135–154, Jan 1979.

Frank M. Rieger, Valentí Bosch-Ramon, and Peter Duffy. Fermi acceleration in astrophysical jets. Astrophy. Space Sci., 309(1-4):119–125, Jun 2007.

A. A. Schekochihin, S. C. Cowley, W. Dorland, G. W. Hammett, G. G. Howes, E. Quataert, and T. Tatsuno. Astrophysical Gyrokinetics: Kinetic and Fluid Turbulent Cascades in Magnetized Weakly Collisional Plasmas. The Astrophysical Journal Supplement Series, 182:310–377, May 2009. ISSN 0067-0049.

Alexander A. Schekochihin. MHD turbulence: A biased review. Journal of Plasma Physics, 88(5): 155880501, October 2022. ISSN 0022-3778, 1469-7807.

R. Schlickeiser. A viable mechanism to establish relativistic thermal particle distribution functions in cosmic sources. Astronomy and Astrophysics, 143:431–434, February 1985. ISSN 0004-6361.

Reinhard Schlickeiser. Cosmic-ray transport and acceleration. I - Derivation of the kinetic equation and application to cosmic rays in static cold media. II - Cosmic rays in moving cold media with application to diffusive shock wave acceleration. The Astrophysical Journal, 336:243–293, January 1989. ISSN 0004-637X.

I. S. Shklovskii. Akademiia Nauk SSSR Doklady, 90:983, January 1953.

S. Siegert, R. Friedrich, and J. Peinke. Analysis of data sets of stochastic systems. Physics Letters A, 243(5-6):275–280, July 1998. ISSN 03759601.

L. Sironi and A. Spitkovsky. Particle Acceleration in Relativistic Magnetized Collisionless Electron-Ion Shocks. Astrophys. J., 726:75, January 2011.

Lorenzo Sironi and Anatoly Spitkovsky. Relativistic reconnection: An efficient source of non-thermal particles. Astrophys. J. Lett., 783(1):L21, 2014.

J Skilling. Cosmic ray streaming—i effect of alfvén waves on particles. Monthly Notices of the Royal Astronomical Society, 172(3):557–566, 1975.

A. Spitkovsky. Particle Acceleration in Relativistic Collisionless Shocks: Fermi Process at Last? Astrophys. J. Lett., 682:L5, July 2008.

Danny Summers and Chun-yu Ma. A model for generating relativistic electrons in the Earth's inner magnetosphere based on gyroresonant wave-particle interactions. Journal of Geophysical Research, 105:2625–2640, February 2000. ISSN 0148-0227.

Xiaochen Sun and Xue-Ning Bai. Particle diffusion and acceleration in magnetorotational instability turbulence. Monthly Notices of the Royal Astronomical Society, 506(1):1128–1147, 2021.

Shuta J. Tanaka and Katsuaki Asano. On the Radio-emitting Particles of the Crab Nebula: Stochastic Acceleration Model. Astrophys. J., 841(2):78, Jun 2017.

J. M. TenBarge, G. G. Howes, W. Dorland, and G. W. Hammett. An oscillating Langevin antenna for driving plasma turbulence simulations. Computer Physics Communications, 185:578–589, February 2014. ISSN 0010-4655.

Christopher Thompson and Omer Blaes. Magnetohydrodynamics in the extreme relativistic limit. Physical Review D, 57(6):3219–3234, March 1998.

Cristian Vega, Stanislav Boldyrev, Vadim Roytershteyn, and Mikhail Medvedev. Turbulence and particle acceleration in a relativistic plasma. The Astrophysical Journal Letters, 924(1):L19, 2022.

Loukas Vlahos and Heinz Isliker. Particle acceleration and heating in a turbulent solar corona. Plasma Phys. Control. Fusion, 61(1):014020, Jan 2019.

Loukas Vlahos, Heinz Isliker, and Fabio Lepreti. Particle acceleration in an evolving network of unstable current sheets. The Astrophysical Journal, 608(1):540, 2004.

G. R. Werner, D. A. Uzdensky, B. Cerutti, K. Nalewajko, and M. C. Begelman. The Extent of Power-law Energy Spectra in Collisionless Relativistic Magnetic Reconnection in Pair Plasmas. Astrophys. J. Lett., 816(1):L8, Jan 2016.

G. R. Werner, D. A. Uzdensky, M. C. Begelman, B. Cerutti, and K. Nalewajko. Non-thermal particle acceleration in collisionless relativistic electron-proton reconnection. Monthly Notices of the Royal Astronomical Society, 473:4840–4861, February 2018. ISSN 0035-8711.

Gregory R. Werner and Dmitri A. Uzdensky. Nonthermal Particle Acceleration in 3D Relativistic Magnetic Reconnection in Pair Plasma. Astrophys. J. Lett., 843(2):L27, Jul 2017.

Gregory R. Werner and Dmitri A. Uzdensky. Reconnection and particle acceleration in 3D current sheet evolution in moderately-magnetized astrophysical pair plasma. Journal of Plasma Physics, 87(6):905870613, December 2021. ISSN 0022-3778, 1469-7807.

Kai Wong, Vladimir Zhdankin, Dmitri A. Uzdensky, Gregory R. Werner, and Mitchell C. Begelman. First-principles Demonstration of Diffusive-advective Particle Acceleration in Kinetic Simulations of Relativistic Plasma Turbulence. The Astrophysical Journal Letters, 893:L7, April 2020. ISSN 0004-637X.

Kane Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. IEEE Transactions on Antennas and Propagation, 14(3):302–307, May 1966. ISSN 1558-2221.

Feng Yuan, Eliot Quataert, and Ramesh Narayan. Nonthermal electrons in radiatively inefficient accretion flow models of sagittarius a*. Astrophys. J., 598(1):301, 2003.

S. Zenitani and M. Hoshino. The Generation of Nonthermal Particles in the Relativistic Magnetic Reconnection of Pair Plasmas. Astrophys. J. Lett., 562(1):L63–L66, Nov 2001.

Jian-Fu Zhang and Fu-Yuan Xiang. Energetic particle acceleration in compressible magnetohydrodynamic turbulence. The Astrophysical Journal, 922(2):209, 2021.

Vladimir Zhdankin. Non-thermal particle acceleration from maximum entropy in collisionless plasmas. Journal of Plasma Physics, 88(3):175880303, June 2022.

Vladimir Zhdankin, Gregory R. Werner, Dmitri A. Uzdensky, and Mitchell C. Begelman. Kinetic Turbulence in Relativistic Plasma: From Thermal Bath to Nonthermal Continuum. Physical Review Letters, 118:055103, February 2017. ISSN 0031-9007.

Vladimir Zhdankin, Dmitri A. Uzdensky, Gregory R. Werner, and Mitchell C. Begelman. Numerical investigation of kinetic turbulence in relativistic pair plasmas - I. Turbulence statistics. Monthly Notices of the Royal Astronomical Society, 474:2514–2535, February 2018a. ISSN 0035-8711.

Vladimir Zhdankin, Dmitri A. Uzdensky, Gregory R. Werner, and Mitchell C. Begelman. System-size Convergence of Nonthermal Particle Acceleration in Relativistic Plasma Turbulence. The Astrophysical Journal Letters, 867:L18, November 2018b. ISSN 0004-637X.

Vladimir Zhdankin, Dmitri A. Uzdensky, Gregory R. Werner, and Mitchell C. Begelman. Electron and Ion Energization in Relativistic Plasma Turbulence. Physical Review Letters, 122:055101, February 2019. ISSN 0031-9007.

Vladimir Zhdankin, Dmitri A. Uzdensky, Gregory R. Werner, and Mitchell C. Begelman. Kinetic turbulence in shining pair plasma: Intermittent beaming and thermalization by radiative cooling. Monthly Notices of the Royal Astronomical Society, 493:603–626, March 2020. ISSN 0035-8711.

Vladimir Zhdankin, Dmitri A Uzdensky, and Matthew W Kunz. Production and persistence of extreme two-temperature plasmas in radiative relativistic turbulence. The Astrophysical Journal, 908(1):71, 2021.

Ellen G. Zweibel. The microphysics and macrophysics of cosmic raysa). Physics of Plasmas, 20(5): 055501, May 2013. ISSN 1070-664X.