

Self-organization in collisionless, high- β turbulence

S. Majeski^{1,2†}, M. W. Kunz^{1,2}, and J. Squire³

¹Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

²Princeton Plasma Physics Laboratory, PO Box 451, Princeton, NJ 08543, USA

³Department of Physics, University of Otago, 730 Cumberland St, North Dunedin, Dunedin 9016, New Zealand

(compiled on 7 May 2024)

The MHD equations, as a collisional fluid model that remains in local thermodynamic equilibrium (LTE), have long been used to describe turbulence in myriad space and astrophysical plasmas. Yet, the vast majority of these plasmas, from the solar wind to the intracluster medium (ICM) of galaxy clusters, are only weakly collisional at best, meaning that significant deviations from LTE are not only possible but common. Recent studies have demonstrated that the kinetic physics inherent to this weakly collisional regime can fundamentally transform the evolution of such plasmas across a wide range of scales. Here we explore the consequences of pressure anisotropy and Larmor-scale instabilities for collisionless, $\beta \gg 1$ turbulence, focusing on the role of a self-organizational effect known as ‘magneto-immutability’. We describe this self-organization analytically through a high- β , reduced ordering of the CGL-MHD equations, finding that it is a robust inertial-range effect that dynamically suppresses magnetic-field-strength fluctuations, anisotropic-pressure stresses, and dissipation due to heat fluxes. As a result, the turbulent cascade of Alfvénic fluctuations continues below the putative viscous scale to form a robust, nearly conservative, MHD-like inertial range. These findings are confirmed numerically via Landau-fluid CGL-MHD turbulence simulations that employ a collisional closure to mimic the effects of microinstabilities. We find that microinstabilities occupy a small ($\sim 5\%$) volume-filling fraction of the plasma, even when the pressure anisotropy is driven strongly towards its instability thresholds. We discuss these results in the context of recent predictions for ion-versus-electron heating in low-luminosity accretion flows and observations implying suppressed viscosity in ICM turbulence.

1. Introduction

1.1. Motivation

Rarely can a problem of astrophysical fluid dynamics be approached without any consideration for the effects of turbulence. In fact, many descriptions of fundamental astrophysical phenomena, such as transport in accretion flows and dynamo amplification of cosmic magnetic fields, are intrinsically reliant upon it. As a result, an abundance of literature exists analyzing the role of turbulence in environments from the solar wind to the intracluster medium (ICM) of galaxy clusters (e.g., Goldstein *et al.* 1995; Brandenburg & Subramanian 2005; Schekochihin *et al.* 2004). Yet many of these studies employ theoretical or numerical methods founded upon the assumption that the plasma that pervades these systems is collisional. X-ray observations of the hot and dilute ICM in

† Email address for correspondence: smajeski@princeton.edu

the Perseus and Coma clusters suggest otherwise, with implied Coulomb-collisional mean free paths typically only $\sim 0.1\text{--}0.01$ times that of the large-scale gradients (e.g., [Kunz et al. 2022](#)). Meanwhile, in-situ measurements of the plasma comprising the solar wind have long revealed that ion Coulomb mean free paths can reach nearly 1 au, allowing significant deviations from local thermodynamic equilibrium (LTE) ([Marsch 2006](#)). These plasmas are frequently modeled as collisional fluids because relaxation of the LTE assumption introduces myriad complications that make their theoretical description and simulation quite difficult. This is especially true in plasmas that possess significant scale separation between their dynamical gradient length scales and the kinetic length scales such as the ion Larmor radius ($r_{L,i}$). Even in high- β plasmas where magnetic fields are energetically weak (with $\beta \doteq 8\pi p/B^2$ the ratio of the thermal and magnetic pressures), the vast length scales characteristic of astrophysical environments like the ICM mean that there can be as much as ten orders of magnitude separating the Coulomb mean free path from the Larmor radius. For fluctuation frequencies that are large compared to the collision frequencies, plasmas approximately conserve the double adiabatic invariants, $p_\perp/\rho B$ and $p_\parallel B^2/\rho^3$, where p_\perp and p_\parallel are the thermal pressures across and along the local magnetic-field direction. When the density (ρ) and magnetic-field strength (B) change, pressure anisotropy $\Delta \doteq \Delta p/p_\parallel \doteq p_\perp/p_\parallel - 1$ results, in turn exciting a plethora of macrophysical and microphysical effects. From kinetic microinstabilities to plasma self-organization, these pressure anisotropy-mediated effects are crucial to our understanding of astrophysical turbulence. For that reason, they are the chief focus of this work.

1.2. Consequences of pressure anisotropy

Perhaps the simplest yet most consequential way in which pressure anisotropy modifies plasma dynamics is through its effect on Alfvén waves. The tension force responsible for these waves' propagation in a collisionless plasma is not only a function of B , but also of Δ . As a result, the propagation speed of an Alfvénic disturbance is the effective Alfvén speed $v_{A,\text{eff}} \doteq v_A \sqrt{1 + \beta \Delta / 2}$, with the pressure anisotropy either enhancing or suppressing wave propagation depending on its sign. In high- β plasmas, only a small amount of anisotropy is required to have a dramatic effect on Alfvénic motions. A notable example of large β values enabling Δ to play this elevated role is in the ‘Alfvén wave interruption’ process of [Squire et al. \(2016, 2017a,c\)](#). Those authors found that, if the amplitude of a long-wavelength shear-Alfvén wave is sufficiently large, then the associated magnetic-field perturbation can adiabatically generate pressure anisotropy satisfying $\Delta \leq -2\beta^{-1}$, causing the Alfvén wave to self-interrupt and cease propagating. This pressure anisotropy also need not come from the Alfvén wave itself, but rather could be produced by other long-wavelength Alfvén waves and/or ion-acoustic waves that interact with the Alfvén wave ([Majeski & Kunz 2024](#)). At Larmor scales, pressure anisotropy plays an additional role as a trigger of kinetic microinstabilities. At high β , these instabilities are typically the mirror and firehose, which can be excited when $\Delta \gtrsim \beta^{-1}$ and $\Delta \lesssim -2\beta^{-1}$, respectively ([Barnes 1966; Hasegawa 1969; Hellinger & Matsumoto 2000](#)). Both instabilities grow rapidly on $r_{L,i}$ -scales, causing sharp kinks in the magnetic field and thereby breaking conservation of the double adiabats. For this reason, their most notable effect on large-scale dynamics is their introduction of an effective collisionality, which isotropizes the thermal pressure back towards marginal instability. This effect has been shown to cause Braginskii-MHD-like behaviour in otherwise collisionless sound waves and suppress the nonlinear saturation of collisionless damping in compressive non-propagating modes, even when their wavelengths far exceed the Larmor scale ([Kunz et al. 2020; Majeski et al. 2023](#)). It has also been suggested that, in the context of ICM turbulence, these microinstabilities might cause so much particle scattering that

the effective Reynolds number could be increased to allow an extended turbulent cascade below the nominal (Coulomb-collisional) viscous scale (St-Onge *et al.* 2020; Zhuravleva *et al.* 2019; Kunz *et al.* 2022).

On the other hand, a competing effect of pressure anisotropy has also been found to *suppress* viscous stress in ICM-relevant plasmas, while holding consequences for many other astrophysical environments. This effect, termed ‘magneto-immutability’, involves the self-organization of Alfvénic turbulence in high- β weakly collisional (Squire *et al.* 2019) and collisionless (Squire *et al.* 2023) plasmas to avoid changes in magnetic-field strength and thus the production of pressure anisotropy. That being said, exactly how this self-organization occurs and how robust it is remains somewhat mysterious.

1.3. Magneto-immutability

As its name suggests, magneto-immutability involves a tendency for a weakly collisional or collisionless turbulent plasma to organize in such a way that fluctuations in the magnetic-field strength become rare (e.g., as compared with fluctuations realized in collisional magnetohydrodynamic (MHD) turbulence). More specifically, it is defined as the suppression of field-aligned gradients in u_{\parallel} and Δp via self-organization (i.e., not through collisional or collisionless damping, or by choice of forcing). To understand how this definition connects to the suppression of changes in B , we briefly review the work of Squire *et al.* (2019), within which magneto-immutability was initially discovered. By studying weakly collisional, Braginskii-MHD turbulence, Squire *et al.* (2019) found that magneto-immutability occurs when the thermal pressure can affect the flow anisotropically. In Alfvénically driven turbulence in which the flow is incompressibly forced, the primary source of pressure anisotropy is through changes in B ; thus, magneto-immutability inherently limits the production of pressure anisotropy by reorganizing turbulent motions to avoid those changes. By studying the incompressible induction equation, $d_t \ln B = \hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{u}$, Squire *et al.* (2019) related these effects to the suppression of field-aligned gradients of u_{\parallel} and, by association, the suppression of field-aligned gradients of Δp . Squire *et al.* (2023) expanded the previous Braginskii-MHD investigation to the fully collisionless regime, finding that magneto-immutability is robust with respect to the closure employed for the pressure equations. This was accomplished by performing simulations using the Chew-Goldberger-Low (CGL; Chew *et al.* 1956) MHD model of a collisionless fluid (referred to as ‘active- Δ ’), and comparing them with simulations using isothermal MHD that passively evolve p_{\perp} and p_{\parallel} in response to the density and magnetic-field fluctuations in the MHD turbulence (referred to as ‘passive- Δ ’ in Squire *et al.* 2023). Prior to these studies of magneto-immutability, it was thought that the only means for strong turbulence to extend to small scales in a weakly collisional, high- β plasma was for the parallel viscous scale associated with the Coulomb collisionality to be significantly reduced via anomalous particle scattering (e.g., by plasma micro-instabilities). However, with magneto-immutability dynamically regulating the level of pressure anisotropy, and therefore the viscous stress, weakly collisional turbulence can form a surprisingly robust, MHD-like, approximately conservative inertial range.

To demonstrate what this looks like qualitatively, we present in figure 1 two snapshots taken from simulations of Alfvénically driven, high- β turbulence: one using ‘active- Δ ’ (top) and one using ‘passive- Δ ’ (bottom), the numerical methodology for which is provided in §3. These simulations are both performed at $\beta_0 = 10$, with the driven magnetic perturbation $\delta B_{\perp} \approx B_0/2$ at the outer scale. Given conservation of the double adiabats, this combination of forcing amplitude and high- β is more than sufficient to generate values of pressure anisotropy that are well beyond the aforementioned mirror and firehose thresholds (the white regions). Yet, only in the ‘passive- Δ ’ simulation do

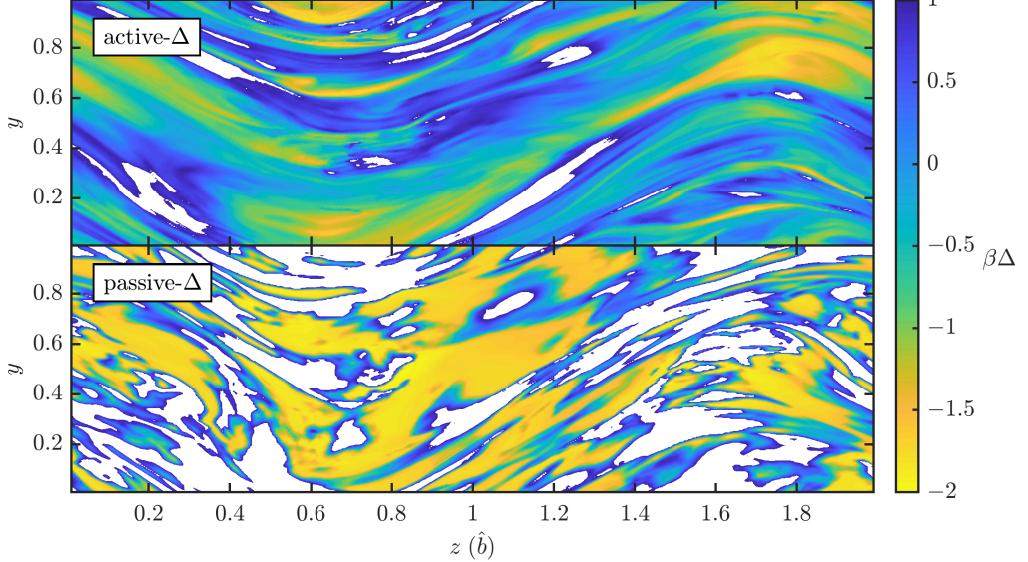


Figure 1: Dimensionless pressure anisotropy, multiplied by β , in an active- Δ and a passive- Δ simulation of driven turbulence. These simulations were performed with an initial background $\beta_0 = 10$, and forced Alfvénically such that the magnetic perturbation amplitude at the outer scale satisfies $\delta B_\perp \approx B_0/2$. White regions lie beyond the mirror and firehose thresholds, $\beta\Delta > 1$ and $\beta\Delta < -2$ respectively. Approximately half of the simulation domain is beyond these thresholds in the passive- Δ run, while only a small percentage is unstable in the active- Δ run, illustrating the effectiveness of magneto-immutability at reducing Δ .

such unstable regions appear to be a regular occurrence. Discerning exactly why this occurs, and in what regimes we can expect this to hold, are the goals of this paper.

1.4. Outline

Squire *et al.* (2019) and Squire *et al.* (2023) demonstrated the profound effect that magneto-immutability can have on turbulence in high- β plasmas. Consequently, it is important to understand just how robust is this self-organization before we can apply it more broadly to turbulent astrophysical plasmas. To that end, in this work we present a theory, based on a reduced model of the CGL-MHD equations, that improves our understanding of several key aspects of magneto-immutability and, in doing so, resolves some limits of their effects on high- β , collisionless turbulence. To prime the reader on how this is accomplished, the following outline summarizes the methodology and key topics of our investigation. In §2.1 we introduce the Landau-fluid CGL-MHD equations, which are employed both analytically and numerically to describe collisionless plasma turbulence. The linear wave solutions of Landau-fluid CGL are discussed in §2.2, with particular attention paid to the difference in timescales between compressive modes in collisionless and collisional plasmas. In §2.3 we introduce an ordering based on $\beta \gg 1$ and the principle of critical balance, which allows us to simplify the CGL-MHD equations and obtain the first analytical description of magneto-immutability in §2.4. This ‘reduced CGL-MHD’ model predicts, for example, that magneto-immutability does not depend on scale, that the associated self-organization process is unaffected to leading order in $\delta B_\perp/B_0$ by heat fluxes caused by field-aligned temperature gradients, and that the pressure anisotropy

exhibits a passive-scalar-like k -space spectrum. In §3 we verify the assumptions used to develop our high- β ordering and test several predictions made by the reduced CGL-MHD equations. This is accomplished using a numerical CGL-MHD solver built into the ATHENA++ framework, described briefly within §3.1 and in more detail in Appendix A of Squire *et al.* (2023). We also explore additional aspects of magneto-immutable turbulence not predicted by the reduced equations, such as the volume-filling fraction of regions unstable to the firehose and mirror (§3.3), how magneto-immutability occurs through the misalignment of the velocity strain with the magnetic field (§3.6, figure 8), and the ability of micro-instability-induced scattering to interfere with self-organization (§3.8). In §4, we summarize the main findings of this work, and discuss the consequences they hold for turbulence in weakly collisional astrophysical environments such as low-luminosity black-hole accretion flows and the ICM.

2. Theoretical description

2.1. The dynamical equations

Because we wish to apply our theory to systems in which the ion Larmor radii are as much as ten orders of magnitude smaller than the collisional mean free paths, we choose to employ a collisionless model that neglects finite-Larmor-radius (FLR) effects, saving micro-scale considerations like the mirror and firehose instabilities for the numerical analysis of §3. The two most suited models for describing a plasma under these assumptions are then the drift kinetic-MHD model of Kulsrud (1983) and the CGL-MHD fluid approach of Chew *et al.* (1956). While the drift-kinetic approach is more accurate in that it self-consistently determines the heat fluxes via moments of the guiding-centre distribution function, we instead employ the simpler CGL-MHD model. The CGL-MHD model by itself provides no closure for the form of the heat fluxes, and so for lack of anything markedly better we adopt the ‘Landau-fluid’ closure introduced by Snyder *et al.* (1997), which is designed to capture the effects of linear Landau damping on the fluid quantities. Fortunately, we find that the turbulence behaves in such a way that the exact form of the heat fluxes is relatively unimportant, suggesting that any errors introduced by applying these approximate heat fluxes are small (indeed, in §3.6 we demonstrate that magneto-immutability occurs in a hybrid-kinetic simulation of high- β Alfvénic turbulence). With that in mind, the CGL-MHD equations are given by (Chew *et al.* 1956):

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (2.1a)$$

$$\rho \frac{d\mathbf{u}}{dt} = -\nabla \left(p_{\perp} + \frac{B^2}{8\pi} \right) + \nabla \cdot \left[\hat{\mathbf{b}} \hat{\mathbf{b}} \left(p_{\perp} - p_{\parallel} + \frac{B^2}{4\pi} \right) \right], \quad (2.1b)$$

$$\frac{dB}{dt} = \mathbf{B} \cdot \nabla \mathbf{u} - \mathbf{B} \nabla \cdot \mathbf{u}, \quad (2.1c)$$

$$p_{\perp} \frac{d}{dt} \ln \frac{p_{\perp}}{\rho B} = -\nabla \cdot (q_{\perp} \hat{\mathbf{b}}) - q_{\perp} \nabla \cdot \hat{\mathbf{b}}, \quad (2.1d)$$

$$p_{\parallel} \frac{d}{dt} \ln \frac{p_{\parallel} B^2}{\rho^3} = -\nabla \cdot (q_{\parallel} \hat{\mathbf{b}}) + 2q_{\perp} \nabla \cdot \hat{\mathbf{b}}, \quad (2.1e)$$

where \mathbf{B} is the magnetic field, \mathbf{u} is the ion flow velocity, and ρ is the mass density. Note that p_{\perp} and p_{\parallel} are defined with respect to the local magnetic-field direction $\hat{\mathbf{b}} = \mathbf{B}/B$. Here, $d/dt \doteq \partial/\partial t + \mathbf{u} \cdot \nabla$ is the convective time derivative. The equations (2.1d) and

(2.1e) demonstrate how, in the absence of heat fluxes, the quantities $p_{\perp}/\rho B$ and $p_{\parallel}B^2/\rho^3$ are conserved in time along the flow of the plasma. Conservation of $p_{\perp}/\rho B$ and $p_{\parallel}B^2/\rho^3$ is the collective result of individual particles conserving both their magnetic moment $\mu \doteq mw_{\perp}^2/2B$ and their parallel action (i.e., bounce invariant) $\mathcal{J} = \oint dw_{\parallel} mw_{\parallel}$, respectively (where $\mathbf{w} \doteq \mathbf{v} - \mathbf{u}$ is the particle velocity peculiar to the fluid frame). Because of the plasma's strong magnetization, the flows of perpendicular/parallel thermal energy, $q_{\perp/\parallel}$, occur exclusively along the local magnetic-field direction. For these quantities we adopt the '3+1 model' of [Snyder et al. \(1997\)](#):

$$q_{\perp} = -\frac{v_{\text{th}\parallel}}{\sqrt{\pi}|k_{\parallel}|} \left[\rho \nabla_{\parallel} \left(\frac{p_{\perp}}{\rho} \right) - p_{\perp} \left(1 - \frac{p_{\perp}}{p_{\parallel}} \right) \frac{\nabla_{\parallel} B}{B} \right], \quad (2.2a)$$

$$q_{\parallel} = -\frac{2v_{\text{th}\parallel}}{\sqrt{\pi}|k_{\parallel}|} \rho \nabla_{\parallel} \left(\frac{p_{\parallel}}{\rho} \right), \quad (2.2b)$$

where $v_{\text{th},\parallel} = \sqrt{2p_{\parallel}/\rho}$ is the parallel thermal speed, and $\nabla_{\parallel} = \hat{\mathbf{b}} \cdot \nabla$ is the field-aligned gradient. The field-aligned wavenumber $|k_{\parallel}|$ in the denominators of (2.2) is meant to be representative of a characteristic parallel scale of the perturbations to ρ , B , and $p_{\perp/\parallel}$ (e.g., [Sharma et al. 2006](#)). This quantity is a stand-in for the magnitude of the magnetic-field-aligned gradient operator, which is generally difficult to calculate as it must be evaluated along the exact perturbed fields at each time step. For the majority of the numerical simulations presented in §3, we take $|k_{\parallel}| = 4\pi/L_{\parallel}$ where L_{\parallel} is the field-parallel outer scale; for the remainder of this section, however, the exact value of $|k_{\parallel}|$ is unimportant. Because the heat fluxes (2.2) are designed to capture only linear collisionless damping, it is implicitly assumed that the perturbations being studied are small enough with respect to the background plasma that nonlinear damping effects can be ignored. We also assume that the background plasma pressure is isotropic, and that the electrons are cold. Effects of a background anisotropy on inertial-range kinetic turbulence have been considered by [Kunz et al. \(2015\)](#) and could be straightforwardly implemented within our model, however they are not necessary for describing magneto-immutability and assessing its robustness. Similarly, electron pressure is not essential to the conclusions we will draw; as long as the electrons are sufficiently collisional and isothermal, a finite electron pressure has no qualitative effect on our results.¹ Finally, the CGL-MHD equations' inability to resolve $r_{L,i}$ -scales inevitably implies that the growth of the mirror and firehose instabilities and their interaction with the plasma particles are not properly captured. In our numerical simulations (§3), we model their effect on the plasma through an anomalous collisionality that isotropizes the pressure to marginally unstable values, but no effort is made to incorporate the effects of this micro-instability scattering into our analytic model for magneto-immutability.

2.2. Relevant properties of collisionless hydromagnetic waves

Developing an intuition of the linear modes and their interactions is particularly beneficial towards understanding the scale-by-scale transfer of energy in a turbulent MHD cascade. It is therefore important to review the linear modes of the Landau-fluid CGL system, so that we can make assumptions and formulate our ordering in light of

¹Evidence of this can be found already in the numerical simulations of [Squire et al. \(2023\)](#), which showed that ion-to-electron temperature ratios of order unity had little effect on signatures of magneto-immutability. If the electrons were to be weakly collisional, however, certain effects caused by electron pressure anisotropy may need to be taken into account, a point we discuss in §4.

these fundamental behaviours. Many of these waves are discussed in greater detail within Appendix B of Majeski *et al.* (2023); here we restrict ourselves to aspects of the linear modes that are relevant to collisionless high- β turbulence.

Our discussion begins with those modes that are least affected by the transition from collisional to collisionless regimes. In particular, linear collisionless Alfvén waves are identical to their collisional counterparts when their wavelengths far exceed $r_{\text{L},i}$, because neither density nor pressure perturbations are generated as they propagate. Meanwhile, the mode that experiences the most minor nonzero change is the fast mode, which has a nearly unchanged phase speed, although it is susceptible to collisionless damping depending on the propagation angle with respect to the background field, $(k_{\parallel}/k_{\perp})$. Consequently, most of the differences between high- β CGL and collisional MHD turbulence do not originate from the behaviour of Alfvén or fast waves. The remaining collisionless hydromagnetic modes are non-propagating modes, ion-acoustic waves, and kinetic entropy modes, which are most easily compared to the slow-magnetosonic and pressure-balanced entropy modes of collisional MHD.² Because there is one more dynamical equation in CGL-MHD than in collisional MHD, there is one additional linear mode solution. However, the kinetic entropy mode is both static and heavily damped at high β , with no clear collisional counterpart (see, e.g., Majeski *et al.* 2023), so it can be excluded from this comparison.

We are then left with the one-to-one comparisons of non-propagating modes with pressure-balanced entropy modes, and ion-acoustic waves with slow magnetosonic waves. Both the non-propagating and pressure-balanced entropy modes have zero phase speed, with the entropy mode being fully static, and the non-propagating mode decaying due to transit-time damping at a rate $\gamma \sim |k_{\parallel}|v_A/\sqrt{\beta}$ that is small at high β . And while entropy modes satisfy perfect isotropic pressure balance (allowing them to remain static), non-propagating modes exist in a state of approximate pressure balance between p_{\perp} and B^2 when $k_{\parallel} \ll k_{\perp}$ (Majeski *et al.* 2023). For these reasons, at high β these modes are rather similar, with the important feature that they are both approximately static. The comparison of ion-acoustic and slow waves, on the other hand, yields far fewer similarities, especially at high β . The most essential difference lies in their dispersion relations, with the ion-acoustic wave having a frequency $\omega \sim k_{\parallel}v_{\text{th}}$, while the slow-wave frequency scales as $k_{\parallel}v_A$ when $\beta \gg 1$ and $\omega \sim k_{\parallel}v_{\text{th}}$ when $\beta \ll 1$. The ion-acoustic wave is also Landau damped at a rate $\gamma \sim |k_{\parallel}|v_{\text{th}}$, meaning that at high β all characteristic timescales of ion-acoustic waves are much shorter than those of slow (and Alfvén) waves. These disparate timescales are fundamental to how collisionless waves interact with an Alfvénic cascade. For example, in collisional MHD turbulence, a timescale separation exists between fast and Alfvén waves because fast waves have a phase speed proportional to k rather than just k_{\parallel} (with $k_{\parallel} \ll k$ in the inertial range of Alfvénic turbulence, see Goldreich & Sridhar (1995)). As a result, it is often argued that the rapid propagation of fast waves decouples them from any Alfvénic dynamics. In the next subsection (§2.3), we will argue that at $\beta \gg 1$, the ion-acoustic wave – the only other compressive ($\nabla \cdot \mathbf{u} \neq 0$) wave – decouples from the Alfvénic dynamics as well because $v_{\text{th}} \gg v_A$.

2.3. An asymptotic ordering for high- β collisionless turbulence

The equations (2.1) are too complicated with which to work analytically, therefore we seek an asymptotic ordering that distills the physics responsible for magneto-

²The non-propagating mode and kinetic entropy mode are equivalent to the ‘+’ and ‘−’ branches of the high- β gyrokinetic dispersion relation from Schekochihin *et al.* (2009, §6.2.2). In that case, those authors simply refer to non-propagating modes as ‘magnetic-field-strength fluctuations’.

immutability. Before doing so, however, it is instructive first to consider the asymptotic ordering used in the collisional reduced MHD (RMHD) employed by [Zank & Matthaeus \(1992\)](#) and [Schekochihin *et al.* \(2009, §2\)](#), from which our ordering borrows heavily:

$$\frac{u_\perp}{v_A} \sim \frac{u_\parallel}{v_A} \sim \frac{\delta B_\perp}{B_0} \sim \frac{\delta B_\parallel}{B_0} \sim \frac{\delta \rho}{\rho_0} \sim \frac{\delta p}{p_0} \sim \frac{k_\parallel}{k_\perp} \doteq \epsilon \ll 1 \sim \beta, \quad \partial_t \sim k_\parallel v_A. \quad (2.3)$$

Here, all of the perturbations belong to a given k_\perp/\parallel ‘shell’ and interactions are taken to be local in k -space, meaning that as perturbations become smaller through the cascade so does the ratio k_\parallel/k_\perp . Arguably the most essential aspect of this ordering is that it enforces ‘critical balance’ (CB) of the Alfvénic fluctuations, viz. $k_\perp u_\perp \sim k_\parallel v_A$. This is a statement that the nonlinear timescale associated with eddy deformation by the Reynolds stress $(k_\perp u_\perp)^{-1}$ is comparable to the linear propagation timescale of Alfvén waves $(k_\parallel v_A)^{-1}$, as is expected for strong turbulence ([Goldreich & Sridhar 1995](#)). As a result, the evolutionary timescale at each wavenumber is expected to be roughly $\partial_t \sim k_\parallel v_A$, since linear and nonlinear timescales are equivalent. This is also true of MHD slow modes, given that they propagate at v_A or slower, allowing for the formation of a similarly strong compressive cascade. As mentioned in §2.2 however, the enhanced propagation speed of fast modes prevents them from being effectively coupled to the Alfvénic motions, and as such they are ordered out of RMHD. Our ordering shares most of these assumptions, yet differs in a few key ways.

First, we incorporate the largeness of the background β directly into the ordering, so that $\beta \sim \epsilon^{-1}$. This may seem odd because β , being a background quantity, is constant across all scales even though ϵ becomes smaller as the fluctuations cascade anisotropically to larger k . However, we are concerned with plasmas in which perturbations near the outer scale often exceed β^{-1} in relative amplitude. As a result, it would be incorrect to eliminate terms of order $(u_\perp/v_A)^2$ but keep those of order $u_\perp/\beta v_A$. Ordering $\beta \sim \epsilon^{-1}$ allows us to retain both terms. We also assume that the Alfvénic component of our turbulence is critically balanced, as in collisional RMHD. This is not necessarily a given, however there is evidence from the work of [Squire *et al.* \(2023\)](#) that immutability does not interfere with critical balance, a feature that we reproduce through simulations of our own in §3. Next, we assume that, as a result of the plasma being collisionless, δp is replaced in the ordering by δp_\perp and δp_\parallel , with similar amplitudes to (2.3). This is a straightforward assumption, yet its implications are more nuanced. Doing so replaces slow magnetosonic waves with ion-acoustic waves, and because we include β in the overall ordering, this implies that nonlinear mixing of ion-acoustic waves by Alfvén waves is weak when local in k -space, since $k_\parallel v_{th} \gg k_\perp u_\perp$. We will therefore assume that ion-acoustic waves, like fast waves, decouple from the Alfvénic turbulence. This assumption is supported by analytical calculations and simulations of the interaction of Alfvén and ion-acoustic waves in high β plasmas, using the same CGL-MHD code employed in this work ([Majeski & Kunz 2024](#)). Ordering out ion-acoustic waves therefore enables us to continue with the assertion that, for all variables, time derivatives are of order $\partial_t \sim k_\parallel v_A$.

Our final major assumption, which follows from those before, is that density fluctuations can be neglected in determining the leading-order dynamics (i.e. $\delta \rho/\rho_0 \sim \epsilon^2$ at most). In Alfvénic turbulence, high values of β naturally inhibit the magnitude of density fluctuations by reducing the sonic Mach number (u/v_{th}) of the forcing. Furthermore, any forcing must be associated with timescales that are sonic or faster if it is to produce substantial compressive fluctuations, a rather rare occurrence in the systems with which we are concerned. As a result, density fluctuations are expected to be small in amplitude at the outer scale. We are able to apply this assumption throughout the inertial range because the only linearly compressive waves – ion-acoustic and fast modes – have been

ordered out of the dynamics. This assumption of small $\delta\rho$ allows us to consider only perturbations to the temperatures $\delta T_{\perp/\parallel}$ rather than the pressures $\delta p_{\perp/\parallel}$.

The collisionless, high- β ordering that results from the above considerations and assumptions is

$$\frac{u_{\perp}}{v_A} \sim \frac{u_{\parallel}}{v_A} \sim \frac{\delta B_{\perp}}{B_0} \sim \frac{\delta B_{\parallel}}{B_0} \sim \frac{\delta T_{\perp}}{T_0} \sim \frac{\delta T_{\parallel}}{T_0} \sim \frac{1}{\beta} \sim \frac{k_{\parallel}}{k_{\perp}} \doteq \epsilon, \quad \partial_t \sim k_{\parallel} v_A. \quad (2.4)$$

Note that the choice $\delta T_{\perp/\parallel}/T_0 \sim \beta^{-1}$ means that $\beta\Delta \sim 1$, or in other words, the anisotropic pressure-stress is present at the same order as the flow inertia.³ This is important, because in order to demonstrate that magneto-immutability avoids significant Δp stress, Δp must first be made large enough to disrupt the turbulence in the absence of immutability.⁴ Some of the assumptions that yield (2.4), such as $\delta\rho/\rho_0$ being negligible, are only justified qualitatively, therefore we test them against numerical simulations in §3 to establish confidence in the relevance of the ordering. We are now prepared to apply the ordering to (2.1) and obtain the high- β reduced CGL-MHD equations.

2.4. The high- β reduced CGL-MHD equations

In this section we formulate equations that describe the leading-order evolution for each of the fields present in the ordering (2.4). Doing so involves expanding these quantities in (fractional) orders of ϵ , for example with T_{\parallel} :

$$T_{\parallel} = T_0 + \delta T_{\parallel}^{(1)} + \delta T_{\parallel}^{(3/2)} + \delta T_{\parallel}^{(2)} + \dots, \quad (2.5)$$

where the parenthetical superscript represents the order of each term in powers of ϵ . Note that beyond the leading-order perturbation, which is of order ϵ , half-integer orders must be used for all higher-order perturbations. This is required by the heat fluxes, which are proportional to $v_{\text{th}} \sim v_A/\sqrt{\epsilon}$.

Keeping in mind the expectation that density fluctuations are $\mathcal{O}(\epsilon^2)$ at most and thus do not affect the leading order dynamics, the continuity equation (2.1a) simplifies to $\nabla \cdot \mathbf{u} = 0$. If ordered according to (2.4), this becomes

$$\nabla \cdot \mathbf{u} = \nabla_{\perp} \cdot \mathbf{u}_{\perp} + \nabla_{\parallel} u_{\parallel} \approx \nabla_{\perp} \cdot \mathbf{u}_{\perp}^{(1)} = 0. \quad (2.6)$$

The divergence-free condition for \mathbf{B} similarly yields $\nabla_{\perp} \cdot \delta \mathbf{B}_{\perp}^{(1)} = 0$ to leading order. These conditions on $\mathbf{u}_{\perp}^{(1)}$ and $\delta \mathbf{B}_{\perp}^{(1)}$ assert that the perpendicular flow and magnetic field perturbations are dominated by Alfvén waves, which are naturally incompressible. Given that large-scale collisionless Alfvén waves are linearly identical to those of collisional plasmas, it is no surprise that this aspect of the reduced equations is unchanged from standard RMHD. Continuing with the induction equation (2.1c), the leading-order

³A separate, strictly Alfvénic ordering could be considered in which $(u_{\perp}/v_A)^2 \sim \Delta$, as one might expect from a pure shear-Alfvén eigenmode. With $\beta^{-1} \sim u_{\perp}/v_A$, this would be relatively uninteresting as, with or without immutability, the Δp -stress would be too weak to affect the Alfvénic turbulence to leading order. Instead, it would be necessary to choose $\beta^{-1} \sim (u_{\perp}/v_A)^2$. However, this would introduce ambiguity into ordering u_{\parallel} and δB_{\parallel} , which in our turbulence and that of Squire *et al.* (2023) appear to be comparable to their perpendicular counterparts.

⁴This is akin to assuming that the Reynolds number $\text{Re} \equiv \rho u L / \mu_{\text{visc}}$ in hydrodynamic turbulence (with μ_{visc} the fluid viscosity) satisfies $\text{Re} \sim 1$, which prohibits the continuation of a turbulent cascade.

parallel and perpendicular components become

$$\frac{d\delta B_{\parallel}^{(1)}}{dt^{(0)}} = B_0 \hat{\mathbf{b}}^{(0)} \cdot \nabla u_{\parallel}^{(1)}, \quad (2.7a)$$

$$\frac{d\delta \mathbf{B}_{\perp}^{(1)}}{dt^{(0)}} = B_0 \hat{\mathbf{b}}^{(0)} \cdot \nabla \mathbf{u}_{\perp}^{(1)}, \quad (2.7b)$$

where

$$\frac{d}{dt^{(0)}} \doteq \frac{\partial}{\partial t} + \mathbf{u}_{\perp}^{(1)} \cdot \nabla_{\perp} \quad \text{and} \quad \hat{\mathbf{b}}^{(0)} \doteq \hat{\mathbf{z}} + \frac{\delta \mathbf{B}_{\perp}^{(1)}}{B_0}, \quad (2.8)$$

with $\hat{\mathbf{z}}$ being the direction of the background magnetic field \mathbf{B}_0 . Once again, these equations are equivalent to those obtained for the evolution of the magnetic field in standard RMHD. Equation (2.7a) states that leading-order changes to the magnetic-field strength are generated through field-aligned shear in u_{\parallel} , which we will see is reduced in magneto-immutable turbulence. That this has not been ordered out directly by (2.4) hints that magneto-immutability must be achieved through self-organization.

Next we address the momentum equation. Beginning with the perpendicular direction, the first two orders are trivially

$$\nabla_{\perp} \delta T_{\perp}^{(1)} = 0 \quad \text{and} \quad \nabla_{\perp} \delta T_{\perp}^{(3/2)} = 0 \quad (2.9)$$

as a result of β being large. Given that $k_{\parallel} \ll k_{\perp}$ at all scales, these perturbations can have no parallel gradients either, thus $\delta T_{\perp}^{(1)} = \delta T_{\perp}^{(3/2)} = 0$. The next two orders give perpendicular pressure balance,

$$\frac{\delta T_{\perp}^{(2)}}{T_0} = \frac{1}{\beta} \frac{\delta B_{\parallel}^{(1)}}{B_0} \quad \text{and} \quad \frac{\delta T_{\perp}^{(5/2)}}{T_0} = \frac{1}{\beta} \frac{\delta B_{\parallel}^{(3/2)}}{B_0}, \quad (2.10)$$

analogous to that of RMHD and observed within the simulations of [Squire et al. \(2023\)](#). The subsequent order dictates the evolution of \mathbf{u}_{\perp} :

$$\begin{aligned} \rho_0 \frac{d\mathbf{u}_{\perp}^{(1)}}{dt^{(0)}} &= -\nabla_{\perp} P_{\text{total}}^{(2)} + \frac{B_0}{4\pi} \frac{\partial \delta \mathbf{B}_{\perp}^{(1)}}{\partial z} + \frac{1}{4\pi} \delta \mathbf{B}_{\perp}^{(1)} \cdot \nabla_{\perp} \delta \mathbf{B}_{\perp}^{(1)} \\ &\quad - \frac{\beta}{8\pi} B_0 \frac{\partial}{\partial z} \left(\delta \mathbf{B}_{\perp}^{(1)} \frac{\delta T_{\parallel}^{(1)}}{T_0} \right) - \frac{\beta}{8\pi} \delta \mathbf{B}_{\perp}^{(1)} \cdot \nabla_{\perp} \left(\delta \mathbf{B}_{\perp}^{(1)} \frac{\delta T_{\parallel}^{(1)}}{T_0} \right). \end{aligned} \quad (2.11)$$

Here, $P_{\text{total}}^{(2)}$ represents the combined thermal and magnetic pressures evaluated at second order, which can be determined as a whole by enforcing $\nabla_{\perp} \cdot \mathbf{u}_{\perp} = 0$. Note that the thermal contribution to $P_{\text{total}}^{(2)}$ depends on the $\mathcal{O}(\epsilon^2)$ density fluctuations that we have so far ignored, however, they do not need to be determined independently of the other contributions to $P_{\text{total}}^{(2)}$. Because of perpendicular pressure balance (2.9), the leading-order pressure anisotropy is dominated by the contribution from $\delta T_{\parallel}^{(1)}/T_0$. The top line of (2.11) contains those terms that are already present in RMHD, while the bottom line incorporates feedback from the pressure anisotropy onto the perpendicular momentum. The contribution from pressure anisotropy remains because our ordering assumes $\beta\Delta \sim 1$, although this holds different meanings at different scales. Near the outer scale, the anisotropy is dominated by fluctuations of amplitude $\Delta \sim \beta^{-1}$. As the cascade continues to smaller scales however, fluctuations in Δ become smaller in amplitude and no longer compete with the magnetic tension, thus dropping out of the leading order of (2.11). As a result, the pressure anisotropy that shows up in (2.11) is really only that of the largest

scales, which the high- k Alfvén modes feel as a nearly constant modified background $v_{A,\text{eff}}$.

It is tempting to identify the pressure-anisotropy-related terms in (2.11) with magneto-immutability, however, equation (2.7a) suggests that magneto-immutability will primarily be mediated by u_{\parallel} . Therefore the parallel component of (2.1b) must be examined, for which the leading two orders are simply

$$\hat{\mathbf{b}}^{(0)} \cdot \nabla \delta T_{\parallel}^{(1)} = (\hat{\mathbf{b}} \cdot \nabla \delta T_{\parallel})^{(1)} = 0 \quad (2.12)$$

and

$$\hat{\mathbf{b}}^{(0)} \cdot \nabla \delta T_{\parallel}^{(3/2)} + \frac{\delta \mathbf{B}_{\perp}^{(3/2)}}{B_0} \cdot \nabla_{\perp} \delta T_{\parallel}^{(1)} = (\hat{\mathbf{b}} \cdot \nabla \delta T_{\parallel})^{(3/2)} = 0. \quad (2.13)$$

These equations result directly from ordering out ion-acoustic waves. If the ordering were such that $\partial_t u_{\parallel} \sim \epsilon k_{\parallel} v_{\text{th}}^2$, then the inertial term would be of the same order as the field-aligned temperature gradient (2.12), which would then be non-zero. Linearization of the resultant reduced equations would yield an ion-acoustic-like eigenfrequency proportional to $k_{\parallel} v_{\text{th}}$, however in our reduced equations no such eigenfrequency can be obtained. This is analogous to pressure balance in the perpendicular equations, which orders out fast modes by virtue of the assumption that the time derivative of u_{\perp} is not proportional to kv_{th} . Unlike (2.9), equations (2.12) and (2.13) involve both parallel and perpendicular gradients and alignment with $\delta \mathbf{B}_{\perp}$. Hence, it is not necessarily true that $\delta T_{\parallel}^{(1)} = \delta T_{\parallel}^{(3/2)} = 0$ as is the case for $\delta T_{\perp}^{(1)}$ and $\delta T_{\perp}^{(3/2)}$, but only that the field-aligned gradients of δT_{\parallel} are negligible. Also of note is that we have yet to make any mention of the heat fluxes, implying that this minimal variation of δT_{\parallel} along field lines is dynamic, rather than diffusive.⁵

For our purposes, there is no need to obtain the higher-order contributions to the parallel momentum equation, and so we continue applying our ordering by examining the evolution equations for the double adiabats. The first step in doing so is to consider the heat fluxes (2.2) in light of (2.9), (2.12), and (2.13), which can be rewritten in the following more transparent manner:

$$q_{\perp} = -\frac{v_{\text{th}} \rho}{\sqrt{\pi} |k_{\parallel}|} \left(\hat{\mathbf{b}} \cdot \nabla T_{\perp} + \frac{\beta_{\perp} \Delta}{2} \frac{\mathbf{B} \cdot \nabla B}{4\pi\rho} \right), \quad (2.14a)$$

$$q_{\parallel} = -\frac{2v_{\text{th}} \rho}{\sqrt{\pi} |k_{\parallel}|} \left(\hat{\mathbf{b}} \cdot \nabla T_{\parallel} \right). \quad (2.14b)$$

Here $\beta_{\perp} \doteq 8\pi p_{\perp}/B^2$ and $\beta_{\perp} \Delta \sim 1$, so the right-most term in (2.14a) is smaller than $p_{\perp} d_t \ln(T_{\perp}/B)$ by roughly $\epsilon^{3/2}$, and can therefore be neglected. What remains are heat fluxes proportional to the field-aligned temperature gradients $\hat{\mathbf{b}} \cdot \nabla T_{\perp/\parallel}$. However, these gradients are zero not just to leading order, but also to next order thanks to (2.9), (2.12), and (2.13). As a result, the $q_{\perp/\parallel}$ can be neglected entirely to the leading two orders at which they would contribute to (2.1d) and (2.1e). Thus, the surprisingly simple outcome is that

$$\frac{d\delta B_{\parallel}^{(1)}}{dt^{(0)}} = \frac{d\delta T_{\parallel}^{(1)}}{dt^{(0)}} = 0. \quad (2.15)$$

⁵Being dynamic, the suppression of $\hat{\mathbf{b}} \cdot \nabla u_{\parallel}$ is independent of the manner in which Δp is generated. It is possible then that this approach might also apply to the weakly collisional model of Braginskii-MHD, in which magneto-immutability was initially discovered by Squire *et al.* (2019). We investigate this possibility and some consequences of it within Appendix B.

This result is essentially the namesake of magneto-immutability. In order to satisfy both μ conservation and perpendicular pressure balance in a high- β plasma, the flow *must* self-organize to enforce $(\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u})^{(1)} = 0$, so that (2.7a) and (2.15) are in agreement. As a result, the magnetic field only experiences convective changes in its strength to leading order. In describing magneto-immutability qualitatively, Squire *et al.* (2023) also drew several comparisons with incompressibility in collisional MHD turbulence, and here we may draw another more direct comparison. In a collisional high- β plasma, pressure balance requires the *isotropic* pressure perturbation be reduced in order to match the magnetic pressure. Conservation of the single adiabat p/ρ^γ then implies that $d_t \delta\rho$ must be zero to leading order in the corresponding high- β collisional reduced MHD. Thus both magneto-immutability and incompressibility in high- β magnetized turbulence result from a combination of pressure balance and adiabatic invariance, only here the involvement of B in the collisionless invariants necessitates magneto-immutability.

Provided proper initial and boundary conditions, we have successfully closed the system of equations. Collecting the full high- β , reduced CGL-MHD equations (and dropping the orders), we have

$$\nabla_\perp \cdot \mathbf{u}_\perp = \nabla_\perp \cdot \delta\mathbf{B}_\perp = 0, \quad (2.16a)$$

$$\frac{d\delta\mathbf{B}_\perp}{dt} = B_0 \hat{\mathbf{b}} \cdot \nabla \mathbf{u}_\perp, \quad (2.16b)$$

$$\rho_0 \frac{d\mathbf{u}_\perp}{dt} = -\nabla_\perp P_{\text{total}} + \left(1 + \frac{\beta}{2} \frac{\Delta p}{p_0}\right) \left(\frac{B_0 \hat{\mathbf{z}} + \delta\mathbf{B}_\perp}{4\pi}\right) \cdot \nabla \delta\mathbf{B}_\perp, \quad (2.16c)$$

$$\frac{d\Delta p}{dt} = \frac{d\delta B_\parallel}{dt} = \hat{\mathbf{b}} \cdot \nabla u_\parallel = \hat{\mathbf{b}} \cdot \nabla \Delta p = 0. \quad (2.16d)$$

In (2.16d) and (2.16c), we have replaced $\delta T_\parallel/T_0$ with $-\Delta p/p_0$ without consequence, because $\Delta p/p_0 = \delta p_\parallel/p_0$ up to order $\epsilon^{3/2}$. As in collisional RMHD, the divergence of \mathbf{u}_\perp can be employed to determine P_{total} , and the convective derivative and field-parallel gradient become

$$\frac{d}{dt} = \frac{\partial}{\partial t} + \mathbf{u}_\perp \cdot \nabla_\perp \quad \text{and} \quad \hat{\mathbf{b}} = \hat{\mathbf{z}} + \frac{\delta\mathbf{B}_\perp}{B_0}, \quad (2.17)$$

respectively. Note that the reduced order of $\nabla_\parallel \delta T_\parallel$ compared to δT_\parallel itself allowed the anisotropic pressure to be pulled outside of the gradient operator in (2.16c) (since $\nabla_\parallel \Delta p$ becomes next order as well). In this sense, the pressure anisotropy is felt only as a modification to the effective Alfvén speed, thereby hindering its ability to interfere with the Alfvénic cascade as an anisotropic pressure-stress that could cause turbulent motions to damp into thermal energy. This is a central characteristic of magneto-immutability, and one of significant consequence. Without it, the turbulent cascade would cease long before reaching kinetic scales, dramatically modifying the transport properties of the plasma.

2.5. Features of the reduced system

We now discuss some properties of the reduced system (2.16), beginning with the quantities that its turbulent cascade conserves. In the same manner as Kunz *et al.* (2015), this analysis is made easier by defining Elsässer variables that incorporate the modification to the Alfvén speed by the pressure anisotropy (Elsässer 1950):

$$z^\pm = \mathbf{u}_\perp \pm v_{A,\text{eff}} \frac{\delta\mathbf{B}_\perp}{B_0}. \quad (2.18)$$

However, unlike in [Kunz et al. \(2015\)](#), here $v_{A,\text{eff}}$ is understood to be a function of both position and time. It is a straightforward process to show that the reduced system (2.16) requires the cascades of Elsässer energies to obey

$$\frac{\partial W_{\text{AW}}^{\pm}}{\partial t} \doteq \frac{\partial}{\partial t} \int d^3x \frac{1}{2} |\mathbf{z}^{\pm}|^2 = \mp \int d^3x \frac{P_{\text{total}}}{\rho_0} \frac{\partial v_{A,\text{eff}}}{\partial z}. \quad (2.19)$$

This is notably distinct from the Elsässer cascades of Alfvénic energy in collisional MHD and reduced kinetic-MHD (RKMHD), which satisfy $\partial_t W_{\text{AW}}^{\pm} = 0$ individually ([Schekochihin et al. 2009; Kunz et al. 2015](#)). Here, only the total Alfvénic fluctuation energy $W_{\text{AW}}^+ + W_{\text{AW}}^-$ is conserved. Meanwhile, the cross helicity $W_{\text{AW}}^+ - W_{\text{AW}}^-$, which measures the imbalance in forward-propagating and backward-propagating Alfvén waves, is not conserved. The fact that $W_{\text{AW}}^+ + W_{\text{AW}}^-$ is conserved proves that the pressure anisotropy, having reduced-order parallel gradients, cannot effectively thermalize the energy contained within Alfvénic fluctuations ([Squire et al. 2023](#), also see figure 7 here). That being said, the non-conservation of cross-helicity means that Δp is still able to redistribute energy between forward- and backward-propagating waves ([Majeski & Kunz 2024](#)). The other conserved quantities of the magneto-immutable cascade are rather trivial to detect from (2.16):

$$\frac{\partial}{\partial t} \int d^3x \frac{B_0 \delta B_{\parallel}}{4\pi} = 0 \quad \text{and} \quad \frac{\partial}{\partial t} \int d^3x \delta p_{\parallel} = 0, \quad (2.20)$$

where we have chosen to configure the constancy of δB_{\parallel} in units of energy for consistency with the other conserved quantities. Equations (2.20) represent, respectively, cascades of non-propagating modes and kinetic entropy modes. The reason for the respective associations of each conserved quantity comes down to the eigenvectors of each mode. Non-propagating modes, being in a state of near-perpendicular-pressure balance, have $\delta B_{\parallel}/B_0 \gg \delta p_{\perp}/p_0$ at high β . However, they also exhibit $\delta p_{\perp} < \delta p_{\parallel}$, which is inconsistent with the fact that $\delta p_{\parallel} \gg \delta p_{\perp}$ in magneto-immutable turbulence. Thus, the passive cascade of δp_{\parallel} can most easily be attributed to kinetic entropy modes, which vanish due to collisional damping in the MHD limit ([Majeski et al. 2023](#)). To leading order, there is no conserved quantity for compressive fluctuations (namely those associated with u_{\parallel}), because ion-acoustic and fast waves have been ordered out of the dynamics.

To further highlight the features that make magneto-immutable turbulence unique, we can compare the reduced equations (2.16) to the standard reduced MHD system, which is obtained by applying (2.3) to the collisional MHD equations. The dynamics of the Alfvénic fluctuations $\delta \mathbf{B}_{\perp}$ and \mathbf{u}_{\perp} are almost entirely unchanged from equations (2.16a)–(2.16c), with the only difference being that $\Delta p = 0$ in collisional MHD ([Zank & Matthaeus 1992; Schekochihin et al. 2009](#)). This suggests that the differences between the two models originate in the compressive cascade, a component of the turbulence that plays a passive role in collisional RMHD. The equations of δB_{\parallel} and u_{\parallel} in RMHD turbulence are ([Schekochihin & Cowley 2007](#)):

$$\frac{d}{dt} \frac{\delta B_{\parallel}}{B_0} = \frac{1}{1 + 2/(\gamma\beta)} \hat{\mathbf{b}} \cdot \nabla u_{\parallel} \approx \hat{\mathbf{b}} \cdot \nabla u_{\parallel} \quad (\beta \gg 1), \quad (2.21a)$$

$$\frac{d}{dt} \frac{u_{\parallel}}{v_A} = v_A \hat{\mathbf{b}} \cdot \nabla \frac{\delta B_{\parallel}}{B_0}, \quad (2.21b)$$

where $\gamma = 5/3$ is the single adiabatic index for a monoatomic gas. These equations describe slow-magnetosonic waves, which are well known to comprise the passively advected compressive cascade in MHD turbulence, something that is clearly not present in

(2.16). As discussed in the context of conserved quantities, the non-Alfvénic cascades can be attributed to non-propagating and kinetic entropy modes, with no truly compressive ($\nabla \cdot \mathbf{u} \neq 0$) cascade of any kind present to leading order.

When constructing the ordering (2.4) under the assumption $\delta B_{\parallel}/B_0 \sim \beta^{-1}$, we essentially guaranteed that realistic turbulent cascades eventually pass out of the parameter space within which the equations (2.16) are strictly valid. Yet, these equations are nonetheless accurate even when $\delta B_{\parallel}/B_0 \ll \beta^{-1}$ (so long as $\beta \gg 1$). To prove this, in Appendix A we show that the signatures of magneto-immutability, *viz.* $\nabla_{\parallel}\Delta p$ and $\nabla_{\parallel}u_{\parallel}$ suppression, can be recovered via a subsidiary high- β ordering of the RKMHD equations, which assume that $\delta B_{\parallel}/B_0 \ll \beta^{-1}$. Indeed, with the assumption of Landau-fluid heat fluxes we once again obtain the reduced system (2.16).⁶ Note, however, that at these small scales where $\delta B_{\parallel}/B_0 \ll \beta^{-1}$, such lack of parallel gradients in u_{\parallel} and Δp need only be enforced passively, as they have already been suppressed by larger scales obeying the ordering (2.4). In this sense, magneto-immutability exists in a somewhat weaker, watchdog-like state. If k_{\parallel} were somehow introduced to u_{\parallel} or Δp at small scales – for instance, by field-line wandering or magnetic reconnection (Meyrand *et al.* 2019) – it would be suppressed by magneto-immutable self-organization. However, if these parallel gradients were not created at small scales, they simply would remain negligible with Δp cascading passively. To contrast this with collisional RMHD or $\beta \sim 1$ RKMHD, any re-introduced parallel gradients in u_{\parallel} and Δp would be allowed to persist in the absence of dissipative effects. Unlike u_{\parallel} , the evolution of small-scale u_{\perp} in high- β collisionless turbulence becomes quite similar to that of collisional RMHD. The u_{\perp} fluctuations only feel a non-local modification to $v_{A,\text{eff}}$ produced by large scale Δp satisfying $\beta\Delta \sim 1$, thus u_{\perp} and δB_{\perp} effectively decouple from the u_{\parallel} and δB_{\parallel} cascades. Furthermore, the gradients in $v_{A,\text{eff}}$ are negligible at these small scales, allowing the energetic coupling between $|z^+|^2$ and $|z^-|^2$ that appears in (2.19) to weaken. Thus, the Alfvénic cascade approaches that of RKMHD with a constant background anisotropy, as described by Kunz *et al.* (2015).

Equations (2.16) leave us with several new testable predictions. First and foremost is that magneto-immutability, as a reduction of $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ and $\hat{\mathbf{b}}\cdot\nabla\Delta p$, is independent of scale. In the initial studies of Squire *et al.* (2019) and Squire *et al.* (2023), the authors concluded that the pressure anisotropy being driven by turbulent fluctuations must be competitive with the background magnetic tension (i.e. $\beta\Delta \sim 1$) in order for the suppression of $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ and $\hat{\mathbf{b}}\cdot\nabla\Delta p$ to occur. However, both of these signatures of magneto-immutability persist throughout the cascade simply as a result of $\beta \gg 1$, even when $\beta\Delta_k \ll 1$ at some large wavenumber k . Note, however, that because the $\delta p_{\perp/\parallel}$ and δB_{\parallel} are passively advected, the outer scale is important in that the compressive perturbations must be seeded with sufficiently large amplitude there.⁷ The second testable prediction is that magneto-immutability is, to leading order, independent of the heat fluxes. This is owed to the dynamical, rather than diffusive, manner in which minimal field-parallel variation of δT_{\parallel} is achieved, forcing the heat fluxes to contribute only to the next-order pressure anisotropy. Both magneto-immutability's independence of the heat fluxes, and

⁶In this case, note that $\beta\Delta$ in the u_{\perp} equation is not the local pressure anisotropy, but a background anisotropy set by large-scale $\beta\Delta \sim 1$ fluctuations, analogous to the anisotropy considered within Kunz *et al.* (2015).

⁷For random forcing at large scales, this is a reasonable assumption. However, if one wished to simulate turbulence using (2.16) and begin forcing somewhere in the inertial range, they would need to take care to properly initialize the amplitude of compressive fluctuations, or else the ordering may not be satisfied.

the field-parallel spreading of δT_{\parallel} in the absence of heat fluxes can therefore be validated directly by well-constructed CGL-MHD simulations. Equation (2.15) also predicts that δT_{\parallel} and Δp behave as passive scalars, thereby adopting the statistics of the flow that nonlinearly mixes them (Biskamp 2003). Although not addressed analytically, it is also implied that a sufficiently large scattering frequency in the pressure equations could disable magneto-immutability. It is essential that the δp_{\perp} and δp_{\parallel} evolve independently enough that δp_{\perp} can be suppressed by pressure balance to give $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u} \approx 0$, while $\hat{\mathbf{b}}\cdot\nabla\delta p_{\parallel}$ is suppressed by the u_{\parallel} momentum equation. This arrangement can still be preserved in the presence of a small amount of scattering ($\nu \ll k_{\parallel}v_A$), as the scattering term would remain next order in the $p_{\perp/\parallel}$ equations and the resultant reduced system would be unchanged. However, a scattering rate on the order of or larger than $k_{\parallel}v_A$ would force the pressures to evolve together, and the two magneto-immutability criteria would not be able to be met independently (see Appendix B for more). This point must be borne in mind when considering the effects of microinstabilities, because if their volume-filling fraction is large enough, then the consequent scattering may cause the plasma to become MHD-like.

3. Numerical simulations

In this section we describe simulations performed for the purpose of verifying the predictions of the reduced equations (2.16), organized in a manner that follows the layout of §2. As an overview, we begin by describing the numerical methods and the simulation setup in §3.1, followed in §3.2 by a summary of the key numerical diagnostics employed. Next, §3.3 and §3.4 provide evidence from the simulations that verify the assumptions that led to the ordering (2.4), namely, that (i) the vast majority of the plasma behaves effectively collisionlessly; (ii) that ion-acoustic waves are weakly mixed by the Alfvénic turbulence; (iii) and that density fluctuations are sufficiently small that they may be neglected. In §3.5, we re-confirm other aspects of the turbulence that were seen previously by Squire *et al.* (2023) and which influenced our ordering, such as the perpendicular balance of the thermal and magnetic pressures, and the critically balanced scaling of the Alfvénic fluctuations. Once the assumptions that led to (2.4) are confirmed, we test the equations (2.16) and their consequences in §3.6 and §3.7. This includes the misalignment of $\hat{\mathbf{b}}$ and the flow rate-of-strain (figure 8), the predicted scale independence of magneto-immutability, the reduction of $\nabla_{\parallel}\Delta p$, and the relative insensitivity of magneto-immutability to the magnitude of the heat flux. Lastly, in §3.8 we discuss the ability of microinstability-induced scattering (ν_{lim}) to interfere in the self-organization process.

3.1. Problem setup and method of solution

To assess the claims set forth in §2, we perform a suite of driven CGL-MHD turbulence simulations with a variety of parameters tuned to address each individual prediction or assumption. We employ a modified version of the ATHENA++ MHD code (Stone *et al.* 2020) that solves the CGL-MHD system (2.1) closed with the Landau-fluid heat fluxes (2.2). This code allows the parallel wavenumber $|k_{\parallel}|$ in the Landau-fluid heat fluxes (2.2) to be specified freely, and incorporates the effects of mirror and firehose instabilities (when excited) through a collisional closure. Unless capped by some maximal value, the Landau-fluid $q_{\perp/\parallel}$ could grow very large at small scales because of the numerical simplification that $|k_{\parallel}|$ is chosen to be constant. To prevent this from occurring, the heat fluxes are not allowed to surpass a maximal ‘free-streaming’ value of $\approx v_{\text{th}}p_{\perp/\parallel}$ (Hollweg 1974; Cowie

& McKee 1977). Exact details of how this limitation is implemented in the code can be found in §3.1.1 of Squire *et al.* (2023). The collisional microinstability closure uses a limiting scattering frequency ν_{lim} (also specified by the user), which activates only in regions of the domain that exceed the mirror ($\beta\Delta > 1$) or firehose ($\beta\Delta < -2$) thresholds (e.g., Sharma *et al.* 2006). Once activated, the pressures $p_{\perp/\parallel}$ are driven back towards the instability thresholds, rather than to complete pressure isotropy, at a rate set by ν_{lim} . The ability to choose $|k_{\parallel}|$ and ν_{lim} provides considerable freedom to explore how various collisionless effects change the behaviour of turbulence in this system. Further details about the CGL-MHD solver and microinstability closure can be found in Appendix A of Squire *et al.* (2023). As first described in §1.3, we supplement these ‘active- Δ ’ CGL-MHD simulations with a set of ‘passive- Δ ’ simulations, which are performed in isothermal MHD but evolve the pressure anisotropy passively using equations (2.1d) and (2.1e) given the simulated MHD fields. Such passive- Δ simulations are useful for comparing MHD-like turbulence with CGL-MHD turbulence.

Given that the particles’ Larmor scales are infinitesimally small in our model equations, the physical dimensions of our simulations are arbitrary. All simulations are performed in a fully periodic domain with dimensions $[L_x, L_y, L_z] = [1, 1, 2]$, where $L_z = L_{\parallel}$ is aligned with the background magnetic field $\mathbf{B}_0 = \hat{\mathbf{z}}$ and $L_x = L_y = L_{\perp}$ are the perpendicular dimensions. The ‘standard’ resolutions employed for these simulations are $n_{\perp} = 192$ and $n_{\parallel} = 2n_{\perp}$, however higher- and lower-resolution simulations are performed (and explicitly referred to) for the sake of convergence and scale-dependence tests (see figure 11). Although the initial magnetic field \mathbf{B}_0 remains the same in all simulations, the initially isotropic thermal pressure $p_0 = \beta_0 B_0^2 / 8\pi$ is varied by choosing β_0 to be either 1, 10, or 100. For $|k_{\parallel}|$ and ν_{lim} , we make use of a set of ‘standard’ values in which $|k_{\parallel}| = 4\pi/L_{\parallel}$ represents the wavenumber of compressive fluctuations near the outer scale and $\nu_{\text{lim}} = 10^{10}v_A/L_{\perp}$ yields a hard-wall limiter that prevents the pressure anisotropy from straying far beyond its microinstability thresholds. Specific instances in which $|k_{\parallel}|$ and ν_{lim} are modified from their standard values are noted on a case-by-case basis. All simulations are run until a final time of at least $t_f = 10L_{\perp}/v_A$ to ensure that a steady-state fluctuation level is achieved within the turbulence.

The turbulence is forced exclusively through sinusoidal perturbations to the flow velocity \mathbf{u} using an Ornstein–Uhlenbeck correlated process (Uhlenbeck & Ornstein 1930), the strength of which is input numerically as the rate of change of the total kinetic energy in the domain $d_t \mathcal{E}_K$. Our fiducial runs employ $d_t \mathcal{E}_K = 0.32\rho_0 v_A^2 L_{\perp}^3$ and a correlation time of $t_{\text{corr}} = L_{\parallel}/v_A$. This choice of t_{corr} assumes Alfvénically correlated forcing, and the energy injection rate corresponds to an outer-scale magnetic perturbation amplitude of $\delta B_{\perp/\parallel} \approx B_0/2$ in steady state. For all simulations, the sinusoidal wavenumbers at which we force are limited to $k \in 2\pi/L_{\parallel} \times [1, 3]$, over which the power distribution scales as k^{-2} . In this study, we frequently vary the mode of \mathbf{u} forcing between Alfvénic and Random forcing. Random forcing is as it sounds: \mathbf{u} is perturbed randomly in all directions without any special conditions relating to the directions of \mathbf{B} and \mathbf{k} (although it is still time correlated). In Alfvénic forcing however, we perturb only $\mathbf{u} \perp \hat{\mathbf{z}}$, and do so in a manner that enforces incompressibility, $\nabla_{\perp} \cdot \mathbf{u}_{\perp} = 0$, at the outer scale (below the outer scale, however, \mathbf{u} does become slightly compressible as a result of the nonlinear amplitudes).

3.2. Numerical diagnostics

To analyze the simulations introduced in §3.1, we make use of several numerical diagnostics, three of which we describe here because they are either used very frequently in our analysis or are particularly tailored to the subject of this work. These three are

the Fourier spectra, the Δp energy transfer function of Squire *et al.* (2023), and a novel scale-by-scale $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ alignment diagnostic based off of the work of St-Onge *et al.* (2020).

The most frequently used diagnostic is a bin-averaged Fourier spectrum, defined for a field χ as

$$E_\chi(k) = \frac{1}{\delta k} \sum_{|\mathbf{k}|=k} |\mathcal{F}[\chi](\mathbf{k})|^2, \quad (3.1)$$

where $\mathcal{F}[\chi](\mathbf{k})$ is the three-dimensional Fourier transform of χ , and $k = \sqrt{k_x^2 + k_y^2 + k_z^2}$ is the magnitude of \mathbf{k} (or $k_\perp = \sqrt{k_y^2 + k_z^2}$ for \mathbf{k}_\perp). The spectra are calculated for k bins, thus each bin contains contributions from several wavenumbers in the Fourier spectrum, and δk represents the bin width in k or k_\perp . Note that in figure 6(b), when calculating the spectra of $\nabla_\perp/\|u_\perp/\|$, the gradients of \mathbf{u} are calculated with respect to the *local* (in configuration space) field, rather than using the same efficiency-motivated $\mathbf{k}_\perp = k_x \hat{\mathbf{x}} + k_y \hat{\mathbf{y}}$ assumption that the spectra employ.

Energy transfer functions illustrate the amount of turbulent kinetic energy transferred into or out of a given k -shell as a result of specific interaction terms in the model equations (e.g., Grete *et al.* 2017). In this work, we are chiefly interested in the effects of pressure anisotropy on the turbulent cascade, thus the transfer function we make use of is designed to capture the energy removed from the flow \mathbf{u} at a given k_\perp due to the anisotropic pressure-stress $\nabla \cdot (\hat{\mathbf{b}}\hat{\mathbf{b}}\Delta p)$. We use the same definition given in Squire *et al.* (2023), where

$$\mathcal{T}_{\Delta p}(k_\perp) = \sum_{q_\perp} \int d^3x \langle \sqrt{\rho}\mathbf{u} \rangle_{k_\perp} \cdot \frac{\mathbf{B}}{\sqrt{4\pi\rho}} \cdot \nabla \left\langle \frac{\Delta p}{B^2} \mathbf{B} \right\rangle_{q_\perp}, \quad (3.2)$$

and $\langle \cdot \rangle_k$ indicates that the quantity has been Fourier transformed, filtered by wavenumber, and returned to real space. There are other definitions for this transfer function (see Arzamasskiy *et al.* 2023), however, as remarked by Squire *et al.* (2023), this particular formulation represents the Δp -stress as a damping of kinetic energy. This is well suited to the current study as we wish to understand how magneto-immutability permits the continuation of a turbulent cascade that would have otherwise been damped away near the viscous scale.

Finally, the $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ alignment diagnostic is a method for visualizing the effects of magneto-immutability's organization on the turbulent flow. It is designed to calculate the cosine of an angle θ that is representative of the alignment between the rate-of-strain tensor $\nabla\mathbf{u}$ and the (spatially) local magnetic field direction $\hat{\mathbf{b}}$, on a scale-by-scale basis in k_\perp . Initially used by St-Onge *et al.* (2020) without separation by scale for a study of the incompressible fluctuation dynamo under the action of Braginskii viscous stresses, it is based off of analytical theory from Kazantsev (1968). In the fluctuation dynamo, changes to the magnetic-field strength are mediated by $\nabla\mathbf{u}$ through its three eigenvalues, which are associated with field-line compressing motions, field-line stretching motions, and the incompressibility constraint (ensuring that the overall flow stretches as much as it compresses). The compressing and stretching motions result in changes to $|B|$, thus the angles between the compressing and stretching eigenvectors and the local magnetic-field direction dictate how efficient the flow is at changing the magnetic field strength.⁸ In practice, we must ensure that the eigenvectors and eigenvalues are real before dotting into $\hat{\mathbf{b}}$, therefore we diagonalize $(\nabla\mathbf{u} + \nabla\mathbf{u}^\top)/2$ with \top denoting the transpose. Once

⁸The incompressibility eigenvalue, sometimes called the ‘null’ eigenvalue, is distinguishable from the other two by having the smallest absolute magnitude.

the eigenvectors are obtained, the cosine of their angle with $\hat{\mathbf{b}}$ is calculated, yielding a compressing $\cos \theta$ and a stretching $\cos \theta$ at each grid point within the domain. These cosines are then compiled into a probability distribution $\mathcal{P}(|\cos \theta|)$, representing the likelihood that the alignment angle cosine takes on a specific value at a given point in time. We have found that the compressing and stretching distributions are qualitatively identical for all simulations performed in this work, which is expected for Alfvénic turbulence (as opposed to the fluctuation dynamo studied in St-Onge *et al.* 2020). For that reason, it suffices to only show one of the two, which we choose to be the stretching angle. To exhibit the $\mathcal{P}(|\cos \theta|)$ as a function of scale, we first Fourier transform \mathbf{u} , filter it using masks in $k_{\perp} = \sqrt{k_x^2 + k_y^2}$ bins, and then transform it back to real space before performing all of the aforementioned operations to obtain $\mathcal{P}(|\cos \theta|)$. The $\mathcal{P}(|\cos \theta|)$ are normalized to unity at each individual k_{\perp} , rather than \mathcal{P} being a distribution in both $|\cos \theta|$ and k , so as to avoid weighting larger wavenumbers less than smaller ones. When magneto-immutability is active, our expectation that the flow organizes to avoid changes in $|B|$ implies that we should see a small $|\cos \theta|$ between the eigenvectors of $\nabla \mathbf{u}$ and $\hat{\mathbf{b}}$, whereas a non-immutable cascade should have $|\cos \theta| \sim 1$. Therefore, we can compare this cosine between passive- and active- Δ simulations to detect whether the flow behaves in a different fashion so as to produce a cascade with minimal variation of $|B|$.

3.3. The effective equation of state

The most fundamental assumption underlying all of the conclusions of §2 is that the plasma behaves in a sufficiently collisionless manner to be described by the CGL-MHD equations. For this to be true, the portion of the plasma having pressure anisotropy that is microphysically unstable, and therefore subject to an effective collisionality associated with particle scattering off Larmor-scale fluctuations (i.e. that with $|\Delta| \gtrsim \beta^{-1}$), must constitute a small fraction of the total volume. In their initial study of magneto-immutability in the CGL-MHD system, Squire *et al.* (2023) found that magneto-immutability suppresses the overall level of fluctuations in the pressure anisotropy, thereby reducing the fraction of the plasma that is unstable. To establish the impact that this suppression has on the plasma's effective collisionality, we plot in figure 2(a) the relationship between the fluctuations in the parallel and perpendicular pressures and the (relatively small) fluctuations in the density, averaged in time over the interval $t v_A / L_{\perp} = [8, 10]$. In an MHD-like plasma, both δp_{\perp} and δp_{\parallel} would scale with $\delta \rho$ with a slope of $5/3$, the single adiabatic index characteristic of a collisional plasma (black dotted lines). Neither is particularly well aligned with the single adiabatic index slope, with δp_{\parallel} having a steeper slope, and δp_{\perp} a shallower one. The significant differences between the pressure slopes suggest that the $\delta p_{\perp/\parallel}$ evolve in an independent manner as is expected for a predominantly collisionless plasma.

This non-MHD equation of state is explained by figure 2(b), which compares the fraction of the domain that is unstable to micro-instabilities between active- (solid curves) and passive- Δ (dashed curves) simulations as a function of time. In the steady state of the passive- Δ simulations, nearly half of the domain lies beyond the instability thresholds, meaning that much of the plasma volume is experiencing the large scattering rate ν_{lim} , and a collisionless model would not be a good description for a significant portion of the turbulent dynamics. Conversely, the active- Δ simulations are rarely beyond 10% unstable, with an average falling closer to just 5% in steady state. Importantly, this reduction in the unstable fraction appears to be independent of β , suggesting that higher values of β are not likely to result in a significantly more collisional plasma. Infrequent excursions of the unstable fraction beyond 10% do occur in all simulations, which could

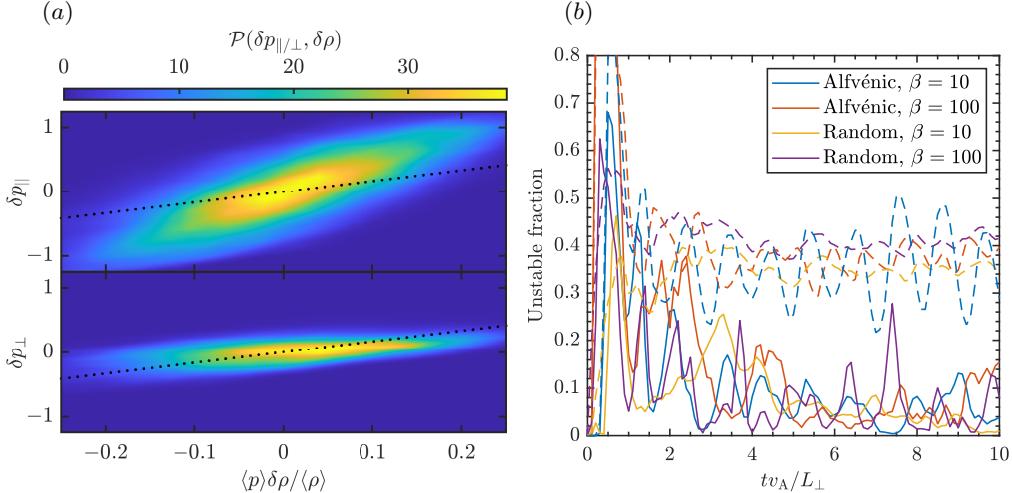


Figure 2: (a) Probability distributions of fluctuations in p_{\parallel} , p_{\perp} , and ρ , for $\beta = 10$ Alfvénically driven turbulence. The black dotted lines represent a slope of $5/3$, the expectation for single-adiabatic, collisional MHD. Neither δp_{\parallel} nor δp_{\perp} appear to align with the MHD prediction or with each other. (b) The fraction of the domain whose pressure anisotropy lies beyond the microinstability thresholds as a function of time for all active- Δ (solid) and passive- Δ (dashed) simulations. The passive- Δ simulations typically have $\sim 50\%$ of their volume unstable, whereas the unstable fraction of the active- Δ simulations typically remains below 10%.

either be the result of intermittency or the randomness of the forcing. Though generally short lived, these brief events do have the ability to change the statistics of the turbulence. Therefore, all statistical measurements we report are obtained by averaging over a time interval of no less than $2L_{\perp}/v_A$, taken beyond $t v_A / L_{\perp} = 6$ to ensure that the turbulent statistics have reached an approximate steady state.

3.4. Compressive forcing and ion-acoustic fluctuations

With the knowledge that the turbulence is taking place within a predominantly collisionless plasma, we surmise that the compressive wave fluctuations will behave not as collisional fast and slow magnetosonic modes, but rather as collisionless fast and ion-acoustic waves. As such, we evaluate the feasibility of an ion-acoustic wave cascade at high β , given our expectation that ion-acoustic waves are not effectively mixed and cascaded by the Alfvénic fluctuations in such plasmas. Conveniently, the only waves that actively make $\nabla \cdot \mathbf{u} \neq 0$ in a collisionless plasma are the fast wave and the ion-acoustic wave (non-propagating modes are pressure balanced and approximately incompressible). Thus, the presence of ion-acoustic waves in the inertial range of these turbulence simulations can be diagnosed by the compressive flow spectrum of $\hat{\mathbf{k}} \cdot \mathbf{u}_k$, which is provided in figures 3(a) and (b).

In figure 3(a), we explore the dependence of the compressive flow spectrum on the correlation time of the forcing. When the correlation time of the forcing is Alfvénic (as is expected for the astrophysical turbulence we are concerned with here), very little energy is present in compressive fluctuations, and the power-law index of the $\nabla \cdot \mathbf{u}$ spectrum is rather steep. This is true in both the randomly driven and Alfvénically driven setups, with very little difference between the two spectra. Only when the correlation

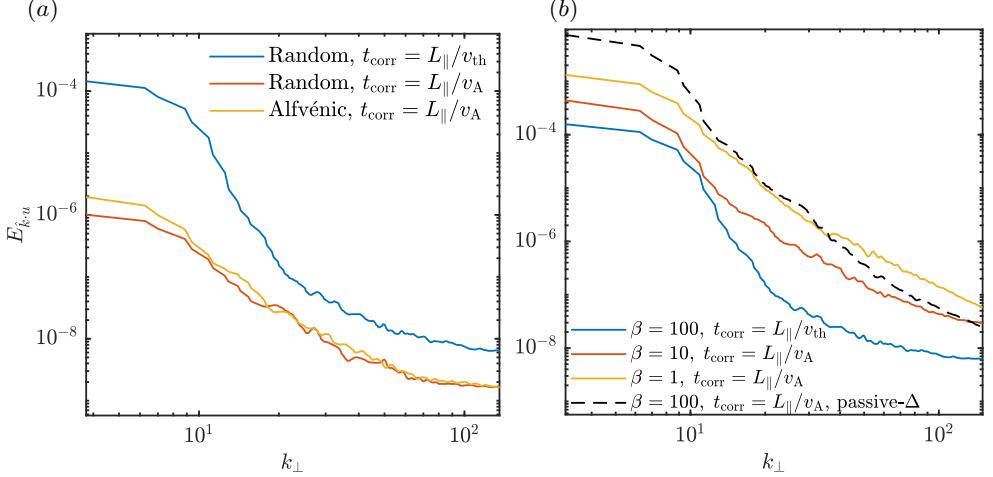


Figure 3: Spectra of compressive velocity fluctuations versus perpendicular wavenumber. (a) The compressive spectra at $\beta = 100$ are studied as a function of the type of forcing. When the correlation time is Alfvénic, random and Alfvénic forcing both produce similar results, with little energy in the compressive flow. If the correlation time is instead sonic, then an increase of two orders of magnitude is seen in the compressive flow energy. (b) Compressive spectra of randomly driven turbulence at different values of the plasma β . At $\beta = 100$ the forcing is sonic, leading to significant compressive fluctuations at the outer scale, but the spectrum is extremely steep. As β is reduced, the spectra become decreasingly steep, qualitatively approaching the passive- Δ result (black dashed).

time of the forcing is decreased to be sonic ($t_{\text{corr}} \sim L_{\parallel}/v_{\text{th}}$), can a substantial amount of energy enter into compressive fluctuations. This is likely because, in Alfvénically correlated forcing, the timescale associated with the randomization of the forcing is too slow for the forced wavenumbers to significantly drive ion-acoustic waves. Indeed, the β dependence of this conclusion is captured by figure 3(b), where we plot these compressive spectra for only randomly driven simulations at $\beta = 1, 10$ and 100 , with the last being sonically correlated. In effect, the $\beta = 1$ simulation is also sonically correlated because $v_{\text{th}} = v_{\text{A}}$, but because a large difference between the Alfvén and ion-acoustic wave speeds is not present, the mixing of ion-acoustic waves by Alfvén waves is stronger, bringing the spectrum much closer to that of the dashed passive- Δ simulation, performed at $\beta = 100$. The spectrum of the $\beta = 100$ CGL-MHD simulation, although sonically correlated to generate substantial outer-scale compressive fluctuations, has an extremely steep spectrum. This steep spectrum implies that very little of the compressive mode energy driven at the outer scale penetrate into the inertial range, supporting our claim that ion-acoustic waves may be ordered out of the dynamics. The spectrum does eventually become less steep, however this is likely due to nonlinearities from Alfvénic fluctuations, given the very small amount of energy contained in $E_{\hat{k} \cdot u}$ at such high k .

The consequences of the lack of compressive modes in the turbulent inertial range can be seen in the statistics of the density fluctuations. Figure 4(a) shows the probability distribution of various values of the density within the simulation domain for each of the Alfvénically correlated simulations, both active (solid) and passive (dashed). As β is increased, the density takes on fewer values that deviate significantly from the background, with the Alfvénically driven simulations being only slightly narrower than the randomly driven runs. These narrow distributions are likely a result of the Mach

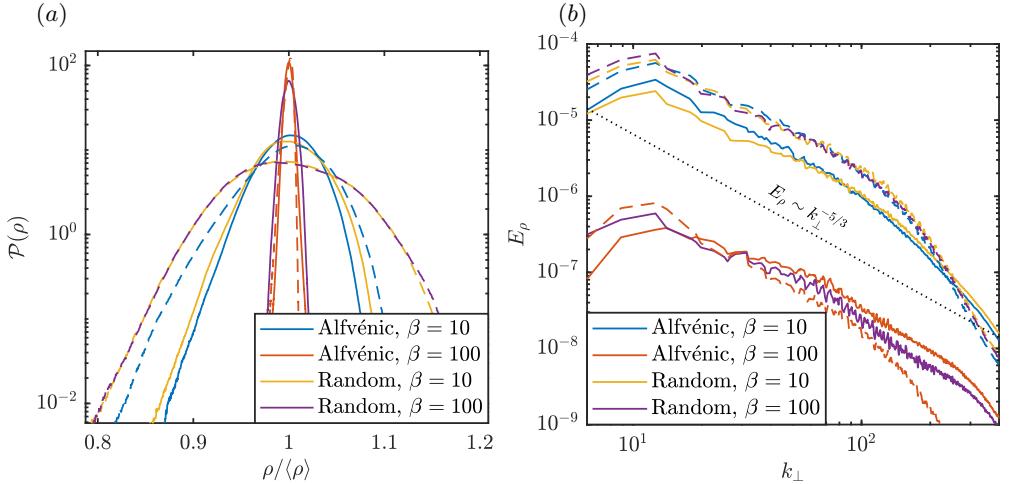


Figure 4: (a) Probability distribution of ρ within the domain for all Alfvénically correlated simulations, active (solid) and passive (dashed). At higher β , the decrease in Mach number leads to weaker density fluctuations being driven at the outer scale, and less overall variation in ρ . Randomly driven simulations exhibit somewhat increased variation in ρ , but the dominant parameter is β . (b) Density fluctuation spectra for all Alfvénically correlated simulations. While the overall amplitudes are decreased with increasing β , the spectra remain strong with spectral indices near $-5/3$. This may indicate that the density is passively advected in the absence of fast and ion-acoustic wave cascades. The apparent break in power law behavior of the $\beta = 100$ spectrum is a result of our particular choice for ν_{lim} , which is discussed in more detail within §3.8. In both panels (a) and (b), most all of the passive simulations exhibit larger density fluctuations, with only the $\beta = 100$ Alfvénically driven passive run having $\delta\rho$ as small as its active counterpart.

number becoming smaller, given that our simulations drive an approximately fixed amplitude of u/v_A at the outer scale. How the level of these fluctuations depends on scale is more clear in the spectrum of density fluctuations, figure 4(b). Interestingly, even though these fluctuations are clearly $\lesssim \epsilon^2 \rho_0$, making them too small to affect the dynamics to leading order, and ion-acoustic wave mixing by the Alfvénic cascade is weak, their spectra appear to follow a near $-5/3$ power law, as given by the dotted line. We expect that this is accounted for by the presence of kinetic entropy and non-propagating modes, which, given that they have no real frequency, should be well mixed by Alfvénic fluctuations (see §2.2 and §2.4). As they do not feed back on the overall dynamics due to their small amplitudes, their spectrum is probably captured by passive advection ($d_t \delta\rho = 0$), hence the $-5/3$ index.⁹ The $\beta = 100$ simulations do appear to show a break from the single power law behavior of the other simulations, although this is likely an effect of our choice of ν_{lim} , which is discussed in more detail within §3.8. Note that in figures 4(a) and (b), all but one of the passive simulations possess density fluctuations larger than their active counterparts; only the $\beta = 100$ Alfvénically

⁹Although none of the simulations we performed produced density fluctuations exceeding $\delta\rho \sim \epsilon^2 \rho_0$, it would certainly be possible to achieve $\delta\rho \sim \epsilon^{3/2} \rho_0$ or $\epsilon \rho_0$ by explicitly seeding larger amplitude kinetic entropy or non-propagating modes at the outer scale. To explore this possibility, the consequences of such significant density fluctuations for immutability are discussed in depth within Appendix C.

driven passive simulation possesses $\delta\rho$ as small as its active counterpart. The fact that the choice between random and Alfvénic forcing makes a difference in the passive simulations at $\beta = 100$, but not the active simulations, highlights the separation of compressive timescales between MHD and CGL. In CGL, the type of forcing makes little difference unless it is sonically correlated so that it may excite ion-acoustic modes. On the other hand, the correlation time need not be sonic to excite MHD slow modes, so random forcing can in fact drive larger density fluctuations in the passive simulation.

3.5. Comparison with previous work

Certain aspects of our ordering (2.4) and reduced equations (2.16) explain features of high- β CGL turbulence that were already observed in the Alfvénically driven simulations of Squire *et al.* (2023), but not yet fully understood. Here we reproduce some of the key results of Squire *et al.* (2023), confirming that our simulations explore the same effects and qualifying the extent to which the compressibility of forcing matters in a system where ion-acoustic and fast modes are not effectively cascaded. All simulations shown within this section are performed at $\beta = 10$ with Alfvénically correlated forcing and standard resolution (as defined in §3.1).

In figure 5(a), the kinetic and magnetic energy spectra from the randomly (solid) and Alfvénically (dash-dotted) forced simulations reveal inertial ranges that are close to the $-5/3$ power law expected in MHD turbulence. Although the individual spectral slopes deviate very slightly above or below this exact value, overall there is little qualitative difference between the turbulence resulting from the two modes of forcing. This indicates that the pressure-anisotropy stress does not effectively remove energy from the cascade, at least not to the extent that would naïvely be expected when $\beta\Delta \sim 1$ (Squire *et al.* 2023). Upon careful inspection, the power-law index of the kinetic energy spectrum in the randomly driven simulation is steeper than that of the Alfvénically driven simulation, although the difference is small, likely originating from next-order effects not captured by (2.16). In figure 5(b), the field-perpendicular and parallel scales of the Alfvénic flow and magnetic-field perturbations are given. For \mathbf{u}_\perp , the flow is calculated with respect to the exact local magnetic field, however \mathbf{B}_\perp is simply $B_y\hat{\mathbf{y}} + B_z\hat{\mathbf{z}}$. The dash-dotted and solid coloured lines once again represent the Alfvénic and randomly driven simulations, while the black dotted line represents the $(l_\parallel/l_0) \sim (l_\perp/l_0)^{2/3}$ relationship predicted for critically balanced MHD turbulence. This appears to indicate that cascades of δB_\perp and u_\perp are critical balanced, another foundational assumption of our ordering (2.4). Details of how these characteristic eddy sizes are measured can be found within §3.2.2 of Squire *et al.* (2023) or §6.3 of Cho & Lazarian (2009).

The most easily verified predictions of (2.16) are that the parallel pressure dominates the pressure anisotropy, and that this results from the perpendicular pressure being balanced by magnetic pressure. In figure 6(a), the thermal and magnetic pressure spectra are shown, depicting both of these features.¹⁰ The perpendicular pressure, although larger at the outer scale, approaches and ultimately matches the magnetic pressure in the inertial range, as predicted by (2.10); as stipulated in the reduced model, this empirical result requires the leading-order δp_\perp perturbation to be $\mathcal{O}(\epsilon^2)$. By contrast, fluctuations in p_\parallel are much larger in amplitude, and all appear to follow the spectral index of $-5/3$ typical of passive advection, with some minor variation. The p_\parallel spectrum appears to have a slightly shorter inertial range than the p_\perp and B^2 spectra, a feature which may be due

¹⁰Note that the magnetic pressure spectrum is the spectrum of B^2 having units of energy squared times length, not to be confused with the magnetic fluctuation spectrum, which is the spectrum of B , and has units of energy times length.

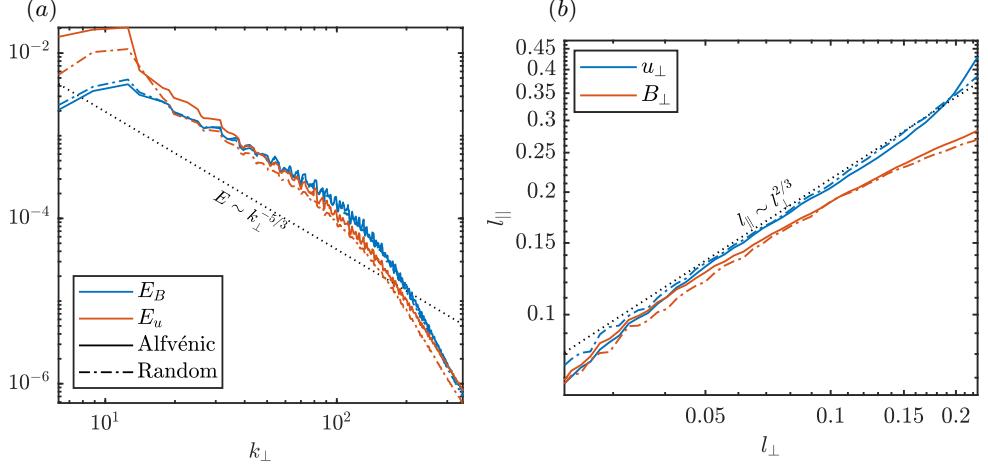


Figure 5: (a) Kinetic energy spectra of $\beta = 10$, Alfvénically correlated simulations with random (solid) and Alfvénic (dash-dotted) forcing. Only slight qualitative differences are visible between each type of forcing, reflecting the lack of a strong ion-acoustic cascade regardless of forcing. (b) Characteristic turbulent eddy sizes along and across the local magnetic field for Alfvénic variables. Both the randomly and Alfvénically driven simulations closely follow the scaling relationship predicted by critical balance in standard MHD, $(l_\parallel / l_0) \sim (l_\perp / l_0)^{2/3}$, represented by the black dotted line.

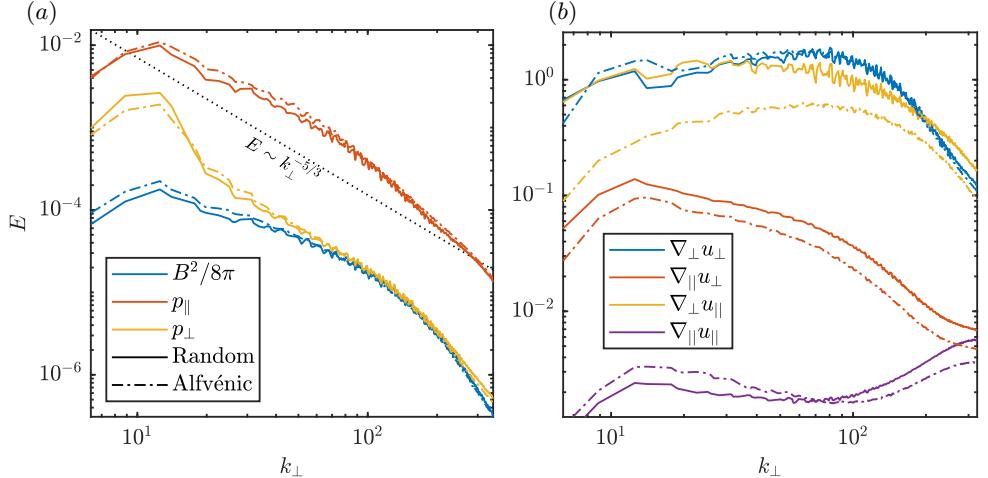


Figure 6: (a) Spectra of the perpendicular-thermal, parallel-thermal, and magnetic pressures for randomly (solid) and Alfvénically (dash-dotted) driven $\beta = 10$ simulations. The difference between p_\parallel and p_\perp , combined with the rough equivalence of B^2 and p_\perp , reflects perpendicular pressure balance (Squire *et al.* 2023). (b) The rate-of-strain spectra, showing suppression of $\nabla_\parallel u_\parallel$ as predicted by (2.16). The most noticeable difference between forcing modes occurs here in the spectra of $\nabla_\perp u_\parallel$.

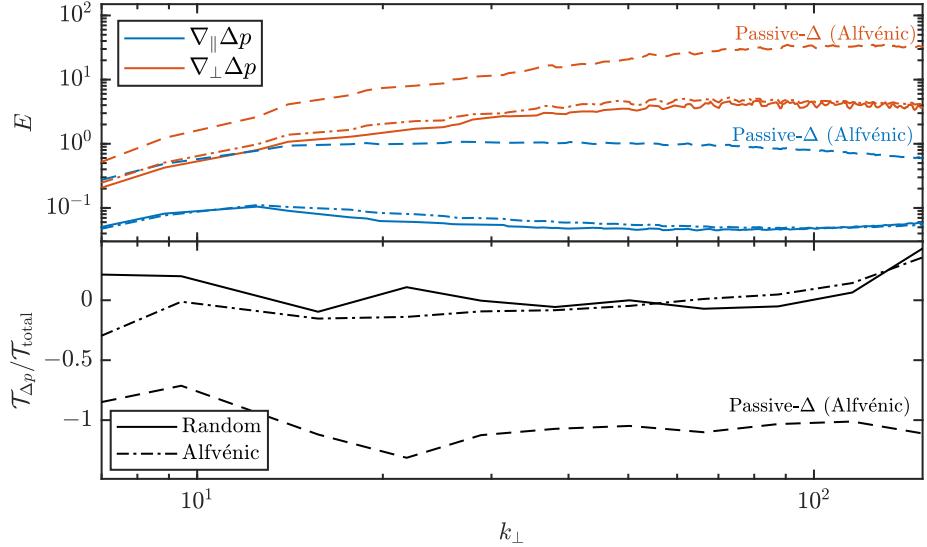


Figure 7: Top panel: Spectra of the field-parallel and -perpendicular gradients of the pressure anisotropy for $\beta = 10$ simulations that are either active and Alfvénically forced, active and Randomly forced, or passive and Alfvénically forced. Both of the active- Δ simulations show a significant decrease in $\nabla_{\parallel} \Delta p$ with respect to the passive simulation, as well as greater difference between $\nabla_{\perp} \Delta p$ and $\nabla_{\parallel} \Delta p$. Bottom panel: Transfer rate into and out of the turbulent flow due to the anisotropic pressure-stress, normalized to the total cascade rate, for the same simulations as the top panel. Both active simulations show significant suppression of the pressure stress resulting from the reduction of $\nabla_{\parallel} \Delta p$ seen in the top panel. The passive simulation predicts a stress that is capable of damping turbulent motions entirely, since $\mathcal{T}_{\Delta p} \sim \mathcal{T}_{\text{total}}$ across the full inertial range.

to the fact that the heat fluxes are stronger for p_{\parallel} than for p_{\perp} . Although the $q_{\perp/\parallel}$ are nominally ordered out by the reduction in $\nabla_{\parallel} \delta T_{\perp/\parallel}$, it is possible that as the turbulence approaches grid scales magneto-immutability weakens somewhat from the effects of finite resolution. We find that in higher resolution simulations this steepening in $E_{p_{\parallel}}$ trends with the grid scale, therefore it likely a numerical artifact and has no impact on our physical understanding of the p_{\parallel} cascade. Figure 6(b) depicts the rate-of-strain spectra of the turbulent flow, broken up into field-perpendicular and parallel gradients, as well as perpendicular and parallel flows. As expected from (2.16), parallel gradients of u_{\parallel} are dramatically suppressed with respect to all other elements of the rate-of-strain tensor (Squire *et al.* 2023). Importantly, this includes $\nabla_{\perp} u_{\parallel}$, which emphasizes that it is not simply that u_{\parallel} itself is being reduced, only its gradients along the magnetic field. Between figures 5 and 6, the most apparent difference between random and Alfvénic driving exists in the spectra of $\nabla_{\perp} u_{\parallel}$. All other gradients of the flow velocity are similar between each mode of forcing. However, the randomly driven simulation features a spectrum of $\nabla_{\perp} u_{\parallel}$ that is roughly constant with k_{\perp} , unlike in the Alfvénically driven run where it is increasing with k_{\perp} , in accordance with standard MHD scalings (Squire *et al.* 2023). This highlights the fact that, even though the forcing has a clear effect on u_{\parallel} , it has no apparent effect on the suppression of $\nabla_{\parallel} u_{\parallel}$, given that it is equally suppressed for both random and Alfvénic forcing.

One of the most significant consequences of magneto-immutability is the reduction of the anisotropic pressure stress via suppression of $\nabla_{\parallel} \Delta p$. To compare the degree to which

this is achieved in random versus Alfvénically forced turbulence, we plot the $\nabla_{\perp/\parallel}\Delta p$ spectra and the Δp transfer function $\mathcal{T}_{\Delta p}(k_{\perp})$ of our $\beta = 10$ simulations in figure 7. In the top panel, the active- Δ simulations exhibit clear suppression of $\nabla_{\parallel}\Delta p$, especially with respect to $\nabla_{\parallel}\Delta p$ from the Alfvénically driven, $\beta = 10$ passive run. That being said, the difference between the ratio $\nabla_{\parallel}\Delta p/\nabla_{\perp}\Delta p$ in the passive and active runs, which is only a factor of ≈ 2 , is more subtle. Although the reduced equations only explicitly constrain the parallel gradient of Δp through (2.12), the overall production of pressure anisotropy is also reduced by suppressing changes in the magnetic-field strength. As a result, perpendicular gradients of Δp are also smaller simply because of the smaller overall magnitude of Δp fluctuations.

The results of $\nabla_{\parallel}\Delta p$ suppression are seen in the energy transfer due to the pressure anisotropy $\mathcal{T}_{\Delta p}(k_{\perp})$ on the bottom panel of figure 7 (equation (3.2)). $\mathcal{T}_{\Delta p}(k_{\perp})$ is normalized to the total energy transfer rate, which we approximate by the Kolmogorov cascade rate $\mathcal{E}_K/\tau_{k_0} \sim \mathcal{E}_K(2\pi u_{\text{rms}}/L_{\perp})$ where τ_{k_0} is the outer-scale turnover time. For the passive run, $\mathcal{T}_{\Delta p}/\mathcal{T}_{\text{total}}$ being ~ 1 means that if the turbulence cascaded according to MHD, the anisotropic pressure-stress would remove all of the energy from the turbulent fluctuations, ending the inertial range as soon as it begins. However with immutability, this pressure stress is much smaller, permitting the inertial range to be relatively conservative. Overall, very little difference exists between the randomly and Alfvénically driven simulations.

3.6. Organizing the turbulence

A key aspect of why magneto-immutability constitutes self-organization is that (2.4) does not explicitly order $\nabla_{\parallel}u_{\parallel}$ as small; rather it becomes necessary for $\nabla_{\parallel}u_{\parallel}$ to be suppressed in order to satisfy both perpendicular pressure balance and μ conservation. Given that $\hat{\mathbf{b}}$ is always $\mathcal{O}(1)$, and figure 6(b) demonstrates that u_{\parallel} is significant, it can only be the alignment angle between the magnetic field and the flow rate-of-strain tensor that suppresses variations in the magnetic-field strength. To probe this organization, we analyze our simulations using the alignment angle diagnostic described in §3.2. Examples of the scale-dependent $\mathcal{P}(|\cos\theta|)$ for both passive and active Alfvénically driven simulations at $\beta = 10$ are shown in figure 8. Recall that the dynamics in the passive simulations are simply isothermal MHD, and so the passively evolved pressure anisotropy does not affect $\mathcal{P}(|\cos\theta|)$. In figure 8(a), it is clear that the peak of the distribution from the active run closely follows $\cos\theta \approx 0$ until dissipation scales are reached at large k_{\perp} . This indicates a near-complete misalignment of the rate-of-strain eigenvectors and $\hat{\mathbf{b}}$. By contrast, the passive (MHD) run distribution in figure 8(b) is peaked near $\cos\theta \approx 0.6$, which yields an $\mathcal{O}(1)$ dot product between $\hat{\mathbf{b}}$ and $\nabla\mathbf{u}$, thereby permitting more change to $|B|$. In figure 9, we show the dominant alignment angle as a function of k_{\perp} for all active (solid) and passive (dashed) simulations by tracking the peak of $\mathcal{P}(|\cos\theta|)$ across k_{\perp} . Without exception, all active simulations exhibit $\cos\theta$ near 0–0.2, while all passive simulations fall between $\cos\theta \approx 0.5$ and 0.7.

Many of the diagnostics shown elsewhere within this work and in Squire *et al.* (2023) demonstrate magneto-immutability by comparing with passive simulations that possess exactly identical parameters, such as forcing and resolution. Unfortunately, for particle-in-cell simulations (or reality), there exist no corresponding passive- Δ simulations that can so accurately represent the MHD equivalent turbulence. With the $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ alignment diagnostic, however, there is no need to compare with an exactly analogous passive- Δ simulation to determine whether magneto-immutability is at work, as small values of $|\cos\theta|$ alone are sufficient indication. This allows us to search for magneto-immutability within a kinetic framework that self-consistently determines the heat fluxes, which the

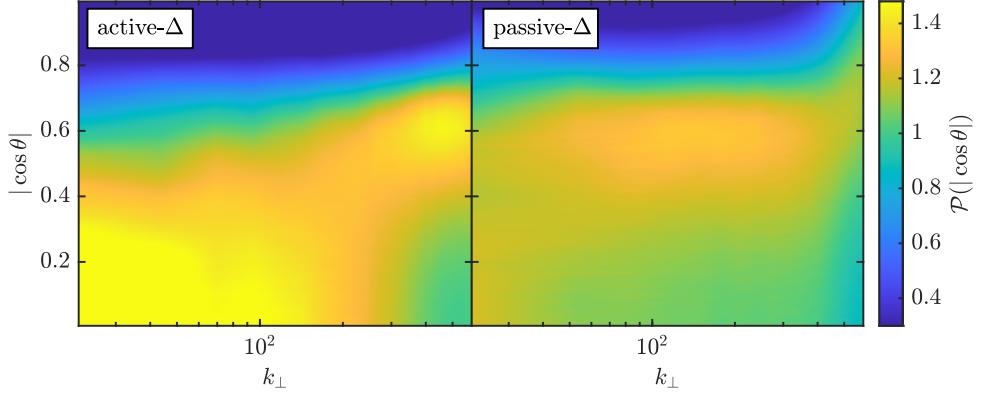


Figure 8: Probability distributions of the cosine of the angle between the rate-of-strain stretching eigenvector and the local magnetic-field direction, for $\beta = 10$, Alfvénically driven turbulence. Distributions are calculated and normalized individually within each k_{\perp} bin. The distributions of the compressing eigenvector cosines are qualitatively indistinguishable from their stretching counterparts for both the active and passive runs. The active- Δ simulation yields $\cos \theta \approx 0$ throughout the inertial range, indicating that motions in the flow that would normally lead to magnetic-field growth are misaligned with \hat{b} , rendering them incapable of significantly perturbing $|B|$. By contrast, the passive- Δ simulation has its peak probability around $\cos \theta \approx 0.6$, which produces an $\mathcal{O}(1)$ dot product between the two vectors and allows significant changes in the magnetic-field strength.

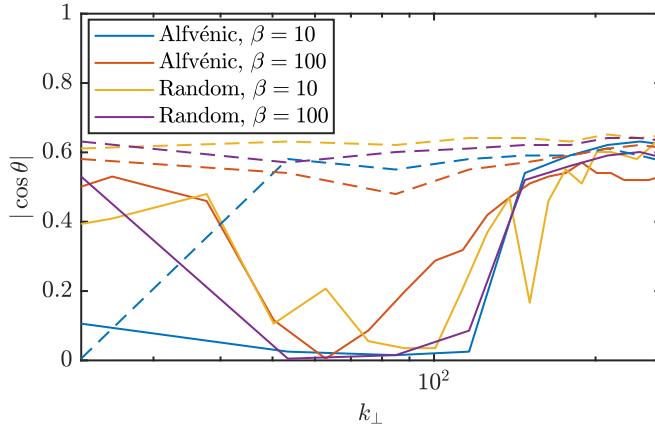


Figure 9: Peak cosines of the $|\cos \theta|$ probability distributions versus k_{\perp} for various Alfvénically correlated simulations, both passive (dashed) and active- Δ (solid). All active simulations exhibit large misalignment in the inertial range, while all passive simulations show greater alignment, permitting larger changes to $|B|$. There is no clear trend with β or forcing mode in the scale at which the active-simulation misalignment ends and begins to resemble the passive simulations. For most this occurs near numerical dissipation scales, which is equivalent for all simulations shown. The lack of misalignment at larger scales is likely due to the forcing not respecting magneto-immutability.

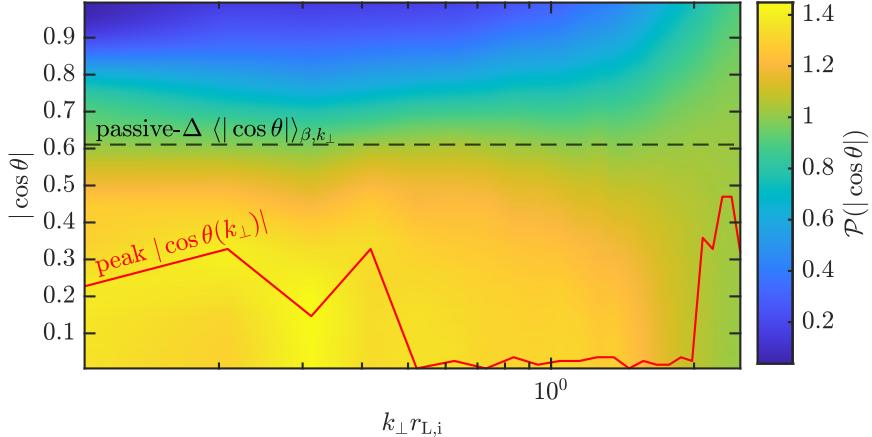


Figure 10: Distribution of $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ alignment angle θ as a function of k_{\perp} for the $\beta = 16$, Alfvénically driven and correlated simulation of Arzamasskiy *et al.* (2023). This simulation was performed within a $(120.5r_{L,i})^2 \times 241r_{L,i}$ box, at a resolution of $384^2 \times 768$. The peak alignment angle cosine (red line) is seen to be very close to 0 in the inertial range, in contrast to the average cosine measured from our passive- Δ (isothermal MHD) simulations (black dashed line). As with our CGL simulations, this misalignment is weaker near the outer scale, where immutability struggles to coexist with the forcing.

Landau-fluid heat fluxes can only approximate. In figure 10, we show the calculation of $|\cos \theta|$ for a $\beta = 16$, Alfvénically driven hybrid-kinetic turbulence simulation initially described within Arzamasskiy *et al.* (2023). The code used to perform this simulation was the `Pegasus++` hybrid kinetic-ion fluid-electron particle-in-cell code (Kunz *et al.* 2014). The simulation shown resolves scales both above and below the ion Larmor scale, additionally incorporating a nonzero electron temperature $T_e = T_i$. The peak of the $|\cos \theta|$ distribution is traced by the red line, while a comparison with the inertial-range average of all of our passive simulations is provided by the dashed black line. The misalignment is distinctly stronger for the hybrid-kinetic simulation than for the passive- Δ simulations, lasting until just past the ion Larmor scale. In Arzamasskiy *et al.* (2023), the authors detected significant viscous dissipation near the outer scale, suggesting that immutability is unable to coexist with the forcing, a fact reflected in the imperfect misalignment demonstrated at the largest scales within figure 10. This outer-scale heating can also be seen to some degree in figure 7, as well as in most of the simulations of Squire *et al.* (2023) (who used a different forcing scheme to that used here).

In the CGL simulations depicted by figures 8 and 9, there is no obvious trend of the scale (in k_{\perp}) at which immutability ceases to affect $\cos \theta$; most simulations trend toward MHD-like alignment at roughly the same scale. To understand the meaning of this scale, we compare the alignment angle cosines with the kinetic energy spectra of the $\beta = 10$, Alfvénically driven simulation at several different resolutions in figure 11. The resolutions are given for the coordinates perpendicular to \mathbf{B}_0 of each run, with $n_{\parallel} = 2n_{\perp}$. A clear trend with resolution exists, meaning that the alignment angle cosine only changes when the flow reaches dissipation scales and the turbulence no longer follows the ordering (2.4). Notably, this misalignment persists to smaller scales than those satisfying $\delta B_{\parallel}/B_0 \gtrsim \beta^{-1}$, where the change in the magnetic field strength due to individual turbulent fluctuations is large enough to generate $\beta\Delta \gtrsim 1$ (in the absence of immutability). This implies that, at least as far as the resolutions we can probe go, magneto-immutability is not an outer-

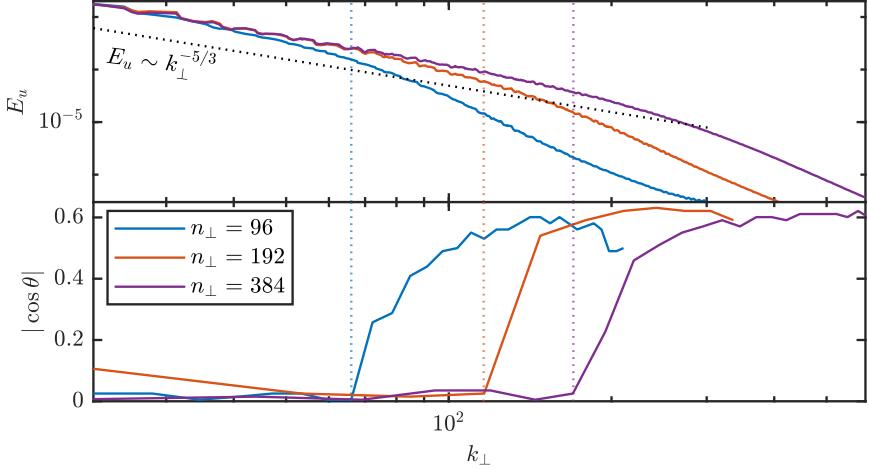


Figure 11: Top panel: Kinetic energy spectra for Alfvénically driven $\beta = 10$ simulations at 3 different resolutions across \mathbf{B}_0 , all with $n_{\parallel} = 2n_{\perp}$. The dotted black line represents a $k_{\perp}^{-5/3}$ power law. Bottom panel: Alignment angle cosines as a function of k_{\perp} for each of the three resolutions, with vertical dotted lines marking the transition away from misalignment of $\nabla \mathbf{u}$ with $\hat{\mathbf{b}}$. It appears that the transition trends with the dissipation scale of the turbulence, likely as a result of departure from the ordering (2.4).

scale effect but rather persists throughout the inertial range regardless of how small $\beta \Delta$ is at a given k_{\perp} , albeit in the somewhat modified sense discussed in §2.5.

3.7. The role of heat fluxes

Not only is the Δp stress suppressed by the reduction of $\nabla_{\parallel} \delta T_{\parallel}$ and $\nabla_{\parallel} \Delta p$, the heat fluxes are as well. Importantly, as predicted by (2.16), the $\nabla_{\parallel} \delta T_{\perp/\parallel}$ reductions are not diffusive (i.e., caused by the strong heat fluxes at high- β), but rather dynamical, originating from the momentum equation. For that reason, we should be able to artificially enhance or suppress the heat fluxes in our simulations – for example, by adjusting the parameter $|k_{\parallel}|$ – and observe little effect on the signatures of magneto-immutability. In figure 12 we gather plots of $\cos \theta$ and the $\nabla_{\parallel} \Delta p$ spectra, expressed as a function of k_{\perp} , for simulations where $|k_{\parallel}|$ is increased or decreased by a factor of 100. In all runs depicted, $\beta = 10$ and the forcing is Alfvénic, with a corresponding passive simulation given as an orange dotted line, given that it was performed with the nominal Landau wavenumber of $|k_{\parallel}| = 2\pi$. The simulation with $|k_{\parallel}|$ increased by 100 (purple curve in each plot) is effectively double-adiabatic MHD because of how weak the heat fluxes are.

In the top panel of figure 12, little variation is seen in the alignment angle cosine as a function of k_{\perp} , with only the double-adiabatic run being misaligned slightly further in k_{\perp} than the others. Considering that the heat fluxes are 10^4 times stronger for the blue curve than the purple curve, the extension of misalignment by less than a factor of 1.5 in k_{\perp} is an extremely small difference. In the bottom panel, the field-parallel gradients of the pressure anisotropy are shown, featuring strong suppression of $\nabla_{\parallel} \Delta p$ in the active- Δ simulations as compared to the passive run. Although the active simulations are all more similar to each other than the passive one, parallel gradients of the double adiabatic run are moderately larger, by a factor of ~ 2 . The comparable reduction of $\nabla_{\parallel} \Delta p$ for all values of $|k_{\parallel}|$ shown implies that the heat fluxes have little effect on whether or not immutability is able to effectively avoid Δp -stress. Note that the results of figure 10

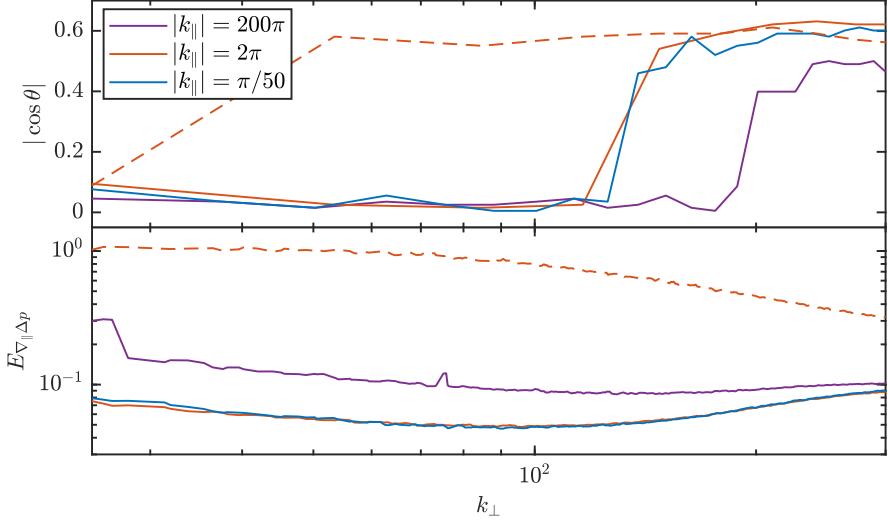


Figure 12: Top panel: Alignment angle cosine versus k_\perp for $\beta = 10$, Alfvénically driven simulations with $|k_\parallel|$ in (2.2) increased (purple), unmodified (orange), and decreased (blue) by a factor of 100. Given a change in heat flux strength of 10^4 , the difference between strong (blue) and weak (purple) heat fluxes appears to have little effect on θ . An orange dashed line represents the passive- Δ equivalent of the unmodified heat flux run. Bottom panel: The spectra of $\nabla_\parallel \Delta p$ for each run. The large suppression of parallel gradients in the pressure anisotropy reflect the suppression of $\nabla_\parallel T_{\parallel/\perp}$, as predicted by (2.16). This limits the ability of heat fluxes to play a role in the turbulence, and importantly does not rely on their strength, with even the $k_\parallel = 200\pi$ simulation, which is effectively double-adiabatic (purple), showing a comparable reduction. The reduced $\nabla_\parallel \Delta p$ also implies that the heat fluxes do not interfere in the avoidance of significant Δp -stress.

suggest that this heat flux suppression in figure 12 extends to the kinetic simulations of Arzamasskiy *et al.* (2023), where the $q_{\perp/\parallel}$ are not approximated via the Landau-fluid form. In order for $\hat{b} \cdot \nabla u_\parallel$ to be relegated to next order in (2.15), the heat fluxes must be negligible, a fact also made clear by (A 5) of Appendix A. Therefore, to observe the kind of misalignment seen in figure 10, it must be that the fully kinetic heat fluxes, like the Landau-fluid approximations, have little effect on the leading-order dynamics. Unfortunately, unless a closure for $q_{\perp/\parallel}$ is assumed this cannot be proven in general. Nonetheless, these results still further the favourable comparison between Landau-fluid CGL and kinetic simulations of high- β collisionless turbulence.

3.8. The role of micro-instabilities

The only portions of this study that are not described by our ordered equations 2.16 are the role of microinstability limiters. Within the CGL model, dependence on this physics arises via our choice of ν_{lim} . As discussed in §3.3, the sizes and distributions of micro-unstable patches are highly intermittent and difficult to predict; however, by varying ν_{lim} we can obtain useful information about how they interact with the turbulence. For all other simulations outside of this section, the instability scattering rate is fixed to $\nu_{\text{lim}} = 10^{10} v_A / L_\perp$, essentially providing a hard wall on the maximum possible pressure anisotropy. Given that the microinstabilities being accounted for grow at $r_{L,i}$ scales,

which are formally zero in the CGL-MHD system, it seems reasonable that they might scatter particles at a rate much faster than any of the dynamics being studied. For this reason, and for simplicity, such ‘hard-wall’ limiters have been employed frequently in past studies of pressure-anisotropic turbulence far above $r_{L,i}$ scales (see *e.g.*, [Sharma et al. 2006](#); [Squire et al. 2023](#); [Squire et al. 2019](#); [Santos-Lima et al. 2014](#)). However, studies focused on the mirror and firehose instabilities in the absence of background turbulence have found that the scattering rate induced typically follows the relationship $\nu \sim \beta \hat{b} \cdot \nabla \mathbf{u}$ ([Kunz et al. 2014](#); [Riquelme et al. 2018](#)), which in a turbulent environment takes on different values at different scales. Unfortunately, this is expensive to implement directly as it requires measuring the flow shear locally at each time step. To investigate the consequences of different choices for ν_{lim} , we perform Alfvénically driven simulations at $\beta = 100$, where the relative forcing of Δp is larger than at $\beta = 10$, and vary ν_{lim} between 20, 200, and $10^{10} v_A / L_\perp$.

The effects of these variations in ν_{lim} on the kinetic spectra are shown in figure 13(a). Interestingly simulations with lower scattering rates, where the pressure anisotropy is allowed to make larger excursions beyond the microinstability thresholds, lead to less steep spectral indices that are much more in line with the $-5/3$ expectation of our ordering and a conservative cascade. On the other hand, the simulation with a hard-wall limiter ν_{lim} , as used in all other simulations throughout this work, exhibits a spectrum that is slightly steeper and appears to be dissipated earlier in k_\perp . This steepening implies an increase in damping from the Δp stress, as some of the energy in the turbulent motions is being removed from the cascade.¹¹ The reason for this steepening and apparent dissipation can be seen in figures 13(b) and (c), which present PDFs of the measured values of $\beta \Delta$ and $\hat{b} \cdot \nabla \mathbf{u}$ in each run. These diagnostics essentially represent how effective magneto-immutability is at regulating the overall magnitude and production rate of pressure anisotropy. In the hard-wall-limited passive simulation (purple dashed), the distribution of $\hat{b} \cdot \nabla \mathbf{u}$ is less peaked in the stable regions ($-2 < \beta \Delta < 1$) than any of the active runs, with values of $\beta \Delta$ that are effectively pinned to the mirror and firehose limiters, and no local peak in $\beta \Delta$ between the limiters. The hard-wall-limited active simulation is the least peaked of the active simulations at $\hat{b} \cdot \nabla \mathbf{u} = 0$, also having no local peak in $\beta \Delta$ between the instability limiters. By comparison, lower values of ν_{lim} appear to better focus the distribution of $\hat{b} \cdot \nabla \mathbf{u}$ around 0, and at least in the case of $\nu_{\text{lim}} = 200 v_A / L_\perp$, maximize the proportion of $\beta \Delta$ that lies between the instability limiters.

We hypothesize that this weakening of magneto-immutability occurs because whenever a strong limiter scattering rate ν_{lim} is activated, the magneto-immutable orderings of (2.1e) and (2.1d) are broken. In an MHD plasma with ν^{-1} being much less than any dynamical timescale, this does not lead to any dissipation because the scattering drives the anisotropy to zero, and hence there is no anisotropic pressure stress. However, when the scattering is induced by microinstabilities, it only drives the anisotropy to $\beta \Delta \sim 1$ levels. This is visible in the hard-wall-limited curve of figure 13(b), where the distribution of $\beta \Delta$ is sharply peaked at the instability thresholds (although the proportion of points $\mathcal{P}(\Delta) d\Delta$ at the thresholds is still not large). As a result, Δ still remains dynamically important, but it is no longer capable of organizing to avoid Δp -stress, because the collisional term makes the dynamical equations different from those that support magneto-immutability. Therefore, the more severely the ordering is violated

¹¹The steepening in the hard-wall-limited simulation flow spectrum is also the most probable cause of the features observed in the density spectra of figure 4(b) at $\beta = 100$, since the density fluctuations are likely passively cascaded.

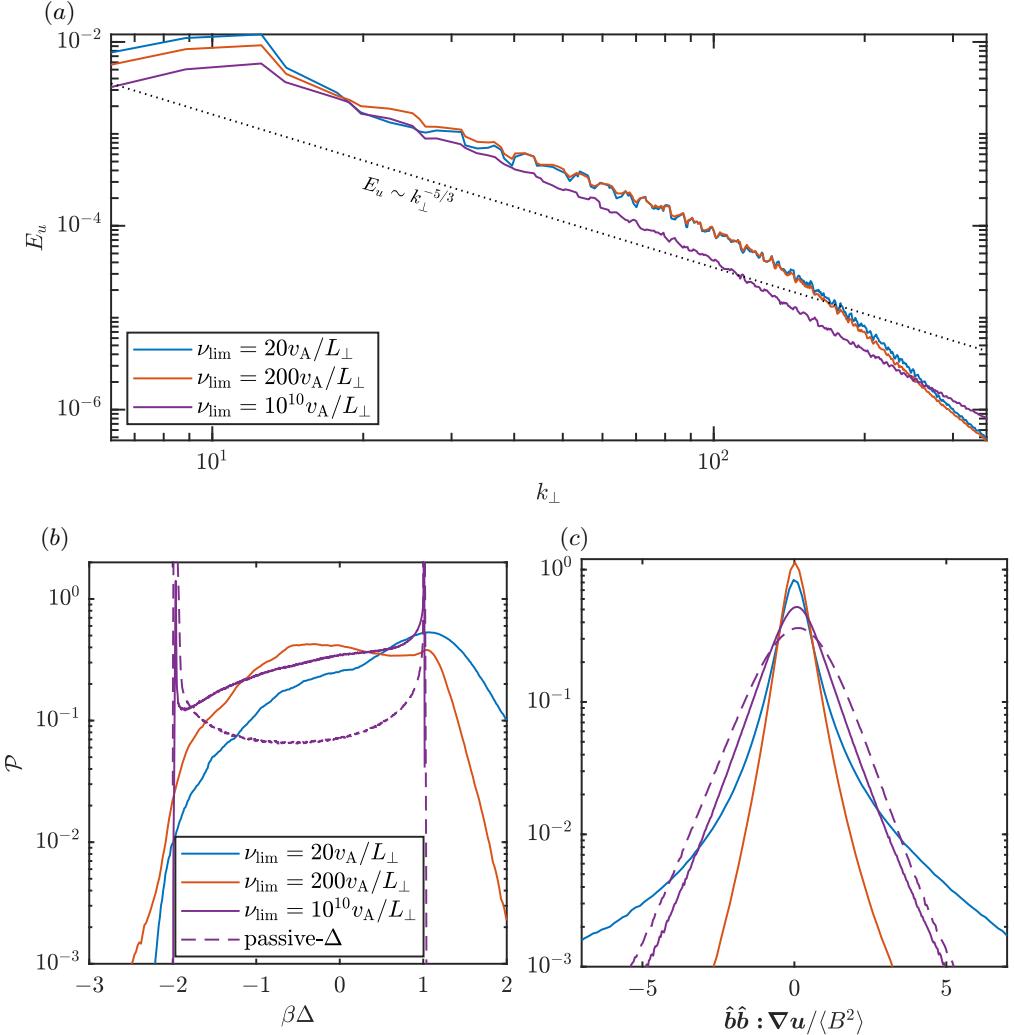


Figure 13: (a) Kinetic energy spectra for $\beta = 100$, Alfvénically driven simulations with $\nu_{\text{lim}} \in [20, 200, 10^{10}]v_A/L_\perp$. Significant spectral steepening is observed for the simulation with hard-wall limiters ($\nu_{\text{lim}} = 10^{10}v_A/L_\perp$, the default in all other simulations), as well as a slightly earlier apparent dissipation scale. (b) Probability distribution of the values of $\beta\Delta$ for each simulation. The hard-wall-limited simulation exhibits significant peaks with cutoffs near the mirror and firehose thresholds, while lower ν_{lim} simulations extend beyond the cutoffs. The $\nu_{\text{lim}} = 200v_A/L_\perp$ simulation is the only run having a distribution with a global maximum between the instability thresholds. (c) Probability distribution of the values of $\hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{u}$ for each simulation, demonstrating that the width of the distribution reflects the strength of magneto-immutable organization. The distribution from the hard-wall-limited simulation is the closest of the active simulations to the passive, non-immutable simulation, with $\nu_{\text{lim}} = 200v_A/L_\perp$ being the narrowest, suggesting that there is an intermediate value of ν_{lim} that allows magneto-immutability to act most effectively.

by the instability limiters, the more the anisotropic pressure-stress is allowed to affect the cascade of turbulent energy. Interestingly then, the results of figure 13 imply that throughout this work, many of the signatures of magneto-immutability we detect would become even stronger if a lower (and more realistic) value of ν_{lim} were used. That being said, it appears that if ν_{lim} is too small, then the organization can again become somewhat less efficient. This drop in efficiency may be a result of the amplitude of $\beta\Delta$ exceeding the values expected by the ordering, possibly upsetting the $\beta\Delta \sim 1$ assumption.

In order to choose the most physically relevant value of ν_{lim} , it is instructive to compare our results with that of the high- β , hybrid-kinetic simulations of Arzamasskiy *et al.* (2023). The authors found that microinstability scattering in their simulations appeared to follow the relationship $\nu \sim \beta(\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u})_L$, where the L subscript denotes that the quantity is estimated at the outer scale. We can estimate this rate in the simulations of figure 13 from the RMS velocity $\langle u \rangle_{\text{RMS}} \approx 0.75v_A$ to find $\nu \sim 235v_A/L$. Interestingly, this nearly coincides with the value of ν_{lim} in our survey that is minimally disruptive of immutability. Whether this is true at all values of β or not remains unclear without a more comprehensive study of the interaction of micro-instabilities and magneto-immutability. However, in light of our understanding of magneto-immutable turbulence as a consequence of critical balance, it is not surprising that Arzamasskiy *et al.* (2023) found the scattering rate to be determined by motions at the outer scale. The nature of Δp evolution in the reduced system (2.16), implies that Δp is only effectively being driven at the outer scale, where the forcing prevents self-organization of the turbulent cascade. Below the outer scale, pressure anisotropy is not actively produced but rather cascades passively to smaller values, having already been reduced to lie between the micro-instability thresholds. As a result, it is likely that the most physically realistic choice of ν_{lim} is also approximately that which best supports a strong, magneto-immutable cascade.

4. Summary and Discussion

This work has investigated high- β collisionless turbulence through analytical and computational means, with special attention paid to the self-organization of the magnetic field and bulk flow via ‘magneto-immutability’. We introduced a new asymptotic ordering that explicitly makes use of $\beta \gg 1$, and yields a set of reduced CGL-MHD equations that not only reproduces previously found characteristics of magneto-immutability, but also makes several new predictions. Numerical simulations of the full CGL-MHD equations with Landau-fluid heat fluxes and microinstability limiters were then employed to verify the assumptions of our ordering and to test the new predictions. The most important conclusions drawn by this study are:

- (i) Magneto-immutability, as defined by the suppression of both $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ and $\nabla_{\parallel}\Delta p$ through self-organization, is an inertial-range effect in high- β turbulence that satisfies the ordering (2.4).
- (ii) By suppressing $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$, magneto-immutability reduces the fraction of the plasma that has pressure anisotropy beyond the mirror/firehose thresholds, in turn preventing the plasma from becoming effectively collisional through microinstability-induced scattering.
- (iii) Despite this lack of scattering, the combination of heat flux, Δp -stress, and microinstability suppression allows high- β collisionless turbulence to behave almost en-

- tirely as a fluid. This is remarkable given that this parameter regime of plasma physics is particularly susceptible to collisionless effects.
- (iv) Magneto-immutability is relatively insensitive to the strength of heat fluxes, because of the dynamical reduction of the field-aligned temperature gradients $\nabla_{\parallel}\delta T_{\perp/\parallel}$.
 - (v) The suppression of $\hat{\mathbf{b}}\hat{\mathbf{b}}:\nabla\mathbf{u}$ is achieved through a local misalignment between $\hat{\mathbf{b}}$ and the eigenvectors of the $\nabla\mathbf{u}$ tensor, rather than through a suppression of the overall rate-of-strain (e.g., figure 8). Importantly, this misalignment is shown to extend beyond CGL-MHD to the hybrid-kinetic simulations of Arzamasskiy *et al.* (2023) (see figure 10).
 - (vi) No strong cascade of ion-acoustic waves exists in high- β collisionless turbulence; these modes can be driven effectively only at the outer scale by sonically correlated (or faster) forcing. This means that there is little dependence of the turbulence on details of the forcing in these plasmas, so long as its correlation time remains sufficiently slow. The spectrum of density fluctuations is determined by non-propagating modes, which adopt the statistics of the Alfvénically turbulent flow that mixes them.

In their study of magneto-immutability within the Landau-fluid CGL system, Squire *et al.* (2023) drew comparisons with the hybrid-kinetic simulations of Arzamasskiy *et al.* (2023), finding that these comparisons reflected well on the ability of the fluid model to capture the essential dynamics of collisionless high- β turbulence. The theory and simulations presented in this work further solidify those conclusions, especially with respect to the rate-of-strain alignment diagnostic, for which we have demonstrated qualitatively similar results between CGL-MHD and hybrid-kinetic simulations (§3.6). That being said, this work has touched on new aspects of high- β collisionless turbulence that could benefit from further comparison with well-tuned kinetic simulations. One such aspect is the microinstability scattering rate ν_{lim} , and how that compares to one that minimally disrupts magneto-immutability while still regulating the overall anisotropy (figure 13). Calculations of ν_{eff} from kinetic simulations of Alfvénic turbulence have been performed in Arzamasskiy *et al.* (2023), finding good agreement with a Braginskii-based estimate that depends on the outer-scale rate of strain. Future investigations may then benefit from a comparison between CGL-MHD simulations with an evolving, rather than fixed, ν_{lim} , and kinetic simulations that capture the scale dependence of ν_{eff} . Not only could the effectiveness of the collisional microinstability closure be studied more directly, but it would also permit a comparison of the relationship between intermittency and the unstable fraction, a topic with broad implications for cosmic-ray propagation through high- β turbulence (Reichherzer *et al.* 2023).

While we have made frequent reference to the importance of magneto-immutability in interpreting ICM observations, there are numerous yet-to-be-investigated ways in which magneto-immutability might affect turbulence in other high- β plasmas. For example, Kempinski *et al.* (2019) showed that incompressible turbulence driven by the magnetorotational instability (MRI; Balbus & Hawley 1991; Hawley & Balbus 1991) when subject to Braginskii viscosity (Balbus 2004) can self-organize so as to reduce the total (fluctuation- plus Keplerian-shear-produced) parallel rate of strain, thereby reducing the average pressure anisotropy in the plasma despite efficient angular-momentum transport by robust Reynolds and Maxwell stresses. That study could be extended using our Landau-fluid CGL-MHD model, exploring further the impact of magneto-immutability on the transport and turbulent cascade while making contact with previously published studies of collisionless MRI turbulence and transport that used Landau-fluid CGL-MHD (Sharma *et al.* 2006), hybrid-kinetic (Kunz *et al.* 2016), and pair-plasma kinetic (Bacchini

(*et al.* 2022; Sandoval *et al.* 2024) simulations. A particularly timely extension would be to investigate the compressive part of this magnetorotationally driven cascade in its inertial range and its role in plasma heating and angular-momentum transport. For example, a recent study by Kawazura *et al.* (2022) used a set of reduced-MHD equations tailored for the ‘shearing sheet’ to perform a local study of a magnetorotationally unstable accretion disc having a predominantly azimuthal mean magnetic field. Those authors found that compressive modes comprise a larger portion of the bulk kinetic energy than Alfvénic fluctuations, a result that they have recently confirmed via large-scale incompressible MHD simulations (Kawazura & Kimura 2024). They drew a link between this dominant compressive component and the heating of particles through the putative Landau damping of these fluctuations. However, their reduced model was based on the MHD equations, and it is known that the MRI in weakly collisional and collisionless plasmas (such as protogalaxies and low-luminosity accretion flows) is different than its MHD counterpart (e.g., Quataert *et al.* 2002; Balbus 2004; Squire *et al.* 2017b). With the MRI driving turbulence on Alfvénic (rather than sonic) timescales, and given the results of figures 3(a) and (b), it is not obvious that greater-than-unity ratios of compressive to Alfvénic energy could be achieved in high- β , collisionless accretion flows.

In constructing the reduced system (2.16), we made the simplifying assumption that the electrons are cold and isothermal. For the purpose of understanding how magneto-immutability behaves in various astrophysical environments where this is not necessarily the case, it is worth exploring the consequences of relaxing these assumptions. If the isothermal electron assumption held but we allowed them to be warm, say $T_e \sim T_i$, magneto-immutability would be unaffected so long as the density fluctuations remained smaller than $\mathcal{O}(\epsilon)$, which is the most likely case for high- β turbulence (this was tested using the same CGL-MHD code in Squire *et al.* 2023, with the authors finding little effect). If however the density fluctuations were larger (approaching $\sim \epsilon \rho_0$), some modifications would be made to the signatures of immutability. First, the perpendicular pressure balance would not lead to the suppression of δp_\perp , but rather the sum of $\delta p_\perp + T_e \delta \rho / m_i \approx 0$. The fluctuation $\delta p_\perp / \rho_0$ could then remain $\mathcal{O}(\epsilon)$, and suppression of $\hat{\mathbf{b}} \cdot \nabla u_\parallel$ would not be guaranteed. On the other hand, the u_\parallel momentum equation would still yield $\hat{\mathbf{b}} \cdot \nabla \Delta p \approx 0$ to leading order, thus immutability’s tendency to reduce anisotropic pressure stress would be unaffected. The picture is much less straightforward if we also relax the assumption of isothermal electrons. Collisionless electrons are necessary for modelling environments such as high- β radiatively inefficient accretion flows (Quataert 2003; Sharma *et al.* 2007), but the effects of electron pressure anisotropy and microinstabilities on magneto-immutability are beyond the scope of this work. Future efforts on this topic would not only yield interesting results on high- β turbulence, but also motivate cost-effective ways to model fully collisionless astrophysical plasmas, akin to what has been done between the Landau-fluid CGL approach and hybrid-kinetic particle-in-cell (Squire *et al.* 2023; Arzamasskiy *et al.* 2023).

Acknowledgments

The authors are grateful to A. Bott, S. Cowley, P. Kempinski, E. Quataert, and A. Schekochihin for valuable conversations.

Funding

S.M. and M.W.K. were supported in part by NSF CAREER Award No. 1944972. J.S. was supported by Rutherford Discovery Fellowship RDF-U001804 and Marsden Fund grant MFP-U002221, which are managed through the Royal Society Te Apārangi. High-

performance computing resources were provided by the PICSciE-OIT TIGRESS High Performance Computing Center and Visualization Laboratory at Princeton University.

Declaration of Interests

The authors report no conflict of interest.

Appendix A. Comparison with reduced kinetic MHD

The ordering (2.4) that yields the reduced high- β CGL-MHD equations is based heavily upon that employed by Schekochihin *et al.* (2009) and Kunz *et al.* (2015) to derive the reduced kinetic MHD (RKMHD) equations, often used to describe collisionless Alfvénic turbulence at long wavelengths. As a result, key signatures of immutability, like $\hat{\mathbf{b}} \cdot \nabla \Delta p$ and $\hat{\mathbf{b}} \hat{\mathbf{b}} : \nabla \mathbf{u}$ suppression, can in fact be obtained by applying a high- β subsidiary ordering to RKMHD. However, other aspects of the reduced system (2.16), such as the non-local influence of Δp and coupled Elsässer energy cascades, cannot be recovered. To better understand the relationship between RKMHD and our reduced equations, then, in this appendix we explore just how well magneto-immutability can be recovered from RKMHD, discussing the key differences from (2.16) as a result of primary versus subsidiary ordering of β .

The RKMHD equations, as given by equations (155)–(160) of Schekochihin *et al.* (2009), are simplified by our assumptions of zero electron temperature and collision frequency to yield the following:

$$\frac{\partial \Psi}{\partial t} = v_A \hat{\mathbf{b}} \cdot \nabla \Phi, \quad (\text{A } 1a)$$

$$\frac{d}{dt} \nabla_{\perp}^2 \Phi = v_A \hat{\mathbf{b}} \cdot \nabla \nabla_{\perp}^2 \Psi, \quad (\text{A } 1b)$$

$$\frac{d}{dt} \left(\delta f - \frac{v_{\perp}^2}{v_{\text{th}}^2} \frac{\delta B_{\parallel}}{B_0} F_0 \right) + v_{\parallel} \hat{\mathbf{b}} \cdot \nabla \delta f = 0, \quad (\text{A } 1c)$$

$$\frac{\delta n}{n_0} = -\frac{1}{2n_0} \int d^3 \mathbf{v} \left(\frac{v_{\perp}^2}{v_{\text{th}}^2} \frac{\beta}{\beta+1} - 2 \right) \left(\delta f - \frac{v_{\perp}^2}{v_{\text{th}}^2} \frac{\delta B_{\parallel}}{B_0} F_0 \right), \quad (\text{A } 1d)$$

$$\frac{\delta B_{\parallel}}{B_0} = -\frac{\beta}{\beta+1} \frac{1}{2n_0} \int d^3 \mathbf{v} \frac{v_{\perp}^2}{v_{\text{th}}^2} \left(\delta f - \frac{v_{\perp}^2}{v_{\text{th}}^2} \frac{\delta B_{\parallel}}{B_0} F_0 \right), \quad (\text{A } 1e)$$

where the full particle distribution function is $f(v_{\perp}, v_{\parallel}) = F_0(v_{\perp}, v_{\parallel}) + \delta f(v_{\perp}, v_{\parallel})$ with F_0 a background Maxwellian distribution, $v_{\perp/\parallel}$ the particle velocities, and

$$\mathbf{u}_{\perp} = \hat{\mathbf{z}} \times \nabla_{\perp} \Phi \quad \text{and} \quad \frac{\delta \mathbf{B}_{\perp}}{\sqrt{4\pi\rho_0}} = \hat{\mathbf{z}} \times \nabla_{\perp} \Psi. \quad (\text{A } 2)$$

The convective derivative d/dt and field-aligned gradient $\hat{\mathbf{b}} \cdot \nabla$ are the same as in (2.17). As a result of our assumptions, equations (A 1d) and (A 1e) have both become equivalent to perpendicular pressure balance, which is more easily seen when the velocity integrals are performed:

$$\frac{\delta p_{\perp}}{p_0} = -\frac{2}{\beta} \frac{\delta B_{\parallel}}{B_0}. \quad (\text{A } 3)$$

Just like in (2.10) then, $\beta \gg 1$ dictates that δp_{\perp} does not contribute to the leading-order pressure anisotropy Δp . Next we derive the parallel momentum equation, which in high- β

reduced CGL-MHD, leads to the suppression of viscosity. Taking $\int d^3v m_i v_{\parallel}$ of (A 1c) leads to

$$\frac{d}{dt}(\rho u_{\parallel}) = -\hat{\mathbf{b}} \cdot \nabla \delta p_{\parallel}. \quad (\text{A } 4)$$

Given that RKMHD has time derivatives that scale as $d_t \sim k_{\parallel} v_A$, the leading order of this equation is $\hat{\mathbf{b}} \cdot \nabla \delta p_{\parallel} \approx 0$; as a result, the anisotropic pressure stress is zero to leading order once again. This should come as no surprise because the CGL-MHD parallel momentum equation is exactly the v_{\parallel} moment of the full kinetic-MHD equation for f (indeed, the only difference between the fluid CGL-MHD model and kinetic MHD is in the higher moments of f , such as $v_{\perp/\parallel}^2$). This is clear if we take $\int d^3v m_i v_{\parallel}^2 / 2$ of (A 1c) to get the equation for the μ adiabat:

$$\frac{d\delta p_{\perp}}{dt} - 2 \frac{p_0}{B_0} \frac{d\delta B_{\parallel}}{dt} = -\hat{\mathbf{b}} \cdot \nabla \left[\frac{1}{2} \int d^3v m_i v_{\parallel} v_{\perp}^2 \delta f \right]. \quad (\text{A } 5)$$

In our model, the right-hand-side of (A 5) is approximated by the heat flux q_{\perp} of (2.2). If we were to ignore q_{\perp} , the reduction of δp_{\perp} from (A 3) would yield $d_t \delta B_{\parallel} \approx 0$, and thus $\hat{\mathbf{b}} \cdot \nabla u_{\parallel} \approx 0$, another signature of magneto-immutability. If the heat fluxes were nonzero, we could still achieve this reduction in $\hat{\mathbf{b}} \cdot \nabla u_{\parallel}$ if $q_{\perp} \propto \hat{\mathbf{b}} \cdot \nabla T_{\perp}$, or $q_{\perp} \propto \hat{\mathbf{b}} \cdot \nabla p_{\perp}$ given small density fluctuations from low Mach number forcing (as in our reduced CGL-MHD model).

However, this reduction of $\hat{\mathbf{b}} \cdot \nabla u_{\parallel}$ and $\hat{\mathbf{b}} \cdot \nabla \Delta p$ is of little consequence to the turbulent evolution, because in this model, the pressure anisotropy stress on the turbulent flow has already been ordered out. The momentum equation for \mathbf{u}_{\perp} is written in terms of the potential Φ in (A 1b), where it is clear that only one characteristic velocity is present – the background Alfvén speed. As RKMHD does not include $\beta^{-1} \sim \epsilon$ in the primary ordering, the effects of anisotropic pressure on the evolution of \mathbf{u}_{\perp} are lost and cannot be recovered through a subsidiary ordering, and so there is no way to enforce $\beta \Delta \sim 1$. This subsidiary ordering would then certainly misrepresent the outer scale of our CGL-MHD turbulence simulations. However, would RKMHD suffice when the fluctuations in Δp_k at some large wavenumber k far from the outer scale become too small to satisfy $\beta \Delta_k \sim 1$? It may seem reasonable to apply the Schekochihin *et al.* (2009) RKMHD to the deep inertial range of such turbulence, however as discussed in §2.5, this would still miss a possibly important effect: Δp can act very non-locally in k -space through its modification of $v_{A,\text{eff}}$. Small-scale Alfvénic fluctuations in the high- β reduced CGL-MHD model are subject to a background effective Alfvén speed set by the turbulence and consequent pressure anisotropies at the largest scales. Because of this, patches of the turbulence may evolve somewhat uniquely, or may vary in their ability to interact with cosmic rays, for example (Marcowith *et al.* 2021). In this situation, the Kunz *et al.* (2015) model of RKMHD that includes pressure anisotropy in the background particle distribution could more accurately capture these effects. The slowly-evolving, large-scale motions would provide the background pressure anisotropy upon which the anisotropic RKMHD could be evolved, and as the model otherwise includes the same assumptions that lead to immutability signatures in isotropic RKMHD, it also captures the reduction of $\hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{u}$ and $\hat{\mathbf{b}} \cdot \nabla \Delta p$ at high β .

Appendix B. Magneto-immutability and the Braginskii viscous stress

The initial discovery of magneto-immutability in Squire *et al.* (2019) came from an investigation of weakly collisional Braginskii-MHD turbulence, rather than turbulence

with a collisionless model as is studied in this work. While the Braginskii closure for Δp differs dramatically from that of our collisionless Landau-fluid CGL model,¹² the magneto-immutable suppression of the Δp -stress in our reduced CGL approach comes only from the momentum equation, which is shared by both models. It is therefore within reason to suspect that the mechanism for viscosity suppression also originates from the momentum equation in Braginskii-MHD. For that reason, in this appendix we derive the condition for viscous stress reduction in Braginskii-MHD by assuming that the cause is the same as that leading to (2.12) (i.e., $\hat{\mathbf{b}} \cdot \nabla \delta T_{\parallel} = 0$ to leading order), essentially obtaining a threshold for realizing magneto-immutable behaviour in the weakly collisional limit.

Unlike the CGL-MHD model, Braginskii-MHD does not evolve the pressure anisotropy directly from conservation of the double adiabatic invariants. Instead, it assumes that the rate of scattering is sufficiently rapid ($\nu \gg k_{\parallel} v_{\text{th}}$) that a balance is struck between production of anisotropy via changes in B and ρ and its depletion through pitch-angle scattering. As a result, the leading-order perturbation to p is isotropic, and Δp only arises at next order in $k_{\parallel} v_{\text{th}}/\nu$. This allows the double-adiabatic equations (2.1d) and (2.1e) to be replaced with (Braginskii 1965):¹³

$$\frac{d}{dt} \ln \frac{p}{\rho^{5/3}} = \frac{2}{\nu} \left[\left(\hat{\mathbf{b}} \hat{\mathbf{b}} - \frac{\mathbf{I}}{3} \right) : \nabla \mathbf{u} \right]^2 \quad \text{and} \quad \Delta p = \frac{3p}{\nu} \left(\hat{\mathbf{b}} \hat{\mathbf{b}} - \frac{\mathbf{I}}{3} \right) : \nabla \mathbf{u}. \quad (\text{B } 1)$$

To achieve suppression of parallel viscous forces through magneto-immutability in the same manner as realized in our reduced CGL-MHD model, we seek to ensure that the leading order of the u_{\parallel} momentum equation becomes $\hat{\mathbf{b}} \cdot \nabla \Delta p \approx 0$. As with the reduced CGL-MHD model, we make the simplifying assumption that both density fluctuations and $\nabla \cdot \mathbf{u}$ are negligible, and apply the ordering (2.3) to the Braginskii-MHD equations. Note, however, that we make no assumption regarding the size of β , as we will instead derive a β -dependent criterion for magneto-immutability to take effect in Braginskii-MHD. Therefore, we will simply have to assume that Δp cannot be neglected in the momentum equation, so that it can later inform us of how large β needs to be in order to suppress Δp 's parallel gradients (otherwise magneto-immutability would be impossible to recover). Reducing the momentum equation produces the following leading-order equation for u_{\perp} , which still describes pressure balance, but in this case the isotropic pressure dominates to leading order

$$\frac{\delta p}{p_0} \approx -\frac{\beta}{2} \frac{\delta B_{\parallel}}{B_0}, \quad (\text{B } 2)$$

only in this case it is struck with the isotropic pressure perturbation. The leading order of the parallel momentum equation can then be written as

$$\rho_0 \frac{du_{\parallel}}{dt} \approx -\hat{\mathbf{b}} \cdot \nabla \Delta p + B_0 \hat{\mathbf{b}} \cdot \nabla \delta B_{\parallel}. \quad (\text{B } 3)$$

¹²Note that the Landau-fluid CGL-MHD equations, given a uniform scattering rate, do reproduce the Braginskii-MHD model in the collisional limit $\nu \gg k_{\parallel} v_{\text{th}}$. This scattering rate must not only isotropize the pressures, but also suppress the heat fluxes, using the approach given in [Sharma et al. \(2003\)](#).

¹³Another regime of Braginskii-MHD can be obtained by instead assuming that $\nu \gg k_{\parallel} v_A$, which is a considerably weaker criterion than $\nu \gg k_{\parallel} v_{\text{th}}$ at high β . In this limit, however, the heat fluxes are in the collisionless regime and not ordered out of the system, and so they must be taken into account. It may be the case that their effects are suppressed by magneto-immutability regardless, although further investigation would be needed to confirm such a conclusion.

To compare the sizes of each term, we now substitute for Δp the weakly collisional closure (B 1), which, if ordered according to (2.3), yields $\Delta p \approx (3p_0/\nu)\hat{\mathbf{b}} \cdot \nabla u_{\parallel}$. Substituting this into (B 3) yields

$$\rho_0 \frac{du_{\parallel}}{dt} \approx -\frac{3\beta}{2\nu} B_0^2 \hat{\mathbf{b}} \cdot \nabla (\hat{\mathbf{b}} \cdot \nabla u_{\parallel}) + B_0 \hat{\mathbf{b}} \cdot \nabla \delta B_{\parallel}. \quad (\text{B } 4)$$

In (B 4), the Δp -stress takes on its familiar viscous form. The left-hand side and the final term on the right-hand-side are both of order $\sim \epsilon k_{\parallel} B_0^2$, while the viscous term is $\sim \epsilon k_{\parallel} B_0^2 (\beta k_{\parallel} v_A / \nu)$. Thus, if the viscosity is to dominate, forcing the plasma to self-organize in order to avoid it, we require

$$\beta \frac{k_{\parallel} v_A}{\nu} \gg 1. \quad (\text{B } 5)$$

If this criterion is met, then $\hat{\mathbf{b}} \cdot \nabla \Delta p$ will be suppressed and the Alfvénic cascade will not be strongly damped by viscous stress. By design, this is precisely the regime within which magneto-immutability was studied in Squire *et al.* (2019) and Squire *et al.* (2023). Equation (B 5) also implies that magneto-immutability in Braginskii-MHD, unlike its CGL-MHD counterpart, *is* scale dependent. Consider a scenario in which outer-scale motions are too slow to produce Δp faster than it can be eroded by the scattering rate, and the criterion (B 5) is not satisfied. As the cascade progresses, the eddy turnover times get shorter at smaller scales, and the production of pressure anisotropy occurs at a faster rate. Eventually, when the generation of pressure anisotropy occurs on timescales small enough that it competes with the scattering, magneto-immutability can step in to regulate its magnitude. At some point, however, the scales will become sufficiently small that the collisional assumption of $\nu/k_{\parallel} v_{\text{th}} \gg 1$ becomes inadequate and a fully collisionless model must be used.

Appendix C. High- β reduced CGL with large density fluctuations

One of the fundamental assumptions we make within this work is that density fluctuations are $\mathcal{O}(\epsilon^2)$ or smaller. Physically, this is motivated by the difficulty of driving both high Mach number and sonically correlated turbulence in astrophysical high- β plasmas. Indeed for all of the simulations performed within the scope of this work, in no circumstances did $\delta\rho$ exceed $\epsilon^2 \rho_0$, a necessary condition for obtaining the excellent agreement between our predictions and the simulation results. Nonetheless, it is worthwhile to at least consider the consequences of $\delta\rho \sim \epsilon^{3/2} \rho_0$ or $\epsilon \rho_0$, and how that would affect the conclusions we have reached so far.

We begin with $\delta\rho \sim \epsilon^{3/2} \rho_0$, after which $\delta\rho \sim \epsilon \rho_0$ is a relatively simple extension. All non- $\delta\rho$ oriented aspects of the ordering (2.4) can be used once again, although we will drop $\delta T_{\perp/\parallel}$ in favor of $\delta p_{\perp/\parallel}$ given the enhancement of $\delta\rho$. Among other things, this means that the pressure balance of equation (2.10) becomes

$$\frac{\delta p_{\perp}^{(2)}}{p_0} = -\frac{2}{\beta} \frac{\delta B_{\parallel}^{(1)}}{B_0} \quad \text{and} \quad \frac{\delta p_{\perp}^{(5/2)}}{p_0} = -\frac{2}{\beta} \frac{\delta B_{\parallel}^{(3/2)}}{B_0}, \quad (\text{C } 1)$$

and equations (2.12) and (2.13) are the same but with δp_{\parallel} swapped for δT_{\parallel} :

$$\hat{\mathbf{b}}^{(0)} \cdot \nabla \delta p_{\parallel}^{(1)} = (\hat{\mathbf{b}} \cdot \nabla \delta p_{\parallel})^{(1)} = 0, \quad (\text{C } 2a)$$

$$\hat{\mathbf{b}}^{(0)} \cdot \nabla \delta p_{\parallel}^{(3/2)} + \frac{\delta B_{\perp}^{(3/2)}}{B_0} \cdot \nabla_{\perp} \delta p_{\parallel}^{(1)} = (\hat{\mathbf{b}} \cdot \nabla \delta p_{\parallel})^{(3/2)} = 0. \quad (\text{C } 2b)$$

This shows that the suppression of the anisotropic pressure-stress, being independent of the magnitude of density fluctuations, is a particularly robust aspect of immutability. Continuing, the equations that evolve $\delta\mathbf{B}_\perp$ and \mathbf{u}_\perp are unaffected, as is the continuity equation. Note that although we are allowing the density fluctuations to be larger, we will still employ the assumption that $d_t \delta\rho^{(3/2)}$ is negligible, owed to the fact that these density fluctuations are still likely the product of non-propagating modes. As such, in evaluating the p_\perp and p_\parallel equations, we can approximate

$$\frac{d}{dt} \ln \frac{p_\perp}{\rho B} \approx -\frac{d}{dt} \ln B \quad \text{and} \quad \frac{d}{dt} \ln \frac{p_\parallel B^2}{\rho^3} \approx \frac{d}{dt} \ln p_\parallel B^2, \quad (\text{C } 3)$$

where we have again used the fact that p_\perp has no $\mathcal{O}(\epsilon)$ perturbation. Starting with equation (2.1d), the heat flux q_\perp is still 0 at order $\epsilon^{1/2}$, since the density fluctuations are only $\mathcal{O}(\epsilon^{3/2})$ and $\hat{\mathbf{b}} \cdot \nabla \delta p_\parallel^{(1,3/2)} = 0$. However, at order ϵ where left-hand-side of (2.1d) first appears, we have

$$\frac{d}{dt} \frac{\delta B_\parallel^{(1)}}{B_0} = \frac{v_{\text{th}}}{\sqrt{\pi} |k_\parallel|} (\hat{\mathbf{b}}^{(0)} \cdot \nabla)^2 \frac{\delta \rho^{(3/2)}}{\rho_0}. \quad (\text{C } 4)$$

As a result, instead of finding that $\hat{\mathbf{b}} \cdot \nabla u_\parallel = 0$ by comparing this with parallel induction (as we found when $\delta\rho \lesssim \epsilon^2 \rho_0$), we find

$$\hat{\mathbf{b}}^{(0)} \cdot \nabla \left(u_\parallel^{(1)} - \frac{v_{\text{th}}}{\sqrt{\pi} |k_\parallel|} \hat{\mathbf{b}}^{(0)} \cdot \nabla \frac{\delta \rho^{(3/2)}}{\rho_0} \right) = 0. \quad (\text{C } 5)$$

Note that we cannot use this equation to fully determine u_\parallel , given that taking $\delta\rho$ to be smaller would imply $u_\parallel^{(1)} = 0$, which is not the same as the misalignment we predicted and measured in §3.6. The equation for $\delta p_\parallel^{(1)}$ is obtained with ease given (C4):

$$\frac{d}{dt} \frac{\delta p_\parallel^{(1)}}{\rho_0} = -\frac{4v_{\text{th}}}{\sqrt{\pi} |k_\parallel|} (\hat{\mathbf{b}}^{(0)} \cdot \nabla)^2 \frac{\delta \rho^{(3/2)}}{\rho_0}. \quad (\text{C } 6)$$

Therefore, although the suppression of the Δp -stress is preserved by the increase in density fluctuation amplitude, the misalignment of $\hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{u}$ and passive advection of Δp are not necessarily. Instead, the anisotropic pressure fluctuations are expected to collisionlessly damp as they are mixed passively by the Alfvénic turbulence. The exact rate of damping would then be determined by the passively advected density fluctuations that are evolved according to $d_t \delta\rho^{(3/2)} = 0$.

The extension of the above to $\delta\rho \sim \epsilon \rho_0$ is rather simple, with only the following modifications: Instead of the $\mathcal{O}(\epsilon^{1/2})$ contribution to the heat fluxes $q_{\perp/\parallel}$ being 0, they will remain nonzero due to the density perturbation, with (2.1e) and (2.1d) becoming $\hat{\mathbf{b}}^{(0)} \cdot \nabla \delta\rho^{(1)}$ to leading order. Following this, (C6) and (C4) remain the same, but it is less clear whether the assumption $d_t \delta\rho = 0$ can be applied to both $\delta\rho^{(1)}$ and $\delta\rho^{(3/2)}$, or if that can only be said of the leading order. This may still be the case as we expect the only other source of density fluctuations to be Alfvén wave nonlinearities, which should appear at order ϵ^2 , rather than $\epsilon^{3/2}$. However, this is also reliant on discerning to how many orders ion-acoustic wave fluctuations can be ignored, given that some weak mixing may still occur.

- ARZAMASSKIY, L., KUNZ, M.W., SQUIRE, J., QUATAERT, E. & SCHEKOCHIHIHIN, A.A. 2023 Kinetic turbulence in collisionless high- β plasmas. *Phys. Rev. X* **13** (2), 021014.
- BACCHINI, F., ARZAMASSKIY, L., ZHDANKIN, V., WERNER, G.R., BEGELMAN, M.C. & UZDENSKY, D.A. 2022 Fully kinetic shearing-box simulations of magnetorotational turbulence in 2D and 3D. I. Pair plasmas. *Astrophys. J.* **938** (1), 86.
- BALBUS, S.A. 2004 Viscous shear instability in weakly magnetized, dilute plasmas. *Astrophys. J.* **616**, 857–864.
- BALBUS, S.A. & HAWLEY, J.F. 1991 A powerful local shear instability in weakly magnetized disks I. Linear analysis. *Astrophys. J.* **376**, 214.
- BARNES, A. 1966 Collisionless damping of hydromagnetic waves. *Phys. Fluids* **9**, 1483.
- BISKAMP, D. 2003 *Magnetohydrodynamic turbulence*. Cambridge University Press.
- BRAGINSKII, S.I. 1965 Transport processes in a plasma. *Rev. Plasma Phys.* **1**, 205.
- BRANDENBURG, A. & SUBRAMANIAN, K. 2005 Astrophysical magnetic fields and nonlinear dynamo theory. *Phys. Rep.* **417** (1-4), 1–209.
- CHEW, G.F., GOLDBERGER, M.L. & LOW, F.E. 1956 The Boltzmann equation and the one-fluid hydromagnetic equations in the absence of particle collisions. *Proc. Roy. Soc. London Ser. A* **236**, 112.
- CHO, J. & LAZARIAN, A. 2009 Simulations of electron magnetohydrodynamic turbulence. *Astrophys. J.* **701** (1), 236–252.
- COWIE, L. L. & MCKEE, C. F. 1977 The evaporation of spherical clouds in a hot gas. I. Classical and saturated mass loss rates. *Astrophys. J.* **211**, 135–146.
- ELSÄSSER, W.M. 1950 The hydromagnetic equations. *Phys. Rev.* **79** (1), 183.
- GOLDREICH, P. & SRIDHAR, S. 1995 Toward a theory of interstellar turbulence. 2: Strong Alfvénic turbulence. *Astrophys. J.* **438**, 763.
- GOLDSTEIN, M.L., ROBERTS, D.A. & MATTHAEUS, W.H. 1995 Magnetohydrodynamic turbulence in the solar wind. *Ann. Rev. Astron. Astrophys.* **33**, 283–326.
- GRETE, P., O'SHEA, B. W., BECKWITH, K., SCHMIDT, W. & CHRISTLIEB, A. 2017 Energy transfer in compressible magnetohydrodynamic turbulence. *Physics of Plasmas* **24** (9), 092311.
- HASEGAWA, A. 1969 Drift mirror instability of the magnetosphere. *Phys. Fluids* **12**, 2642.
- HAWLEY, J.F. & BALBUS, S.A. 1991 A powerful local shear instability in weakly magnetized disks II. Nonlinear evolution. *Astrophys. J.* **376**, 223.
- HELLINGER, P. & MATSUMOTO, H. 2000 New kinetic instability: oblique Alfvén fire hose. *J. Geophys. Res.* **105**, 10519.
- HOLLWEG, JOSEPH V. 1974 On electron heat conduction in the solar wind. *J. Geophys. Res.* **79** (25), 3845.
- KAWAZURA, Y. & KIMURA, S. S. 2024 Inertial range of magnetorotational turbulence. *arXiv e-prints*.
- KAWAZURA, Y., SCHEKOCHIHIHIN, A.A., BARNES, M., DORLAND, W. & BALBUS, S.A. 2022 Energy partition between Alfvénic and compressive fluctuations in magnetorotational turbulence with near-azimuthal mean magnetic field. *J. Plasma Phys.* **88** (3), 905880311.
- KAZANTSEV, A.P. 1968 Enhancement of a magnetic field by a conducting fluid. *Sov. J. Exp. Theor. Phys.* **26**, 1031.
- KEMPSKI, P., QUATAERT, E., SQUIRE, J. & KUNZ, M.W. 2019 Shearing-box simulations of MRI-driven turbulence in weakly collisional accretion discs. *Mon. Not. Roy. Astron. Soc.* **486** (3), 4013–4029.
- KULSRUD, R.M. 1983 MHD description of plasma. In *Basic Plasma Physics: Selected Chapters, Handbook of Plasma Physics, Volume 1* (ed. A.A. Galeev & R.N. Sudan), p. 1.
- KUNZ, M.W., JONES, T.W. & ZHURAVLEVA, I. 2022 Plasma physics of the intracluster medium. In *Handbook of X-ray and Gamma-ray Astrophysics*, p. 56. Springer Nature Singapore.
- KUNZ, M.W., SCHEKOCHIHIHIN, A.A., CHEN, C.H.K., ABEL, I.G. & COWLEY, S.C. 2015 Inertial-range kinetic turbulence in pressure-anisotropic astrophysical plasmas. *J. Plasma Phys.* **81**, 325810501.
- KUNZ, M.W., SQUIRE, J., SCHEKOCHIHIHIN, A.A. & QUATAERT, E. 2020 Self-sustaining sound in collisionless, high- β plasma. *J. Plasma Phys.* **86**, 905860603.
- KUNZ, M.W., STONE, J.M. & BAI, X.-N. 2014 Pegasus: a new hybrid-kinetic particle-in-cell code for astrophysical plasma dynamics. *J. Comput. Phys.* **259**, 154.

- KUNZ, M. W., STONE, J. M. & QUATAERT, E. 2016 Magnetorotational Turbulence and Dynamo in a Collisionless Plasma. *Phys. Rev. Lett.* **117** (23), 235101.
- MAJESKI, S., KUNZ, M.W. & SQUIRE, J. 2023 Microphysically modified magnetosonic modes in collisionless, high- β plasmas. *J. Plasma Phys.* **89** (3), 905890303.
- MAJESKI, S. & KUNZ, M. W. 2024 On hydromagnetic wave interactions in collisionless, high- β plasmas. *J. Plasma Phys.* **90** (1), 535900101.
- MARCOWITH, A., VAN MARLE, A.J. & PLOTNIKOV, I. 2021 The cosmic ray-driven streaming instability in astrophysical and space plasmas. *Phys. Plasmas* **28** (8), 080601.
- MARSCH, E. 2006 Kinetic Physics of the Solar Corona and Solar Wind. *Liv. Rev. Sol. Phys.* **3** (1), 1.
- MEYRAND, R., KANEKAR, A., DORLAND, W. & SCHEKOCHIHIHIN, A.A. 2019 Fluidization of collisionless plasma turbulence. *Proc. Nat. Acad. Sci.* **116** (4), 1185–1194.
- QUATAERT, E. 2003 Radiatively inefficient accretion flow models of Sgr A*. *Astron. Nachrichten Suppl.* **324**, 435–443.
- QUATAERT, E., DORLAND, W. & HAMMETT, G.W. 2002 The magnetorotational instability in a collisionless plasma. *Astrophys. J.* **577**, 524–533.
- REICHHERZER, P., BOTT, ARCHIE F.A., EWART, R.J., GREGORI, G., KEMPSKI, P., KUNZ, M.W. & SCHEKOCHIHIHIN, A.A. 2023 Efficient micromirror confinement of sub-TeV cosmic rays in galaxy clusters. *arXiv e-prints* p. arXiv:2311.01497.
- RIQUELME, M., QUATAERT, E. & VERSCHAREN, D. 2018 PIC simulations of velocity-space instabilities in a decreasing magnetic field: Viscosity and thermal conduction. *Astrophys. J.* **854** (2), 132.
- SANDOVAL, A., RIQUELME, M., SPITKOVSKY, A. & BACCHINI, F. 2024 Particle-in-cell simulations of the magnetorotational instability in stratified shearing boxes. *Mon. Not. Roy. Astron. Soc.* **530** (2), 1866–1884.
- SANTOS-LIMA, R., DE GOUEVIA DAL PINO, E.M., KOWAL, G., FALCETA-GONÇALVES, D., LAZARIAN, A. & NAKWACKI, M.S. 2014 Magnetic field amplification and evolution in turbulent collisionless magnetohydrodynamics: an application to the intracluster medium. *Astrophys. J.* **781** (2), 84.
- SCHEKOCHIHIHIN, A.A. & COWLEY, S.C. 2007 *Turbulence and magnetic fields in astrophysical plasmas*, pp. 85–115. Dordrecht: Springer Netherlands.
- SCHEKOCHIHIHIN, A.A., COWLEY, S.C., DORLAND, W., HAMMETT, G.W., HOWES, G.G., QUATAERT, E. & TATSUNO, T. 2009 Astrophysical gyrokinetics: kinetic and fluid turbulent cascades in magnetized weakly collisional plasmas. *Astrophys. J. Supp. Ser.* **182**, 310.
- SCHEKOCHIHIHIN, A.A., COWLEY, S.C., TAYLOR, S.F., MARON, J.L. & MCWILLIAMS, J.C. 2004 Simulations of the small-scale turbulent dynamo. *Astrophys. J.* **612** (1), 276–307.
- SHARMA, P., HAMMETT, G.W., QUATAERT, E. & STONE, J.M. 2006 Shearing box simulations of the MRI in a collisionless plasma. *Astrophys. J.* **637**, 952.
- SHARMA, P., HAMMETT, G. W. & QUATAERT, E. 2003 Transition from Collisionless to Collisional Magnetorotational Instability. *Astrophys. J.* **596** (2), 1121–1130.
- SHARMA, P., QUATAERT, E., HAMMETT, G.W. & STONE, J.M. 2007 Electron heating in hot accretion flows. *Astrophys. J.* **667**, 714.
- SNYDER, P.B., HAMMETT, G.W. & DORLAND, W. 1997 Landau fluid models of collisionless magnetohydrodynamics. *Phys. Plasmas* **4**, 3974.
- SQUIRE, J., KUNZ, M.W., ARZAMASSKIY, L., JOHNSTON, Z., QUATAERT, E. & SCHEKOCHIHIHIN, A.A. 2023 Pressure anisotropy and viscous heating in weakly collisional plasma turbulence. *J. Plasma Phys.* **89**, 905890417.
- SQUIRE, J., KUNZ, M.W., QUATAERT, E. & SCHEKOCHIHIHIN, A.A. 2017a Kinetic simulations of the interruption of large-amplitude shear-Alfvén waves in a high- β plasma. *Phys. Rev. Lett.* **119** (15), 155101.
- SQUIRE, J., QUATAERT, E. & KUNZ, M. W. 2017b Pressure-anisotropy-induced nonlinearities in the kinetic magnetorotational instability. *J. Plasma Phys.* **83** (6), 905830613.
- SQUIRE, J., QUATAERT, E. & SCHEKOCHIHIHIN, A.A. 2016 A stringent limit on the amplitude of Alfvénic perturbations in high-beta low-collisionality plasmas. *Astrophys. J.* **830**, L25.
- SQUIRE, J., SCHEKOCHIHIHIN, A.A. & QUATAERT, E. 2017c Amplitude limits and nonlinear

- damping of shear-Alfvén waves in high-beta low-collisionality plasmas. *New J. Phys.* **19**, 055005.
- SQUIRE, J., SCHEKOCHIHIHIN, A.A., QUATAERT, E. & KUNZ, M.W. 2019 Magneto-immutable turbulence in weakly collisional plasmas. *J. Plasma Phys.* **85** (1), 905850114.
- ST-ONGE, D.A., KUNZ, M.W., SQUIRE, J. & SCHEKOCHIHIHIN, A.A. 2020 Fluctuation dynamo in a weakly collisional plasma. *J. Plasma Phys.* **86** (5), 905860503.
- STONE, J.M., TOMIDA, K., WHITE, C.J. & FELKER, K.G. 2020 The Athena++ adaptive mesh refinement framework: design and magnetohydrodynamic solvers. *Astrophys. J. Supp. Ser.* **249** (1), 4.
- UHLENBECK, G.E. & ORNSTEIN, L.S. 1930 On the theory of the Brownian motion. *Phys. Rev.* **36** (5), 823–841.
- ZANK, G.P. & MATTHAEUS, W.H. 1992 The equations of reduced magnetohydrodynamics. *J. Plasma Phys.* **48**, 85–100.
- ZHURAVLEVA, I., CHURAZOV, E., SCHEKOCHIHIHIN, A.A., ALLEN, S.W., VIKHLININ, A. & WERNER, N. 2019 Suppressed effective viscosity in the bulk intergalactic plasma. *Nature Astro.* **3**, 832–837.