

Energy diffusion and advection coefficients in kinetic simulations of relativistic plasma turbulence

Kai W. Wong,¹ Vladimir Zhdankin,^{2,3,4} Dmitri A. Uzdensky^{id},^{1,5★} Gregory R. Werner^{id}¹ and Mitchell C. Begelman^{id}^{6,7}

¹Center for Integrated Plasma Studies, Physics Department, 390 UCB, University of Colorado, Boulder, CO 80309, USA

²Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

³Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

⁴Department of Physics, University of Wisconsin – Madison, Madison, WI 53706, USA

⁵Rudolf Peierls Centre for Theoretical Physics, Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, UK

⁶JILA, University of Colorado and National Institute of Standards and Technology, 440 UCB, Boulder, CO 80309, USA

⁷Department of Astrophysical and Planetary Sciences, University of Colorado, 391 UCB, Boulder, CO 80309, USA

Accepted 2025 August 8. Received 2025 July 30; in original form 2025 February 3

ABSTRACT

Turbulent, relativistic non-thermal plasmas are ubiquitous in high-energy astrophysical systems, as inferred from broad-band non-thermal emission spectra. The underlying turbulent non-thermal particle acceleration (NTPA) processes have traditionally been modelled with a Fokker–Planck (FP) diffusion–advection equation for the particle energy distribution. We test FP-type NTPA theories by performing and analysing particle-in-cell simulations of turbulence in collisionless relativistic pair plasma. By tracking large numbers of particles in simulations with different initial magnetization and system size, we first test and confirm the applicability of the FP framework. We then measure the FP energy diffusion (D) and advection (A) coefficients as functions of particle energy γmc^2 , and compare their dependence to theoretical predictions. At high energies, we robustly find $D \sim \gamma^2$ for all cases. Hence, we fit $D = D_0 \gamma^2$ and find a scaling consistent with $D_0 \sim \sigma^{3/2}$ at low instantaneous magnetization $\sigma(t)$, flattening to $D_0 \sim \sigma$ at higher $\sigma \sim 1$. We also find that the power-law index $\alpha(t)$ of the particle energy distribution converges exponentially in time. We build and test an analytic model connecting the FP coefficients and $\alpha(t)$, predicting $A(\gamma) \sim \gamma \log \gamma$. We confirm this functional form in our measurements of $A(\gamma, t)$, which allows us to predict $\alpha(t)$ through the model relations. Our results suggest that the basic second-order Fermi acceleration model, which predicts $D_0 \sim \sigma$, may not be a complete description of NTPA in turbulent plasmas. These findings encourage further application of tracked particles and FP coefficients as a diagnostic in kinetic simulations of various astrophysically relevant plasma processes like collisionless shocks and magnetic reconnection.

Key words: acceleration of particles – plasmas – relativistic processes – turbulence.

1 INTRODUCTION

Relativistic charged particles with non-thermal power-law energy distributions are inferred, from observed non-thermal radiation spectra, to exist in diverse astrophysical systems such as pulsar wind nebulae, black hole accretion flows and jets, e.g. in active galactic nuclei, and gamma-ray bursts (e.g. Longair 2011). The question of how such non-thermal particle acceleration (NTPA) occurs is a long-standing fundamental physics problem, and has been approached both through analytical theory and numerical simulations. Many of these astrophysical systems are manifestly turbulent, and dissipation of magnetized plasma turbulence is an attractive mechanism for explaining non-thermal power-law tails of the particle energy distribution. If NTPA predominantly comes from the ideal motional electric field, then particle acceleration

requires vigorous plasma motion. Turbulence provides this motion at a range of length-scales through the turbulent cascade, which naturally parallels a non-thermal spectrum of particle energies with corresponding resonant gyro-scales.

While proposed NTPA mechanisms include shocks, magnetic reconnection, and turbulence, 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 energization is captured through the Fokker–Planck (FP) diffusion–advection equation in momentum space

$$\frac{\partial F}{\partial t} = \frac{\partial}{\partial \mathbf{p}} \cdot \left(\mathbf{D}_{\text{pp}} \cdot \frac{\partial F}{\partial \mathbf{p}} \right) - \frac{\partial}{\partial \mathbf{p}} \cdot (A_{\text{p}} F), \quad (1)$$

where $F(\mathbf{p}, t)$ is the (spatially averaged) particle momentum distribution, and $\mathbf{D}_{\text{pp}}(\mathbf{p}, t)$ and $A_{\text{p}}(\mathbf{p}, t)$ are, respectively, the momentum-space diffusion coefficient tensor and advection coefficient vector. Existing theoretical NTPA models are usually based on quasi-

* E-mail: dmitri.uzdensky@physics.ox.ac.uk

linear theory (QLT; e.g. Kulsrud & Ferrari 1971; Schlickeiser 1989; Chandran 2000; Demidem, Lemoine & Casse 2020), and differences between 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, Guessoum & Ramaty 1990; Nayakshin & Melia 1998; Blasi 2000; Summers & Ma 2000; Becker, Le & Dermer 2006; Liu, Petrosian & Mariska 2009; Mertsch & Sarkar 2011; Asano et al. 2014; Bian et al. 2014; Kimura, Murase & Toma 2015; Kundu, Vaidya & Mignone 2021). Numerical simulations offer the opportunity to first test the diffusive nature of particle energization, 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).

Magnetized plasma turbulence has been studied both with magnetohydrodynamic (MHD) simulations and, more recently, with 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; Dmitruk, Matthaeus & Seenu 2004; Kowal, Dal Pino & Lazarian 2012; Lynn et al. 2013, 2014; Kimura et al. 2016; Kimura, Tomida & Murase 2019; Medina-Torrejón et al. 2021; Sun & Bai 2021; Zhang & Xiang 2021; Bresci et al. 2022). This method is also applied to cases where the fields are prescribed and not simulated (e.g. Vlahos, Isliker & Lepreti 2004; Isliker et al. 2017b; Demidem et al. 2020). While less computationally expensive than fully kinetic 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 back-reaction of energetic particles on the electromagnetic fields. Indeed, in extreme cases, the turbulent cascade may be damped by non-thermal particles across a range of scales (Lemoine, Murase & Rieger 2024). In contrast to MHD, first-principles particle-in-cell (PIC) kinetic 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 non-thermal power-law particle energy distribution (e.g. Zhdankin et al. 2017; Comisso & Sironi 2018; Zhdankin et al. 2018b; Comisso & Sironi 2019; Hankla et al. 2022; Näätäilä & Beloborodov 2022; Vega et al. 2022; Meringolo et al. 2023; Vega, Boldyrev & Roytershteyn 2024). Hence, there is an opportunity for more detailed tests of NTPA theories through direct inspection of tracked particles.

Our previous paper, Wong et al. (2020), a numerical analysis of an ensemble of tracked particles in a large 3D PIC simulation, 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. We 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. Subsequent work by Comisso & 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 paper, we 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 magnetization 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 our non-radiative simulations due to continuous turbulent heating, we 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, we measure the power-law index of the non-thermal part of the particle energy distribution as a function of time, and relate it to the FP coefficients with an analytical model. We thereby uncover how the power-law index depends on both initial and instantaneous parameters, which yields a more complete understanding of how non-thermal 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 paper is organized as follows: Section 2 outlines previous analytical research on turbulent NTPA, and presents an analytic model relating the FP coefficients to the evolution of the power-law index of the particle energy distribution. Section 3 describes our PIC simulation setup and parameter scans. Section 4 discusses general features of the turbulent particle energization and their time evolution in our simulations. Section 5 presents our 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 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 7 presents our findings on the energy diffusion coefficient as a function of time and system parameters. Section 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 2. Finally, Section 9 summarizes our results.

2 THEORY

2.1 Fokker–Planck theories of non-thermal particle acceleration

The full momentum-space FP equation is given in (1). There, the diffusion and advection coefficients are, respectively, a tensor (\mathbf{D}_{pp}) and a vector (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, we will now assume the form $A_p = A_p \hat{p}$ and $\mathbf{D}_{pp} = D_{pp} \mathbf{I}$, where A_p and D_{pp} are the scalar advection and diffusion coefficients, respectively, with \hat{p} being the unit vector along p and \mathbf{I} the identity matrix.

For simplicity, we will ignore the pitch-angle dependence of the FP equation (i.e. by averaging over pitch angles), which would require a much deeper level of analysis. Thus, we will limit the analysis to a reduced description of the FP equation for the global particle distribution $f(\gamma, t)$ in energy only

$$\partial_t f = \partial_\gamma (D \partial_\gamma f - A f), \quad (2)$$

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 , where m is the particle rest-mass and c is the speed of light. However, since it is common to state the FP coefficients in momentum space, it is useful to provide, for reference, conversion formulae between the energy-space and momentum-space forms. Beginning with (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), we obtain

$$\frac{\partial F}{\partial t} = \frac{1}{p^2} \frac{\partial}{\partial p} \left(p^2 D_{pp} \frac{\partial F}{\partial p} \right) - \frac{1}{p^2} \frac{\partial}{\partial p} (p^2 A_p F), \quad (3)$$

where dn simplifies to $4\pi p^2 F(p) dp$. This is the form used in, e.g. Schlickeiser (1985, 1989). Finally, in the ultrarelativistic limit $\gamma \gg 1$ that we will focus on in this paper, the particle energy γmc^2 is just directly proportional to its momentum p , i.e. $\gamma = p/mc$; we can then substitute $f(\gamma) = 4\pi p^2 F(p)$ and thus write

$$\frac{\partial f}{\partial t} = \frac{\partial}{\partial \gamma} \left[D_{pp} \left(\frac{\partial f}{\partial \gamma} - 2 \frac{f}{\gamma} \right) \right] - \frac{\partial}{\partial \gamma} (A_p f). \quad (4)$$

Comparing this to (2), we 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 $A = A_p + 2D_{pp}/\gamma$.

2.2 Analytical model of the non-thermal power-law tail

Previous PIC numerical simulation studies have shown the formation and gradual evolution of a relativistic non-thermal 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. We wish to understand the relationship between the FP coefficients and the time evolution of the power-law index of the non-thermal section of $f(\gamma)$. Hence, we consider the FP equation (2), where f , D , and A are functions of γ and t . We will focus here only on the non-thermal 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 $f(\gamma, t) = K(t)\gamma^{-\alpha(t)}$.

Here, K should not be viewed as a normalization factor, since the number of particles in the power law does not have to be conserved. Instead, we 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 by Zhdankin et al. (2017, 2018b). Particles moving in and out of these reservoirs can change the total number of particles in the non-thermal tail.

Theoretical models based on QLT (e.g. Kulsrud & Ferrari 1971; Melrose 1974; Skilling 1975; Blandford & Eichler 1987; Chandran 2000; Cho & Lazarian 2006; Demidem et al. 2020) and generalized Fermi acceleration (Lemoine 2019, 2021, 2025) commonly predict $D(\gamma, t) \propto \gamma^2$ in the non-thermal range. Hence, for this model, we assume

$$D(\gamma, t) = D_0(t)\gamma^2. \quad (6)$$

In the absence of an advection coefficient, this diffusion coefficient causes the mean energy $\bar{\gamma}$ of a collection of particles to increase at the instantaneous rate of $\dot{\bar{\gamma}} \equiv \partial_t \bar{\gamma} = 2D_0\bar{\gamma}$: 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 ($\dot{\bar{\gamma}}/\bar{\gamma} = \text{const}$).

While we could also assume a functional form for $A(\gamma, t)$, specifying three independent forms for f , D , and A will almost certainly be inconsistent due to overconstraining the problem. Instead, since A is by far the most uncertain theoretically, we use (2) with the above adopted functional forms (5) and (6) for $f(\gamma, t)$ and $D(\gamma, t)$, respectively, to solve for $A(\gamma, t)$ analytically, and then later compare it to our numerical measurements of A in Section 8. That is, we infer A from consistency conditions on the existence of a power-law range in the energy distribution.

First, we integrate the FP equation (2) from γ to ∞ . Keeping in mind (5) and (6), we assume $\alpha > 1$ and $Af \rightarrow 0$ as $\gamma \rightarrow \infty$, so that all the upper integral limits can be discarded (note that we do not need the stronger condition of $\alpha < 2$ that would cause energy to be independent of the upper limit; we thus allow energy to be dominated by the highest energy particles). The non-advection terms then evaluate to:

$$\int_{\gamma}^{\infty} \partial_{\gamma'} (D \partial_{\gamma'} f) d\gamma' = -D \partial_{\gamma} f = -D_0 \gamma^2 (-\alpha/\gamma) f = D_0 \alpha \gamma f, \\ \int_{\gamma}^{\infty} \partial_t f d\gamma' = -\partial_t \left(K \frac{\gamma^{1-\alpha}}{1-\alpha} \right) = -\frac{\gamma f}{1-\alpha} \left(-\dot{\alpha} \log \gamma + \frac{\dot{K}}{K} + \frac{\dot{\alpha}}{1-\alpha} \right). \quad (7)$$

We collect into (2) and eliminate the common factors of f to get

$$A = -\frac{\gamma}{1-\alpha} \left(-\dot{\alpha} \log \gamma + \frac{\dot{K}}{K} + \frac{\dot{\alpha}}{1-\alpha} + D_0 (\alpha - \alpha^2) \right). \quad (8)$$

This can be recast in a convenient, compact form as

$$A(\gamma, t) = A_0(t) \gamma \log [\gamma/\gamma_A^*(t)], \quad (9)$$

with:

$$A_0(t) = \frac{\dot{\alpha}}{1-\alpha} = -\frac{d \log(\alpha-1)}{dt}, \quad (10) \\ \gamma_A^*(t) = \exp \left[\frac{\dot{K}/K + D_0 (\alpha - \alpha^2)}{\dot{\alpha}} + \frac{1}{1-\alpha} \right], \quad (11)$$

where all right-hand side variables in (10) and (11) are functions of t only. The result (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 (9), with (10) and (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 (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 (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, (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 $A_{\min} = -A_0 \gamma_A^*/e$. We 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 $D \propto \gamma^2$. Finally, an equivalent statement to (9) is that A/γ is linear in $\log \gamma$, which we will use in Section 8.

Up to this point, we have introduced fixed γ -dependences for the particle energy distribution f and FP coefficients D and A , parametrized 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 Section A for some details on tracked particles).

We later find in Section 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; hence an integral comparison can be less noisy and more illuminating. The tradeoff for this smoothing is that we will have to select the integration constants (somewhat arbitrarily). To enable this, we rearrange the A -variables equations (10) and (11) into time integrals for the power-law variables $K(t)$ and $\alpha(t)$. From (10),

$$\frac{\alpha(t) - 1}{\alpha(t_0) - 1} = \exp \left[- \int_{t_0}^t A_0(t') dt' \right], \quad (12)$$

where t_0 is an arbitrary reference time and $\alpha(t_0)$ is an integration constant. Similarly, (11) rearranged for \dot{K}/K integrates to

$$\frac{K(t)}{K(t_0)} = \exp \left\{ \int_{t_0}^t \left[D_0 (\alpha^2 - \alpha) + \dot{\alpha} \left(\log \gamma_A^* + \frac{1}{\alpha - 1} \right) \right] dt' \right\}. \quad (13)$$

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 (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. We also note that the exponential form of (12) has a suggestive (although not quite fully compatible) similarity to the exponential fit to $\alpha(t)$ introduced in Section 5.

To compute the advection coefficient $A(\gamma, t)$ in our PIC simulations, we will first measure the energy-dependent average rate of particle energy gain $M \equiv A + \partial_\gamma D$ (see Section 6.1). Later, we 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, we derive analogues of (9)–(11) for M using $\partial_\gamma D = 2D_0\gamma$ which comes from (6):

$$M(\gamma, t) = M_0 \gamma \log(\gamma/\gamma_M^*), \quad (14)$$

where $M_0 = A_0$ is the same as in (10) (it is convenient to have a different symbol for this pre-factor as later fits to A and M will not necessarily result in the same value) and

$$\gamma_M^*(t) = \exp \left[\frac{\dot{K}/K + D_0(-\alpha^2 + 3\alpha - 2)}{\dot{\alpha}} + \frac{1}{1 - \alpha} \right], \quad (15)$$

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

$$\frac{\alpha(t) - 1}{\alpha(t_0) - 1} = \exp \left[- \int_{t_0}^t M_0(t') dt' \right], \quad (16)$$

$$\frac{K(t)}{K(t_0)} = \exp \left\{ \int_{t_0}^t \left[D_0(\alpha^2 - 3\alpha + 2) + \dot{\alpha} \left(\log \gamma_M^* + \frac{1}{\alpha - 1} \right) \right] dt' \right\} \quad (17)$$

In summary, by adopting an Ansatz that the particle energy distribution is a time-dependent power law $f(\gamma, t) = K(t)\gamma^{-\alpha(t)}$ (5) and that the energy diffusion coefficient is $D(\gamma, t) = D_0(t)\gamma^2$ (6), we obtain a prediction for the energy advection coefficient of $A(\gamma, t) = A_0(t)\gamma \log[\gamma/\gamma_A^*(t)]$ (9), where A_0 is a function of α and $\dot{\alpha}$ (10) and γ_A^* is a function of \dot{K}/K , D_0 , α , and $\dot{\alpha}$ (11).

Our 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. For example, it may be a natural consequence of particle segregation based on acceleration rate (Lemoine & Malkov 2020). 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 SIMULATIONS

3.1 Numerical setup

We analyse 3D PIC simulations of externally driven turbulence in collisionless relativistic pair plasma performed with our code ZELTRON (Cerutti et al. 2013). The system is a periodic cube of side length L and a guide magnetic field $B_0\hat{z}$. The plasma, with total (electrons and positrons combined) charged particle density n_0 , is initialized from a uniform isotropic Maxwell–Jüttner relativistic thermal distribution. We fix the initial temperature to $T_0 \equiv \theta_0 mc^2 = 100 mc^2$, corresponding to an average Lorentz factor of $\gamma_{\text{avg}0} \approx 3\theta_0 = 300$.

We define the (time-dependent) characteristic Larmor radius $\rho_e(t) \equiv \gamma_{\text{avg}} mc^2 / e B_{\text{rms}}$ and skin depth $d_e(t) \equiv (\gamma_{\text{avg}} mc^2 / 4\pi n_0 e^2)^{1/2}$, where $\gamma_{\text{avg}}(t)$ is the instantaneous average Lorentz factor and $B_{\text{rms}}(t)$ 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’) magnetization $\sigma \equiv 3B_{\text{rms}}^2 / 16\pi n_0 \gamma_{\text{avg}} mc^2 = (3/4)(d_e/\rho_e)^2$, which is the ratio of the magnetic enthalpy $B_{\text{rms}}^2/4\pi$ to the relativistic plasma enthalpy $(4/3)n_0 \gamma_{\text{avg}} mc^2$; the system size relative to the Larmor radius L/ρ_e ; and the system size relative to the skin depth L/d_e . For the initial ultrarelativistic Maxwell–Jüttner distribution, the initial magnetization $\sigma_0 = B_0^2/16\pi n_0 T_0$ 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 (TenBerge 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 δB_{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 ultrarelativistic Maxwell–Jüttner distribution, $d_e/\lambda_D = \sqrt{3}$). More precisely, we 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 $N_{\text{ppc}} = 64$. 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 time-scale set by L/v_A), with $t_f \sim 20L/c$ being most common.

Our 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 us from 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. Prior to performing the analysis described in this paper, we process the tracked particle data to remove energy oscillations associated with $\mathbf{E} \times \mathbf{B}$ drift; this procedure is described in Appendix A.

3.2 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 summarized in Table 1. They are grouped as follows:

(i) A system-size scan at fixed $\sigma_0 = 3/8$ (where $\rho_{e0} = \sqrt{2}d_{e0}$) with L/ρ_{e0} ranging from 171 to 1024, corresponding to N varying from 256 to 1536. This group contains our largest simulation with $L/\rho_{e0} = 1024$, performed on an $N^3 = 1536^3$ -cell grid.

(ii) Sets of smaller simulations with varying σ_0 : one with $N = 768$ and $\sigma_0 \in \{3/2, 3/4, \dots, 3/128\}$, and another with $N = 384$ and $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$. As mentioned above, we set $\Delta x = \min(\rho_{e0}, \sqrt{2}d_{e0})/1.5$ to resolve the smaller of the two kinetic scales while maximizing 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.

(iii) 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 $L/\rho_{e0} = 256$ ($N = 384$).

Let us discuss briefly some complications caused by having three initial dimensionless ratios: L/ρ_{e0} , L/d_{e0} , and $d_{e0}/\rho_{e0} \propto \sigma_0^{1/2}$. 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, we fix d_{e0}/ρ_{e0} . However, the σ_0 scan is more complicated:

Table 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 $N = 384$	384	3/2	256	1.5	181	2.1	7.8	64
	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 $N = 768$	768	3/2	512	1.5	362	2.1	22.3	64
	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 scan	256	3/8	171	1.5	241	1.1	16.7	64
	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

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 pre-determined 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 magnetization or to the shorter inertial range.

A full 2D 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 our simulations are in this asymptotic large-system 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 our study with respect to the σ_0 scans. Accordingly, we 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.

4 SYSTEM OVERVIEW

Before beginning the main analysis, we first provide a qualitative overview of how the driven turbulent system evolves. We 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.

Fig. 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 non-thermal tail after a few Alfvén crossing times. This non-thermal segment begins near the average particle energy and extends to roughly the system-size limit $\gamma_{\text{max}} \equiv LeB_0/2mc^2 = L\gamma_{\text{avg}}/2\rho_{e0}$, where the gyroradius reaches the system size. For example, our typical runs with $L/\rho_{e0} = 512$ have $\gamma_{\text{max}} \simeq 7.7 \times 10^4$, while our largest simulation has $L/\rho_{e0} = 1024$ and $\gamma_{\text{max}} \simeq 1.5 \times 10^5$. As the




Figure 1. Evolution of particle energy distribution for simulations with $N = 768$ and different $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$. Five times (separated by equal durations) are shown for each simulation, including the initial time, three intermediate times, and the final time (in order of decreasing colour fade).




Figure 2. Moving average of rate of change of average particle energy ($d\gamma_{\text{avg}}/dt$) (normalized by v_{A0}/L) versus time for simulations with $N = 768$ and $\sigma_0 \in \{3/2, 3/4, \dots, 3/128\}$. The width of the moving average is L/v_{A0} , and this filtering is denoted by angled brackets. Black dashed lines correspond to $\dot{\gamma}_{\text{avg}} \propto \sigma_0$.

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

Fig. 2 shows the plasma heating history, – i.e. the time evolution of the rate of change of the global average particle energy $\dot{\gamma}_{\text{avg}}$, 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 $\sigma_0 \in \{3/2, 3/4, \dots, 3/128\}$. We observe that, aside from order-unity fluctuations, $\dot{\gamma}_{\text{avg}}$ 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 non-thermal particle acceleration), corresponding to a long-term linear increase in $\dot{\gamma}_{\text{avg}}$ with time. The roughly even separation of the $\dot{\gamma}_{\text{avg}}$ lines in Fig. 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 $\dot{\gamma}_{\text{avg}} \propto \sigma_0$ are overlaid on Fig. 2.

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




Figure 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 marker. The latest time in the trajectory is marked by the black dot.

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 (2) being zero. However, the time-dependence of $f(\gamma)$ described in Section 4 complicates the interpretation of the Fokker–Planck coefficients in the context of a steady-state energy spectrum. Nevertheless, we 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 non-thermal 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 non-thermal ranges (see Fig. 4a). For a particle energy distribution at some time t , the local (in γ) power-law index can be defined as $\alpha_{\text{loc}}(\gamma, t) \equiv -\partial \log f / \partial \log \gamma$ (and hence a pure power-law segment is characterized by $\partial \alpha_{\text{loc}} / \partial \gamma = 0$). At a given instant, the local indices $\alpha_{\text{loc}}(\gamma)$ can be summarized into a representative single value, $\alpha(t)$, characterizing the predominant value of $\alpha_{\text{loc}}(\gamma, t)$ in the high-energy non-thermal tail. Moreover, a special value of $\alpha(t)$, which we denote as α_{fin} , is often chosen to represent the late-time asymptotic (or ‘final’) value of $\alpha(t)$. These values $\alpha(t)$ and α_{fin} are main quantities of interest and have observational implications and theoretical significance. For example, the system-size dependence of α_{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, we investigate the time dependence of α , develop a prescription for measuring α_{fin} , and explore their dependence on system parameters such as the system size and initial magnetization.

5.1 Fitting the power-law index

The process of obtaining $\alpha(t)$ from $\alpha_{\text{loc}}(\gamma, t)$ at a given time t is complicated by the necessity to select the section of $\alpha_{\text{loc}}(\gamma)$ 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 maximize the accuracy of the fit. Furthermore, one must then decide how to obtain the late-time asymptotic value α_{fin} . For example, past studies have selected α_{fin} based on ‘the time with the longest fitted power-law segment’ (Zhdankin et al. 2017) or a ‘time that is logarithmic




Figure 4. Particle energy distribution $f(\gamma)$ (a) and the local power-law index $-\alpha_{\text{loc}} = -\partial \log f / \partial \log \gamma$ (b) at several different times for a representative simulation, having $\sigma_0 = 3/8$ and $L/\rho_{e0} = 512$. In panel (a), dashed lines show the power-law fits. In panel (b), dashed lines are the fitted cubics, crosses mark the geometric mean of γ_{avg} and γ_{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.

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, we describe a procedure to distill a unique $\alpha(t)$ value for a given distribution $f(\gamma, t)$ that is well suited to the particle energy distributions produced by our simulations and reasonably insensitive to fitting choices. This scheme yields values 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 non-thermal tail has fully developed and a high-energy pileup has appeared, there is a local minimum in the local power-law index $-\alpha_{\text{loc}}(\gamma)$ (Fig. 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 $-\alpha_{\text{loc}}(\gamma)$ 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 $-\alpha_{\text{loc}}(\gamma, t)$ as the overall (time-dependent) power-law index $\alpha(t)$ characterizing the non-thermal part of the spectrum. We therefore construct an automatic procedure to extract the local minimum of $-\alpha_{\text{loc}}(\gamma, t)$. We 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 non-thermal power-law tail and finally a pileup and cutoff at highest energies.

Our procedure works as follows: We locate the local minimum by fitting a cubic, in $\log \gamma$ space, (Fig. 4b, dashed lines) to the γ -range within a factor of five of the geometric mean of the average energy γ_{avg} and the system-size limited energy γ_{max} (Fig. 4b, crosses). The

Lorentz factor corresponding to the desired local minimum (Fig. 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); points where the corresponding energy spectrum density $f(\gamma)$ is less than 10^{-9} of the peak value at that time are excluded due to noise. We 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. Fig. 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 Fig. 4(a), it very closely fits the centre of the power-law range, and tends to trim off the late-time pileup at γ_{max} . Also, as it fits a turning point in $\alpha_{\text{loc}} \equiv -\partial \log f / \partial \log \gamma$, it can be said to fit the straightest section of the power-law section of $f(\gamma)$. The main idea of this procedure is to adopt the first local minimum value of $-\alpha_{\text{loc}}$ as the preferred single-value α for a given energy spectrum; the exact method for locating this local minimum is secondary.

5.2 Application to the system-size scan and statistical samples

We apply our procedure for identifying the single-value power-law index α on $f(\gamma, t)$ -snapshots from the system-size scan simulations with fixed $\sigma_0 = 3/8$ and varying $L/\rho_{e0} \in \{1024, 683, 512, 341, 256, 171\}$. We ignore data points before $t = 2L/v_{A0}$, when turbulence has not yet fully developed. Fig. 5 shows the resulting $\alpha(t)$. The vertical line segments on each plotted point are obtained as follows: expand a symmetric α_{loc} -range around α until it covers a surrounding contiguous γ -range of at least half a decade; the vertical lines then indicate the actual α_{loc} -range included in this γ -range (neither of which must be symmetric). These ‘error bars’ indicate the range of surrounding α_{loc} values and should not be overly interpreted in a statistical sense. A strongly single-sided ‘error bar’ indicates that $-\alpha$ lies in a local minimum of $-\alpha_{\text{loc}}(\gamma)$. We discard any data point with an ‘error-bar’ range greater than 40 per cent of α as this indicates an insufficiently converged power-law segment.




Figure 5. Time evolution of power-law indices α for simulations with different system size L/ρ_{e0} but the same $\sigma_0 = 3/8$, with exponential fits (dashed lines). Data points for each simulation are slightly time-shifted for visibility.




Figure 6. Dependence of the three exponential fit parameters (α_∞ , Δ_α , and τ normalized to L/v_{A0}) of (18) on the relative system size L/ρ_{e0} for simulations with $\sigma_0 = 3/8$. The fit by (19) is shown in a dashed line.

Fig. 5 shows that $\alpha(t)$ decays with time in every simulation. Hence, we fit each $\alpha(t)$ with a three-parameter exponential decay:

$$\alpha(t) = \alpha_\infty + \Delta_\alpha \exp(-t/\tau), \quad (18)$$

and find that the fits are excellent. Although having just warned against statistical interpretation of the ‘error bars’, we 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 α_{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 time-scale 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 non-thermal power-law range would eventually be squeezed from both sides until it disappears entirely.

We use the fitted exponential parameters to quantify and compare the per-simulation time-evolution characteristics as functions of the relative system size L/ρ_{e0} . Fig. 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. This matches the basic observation in Fig. 5 that $\alpha(t)$ tends to shift upwards in magnitude with increasing L/ρ_{e0} , but otherwise behaves similarly between simulations. The convergence time-scale τ being a few Alfvén times matches the natural expectation for the system to evolve on Alfvénic time-scales. The α_∞ data suggests that the simulations are not quite 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

Table 2. Statistical variation in $\alpha(t)$ exponential fit parameters for sets of repeated simulations with the same physical parameters, listed as mean \pm standard deviation. Both sets had $\sigma_0 = 3/8$ and $N_{\text{ppc}} = 32$.

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




Figure 7. The time evolution of the power-law index α for simulations with varying σ_0 and L/ρ_{e0} , each fit by an exponential of the form (18) (solid and dashed lines). Fig. 8 presents the best-fitting parameters.

a function of L/ρ_{e0} is

$$\alpha_\infty = 0.73 \log_{10}(L/\rho_{e0}) + 0.72, \quad (19)$$

and this is displayed in Fig. 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, we analyse two sets of computationally cheaper simulations where the same system parameters were used multiple times with different random seeds. One set used seven simulations with $\sigma_0 = 3/8$, $N_{\text{ppc}} = 32$, 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 2. The scatter in α_∞ is only a few per cent, whereas the variation levels of Δ_α and τ are both substantial. Therefore, it is plausible that the α_∞ -trend observed in Fig. 6 is real, but the variations in Δ_α and τ are noise.

5.3 Application to the σ_0 scan

We repeat the above analysis on the $N=768$ σ_0 -scan and the $N=384$ σ_0 -scan and combine the results. Fig. 7 shows the $\alpha(t)$ with exponential fits while Fig. 8 shows the exponential fit 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 $3L/v_{A0}$ but displays a large scatter for $\sigma_0 \geq 3/4$. These high- σ_0 values of τ should not be given too much credence. Indeed, one can observe in Fig. 7 that ($N = 384$, $\sigma_0 = 3/2$) has few points, while ($N = 768$, $\sigma_0 = 3/2$) covers only a narrow range of α -values, leading to a curve too shallow to reliably characterize the exponential. Furthermore, Table 2 shows that τ already varies substantially for $\sigma_0 = 3/8$ and one expects even

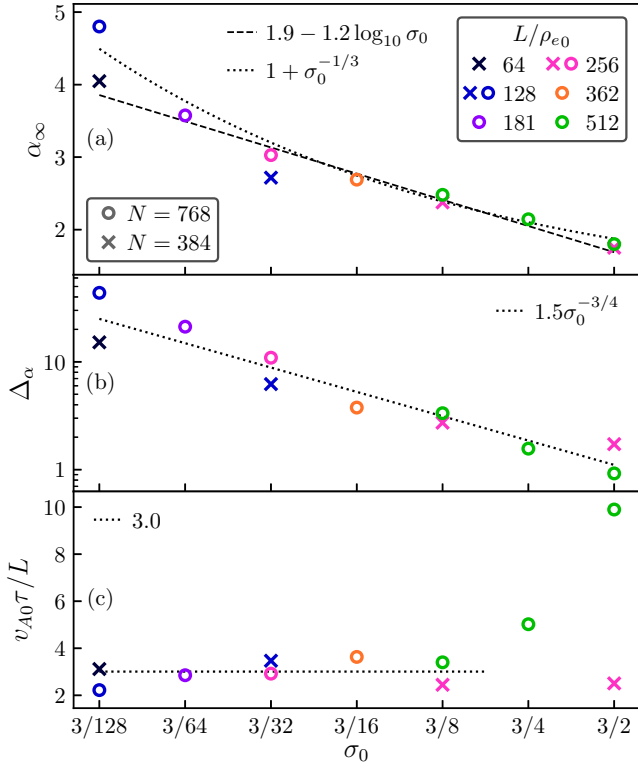


Figure 8. Dependence of the three exponential fit parameters (α_∞ , Δ_α , and τ normalized to L/v_{A0}) of (18) on σ_0 . Data points corresponding to different values of L/ρ_{e0} are indicated by different colours. Fits discussed in the text are shown in dotted or dashed lines.

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 characterize the effect of L/ρ_{e0} with this data set. However, we note that for the four σ_0 values for which there are both $N = 768$ and $N = 384$ simulations, α_∞ is slightly lower in the $N = 384$ case, which is consistent with the Section 5.2 findings. Fig. 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 dependences of α_∞ and Δ_α on the system parameters observed in our simulations, we offer the following simple and practical empirical fits for these functions. First, 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 Section 5.2 shows a significant trend for α_∞ against L/ρ_{e0} . Consequently, we 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.

We find that the best log-linear fit of α_∞ as a function of σ_0 is:

$$\alpha_\infty \approx 1.9 - 1.2 \log_{10} \sigma_0. \quad (20)$$

Alternatively, the offset power-law

$$\alpha_\infty \approx 1 + \sigma_0^{-1/3} \quad (21)$$

is also a good fit (the motivation for the constant term being equal to 1 is that this is a physical lower limit for α). These fits are displayed in Fig. 8(a).

We also find that

$$\Delta_\alpha \approx 1.5 \sigma_0^{-3/4} \quad (22)$$

is a good fit for Δ_α , as shown in Fig. 8(b).

We stress that more simulations would be needed to fully map out the dependence of the fit parameters over the 2D (σ_0 , L/ρ_{e0}) parameter space, and also to further quantify the random variation of repeated simulation runs with the same initial system parameters. We leave these further simulations required for more comprehensive results to future work.

6 TESTS OF DIFFUSIVE PARTICLE ACCELERATION

6.1 Test procedure

This Section details our procedure to measure the FP coefficients and their subsequent use in testing whether an energy diffusion–advection equation (2) is appropriate to model particle acceleration in these turbulence simulations. This was found to be true for a single simulation in our previous study (Wong et al. 2020), and we improve on it here with more comprehensive tests on a larger parameter space.

The tracked particles described in Section A are integral to this analysis. First, we 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, as described in Section A. Then, for each simulation, we select approximately 100 evenly distributed time instants, and at each instant, divide particles into energy bins with edges at 10 percent intervals, at $\gamma = 10(1.1^0, 1.1^1, 1.1^2, \dots)$. We 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, we compute the following quantities as functions of elapsed time $\Delta t \equiv t - t_0$: the energy standard deviation $\delta\gamma_{\text{rms}}$, energy variance $\overline{\delta\gamma^2} \equiv \overline{\delta\gamma_{\text{rms}}^2}$, 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, Friedrich & Peinke 1998, Friedrich, Siegert & Peinke 1999). 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: $\lim_{\Delta t \rightarrow 0} \overline{\delta\gamma^2} \propto \Delta t$. Accordingly, a primary conraindication to FP model would be if the energy trajectories exhibit anomalous diffusion instead, e.g. where the energy histories are dominated by rare but large jumps such that the energy variance intrinsically scales non-linearly with elapsed time. This has been suggested to be the case in plasma turbulence, for instance, in Isliker, Vlahos & Constantinescu (2017a) and Isliker et al. (2017b). 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 particle energy, this can cause a non-linear 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; we will give some concrete examples shortly, see also Friedrich et al. 1999). Therefore, we take the approach of a consistency check: we measure the FP coefficients assuming standard

Brownian diffusion, and then check whether a numerical solution of the FP equation using the measured coefficients reproduces the PIC simulation's particle-energy distribution evolution both globally and for individual bins (Section 6.3). If the agreement is good, we can be confident that even 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 of primary interest).

To measure the FP coefficients from the bin moments, we use the following idealized equations for the latter's time evolution, which are obtained from the FP equation (2) with a $\delta(\gamma - \gamma_0)$ initial condition and in the infinitesimal elapsed time approximation:

$$\Delta\bar{\gamma}(\gamma_0, \Delta t) = [\partial_\gamma D|_{\gamma_0} + A(\gamma_0)]\Delta t \equiv M(\gamma_0)\Delta t \quad (23)$$

$$\overline{\delta\gamma^2}(\gamma_0, \Delta t) = 2D(\gamma_0)\Delta t. \quad (24)$$

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. We also use linear fits to the moments to reduce noise, rather than measuring at a single point, and so, for the reason just mentioned, we 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 non-linear time dependence of the bin moments (Gottschall & Peinke 2008).

To limit the impact of the non-linear moments evolution from finite-time effects, we 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 $|\Delta\bar{\gamma}|/\gamma_0 = 0.1$. 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 well 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 arbitrary intercept (we show in Section 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's slope and hence at no particular value of $\delta\gamma_{\text{rms}}/\gamma_0$ or $|\Delta\bar{\gamma}|/\gamma_0$. However, there is little we 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, we 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.

We must also consider the FP coefficients' time dependence because these simulations are not true steady states. This is due to overall system heating described in Section 4. This heating occurs on the Alfvén time-scale, so we cap the fit time to $\Delta t < L/v_{A0}$. Finally, we enforce a minimum duration of $0.1L/c$, for the rare cases where the above limits result in a zero or negligible fit duration.

We use the above fit-range prescription to fit $\Delta\bar{\gamma}$ and $\overline{\delta\gamma^2}$ to linear functions of elapsed time, then read off D and M from the fitted slopes using (23) and (24). We repeat this for all bins to obtain D and M for every point in our 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, we obtain the advection coefficient using $A = M - \partial_\gamma D$, according to (23).

The rest of this section examines the bin energy moments (Section 6.2) and the consistency check results (Section 6.3), leaving detailed analysis of the FP coefficients to Sections 7 and 8.

6.2 Bin energy moments

Fig. 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 initialized at $v_{A0}t_0/L = 6$ and range from thermal ($\gamma_0/\gamma_{\text{pk}} \approx 1$) to non-thermal ($\gamma_0/\gamma_{\text{pk}} \gg 1$) energies. The plots show the linear fits that will be used to obtain D via (24), which seem reasonable considering the heavy late-time bias inherent in a logarithmic scale. They also show contours of constant $\delta\gamma_{\text{rms}}/\gamma_0$ and $|\Delta\bar{\gamma}|/\gamma_0$ (see caption for details), which give a sense of the relative evolution time-scales as a function of initial bin energy.

Generally, while $\delta\gamma_{\text{rms}}/\gamma_0 \ll 1$, $\overline{\delta\gamma^2}$ increases roughly linearly with Δt , suggestive of classical diffusion. For low initial magnetizations $\sigma_0 \lesssim 3/32$, there is substantial energy subdiffusion at thermal energies $\gamma_0 \lesssim \gamma_{\text{pk}}$, with power-law indices ~ 0.7 – 0.9 depending on σ_0 , which tends to become pronounced after intermediate times when $\delta\gamma_{\text{rms}}/\gamma_0 > 0.3$. Section 6.3 shows that FP evolution reproduces this. There is also superdiffusion at late times, once $\delta\gamma_{\text{rms}} \sim \gamma_0$, which is particularly apparent at non-thermal energies and at higher σ_0 . As mentioned in Section 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 $\delta\gamma_{\text{rms}}/\gamma_0$. A vertical contour implies, through (24), that $D \propto \gamma^2$; there are such near-vertical segments at high energies in all simulations. Similarly, contours slanting from lower left to upper right, or from lower right to upper left, mean, respectively, a shallower or steeper scaling than γ^2 .

Fig. 10 shows the change in mean bin energy as a function of elapsed time, $\Delta\bar{\gamma}(\Delta t)$. The traces are generally linear. The particle bin with minimum $\Delta\bar{\gamma}(\Delta t)$ slope for each given simulation (including negative slope as in the case of $\sigma_0 = 3/128$) tends to occur at central energies $\gamma_0/\gamma_{\text{pk}} \approx 2$ – 4 rather than the lowest energies. For the lowest $\sigma_0 = 3/128$ simulation, the $\gamma_0/\gamma_{\text{pk}} \leq 1/2$ bins have flattening curves. This occurs at long times and is likely due to coefficient energy dependence; we will check in detail for the $\gamma_0/\gamma_{\text{pk}} = 1/2$ bin in the next subsection.

6.3 Consistency checks

In this subsection, we test whether 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.

We choose a bin which shows subdiffusion, corresponding to the $\gamma_0/\gamma_{\text{pk}} \approx 1/2$ line in the $\sigma_0 = 3/128$ panel of Fig. 9, and set the initial condition to a matching narrow rectangular distribution. Fig. 11 shows the subsequent energy distribution evolution along with the nearby FP coefficient values, while Fig. 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. We note that the coefficients are remarkably stable over this time period, as is consistent with the slow dynamics of low-magnetization simulations, and so coefficient time dependence is likely unimportant.

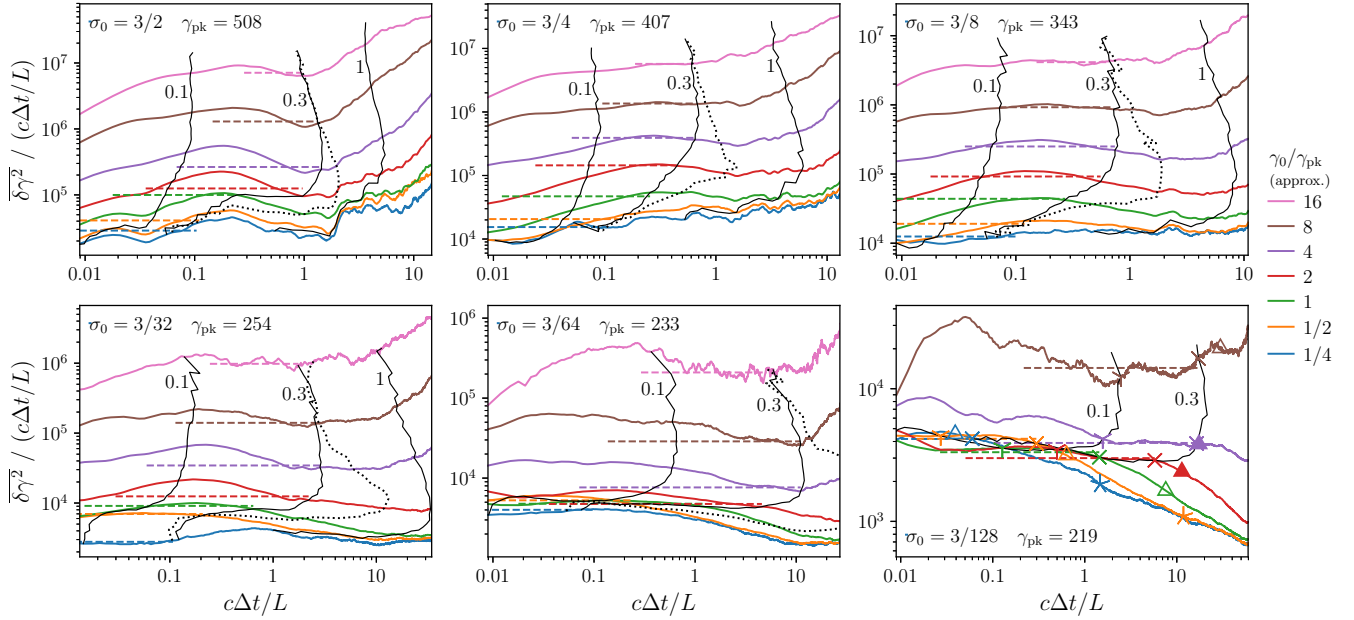


Figure 9. Bin energy variance $\overline{\delta\gamma^2}(\Delta t)$, compensated by $c\Delta t/L$, where each panel corresponds to a simulation with different σ_0 (indicated on the top left of the panel). All simulations have $N = 768$ and the bins are initialized at $v_{A0}t_0/L = 6$. Coloured dashed lines show the fits we use to obtain D ; the line extent shows the fit range. Black solid lines are contours of constant $\delta\gamma_{\text{rms}}/\gamma_0$ annotated by the corresponding $\delta\gamma_{\text{rms}}/\gamma_0$ 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 we 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 $\delta\gamma_{\text{rms}}/\gamma_0 = \{0.1, 0.3, 1\}$ respectively.

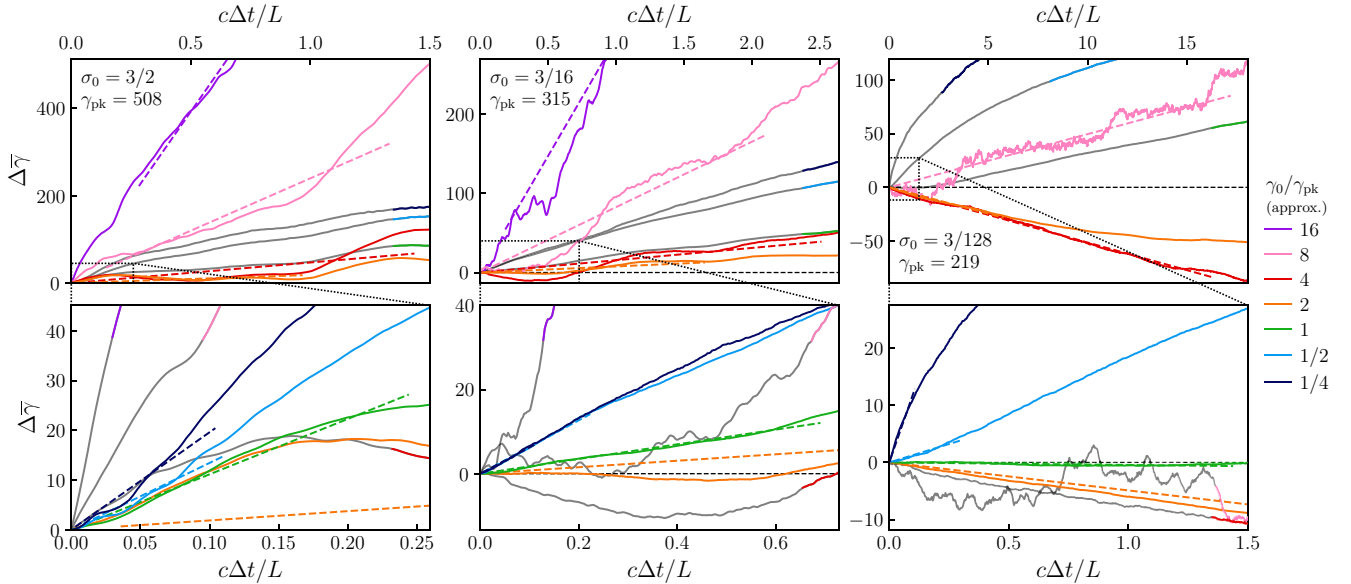


Figure 10. Change in bin mean energy $\Delta\overline{\gamma}(\Delta t)$ for $\sigma_0 \in \{3/2, 3/16, 3/128\}$ simulations (in order from left to right column), with $N = 768$ and bins initialized at $v_{A0}t_0/L = 6$. The top row is best for viewing the lines for the high-energy bins while the bottom row presents the same data but zooms in on the black dotted box in the lower left corner, which is better suited for viewing the low-energy bins. 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.

To get a better idea of what aspect of the FP coefficients' energy dependence results in this subdiffusion, we compare to the Ornstein–Uhlenbeck process mentioned in Section 6.1. This should be a reasonable model because in the vicinity of this relatively low-energy 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 $\Delta t \sim |\theta|^{-1}$. At the chosen bin centre energy, $\theta \approx -0.3c/L$, giving a non-linear time-scale of $\sim 2L/c$, which is a reasonable match to when the variance in Fig. 12 becomes significantly non-linear, considering that this estimate heavily simplifies the FP coefficients' energy dependence. Furthermore, Fig. 9 shows that, up to at least $\gamma_0/\gamma_{\text{pk}} = 2$, the subdiffusive departure becomes later as initial bin




Figure 11. Histogram of single-bin tracked particle energy (black lines in bottom panel) for four elapsed times $\Delta t \in \{0.0, 2.8, 8.5, 19.8\}L/c$, compared to the evolution of a matching initial narrow distribution using the FP equation with measured coefficients (solid lines in bottom panel). 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 (top and middle panel), with representative power-law fits (black lines): the negative slope of A produces effective subdiffusion. The bin is taken at $v_{A0}t_0/L = 6$ and has centre energy $\gamma_0/\gamma_{pk} \approx 1/2$. The simulation has $N = 768$ and $\sigma_0 = 3/128$. The bin distribution function is normalized such that it integrates to 1.




Figure 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 (black solid lines) and those obtained by evolving a matching initial narrow distribution using the FP equation with measured coefficients (blue dashed lines). For large elapsed time, the variance evolves with a subdiffusive scaling $\overline{\delta\gamma^2} \propto \Delta t^{2/3}$ (dotted line). The mean change does not have a definite power-law relationship with time ($\propto \Delta t$ dotted line shown for reference). The bin is taken at $v_{A0}t_0/L = 6$ and has centre energy $\gamma_0/\gamma_{pk} \approx 1/2$. The simulation has $N = 768$ and $\sigma_0 = 3/128$.

energy increases, which is consistent with Fig. 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 .

We 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. Fig. 13 presents a comparison of the numerical FP and PIC 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. We 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 non-thermal distribution can be completely modelled by the FP equation (2).

It is worth noting that, without the advection coefficient, the FP-based distributions have a substantially altered shape. They still exhibit an evolving power-law tail at high energies, but with greater overall energy injection and generally harder power laws than the PIC results. As was stated in Wong et al. (2020), but here shown for a range of σ_0 , the advection coefficient therefore plays an important role in maintaining the shape of the particle energy distribution.

7 DIFFUSION COEFFICIENT

Our previous study, Wong et al. (2020), measured the energy diffusion coefficient D in a single large simulation with moderate initial magnetization, finding that D scales as γ^2 in the non-thermal region and somewhat more shallowly at lower energies. This section characterizes $D(\gamma, t)$ in simulations with different parameters; we examine the time and energy dependence of D , and to what extent the aforementioned scaling characteristics persist. Finding universal γ^2 scaling in the non-thermal range, we extract the γ^2 pre-factor D_0 and compare it to NTPA theory predictions.

To investigate the effects of σ_0 (Section 7.1) and L/ρ_{e0} (Section 7.2), we employ the simulations from the $N = 768$ σ_0 -scan and the system-size scan. We separately describe the convergence of $D(\gamma)$ with respect to N_{ppc} (using the smaller $N = 384$ scan) in Appendix B. Please refer to Table 1 and surrounding text for the complete parameter sets.

We can only 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, we discard results from bins with less than 10 particles. Nevertheless, our results are primarily established by measurements at energies for which there are ample statistics, with thousands of tracked particles per bin.

7.1 Magnetization dependence

This subsection uses the $N = 768$ σ_0 -scan to analyse the dependence of $D(\gamma, t)$ on both the initial magnetization σ_0 and the instantaneous magnetization $\sigma(t)$. Fig. 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_{rms}/mc^2 . The length-scales are:

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




Figure 13. Top row: global particle energy distribution $f(\gamma)$ for the PIC simulations at four times (black lines), compared to the numerical solution of the FP equation (solid coloured lines) using measured D and A . The FP evolution was initialized with the initial thermal distribution. Simulations with varying $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ are shown in the different columns, with excellent agreement between FP and PIC over a broad range of energy in most cases. Bottom row: same as top row except with $A = 0$, which results in clear discrepancies between FP and PIC.

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

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

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

We observe that $D(\gamma)$ can be divided into two segments by the beginning of the $f(\gamma)$ non-thermal power law. In the non-thermal power law, we find a universal scaling $D \propto \gamma^2$ for all σ_0 (flattening slightly at highest energies, perhaps due to encroaching on γ_{\max}). The beginning of this $D \propto \gamma^2$ appears to coincide with the start of the $f(\gamma)$ non-thermal power law, rather than other energy scales such as γ_{avg} , as evidenced by the $\sigma_0 = 3/128$ plot of Fig. 14, where γ_{avg} occurs at a substantially lower energy than the start of the $f(\gamma)$ non-thermal power law. The low-energy scaling of $D(\gamma)$ 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 non-thermal 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 $D(\gamma)$. We leave such work to future studies.

Putting aside the low-energy behaviour, we now focus on the pre-factor of the non-thermal scaling $D_0 \equiv D(\gamma)/\gamma^2$, its time evolution,

and its dependence on σ_0 and $\sigma(t)$. The value of D_0 and its relation to other system parameters is important for testing various particle-acceleration theories, including the ones described in Section 2.2. In particular, standard second-order Fermi acceleration theories (see Section 2.1) yield $D_0 \propto u_A^2$, where $u_A \equiv v_A(1 - v_A^2/c^2)^{-1/2}$, which becomes $D_0 \propto v_A^2 \propto \sigma$ in the non-relativistic turbulence limit $v_A \ll c$. However, a more sophisticated modern analysis by Demidem et al. (2020), also based on quasi-linear theory but accounting for resonance broadening of Alfvén modes, yields $D_0 \propto v_A^3 \sim \sigma^{3/2}$ for Alfvénic turbulent cascades, while still retaining the conventional $D_0 \propto v_A^2 \propto \sigma$ scaling prediction for compressional fast and slow magnetosonic cascades. Furthermore, since our simulations are inhomogeneous and time-dependent, $v_A/c \equiv [\sigma/(\sigma + 1)]^{1/2}$ is only a characteristic definition; we do not know a priori whether to emplace σ_0 or $\sigma(t)$ here for the purpose of comparison with Fermi acceleration models. A careful analysis of first-principles PIC simulation results can help resolve this important issue. Finally, D_0 is also important because D_0^{-1} has a simple intuitive interpretation as the diffusive acceleration time for a particle in the non-thermal 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 23).

We 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 $D(\gamma, t)$. Fig. 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, $\delta B \sim B_0$ is by choice). Thereafter, as the plasma heats over time in our simulations, the relativistic enthalpy density increases, and this causes the magnetization (being inversely proportional to the enthalpy) to decline. We refer to section 4.2 of Zhdarkin et al. (2018a) for further details about $\sigma(t)$ time

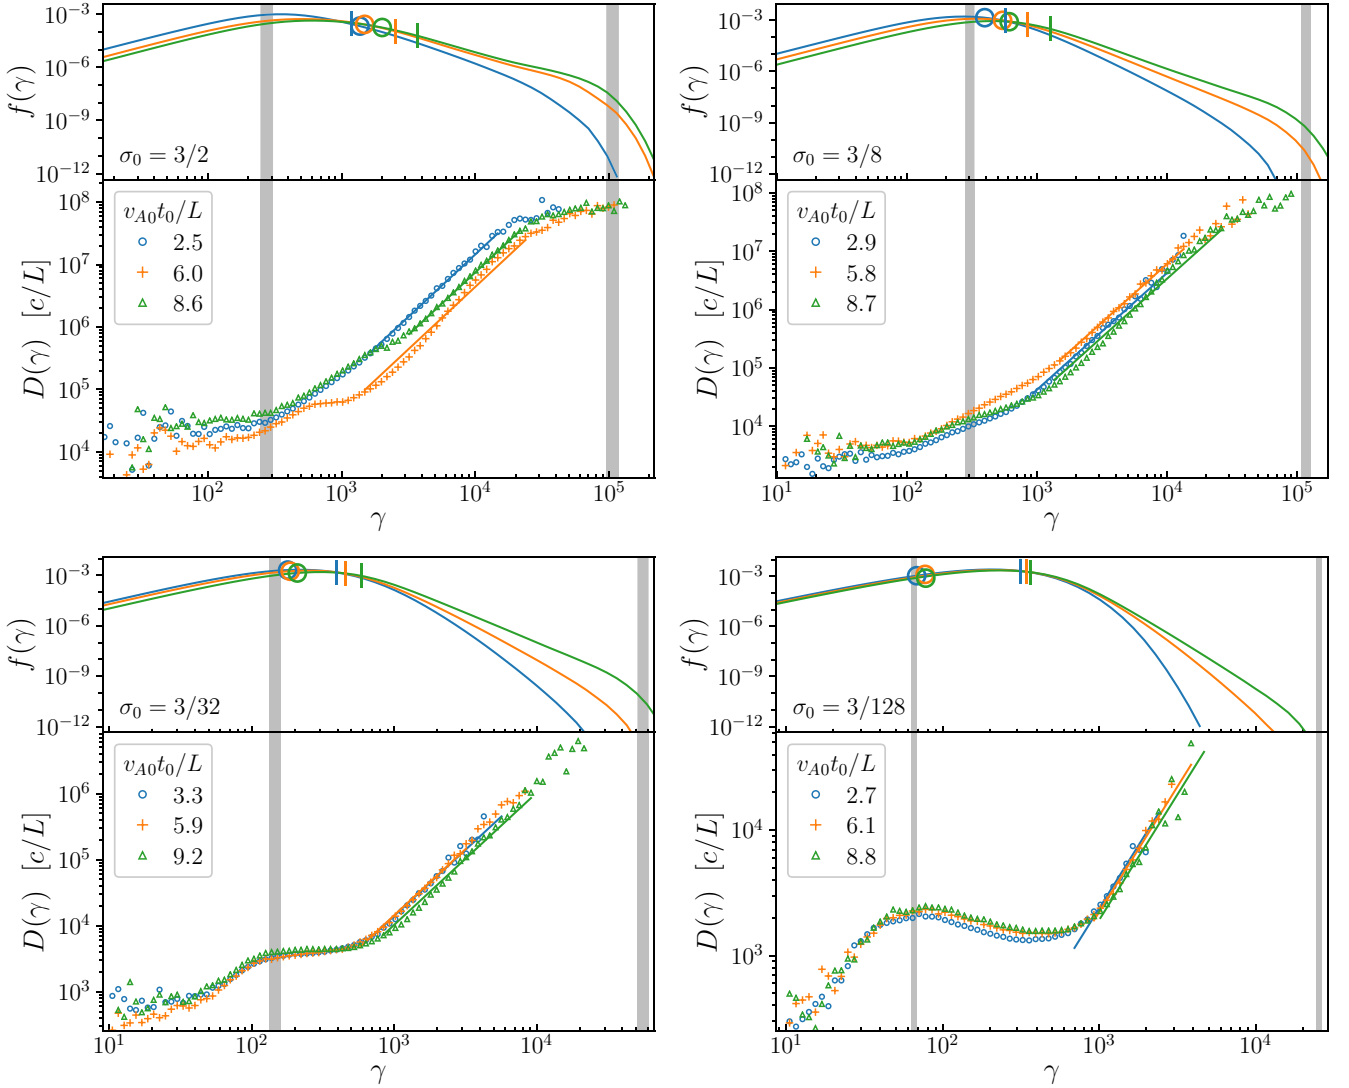


Figure 14. The diffusion coefficient D (lower subpanels) and particle energy distribution f (upper subpanels) as a function of energy γ for $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ (as labelled in each panel) and three different times (indicated in legend), displaying the relevant plasma and simulation length scales (converted to particle energies via the Larmor radius) as well as the $D \sim \gamma^2$ fits to the non-thermal 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 ρ_e (small vertical bars); and the skin-depth d_e (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 γ_{\max} (vertical grey strip at higher energy). For these respective quantities, the strips cover the 10th to 90th percentile values recorded between the times corresponding to the first and last plotted t_0 instance in each plot.

dynamics. The early-time 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 of this σ_0 -dependent decline of σ , simulations beginning at $\sigma_0 \geq 3/8$ reach similar values of 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, so 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 $\sigma(t)$, we obtain $D_0(\sigma)$, and plot this in Fig. 16. We exclude points earlier than two Alfvén crossing times, before which turbulence has not fully developed. We also exclude points after the normalized 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 our 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).




Figure 15. (a) The diffusion coefficient pre-factor $D_0(t)$, and (b) the instantaneous magnetization $\sigma(t)$, 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 .




Figure 16. The diffusion coefficient pre-factor D_0 versus instantaneous magnetization $\sigma(t)$ from all simulations in the $N = 768$ σ_0 -scan, along with lines of $D_0 \propto v_A^3$ (solid), $D_0 \propto \sigma^{3/2}$ (dashed), and $D_0 \propto \sigma$ (dotted). The $D_0(\sigma)$ scaling is roughly $D_0 \propto \sigma^{3/2} \sim (v_A/c)^3$ at low $\sigma \lesssim 0.5$ and flattens somewhat at higher σ , to around $D_0 \propto \sigma$.

Here, we take advantage of the time dependence resulting from our simulation setup. As the time evolution of $\sigma(t)$ from simulations with different σ_0 overlap, we ask whether D_0 also coincides at these times. Indeed, we 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.

We find the $D_0(\sigma)$ scaling to be consistent with $D_0 \propto \sigma^{3/2}$ at low $\sigma \lesssim 0.5$ (i.e. in the weakly relativistic regime), with the dependence flattening somewhat at high $\sigma \gtrsim 0.5$, to about $D_0 \propto \sigma$. Fig. 16 illustrates this behaviour with two fits: $D_0 = 0.43(v_A/c)^3 c/L = 0.43[\sigma/(1+\sigma)]^{3/2} c/L$ and $D_0 = 0.32\sigma^{3/2} c/L$; we see that they agree with the data very well at low $\sigma \lesssim 0.5$, and provide a reasonably good agreement over the entire σ -range, although the $D_0 \propto \sigma^{3/2}$ scaling deviates noticeably at high σ . Thus, our numerical results provide clear evidence against the standard Fermi-acceleration theory's prediction of a linear $D_0 \propto \sigma$ scaling for $\sigma \lesssim 0.5$, although this scaling may still be applicable for $\sigma \gtrsim 0.5$. Our preferred scaling $D_0 \propto v_A^3$, which becomes $D_0 \propto \sigma^{3/2}$ in the low- σ (non-relativistic turbulence) limit, is consistent with the theoretical prediction by Demidem et al. (2020) for Alfvénic turbulence, in which the extra power of v_A (compared to the standard Fermi-acceleration theory) was attributed to resonance broadening. Thus, our results provide direct numerical support for this theory. It is interesting to note, however, that the $D_0 \propto v_A^3$ scaling was predicted by Demidem et al. (2020) only for Alfvénic turbulence, which is essentially incompressible, whereas for the case of turbulence dominated by compressional fast and slow magnetosonic waves that study obtained a traditional $D_0 \propto v_A^2$ scaling. We speculate that perhaps the reason why the $D_0 \propto \sigma^{3/2}$ scaling is observed only at low σ in our simulations is that only in this non-relativistic regime the turbulence is predominantly subsonic ($v_{\text{rms}} \sim v_A \ll c_s \sim c$, where v_{rms} is the rms fluid bulk velocity and c_s is the speed of sound) and hence essentially incompressible, mediated mostly by the Alfvénic cascade. In contrast, at higher σ turbulent motions are relativistic ($v_{\text{rms}} \sim v_A \sim c_s \sim c$) and inevitably compressive, with a greater (perhaps even dominant) role played by fast and slow magnetosonic waves, which leads to the restoration of the traditional $D_0 \propto \sigma$ scaling. In principle, more simulations at higher σ would help test this hypothesis and elucidate the high- σ scaling; in practice, however, the high- σ regime (with strong fluctuations, $\delta B_{\text{rms}}/B_0 \sim 1$ and $v_{\text{rms}} \sim c$) is difficult to maintain for any substantial period of time, because the magnetic energy would rapidly convert to thermal energy, quickly decreasing σ before the turbulence can fully develop (Zhdankin et al. 2018a).

Note that the normalization $D_0(L/c)/(v_A/c)^3 \approx 0.43$ of the scaling fit in Fig. 16 can be compared to the expectations from Demidem et al. (2020) (their equation 38), which evaluates (up to a correction that is logarithmic in scale separation) to ≈ 0.6 for a purely Alfvénic cascade with amplitude $\delta B_{\text{rms}}/B_0 \approx 1$. Thus, the results are in reasonable agreement. A precise theoretical comparison would require measuring the partition between different modes (Alfvén, fast, and slow) and taking into account corrections due to finite system size.

7.2 System-size dependence

Fig. 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 may 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 non-thermal range between the simulations of different L/ρ_{e0} is good, showing uniform $D \propto \gamma^2$ scaling. This indicates that $D(\gamma)$




Figure 17. Diffusion coefficient $D(\gamma)$ at $t_0 = 10L/c$ for each simulation in the system-size scan, where $\sigma_0 = 3/8$. Reference power-law scalings are shown in black dashed lines, and dashed vertical lines indicate γ_{\max} for each case.

in the general parameter space explored in this paper is reasonably well converged with respect to L/ρ_{e0} for the non-thermal region, providing confidence in our results. With regards to the extent of the $D \propto \gamma^2$ scaling, we would naively anticipate it to be proportional to the inertial range, and hence to the system size L/ρ_{e0} . However, Fig. 17 shows that the energy range for which there are good D measurements is almost unchanged with L/ρ_{e0} . This is because the highest energy for which D can be measured reliably in our study is limited by tracked particle statistics due to the $f \propto \gamma^{-\alpha}$ falloff rather than by the extent of the inertial range.

The low-energy behaviour, roughly below $\gamma_{\text{pk}} \sim 10^3$, varies more substantially between the different L/ρ_{e0} , but is still broadly consistent with having a shallower power-law scaling than in the non-thermal range. There is also a weak trend of higher- L/ρ_{e0} values corresponding to larger low-energy diffusion coefficients. The location of the spectral break between the two power laws also varies somewhat and does not show a particular pattern, but is roughly at $\gamma \sim 10^3$, a little bit above γ_{pk} . We leave further investigation of the low-energy range to future work.

Fig. 18 presents $D_0(t)$ and $\sigma(t)$ for the system-size scan, showing that L/ρ_{e0} has only a weak effect on the respective quantities, with lines from different simulations largely overlapping amidst the fluctuations. However, it is hard to tell if this is true due to time fluctuations in the traces. 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 $D_0(t)$ or $\sigma(t)$ for simulations between which only the system size L varies. We leave such an investigation for future work. As in Section 7.1, the $D_0(t)$ and $\sigma(t)$ are qualitatively similar to each other, including the initial increase, subsequent decay, and dependences on σ_0 and N .

7.3 Summary

The diffusion coefficient as a function of particle energy $D(\gamma)$ has broadly similar behaviour in the non-thermal range over simulations with different σ_0 and L/ρ_{e0} . The high-energy scaling of $D \propto \gamma^2$ is universal, while the low-energy scaling varies, with the strongest qualitative dependence on σ_0 . The pre-factor D_0 of the $D \propto \gamma^2$ scaling depends mostly on instantaneous magnetization $\sigma(t)$ rather than σ_0 . For $\sigma \lesssim 0.5$, which is most of our explored range, the




Figure 18. Time evolution of (a) diffusion coefficient prefactor $D_0(t)$ and (b) magnetization $\sigma(t)$, for the system-size scan, where the simulations have the same $\sigma_0 = 3/8$ but different $L/\rho_{e0} \in \{1024, 683, 512, 341, 256, 171\}$. There is no clear systematic variation of these quantities with system size.

D_0 scaling is consistent with $D_0 \propto v_A^3$ or $D_0 \propto \sigma^{3/2}$, in contrast with the common theoretical prediction of $D_0 \propto \sigma$. The convergence of D in the non-thermal range with increasing system size L/ρ_{e0} is excellent, which gives us confidence in the rest of our results.

8 ADVECTION COEFFICIENT

8.1 Numerical measurement of advection coefficient

In Wong et al. (2020), we observed that the particle energy-advection coefficient A was positive at $\gamma < \gamma_{\text{pk}}$ and negative at $\gamma > \gamma_{\text{pk}}$. However, the data at highest non-thermal energies in that study seemed too noisy to analyse in detail. Now, our model in Section 2.2 suggests a concrete energy dependence $A = A_0 \gamma \log(\gamma/\gamma_A^*)$ (9) in the non-thermal range, and also predicts relationships between the FP coefficient model parameters and the power-law index. The present section examines the advection coefficient in the context of that model.

First, we 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 $\gamma_0 = 10^4$. Fig. 19 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 non-thermal 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 spectral 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 at early times in the $\sigma_0 = 3/8$ and $\sigma_0 = 3/32$ simulations. However, in these cases

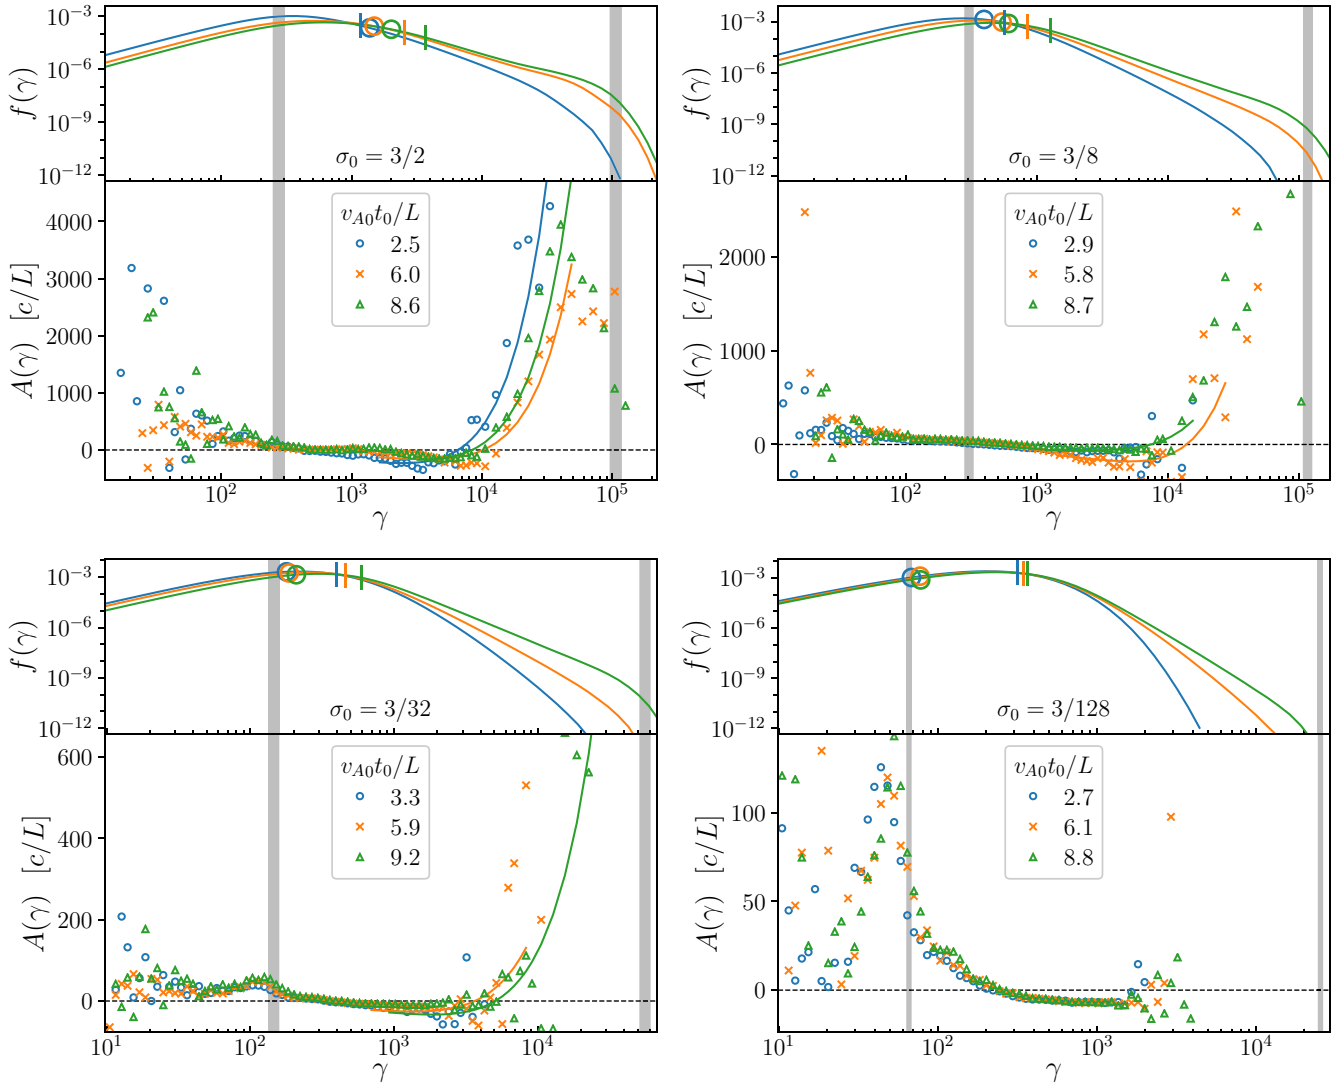


Figure 19. Lower subpanels show the energy dependence of the advection coefficient $A(\gamma, t)$ (markers) at three different times (shown in the legend) for PIC simulations with $\sigma_0 \in \{3/2, 3/8, 3/32, 3/128\}$ (as labelled in each panel), with fits of $A = A_0 \gamma \log(\gamma/\gamma_A^*)$ (solid lines). Noise at high energies has been reduced by combining every two bins after $\gamma = 10^4$ and averaging A over a time $0.5L/v_{A0}$. Upper subpanels show the particle energy distribution at corresponding times, for reference. See the Fig. 14 caption for a description of the energy-scale markings.

the lower-energy data still follows the aforementioned pattern. The first zero crossing, from positive to negative near γ_{pk} , was the initial observation of Wong et al. (2020): that the advection coefficient tends to gather particles together in energy space. This applies to thermal energies and so out of the Section 2.2 model’s validity range. The second zero crossing, from negative to positive, is a key feature of the Section 2.2 model that is borne out in these simulation results. This crossing occurs in the power-law range, but we note that the model does not predict nor require anything about its location. This sign change is unlikely to be an artefact of the particle pileup at the system size limit γ_{max} because it occurs at energies roughly an order of magnitude lower than γ_{max} (which is depicted as the rightmost grey bar in Fig. 19).

Fig. 19 also shows fits of $A(\gamma)$ to (9), limited to the non-thermal range, for the cases with sufficient high-energy data as identified above. The fits are consistent with the measurements inside the

model’s validity range of the power-law tail. Deviations occur outside that range, near γ_{pk} and γ_{max} , which is expected.

We could now examine the trends and tendencies of the A model equation (9) parameters A_0 and γ_A^* , and test their relationships (10)–(13) with the power-law parameters α and K . However, we will first switch to the related variable M , the energy-dependent average acceleration rate, defined in (23). While the FP equation (2) is in terms of A , it is more convenient to use M because, by (23), M comes from $\Delta\bar{\gamma}$ measurements only, while $A = M - \partial_\gamma D$ is contaminated by noise in $\partial_\gamma D$, with the γ -derivative particularly amplifying the noise.

Furthermore, we rewrite $M = M_0 \gamma \log(\gamma/\gamma_M^*)$ (equation 14) into

$$M/\gamma = M_0 \log \gamma - M_0 \log \gamma_M^*. \quad (25)$$

This now simply states that M/γ should be linear in $\log \gamma$, with the x -intercept at $\log \gamma_M^*$. This form is convenient for visual inspection,

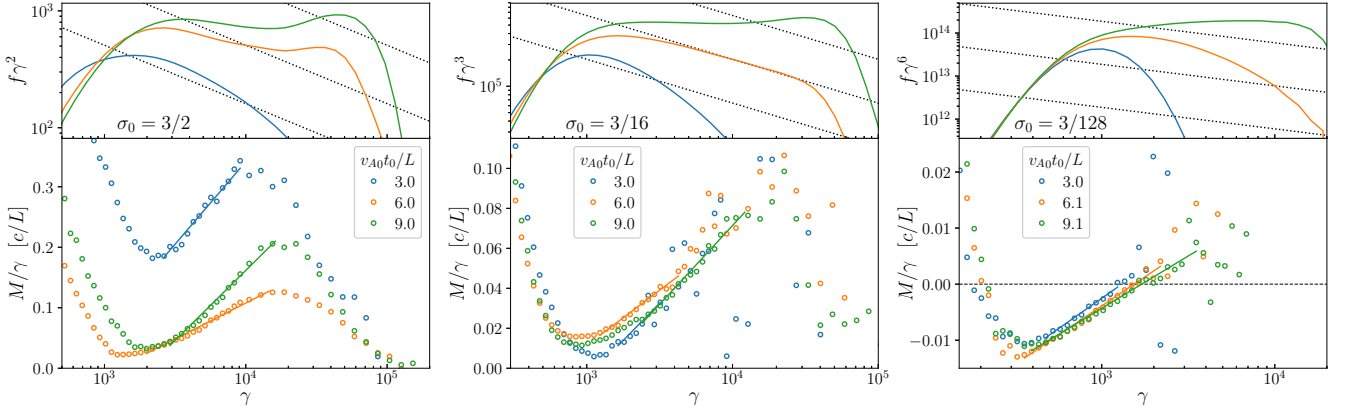


Figure 20. Particle energy distribution (upper subpanels) and M/γ (lower subpanels) as functions of energy γ at three times $t \approx \{3, 6, 9\}L/v_{A0}$ for PIC simulations with different initial magnetization $\sigma_0 \in \{3/2, 3/16, 3/128\}$ (from left to right). Fits by $\log \gamma$ are shown in the lower subpanels (solid lines). Upper subpanels are compensated by powers of energy such that the last particle energy distribution shown is roughly horizontal. Dotted lines are power laws with one-half lower index than the compensated value in each case, for reference.

and also for fitting because linear regression is robust and well understood.

Fig. 20 shows M/γ versus γ on linear-log scales, so that a linear equation in $\log \gamma$ such as (25) appears as a straight line, presented for a selection of simulations from the $N = 768$ σ_0 -scan at measurement times $t_0 \approx \{3, 6, 9\}L/v_{A0}$. We also restrict the plot view to the high-energy region above γ_{pk} . In all cases, M/γ has a remarkably good linear segment spanning about a decade of energies in the power-law region. This is followed by a decline in M/γ at energies near the system-size limit, adulterated by varying amounts of noise: less for high σ_0 and more for low σ_0 . Surprisingly, even the lowest $\sigma_0 = 3/128$ simulation plot shows a high-quality straight-line region despite the paucity of high-energy particles (in contrast to the inconclusive nature of the corresponding region in Fig. 19). The linear M/γ region also appears clearly even before the $f(\gamma)$ power law has fully formed, as seen at the early $t_0 \approx 3L/v_{A0}$ time instances. The presence of this high-quality linear segment over a wide range of initial magnetization values gives strong qualitative support to the Section 2.2 model.

We now examine the coefficients M_0 and γ_M^* , which we extract via (25) from the linear fits (as shown in Fig. 20) to the straight-line region in M/γ . Fig. 21 displays $M_0(t)$ and $\gamma_M^*(t)$ for each simulation in the $N = 768$ σ_0 -scan. We find that M_0 is higher for higher- σ_0 simulations, and tends to decrease over time. As γ_M^* is too noisy to show all the points, we use LOESS (Cleveland 1981) to generate a smoothed curve, with shaded intervals to approximately indicate the point scatter. This is a visual aid only with no intended statistical meaning. The γ_M^* values are around $\gamma = 10^3$, which is typically in the early part of the non-thermal section, and γ_M^* tends to increase over time, as one might expect from overall simulation heating.

Section 2.2 does not obtain any specific 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 -pre-factor) D_0 with instantaneous σ (see Section 7.1). Nevertheless, by analogy to the diffusion coefficient, we plot M_0 against σ in Fig. 22. The scaling is consistent with $M_0 \propto \sigma$, or slightly shallower. The increasing trend matches the basic qualitative expectation from $M_0 = A_0 = -d \log(\alpha - 1)/dt$ (equation 10): as we expect α to evolve faster at higher σ (in order to reach harder power laws), M_0 should increase with σ . We leave more detailed interpretation to future investigation.




Figure 21. (a) Connected circles display M_0 obtained by fitting $M(\gamma)$ to (25), for varying σ_0 indicated in the legend. These show that M_0 is higher for higher σ_0 simulations, and also that it tends to decrease over time. (b) Solid lines are smoothed γ_M^* values (obtained by fitting $M(\gamma)$ to equation 25) while the adjacent shaded regions indicate approximately the spread of unsmoothed measurements. We display the original unsmoothed points for one simulation, to help visualize how the shaded areas relate to the point scatter. We employ LOESS (Cleveland 1981) to generate the smoothed curves and confidence intervals, but this is a visual aid only, without any statistical purpose. The γ_M^* values are around $\gamma = 10^3$, which is typically in the early part of the non-thermal section, and γ_M^* tends to increase over time, as one might expect from overall simulation heating.

8.2 Comparison to analytical model of power-law evolution

Section 2.2 derived relationships between the FP coefficient parameters D_0 , M_0 , and γ_M^* and the power-law parameters K and




Figure 22. The scaling of M_0 versus instantaneous magnetization σ , for PIC simulations with varying initial magnetization σ_0 (indicated in legend), showing a positive correlation independent of σ_0 . The relationship between M_0 and σ is consistent with a linear scaling (particularly for $\sigma_0 \geq 3/64$; dotted line) or a slightly shallower power law ($\sigma^{0.86}$, dashed line).

α . These relations are the differential equations (10) and (15) and their equivalent integral forms (16) and (17). In this subsection, we test the consistency of these formulae with the acquired data. While equivalent, we prefer and use the integral forms. This is because A_0 and γ^* can both fluctuate 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. That is, differential equation comparisons may mislead in terms of how much FP coefficients' time fluctuations affect the power-law index.

The method for obtaining D_0 was described in Section 7; for M_0 and γ_M^* , above; and for α , in Section 5. The procedure for finding the power-law pre-factor K derives from extending the α measurement process. That process identifies, at each instance of time, a preferred point on the local power-law index curve from which we obtain α . To obtain K , we substitute that point's (γ, f) coordinates into $f = K\gamma^{-\alpha}$ (5) to yield $K = f\gamma^\alpha$. We note that K suffers amplified noise from α fluctuations due to the long lever arm in its definition as $f(\gamma = 1)$, and future work may benefit from defining it relative to a more central point.

Fig. 23 compares α and K measured from the particle energy distribution with those obtained through the time integrals (16) and (17) which are functions of the FP coefficient model parameters D_0 , A_0 , and γ_A^* . This figure displays vertical 'error bars' for α indicating a range of nearby values due to the energy dependence of the local power-law index $\alpha_{\text{loc}}(\gamma, t) \equiv -\partial \log f / \partial \log \gamma$. These α bars are the same as those in Fig. 7 and their calculation method is specified in Section 5.2. The K bars follow the same simple inversion procedure described just above, except using a range of α corresponding to the α bars. We 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 5), which is essentially the 'inflection time' discussed by Zhankin et al. (2018b). Substantially before this time, the power law is not fully formed, and so we focus mainly on the vicinity of this moment and later.

We find in general that the integral-derived curves are consistently steeper than the α and K points measured from the power-law slope. This might be explainable by examining the 'error bars'. After the




Figure 23. Integral formulas computing $\alpha(t)$ and $K(t)$ from the fitted FP coefficient model parameters D_0 , M_0 , and γ_M^* (solid curves) match reasonably well to α and K measured directly from the particle energy distribution (points with vertical bars) at varying times. The vertical bars are approximate possible ranges for each parameter due to the particle energy distribution local power-law slope having γ dependence; see text for details.

inflection point, these are consistently unidirectional because the Section 5.1 α -measurement method chooses a local extreme of α_{loc} and so nearby values are all to one side. The bar direction is towards lower α magnitude, corresponding to a flatter power law; this is because including the high-energy pileup in any kind of averaging fit would decrease the apparent overall power-law index. Hence, we observe that beyond the constructed intersection point, the Fig. 23 integral curves all pass on the bar side of the α and K points. This then implies that the model of Section 2.2 is sensitive in part to the high-energy pileup, and has an effective power-law index somewhat smaller in magnitude than that obtained from the Section 5.1 fitting. This could be further examined in future work by using a power-law fitting method that trims the high-energy pileup less aggressively.

These comparisons between the α and K (power-law based) and FP measurements stretch the data quality, with substantial uncertainty on both sides. The uncertainty for α is due to the limited extent of the power law, and for the FP coefficients, due to the limited number of high-energy tracked particles. Hence, these consistent results only indicate that the model is quantitatively plausible. It is nevertheless 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.

Overall, this section's results provide confidence in the basic validity of the Section 2.2 model. In addition to the reasonable quantitative agreement just discussed, the remarkably clear linear segment in M/γ versus $\log \gamma$ in Fig. 20 gives strong qualitative support. It is also encouraging that the assumptions underlying the Section 2.2 model are very simple.

9 CONCLUSIONS

In this paper, we investigated the evolution of non-thermal particle populations in first-principles PIC simulations of driven relativistic pair plasma turbulence through the lens of the FP energy diffusion–advection model. This was accomplished by examining the behaviour of particle energies by tracking large numbers of simulated particles, and measuring the energy diffusion and advection coefficients (D and A) as functions of relativistic particle energy and time. Our main conclusions are:

- (i) The power-law index α of the non-thermal particle energy distribution converges exponentially with time, as described by equation (18).
- (ii) Particle energies exhibit diffusive–advective behaviour with respect to time over broad ranges of system parameters.
- (iii) The energy diffusion coefficient $D(\gamma)$ generally scales as the square of particle energy γmc^2 (or momentum γmc) in the non-thermal range: $D(\gamma) = D_0 \gamma^2$.
- (iv) The coefficient D_0 depends primarily on the instantaneous magnetization σ rather than the initial magnetization, and this relationship is consistent with $D_0 \propto (v_A/c)^3 \sim \sigma^{3/2}$ at low $\sigma \lesssim 1$, in agreement with the Alfvénic case in (Demidem et al. 2020).
- (v) The energy advection coefficient $A(\gamma)$ plays a significant role, and we constructed an analytical model relating it to the evolution of the non-thermal power-law index $\alpha(t)$. This predicts a scaling of $A = A_0 \gamma \log(\gamma/\gamma_A^*)$, which is borne out by the simulations.
- (vi) When fed back into the FP equation, the diffusion and advection coefficients successfully reproduce the evolution of the particle energy distribution obtained from the PIC simulations for all particle energies and over the entire investigated parameter space. This confirms that the FP model provides an adequate description of particle acceleration.

This work opens various avenues for further exploration. The diffusion coefficient scaling is somewhat different from the expected scaling of $D_0 \propto \sigma$. This could be further investigated through simulation and analytical theory. The new scaling of $A = A_0 \gamma \log(\gamma/\gamma_A^*)$ has significant implications for the evolution of the non-thermal energy spectrum (as was shown in Fig. 13). This may be of interest as non-thermal particle acceleration models commonly assume that the momentum-space advection coefficient is zero (leading to $A = 2D/\gamma$). While we have focused on the behaviour of the FP coefficients in the non-thermal region, we also measured D and A at thermal and subthermal energies, finding qualitatively different behaviour. This low-energy range may influence non-thermal particle acceleration through the injection of intermediate-energy particles for stochastic acceleration, and may be highly sensitive to the intermittency of turbulence (e.g. Comisso & Sironi 2019; Vega, Boldyrev & Roytershteyn 2023; Davis, Comisso & Giannios 2024). In any case, further investigation may shed light on other physical phenomena relevant at those energies and scales.

This study of magnetized turbulent non-thermal particle acceleration enhances our understanding of widespread and long-studied plasma physical processes, and has profound implications for space, solar, and astrophysical systems. Apart from turbulence, other plausible mechanisms of non-thermal particle acceleration include magnetic reconnection and shocks. Our analysis methods may be adapted to kinetic simulations of these different processes in order to test theories of particle acceleration in those respective environments.

In future work, our methods may also be applied to model the asymptotic particle distribution function in PIC simulations of turbulence with strong radiative cooling (Zhdankin et al. 2020;

Zhdankin, Uzdensky & Kunz 2021; Grošelj et al. 2024) or escaping particles (Gorbunov, Grošelj & Bacchini 2025), in which a true statistical steady state is achieved. Radiative cooling (in the optically thin regime) would add a negative advection term to the Fokker–Planck equation, which may inhibit particles from being transported through the $D \propto \gamma^2$ range (preventing them from reaching γ_{\max}). Particle injection and escape would appear as sources and sinks in the Fokker–Planck model; inferring these terms from particle trajectories would require other methods to be developed. Since the properties of the turbulent fluctuations may be insensitive to radiative cooling (see, however, Grošelj et al. 2024) or particle escape (which influence kinetic rather than fluid properties), we anticipate that the core FP model described in this work will still be essentially applicable to these scenarios.

ACKNOWLEDGEMENTS

We are grateful to Martin Lemoine for encouraging and insightful discussions of turbulent particle acceleration. We also would like to thank the manuscript’s reviewer, Daniel Grošelj, for valuable comments that helped improve the paper. This work was primarily supported by U.S. National Science Foundation (NSF) grant AST 1806084, which is gratefully acknowledged. DAU, MCB, and GRW also acknowledge partial support from National Aeronautics and Space Administration (NASA) Astrophysics Theory Program grant 80NSSC22K0828. VZ acknowledges support from NSF grant PHY-2409316. This research was also supported in part by grant NSF PHY-2309135 to the Kavli Institute for Theoretical Physics (KITP), through the participation of VZ and DAU at the KITP programmes ‘Interconnections between the Physics of Plasmas and Self-gravitating Systems’ and ‘Relativistic Plasma Physics: From the Lab to the Cosmos’. This work used Stampede2 at the Texas Advanced Computer Center (TACC) through allocation TG-PHY160032 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) programme, which was supported by U.S. National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

DATA AVAILABILITY

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

REFERENCES

- Asano K., Takahara F., Kusunose M., Toma K., Kakuwa J., 2014, *ApJ*, 780, 64
- Becker P. A., Le T., Dermer C. D., 2006, *ApJ*, 647, 539
- Bian N. H., Emslie A. G., Stackhouse D. J., Kontar E. P., 2014, *ApJ*, 796, 142
- Blandford R., Eichler D., 1987, *Phys. Rep.*, 154, 1
- Blasi P., 2000, *ApJ*, 532, L9
- Bresci V., Lemoine M., Gremillet L., Comisso L., Sironi L., Demidem C., 2022, *Phys. Rev. D*, 106, 023028
- Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2013, *ApJ*, 770, 147
- Chandran B. D., 2000, *Phys. Rev. Lett.*, 85, 4656
- Cho J., Lazarian A., 2006, *ApJ*, 638, 811
- Cleveland W. S., 1981, *Am. Stat.*, 35, 54
- Comisso L., Sironi L., 2018, *Phys. Rev. Lett.*, 121, 255101
- Comisso L., Sironi L., 2019, *ApJ*, 886, 122
- Davis Z., Comisso L., Giannios D., 2024, *ApJ*, 964, 14
- Demidem C., Lemoine M., Casse F., 2020, *Phys. Rev. D*, 102, 023003
- Dmitruk P., Matthaeus W., Seenu N., Brown M. R., 2003, *ApJ*, 597, L81

- Dmitruk P., Mattheaus W., Seenu N., 2004, *ApJ*, 617, 667
- Fermi E., 1949, *Phys. Rev.*, 75, 1169
- Friedrich R., Siegert S., Peinke J., 1999, in Müller S. C., Parisi J., Zimmermann W., eds, *Transport and Structure*. Springer-Verlag, Berlin, p. 273
- Gorbanov E. A., Grošelj D., Bacchini F., 2025, *Phys. Rev. Lett.*, 135, 065201
- Gottschall J., Peinke J., 2008, *New J. Phys.*, 10, 083034
- Grošelj D., Hakobyan H., Beloborodov A. M., Sironi L., Philippov A., 2024, *Phys. Rev. Lett.*, 132, 085202
- Hankla A. M., Zhdankin V., Werner G. R., Uzdensky D. A., Begelman M. C., 2022, *MNRAS*, 509, 3826
- Islaker H., Vlahos L., Constantinescu D., 2017a, *Phys. Rev. Lett.*, 119, 045101
- Islaker H., Pisokas T., Vlahos L., Anastasiadis A., 2017b, *ApJ*, 849, 35
- Kimura S. S., Murase K., Toma K., 2015, *ApJ*, 806, 159
- Kimura S. S., Toma K., Suzuki T. K., Inutsuka S.-i., 2016, *ApJ*, 822, 88
- Kimura S. S., Tomida K., Murase K., 2019, *MNRAS*, 485, 163
- Kowal G., Dal Pino E. M. d. G., Lazarian A., 2012, *Phys. Rev. Lett.*, 108, 241102
- Kulsrud R. M., Ferrari A., 1971, *Ap&SS*, 12, 302
- Kundu S., Vaidya B., Mignone A., 2021, *ApJ*, 921, 74
- Lemoine M., 2019, *Phys. Rev. D*, 99, 083006
- Lemoine M., 2021, *Phys. Rev. D*, 104, 063020
- Lemoine M., 2022, *Phys. Rev. Lett.*, 129, 215101
- Lemoine M., 2025, *Phys. Rev. E*, 112, 015205
- Lemoine M., Malkov M. A., 2020, *MNRAS*, 499, 4972
- Lemoine M., Murase K., Rieger F., 2024, *Phys. Rev. D*, 109, 063006
- Liu W., Petrosian V., Mariska J. T., 2009, *ApJ*, 702, 1553
- Longair M. S., 2011, *High Energy Astrophysics*. Cambridge Univ. Press, Cambridge
- Lynn J. W., Quataert E., Chandran B. D., Parrish I. J., 2013, *ApJ*, 777, 128
- Lynn J. W., Quataert E., Chandran B. D., Parrish I. J., 2014, *ApJ*, 791, 71
- Medina-Torrejón T. E., Dal Pino E. M. d. G., Kadowaki L. H., Kowal G., Singh C. B., Mizuno Y., 2021, *ApJ*, 908, 193
- Melrose D., 1974, *Sol. Phys.*, 37, 253
- Meringolo C., Cruz-Osorio A., Rezzolla L., Servidio S., 2023, *ApJ*, 944, 122
- Mertsch P., Sarkar S., 2011, *Phys. Rev. Lett.*, 107, 091101
- Miller J. A., Guessoum N., Ramaty R., 1990, *ApJ*, 361, 701
- Nättilä J., Beloborodov A. M., 2022, *Phys. Rev. Lett.*, 128, 075101
- Nayakshin S., Melia F., 1998, *ApJS*, 114, 269
- Schlickeiser R., 1985, *A&A*, 143, 431
- Schlickeiser R., 1989, *ApJ*, 336, 243
- Siegert S., Friedrich R., Peinke J., 1998, *Phys. Lett. A*, 243, 275
- Skilling J., 1975, *MNRAS*, 172, 557
- Summers D., Ma C.-y., 2000, *J. Geophys. Res.*, 105, 2625
- Sun X., Bai X.-N., 2021, *MNRAS*, 506, 1128
- TenBarge J. M., Howes G. G., Dorland W., Hammett G. W., 2014, *Comput. Phys. Commun.*, 185, 578
- Vega C., Boldyrev S., Roytershteyn V., Medvedev M., 2022, *ApJ*, 924, L19
- Vega C., Boldyrev S., Roytershteyn V., 2023, *ApJ*, 949, 98
- Vega C., Boldyrev S., Roytershteyn V., 2024, *ApJ*, 971, 106
- Vlahos L., Islaker H., Lepreti F., 2004, *ApJ*, 608, 540
- Werner G. R., Uzdensky D. A., Begelman M. C., Cerutti B., Nalewajko K., 2018, *MNRAS*, 473, 4840
- Wong K., Zhdankin V., Uzdensky D. A., Werner G. R., Begelman M. C., 2020, *ApJ*, 893, L7
- Zhang J.-F., Xiang F.-Y., 2021, *ApJ*, 922, 209
- Zhdankin V., 2022, *J. Plasma Phys.*, 88, 175880303
- Zhdankin V., Werner G. R., Uzdensky D. A., Begelman M. C., 2017, *Phys. Rev. Lett.*, 118, 055103
- Zhdankin V., Uzdensky D. A., Werner G. R., Begelman M. C., 2018a, *MNRAS*, 474, 2514
- Zhdankin V., Uzdensky D. A., Werner G. R., Begelman M. C., 2018b, *ApJ*, 867, L18
- Zhdankin V., Uzdensky D. A., Werner G. R., Begelman M. C., 2020, *MNRAS*, 493, 603
- Zhdankin V., Uzdensky D. A., Kunz M. W., 2021, *ApJ*, 908, 71

APPENDIX A: ENERGY OSCILLATION REMOVAL OF TRACKED PARTICLES

As noted in Section 3.1, there are large (order-unity) oscillations in particle energy due to $\mathbf{E} \times \mathbf{B}$ drift, particularly significant in relativistic turbulence with $E_{\text{rms}} \sim B_0$. These were described in Wong et al. (2020), along with a removal procedure which is necessary to facilitate FP-type analysis. We briefly recount these here.

Consider a charged particle with four-velocity $\gamma \mathbf{v}$ moving in constant uniform electromagnetic fields. We use unprimed variables for the lab frame and primed variables for the $\mathbf{E} \times \mathbf{B}$ drift frame moving with velocity \mathbf{v}_D , given by $c\mathbf{v}_D/(c^2 + v_D^2) = \mathbf{E} \times \mathbf{B}/(E^2 + B^2)$. The particle's lab-frame energy, obtained by inverse Lorentz transform from simple gyration in the drift frame, is given by the equations

$$\gamma(t) = \gamma_D \gamma' \left[1 + \beta_D \frac{v'_\perp}{c} \cos(\omega' t') \right], \quad (\text{A1})$$

$$t = t_0 + \gamma_D \left[t' + \beta_D \frac{v'_\perp}{\omega' c} \sin(\omega' t') \right]. \quad (\text{A2})$$

Here, t is the coordinate time, $\beta_D = v_D/c$, $\gamma_D = (1 - \beta_D^2)^{-1/2}$, v'_\perp is the particle's primed-frame velocity perpendicular to \mathbf{B}' , $\omega' = eB'/\gamma' mc$ is the cyclotron frequency with $B' = B/\gamma_D$, and t_0 is a phase. Since we are considering relativistic particles ($v'_\perp \sim c$) and relativistic turbulence ($E_{\text{rms}} \sim B_0$ and $\beta_D \sim 1$), the oscillation magnitude is comparable to γ . In actuality, the motion and fields are not constant, but the result essentially still applies, and such oscillations would heavily complicate FP analysis.

The oscillation removal procedure is informed by (A1), which indicates that in the idealized case of uniform constant fields, the secular component of the lab-frame energy is $\gamma_D \gamma'$. We add a step to smooth \mathbf{v}_D fluctuations as the particle traverses small-scale fields, which is not strictly necessary but slightly improves the result compared to just $\gamma_D \gamma'$. First, we average \mathbf{v}_D over the smoothed gyroperiod $2\pi\gamma' mc/eB$, where $\gamma' = \gamma_D \gamma (1 - \mathbf{v}_D \cdot \mathbf{v}/c^2)$, and denote this smoothed value by $\langle \mathbf{v}_D \rangle$. We 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 $\langle \gamma_D \rangle \equiv (1 - \langle \mathbf{v}_D \rangle^2/c^2)^{-1/2}$. Here, the Lorentz boost on γ has been done twice: first to find a window length to smooth \mathbf{v}_D , and then again with the smoothed \mathbf{v}_D to obtain the final result.

This procedure extracts the secular component of particle energy, greatly reducing oscillations and allowing us to test the FP picture of NTPA. Hereafter, γ and 'energy' refer to $\langle \gamma_D \rangle \langle \gamma' \rangle$, except for overall particle energy distributions and magnetic energy spectra (where the difference is minor anyway).

APPENDIX B: PARTICLES-PER-CELL CONVERGENCE

In this Appendix, we describe the dependence of the measured diffusion coefficient $D(\gamma)$ on the number of particles per cell N_{ppc} . Fig. B1 shows the dependence of $D(\gamma)$ at fixed $t_0 = 10L/c$ on N_{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_{ppc} decreases, with the beginning occurring at slightly higher γ . Lower energies also show a power-law scaling, which becomes shallower with decreasing N_{ppc} . These two trends are displayed consistently and monotonically with respect to N_{ppc} . Importantly, even this lower-energy segment shows convergence with increased N_{ppc} , with essentially full convergence achieved for $N_{\text{ppc}} \gtrsim 64$ for the entire range of particle energies. Note that the effect of decreasing N_{ppc} in Fig. B1 resembles the trend of increasing system size from Fig. 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.




Figure B1. Diffusion coefficient $D(\gamma)$ at $t_0 = 10L/c$ for the N_{ppc} -scan simulations, which have $L/\rho_{e0} = 256$ and $\sigma_0 = 3/8$. In the non-thermal region, there is excellent convergence. At low energy, D is converged for simulations of high N_{ppc} , but increased by numerical effects when N_{ppc} is too low.

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