An excursion set model of hierarchical clustering: Ellipsoidal collapse and the moving barrier

Ravi K. Sheth[1] & Giuseppe Tormen[2]

[1] *NASA/Fermilab Astrophysics Center, MS209, Batavia IL 60510-0500, USA*
[2] *Dipartimento di Astronomia, 35122 Padova, Italy*
Email: sheth@fnal.gov, tormen@pd.astro.it

arXiv:astro-ph/0105113   7 May 2001

ABSTRACT

The excursion set approach allows one to estimate the abundance and spatial distribution of virialized dark matter haloes efficiently and accurately. The predictions of this approach depend on how the nonlinear processes of collapse and virialization are modelled. We present simple analytic approximations which allow us to compare the excursion set predictions associated with spherical and ellipsoidal collapse. In particular, we present formulae for the universal unconditional mass function of bound objects and the conditional mass function which describes the mass function of the progenitors of haloes in a given mass range today. We show that the ellipsoidal collapse based moving barrier model provides a better description of what we measure in the numerical simulations than the spherical collapse based constant barrier model, although the agreement between model and simulations is better at large lookback times. Our results for the conditional mass function can be used to compute accurate approximations to the local-density mass function which quantifies the tendency for massive haloes to populate denser regions than less massive haloes. This happens because low-density regions can be thought of as being collapsed haloes viewed at large lookback times, whereas high-density regions are collapsed haloes viewed at small lookback times. Although we have only applied our analytic formulae to two simple barrier shapes, we show that they are, in fact, accurate for a wide variety of moving barriers. We suggest how they can be used to study the case in which the initial dark matter distribution is not completely cold.

Key words: galaxies: clustering – cosmology: theory – dark matter.

1 INTRODUCTION

Bond et al. (1991) described an approach which allowed them to use the statistical properties of the initial density fluctuation field to derive an estimate of the number density of collapsed dark matter haloes at later times: the so-called universal, unconditional mass function (Press & Schechter 1974). Lacey & Cole (1993) showed how this model could be extended to estimate the rate at which small objects merge with each other to produce larger objects. This allowed them to estimate the conditional mass function of subhaloes within parent haloes. Lacey & Cole also provided formulae for the distribution of halo formation times; Nusser & Sheth (1999) provided formulae for the distribution of the halo mass at formation. Sheth (1996) and Sheth & Pitman (1997) showed how various higher order statistical properties of the forest of merger history trees associated with the formation of these objects could also be estimated within this approach. Mo & White (1996) and Sheth & Lemson

(1999a) showed how information about the forest of merger history trees could be used to quantify the extent to which dark haloes are biased tracers of the matter distribution. Mo, Jing & White (1997) provided predictions for the higher order moments of the spatial distribution of the haloes, and Sheth (1998) showed how to use the approach to estimate the probability that a randomly placed cell contains a certain amount of mass. Clearly, the approach has been very useful.

The approach combines the simple physics of the spherical collapse model with the assumption that the initial fluctuations were Gaussian and small. The problem of estimating any one of the quantities listed above is reduced to solving a problem associated with the crossing of an appropriately chosen barrier by particles undergoing Brownian motion; the Brownian nature of the motion comes from the assumption of Gaussian initial conditions, and the barrier shape is specified by the spherical collapse model (e.g. Sheth 1998). Hence, this is often called the excursion set approach.

Tormen (1998) reported that the spherical collapse based excursion set predictions did not describe the conditional mass function of subclumps in his simulations well. Motivated by this, Sheth, Mo & Tormen (2001) discussed a simple way of modifying the excursion set approach to incorporate the effects of ellipsoidal, rather than spherical collapse. On average, initially denser regions collapse before less dense ones. This means that, at any given epoch, there is a critical density which must be exceeded if collapse is to occur. In the spherical collapse model, this critical density does not depend on the mass of the collapsed object. However, in their parametrization of ellipsoidal collapse, Sheth, Mo & Tormen showed that, of the set of objects which collapse at the same time, the less massive ones must initially have been denser than the more massive ones, since the less massive ones would have had to hold themselves together against stronger tidal forces. They argued that this could be incorporated into the excursion set approach, simply and effectively, if not rigorously, by changing the barrier shape. In essence, whereas the barrier associated with spherical collapse is one whose height does not depend on distance from the origin of the walk, the one associated with ellipsoidal dynamics increases with distance. They showed that the excursion set approach with a moving barrier was able to provide a good fit to the universal halo mass function.

This paper is devoted to a more detailed discussion of moving barrier models. In general, moving barrier models have a richer structure than the constant barrier model. For example, the approach with spherical dynamics predicts that, at any given time, all the mass in the universe is bound up in collapsed objects, whereas a small fraction of the mass remains unbound in the case of ellipsoidal dynamics. In addition, whereas clustering is strictly hierarchical in the case of spherical dynamics, incorporating ellipsoidal collapse into the excursion set approach results in a model in which fragmentation as well as mergers may occur—the approach predicts that some small haloes fragment before they are subsumed into larger ones. Appendix A describes some of these features, which may (or may not!) provide better approximations to the physics of graviational instability than does the constant barrier model, in more detail. The Appendix also describes how the results of this paper can be used to model the halo mass function in warm dark matter scenarios such as that revisited by Bode, Ostriker & Turok (2001).

Section 2 shows how moving barrier models can be used to make simple analytic estimates of a number of statistical quantities which are routinely measured in numerical cosmological simulations. The primary results of Section 2 are equations (4) and (7), which are accurate for a large class of moving barrier shapes. To illustrate how these formulae work, we show the result of inserting the ellipsoidal collapse moving barrier of Sheth, Mo & Tormen (2001) into these formulae. Section 2.1 presents the number density of bound objects as a function of mass (the unconditional mass function), and Section 2.2 describes a simple efficient algorithm for generating it. Section 2.3 presents the average number of progenitor subhaloes as a function of subhalo mass, for a wide range of specified parent halo masses (the conditional mass function), and Section 2.4 shows how the halo mass function depends on the surrounding density field (the local-density mass function). Comparison with simulations shows that the approach, with the ellipsoidal collapse based mov-

ing barrier shape, is quite accurate. Section 2.5 shows that, at small lookback times, neither the constant nor the moving barrier predictions describe the conditional mass functions in the simulations particularly well, although the agreement at large lookback times is quite good if the spherical collapse constant barrier is used, and even better if the ellipsoidal collapse moving barrier is used.

Section 3 shows the result of considering more complicated moving barrier excursion set models. In particular, it shows the result of considering the full six-dimensional random walk associated with Gaussian random fields, rather than the one-dimensional simplification proposed by Sheth, Mo & Tormen (2001). It shows that their simplification is actually quite accurate. Details associated with the calculations in this section are presented in Appendix B.

Section 4 discusses some simple implications of our findings. Although, in this paper, we concentrate on the moving barrier derived by Sheth, Mo & Tormen (2001), we think it worth stressing that our analytic formulae are more general: they are accurate for a wide variety of moving barrier shapes.

2   THE MOVING BARRIER MODEL

As discussed in the introduction, we will mainly be interested in the first crossing distributions of uncorrelated Brownian motion random walks. Following Bond et al. (1991) and Lacey & Cole (1993), these first crossing distributions can be used to provide useful approximations to what have come to be called the conditional and unconditional mass functions of the dark halo distribution. The results of this section should be thought of as generalizations of the results in Lacey & Cole (1993). Whereas they restricted their attention to a barrier of fixed height, this section presents analytic formulae which approximate the barrier crossing distribution for a wide class of moving barriers.

Before we begin, we think a word on notation is helpful. We have chosen to present our formulae for the first crossing distributions using the same notation as Lacey and Cole (1993). This means that we use the symbol S to represent the variance in the density fluctuation field when smoothed on some scale, which is usually denoted σ^2. However, some of our formulae can be written in terms of the scaled variable $(\sigma/\sigma_*)^2 \equiv S/S_*$, for some suitably defined S_*. When this is possible, we also write our formulae in terms of $\nu \equiv S_*/S$.

2.1   The unconditional mass function

To illustrate how our formulae work, rather than use the spherical collapse barrier of constant height, we will use the moving barrier shape derived by Sheth, Mo & Tormen (2001)—the one associated with ellipsoidal collapse:

$$B(\sigma^2, z) = \sqrt{a}\,\delta_{\rm sc}(z)\left[1 + \beta\,(a\,\nu)^{-\alpha}\right], \qquad (1)$$

where $\nu \equiv [\delta_{\rm sc}(z)/\sigma(m)]^2$. In the expression above, $\delta_{\rm sc}(z)$ is the critical overdensity required for spherical collapse at z, extrapolated using linear theory to the present time (it is approximately $1.686(1 + z)$ if $\Omega = 1$), and $\sigma(m)$ is the rms of the initial density fluctuation field when it is smoothed on a scale which contains mass m, extrapolated using linear

theory to the present time. The parameters $\beta \approx 0.485$ and $\alpha \approx 0.615$ come from ellipsoidal dynamics (the spherical collapse model has $\alpha = 0$ and $\beta = 0$), and the value $a \approx 0.7$ comes from normalizing the model to simulations (as we discuss later, it may be more accurate to set $a \approx 0.75$ or so). The spherical collapse model sets $B(\sigma^2, z) = \delta_{sc}(z)$; because B is then independent of σ, spherical collapse is said to be associated with a barrier of constant height.

As Sheth, Mo & Tormen noted, $B(\sigma^2, z)$ scales similarly to the natural scaling associated with random walks: multiplying the barrier height by c and rescaling σ by the same factor results in a barrier of the same shape. They exploited this property as follows: first, they studied what happens when $z = 0$. Whereas an analytic expression for the first crossing distribution of a barrier of constant height has been known for some one hundred years or so, there is as yet no analogous solution for the ellipsoidal collapse moving barrier. (One of the main results of this section is to provide a good analytic approximation to this solution.) So, Sheth, Mo & Tormen simulated a large number of random walks and computed the distribution of first crossings, $f(S)$, of the barrier $B(\sigma^2 = S, z = 0)$ numerically. The scaling of the barrier shape meant that their numerical solution could be scaled to represent the first crossing distribution at any other time also. Therefore, when providing an analytic fit to their simulated first crossing distribution, they expressed it in scaled variables:

$$\nu f(\nu) = A \left[1 + (a\nu)^{-p} \right] \left(\frac{a\nu}{2} \right)^{1/2} \frac{e^{-a\nu/2}}{\sqrt{\pi}}, \qquad (2)$$

where f is the distribution of first crossings, ν and a are the variables defined above, $p = 0.3$ and A is determined by requiring that the integral of $f(\nu)$ over all ν give unity. The distribution associated with the spherical collapse constant barrier is got by setting $a = 1$, $p = 0$ and $A = 1/2$.

In the excursion set approach, the average comoving number density of halos of mass m, often called the universal or unconditional halo mass function $n(m, z)$, is related to the first crossing distribution (this is why the shape of the barrier influences the shape of the mass function) by

$$\nu f(\nu) \equiv m^2 \frac{n(m, z)}{\rho} \frac{d \ln m}{d \ln \nu}, \qquad (3)$$

where ρ denotes the average comoving density of the background. Because the first crossing distribution evolves self-similarly, so does the halo mass function. Fig. 2 in Sheth & Tormen (2001) shows that, in the GIF simulations of clustering in SCDM, OCDM and ΛCDM cosmologies, $n(m, z)$ is well fit by this expression (at least over the range $z = 0$ to $z = 4$). That is to say, the excursion set prediction that the universal unconditional mass function evolves self-similarly, is actually a very good approximation.

Although this fitting function (equation 2) for the first crossing distribution is extremely useful, when we consider conditional mass functions in the next section, it will turn out that we need to compute a new fitting function for each parent mass range of interest. This is, in principle, a computational bottleneck. For this reason, we think it more useful to provide a formula for the first crossing distribution which is more amenable for use in what is to follow.

We have analytic formulae for the first crossing distribution only in the case of constant (e.g. Bond et al. 1991)
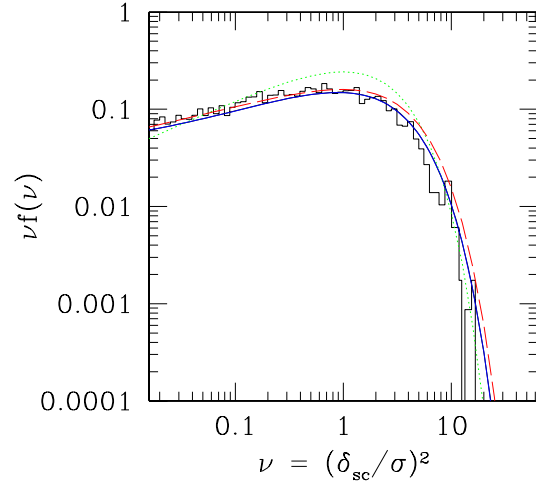
Figure 1. First crossing distributions and the universal unconditional halo mass function. Histogram shows the distribution obtained by simulating random walks which are absorbed on the ellipsoidal collapse moving barrier; solid curve shows our analytic approximation to this distribution (equation 4). Dashed curve shows the distribution which fits the halo mass function in cosmological simulations (equation 2), and dotted curve shows the distribution associated with spherical collapse.

and linear (Sheth 1998 and Appendix A of this paper) barriers. These solutions suggest the following approximation which we have found to work rather well. For a wide range of moving barrier shapes $B(S)$, the first crossing distribution is well approximated by

$$f(S)\, dS = |T(S)| \exp \left(-\frac{B(S)^2}{2S} \right) \frac{dS/S}{\sqrt{2\pi S}}, \qquad (4)$$

where $T(S)$ denotes the sum of the first few terms in the Taylor series expansion of $B(S)$:

$$T(S) = \sum_{n=0}^{5} \frac{(-S)^n}{n!} \frac{\partial^n B(S)}{\partial S^n}.$$

Notice that this expression gives the correct answer in the case of constant and linear barriers (in which only the first, or the first two terms of the series are non-zero). For the ellipsoidal barrier shape, we find we get reasonable accuracy to the numerical result if we truncate this expansion at $n = 5$. The accuracy of this formula increases as the distance between the start of the walk and the barrier height at that initial position, increases (this distance is $\sqrt{a}$ for the ellipsoidal collapse barrier of equation 1).

To illustrate, Fig. 1 shows the result of setting the barrier $B(S)$ to equal that given by equation (1) at $z = 0$, then generating a large ensemble of random walks, and so constructing the first crossing distribution. The histogram shows this Monte-Carlo distribution, and the solid curve shows the approximation presented in equation (4):

$$\nu f(\nu) = \sqrt{\frac{a\nu}{2\pi}} e^{-a\nu[1 + \beta(a\nu)^{-\alpha}]^2/2} \left(1 + \right.$$

$$\frac{\beta}{(a\nu)^\alpha}\left[1 - \alpha + \frac{\alpha(\alpha-1)}{2} + \cdots\right]\right)$$

$$\approx \left(1 + \frac{0.094}{(a\nu)^{0.6}}\right)\sqrt{\frac{a\nu}{2\pi}}e^{-a\nu[1+0.5(a\nu)^{-0.6}]^2/2}. \quad (5)$$

These should be compared with the distribution of first crossings of a barrier of constant height set equal to $\delta_{sc}(z = 0)$ (dotted curve), and equation (2) which fits the GIF simulations (dashed curve). With the exception of the dotted curve (the one associated with the constant barrier spherical collapse model) the other three curves are in reasonably good agreement. We will exploit this fact in the next subsection, when we derive a simple expression for the conditional mass function associated with ellipsoidal collapse.

Before we do so, however, we think it useful to point out that there are some generic properties of moving barrier models which arise from the scaling of the barrier shape. The fact that the barrier shape scales in the way it does means that moving barrier models in which the barrier height increases with S are somewhat more complicated than the constant barrier model, so they may be used to model a wider variety of physical processes. For example, our final expression for the moving barrier crossing distribution is not normalized to unity (notice that the solid curve is always below the dashed one). This is because, if the barrier height increases more steeply than linearly with σ then not all walks cross the barrier. Appendix A uses a simple analytic moving barrier model to illustrate some of these features.

2.2   Generating the unconditional mass function

Before we move on to study the conditional mass function, we thought it useful to describe a fast and simple algorithm for generating random numbers which can be used to provide the correct distribution of halo masses. To generate numbers drawn from the spherical collapse mass function (equation 2 with $A = 1/2$ and $p = 0$) is easy because, in this case, $f(\nu)$ can be got by generating a Chi-squared random variate. This means that, in the case of spherical collapse, the unconditional mass function can be generated quickly and easily because generating Gaussian random variates is easy. In particular, one simply generates a Gaussian random variate x, and then sets $a\nu = x^2$.

The ellipsoidal collapse mass function (equation 2 with $A = 0.32218$ and $p = 0.3$) cannot be transformed to a Gaussian, so constructing an algorithm for generating it is not so straightforward. We have found that first generating a Gaussian variate x, and then setting $a\nu = |x|^{3.6}/(1 + |x|^{1.6})$ is accurate to within a percent or so over the range $0.01 \leq \nu \leq 100$. The speed of this algorithm compensates for the fact that it does not exactly produce variates drawn from the ellipsoidal collapse mass function.

2.3   The conditional mass function

As stated above, we do not know of an analytic expression for the first crossing distribution associated with barriers which have the form given in equation (1). However, we do have two reasonably accurate fitting formulae—equations (2) and (4)—to this distribution. One might have thought that we could use them to make an estimate of the conditional mass function as follows.
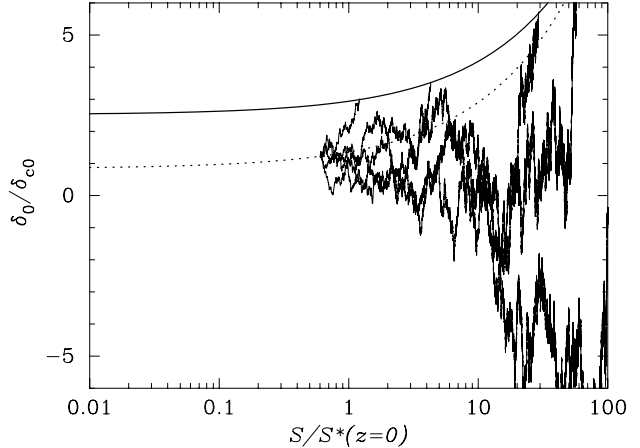
Figure 2. Examples of random walks used to construct the conditional mass functions associated with the ellipsoidal collapse moving barriers at $z = 0$ (dotted curve) and $z = 2$ (solid curve).

Bond et al. (1991) and Lacey & Cole (1993) argued that conditional mass functions could be estimated by considering the successive crossings of boundaries associated with different redshifts. The first crossing of two constant barriers of different heights has an analytic solution, so they were able to provide analytic estimates for the conditional mass function associated with the spherical collapse model. Such a formula is very useful, because, once the conditional mass function is known, the forest of merger history trees can be constructed using the algorithm described by Sheth & Lemson (1999b), from which the nonlinear stochastic biasing associated with this mass function can be derived using the logic of Mo & White (1996) and Sheth & Lemson (1999a).

In the constant barrier case, the conditional mass function is computed by considering walks which start from $[\sigma^2 = S, \delta_{sc}(z_0)]$ rather than from the origin, and then intersect the constant barrier $\delta_{sc}(z_1)$ at, say, s. This is easily computed, because, despite the shift in the origin, the second barrier is still one of constant height. Since the first crossing distribution of a barrier of constant height is known, (recall it was just such a distribution which was associated with the universal, unconditional mass function), the conditional mass function can also be written analytically. Essentially, $f(s|S)$ has the same form as the unconditional mass function $f(S)$, but with the change of variables: $\delta_{sc}(z) \to \delta_{sc}(z_1) - \delta_{sc}(z_0)$ and $S \to s - S$.

One might have wondered if the same change of variables in equation (2) provides a good description of the conditional mass function associated with the ellipsoidal collapse moving barrier. Unfortunately, because the barrier shape is not linear in S, changing the origin of the coordinate system does not yield a barrier of exactly the same functional form. Specifically, the shape of

$$B(s, z_1) - B(S, z_0) = \sqrt{a}\,\delta_1\left[1 + \beta\,s^\alpha/(a\,\delta_1^2)^\alpha\right]$$
$$- \sqrt{a}\,\delta_0\left[1 + \beta\,S^\alpha/(a\,\delta_0^2)^\alpha\right],$$

where $\delta_1 \equiv \delta(z_1)$ etc., can be written as a constant plus a term which scales as $(s - S)^\alpha$ only when α equals zero or one. This means that, formally, the solution to the two barrier problem associated with ellipsoidal dynamics is not

given by a simple rescaling of the unconditional ellipsoidal collapse mass function. Therefore, we cannot simply rescale the fitting function of equation (2) to get a reliable estimate of the conditional mass function: the two barrier problem associated with moving barriers must, in general, be solved numerically.

This is discouraging because it means that, in principle, we must find a different fitting function for each choice of condition, because each condition corresponds to a different origin, say, (B_0, S_0), and so to a slightly different barrier shape. Of course, the result can be generated relatively quickly in at least two ways. The first is to solve the integral equation associated with this barrier numerically (Monaco 1997b; Sheth 1998). The second is to simply simulate the random walk trajectories and so construct the first crossing distribution directly.

Fig. 2 shows an example of what is involved in solving the two barrier problem numerically in this way. The smooth dotted and solid curves show the ellipsoidal collapse barrier $B(S, z)$ of equation (1) scaled to $z = 0$ and $z = 2$, respectively. Jagged curves show a few representative random walk trajectories: they start at the barrier position $B(S, z = 0)$, where $S(M)$ is given by the GIF SCDM power spectrum, and $M/M_* = 2$. These random walks are followed until they first cross the barrier $B(s, z = 2)$. The value of S at which this happens is stored and used to make plots like those shown below. The random walks were generated by the same Monte–Carlo code that was used to generate Fig. A1, except that there the barrier shape was given by equation (A1). Fig. A1 shows that this Monte–Carlo code works correctly. One possible approach to the excursion set conditional mass function, then, is to simply generate it numerically, as the need arises.

Before providing a detailed comparison of the conditional mass functions generated using this Monte–Carlo model and those in numerical simulations, it is useful to study a simple limiting case. Fig. 2 shows that for $S/S_* < 0.5$, the height of the barrier is approximately constant. At small S, the only difference between the barrier at two redshifts, and the spherical collapse constant barriers, arises from the factor of $a = 0.707$. This has the following consequence. At small lookback times (small redshift differences), most random walk trajectories will intersect the barrier before they have travelled very far along the S axis. For these trajectories, the barrier is, effectively one of constant height. This means that the conditional mass function for massive haloes at small lookback times will have the same shape as that predicted by the constant barrier, with one small difference. The factor of $a = 0.707$ has the effect of slightly reducing (by a factor of $\sqrt{a}$) the effective redshift difference relative to the original constant barrier model. As a result, the GIF barrier suggests that massive haloes at small lookback times will be slightly more massive than the original constant barrier model predicts. Since $\sqrt{0.707}$ is close to unity, this effect will be small. In practice, we only expect the barrier predictions to differ significantly from those of the constant barrier for small haloes, or at large lookback times (high redshift). This is encouraging, because these are precisely the regimes in which simulations suggest that the constant barrier model is inaccurate (Tormen 1998).

In addition to Monte-Carloing the conditional mass functions, we can use the results of the previous subsec-

tion to derive a simple analytic expression for their shape. We can do this because our barrier crossing formula is reasonably accurate for a rather wide range of barrier shapes. A glance at Fig. 2 shows that the barrier shapes associated with the conditional mass functions are not likely to be too different from those associated with the unconditional function, so we should be able to use equation (4) to approximate most of the conditional mass functions we will be interested in. In practice, this can be done by simply making the appropriate replacements $B \to B(s) - B(S)$ and $S \to s - S$ in equation (4). At the risk of being repetitive, our approximation for the conditional mass function is:

$$\mathcal{N}(m, \delta_1 | M, \delta_0) \, dm \equiv \frac{M}{m} f(m, \delta_1 | M, \delta_0) \, dm \qquad (6)$$

where $f(m|M) \, dm = f(s|S) \, ds$ with

$$f(s|S) \, ds = \frac{|T(s|S)|}{\sqrt{2\pi(s - S)}} \exp\left(-\frac{[B(s) - B(S)]^2}{2(s - S)}\right) \frac{ds}{s - S}, \quad (7)$$

and

$$T(s|S) = \sum_{n=0}^{5} \frac{(S - s)^n}{n!} \frac{\partial^n [B(s) - B(S)]}{\partial s^n}.$$

Fig. 3 compares this approximation with the actual numerical Monte-Carlo distribution, and compares both with the actual distribution measured in the cosmological simulations. (We will show a comparison with what one gets by rescaling equation 2 shortly.) The histograms (without error bars) show the conditional mass functions generated using our Monte–Carlo code, for parent haloes in the mass range $1 \leq M/M_* \leq 2$ (upper curves) and $16 \leq M/M_* \leq 32$ (lower curves); the upper curves have been shifted upwards by a factor of ten for clarity. The smooth curves show the associated analytic approximation; they describe the results of our numerical walks reasonably well. Therefore, in the comparisons to follow, we will sometimes only show the analytic curves. Recall that the analytic formula should be most accurate for the high-redshift progenitors of massive parents, and least accurate when the parents are not very massive. The figure shows some evidence that this is true. Both these curves should be compared with the symbols which have error bars—they show the conditional mass functions measured in the ΛCDM simulations. The figure is quite encouraging: our excursion set predictions are in reasonably good agreement with the cosmological simulations. In addition, the figure shows that the analytic approximation (equation 7) is reasonably accurate.

Fig. 4 shows a similar comparison in the case of SCDM. In this case, we have chosen to not show the random walk histogram, since it is quite well fit by our formula. Instead, the figure shows the conditional mass functions measured in the SCDM simulations (symbols with error bars) and the various analytic approximations discussed earlier (smooth curves). In order of accuracy, these are equation (7) (solid), the distribution associated with rescaling the unconditional mass function of equation (2) (dashed), and the conditional distribution associated with the constant barrier, spherical collapse model of Lacey & Cole (1993) (dotted). (The finite mass resolution of the simulations means that the distributions are artificially truncated at low masses.)

Notice that the spherical collapse based dotted curves are often quite different from the N-body simulation sym-
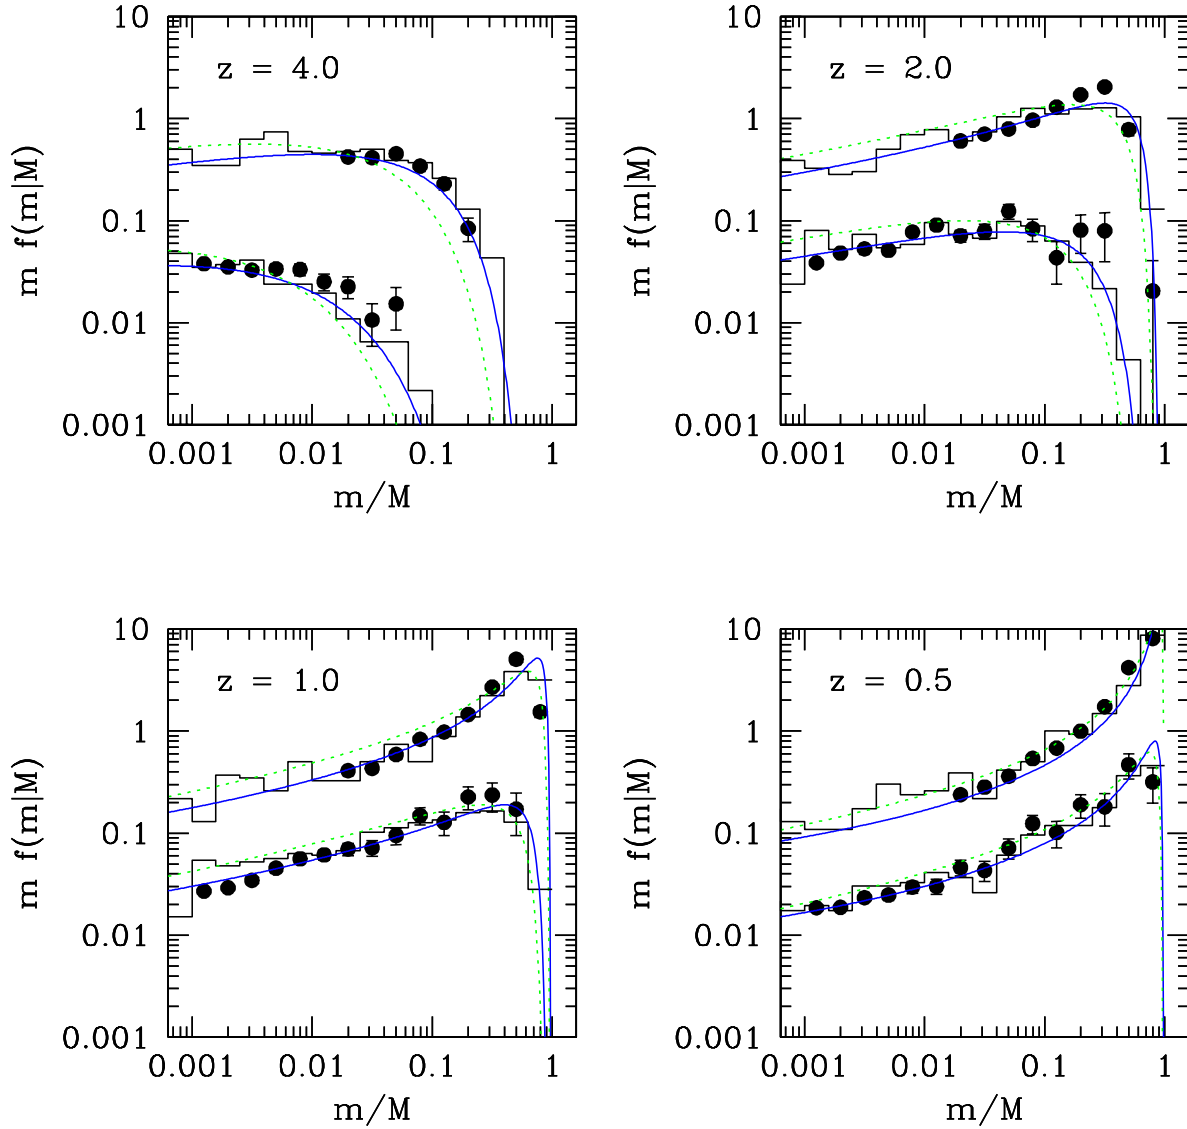
Figure 3. Excursion set conditional mass functions at z, for parent haloes identified at the present time. The two sets of curves in each panel are for parent haloes in the mass range $1 \leq M/M_* \leq 2$ (upper curves) and $16 \leq M/M_* \leq 32$ (lower curves); the upper curves have been shifted upwards by a factor of ten for clarity. Symbols with error bars show the distributions measured in the ΛCDM simulations, histograms show the result of generating the first crossing distribution by simulating an ensemble of 10^4 random walks, smooth solid curves show the analytic approximation discussed in the text, and dotted curves show the distribution associated with barriers of constant height.

bols. This discrepancy is similar to that first noticed by Tormen (1998); haloes in the simulations seem to hold themselves together at higher redshift than the spherical collapse model predicts. Notice also that, whereas the dashed curves we get by rescaling the unconditional halo mass function are certainly better fits to the cosmological simulations, the solid curves, in which the relation between the excursion set

model and the conditional mass function are accounted for more carefully, are almost always even more accurate.

2.4 Dependence on local density

Following Mo & White (1996), knowledge of the conditional mass function allows one to estimate how the distribution of dark matter haloes today depends on the average density in
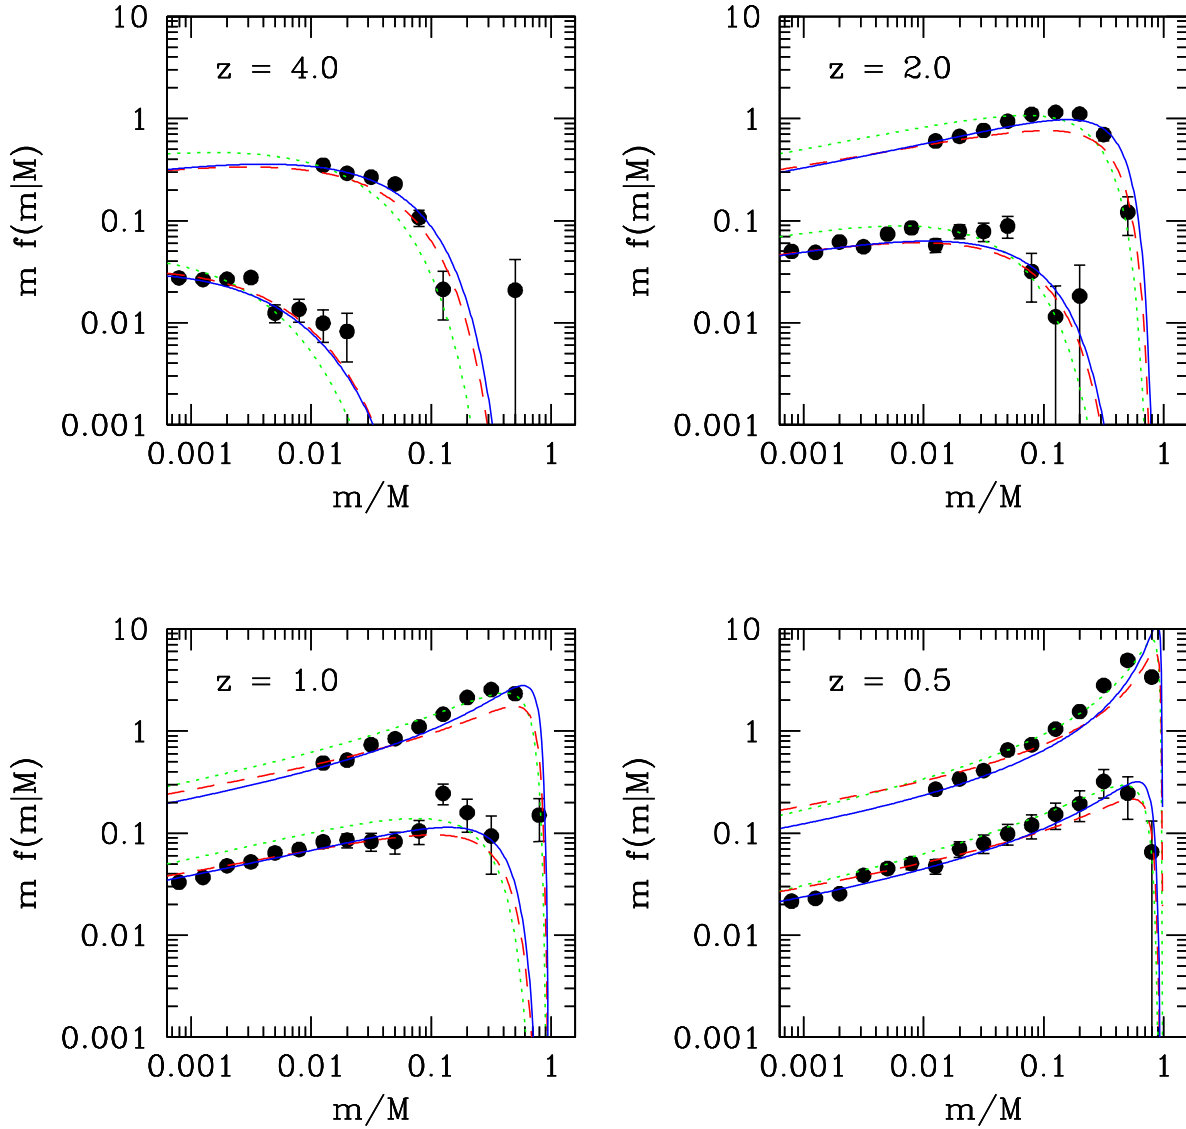
Figure 4. Conditional mass functions in the SCDM simulations (symbols with error bars). The two sets of curves in each panel show parent haloes with masses in the range $1 \leq M/M_* \leq 2$ (upper curves) and $16 \leq M/M_* \leq 32$ (lower curves); the upper curves have been shifted upwards by a factor of ten for clarity. Dotted curves show the spherical collapse prediction (Lacey & Cole 1993), dashed curves show the distribution one gets by rescaling the unconditional mass function (equation 2), and solid curves show our analytic approximation to the random walk with moving-barrier simulations (equation 7).

which the haloes are. In essence, they argued that a dense region should be thought of as an object which will collapse and form a virialized halo in the near future. This means that the haloes in it today can be thought of as 'progenitor subhaloes' viewed at 'low redshift'. In contrast, it will be a much longer time before an underdense region collapses (if it collapses at all), so the haloes within it today are like the progenitor haloes viewed at high redshift. In hierarchical models, massive haloes form later, and less massive haloes

form earlier. The discussion above means that one expects haloes in dense regions to be more massive, on average, than in underdense regions. The precise dependence of halo mass on local density depends on the precise relation between the local density today and the effective 'redshift'. Mo & White used the spherical collapse model to provide this relation. They provided a simple fitting formula to this relation in an Einstein-deSitter universe. Lemson & Kauffmann (1999) and Sheth & Lemson (1999a) showed that this provided a
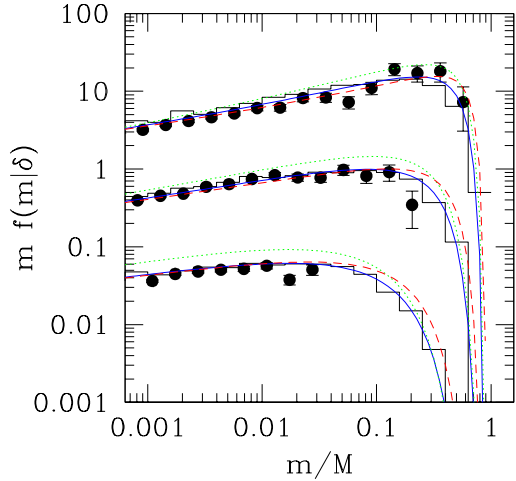
Figure 5. Mass functions as a function of local density in the SCDM simulations (symbols with error bars), plotted as a function of relative mass, so as to resemble the conditional mass functions presented earlier. Dotted curves show the spherical collapse prediction, dashed curves show the distribution one gets by rescaling the unconditional mass function (equation 2), and solid curves show our analytic approximation to the random walk with moving-barrier simulations (equation 7). The curves have been offset upwards by a factor of ten and a hundred, in the case of the middle and topmost curves, respectively. The upper most curves show the densest cells.

reasonably good description of how, in their simulations, the density of haloes depended on local density. We have checked that the following simple modification to their formula is reasonably accurate for all cosmologies of interest:

$$\delta_0(\delta, z_0) = \frac{\delta_{sc}(z_0)}{1.68647} \times \left[1.68647 - \frac{1.35}{(1+\delta)^{2/3}} \right.$$
$$\left. - \frac{1.12431}{(1+\delta)^{1/2}} + \frac{0.78785}{(1+\delta)^{0.58661}} \right], \quad (8)$$

where $M/\bar{\rho}V \equiv (1 + \delta)$ is the nonlinear density of a region containing mass M within the volume V at z_0, and $\delta_{sc}(z_0)$ denotes the critical density for spherical collapse at z_0. The number density of m-haloes in regions of nonlinear density δ is got from the conditional mass function $\mathcal{N}[m, \delta_{sc}(z_1)|M, \delta_0(\delta, z_0)]$ of equation (6).

The previous section showed that, in fact, the conditional mass functions are better fit by ellipsoidal collapse based curves. Therefore, one might reasonably expect that the same will be true for $n(m|\delta)$. To emphasize how similar $n(m|\delta)$ is to the conditional mass function, we have chosen to do the following. We have divided the simulation volume up into into cubes, each 10Mpc/h on a side. We then divided the cubes into three classes: the densest, and least dense ten percent of the cells, and the ten percent around the median density. Figs. 5 and 6 show $m^2 n(m|\delta)$ for the haloes in the cells, plotted as a fraction of the mass of a cell. Symbols
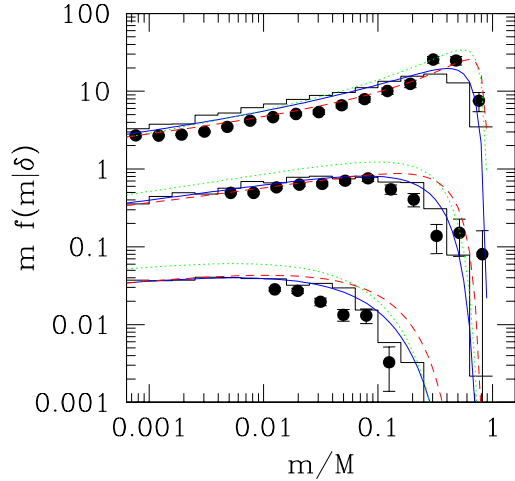

Figure 6. Same as the previous figure, but for ΛCDM.

with error bars show the measurements in the cosmological simulations, histograms show the associated random walk distributions, solid curves show our analytic approximation to the random walks, dashed curves show the result one gets by rescaling the unconditional mass function, and dotted curves are for the spherical collapse, constant barrier model. The curves for the three types of cells have been offset from each other for clarity; the lowest density cells are the lowest curves, and they have not been shifted, average density cells have been shifted upwards by a factor of ten, and the densest cells have been shifted upwards by a factor of a hundred. In all cases, the random walk model provides the best fit to the simulation data.

Notice how similar these curves appear to the conditional mass functions presented earlier. Because most of haloes in the densest cells are a significant fraction of the total mass in the cell, the mass function in dense cells looks very like the low redshift conditional mass functions. In contrast, the mass function in underdense cells looks much more like the high redshift conditional mass functions. This is precisely what the model predicts.

Figs. 7 and 8 show what this trend means for the actual number density of haloes in dense and less dense regions. The various symbols and curves are the same as in the previous two figures; the only difference is that now the x-axes have been multiplied by the total mass in the cell to show physical, rather than relative masses. Clearly, less dense cells have essentially no massive haloes. In addition, the ratio of massive to less massive haloes is higher in denser cells. In particular, note that the density of less massive haloes in dense regions is actually smaller than the density of less massive haloes in underdense regions. One might have thought that dense regions simply have more haloes on average. For example, one might have thought that $n(m|\delta) = (1 + \delta) n(m)$. Our figures show that this is wrong, but that a good estimate of $n(m|\delta)$ can, nevertheless, be computed analytically. We conclude this section with the
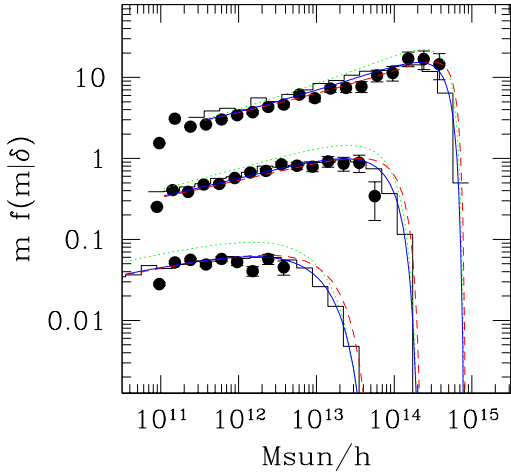
Figure 7. Same as Fig. 5, but now presented as a function of physical, rather than relative, mass. Massive haloes occupy the densest cells (upper most curves).

observation that our moving barrier based formulae provide a more accurate fit to the simulations than does the spherical collapse based constant barrier model.

2.5 Rescaling the conditional mass function

In the excursion set model with a constant barrier height, the unconditional mass function, when expressed as a function of $\nu \equiv \delta_c^2/\sigma^2(m)$, is expected to be a universal function which is independent of redshift, cosmology or initial power spectrum. In addition, if the conditional mass function is expressed in units of $(\delta_{c1} - \delta_{c0})^2/(s - S)$, then it is expected to have the same shape as the unconditional mass function.

The previous section we noted that, although the unconditional mass function is a universal function of ν, this function is not the one predicted by the constant barrier model. We argued that if we interpret the unconditional mass function as coming from a moving barrier, then we no longer expect the conditional mass function to be a universal function of ν. Figs. 9 and 10 show this explicitly: they show the result of applying this rescaling to the conditional mass functions in the SCDM and the ΛCDM simulations we presented earlier.

The four panels in each figure show the conditional mass function at each of the four redshifts we have been studying so far. The symbols show the result of rescaling the conditional mass functions in the simulations for parent haloes with mass in the range 1-2 (filled circles) and 8-32 M_* (open triangles) at $z = 0$. Notice that the symbols in each panel do not overlap exactly—at fixed z, the conditional mass functions for different parent haloes do not rescale exactly. In addition, the band traced out by the symbols at $z = 4$ is quite different from the band traced out at $z = 0.5$; the mass functions at different output times do not rescale either. Both
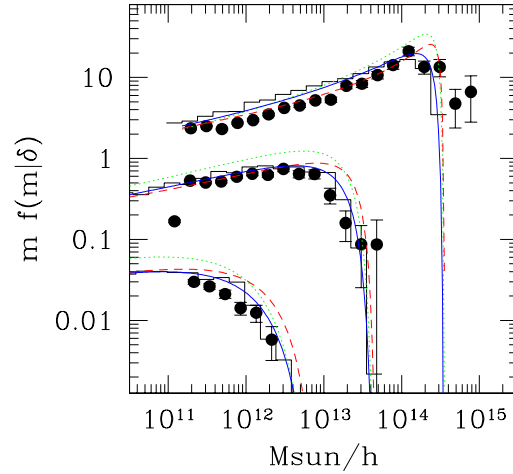

Figure 8. Same as the previous figure, but for ΛCDM.

these findings illustrate our main point: the conditional mass function is not a universal function of the scaling variable ν.

The dotted curves which are the same in all the panels show the predictions of the constant barrier model; they do not provide a good fit at any time for any mass range. The dashed curves which are also the same in all the panels show the result of assuming that, upon rescaling, the conditional mass function will have the same shape as the unconditional mass function; although they provide a good fit at large z, they are increasingly in error at small z. This is true both for the SCDM and the ΛCDM models.

The solid lines in the various panels show the predictions of our moving barrier model. In this case, the predictions depend both on the parent mass range, and on the redshift at which the progenitors are identified: we have chosen to show the predictions for the 1-2M_* haloes only. Whereas the model is in reasonable agreement with the simulations at large z, it has the wrong shape an small z. This is not terribly surprising, because our formula for the first crossing distribution of the moving barrier model was only supposed to work at large z, but it is unfortunate that the disagreement at low z is so bad! One might have thought that the actual first crossing distribution may be in good agreement with the GIF simulations, and that it is only the analytic approximation which is in error at small z. Unfortunately, this is not so. The histograms in Fig. 10 show the result of simulating an ensemble of random walks to construct the conditional mass functions; although they are in slightly better agreement with the simulations, they are still quite different.

The discrepancy between the simulations and our moving barrier model predictions are most pronounced when the subclump mass is $m/M \geq 1/2$. This suggests that our model is unable to describe the histories of clumps at small lookback times. At small lookback times, one might worry that the spherical overdensity groupfinder we use to identify the subclumps in the simulations might find a different set of objects than a friends-of-friends algorithm. Plots of
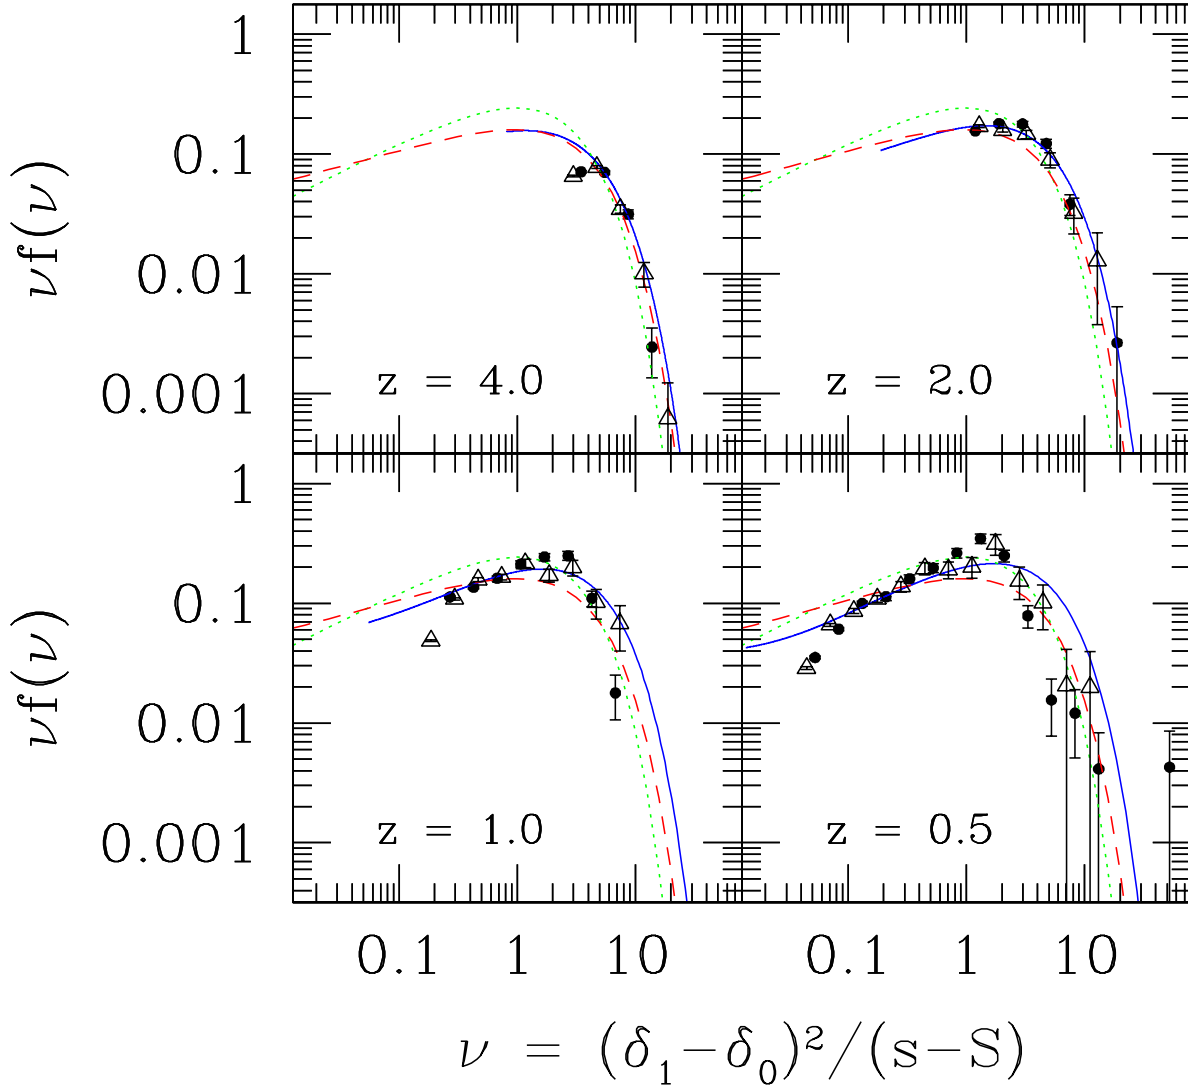
Figure 9. Rescaled conditional mass functions in the SCDM model. Panels show the different redshift bins we studied earlier. Symbols in each panel show the different mass ranges we considered. Dotted line shows the constant barrier prediction (in these variables it is the same as the unconditional mass function), dashed line shows the result of rescaling the actual unconditional mass function, and solid curves show the result of rescaling the moving barrier predictions.

the rescaled conditional mass function constructed using a friends-of-friends algorithm look very similar to the spherical overdensity results presented above—the discrepancy between model and simulations is independent of the choice of groupfinder.

In addition, because the theory assumes that mass is conserved—all the mass of a subclump becomes part of the final halo—whereas this is not true in the simulations: some of the particles which make up the final object may have come from a subclump which merged along with most of

its particles into a different object. This means that there is some freedom associated with how we decide if a clump at an early time should be counted as a subclump of a halo at the final time. We have tried two schemes for idenfifying subclumps: a progenitor is a clump which donates at least half its mass to the final object, or which donates at least one particle to the final object. Once we have made this decision, we must also decide what we wish to count as the mass of the parent object: two natural choices are the mass at the final time (which, by definition is fixed for all earlier redshifts), or
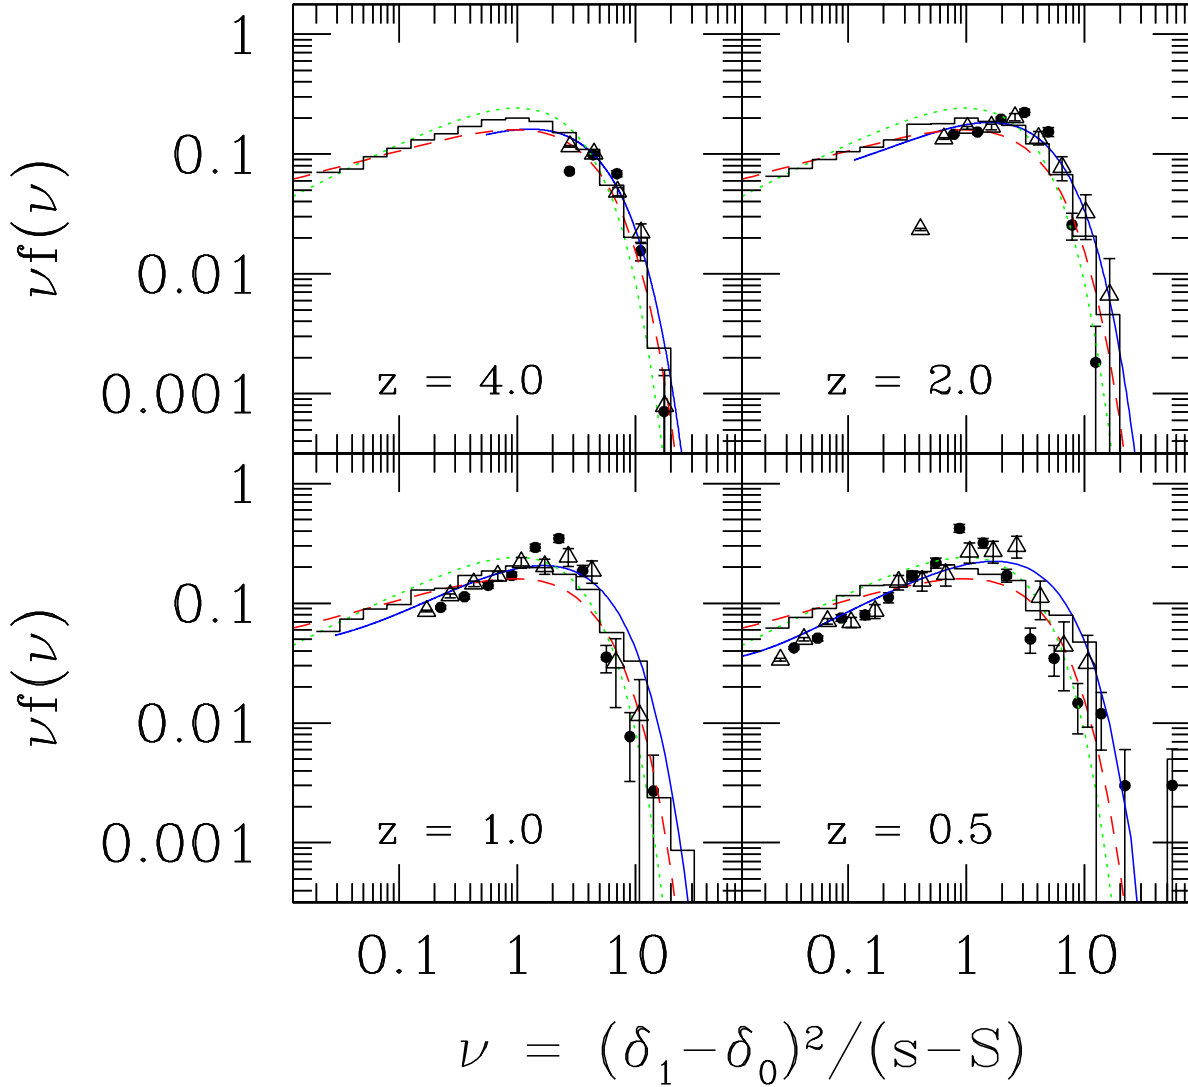
**Figure 10.** Same as the previous figure, but for ΛCDM.

the mass which is got by summing up the masses of all the progenitor subclumps (which may depend on redshift). The figures above are for the case in which a subclump is any clump which donates at least one particle to the final object, and the parent mass is defined as the mass at the final time (so it is independent of redshift). Although the actual conditional mass functions depend slightly on which combination of the above choices we make, the generic results shown above are independent of this choice.

Before moving on, we think it worth noting that the discrepancies between the SCDM simulations and the dotted or dashed curves are qualitatively similar to the discrepancies in the ΛCDM case. This suggests that one should be able to find a model which can account for these discrepancies in a way which is independent of power-spectrum, redshift or cosmology. Our moving barrier model is just not up to the task. The next section studies why.

3   AN EXTENSION

In the ellipsoidal collapse model envisaged by Bond & Myers (1996) and implemented by Sheth, Mo & Tormen (2001), the collapse of a patch is determined by the surrounding shear field. In a Gaussian random field, the field around a
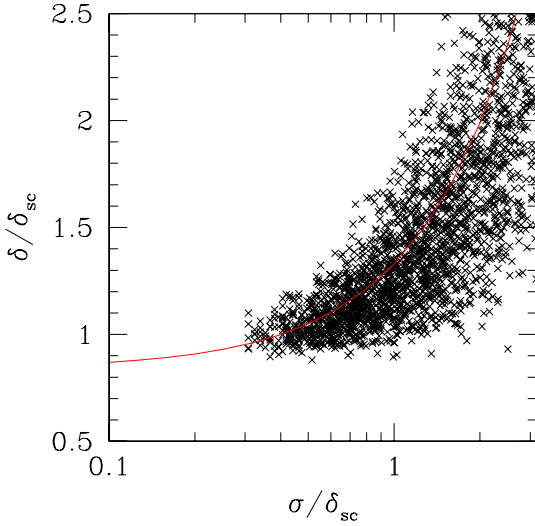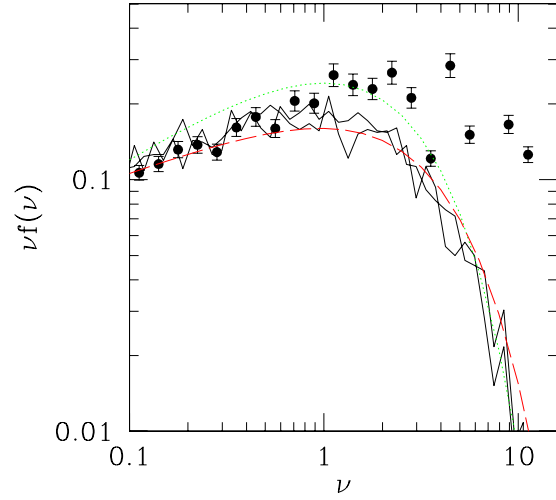
Figure 11. Distribution of overdensities and scales at which the six-dimensional random walks crossed the ellipsoidal collapse barrier. Solid curve shows the approximation used by Sheth, Mo & Tormen (2001) in their, considerably simpler, one-dimensional random walks.

Figure 12. Scaled unconditional (jagged curves) and conditional (filled circles) mass functions associated with six-dimensional random walks which cross the ellipsoidal collapse barrier. Dashed line shows the unconditional mass function which fits the mass function of bound objects in simulations of clustering well. As discussed in the previous section, the dashed curve is well described by the distribution of first crossings by one-dimensional random walks, of a barrier which is associated with ellipsoidal collapse. Dotted line shows the corresponding spherical collapse prediction.

patch may differ from patch to patch. Appendix B provides a simple prescription for choosing a set of patches which have the correct ensemble averaged properties—in essence, this requires studying the first crossing distribution of six-dimensional random walks.

Because a six-dimensional walk is computationally expensive, rather than choosing the distribution of initial patches from this exact statistical distribution, Sheth, Mo & Tormen suggested it should be a good approximation to use an appropriately chosen mean value, and neglect the scatter around this value. This allowed them to reduce what is a six-dimensional random walk to a one-dimensional walk. It is this one-dimensional walk which we have considered so far. One might worry, however, that the discrepancy between model predictions and simulation results we found in the previous section may actually be due to our neglect of the scatter around the average value. The main purpose of this section is to study this possibility.

Fig. 11 shows the result of generating an ensemble of four thousand six-dimensional random walks associated with Gaussian random fields as described in the Appendix. The walks are stopped when they cross the barrier associated with the ellipsoidal collapse model of Bond & Myers (1996). The crosses show the values of δ and σ at which the six-dimensional walks crossed the ellipsoidal collapse barrier $\delta_{\mathrm{crit}}(e,p)$: we actually used the simple fit, equation 3 in Sheth, Mo & Tormen (2001), to the critical density required for collapse $\delta_{\mathrm{crit}}(e,p)$. The solid curve shows the approximation used by Sheth, Mo & Tormen; it provides a rather good description of how δ_{crit} increases with increasing σ. The Appendix describes the reason for this in more detail. For now, we simply note that because the solid curve provides a reasonably good description of the crosses, the first crossing

distributions of the six-dimensional walks considered here are unlikely to be very different from the first crossing distributions associated with the (considerably simpler) one-dimensional walks studied in the previous sections of this paper.

The two jagged curves in Fig. 12 show this explicitly. They show the first crossing distributions associated with the six-dimensional walks which cross the $z = 0$ and $z = 0.5$ six-dimensional ellipsoidal collapse barriers. They have been rescaled similarly to how the unconditional mass functions in simulations rescale: $\nu = (\delta_{\mathrm{sc}}/\sigma)^2$. After rescaling, the two curves appear similar, as they should; note that they are reasonably like the dashed curve, which shows equation (2), and they are rather different from the dotted curve which shows the spherical collapse prediction.

Recall that the dashed curve is very similar to the mass function one gets by simulating one-dimensional random walks. Because the jagged curves are in quite good agreement with the dashed curve, the first crossing distributions associated with the six-dimensional walks are actually rather similar to those associated with the (considerably simpler) one-dimensional walks studied in the previous sections. This agrees with what our conclusions from Fig. 11. The Appendix describes the reason for this in more detail.

Having shown that the unconditional mass functions associated with the six-dimensional random walks are in reasonable with numerical simulations, we now turn to the conditional mass functions—the test which our one-dimensional random walks failed. The solid symbols with error bars

show the conditional mass functions associated with the six-dimensional walks, expressed in the scaled units of the previous section: $\nu = 0.707\,\delta_{\rm sc}^2\,0.5^2/(s-S)$, where s denotes the value of σ^2 at which the $z = 0.5$ six-dimensional boundary was crossed, and S is the scale on which the lower $z = 0$ six-dimensional boundary was crossed. Note the excess of points at large values of ν; these conditional mass functions are quite different from the conditional mass functions in simulations. Indeed, the discrepancy between the excursion set predictions and the simulations of hierarchical clustering has got worse!

There are at least two reasons why the excursion set approach with one-dimensional random walks may fail at small lookback times. The first is the excursion set neglect of correlations between scales (Peacock & Heavens 1990; Bond et al. 1991). At large lookback times most subclumps are a small fraction of the mass of the parent halo, so the smoothing scale associated with the subclumps is sufficiently different from that of the parent that the neglect of correlations between the two scales is probably justified. At smaller lookback times the parent and subclump scales are not so well separated, so the neglect of correlations is a more likely to be a bad approximation. This is one possible reason for the agreement at large lookback times despite the discrepancy at low redshift. The second possibility is that the one-dimensional parametrization of ellipsoidal collapse outlined by Sheth, Mo & Tormen (2001) is too simple. The results of this section suggest that, in fact, it is the first possibility which is the cause of the discrepancy.

4 DISCUSSION

Sheth, Mo & Tormen (2001) argued that a simple modification to the original excursion set approach was enough to improve agreement between the predictions of the approach and numerical simulations. The modification they suggested was to the value of the linearly extrapolated critical overdensity $\delta_{\rm c}$ associated with the collapse of an object. The spherical collapse model assumes that this value is independent of the mass of the collapsed object, whereas ellipsoidal collapse makes $\delta_{\rm c}$ depend on m. In the context of the excursion set approach, this corresponds to studying the first crossing statistics of a set of moving, rather than constant barriers. We also argued that a moving barrier also provides a simple way in which the excursion set appraoch can be extended to apply to models in which the initial dark matter distribution is not completely cold.

Because moving barrier models are so useful, we provided analytic approximations for the required first crossing distributions (equations 4 and 7), and showed that they were reasonably accurate. Although our formulae for the conditional mass functions (solid curves in Figs. 3–8) are slightly different from, and usually more accurate than, those one obtains by a simple rescaling of the unconditional mass function (dashed curves in the same figures), this simple rescaling of the unconditional mass function is still more accurate than what one gets if one rescales the constant barrier formulae (dotted curves). We showed that the predicted unconditional and conditional mass functions were in reasonably good agreement with results from numerical cosmological simulations (Figs. 3–8), and we also provided a simple effi-

cient algorithm which allows one to generating masses which have the correct universal mass function (Section 2.2).

However, we showed that neither the constant nor the moving barrier models were able to describe the simulation results at small lookback times (Figs. 9 and 10). This means that our results for the conditional mass functions cannot be used to generate realizations of the forest of merger history trees. We argued that this discrepancy was most likely due to the excursion set approach's neglect of correlations between scales (e.g. Peacock & Heavens 1990; Bond et al. 1991; Monaco 1997b). While this neglect is a bad approximation at small lookback times, it is reasonably accurate at large lookback times. This is why the excursion set approach is able to provide a reasonably good description of clustering at high redshift, even though it is inaccurate at small redshifts.

Before concluding, we will consider how some of our results are related to other work in the literature. Recently, Jenkins et al. (2001) showed that, although the mass functions in their simulations scaled in accordance with the excursion set prediction, our equation (2) slightly overestimated the unconditional mass functions in their simulations. We thought it would be interesting to show the various approximations to the mass function which we presented in this paper on one plot. The dot-dashed curve in Fig. 13, with cutoffs at low and high masses shows the fitting function Jenkins et al. proposed, which fits their simulations well, the dashed curve shows the fitting formula of equation (2) with $a = 0.707$ (following Sheth & Tormen 1999), the histogram shows the distribution one gets by simulating random walks with the ellipsoidal collapse moving barrier (equation 1, following Sheth, Mo & Tormen 2001), the solid curve shows our approximate formula for this first crossing distribution (equation 4), and the dotted curve shows the spherical collapse, constant barrier prediction (Press & Schechter 1974; Bond et al. 1991). Simulations currently available do not probe the regime where $\nu \leq 0.3$ or so (the Jeans mass is at about $\nu \approx 0.03$).

The upper set of curves show the residuals between our formulae and the one provided by Jenkins et al., in the regime to the right of the low-mass cut-off (marked by an arrowhead). In addition to the previously mentioned formulae, we have included a new dashed curve which shows the result of changing a in equation (2) from 0.707 to 0.75. This simple change appears to be all that is necessary to reduce the discrepancy between it and the simulations substantially. Our formula differs dramatically from the one Jenkins et al. propose at small masses. We hope that simulations in the near future will be able to address which low mass behaviour is correct.

To illustrate that our formulae really do work for a large class of barrier shapes, Fig. 14 shows the first crossing distribution associated with the barrier discussed by Monaco (1997a,b):

$$B(S) = \delta_{\rm sc}\left[1.82/\delta_{\rm sc} - 0.69\,\sqrt{S/S_*}\right],$$

where, as throughout this paper, $S_* \equiv \delta_{\rm sc}^2$. The height of this barrier decreases with S, so it really is quite different from ours (Sheth, Mo & Tormen 2001 discuss the physical reason why). (We chose not to present results for the barrier shape studied by Del Popolo & Gambera (1998) because
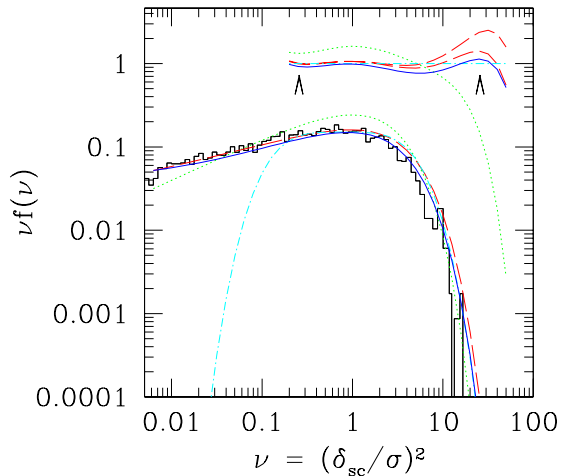
Figure 13. Comparison of the various mass functions described in this paper, with the one which fits the Hubble volume simulations presented in Jenkins et al (2001). Upper curves show the residuals between our analytic formulae and the Jenkins et al. fitting formula.
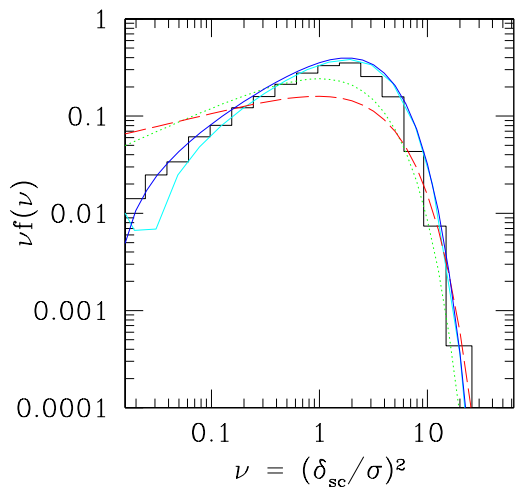
Figure 14. The unconditional mass function associated with a model of collapse which was presented in Monaco (1997a,b). Whereas the height of the barrier studied in our paper increases with distance, the one proposed by Monaco decreases with distance. Despite the dramatic differences between the two barriers, inserting this shape into our equation (4) provides a good fit (solid curve) to the exact result (histogram).

their shape is not so different from ours, whereas Monaco's really is quite different. The fact that this barrier decreases with S means that all walks are guaranteed to cross it, and that there is no fragmentation associated with this barrier shape. (This is in contrast to barriers whose height increases sufficiently strongly with S; for the Sheth, Mo & Tormen 2001 barrier shape studied in the main text, unbound mass and fragmentation are features which are formally possible but rare in practice.) The histogram shows the numerical Monte-Carlo first crossing distribution, and the two solid curves shows our analytic approximation, computed by inserting this barrier shape into our equation (4). The curve which provides a slightly worse fit to the histogram shows the result of using the first five terms in the series (as we did for the other figures in this paper); the other curve shows what happens if we use the first ten terms instead. Just for comparison, the dotted and dashed curves show the spherical collapse prediction, and the one which actually fits the cosmological simulations (equation 2). The figure shows that we are able to describe the first crossing distribution of this barrier shape well. This means that one could, in principle, use our formula, with Monaco's barrier, to study the conditional mass functions associated with his parametrization of nonlinear collapse—we have not pursued this further, although comparison of this first crossing distribution with the $z = 0.5$ panels in Figs. 9 and 10 suggest that this might be a useful exercise.

In summary, we have provided a formula which describes the first crossing distribution of independent uncorrelated Brownian motion random walks, for a wide class of moving barriers. This formula can be used to provide simple accurate predictions for a number of statistical quantities associated with the formation and clustering of dark matter haloes, all within the same formalism.

REFERENCES

Bode P., Ostriker J. P., Turok N., 2001, ApJ, to appear

Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, ApJ, 379, 440

Bond J. R., Myers S., 1996, ApJS, 103, 1

Chiueh T., Lee J., 2001, ApJ, accepted

Del Popolo A., Gambera M., 1998, AA, 337, 96

Heavens A. F., Peacock J. A., 1988, MNRAS, 232, 339

Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, MNRAS, 321, 372

Lacey C., Cole S., 1993, MNRAS, 262, 627
Lacey C., Cole S., 1994, MNRAS, 271, 676
Lemson G., Kauffmann G., 1999, MNRAS, 302, 111
Mo H. J., White S. D. M., 1996, MNRAS, 282, 347
Mo H. J., Jing Y., White S. D. M., 1997, MNRAS, 284, 189
Monaco P., 1997a, MNRAS, 287, 753
Monaco P., 1997b, MNRAS, 290, 439
Nusser A., Sheth R. K., 1999, MNRAS, 303, 685
Peacock J. A., Heavens A. F., 1990, MNRAS, 243, 133
Press W., Schechter P., 1974, ApJ, 187, 425
Sheth R. K., 1996, MNRAS, 281, 1277
Sheth R. K., Pitman J., MNRAS, 289, 66
Sheth R. K., 1998, MNRAS, 300, 1057
Sheth R. K., Lemson G., 1999a, MNRAS, 304, 767
Sheth R. K., Lemson G., 1999b, MNRAS, 305, 946
Sheth R. K., Tormen G., 1999, MNRAS, 308, 119
Sheth R. K., Mo H. J., Tormen G., 2001, MNRAS, 323, 1
Tormen G., 1997, MNRAS, 290, 411
Tormen G., 1998, MNRAS, 297, 648

APPENDIX A: THE LINEAR BARRIER

Suppose that the barrier shape increases linearly with increasing variance $S = \sigma^2$:

$$B(S, z) = \delta_c(z) \left(1 + S/S_*(z) \right), \qquad (A1)$$

where $\delta_c(z) = \delta_{c0}(1 + z)$ and $S_*(z) = S_{*0}(1 + z)^2$ if $\Omega = 1$. These scalings with z are what is required by the self-similarity of Brownian motion. In the Bond et al. (1991) formulation of the constant barrier problem, it is customary to express the S axis in the units it had at some fiducial time, say $z = 0$, and to study the successive crossings of barriers having different values of z. In effect, the constant height Press–Schechter barrier has $S_{*0} = \infty$, so in that case only the scaling of the y-axis was apparent.

This linear barrier shape is motivated by the observation that the GIF simulations have fewer low mass haloes relative to high mass ones, as compared to what is predicted by the constant barrier model. This means that, at least for some range of S, the moving barrier must have a positive slope, since this would make it relatively easier to cross at small S (large mass) than at large S (small mass), as compared to a barrier of constant height.

An additional reason for considering this linear barrier is the following. Recent work (Bode, Ostriker & Turok 2001) suggests that the halo mass function in which the dark matter is warm initially has even fewer low mass halos than cold dark matter based ellipsoidal collapse models predict. In the context of the approach outlined by Sheth, Mo & Tormen (2001), the physical reason for this is relatively simple: low mass haloes do not form because they are hotter initially, so a larger overdensity is required to hold them together against the thermal pressure which prevents collapse, or against the stronger shearing from the velocity field. This suggests that the critical density required for collapse by the present time, $\delta_{ec}(m)$, should increase even more strongly with decreasing m than it does when the dark matter is cold. The warm dark matter model is not particularly well motivated, and its free parameters have not yet been fixed, it seems premature to provide a detailed $\delta_{ec}^{WDM}(m)$ relation at the present time. For this reason, the linear barrier considered in this appendix should be thought of as an example of what happes

when the barrier height increases even more steeply with decreasing mass than it does in the Sheth, Mo & Tormen cold dark matter models.

In the constant barrier model, all random walks were guaranteed to cross the barrier. This is because the rms height of random walks at S is proportional to $\sqrt{S}$, so at sufficiently large S, all walks will have crossed the constant barrier. As a result, the associated first crossing distribution is normalized to unity: since each random walk is associated with a volume element in the initial Lagrangian space, this is usually interpretted as meaning that, in the constant barrier model, all the mass in the universe is bound up in collapsed objects of some mass, however small. In contrast, the linear barrier (A1) increases to arbitrarily high values at high S. Because the rms height of the random walk grows more slowly than the rate at which the barrier height increases, there is no guarantee that all random walk trajectories will intercept this barrier. Indeed, only a fraction $e^{-2\delta_c^2(z)/S_*(z)}$ of them will. So, in the linear barrier model, not all initial volume elements are associated with bound haloes. Since not all particles in numerical cosmological simulations are associated with bound haloes anyway (the fraction of unbound mass is typically on the order of $\sim$10%, though how much of this is a consequence of limited resolution in the simulations is uncertain), this feature of the moving barrier model may or may not be a good thing. In any case, this is one qualitative difference between the moving boundary model and a model with constant barrier height. (Readers who dislike this feature of the linear model are invited to patch a constant barrier of height δ_{c0} at small S to a linearly increasing one at intermediate S to another constant one (but now at at a greater height, of course) at a large S of their choice, and to compute the associated mass function! This can be done relatively easily using the results presented below.)

In addition to the question of normalization, barriers whose heights increase with S more rapidly than $S^{0.5}$ will have mass functions in which the low mass end is depleted relative to the constant barrier case. In the linear barrier model studied below, the mass function has an exponential cutoff at both low and high masses. Because the barrier shape associated with ellipsoidal collapse grows only as $S^{0.6}$ both these effects are extremely weak. Nevertheless, it is worth bearing these features in mind as bigger and better simulations become available.

The second important qualitative difference between a moving barrier model and the original constant barrier model is the following: whereas the y-intercept $B(0, z)$ increases as z increases, the slope decreases as $(1 + z)^{-1}$. This means that it is possible for barriers to intersect at finite values of S. For example, two linear barriers $B(S, z_0)$ and $B(S, z_1)$, with $z_1 > z_0$, will intersect at that critical value of S at which $B(S, z_0) = B(S, z_1)$: $S_{01}/S_{*0} = (1 + z_0)(1 + z_1)$. This means that all trajectories which first cross $B(z_0)$ at $S > S_{01}$ must have crossed $B(z_1)$ at a smaller value of S. The logic of Lacey & Cole (1993) then says that all haloes at z_0 that are less massive than the associated critical mass M_{01} were formed by fragmentation of a halo which, at $z_1 > z_0$, was more massive.

Again, this may or may not be a good thing. Simulations show that $\sim$20% of the total mass ever associated with progenitors of a halo does not find its way to the parent halo (Tormen 1998). Presumably this reflects the fact that while
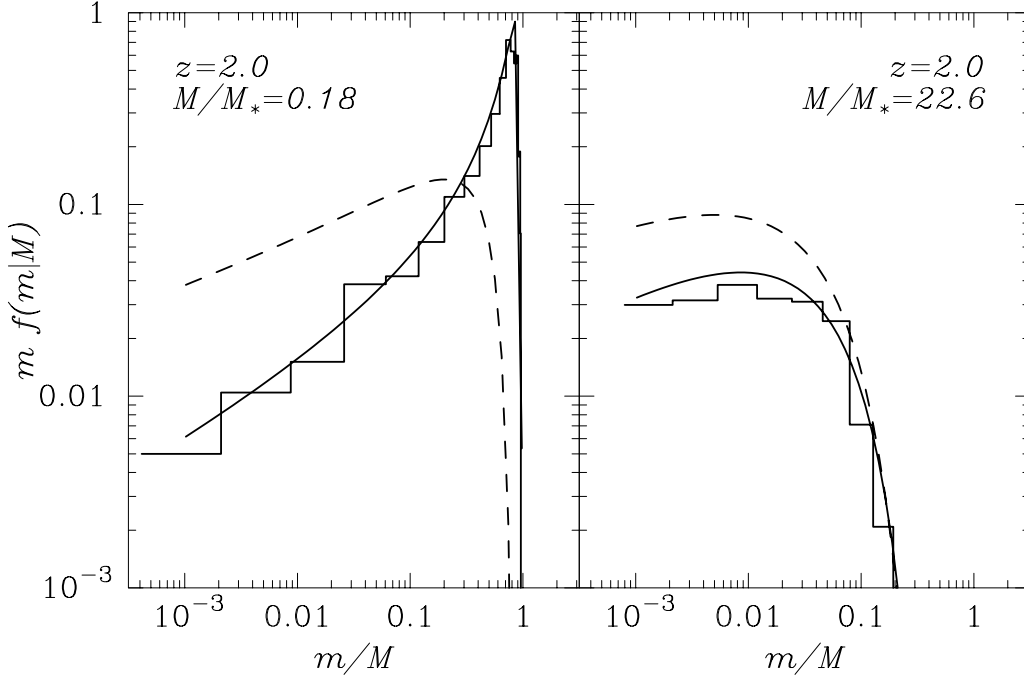
Figure A1. The conditional mass functions associated with the linear barrier. Smooth solid curves show the analytic linear barrier prediction, dashed curves show the conditional constant barrier distribution for comparison, and histograms show the result of our numerical Monte–Carlo calculation.

small haloes may fragment, the fragmented mass is not a very large fraction of the total mass. Simulations also show that some subclumps pass through the virial radius of a given parent halo a number of times, each time depositing some fraction of their mass, before they finally become part of the parent halo. It may be that a moving barrier model is able to incorporate and perhaps even quantify these effects.

The first crossing distribution of trajectories that do cross the linear barrier is Inverse Gaussian:

$$f(S, z) \, dS = \frac{B(0, z)}{\sqrt{2\pi S}} \, \exp\left(-\frac{B^2(S, z)}{2S}\right) \frac{dS}{S}, \qquad (A2)$$

(Whitmore 1978; also see Sheth 1998). The associated unconditional mass function is got by inserting this in equation (3). Whereas the mass function in the constant barrier case has an exponential cutoff at the high mass end, the mass function in this linear barrier model has exponential cutoffs at both low and high masses. Recall this arises from the fact that barrier height increases so rapidly at large S. Therefore, the mass function associated with this linear barrier has fewer small mass objects than the constant barrier predicts, in qualitative agreement with the GIF simulations. (The agreement is certainly not quantitative, but the linear barrier is only used here for illustrative purposes.)

Since a linear barrier is linear whatever the origin of the coordinate system, the solution to the two barrier problem is also Inverse Gaussian. That is, given $z_1 > z_0$ and given that the first crossing of $B(z_0)$ occured at $S_0 \leq S_{01}$, the probability that the first crossing of $B(z_1)$ occurs in the range dS_1 about S_1 is given by equation (A2), with the substitutions $S \rightarrow (S_1 - S_0)$, and $B \rightarrow B_{10}$, where

$$B_{10}(S_1 - S_0) = B(S_1, z_1) - B(S_0, z_0).$$

As was the case for the unconditional mass function, the height of this barrier diverges as $S_1 - S_0 \rightarrow \infty$, so not all trajectories intersect it. Again, it seems reasonable to associate the fraction that do not with the fraction of the parent halo mass that is not associated with bound subclumps. (Those readers who computed the mass functions associated with the patched constant–linear–constant barriers may disregard the previous two sentences, provided they first compute the associated two-patched-barrier problem!)

The smooth solid curves in Fig. A1 show this conditional distribution for a few representative choices of the parent halo mass M. The parents were assumed to have formed at $z_0 = 0$, and the progenitor distributions are shown at an earlier redshift z_1. The underlying power spectrum was chosen to be the same as the GIF SCDM power spectrum. It is also straightforward to solve this problem numerically: for the values of M shown in the figures, the histograms show the distribution generated by simulating the crossings of the higher linear barrier by 10^3 random Brownian motion trajectories that started at the initial positions $B(S, z = 0)$, where $S(M)$ is given by the GIF power spectrum. This figure has been included mainly to show that our Monte–Carlo code works, because we will use the code in the next section. We have also verified that the numerical code gives the correct conditional and unconditional mass functions when the barrier heights are constant.

For completeness, we also show the bias relations associated with this model. Following Mo & White (1996) (also see Sheth & Tormen 1999), the mean Lagrangian bias between haloes and mass is

$$\delta_h^L(1|0) = \frac{f(S_1, z_1|S_0, z_0)}{f(S_1, z_1)} - 1. \qquad (A3)$$

The limit $M_0 \gg M_1$ and $|\delta_0| \ll \delta_1$ is sometimes called the peak-background split. In this limit

$$\delta_{\rm h}^{\rm L}(1|0) \to (\nu_1/\delta_1)\,\delta_0, \quad \text{where} \quad \nu_1 \equiv \delta_c^2(z_1)/S_1$$

for the linear barrier model. Massive haloes have small values of S, so they have large values of ν. Small haloes have $\nu \sim 0$. Thus, in this limit, less massive haloes are unbiased relative to the mass, whereas massive haloes are positively biased. For comparison, the corresponding limit for the constant barrier is

$$\delta_{\rm h}^{\rm L}(1|0) \to \delta_0\,(\nu_1 - 1)/\delta_1.$$

Thus, the predictions of the linear and constant barriers are similar for massive haloes, but they differ for less massive ones. In particular, less massive haloes in Lagrangian space are anti-biased in the constant barrier model, whereas they are unbiased in the linear barrier model.

The halo-to-mass bias in the evolved Eulerian space for the constant barrier can be computed by expanding

$$\delta_{\rm h}^{\rm E}(1|0) = (1 + \delta)\,\left[1 + \delta_{\rm h}^{\rm L}(1|0)\right] - 1 \tag{A4}$$

to lowest order in δ (Mo & White 1996). In this limit, $\delta \approx \delta_0$, so $\delta_{\rm h}^{\rm E}(1|0) = (1 + [\nu_1 - 1]/\delta_1)\,\delta$ for the constant barrier. The same logic gives $\delta_{\rm h}^{\rm E}(1|0) = (1 + \nu_1/\delta_1)\,\delta$ for the linear barrier. Again, the predictions of the constant and linear barriers agree for massive haloes, but are different for less massive ones; the haloes in the linear barrier model are more positively biased. This is encouraging, because, as mentioned in the introduction, this is in qualitative agreement with the trend seen in numerical simulations.

The rate of increase of the ellipsoidal collapse barrier is shallower than for the linear barrier discussed here. Since it is intermediate between the linear barrier and the constant spherical collapse barrier, we might reasonably expect the large scale bias factor of less massive haloes in the ellipsoidal collapse model to be greater than that associated with the spherical collapse model. Since the ellipsoidal collapse barrier rises less steeply than linear, the modified large scale bias will be somewhat less than the linear barrier prediction. Fig. 4 of Sheth, Mo & Tormen (2001) shows that this is, indeed, the case.

We argued at the start of this section that warm dark matter models can be parametrized by making the critical density for collapse depend more strongly on mass than when the dark matter is cold. In this respect, a comparison of the linear barrier formulae given here with the results for the barrier written down by Sheth, Mo & Tormen (2001) shows why, generically, warm dark matter models are expected to have fewer low mass haloes and, consequently, different bias relations, particularly at the low mass end, than cold dark matter models. This is in qualitative agreement with the numerical simulations of Bode, Ostriker & Turok (2000).

APPENDIX B: DISTRIBUTION OF DENSITY AND ANGULAR MOMENTUM OF A PATCH IN A GAUSSIAN RANDOM FIELD

Let $d_{ij} = \nabla_{ij}\phi$, where ϕ is the initial potential, denote the various components of the deformation tensor $\mathcal{D}$ (here $1 \le i \le 3$ and similarly for j). Following, e.g. Bardeen et al.

(1986), these components are

$$
\begin{aligned}
d_{11} &= (-y_1 - 3y_2/\sqrt{15} - y_3/\sqrt{5})/3 \\
d_{22} &= (-y_1 + 2y_3/\sqrt{5})/3 \\
d_{33} &= (-y_1 + 3y_2/\sqrt{15} - y_3/\sqrt{5})/3 \\
d_{12} &= y_4/\sqrt{15} \\
d_{23} &= y_5/\sqrt{15} \\
d_{13} &= y_6/\sqrt{15}
\end{aligned} \tag{B1}
$$

where the y_is are independent Gaussian variates with zero mean and variance σ^2. Poisson's equation says that the trace of this matrix, $\mathrm{Tr}(\mathcal{D})$ equals the overdensity δ. Thus,

$$\delta = d_{11} + d_{22} + d_{33} = -y_1; \tag{B2}$$

the final expression shows explicitly that δ is a Gaussian random variate.

The eigenvalues of this matrix are the roots of the characteristic equation

$$P(\lambda) = \mathrm{Det}[\mathcal{D} - \lambda\mathcal{I}] = -a_0 - a_1\lambda - a_2\lambda^2 - \lambda^3, \tag{B3}$$

where $\mathcal{I}$ is the identity matrix, and the a_ks are various combinations of the d_{ij}s got by expanding the expression above and ordering by powers of λ. Since $\mathcal{D}$ is a 3×3 real symmetric matrix, $P(\lambda)$ is a cubic with three real roots which satisfy

$$
\begin{aligned}
\lambda_1 + \lambda_2 + \lambda_3 &= -a_2 \\
\lambda_1\lambda_2 + \lambda_2\lambda_3 + \lambda_1\lambda_3 &= a_1 \\
\lambda_1\lambda_2\lambda_3 &= -a_0.
\end{aligned} \tag{B4}
$$

Because rotations leave the trace unchanged, the overdensity δ is the sum of the three eigenvalues, so $-a_2 = \delta = -y_1$. In addition, the square of the angular momentum is proportional to

$$
\begin{aligned}
r^2 &= \frac{(\lambda_1 - \lambda_2)^2}{2} + \frac{(\lambda_2 - \lambda_3)^2}{2} + \frac{(\lambda_1 - \lambda_3)^2}{2} \\
&= a_2^2 - 3a_1
\end{aligned} \tag{B5}
$$

(e.g. Heavens & Peacock 1988; Catelan & Theuns 1998). This shows that to get r^2 we don't need to solve the cubic— we just need to read off the appropriate coefficients of the characteristic equation. Thus, we find that

$$
\begin{aligned}
r^2 &= \delta^2 + 3(d_{12}^2 + d_{13}^2 + d_{23}^2 - d_{11}d_{22} - d_{11}d_{33} - d_{22}d_{33}) \\
&= (y_2^2 + y_3^2 + y_4^2 + y_5^2 + y_6^2)/5
\end{aligned} \tag{B6}
$$

Although the first line suggests that r^2 is coupled to δ, the final expression shows that it is not. In particular, the expressions above show that δ is distributed as a Gaussian, and r^2 is an independent variate drawn from a Chi-square distribution with five degrees of freedom, $\chi_5^2(\sigma)$. The fact that the overdensity and the square of the angular momentum are independent does not seem to have been noticed before. A $\chi_5^2(\sigma)$ distribution is rather similar in shape to a Lognormal, so this provides a simple way to see why spins of peaks in Gaussian random fields are also approximately lognormal (Heavens & Peacock 1988).

The results above can be used to generalize the excursion set algorithm studied in the main text. Set $n = 0$ and $y_i(n) = 0$ for $1 \le i \le 6$. Thereafter, at each step labeled by n, choose six, rather than one, independent Gaussian random variates g_i, each with variance s. For each $i = 1, 6$ set
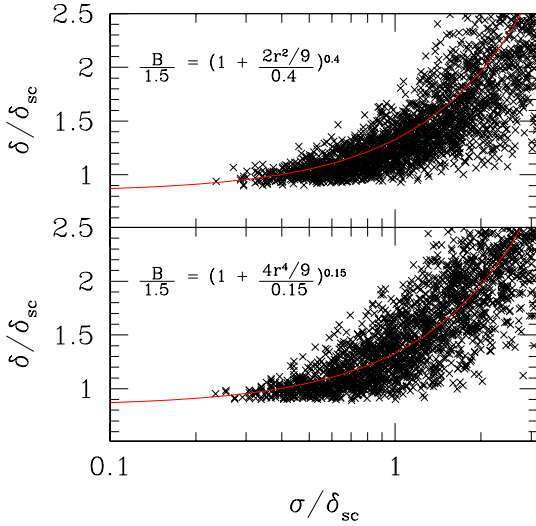
Figure B1. Comparison of two prescriptions for collapse, both of which give good agreement with the unconditional mass function. Crosses show the actual values of collapse densities and scales, and solid curves show the approximation assumed by Sheth, Mo & Tormen (2001).

$y_i(n) = y_i(n-1) + g_i$, and use these to compute δ and r^2. These can be used to give the values of the overdensity and the angular momentum for the scale on which the variance is $\sigma^2 \propto ns$ (recall that n is the number of steps taken). Now check to see if $\delta = -y_1(n)$ exceeds a critical value, say $\delta_{\rm crit}(\sigma, r^2)$. If it does, the six-dimensional walk stops at this scale. If not, the walk continues to smaller scales. Because each of the g_is is chosen independently of the values of the y_is or of σ, the walk takes independent steps in the six-dimensional space; it is in this sense that this algorithm generalizes the one-dimensional excursion set random walk studied in the main body of this paper.

The analysis above shows clearly that the one-dimensional random walk approach of Sheth, Mo & Tormen (2001) corresponds to the following approximation. Replace the dependence of $\delta_{\rm crit}(\sigma, r^2)$ on the random variate r^2 by a dependence on its average value $\langle r^2 \rangle \propto \sigma^2$. This means that the critical density for collapse is a function of σ alone, $\delta_{\rm crit}(\sigma)$. As a result, the random walk in six-dimensions can be reduced to a walk in one-dimension only, thereby greatly reducing the complexity of the problem.

Recently just such a six-dimensional random walk algorithm has been used by Chiueh & Lee (2001), although they did not notice the considerable simplifications which follow from the algebra presented above. They simulated an ensemble of six-dimensional random walks, and set the parameters of the barrier to be crossed by requiring that the resulting first crossing distribution give the unconditional mass function. In particular, they showed that $\delta_{\rm crit} = 1.5[1 + (2r^2/3)^2/0.15]^{0.15}$ provided a good fit to the required critical value of the overdensity.

The algebra above allows one to see what such an approach implies. To do this, suppose the barrier is $\delta_{\rm crit} =$

$1.686(1 + r^2)$. The dependence on r^2 means that if the particle has walked to $\delta = 1.686$, it will still not have crossed, because r^2 is always certainly greater than zero. So, to cross, the particle needs to have some $\delta > 1.686$. How much greater? This depends on the typical value of r^2. Because r^2 is drawn from a χ_5^2 distribution, $\langle r^2 \rangle \sim \sigma^2$. Now suppose that the χ_5^2 distribution is very sharply peaked at its mean value (it quite well peaked, but taking the extreme case helps to see the argument). This means that the barrier shape is something like $\delta_{\rm crit} = 1.686(1 + \sigma^2)$. The fact that a χ_5^2 distribution is not very sharply peaked at its most probable value simply means that sometimes when $\delta = 1.686(1 + \sigma^2)$ the particle will still be less than $\delta_{\rm crit}$, so the walk must go on. Of course, sometimes $r^2 < \sigma^2$, and in this case the walk will stop even if $\delta < 1.686(1 + \sigma^2)$. So, this means that we can think of the dependence on r^2 as making the critical value of the boundary height, when expressed as a function of σ^2 (the way Sheth, Mo & Tormen 2001 did) a little fuzzy. So, provided the χ_5^2 distribution is not too broad, the considerably simpler one-dimensional random walk approach of Sheth, Mo & Tormen (2001) should be a reasonable approximation.

Fig. B1 shows the result of doing this for two choices of the barrier shape, both of which produce first crossing distributions which, when inserted into equation (3), give mass functions of bound objects which are similar to the one in simulations of hierarchical clustering. The upper panel was constructed using a barrier whose height increased linearly with r^2, and the lower panel shows results for the barrier shape used by Chiueh & Lee (2001), which increases as r^4. The crosses show the values of δ and σ at which each six-dimensional random walk crossed the barrier. The solid curve shows the approximation used by Sheth, Mo & Tormen (2001); it provides a reasonable description of the increase of $\delta_{\rm crit}$ with σ.

Determining the barrier shape by requiring agreement with the clustering simulations is unsatisfying, especially in view of the fact that the two different boundaries given above provide equally adequate approximations to the mass function. For this reason, the main text shows the result of combining the six-dimensional walk described above with the ellipsoidal collapse model of Bond & Myers (1996). This was relatively easy to do, because a simple fitting function for how the critical collapse boundary associated with this ellipsoidal collapse depends on the initial shear field has been given by Sheth, Mo & Tormen (their equation 3 and their Fig. 1).