

Non-thermal particle acceleration and power-law tails via relaxation to universal Lynden-Bell equilibria

R. J. Ewart^{1,2†}, M. L. Nascac^{1,3}, and A. A. Schekochihin^{1,4}

¹Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford, OX1 3PU, UK

²Balliol College, Oxford, OX1 3BJ

³St John's College, Oxford, OX1 3JP

⁴Merton College, Oxford, OX1 4JD

(Received xx; revised xx; accepted xx)

Collisionless and weakly collisional plasmas often exhibit non-thermal quasi-equilibria. Among these quasi-equilibria, distributions with power-law tails are ubiquitous. It is shown that the statistical-mechanical approach originally suggested by Lynden-Bell (1967) can easily recover such power-law tails. Moreover, we show that, despite the apparent diversity of Lynden-Bell equilibria, a generic form of the equilibrium distribution at high energies is a 'hard' power-law tail $\propto \varepsilon^{-2}$, where ε is the particle energy. The shape of the 'core' of the distribution, located at low energies, retains some dependence on the initial condition but it is the tail (or 'halo') that contains most of the energy. Thus, a degree of universality exists in collisionless plasmas.

1. Introduction

It is well known that the ultimate fate of a homogeneous collisional plasma is to become a Maxwellian. This result was first inferred for neutral particles by Maxwell (1860) on statistical grounds and given solid dynamical foundation by Boltzmann (1896) with his collision integral. Plasma physics was to wait for Landau (1936) and later Balescu (1960) and Lenard (1960) to be equipped with its own collision integral, and the resulting universality. Nevertheless, distributions with power-law tails, a far cry from Maxwellian equilibria, are observed 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.* 2018; Hartouni *et al.* 2022).

That such non-Maxwellian distributions emerge should perhaps come as no surprise. In a plasma, the timescale of relaxation to Maxwellian equilibrium is associated with two-body Coulomb collisions, but, due to the long-range nature of the forces involved, the plasma may evolve, exchanging energy between fields and particles, on much shorter timescales. Indeed, in the absolute absence of collisions, the Vlasov equation has an infinite set of non-linearly stable equilibria: all distributions that are monotonically decreasing functions of particle energy are certainly stable (Gardner 1963). However, the set of all monotonically decreasing functions of energy is very large, and is certainly not exhausted by the Maxwellian equilibrium, which depends on only two parameters. It is, therefore, an outstanding challenge to determine whether any of these stable collisionless

† Email address for correspondence: robert.ewart@physics.ox.ac.uk

equilibria are naturally favoured by the dynamics of the system in a way that is not sensitive to initial conditions, i.e., whether a degree of universality exists in collisionless plasmas. It is certainly the case that nature appears to prefer distributions with power-law tails, and 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, 2022*b*).

In addition to suggesting dynamical paths towards distributions with power-law tails, 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 (or α -structural entropy: see 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 correlations that therefore should be non-extensive (see, e.g., Livadiotis & McComas 2009; Pierrard & Lazar 2010 and references therein). While this model produces good fits to observed distributions, it 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 some recent progress suggesting that this additional physics might be deducible from free-energy considerations: Zhdankin 2022*a,b*).

An early attempt to tackle the question of entropy in collisionless systems was made by Lynden-Bell (1967). Let us consider a system of N particles with canonical positions \mathbf{r}_i and momenta \mathbf{p}_i that evolve subject to a Hamiltonian $\mathcal{H}(\mathbf{r}_i, \mathbf{p}_i)$. Such a system can be said to be ‘collisionless’ if the evolution equation for the single-particle distribution function $f(\mathbf{r}, \mathbf{p})$ is well approximated by an effective Hamiltonian $\mathcal{H}^{\text{eff}}(\mathbf{r}, \mathbf{v})$ acting on a single particle (i.e., if the mean-field dynamics are a sufficiently good approximation to the true dynamics), viz.,

$$\frac{\partial f}{\partial t} + \frac{\partial \mathcal{H}^{\text{eff}}}{\partial \mathbf{p}} \cdot \frac{\partial f}{\partial \mathbf{r}} - \frac{\partial \mathcal{H}^{\text{eff}}}{\partial \mathbf{r}} \cdot \frac{\partial f}{\partial \mathbf{p}} = 0. \quad (1.1)$$

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 (1.1) as the collisionless Vlasov equation for a plasma, which could be electrostatic or electromagnetic in a non-relativistic or relativistic regimes. The collisionless dynamics described by (1.1) conserve an infinite number of invariants, equivalent to conserving the volume of level sets of the distribution function $f(\mathbf{r}, \mathbf{p})$ in phase space. 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 $f(\mathbf{r}, \mathbf{p})$ 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—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.

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 2006; 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; Ewart *et al.* 2022). However, the main strength of the theory is also its weakness. Unlike in the non-extensive entropy formulations, 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 depend (seemingly, in a complicated way) on an infinite family of invariants (sometimes referred to as ‘Casimirs’). 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.

To see just how non-universal Lynden-Bell equilibria can be, one only needs to consider the relation between the Lynden-Bell equilibria and the aforementioned 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†. 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 paper, 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.

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 an 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

† This takes a surprising amount of work to show formally: see Appendix A.

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 (this is qualitatively similar to the formalism of ‘superstatistics’; cf. Beck & Cohen 2003). 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 paper is organised as follows. In Section 2, we will review briefly the Lynden-Bell formalism, to state the problem and establish notation. We will then proceed to perform a systematic exploration of the nature of the Lynden-Bell equilibria. In Section 3.1, 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’ distribution with a tail). In Section 3.2, 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 3.3, 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 4, 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, virtually all reasonable 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 5, we summarise our findings and discuss their implications for real (observed) plasmas.

2. Lynden-Bell’s statistical mechanics

In this section, we present a brief re-derivation of Lynden-Bell’s equilibria as applied to a homogeneous (plasma) system. We begin with the collisionless Vlasov equation (1.1) evolving a single species of particles. As well as particle number, momentum, and energy, (1.1) conserves an infinite number of ‘Casimir’ invariants, e.g., the volume of phase space where the exact phase-space density is greater than a given value η :

$$\Gamma(\eta) = \iint d\mathbf{r} d\mathbf{p} H(f(\mathbf{r}, \mathbf{p}) - \eta) = \text{const}, \quad (2.1)$$

where $H(x)$ 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, which prompted Lynden-Bell to consider the exact phase-space density $f(\mathbf{r}, \mathbf{p})$ as a random field. Therefore, one may introduce the probability density $P(\mathbf{r}, \mathbf{p}, \eta)$ for the exact phase-space density $f(\mathbf{r}, \mathbf{p})$ to take the value η at position (\mathbf{r}, \mathbf{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. The mean phase-space density is then

$$\langle f \rangle(\mathbf{p}) = \int d\eta \eta P(\mathbf{p}, \eta). \quad (2.2)$$

Here, we have applied the intuition that the steady-state distribution function $P(\mathbf{p}, \eta)$ 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, $P(\mathbf{p}, \eta)$ will maximise the Gibbs-Shannon entropy

$$S = - \iint d\mathbf{p} d\eta P(\mathbf{p}, \eta) \ln P(\mathbf{p}, \eta). \quad (2.3)$$

Note that the integral in η must run over all the possible values, including $\eta = 0$ (the empty waterbag).

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

$$\int d\eta P(\mathbf{p}, \eta) = 1. \quad (2.4)$$

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

$$\iint d\mathbf{p} d\eta \varepsilon(\mathbf{p}) \eta P(\mathbf{p}, \eta) = E = \text{const}, \quad (2.5)$$

where $\varepsilon(\mathbf{p})$ is the energy of a particle as a function of its momentum p , which, unless otherwise specified, we will keep general to allow our particles to be relativistic. Next, we enforce the conservation of the Casimir invariants (2.1) by requiring that the volume-integrated probability of each waterbag stays constant, viz.,

$$\int d\mathbf{p} P(\mathbf{p}, \eta) = \rho(\eta) = \text{const}. \quad (2.6)$$

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.,

$$\rho(\eta) = \frac{1}{V} \iint d\mathbf{r} d\mathbf{p} \delta(\eta - f(\mathbf{r}, \mathbf{p}, t = 0)) = -\frac{1}{V} \frac{d\Gamma}{d\eta}, \quad (2.7)$$

where V is the system's volume. *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 (2.3) subject to the constraints (2.4), (2.5), and (2.6)[†] is

[†] We note that, while we have endowed the invariants (2.4), (2.5), and (2.6) 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 timescale are (2.4), (2.5), and (2.6).

equivalent to maximising, unconditionally, the functional

$$S[P(\mathbf{p}, \eta)] - \int d\mathbf{p} \lambda(\mathbf{p}) \left[\int d\eta P(\mathbf{p}, \eta) - 1 \right] - \beta \left[\iint d\mathbf{p} d\eta \varepsilon(\mathbf{p}) \eta P(\mathbf{p}, \eta) - E \right] + \beta \int d\eta \eta \mu(\eta) \left[\int d\mathbf{p} P(\mathbf{p}, \eta) - \rho(\eta) \right], \quad (2.8)$$

where $\lambda(\mathbf{p})$, β and $-\beta\eta\mu(\eta)$ are Lagrange multipliers. By analogy with textbook statistical mechanics, we will sometimes refer to $\mu(\eta)$ as ‘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

$$P(\mathbf{p}, \eta) = \frac{e^{-\beta\eta[\varepsilon(\mathbf{p})-\mu(\eta)]}}{\int d\eta' e^{-\beta\eta'[\varepsilon(\mathbf{p})-\mu(\eta')]}} \quad (2.9)$$

where $\lambda(\mathbf{p})$ has been computed explicitly to arrange for correct normalisation (2.4), whereas β and $\mu(\eta)$ must be chosen in such a way as to satisfy the constraints (2.5) and (2.6). We note that, despite (2.6), the mean phase-space density $\langle f \rangle$, given by (2.2), will, in general, not have the same level sets (2.1) as the exact one f (since $\langle f \rangle$ is an averaged quantity).

The similarity between the Lynden-Bell equilibria (2.9) and the Fermi–Dirac distribution is immediately apparent. 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 (2.7)), solve two coupled integral equations (2.5) and (2.6) with $P(\mathbf{p}, \eta)$ given by (2.9), determine β and $\mu(\eta)$, and substitute back into (2.9). Before considering the numerical solutions of this problem in section 4, we will first seek to understand the system analytically.

3. Theory: degenerate and non-degenerate equilibria

3.1. Lynden-Bell equilibria as excited Gardner distributions

Before leaping headlong into the analytical solution, it is instructive to understand the properties of the waterbag-content function (2.7). To get a feel for waterbag contents of typical initial conditions, one could compute the integral (2.7) 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 (2.7), any volume-preserving transformation of the coordinates (\mathbf{r}, \mathbf{p}) , including those transformations that ‘splice’ the phase space discontinuously, will leave the waterbag content unchanged. This is unsurprising because the true, 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, 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, called the Gardner distribution, that has that waterbag




Figure 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.

content but is a monotonically decreasing function of energy and, therefore, has the minimum possible energy associated with that waterbag content (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 (2.9) for a particular choice of β and $F(\eta)$, although the proof is technical and left to Appendix A.

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.

3.2. 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 η_{\min} . This mathematical convenience will turn out to be a physical necessity. Thankfully, while it is mathematically and physically important that the cutoff η_{\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.,

$$f_G(\mathbf{p}) = \begin{cases} \eta_{\max} e^{-\varepsilon(\mathbf{p})/\varepsilon_0} & \text{for } \varepsilon(\mathbf{p}) < \varepsilon_0 \ln \frac{\eta_{\max}}{\eta_{\min}}, \\ 0 & \text{for } \varepsilon(\mathbf{p}) > \varepsilon_0 \ln \frac{\eta_{\max}}{\eta_{\min}}. \end{cases} \quad (3.1)$$

Besides η_{\min} , the parameters of this distribution are the energy scale ε_0 and the maximum phase-space density η_{\max} . The latter is straightforwardly related to the particle density, e.g., $n_0 = \eta_{\max} (2\pi m \varepsilon_0)^{3/2}$ in the limit $\eta_{\min}/\eta_{\max} \rightarrow 0$ for a 3D, non-relativistic plasma, where $\varepsilon(\mathbf{p}) = p^2/2m$. The waterbag content of the Gardner distribution for such a plasma is then, from (2.7),

$$\rho(\eta) = \Gamma_{\text{free}} \delta(\eta) + \begin{cases} \frac{2n_0}{\sqrt{\pi} \eta_{\max} \eta} \left(\ln \frac{\eta_{\max}}{\eta} \right)^{1/2} & \text{for } \eta_{\min} < \eta < \eta_{\max}, \\ 0 & \text{otherwise,} \end{cases} \quad (3.2)$$

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 equilibria 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 (3.2). When it comes to solving (2.6) with $P(\mathbf{p}, \eta)$ given by (2.9), in order to find $\mu(\eta)$, this delta function can be accommodated by writing the chemical potential as

$$e^{\beta \eta \mu(\eta)} = \eta_{\text{ref}} \delta(\eta) + \begin{cases} \eta_{\text{ref}} F(\eta) & \text{for } \eta_{\min} < \eta < \eta_{\max}, \\ 0 & \text{otherwise,} \end{cases} \quad (3.3)$$

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 (2.9)—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 (3.3) results in the following expression for the Lynden-Bell equilibrium (2.9):

$$P(\mathbf{p}, \eta) = \frac{\delta(\eta) + e^{-\beta \eta \varepsilon(\mathbf{p})} F(\eta)}{1 + \int_{\eta_{\min}}^{\eta_{\max}} d\eta' e^{-\beta \eta' \varepsilon(\mathbf{p})} F(\eta')}. \quad (3.4)$$

In (3.4), the first, $\delta(\eta)$, term in the numerator accounts for the probability density of finding phase space to be empty at a given location, whereas the second term accounts




Figure 2: (a) Three example Gardner distribution functions 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 η_{\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 η_{\min} is justified in Section 3.5. 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.

for non-empty waterbags. Already Γ_{free} has dropped out of the calculation, as it must, and it is safe to let $\Gamma_{\text{free}} \rightarrow \infty$. 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 (3.4) 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 (3.2) is the η^{-1} power-law behaviour. While the specific form of the logarithmic factor in (3.2) 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 $\propto \exp[-(\varepsilon/\varepsilon_0)^\sigma]$, for some $\sigma > 0$, one will find a waterbag content with an η^{-1} asymptotic at $\eta \ll \eta_{\max}$.

To see this, we note that, since $f_G(\varepsilon)$ is monotonically decreasing with energy, it has a well-defined inverse $f_G^{-1}(\eta)$. In terms of this inverse, one can explicitly express the waterbag content (2.7) as

$$\rho(\eta) = \int d\varepsilon g(\varepsilon) \delta(\eta - f_G(\varepsilon)) = -g(f_G^{-1}(\eta)) \frac{df_G^{-1}}{d\eta} \implies \frac{df_G}{d\varepsilon} = -\frac{g(\varepsilon)}{\rho(f_G(\varepsilon))}, \quad (3.5)$$

where $g(\varepsilon)$ is the density of states in energy. The first of these expressions 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 $f_G(\varepsilon)$, the dominant way in which the inverse function will depend on η will be logarithmic, which, after differentiation in (3.5), will give an η^{-1} asymptotic multiplied by some logarithmic function of η . To illustrate this, in Figure 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, let us present a somewhat more rigorous argument in support of the η^{-1} scaling of $\rho(\eta)$. Let us consider the following integrals of the waterbag content (equivalent to the Casimir invariants introduced by Zhdankin 2022a):

$$\int_{\eta_{\min}}^{\eta_{\max}} d\eta \eta^\gamma \rho(\eta) = \int_{\eta > \eta_{\min}} d\mathbf{p} f_G^\gamma(\mathbf{p}), \quad (3.6)$$

where we have used (2.7), and the integral in momentum space is taken over everywhere where $\eta > \eta_{\min}$. Clearly, for $\gamma = 0$, (3.6) is the volume of the occupied momentum space. This must diverge as $\eta_{\min} \rightarrow 0$, and this divergence must be logarithmic for a Gardner distribution with any form of exponential tail. In contrast, for $\gamma > 0$, the integral of a power of an exponentially decaying function is highly insensitive to its cutoff at low values, and thus the integral of $\rho(\eta)$ must likewise be insensitive to η_{\min} as $\eta_{\min} \rightarrow 0$. The class of functions $\rho(\eta)$ that satisfy these two conditions are ones of the form

$$\rho(\eta) = \frac{n_0}{\eta_{\max} \eta} G\left(\frac{\eta}{\eta_{\max}}\right) \quad \text{for } \eta_{\min} < \eta < \eta_{\max}, \quad (3.7)$$

where $G(x)$ is a dimensionless function for which

$$\lim_{x \rightarrow 0} x^\epsilon G(x) = \begin{cases} 0 & \text{if } \epsilon > 0, \\ \infty & \text{if } \epsilon < 0. \end{cases} \quad (3.8)$$

This concludes the argument that Gardner distributions with exponential tails have waterbag contents with a universal low- η asymptotic.

It is a simple extension of this argument to see what the waterbag content will be for Gardner distributions with non-exponential tails at high energies (low η). Should the Gardner distribution have compact support even in the absence of η_{\min} , i.e., the distribution goes to zero at some finite energy such that (3.6) with $\gamma = 0$ is finite as $\eta_{\min} \rightarrow 0$, then the low- η asymptotic of $\rho(\eta)$ will be of the form $\eta^{-\delta}$ with some $\delta < 1$. If, instead, the Gardner distribution only has a finite number of moments, i.e., if it decays as a power law in energy, then there will exist some $\gamma > 0$ for which (3.6) will diverge as $\eta_{\min} \rightarrow 0$, implying that $\rho(\eta)$ must have the asymptotic $\eta^{-\delta}$ with $\delta > 1$. Hence we may write the waterbag content of these stranger Gardner distribution functions as

$$\rho(\eta) = \frac{n_0}{\eta_{\max}^{2-\delta} \eta^\delta} G\left(\frac{\eta}{\eta_{\max}}\right) \quad \text{for } \eta_{\min} < \eta < \eta_{\max}. \quad (3.9)$$

The explicit value of δ can be found from (3.5). 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.,

$$g(\varepsilon) = A\varepsilon^a, \quad (3.10)$$

where A is an appropriate constant with dimensions $[n_0/\eta_{\max}\varepsilon^{1+a}]$ and a is a real number. The assumption (3.10) is not too restrictive as it can capture both non-relativistic and ultra-relativistic systems of any dimensionality. With the assumption (3.10), we may now use (3.5) to link the value of δ to the high-energy asymptotic of the Gardner distribution. If the Gardner distribution has the power-law tail $f_G(\mathbf{p}) \propto \varepsilon(\mathbf{p})^{-(x+a)}$ at high energies, then, via (3.5), one finds $\delta = 1 + (a+1)/(a+\chi)$. If, instead, the Gardner distribution

goes to zero at some finite energy ε_{\max} so that $f_G(\mathbf{p}) \propto [\varepsilon_{\max} - \varepsilon(\mathbf{p})]^\chi$ near $\varepsilon = \varepsilon_{\max}$, then $\delta = 1 - 1/\chi$.

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 (2.5) and (2.6) in light of these facts.

3.3. 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 A).

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.,

$$D(\varepsilon) = \int_{\eta_{\min}}^{\eta_{\max}} d\eta e^{-\beta\eta\varepsilon} F(\eta) \ll 1. \quad (3.11)$$

We shall call $D(\varepsilon)$ the ‘degeneracy parameter’ since, from (3.4), the probability that a position in phase space with a given energy is non-empty is given by the quotient $D(\varepsilon)/[1 + D(\varepsilon)]$. In the limit (3.11), the distribution function (3.4) can be approximated by

$$P(\mathbf{p}, \eta) \simeq \delta(\eta) [1 - D(\varepsilon(\mathbf{p}))] + e^{-\beta\eta\varepsilon(\mathbf{p})} F(\eta). \quad (3.12)$$

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 (3.12) of the Lynden-Bell distribution makes computing the momentum-space integral in (2.6) and determining the fugacity $F(\eta)$ a simple matter. Substituting (3.12) into (2.6) and using the explicit form (3.10) of the density of states, we find the fugacity in the non-degenerate limit to be

$$F(\eta) = \frac{(\beta\eta)^{1+a}}{A\Gamma(a+1)} \rho(\eta), \quad \eta_{\min} \leq \eta \leq \eta_{\max}, \quad (3.13)$$

where $\Gamma(a+1)$ is the Gamma function. Substituting (3.13) back into (3.12) and using (2.2) finally gives us an expression for the mean phase-space density (although still in terms of the as yet unspecified parameter β):

$$\langle f \rangle(\mathbf{p}) = \frac{\beta^{1+a}}{A\Gamma(a+1)} \int_{\eta_{\min}}^{\eta_{\max}} d\eta \eta^{2+a} \rho(\eta) e^{-\beta\eta\varepsilon(\mathbf{p})}. \quad (3.14)$$

From (3.14), 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 3.2, 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 (3.7), we can make a convenient change

of variables in (3.14), $x = \beta\eta\varepsilon(\mathbf{p})$, to find

$$\langle f \rangle(\mathbf{p}) = \frac{n_0}{A\beta\Gamma(a+1)\eta_{\max}\varepsilon(\mathbf{p})^{2+a}} \int_{\beta\eta_{\min}\varepsilon(\mathbf{p})}^{\beta\eta_{\max}\varepsilon(\mathbf{p})} dx x^a G\left(\frac{x}{\beta\eta_{\max}\varepsilon(\mathbf{p})}\right) e^{-x}. \quad (3.15)$$

This form exposes the fact that there is a natural power-law behaviour at energies such that

$$\frac{1}{\beta\eta_{\max}} \ll \varepsilon(\mathbf{p}) \ll \frac{1}{\beta\eta_{\min}}. \quad (3.16)$$

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. At $\varepsilon(\mathbf{p}) \gg 1/\beta\eta_{\min}$, the lower limit of the integral in (3.15) imposes an exponential cutoff on $\langle f \rangle$. The distribution function $N(\varepsilon)$ of particle energies corresponding to (3.15) is obtained by multiplying the mean phase-space density $\langle f \rangle(\mathbf{p})$ by the density of states:

$$N(\varepsilon) = g(\varepsilon)\langle f \rangle(\mathbf{p}) = \frac{n_0}{\beta\Gamma(a+1)\eta_{\max}\varepsilon^2} \int_{\beta\eta_{\min}\varepsilon}^{\beta\eta_{\max}\varepsilon} dx x^a G\left(\frac{x}{\beta\eta_{\max}\varepsilon}\right) e^{-x}. \quad (3.17)$$

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 (3.9). Following all the same steps as before from (3.14) onwards, but using (3.9) in place of (3.11), one arrives at

$$N(\varepsilon) = \frac{n_0}{\beta\Gamma(a+1)\eta_{\max}^{2-\delta}\varepsilon^{3-\delta}} \int_{\beta\eta_{\min}\varepsilon}^{\beta\eta_{\max}\varepsilon} dx x^{a+1-\delta} G\left(\frac{x}{\beta\eta_{\max}\varepsilon}\right) e^{-x}. \quad (3.18)$$

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 (3.9), which is related to the Gardner distribution of that Lynden-Bell equilibrium via (3.5), as explained at the end of Section 3.2. For Gardner distributions that already have power-law tails, the resulting Lynden-Bell equilibria have shallower (ultra-'hard') power-law tails that strongly diverge in energy, giving the cutoff η_{\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.

Presently, however, we shall return to the Lynden-Bell equilibria (3.17) arising from exponential Gardner distributions, for which we complete the calculation.

3.4. Calculation of β and inevitability of partial degeneracy

Both (3.16) and (3.17) 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 (3.11). To find β , we must compute the energy of our Lynden-Bell equilibrium (3.12) according to (2.5). Equivalently, from (3.15),

$$E = \int d\mathbf{p} \varepsilon(\mathbf{p}) \langle f \rangle(\mathbf{p}) = \frac{a+1}{\beta} \int_{\eta_{\min}}^{\eta_{\max}} d\eta \rho(\eta). \quad (3.19)$$

Therefore, in the non-degenerate limit,

$$\beta = \frac{a+1}{E} \int_{\eta_{\min}}^{\eta_{\max}} d\eta \rho(\eta). \quad (3.20)$$

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 (3.11) of non-degeneracy, we evaluate $D(\varepsilon)$ at $\varepsilon \rightarrow 0$, 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 (3.11), (3.12) and (3.20):

$$E \gg (a+1) \left[\int_{\eta_{\min}}^{\eta_{\max}} d\eta \frac{\eta^{a+1} \rho(\eta)}{A\Gamma(a+1)} \right]^{1/(a+1)} \int_{\eta_{\min}}^{\eta_{\max}} d\eta \rho(\eta). \quad (3.21)$$

Since $\rho(\eta)$ is a function only of the Gardner distribution, the right-hand side of (3.21) must scale with the energy density E_G of the Gardner distribution that has the same waterbag content $\rho(\eta)$, but will always have $E_G \leq E$. For example, for the Gardner distribution (3.1), the right-hand side of (3.21) 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 of the ground state. Note, however, that (3.21) contains an integral of the waterbag content $\rho(\eta)$ with no weighting by η , which, by (2.4) and (2.6), is the total phase volume occupied by non-empty waterbags. Since $\rho(\eta) \propto \eta^{-1}$ at small η , this will diverge, albeit logarithmically slowly, as the minimum waterbag density η_{\min} is taken to zero. Indeed, for our example (3.1), (3.21) becomes

$$E \gg \frac{4}{3\sqrt{\pi}} n_0 \varepsilon_0 \left(\ln \frac{\eta_{\max}}{\eta_{\min}} \right)^{3/2} \left[1 + O\left(\frac{\eta_{\min}}{\eta_{\max}} \right) \right]. \quad (3.22)$$

This means that the condition (3.21) 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 η_{\min} . This is a manifestation of the fact that, as η_{\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 η_{\min} is kept finite.

This is not the only place where the finiteness of the cutoff η_{\min} has raged against the dying of the light. The same effect is manifest in the power law of ε^{-2} appearing in (3.17), 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 (3.19), where the integral of $\rho(\eta)$ has the same logarithmic divergence with $\eta_{\min} \rightarrow 0$ as it did in the right-hand side of (3.21). This then makes its way into the expression (3.20) for β , so, formally in the limit $\eta_{\min} \rightarrow 0$, $\beta \rightarrow \infty$ always!

Since we must keep η_{\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.

3.5. 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. 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; Ewart *et al.* 2022). Ewart *et al.* (2022) argued, 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

$$\Delta\Gamma \sim \frac{1}{\eta_{\text{eff}}} (n_0 \lambda_D^3)^\alpha, \quad \frac{2}{3} < \alpha < 1, \quad (3.23)$$

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

$$\eta_{\min} = \Delta\Gamma^{-1} \sim \eta_{\max} (n_0 \lambda_D^3)^{-\alpha}. \quad (3.24)$$

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 (3.24) might seem damningly small. It is, in fact, ideal. The existence of a broad power-law tail in (3.17) required a scale separation between η_{\max} and η_{\min} . The estimate (3.24) 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 a ε^{-2} power law, we are saved by the fact that only the logarithm of the ratio η_{\max}/η_{\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 only a weak dependence on the value of this cutoff, via the so-called Coulomb logarithm (e.g., Helander & Sigmar 2005).

Nevertheless, the presence of the logarithmic divergence with η_{\min} in the expression for β , signposted at the end of section 3.4, will make it difficult to satisfy the non-degeneracy approximation. There is good reason to suppose, however, that its breakdown may only be partial. We note that evaluating (3.11) 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.

4. Numerical results: partially degenerate equilibria

In this section, we shall recover the analytically predicted power-law tail in (3.17) by solving the constraint equations (2.5) and (2.6) for the Lynden-Bell equilibria (3.4). The numerical scheme for this is documented in detail in Appendix B, 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 $\varepsilon = p^2/2m$, although we anticipate from Section 3 that these results extend, qualitatively, to general regimes.

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




Figure 3: Numerically computed Lynden-Bell equilibria for a range of $\eta_{\max,\sigma}/\eta_{\min}$ and with $\rho(\eta)$ given by (4.1) with $\sigma = 2$. The energy density is equal $10E_G$ in all cases. (a) The numerically computed fugacity $F(\eta)$ (solid lines) compared with the analytical solution (3.13) obtained in the non-degenerate limit (dashed lines). (b) The resulting distribution function $N(\epsilon)$, with the universal power law $\propto \epsilon^{-2}$ shown for reference, cf. (3.17). Overplotted in solid colour (with the value range shown on the right) is the level of degeneracy $D(\epsilon)/[1+D(\epsilon)]$ (the probability that a given energy is occupied by a non-empty waterbag) as a function of energy; $D(\epsilon)$ is defined in (3.11).

defined by

$$\rho(\eta) = \frac{p_{\text{th},\sigma}^3}{\sigma\eta} \left(\ln \frac{\eta_{\max,\sigma}}{\eta} \right)^{(3-\sigma)/\sigma}, \quad \eta_{\min} < \eta < \eta_{\max,\sigma}. \quad (4.1)$$

This defines the family of Gardner distributions

$$f_{G,\sigma}(\mathbf{p}) = \begin{cases} \eta_{\max,\sigma} e^{-(p/p_{\text{th},\sigma})^\sigma} & \text{for } p < p_{\text{th},\sigma} \left(\ln \frac{\eta_{\max,\sigma}}{\eta_{\min}} \right)^{1/\sigma}, \\ 0 & \text{for } p > p_{\text{th},\sigma} \left(\ln \frac{\eta_{\max,\sigma}}{\eta_{\min}} \right)^{1/\sigma}, \end{cases} \quad (4.2)$$

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

$$n_0 = \frac{4\pi}{\sigma} \Gamma\left(\frac{3}{\sigma}\right) p_{\text{th},\sigma}^3 \eta_{\max,\sigma}, \quad E_G = \frac{\Gamma(5/\sigma)}{\Gamma(3/\sigma)} n_0 \frac{p_{\text{th},\sigma}^2}{2m}. \quad (4.3)$$

Since these Gardner distributions represent minimum-energy states, we will be able to scan in the energy density of the system for all $E \geq E_G$, 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 4.2). 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.

4.1. 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 3. Figure 3 shows the result of such a scan for $\sigma = 2$ and $E = 10E_G$. By comparing the exact (numerically calculated) fugacity to the theoretical prediction (3.13) obtained in the absence of phase-space degeneracy, we see that the agreement is nearly perfect when η_{\max} is close to η_{\min} , e.g., when $\eta_{\max,2}/\eta_{\min} = 10$. This is as expected, since the non-degenerate limit is valid when (3.21) holds, which it does, as can be confirmed from the solid red colour in panel (b), showing the probability $D(\varepsilon)/[1 + D(\varepsilon)]$ (with $D(\varepsilon)$ defined in (3.11)) 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 (3.16) necessary to resolve a power-law tail in energies. To see a power-law tail, one must increase η_{\max}/η_{\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 η_{\max}/η_{\min} that we argued in Section 3.5 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 (3.13). 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.

4.2. Energisation of particles and power-law tails

Let us now scan in the energies densities E of the distribution and again look for power-law tails and assess the effect of degeneracy. Figure 4 shows the results of such a scan, again with $\rho(\eta)$ specified by (4.1) 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

$$\langle f \rangle_{\eta_1 < \eta < \eta_2}(\mathbf{p}) = \int_{\eta_1}^{\eta_2} d\eta \eta P(\mathbf{p}, \eta). \quad (4.4)$$

As one would anticipate, when E is only slightly larger than E_G , the effects of phase-space degeneracy are most prominent: the densest portions of the phase 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 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., Asslani *et al.* 2012 and Ewart *et al.* 2022), this hides key universal features of systems with a continuum of level sets. To showcase this universality, in Figure 5, we plot the Lynden-Bell equilibria for two different waterbag contents, $\sigma = 2$ and $\sigma = 8$ in (4.1), and a range of energy densities E . The ε^{-2} power-law tails of these equilibria are immediately apparent, as predicted in (3.17).

Figure 5 shows how these power-law tails become more prominent as one adds more energy to the Gardner distribution. At $E = E_G$, one has a highly non-universal Gardner equilibrium. As a small amount of energy $E - E_G \ll E_G$ 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’,

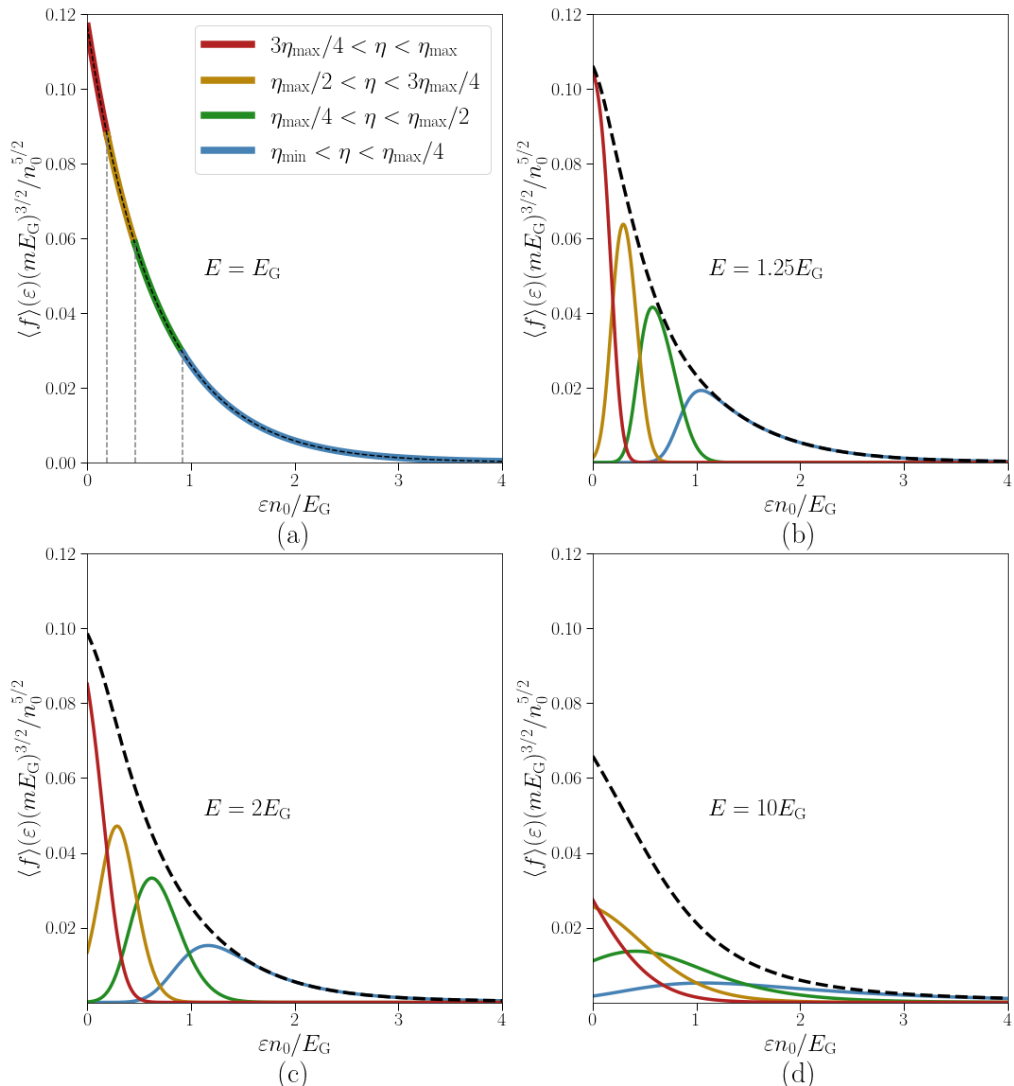


Figure 4: Numerically computed Lynden-Bell equilibria for a range of energy densities E (in multiples of the energy density E_G of the underlying Gardner distribution) with $\rho(\eta)$ given by (4.1) with $\sigma = 2$ and $\eta_{\max}/\eta_{\min} = 10^6$. In each plot, the dashed line is the mean phase-space 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.

which has energy density $\sim E_G$, 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 $\sim E - E_G$. 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

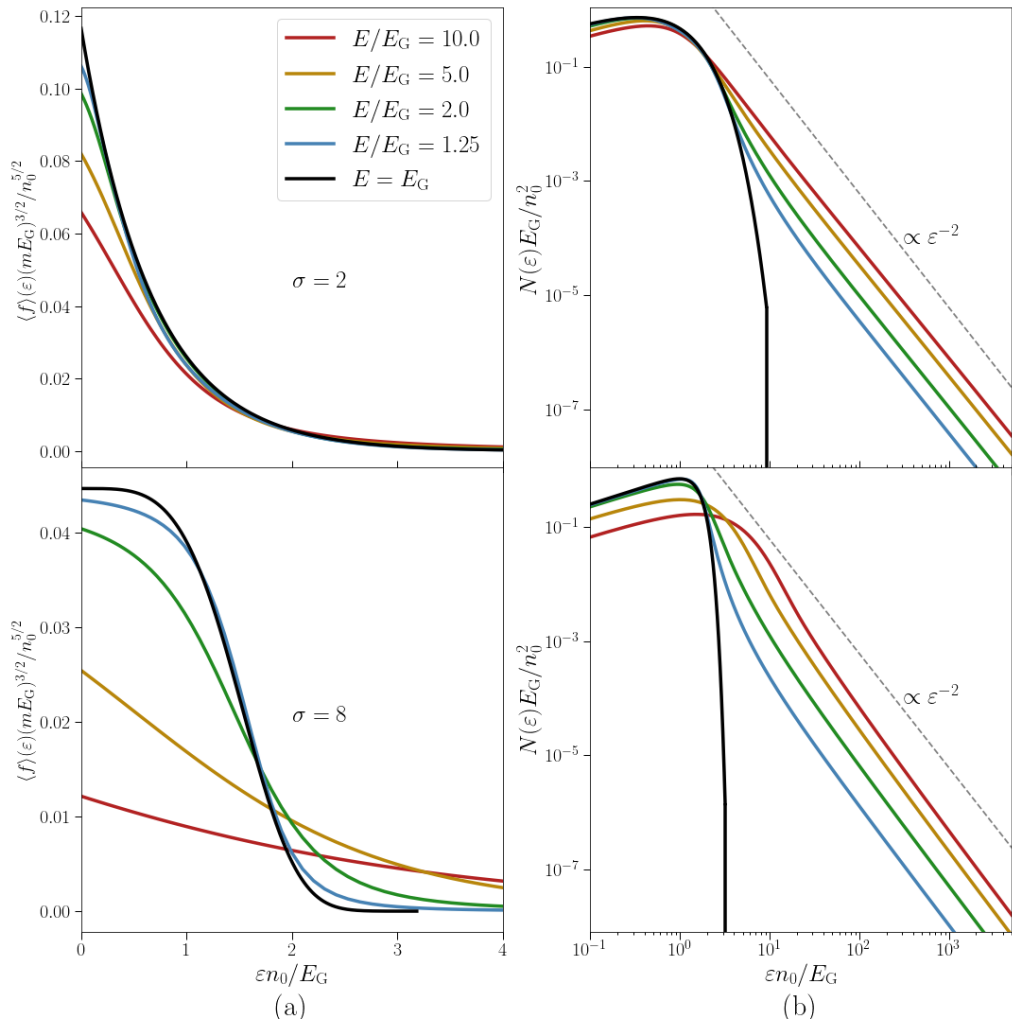


Figure 5: Numerically computed Lynden-Bell equilibria with waterbag content given by the $\sigma = 2$ (top) and $\sigma = 8$ (bottom) cases of (4.1), $\eta_{\max, \sigma} / \eta_{\min} = 10^6$. The distributions are plotted in linear (a) and logarithmic (b) scales, for a range of ratios of E/E_G . The small deviations from the ϵ^{-2} tail can be attributed to the logarithmic corrections arising from the x integral in (3.17).

into the core, but both thermally broaden. At $E \gg E_G$, the asymptotic solution (3.17) with a power law in the range (3.16) suggests that the system strives for a state in which the halo has a factor $\sim \ln(\eta_{\max} / \eta_{\min})$ more energy than the core, so the vast majority of the total energy is contained in the universal power-law tail.

5. Conclusion

5.1. Summary

The Lynden-Bell (1967) equilibria are the natural maximum-entropy states for systems, such as a plasma described by the collisionless Vlasov equation (1.1), which conserve not only density, momentum, and energy, but also an infinite family of further invariants (2.1).

These additional invariants are due to the conservation of phase volume, encoded by the ‘waterbag content’ function $\rho(\eta)$ given by (2.7) (equivalent to the Casimir invariants (3.6)), 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 (2.2) in the form of the Lynden-Bell equilibrium (2.9) coupled with the constraints (2.5) and (2.6). In this paper, 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 3.1, 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 3.2, 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 (3.7)). In Section 3.3, 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 (3.14) 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 (3.7), this gave rise to a universal power-law tail (3.17) 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 4, 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 (4.1) gave rise to the same power-law tail, the numerical solution also showed how this power-law tail formed. We found (Figure 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).

5.2. *Limitations and applications*

While the Lynden-Bell statistical mechanics presents an attractive scenario for the universal generation of power-law tails, which could perhaps offer insight into the distribution of particles in such astrophysical systems as cosmic rays and the solar wind, there are important limitations to this story, some of which we will now discuss.

Chief amongst these limitations is the absolute neglect of collisions. The validity of the Lynden-Bell statistics rests on the assumption that such an equilibrium can be reached long before the conservation of phase volume is broken. Conventional (linear) estimates for this timescale (Su & Oberman 1968), associated with the true Coulomb-collision frequency, indicate that such a relaxation should be possible. However, recent progress (Beraldo e Silva *et al.* 2017; Zhdankin 2022*a*; Nastac *et al.* 2023) has indicated that in turbulent systems, the conservation of phase volume may be broken on fast, collision-independent timescales. While it is possible that this is damning evidence against the existence (or persistence) of Lynden-Bell equilibria, it is also possible that it points to interesting interplay between the Lynden-Bell relaxation and the collisional effects. This is especially interesting because true Coulomb collisions affect the most energetic particles least strongly, perhaps implying that the high-energy tails engendered by the collisionless Lynden-Bell statistics could coexist with some form of partially collisionally relaxed (or at least collisionally affected) equilibria at lower energies.

Another interesting concern is the effect of multiple species. Whereas we have confined ourselves to a single-species problem, this is far from sufficient for real systems (except perhaps when one deals with a dilute admixture, like in the case of cosmic rays). Ewart *et al.* (2022) suggested a number of interesting relaxation effects that arise from collisionless relaxation in a multi-species plasma. However, central to assessing the importance, or indeed reality, of those effects was the need to know how to calculate the correlation volumes $\Delta\Gamma$ (see Section 3.5) for different species—because $\Delta\Gamma$ determined the effective interaction rates between and within the species. This is an area of much uncertainty and ongoing research. If a quantitative handle on this question is achieved, it would provide powerful insight into, e.g., the relative heating rates of different species in collisionless plasmas (since the Lynden-Bell equilibria equilibrate β , not the ‘true’ temperature: see Ewart *et al.* 2022).

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 (Becker Tjus & Merten 2020; Hopkins *et al.* 2022; Kempfski & 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). In collisionless plasmas, such a framework is lacking as both the turbulence and the nature of the underlying equilibrium are mysterious and indeed it is not even guaranteed that they can be understood without detailed reference to each system’s particular initial circumstances. If Lynden-Bell’s statistical mechanics proves to be a viable collisionless substitute for Maxwell’s, a path could be charted towards a theory of the dynamics and thermodynamics of collisionless plasmas possessing a modicum of universality.

Acknowledgements

We would like to thank Georgia Acton, Michael Barnes, Archie Bott, Andrew Brown, Steve Cowley, Jean-Baptiste Fouvry, Chris Hamilton, Per Helander, David Hosking, Matt Kunz, Sid Parameswaran, Juan Ruiz Ruiz and Luis Silva for illuminating discussions. RJE’s work was supported by a UK EPSRC studentship. MLN was supported by a Clarendon Scholarship. The work of AAS was supported in part by grants from STFC (ST/W000903/1) and EPSRC (EP/R034737/1), as well as by the Simons Foundations via a Simons Investigator award.

Declaration of Interests: The authors report no conflicts of interest.

Appendix A. 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 (3.4) 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 (3.4) when $E = E_G$, 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 $\langle f \rangle = f_G$.

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

$$f_{\text{FD}}(\varepsilon) = \frac{e^{-\beta(\varepsilon-\mu)}}{1 + e^{-\beta(\varepsilon-\mu)}}. \quad (\text{A1})$$

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 numerator and denominator approximately cancel to give unity. Borrowing this intuition, we anticipate that our solution should have $\beta \rightarrow \infty$. It is clear what kind of solution one must search for: the degeneracy parameter $D(\varepsilon)$ defined by (3.11) must be large wherever $f_G(\varepsilon) \neq 0$ and zero wherever $f_G(\varepsilon) = 0$. Furthermore, the dominant contribution to $D(\varepsilon)$ in the integral (3.11) must come from $\eta = f_G(\varepsilon)$. 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

$$F(\eta) = \exp \left[\beta \int_{\eta_{\min}}^{\eta} d\eta' f_G^{-1}(\eta') + C \right] \quad (\text{A } 2)$$

and prove that, via (3.4) and (2.6), it recovers the Gardner distribution with the correct waterbag content $\rho(\eta)$ when $\beta \rightarrow \infty$, for a suitable choice of the constant C .

In the denominator of (3.4), we must evaluate the integral

$$D(\varepsilon) = \int_{\eta_{\min}}^{\eta_{\max}} d\eta e^{-\beta\eta\varepsilon + \ln F(\eta)}. \quad (\text{A } 3)$$

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 $\eta_{\text{stat}}(\varepsilon)$, is given by the solution to

$$\beta \frac{d}{d\eta} \left[\eta\varepsilon - \int_{\eta_{\min}}^{\eta} f_G^{-1}(\eta') d\eta' \right] = 0 \implies \eta_{\text{stat}}(\varepsilon) = f_G(\varepsilon), \quad (\text{A } 4)$$

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

$$\begin{aligned} D(\varepsilon) &\simeq \int_{-\infty}^{\infty} d\eta \exp \left\{ -\beta f_G(\varepsilon)\varepsilon + \frac{\beta}{2} \left. \frac{df_G^{-1}}{d\eta} \right|_{\eta=f_G(\varepsilon)} [\eta - f_G(\varepsilon)]^2 \right\} F(f_G(\varepsilon)) \\ &= e^{-\beta f_G(\varepsilon)\varepsilon} F(f_G(\varepsilon)) \sqrt{\frac{2\pi}{\beta}} \left[- \left. \frac{df_G^{-1}}{d\eta} \right|_{\eta=f_G(\varepsilon)} \right]^{-1/2}. \end{aligned} \quad (\text{A } 5)$$

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 $O(1/\beta)$. 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 $\varepsilon = f_G^{-1}(\eta_{\min})$, where the dominant contribution comes precisely from the limit of integration. This is, however, fine for $\beta \rightarrow \infty$ because the intervals in ε where this approximation is bad shrink as $O(\beta^{-1/2})$.

If we further demand that the constant C in (A 2) is chosen so that

$$e^{-\beta\eta_{\min}f_G^{-1}(\eta_{\min}) + \beta C} \sqrt{\frac{2\pi}{\beta}} \left[- \left. \frac{df_G^{-1}}{d\eta} \right|_{\eta=\eta_{\min}} \right]^{-1/2} = 1, \quad (\text{A } 6)$$

then, neglecting terms $O(1/\beta)$, $D(\varepsilon)$ satisfies

$$1 + D(\varepsilon) \simeq \begin{cases} D(\varepsilon) & \text{for } \varepsilon < f_G^{-1}(\eta_{\min}), \\ 1 & \text{for } \varepsilon > f_G^{-1}(\eta_{\min}). \end{cases} \quad (\text{A } 7)$$

Calculating the mean phase-space density $\langle f \rangle(\mathbf{p})$ from (2.2) and (3.4), we find

$$\langle f \rangle(\varepsilon) \simeq \begin{cases} \frac{1}{D(\varepsilon(\mathbf{p}))} \int_{\eta_{\min}}^{\eta_{\max}} d\eta \eta e^{-\beta\eta\varepsilon + \ln F(\eta)} & \text{for } \varepsilon < f_G^{-1}(\eta_{\min}), \\ 0 & \text{for } \varepsilon > f_G^{-1}(\eta_{\min}). \end{cases} \quad (\text{A } 8)$$

Let us prove that $\langle f \rangle(\mathbf{p}) = f_G(\varepsilon(\mathbf{p}))$. The integral in (A 8) can again be computed by the method of steepest descent. The exponent is the same as in (A 3) and so will have

the same maximum, at $\eta_{\text{stat}} = f_G(\varepsilon)$ (up to a small $O(1/\beta)$ correction due to the factor of η in (A 8)). Expanding the integral around this maximum and neglecting all terms that are $O(1/\beta)$ causes the factor of η in the integral to be replaced by $\eta_{\text{stat}}(\varepsilon)$. After this, the remainder of the integral has the same form as (A 5), which cancels with the denominator of (A 8), leaving only the factor of $\eta_{\text{stat}} = f_G(\varepsilon)$, q.e.d.

Thus, the fugacity (A 2) correctly recovers the Gardner distribution $f_G(\varepsilon)$ in the limit of $\beta \rightarrow \infty$. 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 (2.6) for the fugacity (A 2), to show that it is the same waterbag content as that of f_G :

$$\rho(\eta) \simeq \int_0^\infty d\varepsilon g(\varepsilon) \exp \left\{ -\beta\varepsilon [\eta - f_G(\varepsilon)] + \beta \int_{f_G(\varepsilon)}^\eta f_G^{-1}(\eta) d\eta \right\} \cdot \sqrt{\frac{\beta}{2\pi}} \left[-\frac{df_G^{-1}}{d\eta} \Big|_{\eta=f_G(\varepsilon)} \right]^{1/2}. \quad (\text{A } 9)$$

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

$$\begin{aligned} \rho(\eta) &\simeq \int_{-\infty}^\infty d\varepsilon g(f_G^{-1}(\eta)) \exp \left\{ \frac{\beta}{2} \frac{df_G}{d\varepsilon} \Big|_{\varepsilon=f_G^{-1}(\eta)} [\varepsilon - f_G^{-1}(\eta)]^2 \right\} \sqrt{\frac{\beta}{2\pi}} \left[-\frac{df_G^{-1}}{d\eta} \right]^{1/2} \\ &= -g(f_G^{-1}(\eta)) \frac{df_G^{-1}}{d\eta}. \end{aligned} \quad (\text{A } 10)$$

This is exactly the expression (3.5) that we desire. This completes the proof that any Gardner distribution can be written as the $\beta \rightarrow \infty$ limit of the Lynden-Bell statistics, implying that one could, in principle, have discovered Gardner restacking just by analysing the Lynden-Bell equilibria at $\beta \rightarrow \infty$.

By retaining terms that are small in $O(1/\beta)$, 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 $\varepsilon = f_G^{-1}(\eta_{\text{min}})$, 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.

Appendix B. Numerical method for solving for 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 waterbag distribution given by (3.4), the constraints (2.5) and (2.6) are satisfied to sufficient accuracy. For the numerical solutions given in Section 4, 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 (3.4) 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

$$F^{(n+1)}(\eta) = \rho(\eta) \left[4\pi \left(\frac{2}{m} \right)^{3/2} \int_0^\infty d\varepsilon \varepsilon^{1/2} \frac{e^{-\beta^{(n)}\eta\varepsilon}}{1 + \int_{\eta_{\min}}^{\eta_{\max}} d\eta' e^{-\beta^{(n)}\eta'\varepsilon} F^{(n)}(\eta')} \right]^{-1}. \quad (\text{B1})$$

Ignoring for a moment how one iterates $\beta^{(n)}$, we note that if the iteration (B1) converges, then by definition (2.6) 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 (2.5). 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

$$\eta_j = \eta_{\min} + \left(\frac{j}{N_\eta} \right)^q (\eta_{\max} - \eta_{\min}), \quad (\text{B2})$$

where the number $q > 1$ is chosen depending on the waterbag content: for (4.1), $q = 3$ 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 4, the initial guess was set to the analytical solutions (3.13) and (3.20) 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 (B1). The η' integral in (B1) 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 (3.11) becomes smaller than 10^{-5} . This allows the energy integral in (B1) 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 (rather than energy, although this distinction is unimportant) grid linearly spaced with a spacing of 10^{-3} in units such that the $E_G = 1$. 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

$$\epsilon_F = \left[\frac{1}{\eta_{\max} - \eta_{\min}} \sum_{\eta_j} (\eta_{j+1} - \eta_j) \left(\frac{F^{(n+1)}}{F^{(n)}} - 1 \right)^2 \right]^{1/2} < 10^{-3}, \quad (\text{B3})$$

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 $E_G = 1$, 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.

REFERENCES

- ABEL, I. G., PLUNK, G. G., WANG, E., BARNES, M., COWLEY, S. C., DORLAND, W. & SCHEKOCHIHIN, A. A. 2013 Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport and energy flows. *Rep. Prog. Phys.* **76**, 116201.
- AMATO, E. & CASANOVA, S. 2021 On particle acceleration and transport in plasmas in the galaxy: theory and observations. *J. Plasma Phys.* **87**, 845870101.
- ARAD, I. & JOHANSSON, P. H. 2005 A numerical comparison of theories of violent relaxation. *Mon. Not. R. Astron. Soc.* **362**, 252.
- ASSLLANI, M., FANELLI, D., TURCHI, A., CARLETTI, T. & LEONCINI, X. 2012 Statistical theory of quasistationary states beyond the single water-bag case study. *Phys. Rev. E* **85**, 021148.
- BALESCU, R. 1960 Irreversible processes in ionized gases. *Phys. Fluids* **3**, 52.
- BECK, C. & COHEN, E. 2003 Superstatistics. *Physica A* **322**, 267.
- BECKER TJUS, J. & MERTEN, L. 2020 Closing in on the origin of galactic cosmic rays using multimessenger information. *Phys. Rep.* **872**, 1.
- BELL, A. R. 1978 The acceleration of cosmic rays in shock fronts - I. *Mon. Not. R. Astron. Soc.* **182**, 147.
- BENDER, C. M. & ORSZAG, S. A. 1978 *Advanced Mathematical Methods for Scientists and Engineers*. McGraw-Hill.
- BERALDO E SILVA, L., DE SIQUEIRA PEDRA, W., SODRÉ, L., PERICO, E. L. D. & LIMA, M. 2017 The arrow of time in the collapse of collisionless self-gravitating systems: non-validity of the Vlasov–Poisson equation during violent relaxation. *Astrophys. J.* **846**, 125.
- BIRN, J., ARTEMYEV, A. V., BAKER, D. N., ECHIM, M., HOSHINO, M. & ZELENYI, L. M. 2012 Particle acceleration in the magnetotail and aurora. *Space Sci. Rev.* **173**, 49.
- BOLTZMANN, L. 1896 *Vorlesungen über Gastheorie*. Leipzig: J. A. Barth.
- CAPRIOLI, D. & SPITKOVSKY, A. 2014 Simulations of ion acceleration at non-relativistic shocks. i. acceleration efficiency. *Astrophys. J.* **783**, 91.
- CHAVANIS, P.-H. 2004 Generalized thermodynamics and kinetic equations: Boltzmann, Landau, Kramers and Smoluchowski. *Physica A* **332**, 89.
- CHAVANIS, P.-H. 2006 Quasi-stationary states and incomplete violent relaxation in systems with long-range interactions. *Physica A* **365**, 102.
- CHAVANIS, P.-H. 2022 Kinetic theory of collisionless relaxation for systems with long-range interactions. *Physica A* **606**, 128089.
- CHAVANIS, P. H., SOMMERIA, J. & ROBERT, R. 1996 Statistical mechanics of two-dimensional vortices and collisionless stellar systems. *Astrophys. J.* **471**, 385.
- CHEN, C., BALE, S., BONNELL, J., BOROVNIKOV, D., BOWEN, T., BURGESS, D., CASE, A., CHANDRAN, B., DE WIT, T. D., GOETZ, K. & OTHERS 2020 The evolution and role of solar wind turbulence in the inner heliosphere. *Astrophys. J. Suppl. Ser.* **246**, 53.
- COMISSO, L. & SIRONI, L. 2018 Particle acceleration in relativistic plasma turbulence. *Phys. Rev. Lett.* **121**, 255101.
- COMISSO, L. & SIRONI, L. 2022 Ion and electron acceleration in fully kinetic plasma turbulence. *Astrophys. J. Lett.* **936**, L27.
- CRUMLEY, P., CAPRIOLI, D., MARKOFF, S. & SPITKOVSKY, A. 2019 Kinetic simulations of mildly relativistic shocks – I. Particle acceleration in high Mach number shocks. *Mon. Not. R. Astron. Soc.* **485**, 5105.
- CRUZ, F., ALBERTAZZI, B., BAMFORD, R., BELL, A. R., CROSS, J. E., FRASCHETTI, F., GRAHAM, P., HARA, Y., KOZLOWSKI, P. M., KURAMITSU, Y., LAMB, D. Q., LEBEDEV,

- S., MARQUES, J. R., MINIATI, F., MORITA, T., OLIVER, M., REVILLE, B., SAKAWA, Y., SARKAR, S., SPINDLOE, C., TRINES, R., TZEFERACOS, P., SILVA, L. O., BINGHAM, R., KOENIG, M. & GREGORI, G. 2018 Electron acceleration by wave turbulence in a magnetized plasma. *Nat. Phys.* **14**, 475.
- DODIN, I. & FISCH, N. 2005 Variational formulation of the gardner's restacking algorithm. *Phys. Lett. A* **341**, 187.
- DUDÍK, J., DZIFČÁKOVÁ, E., MEYER-VERNET, N., DEL ZANNA, G., YOUNG, P. R., GIUNTA, A., SYLWESTER, B., SYLWESTER, J., OKA, M., MASON, H. E., VOCKS, C., MATTEINI, L., KRUCKER, S., WILLIAMS, D. R. & MACKOVJAK, Š. 2017 Nonequilibrium processes in the solar corona, transition region, flares, and solar wind. *Solar Phys.* **292**, 100.
- ERGUN, R. E., AHMADI, N., KROMYDA, L., SCHWARTZ, S. J., CHASAPIS, A., HOILJOKI, S., WILDER, F. D., CASSAK, P. A., STAWARZ, J. E., GOODRICH, K. A., TURNER, D. L., PUCCI, F., POUQUET, A., MATTHAEUS, W. H., DRAKE, J. F., HESSE, M., SHAY, M. A., TORBERT, R. B. & BURCH, J. L. 2020 Particle acceleration in strong turbulence in the Earth's magnetotail. *Astrophys. J.* **898**, 153.
- EWART, R., BROWN, A., ADKINS, T. & SCHEKOCHIHIN, A. 2022 Collisionless relaxation of a Lynden-Bell plasma. *J. Plasma Phys.* **88**, 925880501.
- FERRIÈRE, K. 2019 Plasma turbulence in the interstellar medium. *Plasma Phys. Control. Fusion* **62**, 014014.
- FISK, L. A. & GLOECKLER, G. 2014 The case for a common spectrum of particles accelerated in the heliosphere: Observations and theory. *J. Geophys. Res. Space Phys.* **119**, 8733.
- GARDNER, C. S. 1963 Bound on the energy available from a plasma. *Phys. Fluids* **6**, 839.
- GLOECKLER, G., FISK, L. A., MASON, G. M. & HILL, M. E. 2008 Formation of power law tail with spectral index-5 inside and beyond the heliosphere. *AIP Conf. Proc.* **1039**, 367.
- HARTOUNI, E., MOORE, A., CRILLY, A., APPELBE, B., AMENDT, P., BAKER, K., CASEY, D., CLARK, D., DÖPPNER, T., ECKART, M., FIELD, J., GATU-JOHNSON, M., GRIM, G., HATARIK, R., JEET, J., KERR, S., KILKENNY, J., KRITCHER, A., MEANEY, K. & ZYLSTRA, A. 2022 Evidence for suprathermal ion distribution in burning plasmas. *Nat. Phys.* **19**, 1.
- HAVRDA, J. & CHARVÁT, F. 1967 Quantification method of classification processes. Concept of structural α -entropy. *Kybernetika* **3**, 30.
- HELANDER, P. 2017 Available energy and ground states of collisionless plasmas. *J. Plasma Phys.* **83**, 715830401.
- 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 Standard self-confinement and extrinsic turbulence models for cosmic ray transport are fundamentally incompatible with observations. *Mon. Not. R. Astron. Soc.* **517**, 5413.
- KADOMTSEV, B. B. & POGUTSE, O. P. 1970 Collisionless relaxation in systems with Coulomb interactions. *Phys. Rev. Lett.* **25**, 1155.
- KAWAZURA, Y., SCHEKOCHIHIN, A., BARNES, M., DORLAND, W. & BALBUS, S. 2022 Energy partition between Alfvénic and compressive fluctuations in magnetorotational turbulence with near-azimuthal mean magnetic field. *J. Plasma Phys.* **88**, 905880311.
- KEMPSKI, P. & QUATAERT, E. 2022 Reconciling cosmic ray transport theory with phenomenological models motivated by Milky-Way data. *Mon. Not. R. Astron. Soc.* **514**, 657.
- KOLMES, E. J. & FISCH, N. J. 2020 Recovering Gardner restacking with purely diffusive operations. *Phys. Rev. E* **102**, 063209.
- KOLMES, E. J., HELANDER, P. & FISCH, N. J. 2020 Available energy from diffusive and reversible phase space rearrangements. *Phys. Plasmas* **27**, 062110.
- KUNZ, M. W., STONE, J. M. & QUATAERT, E. 2016 Magnetorotational turbulence and dynamo in a collisionless plasma. *Phys. Rev. Lett.* **117**, 235101.
- LANDAU, L. D. 1936 Transport equation in the case of Coulomb interaction. *Zh. Eksp. Teor. Fiz.* **7**, 203.
- LENARD, A. 1960 On Bogoliubov's kinetic equation for a spatially homogeneous plasma. *Ann. Phys.* **10**, 390.

- LESUR, G. R. J. 2021 Magnetohydrodynamics of protoplanetary discs. *J. Plasma Phys.* **87**, 205870101.
- LEVIN, Y., PAKTER, R., RIZZATO, F. B., TELES, T. N. & BENETTI, F. P. 2014 Nonequilibrium statistical mechanics of systems with long-range interactions. *Phys. Rep.* **535**, 1.
- LEVIN, Y., PAKTER, R. & TELES, T. N. 2008 Collisionless relaxation in non-neutral plasmas. *Phys. Rev. Lett.* **100**, 040604.
- LIVADIOTIS, G., DESAI, M. I. & WILSON III, L. B. 2018 Generation of kappa distributions in solar wind at 1 au. *Astrophys. J.* **853**, 142.
- LIVADIOTIS, G. & MCCOMAS, D. 2009 Beyond kappa distributions: exploiting Tsallis statistical mechanics in space plasmas. *J. Geophys. Res.* **114**, A11105.
- LUCEK, E., CONSTANTINESCU, D., GOLDSTEIN, M., PICKETT, J., PINCON, J.-L., SAHRAOUI, F., TREUMANN, R. & WALKER, S. 2005 The magnetosheath. *Space Sci. Rev.* **118**, 95.
- LYNDEN-BELL, D. 1967 Statistical mechanics of violent relaxation in stellar systems. *Mon. Not. R. Astron. Soc.* **136**, 101.
- MAILLOUX, J. & OTHERS 2022 Overview of JET results for optimising ITER operation. *Nucl. Fusion* **62**, 042026.
- MAXWELL, J. C. 1860 V. Illustrations of the dynamical theory of gases.—Part I. On the motions and collisions of perfectly elastic spheres. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **19**, 19.
- MONCUQUET, M., MEYER-VERNET, N., ISSAUTIER, K., PULUPA, M., BONNELL, J. W., BALE, S. D., DE WIT, T. D., GOETZ, K., GRITON, L., HARVEY, P. R., MACDOWALL, R. J., MAKSIMOVIC, M. & MALASPINA, D. M. 2020 First in situ measurements of electron density and temperature from quasi-thermal noise spectroscopy with Parker Solar Probe/FIELDS. *Astrophys. J. Suppl. Ser.* **246**, 44.
- NASTAC, M. L., EWART, R. J., SENGUPTA, W., SCHEKOCIHIN, A. A., BARNES, M. & DORLAND, W. 2023 Irreversible stochastic heating via phase-space entropy cascade in nearly collisionless plasma turbulence. *In preparation*.
- OKA, M., BIRN, J., BATTAGLIA, M., CHASTON, C. C., HATCH, S. M., LIVADIOTIS, G., IMADA, S., MIYOSHI, Y., KUCHAR, M., EFFENBERGER, F., ERIKSSON, E., KHOTYAINTEV, Y. V. & RETINÒ, A. 2018 Electron power-law spectra in solar and space plasmas. *Space Sci. Rev.* **214**.
- OKA, M., KRUCKER, S., HUDSON, H. S. & SAINT-HILAIRE, P. 2015 Electron energy partition in the above-the-looptop solar hard X-ray sources. *Astrophys. J.* **799**, 129.
- PIERRARD, V. & LAZAR, M. 2010 Kappa distributions: theory and applications in space plasmas. *Sol. Phys.* **267**, 153.
- ROBERT, R. & SOMMERIA, J. 1991 Statistical equilibrium states for two-dimensional flows. *J. Fluid Mech.* **229**, 291.
- SCHEKOCIHIN, A. A., COWLEY, S. C., DORLAND, W., HAMMETT, G. W., HOWES, G. G., QUATAERT, E. & TATSUNO, T. 2009 Astrophysical gyrokinetics: kinetic and fluid turbulent cascades in magnetized weakly collisional plasmas. *Astrophys. J. Suppl. Ser.* **182**, 310.
- SEVERNE, G. & LUWEL, M. 1980 Dynamical theory of collisionless relaxation. *Astrophys. Space Sci.* **72**, 293.
- SIRONI, L. & SPITKOVSKY, A. 2010 Particle acceleration in relativistic magnetized collisionless electron-ion shocks. *Astrophys. J.* **726**, 75.
- SIRONI, L. & SPITKOVSKY, A. 2014 Relativistic reconnection: an efficient source of non-thermal particles. *Astrophys. J.* **783**, L21.
- SU, C. H. & OBERMAN, C. 1968 Collisional damping of a plasma echo. *Phys. Rev. Lett.* **20**, 427.
- TSALLIS, C. 1988 Possible generalization of Boltzmann-Gibbs statistics. *J. Stat. Phys.* **52**, 479.
- UZDENSKY, D. A. 2022 Relativistic non-thermal particle acceleration in two-dimensional collisionless magnetic reconnection. *J. Plasma Phys.* **88**, 905880114.
- VERSCHAREN, D., KLEIN, K. G. & MARUCA, B. A. 2019 The multi-scale nature of the solar wind. *Living Rev. Sol. Phys.* **16**, 5.
- WERNER, G. R. & UZDENSKY, D. A. 2021 Reconnection and particle acceleration in three-

- dimensional current sheet evolution in moderately magnetized astrophysical pair plasma. *J. Plasma Phys.* **87**, 905870613.
- ZHDANKIN, V. 2021 Particle energization in relativistic plasma turbulence: solenoidal versus compressive driving. *Astrophys. J.* **922**, 172.
- ZHDANKIN, V. 2022a Generalized entropy production in collisionless plasma flows and turbulence. *Phys. Rev. X* **12**, 031011.
- ZHDANKIN, V. 2022b Non-thermal particle acceleration from maximum entropy in collisionless plasmas. *J. Plasma Phys.* **88**, 175880303.
- ZHDANKIN, V., UZDENSKY, D. A., WERNER, G. R. & BEGELMAN, M. C. 2019 Electron and ion energization in relativistic plasma turbulence. *Phys. Rev. Lett.* **122**, 055101.
- ZHDANKIN, V., WERNER, G. R., UZDENSKY, D. A. & BEGELMAN, M. C. 2017 Kinetic turbulence in relativistic plasma: from thermal bath to nonthermal continuum. *Phys. Rev. Lett.* **118**, 055103.
- ZHURAVLEVA, I., CHURAZOV, E., SCHEKOCHIHIN, A. A., ALLEN, S. W., ARÉVALO, P., FABIAN, A. C., FORMAN, W. R., SANDERS, J. S., SIMIONESCU, A., SUNYAEV, R., VIKHLININ, A. & WERNER, N. 2014 Turbulent heating in galaxy clusters brightest in X-rays. *Nature* **515**, 85.