Most massive haloes with Gumbel statistics

O. Davis, J. Devriendt, S. Colombi, J. Silk, C. Pichon

Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

Mon. Not. R. Astron. Soc. **413,** 2087–2092 (2011) doi:10.1111/j.1365-2966.2011.18286.x

Most massive haloes with Gumbel statistics

O. Davis,[1⋆] J. Devriendt,[1] S. Colombi,[2] J. Silk[1] and C. Pichon[1,2]

[1]*Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH*
[2]*Institut d'Astrophysique de Paris, CNRS UMR 7095 & UPMC, 98bis bd Arago, F-75014 Paris, France*

Accepted 2011 January 2. Received 2010 December 10; in original form 2010 October 17

ABSTRACT

We present an analytical calculation of the extreme value statistics for dark matter haloes – i.e., the probability distribution of the most massive halo within some region of the universe of specified shape and size. Our calculation makes use of the counts-in-cells formalism for the correlation functions, and the halo bias derived from the Sheth–Tormen mass function.

We demonstrate the power of the method on spherical regions, comparing the results to measurements in a large cosmological dark matter simulation and achieving good agreement. Particularly good fits are obtained for the most likely value of the maximum mass and for the high-mass tail of the distribution, relevant in constraining cosmologies by observations of most massive clusters.

Key words: methods: analytical – methods: statistical – dark matter – large-scale structure of Universe.

1 INTRODUCTION

Extreme value (or Gumbel) statistics are concerned with the extrema of samples drawn from random distributions. If a large number of samples are drawn from some distribution, the Central Limit Theorem states that their respective means will follow a distribution which tends, in the limit of large sample size, to a member of the family of normal distributions. Analogously, the maximum (or minimum) values u of the samples will have a distribution whose large sample size limit – where such a stable limit exists[1] – is a member of the family of Generalized Extreme Value (GEV) distributions as detailed by Gumbel (1958):

$$-\ln P_{\rm GEV}(y) = (1+\gamma y)^{-1/\gamma}, \quad y = (u-\alpha)/\beta. \tag{1}$$

The shape parameter γ is sensitive to the underlying distribution from which the maxima are drawn, while α and β are position and scale parameters.

Despite their wide use in other fields, extreme value statistics have historically seen very little application in astrophysics; some exceptions are the work of Bhavsar & Barrow (1985) on the brightest galaxies in clusters and the study of Coles (1988) on the hottest hot spots of the cosmic microwave background (CMB) temperature fluctuations.

The past year or so, however, has shown a resurgence of interest in the application of extreme value statistics to cosmology and questions of extreme structures, as revealed either in the clustering of galaxies (Antal et al. 2009), the prevalence of massive clusters (Holz & Perlmutter 2010; Cayón, Gordon & Silk 2010) or the temperature extrema of the CMB (Mikelsons, Silk & Zuntz 2009).

In this paper, we are interested in the dark matter haloes of massive galaxy clusters. The number density of extremely massive clusters is indeed a sensitive probe of the effects of the underlying cosmological model and laws of physics on large scales (Mantz et al. 2010a). These include for instance the equation of state of dark energy (Mantz et al. 2008), the possibility of modified gravity (Rapetti et al. 2010) the physical properties of neutrinos (Mantz, Allen & Rapetti 2010b) and primordial non-Gaussian density fluctuations (Cayón et al. 2010). Although the majority of the previously mentioned studies focus on the growth of massive clusters, simply knowing the mass of the most massive cluster in a survey can already provide strong constraints on cosmology (Holz & Perlmutter 2010).

In the present work, we outline an analytical derivation of the extremal halo mass distribution in standard cosmologies with Gaussian initial conditions. Rather than taking a phenomenological approach, we aim to predict the distribution of the most massive halo in a region for any specified combination of power spectrum, cosmological parameters and region size and shape. The paper is organized as follows. In Section 2 we outline the basics of our method for obtaining an analytical expression of the Gumbel distribution of most massive clusters masses and make an explicit link with equation (1). In Section 3, the theoretical predictions are checked against measurements in a very large N-body cosmological simulation. Finally,

⋆E-mail: olaf.davis@astro.ox.ac.uk

[1] Although it can be shown that where a stable limiting distribution exists it will take the form (1), certain pathological distributions give no such limit. For our purposes it is sufficient to note that the limit indeed exists for distributions which are of exponential type, meaning the cumulative distribution function F obeys $\lim_{x\to\infty} {\rm d}/{\rm d}x\{(1-F(x))/F'(x)\} = 0$, and that this class includes all physical distributions relevant to our applications.

Section 4 follows with a short summary of the main results and conclusions.

2 THEORY

Consider a large patch of the universe, which can be thought of as representing the space covered by a volume-limited sample of clusters, and denote by $m_{\rm max}$ the mass of the most massive dark matter halo in that patch. We wish to study analytically the Gumbel statistics, that is the probability distribution function $p_{\rm G}(m_{\rm max}){\rm d}m_{\rm max}$ of the values taken by $m_{\rm max}$ if we sample a large number of such patches. Obviously, this distribution will depend on the size and shape of the patch, as well as its redshift.

2.1 General expression of the Gumbel statistics

Let us define the cumulative Gumbel distribution by

$$P_{\rm G}(m) \equiv {\rm Prob.}(m_{\rm max} \leq m) \equiv \int_0^m p_{\rm G}(m_{\rm max}){\rm d}m_{\rm max}. \quad (2)$$

Such a probability is also the probability $P_0(m)$ that the patch is empty of haloes of mass above the threshold m (Colombi et al. 2011), hence

$$p_{\rm G}(m) = \frac{{\rm d}P_0}{{\rm d}m}. \quad (3)$$

Note that this assumes that there are no significant edge effects, i.e. the boundaries of the catalogue do not cross too many clusters. This effect is negligible if the patch is large compared to the halo size (and sufficiently compact).

If haloes are unclustered then the void probability follows simply from Poisson statistics,

$$P_0(m) = \exp(-nV), \quad (4)$$

where V is the volume of the patch and $n = n(>m)$ the mean density of haloes above mass m, with the appropriate spatial average with redshift made implicit $n(>m) = \langle n[>m, z(x)]\rangle_{x\in V}$.

We expect the Poisson limit to be reached for patches of size above a few hundred Mpc, where the matter distribution becomes homogeneous. Below this patch size, however, haloes are significantly clustered. In that case, the calculation of the void probability can be performed using a standard count-in-cell formalism if the connected N-point correlations functions, $\xi_N^{\rm h}(x_1, \ldots, x_N)$, of haloes above the threshold are known (e.g. Szapudi & Szalay 1993; Balian & Schaeffer 1989). The superscript h in the previous expression indicates halo correlation functions, while the naked ξ_N refer to correlations of the underlying matter density field. Since deviations from Poisson behaviour occur only for moderate patch sizes, the complex lightcone effects on the correlations induced by the evolution of clustering with redshift inside the patch (e.g. Matsubara, Suto & Szapudi 1997) can safely be neglected.

In particular, one can define the averaged correlations over a patch of volume V:

$$\bar{\xi}_N^{\rm h} \equiv \frac{1}{V^N}\int_V {\rm d}^3x_1 \ldots {\rm d}^3x_N \xi_N^{\rm h}(x_1, \ldots, x_N), \quad (5)$$

and the typical number of haloes above the threshold m in excess to the average in overdense patches as

$$N_{\rm c} \equiv nV\bar{\xi}_2^{\rm h}. \quad (6)$$

In the high-m limit, the void probability can be written

$$P_0(m) = \exp\left[-nV\sigma(N_{\rm c})\right], \quad (7)$$

where the function $\sigma(y)$ reads

$$\sigma(y) = \left(1 + \frac{1}{2}\theta\right){\rm e}^{-\theta}, \quad \theta{\rm e}^{\theta} = y, \quad (8)$$

(Bernardeau & Schaeffer 1999). Note that, as pointed out by these authors, this expression for $\sigma(y)$ follows from a specific hierarchical behaviour of higher order correlation functions of very massive haloes at large separations, $\bar{\xi}_N^{\rm h} \simeq N^{N-2}(\bar{\xi}_2^{\rm h})^{N-1}$.

We now proceed to specify the cumulative halo number density and the average two-point correlation function of haloes in order to fully determine the Gumbel statistics.

2.2 Halo number density

The number density $n(m, z)$ of haloes at a given mass m and redshift z, a.k.a. the halo mass function, we adopt is the one calculated by Sheth & Tormen (1999). It is based on a modification of the original model of Press & Schechter (1974), which links the statistics of the initial matter density field to the distribution of virialized dark matter haloes through a spherical top hat description of their gravitational collapse. As a result, this mass function can be expressed as a universal function of $\nu \equiv (\delta_{\rm c}/\sigma(m, z))^2$, where $\sigma(m, z)^2$ is the variance of the initial density field smoothed over spheres of radius $R(m)$ containing an average average mass m linearly extrapolated to z, and $\delta_{\rm c}$ is the critical overdensity threshold needed to turn an initial spherical top hat density perturbation into a collapsed halo at redshift z. More specifically, the number density of mass-m haloes is given by

$$m^2\frac{n(m, z)}{\bar{\rho}_{\rm m}}\frac{{\rm d}\log m}{{\rm d}\log \nu} = A\left(1 + (a\nu)^{-p}\right)\left(\frac{a\nu}{2\pi}\right)^{1/2}{\rm e}^{-a\nu/2} \quad (9)$$

with $\bar{\rho}_{\rm m} \equiv \Omega_{\rm m}\bar{\rho}$ the averaged matter density of the Universe. The shape of this mass function is parametrized by a and p, and A is simply a normalization factor.

2.3 Halo correlation functions

At sufficiently large separations, the two-point correlation of haloes of mass m can be related to that of the matter density through the bias function

$$\xi_2^{\rm h}(x_1, x_2, z) = b(m, z)^2\xi_2(x_1, x_2, z), \quad (10)$$

where ξ_2 is the linear dark matter density autocorrelation at redshift of interest. The function $b(m, z)$ can be estimated analytically using the Press–Schechter formalism (Mo & White 1996). Here, to remain consistent with equation (9), we use the expression for the bias of Sheth & Tormen (1999)

$$b(m, z) = 1 + \frac{a\nu - 1}{\delta_{\rm c}} + \frac{2p/\delta_{\rm c}}{1 + (a\nu)^p}. \quad (11)$$

This result is valid in the regime where the separation $x = |x_2 - x_1|$ is large enough compared to the mass scale $R(m)$. This has been tested successfully against N-body simulations by Mo, Jing & White (1996) (see, however, e.g. Tinker et al. 2010, for possible improvements on equation 11).

We obtain the bias of haloes exceeding mass threshold m by calculating the weighted average

$$b(>m, z) = \frac{\int_m^\infty b(m', z)n(m', z){\rm d}m'}{\int_m^\infty n(m', z){\rm d}m'}, \quad (12)$$

and hence the averaged two-point correlation function for haloes above the threshold,

$$\bar{\xi}_2^{\rm h}(>m, z) = b(>m, z)^2\bar{\xi}_2(z). \quad (13)$$

Recall that this equation should be valid in the regime where the patch size is large compared to $R(m)$,

$$L \gg R(m) = \left(\frac{3\pi m}{4\bar{\rho}_{\rm m}}\right)^{1/3}, \tag{14}$$

but small enough that light cone effects on the clustering inside it are negligible.

2.4 Generalized extreme value parametrization

The method outlined above allows us to compute the complete distribution function of of the most massive clusters. However, due to its complexity and the necessity of computing some of the integrals numerically, it does not provide us with a neat analytic parametrization of the distribution. Therefore, in order to achieve such a parametrization, we turn to the general theory of extremes, equation (1), using $u = \log_{10} m$ as the random variable. In order to calculate the parameters γ, α and β, we perform Taylor expansions of the analytic $P_{\rm GEV}$, and P_0 as computed by our method in the Poisson regime (equation 4). This Taylor expansion is performed about the peaks of the two distributions $\mathrm{d}P_0/\mathrm{d}u$ and $\mathrm{d}P_{\rm GEV}/\mathrm{d}u$. Equating the first two terms in these expansions give us expressions for the three parameters:

$$\begin{aligned}\gamma &= n(>m_0)V - 1, \quad \beta = \frac{(1+\gamma)^{(1+\gamma)}}{n(m_0)Vm_0\ln 10},\\ \alpha &= \log_{10} m_0 - \frac{\beta}{\gamma}\left[(1+\gamma)^{-\gamma} - 1\right],\end{aligned} \tag{15}$$

where m_0 is the mass at which the distribution $\mathrm{d}P_0/\mathrm{d}z = \ln(10)\, m\, p_{\rm G}(m)$ peaks – hence close to the most likely value of m – and is given implicitly by

$$\begin{aligned}&A\frac{\bar{\rho}_{\rm m}V}{m_0}\sqrt{\frac{a}{2\pi\nu_0}}\mathrm{e}^{-a\nu_0/2}\left(1+(a\nu_0)^{-p}\right) =\\ &\frac{a}{2} + \frac{1}{2\nu_0} + \frac{ap(a\nu_0)^{-(p+1)}}{1+(a\nu_0)^{-p}} - \frac{\nu_0''}{\nu_0'^2},\end{aligned} \tag{16}$$

where ν_0, ν_0' and ν_0'' are ν and its derivatives with respect to m evaluated at $m = m_0$.

These equations, then, allow us to neatly summarize the information contained in the extreme value distribution with the single parameter γ which describes its shape. This statistic has the potential to be used as a tool to compare models with data or with each other, as Mikelsons et al. (2009) proposed for the CMB.

3 NUMERICAL EXPERIMENT

To test our halo mass Gumbel distribution we compare the analytical result to measurements on the Horizon 4Π Simulation (Teyssier et al. 2009), a large cosmological dark matter simulation performed using the RAMESES N-body code (Teyssier 2002). The simulation followed the evolution of a cubic piece of the universe $2\,h^{-1}$ Gpc on a side containing 4096^3 particles, i.e. with a particle mass of $7.7\,h^{-1} \times 10^9\,\mathrm{M}_\odot$. Initial conditions were based on the *3-yr Wilkinson Microwave Anisotropy Probe (WMAP3)* results (Spergel et al. 2007), with the Hubble constant, density and characteristic parameters of the power spectrum given by $(h, \Omega_\Lambda, \Omega_{\rm m}, \Omega_{\rm b}, \sigma_8, n_{\rm s}) = (0.73, 0.76, 0.24, 0.042, 0.77, 0.958)$. Haloes in the simulation were identified at present time, $z = 0$, using a 'Friends-of-Friends' algorithm (e.g. Zeldovich, Einasto & Shandarin 1982; Davis et al. 1985) with a standard linking length parameter value given by 0.2 times the mean interparticle distance.

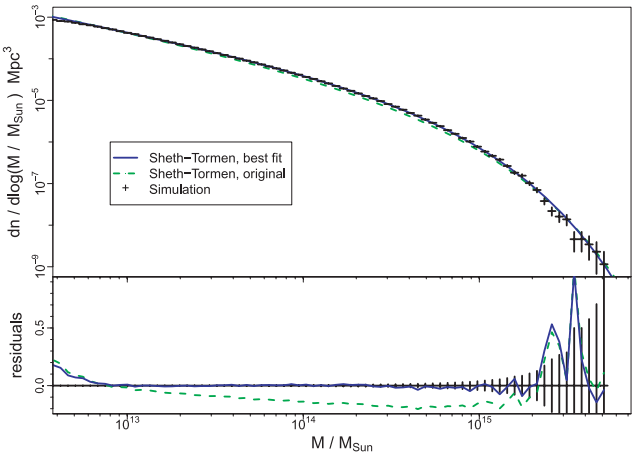

Figure 1. The upper panel shows the mass function of haloes in the simulation (points), compared to the Sheth–Tormen mass function with (p, a) equal to (0.3, 0.707) and (−0.19, 0.777) (solid blue and dashed green lines, respectively). The lower panel shows the residuals of the two theoretical curves compared to the data, i.e. (theory−data)/data.

3.1 Fit of the mass function

Any discrepancies between our derived Gumbel distribution and the true distribution can be thought of as arising from one of two sources: either inaccuracies in our chosen mass function, or inaccuracies due to the various assumptions made in proceeding from the mass function to $p_{\rm G}$. In order to quantify the respective contributions of each of these two sources, we repeat our calculations with two sets of parameters for the mass function: (i) once taking the parameters used in Sheth & Tormen (1999), $(p, a) = (0.3, 0.707)$, and, since the Sheth & Tormen form is its standard parametrization is known to perform only approximately (e.g. Warren et al. 2006; Jenkins, Frenk & White 2001), (ii) once with a best fit for a and p to the simulation's mass function, leading to $(p, a) = (-0.19, 0.777)$. For this latter, we also weight bins by their mass, since it is the high-mass end of the distribution which is of interest to us. Fig. 1 shows both these mass functions along with that measured in the simulation.

3.2 Results

Fig. 2 shows the distribution $p_{\rm G}(m_{\rm max})$ calculated as above, both for Poisson statistics (equation 4, red dots) and incorporating full clustering (equation 7, blue solid lines) using our best-fitting mass function. The full clustering calculation is also shown for the original Sheth–Tormen parameters (green dashes). Points show measurements from the Horizon 4Π simulation for comparison.

Fitting a Gumbel distribution equation (1) to the data presented in the four panels of this figure yields a single value of γ around $-0.21 \pm^{0.02}_{0.01}$ with reduced $\chi^2 \leq 1.1$, whereas the analytic prediction presented in Section 2.4 gives $-0.14 \leq \gamma \leq -0.1$. This lack of agreement has its root in the fact that γ is very sensitive to the higher order (skewness and kurtosis) moments of the data distribution around its peak, and these are poorly captured by assuming Poisson statistics, even on $100\,h^{-1}$ Mpc scales.

However, for a patch (in this case a sphere of radius L) of size $L = 100\,h^{-1}$ Mpc, Fig. 2(a) shows that $p_{\rm G}(m_{\rm max})$ measured in the data is not as badly described by Poisson theoretical results as it would have seemed from the value of γ alone: we are closing in on the so-called 'scale of homogeneity' above which the matter

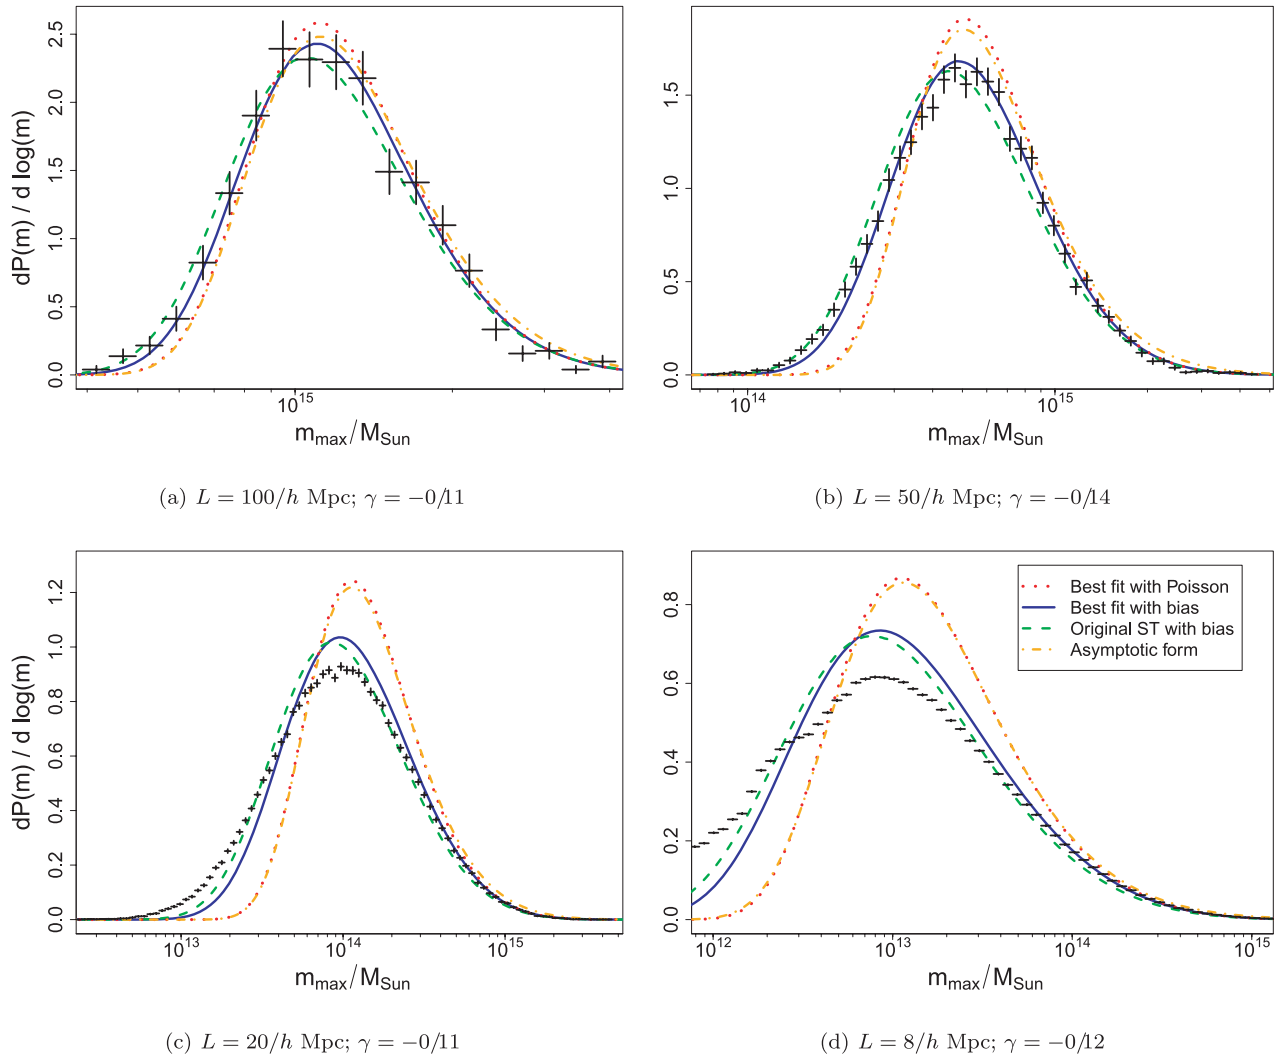

(a) $L = 100/h$ Mpc; $\gamma = -0.11$

(b) $L = 50/h$ Mpc; $\gamma = -0.14$

(c) $L = 20/h$ Mpc; $\gamma = -0.11$

(d) $L = 8/h$ Mpc; $\gamma = -0.12$

Figure 2. Distribution of largest cluster masses $m_{\max}$. The y-axis is probability density per log mass. Points are measurements from the simulation with Poisson error bars for each mass bin. The solid blue line is the theoretical result using the best-fitting mass function ($p = -0.19, a = 0.777$) and full halo clustering. Green dashes have instead the original Sheth–Tormen parameters ($p = 0.3, a = 0.707$). Not surprisingly, they do not agree as well with the measurement as the solid blue line; red dots have the best-fitting values but assume haloes are Poisson distributed. The orange dot–dashed line is the GEV distribution, with parameters calculated as explained in Section 2.4 and assuming Poisson statistics. All calculations use spherical patches, with radius $L = 100, 50, 20, 8\,h^{-1}$ Mpc, respectively.

distribution is essentially unclustered. Reducing the patch size to $50\,h^{-1}$ Mpc (Fig. 2b) causes the full clustering and Poisson curves to diverge. As expected, only the calculations incorporating clustering remain a good match to the simulation on these smaller scales.

Decreasing L further (Figs 2c and 2d) causes even the calculations including clustering to diverge from the data as our approximations – in particular the expression for the function $\sigma(y)$ in Section 2.1 and the condition (14) – fail outside the large-L limit. For instance, we find $R(10^{13}\,\mathrm{M}_\odot) = 3\,\mathrm{Mpc}\,h^{-1}$ and $R(10^{14}\,\mathrm{M}_\odot) = 6.4\,\mathrm{Mpc}\,h^{-1}$ which is a significant fraction of the respective patch sizes of 8 and 20 Mpc h^{-1}. Despite these limitations, the description of the data is still significantly better than that of the Poisson calculation and remains impressive at the high-mass end. This is excellent news as it is this high-mass tail of the distribution which is of interest for assessing the significance of rare events such as surprisingly massive clusters observed in X-ray or redshift surveys. Indeed the lower-mass tail of the distribution for which the prediction fails most significantly lies at masses below $\sim 10^{13}\,\mathrm{M}_\odot$, which corresponds to haloes containing one to a few galaxies rather than tens or hundreds, and therefore are of limited interest in the search for the most massive cluster.

Moreover, note that the position of the peak of the probability distribution function – i.e., the most likely value of $\log_{10} m_{\max}$ – is fairly accurately predicted by the theory even when the shape of the curve begins to diverge from that of the data. Fig. 3 shows this most likely value, $\log_{10} \hat{m}_{\max}$, as a function of L for both theoretical estimates and the simulation data (central line and points in the figure). We also show in this figure the 95 per cent confidence region on $\log_{10} m_{\max}$ (upper and lower lines and bars). This too is well fitted by the theory, particularly for the upper limit, and we emphasize that this is a crucial test of the theory's ability to give significance values for observations of specific overly massive clusters.

In addition to the four values of L shown in Figs 2 and 3 has a final simulation point at $L = 500\,h^{-1}$ Mpc. Here the 95 per cent confidence region is poorly fit, because the patch is too large compared

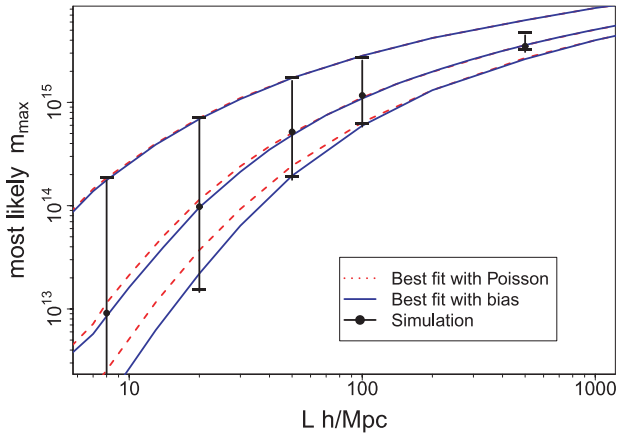

Figure 3. The most likely value of $\log_{10} m_{\rm max}$ (middle line and points) and the 95 per cent confidence limits (upper and lower line and bars). Dashed red and solid blue lines are analytic results for the Poisson limit and full clustering calculations, respectively, and points are simulation values for L corresponding to the four panels of Fig. 2 plus $L = 500\,h^{-1}$ Mpc.

to the simulation box to get good statistics. However, the measured value of $\log_{10} m_{\rm max}$ is still in good agreement with the theory.

3.2.1 *Senstivity to the mass function*

Worthy of note is the close similarity of the curves for the two sets of parameters in the Sheth–Tormen mass function (solid blue and dashed green lines in Fig. 3). Although the mass functions themselves differ by a relatively large amount compared to the best-fitting function's agreement with the simulation points (Fig. 1, lower panel), this translates to only a modest change in the distribution of $m_{\rm max}$. Similar calculations performed with the mass functions of Tinker et al. (2008) and Jenkins et al. (2001) lead us to conclude that the extreme value distribution is fairly robust to the choice of any reasonable analytic fit to the mass function.

3.2.2 *Redshift variation*

While the above calculations use patch sizes L small enough that the redshift evolution within the patch is negligible, it is also interesting to calculate $m_{\rm max}$ for a larger region with significant Δz. As noted in Section 2.1, redshift variation can be taken into account by a weighted spatial average provided the Poisson approximation holds, which is the case for such large patches. In particular, averaging the number density of haloes over the range $z = 0$ to ∞ gives a value for the expected largest mass cluster in the entire observable universe.

We performed this calculation assuming Poisson statistics, and found $m_{\rm max} = 4.6 \pm {}^{1.2}_{0.6} \times 10^{15}\,{\rm M}_\odot$ at the $1-\sigma$ confidence level. We note that this is in fair agreement with the similar calculation performed by Holz & Perlmutter (2010), who obtained $m_{\rm max} = 3.8 \pm {}^{0.6}_{0.5} \times 10^{15}$ using a *WMAP*7 cosmology and the mass function of Tinker et al. (2008).

4 CONCLUSION AND DISCUSSION

We have presented an analytic prediction of the the probability distribution of $m_{\rm max}$, the most massive dark matter halo/galaxy cluster in a specified region of the universe, making use of the counts-in-cells formalism. Our calculation, valid for Gaussian initial conditions, is numerically consistent with that proposed by Holz & Perlmutter (2010) when performed assuming such massive haloes are Poisson distributed spatially. However, the work presented in this paper improves on the calculation performed by these authors in two aspects. (i) Our results are given in a fully explicit analytic form and (ii) they include the contribution of clustering of haloes.

We also compared our analytic predictions to measurements from a large ($2000\,h^{-1}$ Mpc on a side) and well resolved (particle mass $7.7\,h^{-1} \times 10^9\,{\rm M}_\odot$) cosmological N-body simulation at zero redshift. We achieve remarkable agreement with the simulation in the area of parameter space in which our formalism is expected to be valid, namely patch radii above a few tens of h^{-1} Mpc. More surprisingly, even outside this range of scales the high-mass tail of the distribution is well fit by our 'fully clustered' theoretical estimate, as is the most likely value of $\log_{10} m_{\rm max}$.

This unexpected success over a wide range of scales warrants the application of the formalism to quantify the statistical significance of individual clusters observed in surveys. By applying our method to a patch of shape, size and redshift equivalent to a real survey we can obtain the Gumbel distribution and hence a likelihood for the observed value of $m_{\rm max}$. Moreover, we are quite confident that our method can be extended to non-standard cosmologies such as those including initial non-Gaussianities. It could therefore provide a measure of the evidence for such cosmologies from existing surveys of the most massive clusters as advocated in Cayón et al. (2010). We plan to tackle this exciting prospect in the very near future.

In addition to the full likelihood curve of $m_{\rm max}$ we are able, via the analytic GEV formalism, to produce a summary of the distribution in the form of the three parameters a, b and γ. The latter in particular has been proposed previously as a statistic for use in model comparison (e.g. Mikelsons et al. 2009). Although the work described in this paper uses a single cosmology and power spectrum and produces a roughly constant value of γ, the results of Colombi et al. (2011) suggest that γ should be quite sensitive to the shape of the power spectrum. If the effect on γ of the clustering can be fully quantified, we therefore expect to be able to use it as a statistic for direct comparison of models with observation. Likewise a, which is closely related to the peak m_0 of the distribution, may prove a useful statistic since we have demonstrated that it is well predicted by our theory.

ACKNOWLEDGMENTS

The authors are grateful to an anonymous reviewer whose comments led to improvements in this paper.

JD's research is supported by the Oxford Martin School, Adrian Beecroft and STFC.

CP acknowledges support from a Leverhulme visiting professorship at the Astrophysics department of the University of Oxford.

REFERENCES

Antal T., Sylos Labini F., Vasilyev N. L., Baryshev Y. V., 2009, Europhys. Lett., 88, 59001
Balian R., Schaeffer R., 1989, A&A, 220, 1
Bernardeau F., Schaeffer R., 1999, A&A, 349, 697
Bhavsar S. P., Barrow J. D., 1985, MNRAS, 213, 857
Cayón L., Gordon C., Silk J., 2010, preprint (arXiv:1006.1950)
Coles P., 1988, MNRAS, 231, 125
Colombi S., Davis O., Devriendt J. E. G., Silk J., 2011, MNRAS, in press (arXiv:1102.5707)
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Gumbel E., 1958, Statistics of Extremes. Colombia Univ. Press

Holz D. E., Perlmutter S., 2010, preprint (arXiv:1004.5349)
Jenkins A., Frenk C. S., White e. A., 2001, MNRAS, 321, 372
Mantz A., Allen S. W., Ebeling H., Rapetti D., 2008, MNRAS, 387, 1179
Mantz A., Allen S. W., Rapetti D., Ebeling H., 2010a, MNRAS, 406, 1759
Mantz A., Allen S. W., Rapetti D., 2010b, MNRAS, 406, 1805
Matsubara T., Suto Y., Szapudi I., 1997, ApJ, 491, L1
Mikelsons G., Silk J., Zuntz J., 2009, MNRAS, 400, 898
Mo H. J., White S. D. M., 1996, MNRAS, 282, 347
Mo H. J., Jing Y. P., White S. D. M., 1996, MNRAS, 282, 1096
Press W. H., Schechter P., 1974, ApJ, 187, 425
Rapetti D., Allen S. W., Mantz A., Ebeling H., 2010, MNRAS, 406, 1796
Sheth R. K., Tormen G., 1999, MNRAS, 308, 119
Spergel D. N. et al., 2007, ApJS, 170, 377
Szapudi I., Szalay A. S., 1993, ApJ, 408, 43
Teyssier R., 2002, A&A, 385, 337
Teyssier R. et al., 2009, A&A, 497, 335
Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709
Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlober S., 2010, ApJ, 724, 878
Warren M. S., Abazajian K., Holz D. E., Teodoro L., 2006, ApJ, 646, 881
Zeldovich I. B., Einasto J., Shandarin S. F., 1982, Nat, 300, 407

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