Universal equilibria, phase-space structure of collisionless plasma systems, and turbulence in non-Maxwellian plasmas

Robert J. Ewart

Balliol College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2024

Abstract

A fundamental tenet of thermodynamics is that chaotic systems will relax to maximum-entropy states. In plasmas, the chaos is conventionally provided by inter-particle collisions and the universal maximum-entropy equilibrium is a Maxwellian distribution. However, in collisionless plasmas, the chaotic state is due to collective turbulent dynamics that, while certainly disordered, does not maximise the conventional entropy which would give rise to the Maxwellian. The key result of this thesis is the identification and numerical verification of a class of universal equilibria towards which collisionless plasmas relax. These are still maximum-entropy equilibria, but of a different thermodynamics—that proposed by Lynden-Bell (1967) to describe chaotic relaxation in a collisionless system which conserves an infinite family of invariants, known as Casimir invariants (Ye &Morrison, 1992). We show theoretically that, in sufficiently turbulent plasmas, the Lynden-Bell equilibria become highly universal, forming distributions possessing a power-law tail with exponent -2. The conservation of the Casimir invariants however, is not precise due to the presence of a small amount of collisionality. As a result, the conservation laws on short times endow the system with a partial 'memory' of its prior conditions, but the development of a turbulent cascade to small scales, which breaks the precise conservation of phase volume, make this short-term memory imprecise on long times scales. The equilibria are still determined by the short-time collisionless invariants, but the invariants themselves are driven to a universal form by the nature of the turbulence. We demonstrate this effect numerically using the largest, longest run, and best resolved simulations of the two-stream instability to date, showing strong agreement with our theory.

Finally we derive quasilinear collision integrals which, by virtue of possessing an H-theorem, naturally push the system towards a Lynden-Bell equilibrium. This acts as the dynamical counterpart to the thermodynamic formalism we have developed. The collision integrals we derive are similar in form to those derived in collisional and collisionless contexts (Landau, 1936; Balescu, 1960; Lenard, 1960; Kadomtsev &Pogutse, 1970; Severne &Luwel, 1980; Chavanis, 2004, 2022), relying on the scattering of particles off decorrelated fluctuations which, however, are statistically aware of the conservation of the collisionless invariants. We show that relaxation mediated by such a collision integral gives rise to a number of novel effects such as anomalous interspecies drag, and spontaneous current generation.

Acknowledgements

I would first like to thank my supervisor, Alex Schekochihin. Through prolonged chats, frenzied outbursts of creativity, helpful criticism, and criticismful help, Alex has never failed to inspire, uplift, and teach. What can be said to be true physics (and half-decent English) in this thesis is, in the greatest degree, due to his guidance.

There are many collaborators without whom this work would have been completely impossible. Toby Adkins, Andrew Brown, Pablo Bilbao, Michael Nastac, Luis Silva, and Thales Silva. I would especially like to thank Pablo for dragging me, kicking and screaming, through my earliest numerical simulations, and Michael, whose constant commitment to precision (and general worrying) has been a much needed remedy to my more mercurial tendencies. I am truly grateful for their friendship.

There has also been work, the fruits of which did not make it into this thesis, with collaborators that deserve thanks. Archie Bott, Alessandro Geraldini, Barry Ginat, Matt Kunz, Ralf Mackenbach, Felix Parra, and Patrick Reichherzer. Their collaboration so often proved a pleasant change of scenery, and made my DPhil more than the sum of its parts.

I owe a great thanks to the friends that have supported and distracted me throughout my DPhil. Through the covid months, Aston, Jairui, Javier, Justin, Hrushi, Kristine, Marcel, Sophie, and Xinran, and, in more sluggish times Esthy, Georgia, Rachel, and Simone. I also owe thanks to the various friends from before my DPhil, who continue to tolerate my moods despite knowing better. Alex C and Alex H, Jojo, Hannah, Katya, Matt, Sam, Steph, Teb, Vidy, and, of course, Ellen (who deserves to be credited as much more than a friend). If any of them find themselves reading through this manuscript and spot some of my personality clinging to the pages, they should not worry. It is on their account that I have any at all to give.

I would like to thank my siblings, Loren, William, and Ellie. They are the best, and I wouldn't change them for the world. Lastly, I want to thank my parents, Colin and Denyse, for their boundless love, support, and guidance, throughout my whole life. This thesis is dedicated to them.

Contents

List of Figures

'Tis a merry thing to see
At our tasks how glad are we,
When at home we sit and find
Entertainment to our mind.

— From 'Pangur Bán', English translation by Robin Flower

1 Introduction

The problem of how a plasma will evolve, and to what equilibrium state, if any, it evolves towards, is an old one, central to kinetic theory and fundamental to our understanding of the large-scale properties of many astrophysical systems. In neutral gases, it was argued by Maxwell (1860) that the distribution function of particles undergoing elastic collisions would relax to what is now known as the Maxwellian (or Maxwell–Boltzmann) distribution but it was not until Boltzmann (1896) used the *Stosszahlansatz* to derive his collision integral that Maxwell's prediction was put on dynamical footing. Neutral particles, of course, greatly simplify the problem due to their short-range interactions. For long-range interactions, the same dynamical footing would wait for Landau (1936), by considering the Coulomb cross section of charged particles, and would later be formalised by Balescu (1960) and Lenard (1960). From such collisional theories of plasma relaxation, one might naively suppose that the universe is full to the brim with plasma in Maxwell–Boltzmann equilibrium. There are however, many examples in our Universe of plasmas that are not in Maxwellian equilibrium. This is true in many settings, from cosmic rays populating our Galaxy and beyond (Becker Tjus &Merten, 2020), down through the solar neighbourhood (Oka et al., 2018; Wilson et al., 2022), and all the way to Earth-based experiments (Cruz et al., 2018a; Magee et al., 2019; Hartouni et al., 2022). It is not at all difficult however, to motivate why this should be the case. The

long-range interactions present in plasmas allow for an evolution driven by mean (i.e., generated on length scales far larger than the interparticle separation) electric and magnetic fields that the plasma inherently generates. In systems with large-scale (system-size) gradients, these fields may be driven unstable, triggering not just a zoo but an entire ecosystem of plasma instabilities (see, e.g., Krall &Trivelpiece 1973; Bott et al. 2024), with fluctuations growing and reaching amplitudes at which they can react back upon their progenitors, potentially altering the mean distribution and relaxing the system towards some semblance of stability, but by no means Maxwellianity. Despite all this, the distributions that we observe do possess some universal features or, at least, fall into a finite number of universality classes (such as the power-law tails distribution seen in numerical simulations of turbulent astrophysical systems, e.g., Sironi &Spitkovsky (2010); Kunz et al. (2016); Zhdankin et al. (2017), or the so-called "kappa" distributions and flat-topped distributions seen in the solar wind, e.g., Marsch 2006; Pierrard &Lazar 2010; Fisk &Gloeckler 2014; Verscharen et al. 2019).

Motivated by the success of maximum-entropy formulations in predicting the relaxed state of collisional plasmas, there have been multiple attempts to justify the ubiquity of such distributions from a thermodynamic point of view. This is, however, entangled with the question of whether the standard Gibbs–Shannon entropy is applicable to systems with long-range interactions and, if not, then what entropy should be used. Naturally, many entropies have emerged to fill this niche. A popular contender is the Tsallis (1988) entropy (originally called α-structural entropy Havrda &Charvát 1967). The Tsallis entropy was designed to be a non-additive version of the Gibbs–Shannon entropy as a way to model systems with long-range correlations that therefore should not have the additive property inherent to Shannon's entropy[1](see, e.g., Livadiotis &McComas 2009; Pierrard &Lazar 2010 and references therein). While this model produces good fits to observed distributions, it

[1]i.e., that the entropy of two independent systems should split. If σ_1 are the states of system 1 with probabilities p_{σ_1} and σ_2 the states of system 2 with probabilities p_{σ_2} then the probability of joint system being in state (σ_1, σ_2) is $P_{(\sigma_1,\sigma_2)} = p_{\sigma_1} p_{\sigma_2}$ (provided the systems are independent) and the Shannon entropy is naturally additive, viz.,

$$-\sum_{\sigma_1,\sigma_2} P_{(\sigma_1,\sigma_2)} \ln P_{(\sigma_1,\sigma_2)} = -\sum_{\sigma_1} p_{\sigma_1} \ln p_{\sigma_1} - \sum_{\sigma_2} p_{\sigma_2} \ln p_{\sigma_2} \tag{1.1}$$

has a free parameter that is needed to quantify the degree of the non-extensivity and cannot be determined without fitting data, or additional input of physics currently lacking (note, however, some recent progress suggesting that this additional physics might be deducible from free-energy considerations: Zhdankin 2022a,b).

An early attempt to tackle the question of entropy in collisionless systems was made by Lynden-Bell (1967). He proposed to consider a system of N particles with canonical positions $\boldsymbol{x}_i$ and momenta $\boldsymbol{p}_i$ that evolve subject to a Hamiltonian $\mathcal{H}(\boldsymbol{x}_i, \boldsymbol{p}_i)$. From such a system one may derive (via, e.g., the BBGKY hierarchy; Karder 2007) an evolution equation for the single-particle distribution function $f(\boldsymbol{x}, \boldsymbol{v})$ from the distribution of N particles. Formally, when we say such a system is 'collisionless' we mean that the evolution equation for the single-particle distribution function, is well approximated by an effective Hamiltonian $\mathcal{H}^{\rm eff}(\boldsymbol{x}, \boldsymbol{p})$ acting on a single particle (i.e., that the mean-field dynamics are a sufficiently good approximation to the true dynamics), viz.,

$$\frac{\partial f}{\partial t} + \frac{\partial \mathcal{H}^{\rm eff}}{\partial \boldsymbol{p}} \cdot \frac{\partial f}{\partial \boldsymbol{x}} - \frac{\partial \mathcal{H}^{\rm eff}}{\partial \boldsymbol{x}} \cdot \frac{\partial f}{\partial \boldsymbol{p}} = 0. \tag{1.2}$$

In his original treatment, Lynden-Bell focused on relaxation of stellar systems, but the spirit of his statistical mechanics is the same for all collisionless systems, including plasmas. While keeping the calculations as general as possible, one can think of (1.2) as the collisionless Vlasov equation for a plasma, which could be electrostatic or electromagnetic in a non-relativistic or relativistic regime. The collisionless dynamics described by (1.2) conserve an infinite number of invariants, namely the integral of any function $G(f)$ of the distribution function over phase space $(\boldsymbol{x}, \boldsymbol{p})$

$$\int \mathrm{d}\boldsymbol{x}\, \mathrm{d}\boldsymbol{p}\, G(f) = \text{const}, \tag{1.3}$$

equivalent to conserving the volume of level sets of the distribution function $f(\boldsymbol{x}, \boldsymbol{p})$. Thus, the dynamics can be viewed as an extremely complicated rearranging of the elements of phase space, which, however much they are distorted and stirred, will keep the same level sets (often referred to as 'waterbags', in analogy with parcels of incompressible fluid). Lynden-Bell posited that, after a short time, the exact phase-space density $f(\boldsymbol{x}, \boldsymbol{p})$ would become so chaotic that it could be treated as a random field and that any measurement of it—in practice, of a coarse-grained

version of it[2]—was in fact a measurement of the mean phase-space density. This allowed the construction of a statistical mechanics, with an entropy closely related to the Gibbs–Shannon entropy, that encoded an infinite number of invariants and thus predicted the steady states from a given initial condition. These steady states are the Lynden-Bell equilibria. When they were first developed, they were applied to the (somewhat pathological) case of a self-gravitating system, and lacked the computational resources to be fully solved or properly explored. Their solution, and adaptation to the (more well-behaved) context of plasma physics, will be an overarching theme of the following chapters. In the process of doing so, we will uncover certain important nuances and pitfalls that could not have been foreseen at the inception of the theory. As a guide for the remainder of this text, we will present the derivation of the Lynden-Bell equilibria as it existed in 1967, and discuss some hanging questions that arises from it, which will naturally motivate the subsequent chapters.

1.1 Lynden-Bell's statistical mechanics

In this section, we present a brief re-derivation of Lynden-Bell's theory as applied to homogeneous systems. We begin with the collisionless Vlasov equation (1.2) evolving a single species of particles. As well as particle number, momentum, and total energy (i.e., of fields and particles), (1.2) conserves an infinite number of 'Casimir' invariants (1.3) which can be encapsulated by the volume of phase space where the exact phase-space density is greater than a given value η:

$$\Gamma(\eta) = \iint \mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{p}\,\Theta(f(\boldsymbol{x},\boldsymbol{p}) - \eta) = \text{const}, \tag{1.4}$$

where $\Theta(x)$ is the Heaviside function (unity for $x > 0$ and zero otherwise). As discussed above, despite the existence of these invariants, the system's evolution can still be highly chaotic. This is precisely what is seen in figure 1.1, which illustrates the paradigmatic evolution of a turbulent collisionless system: the plasma, in this case initialised as two cold counter-propagating electron beams, goes unstable and is mixed until a myopic eye cannot distinguish the forest (mean phase-space density)

[2]In this work we will apply the Lynden-Bell theory to the distribution function coarse grained over all positions, assuming statistical homogeneity. In fact this simply requires that you average over sufficiently many structures to achieve good statistics.

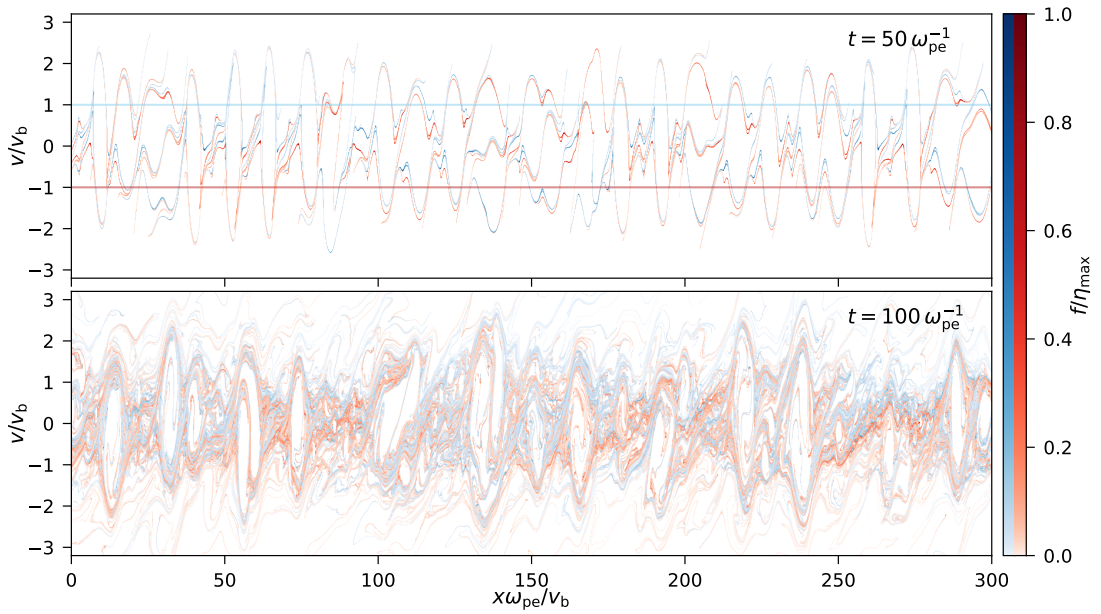

Figure 1.1: Self-consistent particle-in-cell simulation of the early-time temporal evolution of the phase-space density $f(x, v)$ for the electron-only two-stream instability. The two times, $t = 50\,\omega_{pe}^{-1}$ and $t = 100\,\omega_{pe}^{-1}$, are shown in the top and bottom panels, respectively, along with the initial condition (two horizontal lines) overlaid in the top panel. The leftward going beam is drawn in blue and the rightward in red to illustrate the initial beams becoming ever more tangled and chaotic but approximately retaining their phase volume (the full discussion of these simulations, and the sense in which the phase-volume conservation is only approximate, will be discussed in chapter 2).

from the trees (exact phase-space density). Because of this disorder, and in light of the conservation of Casimir invariants (1.4), it is helpful to consider the exact phase-space density $f(\boldsymbol{x}, \boldsymbol{v})$ to be a random quantity and to describe the statistical state of the plasma by the probability density $P(\boldsymbol{x}, \boldsymbol{p}, \eta)$ for the exact phase-space density $f(\boldsymbol{x}, \boldsymbol{p})$ to take the value η at position $(\boldsymbol{x}, \boldsymbol{p})$ (Robert &Sommeria, 1991; Chavanis et al., 1996). The distinct values of η will be referred to as 'waterbags' since this term conjures up the correct mental image: parcels of phase space of a certain density that can be distorted and moved, but not rarefied, compressed, or superimposed. This is indeed what is seen in figure 1.1, with the streams being stretched and wound around one another by the turbulence but (approximately) maintaining their total volume. The mean phase-space density is then

$$\langle f \rangle(\boldsymbol{p}) = \int \mathrm{d}\eta\, \eta P(\boldsymbol{p}, \eta). \tag{1.5}$$

Here, we have applied the intuition that the steady-state distribution function $P(\boldsymbol{p}, \eta)$ will be homogeneous in space (this contrasts with Lynden-Bell's original treatment, which focused on gravitationally bound, and therefore inhomogeneous, systems). Lynden-Bell's statistical mechanics amounts to positing that, before the onset of 'true' collisions, $P(\boldsymbol{p}, \eta)$ will maximise the Gibbs–Shannon entropy

$$S = -\iint \mathrm{d}\boldsymbol{p}\, \mathrm{d}\eta\, P(\boldsymbol{p}, \eta) \ln P(\boldsymbol{p}, \eta). \tag{1.6}$$

Note that the integral in η must run over all the possible values, including $\eta = 0$ (the empty waterbag). The choice of Gibbs–Shannon entropy, although unique by Shannon's theorem (Shannon, 1948), can be further justified to the distrustful reader by a combinatoric argument, which we present in appendix A.

Equipped with an entropy ripe for maximisation, we must decide upon a set of reasonable constraints under which to maximise it. Naturally, since $P(\boldsymbol{p}, \eta)$ is a probability-density function in η at a given $\boldsymbol{p}$, its integral in η at any $\boldsymbol{p}$ must equal unity:

$$\int \mathrm{d}\eta\, P(\boldsymbol{p}, \eta) = 1. \tag{1.7}$$

As well as this, we fix the total energy density of the system:

$$\iint \mathrm{d}\boldsymbol{p}\, \mathrm{d}\eta\, \varepsilon(\boldsymbol{p}) \eta P(\boldsymbol{p}, \eta) = E = \text{const}, \tag{1.8}$$

where $\varepsilon(\boldsymbol{p})$ is the energy of a particle as a function of its momentum $\boldsymbol{p}$. Note that, within this formalism, one could include the interaction energy of particles with coherent stationary fields (electromagnetic, gravitational, etc.), so that in its most general form ε would be a function of both position and momentum, which could need to be solved self-consistently with P (see, e.g., Gruzinov et al. 2020). In this thesis, we will neglect this rich complexity, assuming instead that, in the relaxed state, the energy of the fields has decayed to a negligible fraction of the total energy and, in the process of decaying, has mediated the relaxation of the distribution function[3].

Next, we enforce the conservation of the Casimir invariants (1.4) by requiring that the volume-integrated probability of each waterbag stays constant, viz.,

$$\int \mathrm{d}\boldsymbol{p}\, P(\boldsymbol{p}, \eta) = \rho(\eta) = \text{const}. \tag{1.9}$$

[3]Indeed, in the simulations of the two-stream instability, shown in chapter 3, this is justified by the observation that the electric-field energy decays with time and, despite the strong stirring of the initial instability, never exceeds 10% of the total energy

The function $\rho(\eta)$ will be referred to as the 'waterbag content' and is determined by initial conditions. The waterbag content of the initial condition can be read off by integrating over all portions of phase space where the initial exact phase-space density is equal to a particular value, viz.,

$$\rho(\eta) = \frac{1}{V}\iint \mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{p}\,\delta\Big(\eta - f(\boldsymbol{x},\boldsymbol{p},t=0)\Big) = -\frac{1}{V}\frac{\mathrm{d}\Gamma}{\mathrm{d}\eta}, \tag{1.10}$$

where V is the system's spatial volume[4]. *A priori*, in Lynden-Bell's statistical mechanics, the degree of universality of the equilibrium distribution is determined by $\rho(\eta)$: all initial conditions with the same waterbag content and energy lead to the same equilibrium.

Maximising the entropy (1.6) subject to the constraints (1.7), (1.8), and (1.9)[5] is equivalent to maximising, unconditionally, the functional

$$S[P(\boldsymbol{p},\eta)] - \int \mathrm{d}\boldsymbol{p}\,\lambda(\boldsymbol{p})\left[\int \mathrm{d}\eta\,P(\boldsymbol{p},\eta) - 1\right] - \beta\left[\iint \mathrm{d}\boldsymbol{p}\,\mathrm{d}\eta\,\varepsilon(\boldsymbol{p})\eta P(\boldsymbol{p},\eta) - E\right] + \beta\int \mathrm{d}\eta\,\eta\mu(\eta)\left[\int \mathrm{d}\boldsymbol{p}\,P(\boldsymbol{p},\eta) - \rho(\eta)\right], \tag{1.11}$$

where $\lambda(\boldsymbol{p})$, β and $-\beta\eta\mu(\eta)$ are Lagrange multipliers. By analogy with textbook statistical mechanics, we will sometimes refer to $\mu(\eta)$ as the 'chemical potential' (which it is, being the Lagrange multiplier that fixes the number of particles in waterbag η). Doing so, we find the Lynden-Bell equilibria

$$P(\boldsymbol{p},\eta) = \frac{e^{-\beta\eta[\varepsilon(\boldsymbol{p})-\mu(\eta)]}}{\int \mathrm{d}\eta'\,e^{-\beta\eta'[\varepsilon(\boldsymbol{p})-\mu(\eta')]}}, \tag{1.12}$$

[4]It is important to note that this expression for the waterbag content (1.10) is only meaningful for continuous distribution functions f. For a system of N particles, such a distribution function can be found averaging the Klimontovich distribution over sufficiently small volumes in phase space. These volumes should therefore be sufficiently large that they contain many particles and define a continuous phase-space density f but sufficiently small so that macroscopic structure is not erroneously averaged. It is ultimately this small-scale averaging that will drive the evolution of $\rho(\eta)$ we will find in chapter 3

[5]We note that, while we have endowed the invariants (1.7), (1.8), and (1.9) with special significance as constraints, there may be situations where additional invariants are necessary. For instance, in strongly magnetised plasmas, relaxation may occur before the conservation of particles' magnetic moments are broken. In such cases, further invariants would be necessary and would alter the character of the solution (cf. Helander 2017). Here we shall consider only systems where the fields driving the relaxation may be arbitrary, but the only quantities conserved on the relaxation time scale are (1.7), (1.8), and (1.9).

where $\lambda(\boldsymbol{p})$ has been computed explicitly to arrange for the correct normalisation (1.7), whereas β and $\mu(\eta)$ must be chosen in such a way as to satisfy the constraints (1.8) and (1.9). We note that, despite (1.9), the mean phase-space density $\langle f \rangle$, given by (1.5), will, in general, not have the same level sets (1.4) as the exact one f (since $\langle f \rangle$ is an averaged quantity). The equilibria (1.12) are both homogeneous and isotropic: an inevitable consequence of the system having no preferred position or direction.

The similarity between the Lynden-Bell equilibria (1.12) and the Fermi–Dirac distribution is immediately apparent[6]. This should come as no surprise, because phase-volume conservation functions analogously to Pauli's exclusion principle: pieces of the same waterbag, or different waterbags, cannot cohabit in phase space. The equilibria, therefore, have degeneracy effects incorporated within them.

The prescription for computing Lynden-Bell equilibria is now clear: given an initial condition, with the initial energy density E and waterbag content $\rho(\eta)$ (determined by (1.10)), solve two coupled integral equations (1.8) and (1.9) with $P(\boldsymbol{p}, \eta)$ given by (1.12), determine β and $\mu(\eta)$, and substitute back into (1.12). This is the prescription laid out by Lynden-Bell in 1967, which can only be solved analytically in a handful of cases. The first algorithm to systematically solve such general equilibria was presented in Ewart et al. (2023) and will be discussed in detail in chapter 2 and appendix C. Before this, however, we will immediately remark on the open questions that this formalism leaves, which will be addressed in the subsequent chapters.

1.2 Universality of Lynden-Bell equilibria

As will be pointed out numerous times in the ensuing chapters, statistical mechanics, once applied, represents an app—a device into which one plugs information and receives information in return. For collisional systems, such statistical mechanics is trivial—the app requires only to be given the total energy of the system and its density—and the outcome is starkly universal: all maximum-entropy distributions will be Maxwellian. For the statistical mechanics of collisionless systems, the new piece of information that must be provided arrives in the form of $\rho(\eta)$ and it is not

[6]Indeed, the Fermi–Dirac distribution can be thought of as the special case of a two-level-set system, which further reduces to the Maxwell–Boltzmann distribution when degeneracy is neglected (see, e.g., Chavanis 2006a, or chapter 4, for details).

at all clear that the resulting equilibria will be universal in any way—precisely for the reason that additional conservation laws inherently imply that the resulting equilibria have extra parameters. Famously, only four parameters are required to describe an elephant, five to wiggle its trunk (Dyson, 2004); it is clear that with the continuous set of parameters provided by phase-volume conservation, encoded in $\rho(\eta)$, one could make the elephant write an entire thesis. We therefore tackle the question of how universal these equilibria are in chapter 2. It will turn out that the strong resemblance between the Lynden-Bell equilibria and the Fermi–Dirac distribution allows one to show that these equilibria will indeed possess classes of universality. Like the Fermi–Dirac distribution, there are two relevant limits of the Lynden-Bell equilibria: degenerate and non-degenerate. In the degenerate case, as in the Fermi–Dirac distribution, the system is frozen into its ground state. The fact that such non-trivial ground states (known as Gardner states, see, e.g., Gardner 1963; Dodin &Fisch 2005; Helander 2017) exist, and are in one-to-one correspondence with $\rho(\eta)$, is guaranteed by phase-volume conservation (see, e.g., Dodin &Fisch 2005; Helander 2017; Mackenbach et al. 2022). These ground states will be highly non-universal, precisely for the reasons given above. However, in chapter 2, we will show analytically that, at energies much larger than the ground state energy, the Lynden-Bell equilibria become sensitive to only a single property of $\rho(\eta)$: its low-η asymptotic. We further show that this low-η asymptotic is universal to a wide variety of ground states (all those which decay exponentially at large velocities), and the the corresponding excited states have a universal power-law tail in energy scaling like ε^{-2}.

1.3 Turbulent amnesia

Of course, there is no such thing as a truly collisionless system. Ultimately the finite number of particles in any system means that (1.2) will never be perfectly valid and molecular chaos, however small, will eventually reassert itself. Typically this is included in the form of a collisional term $C[f]$ appended to (1.2) (cf. Helander &Sigmar 2005) and is associated with a typical "collision frequency" ν_{coll}. It is these collisions which would push the mean distribution function of a quiescent plasma (meaning with no spatial variation) towards a Maxwellian equilibrium on

a time scale like $\nu_{\rm coll}^{-1}$. It would therefore be tempting to argue (naively) that the collisionless dynamics creates a Lynden-Bell equilibrium on times $t \ll \nu_{\rm coll}^{-1}$ which is slowly eroded to a Maxwellian on times $t \sim \nu_{\rm coll}^{-1}$. This is only half true. In a turbulent plasma—precisely the kind which creates the strong mixing envisaged by Lynden-Bell—one must inherently create very small-scale structure in position and velocity space, plain to see from figure 1.1. Collisions, being a diffusion in phase space, naturally erase such small-scale structures, altering the previously conserved quantities $\rho(\eta)$. Furthermore, since this small-scale structure is formed on dynamical, not collisional, time scales, it is possible for the conservation of $\rho(\eta)$ to be broken on time scales much shorter than $\nu_{\rm coll}^{-1}$. Estimates of this effect based on linear theory (cf. Su &Oberman 1968) would suggest it occurs at time $t \sim \omega_{\rm pe}^{-2/3}\nu_{\rm coll}^{-1/3}$ (where $\omega_{\rm pe}$ is the electron plasma frequency), while more recent theories of a nonlinear cascade in phase space (cf. Nastac et al. 2024b) suggest that collisional effects could be active on time scales as early as $t \sim \omega_{\rm pe}^{-1}\ln(\omega_{\rm pe}/\nu_{\rm coll})$— an extremely weak dependence on collisionality. It may therefore seem fruitless to hope that the Lynden-Bell equilibria could ever be reached, since the conservation of $\rho(\eta)$ is always in jeopardy. However, this too would be naive. We show in chapter 2, that the breaking of $\rho(\eta)$ does not preclude the system from reaching a Lynden-Bell equilibrium, it simply changes the Lynden-Bell equilibrium that the system reaches. The waterbag content $\rho(\eta)$, therefore, no longer represents the perfect memory of the system's initial conditions, but rather an altered short-term memory, engendered by the very mixing which relaxes the system towards its Lynden-Bell equilibrium in the first place. We demonstrate this phenomenon, which we refer to as 'turbulent amnesia', with first principles numerical simulations of one of the most canonical instabilities in plasma physics: the two-stream instability. In these simulations, which we believe to be the largest, longest run, and best resolved simulations of the two-stream instability to date, we demonstrate that the Lynden-Bell equilibria are reached and that the low-η asymptotic conjectured on theoretical grounds in chapter 2 is developed, giving rise to the predicted ε^{-2} tail and indicating that a small amount of collisionality actually makes the evolution more universal.

1.4 Dynamical processes driving relaxation.

Having demonstrated, at least within the context of the two-stream instability, that it is possible for systems to relax to the universal equilibria predicted in chapter 2, another natural question is 'how is this relaxation mediated?'. This is the same question Boltzmann asked of Maxwell—natural on the grounds that thermodynamic approaches offer little insight into how the disorder, which they require, is generated by the system. The task is therefore a simple one: construct an evolution equation for the mean distribution function (in this context referred to as a 'collision integral') for which there exists an H-theorem which guarantees the system's relaxation to a maximum-entropy state. In chapter 4 we will show that this can be done for a collisionless system, deriving a 'collisionless collision integral' which relaxes the system towards a Lynden-Bell equilibrium. The derivation, which is quasilinear, is very similar to that of the conventional collision integrals with the main task being to establish a closure for the average fluctuation level of the system in terms of the mean phase-space density $\langle f \rangle$. In physical terms, the injection of entropy in all collision integrals comes from some version of a molecular-chaos assumption. Particles scatter off of some object (e.g., other particles, waves, etc.) with which they are not correlated. In our case, the challenge is to describe fluctuations that, despite being uncorrelated, are inherently aware of the conserved quantities (1.3). We develop two such closure schemes by which such collision integrals can be derived and compare the resulting expression to the standard collision integrals of Landau, Balescu, and Lenard. Lastly, we demonstrate that relaxation mediated by these collision integrals has interesting and nontrivial behaviour on its journey to the Lynden-Bell equilibrium exhibiting features such as anomalous interspecies drag and spontaneous current generation.

If I were confined to know
tracks of orbits as they go,
glancing off one, two and fro,
through long-range interaction.

Then, at least, it seems to me,
by the force that I can see
there's no such thing, as entropy
just loss of information.

— From the author, on Maxwell's demon

2

Universal Lynden-Bell equilibria

This chapter is adapted from Ewart et al. 2023, J. Plasma Phys. ***89****. 905890516*

2.1 Introduction

Since its genesis, Lynden-Bell's theory (often referred to as the theory of 'violent relaxation') has received continued attention both thermodynamically (Chavanis et al., 1996; Arad &Johansson, 2005; Chavanis, 2006b,a; Levin et al., 2008, 2014) and dynamically, viz., effective 'collisionless collision integrals' have been proposed that recovered Lynden-Bell equilibria as their fixed points (Kadomtsev &Pogutse, 1970; Severne &Luwel, 1980; Chavanis, 2004, 2022). However, the main strength of the theory is also its weakness. Unlike in the non-extensive entropy formulations (e.g., Tsallis 1988; Zhdankin 2022b) discussed in chapter 1, there is no *ad-hoc* parameter in the Lynden-Bell theory: equilibria are uniquely determined by the 'waterbag content' of the initial conditions. However, this necessarily means that the equilibria 1.12 depend (seemingly, in a complicated way) on an infinite family of invariants. This has limited any actual calculations with Lynden-Bell equilibria to simplified situations with only a small number of level sets (in practice, between one and three, e.g., Assllani et al. 2012). At any rate, the intricate dependence on an infinite family of invariants might not appear to be a step towards general power-law tails or any

other meaningful form of universality. Nevertheless, collisionless systems naturally produce universal features. Most notably, distributions with power-law tails appear in a myriad of plasma systems including cosmic rays (Becker Tjus &Merten, 2020; Amato &Casanova, 2021), the solar corona and solar flares (Dudík et al., 2017; Oka et al., 2018), the solar wind (Gloeckler et al., 2008; Fisk &Gloeckler, 2014; Livadiotis et al., 2018; Moncuquet et al., 2020), the Earth's magnetosheath (Birn et al., 2012; Ergun et al., 2020), and laser plasmas (Cruz et al., 2018b; Hartouni et al., 2022). By way of justifying these observations, direct numerical simulations have indicated that power-law tails are the natural result of a number of dynamical processes including relativistic and non-relativistic shocks (Sironi &Spitkovsky, 2010; Caprioli &Spitkovsky, 2014; Crumley et al., 2019), magnetic reconnection (Sironi &Spitkovsky, 2014; Werner &Uzdensky, 2021; Uzdensky, 2022), and various types of plasma turbulence (Kunz et al., 2016; Comisso &Sironi, 2018, 2022; Zhdankin et al., 2017, 2019; Zhdankin, 2021, 2022b).

To see just how non-universal Lynden-Bell equilibria can be, one only need to consider the set of all monotonically decreasing functions of energy (Gardner distributions). Should the initial distribution be a monotonically decreasing function of particle energy, then it is a Gardner distribution and there are no possible rearrangements of the phase volume that do not increase energy. Hence the only state available via collisionless dynamics is this Gardner distribution, which must therefore be its own Lynden-Bell equilibrium[1]. But since *any* monotonically decreasing function of energy is a Gardner distribution, these minimum-energy states are clearly highly non-universal. However, this is only a good intuition for systems where the number of level sets is small or where phase-volume conservation conspires with energy conservation to render much of the phase space inaccessible to the system (as is the case for Gardner distributions). In this chapter, by solving for the full Lynden-Bell equilibria numerically (as well as analytically, in a tractable limit), we will show that most Lynden-Bell equilibria are much more generic. Namely, we will show that, in the limit of a continuum of level sets, and for energies sufficiently greater than the ground-state energy (the energy of the corresponding Gardner distribution), the Lynden-Bell equilibria exhibit power-law tails at high energies, typically with a scaling of ε^{-2}, where ε is the particle energy.

[1]This takes a surprising amount of work to show formally: see appendix B.

The physical argument for these power-law tails is as follows. Phase-volume conservation effectively makes the particles occupying each waterbag (level set of the distribution) behave as if they were members of a separate species, which can communicate with the other waterbags only via the equilibration of some effective 'temperature' subject to competition for the same volumes of phase space. In essence, this turns the system into a ensemble of many different fermionic species, all of which exclude each other. When the system is in its minimum-energy (ground) state, the Gardner distribution, the competition for phase volume is the overpowering factor, giving the distribution a highly non-universal shape. However, when the energy of the system is increased, more of the phase space becomes accessible, so, as in Fermi–Dirac statistics, the competition for the same phase volume becomes weaker. For sufficiently large energies, the competition for any volume of phase space is minuscule. In this limit, each waterbag will form its own Maxwellian distribution, in thermal equilibrium with all other waterbags. However, despite these Maxwellian equilibria having the same effective temperature, they will have different thermal spreads because waterbags of larger phase-space density 'cost' more energy to be placed at a given momentum p in phase space. The true distribution function is recovered by summing up (in the limit of many waterbags, integrating) the contributions from each of these Maxwellians to the mean phase-space density. This procedure naturally gives rise to a power-law tail that depends on the relative weighting for each Maxwellian[2]. This weighting turns out to depend only weakly on the level sets of the initial condition. For a wide class of initial conditions, the resulting Lynden-Bell equilibria turn out to have the same universal power-law tail, ε^{-2}.

The rest of this chapter is organised as follows. In Section 2.2.1, we will argue that to each initial condition, one can uniquely assign a Gardner distribution function with the same Casimir invariants (waterbag content). All Lynden-Bell equilibria can thus be viewed as the result of adding some amount of energy to a Gardner distribution with the same waterbag content and letting it reach a maximum-entropy state (one can think of this approach as describing how adding energy to a population of collisionless particles causes them to form a 'non-thermal'

[2]This is similar to how Zipf's law arises in systems where one marginalises over a 'hidden variable' (cf. Mora &Bialek 2011; Schwab et al. 2014; Aitchison et al. 2016), and to the formalism of 'superstatistics'; cf. Beck &Cohen 2003; Chavanis 2006a; Davis et al. 2023)

distribution with a tail). In Section 2.2.2, we will show that the function describing the waterbag content of a large class of Gardner distributions has a relatively generic form, which will contribute to the universality of the resulting equilibria. In Section 2.2.3, we will solve for the Lynden-Bell equilibria in the limit where the energy of the system far exceeds the energy of the corresponding Gardner distribution. This will ensure that competition for volumes of phase space can be neglected. This makes the problem analytically tractable, and the resulting analytical solution will exhibit the universal power-law tail $\propto \varepsilon^{-2}$ at high energies. In Section 2.3, by solving for the Lynden-Bell equilibria numerically, we will show that the qualitative features of this analytical solution are retained even for energies that are of the same order as the energy of the Gardner distribution. Therefore, a large class of Lynden-Bell equilibria display a universal power-law tail. This tail contains much of the distribution's energy, whereas the low-energy 'core' retains some dependence on the initial conditions. In Section 2.4, we summarise our findings and discuss their implications for real (observed) plasmas.

2.2 Theory: degenerate and non-degenerate equilibria

2.2.1 Lynden-Bell equilibria as excited Gardner distributions

Just as it is only meaningful to consider a Maxwellian with a positive energy, it is only meaningful to solve the Lynden-Bell equilibria (1.12) subject to reasonable choices of the constraints (1.8) and (1.9). It is therefore instructive to understand the properties of the waterbag-content function (1.9). To get a feel for waterbag contents of typical initial conditions, one could compute the integral (1.10) for a range of examples (this is relatively simple due to the presence of the delta function). One quickly discovers that many different initial conditions have similar waterbag contents, just as many different initial conditions can have the same energy. To see this, we note that, from (1.10), any volume-preserving transformation of the coordinates $(\boldsymbol{x}, \boldsymbol{p})$, including those transformations that 'splice' the phase space discontinuously, will leave the waterbag content unchanged. This is unsurprising because the true,

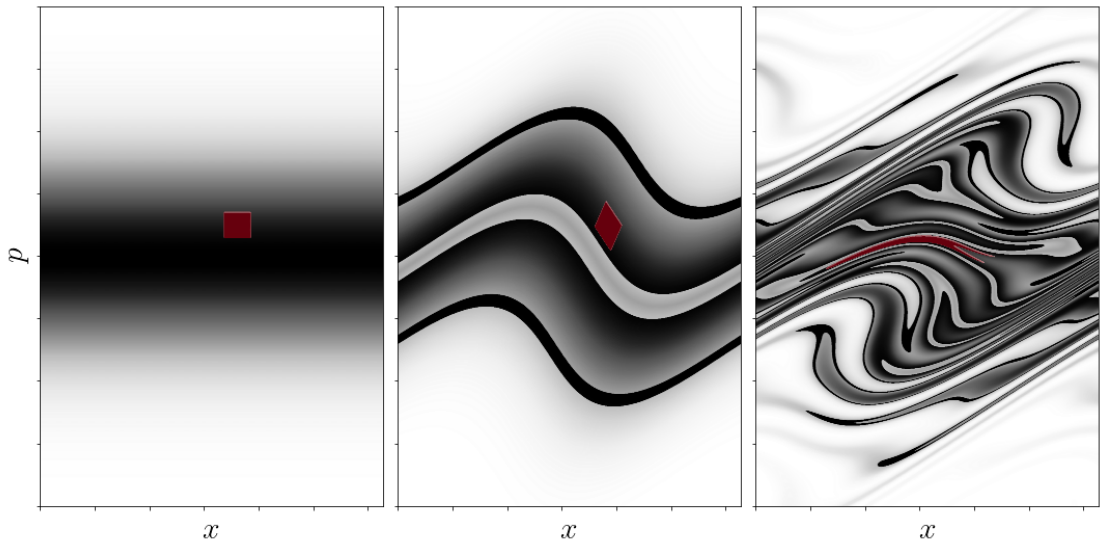

Figure 2.1: A cartoon contour plot in phase space of three possible distribution functions, all of which possess identical waterbag contents. Panel (a) shows the Gardner distribution function corresponding to this waterbag content. Panels (b) and (c) show distributions, at different (higher) energies, which can be reduced to the Gardner distribution by deforming and splicing the phase space incompressibly. A small patch of phase space is highlighted in red between plots to show the effect of the deformation.

incompressible, flow of probability in phase space is precisely one such volume-preserving transformation. It is this freedom that implies that vastly different initial conditions can possess identical, or similar, waterbag contents. A cartoon illustrating this is given in Figure 1.1, showing how seemingly complex distributions have the same waterbag content as very simple distributions. There will be families of initial conditions that have the same waterbag content, but different energies.

For every family of initial conditions possessing the same waterbag content, there will be a unique distribution function that has that waterbag content but is a monotonically decreasing function of energy and, therefore, has the minimum possible energy associated with that waterbag content. Such a distribution function, for which the exact phase-space density satisfies $f(\boldsymbol{x}, \boldsymbol{p}) = f_{\mathrm{G}}(\varepsilon(\boldsymbol{p}))$, is known as the Gardner distribution (Gardner, 1963; Helander, 2017). The sequence of deformations of the distribution function to map an initial condition to its Gardner distribution function is often referred to as a 'restacking', as it amounts to a reordering of phase-space elements into their minimum-energy configuration (Dodin &Fisch, 2005; Kolmes et al., 2020; Kolmes &Fisch, 2020). Gardner distributions can be viewed as

'ground states' associated with a given waterbag content (e.g., Helander 2017), since no more energy can be extracted from such a distribution without violating phase-volume conservation. This fact intuitively guarantees that any initial condition that is a Gardner distribution is its own Lynden-Bell equilibrium since no other states are available to the system. Indeed, one can show that any Gardner distribution can be reconstructed from (1.12) for a particular choice of β and $\mu(\eta)$, although the proof is technical and left to appendix B.

A generic initial condition can then be viewed as equivalent to taking some Gardner distribution and driving it out of equilibrium by the injection of some energy without changing the waterbag content. The Lynden-Bell equilibria are then simply the collisionless, phase-volume preserving, entropy-maximising equilibria of these higher-energy states, making them the natural excited states of Gardner distributions. Therefore, to capture the set of all possible waterbag contents, we need only study the set of all these 'ground states', to which we would then add energy—the first step in the direction of universal outcomes. Physically, this approach is equivalent to asking to what distribution a population of collisionless particles will relax once a certain amount of energy is injected into it—in a manner of speaking, a 'thermodynamical' approach to 'non-thermal' particle acceleration.

2.2.2 Waterbag content of Gardner distributions

Having stated the problem in this way, we now consider the waterbag content associated with Gardner distributions. In what follows, we will consider Gardner distribution functions that are truncated at some minimum phase-space density $\eta_{\min}$. This mathematical convenience will turn out to be a physical necessity. Thankfully, while it is mathematically and physically important that the cutoff $\eta_{\min}$ be finite, it will only appear logarithmically in the outcomes of our calculations, making them highly insensitive to its actual value—yet another theory where the need for a cutoff is unavoidable but non-lethal.

As a prototypical example, we compute $\rho(\eta)$ for a particular Gardner distribution: a truncated Maxwellian, viz.,

$$f_{\mathrm{G}}(\boldsymbol{p}) = \begin{cases} \eta_{\max} e^{-\varepsilon(\boldsymbol{p})/\varepsilon_0} & \text{for} \quad \varepsilon(\boldsymbol{p}) < \varepsilon_0 \ln \dfrac{\eta_{\max}}{\eta_{\min}}, \\ 0 & \text{for} \quad \varepsilon(\boldsymbol{p}) > \varepsilon_0 \ln \dfrac{\eta_{\max}}{\eta_{\min}}. \end{cases} \tag{2.1}$$

Besides $\eta_{\min}$, the parameters of this distribution are the energy scale ε_0 and the maximum phase-space density $\eta_{\max}$. The latter is straightforwardly related to the particle's spatial density, e.g., $n_0 = \eta_{\max}(2\pi m\varepsilon_0)^{3/2}$ in the limit $\eta_{\min}/\eta_{\max} \to 0$ for a 3D, non-relativistic plasma, where $\varepsilon(\boldsymbol{p}) = p^2/2m$. The waterbag content of the Gardner distribution for such a plasma is then, from (1.10),

$$\rho(\eta) = \Gamma_{\text{free}}\delta(\eta) + \begin{cases} \dfrac{2n_0}{\sqrt{\pi}\eta_{\max}\eta}\left(\ln\dfrac{\eta_{\max}}{\eta}\right)^{1/2} & \text{for} \quad \eta_{\min} < \eta < \eta_{\max}, \\ 0 & \text{otherwise,} \end{cases} \tag{2.2}$$

where Γ_{free} is the total volume of the momentum space that is unoccupied (i.e., where the exact phase-space density is zero). Of course, in reality, momentum space is unbounded and so Γ_{free} is infinite. Formally, we are solving for the waterbag content and Lynden-Bell equilibrium in a momentum space of large, but finite, volume and will take this volume to infinity at the end of the calculation—of course, nothing physical will depend on Γ_{free} as it becomes large.

The presence of a $\delta(\eta)$ term in the expression for $\rho(\eta)$ is a generic feature, not restricted to the specific example (2.2). When it comes to solving (1.9) with $P(\boldsymbol{p}, \eta)$ given by (1.12), in order to find $\mu(\eta)$, this delta function can be accommodated by writing the chemical potential as

$$e^{\beta\eta\mu(\eta)} = \eta_{\text{ref}}\delta(\eta) + \begin{cases} \eta_{\text{ref}}F(\eta) & \text{for} \quad \eta_{\min} < \eta < \eta_{\max}, \\ 0 & \text{otherwise,} \end{cases} \tag{2.3}$$

where η_{ref} is some reference constant that must have dimensions of phase-space density. Its value is unimportant because $e^{\beta\eta\mu(\eta)}$ can always be rescaled by a constant without changing the Lynden-Bell equilibrium (1.12)—in essence, η_{ref} is a gauge choice for the function $\mu(\eta)$. By analogy to textbook statistical mechanics, the function $F(\eta)$ will be referred to as the 'fugacity' of the distribution. The form (2.3) results in the following expression for the Lynden-Bell equilibrium (1.12)[3]:

$$P(\boldsymbol{p}, \eta) = \frac{\delta(\eta) + e^{-\beta\eta\varepsilon(\boldsymbol{p})}F(\eta)}{1 + \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta'\, e^{-\beta\eta'\varepsilon(\boldsymbol{p})}F(\eta')}. \tag{2.4}$$

[3]It should be noted that (2.4) is essentially just a relabelling of (1.12), which is more useful due to the importance given to $\eta = 0$. However, for compact notation it will occasionally be useful, especially in chapter 4, to use (1.12)

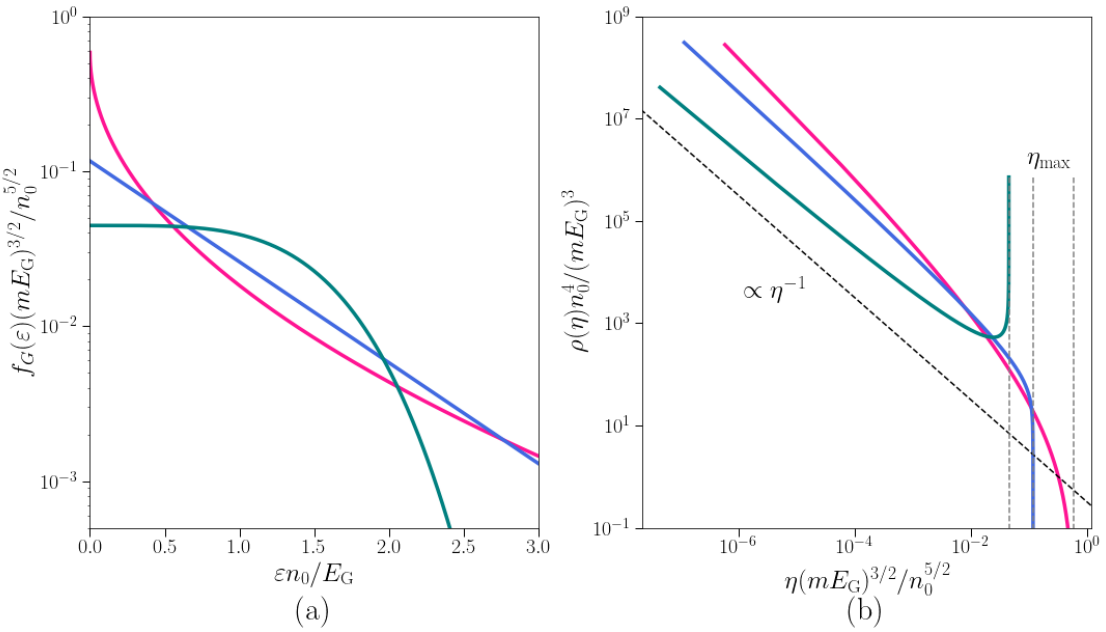

Figure 2.2: (a) Three example Gardner distribution functions (the phase-space density here is plotted as a function of energy) and (b) their corresponding waterbag contents. All three distributions were chosen to have the same particle density n_0 and energy density E_G. The maximum phase-space density $\eta_{\max}$ of the distribution sets the upper cutoff of the waterbag content in η, shown by the dashed vertical lines in (b). The lower cutoff $\eta_{\min}$ is justified in Section 2.2.5. We see that large differences at low ε only change the behaviour of $\rho(\eta)$ significantly at $\eta \sim \eta_{\max}$. For $\eta \ll \eta_{\max}$, all three waterbag contents asymptote to a universal η^{-1} scaling.

In (2.4), the first, $\delta(\eta)$, term in the numerator accounts for the probability density of finding phase space to be empty at a given location (this part of the phase space is referred to, aptly, as the 'vacuum' by Chavanis 2006a), whereas the second term accounts for non-empty waterbags. Already Γ_{free} has dropped out of the calculation, as it must, and it is safe to let $\Gamma_{\text{free}} \to \infty$. Likewise, it is immediately obvious that the reference phase-space density η_{ref} has cancelled, as it also must. Thus, the Lynden-Bell equilibrium distribution (2.4) depends only on the Lagrange multiplier β and the fugacity $F(\eta)$ (which themselves depend on $\rho(\eta)$ for $\eta > 0$ and the energy density E).

The key feature of the example (2.2) is the η^{-1} power-law behaviour. While the specific form of the logarithmic factor in (2.2) was set by our (non-universal) choice of a Maxwellian Gardner distribution, the η^{-1} behaviour at small η is relatively universal. For any exponentially decaying Gardner distribution, i.e., for

any distribution that at large energies can be approximated by $\propto \exp\left[-\left(\varepsilon/\varepsilon_0\right)^\sigma\right]$, for some $\sigma > 0$ (or indeed can be bounded between two such functions), one will find a waterbag content with an η^{-1} asymptotic at $\eta \ll \eta_{\max}$.

To see this, we note that, since $f_{\rm G}(\varepsilon)$ is monotonically decreasing with energy, it has a well-defined inverse $f_{\rm G}^{-1}(\eta)$. In terms of this inverse, one can explicitly express the waterbag content (1.10) as

$$\rho(\eta) = \int \mathrm{d}\varepsilon\, g(\varepsilon)\delta\left(\eta - f_{\rm G}(\varepsilon)\right) = -g\Big(f_{\rm G}^{-1}(\eta)\Big)\frac{\mathrm{d}f_{\rm G}^{-1}}{\mathrm{d}\eta} \implies \frac{\mathrm{d}f_{\rm G}}{\mathrm{d}\varepsilon} = -\frac{g(\varepsilon)}{\rho\Big(f_{\rm G}(\varepsilon)\Big)}, \quad (2.5)$$

where $g(\varepsilon)$ is the density of states in energy, defined by the equation

$$\int \mathrm{d}\boldsymbol{p}\,(...) = \int \mathrm{d}\varepsilon\, g(\varepsilon)(...). \quad (2.6)$$

The first equality in (2.5) is a straightforward way to calculate the waterbag content of a given Gardner distribution, while the second is an equation from which the Gardner distribution can be constructed given knowledge of the system's waterbag content (cf. Dodin &Fisch 2005; Helander 2017). It is immediately clear why the η^{-1} scaling should arise in $\rho(\eta)$. For any exponentially decaying $f_{\rm G}(\varepsilon)$, the inverse function will be some logarithmic function of η, which, after differentiation in (2.5), will give an η^{-1} asymptotic multiplied by some logarithmic function of η. To illustrate this, in Figure 2.2, we give three examples of starkly different Gardner distributions that, despite their differences, all possess waterbag contents which scale as η^{-1} at low η.

To convince a doubtful reader, we consider an alternative argument in support of the η^{-1} scaling of $\rho(\eta)$. First, let us imagine a system in which we are allowed to vary $\eta_{\min}$ freely while leaving the exact phase-space density $f(\boldsymbol{v})$ otherwise unchanged, as if there were some true distribution that had $\eta_{\min} = 0$ and we were examining successive approximations to it (e.g., the difference between a truncated Maxwellian (2.1) and a true Maxwellian). Let us consider the following integrals of such distribution functions (similar to the Casimir invariants considered by Zhdankin 2022a):

$$\frac{1}{V}\iint_{f(\boldsymbol{x},\boldsymbol{p})>\eta_{\min}} \mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{p}\,\left[f(\boldsymbol{x},\boldsymbol{p})\right]^\gamma = \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\,\eta^\gamma\rho(\eta), \quad (2.7)$$

where we have used (1.10), and the phase-space integral is taken over the volumes where $\eta > \eta_{\min}$. We have also used the property that the process of varying $\eta_{\min}$

while otherwise leaving $f(\boldsymbol{x},\boldsymbol{p})$ unchanged only changes the integration limits of $\rho(\eta)$, without changing $\rho(\eta)$ itself. The idea now is to vary the values of $\eta_{\min}$ and γ and use what we know about $f(\boldsymbol{x},\boldsymbol{p})$ to deduce the form of $\rho(\eta)$. Clearly, for $\gamma = 1$, (2.7) is the particle density of the truncated $f(\boldsymbol{x},\boldsymbol{p})$, which must be finite. This tells us that $\rho(\eta)$ must integrate to a finite value when multiplied by η. Furthermore, should f be any exponentially decaying function, then there would be a characteristic momentum scale above which the distribution is suppressed. This means that, as $\eta_{\min}$ is taken to zero, both sides of (2.7) must converge for $\gamma = 1$. This is effectively a statement that the amount of probability contained beyond a few standard deviations is small. However, for an exponentially decaying phase-space density f and any positive power $\gamma > 0$, f^γ is also exponentially decaying, so the same argument applies. Therefore, for exponentially decaying phase-space densities, $\rho(\eta)$ must be such that, when multiplied by any positive power of η, it integrates to some finite value and is largely independent of the choice of lower cutoff $\eta_{\min}$ of the integral. However, for $\gamma = 0$, (2.7) becomes the (spatially averaged) momentum-space volume occupied by the truncated distribution: its support. For exponentially decaying distributions, which do not have compact support without truncation, this quantity will continue to grow without bound as $\eta_{\min}$ is decreased. Therefore, $\rho(\eta)$ integrated with no powers of η must diverge as $\eta_{\min} \to 0$. It is obvious that η^{-1} is a function that has all these properties, but, more generally, $\rho(\eta)$ could be any function of the form

$$\rho(\eta) = \frac{n_0}{\eta_{\max}\eta} G\left(\frac{\eta}{\eta_{\max}}\right) \quad \text{for} \quad \eta_{\min} < \eta < \eta_{\max}, \tag{2.8}$$

where $G(x)$ is a dimensionless function whose dependence on x is weaker than any power law, viz.,

$$\lim_{x\to 0} x^{\gamma-1} G(x) = \begin{cases} 0 & \text{if} \quad \gamma > 1, \\ \infty & \text{if} \quad \gamma < 1. \end{cases} \tag{2.9}$$

Note that the exact limit at $\gamma = 1$ cannot be determined by this argument; it is, in fact, dependent on the exact details of the exponential decay, but we will not require it for any calculations. This concludes the argument that Gardner distributions with exponential tails have waterbag contents with a universal low-η asymptotic.

It is a straightforward extension of this argument to work out what the waterbag content will be for Gardner distributions with non-exponential tails at high energies (low η). Suppose that, instead of being exponentially decaying, the Gardner distribution behaves as a power law at large momenta. Then there is a choice of $\gamma > 0$ for which f^γ, multiplied by the density of states in $|\boldsymbol{p}|$, decays slower than $|\boldsymbol{p}|^{-1}$, implying that the integral (2.7) will diverge as $\eta_{\min}$ is decreased. This means that $\rho(\eta)$ multiplied by η^γ for some $\gamma > 0$ must still have a divergent integral, so it must take the form

$$\rho(\eta) = \frac{n_0}{\eta_{\max}^{2-\delta}\eta^\delta} G\left(\frac{\eta}{\eta_{\max}}\right) \quad \text{for} \quad \eta_{\min} < \eta < \eta_{\max}, \tag{2.10}$$

with some $\delta > 1$. It turns out that (2.10) also holds, but with $\delta < 1$, for phase-space densities that go to zero algebraically even when $\eta_{\min} = 0$[4]. This is because the integral (2.7) will be over a finite momentum-space volume even as $\eta_{\min}$ is taken to zero, so, while $\gamma < 0$ will make f^γ diverge is this finite domain, that divergence will still be integrable if γ is chosen sufficiently small and negative.

The explicit value of δ can be found from (2.5). To enable an explicit calculation, we will assume henceforth that the density of states $g(\varepsilon)$ is related to energy by a simple power law, viz.,

$$g(\varepsilon) = A\varepsilon^a, \tag{2.11}$$

where A is an appropriate constant with dimensions $[n_0/\eta_{\max}\varepsilon^{1+a}]$, and a is a real number. The assumption (2.11) is not too restrictive as it can capture both non-relativistic and ultra-relativistic systems of any dimensionality. With the assumption (2.11), we may now use (2.5) to link the value of δ to the high-energy asymptotic of the Gardner distribution. If the Gardner distribution has the power-law tail $f_{\mathrm{G}}(\boldsymbol{p}) \propto \varepsilon(\boldsymbol{p})^{-(\chi+a)}$ at high energies, then, via (2.5), one finds $\delta = 1 + (a+1)/(a+\chi)$. If, instead, the Gardner distribution goes to zero at some finite energy $\varepsilon_{\max}$ so that $f_{\mathrm{G}}(\boldsymbol{p}) \propto [\varepsilon_{\max} - \varepsilon(\boldsymbol{p})]^\chi$ near $\varepsilon = \varepsilon_{\max}$, then $\delta = 1 - 1/\chi$.

Having catalogued the possible Gardner distributions and their corresponding $\rho(\eta)$, an important open question is now what Gardner distributions are the

[4]To apply this argument more generally to functions that have compact support in momentum space but do not go to zero algebraically (e.g., step or bump functions), one can represent them as the limit of a sequence of functions that do go to zero algebraically.

most common. Within the domain of numerical experiments, this would appear to be decided by the whims of the numerical experimenter. One often considers systems that are initially close to a Maxwellian equilibrium but which develop kinetic instabilities as a result of large-scale inhomogeneities (e.g., mixing of counter-propagating flows, large-scale fluid motions, collapsing distributions of matter, etc.). By virtue of the initial proximity to a Maxwellian equilibrium, such systems naturally have an exponential Gardner distribution. However, it would seem reasonable to conjecture that in nature, as the result of weakly collisional effects, exponential Gardner distribution may arise naturally in all systems. This intuition, by no means a theorem, comes from the fact that the only processes that can change the Gardner distribution are, by definition, collisional, and collisional processes naturally smooth the system to a Maxwell–Boltzmann distribution. In chapter 3, we will show numerically that indeed weak collisional effects do cause the system to develop an exponential Gardner distribution but, as yet, there is no well-developed theory for the evolution of $\rho(\eta)$.

We are now safe in the knowledge that the exponentially decaying Gardner distributions have $\rho(\eta) \propto \eta^{-1}$, or, in more exotic cases, $\rho(\eta) \propto \eta^{-\delta}$. We may now return to the task of solving (1.8) and (1.9) in light of these facts.

2.2.3 Non-degenerate Lynden-Bell equilibria

Taking inspiration from the similarity between Fermi–Dirac and Lynden-Bell statistics, we may expect that there are two analytically tractable limits: degenerate ('cold') and non-degenerate ('hot'). In this section, we will explore the latter limit, which will turn out to be far more useful than the former (which is, nevertheless, also treated, for completeness, in appendix B).

We define the non-degenerate limit as one in which the probability of finding the exact phase-space density to be non-zero is small, viz.,

$$D(\varepsilon) = \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta \, e^{-\beta\eta\varepsilon} F(\eta) \ll 1. \tag{2.12}$$

We shall call $D(\varepsilon)$ the 'degeneracy parameter' since, from (2.4), the probability that a position in phase space with a given energy is non-empty is given by the

quotient $D(\varepsilon)/\left[1+D(\varepsilon)\right]$. In the limit (2.12), the distribution function (2.4) can be approximated by

$$P(\boldsymbol{p},\eta) \simeq \delta(\eta)\left[1-D\Big(\varepsilon(\boldsymbol{p})\Big)\right] + e^{-\beta\eta\varepsilon(\boldsymbol{p})}F(\eta). \tag{2.13}$$

The effect of this simplification is that the competition for any particular sub-volume of phase space is so weak that the waterbags are free to arrange themselves as Maxwellians η by η. The waterbags with lower η cost less energy to be placed at larger momenta—therefore, they have a larger thermal spread. In the non-relativistic limit, this is equivalent to the intuition that particles belonging to the waterbags with higher phase-space densities behave as though they have larger masses.

The approximate form (2.13) of the Lynden-Bell distribution makes computing the momentum-space integral in (1.9) and determining the fugacity $F(\eta)$ a simple matter. Substituting (2.13) into (1.9) and using the explicit form (2.11) of the density of states, we find the fugacity in the non-degenerate limit to be

$$F(\eta) = \frac{(\beta\eta)^{1+a}}{A\Gamma(a+1)}\rho(\eta), \quad \eta_{\min} \le \eta \le \eta_{\max}, \tag{2.14}$$

where $\Gamma(a+1)$ is the gamma function. Substituting (2.14) back into (2.13) and using (1.5) finally gives us an expression for the mean phase-space density (although still in terms of the as yet unspecified parameter β):

$$\langle f\rangle(\boldsymbol{p}) = \frac{\beta^{1+a}}{A\Gamma(a+1)}\int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, \eta^{2+a}\rho(\eta)e^{-\beta\eta\varepsilon(\boldsymbol{p})}. \tag{2.15}$$

From (2.13), it might seem as though, by diverse choices of $\rho(\eta)$, a wide variety of distribution functions $\langle f\rangle$ can be obtained. However, as we showed in section 2.2.2, a diversity of choices of waterbag contents is exactly what we do not have. Instead, fairly generic Gardner distributions with any form of exponential tails possess waterbag contents that are highly universal at low η. Using (2.8), we can make a convenient change of variables in (2.13), $x = \beta\eta\varepsilon(\boldsymbol{p})$, to find

$$\langle f\rangle(\boldsymbol{p}) = \frac{n_0}{A\beta\Gamma(a+1)\eta_{\max}\varepsilon(\boldsymbol{p})^{2+a}}\int_{\beta\eta_{\min}\varepsilon(\boldsymbol{p})}^{\beta\eta_{\max}\varepsilon(\boldsymbol{p})} \mathrm{d}x\, x^{1+a}G\left(\frac{x}{\beta\eta_{\max}\varepsilon(\boldsymbol{p})}\right)e^{-x}. \tag{2.16}$$

Despite appearing complex, most of this expression is superfluous and the only key ingredient is the power-law dependence on energy outside the front of the integral. To see this, we note that there is a natural range of energies

$$\frac{1}{\beta\eta_{\max}} \ll \varepsilon(\boldsymbol{p}) \ll \frac{1}{\beta\eta_{\min}}, \tag{2.17}$$

in which the upper and lower limits of the integral in (2.16) will be large and small respectively. This corresponds to the range of energies that are well within the thermal spread of the least dense waterbags, but far outside the thermal spread of the densest ones. Because of the exponential term e^{-x} suppressing large x in the integrand of (2.16), and the polynomial x^{1+a} suppressing low values of x, the integral will be largely insensitive to the value of the integration limits in this energy range. Thus, the only energy dependence of the integral could come from the argument of G. But we have argued that G is a weak function (essentially meaning logarithmic, cf. (2.9)) of its argument, so its energy dependence in the integral will also be weak. This leaves only the power-law energy dependence outside the integral which is universal[5]. The distribution function $N(\varepsilon)$ of particle energies corresponding to (2.16) is obtained by multiplying the mean phase-space density $\langle f\rangle(\boldsymbol{p})$ by the density of states:

$$N(\varepsilon) = g(\varepsilon)\langle f\rangle(\boldsymbol{p}) = \frac{n_0}{\beta\Gamma(a+1)\eta_{\max}\varepsilon^2}\int_{\beta\eta_{\min}\varepsilon}^{\beta\eta_{\max}\varepsilon} \mathrm{d}x\, x^{1+a} G\left(\frac{x}{\beta\eta_{\max}\varepsilon}\right) e^{-x}. \quad (2.18)$$

This shows that the non-degenerate Lynden-Bell equilibria express a natural power-law tail and, furthermore, that this power law is independent of the type of plasma system under consideration. Note that the origin of the ε^{-2} scaling found here is entirely different than the ε^{-2} arising from particle acceleration in shocks (Bell, 1978).

Consider now what happens if the Gardner distribution does not have an exponential tail. Then $\rho(\eta)$ can be written as (2.10). Following all the same steps as before from (2.13) onwards, but using (2.10) in place of (2.8), one arrives at

$$N(\varepsilon) = \frac{n_0}{\beta^{2-\delta}\Gamma(a+1)\eta_{\max}^{2-\delta}\varepsilon^{3-\delta}}\int_{\beta\eta_{\min}\varepsilon}^{\beta\eta_{\max}\varepsilon} \mathrm{d}x\, x^{2+a-\delta} G\left(\frac{x}{\beta\eta_{\max}\varepsilon}\right) e^{-x}. \quad (2.19)$$

Thus, the resulting Lynden-Bell equilibrium again displays a power-law tail. The power law's exponent is set by the particular value of δ in (2.10), which is related to the Gardner distribution of that Lynden-Bell equilibrium via (2.5), as explained at the end of Section 2.2.2. For Gardner distributions that already have power-law tails, the resulting Lynden-Bell equilibria have shallower (ultra-'hard') power-law tails that

[5]For energies much lower than $(\beta\eta_{\max})^{-1}$ the upper integration limit begins to suppress the integral and it is easy to see from (2.15) that the distribution function will be finite at zero energy. For energies much larger than $(\beta\eta_{\min})^{-1}$ the lower integration limit becomes so large that the exponential e^{-x} begins to suppress the whole integral, leading to an exponential fall off.

strongly diverge in energy, giving the cutoff $\eta_{\min}$ pivotal importance. For Gardner distributions that decay faster than any exponential, the Lynden-Bell equilibria have 'soft' power-law tails, with total energy depending only very weakly on the cutoff[6].

Presently, however, we shall return to the Lynden-Bell equilibria (2.18) arising from exponential Gardner distributions, for which we complete the calculation.

2.2.4 Calculation of β and inevitability of partial degeneracy

Both (2.17) and (2.18) still depend on the as yet unknown parameter β, which, as well as fixing the energy, will determine the accuracy of the non-degeneracy approximation (2.12). To find β, we must compute the energy of our Lynden-Bell equilibrium (2.13) according to (1.8). Equivalently, from (2.13),

$$E = \int \mathrm{d}\boldsymbol{p}\, \varepsilon(\boldsymbol{p}) \langle f \rangle(\boldsymbol{p}) = \frac{a+1}{\beta} \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, \rho(\eta). \tag{2.20}$$

Therefore, in the non-degenerate limit,

$$\beta = \frac{a+1}{E} \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, \rho(\eta). \tag{2.21}$$

We see that β decreases with increasing total energy of the distribution function; this is natural if β is viewed as an inverse thermodynamic temperature.

To see how this affects the underlying assumption (2.12) of non-degeneracy, we evaluate $D(\varepsilon)$ at $\varepsilon \to 0$, where the degeneracy effect is strongest, since $D(\varepsilon) \leq D(0)$. Requiring $D(0) \ll 1$ gives the following condition on the total energy of the distribution, via (2.12), (2.14) and (2.21):

$$E \gg (a+1) \left[\int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, \frac{\eta^{a+1} \rho(\eta)}{A \Gamma(a+1)} \right]^{1/(a+1)} \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, \rho(\eta). \tag{2.22}$$

Since $\rho(\eta)$ is a function only of the Gardner distribution, the right-hand side of (2.22) must scale with the energy density E_{G} of the Gardner distribution that has the same waterbag content $\rho(\eta)$, but will always have $E_{\mathrm{G}} \leq E$. For example, for the Gardner distribution (2.1), the right-hand side of (2.22) should be proportional to $n_0 \varepsilon_0$. Thus, just like in Fermi–Dirac statistics, the non-degeneracy approximation becomes more accurate as the distribution's energy density E begins to dwarf the energy density E_{G}

[6]The fact that certain choices of fugacity $F(\eta)$ give rise to Lynden-Bell equilibria that have power laws with various exponents was first noted, in connection to superstatistics, by Chavanis (2006a).

of the ground state. Note, however, that (2.22) contains an integral of the waterbag content $\rho(\eta)$ with no weighting by η, which, by (1.7) and (1.9), is the total phase volume occupied by non-empty waterbags. Since $\rho(\eta) \propto \eta^{-1}$ at small η, this will be large, depending, albeit logarithmically, on the minimum waterbag density $\eta_{\min}$. Indeed, for our example (2.1), (2.22) becomes

$$E \gg \frac{4}{3\sqrt{\pi}} n_0 \varepsilon_0 \left(\ln \frac{\eta_{\max}}{\eta_{\min}} \right)^{3/2} \left[1 + \mathcal{O}\left(\frac{\eta_{\min}}{\eta_{\max}} \right) \right]. \tag{2.23}$$

This means that the condition (2.22) requires the energy of the distribution to be much greater not just than the energy of the corresponding Gardner distribution, but than the Gardner energy multiplied by a polylogarithmic function of $\eta_{\min}$. This is a manifestation of the fact that, as $\eta_{\min}$ is taken to zero, more of the phase space is pervaded by low-density waterbags, making true non-degeneracy harder to achieve. It is thus impossible to achieve a non-degenerate limit unless $\eta_{\min}$ is kept finite.

This is not the only place where the finiteness of the cutoff $\eta_{\min}$ has raged against the dying of the light. The same effect is manifest in the power law of ε^{-2} appearing in (2.18), which would have led to a logarithmically divergent mean particle energy were it not for the exponential cutoff at $\varepsilon \sim 1/\beta\eta_{\min}$. This is obvious in (2.20), where the integral of $\rho(\eta)$ has the same logarithmic divergence with $\eta_{\min} \to 0$ as it did in the right-hand side of (2.22). This then makes its way into the expression (2.21) for β, so, formally in the limit $\eta_{\min} \to 0$, $\beta \to \infty$ always!

Since we must keep $\eta_{\min}$ finite, it is of substantial importance to understand the physical significance of it and the extent to which one should be prepared to accept one's equilibrium's dependence on its value.

2.2.5 Physical meaning of minimum waterbag density

A Lynden-Bell equilibrium is essentially the thermal equilibrium of a collection of correlated blobs in phase space. This is to say that, inherent to the idea of computing the mean phase-space density, we have assumed that an exact phase-space density, i.e., a finite value of η, is a meaningful concept. But, of course, an exact phase-space density is a fiction, since a plasma is composed of many discrete particles, and a phase-space density is only an average occupation number of particles' positions in phase space. The only sense in which an exact, continuous, phase-space density can

be meaningful in a collisionless plasma then is if, within a small enough phase-space volume $\Delta\Gamma$, many particles can be considered to move as a collective entity: a waterbag[7]. Then, on the scale of $\Delta\Gamma$, the system is composed of many waterbags with some 'exact' phase-space density, whereas on scales much larger than $\Delta\Gamma$, the system can attain a mean phase-space density. This 'correlation volume' provides a natural way to introduce the minimum non-zero phase-space density: clearly that should correspond to a single particle sitting in $\Delta\Gamma$, giving $\eta_{\min} = \Delta\Gamma^{-1}$.

Determining the value of $\Delta\Gamma$, is, however, a non-trivial challenge (see, e.g., discussions in Kadomtsev &Pogutse 1970; Chavanis 2022). In chapter 4 we will argue, on the grounds that any meaningful collisionless-relaxation rate must be smaller than the plasma frequency but larger than the rate at which collisions break phase-volume conservation, that a reasonable constraint to place on the correlation volume is

$$\Delta\Gamma \sim \frac{1}{\eta_{\text{eff}}} \left(n_0 \lambda_D^3\right)^{\alpha}, \quad \frac{2}{3} < \alpha < 1, \tag{2.24}$$

where η_{eff} is some typical phase-space density, which we can here estimate by $\eta_{\max}$, and λ_D is the Debye length. This gives us an estimate for the minimum waterbag density:

$$\eta_{\min} = \Delta\Gamma^{-1} \sim \eta_{\max} \left(n_0 \lambda_D^3\right)^{-\alpha}. \tag{2.25}$$

Since the typical number of particles in a Debye sphere can be as large as 10^6 to 10^8 in collisionless plasma environments (such as the solar wind or interstellar gas: see, e.g., Verscharen et al. 2019 and Ferrière 2019), the estimate (2.25) might seem damningly small. It is, in fact, ideal. The existence of a broad power-law tail in (2.18) required a scale separation between $\eta_{\max}$ and $\eta_{\min}$. The estimate (2.25) certainly provides this separation, tied to the plasma parameter. As for the breakdown of the non-degenerate approximation and the marginal divergence of the mean particle energy of a distribution with an ε^{-2} power law, we are saved by the fact that only the logarithm of the ratio $\eta_{\max}/\eta_{\min}$ will appear, which, while large, can only ever be in the range of $10 - 30$. This is somewhat reminiscent of the situations in the conventional theory of Coulomb collisions in plasmas, where forcible introduction of a phase-space cutoff into the collision integral results in

[7]This phase-space volume $\Delta\Gamma$ will reappear in chapter 4 in the form of a correlation volume in phase-space

only a weak dependence on the value of this cutoff, via the so-called Coulomb logarithm (see, e.g., Helander &Sigmar 2005).

Nevertheless, the presence of the logarithmic divergence with $\eta_{\min}$ in the expression for β, signposted at the end of section 2.2.4, will make it difficult to satisfy the non-degeneracy approximation[8]. There is good reason to suppose, however, that its breakdown may only be partial. We note that evaluating (2.12) at $\varepsilon = 0$ is tantamount to requesting non-degeneracy everywhere in the distribution. Since the degeneracy parameter $D(\varepsilon)$ decreases with increasing energy, it is reasonable to expect (and indeed we will see) some degeneracy at low energies, which gives way to non-degeneracy at higher energies, where our power-law tails will be recovered. As ever, the true solution lies on the cusp of asymptotic theory, and to go any further we must resort to numerical computation.

2.3 Numerical results: partially degenerate equilibria

In this section, we shall recover the analytically predicted power-law tail in (2.18) by solving the constraint equations (1.8) and (1.9) for the Lynden-Bell equilibria (2.4). The numerical scheme for this is documented in detail in appendix C, amounting to an iterative method coupled to a 1D root finder. For these numerical results, we have restricted ourselves to a 3D, non-relativistic plasma with $\varepsilon(\boldsymbol{p}) = p^2/2m$, although we anticipate from Section 2.2 that these results extend, qualitatively, to general regimes.

To capture a broad range of initial conditions, we consider a family of waterbag contents defined by

$$\rho(\eta) = \frac{4\pi p_{\mathrm{th},\sigma}^3}{\sigma\eta}\left(\ln\frac{\eta_{\max,\sigma}}{\eta}\right)^{(3-\sigma)/\sigma}, \quad \eta_{\min} < \eta < \eta_{\max,\sigma}, \tag{2.26}$$

with $\sigma > 0$. This defines the family of Gardner distributions

$$f_{\mathrm{G},\sigma}(\boldsymbol{p}) = \begin{cases} \eta_{\max,\sigma} e^{-\left(p/p_{\mathrm{th},\sigma}\right)^\sigma} & \text{for} \quad p < p_{\mathrm{th},\sigma}\left(\ln\dfrac{\eta_{\max,\sigma}}{\eta_{\min}}\right)^{1/\sigma}, \\ 0 & \text{for} \quad p > p_{\mathrm{th},\sigma}\left(\ln\dfrac{\eta_{\max,\sigma}}{\eta_{\min}}\right)^{1/\sigma}, \end{cases} \tag{2.27}$$

[8]Although, notably, the fully non-degenerate limit can be naturally recovered by numerical noise in particle-in-cell (PIC) codes: see appendix D.

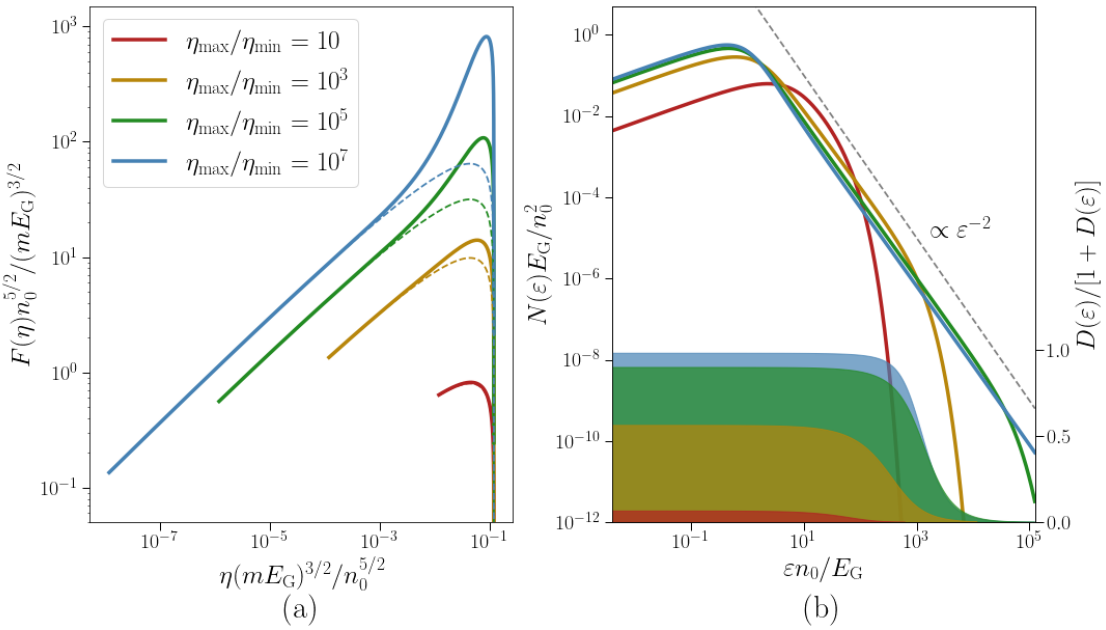

Figure 2.3: Numerically computed Lynden-Bell equilibria for a range of $\eta_{\max,\sigma}/\eta_{\min}$ and with $\rho(\eta)$ given by (2.26) with $\sigma = 2$. The energy density is equal to $10E_{\mathrm{G}}$ in all cases. (a) The numerically computed fugacity $F(\eta)$ (solid lines) compared with the analytical solution (2.14) obtained in the non-degenerate limit (dashed lines). (b) The resulting distributions $N(\varepsilon)$ of particle energies, with the universal power law $\propto \varepsilon^{-2}$ shown for reference, cf. (2.18). Overplotted in solid colour (with the value range shown on the right) is the level of degeneracy $D(\varepsilon)/[1+D(\varepsilon)]$ (the probability that a given energy is occupied by a non-empty waterbag) as a function of energy; $D(\varepsilon)$ is defined in (2.12).

which, in the limit of $\eta_{\min}/\eta_{\max,\sigma} \to 0$, have particle densities n_0 and energy densities E_{G} that satisfy

$$n_0 = \frac{4\pi}{\sigma}\Gamma\left(\frac{3}{\sigma}\right)p_{\mathrm{th},\sigma}^3\eta_{\max,\sigma}, \quad E_{\mathrm{G}} = \frac{\Gamma(5/\sigma)}{\Gamma(3/\sigma)}n_0\frac{p_{\mathrm{th},\sigma}^2}{2m}. \tag{2.28}$$

Since these Gardner distributions represent minimum-energy states, we will be able to scan in the energy density of the system for all $E \geq E_{\mathrm{G}}$, imagining some initial distribution of particles, with waterbag content $\rho(\eta)$ and energy density E_{G}, being accelerated to the energy density E, and then seeking a maximum-entropy state (see Section 2.2.1). We can also vary σ in order to see the effects of the shape of the underlying Gardner distribution on the resulting Lynden-Bell equilibria.

2.3.1 Degrees of degeneracy

Let us first scan in $\eta_{\max,\sigma}/\eta_{\min}$ in order to show that we can indeed recover the fully non-degenerate limit solved in Section 2.2. Figure 3.3 shows the result of such a scan for $\sigma = 2$ and $E = 10E_{\rm G}$. By comparing the exact (numerically calculated) fugacity to the theoretical prediction (2.14) obtained in the absence of phase-space degeneracy, we see that the agreement is nearly perfect when $\eta_{\max}$ is close to $\eta_{\min}$, e.g., when $\eta_{\max}/\eta_{\min} = 10$. This is as expected, since the non-degenerate limit is valid when (2.22) holds, which it does, as can be confirmed from the solid red colour in panel (b), showing the probability $D(\varepsilon)/\left[1 + D(\varepsilon)\right]$ (with $D(\varepsilon)$ defined in (2.12)) that a region of phase space is occupied. However, at $\eta_{\max}/\eta_{\min} = 10$, there is an insufficient range of waterbag levels to achieve the scale separation (2.17) necessary to resolve a power-law tail in energies. To see a power-law tail, one must increase $\eta_{\max}/\eta_{\min}$ to higher values, but this comes at the price of increasing the degeneracy of the system and, hence, undermining the asymptotic regime in which the tail was derived in the first place. For the values of $\eta_{\max}/\eta_{\min}$ that we argued in Section 2.2.5 to be realistic, the system becomes strongly degenerate at low energies, which can again be seen from the solid colours in panel (b). In spite of this, at high energies, the degeneracy falls away, and, correspondingly, $F(\eta)$ at low η still agrees well with the non-degenerate approximation (2.14). All this conspires to ensure that, even formally outside the non-degenerate limit, the power-law tail $N(\varepsilon) \propto \varepsilon^{-2}$ is still manifestly present.

2.3.2 Energisation of particles and power-law tails

Let us now scan in the energy densities E of the distribution and again look for power-law tails and assess the effect of degeneracy. Figure 2.4 shows the results of such a scan, again with $\rho(\eta)$ specified by (2.26) with $\sigma = 2$. Plotted underneath the mean phase-space densities are the contributions from a number of finite ranges of exact phase-space density defined by

$$\langle f \rangle_{\eta_1<\eta<\eta_2}(\boldsymbol{p}) = \int_{\eta_1}^{\eta_2} \mathrm{d}\eta\, \eta P(\boldsymbol{p}, \eta). \tag{2.29}$$

As one would anticipate, when E is only slightly larger than $E_{\rm G}$, the effects of phase-space degeneracy are most prominent: the densest portions of the phase

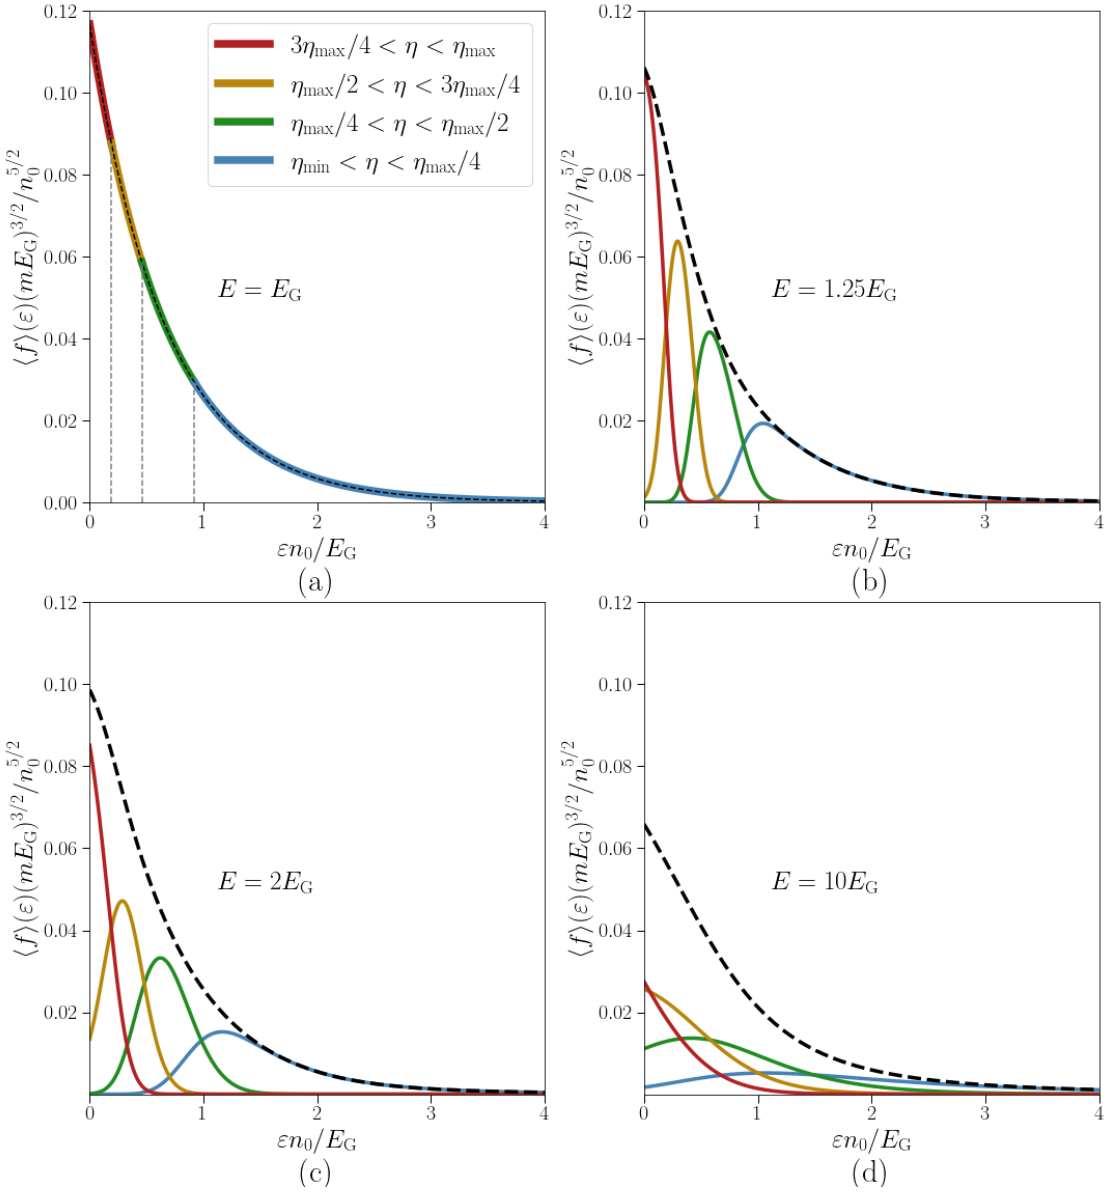

Figure 2.4: Numerically computed Lynden-Bell equilibria for a range of energy densities E (in multiples of the energy density E_{G} of the underlying Gardner distribution) with $\rho(\eta)$ given by (2.26) with $\sigma = 2$ and $\eta_{\max}/\eta_{\min} = 10^6$. In each plot, the dashed line is the mean phase-space density, while the underplotted solid lines are the contributions from four distinct ranges of exact phase-space density as functions of energy. Note that, while the exact phase-space densities have been grouped into four, this is not the same as solving for a four-waterbag Lynden-Bell equilibrium, as each grouping is still composed of a continuum of waterbags.

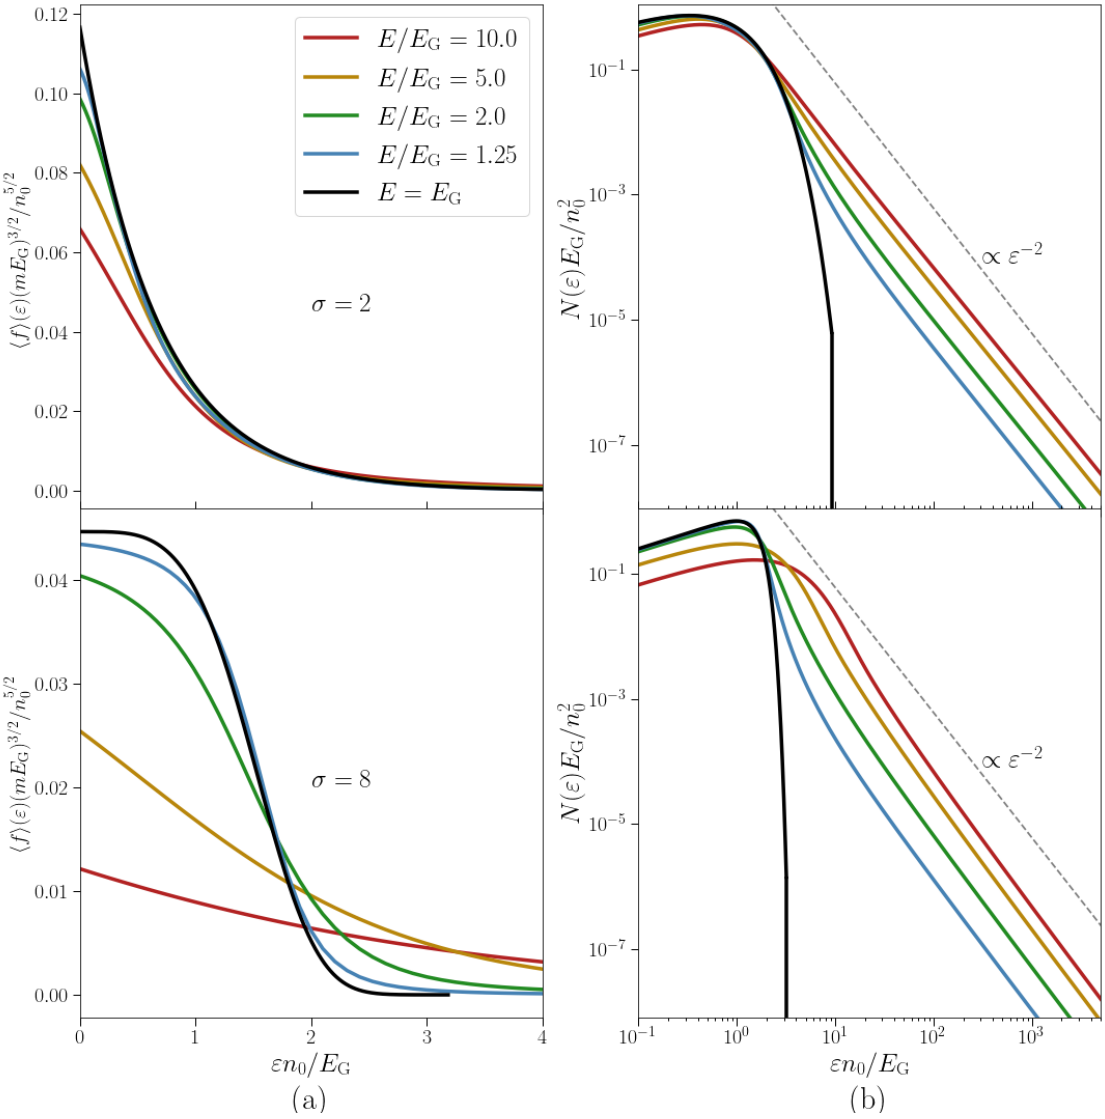

Figure 2.5: Numerically computed Lynden-Bell equilibria with waterbag content given by the $\sigma = 2$ (top) and $\sigma = 8$ (bottom) cases of (2.26), $\eta_{\max,\sigma}/\eta_{\min} = 10^6$. (a) The phase-space densities shown in linear scale, (b) the corresponding distributions of particle energies in logarithmic scale, for a range of ratios of E/E_G. The small deviations from the ε^{-2} tail can be attributed to the logarithmic corrections arising from the x integral in (2.18).

space clog up the lowest energies, forcing less dense portions to higher energies. As the total energy density is increased, we see that the contribution from each range of waterbags spreads out. This is because the increased energy allows dense portions of phase space to be promoted to larger energies, making room at lower energies for less dense portions of phase space to fill.

While the solutions plotted in Figure 2.4 might appear qualitatively similar to the Lynden-Bell equilibria obtained in numerical experiments with a small discrete number of level sets (see, e.g., the three waterbag case in figure 4.3, or Assllani et al. 2012 for examples), this hides key universal features of systems with a continuum of level sets. To showcase this universality, in Figure 2.5, we plot the Lynden-Bell equilibria for two different waterbag contents, $\sigma = 2$ and $\sigma = 8$ in (2.26), and a range of energy densities E. The ε^{-2} power-law tails of these equilibria are immediately apparent, as predicted in (2.18).

Figure 2.5 shows how these power-law tails become more prominent as one adds more energy to the Gardner distribution. At $E = E_{\mathrm{G}}$, one has a highly non-universal Gardner equilibrium. As a small amount of energy $E - E_{\mathrm{G}} \ll E_{\mathrm{G}}$ is added to E_{G}, the mean phase-space density at low energies is largely unaffected, while the power-law tail grows from the lowest-density waterbags. Thus, for energies close to the energy of the underlying Gardner distribution, the Lynden-Bell equilibria have a 'core-halo' structure: the 'core', which has energy density $\sim E_{\mathrm{G}}$, is comprised of dense waterbags, which the system does not have sufficient energy to excite, whereas the 'halo' is the tail comprised of those less dense waterbags that are capable of sampling a larger portion of phase space and thus arrange themselves into a universal ε^{-2} power law, containing the excess energy density $\sim E - E_{\mathrm{G}}$. As the energy of the distribution is further increased, more waterbags have sufficient energy to sample a larger range of phase space, the halo continues to eat into the core, but both thermally broaden. At $E \gg E_{\mathrm{G}}$, the asymptotically non-degenerate solution (2.18) with a power law in the energy range (2.17) suggests that the system strives for a state in which the halo has much more energy than the core. In this limit, one expects the transition between the core and halo to occur at $\varepsilon \sim 1/\beta\eta_{\max}$. Since the halo should be exponentially suppressed at $\varepsilon \gtrsim 1/\beta\eta_{\min}$, one can compute the ratio of the core energy to the halo energy. Owing to the ε^{-2} tail, this ratio will be proportional to $\ln(\eta_{\max}/\eta_{\min})$ raised to some power, which depends on the

specific functional form of the Gardner distribution. Whatever this power is, the vast majority of the total energy will be contained in the universal power-law tail for any system whose energy is much larger than that of its Gardner state.

2.4 Summary

The Lynden-Bell (1967) equilibria are the natural maximum-entropy states for systems, such as a plasma described by the collisionless Vlasov equation (1.2), which conserve not only density, momentum, and energy, but also an infinite family of further invariants (1.3). These additional invariants are due to the conservation of phase volume, encoded by the 'waterbag content' function $\rho(\eta)$ given by (1.10) (equivalent to the Casimir invariants (2.7)), which measures the amount of phase volume where the exact phase-space density takes the value η (a 'waterbag'), per unit η. Maximising entropy subject to all these conservation laws then gives the mean phase-space density (1.5) in the form of the Lynden-Bell equilibrium (1.12) coupled with the constraints (1.8) and (1.9). In this chapter, we have solved these constraint equations numerically in the general case, as well as analytically in a tractable limit (which turned out to be the practically relevant one). We have been able to show that, despite their apparent dependence on non-universal initial conditions, Lynden-Bell equilibria generically exhibit power-law tails, and, in particular, that a broad class of initial conditions will give rise to the distribution of particles' energies scaling as ε^{-2} at high ε.

To study the Lynden-Bell equilibria systematically, we first considered what values the invariants of the system, the energy density E and waterbag content $\rho(\eta)$, could take. This led us to the concept of a Gardner distribution function, which is any monotonically decreasing function of the particle's energy. In Section 2.2.1, we argued that to each possible initial condition one could assign a unique Gardner distribution, with the same waterbag content (and, therefore, the same Casimir invariants) as the initial condition, but a different energy; the Gardner distribution having, by definition, the lowest possible energy of all distributions with a given waterbag content $\rho(\eta)$. The Lynden-Bell equilibria at higher energies and the same $\rho(\eta)$ can, therefore, be thought of as excited states of this Gardner distribution.

In Section 2.2.2, we argued that the typical $\rho(\eta)$ would have a fairly generic power-law form at low η, and, in particular, that it would scale as η^{-1} for a wide class of initial conditions (see (2.8)). In Section 2.2.3, we were able to find the Lynden-Bell equilibria analytically provided the energy of the system was sufficiently large for the competition of waterbags for phase space to be ignorable. In this 'non-degenerate' limit, particles belonging to each waterbag arrange themselves into a separate Maxwellian equilibrium with an effective 'temperature' inversely proportional to the phase-space density of that waterbag: denser portions of phase space are energetically costlier to move to higher energies. The resulting mean phase-space density (2.13) was found by integrating the contributions of all waterbags, each with their own Maxwellian distribution weighted by the amount of phase space which that waterbag occupied. Since the amount of each waterbag had a universal form (2.8), this gave rise to a universal power-law tail (2.18) scaling as ε^{-2} at high particle energies.

However, the non-degenerate limit required the system's energy to be asymptotically larger than the energy of the corresponding Gardner distribution. Formally, this turned out to be a very stringent limitation, and indeed one that could hardly ever be strictly fulfilled. Our analytical results were rescued by the argument, confirmed by the numerical solutions presented in Section 2.3, that the effects of phase-space degeneracy were confined to the low-energy part of the distribution. The universal ε^{-2} tail was numerically confirmed to be a robust feature of the generic Lynden-Bell equilibria. As well as ascertaining that a range of different initial conditions (2.26) gave rise to the same power-law tail, the numerical solution also showed how this power-law tail formed. We found (Figure 2.5) that at energies comparable to the Gardner energy, the Lynden-Bell equilibria had a 'core-halo' structure. The halo, consisting of the ε^{-2} tail, was formed from low-density waterbags, which had sufficient energy to explore large portions of phase space, while the non-universal core was made up of denser waterbags, which did not have sufficient energies to be excited. As the energy of the Lynden-Bell equilibrium was increased, the halo ate its way into the core, making the distribution less and less degenerate and more universal in its shape. This behaviour is perhaps reminiscent of the measurements of the 'non-thermal fraction' of particles in the solar wind (see, e.g., Pierrard &Lazar 2010; Oka et al. 2015). This theory must still, however, contest with the question of whether it can be realised in a turbulent

collisionless plasma with a very small, but no longer vanishing, collisionality. It is this question that we will address in the next chapter.

Gibbs and Shannon strain an eye,
orbs around them, blinding fly,
No mortal mind could ever try,
my feats of fast summation.

Yet the humans, still they yearn,
a different kind of truth they learn,
one that I can only spurn,
born of crude approximation.

— From the author, on Maxwell's demon

3

Relaxation to universal equilibria in turbulent plasma

This chapter is adapted from Ewart et al. 2024, arxiv:2409.01742

3.1 Introduction

When Lynden-Bell's theory of collisionless relaxation first came to light, the capacity to test collisionless physics was still in its infancy—the highest resolution simulations of the time, while ground breaking (e.g., Roberts &Berk 1967; Morse &Nielson 1969), could scarcely be called collisionless by modern standards (cf. Sironi &Spitkovsky (2014)) and even Landau damping had only been experimentally confirmed three years earlier (Malmberg &Wharton, 1964). Presently, they can be tested, but the difficulty of solving the equilibria numerical or resolving them within simulation has limited studies to those with only a few level sets and has been met with mixed success (see, e.g., Arad &Johansson 2005; Arad &Lynden-Bell 2005; Levin et al. 2008; Assllani et al. 2012; Levin et al. 2014). As a result, the universal equilibria proposed in chapter 2, and the ever advancing power of computers, makes this an ample question to pose once more: 'can Lynden-Bell equilibria be realised?'. A conceptually grave concern in chapter 2, however, is the reliance on the *precise* conservation of

the Casimir invariants. In any real system, their conservation was only ever going to be approximate, with molecular chaos eventually reasserting itself. Aggravating this concern further, it has been demonstrated both theoretically (e.g., Schekochihin et al. 2008; Eyink 2018; Nastac et al. 2024b) and numerically (e.g., Tatsuno et al. 2009; Zhdankin 2022a) that even for systems with nominally weak interparticle collisions—such that the time scale for relaxation to a Maxwellian equilibrium is much longer than dynamical times—the Casimir invariants are broken on time scales competitive with the evolution of the system (provided there is a sufficient level of turbulence to stir it), rendering their status as invariants questionable at best.

In this chapter, we show that turbulent plasmas do achieve Lynden-Bell equilibria, and that, rather than doing this in spite of the breaking of Casimir invariants, they manage it *in tandem with* this. We further show that the Lynden-Bell equilibrium that is achieved has a universal high-energy asymptotic. We carry out this study the aid of the particle-in-cell (PIC) code OSIRIS (Fonseca et al., 2002), applied to one of the most classically studied turbulent collisionless systems: the two-stream instability. The very earliest simulations (Roberts &Berk, 1967; Morse &Nielson, 1969) showed vividly that this instability leads a system composed of long-lived structures in phase space (cf. 'BGK modes' Bernstein et al. 1957), known as electron 'holes' (see figure 3.1). These holes, since studied analytically, numerically, and observationally (see, e.g., Hutchinson 2017, 2024 and references therein) move around, merge, and generically represent a relaxing turbulent state. It is precisely this turbulence that drives the system's relaxation towards the Lynden-Bell equilibrium. However, the same turbulence generates small velocity- and position-space structures (in the manner of a turbulent cascade predicted by Nastac et al. 2024a), which cause the Casimir invariants to evolve with time. In the course of this evolution, the underlying Lynden-Bell equilibrium to which the system wants to relax is gradually changed and the system adjusts to reach this evolving target equilibrium. Thus, the system's precise memory of initial conditions is replaced with a 'turbulent amnesia': turbulent fluctuations are perpetually trying to push the system towards its Lynden-Bell equilibrium, but the goalposts are continually moved by the breaking of Casimir invariants driven by those same fluctuations. Eventually the Casimir invariants themselves reach a steady state—$\rho(\eta) \sim \eta^{-1}$ proposed as universal in chapter 2—which causes the system to converge to a

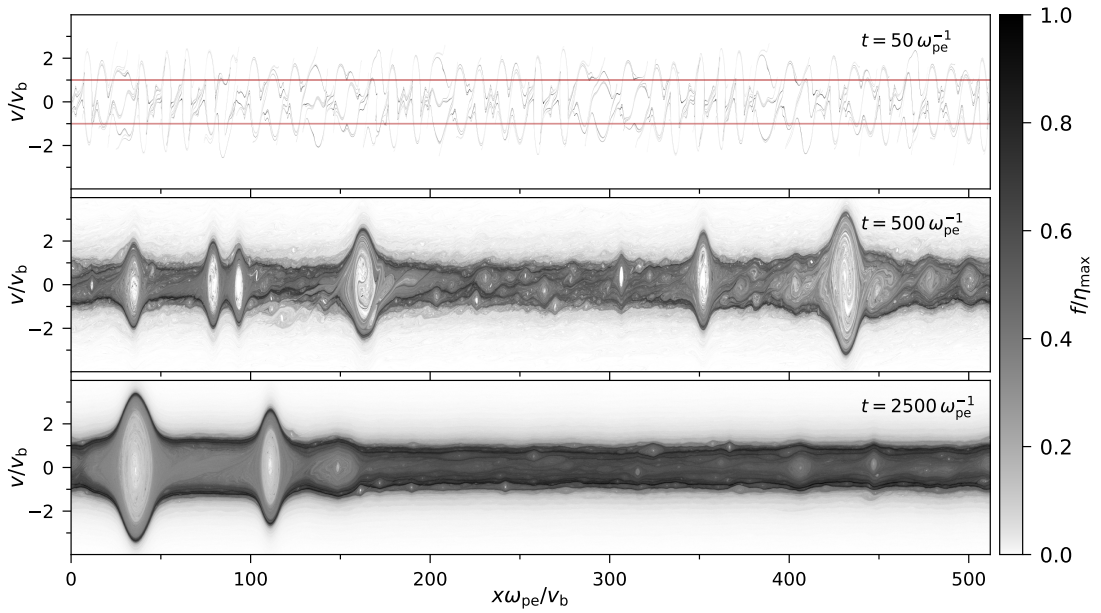

Figure 3.1: Self-consistent particle-in-cell simulation of the temporal evolution (from top to bottom) of the phase-space density $f(x, v)$ for the electron-only two-stream instability, visualised across three time snapshots: $50\,\omega_{pe}^{-1}$, $500\,\omega_{pe}^{-1}$, and $2500\,\omega_{pe}^{-1}$ (where ω_{pe}^{-1} is the electron plasma frequency). The two counter-propagating beams, of speed v_b, (top in red) rapidly go unstable generating the phase-space structure shown at $50\,\omega_{pe}^{-1}$. As the system evolves (middle and bottom snapshots), the formation and merger of turbulent electron holes is clearly evident, which drives the mixing of the phase space and relaxation towards a collisionless equilibrium. The colour scale is normalised to the peak value of $f(x, v)$ in each snapshot.

final, universal Lynden-Bell equilibrium. This equilibrium exhibits a particle-energy distribution possessing a power-law tail with exponent -2, confirming the existence of universal equilibria in strongly turbulent relaxing plasmas.

3.2 Lynden-Bell equilibria and Turbulent amnesia

The formalism of Lynden-Bell is completely prescriptive app: in specifying the initial condition, the energy of the system E and the waterbag content $\rho(\eta)$, via(1.10), are given. Strict adherence to Lynden-Bell's theory would then imply that the system should relax towards the appropriately solved equilibrium (2.4). In chapter 2, this fact led us to investigate the universality of the resulting equilibria. Here, we will ask whether the app in itself should be modified. In any real situation, including

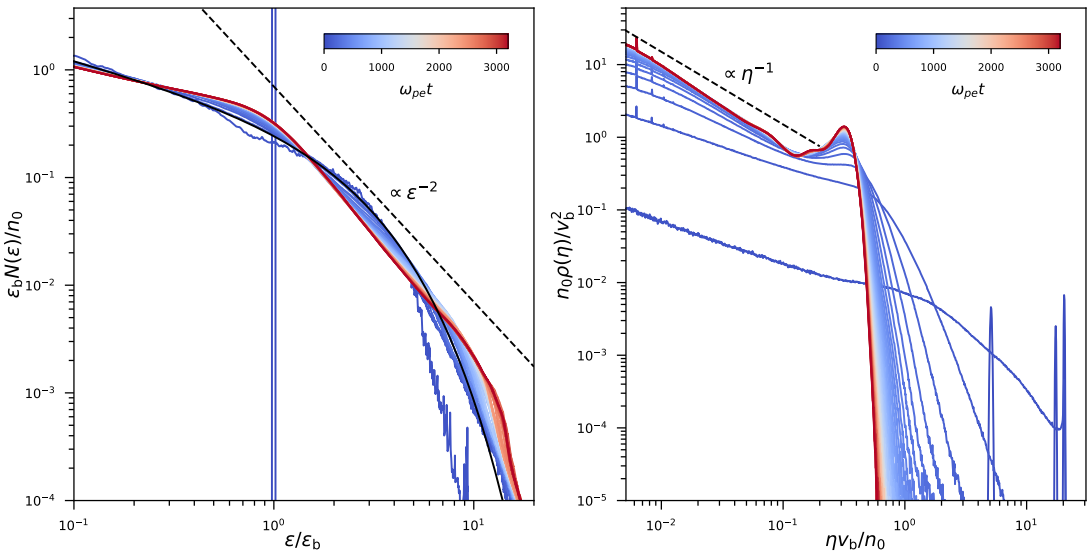

Figure 3.2: Left panel: the evolution of the mean distribution function $N(\varepsilon)$ of particle energies $\varepsilon = mv^2/2$ for the electron two-stream instability. The energies are normalised to the initial beam-particle energy $\varepsilon_\mathrm{b} = mv_\mathrm{b}^2/2$; n_0 is the mean particle number density. The solid black line shows a Maxwellian of the same energy as the initial condition. The lines shown in shades of the blue-to-red colour palette are the distributions from the initial condition (dark blue) to $t = 3200\,\omega_\mathrm{pe}^{-1}$ (dark red) at successive time intervals of $50\,\omega_\mathrm{pe}^{-1}$. Right panel: the evolution of the 'waterbag content' $\rho(\eta)$, over time, computed from (1.10) for the fine-grained phase-space density $f(x, v)$, as detailed in appendix E. The colour scheme for time evolution is the same as in the left panel. The universal asymptotics of $N(\varepsilon)$ and $\rho(\eta)$ are plotted as dashed lines showing the development of a power-law tail (3.9) and the agreement with the predicted scaling of the waterbag content (3.7).

any numerical simulation, one will find that the mean phase-space density will continue evolving even after the initial instability is quenched—and the putative collisionless invariants $\rho(\eta)$ will evolve as well. This is manifest in figure 3.2, showing the energy distribution $N(\varepsilon)$ of particles, and their waterbag content $\rho(\eta)$[1] for a plasma undergoing the nearly collisionless two-stream instability depicted in figure 3.1. It is, of course, clear why $\rho(\eta)$ should change: no truly collisionless plasma (or simulation) can exist, the Liouville equation is never perfectly satisfied, and so the memory of the initial conditions encoded by $\rho(\eta)$ cannot be preserved

[1]Note that, at the earliest time step, there are three peaks in $\rho(\eta)$ which rapidly become a single peak with a low-η tail, in disagreement with the initial condition (3.1). This is because $\rho(\eta)$ is computed on a grid in position and velocity space (see appendix E for details) which have a smaller width than the beam width v_b in velocity space. As such the cells on either side of the beam are only partially filled, giving two extra possible phase-space densities.

forever. Why $\rho(\eta)$ evolves relatively fast, even for nearly collisionless systems, will be discussed below. For the purposes of formulating a theoretical scenario and verifying it numerically, it suffices to know that $\rho(\eta)$ does evolve. We therefore propose an amendment to the Lynden-Bell theory: the principle of 'turbulent amnesia'. Under this scheme, the collisionless dynamics of the plasma push the system towards the Lynden-Bell equilibrium computed using the time-evolving $\rho(\eta)$, which is changed as the turbulence scrambles the system's long-term memory of its prior conditions. This is somewhat analogous to the way in which collisional plasmas undergoing heating will pass through a sequence of Maxwellian distributions with distinct temperatures even though the energy is not a conserved quantity: systems still strive to maximise entropy rapidly even when their invariants are imperfect. We shall return to the theoretical justification of this scenario once we have established its validity numerically.

3.3 Numerical verification

Since plasmas expected to relax to Lynden-Bell equilibria should be as collisionless as possible, the best testbed currently feasible for such relaxation is 1D-1V simulations, which afford the largest number of particles per cell, and, therefore, the lowest collisionality due to the reduced noise levels (Birdsall &Langdon, 1985; Touati et al., 2022). As well as this, one can only expect the Lynden-Bell relaxation to be relevant when the plasma is sufficiently turbulent because inherent in the expectation of rapid maximization of entropy is the assumption of near-perfect mixing in phase space (we shall return below to the question of how perfect it really is). We therefore test our relaxation scenario for two-beam plasmas in 1D—violently unstable situations for which the saturated state is not likely to be obtained via quasilinear theory (cf. Banik et al. 2024)[2]. The beam instabilities that we study are described in detail in appendix E the electron-only two-stream instability and the electron-positron two-stream instability. In the interest of simplicity, we have neglected the ion dynamics in the electron-only two-stream instability. Undoubtedly, this will prove to be an oversimplification at late times and on long length scales

[2]Remarkably, it was conjectured (but not checked) already in the report of the first-ever nonlinear simulation of the two-stream instability that the equilibria eventually reached by the plasma could be explained with Lynden-Bell statistics (Roberts &Berk, 1967).

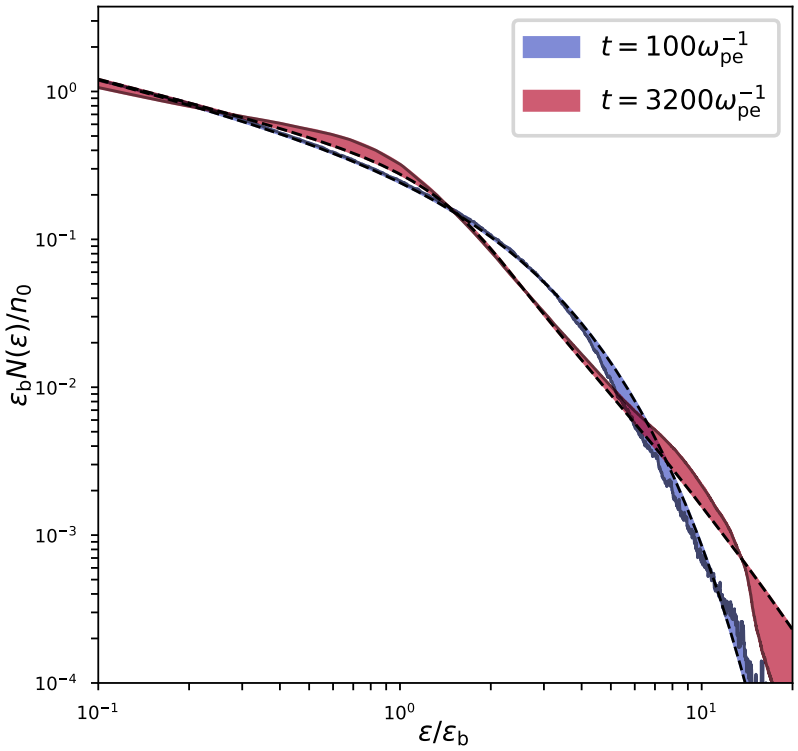

Figure 3.3: Comparison of the directly measured mean phase-space density of the system (solid lines) and the Lynden-Bell equilibrium obtained by taking the first η moment (1.5) of (2.4) (dashed black lines) at two different times. At early times (blue), the mean phase-space density closely follows the Maxwellian associated with $\rho(\eta)$ computed from the initial condition (the two beams, which have a single value of η), while at late times (red), it is better fit by the Lynden-Bell equilibrium associated with the evolved $\rho(\eta)$ shown in figure 3.2 (right panel). The difference between the Lynden-Bell equilibrium and the simulated distribution is highlighted.

(missing the possibility of a plethora of ion-scale physics: see, e.g., Hutchinson 2024), but it is the simplest possible framework to demonstrate our theory.

In figures 3.1 and 3.2, we see that, at early times, $\rho(\eta)$ is well approximated by a single delta function, in line with the initial condition

$$\rho(\eta, t = 0) = 2\Delta v_b \delta(\eta - \eta_{max}), \tag{3.1}$$

where Δv_b is the beam width. At these early times, the mean phase-space density, also shown in figure 3.2, becomes very nearly Maxwellian. As time progresses, however, the mean phase-space density deviates from the Maxwellian equilibrium, forming an energy distribution with a power-law tail $N(\varepsilon) \sim \varepsilon^{-2}$, while its waterbag

content develops a low-η asymptotic $\rho(\eta) \sim \eta^{-1}$. For the measured values of $\rho(\eta)$ and total kinetic energy, we can now treat equation (2.4) as a black box (the numerical method for this is detailed in appendix C) and compare its output to the measured mean phase-space density. Such a comparison for early and late times is shown in figure 3.3 for the the electron-only two-stream instability. We have done this for initial beam widths Δv_b between 2% and 50% of the beam speed v_b and found good agreement with (2.4) up to around 25% of v_b, beyond which point the instability is insufficiently violent to drive strong relaxation over the times and domain sizes that were simulated.

Quantitatively, this is the fundamental result of this chapter (and thesis). Viewed as a purely thermodynamic tool, the Lynden-Bell statistical mechanics correctly predicts the relaxation of the mean phase-space density in a two-stream unstable plasma. Qualitatively, however, it is possible to deduce how this thermodynamic tool operates, and why its output may be generic to many turbulent relaxing systems.

3.3.1 Universal Lynden-Bell equilibria

We first note that the mean distribution (2.4) has a form very similar to the Fermi–Dirac distribution, owing to the analogy between the Pauli exclusion principle and the incompressibility of phase space: two fermions cannot be in the same quantum state and two level sets of phase space cannot be forced together. In a further analogy, the Lynden-Bell equilibrium can be roughly construed as the competition between two pieces of physics: the tendency of each level set to form a Maxwellian equilibrium η by η (as in the numerator of (2.4)) and the incompressibility of phase space (expressed by the denominator of (2.4) being greater than unity).

When the incompressibility condition wins this competition (which is the classical analogue of the system being nearly degenerate), the system is very close to its ground state (Gardner state). The 'Gardner-restacked' minimum-energy counterpart to any given f is a distribution function $f_\mathrm{G}(\varepsilon)$ that has $\rho(\eta)$ identical to that of f, but is a monotonically decreasing function of solely the particle energy $\varepsilon = m|\boldsymbol{v}|^2/2$. In this way, any further reduction of energy would require the distribution function to be larger at lower velocities, impossible without the compression of phase space,

which is forbidden. Therefore, the ground states $f_{\rm G}$ can be computed implicitly from—and are in one-to-one correspondence with— the waterbag content $\rho(\eta)$:

$$\rho(\eta) = \int \mathrm{d}\boldsymbol{v}\, \delta\left(f_{\rm G} - \eta\right) \overset{\rm 1D}{=} -2\frac{\mathrm{d}f_G^{-1}}{\mathrm{d}\eta}. \tag{3.2}$$

The last equality is correct in 1D; in higher dimensions, it would involve the density of states.

This ground state, once computed, sets an important energy scale of the system: the energy of the Gardner distribution $E_{\rm G}$. At energies much larger than this (such as in strongly turbulent systems), one should expect the effect of phase-volume exclusion to become sub-dominant to the effect that pushes each level set towards a Maxwellian—equivalent to the denominator of (2.4) being approximately unity. Then one should anticipate the solution to have the approximate form

$$P(\boldsymbol{v}, \eta > 0) \approx e^{-\beta\eta\varepsilon(\boldsymbol{v})} F(\eta). \tag{3.3}$$

The fugacity $F(\eta)$ can then immediately be deduced from (1.10):

$$F(\eta) \approx \rho(\eta) \left[\int \mathrm{d}\boldsymbol{v}\, e^{-\beta\eta\varepsilon(\boldsymbol{v})}\right]^{-1} \overset{\rm 1D}{=} \sqrt{\frac{\beta\eta m}{2\pi}}\rho(\eta). \tag{3.4}$$

The mean phase-space density can now be computed from (3.3) and (3.4) via (1.5): in 1D-1V,

$$\langle f\rangle(v) \approx \int_{\eta_{\rm min}}^{\eta_{\rm max}} \mathrm{d}\eta \sqrt{\frac{\beta\eta m}{2\pi}}\eta\rho(\eta) e^{-\beta\eta\varepsilon(v)}. \tag{3.5}$$

The meaning of equation (3.5) is physically transparent. Each level set relaxes to a Maxwellian distribution whose temperature is inversely proportional to its phase-space density: less dense portions behave as though they were lighter particles and more dense portions as though they were heavier. The resulting mean phase-space density is, therefore, a superposition of many Maxwellians with relative abundances set by $\rho(\eta)$. At early times in our simulation, $\rho(\eta)$ is approximately a delta function in η, as in (3.1) (by design). This delta function selects from the integral in (3.5) a single η—so the mean phase-space density becomes an actual Maxwellian, viz.,

$$\langle f\rangle(v) = \sqrt{\frac{\beta\eta_{\rm max} m}{2\pi}} 2\Delta v_{\rm b} \eta_{\rm max} e^{-\beta\eta_{\rm max}\frac{1}{2}mv^2} \tag{3.6}$$

This is precisely what is seen in figures 3.2 and 3.3. We further note that this Maxwellianisation is a *fundamentally collisionless effect*: it is an entropic property

of collisionless plasma that turbulence should initially want to push mono-energetic beams towards a Maxwellian equilibrium. That this is a collisionless effect is also obvious as the mean phase-space density is inhomogeneous and continues to evolve after reaching a Maxwellian (the behaviour that would be forbidden by collisional dynamics).

As the collisionless system evolves, $\rho(\eta)$ changes form, as seen in figure 3.2, developing a low-η asymptotic for which $\rho(\eta) \sim \eta^{-1}$. We may therefore assume that $\rho(\eta)$ has a time-asymptotic limit

$$\rho(\eta) \to \eta^{-1} v_{\rm b} G\left(\frac{\eta}{\eta_{\max}}\right), \tag{3.7}$$

where G is the dimensionless function (as in (2.8)) that only has strong dependence on η near the maximum phase-space density $\eta = \eta_{\max}$. Then (3.5) can be written as

$$\langle f\rangle(v) = \frac{v_{\rm b}}{\varepsilon(v)^{3/2}} \int_{\beta\eta_{\min}\varepsilon(v)}^{\beta\eta_{\max}\varepsilon(v)} \mathrm{d}\bar{\eta} \sqrt{\frac{m\bar{\eta}}{2\beta\pi}} G\left(\frac{\bar{\eta}}{\beta\eta_{\max}\varepsilon(v)}\right) e^{-\bar{\eta}}, \tag{3.8}$$

where we have changed the integration variable to $\bar{\eta} = \beta\eta\varepsilon(v)$. Therefore, for velocities such that $\beta\eta_{\max}\varepsilon(v) \gg 1$ and $\beta\eta_{\min}\varepsilon(v) \ll 1$, the $\bar{\eta}$ integral will be a weak function of v, giving $\langle f\rangle$ a universal power-law tail. Recasting this asymptotic as a particle-energy distribution, we get, as in chapter 2

$$N(\varepsilon) \propto \varepsilon^{1/2}\langle f\rangle \sim \varepsilon^{-2} \tag{3.9}$$

There is, of course the subtlety associated with what the value of β turns out to be and, therefore, whether neglecting degeneracy effects was ever a good approximation, which we will address shortly. However, for this case, the ill-justified assumption of neglecting phase-space degeneracy, while quantitatively incorrect, was sufficient to recover the headline result (3.9). This occurs because the denominator of (2.4) is substantially different from unity only at low velocities where high-density level sets crowd each other out in the phase space, whereas at larger velocities, this effect is unimportant—but this is precisely where we expect the ε^{-2} power law to emerge.

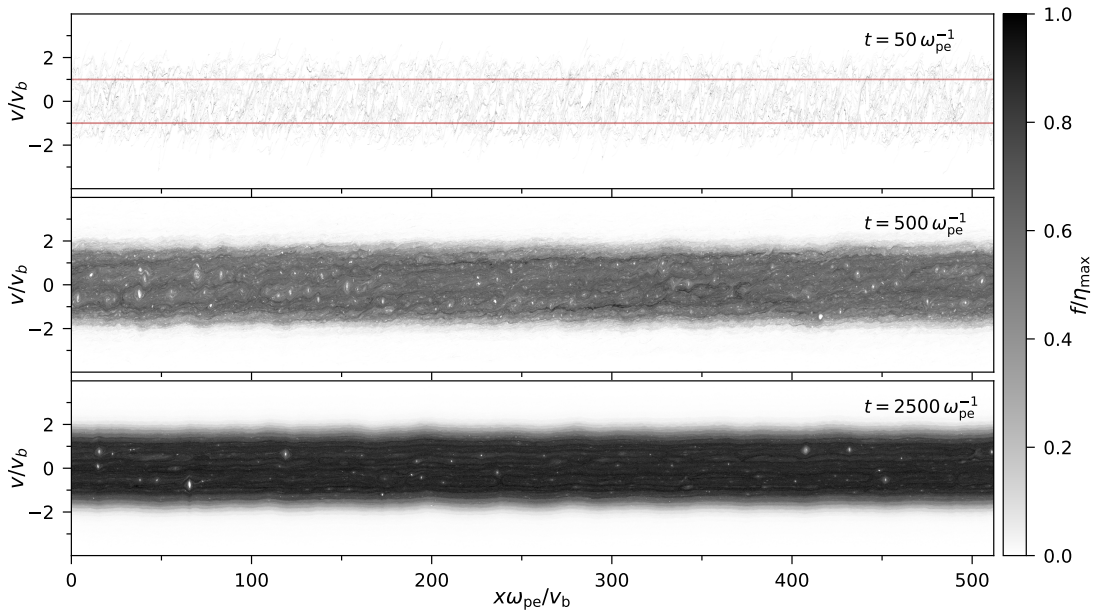

Figure 3.4: Same as figure 3.1 but for the electron-positron two-stream instability, showing the distribution of electrons (which is statistically identical to that of the positrons). While the mixing of the distribution still occurs, the formation of large-scale turbulent structures present in figure 3.1 is entirely absent here, so the resulting distribution is much closer to its ground state. Details of the simulation set-up and parameters are documented in appendix E.

3.3.2 Universality of waterbag content

Thus, the universal high-energy tail $N(\varepsilon) \propto \varepsilon^{-2}$ is a direct consequence of the waterbag content of the system tending towards a $\rho(\eta) \sim \eta^{-1}$ asymptotic at low η—as indeed it does in our numerical experiment. While we do not know how to prove formally that this must happen, it can be justified qualitatively in the following way.

As follows from (1.10) and (1.3), all Casimir invariants of the system can be recovered as moments of $\rho(\eta)$. In particular, the zeroth moment of $\rho(\eta)$ is the volume of phase space in which the phase-space density f takes a non-zero value. This will clearly be finite for our initial set-up with two beams. As the turbulence stirs the plasma, f becomes highly filamented, developing ever sharper gradients in phase space. As collisions (equivalently, particle noise Touati et al. 2022) smooth out these sharp gradients of f, the volume in which f is non-zero should grow (this is manifest in figure 3.1). However, the first moment of $\rho(\eta)$ (which is the particle number) must stay fixed, implying that $\rho(\eta)$ must decrease at large η and increase at small η:

most of the non-empty phase space must be occupied by relatively low phase-space densities. This is manifestly (although not uniquely) satisfied by $\rho(\eta) \sim \eta^{-1}$.

Another argument to the same effect is as follows. The energy of the Gardner distribution corresponding to $\rho(\eta)$ (the ground state defined by (3.2)), can only increase under the action of collisional phase-space diffusion (see, e.g., Tremaine et al. (1986)). It seems reasonable to conjecture, although not easy to prove rigorously, that f_{G} should become more generic as it is thus heated. As was shown in Chapter 2, for a wide class of f_{G}—all functions with any form of exponential decay at high energies[3]—have waterbag content with the asymptotic $\rho(\eta) \sim \eta^{-1}$ as $\eta \to 0$.

3.3.3 Degenerate equilibria

While the formation of the $\rho(\eta) \sim \eta^{-1}$ asymptotic may be generic, it is not the only ingredient necessary to achieve the ε^{-2} power law. This is evidenced in our second numerical experiment: the (non-relativistic) electron-positron two-stream instability—a purely numerical invention (although a reality in the relativistic set-up Arrowsmith et al. 2024; Qu &Fisch 2024), but an interesting case study because it has identical linear physics to the electron-only two-stream instability (up to a rescaling of time) but exhibits a vastly different saturation scenario. In the lower right panel of figure 3.5, we see that the turbulence again pushes the waterbag content towards the asymptotic $\rho(\eta) \sim \eta^{-1}$. However, the mean phase-space density does not form an ε^{-2} tail because the assumption of non-degeneracy (3.3) is completely violated. This can be seen in the upper right panel of figure 3.5. Initially, as the beams are two thin slivers in phase space, the Gardner restacking would amount to just placing the beams at $v = 0$. The resulting ground state would have an energy E_{G} much smaller than the total kinetic energy E, making the system highly non-degenerate. However, as the turbulence breaks $\rho(\eta)$ conservation, the energy of the underlying Gardner distribution grows, causing the system to become more degenerate. In the case of the electron-only two-stream instability, the formation and persistence of coherent structures (phase-space holes) seen in figure 3.1 causes the growth of the Gardner energy to saturate. In contrast, for the electron-positron instability, the holes that do form fail to merge into large-scale,

[3]To be more mathematically precise, this means any function that does not have compact support but decays faster than any power law.

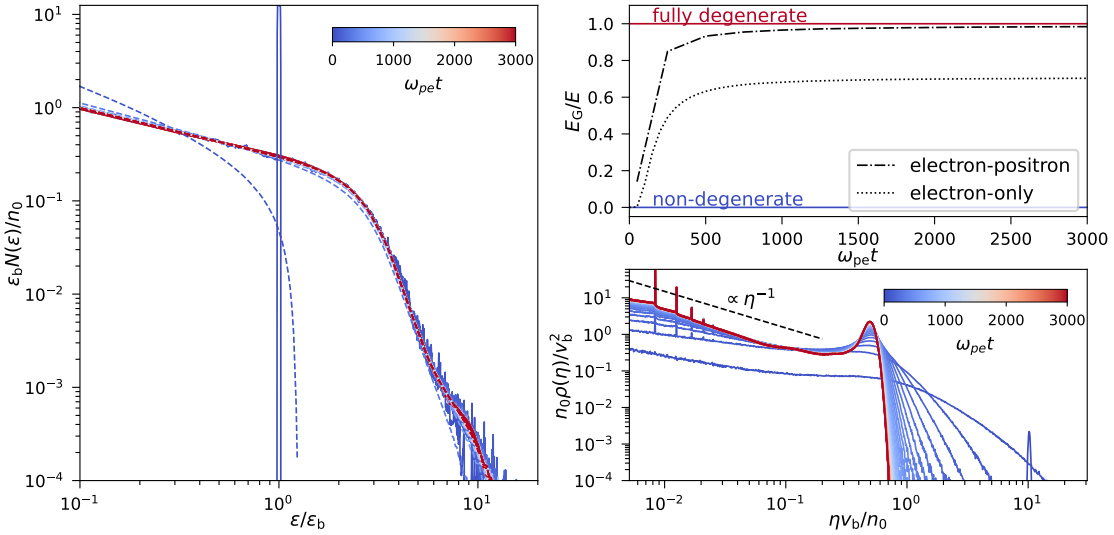

Figure 3.5: Left panel: comparison of the Gardner distribution function of electrons (3.2) (dashed lines) computed from the measured waterbag content $\rho(\eta, t)$ and the mean phase-space density of electrons (solid lines) at a range of times during the evolution of the electron-positron two-stream instability. Upper right panel: ratios of the Gardner energy E_G to the total system energy E during the evolution of the electron-positron (dot dashed) and electron-only (dotted) two-stream instabilities. The degenerate and non-degenerate limits are highlighted in red ($E_G = E$) and blue ($E_G = 0$), respectively. Lower right panel: evolution of the waterbag content $\rho(\eta)$ (of electrons) over time, as in figure 3.2, but this time for the electron-positron two-stream instability. Thus, while the electron-positron system still achieves a universal waterbag content, its Gardner energy has grown so close to its actual energy that it has been frozen into its ground state.

large-amplitude structures (see figure 3.4). As a result, the Gardner energy increases to meet the system's total energy and the evolution freezes in a degenerate state (a phenomenon somewhat similar to 'incomplete relaxation' Chavanis 2006b). This again amounts, of course, to the system reaching its Lynden-Bell equilibrium, but in a much more trivial manner: the distribution simply becomes the Gardner distribution, as seen in the left panel of figure 3.5.

3.4 Breaking of Casimir invariants and the phase-space cascade

Let us now turn our attention to the question of why $\rho(\eta)$ is able to evolve relatively quickly (and how quickly) even in low-collisionality systems. This is because

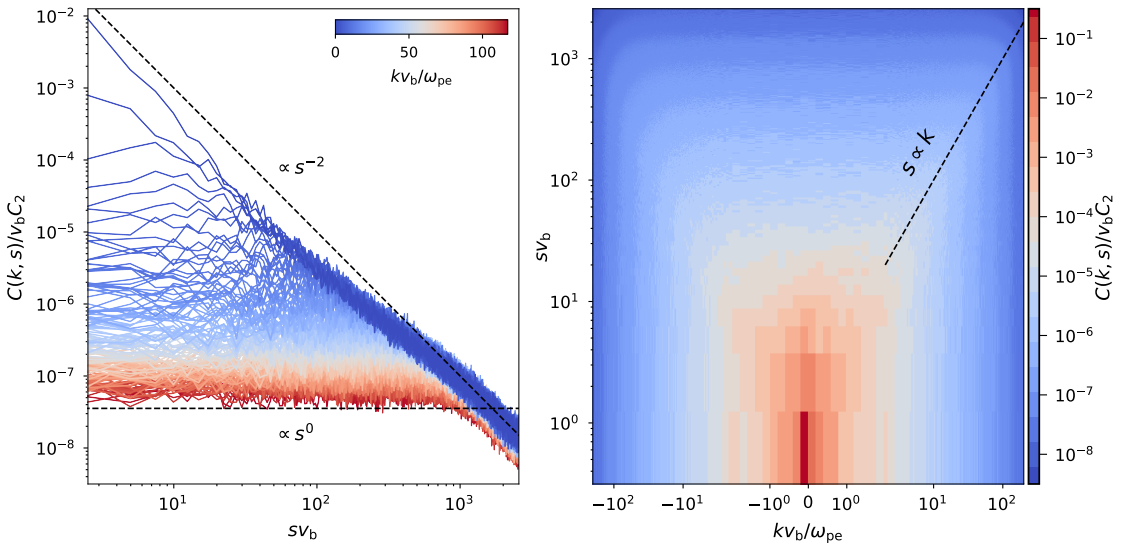

Figure 3.6: Left panel: cuts of $C(k, s)$ as a function of the velocity wave number s (see (3.11)). This spectrum was taken at time $t = 250\,\omega_{\mathrm{pe}}^{-1}$ during the evolution of the electron-only two-stream instability. The lines shown in the blue-to-red colour palette are $C(k, s) = |f_{k,s}|^2$ for different values of k (plotted in increments of $0.8\,\omega_{\mathrm{pe}}/v_{\mathrm{b}}$), representing the velocity-space structure of fluctuations of different scales (large scales: dark blue; small scales: dark red). The theoretically expected asymptotic scalings (3.13) are shown in dashed black lines. Right panel: the same spectrum $C(k, s)$, plotted in 2D to show the so-called 'critical-balance' line $s \sim \gamma^{-1}k$, separating the two asymptotics in (3.13). Physically, the velocity-space linear phase mixing dominates at $s \ll \gamma k$ and the position-space nonlinear mixing of the phase-space density by the electric field dominates at $s \gg \gamma k$.

the thorough mixing in phase space, which is required for the system to reach the Lynden-Bell maximum-entropy state, generates progressively smaller scales in velocity and position space (as seen in figure 3.1) until, at sufficiently small scales, collisions, however small is their rate, act to smooth the distribution function, altering $\rho(\eta)$. For electrostatic plasmas, this process has been shown theoretically (Nastac et al., 2024b) and numerically (Nastac et al., 2024a) to be able to be described as a turbulent cascade through phase space of a representative Casimir invariant (cf. Knorr 1977; Diamond et al. 2010; Servidio et al. 2017; Eyink 2018)

$$C_2 = \int \mathrm{d}\eta\, \eta^2 \rho(\eta) = \frac{1}{L} \int \mathrm{d}x\, \mathrm{d}v\, f^2, \tag{3.10}$$

where L is the domain size. The distribution of this invariant across spatial and velocity scales can be quantified by its spectrum $C(k,s) = |f_{k,s}|^2$, where

$$f_{k,s} = \frac{1}{L}\int \mathrm{d}x\,\mathrm{d}v\, e^{-ikx+isv} f(x,v) \tag{3.11}$$

is the Fourier transform of the phase-space density in position and velocity space. Since C_2 is a quadratic norm of the phase-space density, the contributions to it from the (spatial) mean (1.5) and perturbed $\delta f = f - \langle f \rangle$ parts of f add:

$$C_2 = C_{2,0} + \delta C_2 = \int \mathrm{d}v\, \langle f \rangle^2 + \frac{1}{2\pi}\sum_{k\neq 0}\int \mathrm{d}s\, C(k,s). \tag{3.12}$$

It is not hard to see that the relaxation of the initially unstable state will lead to $C_{2,0}$ decreasing and, therefore, to δC_2 receiving the balance of the C_2 density. This results in an approximately constant flux of C_2 towards small scales (large k and s). A cascade theory in the spirit of Kolmogorov (1941) leads to the following asymptotic form of the spectrum Nastac et al. (2024a):

$$C_{k,s} \propto \begin{cases} s^{-2}, & k \ll \gamma s, \\ k^{-2}, & k \gg \gamma s, \end{cases} \tag{3.13}$$

where γ is a typical shearing rate in phase space set by the amplitude of the electric field. Since the electric field $E = -\partial\varphi/\partial x$ is determined, via Poisson's equation

$$\nabla^2\varphi = 4\pi e\left(\int \mathrm{d}v\, f - n_0\right), \tag{3.14}$$

by the perturbed electron density (for the experiment with static ions), the spectrum of the electric fluctuations at small spatial scales can be determined from the $s \to 0$ asymptotic of (3.13):

$$|E_k|^2 = \frac{16\pi^2 e^2}{k^2} C_{k,s\to 0} \propto k^{-4}. \tag{3.15}$$

Because this spectrum is quite steep, the phase-space density, at any scale, is predominantly stirred by the electric field at the outer scale Batchelor (1959), which, for the two-stream instability, are of the order of (several times) the Debye length $\lambda_\mathrm{D} \sim v_\mathrm{b}/\omega_\mathrm{pe}$. It is because of this large-scale dominance that the shearing rate γ in (3.13) can be assumed to be a scale-independent constant (cf. Nastac et al. 2024a). This also implies that the typical time scale at which f will be mixed all the

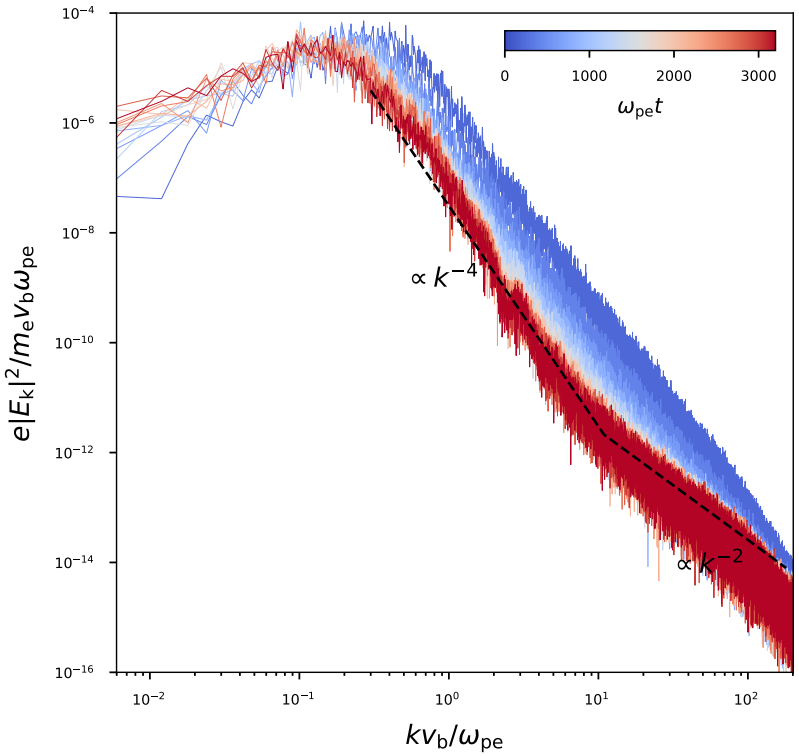

Figure 3.7: The spectrum of the electric field for a range of times during the evolution of the electron-only two-stream instability. The lines shown in the blue-to-red colour palette represent the times from $t = 200\,\omega_{pe}^{-1}$ (dark blue) to $t = 3200\,\omega_{pe}^{-1}$ (dark red), plotted in increments of $200\,\omega_{pe}^{-1}$. The dashed black lines represent the theoretical prediction of a phase-space cascade (3.15), $|E_k|^2 \propto k^{-4}$, which gives way to $|E_k|^2 \propto k^{-2}$, the floor due to the Poisson shot noise of discrete particles at large k (cf. Rostoker (1961); Nastac et al. (2024a)).

way to the phase-space scale l_c where collisions—equivalently, particle noise—start erasing the small-scale structure is (Nastac et al., 2024a)

$$\tau_c \sim \gamma^{-1} \ln \frac{\lambda_D}{l_c}. \tag{3.16}$$

Even without directly computing l_c, it is clear that $l_c \ll \lambda_D$, provided the noise floor is sufficiently low (equivalent to the plasma being weakly coupled, i.e., $n_0 \lambda_D \gg 1$). Thus, (3.16) tells us that $\rho(\eta)$ will change on time scales that are only logarithmically longer than the dynamical time scale ($\sim \gamma^{-1}$) of collisionless relaxation. This provides a modicum of justification for our scheme—confirmed by the numerical experiment with the electron two-stream instability—of evolving the mean phase-

space density as an 'instantaneous' Lynden-Bell equilibrium coupled to a time-dependent waterbag content $\rho(\eta, t)$.

The above argument is supported by the excellent agreement between the theoretical predictions (3.13) and (3.15) and the spectra measured in our numerical simulation of the electron two-stream instability. The phase-space spectrum $C(k, s)$ and the electric-energy spectrum $|E_k|^2$ shown in figures 3.6 and 3.7, respectively, approach their theoretically predicted asymptotic forms around the time when the electron holes, vividly displayed in figure 3.1, begin to move around and merge. It is these dynamics that provide the vigorous mixing that leads to the phase-space cascade and ultimately pushes $\rho(\eta)$ and, therefore, $\langle f \rangle(v)$ towards their universal forms discussed and measured above.

3.5 Summary

In this chapter, we have shown that statistical mechanics can be used to predict classes of universal equilibria for relaxing collisionless plasmas. We have tested this proposition on the example of an electrostatic plasma destabilised by one of the simplest, best-studied instabilities in plasma physics, the two-stream instability, and found good agreement with the theory.

The theory is based on the thermodynamic approach first proposed by Lynden-Bell (1967), maximising the entropy (1.6) subject to the conservation of phase volume (1.3). The additional constraints arising from the latter can be captured by tracking the level sets of the phase-space density via the 'waterbag content' $\rho(\eta)$ (1.10) and endowing the resulting equilibrium (2.4) with some memory of the plasma's earlier state. For the case of the electron two-stream instability, this theory predicts the formation of a Maxwellian distribution (despite the purely collisionless dynamics). This is indeed achieved at early times (see figures 3.2 and 3.3). However, unlike in a collisional regime, the system continues evolving after reaching a Maxwellian. This further evolution is driven by phase-space turbulence stirred up by electron holes generated in the early stages of the instability. We show that this turbulence drives a phase-space cascade of the form predicted by Nastac et al. (2024a,b), giving rise to small-scale structure in both velocity and position space (see figures 3.6 and 3.7). Since no system is truly collisionless (there is

always a finite number of particles), this small-scale structure causes the collisionless invariants $\rho(\eta)$ to be broken (see figures 3.2 and 3.5), causing the Lynden-Bell equilibrium of the system to change and the system to continue evolving on time scales that we estimate, in (3.16), to be only logarithmically longer than the dynamical relaxation times. At these later times, the waterbag content develops a low-η asymptotic $\rho(\eta) \sim \eta^{-1}$, which we argue to be universal, being the natural $\rho(\eta)$ associated with systems that have smooth ground-state phase-space densities. We show that this is indeed what happens in the numerical experiments featuring both the electron-only two-stream instability and the electron-positron two-stream instability. The details of the resulting equilibria in the presence of this universal waterbag content then depend strongly on the amount of energy the well-mixed system is able to retain relative to its ground state (defined by (3.2)). In the case of the electron-only two-stream instability, the saturated energy of the system is larger than the energy of the ground state by a factor of order unity and, as a result, the corresponding Lynden-Bell equilibrium (2.4) is approximately non-degenerate, featuring a distribution of particle energies that has a universal power-law tail $\propto \varepsilon^{-2}$. In contrast, the case of electron-positron two-stream instability exemplifies systems where the non-conservation of $\rho(\eta)$ causes the system to freeze in its ground state, achieving a fully degenerate Lynden-Bell equilibrium (cf. Hosking et al. 2024).

Thus, we have two examples representing what are likely to be two equivalence classes of universal collisionless equilibria: those that, given initial energy E, relax to (approximately) non-degenerate Lynden-Bell distributions such that the corresponding ground state's energy E_{G} is a finite fraction of (or, better still, much smaller than) the system's energy E—and those for which $E_{\mathrm{G}} \approx E$ in the final state, which is, therefore, a fully degenerate Lynden-Bell equilibrium. It is the former class that features the universal high-energy tail $\propto \varepsilon^{-2}$. What appears to help such a state into existence is the emergence of highly non-equilibrium structures—in the case of the electron two-stream instability, electron holes—that can store a certain amount of energy and engage in long-time nonlinear dynamics (moving around, merging) that stir up the plasma and keep it turbulent, rather than decaying into a ground state with no available energy. In a certain sense, this is similar to a non-equilibrium, driven system, where there is continuous injection of δC_2 (see (3.10)), and, presumably, other moments of $\rho(\eta)$, into small scales—which

is why the turbulent spectra that we have observed (figures 3.6 and 3.7) can be predicted by theories that assume continuous driving (Nastac et al., 2024b,a).

This argument contains an apparent internal contradiction: non-degenerate Lynden-Bell equilibria emerge thanks to dynamics that must clearly be inimical to the perfect satisfaction of the hypothesis of perfect mixing. Indeed, it seems unlikely that, with phase-space holes roaming the system, all parts of the phase space can be plausibly assumed equally accessible. It appears, however, that enough of it is accessible for the maximum-entropy principle to assert itself in a theoretically computable way, and that this partial accessibility is a compromise that allows the system to remain turbulent, and converge to a statistical state that is both interesting and has a modicum of universality.

"Those that fix their eager eyes
Ever on the nearest prize
Well may venture to despise
Loftier aspirations.
Pedantry is in demand!
Buy it up at second-hand,
Seek no more to understand
Profitless speculations."

— From 'Lines written under the conviction that it is not wise to read mathematics in November after one's fire is out' by James Clerk Maxwell

4

Collisionless collision integrals

This chapter is adapted from Ewart et al. 2022, J. Plasma Phys. ***88****. 925880501*

4.1 Introduction

It is a common theme of statistical mechanics that it should be a scoundrel's last resort. This is to say that when all hope of understanding the system dynamically is gone, one can turn to thermodynamics to gain some insight, which is often sufficient to then, once again, make dynamical progress. It is for the reason that Landau's collision integral post-dates Maxwell's statistical insight by some 76 years. In a similar vein, it was not long after Lynden-Bell's theory, with its strong analogy to Fermi–Dirac statistics, that Kadomtsev &Pogutse (1970) derived a collision integral for the evolution of a plasma that evolved certain simplified initial conditions towards Fermi–Dirac distributions. Kadomtsev and Pogutse's calculation followed the ethos of the Balescu–Lenard collision integral, but deviated by assuming that the underlying exact distribution function was a piecewise constant function that was everywhere equal either to 0 or to a single constant, the 'single-waterbag' limit of the Lynden-Bell formalism. While taking a step towards verifying that Lynden-Bell's equilibria were relevant entities, Kadomtsev and Pogutse's assumption of

the distribution function only taking two possible values was a very restrictive one, as it limits the system to only two choices of η's.

In this chapter, we will show how Kadomtsev and Pogutse's collision integral can be generalised to the 'multi-waterbag' case, which capture all conceivable initial conditions and leads to the full range of Lynden-Bell equilibria. Our results recover, or generalise, the previous extensions of the Kadomtsev–Pogutse result to multiple waterbags for gravitational and fluid-dynamical kinetic systems, due to Severne &Luwel (1980) and Chavanis (2004, 2005). They also allow one to recover very straightforwardly the Balescu–Lenard integral for collisional plasmas as well as Coulomb collision integrals for quantum plasmas consisting of fermions and bosons.

This chapter is structured as follows. In section 4.2, we outline a quasilinear derivation of a 'collisionless collision integral', which mimics that of Kadomtsev &Pogutse (1970). This will lead us to a general collision integral in terms of an unknown second-order correlator of the exact phase-space density, for which a closure is required. In section 4.2.5, we make the first step towards this closure by assuming short correlation lengths in phase space (the 'microgranulation ansatz', the collisionless version of Boltzmann's *Stosszahlansatz*). In section 4.3, we will make the 'waterbag closure' that will lead to a closed collision integral. To do this, we will first discuss, in section 4.3.1, the 'single-waterbag' closure used by Kadomtsev and Pogutse, and the difficulty of generalising it to multiple waterbags; then, in section 4.3.2, we will show how one can use one's incomplete knowledge of the system to find a 'multi-waterbag' closure (cf. Chavanis 2005). Having derived a multi-waterbag collision integral, we will show in section 4.3.3 that its fixed points are the Lynden-Bell equilibria. In section 4.3.3, we will show that these fixed points are stable by proving that the multi-waterbag collision integral has an H-theorem, viz., that the system will increase Lynden-Bell's entropy (1.6). In section 4.4, we will show that the resolution of the closure problem in section 4.3.2 is actually unnecessary if one is willing to consider the kinetics of a different set of objects: the kinetics of $P(\boldsymbol{x}, \boldsymbol{v}, \eta)$. This 'hyperkinetics' treats the waterbags as the fundamental objects and describes the evolution of $P(\boldsymbol{x}, \boldsymbol{v}, \eta)$ in seven-dimensional phase space, as opposed to the usual six-dimensional phase space of conventional kinetics. In section 4.4.1, we derive the 'hyperkinetic collision integral' (cf. Severne &Luwel 1980), which is pleasingly similar in form to the previous collision integral but has

no need for its thermodynamically motivated closure. In section 4.4.2, we prove an H-theorem for the hyperkinetic collision integral, whence it follows that the latter also pushes the distribution function towards the Lynden-Bell equilibria. In section 4.5.5, we relate the two collision integrals by proving their equivalence over a broad range of initial conditions. Their equivalence however, is not guaranteed for all initial conditions, as the thermodynamically motivated closure can sometimes be too restrictive. In section 4.5, we compare the 'effective collisions' described by the hyperkinetic collision integral to the 'true' collisions described by the Balescu–Lenard collision integral. Namely, we show that, should the microgranulation ansatz hold, the effective collision rate and the energy stored in the fluctuating electric fields will be generically much larger than for true collisions, which gives a method by which the microgranulation ansatz can be tested. A number of caveats with regard to the existence of Lynden-Bell plasmas are discussed in section 4.5.4. Inter-species interactions in Lynden-Bell plasmas are studied in section 4.6: while isotropisation of the electrons and the relaxation of the temperatures of the Lynden-Bell equilibria are a relatively straightforward generalisation of what happens due to standard collisional relaxation (sections 4.6.3 and 4.6.4), the relaxation of the distributions' mean momentum turn out to contain a somewhat surprising effect of spontaneous generation of electron current (section 4.6.3). In section 4.7, we summarise the narrative presented in this chapter and discuss its limitations.

4.2 Derivation of collision integrals from quasi-linear theory

The calculation contained in this section is fundamentally a textbook one, although in practice it may be difficult to find a textbook where it is presented in quite this form, which we consider to be the most transparent.

To understand the relaxation of a distribution function to equilibrium, we begin by considering the evolution of a plasma consisting of multiple species of particles, indexed by α, with mass m_α and charge q_α. The distribution function f_α of these particles evolves according to the collisionless Vlasov equation

$$\frac{\partial f_\alpha}{\partial t} + \boldsymbol{v} \cdot \boldsymbol{\nabla} f_\alpha - \frac{q_\alpha}{m_\alpha}(\boldsymbol{\nabla}\varphi) \cdot \frac{\partial f_\alpha}{\partial \boldsymbol{v}} = 0, \tag{4.1}$$

where the potential φ is determined by Poisson's equation

$$-\nabla^2\varphi = 4\pi \sum_\alpha q_\alpha \int \mathrm{d}\boldsymbol{v}\, f_\alpha. \tag{4.2}$$

For simplicity, we shall limit ourselves to the consideration of the electrostatic case only, forbidding the plasma to host any magnetic field.

In principle, (4.1) and (4.2) already contain the information necessary to evolve the distribution towards equilibrium. Of course, Michelangelo's David was wholly contained within a block of marble, which did not, however, provide great insight into what could lie beneath (Coonin, 2014). The aim of this calculation is then to discern what information can be cut away from from (4.1) and (4.2) to leave only that which is necessary to understand the relaxation of the mean distribution. If we wished to answer the question of collisionless relaxation with complete generality, then the most likely answer is that no information can be cut away. We will therefore specialise to the following simplified, but important, case: a system that is on average uniform in space and for which deviations from homogeneity occur only as small perturbations. In such a regime, it is natural to write f_α as a sum of Fourier modes

$$f_\alpha(\boldsymbol{x}, \boldsymbol{v}) = \sum_{\boldsymbol{k}} f_{\boldsymbol{k}\alpha}(\boldsymbol{v}) e^{i\boldsymbol{k}\cdot\boldsymbol{x}}. \tag{4.3}$$

The evolution of the mean part of the $\boldsymbol{k} = 0$ mode of the distribution function, $f_{0\alpha} = \langle f_{\boldsymbol{k}=0,\alpha}\rangle$, is then

$$\frac{\partial f_{0\alpha}}{\partial t} = \frac{\partial}{\partial \boldsymbol{v}} \cdot \left[\frac{q_\alpha}{m_\alpha} \sum_{\boldsymbol{k}} \boldsymbol{k} \mathrm{Im}\langle \varphi^*_{\boldsymbol{k}} f_{\boldsymbol{k}\alpha}\rangle \right]. \tag{4.4}$$

Here the averages can be rationalised by having many copies of the system, which only differ from one another in microscopic detail. After these copies are allowed to evolve forward to reach the present time, they will each have different values of $f_{\boldsymbol{k}\alpha}$ owing to the initial differences. The angle brackets therefore represent ensemble averages of the system, and the restriction of statistical homogeneity is that the average of any $f_{\boldsymbol{k}\alpha}$ is zero.

From (4.4), we see that to work out the evolution of $f_{0\alpha}$, we must know the remaining fluctuating part of the distribution $f_{\boldsymbol{k}\alpha}$. Therefore, we consider the linearised Vlasov equation for $f_{\boldsymbol{k}\alpha}$:

$$\frac{\partial f_{\boldsymbol{k}\alpha}}{\partial t} + i\boldsymbol{k}\cdot\boldsymbol{v} f_{\boldsymbol{k}\alpha} = i\frac{q_\alpha}{m_\alpha}\varphi_{\boldsymbol{k}}\boldsymbol{k}\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}}, \tag{4.5}$$

where Poisson's equation (4.2) becomes

$$\varphi_{\boldsymbol{k}} = \sum_{\alpha} \frac{4\pi q_\alpha}{k^2} \int \mathrm{d}\boldsymbol{v}\, f_{\boldsymbol{k}\alpha}. \tag{4.6}$$

Note that in a homogeneous system, there can be no mean electric field. Provided that the fluctuations have much smaller amplitudes than $f_{0\alpha}$, (4.4) and (4.5) imply that $f_{\boldsymbol{k}\alpha}$ evolves much faster than the mean distribution. Therefore, the programme for deriving the evolution of $f_{0\alpha}$ becomes the linear one (e.g., Kadomtsev 1965): find the evolution of the fluctuating $f_{\boldsymbol{k}\alpha}$ subject to a constant mean distribution function $f_{0\alpha}$, then evolve $f_{0\alpha}$ using this $f_{\boldsymbol{k}\alpha}$, with the knowledge that, as $f_{0\alpha}$ varies, the fluctuations will constantly adjust themselves.

4.2.1 Linear theory for $f_{\boldsymbol{k}\alpha}$

To solve for the evolution of $f_{\boldsymbol{k}\alpha}$ from (4.5) and (4.6), we follow Landau (1936) and introduce the Laplace transform

$$\hat{\varphi}(p) = \int_0^\infty e^{-pt} \varphi(t) \mathrm{d}t\,, \tag{4.7}$$

and similarly for $f_{\boldsymbol{k}\alpha}(t)$. We now take the Laplace transform of (4.5) and (4.6) to get

$$\hat{f}_{\boldsymbol{k}\alpha}(p) = i\frac{q_\alpha}{m_\alpha} \frac{\hat{\varphi}_{\boldsymbol{k}}(p)}{p + i\boldsymbol{k}\cdot\boldsymbol{v}} \boldsymbol{k}\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} + \hat{h}_{\boldsymbol{k}\alpha}(p), \tag{4.8}$$

$$\hat{\varphi}_{\boldsymbol{k}}(p) = \sum_{\alpha'} \frac{4\pi q_{\alpha'}}{k^2 \epsilon_{\boldsymbol{k}}(p)} \int \mathrm{d}\boldsymbol{v}'\, \hat{h}_{\boldsymbol{k}\alpha'}(p), \tag{4.9}$$

where the dielectric function has emerged, defined by

$$\epsilon_{\boldsymbol{k}}(p) = 1 - i\sum_{\alpha'} \frac{4\pi q_{\alpha'}^2}{m_{\alpha'} k^2} \int \mathrm{d}\boldsymbol{v}' \, \frac{1}{p + i\boldsymbol{k}\cdot\boldsymbol{v}'} \boldsymbol{k}\cdot\frac{\partial f_{0\alpha'}}{\partial \boldsymbol{v}'}, \tag{4.10}$$

while information about the initial condition enters via

$$\hat{h}_{\boldsymbol{k}\alpha}(p) = \frac{g_{\boldsymbol{k}\alpha}(\boldsymbol{v})}{p + i\boldsymbol{k}\cdot\boldsymbol{v}}. \tag{4.11}$$

where $g_{\boldsymbol{k}\alpha}(\boldsymbol{v}) = f_{\boldsymbol{k}\alpha}(t = 0, \boldsymbol{v})$.

The time-dependent solution is given by the inverse Laplace transforms:

$$\varphi_{\boldsymbol{k}}(t) = \frac{1}{2\pi i} \int_{-i\infty+\sigma}^{i\infty+\sigma} \mathrm{d}p\, e^{pt} \hat{\varphi}_{\boldsymbol{k}}(p), \tag{4.12}$$

$$f_{\boldsymbol{k}\alpha}(t) = \int_{-i\infty+\sigma}^{i\infty+\sigma} \frac{\mathrm{d}p}{2\pi i} e^{pt} \left[i\frac{q_\alpha}{m_\alpha} \frac{\hat{\varphi}_{\boldsymbol{k}}(p)}{p + i\boldsymbol{k}\cdot\boldsymbol{v}} \boldsymbol{k}\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} + \hat{h}_{\boldsymbol{k}\alpha}(p) \right]. \tag{4.13}$$

where σ must be chosen so that for all p with $\mathrm{Re}(p) > \sigma$ the integrands are analytic functions.

4.2.2 Quasilinear evolution of $f_{0\alpha}$

Having computed $\varphi(t)$ and $f_{\boldsymbol{k}\alpha}(t)$ we are now in a position to determine the evolution of $f_{0\alpha}$. Substituting (4.12) and (4.13) into (4.4) we obtain the earliest form of our collision integral for the evolution of $f_{0\alpha}$:

$$\frac{\partial f_{0\alpha}}{\partial t} = -\frac{\partial}{\partial \boldsymbol{v}} \cdot \frac{q_\alpha}{m_\alpha} \sum_{\boldsymbol{k}} \boldsymbol{k} \mathrm{Im} \iint \frac{\mathrm{d}p\, \mathrm{d}p'}{(2\pi)^2} e^{(p+p')t} \Big[i \frac{q_\alpha}{m_\alpha} \frac{\langle \hat{\varphi}_{\boldsymbol{k}}(p) \hat{\varphi}^*_{\boldsymbol{k}}(p'^*) \rangle}{p + i\boldsymbol{k}\cdot\boldsymbol{v}} \boldsymbol{k} \cdot \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} + \langle \hat{h}_{\boldsymbol{k}\alpha}(p, \boldsymbol{v}) \hat{\varphi}^*_{\boldsymbol{k}}(p'^*) \rangle \Big]. \quad (4.14)$$

Note that, to avoid confusion in (4.14), both inverse Laplace transforms are taken along the contour running from $-i\infty + \sigma$ to $i\infty + \sigma$. Since the p' contour in (4.14) comes from a complex conjugation, this results in an overall minus sign appearing.

Presently, (4.14) is still in need of information, as it still depends on the averages of $\hat{\varphi}_{\boldsymbol{k}}$ and $\hat{h}_{\boldsymbol{k}}$. The eventual aim is to 'close' this collision integral: to replace these correlators with expressions involving only the mean distribution $f_{0\alpha}$. This will give a differential equation for each $f_{0\alpha}$ in terms of other $f_{0\alpha'}$ only, which, in principle, can then be solved. Such a closure will come from some model of the averages that we have yet to compute, and will be enabled by assumptions introduced in sections 4.2.5 and 4.3.

First we will manipulate (4.14) into a more agreeable form by rewriting the averages $\langle \hat{\varphi}_{\boldsymbol{k}}(p) \hat{\varphi}^*_{\boldsymbol{k}}(p'^*) \rangle$ and $\langle \hat{h}_{\boldsymbol{k}\alpha}(p, \boldsymbol{v}) \hat{\varphi}^*_{\boldsymbol{k}}(p'^*) \rangle$ as correlators solely of $\hat{h}_{\boldsymbol{k}\alpha}$. For the first term in (4.14), we get, using (4.9),

$$i \frac{q_\alpha^2}{m_\alpha^2} \frac{\langle \varphi_{\boldsymbol{k}}(p) \varphi^*_{\boldsymbol{k}}(p'^*) \rangle}{p + i\boldsymbol{k}\cdot\boldsymbol{v}} \boldsymbol{k} \cdot \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} = i \frac{q_\alpha^2}{m_\alpha^2} \sum_{\alpha'\alpha''} \frac{16\pi^2 q_{\alpha'} q_{\alpha''}}{k^4 \epsilon_{\boldsymbol{k}}(p) \epsilon^*_{\boldsymbol{k}}(p'^*)} \iint \mathrm{d}\boldsymbol{v}'\, \mathrm{d}\boldsymbol{v}'' \frac{\langle \hat{h}_{\boldsymbol{k}\alpha''}(p, \boldsymbol{v}'') \hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*, \boldsymbol{v}') \rangle}{p + i\boldsymbol{k}\cdot\boldsymbol{v}} \boldsymbol{k} \cdot \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}}. \quad (4.15)$$

The second term in (4.14), again via (4.9), becomes

$$\begin{aligned} \frac{q_\alpha}{m_\alpha} \langle \hat{h}_{\boldsymbol{k}\alpha}(p) \varphi^*_{\boldsymbol{k}}(p'^*) \rangle &= \sum_{\alpha'} \frac{4\pi q_\alpha q_{\alpha'}}{m_\alpha k^2 \epsilon^*_{\boldsymbol{k}}(p'^*)} \int \mathrm{d}\boldsymbol{v}\, \langle \hat{h}_{\boldsymbol{k}\alpha}(p, \boldsymbol{v}) \hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*, \boldsymbol{v}') \rangle \\ &= \sum_{\alpha'} \frac{4\pi q_\alpha q_{\alpha'}}{m_\alpha k^2 \epsilon_{\boldsymbol{k}}(p) \epsilon^*_{\boldsymbol{k}}(p'^*)} \int \mathrm{d}\boldsymbol{v}' \Bigg[\langle \hat{h}_{\boldsymbol{k}\alpha}(p, \boldsymbol{v}) \hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*, \boldsymbol{v}') \rangle \\ &\quad - i \sum_{\alpha''} \frac{4\pi q_{\alpha''}^2}{m_{\alpha''} k^2} \int \mathrm{d}\boldsymbol{v}'' \frac{\langle \hat{h}_{\boldsymbol{k}\alpha}(p, \boldsymbol{v}) \hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*, \boldsymbol{v}') \rangle}{p + i\boldsymbol{k}\cdot\boldsymbol{v}''} \boldsymbol{k} \cdot \frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''} \Bigg]. \end{aligned} \quad (4.16)$$

At the last step, seemingly gratuitously, we multiplied and divided by the dielectric function (4.10). This will prove to be a useful tactic, as it separates this correlator in two parts: the second, which bears strong resemblance to (4.15), and the first, which will later vanish under certain assumptions. Since this first term is destined to vanish, we will continue as though it has already done so and confirm its disappearance at the end of section 4.2.5. The full evolution equation for the mean distribution function of species α is then

$$\frac{\partial f_{0\alpha}}{\partial t} = -\frac{q_\alpha}{m_\alpha}\frac{\partial}{\partial \boldsymbol{v}}\cdot\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\cdot \mathrm{Re}\iint\frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2}e^{(p+p')t}\sum_{\alpha'\alpha''}\frac{16\pi^2 q_{\alpha'}q_{\alpha''}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)}\iint \mathrm{d}\boldsymbol{v}'\,\mathrm{d}\boldsymbol{v}''$$
$$\left[\frac{\langle\hat{h}_{\boldsymbol{k}\alpha''}(p,\boldsymbol{v}'')\hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*,\boldsymbol{v}')\rangle}{p+i\boldsymbol{k}\cdot\boldsymbol{v}}\frac{q_\alpha}{m_\alpha}\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \frac{\langle\hat{h}_{\boldsymbol{k}\alpha}(p,\boldsymbol{v})\hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*,\boldsymbol{v}')\rangle}{p+i\boldsymbol{k}\cdot\boldsymbol{v}''}\frac{q_{\alpha''}}{m_{\alpha''}}\frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''}\right]. \quad (4.17)$$

This can be written more compactly as

$$\frac{\partial f_{0\alpha}}{\partial t} = \frac{\partial}{\partial \boldsymbol{v}}\cdot\sum_{\alpha''}\int \mathrm{d}\boldsymbol{v}''\left[\mathsf{D}^\alpha_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'')\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \mathsf{D}^\alpha_{\alpha''\alpha}(\boldsymbol{v}'',\boldsymbol{v})\cdot\frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''}\right], \quad (4.18)$$

where the 'diffusion kernel' is

$$\mathsf{D}^\alpha_{\mu\nu}(\boldsymbol{w},\boldsymbol{v}) =$$
$$-\sum_{\nu'}\frac{16\pi^2 q_\mu^2 q_\nu q_{\nu'}}{m_\alpha m_\mu}\mathrm{Re}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\iint\frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2}\frac{e^{(p+p')t}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)}\int\mathrm{d}\boldsymbol{v}'\,\frac{\langle\hat{h}_{\boldsymbol{k}\nu}(p,\boldsymbol{v})\hat{h}^*_{\boldsymbol{k}\nu'}(p'^*,\boldsymbol{v}')\rangle}{p+i\boldsymbol{k}\cdot\boldsymbol{w}}. \quad (4.19)$$

4.2.3 Simplification of the diffusion kernel

Even before we make any assumptions about the nature of the initial condition to decompose the averages, it is possible to simplify the diffusion kernel by appealing to the separation of time scales between the mean and the fluctuations. To do so, we will carry out the p and p' integrals. First we rewrite (4.19) as follows:

$$\mathsf{D}^\alpha_{\mu\nu}(\boldsymbol{w},\boldsymbol{v}) = \sum_{\nu'}\frac{16\pi^2 q_\mu^2 q_\nu q_{\nu'}}{m_\alpha m_\mu}\mathrm{Re}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\int\mathrm{d}\boldsymbol{v}\,'\langle g_{\boldsymbol{k}\nu}(\boldsymbol{v})g^*_{\boldsymbol{k}\nu'}(\boldsymbol{v}')\rangle I_{\boldsymbol{k}}(\boldsymbol{v},\boldsymbol{v}',\boldsymbol{w}), \quad (4.20)$$

where

$$I_{\boldsymbol{k}}(\boldsymbol{v},\boldsymbol{v}',\boldsymbol{w}) = -\iint\frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2}\frac{e^{(p+p')t}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)}\frac{1}{(p+i\boldsymbol{k}\cdot\boldsymbol{v})(p+i\boldsymbol{k}\cdot\boldsymbol{w})(p'-i\boldsymbol{k}\cdot\boldsymbol{v}')}. \quad (4.21)$$

Since this is a contour integral of a holomorphic function (note that $\epsilon^*_{\boldsymbol{k}}(p'^*)$ is a holomorphic function of p'), these integrals can be computed by deforming the p

and p' contours far into the left-hand plane, where the real part of e^{pt} will suppress the integral. All that will remain from this is the contribution from the poles of the integral. Generally, as well as the 'ballistic poles' on the imaginary line, the dielectric function can have poles in the left and right halves of the complex plane. Poles of the dielectric function in the left half of the complex plane, however, correspond to decaying modes, which we will neglect (see figure 4.1). Poles of the dielectric function in the right half of the complex plane correspond to linear instabilities of the distribution function and in principle cannot be neglected. However, we may restrict ourselves to linearly-stable distribution functions with the proviso that we take our initial condition to be the distribution function after all instabilities have vanished and that the initial evolution with instabilities growing and saturating must be treated by a different theory. Within these assumptions the result of the contour integration is

$$I_{\boldsymbol{k}}(\boldsymbol{v},\boldsymbol{v}',\boldsymbol{w}) = \frac{i}{\epsilon_{\boldsymbol{k}}(-i\boldsymbol{k}\cdot\boldsymbol{v})\epsilon^*_{\boldsymbol{k}}(-i\boldsymbol{k}\cdot\boldsymbol{v}')}\frac{e^{-i\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}')t}}{\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{w})}\left[1-\frac{\epsilon_{\boldsymbol{k}}(-i\boldsymbol{k}\cdot\boldsymbol{v})}{\epsilon_{\boldsymbol{k}}(-i\boldsymbol{k}\cdot\boldsymbol{w})}e^{i\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{w})t}\right]. \tag{4.22}$$

We have previously assumed a separation of time scales between the mean distribution function and the perturbed distribution function (known generally as the Bogoliubov ansatz: see Swanson 2008). Therefore, the perturbation can be allowed to evolve to large t before we consider its effect on the mean distribution function. Accordingly, we take the limit $t \to \infty$ in (4.21):

$$I_{\boldsymbol{k}}(\boldsymbol{v},\boldsymbol{v}',\boldsymbol{w}) = \pi\frac{e^{-i\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}')t}}{\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}\epsilon^*_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}'}}\delta\left(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{w})\right), \tag{4.23}$$

where $\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}} \equiv \epsilon_{\boldsymbol{k}}(-i\boldsymbol{k}\cdot\boldsymbol{v})$. Substituting (4.23) back into (4.20) gives a simplified expression for the diffusion kernel:

$$\mathsf{D}^{\alpha}_{\mu\nu}(\boldsymbol{w},\boldsymbol{v}) = \sum_{\nu'}\frac{16\pi^3 q_\mu^2 q_\nu q_{\nu'}}{m_\alpha m_\mu}\mathrm{Re}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\delta(\boldsymbol{k}\cdot(\boldsymbol{w}-\boldsymbol{v}))\int \mathrm{d}\boldsymbol{v}'\,\frac{e^{-i\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}')t}\langle g_{\boldsymbol{k}\nu}(\boldsymbol{v})g^*_{\boldsymbol{k}\nu'}(\boldsymbol{v}')\rangle}{\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}\epsilon^*_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}'}}. \tag{4.24}$$

Note that it is in obtaining this expression that the quasilinear approximation has first been truly used. In the more general expression (4.19), $\hat{h}_{\boldsymbol{k}\nu}(p,\boldsymbol{v})$, instead of being given by (4.11), can always be assumed to contain the nonlinear contributions neglected in (4.5).

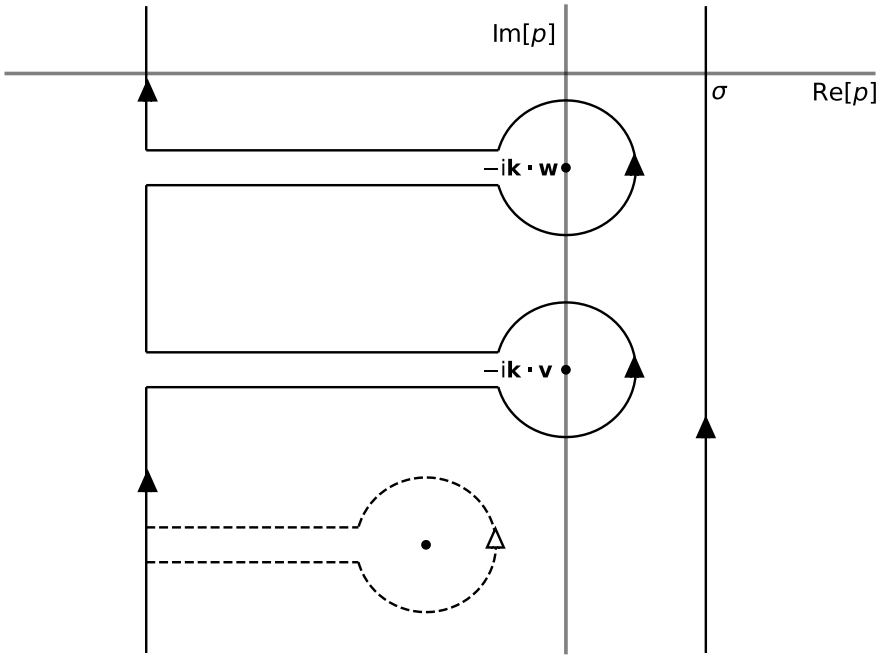

Figure 4.1: The contours used for the p integration in (4.21). The right-hand contour is the original contour, which is calculated by deforming it around the poles of (4.22) onto the left-hand contour. Poles of the dielectric function would correspond to decaying modes but are ignored (shown in the dotted line).

4.2.4 Conservation laws of the general quasilinear collision integral

Due to the symmetries that are possessed by (4.24), it is now possible to show that the collision integral (4.18) with the diffusion kernel (4.24) conserves the particle number, total mean momentum and total mean energy between the species.

The number of particles of species α

$$N_\alpha = V \int \mathrm{d}\boldsymbol{v}\, f_{0\alpha}, \tag{4.25}$$

is trivially conserved because the right-hand side of (4.18) is a full derivative with respect to $\boldsymbol{v}$.

The total momentum of the system is given by

$$P = V \sum_\alpha \int \mathrm{d}\boldsymbol{v}\, m_\alpha \boldsymbol{v} f_{0\alpha}. \tag{4.26}$$

Taking the time derivative of (4.26) using (4.18), we get

$$\begin{aligned}\frac{\mathrm{d}P}{\mathrm{d}t} &= V\sum_{\alpha\alpha''}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\, m_\alpha \boldsymbol{v}\frac{\partial}{\partial \boldsymbol{v}}\cdot\left[\mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'')\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \mathsf{D}^{\alpha}_{\alpha''\alpha}(\boldsymbol{v}'',\boldsymbol{v})\cdot\frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''}\right]\\ &= -V\sum_{\alpha\alpha''}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\,\left[m_\alpha \mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'') - m_{\alpha''}\mathsf{D}^{\alpha''}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'')\right]\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}}\\ &= 0.\end{aligned} \tag{4.27}$$

In the second equality, we integrated by parts and swapped the α, α'' indices, as well as the arguments $\boldsymbol{v}$ and $\boldsymbol{v}''$ in the second term of the integral. This expresses the condition for momentum conservation as the symmetry

$$m_\alpha \mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'') = m_{\alpha''}\mathsf{D}^{\alpha''}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}''). \tag{4.28}$$

From (4.24), and indeed already from (4.20), we see that this is manifestly satisfied, hence the total momentum is conserved but particles of different species may exchange momentum.

We apply much the same procedure to the total energy

$$E = V\sum_{\alpha}\int \mathrm{d}\boldsymbol{v}\,\frac{1}{2}m_\alpha|\boldsymbol{v}|^2 f_{0\alpha}(\boldsymbol{v}). \tag{4.29}$$

Taking the time derivative of (4.29), we get

$$\begin{aligned}\frac{\mathrm{d}E}{\mathrm{d}t} &= V\sum_{\alpha\alpha''}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\,\frac{1}{2}m_\alpha|\boldsymbol{v}|^2\frac{\partial}{\partial \boldsymbol{v}}\cdot\left[\mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'')\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \mathsf{D}^{\alpha}_{\alpha''\alpha}(\boldsymbol{v}'',\boldsymbol{v})\cdot\frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''}\right]\\ &= -V\sum_{\alpha\alpha''}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\,\left[m_\alpha \boldsymbol{v}\cdot\mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'') - m_{\alpha''}\boldsymbol{v}''\cdot\mathsf{D}^{\alpha''}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'')\right]\cdot\frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}}\\ &= 0.\end{aligned} \tag{4.30}$$

Again the swapping of the indices and velocities allows the condition of total-energy conservation to be cast as the symmetry

$$m_\alpha \boldsymbol{v}\cdot\mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}'') = m_{\alpha''}\boldsymbol{v}''\cdot\mathsf{D}^{\alpha''}_{\alpha\alpha''}(\boldsymbol{v},\boldsymbol{v}''), \tag{4.31}$$

which is satisfied by (4.24) due to the delta function that emerged from the approximation relating to the separation of time scales. Essentially what this means is that when the mean distribution function is considered to be linearly stable and to change slowly compared to the fluctuations, the energy E of this distribution function cannot change.

While it is encouraging that we have a collision integral in a general form into which conservation laws are hard-wired, without a general form of the correlator $\langle g_{\alpha\boldsymbol{k}}(\boldsymbol{v})g^*_{\nu\boldsymbol{k}}(\boldsymbol{v}')\rangle$ it is still not possible to determine the evolution of $f_{0\alpha}$. We will now introduce the approximations necessary to arrive at collision integrals in a closed form.

4.2.5 Microgranulation ansatz

The collision integral (4.18), with D given by (4.24), expresses the evolution of the mean distribution function for each species in terms of correlators of the form $\langle g_{\boldsymbol{k}\alpha}(\boldsymbol{v})g^*_{\boldsymbol{k}\alpha'}(\boldsymbol{v}')\rangle$. Despite $g_{\boldsymbol{k}\alpha}$ first entering the calculation as an initial condition, we must now interpret it in a slightly different way.

Suppose we begin at time t_0 with a fluctuation in the distribution function in the form of an initial condition $g_{\boldsymbol{k}\alpha} = f_{\boldsymbol{k}\alpha}(t = t_0)$. We then evolve this perturbation, $f_{\boldsymbol{k}\alpha}(t)$, according to the linearised Vlasov equation (4.5). This evolution is fast, and we assume that the mean distribution, $f_{0\alpha}(t_0)$, remains constant during it. But of course, the mean distribution does evolve, albeit on a longer time scale. So if we pick some later time t_1, such that $t_1 - t_0$ is comparable to that longer time scale, we will find a slightly altered mean distribution function, $f_{0\alpha}(t_1)$. It would then make sense to evolve the perturbation $f_{\boldsymbol{k}\alpha}$ using this new mean distribution. We therefore 'reset' the initial condition to $g_{\boldsymbol{k}\alpha} = f_{\boldsymbol{k}\alpha}(t = t_1)$ and restart the evolution of $f_{\boldsymbol{k}\alpha}(t)$ from this 'new' initial state.

This is very similar to how one might do this numerically: we evolve linearly under a given $f_{0\alpha}$ and initial condition; then, after a small amount of time we, ratchet-like, restart the system, declaring the new initial conditions of this next iteration to be the final state of the previous iteration. The ideal model for the correlation function of $g_{\boldsymbol{k}\alpha}$ would then be one that continually updated to be the correct one for a given mean distribution function. The determination of the steady state phase-space correlation function associated with a given $f_{0\alpha}$ is a difficult problem, not in general solved (cf. Adkins &Schekochihin 2018). Instead, following Kadomtsev &Pogutse (1970), Balescu (1960) and Lenard (1960), we will make some simplifying assumptions about the correlations of $g_{\boldsymbol{k}\alpha}$. The first step towards doing this is to assume that these correlations are very short-distanced in phase space, i.e., that the correlator $\langle g_\alpha(\boldsymbol{x},\boldsymbol{v})g_{\alpha'}(\boldsymbol{x}',\boldsymbol{v}')\rangle$ is zero unless the points $(\boldsymbol{x},\boldsymbol{v})$ and $(\boldsymbol{x}',\boldsymbol{v}')$ lie

very close to each other. We also assume that the perturbed distribution functions of different species are uncorrelated. Mathematically, we express these assumptions in the form of the *microgranulation ansatz*:

$$\langle g_\nu(\boldsymbol{x},\boldsymbol{v})g_{\nu'}(\boldsymbol{x}',\boldsymbol{v}')\rangle = \Delta\Gamma_\nu\delta_{\nu\nu'}\delta(\boldsymbol{x}-\boldsymbol{x}')\delta(\boldsymbol{v}-\boldsymbol{v}')\langle g_\nu^2\rangle(\boldsymbol{v}). \tag{4.32}$$

Here the remaining correlator $\langle g_\nu^2\rangle(\boldsymbol{v})$ is assumed to be spatially independent due to the statistical homogeneity of the system. The new parameter $\Delta\Gamma_\nu$ is the 'volume' of phase space over which the distribution function of a given species has correlated fluctuations. The ansatz (4.32) does not yet constitute a closure as we have not specified the correlator $\langle g_\nu^2\rangle(\boldsymbol{v})$ in terms of $f_{0\nu}(\boldsymbol{v})$. But before this is done, let us implement the ansatz (4.32) to simplify the diffusion kernel (4.24) again.

The Fourier-space correlation function that appears in (4.24) is easily computed from (4.32):

$$\begin{aligned}\langle g_{\boldsymbol{k}\nu}(\boldsymbol{v})g^*_{\boldsymbol{k}\nu'}(\boldsymbol{v}')\rangle &= \iint\frac{\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{x}'}{V^2}e^{-i\boldsymbol{k}\cdot(\boldsymbol{x}-\boldsymbol{x}')}\langle g_\nu(\boldsymbol{x},\boldsymbol{v})g_{\nu'}(\boldsymbol{x}',\boldsymbol{v}')\rangle\\ &= \frac{\Delta\Gamma_\nu}{V}\delta_{\nu\nu'}\langle g_\nu^2\rangle(\boldsymbol{v})\delta(\boldsymbol{v}-\boldsymbol{v}').\end{aligned} \tag{4.33}$$

Thus, the microgranulation ansatz simplifies (4.24) to

$$\mathsf{D}^\alpha_{\mu\nu}(\boldsymbol{w},\boldsymbol{v}) = \frac{16\pi^3q_\mu^2q_\nu^2\Delta\Gamma_\nu}{m_\alpha m_\mu}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{w}-\boldsymbol{v}))}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\langle g_\nu^2\rangle(\boldsymbol{v}). \tag{4.34}$$

This turns the collision integral (4.18) into the following form, written in terms of an as yet undetermined correlator $\langle g_\alpha^2\rangle(\boldsymbol{v})$:

$$\begin{aligned}\frac{\partial f_{0\alpha}}{\partial t} = \sum_{\alpha''}\frac{16\pi^3q_\alpha^2q_{\alpha''}^2}{m_\alpha V}\frac{\partial}{\partial\boldsymbol{v}}\cdot\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\cdot\int\mathrm{d}\boldsymbol{v}''\,\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\\ \left[\frac{\Delta\Gamma_{\alpha''}}{m_\alpha}\langle g_{\alpha''}^2\rangle(\boldsymbol{v}'')\frac{\partial f_{0\alpha}}{\partial\boldsymbol{v}} - \frac{\Delta\Gamma_\alpha}{m_{\alpha''}}\langle g_\alpha^2\rangle(\boldsymbol{v})\frac{\partial f_{0\alpha''}}{\partial\boldsymbol{v}''}\right].\end{aligned} \tag{4.35}$$

In the next section we will show that one can link $\langle g_\alpha^2\rangle(\boldsymbol{v})$ to assumptions about the nature of the exact distribution function, leading finally to a closure in terms of $f_{0\alpha}$.

However, first, let us take care of a piece of unfinished business: we are now in a position to confirm that the first term in (4.16) does indeed vanish. The contribution of that term to (4.14) is

$$-\frac{\partial}{\partial\boldsymbol{v}}\cdot\sum_{\boldsymbol{k}}\boldsymbol{k}\mathrm{Im}\sum_{\alpha'}\frac{4\pi q_\alpha q_{\alpha'}}{m_\alpha k^2}\iint\frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2}\frac{e^{(p+p')t}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)}\int\mathrm{d}\boldsymbol{v}'\,\langle\hat{h}_{\boldsymbol{k}\alpha}(p,\boldsymbol{v})\hat{h}^*_{\boldsymbol{k}\alpha'}(p'^*,\boldsymbol{v}')\rangle \tag{4.36}$$

Therefore, this term will be zero if the quantity

$$\sum_{\alpha'} \frac{4\pi q_\alpha q_{\alpha'}}{m_\alpha k^2} \iint \frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2} \frac{e^{(p+p')t}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)} \int \mathrm{d}\boldsymbol{v}' \, \frac{\langle g_{\boldsymbol{k}\alpha}(\boldsymbol{v}) g^*_{\boldsymbol{k}\alpha'}(\boldsymbol{v}')\rangle}{(p + i\boldsymbol{k}\cdot\boldsymbol{v})\,(p' - i\boldsymbol{k}\cdot\boldsymbol{v}')} \tag{4.37}$$

is real. Taking the complex conjugate of (4.37) and permuting $p \leftrightarrow p'$, we find

$$\sum_{\alpha'} \frac{4\pi q_\alpha q_{\alpha'}}{m_\alpha k^2} \iint \frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2} \frac{e^{(p+p')t}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)} \int \mathrm{d}\boldsymbol{v}' \, \frac{\langle g^*_{\boldsymbol{k}\alpha}(\boldsymbol{v}) g_{\boldsymbol{k}\alpha'}(\boldsymbol{v}')\rangle}{(p' - i\boldsymbol{k}\cdot\boldsymbol{v})\,(p + i\boldsymbol{k}\cdot\boldsymbol{v}')}, \tag{4.38}$$

which is the same as (4.37) if the correlation function satisfies

$$\int \mathrm{d}\boldsymbol{v}' \, \frac{\langle g_{\boldsymbol{k}\alpha}(\boldsymbol{v}) g^*_{\boldsymbol{k}\alpha'}(\boldsymbol{v}')\rangle}{(p + i\boldsymbol{k}\cdot\boldsymbol{v})\,(p' - i\boldsymbol{k}\cdot\boldsymbol{v}')} = \int \mathrm{d}\boldsymbol{v}' \, \frac{\langle g^*_{\boldsymbol{k}\alpha}(\boldsymbol{v}) g_{\boldsymbol{k}\alpha'}(\boldsymbol{v}')\rangle}{(p' - i\boldsymbol{k}\cdot\boldsymbol{v})\,(p + i\boldsymbol{k}\cdot\boldsymbol{v}')}. \tag{4.39}$$

The microgranulation ansatz (4.32), manifestly does satisfy this, so our earlier neglect of the first term in (4.16) is vindicated. Of course, the microgranulation ansatz is quite a simplification of the correlation function of the phase-space density. It is possible that the true correlation function would not have this symmetry and, consequently, give rise to a qualitatively different evolution of the mean distribution function.

4.3 Waterbag representation

The microgranulation moves us a long way to crafting a complete collision integral. However, it still requires that we express the variance of the fluctuations $\langle g_\alpha^2\rangle$ in terms of mean quantities. Of course, we know from (1.3) that the level sets of the distribution must be the same as they were in the initial condition (modulo the caveats present in chapter 3); a fact which limits the ways in which the exact phase-space density can be distinct from the mean. To explore this further, let us consider, as did Lynden-Bell (1967) and Kadomtsev &Pogutse (1970), the single-waterbag distribution—a distribution for which the phase-space density is a piecewise constant function. A single-waterbag distribution is then one for which $f_\alpha(\boldsymbol{x},\boldsymbol{v})$ is piecewise constant and equal either to zero or to a single value η (see figure 4.2), while a multi-waterbag distribution function can take some countable set of values $\{\eta_i\}$.

Before discussing the multi-waterbag case, we will extend Kadomtsev and Pogutse's single-waterbag model to multiple species, recovering their results and highlighting the analytical gain of the waterbag model.

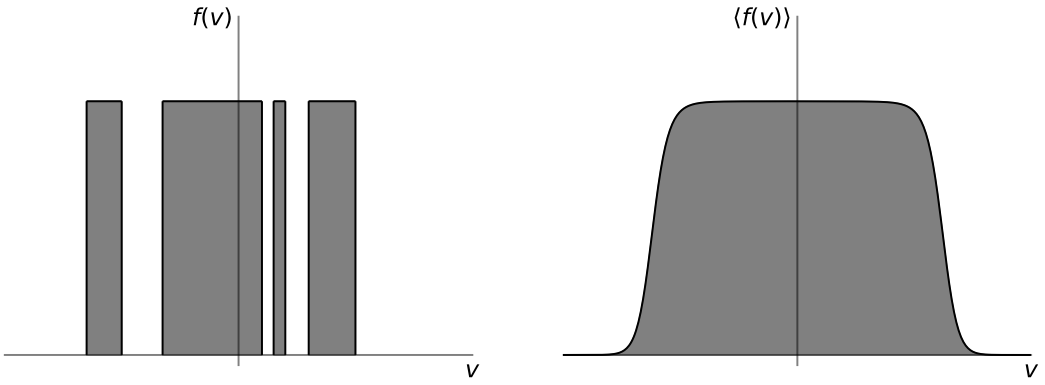

Figure 4.2: A 1D cartoon of the exact and mean distribution functions (phase-space densities) for the single-waterbag model. As discussed in section 4.2.5, the mean distribution should be considered in the ensemble-averaged sense, so that the many realisations of systems like the exact one given in the left panel evolve on average to look like the Fermi–Dirac distribution in the right panel.

4.3.1 Single-waterbag closure

When only one waterbag density η_α for each species α is assumed, this has immediate implications for the mean phase-space density. If the exact phase-space density $f_\alpha(\boldsymbol{x}, \boldsymbol{v})$ is only ever η_α or zero, then its average value is directly related to the probability of a portion of phase space being occupied. We will therefore define $p_\alpha(\boldsymbol{v})$ as the probability that, at a given point $\boldsymbol{v}$, the phase-space density is η_α for particles of species α. This can also be said as '$p_\alpha(\boldsymbol{v})$ is the probability that there is a waterbag of species α at $\boldsymbol{v}$' with the understanding that a waterbag is a patch of phase space of a certain density.

Written in terms of this $p_\alpha(\boldsymbol{v})$, the mean phase-space density $f_{0\alpha}(\boldsymbol{v})$ is then

$$f_{0\alpha}(\boldsymbol{v}) = \langle f_\alpha(\boldsymbol{x}, \boldsymbol{v}) \rangle = \eta_\alpha p_\alpha(\boldsymbol{v}). \tag{4.40}$$

Note that there is no $\boldsymbol{x}$ dependence because our system is assumed to be statistically homogeneous in the position space. Likewise, the squared phase-space density will be η_α^2 with probability p_α or zero otherwise. Therefore,

$$\langle f_\alpha^2 \rangle(\boldsymbol{v}) = \eta_\alpha^2 p_\alpha(\boldsymbol{v}) = \eta_\alpha f_{0\alpha}(\boldsymbol{v}). \tag{4.41}$$

The correlator $\langle g_\alpha^2 \rangle(\boldsymbol{v})$ can then be determined immediately:

$$\langle g_\alpha^2 \rangle(\boldsymbol{v}) = \langle [f_\alpha(\boldsymbol{x}, \boldsymbol{v}) - f_{0\alpha}(\boldsymbol{v})]^2 \rangle = \langle f_\alpha^2 \rangle(\boldsymbol{v}) - f_{0\alpha}^2(\boldsymbol{v}) = [\eta_\alpha - f_{0\alpha}(\boldsymbol{v})]\, f_{0\alpha}(\boldsymbol{v}). \tag{4.42}$$

With this closure, (4.35) becomes the multi-species generalisation of Kadomtsev and Pogutse's collision integral:

$$\frac{\partial f_{0\alpha}}{\partial t} = \sum_{\alpha''} \frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{m_\alpha V} \frac{\partial}{\partial \boldsymbol{v}} \cdot \sum_{\boldsymbol{k}} \frac{\boldsymbol{k}\boldsymbol{k}}{k^4} \cdot \int \mathrm{d}\boldsymbol{v}'' \, \frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2} \left\{ \frac{\Delta\Gamma_{\alpha''}}{m_\alpha} \Big[\eta_{\alpha''} - f_{0\alpha''}(\boldsymbol{v}'')\Big] f_{0\alpha''}(\boldsymbol{v}'') \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \frac{\Delta\Gamma_\alpha}{m_{\alpha''}} \Big[\eta_\alpha - f_{0\alpha}(\boldsymbol{v})\Big] f_{0\alpha}(\boldsymbol{v}) \frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''} \right\}. \quad (4.43)$$

The fixed points of this collision integral are Fermi–Dirac distributions:

$$f_{0\alpha}(\boldsymbol{v}) = \frac{\eta_\alpha}{1 + \exp\left[\beta \Delta\Gamma_\alpha \eta_\alpha \left(\frac{1}{2} m_\alpha |\boldsymbol{v}|^2 - \mu_\alpha\right)\right]}. \quad (4.44)$$

We can see that these are indeed fixed points by noting that, for (4.44),

$$\frac{\Delta\Gamma_{\alpha''}}{m_\alpha} \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} = -\beta \Delta\Gamma_\alpha \Delta\Gamma_{\alpha''} \boldsymbol{v} \Big[\eta_\alpha - f_{0\alpha}(\boldsymbol{v})\Big] f_{0\alpha}(\boldsymbol{v}). \quad (4.45)$$

Thus, the two bracketed terms in (4.43) become identical except for a factor of $\boldsymbol{v}$ or $\boldsymbol{v}''$. This difference vanishes when the bracket is dotted with $\delta\left(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}'')\right)\boldsymbol{k}$, setting the collision integral to zero. The demonstration that the fixed points (4.44) are stable will come in section 4.3.3, where we prove the general H-theorem for the multi-waterbag model, which reduces to the single-waterbag model in the appropriate limit. This H-theorem will guarantee that the mean phase-space density tends towards the maximum the entropy (1.6) subject to the conservation laws proved in section 4.2.4. These conservation laws will therefore dynamically enforce the correct choice of β and μ_α in (4.44).

4.3.2 Multi-waterbag closure

By assuming, as we did in section 4.3.1, that the exact phase-space density was a single-waterbag distribution, we made the closure for $\langle g_\alpha^2 \rangle(\boldsymbol{v})$ very simple. The reason for this boils down to the fact that, for a single-waterbag distribution, knowledge of the mean phase-space density was sufficient to determine all the requisite statistical information about the exact one, namely the probabilities $p_\alpha(\boldsymbol{v})$. In a similar vein, one can ask what information there is to be determined in an N-waterbag system, where the exact phase-space density can take the values $\{\eta_{J\alpha}\}_{J=1,2,\ldots,N}$. As

a direct generalisation of the single-waterbag model, we then define $p_{J\alpha}(\boldsymbol{v})$ to be the probability that the phase-space density of species α will be equal to $\eta_{J\alpha}$ at velocity $\boldsymbol{v}$. In other words, 'the $\eta_{J\alpha}$ waterbag of species α has probability $p_{J\alpha}(\boldsymbol{v})$ of being present at $\boldsymbol{v}$'. Then we can write the mean phase-space density as

$$f_{0\alpha}(\boldsymbol{v}) = \langle f_\alpha(\boldsymbol{x}, \boldsymbol{v})\rangle = \sum_J \eta_{J\alpha} p_{J\alpha}(\boldsymbol{v}), \tag{4.46}$$

and the mean square one as

$$\langle f_\alpha^2\rangle(\boldsymbol{v}) = \sum_J \eta_{J\alpha}^2 p_{J\alpha}(\boldsymbol{v}). \tag{4.47}$$

Now the complication introduced by multiple waterbags becomes obvious. When $N > 1$, knowledge of just $f_{0\alpha}$ cannot uniquely determine the probabilities $p_{J\alpha}(\boldsymbol{v})$ and so does not exactly determine $\langle g_\alpha^2\rangle$ in (4.35). Therefore, it is not possible to close (4.35) without further information. In principle, this information could be extracted. One could construct the evolution equations for the higher-order moments of the exact phase-space density $\langle f_\alpha^2\rangle$, $\langle f_\alpha^3\rangle$,..., $\langle f_\alpha^N\rangle$. The evolution of the i-th such moment would generally depend on the $(i+1)$-st moment. In that way, one would have N analogues of (4.46) and (4.47) for the N unknowns $p_{J\alpha}(\boldsymbol{v})$. This would then allow one to write $\langle f_\alpha^{N+1}\rangle$ as a function of all the lower-order moments and finally close the system, featuring N coupled collision integrals.

In section 4.4, we will show that such a scheme can be made tractable by calculating all such moments in one fell swoop by increasing the dimension of the phase space. First, however, we consider a simple closure that will prove illuminating in understanding the relaxation of collisionless systems. Instead of hoping to gain any further knowledge of the system, we ask what is the most likely assignment of probabilities $p_{J\alpha}(\boldsymbol{v})$ given that the mean phase-space density is $f_{0\alpha}(\boldsymbol{v})$ and nothing else is known. This can be answered in the spirit of statistical inference by maximising the standard Shannon (1948) entropy:

$$S_\alpha = -\int \mathrm{d}\boldsymbol{v} \sum_{J=0}^{N} p_{J\alpha}(\boldsymbol{v}) \ln p_{J\alpha}(\boldsymbol{v}), \tag{4.48}$$

where the sum now includes $J = 0$ as 'the empty waterbag', representing the probability that a given point in phase space is has zero density of particles. The Shannon entropy (4.48) must be maximised subject to the condition that the

mean phase-space density is $f_{0\alpha}(\boldsymbol{v})$, i.e., that the $p_{J\alpha}(\boldsymbol{v})$'s obey (4.46), and that the probabilities $p_{J\alpha}(\boldsymbol{v})$ sum to unity at each $\boldsymbol{v}$. Another constraint on these probabilities is that the total number of particles per unit volume contained within each waterbag,

$$n_{J\alpha} = \eta_{J\alpha} \int \mathrm{d}\boldsymbol{v}\, p_{J\alpha}(\boldsymbol{v}), \tag{4.49}$$

is fixed for all $J \neq 0$, since the phase volume corresponding to each $\eta_{J\alpha}$ must be conserved by the evolution under the Vlasov equation. Such invariants are often called (or are equivalent to) Casimir invariants (cf. Chavanis 2004, 2005, Zhdankin 2022a).

Thus, we maximise the following functional for each species:

$$\begin{aligned} &- \int \mathrm{d}\boldsymbol{v} \sum_J p_{J\alpha}(\boldsymbol{v}) \ln p_{J\alpha}(\boldsymbol{v}) - \int \mathrm{d}\boldsymbol{v}\, \psi_\alpha(\boldsymbol{v}) \left[\sum_J \eta_{J\alpha} p_{J\alpha}(\boldsymbol{v}) - f_{0\alpha}(\boldsymbol{v}) \right] \\ &\quad - \sum_{J\neq 0} \gamma_{J\alpha} \left[\int \mathrm{d}\boldsymbol{v}\, p_{J\alpha}(\boldsymbol{v}) - \frac{n_{J\alpha}}{\eta_{J\alpha}} \right] - \int \mathrm{d}\boldsymbol{v}\, \lambda_\alpha(\boldsymbol{v}) \left[\sum_J p_{J\alpha}(\boldsymbol{v}) - 1 \right] \to \text{ max}, \end{aligned} \tag{4.50}$$

where $\psi_\alpha(\boldsymbol{v})$, $\gamma_{J\alpha}$, and $\lambda_\alpha(\boldsymbol{v})$ are Lagrange multipliers. The result is

$$p_{J\alpha}(\boldsymbol{v}) = \frac{1}{Z_\alpha\left(\psi_\alpha(\boldsymbol{v})\right)} e^{-\psi_\alpha(\boldsymbol{v})\eta_{J\alpha} - \gamma_{J\alpha}}, \tag{4.51}$$

where the 'partition function' of species α is

$$Z_\alpha(\psi_\alpha(\boldsymbol{v})) = 1 + \sum_{J\neq 0} e^{-\psi_\alpha(\boldsymbol{v})\eta_{J\alpha} - \gamma_{J\alpha}}, \tag{4.52}$$

and the Lagrange multipliers $\psi_\alpha(\boldsymbol{v})$ and $\gamma_{J\alpha}$ must be chosen to enforce the constraints (4.46) and (4.49). Analogously to the standard Gibbs (1902) statistical mechanics, (4.46) becomes

$$f_{0\alpha}(\boldsymbol{v}) = \frac{1}{Z_\alpha\left(\psi_\alpha(\boldsymbol{v})\right)} \sum_J \eta_{J\alpha} e^{-\psi_\alpha(\boldsymbol{v})\eta_{J\alpha} - \gamma_{J\alpha}} = -\frac{\partial \ln Z_\alpha}{\partial \psi_\alpha}. \tag{4.53}$$

Thus, the mean phase-space density plays the role that energy does in the regular prescription of statistical mechanics, and $\psi_\alpha(\boldsymbol{v})$ that of inverse temperature, which is local in $\boldsymbol{v}$. In this formalism, therefore,

$$\begin{aligned} \langle g_\alpha^2 \rangle(\boldsymbol{v}) &= \frac{1}{Z_\alpha(\psi_\alpha(\boldsymbol{v}))} \sum_J \eta_{J\alpha}^2 e^{-\psi_\alpha(\boldsymbol{v})\eta_{J\alpha} - \gamma_{J\alpha}} - f_{0\alpha}^2(\boldsymbol{v}) \\ &= \frac{1}{Z_\alpha} \frac{\partial^2 Z_\alpha}{\partial \psi_\alpha^2} - \frac{1}{Z_\alpha^2} \left(\frac{\partial Z_\alpha}{\partial \psi_\alpha} \right)^2 = \frac{\partial^2 \ln Z_\alpha}{\partial \psi_\alpha^2} = -\frac{\partial f_{0\alpha}}{\partial \psi_\alpha}, \end{aligned} \tag{4.54}$$

reminiscent of the heat capacity of a system in the Gibbs ensemble. Such a closure appears to have been first proposed by Chavanis (2005), in the context of geophysical turbulence.[1]

With $\langle g_\alpha^2 \rangle(\boldsymbol{v})$ thus specified, we may substitute (4.54) into (4.35) to get

$$\frac{\partial f_{0\alpha}}{\partial t} = -\sum_{\alpha''} \frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{m_\alpha V} \frac{\partial}{\partial \boldsymbol{v}} \cdot \sum_{\boldsymbol{k}} \frac{\boldsymbol{k}\boldsymbol{k}}{k^4} \cdot \int \mathrm{d}\boldsymbol{v}'' \, \frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2} \left(\frac{\Delta\Gamma_{\alpha''}}{m_\alpha} \frac{\partial \psi_\alpha}{\partial \boldsymbol{v}} - \frac{\Delta\Gamma_\alpha}{m_{\alpha''}} \frac{\partial \psi_{\alpha''}}{\partial \boldsymbol{v}''} \right) \frac{\partial f_{0\alpha}}{\partial \psi_\alpha}(\boldsymbol{v}) \frac{\partial f_{0\alpha''}}{\partial \psi_{\alpha''}}(\boldsymbol{v}''). \quad (4.55)$$

This is the collision integral for a multi-waterbag Lynden-Bell plasma. The instantaneous relationship between $f_{0\alpha}(\boldsymbol{v})$ and $\psi_\alpha(\boldsymbol{v})$ is given by (4.53) with Z_α defined by (4.52) and $\gamma_{J\alpha}$'s set by (4.49) and (4.51). The set of constants $n_{J\alpha}$ in (4.49) is fixed by the initial condition and cannot change during the evolution of $f_{0\alpha}$. This is the way in which phase-volume conservation in a collisionless plasma imprints a signature of the initial distribution (its 'waterbag content') on its otherwise universal evolution towards the Lynden-Bell equilibria.

Note that the closure proposed above amounts to assuming that the system always quickly attains a 'local' equilibrium in phase space given by (4.51) with a $\boldsymbol{v}$-dependent 'inverse phase temperature' $\psi_\alpha(\boldsymbol{v})$. The integral (4.55) then describes the evolution toward a 'global' equilibrium. We shall discuss the plausibility of this assumption in section 4.5.5.

4.3.3 Properties of the multi-waterbag collision integral

Having derived the collision integral (4.55), we now proceed to study its properties. First, we will determine its fixed points, then confirm that they are stable attractors by proving an H-theorem for our collision integral.

Multi-waterbag equilibria

Since $\partial f_{0\alpha}/\partial \psi_\alpha \leq 0$ (and only zero in pathological cases), the integral on the right-hand side of (4.55) can vanish only if its integrand vanishes. Manifestly, it does so if

$$\psi_\alpha(\boldsymbol{v}) = \beta \Delta\Gamma_\alpha \frac{1}{2} m_\alpha |\boldsymbol{v}|^2 \equiv \beta \Delta\Gamma_\alpha \epsilon_\alpha(\boldsymbol{v}), \quad (4.56)$$

[1] For a different collision integral, but one can recover an integral similar to the one that we are about to produce if one applies this closure to one of the collision integrals proposed in Chavanis (2004).

where $\epsilon_\alpha(\boldsymbol{v})$ is the energy of a particle of species α with velocity $\boldsymbol{v}$. Therefore, from (4.51),

$$p_{J\alpha}(\boldsymbol{v}) = \frac{\exp\Big\{ -\beta \Delta\Gamma_\alpha \eta_{J\alpha} \left[\epsilon_\alpha(\boldsymbol{v}) - \mu_{J\alpha}\right] \Big\}}{1 + \sum_{J'\neq 0} \exp\Big\{ -\beta \Delta\Gamma_\alpha \eta_{J'\alpha} \left[\epsilon_\alpha(\boldsymbol{v}) - \mu_{J'\alpha}\right] \Big\}}. \tag{4.57}$$

We have brought this into a form pleasingly similar to the Fermi–Dirac distribution by rewriting $\gamma_{J\alpha} = -\beta\Delta\Gamma_\alpha \eta_{J\alpha}\mu_{J\alpha}$ to define the chemical potential $\mu_{J\alpha}$ of the waterbag J of species α. These are, of course, the multi-species discrete form of (2.4) of which examples are shown in figure 4.3. In this derivation, they have emerged as fixed points of the plasma's dynamical evolution and we will now prove their stability to the collision integral (4.55).

The H-theorem

To prove the stability of the Lynden-Bell equilibria (4.57) to the collision integral (4.55), we will now prove an H-theorem for the collision integral (4.55). Namely, we will prove that there is a functional of $f_{0\alpha}$ that can only be increased by evolution under (4.55). This functional is then the entropy of our system, and if it has a maximum, this is a stable attractor of the evolution, since the system could not depart from this state without lowering the entropy.

Consider the following obvious candidate for entropy[2]:

$$S = -\sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \int \mathrm{d}\boldsymbol{v} \sum_J p_{J\alpha}(\boldsymbol{v}) \ln p_{J\alpha}(\boldsymbol{v}). \tag{4.58}$$

As before, the sum in this definition includes the empty waterbag $J = 0$, whose probability is $p_{0\alpha} = 1 - \sum_{J\neq 0} p_{J\alpha}$. Therefore, (4.58) reduces to the well-known Fermi–Dirac entropy in the single-waterbag case. According to our closure scheme, $p_{J\alpha}(\boldsymbol{v})$ can be written as (4.51). The entropy (4.58) then becomes

$$\begin{aligned} S &= \sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \int \mathrm{d}\boldsymbol{v} \sum_J p_{J\alpha}(\boldsymbol{v}) \Big[\psi_\alpha(\boldsymbol{v}) \eta_{J\alpha} + \gamma_{J\alpha} + \ln Z_\alpha(\psi_\alpha(\boldsymbol{v})) \Big] \\ &= \sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \left\{ \int \mathrm{d}\boldsymbol{v} \Big[f_{0\alpha}(\boldsymbol{v}) \psi_\alpha(\boldsymbol{v}) + \ln Z_\alpha(\psi_\alpha(\boldsymbol{v})) \Big] + \sum_J \gamma_{J\alpha} \frac{n_{J\alpha}}{\eta_{J\alpha}} \right\}. \end{aligned} \tag{4.59}$$

[2]The attentive reader may be wondering about the species-dependent prefactor of $\Delta\Gamma_\alpha$ appearing in (4.58), which was not a priori obvious from (1.6). Effectively smaller correlation volumes for particular species mean there are more permutations of phase space possible, creating more entropy. A more rigorous justification of this can be found in appendix A

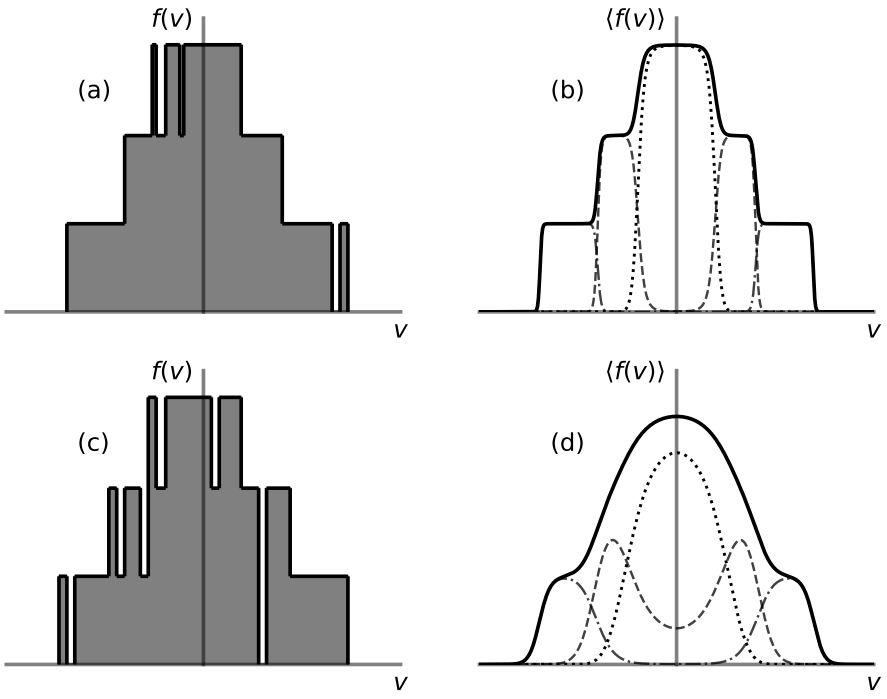

Figure 4.3: A 1D cartoon depicting exact and mean phase-space densities for a three-waterbag system. Panels (b) and (d) show $f_0(v)$ corresponding to the exact $f(v)$ shown in panels (a) and (c), respectively. The dotted, dashed and dot-dashed lines in panels (b) and (d) show to the contributions $\eta_J p_J(v)$ of each of the three waterbag densities to the mean phase-space density. The system in the top two panels has a total energy equal to $E = 1.003E_{\rm min}$, where $E_{\rm min}$ is the minimal energy that the system could have subject to phase-volume conservation. The system in the bottom two panels has $E = 1.05E_{\rm min}$. This allows the second most dense waterbag (dashed line) to intermingle with the densest waterbag (dotted line) giving a substantial contribution to the mean phase-space density at $v = 0$. In contrast, in the top panels, the little energy that the system has above $E_{\rm min}$ permits very little mixing between waterbags, making the mean phase-space density appear very close to a step function (a 'ziggurat').

We now take the time derivative of (4.59):

$$\frac{\mathrm{d}S}{\mathrm{d}t} = \sum_\alpha \frac{V}{\Delta\Gamma_\alpha}\left[\int \mathrm{d}\boldsymbol{v}\left(\frac{\partial f_{0\alpha}}{\partial t}\psi_\alpha + f_{0\alpha}\frac{\partial \psi_\alpha}{\partial t} + \frac{\partial \ln Z_\alpha}{\partial t}\right) + \sum_J \frac{\partial \gamma_{J\alpha}}{\partial t}\frac{n_{J\alpha}}{\eta_{J\alpha}}\right]. \tag{4.60}$$

This can be simplified by taking the time derivative of (4.53) and noting that

$$\begin{aligned}\frac{\partial \ln Z_\alpha}{\partial t} &= -\frac{1}{Z_\alpha}\sum_J e^{-\psi_\alpha\eta_{J\alpha}-\gamma_{J\alpha}}\left(\frac{\partial \psi_\alpha}{\partial t}\eta_{J\alpha} + \frac{\partial \gamma_{J\alpha}}{\partial t}\right)\\ &= -f_{0\alpha}\frac{\partial \psi_\alpha}{\partial t} - \sum_J p_{J\alpha}\frac{\partial \gamma_{J\alpha}}{\partial t}.\end{aligned} \tag{4.61}$$

Substituting this into (4.60) and using (4.49), we get

$$\frac{\mathrm{d}S}{\mathrm{d}t} = \sum_\alpha \frac{V}{\Delta\Gamma_\alpha}\int \mathrm{d}\boldsymbol{v}\,\frac{\partial f_{0\alpha}}{\partial t}\psi_\alpha. \tag{4.62}$$

In this form, we may finally use the collision integral (4.55) and integrate by parts:

$$\begin{aligned}\frac{\mathrm{d}S}{\mathrm{d}t} &= \sum_{\alpha\alpha''}\frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}\sum_{\boldsymbol{k}}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\,\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{k^4|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\frac{\partial f_{0\alpha}}{\partial \psi_\alpha}(\boldsymbol{v})\frac{\partial f_{0\alpha''}}{\partial \psi_{\alpha''}}(\boldsymbol{v}'')\\ &\qquad\left[\frac{\Delta\Gamma_{\alpha''}^2}{m_\alpha^2}\left(\boldsymbol{k}\cdot\frac{\partial \psi_\alpha}{\partial \boldsymbol{v}}\right)^2 - \frac{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}{m_\alpha m_{\alpha''}}\boldsymbol{k}\cdot\frac{\partial \psi_{\alpha''}}{\partial \boldsymbol{v}''}\boldsymbol{k}\cdot\frac{\partial \psi_\alpha}{\partial \boldsymbol{v}}\right]\\ &= \sum_{\alpha\alpha''}\frac{8\pi^3 q_\alpha^2 q_{\alpha''}^2}{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}\sum_{\boldsymbol{k}}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\,\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{k^4|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\frac{\partial f_{0\alpha}}{\partial \psi_\alpha}(\boldsymbol{v})\frac{\partial f_{0\alpha''}}{\partial \psi_{\alpha''}}(\boldsymbol{v}'')\\ &\qquad\left(\frac{\Delta\Gamma_{\alpha''}}{m_\alpha}\boldsymbol{k}\cdot\frac{\partial \psi_\alpha}{\partial \boldsymbol{v}} - \frac{\Delta\Gamma_\alpha}{m_{\alpha''}}\boldsymbol{k}\cdot\frac{\partial \psi_{\alpha''}}{\partial \boldsymbol{v}''}\right)^2 \geq 0.\end{aligned} \tag{4.63}$$

Here, to get to the second equality, we have symmetrised the expression by swapping $\alpha \leftrightarrow \alpha''$ and $\boldsymbol{v} \leftrightarrow \boldsymbol{v}''$. From (4.54), we note that $\partial f_{0\alpha}/\partial\psi_\alpha$ is a negative semi-definite quantity so the integrand in (4.63) is certainly non-negative. This proves that the entropy (4.59) is increased by the collision integral (4.55) unless the integrand is zero, in which case the entropy is conserved. For the latter possibility to be realised, the squared expression in (4.63) must vanish, which it does only for $\psi_\alpha(\boldsymbol{v})$ given by (4.56). Thus, the collision integral (4.55) relaxes the mean phase-space density of the system towards the Lynden-Bell multi-waterbag equilibria (4.57).

4.3.4 Continuous limit of the multi-waterbag formalism

While conceptually enlightening, using a discrete number of waterbag densities is somewhat problematic. At face value, it restricts one to considering initial

conditions that are piecewise-constant functions. In fact, a continuous phase-space density $f_\alpha(\boldsymbol{v})$ can easily be accommodated by making the grid of η arbitrarily fine. We do this by choosing some large number of waterbags densities to fill the interval $[0, \eta_{\alpha\,\max}]$ with spacing $\Delta\eta \to 0$. We can then rewrite what were sums over J as integrals with respect to the waterbag density η:

$$\sum_J \to \frac{1}{\Delta\eta}\int_0^{\eta_{\alpha\,\max}} \mathrm{d}\eta\,. \tag{4.64}$$

Note that the sum above includes the empty waterbag, and so care must be taken to ensure that, when the sums are transformed in this way, the empty waterbag has been included. Now all discrete quantities indexed by J are to be upgraded to continuous functions of η. Namely, we let

$$\gamma_\alpha(\eta) \equiv \gamma_{J\alpha}, \qquad P_{0\alpha}(\boldsymbol{v},\eta) \equiv \frac{p_{J\alpha}(\boldsymbol{v})}{\Delta\eta}, \tag{4.65}$$

the latter function being a probability density with respect to η (hence the normalisation to $\Delta\eta$). Then (4.51) becomes

$$P_{0\alpha}(\boldsymbol{v},\eta) = \frac{1}{Z_\alpha(\psi_\alpha(\boldsymbol{v}))} e^{-\psi_\alpha(\boldsymbol{v})\eta - \gamma_\alpha(\eta)}, \tag{4.66}$$

with the partition function (4.52) redefined as

$$Z_\alpha(\psi_\alpha(\boldsymbol{v})) = \int_0^{\eta_{\alpha\max}} \mathrm{d}\eta\, e^{-\psi_\alpha(\boldsymbol{v})\eta - \gamma_\alpha(\eta)}. \tag{4.67}$$

The Lagrange multiplier $\psi_\alpha(\boldsymbol{v})$ is now determined by the continuous version of (4.46) and (4.53):

$$f_{0\alpha}(\boldsymbol{v}) = \int \mathrm{d}\eta\, \eta P_{0\alpha}(\boldsymbol{v},\eta) = -\frac{\partial \ln Z_\alpha}{\partial \psi_\alpha}. \tag{4.68}$$

To determine $\gamma_\alpha(\eta)$, one must solve the continuous version of the constraint (4.49) fixing the 'waterbag content' of the distribution:

$$\int \mathrm{d}\boldsymbol{v}\, P_{0\alpha}(\boldsymbol{v},\eta) = \frac{n_\alpha(\eta)}{\eta\Delta\eta} \equiv \rho_\alpha(\eta), \tag{4.69}$$

where $n_\alpha(\eta)$ is the continuous generalisation of $n_{J\alpha}$, and the newly defined function $\rho_\alpha(\eta)$ encodes all the information about the species α that must be retained from its initial distribution (an infinite set of Casimir invariants).

Thus, the price of the generalisation to a continuous-waterbag model is having to solve the two coupled integral equations (4.68) and (4.69) for the functions $\psi_\alpha(\boldsymbol{v})$

and $\gamma_\alpha(\eta)$. The collision integral is then again (4.55), its fixed points are (4.68) with $\psi_\alpha(\boldsymbol{v})$ given by (4.56), viz.,

$$f_{0\alpha}(\boldsymbol{v}) = \frac{\int \mathrm{d}\eta\, \eta e^{-\psi_\alpha(\boldsymbol{v})\eta - \gamma_\alpha(\eta)}}{\int \mathrm{d}\eta\, e^{-\psi_\alpha(\boldsymbol{v})\eta - \gamma_\alpha(\eta)}}, \tag{4.70}$$

and the H-theorem (4.63) continues to hold, with the continuous limit of the entropy (4.58):

$$S = -\sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \iint \mathrm{d}\boldsymbol{v}\, \mathrm{d}\eta\, P_{0\alpha}(\boldsymbol{v},\eta) \ln P_{0\alpha}(\boldsymbol{v},\eta). \tag{4.71}$$

4.4 Hyperkinetics

We saw in section (4.3) that, to work out the collisionless evolution of the plasma, we needed to know the probability for the phase-space density $f_\alpha(\boldsymbol{x},\boldsymbol{v})$ to have a certain value η at a velocity $\boldsymbol{v}$. We called that probability $p_{J\alpha}(\boldsymbol{v})$, or, in the continuous limit, $P_{0\alpha}(\boldsymbol{v},\eta)$. This led to a closure problem because these probabilities could not be uniquely determined from $f_{0\alpha}(\boldsymbol{v})$ alone. There was an analogy between the resolution of this closure problem proposed in section 4.3.2 and the resolution of closure problems in fluid theories: it was done by appealing to a local maximisation of entropy—in our case, local in $\boldsymbol{v}$, so the functional form of $P_{0\alpha}(\boldsymbol{v},\eta)$ with respect to η was fixed by (4.66), whereas its dependence on $\boldsymbol{v}$ remained undetermined and encoded by $\psi_\alpha(\boldsymbol{v})$. This fluid analogy is further strengthened by noting that the mean phase-space density was then written in (4.68) as a fluid quantity: the first moment of $P_{0\alpha}(\boldsymbol{v},\eta)$ with respect to η.

All this points to an alternative route to describing collisionless relaxation. The closure problem in fluid theories is resolved by recognising that the system can be described kinetically. Following this logic, instead of considering $f_{0\alpha}(\boldsymbol{v})$ as the core object of our theory, we will consider $P_{0\alpha}(\boldsymbol{v},\eta)$, from which $f_{0\alpha}(\boldsymbol{v})$ can be derived. Thus, we are extending our phase space from 6D, $(\boldsymbol{x},\boldsymbol{v})$ to 7D, $(\boldsymbol{x},\boldsymbol{v},\eta)$, in such a way that the kinetics of $f_{0\alpha}(\boldsymbol{x},\boldsymbol{v})$ will be derivable by taking moments of the new kinetics of $P_\alpha(\boldsymbol{x},\boldsymbol{v},\eta)$ and $P_{0\alpha}(\boldsymbol{v},\eta)$. Like Lynden-Bell's statistical mechanics, this approach, which we shall call 'hyperkinetics', originates from galactic dynamics (as well as fluid mechanics; see Chavanis et al. 1996 and references therein). Its logical conclusion is the hyperkinetic collision integral that we shall now derive—it is a version of the collision integral first derived, in the context of galactic dynamics

and for a model with discrete multiple waterbags, by Severne &Luwel (1980) (who used a somewhat different method).

4.4.1 Hyperkinetic collision integral

To construct the kinetics of the individual waterbags, we first define the 'waterbag distribution function'

$$P_\alpha(\boldsymbol{x},\boldsymbol{v},\eta) = \delta(f_\alpha(\boldsymbol{x},\boldsymbol{v}) - \eta), \tag{4.72}$$

which is the probability density of finding the exact phase-space density $f_\alpha(\boldsymbol{x},\boldsymbol{v})$ to have the value η at the phase-space position $(\boldsymbol{x},\boldsymbol{v})$. The evolution equation for P_α takes the same form as the Vlasov equation:

$$\begin{aligned}\frac{\partial P_\alpha}{\partial t} &= \frac{\partial f_\alpha}{\partial t}\delta'(f_\alpha(\boldsymbol{x},\boldsymbol{v}) - \eta) \\ &= \left(-\boldsymbol{v}\cdot\nabla f_\alpha + \frac{q_\alpha}{m_\alpha}\nabla\varphi\cdot\frac{\partial f_\alpha}{\partial \boldsymbol{v}}\right)\delta'(f_\alpha(\boldsymbol{x},\boldsymbol{v}) - \eta) \\ &= -\boldsymbol{v}\cdot\nabla P_\alpha + \frac{q_\alpha}{m_\alpha}\nabla\varphi\cdot\frac{\partial P_\alpha}{\partial \boldsymbol{v}}.\end{aligned} \tag{4.73}$$

Since we are still concerned with the relaxation to equilibrium, we now wish to find the collision integral for the evolution of P_α in much the same way as we did for f_α in 4.2.

As before, we Fourier-decompose

$$P_\alpha(\boldsymbol{x},\boldsymbol{v},\eta) = P_{0\alpha}(\boldsymbol{v},\eta) + \sum_{\boldsymbol{k}} e^{i\boldsymbol{k}\cdot\boldsymbol{x}} P_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta), \tag{4.74}$$

and expect the ensemble average of the homogeneous part of the hyperkinetic distribution, $P_{0\alpha}$, to be much greater in size than the fluctuating part, sanctioning the quasilinear approach. Therefore, we again linearise our hyperkinetic equation (4.73) to get the fluctuating part of the distribution:

$$\frac{\partial P_{\boldsymbol{k}\alpha}}{\partial t} + i\boldsymbol{k}\cdot\boldsymbol{v}P_{0\alpha} = i\frac{q_\alpha}{m_\alpha}\varphi_{\boldsymbol{k}}\boldsymbol{k}\cdot\frac{\partial P_{0\alpha}}{\partial \boldsymbol{v}}, \tag{4.75}$$

which in turn determines the evolution of the homogeneous part:

$$\frac{\partial P_{0\alpha}}{\partial t} = \frac{\partial}{\partial \boldsymbol{v}}\cdot\left(\frac{q_\alpha}{m_\alpha}\sum_{\boldsymbol{k}}\boldsymbol{k}\mathrm{Im}\langle\varphi^*_{\boldsymbol{k}}P_{\boldsymbol{k}\alpha}\rangle\right). \tag{4.76}$$

The only difference with the calculation in section 4.2 comes from the fact that electric-field perturbations are still made by charge-density perturbations, which are a cumulative effect of all the waterbags. In other words, Poisson's equation (4.6) is

$$\varphi_{\boldsymbol{k}} = \sum_{\alpha} \frac{4\pi q_\alpha}{k^2} \int \mathrm{d}\boldsymbol{v}\, f_{\boldsymbol{k}\alpha}(\boldsymbol{v}) = \sum_{\alpha} \frac{4\pi q_\alpha}{k^2} \iint \mathrm{d}\boldsymbol{v}\, \mathrm{d}\eta\, \eta P_{\boldsymbol{k}\alpha}(\boldsymbol{v}, \eta). \tag{4.77}$$

The derivation of the collision integral for $P_{0\alpha}$ can now be ported over from section 4.2 near verbatim. The only difference is comes from Poisson's equation (4.77): instead of velocity integrals, we now have integrals over both velocities and waterbag densities:

$$\int \mathrm{d}\boldsymbol{v}\, \big(...\big) \to \iint \mathrm{d}\boldsymbol{v}\, \mathrm{d}\eta\, \eta \big(...\big). \tag{4.78}$$

To avoid repetition, we will just note the key equations that are reached throughout the derivation, with reference to the corresponding equations in section 4.2.

The dielectric function, previously (4.10), is now

$$\epsilon_{\boldsymbol{k}}(p) = 1 - i \sum_{\alpha'} \frac{4\pi q_{\alpha'}^2}{m_{\alpha'} k^2} \iint \mathrm{d}\boldsymbol{v}'\, \mathrm{d}\eta'\, \frac{\eta'}{p + i\boldsymbol{k}\cdot\boldsymbol{v}'} \boldsymbol{k} \cdot \frac{\partial P_{0\alpha'}}{\partial \boldsymbol{v}'}. \tag{4.79}$$

The resultant general form of the collision integral, previously (4.18), becomes

$$\frac{\partial P_{0\alpha}}{\partial t} = \frac{\partial}{\partial \boldsymbol{v}} \cdot \sum_{\alpha''} \iint \mathrm{d}\boldsymbol{v}'' \mathrm{d}\eta''\, \eta'' \left[\mathsf{D}^{\alpha}_{\alpha\alpha''}(\boldsymbol{v}, \boldsymbol{v}'', \eta'') \cdot \frac{\partial P_{0\alpha}}{\partial \boldsymbol{v}}\bigg|_{\eta} - \mathsf{D}^{\alpha}_{\alpha''\alpha}(\boldsymbol{v}'', \boldsymbol{v}, \eta) \cdot \frac{\partial P_{0\alpha''}}{\partial \boldsymbol{v}''}\bigg|_{\eta''} \right]. \tag{4.80}$$

The diffusion kernel, via an expression analogous to (4.19), can again be simplified by assuming that the fluctuations evolve much more rapidly than the mean: a calculation identical to that done in section 4.2.3 returns the following analogue of (4.24):

$$\begin{aligned} \mathsf{D}^{\alpha}_{\mu\nu}(\boldsymbol{w}, \boldsymbol{v}, \eta) = \sum_{\nu'} \frac{16\pi^3 q_\mu^2 q_\nu q_{\nu'}}{m_\alpha m_\mu} \\ \mathrm{Re} \sum_{\boldsymbol{k}} \frac{\boldsymbol{k}\boldsymbol{k}}{k^4} \delta(\boldsymbol{k}\cdot(\boldsymbol{w} - \boldsymbol{v})) \iint \mathrm{d}\boldsymbol{v}'\, \mathrm{d}\eta'\, \eta' e^{-i\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}')t} \frac{\langle g_{\boldsymbol{k}\nu}(\boldsymbol{v}, \eta) g^*_{\boldsymbol{k}\nu'}(\boldsymbol{v}', \eta')\rangle}{\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}} \epsilon^*_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}'}}, \end{aligned} \tag{4.81}$$

where $g_{\boldsymbol{k}\nu}(\boldsymbol{v}, \eta) = P_{\boldsymbol{k}\nu}(t = 0, \boldsymbol{v}, \eta)$ is the initial distribution and, as in section 4.2, we use the simplified notation $\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}} \equiv \epsilon_{\boldsymbol{k}}(-i\boldsymbol{k}\cdot\boldsymbol{v})$.

Now we need a closure for the correlation function $\langle g_{\boldsymbol{k}\nu}(\boldsymbol{v}, \eta) g^*_{\boldsymbol{k}\nu'}(\boldsymbol{v}', \eta')\rangle$ in terms of $P_{0\nu}(\boldsymbol{v}, \eta)$. The first step is again the microgranulation ansatz introduced in

section 4.2.5, viz., the assumption that only very near points in phase space are correlated with each other, over a phase-space volume $\Delta\Gamma_\nu$:

$$\langle g_\nu(\boldsymbol{x},\boldsymbol{v},\eta)g_{\nu'}(\boldsymbol{x}',\boldsymbol{v}',\eta')\rangle = \Delta\Gamma_\nu\delta_{\nu\nu'}\langle g_\nu(\eta)g_\nu(\eta')\rangle(\boldsymbol{v})\delta(\boldsymbol{v}-\boldsymbol{v}')\delta(\boldsymbol{x}-\boldsymbol{x}'), \quad (4.82)$$

or, in Fourier space,

$$\langle g_{\boldsymbol{k}\nu}(\boldsymbol{v},\eta)g^*_{\boldsymbol{k}\nu'}(\boldsymbol{v}',\eta')\rangle = \frac{\Delta\Gamma_\nu}{V}\delta_{\nu\nu'}\langle g_\nu(\eta)g_\nu(\eta')\rangle(\boldsymbol{v})\delta(\boldsymbol{v}-\boldsymbol{v}'). \quad (4.83)$$

This is the generalisation of (4.32). Note that, while localising the correlator in the $(\boldsymbol{x},\boldsymbol{v})$ space, we allow for correlations between different values of η. To determine this remaining correlator in terms of $P_{0\nu}$, we use (4.72) to find

$$\begin{aligned}\langle g_\nu(\eta)g_\nu(\eta')\rangle(\boldsymbol{v}) &= \langle[\delta(f_\nu(\boldsymbol{x},\boldsymbol{v})-\eta)-P_{0\nu}(\boldsymbol{v},\eta)][\delta(f_\nu(\boldsymbol{x},\boldsymbol{v})-\eta')-P_{0\nu}(\boldsymbol{v},\eta')]\rangle\\ &= \langle\delta(f_\nu(\boldsymbol{x},\boldsymbol{v})-\eta)\delta(f_\nu(\boldsymbol{x},\boldsymbol{v})-\eta')\rangle - P_{0\nu}(\boldsymbol{v},\eta)P_{0\nu}(\boldsymbol{v},\eta')\\ &= \delta(\eta-\eta')P_{0\nu}(\boldsymbol{v},\eta) - P_{0\nu}(\boldsymbol{v},\eta)P_{0\nu}(\boldsymbol{v},\eta').\end{aligned} \quad (4.84)$$

This formula is a direct generalisation of the single-waterbag closure (4.42), but of course there is no longer any need to assume a single-waterbag distribution. Neither is there any need to assume a local maximisation of entropy, as we did for our multi-waterbag distribution in section 4.3.2: working in the extended phase space $(\boldsymbol{x},\boldsymbol{v},\eta)$ has resolved the closure problem for $\langle f_\alpha^2\rangle$ vs. $f_{0\alpha}$ automatically.

Using (4.83) and (4.84), we can now write the diffusion kernel (4.81) in a closed form:

$$\mathsf{D}^\alpha_{\mu\nu}(\boldsymbol{w},\boldsymbol{v},\eta) = \frac{16\pi^3 q_\mu^2 q_\nu^2\Delta\Gamma_\nu}{m_\alpha m_\mu V}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{w}-\boldsymbol{v}))}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\Big[\eta - f_{0\nu}(\boldsymbol{v})\Big]P_{0\nu}(\boldsymbol{v},\eta), \quad (4.85)$$

where $f_{0\nu}(\boldsymbol{v}) = \int \mathrm{d}\eta'\,\eta' P_{0\nu}(\boldsymbol{v},\eta')$. Finally, substituting (4.85) into (4.80) gives us the generalised hyperkinetic collision integral:

$$\begin{aligned}\frac{\partial P_{0\alpha}}{\partial t} &= \sum_{\alpha''}\frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{m_\alpha V}\frac{\partial}{\partial\boldsymbol{v}}\cdot\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\cdot\int\mathrm{d}\boldsymbol{v}''\,\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\int\mathrm{d}\eta''\,\eta''\\ &\left\{\frac{\Delta\Gamma_{\alpha''}}{m_\alpha}\Big[\eta''-f_{0\alpha''}(\boldsymbol{v}'')\Big]P_{0\alpha''}(\boldsymbol{v}'',\eta'')\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta - \frac{\Delta\Gamma_\alpha}{m_{\alpha''}}\Big[\eta-f_{0\alpha}(\boldsymbol{v})\Big]P_{0\alpha}(\boldsymbol{v},\eta)\left.\frac{\partial P_{0\alpha''}}{\partial\boldsymbol{v}''}\right|_{\eta''}\right\}.\end{aligned} \quad (4.86)$$

To reiterate, the only closure required in the derivation of this collision integral was the microgranulation ansatz (4.82). No need to break higher-order correlators has

arisen, because, within the hyperkinetic formalism, all moments of the phase-space density can be derived from the distribution function $P_{0\alpha}$: $\langle f_\alpha^n \rangle = \int \mathrm{d}\eta\, \eta^n P_{0\alpha}(\boldsymbol{v}, \eta)$. Indeed, if one takes the first moment of (4.86) with respect to η, it is easy to show that one recovers (4.35), which was the general collision integral before a choice of waterbag closure was made in section 4.3.

As well as conserving total energy, momentum, and particle number, (4.86) has two further invariants, which represent the conservation of phase volume,

$$\int \mathrm{d}\boldsymbol{v}\, P_{0\alpha}(\boldsymbol{v}, \eta) = \rho_\alpha(\eta), \tag{4.87}$$

and the conservation of probability

$$\int \mathrm{d}\eta\, P_{0\alpha}(\boldsymbol{v}, \eta) = 1. \tag{4.88}$$

As in (4.69), (4.87) distils to a single function $\rho_\alpha(\eta)$ (the 'waterbag content' of the distribution, or the infinite set of its Casimir invariants) all the information from the initial condition that must be preserved by collisionless evolution. Given these invariants, it is unsurprising that the fixed points of the collision integral (4.86) will again be the Lynden-Bell equilibria, viz.,

$$P_{0\alpha}(\boldsymbol{v}, \eta) = \frac{e^{-\beta \Delta\Gamma_\alpha \eta [\epsilon_\alpha - \mu_\alpha(\eta)]}}{\int \mathrm{d}\eta\, e^{-\beta \Delta\Gamma_\alpha \eta [\epsilon_\alpha - \mu_\alpha(\eta)]}}, \tag{4.89}$$

where $\epsilon_\alpha = m_\alpha |\boldsymbol{v}|^2/2$, and the normalisation (4.88) has been enforced. The parameters β and $\mu(\eta)$ are determined from (4.87) and the energy conservation,

$$V \iint \mathrm{d}\boldsymbol{v}\, \mathrm{d}\eta\, \eta \epsilon_\alpha(\boldsymbol{v}) P_{0\alpha}(\boldsymbol{v}, \eta) = E. \tag{4.90}$$

That (4.89) are indeed fixed points of (4.86) can be confirmed by direct substitution. Just as in section 4.3.3, we must provide an H-theorem to prove their stability.

4.4.2 Hyperkinetic H-theorem

We define the Shannon entropy as before but now upgraded to its continuous variant (4.71). Taking the time derivative of S using (4.86) and integrating by

parts in $\boldsymbol{v}$, we get

$$\begin{aligned}\frac{\mathrm{d}S}{\mathrm{d}t} &= -\sum_{\alpha}\frac{V}{\Delta\Gamma_\alpha}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\eta\left(1+\ln P_{0\alpha}\right)\frac{\partial P_{0\alpha}}{\partial t}\\ &=\sum_{\alpha\alpha''}\frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\sum_{\boldsymbol{k}}\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{k^4|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\\ &\quad\iint \mathrm{d}\eta\,\mathrm{d}\eta''\,\eta''\left\{\frac{\Delta\Gamma_{\alpha''}^2}{m_\alpha^2}\Big[\eta''-f_{0\alpha''}(\boldsymbol{v}'')\Big]\frac{P_{0\alpha''}(\boldsymbol{v}'',\eta'')}{P_{0\alpha}(\boldsymbol{v},\eta)}\left(\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta\right)^2\right.\\ &\quad\left.-\frac{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}{m_\alpha m_{\alpha''}}\Big[\eta-f_{0\alpha}(\boldsymbol{v})\Big]\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}''}\right|_{\eta''}\right\}.\end{aligned}\tag{4.91}$$

First, we notice that the η'' integral of the first bracketed term is

$$\int \mathrm{d}\eta''\,\eta''\Big[\eta''-f_{0\alpha''}(\boldsymbol{v}'')\Big]P_{0\alpha''}(\boldsymbol{v}'',\eta'')=\int \mathrm{d}\eta''\,\Big[\eta''-f_{0\alpha''}(\boldsymbol{v}'')\Big]^2 P_{0\alpha''}(\boldsymbol{v}'',\eta'').\tag{4.92}$$

Secondly, in the second bracketed term, anything that is not multiplied by both η and η'' integrates to zero by (4.88):

$$\int \mathrm{d}\eta\,\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta=\frac{\partial}{\partial\boldsymbol{v}}\int \mathrm{d}\eta\,P_{0\alpha}(\boldsymbol{v},\eta)=0.\tag{4.93}$$

With these insights, the rate of change of entropy becomes

$$\begin{aligned}\frac{\mathrm{d}S}{\mathrm{d}t} &=\sum_{\alpha\alpha''}\frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\sum_{\boldsymbol{k}}\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{k^4|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\\ &\quad\iint \mathrm{d}\eta\,\mathrm{d}\eta''\left\{\frac{\Delta\Gamma_{\alpha''}^2}{m_\alpha^2}\Big[\eta''-f_{0\alpha''}(\boldsymbol{v}'')\Big]^2\frac{P_{0\alpha''}(\boldsymbol{v}'',\eta'')}{P_{0\alpha}(\boldsymbol{v},\eta)}\left(\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta\right)^2\right.\\ &\quad\left.-\frac{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}{m_\alpha m_{\alpha''}}\eta\eta''\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha''}}{\partial\boldsymbol{v}''}\right|_{\eta''}\right\}.\end{aligned}\tag{4.94}$$

Finally, we symmetrise the entire expression by swapping $\alpha\leftrightarrow\alpha''$, $\eta\leftrightarrow\eta''$, $\boldsymbol{v}\leftrightarrow\boldsymbol{v}''$, which allows us to write (4.94) in an explicitly positive-semidefinite form:

$$\begin{aligned}\frac{\mathrm{d}S}{\mathrm{d}t} &=\sum_{\alpha\alpha''}\frac{8\pi^3 q_\alpha^2 q_{\alpha''}^2}{\Delta\Gamma_\alpha\Delta\Gamma_{\alpha''}}\iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\boldsymbol{v}''\sum_{\boldsymbol{k}}\frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{k^4|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\\ &\quad\iint \mathrm{d}\eta\,\mathrm{d}\eta''\left\{\frac{\Delta\Gamma_{\alpha''}}{m_\alpha}\Big[\eta''-f_{0\alpha''}(\boldsymbol{v}'')\Big]\sqrt{\frac{P_{0\alpha''}(\boldsymbol{v}'',\eta'')}{P_{0\alpha}(\boldsymbol{v},\eta)}}\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial\boldsymbol{v}}\right|_\eta\right.\\ &\quad\left.-\frac{\Delta\Gamma_\alpha}{m_{\alpha''}}\Big[\eta-f_{0\alpha}(\boldsymbol{v})\Big]\sqrt{\frac{P_{0\alpha}(\boldsymbol{v},\eta)}{P_{0\alpha''}(\boldsymbol{v}'',\eta'')}}\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha''}}{\partial\boldsymbol{v}''}\right|_{\eta''}\right\}^2\geq 0.\end{aligned}\tag{4.95}$$

Expanding the squared expression in (4.95) does indeed recover (4.94) as all excess terms cancel. The collision integral (4.86) therefore never decreases the Shannon entropy (4.71).

To prove that all initial conditions will eventually reach their corresponding Lynden-Bell equilibrium (4.89), we must show that the rate of entropy growth (4.95) will equal zero if and only if the Lynden-Bell equilibrium has been reached. Owing to the squared expression in (4.95), the entropy growth will equal zero when

$$\frac{1}{\Delta\Gamma_\alpha m_\alpha\Big[\eta - f_{0\alpha}(\boldsymbol{v})\Big]P_{0\alpha}(\boldsymbol{v},\eta)}\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial \boldsymbol{v}}\right|_\eta = \frac{1}{\Delta\Gamma_{\alpha''} m_{\alpha''}\Big[\eta'' - f_{0\alpha''}(\boldsymbol{v}'')\Big]P_{0\alpha''}(\boldsymbol{v}'',\eta'')}\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha''}}{\partial \boldsymbol{v}''}\right|_{\eta''}, \quad (4.96)$$

for all α, α'', η and η'' when $\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}'')=0$. Save for $\boldsymbol{v}$ and $\boldsymbol{v}''$, the left- and right-hand sides of (4.96) are functions of distinct, independent variables. The only solution to (4.96) must, therefore, satisfy

$$\boldsymbol{k}\cdot\left.\frac{\partial P_{0\alpha}}{\partial \boldsymbol{v}}\right|_\eta = -\beta\Delta\Gamma_\alpha m_\alpha\Big[\eta - f_{0\alpha}(\boldsymbol{v})\Big]P_{0\alpha}(\boldsymbol{v},\eta)\boldsymbol{k}\cdot\boldsymbol{v}, \quad (4.97)$$

where β is an as yet undetermined constant, although it is clear it will come to mean the thermodynamic beta shortly. Note that by writing $\boldsymbol{k}\cdot\boldsymbol{v}$ we have assumed, without loss of generality, that we are in the zero-net-momentum frame. It is clear that the Lynden-Bell equilibria (4.89) satisfy the condition (4.97) (as they must, being maximisers of the entropy (4.71)).

To prove that these are the only solutions, we rewrite (4.97) as

$$\left.\frac{\partial \ln P_{0\alpha}}{\partial \boldsymbol{v}}\right|_\eta = -\beta\Delta\Gamma_\alpha\eta m_\alpha\boldsymbol{v} + \beta m_\alpha\Delta\Gamma_\alpha\int \mathrm{d}\eta\,\eta P_{0\alpha}(\boldsymbol{v},\eta)\boldsymbol{v}. \quad (4.98)$$

Then, without loss of generality,

$$P_{0\alpha}(\boldsymbol{v},\eta) = \frac{C_\alpha(\boldsymbol{v},\eta)}{\int \mathrm{d}\eta\, C_\alpha(\boldsymbol{v},\eta)}, \quad (4.99)$$

for which (4.98) becomes

$$\left.\frac{\partial \ln C_\alpha}{\partial \boldsymbol{v}}\right|_{\boldsymbol{v}} + \beta\Delta\Gamma_\alpha\eta m_\alpha\boldsymbol{v} = \beta\Delta\Gamma_\alpha m_\alpha\frac{\int \mathrm{d}\eta\,\eta C_\alpha(\boldsymbol{v},\eta)}{\int \mathrm{d}\eta\, C_\alpha(\boldsymbol{v},\eta)}\boldsymbol{v} + \frac{\partial}{\partial \boldsymbol{v}}\ln\int \mathrm{d}\eta\, C_\alpha(\boldsymbol{v},\eta)\boldsymbol{v}$$
$$= \frac{\partial}{\partial \boldsymbol{v}}\Phi_\alpha(\boldsymbol{v}). \quad (4.100)$$

In (4.100), we have collected all η-dependent terms on the right-hand side and further declared that, since the curl of the left-hand is zero, these terms can be written as the gradient of some function $\Phi_\alpha(\boldsymbol{v})$, which will be determined self-consistently after determining $C_\alpha(\boldsymbol{v})$. With this sleight of hand, (4.100) is solved by

$$C_\alpha(\boldsymbol{v},\eta) = e^{\Phi_\alpha(\boldsymbol{v})} e^{-\beta \Delta\Gamma_\alpha \eta[\epsilon_\alpha(\boldsymbol{v}) - \mu_\alpha(\eta)]}, \tag{4.101}$$

where $\epsilon_\alpha(\boldsymbol{v}) = m_\alpha|\boldsymbol{v}|^2/2$ and the chemical potential $\mu_\alpha(\eta)$ has emerged as an integration constant. Now, substituting this into (4.99), we find that $P_{0\alpha}(\boldsymbol{v},\eta)$ is the Lynden-Bell equilibrium (4.89), proving that it is the only solution for which the entropy does not grow. Note that the determination of $\Phi_\alpha(\boldsymbol{v})$ is unimportant for the calculation of $P_{0\alpha}(\boldsymbol{v},\eta)$ because $C_\alpha(\boldsymbol{v})$ can be freely multiplied by any function of $\boldsymbol{v}$ without changing $P_{0\alpha}(\boldsymbol{v},\eta)$; also, given (4.101), the solvability condition

$$\frac{\partial}{\partial \boldsymbol{v}} \Phi_\alpha(\boldsymbol{v}) = \beta \Delta\Gamma_\alpha m_\alpha \frac{\int \mathrm{d}\eta\, \eta C_\alpha(\boldsymbol{v},\eta)}{\int \mathrm{d}\eta\, C_\alpha(\boldsymbol{v},\eta)} \boldsymbol{v} + \frac{\partial}{\partial \boldsymbol{v}} \ln \int \mathrm{d}\eta\, C_\alpha(\boldsymbol{v},\eta) \boldsymbol{v} \tag{4.102}$$

is satisfied for all functions $\Phi_\alpha(\boldsymbol{v})$. This completes the proof that all initial conditions will reach their Lynden-Bell equilibria.

4.5 Collisionless vs. collisional relaxation

Thus far we have not discussed the absence of true collisions within this formalism. In this context, 'true collisions' are any type of relaxation to equilibrium that does not conserve phase volume and thus releases the stranglehold that the invariants (1.3) had on the evolution. It is for this reason that Lynden-Bell (1967) originally proposed the idea of a violent relaxation: so that steady states could be reached long before the conservation of phase volume was broken (although we now know this to be false from chapter 3). The collision integrals derived above do not describe such a violent, highly nonlinear, regime but rather a quasilinear one, where the mean distribution function evolves slowly compared to the fluctuations. Nevertheless, for our collision integrals to be valid, the rate of relaxation due to them must be much greater than the rate of relaxation due to true collisions. We shall estimate the collisionless relaxation rate in a moment, but first let us show how to recover the collision integral of Balescu (1960) and Lenard (1960) from the Kadomtsev–Pogutse collision integral (4.43), in order to have a 'true collisionality' with which to compare our 'collisionless collision rate'.

4.5.1 Balescu–Lenard collision integral

In reality, a plasma is not a phase fluid but a collection of N particles. The true, exact phase-space density of these particles is the Klimontovich distribution (see, e.g., Klimontovich 1967):

$$f(\boldsymbol{x},\boldsymbol{v}) = \sum_{i=1}^{N} \delta(\boldsymbol{x}-\boldsymbol{x}_i)\delta(\boldsymbol{v}-\boldsymbol{v}_i), \tag{4.103}$$

where $\boldsymbol{x}_i$ and $\boldsymbol{v}_i$ are the particles' instantaneous positions and velocities, respectively. Since each particle thus occupies precisely zero phase volume, it might seem as though phase-volume conservation were a meaningless idea. The reason one can talk about phase-volume conservation at all is that one assumes that particles that are neighbours in phase space move in a similar way, implying that replacing the Klimontovich distribution with a smoothed phase-space density is a reasonable approximation. It is then the phase volumes associated with this smoothed function that are conserved. This is closely related to the microgranulation ansatz (4.32), which posits that, within a phase volume $\Delta\Gamma_\alpha$, the fluctuations of the phase-space density are correlated, implying that particles are moving collectively. A 'true collision' occurs when a single particle is not correlated at all with its neighbouring particles. This can be accommodated within the microgranulation ansatz by assuming that each particle is its own waterbag. Mathematically this is just the single-waterbag model of section 4.3.1 in the limit where the correlation volume $\Delta\Gamma_\alpha$ is made smaller that the interparticle separation in phase space. Then, by assumption there is only one particle in the correlation volume, so

$$\Delta\Gamma_\alpha \eta_\alpha = 1. \tag{4.104}$$

The smoothed distribution function is given by the average occurrence of single particles, hence $f_{0\alpha}(\boldsymbol{v}) \ll \eta_\alpha$. Under these assumptions, the Kadomtsev–Pogutse collision integral (4.43) becomes the Balescu–Lenard collision integral:

$$\frac{\partial f_{0\alpha}}{\partial t} = \sum_{\alpha''} \frac{16\pi^3 q_\alpha^2 q_{\alpha''}^2}{m_\alpha V} \frac{\partial}{\partial \boldsymbol{v}} \cdot \sum_{\boldsymbol{k}} \frac{\boldsymbol{k}\boldsymbol{k}}{k^4} \cdot \int \mathrm{d}\boldsymbol{v}'' \, \frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}''))}{\left|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}\right|^2} \left[\frac{1}{m_\alpha} f_{0\alpha''}(\boldsymbol{v}'') \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \frac{1}{m_{\alpha''}} f_{0\alpha}(\boldsymbol{v}) \frac{\partial f_{0\alpha''}}{\partial \boldsymbol{v}''}\right]. \tag{4.105}$$

This reduction allows one to make a useful comparison between the 'true' collision rate, which is the typical collision rate associated with the Balescu–Lenard integral (4.105), and the effective collision rate associated with the hyperkinetic 'collisionless collision integral' (4.86). Before proceeding to do that, we observe in passing that the formalism developed in sections 4.2 and 4.3 also allows one to derive very efficiently true collision operators for quantum plasmas consisting of fermionic and bosonic species. This is done in appendix F.

4.5.2 Effective collision rates

To keep this discussion as transparent as possible, let us consider only the like-particle effective collision rates. Consider the generic quasilinear collision integral (4.35), with $\alpha'' = \alpha$ and

$$\langle g_\alpha^2 \rangle(\boldsymbol{v}) = \int \mathrm{d}\eta \, [\eta - f_{0\alpha}(\boldsymbol{v})]^2 \, P_{0\alpha}(\boldsymbol{v}, \eta), \qquad f_{0\alpha}(\boldsymbol{v}) = \int \mathrm{d}\eta \, \eta P_{0\alpha}(\boldsymbol{v}, \eta), \tag{4.106}$$

where $P_{0\alpha}(\boldsymbol{v}, \eta)$ is evolved by (4.86). This integral should be compared to its Balescu–Lenard counterpart (4.105), also with $\alpha'' = \alpha$. It is immediately apparent that the difference in the rates of the true and effective collisions is

$$\frac{\nu_{\alpha\alpha}^{\mathrm{eff}}}{\nu_{\alpha\alpha}^{\mathrm{true}}} \sim \frac{\Delta\Gamma_\alpha \langle g_\alpha^2 \rangle}{f_{0\alpha}} \equiv \Delta\Gamma_\alpha \eta_\alpha^{\mathrm{eff}}, \tag{4.107}$$

where $\eta_\alpha^{\mathrm{eff}}$ is the typical deviation of the phase-space density from its mean. The quantity on the right-hand side of (4.107) is the typical variation in the number of particles contained in the correlation volume $\Delta\Gamma_\alpha$. Since this collisionless theory is built upon the assumption that the correlation volume is sufficiently large for its mean phase-space density η to be meaningfully specified, we have inherently assumed that the number of particles contained in a correlation volume is large. However, this does not tell us immediately about the typical deviation of this number from its mean. We therefore aim to compare $\eta_\alpha^{\mathrm{eff}}$ to the typical value of the phase-space density, $n_\alpha / v_{\mathrm{th}\alpha}^3$. To make this comparison quantitative, we estimate (4.107) close to a Lynden-Bell equilibrium. In this case, using (4.54) with $\psi_\alpha = \beta \Delta\Gamma_\alpha \epsilon_\alpha$, we find

$$\Delta\Gamma_\alpha \eta_\alpha^{\mathrm{eff}} = -\frac{1}{\beta} \frac{\partial \ln f_{0\alpha}}{\partial \epsilon_\alpha} \sim \frac{1}{\beta m_\alpha v_{\mathrm{th}\alpha}^2}, \tag{4.108}$$

where $v_{\mathrm{th}\alpha}$ is the typical velocity scale of $f_{0\alpha}$. Since β must be determined from the constraints of energy conservation (4.29) and phase-volume conservation (4.69), (4.108) provides an implicit expression for the $\eta_\alpha^{\mathrm{eff}}$ in terms of E and $\rho_\alpha(\eta)$. For the purposes of order-of-magnitude estimates, the exact calculation of β is unnecessary and it will suffice to consider two relevant limits. In the degenerate case, when the initial condition is very close to the minimum-energy state (the Gardner state; see section 4.3.3), β will be very large and can be computed by a Sommerfeld-like expansion. In this case, naively ordering $\eta_\alpha^{\mathrm{eff}} \sim n_\alpha / v_{\mathrm{th}\alpha}^3$ would be an overestimation. This is because deviations from the mean phase-space density require the system to have sufficient energy to allow two species of waterbag to intermingle (see discussion around 4.3.3). Thankfully, as discussed in section 4.3.3, the system does not need to be energetically very far from its Gardner distribution to become, at least partially, non-degenerate. In the non-degenerate limit, which is by far the most common, waterbags of different phase-space density can freely intermingle, and so we can estimate $\eta_\alpha^{\mathrm{eff}} \sim n_\alpha / v_{\mathrm{th}\alpha}^3$, implying $\Delta\Gamma_\alpha \eta_\alpha^{\mathrm{eff}}$ is on the order of the number of particles in a correlation volume, which we have assumed to be large.

Thus, the effective collision rate is generally much larger than the true-collision rate. Note however, that this fast collisionless relaxation must still be slower that the evolution of the fluctuations. The most straightforward estimate of the typical rate of the latter is the plasma frequency, $\omega_{\mathrm{pe}} = (4\pi e^2 n_{\mathrm{e}}/m_{\mathrm{e}})^{1/2}$ (specialising to electrons for the purposes of this estimate). Thus, we require (and expect) the ordering

$$\nu_{\mathrm{ee}}^{\mathrm{true}} \sim \frac{\omega_{\mathrm{pe}}}{n_{\mathrm{e}} \lambda_{\mathrm{De}}^3} \ll \nu_{\mathrm{ee}}^{\mathrm{eff}} \ll \omega_{\mathrm{pe}}, \tag{4.109}$$

where $\lambda_{\mathrm{De}} \sim \omega_{\mathrm{pe}}^{-1} v_{\mathrm{the}}$ is the Debye length. This places a constraint on the correlation volume:

$$1 \ll \Delta\Gamma_{\mathrm{e}} \eta_{\mathrm{e}}^{\mathrm{eff}} \ll n_{\mathrm{e}} \lambda_{\mathrm{De}}^3. \tag{4.110}$$

If the plasma parameter $n_{\mathrm{e}} \lambda_{\mathrm{De}}^3$ is large (i.e., if the ideal-gas approximation applies), this constraint is not very stringent and still allows the effective collision frequency to be much larger than $\nu_{\mathrm{ee}}^{\mathrm{true}}$.

A more stringent constraint emerges if one works out the time τ_{c} that it takes for exact phase-volume conservation to be broken. This time-scale is substantially less well understood as it depends on the exact rate at which the fluctuating part δf_α

of the exact phase-space density is altered irrevocably by collisions. A typical estimate (see, e.g., Su &Oberman 1968) is[3]

$$\tau_{\mathrm{c}}^{-1} \sim (\nu_{ee}^{\mathrm{true}})^{1/3}\omega_{\mathrm{pe}}^{2/3} \sim \frac{\omega_{\mathrm{pe}}}{(n_{\mathrm{e}}\lambda_{\mathrm{De}}^3)^{1/3}}. \tag{4.111}$$

We may then argue that, for the collisionless relaxation to the Lynden-Bell equilibria to be of any importance it must happen long before phase-volume conservation is broken: (4.110) is then revised to

$$(n_{\mathrm{e}}\lambda_{\mathrm{De}}^3)^{2/3} \ll \Delta\Gamma_{\mathrm{e}}\eta_{\mathrm{e}}^{\mathrm{eff}} \ll n_{\mathrm{e}}\lambda_{\mathrm{De}}^3. \tag{4.112}$$

Depending on just how large the plasma parameter is, this could be a more difficult ordering to satisfy. However, if satisfied, it would make the collisionless relaxation rate far larger than the rate of relaxation due to 'true' collisions. To know just how much larger, we must be able to calculate $\Delta\Gamma_\alpha$ independently. Obviously this, along with a quantitative assessment of the validity of the microgranulation ansatz, requires a full theory of the two-point correlation function of the phase-space density. Without such a theory, certain revealing estimates can, nevertheless, be made.

4.5.3 Energy of fluctuations and the correlation volume

In all the above, the major shortcoming of the theory is the lack of clarity about the size of $\Delta\Gamma_\alpha$. In order to relate $\Delta\Gamma_\alpha$ to something measurable within a plasma, let us calculate the energy E_φ stored in the electric-field fluctuations under the microgranulation ansatz (4.83). This energy is given by

$$E_\varphi = \int \mathrm{d}\boldsymbol{x}\, \frac{|\nabla\varphi|^2}{8\pi} = \frac{1}{2}V\sum_\alpha q_\alpha \sum_{\boldsymbol{k}} \iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\eta\,\eta \mathrm{Re}\langle\varphi_{\boldsymbol{k}}^* P_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta)\rangle. \tag{4.113}$$

where we have integrated by parts and used Poisson's equation (4.77) in the second equality. Only the real part has survived because the summation in $\boldsymbol{k}$ is even. While one could again apply Poisson's equation to the remaining $\varphi_{\boldsymbol{k}}$ in (4.113) and then naively utilise the microgranulation ansatz (4.83), it is worth noting that the correlator in (4.113) is the real part of the same correlator the imaginary part of which appeared in (4.76). The real part of this correlator is equivalent to

[3]Here, we do not use the nonlinear estimate for the time for $\rho(\eta)$ conservation to break (Nastac et al., 2024a) on the grounds that the calculation here is quasilinear. Whether such a quasilinear relaxation can ever be achieved has not been shown.

the hyperkinetic generalisation of the discarded term of (4.16). Recovering the hyperkinetic generalisation of this term, from (4.16) via (4.37), we get

$$\mathrm{Re}\langle\varphi^*_{\boldsymbol{k}}P_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta)\rangle = \mathrm{Re}\sum_\alpha \frac{4\pi q_{\alpha'}}{k^2}\iint\frac{\mathrm{d}p\,\mathrm{d}p'}{(2\pi)^2}\frac{e^{(p+p')t}}{\epsilon_{\boldsymbol{k}}(p)\epsilon^*_{\boldsymbol{k}}(p'^*)} \times\iint \mathrm{d}\boldsymbol{v}'\,\mathrm{d}\eta'\,\eta'\frac{\langle g_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta)g^*_{\boldsymbol{k}\alpha'}(\boldsymbol{v}',\eta')\rangle}{(p+i\boldsymbol{k}\cdot\boldsymbol{v})(p'-i\boldsymbol{k}\cdot\boldsymbol{v}')}, \quad (4.114)$$

with the understanding that, after the microgranulation ansatz, the real part of the other contributions will vanish. Carrying out the p and p' contour integration (see section 4.2.3 for details) and applying the microgranulation ansatz (4.83), we find

$$\mathrm{Re}\langle\varphi^*_{\boldsymbol{k}}P_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta)\rangle = \frac{4\pi q_\alpha \Delta\Gamma_\alpha}{k^2V|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\int\mathrm{d}\eta'\,\eta'\langle g_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta)g^*_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta')\rangle. \quad (4.115)$$

Substituting (4.115) into (4.113) and using the formula (4.84) for the correlation function of $g_{\boldsymbol{k}\alpha}(\boldsymbol{v},\eta)$, we get

$$\begin{aligned}E_\varphi &= \sum_\alpha\sum_{\boldsymbol{k}}\frac{2\pi q_\alpha^2\Delta\Gamma_\alpha}{k^2}\int\mathrm{d}\boldsymbol{v}\,\frac{1}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\iint\mathrm{d}\eta\,\mathrm{d}\eta'\,\eta\eta'\Big[\delta(\eta-\eta')P_{0\alpha}(\boldsymbol{v},\eta)\\ &\qquad - P_{0\alpha}(\boldsymbol{v},\eta)P_{0\alpha}(\boldsymbol{v},\eta')\Big]\\ &= \sum_\alpha\sum_{\boldsymbol{k}}\frac{2\pi q_\alpha^2\Delta\Gamma_\alpha}{k^2}\int\mathrm{d}\boldsymbol{v}\,\frac{\langle g_\alpha^2\rangle(\boldsymbol{v})}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}, \end{aligned}\quad (4.116)$$

where $\langle g_\alpha^2\rangle(\boldsymbol{v})$ is given by (4.115). This formula describes the energy in the fluctuating electric field for a plasma obeying the microgranulation ansatz.

In a Lynden-Bell equilibrium, (4.116) can be further simplified by calculating $\Delta\Gamma_\alpha\langle g_\alpha^2\rangle$ via (4.107) and (4.108). Approximating also $|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2\approx 1$, we have

$$E_\varphi \approx \sum_\alpha\sum_{\boldsymbol{k}}\frac{2\pi q_\alpha^2}{k^2}\int\mathrm{d}\boldsymbol{v}\,\left(-\frac{1}{\beta}\frac{\partial f_{0\alpha}}{\partial\epsilon_\alpha}\right) = \sum_\alpha\frac{4Vq_\alpha^2k_{\max}}{\beta m_\alpha}\int_0^\infty\mathrm{d}v\,f_{0\alpha}(v), \quad (4.117)$$

where $k_{\max}$ is the UV cutoff for the wave-number integal (in 3D). Using (4.108) and making a rough estimate of everything, we have

$$E_\varphi \sim \sum_\alpha \Delta\Gamma_\alpha\eta_\alpha^{\mathrm{eff}}q_\alpha^2n_\alpha k_{\max}V. \quad (4.118)$$

Let us compare these estimates with the energy stored in electric fluctuations associated with true, Coulomb collisions (i.e., with discrete particle noise). The latter can be calculated in a manner analogous to the above but using the Balescu–Lenard collision operator (4.105). To do this, we set

$$P_{0\alpha}(\boldsymbol{v},\eta) = \left[1-\frac{f_{0\alpha}(\boldsymbol{v})}{\eta_\alpha}\right]\delta(\eta)+\frac{f_{0\alpha}(\boldsymbol{v})}{\eta_\alpha}\delta(\eta-\eta_\alpha), \quad (4.119)$$

where $\eta_\alpha = \Delta\Gamma_\alpha^{-1}$ according to (4.104). Inserting (4.119) into (4.116) and collecting only the lowest-order terms in the limit of $f_{0\alpha}/\eta_\alpha \to 0$ gives

$$E_{\varphi,\mathrm{BL}} = \sum_\alpha \sum_{\boldsymbol{k}} \frac{2\pi q_\alpha^2}{k^2} \int \mathrm{d}\boldsymbol{v}\, \frac{f_{0\alpha}(\boldsymbol{v})}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2} \approx \sum_\alpha 4V q_\alpha^2 k_{\max} n_\alpha, \tag{4.120}$$

where the last estimate has been obtained in the same manner as (4.117). From a direct comparison of (4.116) and (4.118) with (4.120), it is clear that, both $\boldsymbol{k}$ by $\boldsymbol{k}$ and overall, a plasma obeying the microgranulation ansatz has more energy stored in the electric field than a collisional plasma, by a factor of $\Delta\Gamma_\alpha \eta_\alpha^{\mathrm{eff}} \gg 1$, the typical number of particles in a correlation volume.[4]

Note that, despite the fluctuation energy being large compared to particle noise, it will still be small compared to the kinetic energy of the particles

$$K = V \sum_\alpha \int \mathrm{d}\boldsymbol{v}\, \frac{m_\alpha |\boldsymbol{v}|^2}{2} f_{0\alpha} \sim V \sum_\alpha m_\alpha v_{\mathrm{th}\alpha}^2 n_\alpha. \tag{4.121}$$

Therefore, using (4.118), we find

$$\frac{E_\varphi}{K} \sim \frac{\sum_\alpha \Delta\Gamma_\alpha^{\mathrm{eff}} \eta_\alpha^{\mathrm{eff}} q_\alpha^2 n_\alpha k_{\max}}{\sum_\alpha m_\alpha v_{\mathrm{th}\alpha}^2 n_\alpha} \sim \frac{\Delta\Gamma_{\mathrm{e}} \eta_{\mathrm{e}}^{\mathrm{eff}}}{n_{\mathrm{e}} \lambda_{\mathrm{De}}^3} k_{\max} \lambda_{\mathrm{De}} \ll 1, \tag{4.122}$$

ignoring at the last step any potential disparities between contributions from different species. Taking cue from our discussion in section 4.5.2, we conclude that this ratio will be small in all conceivable cases of interest (essentially by construction of this quasilinear theory).

Since the distributions evolved by the quasilinear collision integrals are assumed to be linearly stable (see section 4.2.3), their kinetic energy does not change [see (4.30)] and, therefore, neither can the fluctuation energy E_φ change. It is then useful to think of E_φ as a feature of the system that tells us about the size of the correlation volume $\Delta\Gamma_\alpha$ and of the applicability of the collisionless relaxation, via the estimate (4.118).

This calculation does not determine $\Delta\Gamma_\alpha$, however, together with the discussion of the collision time scales, it sets feasible limits on what the allowed values of $\Delta\Gamma_\alpha$ can be. Crucially, we see that a larger value of $\Delta\Gamma_\alpha$ corresponds to a larger energy in

[4] As with the estimates of the collision time scale, we do not expect this factor to be large for distributions sufficiently close to their ground states (Gardner distributions) when $\eta_\alpha^{\mathrm{eff}} \ll n_\alpha / v_{\mathrm{th}\alpha}^3$, but such systems will barely evolve anyway.

the fluctuating electric field. This further suggests that the validity of the Balescu–Lenard collision integral should be called into question when the fluctuations in the electric field are anomalously large compared to (4.120). The Balescu–Lenard collision integral is designed to describe slow relaxation mediated by discrete particle noise. The hyperkinetic collision integral, on the other hand, describes a faster relaxation mediated by correlated volumes of phase space.

4.5.4 Caveats on the existence of a Lynden-Bell plasma

Despite promising universality, previous attempts at the numerical demonstration of systems reaching their Lynden-Bell equilibria have not been met with universal success. Instead, only certain initial conditions will cause a system to reach its Lynden-Bell equilibrium (see, e.g., Arad &Johansson 2005; Levin et al. 2014 and references therein), which naturally forces one to consider the validity of this theory, in particular, its chief assumption: the microgranulation ansatz. From sections 4.5.1-4.5.3, we are now equipped with a better understanding of how the microgranulation ansatz has affected not only the collision integral, but also the fluctuations of the electric field. It is therefore now prudent, in spite of the comfort provided by the estimates (4.112) and (4.118), to question what has been lost and, therefore, what the validity of this closure is.

Should phase-volume conservation survive, we must further ask if it is reasonable to expect that the correlation volume $\Delta\Gamma_\alpha$ should be a constant independent of time. Naturally one might assume that, as fluctuations travel through phase space, they gradually break up and form smaller and smaller structures. Another factor that indicates that $\Delta\Gamma_\alpha$ should be dependent on time is (4.116) itself. Since the kinetic energy of the mean distribution function is also constant, then to conserve the total energy, the electric-field energy must also be constant. By (4.116), this implies that $\Delta\Gamma_\alpha$ must vary in concert with the integral of $\langle g_\alpha^2 \rangle$ over velocities. Alternatively one could take the view that the mean energy of the distribution function is only approximately constant, justified by the smallness of the electric-field energy compared to the kinetic energy as computed in (4.122).

Only one of our results requires that $\Delta\Gamma_\alpha$ not be a function of time: the H-theorem. If $\Delta\Gamma_\alpha$ evolves with time, then the H-theorem is broken because our entropy (4.71) has a prefactor proportional to $\Delta\Gamma_\alpha$. This could be remedied if $\Delta\Gamma_\alpha$

were independent of the species, because it would be an overall multiplicative prefactor that the entropy need not include. This would result in a working, but weaker, H-theorem under which any state could be a steady state if $\Delta\Gamma_\alpha$ decayed sufficiently fast to halt its evolution. Such relaxation has been called 'incomplete violent relaxation', indicating that the system tried to reach its Lynden-Bell equilibrium, but stalled before the relaxation could be completed (Chavanis, 2006b).

Thus, for the phase-volume conservation and the correlation volume to be meaningful features of any complete theory, they must earn their place in it. Since the microgranulation ansatz grants them a privileged position without question, it cannot be trusted without question. Nevertheless, it provides an insight into what interesting effects could be contained in a theory where these features are valid and meaningful. We give one such interesting example in the next section.

4.5.5 Relation between multi-waterbag and hyperkinetic collision integrals

In this section, we aim to draw a relation between the two collision integrals derived above: the multi-waterbag collision integral (4.55) and the hyperkinetic collision integral (4.86). The apparent distinction between these collision integrals is the need for an artificial closure in the multi-waterbag collision integral that is not present in the hyperkinetic collision integral. We will begin to shed light on this by stating the following exact relation between the two collision integrals:

1) The hyperkinetic collision integral recovers the multi-waterbag collision integral if the waterbag distribution function $P_{0\alpha}(\boldsymbol{v},\eta)$ satisfies the closure relation (4.66). To prove this, we observe, as we did already in section (4.4.1), that the first moment of the hyperkinetic collision integral (4.86) with respect to η is (4.35) with $\langle g_\alpha^2\rangle(\boldsymbol{v})$ given by

$$\langle g_\alpha^2\rangle(\boldsymbol{v}) = \int \mathrm{d}\eta \left[\eta - f_{0\alpha}(\boldsymbol{v})\right]^2 P_{0\alpha}(\boldsymbol{v},\eta) = \int \mathrm{d}\eta\, \eta^2 P_{0\alpha}(\boldsymbol{v},\eta) - \left[\int \mathrm{d}\eta\, \eta P_{0\alpha}(\boldsymbol{v},\eta)\right]^2. \tag{4.123}$$

Substituting $P_{0\alpha}(\boldsymbol{v},\eta)$ from (4.66) into (4.123) gives

$$\langle g_\alpha^2\rangle(\boldsymbol{v}) = \frac{1}{Z_\alpha(\psi_\alpha)} \int \mathrm{d}\eta\, \eta^2 e^{-\psi_\alpha(\boldsymbol{v})\eta - \gamma_\alpha(\eta)} - f_{0\alpha}^2(\boldsymbol{v}), \tag{4.124}$$

which is the continuous version of (4.54) used to arrive at the multi-waterbag collision integral. Therefore, if the waterbag distribution function satisfies (4.66), the hyperkinetic collision integral reduces to the multi-waterbag collision integral. However, this is not an automatic guarantee that the waterbag distribution function $P_{0\alpha}$ will 'stay on the closure'. In the event, however, it is possible to prove that it will.

2) If the waterbag distribution function $P_{0\alpha}(\boldsymbol{v},\eta)$ evolving under the hyperkinetic collision integral (4.86) ever satisfies the closure (4.66), then it will continue to satisfy it for all future times. For a given waterbag distribution function $P_{0\alpha}(\boldsymbol{v},\eta)$, let us define $\bar{P}_{0\alpha}(\boldsymbol{v},\eta)$ to be the waterbag distribution function that maximises the entropy (4.71) subject to having the same mean phase-space density $f_{0\alpha}(\boldsymbol{v})$ and waterbag content $\rho_\alpha(\eta)$ as $P_{0\alpha}(\boldsymbol{v},\eta)$, i.e., $\bar{P}_{0\alpha}(\boldsymbol{v},\eta)$ is the waterbag distribution function (4.66). Now, we examine the difference between the entropies of $\bar{P}_{0\alpha}(\boldsymbol{v},\eta)$ and $P_{0\alpha}(\boldsymbol{v},\eta)$:

$$\bar{S} - S = \sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\eta \left(P_{0\alpha} \ln P_{0\alpha} - \bar{P}_{0\alpha} \ln \bar{P}_{0\alpha} \right). \tag{4.125}$$

We first note that this relative entropy is, by definition of $\bar{P}_{0\alpha}$, non-negative and only zero if $P_{0\alpha} = \bar{P}_{0\alpha}$ for all α. Furthermore, since $\bar{P}_{0\alpha}$ is determined by $P_{0\alpha}$ and $\bar{S}$ by $\bar{P}_{0\alpha}$, we may calculate this relative entropy at each time, and take its time derivative. Taking the time derivative of S is straightforward given the collision integral (4.86). The time derivative of $\bar{S}$ can be determined in terms of only ψ_α and $\partial f_{0\alpha}/\partial t$, as the entropy has the form (4.60), the only difference being that the evolution of $f_{0\alpha}$ is now governed by the hyperkinetic collision integral, not the multi-waterbag one. Using (4.62), we may, therefore, write the evolution of the relative entropy as

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\left(\bar{S} - S\right) &= \sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \int \mathrm{d}\boldsymbol{v} \left[\psi_\alpha \frac{\partial f_{0\alpha}}{\partial t} + \int \mathrm{d}\eta\, (1 + \ln P_{0\alpha}) \frac{\partial P_{0\alpha}}{\partial t} \right] \\ &= \sum_\alpha \frac{V}{\Delta\Gamma_\alpha} \iint \mathrm{d}\boldsymbol{v}\,\mathrm{d}\eta \left(\eta\psi_\alpha + \ln P_{0\alpha} \right) \frac{\partial P_{0\alpha}}{\partial t}. \end{aligned} \tag{4.126}$$

Consequently, when the waterbag distribution function satisfies the closure (4.66), the relative entropy will be zero by definition of $\bar{P}_{0\alpha}$ and, by (4.126), its time derivative will also be zero, so the relative entropy will stay zero for all future times. Since the relative entropy is zero only for $P_{0\alpha} = \bar{P}_{0\alpha}$ this proves our second statement. Therefore, it is possible to construct initial conditions for which the

multi-waterbag collision integral (4.55) correctly describes the evolution due to the hyperkinetic collision integral for all future times.

Given that the distribution function $P_{0\alpha}(\boldsymbol{v}, \eta)$ remains on the closure if it begins exactly on the closure, it is natural to ask if the evolution of the distribution function is, in fact, forced towards the closure by the hyperkinetic collision integral. If this was true, it would make the multi-waterbag collision integral (4.55) a valid approximation for the hyperkinetic collision integral after an initial transient period. However, for this to be true, there would have to be two time scales hidden within the hyperkinetic collision operator: a shorter time scale on which the hyperkinetic distribution function approached the closure and a longer time scale over which the closure evolved according to the multi-waterbag collision integral. It is trivially clear that this separation of time scales exists in a one-dimensional system, since the 1D hyperkinetic collision integral forbids the mean phase-space density to change. In this case, the only evolution is towards satisfying the closure (4.66). In a higher number of dimensions, it may therefore be useful to consider two types of collisions: those which do and do not alter the mean phase-space density $f_{0\alpha}(\boldsymbol{v})$. This artificial divide would then describe two processes by which the system raises its entropy: one by making $f_{0\alpha}(\boldsymbol{v})$ approach a Lynden-Bell equilibrium and the other by reordering waterbags without altering $f_{0\alpha}(\boldsymbol{v})$. The latter process is one in which entropy is maximised locally in phase space (i.e., for each $\boldsymbol{v}$). A useful analogy might be the conventional collisional dynamics of gases, which relax quickly, at the collision rate, to a local Maxwellian, and slowly, at the diffusion rate, to the global one.

4.6 Strange relaxation in multispecies Lynden-Bell plasma

Having derived the hyperkinetic collision integral, its conservation laws, steady states and H-theorem, it would seem that relaxation to equilibrium has successfully been turned into yet another app. Like a figure of Greek myth, the collision integral is then destined to be made redundant by its own offspring: the H-theorem. This is because the H-theorem prescribes for each initial condition a steady-state distribution function, which makes evolving the hyperkinetic collision integral in time seem moot. In this section, however, we will show that there are certain

fairly general initial conditions for which the process of relaxation has interesting and non-trivial properties.

For this to be the case, it is clear that there must be multiple time scales within the problem. Based on our estimate (4.107) of effective collision frequency relative to the true-collision frequency, it is clear that one way to make this possible is to consider two species with a large mass ratio: electrons and ions. Borrowing intuition from the conventional collisional theory, this will allow for a separation of time scales whereby particles of the lighter, faster, species will collide most frequently while the heavier, slower, species will collide, and therefore evolve, at a lower rate. This intuitive picture is, however, complicated by the fact that, from our knowledge of the steady states (4.89), we expect the thermal velocities of the particles to be such that

$$\beta \Delta\Gamma_\alpha \eta_\alpha^{\text{eff}} \frac{1}{2} m_\alpha v_{\text{th}\alpha}^2 \sim 1 \implies v_{\text{th}\alpha} \sim \sqrt{\frac{2}{\beta \Delta\Gamma_\alpha^{\text{eff}} \eta_\alpha^{\text{eff}} m_\alpha}}. \tag{4.127}$$

This is to say that, relative to other species, particles of a given species behave as though they had an effective mass $\Delta\Gamma_\alpha \eta_\alpha^{\text{eff}}$ times greater than their true mass. To avoid this complexity, we will restrict ourselves to the case where the typical number of particles within a correlation volume is comparable for electrons and ions, viz.,

$$\Delta\Gamma_{\text{e}} \eta_{\text{e}}^{\text{eff}} \sim \Delta\Gamma_{\text{i}} \eta_{\text{i}}^{\text{eff}}. \tag{4.128}$$

This amounts to assuming that collisionless effects do not override the scale separation imposed by the mass ratio, which is what creates the interesting effects in the conventional collisional theory. Furthermore, we will make the restriction that all velocity scales of $P_{0\alpha}$ do not differ from $v_{\text{th}\alpha}$ given in (4.127) by a factor of $\sqrt{m_{\text{e}}/m_{\text{i}}}$ or more (i.e., that $P_{0\alpha}$ does not have sharp discontinuities associated with extreme degeneracy). This will ensure that velocity derivatives in (4.86) do not override the orderings imposed by the mass ratio, which will be particularly important for inter-species collisions. From (4.86) and the orderings (4.127) and (4.128), it is then apparent that the same-species collision operators only contain a single time scale. This is simply a statement that the only purpose of the same-species collision operator is to increase the entropy of that species independent of the others, and that it can only do that on a single time scale. In contrast, owing to the mass ratio, the inter-species collisions can have multiple time scales. We know from standard collisional theory that these give rise to the rates at which momentum

and temperature are equalised between species. To extract similar, but distinct, features from our new collision integral, we would therefore have to expand it in small mass ratio. Before we do this, however, it is possible to anticipate from simple observations what the interesting new physics will be.

4.6.1 Preview of strange relaxation

To study the interaction between species we will write our collision operator (4.86) as

$$\frac{\partial P_\alpha}{\partial t} = \sum_{\alpha'} C_{\alpha\alpha'}[P_\alpha, P_{\alpha'}], \tag{4.129}$$

in terms of the interspecies collision operator $C_{\alpha\alpha'}$, which can be easily read from elements of the species sum in (4.86). For compactness of notation, we shall henceforth drop 0's from the subscripts of the mean distribution function. For electrons interacting with ions, the dominant effect is the diffusion-like first term of (4.86). Comparing this to the true-collision operator (4.105), we see that (aside from acting on P_e instead of f_e) the new, 'collisionless' feature in this diffusion term is that f_i is replaced by $\Delta\Gamma_\mathrm{i}\langle g_\mathrm{i}^2\rangle$ with $\langle g_\mathrm{i}^2\rangle$ defined by (4.123). Therefore, any effects on the electrons relating to f_i in the standard collisional theory will now be replaced by the analogous effects relating to $\Delta\Gamma_\mathrm{i}\langle g_\mathrm{i}^2\rangle$ (see figure 4.4 for an illustration of this in a single-waterbag distribution). In particular, instead of being isotropised by the ions of density n_i and dragged towards the ion velocity $\boldsymbol{u}_\mathrm{i}$, the electrons will see an 'anomalous density' n_i^a and find themselves dragged towards an 'anomalous velocity' $\boldsymbol{u}_\mathrm{i}^\mathrm{a}$, given by

$$n_\mathrm{i}^\mathrm{a} = \Delta\Gamma_\mathrm{i} \int \mathrm{d}\boldsymbol{v}\,' \langle g_\mathrm{i}^2\rangle(\boldsymbol{v}'), \qquad \boldsymbol{u}_\mathrm{i}^\mathrm{a} = \frac{\int \mathrm{d}\boldsymbol{v}'\, \boldsymbol{v}' \langle g_\mathrm{i}^2\rangle(\boldsymbol{v}')}{\int \mathrm{d}\boldsymbol{v}'\, \langle g_\mathrm{i}^2\rangle(\boldsymbol{v}')}. \tag{4.130}$$

Since ion relaxation is slow compared to the electron one, the ion distribution need not be isotropic, so, in general, $\boldsymbol{u}_\mathrm{i}^\mathrm{a} \neq 0$. Furthermore, there is no need for $\boldsymbol{u}_\mathrm{i}^\mathrm{a}$ to point in the same direction as the ions' mean velocity $\boldsymbol{u}_\mathrm{i}$, and indeed it is even possible that $\boldsymbol{u}_\mathrm{i} = 0$ while $\boldsymbol{u}_\mathrm{i}^\mathrm{a} \neq 0$. This means that collisionless relaxation can lead to spontaneous generation of current from ion anisotropies (e.g., from an ion heat flux). It should be stressed that while this is a drag, the collision operator (4.86) is still guaranteed to conserve the total energy of the system: the gain in electron energy due to the anomalous drag comes at the comes at the price of the energy lost by the ions in turn.

The physical significance of this becomes especially clear for a single-waterbag example shown in figure 4.4. The phase-space exclusion effect makes the ions behave as though they were fermions. Therefore, the probability for an electron to have a 'collision' with an ion whose velocity is $\boldsymbol{v}$ is not just the probability $f_\mathrm{i}(\boldsymbol{v})/\eta$ that an ion can be found at that velocity but the probability that an ion can be found at that velocity $\boldsymbol{v}$ and that the velocity $\boldsymbol{v}+\Delta\boldsymbol{v}$ into which that ion will be scattered by the interaction is not already occupied. Since an ion is deflected by a very small amount in a collision with an electron ($\Delta\boldsymbol{v} \ll \boldsymbol{v}$), this probability is approximately proportional to $f_\mathrm{i}(\boldsymbol{v})\,[\eta - f_\mathrm{i}(\boldsymbol{v})]$, which is precisely $\langle g_\mathrm{i}^2\rangle(\boldsymbol{v})$ from which the anomalous ion velocity and density are defined in (4.130). This means that densely packed portions of the ion phase space become effectively invisible to the electrons. Thus, the mean velocity towards which the electrons are dragged is not the true ion mean velocity but the mean velocity of those ions that move in the less densely occupied regions of the phase space.

After an initial period of this 'strange relaxation' the electrons will converge to their own Lynden-Bell equilibrium moving at the anomalous ion velocity $\boldsymbol{u}_\mathrm{i}^\mathrm{a}$. The ions, however, will not thus far have relaxed significantly except to ensure the conservation of total momentum (and thus alter the mean velocity by a mass-ratio-small amount $\sim -\boldsymbol{u}_\mathrm{i}^\mathrm{a} m_\mathrm{e} n_\mathrm{e}/m_\mathrm{i} n_\mathrm{i}$). Once the ion-ion relaxation time scale is reached, the ions will begin to erase any anisotropy in their distribution function as they proceed towards their own Lynden-Bell equilibrium. As their anisotropy vanishes, the anomalous mean ion velocity will tend to the true mean ion velocity, $\boldsymbol{u}_\mathrm{i}^\mathrm{a} \to \boldsymbol{u}_\mathrm{i}$, and the electrons' mean velocity will doggedly follow it, relaxing towards two Lynden-Bell equilibria of equal velocities but distinct temperatures. Finally, on the longest (ion-electron) interaction time scale, the distributions will equalise their temperatures, reaching the overall maximum-entropy state and completing the relaxation.

This is the physics of the strange relaxation process[5]. In the remainder of this

[5]It should be noted that anomalous currents generated by this relaxation will, for a plasma hosting electromagnetic fields (as opposed to electrostatic fields), lead to the production of waves and potentially large-scale fields: a so-called "battery". These fields, however, will undercut the assumption that the relaxation is driven only by the collision operator, meaning the solution would require a more sophisticated scheme such as that used for the Spitzer–Härm problem (cf. Spitzer &Härm 1953).

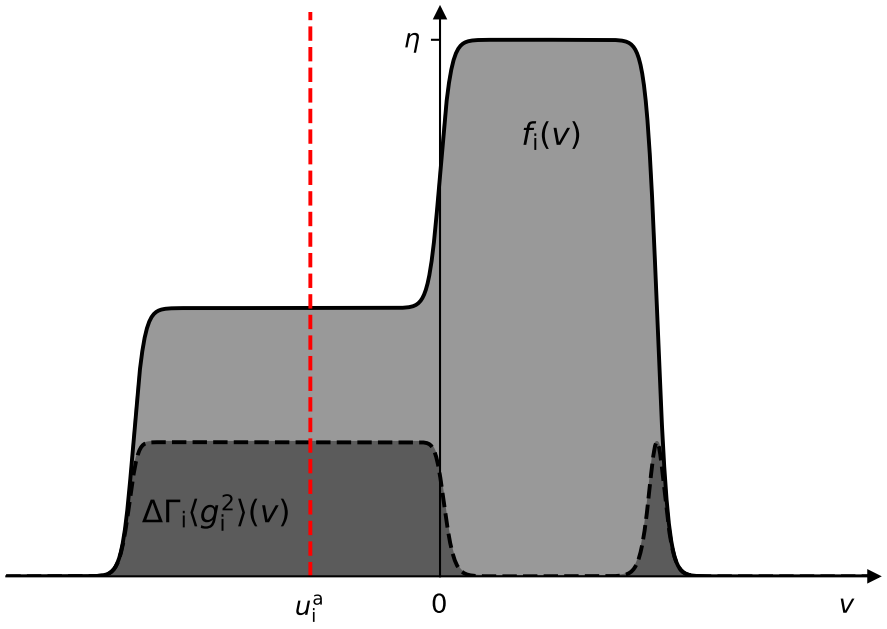

Figure 4.4: A cartoon of a possible anisotropic ion distribution function that possesses an anomalous velocity $\boldsymbol{u}_{\mathrm{i}}^{\mathrm{a}}$. The solid line represents the mean ion distribution function. If the exact ion phase-space density is viewed as a single waterbag (as in the Kadomtsev–Pogutse collision integral (4.43)), then the 'effective distribution' of ions with which the electrons will interact is shown by the dashed black line. The mean ion velocity is zero, but the anomalous velocity is not, as shown by the dashed red line. Note that the height of the effective distribution would be rescaled if η and $\Delta\Gamma_{\mathrm{i}}$ took different values, but this would only affect the rate of relaxation to the anomalous velocity, not the anomalous velocity itself.

section, we demonstrate formally that this is indeed what happens by carrying out the mass-ratio expansion of the electron-ion (section 4.6.3) and the ion-electron (section 4.6.4) collision operators.This calculation follows the standard path well trodden for true collisions (specifically, as presented, e.g., in Parra 2019), but with a few important adjustments.

4.6.2 Landau form of the hyperkinetic collision operator

We will first make a further simplification by reducing our hyperkinetic collision operator (4.86) to the so-called 'Landau form'. Doing so amounts to finding an

approximate expression for the $\boldsymbol{k}$ sum

$$\frac{1}{V}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\frac{\delta\left(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}')\right)}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}, \tag{4.131}$$

which is complicated by the presence of the dielectric function

$$\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}} = 1+\sum_{\alpha'}\frac{4\pi q_{\alpha'}^2}{m_{\alpha'}k^2}\iint \mathrm{d}\boldsymbol{v}'\,\mathrm{d}\eta'\,\frac{\eta'}{\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}')}\boldsymbol{k}\cdot\frac{\partial P_{\alpha'}}{\partial \boldsymbol{v}}. \tag{4.132}$$

Here the $\boldsymbol{v}'$ integral is taken along the Landau contour. To simplify (4.131), the important feature to note is that, for length scales shorter than the Debye length and in the absence of instabilities, the second term in (4.132) will be small and $\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}$ can be approximated by unity. For scales significantly longer than the Debye length, $\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}\sim(\lambda_{\mathrm{De}}k)^{-2}$, which will make their contribution to the sum in (4.131) small in $(\lambda_{\mathrm{De}}k)^4$. Therefore, following Landau (1936), we truncate the $\boldsymbol{k}$ sum at $(k\lambda_{\mathrm{De}})\gtrsim 1$, and approximate the dielectric function by unity. Denoting $\boldsymbol{w}=\boldsymbol{v}-\boldsymbol{v}'$ and integrating in cylindrical coordinates such that $\boldsymbol{k}=k_\parallel\hat{\boldsymbol{w}}+\boldsymbol{k}_\perp$, we get

$$\frac{1}{V}\sum_{\boldsymbol{k}}\frac{\boldsymbol{k}\boldsymbol{k}}{k^4}\frac{\delta(\boldsymbol{k}\cdot\boldsymbol{w})}{|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}|^2}\approx\iiint\frac{\mathrm{d}\boldsymbol{k}}{(2\pi)^3}\frac{\boldsymbol{k}_\perp\boldsymbol{k}_\perp}{k_\perp^4}\delta(k_\parallel w)=\frac{1}{8\pi^2}\frac{1}{w}\left(\mathbb{I}-\frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right)\ln\left(\frac{k_{\max}}{k_{\min}}\right). \tag{4.133}$$

Here we have truncated the $k_\perp$ integral, not only at $k_{\min}^{-1}\sim\lambda_{\mathrm{De}}$ as discussed above, but also at scales shorter than $k_{\max}^{-1}$. In the standard collisional theory, this corresponds to cutting off the integral at the distance of closest approach of two particles and gives rise to the Coulomb logarithm. In our treatment, such a high-k cutoff should instead be associated with the physical length scale of the correlations, which may be thought of as the distance of closest approach for two waterbags. Henceforth, we will denote this generalised Coulomb-logarithm-like quantity by $\Lambda_{\alpha\alpha'}$.

With the approximation (4.133), the interspecies collision operator in (4.86) becomes

$$\begin{aligned}C_{\alpha\alpha'}[P_\alpha,P_{\alpha'}] = {}& \frac{\gamma_{\alpha\alpha'}}{m_\alpha}\frac{\partial}{\partial\boldsymbol{v}}\cdot\int\frac{\mathrm{d}\boldsymbol{v}'}{w}\left(\mathbb{I}-\frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right)\cdot\int\mathrm{d}\eta'\,\eta'\\ &\left\{\frac{\Delta\Gamma_{\alpha'}}{m_\alpha}\left[\eta'-f_{\alpha'}(\boldsymbol{v}')\right]P_{\alpha'}(\boldsymbol{v}',\eta')\left.\frac{\partial P_\alpha}{\partial\boldsymbol{v}}\right|_\eta-\frac{\Delta\Gamma_\alpha}{m_{\alpha'}}\left[\eta-f_\alpha(\boldsymbol{v})\right]P_\alpha(\boldsymbol{v},\eta)\left.\frac{\partial P_{\alpha'}}{\partial\boldsymbol{v}'}\right|_{\eta'}\right\},\end{aligned} \tag{4.134}$$

where $\gamma_{\alpha\alpha'}=2\pi q_\alpha^2q_{\alpha'}^2\Lambda_{\alpha\alpha'}$.

4.6.3 Electron-ion relaxation

Zeroth order: isotropisation of electrons

To lowest order in $\sqrt{m_\mathrm{e}/m_\mathrm{i}}$, the second term in brackets in (4.134) vanishes entirely and the velocity difference between the electrons and ions, $\boldsymbol{w}$, becomes approximately equal to the electron velocity $\boldsymbol{v}$. This leaves the collision operator in the immensely simple form

$$C_\mathrm{ei} = \frac{\gamma_\mathrm{ei} n_\mathrm{i}^\mathrm{a}}{m_\mathrm{e}^2} \frac{\partial}{\partial \boldsymbol{v}} \cdot \left[\frac{1}{v} \left(\mathbb{I} - \frac{\boldsymbol{v}\boldsymbol{v}}{v^2} \right) \cdot \left. \frac{\partial P_\mathrm{e}}{\partial \boldsymbol{v}} \right|_\eta \right], \tag{4.135}$$

where n_i^a is the anomalous ion density given by the first expression in (4.130). This operator is a pitch-angle-scattering (Lorentz) operator (see, e.g., Helander &Sigmar 2005), which causes relaxation on time scales comparable to that of electron-electron interactions. Its effect is to isotropise P_e. Physically this is a consequence of the ions' high mass. Like ping-pong balls bouncing off bowling balls, the electrons bounce off the ions without exchanging any energy and the only effect is to isotropise their distribution. However, to this lowest order, we have only retained terms ordered with the electron thermal velocity, losing any speeds ordered with the ion velocity. To retain velocities of that size, viz., the electrons' mean velocity, we must go to next order.

First order: anomalous drag

To next order, determining the collision integral becomes moderately less trivial. First, we will expand the electron hyperkinetic distribution function in mass ratio and assume that its lowest-order part, denoted by $P_\mathrm{e}^\mathrm{I}(v,\eta)$, is isotropic. We further assume that any deviations from isotropy are due to velocities ordered with the ion thermal velocity, so the anisotropic correction, denoted $P_\mathrm{e}^\mathrm{A}(\boldsymbol{v},\eta)$, will be smaller than $P_\mathrm{e}^\mathrm{I}(v,\eta)$ by a factor of $\sqrt{m_\mathrm{e}/m_\mathrm{i}}$. As well as this, since the $\boldsymbol{v}'$ integral in (4.134) is over ion velocities, $\boldsymbol{v}'$ will be small in $\sqrt{m_\mathrm{e}/m_\mathrm{i}}$ compared to $\boldsymbol{v}$. Therefore, we can now expand the $\boldsymbol{w}$-dependent tensor in (4.134) to first order in $\sqrt{m_\mathrm{e}/m_\mathrm{i}}$, giving

$$\begin{aligned} \frac{1}{w}\left(\mathbb{I} - \frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right) &\approx \frac{1}{v}\left(\mathbb{I} - \frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right) - \boldsymbol{v}' \cdot \frac{\partial}{\partial \boldsymbol{v}}\left[\frac{1}{v}\left(\mathbb{I} - \frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right)\right] \\ &= \frac{1}{v}\left(1 + \frac{3\boldsymbol{v}\cdot\boldsymbol{v}'}{v^2}\right)\left(\mathbb{I} - \frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right) + \frac{\boldsymbol{v}'\boldsymbol{v} + \boldsymbol{v}\boldsymbol{v}' - 2(\boldsymbol{v}\cdot\boldsymbol{v}')\mathbb{I}}{v^3}. \end{aligned} \tag{4.136}$$

Since the anisotropic part of the hyperkinetic distribution function is already small, we only need the lowest-order contribution from (4.136) to act on $P_{\rm e}^{\rm A}(\boldsymbol{v},\eta)$. For the isotropic part $P_{\rm e}^{\rm I}(\boldsymbol{v})$, the lowest-order contribution vanishes (as it must, since any isotropic distribution is a solution of the lowest-order problem) and we must keep the next-order terms. These terms are

$$\begin{aligned}\frac{1}{w}\left(\mathbb{I}-\frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right)\cdot\frac{\partial P_{\rm e}^{\rm I}}{\partial\boldsymbol{v}} &\approx \frac{\boldsymbol{v}'\boldsymbol{v}+\boldsymbol{v}\boldsymbol{v}'-2(\boldsymbol{v}\cdot\boldsymbol{v}')\mathbb{I}}{v^3}\cdot\frac{\boldsymbol{v}}{v}\frac{\partial P_{\rm e}^{\rm I}}{\partial v}\\ &=\frac{1}{v}\left(\mathbb{I}-\frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right)\cdot\frac{\boldsymbol{v}'}{v}\frac{\partial P_{\rm e}^{\rm I}}{\partial v}=\frac{1}{v}\left(\mathbb{I}-\frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right)\cdot\frac{\partial}{\partial\boldsymbol{v}}\left(\frac{\boldsymbol{v}\cdot\boldsymbol{v}'}{v}\frac{\partial P_{\rm e}^{\rm I}}{\partial v}\right),\end{aligned}\tag{4.137}$$

In the final equality, we made use of the isotropy of $P_{\rm e}^{\rm I}$ in order to insert an additional derivative into the product. The effect of this is to cast this term in the form of a pitch-angle-scattering operator. The resulting collision operator is

$$C_{\rm ei}[P_{\rm e},P_{\rm i}]=\frac{\gamma_{\rm ei}}{m_{\rm e}^2}n_{\rm i}^{\rm a}\frac{\partial}{\partial\boldsymbol{v}}\cdot\left\{\frac{1}{v}\left(\mathbb{I}-\frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right)\cdot\frac{\partial}{\partial\boldsymbol{v}}\left[P_{\rm e}^{\rm A}(\boldsymbol{v},\eta)+\frac{\boldsymbol{v}\cdot\boldsymbol{u}_{\rm i}^{\rm a}}{v}\frac{\partial P_{\rm e}^{\rm I}}{\partial v}\right]\right\},\tag{4.138}$$

where $n_{\rm i}^{\rm a}$ and $\boldsymbol{u}_{\rm i}^{\rm a}$ are given by (4.130).

Despite not knowing the exact evolution due to (4.138), we may use the fact that it is a pitch-angle-scattering operator to read off the steady state. This will occur when the expression in the square brackets is isotropic. Without loss of generality, we may define $P_{\rm e}^{\rm A}$ to have zero spherical average (isotropic part), in which case the only way for the term in square brackets in (4.138) to be isotropic is for it to be identically zero. Therefore the fixed point of the collision operator (4.138) is

$$P_{\rm e}(\boldsymbol{v},\eta)=P_{\rm e}^{\rm I}(v,\eta)+P_{\rm e}^{\rm A}(\boldsymbol{v},\eta)=P_{\rm e}^{\rm I}(v,\eta)-\frac{\boldsymbol{v}\cdot\boldsymbol{u}_{\rm i}^{\rm a}}{v}\frac{\partial P_{\rm e}^{\rm I}}{\partial v}\approx P_{\rm e}^{\rm I}(|\boldsymbol{v}-\boldsymbol{u}_{\rm i}^{\rm a}|,\eta),\tag{4.139}$$

the last equality holding with $O(m_{\rm e}/m_{\rm i})$ precision.

Thus, as prophesied in section 4.6.1, the electrons become isotropic around the anomalous ion velocity $\boldsymbol{u}_{\rm i}^{\rm a}$ rather than the true mean ion velocity. Dynamically, this manifests itself as a drag on the electrons with the rate of change of the electron momentum $m_{\rm e}n_{\rm e}\boldsymbol{u}_{\rm e}$ being

$$\begin{aligned}m_{\rm e}n_{\rm e}\frac{{\rm d}\boldsymbol{u}_{\rm e}}{{\rm d}t}&=\frac{\rm d}{{\rm d}t}\int{\rm d}\boldsymbol{v}\,m_{\rm e}\boldsymbol{v}\int{\rm d}\eta\,\eta P_{\rm e}(\boldsymbol{v},\eta)=\int{\rm d}\boldsymbol{v}\,m_{\rm e}\boldsymbol{v}\int{\rm d}\eta\,\eta C_{\rm ei}[P_{\rm e},P_{\rm i}]\\ &=-\frac{\gamma_{\rm ei}n_{\rm i}^{\rm a}}{m_{\rm e}}\int\frac{{\rm d}\boldsymbol{v}}{v}\left(\mathbb{I}-\frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right)\cdot\frac{\partial}{\partial\boldsymbol{v}}\left(f_{\rm e}^{\rm A}+\frac{\boldsymbol{v}\cdot\boldsymbol{u}_{\rm i}^{\rm a}}{v}\frac{\partial f_{\rm e}^{\rm I}}{\partial v}\right)\\ &=-\frac{\gamma_{\rm ei}n_{\rm i}^{\rm a}}{m_{\rm e}}\left[\int\frac{{\rm d}\boldsymbol{v}}{v^2}\left(\mathbb{I}-\frac{\boldsymbol{v}\boldsymbol{v}}{v^2}\right)\frac{\partial f_{\rm e}^{\rm I}}{\partial v}\right]\cdot\boldsymbol{u}_{\rm i}^{\rm a}-\frac{2\gamma_{\rm ei}n_{\rm i}^{\rm a}}{m_{\rm e}}\int{\rm d}\boldsymbol{v}\,\frac{\boldsymbol{v}}{v^3}f_{\rm e}^{\rm A}\\ &=\frac{8\pi\gamma_{\rm ei}n_{\rm i}^{\rm a}f_{\rm e}^{\rm I}(0)}{3m_{\rm e}}\boldsymbol{u}_{\rm i}^{\rm a}-\frac{2\gamma_{\rm ei}n_{\rm i}^{\rm a}}{m_{\rm e}}\int{\rm d}\boldsymbol{v}\,\frac{\boldsymbol{v}}{v}f_{\rm e}^{\rm A},\end{aligned}\tag{4.140}$$

where $f_\mathrm{e}^\mathrm{I}(v)$ and $f_\mathrm{e}^\mathrm{A}(\boldsymbol{v})$ are, respectively, the isotropic and anisotropic parts of the electron phase-space density. Going from the third to the fourth line we used the isotropy of f_e^I to compute the angle integral explicitly. The same cannot be done for the integral over the anisotropic part of the phase-space density. However, if we assume that, to lowest order, the electron phase-space density is isotropic around the mean electron velocity $\boldsymbol{u}_\mathrm{e}$, then

$$f_\mathrm{e}(\boldsymbol{v}) = f_\mathrm{e}^\mathrm{I}(|\boldsymbol{v}-\boldsymbol{u}_\mathrm{e}|) \approx f_\mathrm{e}^\mathrm{I} - \frac{\boldsymbol{v}\cdot\boldsymbol{u}_\mathrm{e}}{v}\frac{\partial f_\mathrm{e}^\mathrm{I}}{\partial v} \implies f_\mathrm{e}^\mathrm{A}(\boldsymbol{v}) = -\frac{\boldsymbol{v}\cdot\boldsymbol{u}_\mathrm{e}}{v}\frac{\partial f_\mathrm{e}^\mathrm{I}}{\partial v}, \tag{4.141}$$

which allows the final integral to be computed, giving the anomalous drag force

$$\boldsymbol{F}_\mathrm{ei} = m_\mathrm{e} n_\mathrm{e}\frac{\mathrm{d}\boldsymbol{u}_\mathrm{e}}{\mathrm{d}t} = -\frac{8\pi\gamma_\mathrm{ei} n_\mathrm{i}^\mathrm{a} f_\mathrm{e}^\mathrm{I}(0)}{3m_\mathrm{e}}\left(\boldsymbol{u}_\mathrm{e}-\boldsymbol{u}_\mathrm{i}^\mathrm{a}\right). \tag{4.142}$$

This is the same as the standard expression for the collisional drag (Spitzer, 1967), but replacing the ion density and mean velocity with the corresponding anomalous variants (4.130).

We note that neither the lowest- nor the first-order approximations of the electron-ion collision operator yet describes the transfer of energy that leads to the equilibration of temperature. To achieve this, and obtain the expected relaxation of temperature, one must go to higher order, which is easiest to do via the ion-electron collision operator.

4.6.4 Ion-electron relaxation: temperature equilibration

The ion-electron collision operator (4.134) is

$$C_\mathrm{ie} = \frac{\gamma_\mathrm{ei}}{m_\mathrm{i}}\frac{\partial}{\partial\boldsymbol{v}}\int\frac{\mathrm{d}\boldsymbol{v}'}{w}\left(\mathbb{I}-\frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right)\cdot\int\mathrm{d}\eta'\,\eta' \left\{\underbrace{\frac{\Delta\Gamma_\mathrm{e}}{m_\mathrm{i}}\left[\eta'-f_\mathrm{e}(\boldsymbol{v}')\right]P_\mathrm{e}(\boldsymbol{v}',\eta')\frac{\partial P_\mathrm{i}}{\partial\boldsymbol{v}}}_{\sim\frac{\Delta\Gamma_\mathrm{e}\eta_\mathrm{e}^\mathrm{eff}}{m_\mathrm{i}v_\mathrm{thi}}P_\mathrm{e}P_\mathrm{i}} - \underbrace{\frac{\Delta\Gamma_\mathrm{i}}{m_\mathrm{e}}\left[\eta-f_\mathrm{i}(\boldsymbol{v})\right]P_\mathrm{i}(\boldsymbol{v},\eta)\frac{\partial P_\mathrm{e}}{\partial\boldsymbol{v}'}}_{\sim\frac{\Delta\Gamma_\mathrm{i}\eta_\mathrm{i}^\mathrm{eff}}{m_\mathrm{e}v_\mathrm{the}}P_\mathrm{e}P_\mathrm{i}}\right\}. \tag{4.143}$$

Naively, this gives an ion-electron relaxation rate that is comparable to the ion-ion relaxation rate, which is smaller than the electron-ion one by a factor of $\sqrt{m_\mathrm{e}/m_\mathrm{i}}$. However, this neglects the fact that the electron distribution will isotropise itself on the electron-ion relaxation time scale. Thus, the lowest-order term in (4.143) will vanish before it ever gets a chance to participate. This renders the ion-electron

relaxation rate smaller than the ion-ion one by another factor of $\sqrt{m_\mathrm{e}/m_\mathrm{i}}$. As in the case of true collisions, this does not preclude the conservation of total momentum because an order-unity velocity change of the electrons only requires an order-mass-ratio adjustment to the ion velocity in order to conserve momentum, and the ion-electron collision operator is still fully capable of achieving a change this small on the electron-ion collision time scale.

Using the fact that the electron distribution will long since have become isotropic to lowest order, we can again carry out a mass-ratio expansion of the ion-electron collision operator. Since the first bracketed term in (4.143) is nominally smaller, the lowest-order contribution only requires the isotropic part of the electron distribution function. This becomes (noting that now the integration is over electron velocities so, $\boldsymbol{v}'' \gg \boldsymbol{v}$)

$$\int \frac{\mathrm{d}\boldsymbol{v}'}{v'}\left(\mathbb{I}-\frac{\boldsymbol{v}'\boldsymbol{v}'}{v'^2}\right)\int \mathrm{d}\eta'\,\eta'\left[\eta'-f_\mathrm{e}(\boldsymbol{v}')\right]P_\mathrm{e}(\boldsymbol{v}',\eta)=\frac{8\pi}{3}\mathbb{I}\int_0^\infty\langle(g_\mathrm{e}^\mathrm{I})^2\rangle(\boldsymbol{v}')v'\mathrm{d}v'\,. \tag{4.144}$$

The second bracketed term in (4.143) requires both the isotropic and anisotropic contributions, leading us to evaluate the integral

$$\begin{aligned}\int \frac{\mathrm{d}\boldsymbol{v}'}{w}\left(\mathbb{I}-\frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right)\cdot\frac{\partial f_\mathrm{e}}{\partial\boldsymbol{v}'}&=\int \frac{\mathrm{d}\boldsymbol{v}'}{w}\left(\mathbb{I}-\frac{\boldsymbol{w}\boldsymbol{w}}{w^2}\right)\cdot\left(\frac{\partial f_\mathrm{e}^\mathrm{A}}{\partial\boldsymbol{v}'}+\frac{\boldsymbol{v}'}{v'}\frac{\partial f_\mathrm{i}^\mathrm{I}}{\partial v'}\right)\\&=\int \frac{\mathrm{d}\boldsymbol{v}'}{v'}\left(\mathbb{I}-\frac{\boldsymbol{v}'\boldsymbol{v}'}{v'^2}\right)\cdot\frac{\partial f_\mathrm{e}^\mathrm{A}}{\partial\boldsymbol{v}'}+\left[\int\frac{\boldsymbol{v}'}{v'^2}\left(\mathbb{I}-\frac{\boldsymbol{v}'\boldsymbol{v}'}{v'^2}\right)\frac{\partial f_\mathrm{e}^\mathrm{I}}{\partial v'}\right]\cdot\boldsymbol{v}\\&=2\int\mathrm{d}\boldsymbol{v}'\,\frac{\boldsymbol{v}'}{v'^2}f_\mathrm{e}^\mathrm{A}(\boldsymbol{v}')-\frac{8\pi f_\mathrm{e}^\mathrm{I}(0)}{3}\boldsymbol{v}.\end{aligned} \tag{4.145}$$

Collecting the contributions (4.144) and (4.145) together, we find that the electron-ion collision operator is, to lowest order,

$$\begin{aligned}C_\mathrm{ie}=-\frac{\gamma_\mathrm{ei}\Delta\Gamma_\mathrm{i}}{m_\mathrm{e}m_\mathrm{i}}\frac{\partial}{\partial\boldsymbol{v}}\cdot\left\{\left[\eta-f_\mathrm{i}(\boldsymbol{v})\right]P_\mathrm{i}(\boldsymbol{v},\eta)\left[2\int\mathrm{d}\boldsymbol{v}'\,\frac{\boldsymbol{v}'}{v'^2}f_\mathrm{e}^\mathrm{A}(\boldsymbol{v}')-\frac{8\pi f_\mathrm{e}^\mathrm{I}(0)}{3}\boldsymbol{v}\right]\right\}\\+\frac{8\pi\gamma_\mathrm{ei}\Delta\Gamma_\mathrm{e}}{3m_\mathrm{i}^2}\int\langle(g_\mathrm{e}^\mathrm{I})^2\rangle(v')v'\mathrm{d}v'\,\frac{\partial^2 P_\mathrm{i}}{\partial\boldsymbol{v}^2}.\end{aligned} \tag{4.146}$$

While not amazingly insightful, it is at least obvious how to verify that, together with the electron-ion collision operator (4.138), this conserves the total momentum of the system, giving the opposite momentum change to (4.140), as it must.

Of course, we need not have stopped at the simplest assumption that the electron distribution is, to lowest order, isotropic. The electron-electron interactions,

which are just as frequent as the electron-ion interactions, will push the electron distribution function towards a Lynden-Bell equilibrium, which will have some mean velocity $\boldsymbol{u}_\mathrm{e}$ that lies close to the anomalous ion velocity $\boldsymbol{u}_\mathrm{i}^\mathrm{a}$ given by (4.130). Such a Lynden-Bell equilibrium for the electrons has the form

$$f_\mathrm{e}(\boldsymbol{v}) = \frac{\int \mathrm{d}\eta\, \eta e^{-\beta_\mathrm{e}\Delta\Gamma_\mathrm{e}\eta\left[\frac{1}{2}m_\mathrm{e}|\boldsymbol{v}-\boldsymbol{u}_\mathrm{e}|^2-\mu_\mathrm{e}(\eta)\right]}}{\int \mathrm{d}\eta\, e^{-\beta_\mathrm{e}\Delta\Gamma_\mathrm{e}\eta\left[\frac{1}{2}m_\mathrm{e}|\boldsymbol{v}-\boldsymbol{u}_\mathrm{e}|^2-\mu_\mathrm{e}(\eta)\right]}}. \tag{4.147}$$

This immediately allows f_e^A in (4.146) to be computed as in (4.141). Using the property of Lynden-Bell equilibria expressed by (4.107) and (4.108), we also find

$$\langle (g_\mathrm{e}^\mathrm{I})^2 \rangle = -\frac{1}{\beta_\mathrm{e}\Delta\Gamma_\mathrm{e} m_\mathrm{e} v}\frac{\partial f_\mathrm{e}^\mathrm{I}}{\partial v}. \tag{4.148}$$

The ion-electron collision operator (4.146) then reduces to

$$C_\mathrm{ie}[P_\mathrm{i}, P_\mathrm{e}] = \frac{8\pi\gamma_\mathrm{ei}\Delta\Gamma_\mathrm{i} f_\mathrm{e}^\mathrm{I}(0)}{3m_\mathrm{i}m_\mathrm{e}}\frac{\partial}{\partial \boldsymbol{v}}\cdot\left\{\frac{1}{\beta_\mathrm{e}\Delta\Gamma_\mathrm{i} m_\mathrm{i}}\frac{\partial P_\mathrm{i}}{\partial \boldsymbol{v}} + (\boldsymbol{v}-\boldsymbol{u}_\mathrm{e})\left[\eta - f_\mathrm{i}(\boldsymbol{v})\right] P_\mathrm{i}(\boldsymbol{v},\eta)\right\}. \tag{4.149}$$

Clearly, this will be stationary for a Lynden-Bell equilibrium if the mean electron velocity $\boldsymbol{u}_\mathrm{e}$ is equal to the mean ion velocity (i.e., if neither distribution has a mean velocity in the frame moving with $\boldsymbol{u}_\mathrm{e}$), and the thermodynamic beta β_i of the ions' Lynden-Bell equilibrium is equal to that of the electrons, β_e.

Because the ion-ion interactions increase the entropy of the ion hyperkinetic distribution function and are faster than the ion-electron interactions, it is certainly true that the ion distribution will be a Lynden-Bell equilibrium. Furthermore, because the anomalous ion velocity $\boldsymbol{u}_\mathrm{i}^\mathrm{a}$ is equal to the mean ion velocity for isotropic distributions like the Lynden-Bell equilibria, the electron-ion collision operator will have seen to it that the mean ion and electron velocities match. However, since the electron-electron, electron-ion and ion-ion collision operators to lowest order do not alter the energies of the two distributions, it is not guaranteed that the electrons and ions will have the same thermodynamic beta. Thus, as anticipated, the final piece of the relaxation puzzle, the relaxation of temperatures, occurs on the ion-electron interaction time scale. By assuming a Lynden-Bell equilibrium for ions with mean velocity equal to the electron mean velocity (i.e., both species stationary in the zero momentum frame), we can calculate the rate of change of

the kinetic energy of the ion distribution function:

$$\begin{aligned}\frac{\mathrm{d}E_\mathrm{i}}{\mathrm{d}t} &= \frac{\mathrm{d}}{\mathrm{d}t}\int \mathrm{d}\boldsymbol{v}\,\frac{1}{2}m_\mathrm{i}v^2\int \mathrm{d}\eta\,\eta P_\mathrm{i}(\boldsymbol{v},\eta)\\ &= \frac{8\pi\gamma_\mathrm{ei}f_\mathrm{e}^\mathrm{I}(0)}{3m_\mathrm{i}m_\mathrm{e}}\left[\frac{3}{\beta_\mathrm{e}}\int \mathrm{d}\boldsymbol{v}\,f_\mathrm{i}(\boldsymbol{v}) - m_\mathrm{i}\Delta\Gamma_\mathrm{i}\int \mathrm{d}\boldsymbol{v}\,|\boldsymbol{v}|^2\langle g_\mathrm{i}^2\rangle(\boldsymbol{v})\right]\\ &= \frac{8\pi\gamma_\mathrm{ei}f_\mathrm{e}^\mathrm{I}(0)}{3m_\mathrm{i}m_\mathrm{e}}\left(\frac{3n_\mathrm{i}}{\beta_\mathrm{e}} + \frac{1}{\beta_\mathrm{i}}\int \mathrm{d}\boldsymbol{v}\,\boldsymbol{v}\cdot\frac{\partial f_\mathrm{i}}{\partial \boldsymbol{v}}\right)\\ &= \frac{8\pi\gamma_\mathrm{ei}f_\mathrm{e}^\mathrm{I}(0)n_\mathrm{i}}{3m_\mathrm{i}m_\mathrm{e}}\left(\frac{1}{\beta_\mathrm{e}} - \frac{1}{\beta_\mathrm{i}}\right),\end{aligned} \tag{4.150}$$

where, in going to the second line we integrated by parts, in going to the third line we exploited the relationship (4.148) between $\langle g_\mathrm{i}^2\rangle$ and f_i, and finally integrated by parts again in going to the last line. This is simply another manifestation of the second law of thermodynamics (guaranteed by the H-theorem): energy will flow from the thermodynamically hotter species to the thermodynamically colder one until temperatures equalise. Note, however, that this energy flow need not produce the equilibration of kinetic energies, which are no longer directly tied to the thermodynamic temperatures, the equilibria being non-Maxwellian.

4.7 Summary

The existence of collision integrals that have the Lynden-Bell (1967) equilibria as their steady-state solutions implies that these equilibria can be reached dynamically. Two collision integrals have arisen: the multi-waterbag collision integral (4.55), and the hyperkinetic collision integral (4.86), both generalising the single-waterbag collision integral first derived by Kadomtsev &Pogutse (1970). It is a key ingredient that these collision integrals are equipped with H-theorems, confirming that the dynamical evolution towards the Lynden-Bell equilibria is made inevitable by the requirement to increase the Lynden-Bell entropy. The derivation of both collision integrals had to contend with a closure problem: the evolution of the mean phase-space density $f_{0\alpha}$ requires knowledge of the correlation function of the exact phase-space density between separate parts of phase space: $\langle f_\alpha(\boldsymbol{x},\boldsymbol{v})f_{\alpha'}(\boldsymbol{x}',\boldsymbol{v}')\rangle$, which is not, in general, known. For both collision integrals, this problem is partially resolved by the microgranulation ansatz (see section 4.2.5), a version of which was first used by Kadomtsev &Pogutse (1970) to derive their collision integral. The

microgranulation ansatz posits that the exact phase-space density of particles of species α is correlated over a small, but non-zero, volume in phase space $\Delta\Gamma_\alpha$, but is perfectly mixed over larger volumes. In contrast, one can derive the Balescu–Lenard integral (4.105) describing 'true' particle collisions by assuming that all particles are statistically independent, i.e., that a particle is only correlated with itself. The effect of the microgranulation ansatz is to reduce the problem of calculating the worrisome two-point correlator $\langle f_\alpha(\boldsymbol{x},\boldsymbol{v})f_{\alpha'}(\boldsymbol{x}',\boldsymbol{v}')\rangle$ to one of calculating a more manageable one-point correlator $\langle f_\alpha^2\rangle(\boldsymbol{v})$.

A closure is still required because the variance of a random quantity cannot, in general, be determined by its mean. Kadomtsev &Pogutse (1970) restricted the exact phase-space density $f_{0\alpha}(\boldsymbol{x},\boldsymbol{v})$ to only two possible values, η_α or 0, which immediately implied $\langle f_\alpha^2\rangle = \eta_\alpha f_{0\alpha}$, removing the closure problem (see section 4.3.1). However, this single-waterbag model is obviously extremely non-general. To move past it, in section 4.3.2, a scheme motivated by statistical mechanics is employed, which can be traced back to the treatment of geophysical turbulence by Chavanis (2005). By maximising the Shannon entropy (4.48) subject to a continuum of constraints (4.69) (the 'waterbag content' of the distribution function, or its Casimir invariants), one finds that the correlator $\langle f_\alpha^2\rangle$ can be written implicitly in terms of $f_{0\alpha}$. This leads to the multi-waterbag collision integral (4.55), which grows the Lynden-Bell entropy (4.58) and has the Lynden-Bell equilibria (4.57) as its only fixed points. Notably, if laterally to our main purpose here, this formalism allows one to recover very efficiently the 'true' collision integrals: Balescu–Lenard (section 4.5.1) and the collision integrals for fermionic and bosonic quantum plasmas (Appendix F).

An alternative route to solving the closure problem for $\langle f_\alpha^2\rangle$ is to dodge it entirely. In section 4.4, instead of trying to determine the evolution of the mean phase-space density $f_{0\alpha}(\boldsymbol{v})$, a 'hyperkinetic' approach is introduced, treating the exact phase-space density $f_\alpha(\boldsymbol{x},\boldsymbol{v})$ as a random field and asking for the probability of finding it to have a particular value η at a particular position in phase space. This method, pioneered by Severne &Luwel (1980) (for discrete waterbags) and having its origins in vortex kinetics and galactic dynamics (Chavanis et al., 1996), resolves the closure problem by calculating $f_{0\alpha}$ and $\langle f_\alpha^2\rangle$ as the first and second moments, respectively, of the 'hyperkinetic distribution function' $P_{0\alpha}(\boldsymbol{v},\eta) = \langle\delta(f_\alpha(\boldsymbol{x},\boldsymbol{v})-\eta)\rangle$. This function

is evolved by the hyperkinetic collision integral (4.86), which also grows the Lynden-Bell entropy (4.86) and has the Lynden-Bell equilibria (4.89) as its only fixed points.

It is immediately apparent that the statistical-mechanical closure (4.66) used to derive the multi-waterbag collision integral (4.55) is simply a special case of the hyperkinetic distribution function. In section 4.5.5, we show that in fact their connection is deeper: we prove that, should the hyperkinetic distribution be placed precisely on the closure (4.66), the two collision integrals would then be equivalent for all future times. This could imply that entropy maximisation local in phase space is an inherent feature of the hyperkinetic collision integral (4.86). For this to be so, there would have to be a shorter effective collision time scale on which the hyperkinetic distribution function relaxed towards the closure (4.66) before continuing its relaxation towards a Lynden-Bell equilibrium on a longer time scale.

While conceptually fascinating, to be put on solid ground, the 'collisionless collision integals' must be verified. While there is substantial evidence of the validity (within certain regimes) of the Balescu–Lenard integral (4.105) describing 'true' Coulomb collisions, it remains to be seen whether the hyperkinetic collision integral, and, more generally and especially, the microgranulation ansatz, are valid. We therefore outline a number potential observable quantities that would be indicative of the collisionless relaxation. Most obviously, the verification of Lynden-Bell equilibria themselves would vindicate the assumption of phase-volume conservation despite the fairly stringent formal limitation on it imposed by (4.112) (the requirement that collisionless relaxation occur before the exact phase-space density is filamented down to collisional scales), and potentially an even more stringent limitation in systems where a dissipative anomaly is present and the phase-volume conservation is broken on a time scale independent of 'true' collisionality (see discussion and references in section 4.5.4). As discussed in section 4.3.3, the Lynden-Bell equilibria can be extremely varied, but there is a direct correspondence between the waterbag content of the initial condition and the ultimate equilibrium, which can clearly be tested. The relaxation process itself would bear telltale signs, should it obey the microgranulation ansatz. In section 4.5, we showed that a plasma obeying the microgranulation ansatz has an effective collision rate much higher than the true-collision rate associated with the Balescu–Lenard collision integral, and that the typical energy stored in the electric fluctuations in such a plasma

is much greater than the fluctuation level arising from Coulomb collisions. This creates a picture of collisionless relaxation mediated not by single-particle collisions but by the effective collisions of larger correlated volumes of phase space, which seed larger perturbations to the electric field.

In section 4.6, this collisionless relaxation is shown to have some curious consequences due to inter-species interactions. Most interestingly, electrons are dragged towards a mean velocity that is not necessarily the mean velocity of the ions, as the case of 'true' collisions, but a certain anomalous velocity associated with an anisotropic ion distribution, which can be non-zero even if the mean ion velocity is zero (e.g., for an ion distribution that carries a heat flux but no net momentum). A current and, therefore, magnetic field, could be spontaneously generated by this mechanism.

Still, thou art blest, compar'd wi' me!
The present only toucheth thee:
But Och! I backward cast my e'e,
On prospects drear!
An' forward tho' I canna see,
I guess an' fear!

— From 'To a mouse' by Robert Burns

5

Conclusion and future work

We have shown, in chapters 1 and 2, that the Lynden-Bell equilibria—the natural maximum-entropy states for a collisionless plasma—are highly universal and generically form distributions with power-law tails. Specific amongst these distributions appears to be those with the exponent -2 in energy, which we show, through numerical simulation in chapter 3, to be the generic profile achieved by the two-stream instability. Given the propensity for distributions with power-law tails in energy to occur in a wide variety of collisionless plasmas—including, but not limited to, the solar wind (e.g., Gloeckler et al. 2008; Fisk &Gloeckler 2014; Yang et al. 2020), the solar neighbourhood (e.g., Oka et al. 2018), and numerous numerical studies (e.g., Spitkovsky 2008; Sironi &Spitkovsky 2014; Werner &Uzdensky 2017; Zhdankin et al. 2017)—it is interesting to speculate whether the adjusted Lynden-Bell formalism proposed (and validated) here could be used to explain some of the observed equilibria. As has been abundantly clear since its inception, the strongest feature of the Lynden-Bell approach has also been its fundamental drawback: it is a thermodynamic theory and, while thermodynamic theories can be remarkably robust and successful, it is also difficult to predict how badly their underlying assumptions must be broken for the theory to fail completely. The idea of an evolving waterbag content $\rho(\eta)$ put forward in chapter 3 offers some promise for understanding how the system should relax. There is, however, currently no theory of how to compute

dynamically the time evolution of $\rho(\eta)$, which is the final piece of the puzzle. Since it is this evolution that determines how far from its ground state the system saturates (cf. the discussion of phase-space degeneracy in chapter 3) theoretical insights into the dynamical evolution of $\rho(\eta)$ are precisely what is required to make this thermodynamic theory into a predictive theory of plasma relaxation.

While the power law of ε^{-2} in energy is roughly consistent with with some of the aforementioned power laws seen in nature and close to, but distinct from, the inferred value for cosmic-ray sources (e.g., Ormes &Freier 1978; Reichherzer et al. 2021) it fails to capture the wide range of power laws seen both in numerical simulations (e.g., Sironi &Spitkovsky 2014; Zhdankin et al. 2017; Werner &Uzdensky 2017) and observationally. Such systems are usually turbulent, possibly inhomogeneous, and invariably magnetised. In principle, the theory favours turbulence; as we saw in chapter 3, it actively thrives on it. The question of inhomogeneous, or magnetised systems however, presents a much more interesting prospect. Lynden-Bell's original theory, by virtue of being applied to gravitating systems, was fundamentally inhomogeneous. This naturally makes the problem more difficult, as the potential (be it gravitational, or generated by another source, see, e.g., Levin et al. 2008) of the system may even need to be solved self consistently with the equilibria and it is not even clear that such solutions exist (cf. Ng &Bhattacharjee 2005). However, this presents a rich avenue of research given all previous studies of this type have been confined to using a small number of level sets.

With regards to the question of magnetised equilibria the most obvious straw to grasp at is that there is no guarantee that the invariants (1.8) and (1.9) are the only invariants respected on the relaxation time scale. In drift-kinetic plasmas for instance, each particle conserves its magnetic moment μ_b. This implies a new conserved function

$$\frac{2\pi}{V}\int \mathrm{d}\boldsymbol{x}\int \mathrm{d}v_\parallel\, B(\boldsymbol{x})\delta\big(\eta - f(\boldsymbol{x}, v_\parallel, \mu_b)\big) = \rho(\eta, \mu_b), \tag{5.1}$$

which would supersede the conservation of the now mundane $\rho(\eta)$. While some Gardner distributions have been studied for such systems (see Helander 2020; Mackenbach et al. 2022), the Lynden-Bell equilibria in them are unexplored and may contain a wealth of interesting physics. This said, in turbulent systems, the conservation law (5.1) may be just as fragile as the conservation of phase volume.

Indeed it has been suggested that the breaking of adiabatic invariance may be essential to understanding the transport properties of non-thermal particles (see Ruszkowski &Pfrommer 2023 and references therein).

From the previous discussion it becomes clear that perhaps the most relevant limitation of Lynden-Bell's statistical mechanics—or indeed of any equilibrium statistical mechanics—in application to observed plasma phenomena is that much of real plasma dynamics are out of equilibrium in a physically essential way: any local relaxation processes, collisional or collisionless, tend to have to be taken into account alongside various 'sources' and 'sinks' of particles and/or energy, e.g., the energisation and escape of cosmic rays (Schlickeiser 1989; Chandran 2000; Becker Tjus &Merten 2020; Hopkins et al. 2022; Kempski &Quataert 2022), the turbulent heating and radiative cooling of the intracluster medium (e.g., Zhuravleva et al. 2014 and references therein) or accretion-disc plasmas (e.g., Lesur 2021; Kawazura et al. 2022 and references therein), a veritable zoo of such processes in the solar wind (e.g., Verscharen et al. 2019; Chen et al. 2020) and the Earth's magnetosphere (e.g., Lucek et al. 2005), the birth of energetic α-particles in fusion reactions and their subsequent slowing down and escape from confined plasmas (e.g., Helander &Sigmar 2005; Mailloux et al. 2022), etc. In plasmas where Coulomb collisions can be assumed to relax the particle distribution quickly to a local Maxwellian, we have a robust analytical framework for handling all these non-equilibrium processes in terms of the evolution of the density, momentum and temperature of that Maxwellian and a separation of the dynamics into that evolution plus the turbulence of small fluctuations around the local equilibrium (e.g., Schekochihin et al. 2009; Abel et al. 2013). It is on this basis that we tackled the question of creating collisionless collision integrals in chapter 4. While those collision integrals may serve as mnemonics for interesting properties of the relaxation of collisionless systems, a number of elements went into their derivation that can be improved. Most important are the quasilinear assumption and the microgranulation ansatz (4.32). For either of these assumptions to be lifted, it is clear that a theory of correlations of fluctuations in turbulent plasmas must be concretely developed. The phase-space cascade theory of Adkins &Schekochihin (2018) and Nastac et al. (2024b), goes some way to achieving this but does not yet include the effects of self-consistency on the fluctuations, which may prove to be important. Should this be achieved, then it opens up the true

possibility of creating nonlinear collision integrals for which Lynden-Bell equilibria are the fixed points. We end this discussion here. On the problem of relaxation in collisionless plasmas there is plenty still to be done, and ample avenues by which to do it, but we have achieved some modicum of closure and universality.

Appendices

A

Justification for the Lynden-Bell entropy

In this appendix we will justify the use of the formula (1.6) for the Lynden-Bell entropy used throughout this thesis. On a surface level, the entropy is natural from Shannon's theorem—$P(\boldsymbol{x}, \boldsymbol{v}, \eta)$ is the probability of the system being in a given microstate, and the entropy is naturally of the form $P \ln P$, integrating over all possible microstates. Here we will derive it through the classical analogue of what lies on Boltzmann's tomb—the logarithm of the number of 'complexions'—in line with the derivation of Jaynes (2003).

We suppose that phase space of the plasma is divided up into very many cells (refered to as macro-cells by Lynden-Bell) and that in each macro-cell is further divided into $\mathcal{M}$ many micro-cells. The micro-cells are taken to be so infinitesimally small that there can be no reasonable variation of the exact phase-space density f over them. The macro-cells however, are taken to be sufficiently large that they contain very many micro-cells, but small enough that there can be little meaningful variation of the *mean* phase-space density $\langle f \rangle$ across them[1]. The cartoon of this set-up is shown in figure A, illustrating the micro- and macro-cells in the left hand panel

[1]Given the scale-invariant nature of the cascade proposed by Nastac et al. (2024b,a) and the subsequent reaching of collisions, it is not entirely clear that such a separation of micro- and macro-cell can ever be well defined, perhaps motivating the turbulent amnesia proposed in chapter 3

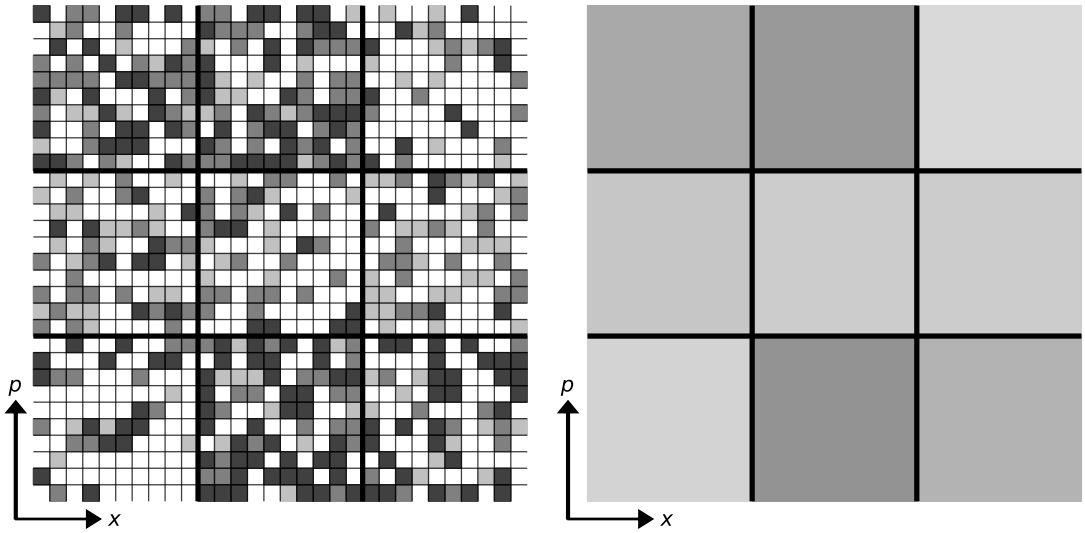

Figure A.1: Cartoon showing the distinction between Lynden-Bell's micro-cells and macro-cells, which are the small and large cells respectively. The left panel shows a single microstate of the system, many of which could give rise to the same macrostate in the right panel, which simply shows the mean phase-space over each macro-cell.

from which we, the imperfect observer, see the mean phase-space density in the right-hand panel. Suppose then that there are a discrete number of (non-empty) level sets of the system $\{\eta_1, \eta_2, ..., \eta_{\max}\}$. Due to phase-volume conservation, the total number of micro-cells of exact phase-space density η_i is a constant and can be labelled $\mathcal{N}_{\eta_i}$. Within any given macro-cell, indexed by $(\boldsymbol{x}, \boldsymbol{p})$, the number of micro-cells which should be filled with exact phase space density η_i, can be labelled as $\mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}$, so that

$$\sum_{(\boldsymbol{x},\boldsymbol{p})} \mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}} = \mathcal{N}_{\eta_i}. \tag{A.1}$$

From this, the probability $P(\boldsymbol{x}, \boldsymbol{p}, \eta_i)$ is simply the the fraction of the marco-cell which is occupied.

$$P(\boldsymbol{x}, \boldsymbol{p}, \eta_i) = \frac{\mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}}{\mathcal{M}}. \tag{A.2}$$

The entropy of a given distribution $P(\boldsymbol{x}, \boldsymbol{v}, \eta_i)$, in line with Jaynes (2003), is simply the (logarithm of) the number of distinct ways W of rearranging the micro-cells which give the same $P(\boldsymbol{x}, \boldsymbol{v}, \eta_i)$ for each macro-cell. This number of ways is given by

$$W = \frac{\prod_{\eta_i} \mathcal{N}_{\eta_i}!}{\prod_{\eta_i} \prod_{(\boldsymbol{x},\boldsymbol{p})} \mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}!} \prod_{(\boldsymbol{x},\boldsymbol{p})} W_{(\boldsymbol{x},\boldsymbol{p})} \tag{A.3}$$

where the left hand factor is the number ways, for each level set η_i, of dealing the required number of micro-cells, $\mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}$, out to each macro-cell while the second term, $W_{(\boldsymbol{x},\boldsymbol{p})}$ is the number of ways of rearranging the micro-cells within each macro-cell and is given by

$$W_{(\boldsymbol{x},\boldsymbol{p})} = \frac{\mathcal{M}!}{\left(\mathcal{M} - \sum_{\eta_i} \mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}\right)!} = \frac{\mathcal{M}!}{\mathcal{N}_{0,\boldsymbol{x},\boldsymbol{p}}!} \tag{A.4}$$

where in the last equality we have identified the number of empty micro-cells in a given macro-cell $\mathcal{N}_{0,\boldsymbol{x},\boldsymbol{p}} = \mathcal{M} - \sum_{\eta_i} \mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}$. Taking the logarithm of (A.3) and using Stirling's approximation, one arrives at

$$\ln W = \sum_{\eta_i \neq 0}^{\eta_{\max}} \mathcal{N}_{\eta_i} \left(\ln \mathcal{N}_{\eta_i} - 1\right) + \sum_{(\boldsymbol{x},\boldsymbol{p})} \mathcal{M} \ln \mathcal{M} - \sum_{\eta_i = 0}^{\eta_{\max}} \sum_{(\boldsymbol{x},\boldsymbol{p})} \mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}} \ln \mathcal{N}_{\eta_i,\boldsymbol{x},\boldsymbol{p}}, \tag{A.5}$$

from which we may use (A.2) to write the entropy in terms of P

$$\ln W = \sum_{\eta_i \neq 0}^{\eta_{\max}} \mathcal{N}_{\eta_i} \left(\ln \mathcal{N}_{\eta_i} - 1\right) - \sum_{\eta_i = 0}^{\eta_{\max}} \sum_{(\boldsymbol{x},\boldsymbol{p})} \mathcal{M} P(\boldsymbol{x},\boldsymbol{p},\eta_i) \ln P(\boldsymbol{x},\boldsymbol{p},\eta_i). \tag{A.6}$$

The first term is simply a constant related to the conservation of the waterbag content, while the second term is the familiar Gibbs–Shannon entropy applied to P, in discrete form. Passing to the continuous limit we find

$$\ln W = \text{const} - \frac{1}{\Delta\eta \Delta\Gamma} \int \mathrm{d}\boldsymbol{x}\, \mathrm{d}\boldsymbol{p}\, \mathrm{d}\eta\, P(\boldsymbol{x},\boldsymbol{p},\eta_i) \ln P(\boldsymbol{x},\boldsymbol{p},\eta_i). \tag{A.7}$$

Giving the Lynden-Bell entropy up to constant pre-multiplying factors. Here $\Delta\eta$ is the η spacing and $\Delta\Gamma$, the volume of a micro-cell. Here we have chosen the same notation as the correlation volume in chapter 4 on the grounds that the volume of a micro-cell is precisely this smallest volume over which the exact phase-space density should be correlated.

B

The degenerate limit of Lynden-Bell's statistics

In the main text, we have made use of the claim that, given the waterbag content $\rho(\eta)$, the Gardner distribution with the same waterbag content represents the ground state of all possible Lynden-Bell equilibria given by (2.4) that have this waterbag content. This is intuitive: should the initial condition of the system be a Gardner distribution, then that is, by definition, the only state available to the system, so it must also be the maximum-entropy state for that choice of $\rho(\eta)$. However, for completeness, and as a test of the Lynden-Bell formalism, it is prudent to check that the Gardner distribution can be recovered for some choice of the fugacity $F(\eta)$ and β. This is the aim of this appendix: to solve explicitly for β and $F(\eta)$ in (2.4) when $E = E_{\mathrm{G}}$, where E_{G} is the energy of a given Gardner distribution f_{G}, and to show that the mean phase-space density obtained by maximising Lynden-Bell's entropy is $\langle f \rangle = f_{\mathrm{G}}$.

To understand how the Gardner distribution will be recovered, we appeal to the familiar Fermi–Dirac distribution

$$f_{\mathrm{FD}}(\varepsilon) = \frac{\eta_{\max} e^{-\beta(\varepsilon-\mu)}}{1 + e^{-\beta(\varepsilon-\mu)}}. \tag{B.1}$$

To work out what this is in the degenerate limit, every textbook notes that, when β is very large, the numerator (and second term in the denominator) is either very

small for $\varepsilon > \mu$, making the expression approximately zero, or very large for $\varepsilon < \mu$, making the exponentials in the numerator and denominator approximately cancel to give $\eta_{\max}$. Borrowing this intuition, we anticipate that our solution should have $\beta \to \infty$. It is clear what kind of solution one must search for: the degeneracy parameter $D(\varepsilon)$ defined by (2.12) must be large wherever $f_{\rm G}(\varepsilon) \neq 0$ and zero wherever $f_{\rm G}(\varepsilon) = 0$. Furthermore, the dominant contribution to $D(\varepsilon)$ in the integral (2.12) must come from $\eta = f_{\rm G}(\varepsilon)$. This is the Lynden-Bell version of the statement that the phase space is completely filled up below the Fermi energy.

To see how this works in practice, let us posit the fugacity in the form

$$F(\eta) = \frac{1}{\bar{\eta}} \exp\left[\beta \int_{\eta_{\min}}^{\eta} \mathrm{d}\eta'\, f_{\rm G}^{-1}(\eta')\right] \tag{B.2}$$

and prove that, via (2.4) and (1.9), it recovers the Gardner distribution with the correct waterbag content $\rho(\eta)$ when $\beta \to \infty$, for a suitable choice of the dimensional constant $\bar{\eta}$.

In the denominator of (2.4), we must evaluate the integral

$$D(\varepsilon) = \frac{1}{\bar{\eta}} \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, e^{-\beta\eta\varepsilon + \ln\bar{\eta}F(\eta)}. \tag{B.3}$$

Since we are working in the limit of large β, this integral will be dominated by the contribution near the maximum (in η at fixed ε) of the exponent and thus can be evaluated by the method of steepest descent (Bender &Orszag, 1978). The location of the maximum of the exponent, which we denote by $\eta_{\rm stat}(\varepsilon)$, is given by the solution to

$$\beta \frac{\mathrm{d}}{\mathrm{d}\eta}\left[\eta\varepsilon - \int_{\eta_{\min}}^{\eta} f_{\rm G}^{-1}(\eta')\mathrm{d}\eta'\right] = 0 \implies \eta_{\rm stat}(\varepsilon) = f_{\rm G}(\varepsilon), \tag{B.4}$$

as we anticipated above. We may now expand the exponent of the integrand in (B.4) around this maximum to approximate the integral by

$$\begin{aligned} D(\varepsilon) &\simeq \int_{-\infty}^{\infty} \mathrm{d}\eta\, \exp\left\{-\beta f_{\rm G}(\varepsilon)\varepsilon + \frac{\beta}{2}\left.\frac{\mathrm{d}f_{\rm G}^{-1}}{\mathrm{d}\eta}\right|_{\eta=f_{\rm G}(\varepsilon)} [\eta - f_{\rm G}(\varepsilon)]^2\right\} F(f_{\rm G}(\varepsilon)) \\ &= e^{-\beta f_{\rm G}(\varepsilon)\varepsilon} F(f_{\rm G}(\varepsilon)) \sqrt{\frac{2\pi}{\beta}} \left[-\left.\frac{\mathrm{d}f_{\rm G}^{-1}}{\mathrm{d}\eta}\right|_{\eta=f_{\rm G}(\varepsilon)}\right]^{-1/2}. \end{aligned} \tag{B.5}$$

As is customary, we have neglected the contributions from higher derivatives of the exponent in the knowledge that they will contribute terms that are smaller by $\mathcal{O}(1/\beta)$. We have also replaced the upper and lower limits of integration by $\pm\infty$,

assuming that the exponential decays sufficiently fast for the presence of integration limits not to be noticed by the integral. This will not be accurate near $\varepsilon = 0$ and $\varepsilon = f_{\mathrm{G}}^{-1}(\eta_{\min})$, where the dominant contribution comes precisely from the limit of integration. This is, however, fine for $\beta \to \infty$ because the intervals in ε where this approximation is bad shrink as $\mathcal{O}\left(\beta^{-1/2}\right)$.

If we further demand that the constant $\bar{\eta}$ in (B.2) is chosen so that

$$\frac{1}{\bar{\eta}} e^{-\beta\eta_{\min} f_{\mathrm{G}}^{-1}(\eta_{\min})} \sqrt{\frac{2\pi}{\beta}} \left[-\left.\frac{\mathrm{d}f_{\mathrm{G}}^{-1}}{\mathrm{d}\eta}\right|_{\eta=\eta_{\min}} \right]^{-1/2} = 1, \tag{B.6}$$

then, neglecting terms of $\mathcal{O}(1/\beta)$, $D(\varepsilon)$ satisfies

$$1 + D(\varepsilon) \simeq \begin{cases} D(\varepsilon) & \text{for} \quad \varepsilon < f_{\mathrm{G}}^{-1}(\eta_{\min}), \\ 1 & \text{for} \quad \varepsilon > f_{\mathrm{G}}^{-1}(\eta_{\min}). \end{cases} \tag{B.7}$$

Calculating the mean phase-space density $\langle f \rangle(\boldsymbol{p})$ from (1.5) and (2.4), we find

$$\langle f \rangle(\varepsilon) \simeq \begin{cases} \dfrac{1}{\bar{\eta} D(\varepsilon(p))} \displaystyle\int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta\, \eta e^{-\beta\eta\varepsilon + \ln \bar{\eta} F(\eta)} & \text{for} \quad \varepsilon < f_{\mathrm{G}}^{-1}(\eta_{\min}), \\ 0 & \text{for} \quad \varepsilon > f_{\mathrm{G}}^{-1}(\eta_{\min}). \end{cases} \tag{B.8}$$

Let us prove that $\langle f \rangle(\boldsymbol{p}) = f_{\mathrm{G}}(\varepsilon(\boldsymbol{p}))$. The integral in (B.8) can again be computed by the method of steepest descent. The exponent is the same as in (B.3) and so will have the same maximum, at $\eta_{\mathrm{stat}} = f_{\mathrm{G}}(\varepsilon)$ (up to a small $\mathcal{O}(1/\beta)$ correction due to the factor of η in (B.8)). Expanding the integral around this maximum and neglecting all terms that are $\mathcal{O}(1/\beta)$ causes the factor of η in the integral to be replaced by $\eta_{\mathrm{stat}}(\varepsilon)$. After this, the remainder of the integral has the same form as (B.5), which cancels with the denominator of (B.8), leaving only the factor of $\eta_{\mathrm{stat}} = f_{\mathrm{G}}(\varepsilon)$. Q.E.D.

Thus, the fugacity (B.2) correctly recovers the Gardner distribution $f_{\mathrm{G}}(\varepsilon)$ in the limit of $\beta \to \infty$. However, it is possible, in principle, to have accidentally chosen a fugacity which, while recovering the correct distribution, has an incorrect waterbag content. To complete the proof, we compute the waterbag content (1.9) for the fugacity (B.2), to show that it is the same waterbag content as that of f_{G}:

$$\begin{aligned} \rho(\eta) \simeq \int_0^\infty \mathrm{d}\varepsilon\, g(\varepsilon) \exp\left\{ -\beta\varepsilon \left[\eta - f_{\mathrm{G}}(\varepsilon)\right] + \beta \int_{f_{\mathrm{G}}(\varepsilon)}^{\eta} f_{\mathrm{G}}^{-1}(\eta) \mathrm{d}\eta \right\} \\ \cdot \sqrt{\frac{\beta}{2\pi}} \left[-\left.\frac{\mathrm{d}f_{\mathrm{G}}^{-1}}{\mathrm{d}\eta}\right|_{\eta=f_{\mathrm{G}}(\varepsilon)} \right]^{1/2}. \end{aligned} \tag{B.9}$$

Again using the method of steepest descent, we observe that the exponent has its stationary point, this time in ε at fixed η, at $\varepsilon = f_{\mathrm{G}}^{-1}(\eta)$, just as we should expect. We can again expand the exponent around this stationary point and neglect $\mathcal{O}(1/\beta)$ terms, giving us

$$\begin{aligned}\rho(\eta) &\simeq \int_{-\infty}^{\infty} \mathrm{d}\varepsilon\, g(f_{\mathrm{G}}^{-1}(\eta)) \exp\left\{\frac{\beta}{2}\left.\frac{\mathrm{d}f_{\mathrm{G}}}{\mathrm{d}\varepsilon}\right|_{\varepsilon=f_{\mathrm{G}}^{-1}(\eta)} \left[\varepsilon - f_{\mathrm{G}}^{-1}(\eta)\right]^2\right\} \sqrt{\frac{\beta}{2\pi}} \left[-\frac{\mathrm{d}f_{\mathrm{G}}^{-1}}{\mathrm{d}\eta}\right]^{1/2} \\ &= -\, g\Big(f_{\mathrm{G}}^{-1}(\eta)\Big) \frac{\mathrm{d}f_{\mathrm{G}}^{-1}}{\mathrm{d}\eta}. \end{aligned} \tag{B.10}$$

This is exactly the expression (2.5) that we desire. This completes the proof that any Gardner distribution can be written as the $\beta \to \infty$ limit of the Lynden-Bell statistics, implying that one could, in principle, have discovered Gardner restacking just by analysing the Lynden-Bell equilibria at $\beta \to \infty$.

By retaining terms that are small in $\mathcal{O}(1/\beta)$, it is possible to analyse the Lynden-Bell equilibria analytically for energies very close but slightly above E_{G}. However, in analogy to Fermi–Dirac statistics, this should just amount to broadening the 'step' in the Gardner distribution around $\varepsilon = f_{\mathrm{G}}^{-1}(\eta_{\mathrm{min}})$, which limits the validity of this expansion to a very small range of energies, and thus makes such an expansion an exercise of infinitesimal utility.

C

Numerical method for solving the Lynden-Bell equilibria

In this appendix, we detail the numerical method by which we solve for the Lynden-Bell equilibria. To reiterate, the objective is, for a given waterbag content $\rho(\eta)$ and energy density E, to compute the function $F(\eta)$ and the parameter β such that, with the distribution function given by (2.4), the constraints (1.8) and (1.9) are satisfied to sufficient accuracy. For the numerical solutions given in Section 2.3, the numerical method was tailored to a 3D, non-relativistic, system where $\varepsilon = p^2/2m$, but the method can be easily extended to any systems with a specified density of states.

The formula (2.4) for the distribution function can be rewritten in such a way as to lend itself naturally to an iterative scheme. Namely, if one denotes the fugacity and the thermodynamic beta at the n^{th} iteration by $F^{(n)}$ and $\beta^{(n)}$, respectively, then the natural iteration for the fugacity is

$$F^{(n+1)}(\eta) = \rho(\eta) \left[2\pi \, (2m)^{3/2} \int_0^\infty \mathrm{d}\varepsilon \, \varepsilon^{1/2} \frac{e^{-\beta^{(n)}\eta\varepsilon}}{1 + \int_{\eta_{\min}}^{\eta_{\max}} \mathrm{d}\eta' \, e^{-\beta^{(n)}\eta'\varepsilon} F^{(n)}(\eta')} \right]^{-1} . \quad \text{(C.1)}$$

Ignoring for a moment how one iterates $\beta^{(n)}$, we note that if the iteration (C.1) converges, then by definition (1.9) is satisfied. This means that a solution with the correct waterbag content has been found, but it may have an incorrect energy, since it does not necessarily satisfy (1.8). However, by converging to a correct

fugacity for a given $\beta^{(n)}$, the problem is essentially reduced to a one-dimensional root-finding problem: finding β such that the energy takes the desired value, for which numerous numerical methods exist. This is the basis for our numerical algorithm, of which we will now give the specific details.

The η domain is discretised into $N_\eta = 10000$ points in preparation for future integration. To ensure that the lowest-density waterbags are well resolved without wasting resolution on the highest-density ones, a non-uniform discretisation in η is used: the j^{th} grid point is given by

$$\eta_j = \eta_{\min} + \left(\frac{j}{N_\eta}\right)^q (\eta_{\max} - \eta_{\min}), \tag{C.2}$$

where the number $q > 1$ is chosen depending on the waterbag content: for (2.26), $q = 3$ was used for $\sigma \leq 3$ and $q = 2$ for $\sigma > 3$, to compromise on the resolution near $\eta = \eta_{\max}$.

An initial guess is chosen for the fugacity and thermodynamic beta, denoted $F^{(0)}(\eta_j)$ and $\beta^{(0)}$, respectively. For the numerical solutions shown in Section 2.3, the initial guess was set to the analytical solutions (2.14) and (2.21) obtained in the non-degenerate limit; other arbitrary choices were tested, all of which converged, albeit usually taking more iterations to do so.

At each further iteration, the fugacity is updated according to (C.1). The η' integral in (C.1) is computed by a second-order midpoint method, interpolating the fugacity linearly between the neighbouring grid points. To compute the ε integral, a preliminary scan is first conducted to find the energy $\varepsilon_{\text{upper}}$ at which the degeneracy parameter (2.12) becomes smaller than 10^{-5}. This allows the energy integral in (C.1) to be split into two parts. The first of them, over energies below $\varepsilon_{\text{upper}}$, must be computed numerically, while the second, over energies above $\varepsilon_{\text{upper}}$, can be approximated by an analytically calculable function of $\varepsilon_{\text{upper}}$ and η. In the region where the integral must be computed numerically, the integration is carried out assuming the denominator to be piecewise linear on a momentum grid (rather than an energy grid, although this distinction is unimportant) linearly spaced with a spacing of 10^{-3} in units such that the $E_{\text{G}} = 1$. The fugacity can thus be iterated at fixed thermodynamic beta.

Once the integrated root-mean-square relative change in the fugacity over a single iteration is

$$\epsilon_F = \left[\frac{1}{\eta_{\max} - \eta_{\min}} \sum_{\eta_j} (\eta_{j+1} - \eta_j) \left(\frac{F^{(n+1)}}{F^{(n)}} - 1 \right)^2 \right]^{1/2} < 10^{-3}, \tag{C.3}$$

the energy of the resulting mean distribution can be computed. The root-finding method that we then use to determine β is an extremely primitive one: interval halving. The energy of the mean phase-space density with the fugacity resulting from above is computed, and $\beta^{(n+1)}$ is then increased or decreased depending on whether the computed energy is too high or too low, respectively. The initial step size in β is $\beta^{(0)}/2$. If the iteration in β passes over the root (i.e., in going from the n^{th} iteration to the $(n+1)^{\text{st}}$, the energy goes from above the correct energy to below it, or vice versa), then the step size in β is halved, so that it eventually homes in on the root correctly. The step size is also halved if the step would otherwise result in a negative value for $\beta^{(n+1)}$. If the computed energy is within a tolerance of 10^{-3} in units where $E_{\text{G}} = 1$, then the iteration in fugacity is allowed to proceed until ϵ_F finally falls below 10^{-5}, at which point the solution is considered converged.

D

Lynden-Bell equilibria and PIC plasmas

While we have shown numerically that the non-degenerate approximation of Lynden-Bell equilibria taken in section (2.2.3) is qualitatively accurate even in systems that are nowhere near complete non-degeneracy, there is one (admittedly contrived) case where it is not just approximately true but represents an exact result. In this appendix, we show that non-degenerate Lynden-Bell equilibria are the natural long-time equilibria of a plasma which is evolved using the PIC algorithm (PIC plasma) in which any given true species is represented by PIC particles with multiple different weights. The intuitive reason for this is that PIC particles behave in a manner analogous to the 'waterbags' central to the idea of Lynden-Bell relaxation. Waterbags are parcels of phase space, therefore containing some inherent number of true particles that move as a collective entity. PIC particles are hard-wired to represent such collections of true particles. To make this comparison more concrete, we will map a 'collisionless collision operator', that describes the relaxation of a system to a Lynden-Bell equilibrium (chapter 4) onto a numerical collision operator describing relaxation in a PIC plasma (Touati et al., 2022)—by mapping the physical picture of waterbags onto that of PIC particles.

The collisionless collision operator relaxes the probability density $P_\alpha(\boldsymbol{v}, \eta)$ of

species α as follows:

$$\frac{\partial P_\alpha}{\partial t} = \sum_{\alpha'} \frac{q_\alpha^2 q_{\alpha'}^2}{m_\alpha} \frac{\partial}{\partial \boldsymbol{v}} \cdot \int \mathrm{d}\boldsymbol{v}' \, \mathsf{Q}(\boldsymbol{v}, \boldsymbol{v}') \cdot \int \mathrm{d}\eta' \, \eta' \\ \left\{ \frac{\Delta\Gamma_{\alpha'}}{m_\alpha} \Big[\eta' - f_{\alpha'}(\boldsymbol{v}')\Big] P_{\alpha'}(\boldsymbol{v}', \eta') \left.\frac{\partial P_\alpha}{\partial \boldsymbol{v}}\right|_\eta - \frac{\Delta\Gamma_\alpha}{m_{\alpha'}} \Big[\eta - f_\alpha(\boldsymbol{v})\Big] P_\alpha(\boldsymbol{v}, \eta) \left.\frac{\partial P_{\alpha'}}{\partial \boldsymbol{v}'}\right|_{\eta'} \right\}, \quad \text{(D.1)}$$

where $\Delta\Gamma_\alpha$ is the typical volume over which a fluctuation in phase space is correlated (for details, see the discussion in section 2.2.5 or in chapter 4) and $\mathsf{Q}(\boldsymbol{v}, \boldsymbol{v}')$ is a tensor containing information about the interaction potential, whose explicit form we will not need here. The derivation of collision operators such as (D.1) is, naturally, subject to a number of approximations and caveats. Chief amongst these is the assumption of an electrostatic, quasilinear system in which phase volume is conserved. A full derivation and discussion of such collision operators can be found in, e.g., Chavanis (2022) or chapter 4.

In a PIC simulation, a given true species of particle may be represented by a number of different macroparticles that have different 'weights'—what this means quantitatively in our language, we shall explain shortly. To describe such a system, we set the distribution P_α to be discrete in η, the latter taking values $\eta_{\alpha,a}$ corresponding to macroparticle 'species' a:

$$P_\alpha(\boldsymbol{v}, \eta) = \sum_a P_{\alpha,a}(\boldsymbol{v}) \delta(\eta - \eta_{\alpha,a}). \quad \text{(D.2)}$$

The collision operator (D.1) then becomes

$$\frac{\partial P_{\alpha,a}}{\partial t} = \sum_{\alpha'} \frac{q_\alpha^2 q_{\alpha'}^2}{m_\alpha} \frac{\partial}{\partial \boldsymbol{v}} \cdot \int \mathrm{d}\boldsymbol{v}' \, \mathsf{Q}(\boldsymbol{v}, \boldsymbol{v}') \cdot \sum_{a'} \\ \left\{ \frac{\Delta\Gamma_{\alpha'} \eta_{\alpha',a'}}{m_\alpha} \Big[\eta_{\alpha',a'} - f_{\alpha'}(\boldsymbol{v}')\Big] P_{\alpha',a'}(\boldsymbol{v}') \frac{\partial P_{\alpha,a}}{\partial \boldsymbol{v}} - \frac{\Delta\Gamma_\alpha \eta_{\alpha',a'}}{m_{\alpha'}} \Big[\eta_{\alpha,a} - f_\alpha(\boldsymbol{v})\Big] P_{\alpha,a}(\boldsymbol{v}) \frac{\partial P_{\alpha',a'}}{\partial \boldsymbol{v}'} \right\}. \quad \text{(D.3)}$$

Here, to reiterate, $f_\alpha(\boldsymbol{v})$ is the mean phase-space density of particles of species α, which can be written as the sum of the mean phase-space densities $f_{\alpha,a}(\boldsymbol{v})$ of different macroparticle 'species':

$$f_\alpha(\boldsymbol{v}) = \sum_a f_{\alpha,a}(\boldsymbol{v}) = \sum_a \eta_{\alpha,a} P_{\alpha,a}(\boldsymbol{v}). \quad \text{(D.4)}$$

Next, one must note that PIC particles, like classical particles, occupy zero phase volume. Therefore, if all PIC particles are assumed decorrelated, $\Delta\Gamma_\alpha = 0$. However,

the phase-space densities $\eta_{\alpha,a}$ are then infinite. Mathematically, this corresponds to taking the limit of $\eta_{\alpha,a} \to \infty$ and $\Delta\Gamma_\alpha \to 0$ in (D.3) and (D.4) while holding the product $\Delta\Gamma_\alpha \eta_{\alpha,a}$—the number of particles in a correlated volume—fixed to the number of 'true' particles $\delta N_{\alpha,a}$ contained in a macroparticle of species (α, a); the quantity $\delta N_{\alpha,a}$ is what is usually called the 'weight' of the macroparticle in the PIC terminology.

Clearly, as $\eta_{\alpha,a}$ is taken to infinity, it is the mean phase-space density $f_{\alpha,a} = \eta_{\alpha,a} P_{\alpha,a}$ that remains finite. Making all these substitutions and taking the appropriate limit, one finds from (D.3) that it relaxes according to

$$\frac{\partial f_{\alpha,a}}{\partial t} = \sum_{\alpha',a'} \frac{q_\alpha^2 q_{\alpha'}^2}{m_\alpha} \frac{\partial}{\partial \boldsymbol{v}} \cdot \int \mathrm{d}\boldsymbol{v}' \, \mathsf{Q}(\boldsymbol{v}, \boldsymbol{v}') \cdot \left(\frac{\delta N_{\alpha'a'}}{m_\alpha} f_{\alpha',a'} \frac{\partial f_{\alpha,a}}{\partial \boldsymbol{v}} - \frac{\delta N_{\alpha,a}}{m_{\alpha'}} f_{\alpha,a} \frac{\partial f_{\alpha',a'}}{\partial \boldsymbol{v}'} \right). \tag{D.5}$$

Modulo the details of the tensor Q (due to the discrete nature of PIC codes), this is the effective collision operator due to numerical noise inherent in the PIC algorithm (Boris &Shanny, 1972; Birdsall &Langdon, 1985; Touati et al., 2022). It is easy to show that, provided the tensor $\mathsf{Q}(\boldsymbol{v}, \boldsymbol{v}')$ is positive definite and symmetric in $(\boldsymbol{v}, \boldsymbol{v}')$, the collision operator (D.5) has an H-theorem with the entropy

$$S = -\sum_{\alpha,a} \frac{1}{\delta N_{\alpha,a}} \int \mathrm{d}\boldsymbol{v} \, f_{\alpha,a}(\boldsymbol{v}) \ln f_{\alpha,a}(\boldsymbol{v}), \tag{D.6}$$

maximised by the equilibria

$$f_{\alpha,a} = N_{\alpha,a} e^{-\beta \delta N_{\alpha,a} \varepsilon(\boldsymbol{v})}, \tag{D.7}$$

where $N_{\alpha,a}$ is a normalisation constant set by the number of macroparticles of each weight (the fugacity of these marcoparticles). The resulting distribution function of particles of true species α is simply a superposition of Maxwellians:

$$f_\alpha(\boldsymbol{v}) = \sum_a f_{\alpha,a}(\boldsymbol{v}) = \sum_a N_{\alpha,a} e^{-\beta \delta N_{\alpha,a} \varepsilon(\boldsymbol{v})}, \tag{D.8}$$

which is manifestly the discrete form of the non-degenerate Lynden-Bell equilibrium (2.13).

Thus, non-degenerate Lynden-Bell equilibria could emerge organically in PIC simulations where multiple macroparticle weights represent the same true particle species. Of course, this is more a numerical artefact than a physical result. The equilibrium towards which such a system is pushed by the numerical noise is

effectively hard-coded by the choice of macroparticle weights. We note finally that the effects of a numerical collision operator such as (D.5) actually extend further than spurious consequences for the steady state. It is shown in chapter 4 that such collision operators could give rise to an anomalous interspecies drag, which again here would be a defect of the numerical method (cf. May et al. 2014).

E

Numerical details of two-stream simulations

Here, we provide some details about the set-ups of our numerical simulations of the relaxation of two-stream instabilities to Lynden-Bell equilibria. The simulations were conducted using the PIC code OSIRIS Fonseca et al. (2002) with an initial condition containing two thin beams:

$$f(\boldsymbol{x}, \boldsymbol{v}) = \begin{cases} \dfrac{n_0}{2\Delta v_\mathrm{b}}, & v_\mathrm{b} - \dfrac{\Delta v_\mathrm{b}}{2} < |v| < v_\mathrm{b} + \dfrac{\Delta v_\mathrm{b}}{2}, \\ 0, & \text{otherwise}, \end{cases} \tag{E.1}$$

where v_b is the beams' velocity and Δv_b is their width. For the simulation of the two-stream electron-positron instability, the same initial condition was used for both species. As OSIRIS is a fully relativistic code, a beam velocity of $v_\mathrm{b} = c/20$ was chosen, so that the simulation was approximately classical with the relativistic gamma factor $\gamma_\mathrm{rel} - 1 \sim 10^{-3}$. Both the simulation of the electron-only and electron-positron two-stream instabilities used a periodic 1D domain of size $25.6 d_\mathrm{e}$ (where d_e is the electron skin depth) consisting of 2^{15} cells. Both linear and quadratic interpolations for particles were checked, with no discernible difference. With these parameters, the energy and momentum were conserved up to diagnostic precision for the duration of the simulation. No smoothing was used on the fields or currents.

For convergence tests, the simulation of the electron-only two-stream instability was repeated on a 1D periodic domain of size $210 d_\mathrm{e}$ with $268,800$ cells, showing good convergence and much cleaner statistics. To our knowledge, this makes this the largest, longest-run, and most collisionless simulation of the electron-only two-stream instability to date. In the main text, figures 3.1 and 3.6 use the results of the smaller electron-only simulation, while figures 3.2, 3.3, 3.5, and 3.7 use the larger simulation. Despite the fact that the peak growth rate of the two-stream instability sits at $k \sim \sqrt{3/8}\lambda_\mathrm{D}^{-1} \sim \sqrt{3/8} d_\mathrm{e}^{-1} c/v_\mathrm{b}$, the domain size was chosen as a trade off between the need for resolution of sub-Debye physics and the desire for the multiple hole mergers seen in figure 3.1 to drive vigorous relaxation of the distribution function towards equilibrium. The number of particles per cell was chosen to be $20,000$ (in the case of the electron-positron simulation, $20,000$ of each species), giving the effective plasma parameter of $n_0 \lambda_\mathrm{D} \sim 10^6$. Naively, this would imply that the time that it would take for collisions to modify the mean distribution function in the absence of turbulence should be $\sim 10^6 \omega_\mathrm{pe}^{-1}$. As a result, our simulations should be considered collisionless as far as the evolution of $\langle f \rangle$ is concerned, despite the fact that $\rho(\eta)$ is modified on much shorter time scales, for the reasons explained in the main text.

In order to compute the waterbag content $\rho(\eta)$ of the exact phase-space density f, the PIC particles must be placed on a grid in (x, v) (since they are otherwise slivers with a finite extent in position space but zero width in velocity space). Therefore, we compute the fine-grained phase-space density f on evenly spaced rectangular bins with a width of 4 cells in position space and $0.0048 v_\mathrm{b}$ in velocity space in the interval between $v = \pm 0.25c$. The integral of the waterbag content

$$\Gamma(\eta) = \frac{1}{L} \int \mathrm{d}x \, \mathrm{d}v \, \Theta(f - \eta) = \int_\eta^\infty \mathrm{d}\eta' \, \rho(\eta') \tag{E.2}$$

is then computed assuming a piecewise-constant f on this fine-grained grid. This quantity is computed on a grid of 1000 logarithmically spaced values of η between the maximum phase-space density η_max (at that time step) and η_min corresponding to one particle per bin. We tested the results of this scheme for different bin sizes. The chosen bin size is a compromise between the need to have small bins in order to resolve fine-grained features of the phase-space density, and the need to have bins that are sufficiently large to capture many particles. Our results were insensitive

to increasing or decreasing the bin dimensions by up to a factor of 16. Due to the shape of the PIC particles in position space, it is also possible to change the lower limit for the η grid (as a fine-grained cell may contain a fraction of a particle). The results are insensitive to changing this value by a factor of 10 in either direction.

From $\Gamma(\eta)$ and the system's energy, one knows everything required for the Lynden-Bell equilibrium to be computed numerically with correct β and $F(\eta)$. This is done using the iterative scheme laid out in appendix C. The scheme is modified given that $\rho(\eta)$ is computed numerically, and therefore is not guaranteed to be smooth. This amounts to changing the second-order integration in η to first-order integration. This iteration continues until a Lynden-Bell equilibrium is found for which the normalised root-mean-square error in $\rho(\eta)$ is below 10^{-4} and the error in the energy (normalised to the Gardner energy) is below 10^{-3}.

F

Collision integrals in quantum plasmas

In this appendix, we show how the scheme used in sections 4.2 and 4.3 can be applied to work out swiftly the 'true' collision integrals for fermions and bosons. It is perhaps unsurprising that this is possible given the close analogy between phase-volume conservation and the exclusion principle. However, the procedure is non-rigorous because sections 4.2 and 4.3 assume point particles occupying definite positions in the $(\boldsymbol{x}, \boldsymbol{v})$ phase space which is not the case for quantum gases. This issue can be handled rigorously by the derivation and expansion of the Uehling &Uhlenbeck (1933) collision operator leading to the quantum Balescu–Lenard collision operator and its Landau siblings (see, e.g., Danielewicz 1980). We instead opt for the cheap and fast route to arrive at the (not-yet-closed) collision operator (4.35). This is written in terms of the unknown correlation volume $\Delta\Gamma_\alpha$ and the correlator $\langle g_\alpha^2 \rangle(\boldsymbol{v})$, which we can resolve by falling back on the quantum nature of our particles. To consider true collisions, we assume that particles are uncorrelated wherever this is compatible with quantum constraints (the exclusion principle). Crucially, however, due to the quantisation of their positions and momenta, particles occupy – and, therefore, must be correlated over – a finite phase space volume, $\Delta\Gamma_\alpha$. We can infer this volume by considering particles in a box of volume V, for which the possible momenta are

$$\boldsymbol{p} = \frac{2\pi\hbar}{V^{1/3}} \left(i_x, i_y, i_z\right), \tag{F.1}$$

where i_x, i_y, i_z are integers. The phase-volume per particle can then be read off from the spacing of their momenta:

$$\Delta\Gamma_\alpha = \left(\frac{2\pi\hbar}{m_\alpha}\right)^3 \equiv \eta_{0\alpha}^{-1}, \tag{F.2}$$

where we have defined the 'quantum density' $\eta_{0\alpha}$ (equivalent in spirit to the 'waterbag density' considered in section 4.3) to be the inverse of this correlation volume. The possible values of the 'exact phase-space density' f_α are then simply $\eta_{0\alpha}$ multiplied by the possible occupation numbers.

To calculate the correlator $\langle g_\alpha^2\rangle(\boldsymbol{v})$ (the variance of the initial condition, understood in the sense discussed in section 4.2.5), we will appeal to the closure scheme in section 4.3.2, which will account for the particle statistics. We consider all possible occupation numbers at velocity $\boldsymbol{v}$, and maximise entropy subject to knowing the mean phase-space density $f_{0\alpha}$ (equal to $\eta_{0\alpha}$ times the mean occupation number). The correlator $\langle g_\alpha^2\rangle(\boldsymbol{v})$ is then $\eta_{0\alpha}^2$ times the variance of the occupation number given this maximum-entropy assignment. We will consider the case where the bosons or fermions have σ_α possible spins (or indeed any quantum number giving rise to degeneracy). Then the partition function becomes

$$Z_\alpha = \sum_{n_1,n_2\ldots n_{\sigma_\alpha}} e^{-\psi_\alpha(\boldsymbol{v})\eta_{0\alpha}\sum_j n_j} = \left[\sum_n e^{-\psi_\alpha(\boldsymbol{v})\eta_{0\alpha}n}\right]^{\sigma_\alpha}, \tag{F.3}$$

where $\psi_\alpha(\boldsymbol{v})$ is a Lagrange-multiplier function chosen as in (4.50) to guarantee the correct mean phase-space density $f_{0\alpha}(\boldsymbol{v})$. The first sum in (F.3) is over all possible occupation numbers n_j of each of the possible spins. Since the effect on the phase-space density due to occupation numbers from different spins is additive, this is then split into the product of σ_α identical sums. The remaining sum is taken over the allowed values of the occupation number at each spin: $n \in \{0, 1\}$ for fermions and $n \in \{0, 1, 2, ...\}$ for bosons. The result is

$$Z_\alpha = \left[1 \pm e^{-\psi_\alpha(\boldsymbol{v})\eta_{0\alpha}}\right]^{\pm\sigma_\alpha}, \tag{F.4}$$

with the '+' sign for fermions and the '−' sign for bosons. As in (4.54), $\langle g_\alpha^2\rangle$ can now be computed thus:

$$\langle g_\alpha^2\rangle = \frac{\partial^2 \ln Z_\alpha}{\partial\psi_\alpha^2} = \left(\eta_{0\alpha} \mp \frac{f_{0\alpha}}{\sigma_\alpha}\right) f_{0\alpha}, \tag{F.5}$$

with the '−' sign for fermions and the '+' sign for bosons.

Using (F.5) in (4.35) gives the desired collision operator for charged particles obeying Bose–Einstein or Fermi–Dirac statistics:

$$\frac{\partial f_{0\alpha}}{\partial t} = \sum_{\alpha'} \frac{16\pi^3 q_\alpha^2 q_{\alpha'}^2}{m_\alpha V} \frac{\partial}{\partial \boldsymbol{v}} \cdot \sum_{\boldsymbol{k}} \frac{\boldsymbol{k}\boldsymbol{k}}{k^4} \cdot \int \mathrm{d}\boldsymbol{v}' \, \frac{\delta(\boldsymbol{k}\cdot(\boldsymbol{v}-\boldsymbol{v}'))}{\left|\epsilon_{\boldsymbol{k},\boldsymbol{k}\cdot\boldsymbol{v}}\right|^2} \left\{ \frac{1}{m_\alpha}\left[1 \mp \frac{f_{0\alpha'}(\boldsymbol{v}')}{\sigma_{\alpha'}\eta_{0\alpha'}}\right] f_{0\alpha'}(\boldsymbol{v}') \frac{\partial f_{0\alpha}}{\partial \boldsymbol{v}} - \frac{1}{m_{\alpha'}}\left[1 \mp \frac{f_{0\alpha}(\boldsymbol{v})}{\sigma_\alpha \eta_{0\alpha}}\right] f_{0\alpha}(\boldsymbol{v}) \frac{\partial f_{0\alpha'}}{\partial \boldsymbol{v}'} \right\}. \quad \text{(F.6)}$$

This clearly recovers the Balescu–Lenard collision operator (4.105) when the system is nowhere near degeneracy (i.e., when $f_\alpha \ll \sigma_\alpha \eta_{0\alpha}$). We further note that the results regarding strange relaxation in section 4.7 are equally applicable (and easier to interpret physically) for the collision operators (F.6). This implies an anomalous resistivity of quantum plasmas resulting from degeneracy prohibiting certain collisions. The modification to the electron-ion friction force in quantum plasmas has been studied before (see, e.g., Daligault 2016; Rightley &Baalrud 2021) in systems close to the relevant equilibrium (i.e., after isotropisation has removed the possibility of $\boldsymbol{u}_\mathrm{i}^\mathrm{a} \neq \boldsymbol{u}_\mathrm{i}$). We note, however, that, owing to the scaling of (F.2) with species mass, ions are conventionally less degenerate than electrons, so achieving $\boldsymbol{u}_\mathrm{i}^\mathrm{a} \neq \boldsymbol{u}_\mathrm{i}$ requires extreme densities.

Bibliography

Abel I. G., Plunk G. G., Wang E., Barnes M., Cowley S. C., Dorland W., Schekochihin A. A., 2013, Rep. Prog. Phys., 76, 116201

Adkins T., Schekochihin A. A., 2018, J. Plasma Phys., 84, 905840107

Aitchison L., Corradi N., Latham P. E., 2016, PLoS Comput. Bio., 12, 1

Amato E., Casanova S., 2021, J. Plasma Phys., 87, 845870101

Arad I., Johansson P. H., 2005, Mon. Not. R. Astron. Soc., 362, 252

Arad I., Lynden-Bell D., 2005, Mon. Not. Roy. Astron. Soc., 361, 385

Arrowsmith C. D., et al., 2024, Nat. Commun., 15, 5029

Assllani M., Fanelli D., Turchi A., Carletti T., Leoncini X., 2012, Phys. Rev. E, 85, 021148

Balescu R., 1960, Phys. Fluids, 3, 52

Banik U., Bhattacharjee A., Sengupta W., 2024, arXiv,, 2408.07127

Batchelor G. K., 1959, J. Fluid Mech., 5, 113

Beck C., Cohen E., 2003, Physica A, 322, 267

Becker Tjus J., Merten L., 2020, Phys. Rep., 872, 1

Bell A. R., 1978, Mon. Not. R. Astron. Soc., 182, 147

Bender C. M., Orszag S. A., 1978, Advanced Mathematical Methods for Scientists and Engineers. McGraw-Hill

Bernstein I. B., Greene J. M., Kruskal M. D., 1957, Phys. Rev., 108, 546

Birdsall C., Langdon A., 1985, Plasma Physics via Computer Simulation. CRC Press, United Kingdom

Birn J., Artemyev A. V., Baker D. N., Echim M., Hoshino M., Zelenyi L. M., 2012, Space Sci. Rev., 173, 49

Boltzmann L., 1896, Vorlesugnen über Gastheorie. Leipzig: J. A. Barth

Boris J., Shanny R., 1972, Proceedings of the 4th Conference on Numerical Simulation of Plasmas. Naval Research Laboratory, Washington

Bott A. F. A., Cowley S. C., Schekochihin A. A., 2024, J. Plasma Phys., 90, 975900207

Caprioli D., Spitkovsky A., 2014, Astrophys. J., 783, 91

Chandran B. D. G., 2000, Phys. Rev. Lett., 85, 4656

Chavanis P.-H., 2004, Physica A, 332, 89

Chavanis P.-H., 2005, Physica D, 200, 257

Chavanis P.-H., 2006a, Physica A, 359, 177

Chavanis P.-H., 2006b, Physica A, 365, 102

Chavanis P.-H., 2022, Physica A, 606, 128089

Chavanis P. H., Sommeria J., Robert R., 1996, Astrophys. J., 471, 385

Chen C., et al., 2020, Astrophys. J. Suppl. Ser., 246, 53

Comisso L., Sironi L., 2018, Phys. Rev. Lett., 121, 255101

Comisso L., Sironi L., 2022, Astrophys. J. Lett., 936, L27

Coonin A., 2014, From Marble to Flesh: The Biography of Michelangelo's David. Florentine Press, B'Gruppo

Crumley P., Caprioli D., Markoff S., Spitkovsky A., 2019, Mon. Not. R. Astron. Soc., 485, 5105

Cruz F., et al., 2018a, Nat. Phys., 14, 475

Cruz F., et al., 2018b, Nat. Phys., 14, 475

Daligault J., 2016, Phys. Plasmas, 23, 032706

Danielewicz P., 1980, Physica A, 100, 167

Davis S., Avaria G., Bora B., Jain J., Moreno J., Pavez C., Soto L., 2023, Phys. Rev. E, 108, 065207

Diamond P. H., Itoh S.-I., Itoh K., 2010, Modern Plasma Physics: Volume 1, Physical Kinetics of Turbulent Plasmas. Cambridge University Press, Cambridge

Dodin I., Fisch N., 2005, Phys. Lett. A, 341, 187

Dudík J., et al., 2017, Solar Phys., 292, 100

Dyson F., 2004, Nature, 427, 297

Ergun R. E., et al., 2020, Astrophys. J., 898, 153

Ewart R., Nastac M., Schekochihin A., 2023, J. Plasma Phys., 89, 905890516

Eyink G. L., 2018, Phys. Rev. X, 8, 041020

Ferrière K., 2019, Plasma Phys. Control. Fusion, 62, 014014

Fisk L. A., Gloeckler G., 2014, J. Geophys. Res. Space Phys., 119, 8733

Fonseca R. A., et al., 2002, in Computational Science — ICCS 2002. Springer Berlin-Heidelberg, p. 342

Gardner C. S., 1963, Phys. Fluids, 6, 839

Gibbs J. W., 1902, Elementary Principles In Statistical Mechanics Developed With Especial Reference To The Rational Foundation Of Thermodynamics. New York: Charles Scribner

Gloeckler G., Fisk L. A., Mason G. M., Hill M. E., 2008, AIP Conf. Proc., 1039, 367

Gruzinov A., Levin Y., Zhu J., 2020, Astrophys. J., 905, 11

Hartouni E., et al., 2022, Nat. Phys., 19, 1

Havrda J., Charvát F., 1967, Kybernetika, 3, 30

Helander P., 2017, J. Plasma Phys., 83, 715830401

Helander P., 2020, J. Plasma Phys., 86, 905860201

Helander P., Sigmar D. J., 2005, Collisional Transport in Magnetised Plasmas. Cambridge: Cambridge University Press

Hopkins P. F., Squire J., Butsky I. S., Ji S., 2022, Mon. Not. R. Astron. Soc., 517, 5413

Hosking D. N., Wasserman D., Cowley S. C., 2024, Metastability of stratified magnetohydrostatic equilibria and their relaxation (`arXiv:2401.01336`)

Hutchinson I. H., 2017, Phys. Plasmas, 24, 055601

Hutchinson I. H., 2024, arXiv,, 2407.08539

Jaynes E. T., 2003, Probability Theory: The Logic of Science. Cambridge: Cambridge University Press

Kadomtsev B. B., 1965, Plasma Turbulence. New York: Academic Press

Kadomtsev B. B., Pogutse O. P., 1970, Phys. Rev. Lett., 25, 1155

Karder M., 2007, Statistical Physics of Particles. Cambridge University Press, Cambridge

Kawazura Y., Schekochihin A., Barnes M., Dorland W., Balbus S., 2022, J. Plasma Phys., 88, 905880311

Kempski P., Quataert E., 2022, Mon. Not. R. Astron. Soc., 514, 657

Klimontovich Y. L., 1967, The Statistical Theory of Non-Equilibrium Processes in a Plasma. MIT Press

Knorr G., 1977, J. Plasma Phys., 17, 533

Kolmes E. J., Fisch N. J., 2020, Phys. Rev. E, 102, 063209

Kolmes E. J., Helander P., Fisch N. J., 2020, Phys. Plasmas, 27, 062110

Kolmogorov A., 1941, Doklady Akad. Nauk SSSR, 30, 301

Krall N. A., Trivelpiece A. W., 1973, Principles of Plasma Physics. McGrawHill, New York

Kunz M. W., Stone J. M., Quataert E., 2016, Phys. Rev. Lett., 117, 235101

Landau L. D., 1936, Zh. Eksp. Teor. Fiz., 7, 203

Lenard A., 1960, Ann. Phys., 10, 390

Lesur G. R. J., 2021, J. Plasma Phys., 87, 205870101

Levin Y., Pakter R., Teles T. N., 2008, Phys. Rev. Lett., 100, 040604

Levin Y., Pakter R., Rizzato F. B., Teles T. N., Benetti F. P., 2014, Phys. Rep., 535, 1

Livadiotis G., McComas D., 2009, J. Geophys. Res., 114, A11105

Livadiotis G., Desai M. I., Wilson III L. B., 2018, Astrophys. J., 853, 142

Lucek E., Constantinescu D., Goldstein M., Pickett J., Pincon J.-L., Sahraoui F., Treumann R., Walker S., 2005, Space Sci. Rev., 118, 95

Lynden-Bell D., 1967, Mon. Not. R. Astron. Soc., 136, 101

Mackenbach R. J. J., Proll J. H. E., Helander P., 2022, Phys. Rev. Lett., 128, 175001

Magee R. M., et al., 2019, Nat. Phys., 15, 281

Mailloux J., et al., 2022, Nucl. Fusion, 62, 042026

Malmberg J. H., Wharton C. B., 1964, Phys. Rev. Lett., 13, 184

Marsch E., 2006, Living Rev. Sol. Phys., 3, 1

Maxwell J. C., 1860, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 19, 19

May J., Tonge J., Ellis I., Mori W. B., Fiuza F., Fonseca R. A., Silva L. O., Ren C., 2014, Phys. Plasmas, 21, 052703

Moncuquet M., et al., 2020, Astrophys. J. Suppl. Ser., 246, 44

Mora T., Bialek W., 2011, J. Stat. Phys., 144, 268

Morse R. L., Nielson C. W., 1969, Phys. Rev. Lett., 23, 1087

Nastac M. L., Ewart R. J., Juno J., Barnes M., Schekochihin A. A., 2024a, in preparation

Nastac M. L., Ewart R. J., Sengupta W., Schekochihin A. A., Barnes M., Dorland W. D., 2024b, Phys. Rev. E, 109, 065210

Ng C. S., Bhattacharjee A., 2005, Phys. Rev. Lett., 95, 245004

Oka M., Krucker S., Hudson H. S., Saint-Hilaire P., 2015, Astrophys. J., 799, 129

Oka M., et al., 2018, Space Sci. Rev., 214

Ormes J., Freier P., 1978, Astrophys. J., 222, 471

Parra F. I., 2019, Collisional Plasma Physics. Lecture Notes for an Oxford MMathPhys course; URL: `http://www-thphys.physics.ox.ac.uk/people/FelixParra/\CollisionalPlasmaPhysics/CollisionalPlasmaPhysics.html`

Pierrard V., Lazar M., 2010, Sol. Phys., 267, 153

Qu K., Fisch N. J., 2024, Phys. Plasmas, 31, 062102

Reichherzer P., Merten L., Dörner J., Tjus J. B., Pueschel M. J., Zweibel E. G., 2021, SN Appl. Sci., 4

Rightley S., Baalrud S. D., 2021, Phys. Rev. E, 103, 063206

Robert R., Sommeria J., 1991, J. Fluid Mech., 229, 291

Roberts K. V., Berk H. L., 1967, Phys. Rev. Lett., 19, 297

Rostoker N., 1961, Nucl. Fusion, 1, 101

Ruszkowski M., Pfrommer C., 2023, Aston. Astrophys. Rev., 31, 4

Schekochihin A. A., Cowley S. C., Dorland W., Hammett G. W., Howes G. G., Plunk G. G., Quataert E., Tatsuno T., 2008, Plasma Phys. Control. Fusion, 50, 124024

Schekochihin A. A., Cowley S. C., Dorland W., Hammett G. W., Howes G. G., Quataert E., Tatsuno T., 2009, Astrophys. J. Suppl. Ser., 182, 310

Schlickeiser R., 1989, Astrophys. J., 336, 243

Schwab D. J., Nemenman I., Mehta P., 2014, Phys. Rev. Lett., 113, 068102

Servidio S., et al., 2017, Phys. Rev. Lett., 119, 205101

Severne G., Luwel M., 1980, Astrophys. Space Sci., 72, 293

Shannon C. E., 1948, The Bell System Technical Journal, 27, 379

Sironi L., Spitkovsky A., 2010, Astrophys. J., 726, 75

Sironi L., Spitkovsky A., 2014, Astrophys. J., 783, L21

Spitkovsky A., 2008, Astrophys. J., 682, L5

Spitzer L., 1967, Physics of Fully Ionized Gases. New York: Interscience Publishers

Spitzer L., Härm R., 1953, Phys. Rev., 89, 977

Su C. H., Oberman C., 1968, Phys. Rev. Lett., 20, 427

Swanson D. G., 2008, Plasma Kinetic Theory. New York: Chapman and Hall

Tatsuno T., Dorland W., Schekochihin A. A., Plunk G. G., Barnes M., Cowley S. C., Howes G. G., 2009, Phys. Rev. Lett., 103, 015003

Touati M., Codur R., Tsung F., Decyk V. K., Mori W. B., Silva L. O., 2022, Plasma Phys. Control. Fusion, 64, 115014

Tremaine S., Henon M., Lynden-Bell D., 1986, Mon. Not. Roy. Astron. Soc., 219, 285

Tsallis C., 1988, J. Stat. Phys., 52, 479

Uehling E. A., Uhlenbeck G. E., 1933, Phys. Rev., 43, 552

Uzdensky D. A., 2022, J. Plasma Phys., 88, 905880114

Verscharen D., Klein K. G., Maruca B. A., 2019, Living Rev. Sol. Phys., 16, 5

Werner G. R., Uzdensky D. A., 2017, Astrophys. J. Lett., 843, L27

Werner G. R., Uzdensky D. A., 2021, J. Plasma Phys., 87, 905870613

Wilson L. B., Goodrich K. A., Turner D. L., Cohen I. J., Whittlesey P. L., Schwartz S. J., 2022, Front. Astron. Space Sci., 9, 369

Yang L., et al., 2020, Astrophys. J. Lett., 896, L5

Ye H., Morrison P. J., 1992, Phys. Fluids B, 4, 771

Zhdankin V., 2021, Astrophys. J., 922, 172

Zhdankin V., 2022a, Phys. Rev. X, 12, 031011

Zhdankin V., 2022b, J. Plasma Phys., 88, 175880303

Zhdankin V., Werner G. R., Uzdensky D. A., Begelman M. C., 2017, Phys. Rev. Lett., 118, 055103

Zhdankin V., Uzdensky D. A., Werner G. R., Begelman M. C., 2019, Phys. Rev. Lett., 122, 055101

Zhuravleva I., et al., 2014, Nature, 515, 85