Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **471,** 3753–3767 (2017)
Advance Access publication 2017 July 19

doi:10.1093/mnras/stx1817

The distribution of density in supersonic turbulence

Jonathan Squire[1,2]★ and Philip F. Hopkins[1]

[1]*TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA*
[2]*Walter Burke Institute for Theoretical Physics, Pasadena, CA 91125, USA*

Accepted 2017 July 17. Received 2017 July 14; in original form 2017 February 24

ABSTRACT

We propose a model for the statistics of the mass density in supersonic turbulence, which plays a crucial role in star formation and the physics of the interstellar medium (ISM). The model is derived by considering the density to be arranged as a collection of strong shocks of width $\sim\mathcal{M}^{-2}$, where $\mathcal{M}$ is the turbulent Mach number. With two physically motivated parameters, the model predicts all density statistics for $\mathcal{M} > 1$ turbulence: the density probability distribution and its intermittency (deviation from lognormality), the density variance–Mach number relation, power spectra and structure functions. For the proposed model parameters, reasonable agreement is seen between model predictions and numerical simulations, albeit within the large uncertainties associated with current simulation results. More generally, the model could provide a useful framework for more detailed analysis of future simulations and observational data. Due to the simple physical motivations for the model in terms of shocks, it is straightforward to generalize to more complex physical processes, which will be helpful in future more detailed applications to the ISM. We see good qualitative agreement between such extensions and recent simulations of non-isothermal turbulence.

Key words: shock waves – turbulence – ISM: kinematics and dynamics – galaxies: star formation.

1 INTRODUCTION

A detailed knowledge of gas statistics in the interstellar medium (ISM) is pivotal for theories of star formation and the stellar initial mass function (see e.g. Krumholz & McKee 2005; Federrath & Klessen 2012; Hennebelle & Chabrier 2013; Hopkins 2013a; Padoan et al. 2014, and references therein). The difficulty in understanding these statistics arises because much of the ISM is in a state of *supersonic turbulence* – a highly chaotic tangle of interacting shocks with structure spanning an enormous range in scale. However, despite a wide range of work on the subject, enabled in large part by the explosive growth in computational power (e.g. Kritsuk et al. 2007; Lemaster & Stone 2008; Molina et al. 2012; Federrath 2013; Pan et al. 2016; Colbrook et al. 2017, and references therein), much less is known about the statistical properties of supersonic turbulence in comparison to its subsonic cousin (Federrath 2013; Pan et al. 2016). For example, despite some promising results (e.g. Fleck 1996; Boldyrev 2002; Boldyrev, Nordlund & Padoan 2002; Padoan et al. 2004; Schmidt, Federrath & Klessen 2008; Aluie 2011; Galtier & Banerjee 2011; Banerjee & Galtier 2013), we currently lack a well-accepted theory for the velocity power spectrum, similar to the standard Kolmogorov phenomenology for subsonic turbulence (Kolmogorov 1941).

An important aspect of supersonic turbulence theory, which is crucial for star formation applications but plays only a minor role in subsonic theories, is the density statistics. The density statistics can be observationally probed through the projected (2D) column density, and various methods having been developed to extract the full 3D statistics from such measurements (see e.g. Vázquez-Semadeni & García 2001; Brunt, Federrath & Price 2010a,b; Burkhart et al. 2010; Brunt & Federrath 2014; Kainulainen, Federrath & Henning 2014). However, although there are certain well-established theoretical results – most importantly the *density variance–Mach number relation*: that the density distribution is approximately lognormal with a variance that increases with Mach number (Passot & Vázquez-Semadeni 1998; Burkhart et al. 2010; Padoan & Nordlund 2011; Price, Federrath & Brunt 2011; Molina et al. 2012; Federrath & Banerjee 2015; Nolan, Federrath & Sutherland 2015; Pan et al. 2016) – we lack detailed understanding of many important issues. For example, the power spectrum of density and its variation with Mach number (Fleck 1996; Beresnyak, Lazarian & Cho 2005; Kim & Ryu 2005; Kowal, Lazarian & Beresnyak 2007; Kritsuk et al. 2007; Konstandin et al. 2016), or other parameters (e.g. turbulent driving or self-gravity; Federrath, Klessen & Schmidt 2009; Federrath & Klessen 2013; Burkhart, Collins & Lazarian 2015) is not well constrained. Further, an important limitation of the density variance–Mach number relation is that the density can be quite *intermittent*, viz., it is not distributed lognormally. This behaviour manifests itself in a significant negative skewness in the probability density function (PDF; Kritsuk et al. 2007; Federrath, Klessen & Schmidt 2008; Burkhart et al. 2009; Schmidt et al. 2009; Federrath et al. 2010; Price & Federrath 2010; Konstandin et al. 2012; Hopkins 2013b), enhancing the

★E-mail: jsquire@caltech.edu

probability of low-density regions while decreasing the probability of high-density regions. The purpose of this work is to propose and examine a simple phenomenological model for the turbulent density field that encompasses all such statistics: the power spectrum, the density PDF and intermittency, and how these vary with turbulent Mach number.

It is worth elaborating on the density intermittency mentioned in the previous paragraph: is this important, or simply a formal nuisance to occupy the idle theorist? For some applications, in particular those that depend on events that occur regularly (i.e. with high probability), the answer is probably that intermittency is not important: so long as the density PDF is approximately lognormal near its peak, high-probability events will be well-characterized purely by the variance. However, many physically interesting properties that one might wish to derive from turbulent PDFs involve rare events, in particular those involving high-density regions. As a simple example, if we were interested in regions with a factor $\sim$100 enhancement in density over the mean at a turbulent Mach number $\mathcal{M} \sim 15$ (e.g. to push the density above the Jeans density and enable gravitational collapse), the probability of finding such a region could easily be overestimated by 1 to 3 orders of magnitude by a lognormal model, compared to a more realistic intermittent model with the same variance. Gaining a better understanding of the density PDF can also be motivated by other more fundamental interests. For instance, averaged over scale l, the PDF of the density field *cannot* be lognormal (note that l could be the grid scale here): by predicting a non-zero probability of finding densities greater than the ratio between the box volume and l^3, a lognormal PDF violates mass conservation (see also Hopkins 2013b). Further, we shall find that the process of understanding the origin of the density intermittency using the properties of shocks leads to various other insights; for example, the origin of the qualitatively different PDF shapes seen in non-isothermal or magnetohydrodynamic (MHD) turbulence.

This work involves an extension of the supersonic density PDF model proposed by Hopkins 2013b (hereafter H13) to multipoint statistics. Specifically, this involves specifying the properties of the PDF *as a function of scale*. In other words: after averaging the density field over scale l (which is some fraction of the box scale L), what is its PDF? Such a model completely specifies the statistical properties of the turbulent density field and contains significantly more information than the 1-point PDF that is usually calculated from numerical studies. For example, it allows the prediction of the density variance–Mach number relation, a similar relation for the intermittency, density power spectra and structure functions. Further, these predictions are made with two physically motivated free parameters that are independent of Mach number $\mathcal{M}$ once this is large. Extending the discussions in H13 and based on She-Leveque intermittency models that have been successful for subsonic turbulence (She & Leveque 1994; She & Waymire 1995; Dubrulle 1994; Castaing 1996; He, Dubrulle & Graner 1998), the model relies on a simple physical picture in which the density field is made up of a series of shocks covering a wide range of scales down to where the turbulence becomes subsonic. Starting from the box scale, at which the density field is simply constant, the shocks add density variance in the form of discrete multiplicative events, where the size of individual events is related to the physical properties of a shock. This predicts a strong relationship between the size of individual events – which controls the intermittency – and the variance of the density PDF, governed by the range of scales (i.e. L/l) and the Mach number. Comparison to various well-known results and trends, as well as our own numerical simulations, illustrates reasonable agreement across a range in Mach numbers. In addition, given the simple physical reasoning used to derive model parameters from isothermal shock properties, we extend the model to more complex and realistic physics – for example, a non-isothermal gas equation of state, or MHD – explaining various aspects of the density PDFs and why these are different from isothermal turbulence.

The remainder of the paper is organized as follows. In Section 2, we outline a few general considerations that will be used to motivate various choices in our model. In Section 3, we explain the model, in particular how the mathematical structure of a (compound) log-Poisson random process can be related to the properties of individual shocks. We take particular care here to outline the choices necessary for various parameters, and how these may be motivated or phenomenologically derived from physical properties of the turbulence. We then outline a variety of predictions of the model – including the density variance–Mach number relation, intermittency predictions and power spectra – and compare these to numerical simulations. This comparison involves both previous results and a variety of new simulations using the Lagrangian meshless finite mass (MFM) method in the GIZMO code (Hopkins 2015a), which we use to directly compute the density PDF as a function of scale. We also directly measure the physical size of shock structures, which forms an important part of our argument, in Appendix B. Overall, model predictions seem to match with numerical results up to numerical uncertainties, although more detailed comparisons will be necessary to understand its successes and failures more completely. We finish with an extension to non-isothermal turbulence (illustrating reasonable qualitative agreement with the simulations of Federrath & Banerjee 2015) and a discussion of MHD, before concluding by reiterating the model's main predictions.

2 GENERAL CONSIDERATIONS

Before continuing, it seems worth enumerating several general points about supersonic turbulence. While some of these are well known, given that each plays some crucial role in the derivation of our model, it is helpful to clearly explicate these ideas early on in our discussion.

(i) Supersonic turbulence is not scale invariant, except in the infinite Mach number limit. This arises because of the importance of the sonic scale $l_{\rm sonic}$ (Vázquez-Semadeni, Ballesteros-Paredes & Klessen 2003), which is the scale at which $v_l = v_{|\boldsymbol{l}|} = \langle |v(\boldsymbol{x}+\boldsymbol{l}) - v(\boldsymbol{x})| \rangle \sim c_s$; i.e. the scale at which the turbulence becomes subsonic (see e.g. Federrath 2016 for more detailed derivation and discussion). For modest Mach numbers common in nature or numerical experiments, the scale separation between $l_{\rm sonic}$ and the driving scale is also modest, challenging the relevance of the concept of a supersonic 'inertial range'. This feature necessarily leads to some important differences in the theoretical treatment of supersonic turbulence in comparison to subsonic turbulence.

(ii) Because density is effectively defined with reference to a volume (it is the mass per unit volume), the density itself and its PDF, are naturally defined with reference to an averaging scale. This density PDF as a function of scale encodes a wide variety of useful statistical information about the density field (see Vázquez-Semadeni & García 2001). By only ever studying the PDF with respect to some arbitrary scale (usually the grid scale in numerical simulations), one risks missing important trends or effects, particularly considering that the sonic scale is often close to the grid scale for Mach numbers of $\sim 10 \rightarrow 20$ at currently available numerical resolutions (see point (i). As an example, in the numerical

simulations in this work, we find that the density PDF becomes significantly *less* intermittent below $l_{\rm sonic}$, presumably because the density field on subsonic scales involves nearly Gaussian fluctuations (Federrath et al. 2010; Konstandin et al. 2012). The density averaged over scale is also relevant for applications, being an important quantity for studying gravitational collapse. For example, a region of (linear) size l is unstable to collapse if l is larger than the Jeans length

$$\lambda_{\rm J} \sim \frac{c_s}{\sqrt{G\rho_l}}. \quad (1)$$

Here, ρ_l is the average density over l, so to understand the statistics of gravitational collapse in a turbulent cloud, we require an understanding of the density PDF as a function of scale.

(iii) A PDF that appears more lognormal does not necessarily imply that the underlying statistics are closer to Gaussian. Instead, the apparent Gaussianity of the PDF may be a result of the suppression of low densities compared to the isothermal case. This point is important for MHD and non-isothermal turbulence.

(iv) The assumption that the shock width is equal to the sonic scale – while useful as a phenomenological tool for deriving density variance–Mach number relations (Price et al. 2011; Padoan & Nordlund 2011; Molina et al. 2012; Federrath & Banerjee 2015) – is inconsistent with lognormal density statistics. In particular, for isothermal shocks with a density contrast $\sim\mathcal{M}^2$, such a model involves *all* of the mass being concentrated in a single shock, which leads to an unphysically intermittent density distribution. In our model, we take the shock width to be some small fixed fraction κ of $l_{\rm sonic}$, and it will transpire that the κ parameter controls the intermittency. The success of density variance estimates using $l_{\rm sonic}$ as the shock width may then be related to κ being approximately universal, even in more complex physical situations (e.g. MHD).

(v) The infinite Mach number limit is not equivalent to Burgers turbulence (where the pressure term is neglected in the Navier–Stokes equations). This is because there are always some regions, no matter how large $\mathcal{M}$, where the pressure forces become important; see Passot & Vázquez-Semadeni (1998).

3 MODEL DESCRIPTION

In this section, we describe the mathematical structure and physical motivation for the density model. We shall denote the scale over which the density is averaged as l, the box scale as L, the PDF of the density averaged over scale l as $\mathcal{P}_l(\rho)$, and assume that the volume average of ρ over the whole box is $\langle\rho\rangle = M/L^3 = 1$. Let us start by considering the density averaged over the scale of the box $l = L$, which by definition is $\rho = 1$ with the PDF $\mathcal{P}_L(\rho) = \delta(\rho - 1)$. The model then provides a description of the density PDF averaged over successively smaller subvolumes of the box, $\mathcal{P}_l(\rho)$, a full knowledge of which effectively provides a full description of the statistical state of the turbulent density field.

This process is described mathematically by a series of N steps, each of which decrease the volume over which ρ is averaged by a factor $\Gamma > 1$, such that $l^3 = L^3/\Gamma^N$. The basic idea is that as one makes a jump down in scale from l to l/Γ, there is some probability, scaling with $\Gamma - 1$, of an 'event' that changes the density by $\delta\rho$. This event will be related to a shock structure that was previously in the larger volume (at scale l) being lost from the smaller volume (at scale l/Γ), thus causing a decrease in density in the volume being considered. This shock is related to velocity and density structures that vary over scale l because smaller structures cannot contribute significantly due to their smaller size, while there are too few larger

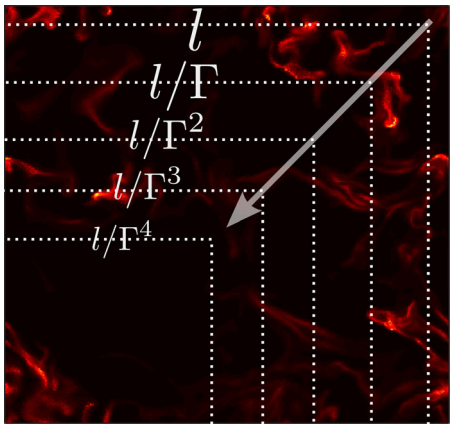

Figure 1. Cartoon of how the average density on scale l can change as l decreases, projected into the 2D plane. The brightness illustrates the density field, which is dominated by thin shock structures with a large proportion of the total mass. In the hypothetical step in scale from l to l/Γ the difference in volume (between the squares) covers a very high density shock region (upper right), implying the average density at scale l/Γ is *less* than that at l, since a significant proportion of the mass is lost in the jump. In contrast, in the jump from l/Γ to l/Γ^2, which entails the same proportional volume change, there is not a significant change in density because all high-density regions remain inside the new volume at scale l/Γ^2.

structures (since these vary over scales $\gg l$). The size of this 'jump' decrease in density $\delta\rho$ is related to the proportion of mass that resides in the shock that is lost in the step, and is itself a random variable. Assuming that most of the mass in the system is tied up in shocks, steps with no event will cause the average density to increase slightly, because the density in the volume that is lost in the step is less than the average density in the system. We graphically illustrate such a process in Fig. 1, showing how in a volume filled with high-density shocks, some steps will cause a large decrease in the mass enclosed by the new volume (i.e. an event), while other steps will not (no event).

3.1 Compound-log-Poisson cascades

The process we have just described is known mathematically as a *compound-log-Poisson* random process. Such processes are well characterized and have been particularly successful in phenomenological turbulence intermittency models for velocity structure functions (Dubrulle 1994; She & Leveque 1994; Castaing 1996; He et al. 1998; Boldyrev 2002; Schmidt et al. 2008; Mallet & Schekochihin 2017). The process occurs in log space because we consider multiplicative changes to the volume and density (which are additive in log space), while it is a Poisson process because the probability of an event is proportional to the size of the jump in volume (with infinitesimally small jumps being considered). If the size of each density jump, $\ln\rho \to \ln\rho - \delta\ln\rho$, is fixed (i.e. if $\delta\ln\rho$ is a number), then the resulting distribution will be log-Poisson,

$$\mathcal{P}(\ln\rho) = \frac{\lambda^{-\ln\rho+\Upsilon}\mathrm{e}^{-\lambda}}{(-\ln\rho+\Upsilon)!}, \quad (2)$$

where Υ accounts for a mean shift in the distribution (required to keep $\langle\rho\rangle = 1$; see below). However, the distribution (2) has the unphysical property of being discrete, since the jump sizes are discrete, and is thus somewhat inconvenient as a model for density. This issue may be circumvented by postulating that the size of each jump $\delta\ln\rho$ is itself a random variable, with probability

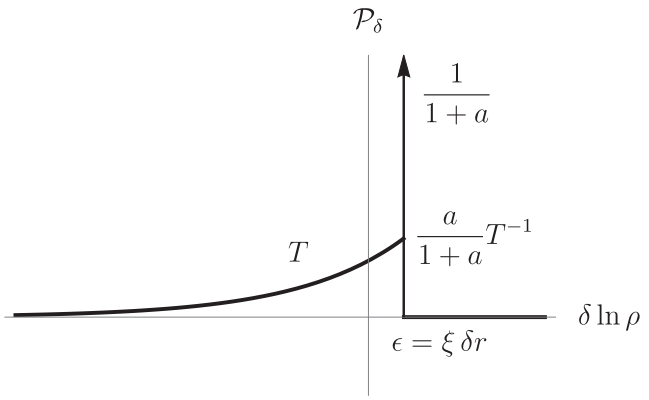

Figure 2. The PDF for a single step in the cascade, which is $a(1+a)^{-1}\mathcal{P}_\delta(x) + (1+a)^{-1}\delta(x)$, where a is a small parameter proportional to the size of the step in volume $\delta r = \ln(\Gamma - 1)$. The maximum ϵ is chosen as $\epsilon = \xi\delta r$, which correctly captures the true maximum of ρ when $\xi = 3$, and otherwise describes how the maximum density changes with scale. When coupled with the constraint $\langle\rho\rangle = 1$, ξ sets the proportionality between a and δr as $a = \xi\delta r(1+T)T^{-1}$.

distribution $\delta\ln\rho \sim \mathcal{P}_\delta(\delta\ln\rho)$. This leads to a 'compound-log-Poisson' distribution, where 'compound' refers to the idea that the distribution is formed as a random process of a random variable. Following H13, we take the jump sizes to be distributed exponentially (Castaing 1996),

$$\mathcal{P}_\delta(\delta\ln\rho) = \begin{cases} T^{-1}\exp[(\delta\ln\rho - \epsilon)/T] & \delta\ln\rho - \epsilon < 0 \\ 0 & \delta\ln\rho - \epsilon > 0 \end{cases}, \tag{3}$$

where T is the mean jump size and ϵ is a constant that is used to ensure $\langle\rho\rangle = 1$; see Fig. 2. This leads to a convenient form for the PDF that matches density PDFs measured from simulations remarkably well (H13).[1]

We now derive the PDF of $\ln\rho$, $\mathcal{P}_l(\ln\rho)$. that arises from this process. Assuming that the parameter T does not depend on scale, $\mathcal{P}_l(\ln\rho)$ is the convolution of n $\mathcal{P}_\delta$ distributions shifted by the total number of steps $\Upsilon = N\epsilon$,

$$\mathcal{P}_l(\ln\rho) = \mathcal{P}_\delta^{\otimes n} = \mathrm{Gamma}(-\ln\rho + \Upsilon; n, T), \tag{4}$$

where $\cdot^{\otimes n}$ denotes the convolution power and $\mathrm{Gamma}(x;\, n,\, T) = x^{n-1}\mathrm{e}^{-x/T}T^{-n}\Gamma(n)^{-1}$ is the Gamma distribution. As one takes the limit $N \to \infty$, $\Gamma \to 1$ with $N(\Gamma - 1) = \lambda$, the number of events n is itself a Poisson-distributed random variable with mean λ (this will depend on scale l and is specified below), so the full PDF for $\ln\rho$ is simply a sum of the PDFs for a given n, weighted by the probability that such an n occurs $n \sim \lambda^n\mathrm{e}^{-\lambda}/n!$. Putting this together, one obtains

$$\begin{aligned}\ln\rho \sim \mathcal{P}_l(\ln\rho) &= \sum_{n=0}^{\infty}\frac{\lambda^n\mathrm{e}^{-\lambda}}{n!}\frac{u^{n-1}}{T\Gamma(n)}\exp(-u) \\ &= T^{-1}\sqrt{\frac{\lambda}{u}}I_1(2\sqrt{u\lambda})\exp[-(\lambda+u)],\end{aligned} \tag{5}$$

where $u \equiv (-\ln\rho + \Upsilon)/T$, $I_1(x)$ is the first-order modified Bessel function of the first kind, and $\mathcal{P}_l(\ln\rho)$ is non-zero only for $u > 0$. We may then fix $\Upsilon = N\epsilon = \lambda T(1+T)^{-1}$ using the constraint $\langle\mathrm{e}^{\ln\rho}\rangle = \langle\rho\rangle = 1$, which leads to the density PDF proposed in H13. The volume-weighted variance is $S_{\ln\rho,\,V} = 2\lambda T^2$, while T – the mean size of the jumps – is an intermittency parameter that skews the distribution towards higher probability at $\ln\rho < 0$. In line with our intuition, small numbers of large density jumps (e.g. large shocks) lead to highly intermittent distributions, while a large number of small jumps leads to distributions that are very close to lognormal. The mass-weighted variance is $S_{\ln\rho,\,M} = 2\lambda T^2(1+T)^{-3}$, so $S_{\ln\rho,\,M} = S_{\ln\rho,\,V}$ in the $T = 0$ (lognormal) limit as expected (H13). Also note that, unlike a lognormal PDF, this distribution has a maximum value $\ln\rho = \Upsilon$ above which $\mathcal{P}_l(\ln\rho) = 0$. This property is entirely physical: the probability of encountering $\rho > \langle\rho\rangle(L/l)^3$ is identically zero, since there is not enough mass in the system.

To relate the mathematical model (5) to the physical properties of supersonic turbulence, we require two extra pieces of information: (i) how the average number of events λ relates to the physical scale l, and (ii) how the parameter T (the size of an event) relates to the physical properties of a shock structure. We tackle these issues in the next two sections.

3.2 Variation with scale and mass conservation

In this section, we relate the average number of events λ to the physical scale l. To explore this, it is helpful to consider the PDF of $\delta\ln\rho$ for a single step, which is illustrated in Fig. 2. This is a mixture distribution of $\mathcal{P}_\delta$ with probability $a/(1+a)$ (there is an event), and a (shifted) δ-function distribution with probability $1/(1+a)$ (there is no event). The parameter a thus determines the mean number of events over a given number of steps, implying the relation between λ and l is controlled by the proportionality constant between a and the (log) differential change in scale between l and l/Γ, $\delta r = \ln\Gamma \approx \Gamma - 1$.

The first constraint we can apply to the single-step PDF is that the total mass must conserved, $\langle\delta\rho\rangle_\delta = \langle\mathrm{e}^{\delta\ln\rho+\epsilon}\rangle_\delta = 1$ (where $\langle\cdot\rangle_\delta$ indicates the mean over the single-step PDF). This leads to the constraint

$$a = \epsilon\frac{1+T}{T}. \tag{6}$$

The second mass-conservation-related constraint we apply is that the *maximum* density scales as $(L/l)^\xi$; i.e. $\epsilon = \xi\delta r$, or

$$\lambda_l = Na = \xi\frac{1+T}{T}\ln\left(\frac{L}{l}\right). \tag{7}$$

This constraint simply states that the increase in density when there is no event is proportional to some power of the change in linear dimension. If $\xi = 3$, this implies that *all* of the mass from the old volume is contained into the new volume when there is no event, which will give a density PDF with the true (physical) density maximum $\rho_{\max} = (L/l)^3$. However, we find empirically (see Section 4) that $\epsilon = 3\delta r$ predicts a PDF that is insufficiently intermittent, and that $\epsilon \approx 1.5\delta r$ provides a closer fit to the data[2] (for the model of T

[1] The choice of an exponential distribution for $\mathcal{P}_\delta(\delta\ln\rho)$ can be motivated as the only choice other than a delta-function distribution (i.e. a standard log-Poisson process) that leads to a single fractal dimension for the most singular structures (He et al. 1998). However, since the system we model is not scale invariant anyway, our motivation for this choice is primarily simplicity – the PDF may be written in a simple closed form without unphysical discrete jumps – and other choices give similar results.

[2] Although it is difficult to motivate a particular value for ξ, a value $\xi < 3$ fits our intuitive picture of a collection of shocks. In particular, even in the case when all of the mass is contained in infinitely thin shocks, most of these will be at some angle to the changing volume and thus a 'no event' step will involve losing a small part of this mass. For example, if such shocks on average intersect one dimension of the changing volume, this would imply

proposed in Section 3.3 below). We thus consider ξ to be a parameter of our model, which should be universal for all turbulence with $\mathcal{M} \gg 1$. Note that a deviation from this general form (i.e. if ϵ was not taken proportional to δr but chosen through some other method) would imply that the geometrical properties and distribution of the shocks was a function of scale. This would be inconsistent with the existence of an inertial range in the $\mathcal{M} \gg 1$ limit, but could in principle occur for $\mathcal{M} \sim 1$.

3.3 Relating T to the size of a shock

In this section, we relate the intermittency parameter T to the physical properties of individual shocks (see also H13 appendix). In the moderate-$\mathcal{M}$ regime, which is both physically relevant and most commonly probed by current numerical simulations, the system is not scale invariant and T must be an increasing function of $\mathcal{M}$, eventually asymptoting to some value $T \lesssim 1$.[3] Here, we propose a simple physically motivated model, which gives a good match to simulations (see Section 4), is easily extendable to other equations of state (see Section 5) and is based on well-accepted ideas put forth in previous literature (Passot & Vázquez-Semadeni 1998; Padoan & Nordlund 2011; Price et al. 2011; Molina et al. 2012). However, a variety of other possibilities exist, some of which could prove similarly successful and also be physically motivated. This is discussed further in Section 3.5.

'Shocks' will be taken as regions that are extended across the volume being considered (i.e. of size l) in two dimensions, and be of small but finite extent in the other dimension (Federrath et al. 2008). Concentrating for now on the simplest isothermal equation of state, the Rankine–Hugoniot conditions for mass and momentum conservation of ρ and v on either side of a shock are

$$\rho_1 v_{c1} = \rho_2 v_{c2}, \quad \rho_1 \left(v_{c1}^2 + c_s^2\right) = \rho_2 \left(v_{c2}^2 + c_s^2\right), \tag{8}$$

where v_{ci} is the velocity on side i in the direction perpendicular to the shock. Solution of equation (8) with $v_{c0} = b\mathcal{M}c_s$, based on some average Mach number $\mathcal{M}$ and average fraction of the velocity in compressive modes b (Padoan, Nordlund & Jones 1997; Passot & Vázquez-Semadeni 1998; Federrath, Klessen & Schmidt 2008), leads to an approximate relation for the density contrast:

$$\frac{\rho_1}{\rho_0} \sim b^2\mathcal{M}^2. \tag{9}$$

As common in previous works (Padoan et al. 1997; Passot & Vázquez-Semadeni 1998; Lemaster & Stone 2008), we use equation (9) to relate the density in a shock to that outside the shock, for each scale in the turbulence.

The important quantity controlling T is the average density jump that occurs when a shock is removed from the volume being considered. This depends on both the density contrast and the volume of shocks, which requires a measure of their physical width $r_{\rm shock}$. Given that there is only one scale in the system – the sonic scale $l_{\rm sonic}$, at which $v \sim c_s$ – we effectively have only one choice: that the shock has width $\kappa l_{\rm sonic}$, with $\kappa < 1$ some arbitrary parameter. Denoting v_l as the approximate velocity difference across scale l and taking $v_l \sim (l/L)^\zeta \mathcal{M}c_s$ with $\zeta \sim 1/2$ (i.e. a velocity power spectrum $E \sim k^{-1-2/\zeta}$; Federrath 2013), one obtains $l_{\rm sonic}/L \sim \mathcal{M}^{-2}$ (for details, see section 5 in Federrath et al. 2016). Given the scaling of the density contrast equation (9), we see that this prescription $r_{\rm shock} \sim \kappa l_{\rm sonic}$, states that a fixed fraction of the mass ($\sim\kappa b^2$) is contained in individual shocks in the high-$\mathcal{M}$ limit.

Denoting the density in the shock $\rho_{\rm shock}$, the density outside the shock $\rho_{\rm out}$ and the average density over both regions $\rho_{\rm av}$, we see from equation (3) that

$$T \approx \ln \frac{\rho_{\rm out}}{\rho_{\rm av}} \tag{10}$$

(where we have neglected the offset ϵ since δr may be taken to be small). If we then take $\rho_{\rm shock} \sim \mathcal{M}^2 \rho_{\rm out}$ (absorbing b into κ),[4] then from mass conservation, viz.

$$\rho_{\rm out}(V - V_{\rm shock}) + \rho_{\rm shock} V_{\rm shock} = V \rho_{\rm av}, \tag{11}$$

where $V = l^3$ is the current volume and $\mathrm{V}_{\rm shock} \sim \kappa\, l_{\rm sonic} l^2$, we obtain

$$T \sim \kappa \left(1 - \mathcal{M}_l^{-2}\right) = \kappa \left(1 - \frac{L}{l}\mathcal{M}^{-2}\right), \tag{12}$$

assuming small κ.

We thus have a simple connection between the mathematical log-Poisson cascade described above and the physical size and density structure of shocks. It is worth noting that, given the isothermal shock-jump relation, making the assumption 'shock width $\sim l_{\rm sonic}$' (without some extra factor $\kappa < 1$) is technically inconsistent with Gaussian statistics, although it is often used to estimate the density variance (Padoan & Nordlund 2011; Molina et al. 2012; Federrath & Banerjee 2015). Systems with a larger proportion of the mass in a small number of shocks will have more intermittent statistics, since T controls the deviation from Gaussianity in equation (5). In Appendix B, we measure the shock sizes from simulations, finding reasonable agreement with the hypothesis that they scale as some fraction of $l_{\rm sonic}$ (see Fig. B1). Finally, we note that since we consider only density changes arising from shocks, we are explicitly neglecting all scales $l < l_{\rm sonic}$ where $\mathcal{M} \lesssim 1$; that is, the property $T(\mathcal{M}_l < 0) = 0$ is a consequence of our focus on shocks as drivers of density change and is not physical. Of course, subsonic motions do cause variation in density, and a model for $T(\mathcal{M} < 1)$ could be added to the supersonic model presented here if so desired (see e.g. Federrath et al. 2010; Nolan et al. 2015 for detailed discussion of the subsonic regime).

3.4 The model

We have now completely specified the full, scale-dependent statistics for the density field. Assuming for the moment that $T = \kappa$ is constant for $l > l_{\rm sonic}$ [this is true for $\mathcal{M} \gtrsim 5$; see equation (12)], we obtain a simple, two-parameter (κ and ξ) model for the density

$\xi = 3 - 1 = 2$, while if they intersect two it would imply $\xi = 3 - 2 = 1$. Thus, ξ is related to the density in unshocked regions, the (fractal) dimension of the shocks (Kowal et al. 2007; Federrath et al. 2009), and their distribution in angle.

[3] Note that $T > 1$ implies most of the mass in the system is contained in a single shocked structure, which might lead one to question whether the system is truly turbulent.

[4] We thus expect κ to differ between compressibly and solenoidally forced turbulence (Federrath 2013). Of course, b and κ are not quite equivalent, relating to the shock density contrast and volume, respectively. Nonetheless, retaining b separately leads to nearly the same relation as equation (12) (with $\kappa \to b^2\kappa$), but with unphysical ($T < 0$) behaviour near $\mathcal{M} = 1$. In any case, κ and b are each based on heuristic ideas, and assigning too much physical relevance to the details of this model is not particularly productive.

PDF as a function of scale:

$$\ln\rho \sim \mathcal{P}_l(\ln\rho) \approx T^{-1}\sqrt{\frac{\lambda}{u}}I_1(2\sqrt{u\lambda})\exp[-(\lambda+u)],$$
$$\lambda = \xi\left(1+\frac{1}{T}\right)\ln\left(\frac{L}{l}\right), \quad T=\kappa,$$
$$u = -\ln\rho + \xi\ln\left(\frac{L}{l}\right) \quad (\text{with } u>0), \tag{13}$$

for $l > l_{\rm sonic} = L\mathcal{M}^{-2}$. For $l < l_{\rm sonic}$, we take $\mathcal{P}_l(\ln\rho) = \mathcal{P}_{l_{\rm sonic}}(\ln\rho)$; that is, we neglect the subsonic contribution, which is minor for $\mathcal{M} \gg 1$ (Federrath et al. 2010). Note that the Mach number dependence is implicit in equation (13) through the dependence on $l_{\rm sonic}$, since higher $\mathcal{M}$ will lead to larger L/l with $l > l_{\rm sonic}$, and thus larger λ. The difference between compressible and solenoidal large-scale motions (e.g. due to forcing; Schmidt et al. 2009; Federrath 2013) is absorbed into the parameter κ that controls the width and density contrast of individual shocks.

Since T is not truly constant, the functional form of the PDF will differ somewhat from equation (13). In fact, even the H13 form [equation (5)] is not produced by the model with scale-dependent T, since the PDF of each step varies with l. While in principle one needs to take the convolution of a series of distributions $\mathcal{P}_\delta(\delta\ln\rho; T)$ with differing T, this is analytically unfeasible. Instead, the true PDF can be well approximated by taking the same algebraic form of the PDF equation (5), with mean and variance calculated from the true random process (Stewart et al. 2006; Hopkins 2015b). This gives

$$T_l = \frac{\int_{\ln l}^{\ln L} T(l')[1+T(l')]\mathrm{d}\ln l'}{\int_{\ln l}^{\ln L}[1+T(l')]\mathrm{d}\ln l'},$$
$$\lambda_l = \frac{\left(\int_{\ln l}^{\ln L}[1+T(l')]\mathrm{d}\ln l'\right)^2}{\int_{\ln l}^{\ln L} T(l')(1+T(l'))\mathrm{d}\ln l'}, \tag{14}$$

where T_l and λ_l are the 'averaged' values to use in the PDF equation (13), while $T(l) = \kappa(1-\mathcal{M}_l^{-2}) = \kappa(1-\mathcal{M}^{-2}L/l)$ is the 'local' value of T taken from the shock model. Although the integrals in equation (14) are straightforward analytically, the added complexity makes these forms inconvenient except for plotting. In any case, at high $\mathcal{M}$ the differences compared to the $T=$ constant assumption of equation (13) are modest, since $\mathcal{M}_l^{-2}$ falls off steeply away for $l > l_{\rm sonic}$.

3.5 Choices made and other possibilities

Through out the previous sections we have endeavoured to formulate a model with as few parameters as possible, based on simple physical considerations. Indeed, equation (13) involves just two physically motivated free parameters, κ and ξ, to describe the variation in a function, $\mathcal{P}_l(\ln\rho)$, across all scales in the system, for a wide range of turbulent Mach numbers. We shall see below (Section 4) that the model works relatively well in comparison to numerical simulations, both for measures that consider the variation in global parameters with physical parameters (e.g. the variance–Mach-number relation) and for measures in individual simulations (e.g. power spectra and the density PDF).

However, we feel it useful to reiterate the choices that have been made throughout the derivation, since some of these can be evaluated directly from simulation (or perhaps observational data). In this way, one might imagine calibrating certain aspects of the model, for example, to improve the accuracy of star formation models. Given the success of the general shape of the PDF (H13), here we consider various aspects of the model that might be modified (and the possible utility in doing so), while retaining the compound-log-Poisson structure.

The single-step PDF: The choice of an exponential PDF for the size of a single jump $\mathcal{P}_\delta$ (Fig. 2) was in part arbitrary, for the sake of convenience. For example, a similar $\mathcal{P}_\lambda$ is obtained with a δ function (this is used in most subsonic intermittency models; She & Leveque 1994; Boldyrev 2002; Hopkins 2015b; Mallet & Schekochihin 2017), or with a shifted Gamma distribution. There are, however, some important properties, which, if not satisfied would cause $\mathcal{P}_l(\ln\rho)$ to look quite different. In particular, the presence of an absolute maximum for $\delta\ln\rho$ is important, since without this $\mathcal{P}_l(\ln\rho)$ extends to infinitely high densities. Further, if this maximum is not equal to ϵ, the value of $\delta\ln\rho$ when there is no event, one has a similar problem, because as $\Gamma\rightarrow 1$ there is a non-zero probability of an arbitrarily large number of events n.

Number of structures encountered: The proportionality between the probability of an event a and the jump size in scale $\delta r = \ln\Gamma$ is an important parameter that controls the level of intermittency for a given variance (i.e. the relation between $S_{\ln\rho}$ and T). In our model this is set by relating a to ϵ using $\langle\rho\rangle = 1$, then setting ϵ to the maximum possible density from a volume change in ξ dimensions $\epsilon = \xi\delta r$. While the proportionality between ϵ and δr simply relies on having a density field that is approximately scale invariant, the value for ξ is less well constrained. It is reasonable to expect $\xi < 3$, although its exact value depends on properties of the density field, such as the density in 'unshocked' regions. We have found empirically that a range $1 \lesssim \xi \lesssim 2$ gives a more accurate match to data, (the exact value is hard to constrain given the significant scatter seen in numerical results; see Fig. 4). This is presumably related to the maximum density being closer to L/l, as opposed to $(L/l)^3$, because compression happens primarily along one dimension (in other words, the maximum possible density is when all of the mass is contained within a single 2D shock). As a further argument for the plausibility of $1 \lesssim \xi \lesssim 2$, we note the recent study of Federrath (2016), which argues for the importance of density filaments that form at the intersection between planar shocks. Such a mix of filamentary and planar high-density regions heuristically suggests a maximum density scaling between L/l and $(L/l)^2$ (i.e. ξ between 1 and 2). It is also reasonable to expect that ξ might be somewhat larger for compressibly forced simulations: because these have more intense shocks and thus (likely) higher maximum densities for the same L/l, we would expect a larger value for ξ.

Variation of T *with scale:* The variation in T with scale, and/or with Mach number, is difficult to constrain precisely, primarily because the system is *not* scale invariant at moderate Mach number. Specifically, although the sonic scale $l_{\rm sonic}$ is the only physically important scale above the viscous scale (which is ideally well into the subsonic regime), and so shock widths should eventually scale with $l_{\rm sonic}$, it is hard to know how far above $l_{\rm sonic}$ it is necessary to go before this occurs. This implies that the shock width could deviate from $\sim\kappa\mathcal{M}^{-2}$; in other words, κ could depend on $\mathcal{M}$ at moderate $\mathcal{M}$. Nonetheless, the general form of T – an increase at low $\mathcal{M}$, followed by a flattening at high $\mathcal{M}$ – is quite robust. Based on fits to numerical simulations and power spectra (see Sections 4.1 and 4.4), we have found that the form (12) slightly underpredicts the increase in T with $\mathcal{M}$; i.e. it should asymptote more slowly to $T=$ const. Physically, this implies that the fraction of the mass contained in

individual shocks should increase with $\mathcal{M}$ at higher values of $\mathcal{M}$ than suggested by $r_{\rm shock} = \kappa l_{\rm sonic}$.

Another possible uncertainly stems from our assumption that the velocity scales as $v_l/c_s \sim \mathcal{M}(l/L)^\zeta$ with $\zeta = 1/2$, and that shocks are 2D structures. If these parameters differ from these fiducial values (as may be the case; see e.g. Federrath et al. 2008; Federrath 2013), this will also change the variation of T with $\mathcal{M}$. The sign of the change is such that an increase in shock dimension,[5] or an increase in ζ, causes T to increase with $\mathcal{M}$ (i.e. T changes from being constant at $\mathcal{M} \gg 1$ to being a slowly increasing function of $\mathcal{M}$, although it must eventually flatten out).

Finally, it worth emphasizing why it can be hard to describe the full density statistics with two parameters. The primary difficulty is that, unlike models that assume lognormal statistics, the density variance and its variation with scale are now entangled with the distribution's intermittency. For example, the prediction for the density variance–Mach number relation depends strongly on T (or κ) and its variation with $\mathcal{M}$. Thus, any attempt to increase the variance by increasing κ creates unphysically intermittent density distributions. One can partially compensate for this by modifying ξ, higher values of which will increase the variance without much changing the intermittency, but this only works to some degree (e.g. $\xi < 3$). Further, it is then necessary to match any measurements to different simulations across a variety of Mach numbers. In this regard, similar models for subsonic turbulence (e.g. She & Leveque 1994) are more easily constrained: there are more known parameters (e.g. the power spectrum), and any model is required to fit only one simulation (there is an inertial range, so there is no requirement for a model that remains accurate across all Mach numbers). In the tests of the following section, we compare directly to data from numerical simulations. The most stringent of these tests is the explicit calculation of $\mathcal{P}_l(\ln\rho)$ in Section 4.3, which probes the variation in both density variance and intermittency with scale.

4 PREDICTIONS AND NUMERICAL COMPARISONS

In this section, we outline the main predictions of the model, comparing these to results from previous works and several numerical simulations. In addition to the basic functional form of the PDF, which is well-known to accurately match simulations (Hopkins 2013b; Federrath 2013; Konstandin et al. 2016), the model's scale dependence implies we can predict measures of the statistical variation with scale, such as power spectra and structure functions. Each of these agrees within uncertainty to results from numerical simulations. In addition, in Appendix B we test the shock width hypothesis $r_{\rm shock} \sim \kappa l_{\rm sonic}$, to check basic consistency with the model for T that was laid out in Section 3.3.

A direct comparison of the model to numerical simulations is hindered somewhat by the significant scatter seen across results in previous literature. This necessarily implies that, whatever fiducial model parameters (κ and ξ) we choose to fit a collection of previous results, any particular simulation will not be perfectly fit by this choice. Thus, when fitting the density PDF of an individual simulation in a way that appears satisfactory 'by eye' (e.g. Figs 3a and 5), we will necessarily be required to modify the fiducial model parameters somewhat. Based loosely on the high-resolution simulations of Federrath (2013), we suggest fiducial parameters of $\kappa \approx 0.2, \xi \approx 1.5$ for solenoidally forced turbulence, and $\kappa \approx 0.45$, $\xi \approx 2$ for compressively forced turbulence (these are also physically reasonable based on the theoretical arguments of Section 3.5). Other results, including our own simulations in Section 4.3, require only minor modifications to this choice across a large range in $\mathcal{M}$.

[5] This can take on non-integer values if the shocks have a fractal structure on scales $l \gg l_{\rm sonic}$ (see Federrath et al. 2009).

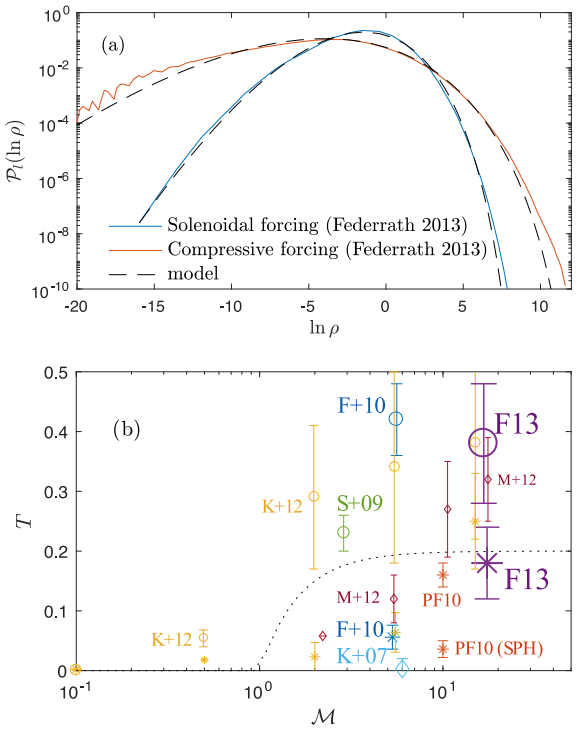

Figure 3. (a) Comparison of the model (dashed black lines) to the volume-weighted density PDF from the highest resolution simulations currently available (blue and red lines): those of Federrath (2013) at 4096^3. The simulations use solenoidal (blue line) or compressive (red line) forcing, and have Mach numbers of 17.4 ± 1.1 and 16.7 ± 1.1, respectively. The fits shown have parameters $T = 0.2$, $S = 4.1$ for the solenoidally forced simulation [this corresponds to model parameters $\kappa \approx 0.24$, $\xi \approx 1.5$ using equation (14)], and $T = 0.38$, $S = 12.5$ for the compressively forced simulation [this corresponds to model parameters $\kappa \approx 0.45$, $\xi \approx 2.1$ using equation (14)]. (b) Measurements of T taken from a variety of simulations as a function of Mach number. Circle markers indicate simulations with compressive forcing, asterisks indicate solenoidal forcing and diamond markers indicate forcing that involves some mix of compressive and solenoidal modes. Marker sizes are scaled by simulation resolution (a simulation at N^3 has a marker size $\propto \sqrt{N}$) so as to heuristically emphasize the most significant data points. In addition to the 4096^3 simulations of Federrath (2013) (purple, labelled F13), this data is taken from H13 based on Federrath et al. (2010) (blue, labelled F+10, 1024^3), Price & Federrath (2010) (red PF10, 512^3; SPH refers to their smoothed-particle hydrodynamics simulation), Konstandin et al. (2012) (yellow K+12, 512^3), Schmidt et al. (2009) (green S+09, 768^3), Kritsuk et al. (2007) (light blue K+07, 1024^3), Molina et al. (2012) (maroon M+12, 256^3). H13 lists the relevant parameters for each simulation in table 1 (see also H13 fig. 3). The dotted line shows the function $T = \kappa(1 - \mathcal{M}^{-2})$ for $\kappa = 0.2$, which was proposed in Section 3 based on the physical size and density contrast of shocks. Note that compressible forcing leads to higher intermittencies (T). With the large scatter in the data, it is unclear whether the simple model for T proposed in Section 3.3 is correct, but there does not appear to be a strong further increase in T with $\mathcal{M}$ for $\mathcal{M} \gtrsim 7$ (see Section 3.5).

4.1 Density PDF: basic form

The shape of the density PDF is shown for illustration purposes in Fig. 3(a), which compares the form (5) to the numerically measured density PDF from the highest resolution isothermal supersonic turbulence simulations yet run, from Federrath (2013). As also shown in Hopkins (2013b) for a range of other simulations, the quality of the agreement with the numerical PDF is impressive, with close-to-perfect agreement seen far into the tails of the distribution, where the deviation from lognormality is very significant.

The variation of the intermittency parameter T with Mach number is the first prediction of the model beyond Hopkins (2013b). In particular, we predict that T should be approximately constant with $\mathcal{M}$ for $\mathcal{M} \gg 1$, but may depend on the degree of compressibility (i.e. the ratio of solenoidal to compressive motions) of the turbulence, because this will change the shock density contrast b and/or shock width (compared to $l_{\rm sonic}$). As shown in Fig. 3(b), these behaviours are indeed observed in simulation data (in so far as the very large uncertainties permit). In particular, T is very small for $\mathcal{M} < 1$, rises rapidly to $\mathcal{M} \sim 3 \rightarrow 5$, then appears to stay constant between 0.15 and 0.4 above this (possibly with some slow increase with $\mathcal{M}$, more simulations are needed to address this more accurately). In addition, the 4096^3 simulations of Federrath (2013) show a large difference in T between solenoidally and compressively forced simulations ($T \approx 0.2$ and $T \approx 0.4$, respectively) as expected from the arguments in Section 3.3. Note that the intermittency (T) seems particularly sensitive to numerical details, such as the numerical method and resolution. For example, fig. 5 of Federrath (2013) illustrates how under-resolved simulations can overestimate T, while the Price & Federrath (2010) SPH simulations and Kritsuk et al. (2007) adaptive-mesh-refinement simulations produce very low intermittencies [see Fig. 3b, points PF10 (SPF) and K+07].

4.2 Density variance–Mach number relation

The density variance–Mach number relation is the simplest prediction of the model that depends on the variation of the PDF with scale (since higher Mach numbers have a larger range of scales $l \gtrsim l_{\rm sonic}$). Phenomenological arguments and simulations in previous literature suggest that this relation is

$$S_{\ln\rho} \approx \ln\left(1 + b^2\mathcal{M}^2\right), \tag{15}$$

with $b \approx 0.3$ solenoidally forced turbulence, or $b \approx 1$ for compressively forced turbulence (see e.g. Padoan et al. 1997; Passot & Vázquez-Semadeni 1998; Lemaster & Stone 2008; Price et al. 2011; Molina et al. 2012; Federrath & Banerjee 2015; Nolan et al. 2015). This 'standard' result does not include a distinction between the volume-weighted variance $S_{\ln\rho,\,V}$ and the mass-weighted variance $S_{\ln\rho,\,M}$, because these are identical when the density is distributed lognormally.

From equation (13) and $S_{\ln\rho,\,V} = 2T^2\lambda$, it is clear that since $\lambda \propto \ln(L/l)$, the model predicts $S_{\ln\rho} \propto \ln(\mathcal{M}^2)$ for $\mathcal{M} \gg 1$ (i.e. once $T \sim$ const.). This is a generic consequence of the constant accumulation of variance as a function of scale in log space, so this form is shared between any compound-log-Poisson model with scale-independent parameters. More precisely, the 'exact' result, using T from equation (12) and equation (14), is

$$S_{\ln\rho,V} = 2\xi\kappa(\kappa+1)\ln(\mathcal{M}^2) + \kappa\xi(1-\mathcal{M}^{-2}) \times [2 + \kappa(3-\mathcal{M}^{-2})], \tag{16}$$

which becomes $S_{\ln\rho,\,V} \approx 2\xi\kappa(1+\kappa)\ln(\mathcal{M}^2)$ for $\mathcal{M} \gg 1$. While this form is not identical to $S_{\ln\rho} \approx \ln(1 + b^2\mathcal{M}^2)$, it can be very similar for $\mathcal{M} \gtrsim 5$ depending on the constant of proportionality. The question then becomes: For a reasonable intermittency (i.e. a value of T that matches that measured from the PDF), does the model prediction $S_{\ln\rho,\,V} \approx 2\xi\kappa(1+\kappa)\ln\mathcal{M}^2$ [or equation (16)] also match the measured $S_{\ln\rho,\,V}$?

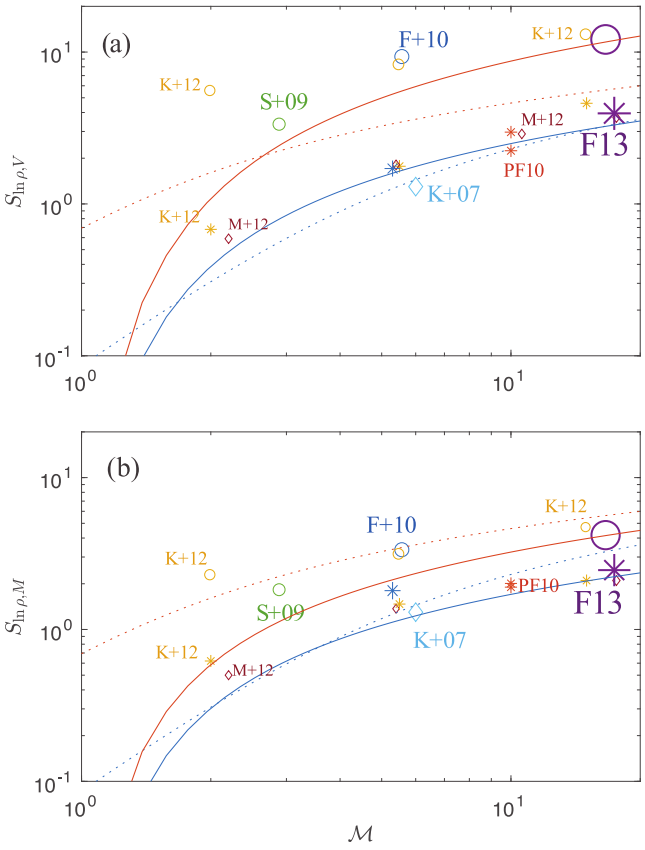

Figure 4. Density variance–Mach number relation (for $\mathcal{M} > 1$) as predicted by the model (solid lines) and as measured from simulations in previous literature (markers). Panel (a) shows the volume-weighted variance $S_{\ln\rho,\,V}$, while panel (b) shows the mass-weighted variance $S_{\ln\rho,\,M}$. We show model predictions for $\xi = 1.5$, $\kappa = 0.2$ (blue, lower curve) or $\xi = 2$, $\kappa = 0.45$ (red, upper curve) to indicate a range of values that might apply to solenoidally and compressively forced turbulence (see Fig. 3). The dotted curves show the standard result $S_{\ln\rho} \approx \ln(1 + b^2\mathcal{M}^2)$ (Padoan et al. 1997) for $b = 0.3$ (blue) and $b = 1$ (red), respectively, for solenoidal and compressive forcing (Federrath et al. 2010). Note that for the standard result, the volume-weighted and mass-weighted variances are the same since the PDF is assumed to be lognormal. The simulation results are as listed in Fig. 3(b) (see H13), with the marker size again scaled by simulation resolution. Although there is a large amount of scatter in the simulation results, the proposed model fits the data at least as well as the standard fit and captures the difference between the volume- and mass-weighted distributions.

In Fig. 4, we compare the model prediction with the standard result $S_{\ln\rho} \approx \ln(1 + b^2\mathcal{M}^2)$ and various previous simulation results for $\mathcal{M} > 1$. This is done for both the volume-weighted variance (Fig. 4a) and the mass-weighted variance (Fig. 4b). We plot $S_{\ln\rho}$ in each case for the model parameters $\kappa = 0.2$, $\xi = 1.5$, and $\kappa = 0.45$, $\xi = 2$, which were chosen as reasonable 'fiducial' values for solenoidally and compressibly forced simulations, respectively.[6] Given the significant scatter, the agreement of the model prediction is decent, and it seems fair to say

[6] Note that the values of T_l from equation (14) are somewhat lower than κ.

that the model relation is of (at least) a similar quality to the standard result. We also see better agreement between the predicted and measured mass-weighted variances $S_{\ln\rho,M} = (1+T)^{-3} S_{\ln\rho,V}$, compared to the standard result. This is unsurprising since the effects of intermittency generically act to reduce $S_{\ln\rho,M}$ compared to $S_{\ln\rho,V}$, and we have already seen that T increases with increasing $\mathcal{M}$ (Fig. 3b). It is also worth reiterating that the model neglects the subsonic contribution to the variance, which is significant for simulations with $\mathcal{M}$ approaching 1, and is the cause of the variance underprediction at low $\mathcal{M}$.

4.3 Density PDF as a function of scale

A more stringent test of model predictions is to compute $\mathcal{P}_l(\ln\rho)$ directly from simulation, viz., bin the density into volumes of size l^3 then compute the density PDF. Unfortunately, so far as we are aware there is no measurement of this in previous literature, despite its physical, as well as theoretical, relevance. We have thus run a variety of isothermal turbulence simulations to make such measurements directly. These simulations use the GIZMO code (Hopkins 2015a; Hopkins & Raives 2016) with the MFM method and 256^3 elements. Although this resolution may be relatively low by modern standards, the Lagrangian nature of the MFM method more accurately captures the high-density shock regions by naturally having higher resolution in such regions (Price & Federrath 2010), and the MFM method has proven very accurate in a wide variety of test problems (Hopkins 2015a; Hopkins & Raives 2016). That said, given the significant dependence of intermittency properties on resolution and numerical method (Price & Federrath 2010; Federrath 2013), it will be important to verify the scaling of these results with resolution. The simulations are forced by a Ornstein–Uhlenbeck process with an equal mix of solenoidal and compressive large-scale modes, as described in Bauer & Springel (2012). Different forcing strengths are used to drive turbulence across a range of Mach numbers. PDFs are calculated by depositing the density field on to a 512^3 uniform grid using a Gaussian kernel for each Lagrangian mesh element (of width $\sigma = \sqrt{3/40}h$, where h is the cell smoothing length; Dehnen & Aly 2012; Hopkins 2015a), then averaging over successively larger volumes to form the PDF as a function of scale. This grid-based method agrees with the volume-weighted PDF calculated directly from the Lagrangian mesh for the finest $N = 512$ grid.[7]

Fig. 5 compares results from simulations at (a) $\mathcal{M} \approx 12$, (b) $\mathcal{M} \approx 40$ and (c) $\mathcal{M} \approx 7$ with model predictions (shown with dashed lines). In each case, to specify the model, we use the measured $\mathcal{M}$, $\xi = 1.3$, and choose κ to match the variance on the smallest supersonic scales.[8] We then compare the model predictions for the intermittency and variance on scales $l > l_{\rm sonic}$ with that measured by binning the density on the same scale (for three values of l). More precisely, in the $\mathcal{M} \approx 12$ simulation (Fig. 5a), $l_{\rm sonic} \approx L/150$, so we compare data on scales $l = L/128$, $l = L/32$ and $l = L/4$; in the $\mathcal{M} \approx 40$ simulation (Fig. 5b), $l_{\rm sonic} \approx L/1600$, so we compare data on scales $l = L/512$, $l = L/32$ and $l = L/4$; and in the $\mathcal{M} \approx 7$ simulation (Fig. 5c), $l_{\rm sonic} \approx L/50$, so we compare data on scales $l = L/32$, $l = L/8$ and $l = L/4$. This method is chosen to explicitly remove the subsonic scales from the comparison, since these are not included in the model.

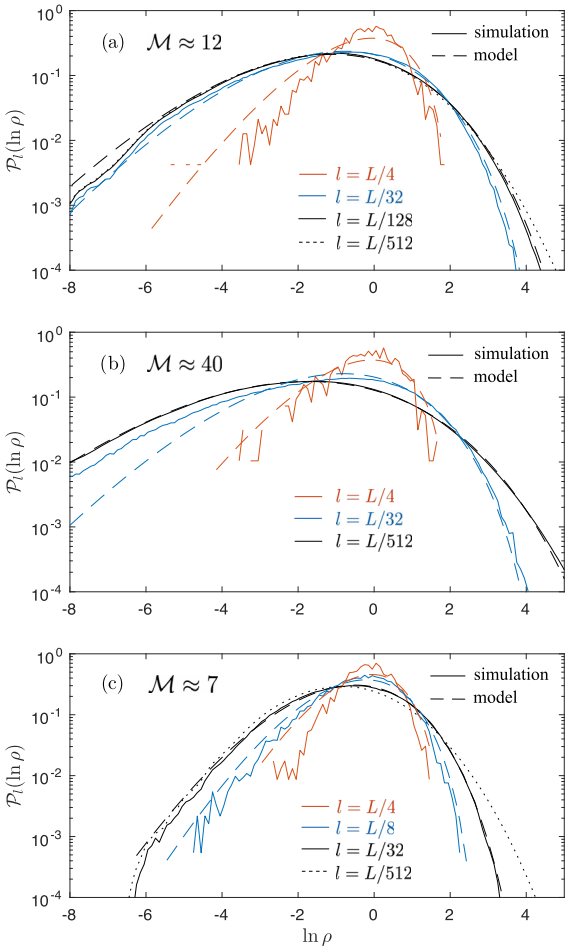

Figure 5. Density PDF $\mathcal{P}_l(\ln\rho)$ for (a) $\mathcal{M} \approx 12$, (b) $\mathcal{M} \approx 40$ and (c) $\mathcal{M} \approx 7$ turbulence. In each panel, solid curves show the measured PDFs of the density averaged over a variety of scales, while dashed curves show model predictions for the same parameters. In each case, we only compare the model with simulations on scales $l > l_{\rm sonic} \approx \mathcal{M}^{-2}$; specifically for (a) $\mathcal{M} \approx 12$ with $l_{\rm sonic} \approx L/150$, we take $l = L/128$ (black, widest curve), $l = L/32$ (blue, middle curve) and $l = L/4$ (red, narrowest curve); for (b) $\mathcal{M} \approx 40$ with $l_{\rm sonic} \approx L/1600$, we take $l = L/512$ (black), $l = L/32$ (blue) and $l = L/4$ (red); and for (c) $\mathcal{M} \approx 7$ with $l_{\rm sonic} \approx L/50$, we take $l = L/32$ (black), $l = L/8$ (blue) and $l = L/4$ (red). In this way the subsonic contributions to the variance, which are not included in the model, are explicitly removed from the simulation results. We take $\kappa = 0.24$ (a and b) and $\kappa = 0.18$ (c), with $\xi = 1.3$ in all cases, illustrating its success across a range of $\mathcal{M}$ with little change to κ (see text for discussion). In each panel, we also show the PDF on the smallest scales measured, $l = L/512$, and the difference between this and the solid black curve explicitly illustrates the contributions from subsonic scales in (a) and (c). Although there seems to be a slight underprediction of the intermittency on the largest scales $l = L/4$, the numerical PDF is likely also affected by discretization and insufficient statistics, which would tend to raise the low-density tail.

[7] There is some discrepancy at the lowest densities. This is expected because the density field deposited using the Gaussian kernel contains regions of lower density (the regions in between mesh elements) than that of the lowest density mesh elements. To find the true density in such regions one should use the 'gather' method for constructing gridded data, as this is actually used in the simulation (Hopkins 2015a); however, the Gaussian kernel method we use probably provides a better representation of the true field than the PDF from Lagrangian data, which is effectively undersampling the lowest density regions.

[8] This value of $\xi = 1.3$ works slightly better for these simulations than the $\xi = 1.5$ plotted in Fig. 4. As mentioned above (see introduction to Section 4) the significant scatter across simulations and numerical methods makes it inevitable that ξ (or κ) should have to be modified somewhat to fit an particular simulation set.

The agreement of the model to simulation is seen to be relatively good. In particular, identical model parameters ($\kappa = 0.24$ and the measured $\mathcal{M}$) give very good fits to $\mathcal{M} \approx 12$ and $\mathcal{M} \approx 40$ across a wide range of scales in the system. A similar, though slightly lower, value of κ ($\kappa = 0.18$) gives a very good fit at lower $\mathcal{M} \approx 7$. We attribute this difference in κ to the slight underestimation of the increase in T with $\mathcal{M}$ in the model [equation (12)]. Although the model possibly overpredicts the variance at the largest scale ($l = L/4$) in each case, it is worth noting that this scale is very close to the driving (at $\sim L/2$) and may be influenced by this. Further, the statistics at $l = L/4$ in each case are somewhat undersampled (there are only 64 values per density snapshot), and more values would tend to increase the low-density tail.

Finally, it is worth briefly mentioning the contribution of the subsonic scales to the full 512^3 density PDF (dotted line in each panel of Fig. 5). As can be seen from Fig. 5(a) and (c), the subsonic scales have the effect of decreasing the intermittency (i.e. making the distribution more lognormal) by contributing to the high-density tail. This should be expected, since the subsonic contribution will involve large numbers of small events (i.e. small T). While the effect appears more significant in the $\mathcal{M} \approx 7$ simulation (compared to the larger $\mathcal{M}$ cases), in fact, the absolute increase in the density maximum – i.e. the difference between $\rho_{\rm max}$ with and without the subsonic contributions – is about the same at $\mathcal{M} \approx 7$ and $\mathcal{M} \approx 12$. It appears larger at $\mathcal{M} \approx 7$ due to the smaller contribution to the variance from supersonic motions. This justifies our neglect of subsonic scales in the model, which is primarily intended for study of the $\mathcal{M} \gg 1$ limit. Because the subsonic contribution occurs on the very smallest scales of any simulation, and will thus presumably be affected by the numerical method, it is possible that they play a role in the wide scatter seen between different simulations in both the density variance–Mach number relation and the intermittency (see Figs 4 and 3). See Federrath et al. (2010) for further discussion.

4.4 Spectrum

As shown in Appendix A, the density power spectrum $\phi_\rho(k)$ is related to the variation in the second-order statistics (variance and mean) of the PDF with scale. In particular, for some variable s, the 1D power spectrum is

$$\phi_s(k) \sim \frac{\mathrm{d}}{\mathrm{d}k}\left(S_l + \bar{s}_l^2\right), \tag{17}$$

where S_l and $\bar{s}_l^2$ are the variance and mean of $\mathcal{P}_l(s)$. From equation (5), the volume-weighted variance of ρ can be calculated as (H13),

$$S_\rho = \exp\left(\frac{S_{\ln\rho}}{(2T+1)(T+1)}\right) - 1. \tag{18}$$

Then, using $\langle\rho\rangle = 1$, the simplified form for S, $S_{\ln\rho} = 2\xi T(1+T)\ln(L/l)$ (i.e. neglecting the scale variation of T),[9] and approximating $T \approx \kappa$ for $\mathcal{M} \gg 1$, one obtains the scaling

$$S_{\rho,l} + \bar{\rho}_l^2 \sim \left(\frac{l}{L}\right)^{-\frac{2\xi\kappa}{2\kappa+1}}. \tag{19}$$

The model thus predicts the ρ power spectrum

$$\phi_\rho(k) \sim k^{-\nu} \qquad \nu = \frac{1+2\kappa(1-\xi)}{1+2\kappa}, \tag{20}$$

in the limit $\mathcal{M} \gg 1$ for $1/k \gg l_{\rm sonic}$. This spectrum is somewhat less steep than k^{-1} (it is k^{-1} for $\kappa \to 0$ or $\xi \to 0$) and becomes less steep with increasing κ or ξ – for example, the parameters used in Fig. 5 give $\phi_\rho(k) \sim k^{-0.56}$. The same procedure for the power spectrum of $\ln\rho$, using $\langle\ln\rho\rangle = -S_{\ln\rho}(1+T)^{-1}/2$ leads to

$$\phi_{\ln\rho}(k) \sim \frac{2\xi\kappa}{9}\left(k^{-1} + \frac{3\ln k}{k}\right) + \mathcal{O}(\kappa)^2, \tag{21}$$

implying the power spectrum of $\ln\rho$ is not expected to be a power law but is close to $\sim k^{-1}$.

An important difference compared to some previous models of the supersonic density power spectrum (Saichev & Woyczynski 1996; Kim & Ryu 2005; Konstandin et al. 2016) is that we do not predict a density spectrum that approaches k^0 for $\mathcal{M} \gg 1$ (but see also Fleck 1996). Instead, our prediction is that the spectrum approaches some power law that depends relatively strongly on the intermittency of the density distribution (through κ). We thus predict a steeper spectrum for solenoidal compared to compressive forcing (because measured values of T are larger for compressive forcing); e.g. using the values $\kappa \approx 0.2$, $\xi \approx 1.5$, and $\kappa \approx 0.45$, $\xi \approx 2$ suggests that the density spectra scale as $\sim k^{-0.6}$ and $\sim k^{-0.05}$ for solenoidal and compressive forcing, respectively. We also predict that the $\ln\rho$ spectrum should not depend on the forcing so strongly, although it is also not a power law. These predictions are valid only for scales well above the sonic scale, since model parameters change significantly as $\mathcal{M}_l \to 1$.

With currently available simulation data, these predictions are difficult to verify or disprove. In particular, density power spectra are seen to depend significantly on resolution and numerical method (see e.g. Konstandin et al. 2016, fig. 8), and the true scaling remains somewhat unclear. Nonetheless, a number of works have studied spectra across a range of $\mathcal{M}$ (e.g. Kim & Ryu 2005; Kowal et al. 2007; Kritsuk et al. 2007; Federrath et al. 2009; Burkhart et al. 2010; Federrath & Klessen 2013; Konstandin et al. 2016), and we see reasonable agreement with most of these results. For example, Burkhart et al. (2010) finds that the spectrum saturates to $\sim k^{-0.5}$ for (solenoidally driven) MHD turbulence at high $\mathcal{M}$, while Konstandin et al. (2016) find a spectrum between $\sim k^{-1}$ and $\sim k^{-0.5}$ (depending on resolution) at $\mathcal{M} \gtrsim 10$ with mixed forcing, but see steeper spectra at lower $\mathcal{M}$. Similarly, the simulations presented in Fig. 5 have density power spectra consistent with predictions at high $\mathcal{M}$ ($\sim k^{-0.6}$ at $\mathcal{M} \approx 40$ and $\sim k^{-0.7}$ at $\mathcal{M} \approx 12$; comparable to the prediction at $\kappa = 0.24$, $\xi = 1.3$ of $\sim k^{-0.56}$) but are a little steeper at $\mathcal{M} \approx 7$ ($\sim k^{-0.9}$). This observed steeper spectrum at low $\mathcal{M}$ is actually expected from the model: the decrease in T with scale (i.e. before the plateau in Fig. 3b) is not taken into account in equation (20) and acts to steepen the spectrum significantly.[10] Nonetheless, based on simulation data, the 'true' density spectrum is likely to be somewhat steeper than predicted at moderate $\mathcal{M}$, which is possibly due to a slower increase in T with $\mathcal{M}$ than suggested by equation (12) (see below).

Perhaps the biggest possible inconsistency between the model and previous simulation results is seen in Federrath et al. (2009)

[9] The power spectrum can be derived analytically using the full integrals, equation (14); however, the resulting expressions are very complicated and no longer follow a power law at high k.

[10] For example, using $\kappa = 0.2$, $\xi = 1.5$ at $\mathcal{M} = 3$, one finds that the full predicted ρ spectrum [i.e. using the result (16) for $S_{\ln\rho}$] has a spectral slope $\sim k^{-2}$ at $l \approx 2l_{\rm sonic}$, although the spectrum is also not a power law (of course, at this point the subsonic scales also become important).

and Federrath & Klessen (2013), which each compare the density spectra in solenoidally and compressibly forced simulations. At low $\mathcal{M} \lesssim 5$, it is found that compressible forcing leads to a steeper density spectrum, while at higher $\mathcal{M}$, the spectrum is approximately independent of the forcing. While this may appear to be at odds with our predictions, it is worth noting that at lower $\mathcal{M}$ there is very little supersonic range (where $l > l_{\rm sonic}$) with which to compute a spectrum, and their results may be influenced by the subsonic scales and/or the driving routine directly (the simulations do not show obvious power laws; see e.g. Federrath et al. 2009, fig. 4). At higher $\mathcal{M}$, the results in Federrath & Klessen (2013) again lack well-defined inertial ranges in the density spectra (see their fig. 7), and the observed spectral scaling of $\sim k^{-0.1}$ for solenoidal turbulence at $\mathcal{M} \approx 50$ disagrees with our simulations (where $\sim k^{-0.6}$ was seen at $\mathcal{M} \approx 40$). Further simulations seem necessary to understand whether there is a meaningful difference in spectral scaling between solenoidally and compressively driven turbulence (as suggested by the model), or not (as suggested by Federrath & Klessen 2013). We reiterate that our basic prediction (for $\mathcal{M} \gg 1$) is for a flatter density spectrum in turbulence with a higher intermittency, and directly comparing measured intermittencies and spectra might be the most robust way to test this.

The results for the spectrum of $\ln \rho$ are even less well known, to our knowledge appearing in previous literature only in Beresnyak et al. (2005) and Kowal et al. (2007) (for MHD), and in Federrath et al. (2010). Federrath et al. (2010) reports $\ln \rho$ spectra of $\sim k^{-1.6}$ and $\sim k^{-2.3}$ for solenoidally and compressively forced turbulence respectively at $\mathcal{M} \approx 5.5$, while Kowal et al. (2007) report $\sim k^{-1.5}$ for $\mathcal{M} \approx 7$ (or perhaps flatter at the largest scales). However, because these simulations each have a relatively modest Mach number, a spectrum steeper than k^{-1} should be expected: they are in the regime where $T \neq$ const., which implies that the accumulation of variance with $\ln(l)$ is faster than linear. In Beresnyak et al. (2005), although at higher Mach number ($\mathcal{M} \approx 10$), they do not see a clear power law (it may be somewhere between $\sim k^{-1}$ and $\sim k^{-5/3}$), and are in the strongly magnetized regime where our results do not apply directly. Our simulations give somewhat unclear results, although also appear to show steeper $\ln \rho$ spectra than predicted (somewhere between k^{-1} and k^{-2} for the largest scales).

Overall, these results may suggest that the proposed $T(\mathcal{M})$ scaling, $T \sim \kappa(1 - \mathcal{M}_l^{-2})$ [equation (12)], reaches its high-$\mathcal{M}$ asymptotic value at a value of $\mathcal{M}$ that is somewhat too low[11] (see Section 3.5). If this effect was larger with compressible forcing (which might be expected, since $T(\mathcal{M} \gg 1)$ is larger), this could also explain the steeper spectra in the compressively forced case for low $\mathcal{M}$, as seen in Federrath et al. (2009) and Federrath & Klessen (2013). For the sake of example, if we assumed a different model with a faster increase in T, $\kappa \sim \mathcal{M}_l^2 \sim (l/L)\mathcal{M}^2$ at moderate $\mathcal{M}$, this would steepen the $\ln \rho$ power law to $\sim k^{-2}$, and the power law of ρ by a factor $\sim k^{-1}$. Study of higher resolution simulations at higher $\mathcal{M}$ (e.g. the simulations of Federrath 2013) is required to better assess our predictions. Nonetheless, the general arguments presented here should provide a useful framework for interpreting future results, even if the agreement is not perfect.

5 EXTENSIONS

Given the simple relation between the intermittency parameter T and the physical properties of shocks, it is possible to straightforwardly extend the model to situations with more complex physics. Here, we briefly consider turbulence with a non-isothermal polytropic equation of state, as relevant for various phases of the ISM (see e.g. Audit & Hennebelle 2005; Gazol & Kim 2013; Federrath & Banerjee 2015, and references therein). Other extensions – e.g. to supersonic MHD turbulence – are also possible using similar ideas. Note that the discussion and results in this section are intended to be of a qualitative nature. While the simple extensions we propose do give a reasonable match to simulation results, the purpose of the analysis is as much to illustrate the applicability of the shock model in Section 3, as to provide useful models for turbulent PDFs. With this in mind, some of the ideas discussed can likely be applied more rigorously if so desired; for example, to derive scalings for the low- or high-density tails of the PDF.

The key idea of the method – which is effectively that proposed in Passot & Vázquez-Semadeni (1998) extended to non-lognormal isothermal PDFs – is to assume that the primary effect of the non-isothermal equation of state is to modify the sound speed with the density. This in turn modifies the local Mach number and causes the shock density contrast to depend on the local value of the density. As shown in Federrath & Banerjee (2015), this behaviour is indeed seen as a correlation between ρ and $\mathcal{M}$ in the turbulent joint Mach number–density PDF, as well as leading to a useful estimate for the modified density variance–Mach number relation when applied to the shock-jump condition.

More precisely, for a polytropic equation of state $p/p_0 = (\rho/\rho_0)^\gamma$ (p is the pressure), one takes

$$\mathcal{M} \propto c_s^{-1} \sim \left(\frac{p}{\rho}\right)^{-1/2} \sim \rho^{(1-\gamma)/2} = \exp\left(\frac{1-\gamma}{2}\ln\rho\right), \qquad (22)$$

which is then used in the shock-density-contrast relation $(\rho_1/\rho_0) \sim b^2\mathcal{M}^2$. We thus see that with $\gamma < 1$, the system will have higher contrast shocks (compared to isothermal expectations) at high densities and lower contrast shocks at low densities, while $\gamma > 1$ leads to the opposite behaviour. Within our model, this causes the mean jump size $\langle \delta \ln \rho \rangle - \epsilon = T$ to depend on the local value of the density through the replacement of $\mathcal{M}_l$ with $\mathcal{M}_l \exp[(1-\gamma)/2\ \ln\rho]$, or

$$T \sim \kappa\left(1 - \frac{L}{l}\mathcal{M}^{-2}\mathrm{e}^{(\gamma-1)\ln\rho}\right), \qquad (23)$$

(with $T = 0$ if $\mathcal{M}_l^{-2}\mathrm{e}^{(\gamma-1)\ln\rho} > 1$).[12] It is clear that this form of T will decrease the low-density tail for $\gamma < 1$ due to the fast increase in $\mathrm{e}^{(\gamma-1)\ln\rho}$. In contrast, for $\gamma > 1$ the low-density tail will increase in probability, because the cascade can proceed further (go to smaller scales) before $T \to 0$, meaning individual $\delta \ln \rho$ jumps are larger.

[11] In other words, the dotted line in Fig. 3(b) should increase more slowly at low $\mathcal{M}$.

[12] Note that we have neglected a potentially important effect here, which is the change in shock contrast and width with γ due to the differing sound speeds on either side of the shock. The density contrast is derived in the form of a transcendental equation in Federrath & Banerjee (2015); however, the complexity of these expressions, as well as the necessity of deriving the density jump based on the shock width (which also must scale with the density), leads to complex systems of transcendental equations that are difficult to use in the model. The method is thus more similar to Passot & Vázquez-Semadeni (1998) as opposed to Federrath & Banerjee (2015). Although working this out correctly will certainly change the functional dependence of T on $\mathcal{M}$, as well as adding γ dependence into the shock-jump size (κ or some similar parameter), the key differences compared to an isothermal equation of state – in particular the reduction in the size of jumps at low (high) density for $\gamma < 1$ ($\gamma > 1$) – are retained in the much simplified version.

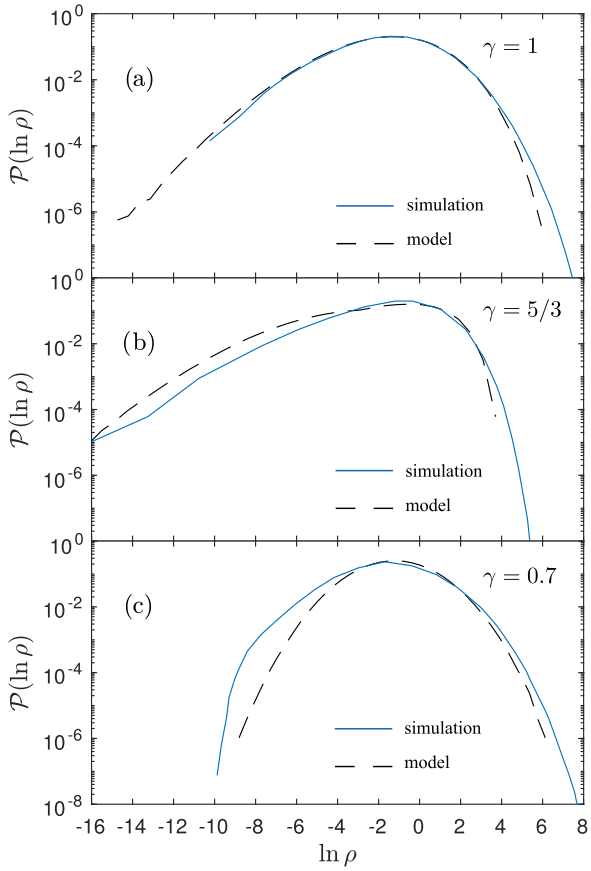

Figure 6. Predicted density PDF for non-isothermal polytropic turbulence (dashed black line) compared to the simulations of Federrath & Banerjee (2015) (solid blue line). The Mach numbers are $\mathcal{M} = 11.6$ ($\gamma = 1$), $\mathcal{M} = 13.3$ ($\gamma = 5/3$) and $\mathcal{M} = 8.4$ ($\gamma = 0.7$), and the Federrath & Banerjee (2015) simulations were run at resolutions of 2048^3 ($\gamma = 5/3$ and 0.7) or 1024^3 ($\gamma = 1$). For the model, there are *no* free parameters used to fit the non-isothermal PDFs: we use the same value of $\kappa = 0.26$ in each case (this was chosen to match the $\gamma = 1$ distribution) and the physical value of the Mach number listed above. Although the fits here are not perfect, the method does a good job at capturing the qualitative change in the PDF shape considering there are no free parameters. In addition, the low-density regions where the largest discrepancies are seen have more significant numerical error bars and resolution dependence (see Federrath & Banerjee 2015, fig. 4), and there is presumably some contribution to the high densities from the subsonic scales (see e.g. Fig. 5a).

Results are illustrated in Fig. 6, which shows the comparison of this model to the volume-weighted PDF data from Federrath & Banerjee (2015). Despite a variety of limitations (see below), we see that the qualitative trends for the PDFs are captured well. In particular, at $\gamma > 1$ we see a faster fall off at high densities and a long tail at low densities with a slope that generally matches the simulation PDF, while at $\gamma < 1$ the lower densities are significantly reduced (although the prediction is too severe, cutting off at somewhat higher densities than in the simulation). Given the possible resolution dependence of the low-probability regions (see Federrath & Banerjee 2015, fig. 4, second row), the overall agreement is encouraging.

To calculate the illustrated PDFs, we use a simple Monte Carlo method with the prescription for T taken from equation (23), using the same value of $\kappa = 0.26$ and $\xi = 1.5$ in each case, with $\mathcal{M}$ as quoted in Federrath & Banerjee (2015). We thus have no free parameters to aid in the fitting for the non-isothermal PDFs in Fig. 6, and the fit could potentially be significantly improved by optimizing over κ. (Note that, based on physical arguments, κ *should* be modified somewhat with γ, becoming smaller with increasing γ due to the change in shock-jump condition calculated in Federrath & Banerjee 2015.) For consistency with the simulations, we retain scales down to $L/l \sim 1024$ in the Monte Carlo cascade (estimating twice the grid scale as the minimum resolvable scale). In addition, the subsonic scales are included in the illustrated PDFs making a direct comparison difficult, and at this Mach number $\mathcal{M} \sim 12$ these will have a minor but observable contribution to the high-density probability (e.g. compare the $L/l = 512$ and $L/l = 128$ curves in Fig. 5a, which explicitly shows the subsonic contribution at a similar Mach number). It is interesting to note that the low-density tail of the $\gamma > 1$ PDF may flatten further with resolution beyond 2048^3 (this is seen in our model if a wider range of scales is kept): even at low velocities, very low density regions remain supersonic with strong shocks that cause large-density contrasts (see also Federrath & Banerjee 2015, fig. 4).

Finally, we note that similar ideas can be applied to MHD turbulence. A simple method, used in Padoan & Nordlund (2011), Molina et al. (2012) and Federrath & Banerjee (2015) for the density variance–Mach number relation, is again to consider how the Mach number is altered by the influence of the magnetic field on the total pressure,

$$\mathcal{M} \sim c_s^{-1} \sim \frac{\rho}{p_{\rm gas} + p_{\rm mag}} \sim c_{s0}^{-1}(1 + \beta^{-1})^{-1/2}, \tag{24}$$

where $\beta = p_{\rm gas}/p_{\rm mag}$ is the ratio of thermal to magnetic pressure and c_{s0} is the sound speed without the magnetic field. To apply this form of $\mathcal{M}$ to the density PDF model, we need a prescription for how B changes with ρ. While this remains uncertain, it unequivocally depends on turbulence parameters (e.g. Alfvén–Mach number and β; see Lithwick & Goldreich 2001; Cho & Lazarian 2003; Banerjee et al. 2009; Burkhart et al. 2009). As an example, taking $B \sim \rho^{1/2}$ (Crutcher 1999; Banerjee et al. 2009; Molina et al. 2012), we obtain a PDF of exactly the same functional form as the isothermal PDF, with shock sizes reduced by $1 + \beta^{-1}$ (this is effectively identical to the model of Molina et al. 2012). In contrast, a scaling $B \sim \rho^{\chi}$ with $\chi < 1/2$ acts to decrease the PDF at low densities, which is indeed seen in simulations (Molina et al. 2012). A similar effect would be seen if $B(\rho)$ became constant below some density threshold[13] (i.e. if the relation between B and ρ was not a power law, but a more sudden change). An interesting consequence of this is that the increased lognormality observed in MHD turbulence is probably not due to increased Gaussianity in the underlying turbulence. Instead, we may be seeing suppression of the low-density tail of a compound-log-Poisson distribution, which causes the PDF to appear lognormal even though the underlying turbulence could have similar intermittency properties (this is the same effect as in non-isothermal turbulence with $\gamma < 1$; see Fig. 6c). While there are many interesting (and astrophysically relevant) issues to explore here, we postpone such studies to future work due to the uncertainties regarding the scaling of B with ρ.

[13] This form is suggested by the observations of Crutcher et al. (2010), who report a lower density bound below which the density and magnetic field are uncorrelated. It is also expected on physical grounds because the turbulence will become Alfvénic in character ($v_A \propto B/\rho^{1/2} > v_l$) below some density (Lithwick & Goldreich 2001; Cho & Lazarian 2003).

6 DISCUSSION AND CONCLUSIONS

In this paper, we propose a simple phenomenological model to describe the distribution of density in supersonic turbulence. Given the turbulent Mach number $\mathcal{M}$ and two free parameters (κ and ξ) that relate to the physical properties of shocks, the model predicts the PDF of the density averaged over scale l, $\mathcal{P}_l(\ln\rho)$ [equation (13)]. Because $\mathcal{P}_l(\ln\rho)$ completely specifies the statistics of the density field, the model predicts all relevant statistic quantities of the density field: the density variance–Mach number relation, the density PDF and intermittency, power spectra and structure functions. We see reasonable agreement between model predictions, results from previous literature, and our own set of simulations. The model is also straightforwardly extendable to more complex gas physics (e.g. varied gas equations of state, or MHDs) and shows decent agreement to recent simulations of non-isothermal turbulence (Federrath & Banerjee 2015).

The main predictions and results are summarized as follows:

(i) The gas density, averaged across scale l, is distributed according to the PDF suggested in H13; see equation (5). The intermittency parameter T controls the deviation from lognormality ($T = 0$ describes a lognormal distribution). This form of the PDF matches numerical measurements very well across many orders of magnitude (see Fig. 3). The intermittency and variance of the density PDF change with scale and Mach number.

(ii) The density is arranged into a random collection of shocks across all scales (above the scale at which the velocity becomes subsonic). The physical size of the shocks and their relative density contrast are controlled by the model parameter $\kappa \sim r_{\rm shock}/l_{\rm sonic}$, which sets T in the density PDF. Larger and higher density shocks create a density distribution that is more intermittent.

(iii) Mathematically, the density is constructed via a compound-log-Poisson process. The size of each individual event (shock) is distributed according to an exponential distribution (see Fig. 2).

(iv) The number of shocks encountered across some range in scales is set by the maximum density that is possible if the gas is compressed in ξ dimensions, where ξ is effectively a model parameter (its maximum is $\xi = 3$). Empirically, we find that $1 \lesssim \xi \lesssim 2$ gives a reasonable match to simulation data (depending on forcing and the numerical scheme), but, given the significant scatter in previous results and between numerical methods (see Fig. 4) this estimate is approximate.

(v) The parameter κ also differs between compressibly and solenoidally forced turbulence, because the shocks are more intense with compressive forcing (Federrath 2013). Based on numerical PDFs (see also Appendix B for a more direct measurement) κ ranges from $\sim$0.2 to $\sim$0.5 as the compressive fraction is increased (but could also be lower in some cases; see Pan et al. 2016).

(vi) The density variance–Mach number relation is similar to the standard result $S_{\ln\rho} \approx \ln(1 + b^2\mathcal{M}^2)$ for values of κ that match the observed intermittency.

(vii) The model predicts a density power spectrum for $\mathcal{M} \gg 1$ between $\sim k^{-1}$ and $\sim k^0$, depending on κ and ξ [see equation (20)], thus predicting a spectrum that is directly related to the intermittency. The power spectrum does not necessarily approach k^0 in the $\mathcal{M} \rightarrow \infty$ limit.

(viii) We neglect the influence of subsonic motions on the density PDF, since these are negligible at high $\mathcal{M}$ and stem from different physical processes. However, subsonic contributions to the PDF can reduce the intermittency (see Fig. 5) and may be responsible for some of the scatter seen across numerical results (see Fig. 3).

(ix) Extensions to the assumption of isothermal neutral gas may be included by considering the sound speed, and thus local shock size, to be a function of local gas density (as in Passot & Vázquez-Semadeni 1998). This leads to density PDFs that agree well with those observed in simulations with a non-isothermal equation of state (see Fig. 6).

More generally – particularly considering the very significant scatter between simulations reported in previous literature (see e.g. Figs 3 and 4) – the model can be seen as a framework for understanding density statistics and intermittency, describing how different statistical measures might be compared to provide interesting information about the underlying structures. This is particularly true in the moderate-$\mathcal{M}$ regime, which is most relevant physically and easiest to study numerically, but is not blessed with a true inertial range (due to the proximity of $l_{\rm sonic}$ to the scales of interest). For example, as mentioned throughout the text, our simple model for the shock size $r_{\rm shock} \sim \kappa l_{\rm sonic}$, probably underestimates the increase in T with $\mathcal{M}$, which could stem from the mass-fraction contained in shocks increasing somewhat from $\mathcal{M} \sim 1$ to $\mathcal{M} \gg 1$ (i.e. κ increasing with $\mathcal{M}$). In this vein, it is prudent to carry out further tests of the model using higher resolution numerical simulations (ideally with differing numerical methods), in particular, direct measurements of $\mathcal{P}_l(\ln\rho)$ (as in Fig. 5) across a range of $\mathcal{M}$.

Of course, real turbulence in the ISM involves a wide variety of other physical effects, which could strongly modify the ideal isothermal behaviour discussed through most of this work. For example, magnetic fields, self-gravity, dust and more complicated radiation physics could all play key roles in some cases, and the ISM can hardly be considered a homogeneous medium. Given the physical motivations behind various choices in the model, some of these features can be included in extended versions of the model (see Section 5), albeit heuristically. Such extensions could be interesting to study in future work and potentially important for making astrophysically relevant predictions. The model could also form the basis for a description of other physical effects that are strongly influenced by turbulence. Star formation (Hennebelle & Chabrier 2013; Hopkins 2013a; Padoan et al. 2014) is an obvious example for such applications, but there are also a variety other possibilities; for instance, the dynamics of dust grains, which are key in controlling the distribution of metals in the ISM (Draine 2003) and strongly affected by turbulence (see e.g. Hopkins 2016; Hopkins & Lee 2016; Pumir & Wilkinson 2016; Lee, Hopkins & Squire 2017; Monceau-Baroux & Keppens 2017).

ACKNOWLEDGEMENTS

It is a pleasure to thank A. Beresnyak, B. Burkhart and A. Schekochihin, as well as the anonymous Referee, for helpful comments and discussion. JS was funded in part by the Gordon and Betty Moore Foundation through Grant GBMF5076 to Lars Bildsten, Eliot Quataert and E. Sterl Phinney. Support for PFH was provided by NASA Astrophysics Theory Grant NNX14AH35G & National Science Foundation Collaborative Research Grant #1411920 and CAREER grant #1455342. Numerical calculations were run on Caltech cluster 'Zwicky' (NSF MRI award #PHY-0960291) and XSEDE allocation TG-AST130039 supported by the NSF.

REFERENCES

Aluie H., 2011, Phys. Rev. Lett., 106, 174502
Audit E., Hennebelle P., 2005, A&A, 433, 1

Banerjee S., Galtier S., 2013, Phys. Rev. E, 87, 013019
Banerjee R., Vázquez-Semadeni E., Hennebelle P., Klessen R. S., 2009, MNRAS, 398, 1082
Bauer A., Springel V., 2012, MNRAS, 423, 3102
Beresnyak A., Lazarian A., Cho J., 2005, ApJ, 624, L93
Boldyrev S., 2002, ApJ, 569, 841
Boldyrev S., Nordlund Å., Padoan P., 2002, Phys. Rev. Lett., 89, 031102
Brunt C. M., Federrath C., 2014, MNRAS, 442, 1451
Brunt C. M., Federrath C., Price D. J., 2010a, MNRAS, 403, 1507
Brunt C. M., Federrath C., Price D. J., 2010b, MNRAS, 405, L56
Burkhart B., Falceta-Gonçalves D., Kowal G., Lazarian A., 2009, ApJ, 693, 250
Burkhart B., Stanimirović S., Lazarian A., Kowal G., 2010, ApJ, 708, 1204
Burkhart B., Collins D. C., Lazarian A., 2015, ApJ, 808, 48
Castaing B., 1996, J. Phys. II, 6, 105
Cho J., Lazarian A., 2003, MNRAS, 345, 325
Colbrook M. J., Ma X., Hopkins P. F., Squire J., 2017, MNRAS, 467, 2421
Crutcher R. M., 1999, ApJ, 520, 706
Crutcher R. M., Wandelt B., Heiles C., Falgarone E., Troland T. H., 2010, ApJ, 725, 466
Davidson P. A., 2015, Turbulence, An Introduction for Scientists and Engineers. Oxford Univ. Press, Oxford
Dehnen W., Aly H., 2012, MNRAS, 425, 1068
Draine B. T., 2003, ARA&A, 41, 241
Dubrulle B., 1994, Phys. Rev. Lett., 73, 959
Federrath C., 2013, MNRAS, 436, 1245 (F13)
Federrath C., 2016, MNRAS, 457, 375
Federrath C., Banerjee S., 2015, MNRAS, 448, 3297
Federrath C., Klessen R. S., 2012, ApJ, 761, 156
Federrath C., Klessen R. S., 2013, ApJ, 763, 51
Federrath C., Klessen R. S., Schmidt W., 2008, ApJ, 688, L79
Federrath C., Klessen R. S., Schmidt W., 2009, ApJ, 692, 364
Federrath C., Roman-Duval J., Klessen R. S., Schmidt W., Mac Low M. M., 2010, A&A, 512, A81
Federrath C. et al., 2016, ApJ, 832, 143
Fleck R. C., Jr, 1996, ApJ, 458, 739
Galtier S., Banerjee S., 2011, Phys. Rev. Lett., 107, 134501
Gazol A., Kim J., 2013, ApJ, 765, 49
He G., Dubrulle B., Graner F., 1998, Phys. Lett. A, 245, 419
Hennebelle P., Chabrier G., 2013, ApJ, 770, 150
Hopkins P. F., 2013a, MNRAS, 430, 1653
Hopkins P. F., 2013b, MNRAS, 430, 1880 (H13)
Hopkins P. F., 2015a, MNRAS, 450, 53
Hopkins P. F., 2015b, MNRAS, 455, 89
Hopkins P. F., 2016, MNRAS, 456, 2383
Hopkins P. F., Lee H., 2016, MNRAS, 456, 4174
Hopkins P. F., Raives M. J., 2016, MNRAS, 455, 51
Kainulainen J., Federrath C., Henning T., 2014, Science, 344, 183
Kim J., Ryu D., 2005, ApJ, 630, L45
Kolmogorov A., 1941, Akade. Nauk SSSR Dokl., 30, 301
Konstandin L., Girichidis P., Federrath C., Klessen R. S., 2012, ApJ, 761, 149
Konstandin L., Schmidt W., Girichidis P., Peters T., Shetty R., Klessen R. S., 2016, MNRAS, 460, 4483
Kowal G., Lazarian A., Beresnyak A., 2007, ApJ, 658, 423
Kritsuk A. G., Norman M. L., Padoan P., Wagner R., 2007, ApJ, 665, 416
Krumholz M. R., McKee C. F., 2005, ApJ, 630, 250
Lee H., Hopkins P. F., Squire J., 2017, MNRAS, 469, 3532
Lemaster M. N., Stone J. M., 2008, ApJ, 682, L97
Lithwick Y., Goldreich P., 2001, ApJ, 562, 279
Mallet A., Schekochihin A. A., 2017, MNRAS, 466, 3918
Molina F. Z., Glover S. C. O., Federrath C., Klessen R. S., 2012, MNRAS, 423, 2680
Monceau-Baroux R., Keppens R., 2017, A&A, 600, A134
Monin A. S., Yaglom A. M., 1975, Statistical Fluid Mechanics: Mechanics of Turbulence, Vol. 2. The MIT Press, Cambridge, Massachusetts
Nolan C. A., Federrath C., Sutherland R. S., 2015, MNRAS, 451, 1380
Padoan P., Nordlund A., 2011, ApJ, 730, 40
Padoan P., Nordlund A., Jones B. J. T., 1997, MNRAS, 288, 145
Padoan P., Jimenez R., Nordlund Å., Boldyrev S., 2004, Phys. Rev. Lett., 92, 191102
Padoan P., Federrath C., Chabrier G., Evans N. J., II, Johnstone D., Jørgensen J. K., McKee C. F., Nordlund Å., 2014, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, Protostars and Planets VI. University of Arizona Press, Tucson, p. 77
Pan L., Padoan P., Haugbølle T., Nordlund Å., 2016, ApJ, 825, 30
Passot T., Vázquez-Semadeni E., 1998, Phys. Rev. E, 58, 4501
Price D. J., Federrath C., 2010, MNRAS, 406, 1659 (PF10)
Price D. J., Federrath C., Brunt C. M., 2011, ApJ, 727, L21
Pumir A., Wilkinson M., 2016, Ann. Rev. Condens. Matter Phys., 7, 141
Saichev A. I., Woyczynski W. A., 1996, SIAM J. Appl. Math., 64, 1008
Schmidt W., Federrath C., Klessen R., 2008, Phys. Rev. Lett., 101, 194505
Schmidt W., Federrath C., Hupp M., Kern S., Niemeyer J. C., 2009, A&A, 494, 127
She Z.-S., Leveque E., 1994, Phys. Rev. Lett., 72, 336
She Z. S., Waymire E. C., 1995, Phys. Rev. Lett., 74, 762
Stewart T., Strijbosch L., Moors J., Batenburg P., 2006, A simple approximation to the convolution of gamma distributions, Discussion paper. Center for Economic Research, Tilburg University
Vázquez-Semadeni E., García N., 2001, ApJ, 557, 727
Vázquez-Semadeni E., Ballesteros-Paredes J., Klessen R. S., 2003, ApJ, 585, L131

APPENDIX A: DERIVATION OF THE DENSITY STRUCTURE FUNCTIONS AND POWER SPECTRA FROM $\mathcal{P}_l(\ln\rho)$

In this appendix, we illustrate how to derive structure functions and power spectra from the scale variation of the PDF $\mathcal{P}_l(\ln\rho)$. This derivation is completely general and could apply to any statistical field for which one had $\mathcal{P}_l$.

Consider the variable s, denoting its average over scale l as s_l. We take s_l distributed according to the PDF $s_l \sim \mathcal{P}_l(s_l)$, with mean $\bar{s}_l = \int s_l \mathcal{P}_l(s_l)\,\mathrm{d}s_l$ and variance $S_l = \int s_l^2 \mathcal{P}_l(s_l)\,\mathrm{d}s_l - \bar{s}_l^2$. The first observation is that the isotropic autocorrelation of $s_{l'}$

$$R_{s_{l'}}(l) = R_{s_{l'}}(|\boldsymbol{l}|) = \langle s_{l'}(\boldsymbol{x}+\boldsymbol{l}) s_{l'}(\boldsymbol{x})\rangle, \tag{A1}$$

is the same as $R_{s_{l''}}(l)$, so long as $l'' < l$ and $l' < l$. We also note that if $l' > l$,

$$R_{s_{l'}}(l) = S_l + \bar{s}_l^2, \tag{A2}$$

which follows because $s_{l'}$ is constant on scales less than l'. This implies

$$R_s(l) = S_l + \bar{s}_l^2, \tag{A3}$$

relating direct measurements of s to $\mathcal{P}_l$.

The second-order structure function is now easily calculated as

$$\langle \Delta s^2\rangle = \langle [s(\boldsymbol{x}+\boldsymbol{l}) - s(\boldsymbol{x})]^2\rangle = 2\langle s^2\rangle - 2\left(S_l + \bar{s}_l^2\right). \tag{A4}$$

However, the non-locality of equation (A4) (it depends on the smallest scales through $\langle s^2\rangle$) is inconvenient, and it is more helpful to consider its derivative,

$$\frac{1}{2}\frac{\mathrm{d}\langle \Delta s^2\rangle}{\mathrm{d}l} = -\frac{\mathrm{d}}{\mathrm{d}l}\left(S_l + s_l^2\right). \tag{A5}$$

Equation (A5) is more useful than equation (A4) because it lacks dependence on the smallest or largest scales in the system.

The 3D power spectrum $\Phi_s(\boldsymbol{k})$ is related to the autocorrelation through the standard Fourier transform. However, since we are

interested in the 1D spectrum $4\pi k^2 \phi_s(k) = \Phi_s(\boldsymbol{k})$, the relation is instead (Davidson 2015)

$$R_s(l) = \int_0^\infty \phi_s(k) \mathrm{sinc}(kl)\, \mathrm{d}k, \tag{A6}$$

where $\mathrm{sinc}\,(x) = x^{-1}\sin(x)$. Approximating $\mathrm{sinc}\,(kl) \approx \Pi(kl/2\pi)$, where $\Pi(x)$ is the top-hat function (1 for $-1 < x < 1$, 0 otherwise), we see that the transform (A6) is related to a filtering operation

$$R_s(l) \approx \int_{2\pi/l}^\infty \phi_s(k)\, \mathrm{d}k, \tag{A7}$$

or

$$\frac{\mathrm{d}\langle \Delta s^2 \rangle}{\mathrm{d}l} \approx \frac{2\pi}{l^2} \phi_s\left(\frac{2\pi}{l}\right). \tag{A8}$$

Neglecting numerical constants (we are only interested in the k-scaling of the power spectrum), we obtain

$$\phi_s(k) \sim \frac{\mathrm{d}}{\mathrm{d}k}\left(S_l + s_l^2\right), \tag{A9}$$

which fits with the intuition that the power spectrum should encode the change in the variance of s with scale.

We note that the relation (A8) [and thus also equation (A9)] is not valid if $\phi_s(k) \sim k^{-\nu}$ with $\nu > 3$, because a second-order structure function cannot be steeper than $\langle \Delta s^2 \rangle \sim l^2$ (see Monin & Yaglom 1975, chapter 13). Similarly, it also cannot be valid for $\nu < 1$ in the limit $k \to \infty$, because $\langle \Delta s^2 \rangle(0)$ would become infinite (equivalently, the integral (A7) does not converge; Monin & Yaglom 1975). In the main text we consider only the large (i.e. supersonic) scales, and so will apply equation (A9) to calculate the spectrum, even when $\phi_s(k) \sim k^{-\nu}$ with $\nu < 1$.[14]

APPENDIX B: SHOCK SIZES

In this appendix, we explicitly test the assumptions about shock width $r_{\rm shock}$ that went into deriving T. In particular, the scaling $r_{\rm shock} \sim \kappa l_{\rm sonic}$ was important for relating the mathematical properties of the model to physical characteristics of the turbulent density field. To test this, we measure $r_{\rm shock}$ from simulation and compare this to $\kappa l_{\rm sonic}$. While not technically a test of the model, this is important to verifying that κ, as measured from the intermittency of the PDF through T, is broadly consistent with the true width of shock structures. In other words, having seen in Section 4 that the model gives decent predictions of turbulent statistics, is our physical interpretation $\kappa \sim r_{\rm shock}/l_{\rm sonic}$ consistent with the properties of shocks seen in simulations?

The Lagrangian nature of the GIZMO code implies that the density is directly proportional to the number density of cells. We thus use a simple counting method to measure $r_{\rm shock}$. This involves defining all cells with $\rho > \rho_{\rm max}/10$ (where $\rho_{\rm max}$ is the maximum density across the current snapshot) as being part of a 'shock', then counting the number of such cells within radius r, $N_<(r)$, of a randomly chosen centre cell. Ideally, if $r < r_s$, then $N_<(r) \sim r^3$, while if $r > r_s$ then $N_<(r) \sim r^\upsilon$, where $\upsilon \approx 2$ is the fractal dimension of the shock (Federrath et al. 2008). We carry out this procedure for 10 000 randomly chosen centre cells per simulation, averaging the results, then averaging these results over time in the statistical steady state of the turbulence. Of course, in the messiness of a true turbulent density field, the transition at $r \sim r_s$ will be relatively smooth, and it is difficult to unambiguously define $r_{\rm shock}$. In addition, at very small scales $N_<(r)$ is adversely affected by the finite number of cells, since $N_<(r) \to 1$ as $r < r_{\rm cell}$ (where $r_{\rm cell}$ is the size of a cell in the shocked region). Note that this method assumes an approximately constant distribution of cells inside the shock, which appears to be the case based on examination of 2D density field slices.

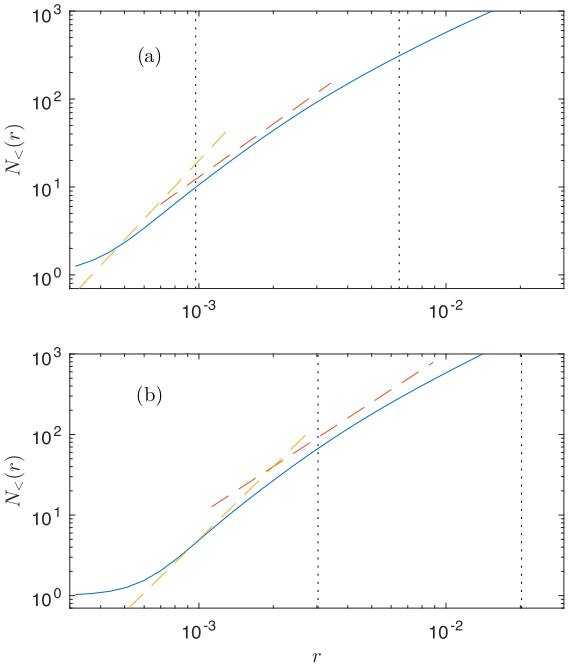

Figure B1. Average of $N_<(r)$ over snapshots and 10 000 randomly chosen centre cells for (a) $\mathcal{M} \approx 12$ and (b) $\mathcal{M} \approx 7$. The dashed red and yellow lines show r^2 and r^3 scalings, respectively. The vertical dashed lines show the inferred sonic scale $l_{\rm sonic} \sim \mathcal{M}^{-2}$ (right) and $\kappa l_{\rm sonic}$ (left) with $\kappa = 0.15$. Although of narrow extent due to the flattening of $N_<(r)$ to $N_<(r) = 1$ for $r < r_{\rm cell}$, there is a region at r where the scaling is substantially steeper than r^2. The transition to $N_<(r) \sim r^2$ scaling occurs significantly below $l_{\rm sonic}$ but in rough agreement with $\kappa l_{\rm sonic}$ for both $\mathcal{M} \approx 12$ and $\mathcal{M} \approx 7$.

Results are shown in Fig. B1 for the $\mathcal{M} \approx 12$ and $\mathcal{M} \approx 7$ simulations from above. In both cases there is a clear flattening to $N_<(r) \sim 2$ (the further flattening at higher r is probably related to the finite extent of the nearly 2D high density regions). The key point is that this flattening occurs well below $r \sim l_{\rm sonic} = \mathcal{M}^{-2}$, in approximate agreement with the estimate $r \sim \kappa l_{\rm sonic}$ with $\kappa \sim 0.2$. Thus, while it is difficult to accurately measure the shock width, our hypothesis that $r_{\rm shock} \sim \kappa l_{\rm sonic}$ is consistent with the data, but $r_{\rm shock} \sim l_{\rm sonic}$ provides a significant overestimate of the width. Future simulations at higher resolution may allow for a more accurate determination of these properties, by more accurately capturing the transition from supersonic to subsonic motions and allowing measurements over a wider range of $\mathcal{M}$. Unfortunately, even with the Lagrangian numerical method, the shock size at $\mathcal{M} \approx 40$ is still too small to see a transition from $N_<(r) \sim r^2$ to $N_<(r) \sim r^3$ (i.e. $r_{\rm cell} < r_{\rm shock}$), so we do not plot this here.

This paper has been typeset from a TEX/LATEX file prepared by the author.

[14] More concretely, we implicitly assume that the integral (A7) converges because $\phi_s(k)$ becomes sufficiently steep ($\nu > 1$) as $k \to \infty$.