



An Empirical Mass Function Distribution

S. G. Murray^{1,2,3} , A. S. G. Robotham¹ , and C. Power^{1,2} ¹ International Centre for Radio Astronomy Research (ICRAR), University of Western Australia, Crawley, WA 6009, Australia² ARC Centre of Excellence for All-sky Astrophysics, Australia³ International Centre for Radio Astronomy Research (ICRAR), Curtin University, Bentley, WA 6102, Australia

Received 2017 July 31; revised 2017 December 14; accepted 2018 January 2; published 2018 February 27

Abstract

The halo mass function, encoding the comoving number density of dark matter halos of a given mass, plays a key role in understanding the formation and evolution of galaxies. As such, it is a key goal of current and future deep optical surveys to constrain the mass function down to mass scales that typically host L_* galaxies. Motivated by the proven accuracy of Press–Schechter-type mass functions, we introduce a related but purely empirical form consistent with standard formulae to better than 4% in the medium-mass regime, 10^{10} – $10^{13} h^{-1} M_\odot$. In particular, our form consists of four parameters, each of which has a simple interpretation, and can be directly related to parameters of the galaxy distribution, such as L_* . Using this form within a hierarchical Bayesian likelihood model, we show how individual mass-measurement errors can be successfully included in a typical analysis, while accounting for Eddington bias. We apply our form to a question of survey design in the context of a semi-realistic data model, illustrating how it can be used to obtain optimal balance between survey depth and angular coverage for constraints on mass function parameters. Open-source Python and R codes to apply our new form are provided at <http://mrpy.readthedocs.org> and <https://cran.r-project.org/web/packages/tggd/index.html> respectively.

Key words: dark matter – galaxies: halos – methods: analytical – methods: statistical

1. Introduction

Our understanding of galaxy evolution is largely hinged upon the galaxies’ dark scaffolding—the dark matter distribution and its network of filaments, halos, and voids. In particular, every galaxy is believed to be housed in a dark matter halo, each the product of a turbulent series of mergers over cosmic time. This history of mergers gives rise to an elegant picture of cosmic structure formation: the so-called hierarchical paradigm (White & Rees 1978). This paradigm has been helpful in allowing predictions of many key characteristics of the halo distribution, including the halo mass function (HMF): the comoving number density of halos as a function of their mass.

Clearly, the evolution of the halo distribution and that of the hosted galaxies is fundamentally interconnected. For this reason, measurement of the HMF and its connection to observed galaxy properties are of fundamental importance. While the abundance of cluster-mass halos is tightly related to the underlying cosmology (e.g., Bocquet et al. 2016; de Haan et al. 2016), the galaxy population is predominantly hosted by halos in the medium-mass regime $\sim 10^{10}$ – $10^{13} h^{-1} M_\odot$. Measurement of the HMF in this regime via dynamical mass estimates is thus a key goal for current surveys such as GAMA (Galaxy and Mass Assembly, Driver et al. 2011) and future deep surveys such as WAVES (Wide Area VISTA Extragalactic Survey, Driver et al. 2016) and MSE (Manukea Spectroscopic Explorer, McConnachie et al. 2016).

The utility of measuring the HMF with dynamical halo masses inferred from groups of galaxies has been shown in Eke et al. (2006), who used the 2PIGG (2dFGRS Percolation-Inferred Galaxy Group) catalog. Using a simple empirical correction to mitigate Eddington bias, they were able to infer the group luminosity function (LF) and mean mass-to-light ratio of galaxies within their sample. Additionally, they were able to infer a best-fit value of σ_8 by comparing to simulation results. However, similar studies with improved analysis

methods and larger surveys have not been forthcoming. Instead, there has been an increased focus on less direct methods of connecting the dark sector to galaxies, such as the halo occupation distribution (HOD; Berlind et al. 2003; Zehavi et al. 2005, 2011). One reason for this focus may be an increased reliance on the standard formalism for theoretical HMFs.

The standard model for the HMF arises from physical arguments proposed by Press & Schechter (1974), and we shall refer to this framework (with its many specific fits) collectively as the “PS formalism” (not to be confused with the specific original form of Press & Schechter 1974) and we shall refer to any fit within this framework as a “PS-type” fit. In this model, the HMF can be written

$$\frac{dn}{dM} = \frac{\rho_0}{M^2} \left| \frac{d \ln \sigma}{d \ln M} \right| f(\sigma), \quad (1)$$

where ρ_0 is the universal mean matter density, and the model is dependent on the underlying cosmology via the mass variance, $\sigma^2(m)$, which is the second moment of the matter power spectrum, smoothed (by an arbitrary filter W) on a physical scale R corresponding to m :

$$\sigma^2(R) = \frac{1}{2\pi^2} \int_0^\infty dk k^2 P(k) W^2(kR). \quad (2)$$

Assuming spheroidal collapse of halos, Press & Schechter (1974) derived a form for $f(\sigma)$ given by

$$f(\sigma) = \sqrt{\frac{2}{\pi}} \frac{\delta_c}{\sigma} \exp\left(-\frac{\delta_c^2}{2\sigma^2}\right), \quad (3)$$

where $\delta_c \approx 1.686$ is the critical overdensity for spherical collapse.

The PS formalism has proven remarkably successful in describing the HMF measured in cosmological simulations, albeit with modifications of $f(\sigma)$ to better capture the

non-spherical nature of halo collapse (e.g., Sheth et al. 2001). However, though it appears to be universal with respect to cosmology and redshift, it has been shown to exhibit non-universality at the $\sim 20\%$ level, especially for halos defined by a spherical-overdensity criterion (Tinker et al. 2008; Watson et al. 2013). Furthermore, the best-fit seems to depend not only on underlying physics and halo definition, but on the algorithm used for locating halos within cosmological simulations (Knebe et al. 2011). Though these uncertainties do not prevent the use of the PS formalism in accurate fits to cluster abundance (de Haan et al. 2016), they motivate the use of a simpler empirical distribution in the context of observed HMFs at lower masses.

In addition to this motivation, the most common descriptions of ensemble galaxy properties are entirely empirical; both LFs and galaxy stellar-mass functions (GSMFs) are parameterized by single- or double-Schechter functions (Schechter 1976), which capture their forms adequately without any first-principle motivation. Thus, for example, in the single-Schechter case, the LF is commonly expressed as

$$\frac{dn}{dL} = \phi_* \left(\frac{L}{L_*} \right)^{-\alpha} \exp \left(-\frac{L}{L_*} \right). \quad (4)$$

Here the model parameters themselves hold explanatory power: L_* is a characteristic turn-over luminosity, and α is a power-law slope for low-luminosity galaxies. Measuring the HMF via the PS formalism obfuscates the relation between parameters and data, rendering the comparison of the dark HMF with the visible galaxy properties less clear. Given the apparently similar forms of the HMF, LF, and GSMF, it would seem beneficial to construct their parameterizations in a similar fashion so as to facilitate comparison; for example, can we relate the evolution of a typical halo mass \mathcal{H}_* with that of the typical stellar mass, M_* , or luminosity L_* ?

Using a simple empirical form defined directly in terms of mass also bears other advantages. It is portable, both in the sense that anyone can easily create their own numerical function for its evaluation and that data described by the form are reduced to a small set of interpretable parameters. It is fast since it is a direct calculation. And, since it is analytic, it affords more extended analyses and methods.

Conversely, the empirical nature of the form means that it is not easily linked with an underlying cosmology. For this reason, it will not be useful for cluster abundance studies, which primarily seek to constrain the power spectrum normalization σ_8 and dark energy parameters. Rather, it is most usefully applied to questions about the connection between dark halos and galaxies.

Indeed, such a form would seem to be appropriately used in empirical low-mass HMF studies in a similar way to how the Schechter function is used for LF or galaxy mass function studies. That is, the results of HMF analyses can simply be expressed as a collection of a small number of conceptually informative parameters with their uncertainties—furthermore, these can be directly compared to parameters of the galaxy population. In this way, the form would be seen as an observer’s tool, rather than a theoretician’s model.

In this paper, we propose an empirical form that is equivalent to a renormalized generalized gamma distribution (hereafter GGD). The renormalized GGD (which throughout we term the MRP, an acronym for the author’s surnames) is a generalization of the Schechter function, in which the sharpness of the exponential cutoff is variable. This similarity enables direct

comparison between the HMF and galaxy ensemble properties (see Section 6). Furthermore, its simple analytical nature offers a number of advantages in estimation of its parameters via hierarchical Bayesian models, which serve to account for potential distortions such as Eddington bias.

After introducing the form and its basic properties in Section 2, we outline a fully general method of parameter estimation with the MRP in Section 3, including arbitrary measurement uncertainties and selection effects. Using these methods, in Section 4, we assess the accuracy of the MRP form with respect to both prediction of the PS formalism and large cosmological simulations.

Following this, we outline two simple ways in which the MRP may be linked to physical parameters: first, in Section 5, we provide general fitting functions that relate the MRP to the underlying cosmology. Following this, in Section 6, we develop a generic form for the stellar-mass halo-mass (SMHM) relation, which connects halos with their stellar-mass content, and which in the case of the MRP can be specified exactly in both mass limits.

After this laying out of methodology and justification of our use of the MRP, in Section 7, we set out to illustrate application of the MRP to questions of survey design, in the process utilizing many of the results of the previous sections. Finally, we conclude in Section 8.

To accompany this paper, we have developed an extensive Python code, called MRPY,⁴ which implements all of the functionality explored throughout the paper. Furthermore, we provide a reduced R version of this code that focuses on the basic statistical applications, called TGGD.⁵

2. The MRP Distribution

The MRP has four parameters, and can be written

$$\begin{aligned} \frac{dn}{d \log_{10} M} &\equiv \frac{dn}{dm} \equiv \phi(m|\boldsymbol{\theta}) \\ &= A\beta 10^{(\alpha+1)(m-h_*)} \exp(-10^{\beta(m-h_*)}) \\ &\equiv g(m|\boldsymbol{\theta}), \end{aligned} \quad (5)$$

where $m = \log_{10} M$ is the logarithmic mass and $\boldsymbol{\theta}$ is the parameter vector ($h_* \equiv \log_{10} \mathcal{H}_*$, α , β , $\ln A$). We purposely set ϕ as the generalized number density as a function of logarithmic mass, while for convenience we denote the specific MRP form as $g(m|\boldsymbol{\theta})$.

We note that g is a generalization of the Schechter function, which has a fixed $\beta = 1$, so that many of the results derived in this paper will be applicable (in a modified form) to the Schechter function also.

Usually, the GGD is defined for $A > 0$, $\beta > 0$, and $\alpha > -1$, where the latter two constraints ensure convergence at the high- and low-mass ends respectively. We dispense with the constraint on α , under the assumption that a turnover occurs well below the observed mass limits to ensure convergence. Indeed, typical HMFs exhibit $\alpha \sim -2$, and we note that for $\alpha \leq -2$ the GGD is divergent in total mass density if not lower-truncated. Note that it is also common for HMFs defined under the PS formalism to be

⁴ Found at <http://mrpy.readthedocs.org>. Note that wherever possible, we provide the exact code to reproduce the figures in this paper as examples at <https://github.com/steven-murray/mrpy/examples>.

⁵ Found at <https://github.com/asgr/tggd> and downloadable from CRAN as `tggd`.

divergent in the lower limit (Jenkins et al. 2001), with the understanding that extra modeling is required below scales that we have currently probed in cosmological simulations.

2.1. Relationship to the PS Formula

The MRP form can be precisely recovered via the PS formalism if we assume the analytic fit of Press & Schechter (1974) and that $\sigma(m)$ can be approximated by a power law, $(m/m_1)^{-\gamma}$. Doing so yields the following equalities:

$$\{\mathcal{H}_*, \alpha, \beta, A\} = \left\{ m_1 \sqrt{a}, \gamma - 2, 2\gamma, \frac{a}{m_1 \sqrt{\pi}} \right\}, \quad (6)$$

where $a = (2/\delta_c^2)^{1/\gamma}$. In practice, the MRP is not so restrictive— α and β are free—so it is able to achieve more precise fits than this approximation would suggest.

We note also that the GGD naturally arises by considering the sum of exponential variates, raised to the power β . However, it is unclear how such a process relates to the HMF. Nevertheless, this composition of the GGD from simpler distributions motivates an expectation that a comprehensible chain of stochastic events leads to the HMF intrinsically having a form close to the GGD.

2.2. Cumulative Distribution

An important quantity is the integral of g over mass, which determines the normalization of the distribution. Two cases are important here. When $\alpha > -1$, which is not typical of the HMF, but may be achieved in practice by a power-law selection function (see the next section), we have

$$\int g(m|\theta) dm = A \mathcal{H}_* \Gamma(z_0) \equiv q_0(\theta). \quad (7)$$

When $\alpha < -1$, the integral is nonconvergent, and we present the result for a distribution truncated at some lower mass:

$$\int_{m_{\min}}^{\infty} g(m|\theta) dm = A \mathcal{H}_* \Gamma(z_0, x) \equiv q(\theta, m_{\min}), \quad (8)$$

where

$$z_n = (\alpha + 1 + n)/\beta, \quad (9)$$

$$x = \left(\frac{M_{\min}}{\mathcal{H}_*} \right)^\beta \equiv 10^{\beta(m-h_*)}, \quad (10)$$

and Γ is the Euler (upper-incomplete) gamma function. We will use z_n and x in this manner throughout the paper, noting that x_m will refer to x with m_{\min} replaced by arbitrary m , and z without a subscript refers to z_0 .

We will typically deal with the latter case in this paper, as it is the more general of the two. It should be noted that for $z > 0$, $\Gamma(z, 0) = \Gamma(z)$, so that q_0 is merely the expected special case of q , under the obvious proviso that the integral converges. Thus we need only specify results for the truncated case, and the special case can be read in.

We note that many numerical libraries only implement $\Gamma(z, x)$ for $z \geq 0$, so that α is required to be > -1 (consistent with the usual definition of the GGD).⁶ Fortunately, Γ admits a stable recurrence relation, which can be used to efficiently calculate the function with $z < 0$, in terms of only positive

⁶ Notable exceptions are the GNU scientific library (GSL), which the R TGGD package interfaces with, and the MPMATH Python library, which our MRPY library utilizes.

values, so long as $|z|$ is not too large (in our case, it should not exceed 3 very often). We present details of the mathematical derivation and numerical implementation of this relation for our work in Appendix A.

2.3. Normalization

While in principle, A can be set arbitrarily to fit given data, there are three values of interest. First, to ensure that the MRP is a valid statistical distribution defined on the support (m_{\min}, ∞) , one can set

$$A = A_1 = 1/\mathcal{H}_* \Gamma(z, x). \quad (11)$$

Alternatively, to normalize the MRP so that its magnitude is equivalent to the expected number density of objects, $\frac{dn}{dm}$, we may make the (standard halo model) assumption that all CDM density is contained within halos (including subhalos) at some scale (Cooray & Sheth 2002). In this case, we force the integral of the mass-weighted function over all masses to converge to the known matter density, $\Omega_m \rho_{\text{crit}}$. We find that

$$A = A_\rho = \Omega_m \rho_{\text{crit}} / k(\theta), \quad (12)$$

where

$$k(\theta) = \mathcal{H}_*^2 \Gamma(z_1), \quad (13)$$

which is defined for $\alpha \geq -2$, as per previous discussion.

We note that, in practice, setting the normalization in this way often leads to systematic errors with magnitude highly dependent on the value of m_{\min} , and thus is not suitable when fitting data. This problem is similar to that faced by the PS formalism (Jenkins et al. 2001).

A third value of interest arises from matching the expected total number of variates in the distribution to that observed. For N halos in a volume V , we have

$$A = A_N = A_1 N / V. \quad (14)$$

Table 1 lists some further properties and derived quantities of MRP. Each of these is defined in our accompanying python code, MRPY.

2.4. Reparameterizations

Reparameterizations of the MRP are available that can reduce the covariance between the parameters (Lawless 1982; Lagos Álvarez et al. 2011). We explore some of these additional parameterizations in more detail in Appendix B. We only comment here that while these may be beneficial in attaining parameter fits, we choose to use the parameterization in Equation (5) because the parameters θ_i have well defined characteristics in this form: A is the normalization, h_* is the turnover mass, α is the low-mass power-law slope, and β controls the sharpness of the exponential cutoff at high mass.

We note here that it is possible to formulate the MRP in terms of the logarithmic mass mode \mathcal{H}_T directly, where \mathcal{H}_T is the mass at which

$$\frac{d}{dm} m g(m|\theta) = 0, \quad (15)$$

and that this parameter is closely related to \mathcal{H}_* (the exact relation is given in Table 1). Conceptually, \mathcal{H}_T is the “typical” mass of a halo for a given distribution, and we expect it to be less correlated with the other parameters than \mathcal{H}_* . It may be beneficial in future studies to adopt this parameterization, but

Table 1
Basic Properties and Derived Quantities of MRP

Property	Value ($z \equiv z_0$)
$n(>m \theta)$	$q(\theta, m)$
$n(>m_l, <m_u \theta)$	$q(\theta, m_l) - q(\theta, m_u)$
$\rho(>m \theta)$	$A\mathcal{H}_*^2\Gamma(z_1, x)$
Mode	$\begin{cases} m_{\min} & \alpha < 0 \\ \mathcal{H}_*\left(\frac{\alpha}{\beta}\right)^{1/\beta} & \alpha > 0 \end{cases}$
n th Raw Moment	$\mathcal{H}_*^n \frac{\Gamma(z_n, x)}{\Gamma(z_0, x)}$
Mean	$\mathcal{H}_* \frac{\Gamma(z_1, x)}{\Gamma(z, x)}$
n th Central Moment	$\mathcal{H}_*^n \sum_{k=0}^n (-1)^{n-k} \binom{n}{k} \frac{\Gamma(z_1, x)^{n-k} \Gamma(z_k, x)}{\Gamma(z_0, x)^{n+1-k}}$
Variance	$\mathcal{H}_*^2 \left(\frac{\Gamma(z+2/\beta, x)}{\Gamma(z, x)} - \left(\frac{\Gamma(z+1/\beta, x)}{\Gamma(z, x)} \right)^2 \right)$
Log. Mass Mode, \mathcal{H}_T	$\mathcal{H}_* \sqrt[\beta]{z+1/\beta}$

Note. Properties and derived quantities of MRP. The mass mode is the logarithmic mass bin with the highest mass density.

we do not here for the sake of simplicity and ease of comparison to the Schechter function.

We also note that the Schechter function affords the same relation, though since $\beta = 1$ and α is generally close to -1 for LFs, \mathcal{H}_* is often very close to \mathcal{H}_T .

3. Parameter Estimation with MRP

In this section, we outline a general method for parameter estimation with the MRP. The key quantity in Bayesian estimation is the log-likelihood, which may be minimized using downhill gradient optimization methods, or traced with Markov-Chain Monte Carlo (MCMC). We follow the usual method of defining the log-likelihood common in cluster cosmology studies (e.g., Bocquet et al. 2016; de Haan et al. 2016).

A realized sample of halos will be drawn from the MRP distribution within a certain volume. In this work, we incorporate into this “volume” the effects of a selection function—that is, the volume is an effective volume at a given locus of the halo properties. The halo properties are not necessarily univariate—i.e., we may be interested in more than just mass—however, in this work, we will restrict ourselves to this univariate distribution. The effective volume is also potentially dependent on the model parameters, θ . Thus we specify the effective volume as $V(m|\theta)$, noting that the selection function for a simulation may be modeled as a simple step-function in m , which would incur the usage of $q(\theta, m_{\min})$ as the number-density cumulant.

Let \mathbf{x} be a set of observed properties of the halos—mass, velocity dispersion, luminosity, occupation number, etc.—and define $\rho(\mathbf{x}|m, \theta)$ as the probability that a halo of mass m in the model with parameters θ is observed with the properties \mathbf{x} . Given an arbitrary binning of the data, and an occupation of those bins that is independent and Poisson-distributed, then going to the limit of zero-width bins, the likelihood of the parameters given a sample of N halos is

$$\ln \mathcal{L}(\theta) = - \int V(m|\theta) \phi(m|\theta) dm + \sum_{i=1}^N \ln \int V(m|\theta) \phi(m|\theta) \rho_i(\mathbf{x}_i|m, \theta) dm. \quad (16)$$

Note that this formula does not include effects from cosmic variance or other systematic uncertainties. We briefly discuss how one may incorporate such effects in Appendix C, but we do not include them in any analyses in this work. We expect that these effects will be dominant only for deep small-area surveys (i.e., those approaching a “pencil-beam” configuration), and we focus on larger-area surveys (or indeed large-volume simulations) in our examples. Care should be taken in any application as to whether this assumption is valid.

We note a couple of special cases here as examples. First, in the case of the MRP where we directly measure the precise mass of a halo (e.g., from an N -body simulation), we have $\rho(\mathbf{x}|m\theta) = \delta(m' - m)$ (with m' the observed mass), and we are left with

$$\ln \mathcal{L}(\theta) = - \int Vg dm + \sum_{i=1}^N \ln(V(m'_i)g(m'_i)), \quad (17)$$

where we have dropped the explicit dependence on m and θ for each function for clarity, and shall do so hereafter unless it promotes ambiguity.

Second, if the selection function is truly a step function (a reasonable approximation in an N -body simulation) with a step at m_{\min} , up to volume V_0 , then we have

$$\ln \mathcal{L}(\theta) = -V_0 q(\theta, m_{\min}) + \sum_{i=1}^N \ln(V_0 g(m'_i)). \quad (18)$$

In practice, typically neither of these special cases are applicable, and the general Equation (16) must be used.

3.1. Acceleration Schemes

It has been widely realized that the bottleneck in the calculation of Equation (16) lies in the evaluation of N integrals in the second term (e.g., de Haan et al. 2016). In fact, the situation as presented here is already somewhat ideal—we have assumed that each halo mass is uncorrelated with all other halos. If this is not the case, then the N one-dimensional integrals become a single N -dimensional integral, which is clearly numerically intractable,⁷ and we shall not pursue such difficulties here.

One interesting way to mitigate this bottleneck is to follow the *fit-and-debias* method of Obreschkow et al. (2018). This method proposes a new likelihood, based on the most likely posterior mass for each observation:

$$\ln \mathcal{L}(\theta) = - \int V\phi dm + \int \ln(V\phi) \sum_{i=1}^N \bar{\rho}_i(\mathbf{x}_i) dm, \quad (19)$$

where

$$\bar{\rho}_i = \frac{V\phi\rho_i(\mathbf{x}_i)}{\int V\phi\rho_i(\mathbf{x}_i) dm}. \quad (20)$$

This likelihood has a maximum at the input parameters $\hat{\theta}$ if $\bar{\rho}$ is defined to be constant w.r.t θ , and is evaluated at $\hat{\theta}$. Thus, this

⁷ In such cases, specifying each “true” mass as a hyper-parameter and forming the model as a hierarchical Bayesian model is a useful way forward. Such models can be reasonably efficiently evaluated using advanced Monte-Carlo step methods such as Hamiltonian Monte Carlo, as implemented, for example, in the STAN language. One requires that the likelihood be analytic to use STAN, which provides another motivation for using a direct form such as the MRP.

method follows an iterative approach in which first ρ_i is evaluated at a given estimate θ_0 , then the likelihood is maximized to yield a better estimate, θ_1 , and so on until convergence.

In this approach, the full N integrals need only be calculated once every method iteration, rather than on every likelihood evaluation. This significantly speeds up calculations for which these integrals are the bottleneck.

One must be aware, however, that the likelihood used here, while having the same maximum point as Equation (16), is not the same likelihood. Thus, the parameter errors derived from the likelihood are not necessarily accurate, and one should resort to the proper likelihood to evaluate these.

3.2. Expected Parameter Covariance

The errors on parameter estimates, given a sample of data, are most robustly evaluated using the chains from an MCMC sampling of the posterior. However, a more efficient estimate of the covariance of parameters, assuming a Gaussian posterior at the maximum likelihood estimate (MLE), can be gained from the Hessian, as $-(H(\theta))^{-1}$. Furthermore, in the process of model-building or survey-construction, one may be more interested in the *expected* covariance on their parameters, rather than the covariance actually obtained from a given set of data.

The expected Hessian gained from sampling a model ϕ with volume V and uncertainties ρ is (see Appendix D.3)

$$H(\theta) = - \int \frac{\frac{dn(x)}{d\theta_i} \Big|_{\hat{\theta}} \frac{dn(x)}{d\theta_j} \Big|_{\hat{\theta}}}{n(x|\hat{\theta})} dx, \quad (21)$$

where $n(x) = \int V\phi\rho(x)dm$.

This expression for the Hessian is beneficial, as it involves only first derivatives of $n(x)$, which improves numerical efficiency and accuracy over second derivatives. Interestingly, Equation (21) is analytically solvable in the simplest case of the MRP with precisely observed masses. We provide this solution in Appendix D.3.3.

4. Accuracy of the MRP

The Schechter function is (usually) a reasonable fit to luminosity or stellar-mass data, up to the uncertainty in the data itself. In the case of the MRP, we have a wealth of information, in the form of first-principles simulations and fits to such simulations using the PS formalism. Thus, our first task is to measure the accuracy of the MRP with respect to the precision of our current knowledge of the theoretical HMF. Clearly, the MRP will be an imperfect model; it will lose information in the reduction of a simulation to five parameters. However, we wish to test whether the MRP is an adequate approximation to theoretical HMFs up to the uncertainty both inherent in the theory, and in relevant data sets.

We approach this in two steps. First, we perform simple comparisons to HMF fits using the PS formalism, allowing us to quickly explore parameter space and assess the accuracy of MRP in different regimes. Second, we compare the MRP to the output of a full N -body simulation, comparing the residuals to those from the best fit using the PS formalism. Our aim is not to show that the MRP performs better than the PS formalism, but that it performs adequately for use with real data.

Throughout this paper, we adopt the fiducial cosmology of Planck Collaboration (2014; hereafter P13).⁸ Likewise, we adopt the fit of Tinker et al. (2008, hereafter T08) as our fiducial mass function.

4.1. Comparison to PS-type Fits

The fit of T08 is very commonly employed in observational studies (e.g., Pacaud et al. 2016; Planck Collaboration et al. 2016), since it provides an accurate fit over redshift (up to $z=2$) and also for varying spherical overdensities. Using a simple χ^2 -minimization in log-log space, we fit T08 for a range of redshifts and halo overdensities, as well as varying lower mass limits. The mass range for each is calculated between constant limits on the mass variance, so as to probe the interesting part of the HMF for each combination of parameters. The limits are $\sigma_{\max} = [4, 3, 2, 1.5]$ and $\sigma_{\min} = 0.5$. The T08 HMFs were produced with the HMF Python code, v3.0.2 (Murray et al. 2013). Results are plotted in Figure 1.

The goodness-of-fit of the MRP varies across redshift and overdensity, but overall even in the worst-case scenario the rms deviation over the fitted mass range is about 2% (for $z=0$ and $\Delta_h = 200$, where Δ_h is the overdensity with respect to the mean matter density ρ_m). The approximate uncertainty in T08 is reported as 5%, so in a global sense, we find the MRP to be a good approximation.

The variation of the residuals with redshift is such that they generally tighten toward higher redshift. This can be ascribed to the fact that the mass function turnover becomes less marked at these higher redshifts—the function is close to a steep power law. This is seen by the background gray curves that represent log-scale HMFs of arbitrary units, to which each MRP was fit. Since the transition from power law to exponential cutoff is the most difficult regime for the MRP to model, it excels when this transition is muted. An increased overdensity criterion adversely affects the residuals at high redshift, but has the opposite effect at low redshift. This can be attributed to the fact that a change in overdensity is predominantly a nonlinear shift of the HMF on the mass-axis. The primary resulting change in the HMF is the power-law slope at low redshift, but the position of the cutoff at high redshift. Thus the explanation for the increased residuals is similar to that of its variation with redshift. In effect, the MRP performs most poorly around the transition from power law to exponential cutoff, but is still typically within 5% over a large parameter space. In addition, it is difficult to ascertain whether the residuals above the transition are robust, since the T08 fit itself is more poorly constrained on these scales due to the inherent Poisson noise and cosmic variance in simulations.

In order to assess whether the MRP form is flexible enough to accommodate various cosmologies and halo finders, we provide Figure 2, which is very similar to Figure 1, except that here the fitting function is that of Warren et al. (2006)—a form fit to friends-of-friends (FOF) halos—and the columns and rows represent Ω_m and σ_8 respectively. We first find that large changes in cosmology have little effect on the ability of the MRP to describe the HMF—the difference in rms between all curves over the same range is less than 1%. Second, we find that the fit to an FOF-based HMF is slightly worse than that to

⁸ Results for the more recent 2015 data release are almost identical, and the choice of P13 is driven by its use in simulations in Section 4.2.

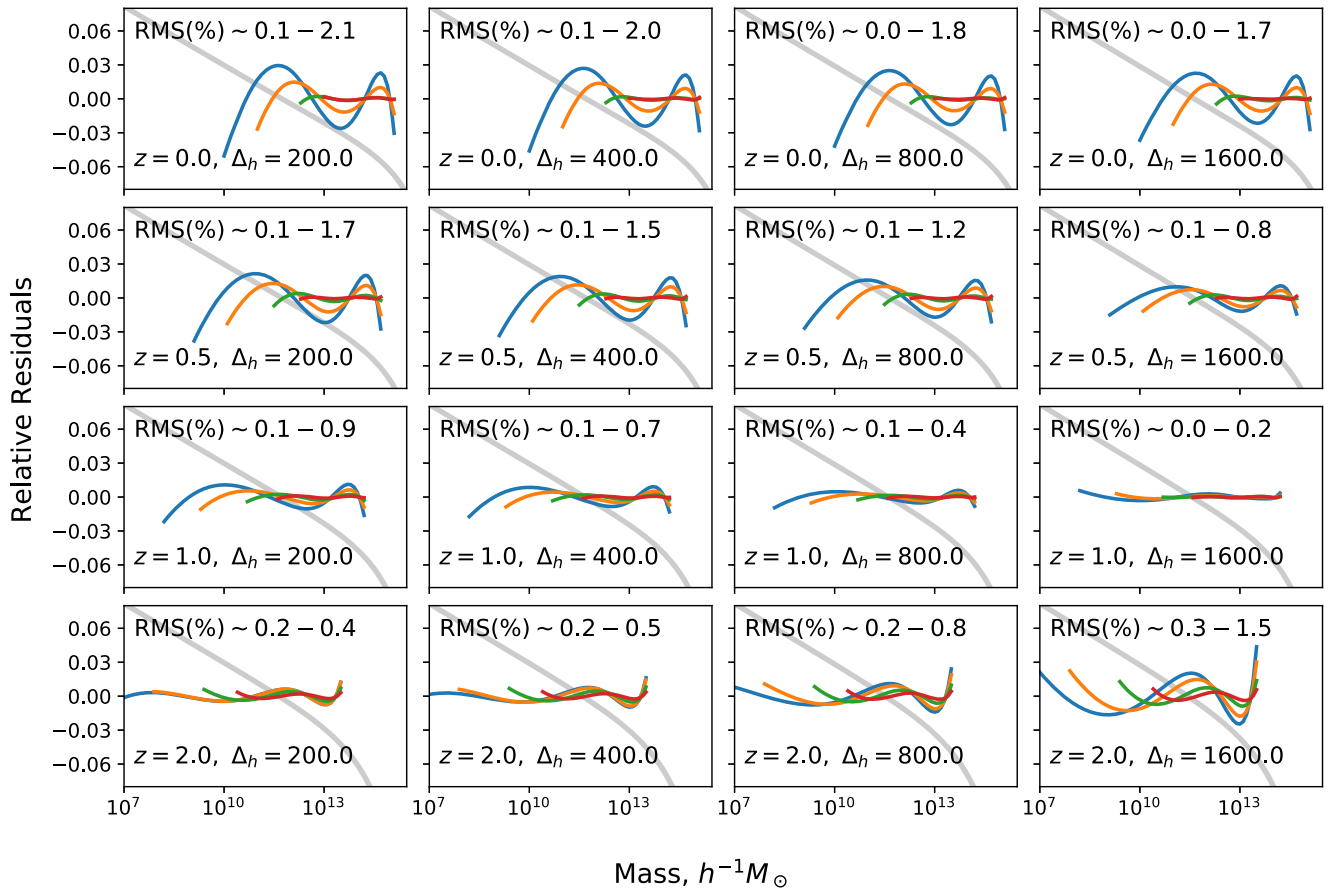


Figure 1. Comparison of MRP form to T08 over several redshifts and overdensities. Each column represents an overdensity corresponding to those reported in T08. Each row represents a redshift, up to $z = 2$. The four lines in each represent a fit over a different mass range. The gray curves in the background of each panel show the log-scale HMF in arbitrary units, in order to show the kind of shape that the MRP was fit to in each case. The rms deviation as a percentage is shown for the best and worst case in each panel. The rms is better than 2% over all overdensities for $z \leq 1$, and is a maximum of 2% for $z = 0$ and $\Delta_h = 200$. Expectedly, residuals are worse for fits over wider mass ranges.

the SO-based T08 function, but only by $\sim 0.5\%$ (we compare to the top-left panel of Figure 1, which should be the most similar to the curves in this figure). Thus we conclude that the MRP is flexible enough to fit HMFs from various cosmologies, redshifts, halo-finders, and halo-definitions with roughly equivalent accuracy—and that to within the quoted uncertainties from PS-type fits.

The most concerning feature of the residuals is an oscillatory pattern in the power-law region of the distribution. This oscillatory pattern is expanded as the lower mass limit is decreased, indicating that it is a model deficiency. We do note, however, that similar oscillatory residuals are seen in standard PS formalism fits to N -body simulations. It is of interest to what extent this feature is a loss of information in MRP as compared to the true physical distribution. This question is best answered by observing residuals of fits to simulations, to which we now turn.

4.2. Comparison to N -body Simulations

The simplest realistic halo catalogs arise from DM-only simulations. Such catalogs contain high-resolution, fully non-linear determinations of the HMF, completely independent of the PS framework assumptions. This enables comparisons of the fidelity of the MRP versus the PS formalism without any possible biases from the PS modeling itself.

In this study, we use the publicly available halos from the New Numerical Galaxy Catalog (ν^2 GC, Ishiyama et al. 2015; hereafter I15). These simulations were run with the P13 cosmology, in several box-sizes and mass resolutions. Specifically, we use ν^2 GC-M and ν^2 GC-H1, which when taken together probe a volume of $560 h^{-3} \text{Mpc}^3$ down to halos of mass $2.75 \times 10^9 h^{-1} M_{\odot}$. As such, it spans the halo mass range relevant for group catalog studies. Details of the simulations are found in Table 2. In particular, we note that though the original catalogs contain halos defined down to 40 particles, we fit only to those with 100 particles or more, to alleviate the effect of Poisson noise on the masses of these highly influential halos. Furthermore, the box size of ν^2 GC-M is small enough to make the highest mass objects quite rare, and thus dominated by cosmic variance and Poisson noise. We set upper limits of $7 \times 10^{14} h^{-1} M_{\odot}$ and $2 \times 10^{13} h^{-1} M_{\odot}$ on the mass ranges respectively.

In order to test the robustness of the MRP to different halo definitions, we use both FOF and Spherical-Overdensity (SO) halos identified in each simulation. The SO halos are defined via the critical overdensity, $\Delta_h = 200\rho_c$ using the ROCKSTAR⁹ code (Behroozi et al. 2013b). Note that ROCKSTAR is a 6D FOF halo-finder, but does produce SO halos as part of its output.

⁹ <https://bitbucket.org/gcfstanford/rockstar>

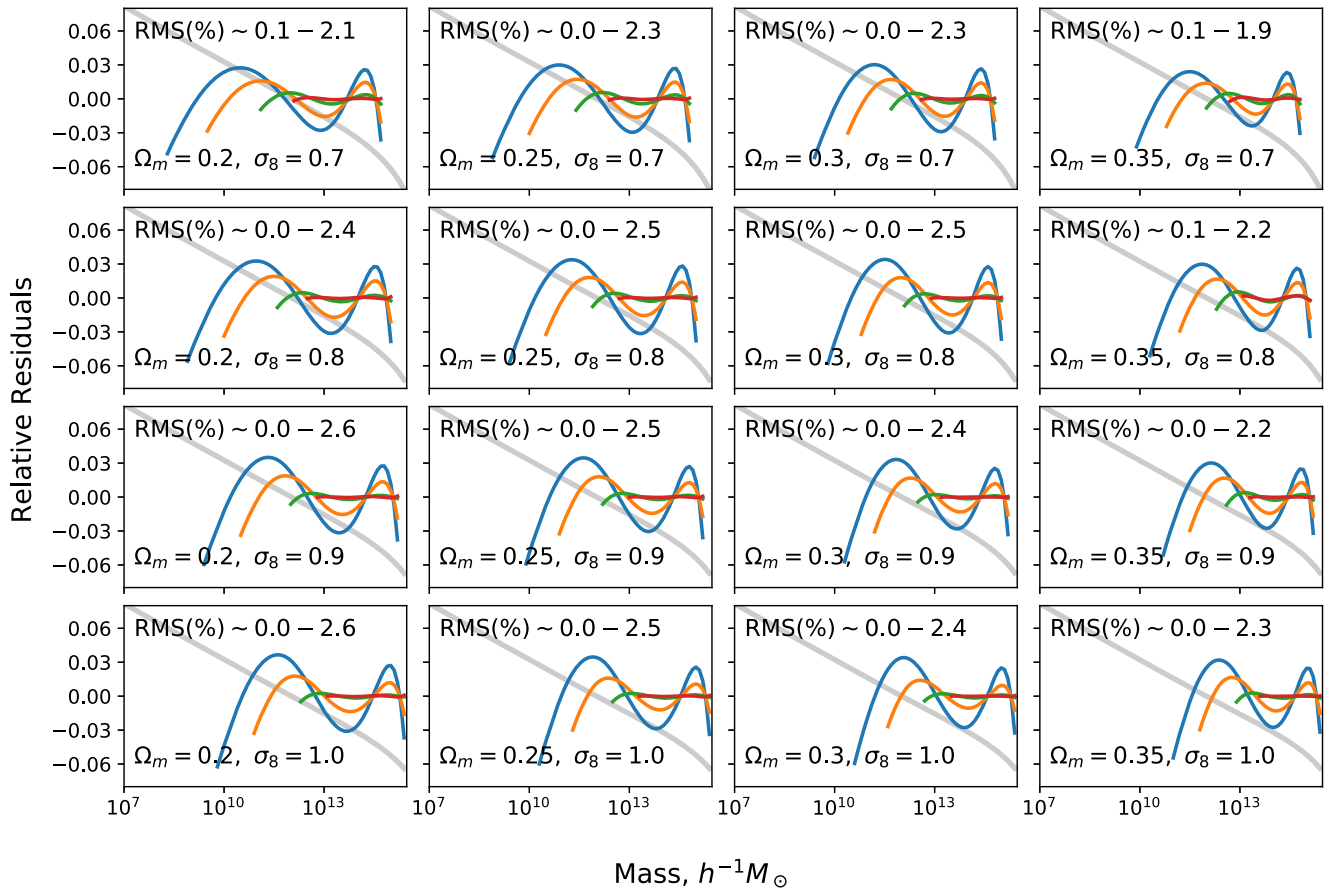


Figure 2. Comparison of MRP form to the fit of Warren et al. (2006) over several values of Ω_m and σ_8 . Everything is the same as in Figure 1, except that here each column represents a value of Ω_m , and each row a value of σ_8 . The rms is better than $\sim 2.6\%$ over all parameters, and is a maximum of 2.6% for $\Omega_m = 0.2$ and $\sigma_8 = 1.0$.

We estimate the parameters of the MRP via the likelihood in Equation (18). However, some care must be taken when simultaneously analyzing multiple data sets with different volumes. In this case, we calculate the log-likelihood for each data set with the same shape parameters, but altering the normalization so that it is relatively correct for each data set. More specifically, given n data sets, each with N_j halos in a volume V_j , we use

$$\ln \mathcal{L} = \sum_{j=1}^n -V_j q(\boldsymbol{\theta}, m_{\min, j}) + N_j \ln V_j + \sum_{i=1}^{N_j} \ln g(\boldsymbol{\theta}, m_i). \quad (22)$$

In our case, this means that over most scales, the ν^2 GC-M box is highly weighted, but at the smallest scales, which it does not probe, the H1 box is effective. In general, since we have manually truncated the distribution at the high-mass end, we should amend q so that it is correctly calculated between the mass limits. However, due to the rapid decline of the distribution, such an amendment has a very small effect, and in this case we neglect it.

The parameter estimation can be performed either with simple likelihood maximization (via downhill-gradient methods), or MCMC. For illustration purposes, we use MCMC in this case. The MCMC sampling used the EMCEE¹⁰ python package (Foreman-Mackey et al. 2012), in which we used 50 “walkers” each for 200 steps of burn-in, and 500 retained steps,

for a total of 25,000 samples. We begin the chains in a small ball around the optimization solution, as recommended in Foreman-Mackey et al. (2012). We verify convergence with the Gelman–Rubin test (Gelman & Rubin 1992), ensuring that the *potential scale reduction factor* $\hat{R} < 1.1$ for all parameters.

We plot the posterior joint-likelihoods of the parameters, for the FOF halos, in Figure 3. We stress that this is a high-resolution run, which will be poorly represented by realistic observed data sets. With that in mind, the joint posterior is a very good approximation to a multivariate normal distribution, which has the benefit of allowing for quick analyses by downhill-gradient methods, combined with analytic determination of the distribution around the solution via the Hessian (see Appendix D.2 for the exact equation for this quantity). There is a high degree of covariance, as expected, between $\ln A$ and h_* .

Because of the high volume of data, the parameters are extremely well constrained, but their precision is not necessarily the best metric of performance. Since the underlying data does not perfectly trace the MRP form, the parameters are likely to be biased, and it is the residuals that are the most telling. The turnover-mass is around $10^{14.52} h^{-1} M_\odot$, and the mass mode is $\sim 10^{13.62} h^{-1} M_\odot$. We find a power-law slope of -1.9 . We also find a cutoff parameter β of 1.19, which is significantly sharper than a Schechter function (in which $\beta \equiv 1$). The log-normalization is -44.4 . Of these, we expect the value of β to be the most inaccurate, as it is principally determined by the high-mass halos, of which there are very few.

¹⁰ <http://dan.iel.fm/emcee/current/>

Table 2
Details of Halo Abundances in the ν^2 GC Boxes Used in This Study

Name	N	L , [h^{-1} Mpc]	m_p , [$h^{-1} M_\odot$]	# Halos	M_{\min} , [$h^{-1} M_\odot$]	M_{\max} , [$h^{-1} M_\odot$]
ν^2 GC-M	4096 ³	560	2.2×10^8	53×10^6	2.2×10^{10}	7×10^{14}
ν^2 GC-H1	2048 ³	140	2.75×10^7	5.5×10^6	2.75×10^9	2×10^{13}

Note. M_{\min} , as given by the public catalogs, is set by a minimum of 40 particles per halo. For our fits, we prefer to use a minimum of 100 particles per halo. Note that this changes the number of halos that are useable for our fits as compared to the number shown here. The SO halos have exactly the same parameters except number of halos.

Figure 4 shows the residuals of this fit as thick solid lines (different colors representing different simulation boxes). The left panel shows the fit to the FOF halos, while the right panel shows the fit to the SO halos. In each case, the figure shows the naively binned HMF from the simulation as a ratio to the best fit.

In addition, the right panel shows the binned FOF HMF as a ratio to the SO fit in gray, to illustrate the magnitude of the difference between the halo-finding schemes. The MRP fit is accurate to within 5% over the mass range $10^{9.7}$ – $10^{14} h^{-1} M_\odot$, which is a far smaller margin than the discrepancy between halo finders. This indicates that an arbitrary choice of halo definition is a greater loss of information than the approximation introduced by the MRP.

Also plotted in Figure 4 are the residuals of the best-fit PS-type mass function from I15 (dashed lines), which was calibrated with the FOF halos (and therefore is only shown in the left panel). The residuals are comparable in magnitude over much of the mass range, with the PS-type function providing a tighter fit at $\sim 10^{12}$ – $10^{13} h^{-1} M_\odot$. At lower masses, the MRP is much more accurate than the PS-type fit. The reason for this difference in accuracy at different scales arises from at least two considerations. First, the MRP (and the PS HMF) are not perfect representations of the HMF, and therefore cannot capture the behavior on all scales. Second, while it is likely that the I15 form was fit to binned counts in log–log-space, and therefore treats all scales equally, our likelihood implicitly fits the HMF in real-space, so that the low-mass halos dominate the fit. Whether one fit is “better” than the other is therefore based on the application at hand. Typically, for plots such as Figure 4, fits in log–log-space will appear more accurate over a larger range.

The fit to the SO halos performs relatively poorly, and this is due to the inconsistency between the two catalogs involved, as can be clearly seen in Figure 4. It is unclear whether this is due to sample variance, or a miscalibration in the SO algorithm, but the results from these catalogs should be treated with caution.

The oscillation that was apparent in the fits to the PS formalism HMFs is still present, indicating that the MRP is not flexible enough to capture the details of the HMF distribution. However, at some level this is clearly also true of the PS formalism, which systematically overpredicts the low-mass end and underpredicts the high-mass end of the HMF in this simulation. The MRP is expected to perform more poorly than the PS formalism when the mass range is increased. This must be remembered when using the MRP for any fit—the truncation scale(s) are an inherent part of the formula, and will to some extent affect the estimation of the other parameters.

Nevertheless, we see that MRP is able to empirically model the HMF over a sufficiently large mass range to be relevant for group catalogs.

5. Dependence on Physical Parameters

While it is envisaged that the MRP will generally be used as a standalone description of halo mass data, it is instructive to determine the dependence of the parameters $\theta \equiv (\mathcal{H}_*, \alpha, \beta, \ln A)$ on the more important physical parameters.

The parameter with the most influence over the HMF is the redshift, which through the growth rate affects the normalization of the power spectrum, and is also explicitly modeled into the HMF fit in several recent studies (e.g., our fiducial model of T08). Along with this, ? showed that the parameters Ω_m and σ_8 have the most significant effects across a broad range of masses. Thus, in this section, we set out to model the dependence of the θ_i on $\phi = (\Omega_m, z, \sigma_8)$.

We note that the choice of base model to fit to is somewhat arbitrary. Fits in the literature vary at the level of tens of percent across interesting mass ranges. Furthermore, there is a dichotomy between fits based on halos found using FOF and SO methods, as we have seen. Given the inherently approximate nature of this dependency modeling, we opt to fit to a chosen fiducial model, and use the uncertainty of this fit (both intrinsically and with respect to other fits) to guide our judgment of the appropriateness of the model.

Since we are interested in modeling both low and high redshifts, we take as our fiducial fit the form of Behroozi et al. (2013a; hereafter B13), which is an empirical modification of T08 that increases accuracy at high redshift (up to the Epoch of Reionization).

When calculating the HMFs for modeling the MRP parameters, we ensure that in each case the wavenumber range is wide enough to encapsulate all relevant information (see Section 2 of Murray et al. 2013 for a brief discussion of relevant limits on the product of the radii and wavenumber), and also that the resolution is high enough to avoid small oscillatory artifacts.

An important consideration is the mass range over which to fit the MRP, as this choice affects the derived parameters (see Figure 1). In order to probe the domains of influence of all three shape parameters, and to provide a standardized system, we opt to define the mass range as constant with respect to the logarithmic mass mode \mathcal{H}_T . This value is calculated explicitly as the zero of the derivative of a quartic spline interpolation of the B13 HMF for each value in the sample of ϕ . Figure 5 shows \mathcal{H}_T as a function of redshift for fiducial values of Ω_m and σ_8 (the dependence on these latter parameters is comparatively negligible). We have verified that \mathcal{H}_T is within the valid mass range of B13 for all samples.

It must be urged that the fitted parameters can change significantly given a different truncation mass with all other parameters fixed. Specifically, if a smaller mass range is of interest, a more accurate fit can generally be attained (see Figure 1). Furthermore, an optimization of the MRP parameters for any input model will achieve a greater accuracy than using

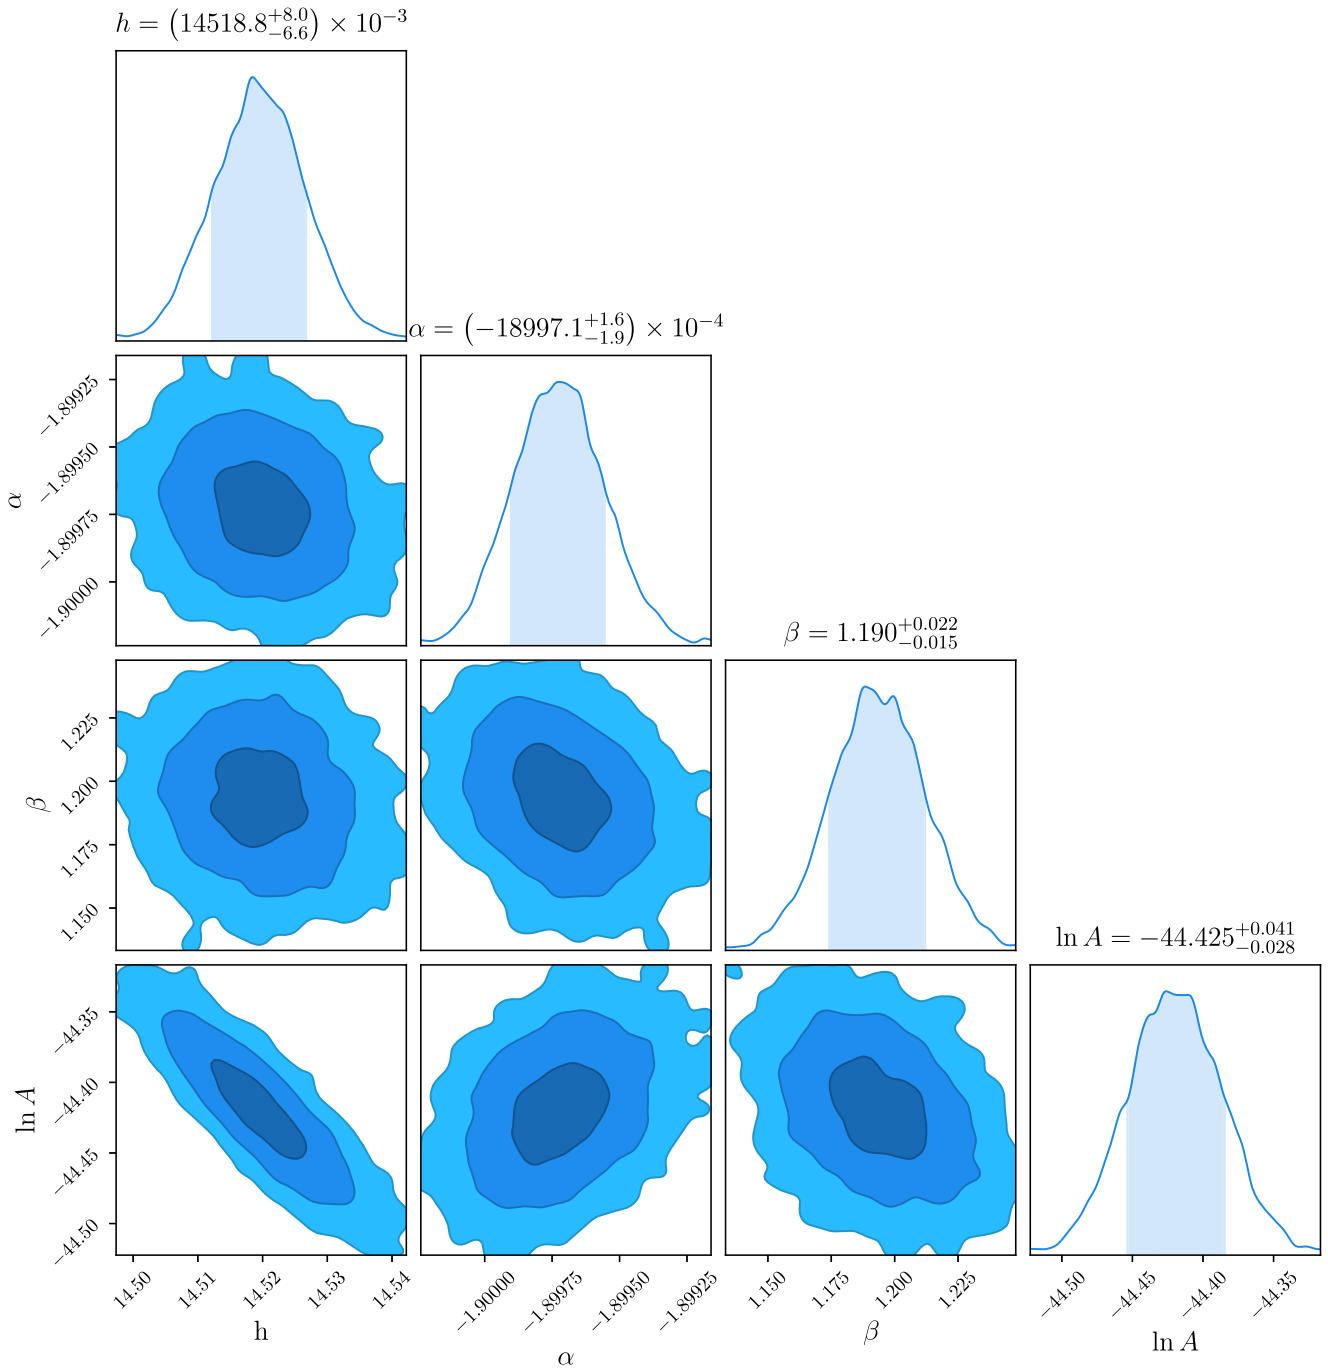


Figure 3. Posterior joint-likelihoods of the MRP parameters estimated simultaneously on both ν^2 GC simulations, using FOF halos. All contours display high degrees of normality. The distributions are rather covariant, which can pose problems for standard sampling techniques.

the models presented in this section (and doing so is not computationally expensive). However, the most important benefit of this section is the insight gained from determining the approximate relationship between physical and MRP parameters, and that remains relatively consistent with variable m_{\min} .

We produce 2000 samples, ϕ_i . The redshifts are drawn uniformly in log-space, so as to lightly increase the weight of lower redshifts. The cosmological parameters are drawn from the normal distributions presented in P13, so as to highly weight values that are more likely to be physically appropriate.

In summary,

$$\begin{aligned} \log(1+z) &\sim U(0, \log 9) \\ \sigma_8 &\sim \mathcal{N}(0.829, 0.012) \\ \Omega_m &\sim \mathcal{N}(0.315, 0.017). \end{aligned}$$

A value θ_i is fit to each model using χ^2 -minimization in log-log space (identical to the procedure for producing Figure 1). To find an optimal parameterization for θ as a function of ϕ , we use the symbolic regression program EUREQA¹¹ (Schmidt &

¹¹ <http://www.nutonian.com/products/eureqa-desktop/>

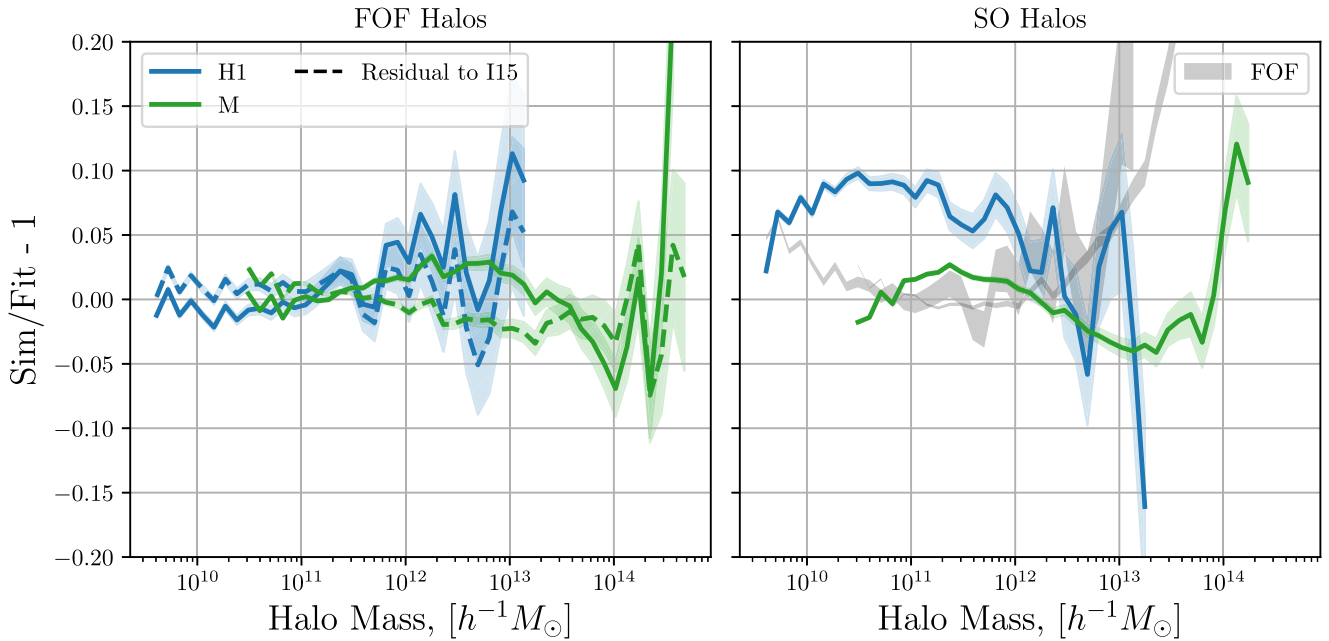


Figure 4. Residuals of fits to ν^2 GC simulations. The solid lines show residuals from the suites of halos against the global MRP fit. The dashed lines show the I15 residuals to the same halos. The left panel shows the FOF halo fit, while the right panel is normalized to the SO halo fit. Shaded regions indicate the 1σ uncertainty regions generated from the Poisson noise estimate. The gray shaded regions in the right-hand panel are the naive FOF binned HMFs normalized to the SO fit, to illustrate the magnitude of the difference between halo-finders.




Figure 5. Logarithmic mass mode \mathcal{H}_r as a function of redshift for the B13 HMF.

Lipson 2009). This program iterates over models using a genetic algorithm, with constraints set by the user, to find underlying trends in input data. To assist the convergence of the program, we first fit a one-dimensional model in redshift only, which is the most effective parameter. Using this model as a baseline, we introduce Ω_m and σ_8 for a full 3D parameterization. Each run of EUREQA produces a series of potential parameterizations of varying complexity and we choose models that offer a balance between accuracy and simplicity.

We present our chosen models in Table 3. Each model is normalized and scaled to its mean and standard deviation, highlighting the primary dependencies among the parameters. A more intuitive presentation of the models can be found in Figure 6, in which the univariate dependencies of θ_i on ϕ_j are shown. Here, σ_8 and Ω_m are varied at $z = 0$, and otherwise take their fiducial values.

In particular, we note that across the board, the most effective parameter is z , with which all three MRP parameters are anti-correlated. The MRP parameter most sensitive to the physical parameters is β , followed by \mathcal{H}_* , with α changing very slowly (mainly with redshift).

The visually poor performance of the models over independent cosmological parameters is due to the low effectiveness of these parameters. In particular, the normalization has no dependence on Ω_m at $z = 0$, which is clearly not the case in detail. This suggests that the “automatic” fits from EUREQA can be improved.

Summarily, the redshift dependence of the parameters indicates a crossover between two behaviors. At low redshifts, there are two regimes—a low-mass power-law tail, and a high-mass cutoff region. At higher redshifts, the high-mass cutoff moves to become so far beyond \mathcal{H}_r that the function resembles a power-law over all relevant scales.

Figure 7 shows the 68% region of relative uncertainties in the models, over the distribution of physical parameters we have employed. The MRP model attained by simply using the formulas from Table 3 produces HMFs within 5% of B13 for almost the entire mass range between $0 \leq z \leq 1$, but at high redshifts can deviate by up to $\sim 10\%$. Given that uncertainties within the B13 HMF itself are of the order of 5%–10%, in addition to uncertainties due to cosmology, choice of fitting function and halo finder, this is accurate enough to trust our general results.

6. The SMHM Relation

The SMHM relation connects galaxies to halos by relating the typical stellar-mass content of a halo to its total mass. Determining the form of this relation has been the subject of numerous studies (e.g., Behroozi et al. 2010; Moster et al. 2010; Mutch et al. 2013; Durkalec et al. 2015) and is important for quantifying the effects of physical processes such as

Table 3
Table Showing Formulae for the Physical Dependence of the MRP Parameters

Parameter	Formula, $\mu + \sigma f$	μ	σ
h_*	$0.058562 + 1.4394\sigma_8 + 0.39111\Omega_m + 0.11159\sigma_8 z$ $+ 0.056010z^2 + 0.42444\sigma_8\Omega_m z - 0.90369z - 0.0029417z^3$	12.214	1.6385
α	$2.6172\Omega_m + 2.06023\sigma_8 + 1.4791 \times 2.2142^{\Omega_m} 0.53400z - 2.7098$ $- 0.19690z$	-1.9097	0.026906
β	$7.5217\sigma_8\Omega_m - 0.18866 - 0.36891z - 0.071716 \times 0.0029092z$ $- 3.4453\Omega_m z 0.71052z$	0.49961	0.12913
$\ln A$	$z + 0.0029187z^3 - 0.11541 - 1.4657\sigma_8 - 0.055025z^2 - 0.24068\sigma_8 z$ $- 0.33620\Omega_m z$	-33.268	7.3593

Note. Parameterizations of θ_i as functions of ϕ . All parameterizations are expressed as rescaled by the mean and standard deviation of the input data, which highlights the typical values and associated sensitivity of each MRP parameter.




Figure 6. Dependence of θ_i (rows) on ϕ (columns). The bottom row shows the logarithmic mass mode, \mathcal{H}_T . Blue markers show best-fit values derived from χ^2 -minimization on the B13 HMF. Red lines show the EUREQA-derived models from Table 3. The lower effectiveness of the cosmological parameters renders their independent models visually poorer, especially in the case of the normalization.




Figure 7. Central 68% regions of the relative error for the 3D parameterizations of the MRP parameters compared to the fits of B13. The x -axis is the mass normalized to \mathcal{H}_T , in which each sample has the same range. Each panel indicates a different redshift range, as marked. The solid line is the median value of each subsample. The uncertainty in the models sub-10% over the entire redshift and mass range.

supernova and AGN heating. Given that the fundamental motivation for defining the MRP form is to enable simpler connections between galaxies and halos, in this section, we

construct a direct formula that ties together the MRP and a double-Schechter parameterization of the GSMF. In this simplistic setting, we can derive accurate formulae involving both MRP and Schechter parameters, which serve to illustrate how a direct parameterization of the HMF can be conceptually useful.

One simple method of estimating the SMHM relation is to use observed GSMFs, along with theoretical HMFs. If using a subhalo-calibrated HMF, there must be at most one galaxy per halo. It is generally also assumed that each (sub)halo contains a galaxy, so that each halo contains exactly one galaxy. If we also assume that the relationship between a halo mass and its corresponding stellar mass is both deterministic and monotonic (an unrealistic assumption, though studies suggest that the scatter in the SMHM relation is small and constant over mass), then we may derive a relationship between m_* and m_h by

$$n_g(>m_*) = n_h(>m_h). \quad (23)$$

Typically, $n_h(>m_h)$ is derived by the PS formalism, and thus can be treated somewhat as a black box in this equation. Though the equation can be solved numerically for $m_*(m_h)$ given any observed GSMF, it is convenient to be able to solve it, at least approximately, analytically in terms of the parameters of the underlying GSMF and the MRP.

The GSMF is commonly accurately parameterized as a double-Schechter function (Baldry et al. 2012), which has the integral

$$n_g(>m_*) = \Phi_1 \Gamma\left(\alpha_1 + 1, \frac{M_*}{\mathcal{M}_*}\right) + \Phi_2 \Gamma\left(\alpha_2 + 1, \frac{M_*}{\mathcal{M}_*}\right), \quad (24)$$

where, without loss of generality, we will assume that $\alpha_2 < \alpha_1$. If we let the HMF be in the form of the MRP, we then arrive at the expression

$$\Phi_1 \Gamma\left(\alpha_1 + 1, \frac{M_*}{\mathcal{M}_*}\right) + \Phi_2 \Gamma\left(\alpha_2 + 1, \frac{M_*}{\mathcal{M}_*}\right) = A \mathcal{H}_* \Gamma(z_h, x_h). \quad (25)$$

Unfortunately, this is analytically unsolvable. However, it is easy to solve by root-finding, as it has a single root and simple derivatives. Thus a numerical procedure is able to produce the empirical SMHM relation.

This procedure is illustrated in Figure 8, where the cumulative HMF is shown in blue, and the cumulative stellar-mass function in orange. Given Equation (25), the

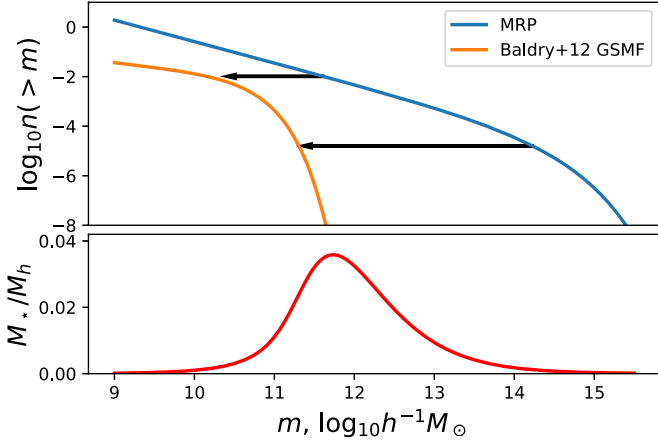


Figure 8. Upper panel: the cumulative MRP corresponding to T08 at $z = 0$ in blue, and the cumulative double-Schechter parameterization of the $z = 0$ GSMF from Baldry et al. (2012) in orange. Identifying the SMHM relation equates to computing the transformation of halo mass to stellar mass, as shown by the arrows. Lower panel: the numerically derived fraction M_*/M_h from the functions in the top panel. There is a strong peak efficiency at $M_h \sim 10^{11.8} h^{-1} M_\odot$.

procedure to calculate the ratio M_*/M_h is to identify the ratio of the curves in the horizontal direction. This can be simply performed by splines, or more robustly by root-finding, to find the red curve in the lower panel. This curve displays the efficiency of producing stellar mass as a function of halo mass, and exhibits a characteristic peak around $10^{12} h^{-1} M_\odot$.

It is convenient to have a closed form for the SMHM relation (or equivalently $f = M_*/M_h$). Several have been proposed in the literature, each of which lies at a different point on the spectrum of simplicity versus flexibility. We list three such parameterizations in Table 4, and show an instance of each in Figure 9, where each is fit against the numerically calculated black curve (we note that this curve is calculated under all the assumptions of strict halo abundance matching as described above, and the goodness-of-fit of these curves is not an indication of their accuracy to observed data or their general usefulness). In this section, we add our own parameterization, based on solving Equation (25).

Though we cannot solve Equation (25) for all m_h , we do well to identify its behavior in the limits. As $x \rightarrow 0$, we have the identity $\Gamma(z, x) \rightarrow -x^z/z$, for $z < 0$ (this can be simply derived from the recursion relation, Equation (46)). The speed of convergence of this limit depends heavily on the value of z (the more negative, the faster the convergence).

In this case, at small mass, the α_2 term dominates the α_1 term (which generally has a positive shape parameter), so we can simply write

$$-\frac{\Phi_2\left(\frac{M_*}{M_*}\right)^{(\alpha_2+1)}}{\alpha_2+1} = -\frac{A\mathcal{H}_*x_h^{z_h}}{z_h}, \quad (26)$$

so that we find that the ratio is a power law:

$$m_* - m_h = \log_{10} K + pm_h, \quad (27)$$

$$p = \frac{\alpha_h - \alpha_2}{\alpha_2 + 1} \quad (28)$$




Figure 9. Numerically calculated SMHM relation (black, largely hidden by thick red line except just below the turning point), with three best-fit models from the literature. The lower panel shows their ratios. All fits are in log-log space.

Table 4
Summary of SMHM Relations in the Literature

Ref.	# par.	Formula
M13	3	$\epsilon \exp\left(-\left[\frac{m_h - m_{\text{peak}}}{\sigma}\right]^2\right)$
M10	4	$2f_0 \left[\left(\frac{M_h}{M_1}\right)^{-\beta} + \left(\frac{M_h}{M_1}\right)^\gamma \right]^{-1}$
B10	5	$\frac{m_*}{M_0 \left(\frac{M_*}{M_1}\right)^\beta 10 \left(\frac{M_*}{M_1}\right)^\delta / \left(1 + \left(\frac{M_*}{M_1}\right)^{-\gamma}\right)^{-1/2}}$

Note. Parameterizations of several SMHM relations from the literature. Note that B10 is in terms of M_* rather than M_h , and so must be numerically inverted (except in limiting cases where it can be analytically inverted). Note also that parameters may have different meanings across parameterizations.

$$\log_{10} K = m_* - \frac{\alpha_h}{\alpha_2 + 1} h_* + \frac{1}{\alpha_2 + 1} \log_{10} \left(\frac{A(\alpha_2 + 1)}{\Phi_2 z_h} \right), \quad (29)$$

Thus we favor a parameterization that is a power law in the lower limit, such as M10 or B10 in Table 4. Specifically, M10 is equivalent in the low-mass limit if $\beta = p$ and $2f_0/M_1^p = K$, while B10 requires $\beta = 1/(p + 1)$ and $M_1(\sqrt{10}/M_0)^{1/\beta} = K$.

Unfortunately, while the high-mass limit affords an analytic asymptotic approximation, we are more interested in scales $m_*/M_* \lesssim 10$ (higher masses are almost nonexistent). We have not been able to derive an approximation in this regime. However, the power-law behavior exemplified by M10 and B10 seems sufficient.

The most severe deficiency remaining is the position of the turning point. A simple way to change the position of the turning point without affecting the behavior in the limits, is to use an extension of M10:

$$\frac{M_*}{M_h} = \frac{w \left(\frac{M_h}{M_1}\right)^{-\delta} + 2f_0}{\left(\frac{M_h}{M_1}\right)^{-\beta} + w \left(\frac{M_h}{M_1}\right)^\gamma + k}, \quad (30)$$

where w controls the peak position, k is able to correct the amplitude of the high-mass power law, and δ adds the flexibility needed to induce an upturn left of the turning point.

This extension obeys the same relations as M10 in terms of the low-mass approximation, but has the difference that the high-mass power law has the slope $\delta - \gamma$.

Using the two fixed parameters from these relations, the resultant best fit is plotted as the red dashed curve in Figure 9 (“TPL,” for “triple power-law”). We find much better agreement in terms of the turning point position, though the precise shape of the maximum is clearly still not reproduced. With the restrictions enforced, the model has the same number of parameters as the B10 form.

7. An Application: Survey Design

In this section, we use many of the results of the previous sections in a survey design application. In particular, we determine a semi-realistic model for the observations of a galaxy survey that determines group masses dynamically. While it is clear that this model ignores some of the complications associated with a realistic galaxy survey, it nonetheless attempts to capture some of this complexity, and show how the MRP in conjunction with the robust Bayesian parameter estimation of Section 3 can be used in realistic situations.

Our key question in this section will be: for a galaxy survey within a reasonably thin redshift slice with a fixed total number of observed galaxies, are the MRP parameters better fit by (i) increasing the sky area, (ii) probing fainter objects, or (iii) making more aggressive quality cuts?

Our broad approach is to first define a model for the underlying distribution of halo masses, their observables, and their measurement uncertainties. This forward model in turn provides the relevant functions for use in Equation (16) to determine the likelihood of a given set of data. Instead of producing synthetic data and generating parameter estimates for it, we will focus on the expected covariance of parameters under our data model (see Section 3.2).

7.1. Data Model

We will here exclusively consider an observation within a solid angle Ω , with no angular selection function, and within a small redshift slice, $z \in (z_0, z_1)$. Specifically, the width of the redshift slice is assumed to admit negligible evolution in the HMF parameters θ . The (mean) redshift of the observation is thus taken to be

$$z = 0.75 \frac{z_1^4 - z_0^4}{z_1^3 - z_0^3} \quad (31)$$

and the total volume is

$$V_0 = \frac{\Omega}{4\pi \text{sr}} [V_c(z_1) - V_c(z_0)], \quad (32)$$

where V_c is the comoving volume out to a given redshift over the whole sky (Hogg 1999). Furthermore, we specify the depth of the survey as a limiting apparent magnitude, m_0 .

We can simply generate the limiting absolute magnitude using the following relation:

$$M_0 = m_0 - (5 \log_{10} D_L(z) + K(z) + 25), \quad (33)$$

where $K(z)$ is the K -correction (Hogg & Baldry 2002) and D_L is the luminosity distance (Hogg 1999).

Now let the LF be denoted $\phi_L(M)$, and for it we employ a single-Schechter function (with parameters given by Loveday

et al. 2012). In addition, let the stellar-mass function (GSMF) be denoted $\phi_*(m_*)$, which will be a double-schechter function (with parameters given by Baldry et al. 2012). Using abundance matching, we can specify the limiting stellar mass by equating

$$\int^{m_{*,0}} \phi_*(m_*) dx_* = \int^{M_0} \phi_L(M) dM, \quad (34)$$

and solving for $m_{*,0}$ (this would typically require numerical splines or root-finding).

Let the underlying HMF be $g(m|\theta)$. Furthermore, let the function $f(m|\theta)$ define the mean ratio of stellar-to-halo mass, M_*/M for a given halo mass (i.e., the SMHM relation). We will use the best-fit TPL model seen in Figure 9. Then the average total stellar-mass content in a halo of mass m is

$$m_*^T = m + \log_{10} f(m|\theta). \quad (35)$$

This mass is, on average, comprised of a number of galaxies drawn from the GSMF. Note that here we are cutting a corner—the SMHM relation is typically defined using abundance matching between the GSMF and the HMF. That is, the SMHM itself assumes a single galaxy of a given mass in any halo. We continue to assume this for the SMHM relation definition, but change this assumption at this step to that of a number of galaxies constituting the given mass in any halo. We do not worry about conditional mass functions here, but identically draw galaxies from the universal GSMF within each halo.

What we desire is the distribution of galaxy mass counts for which the sum totals a given mass, i.e., the distribution of N such that

$$\sum_i^N 10^{X_i} = 10^{m_*^T}, \quad (36)$$

where $X_i \sim \phi_*(m_*)$. The solution to this problem is provided by the hitting time theorem (Hofstad & Keane 2008); however, the solution is rather intractable for our distributions.

Instead, we find an approximate answer by first generating an overall normalization of the GSMF required to generate the mass m_*^T :

$$D = \frac{m_*^T}{\int^{m_*^T} M_* \phi_*(m_*) dm_*}. \quad (37)$$

We then use this normalization to determine the average number of galaxies that make up the mass in the halo:

$$\bar{n}(m_*^T) = D \int_{m_{*,0}}^{m_*^T} \phi_*(m_*) dm_*. \quad (38)$$

What this model ignores is the enforced correlation between masses due to the necessity of summing to x_*^T , and the discrete nature of the distribution. We expect that these effects will be second-order, and for our purposes will be negligible. We make up for the latter in some respects by assuming that the occupation of the halo is Poisson distributed. We can generate the total fraction of observed halo masses, given some threshold for mass estimation, n_{\min} , by calculating the CDF of

the Poisson distribution at n_{\min} :

$$f_{\text{obs}}(m_{\star}^T) = 1 - \frac{\Gamma(n_{\min} + 1, \bar{n}(m_{\star}^T))}{n_{\min}!}. \quad (39)$$

Furthermore, we use the empirical relation of Robotham et al. (2011) to specify mass uncertainties on the dynamically estimated masses:

$$\sigma_i = \max(0.02, 1 - 0.43 \log_{10} N_{\text{FoF},i}), \quad (40)$$

with

$$N_{\text{FoF}} \sim \text{Poiss}(\bar{n}(m_{\star}^T)). \quad (41)$$

Thus we arrive at the following forward model:

$$V(m|\theta) = f_{\text{obs}}(m_{\star}^T(m|\theta)) V_0 \quad (42a)$$

$$m \sim V(m|\theta) \phi(m|\theta) \quad (42b)$$

$$N_{\text{FoF}} \sim \text{Poiss}(\bar{n}(m_{\star}^T(m|\theta)))_{n_{\min}} \quad (42c)$$

$$\sigma = \max(0.02, 1 - 0.43 \log_{10} N_{\text{FoF},i}) \quad (42d)$$

$$m' \sim \mathcal{N}(m, \sigma). \quad (42e)$$

We note that in this forward model we have avoided the fact that the mass-observable is in fact the velocity dispersion of the galaxies, other than incorporating an empirical relation for the uncertainty of this proxy. We note also that two ‘‘observables’’ are present in the model, both the mass (via the velocity proxy) and the number of galaxies in the halo. There is a temptation to neglect the latter in favor of directly using the derived uncertainties, σ , however, this proves to be an unreliable method as it ignores vital information.

Our application here will be to determine the constraints on θ for a range of survey parameters, Ω , m_0 , and n_{\min} .

7.2. Likelihood and Covariance

The basic task is to specify $n(\mathbf{x}'|\theta)$, the expected number density of observables, $\mathbf{x}' = (m', N_{\text{FoF}})$. This clearly requires the three sampling statements from the forward model above, and is merely integrated over underlying mass m :

$$n(\mathbf{x}'|\theta) = \int dm V(m|\theta) g(m|\theta) e^{-(m-m')^2/2\sigma^2(m)} \times \frac{\bar{n}(m)^{N_{\text{FoF}}} e^{-\bar{n}(m)}}{N_{\text{FoF}}!} \frac{1}{1 - \sum_{j=0}^{n_{\min}} \bar{n}(m)^j e^{-\bar{n}(m)}/j!}. \quad (43)$$

The second line of this equation is the truncated Poisson distribution. For clarity, we express the average observed halo occupancy as $\bar{n}(m)$, while noting that the dependency on the underlying mass arises explicitly through Equation (35), and that through this equation it also implicitly depends on θ .

With this expression, the likelihood of a given set of data, according to Equation (16), is

$$\ln L(\theta) = -n_{\theta} + \sum_i \ln n(m'_i, N_{\text{FoF},i}|\theta), \quad (44)$$

where n_{θ} is the expected number of observations, $n_{\theta} \equiv \int V g dm$.

The expected likelihood is gained by converting the sum over observations into a sum over infinitesimal bins in the observable space—i.e., conversion into a density-weighted integral (see Section 3.2). In our case, the observable space is two-dimensional, and one of the dimensions is discrete (i.e.

N_{FoF}). Thus we arrive at an expected likelihood of

$$\ln L = -n_{\theta} + \int dm' \sum_{j=n_{\min}} n'(m', j) \ln n(m', j). \quad (45)$$

We may also use the same reasoning to derive the expected Hessian, using Equation (21).

7.3. Results

We define the parameters of the data model such that observations are between $z \in (0.07, 0.13)$, and use an input θ gained from the formulae of Table 3 at the mean redshift. We then calculate the expected covariance at a range of values of m_0 and n_{\min} , with the sky area Ω varied so that the total number of observed galaxies (including those in systems with fewer than n_{\min} galaxies) is 5×10^4 .

Figure 10 shows the results of this analysis. In particular, the depth of the observation increases to the right in each panel, while the occupation threshold (i.e., the fewest number of galaxies a group must contain for inclusion in the sample) is denoted by the color of the curve.

We first consider the global properties of the uncertainty. We see that other than β , the parameters are typically fit to better than 10%, approaching 1% for h_{\star} and α . The poorer fit of β is expected as it is solely probed by the scant high-mass halos.

As a function of n_{\min} there is perhaps surprising unanimity on the result that a lower truncation yields a tighter posterior. This means that the extra raw statistics brought about by increased sample size outweighs the relative degradation in sample quality that a lowering in truncation brings.

Finally, as a function of depth, there are two groups. First, α and $\ln A$ monotonically increase in precision as the depth increases (though there is a clear flattening of this relationship above $m_0 \sim 20$). Conversely, h_{\star} and β initially increase their precision, but subsequently turn around and yield decreased precision for the deepest surveys. This can be intuitively understood. Initially, as the depth increases, the average quality of the data increases for the same average number of sources. However, at the same time, the range of masses probed is increased. As more and more low-mass halos are added to the sample, they begin to dominate the fit, and the high-mass halos to which β and h_{\star} are sensitive become relatively insignificant. At some point, around $m_0 \sim 20$ –23, the latter effect wins out, and the precision of the fits deteriorates.

This would suggest that an optimal strategy for determining all HMF parameters is to observe down to $m_0 \sim 21$ and use all observed group masses. While this analysis is rather simplistic, it is offered as an example application to show the potential power of more complex forecasting in future survey designs.

8. Conclusion

We have introduced a new analytic formula to describe the HMF over a broad range of scales, based loosely on the existing PS formalism, and closely related to the popular Schechter functional form. Our form is motivated by the desire to concisely and intuitively measure the HMF down to group-scale halos, and provide simple comparisons to galaxy statistics such as the stellar-mass function and LF.

We have shown that the MRP is an accurate model, within the uncertainties inherent in the mass function due to halo finders, halo definition, and simulation resolution.

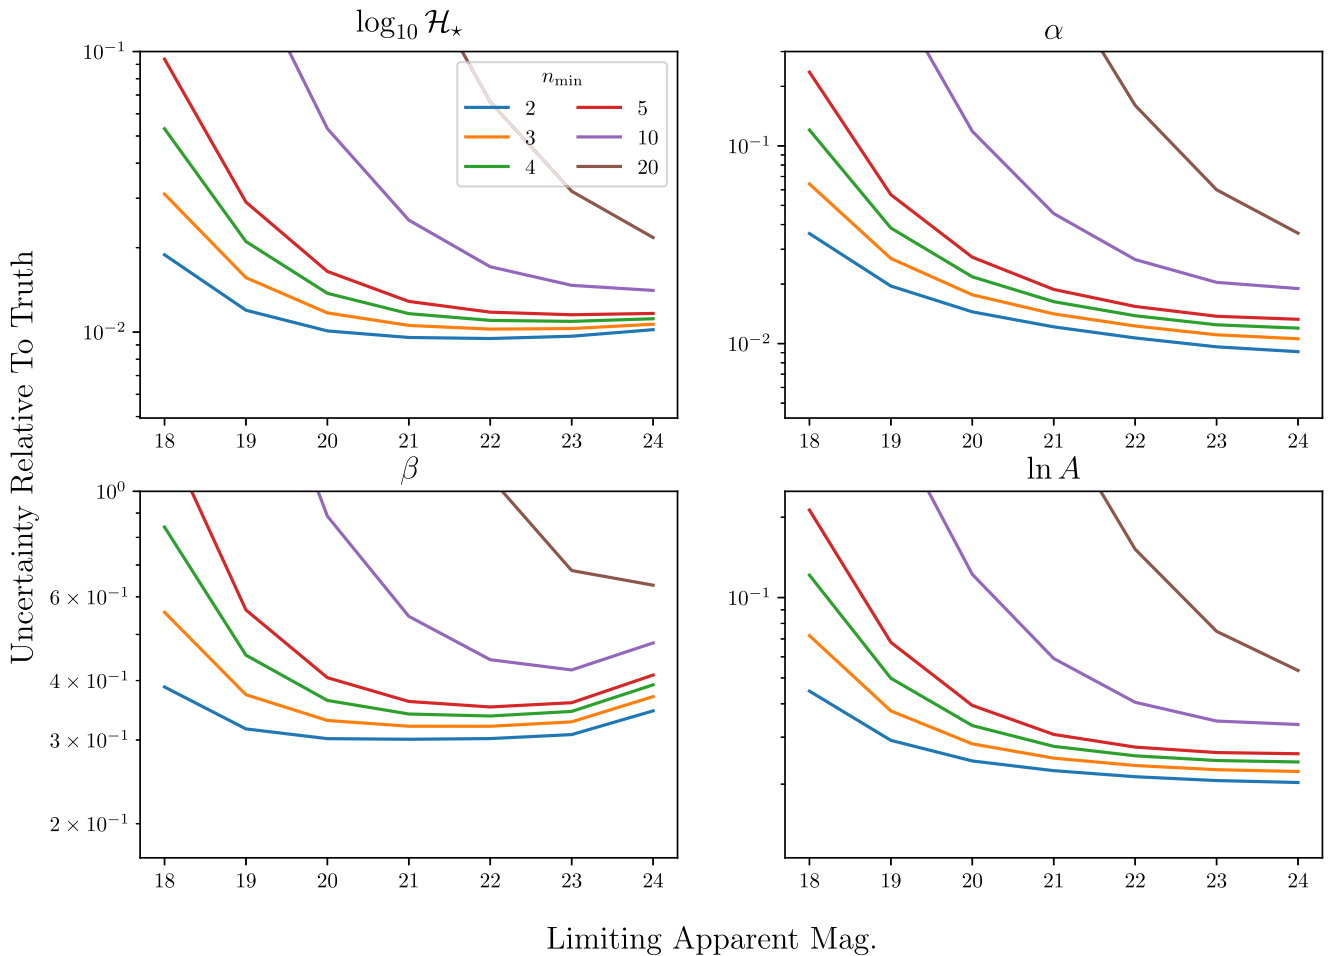


Figure 10. Marginal variance of the MRP parameters. The x -axis is the limiting apparent magnitude, m_0 , while the y -axis is the marginalized standard deviation of the parameter as a fraction of its true value. Different colors represent different threshold occupations.

To exemplify the utility of our functional form, we determined simple fitting functions for the dependence of the MRP parameters on the physical parameters z , Ω_m , and σ_8 , showing that each of the parameters primarily has a negative correlation with redshift. Furthermore, we showed that it is possible to analytically motivate more precise parameterizations of the SMHM relation by using the known properties of the MRP in conjunction with the common double-Schechter form for the stellar-mass function.

We presented in some detail the machinery required to use the MRP form to fit to halo catalogs—both simulated and observed. Importantly, we highlighted the Bayesian methods required to incorporate arbitrary measurement uncertainties into the analysis. We also presented novel results concerning the *expected* likelihood and covariance within a given data model.

Ultimately, we used these results to perform an example application, in which we determined the expected errors on fitted MRP parameters in a complex data model, including mass-dependent selection effects and dynamical mass uncertainties. We showed that in such a model, with a constant number of observed galaxies, the optimal depth of the survey is an apparent magnitude of approximately 21.

We have highlighted several areas in which the work presented here might be extended throughout the course of this paper. In particular, investigations into alternate parametrizations that have

better covariance properties may be beneficial. However, the best extensions will be to use the framework presented to fit actual data sets and begin to relate HMF parameters to those of the galaxy population.

S.G.M. would like to thank Ian McCarthy and Ishiyama Tomoaki for access to and help with simulation data. A.S.G.R. acknowledges support of a UWA postdoctoral research fellowship, as well as an Australian Research Council (ARC) Discovery Project (DP) grant, DP140100395. The Centre for All-Sky Astrophysics (CAASTRO) is an Australian Research Council Centre of Excellence, funded by grant CE11E0090. This research has made use of NASA’s Astrophysics Data System.

Appendix A Recurrence Relation

A problem that arises in the calculation of the incomplete gamma function is that several libraries only implement $\Gamma(z, x)$ for $z > 0$, and we often need it for $z < 0$. One way around this is to use the recurrence relation

$$\Gamma(z, x) = \frac{\Gamma(z + 1, x) - x^z e^{-x}}{z}, \quad (46)$$

In general, for arbitrary negative z , one can use the following extended relation to calculate the incomplete gamma function

Table 5

Alternate Parameterizations Found in the Literature or Presented Here

References	h_*	α	β
GG2 (Lagos Álvarez et al. 2011)	$\mu^{-1/\delta}$	$\nu - 1$	δ
GG3 (Lagos Álvarez et al. 2011)	μ	$\delta\nu - 1$	δ
HT	$\mu\left(\frac{\delta}{\nu+2}\right)^{1/\delta}$	ν	δ

Note. Columns indicate the transformation from the new parameters (μ, ν, δ) to the standard (h_*, α, β). We note that in the HT formula, μ is equivalent to the logarithmic mass mode \mathcal{H}_T .

in terms of only positive shape parameter:

$$\Gamma(z, x) = \Gamma(z) \left(Q(z+n, x) - x^z e^{-x} \sum_{k=0}^{n-1} \frac{x^k}{\Gamma(z+k+1)} \right), \quad (47)$$

where Q is the regularized incomplete gamma, $Q(z, x) = \Gamma(z, x)/\Gamma(z)$. To ensure the use of only positive shape parameters, n should be set to $\lceil -z \rceil$. Note that the accuracy of this formula will decrease for large $|z|$ due to numerical precision, and likewise the evaluation time will increase due to the sum.

Appendix B Reparameterizations

There are several reparameterizations of the GGD to be found in the literature, which could potentially bring better covariance properties than vanilla MRP. In this section, we investigate three of these formulations, along with one that we have devised.

Table 5 summarizes the parameterizations we will investigate here. There is also at least one form in the literature that claims to have better properties than those listed here (Lawless 1982), but is purely defined for $z > 0$, and therefore cannot be used.

The HT formulation has the immediate benefit that two of the parameters have physical meanings: ν is the power-law slope, and $\mu \equiv \mathcal{H}_T$ is the logarithmic mass mode. We expect this to be a very robust scale in any given data in which it exists (it will not exist when the slope is ≤ -2), and should be able to fit almost independently of the other parameters.

In nonhierarchical contexts, we may simply test the expected covariance of any reparameterization by converting the identified covariance of the MRP form to any reparameterization for which we have the derivatives of the transform.

Given a new vector of parameters, $\psi = (\mu, \nu, \delta)$, and a set of three equations $\theta(\psi) = \theta$ relating the parameters, then the Jacobian at any point ψ is given by

$$J_i^\psi = \mathbf{J}^\theta \cdot \frac{\partial \theta}{\partial \psi_i}. \quad (48)$$

We can then calculate the Hessian as

$$H_{ij}^\psi = \mathbf{J}^\theta \cdot \frac{\partial \theta}{\partial \psi^i \partial \psi^j} + \frac{\partial \theta}{\partial \psi^j} \cdot \mathbf{H}^\theta \cdot \frac{\partial \theta}{\partial \psi^i}. \quad (49)$$

To test the covariance properties of these parameterizations with respect to the MRP form, we generate mock halo catalogs by sampling the MRP at a typical point of




Figure 11. Properties of three reparameterizations of the MRP compared to the form used in this paper. The x -axis represents truncation mass. All dotted (dashed, solid) lines refer to GG2 (GG3, HT) respectively. The colors represent different quantities—the red (green, blue) line in the top-panel is the ratio of the standard deviation of μ (ν, δ) to h_* (α, β) at the solution (where the standard deviation is relative to the value at the solution). The black (magenta, cyan) line in the lower panel is the ratio of the correlation coefficient between $\mu - \nu$ ($\mu - \delta, \nu - \delta$) to the corresponding MRP correlations. All ratios are taken as the means of 20 realizations from ideal mock catalogs.

$\theta = (14.5, -1.85, 0.72)$ (corresponding roughly to P13 best-fit values). We do this for a series of differing truncation masses (from $m_{\min} = 10$ to 13) and for each, we produce 20 realizations, calculating the Hessians at the input point. We convert these into the new parameterization space using Equation (49), and then convert the result into covariance and correlation matrices.

Figure 11 shows the mean results of this analysis, with errorbars. Each of the linestyles represents different parameterizations, while the colors represent different parameter combinations. The top panel shows the ratios of variances in the new parameterization to MRP. The variances are all normalized to their parameter values. This indicates which parameterizations are more effective at net constraints on the parameters. We find that, as expected, $\delta \sim \beta$ is not improved for any of the alternate forms. On the other hand, ν is estimated similarly for HT, but is outside the range of the plot for both GG2 and GG3, and so has poorer precision. However, μ varies widely between the new forms. For GG2, it has a far poorer relative deviation, while for GG3, as expected, it remains consistent with unity. For HT, it is significantly more precise—for some truncation masses, it has double the precision.

The double-parameter quantities in the lower panel are the ratios of the correlation coefficients between the new parameterizations and the MRP. While the $\mu - \nu$ correlation swaps sign for GG2 and GG3, it perhaps surprisingly remains consistent with a parity of magnitude. Conversely, for HT, the correlation is rather dependent on m_{\min} , but is always significantly lower than MRP, as expected. We find much the same result for $\mu - \delta$, though HT only outperforms MRP for high m_{\min} . Finally, $\nu - \delta$ is relatively unchanged in all parameterizations except for GG3, in which it performs poorly at low m_{\min} .

We conclude that neither GG2 or GG3 are useful alternatives for the MRP in the context of typical values for the HMF, having no beneficial properties. Conversely, HT is either beneficial or benign in every aspect, and provides an interesting

reformulation for future studies. Its major potential drawback is that it is undefined for $\nu \leq 2$, which may be transgressed in some rare cases for certain values of m_{\min} .

An interesting corollary is that the same situation applies to the popular Schechter function. In that case, $\beta \equiv 1$, so that $\mathcal{H}_T = \mathcal{H}_s(\alpha + 2)$. Since $\alpha \approx -1$, \mathcal{H}_s has often been directly associated with the turnover. However, this is not strictly true, and for common measurements of α between -1.4 and -1.1 , there may be a corresponding benefit in using the true turnover scale.

Appendix C Cosmic Variance and Other Uncertainties

Throughout this paper, we have used a likelihood that is exact when the model prediction of the counts per infinitesimal bin is known, and the distribution from which the bin counts are taken is Poisson. In practice, the model prediction itself contains uncertainty due to cosmic variance, or even due to an uncertain selection function. In this section, we provide two methods of dealing with such uncertainty, both of which unfortunately require some form of ad hoc binning.

In general, the cosmic variance will have a *covariant* effect on the counts—bins with similar masses will tend to be correlated (while the Poisson uncertainty remains independent). One can account for this by employing an arbitrarily fine mass binning such that we have a vector \mathbf{X} of counts, of length N_{bins} . The first method of incorporating cosmic variance relies on integrating over the extra uncertainty. The expected counts per bin is given by $\boldsymbol{\lambda} = \mathbf{V}\phi$, and we have elsewhere assumed that for an infinitesimal bin, the number counts are drawn as $X_i \sim \text{Pois}(\lambda_i)$. However, with cosmic variance, the value of $\boldsymbol{\lambda}$ contains uncertainty. Let $f_{\text{cv}}(\boldsymbol{\lambda}'|\boldsymbol{\lambda})$ be the probability distribution function of values of $\boldsymbol{\lambda}'$, given the prediction from the model $\mathbf{V}\phi$, arising from cosmic variance (for example). Then the probability mass function of source counts, \mathbf{X} , is

$$P(\mathbf{X} = \mathbf{k}) = \int d\boldsymbol{\lambda}' \prod_i^{N_{\text{bins}}} \frac{\lambda_i'^{k_i} e^{-\lambda_i'}}{k_i!} f(\boldsymbol{\lambda}'|\boldsymbol{\lambda}). \quad (50)$$

This is a highly multi-dimensional integral, which will be typically impractical to solve on any iteration.

The second method is to use a hierarchical Bayesian model. This is most easily conceptualized as a generative model with the following steps:

1. Choose MRP parameters θ from their prior distributions.
2. Calculate $\boldsymbol{\lambda}$ for the arbitrary binning employed, which represents the universal background mass function.
3. Generate a realization of $\boldsymbol{\lambda}'$ for the sample at hand (i.e., N_{bins} new parameters), given the cosmic variance distribution $f_{\text{cv}}(\boldsymbol{\lambda}'|\boldsymbol{\lambda})$.
4. Generate true masses using a Poisson distribution with mean λ_i' in each bin.
5. Scatter true masses according to a measurement uncertainty distribution.

Such a model can be constrained by using this process on every iteration. That is, we employ N_{bins} new unknown parameters describing $\boldsymbol{\lambda}'$, which are drawn on every iteration. Equation (16) is then evaluated (necessarily numerically) using $\boldsymbol{\lambda}'(m)$ rather than $\boldsymbol{\lambda}(m)$. While this method involves an ad hoc binning scheme, it can in principle be tested for convergence by increasing the bin resolution (with the caveat that doing so will

require a greater computation time due to the extra parameters it must constrain).

Appendix D Parameter Estimation with MRP

In Section 3, we introduced a robust likelihood for a given choice of parameters, θ , and a sample of data, to be used for parameter estimation (see Equation (16)). In addition, we specified the expected covariance at the MLE given an input model (see Equation (21)). In this appendix, we supply additional details to these likelihoods to aid in practical implementations, and we also add details of a different kind of fit, in which the MRP functional form is directly fit to another model using least-squares minimization in log–log space.

In particular, we present the Jacobian (here interchangeable with the gradient) and Hessian (i.e., the matrix of second-order derivatives) for these likelihoods in the context of the MRP. These latter quantities are useful for accelerating convergence in downhill-gradient methods, and also enable the estimation of parameter covariances at the solution.

D.1. Method I: Binned/Curve Data

The basic method for fitting a model to a curve (or equivalently to binned data) is to use a minimization of χ^2 . In this case, due to the massive dynamic range of the data (both in mass and number density), it is preferable to perform the fit in log-space. Specifically, the log-likelihood of the model is given by

$$\begin{aligned} \ln \mathcal{L}_{\text{bin}}(\theta) &= -\frac{\chi^2}{2} = -\sum_i^N \frac{(\ln g(m_i|\theta) - \ln d_i)^2}{2\sigma_i^2} \\ &\equiv -\sum_i^N \frac{\Delta_{\text{data},i}^2}{2\sigma_i^2}, \end{aligned} \quad (51)$$

where g is defined as in Equation (5), d_i is the value of the i th bin of data, and σ_i is the uncertainty in the i th bin of data.

This likelihood may be generalized by allowing an extra constraint to play a role, namely that of the universal mean density, $\bar{\rho}_0 = \Omega_m \rho_c$. Briefly, this engenders a likelihood of

$$\ln \mathcal{L}_{\text{bin}}(\theta) = -\left(\sum_i^N \frac{\Delta_{\text{data},i}^2}{2\sigma_i^2} + \frac{\Delta_\rho^2}{2\sigma_\rho^2} \right), \quad (52)$$

where

$$\Delta_\rho = \ln A - \ln \bar{\rho}_0 + \ln k(\theta_3) \quad (53)$$

and k is defined (see Equation (13)) as

$$k(\theta_3) = \mathcal{H}_*^2 \Gamma(z_1). \quad (54)$$

In Equation (52), σ_ρ controls the tightness of the constraint. A value of infinity provides no constraint, in which case the likelihood is equivalent to Equation (51). Conversely, a value of zero requires that the parameters be chosen to exactly reproduce $\bar{\rho}_0$. In this case, the estimated parameters are $\theta_3 \equiv (h_*, \alpha, \beta)$, and Equation (53) is used to determine $\ln A$.

We provide this constraint purely for theoretical interest—using it in practice could lead to systematic biases since the behavior of the HMF below m_{\min} is not well determined.

D.1.1. Jacobians and Hessians

We here only derive the Jacobian and Hessian in the case that $\sigma_\rho \neq 0$. In this case, the Jacobian is a four-vector, determined by simple differentiation:

$$\mathbf{J}_{\text{bin}} = - \left(\sum \frac{\Delta_{\text{data},i}}{\sigma_i^2} \nabla \Delta_{\text{data},i} + \frac{\Delta_\rho}{\sigma_\rho^2} \nabla \Delta_\rho \right). \quad (55)$$

The Hessian is written as

$$\mathbf{H}_{lm} = - \left(\sum_i \frac{[\nabla_l \Delta_{\text{data},i} \nabla_m \Delta_{\text{data},i}] + \Delta_{\text{data},i} \nabla_{lm}^2 \Delta_{\text{data},i}}{\sigma_i^2} + \frac{[\nabla_l \Delta_\rho \nabla_m \Delta_\rho] + \Delta_\rho \nabla_{lm}^2 \Delta_\rho}{\sigma_\rho^2} \right). \quad (56)$$

The Jacobian and Hessian are thus merely constructed from the weighed sum of first and second derivatives of $\ln g$ and $\ln k$ (via Δ_{data} and Δ_ρ respectively). We list the values of these terms in Appendix D.4.

D.2. Method II: Sampled Data

The most typical case will be that of sampled data, for which the general likelihood was presented as Equation (16). Since each of the functions V , $\phi \equiv g$, and ρ_i may be dependent on θ , it is not useful to specify the derivatives of the likelihood here—often they will only be determined numerically, and as such it is simpler then to directly take the derivative of the likelihood itself.

Nevertheless, it is perhaps helpful to specify the simplest special case, in which V and ρ are independent of θ . In this case, the jacobian is

$$\mathbf{J}_l = - \int V g_l dm + \sum_{i=1}^{N_{\text{data}}} \frac{\int V g_l \rho(x_i) dm}{\int V g \rho(x_i) dm}, \quad (57)$$

and the Hessian is

$$\mathbf{H}_{lm} = - \int V g_{lm} dm + \sum_{i=1}^{N_{\text{data}}} \frac{\int V g_{lm} \rho(x_i) dm}{\int V g \rho(x_i) dm} - \frac{\int V g_l \rho(x_i) dm \int V g_m \rho(x_i) dm}{\left(\int V g \rho(x_i) dm \right)^2}, \quad (58)$$

where g_{lm} is the (double) derivative of g with respect to θ_l and θ_m .

These equations are true given that uniform priors are used for all parameters. Insertion of a nonuniform prior requires the addition of the derivatives of the logarithm of the prior distribution.

D.2.1. Special Case: Power-law Volume Function

The most complex special case that can be dealt with analytically is that in which the mass uncertainties are zero—i.e., ρ is a delta function—and $V(m) = V_0 \kappa m$. Unfortunately, the most commonly used uncertainty distribution for the masses is a log-normal distribution (i.e., a normal distribution in m), which renders the integral over $g\rho$ nonanalytic. The form of V taken here is reasonably general—it is a power law, which

is a reasonable form to expect physically, and to that we add an optional lower truncation at m_{min} , in the case that κ is too low to ensure convergence in density.

In this case, the form of Vg is equivalent to g with modified parameters, i.e., $V_0 \kappa m g(h_*, \alpha, \beta, \ln A) = g(h_*, \alpha + \kappa, \beta, \ln A + \ln V_0 + \ln 10 \kappa h_*)$, and we will denote the latter simply as \tilde{g} . The Jacobian of this system is thus

$$\mathbf{J}_l = -\tilde{q}_l + \sum_{i=1}^{N_{\text{data}}} \frac{\tilde{g}_l^i}{\tilde{g}^i} \equiv -\tilde{q}_l + \sum_{i=1}^{N_{\text{data}}} \ln \tilde{g}_l^i, \quad (59)$$

and the Hessian is

$$\begin{aligned} \mathbf{H}_{lm} &= -\tilde{q}_{lm} + \sum_{i=1}^{N_{\text{data}}} \frac{\tilde{g}_{lm}^i}{\tilde{g}^i} - \frac{\tilde{g}_l^i \tilde{g}_m^i}{\tilde{g}_i^2} \\ &\equiv -\tilde{q}_{lm} + \sum_{i=1}^{N_{\text{data}}} \ln \tilde{g}_{lm}^i. \end{aligned} \quad (60)$$

The multivariate chain rule may be used to relate \tilde{g}_l to g_l .

D.3. Method III: Expected Data

In Section 3.2, we presented the expected covariance at the MLE for a sample drawn from a given model, and fit with the same model. In this section, we present the general likelihood in this framework, and also some special cases in the context of the MRP.

The principle observation to make is that the sum in Equation (16) is over the samples, which must be drawn from some model distribution. The expected value of this distribution for any observed \mathbf{x}' may be written $n'(\mathbf{x}')$, where the prime represents an observed quantity. In the limit of infinitesimal bin-width, the sum becomes an integral over the observations:

$$\ln \mathcal{L} = - \int V \phi dm + \int n'(\mathbf{x}') \ln \int V \phi \rho(\mathbf{x}') dmd\mathbf{x}'. \quad (61)$$

There are two main special cases of interest, as well as their conjunction, in the context of the MRP.

D.3.1. Special Case I: Power-law Volume Function

Here we revisit the special case of Section D.2.1 to determine the expected likelihood. Here, the likelihood is

$$\ln \mathcal{L} = -\tilde{q} + \int n'(m) \ln \tilde{g}(m) dx, \quad (62)$$

with expected jacobian and Hessian given as

$$\mathbf{J}_l = -\tilde{q}_l + \int n'(m) \ln \tilde{g}_l dm \quad (63)$$

$$\mathbf{H}_{lm} = -\tilde{q}_{lm} + \int n'(m) \ln \tilde{g}_{lm} dm. \quad (64)$$

D.3.2. Special Case II: Correct Data Model

In this case, we assume that the model from which the data was actually drawn is the same as that being fit, such that the MLE of the parameters θ is equivalent to the true parameters of the data. Thus we have

$$n'(\mathbf{x}') = \int V' g' \rho(\mathbf{x}') dm. \quad (65)$$

The application of this formula to the likelihood, Equation (61) is trivial. More interesting is its application to the Jacobian and Hessian at the point of the MLE. Here, we use

the notation n^θ to refer to the total expected number of observations, equivalent to $\int Vgdm$. The Jacobian is

$$J_l = -n_l^\theta + \int n'(x') \frac{n_l(x')}{n(x')} dx', \quad (66)$$

where $n(x')$ is the model-dependent number density of observations at x' (i.e., the same formula as n' with primes removed). At the MLE, $n \equiv n'$, so that only $n_l(x')$ remains within the integral. However, the expected total density of observed objects is the same as the total number of true objects (i.e., the integral of $\rho(x')$ is unity), so that the first and second terms cancel each other, leaving a Jacobian of equivalently zero, as we should expect.

Using a similar method for the Hessian, we arrive at

$$H_{lm} = - \int \frac{n_l(x'|\hat{\theta})n_m(x'|\hat{\theta})}{n(x'|\hat{\theta})} dx', \quad (67)$$

which is just Equation (21). This equation is important, as it specifies a Hessian only in terms of first-order derivatives. Even in cases where the derivative of $n(x')$ must be attained numerically, this provides a significant performance and accuracy boost to the Hessian calculation.

D.3.3. Special Case III: Conjunction of Special Cases

Here we take the conjunction of the previous two special cases, such that $n(x') = \tilde{g}$. This case is very useful, as we can determine a completely analytic solution to both the likelihood and covariance at the MLE. The results can be used to quickly infer the amount of information in the raw MRP model at a given set of parameters, and also to determine optimal power-law slopes for the volume function.

The likelihood at any θ , for data drawn from θ' , is

$$\ln \mathcal{L} = -\tilde{q} + \int \tilde{g}' \ln \tilde{g} dm. \quad (68)$$

For the remainder of the derivation, we omit the tilde notation, since every variable is in the transformed variables. To evaluate the integral, we can use integration by parts. We note that the indefinite integral

$$\int g dm = A \mathcal{H}_* \gamma(z, x) \equiv A(q_0 - q(m)), \quad (69)$$

where γ is the lower-incomplete gamma function, and use it to find

$$\int_{m_{\min}}^{\infty} g' \ln g dm = q' \ln g_{\min} + A' \mathcal{H}_* \int_{m_{\min}}^{\infty} -\gamma(z', x') \frac{d \ln g}{dm} dm, \quad (70)$$

where the subscript min refers to the fact that here g is a function of the (optional) truncation mass m_{\min} , and we have

$$\frac{d \ln g}{dm} = \frac{\alpha + 1}{m} - \frac{\beta x}{m}. \quad (71)$$

Compiling everything, we have

$$\ln \mathcal{L} = -q + q' \ln g_{\min} + A' \mathcal{H}_* ((\alpha + 1)t' - u), \quad (72)$$

where

$$t' = \int_{m_{\min}}^{\infty} (\Gamma(z', x') - \Gamma(z')) \frac{dm}{m} \quad (73)$$

and

$$u = \int_{m_{\min}}^{\infty} (\Gamma(z', x') - \Gamma(z')) \frac{\beta x}{m} dm. \quad (74)$$

For both, we use the substitution $dm = m dx' / (\beta' x')$, and for u we also use integration by parts to finally yield

$$u = \left(\frac{\mathcal{H}_*'}{\mathcal{H}_*} \right)^\beta \left[\Gamma \left(z' + \frac{\beta}{\beta'}, x' \right) - \frac{x'^{\beta/\beta'} q'}{\mathcal{H}_*'} \right] \quad (75)$$

and

$$t' = \frac{\Gamma(z')}{\beta'} \left[y'^{\alpha'+1} \Gamma(z') {}_2F_2 \left(\begin{matrix} z' & z' \\ z'+1 & z'+1 \end{matrix} \middle| -x' \right) + \text{Poly} \Gamma(z') - \beta' \ln y' \right], \quad (76)$$

where ${}_pF_q$ is the generalized hypergeometric function and $y = 10^{m-h_*}$.

To derive the Hessian, it makes the most sense to begin with Equation (72) rather than Equation (67). We find

$$H_{lm} = -q_{lm} + q' \ln g_{\min, lm} - A' \mathcal{H}_*' u_{lm}. \quad (77)$$

Expressions for the derivatives of $\ln g$, q , and u can be found in Section D.4 (note that in this case, the derivatives of $\ln g$ have m replaced with m_{\min}).

D.4. Derivatives of g , q , k , and u

Several of the Jacobians and Hessians listed in this section require the derivatives of $\ln g$, q , $\ln k$, and u . Here we list these as coupled equations. Throughout the following, we denote partial derivatives with respect to any parameter by a subscript, and in each case, the order of parameters in the Jacobian/Hessian is $(h_*, \alpha, \beta, \ln A)$.

We begin with $\ln g$:

$$J_{\ln g} = \left\{ \ln(10)(\beta x_m - \alpha), \ln y, \frac{1}{\beta}(1 - x_m \ln x_m), 1 \right\} \quad (78)$$

$$H_{\ln g} = \begin{pmatrix} -\ln^2(10)\beta^2 x & -\ln 10 & \ln(10)x(1 + \ln x) & 0 \\ \cdot & 0 & 0 & 0 \\ \cdot & \cdot & -\frac{1}{\beta^2}(1 + x \ln^2 x) & 0 \\ \cdot & \cdot & \cdot & 0 \end{pmatrix}. \quad (79)$$

For the remaining, it is helpful if we first define a couple of functions based on the Meijer-G function G :

$$G1(z, x) = G_{2,3}^{3,0} \left(x \middle| \begin{matrix} 1 & 1 \\ 0 & 0 \end{matrix} \middle| z \right) / \Gamma(z, x), \quad (80)$$

$$G2(z, x) = G_{3,4}^{4,0} \left(x \middle| \begin{matrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{matrix} \middle| z \right) / \Gamma(z, x), \quad (81)$$

and

$$\bar{G}(z, x) = G1(z, x) \ln x + 2G2(z, x). \quad (82)$$

Furthermore, we use the following commonly occurring relation:

$$\Phi = yg / \Gamma(z, x). \quad (83)$$

We note that everywhere in the derivatives of q , instances of m are taken to rather be m_{\min} . Keeping this in mind, for the first derivatives of $\ln q$, we have

$$\mathbf{J}_q = q \left\{ \ln(10)(1 + \Phi), \beta^{-1}(\ln x + G1), -z \ln q_\alpha - \frac{\ln y}{\beta} \Phi, 1 \right\} \quad (84)$$

and for the double-derivatives

$$\mathbf{H}_q = \begin{pmatrix} \ln(10)(q_h + q\Phi \ln g_h) & \ln(10)(q_\alpha + q\Phi \ln g_\alpha) & \ln(10)(q_\beta + q\Phi \ln g_\beta) & q_h \\ \cdot & \beta^{-1} \left(-q_\alpha + q_\beta \ln x - q \frac{z\bar{G}}{\beta} \right) & \beta^{-1} \left(-q_\alpha + q_\beta \ln x - q \frac{z\bar{G}}{\beta} \right) & q_\alpha \\ \cdot & \cdot & q \frac{\Phi \ln y}{\beta} \left(\frac{1}{\beta} - \ln g_\beta \right) - z(q_{\alpha,\beta} - \frac{q_\alpha}{\beta}) & q_\beta \\ \cdot & \cdot & \cdot & q \end{pmatrix}. \quad (85)$$

The derivatives of $\ln k$:

$$\mathbf{J}_{\ln k}^3 = \{2 \ln(10), \text{Poly}\Gamma(0, z_k)/\beta - z_k \ln k_\alpha\} \quad (86)$$

and

$$\mathbf{H}_{\ln k}^3 = \begin{pmatrix} 0 & 0 & 0 \\ \cdot & \text{Poly}\Gamma(1, z_k)/\beta^2 & - \left(z_k \ln k_{\alpha,\alpha} + \frac{\ln k_\alpha}{\beta} \right) \\ \cdot & \cdot & \ln k_{\alpha,\alpha} z_k^2 + z_k \frac{\ln k_\alpha}{\beta} \end{pmatrix}. \quad (87)$$

Finally, the derivatives of u :




$$\mathbf{J}_u^3 = \left\{ -\ln(10)\beta u, 0, \frac{u \ln x + G1 \cdot \Gamma}{\beta} \right\}, \quad (88)$$

and

$$\mathbf{H}_u^3 = \begin{pmatrix} -\ln(10)\beta u_h & 0 & -\ln(10)(u + \beta u_\beta) \\ \cdot & 0 & 0 \\ \cdot & \cdot & \frac{u_\beta \ln x + \bar{G} \cdot \Gamma/\beta}{\beta} \end{pmatrix} \quad (89)$$

where Γ , $G1$, and \bar{G} are here functions of $(z + 1, x)$. Note that these formulae for the partial derivatives of u are only strictly correct at the solution (in general, they involve θ' as well as θ). The accompanying implementations calculate the more general form.

ORCID iDs

S. G. Murray  <https://orcid.org/0000-0003-3059-3823>
A. S. G. Robotham  <https://orcid.org/0000-0002-4003-0904>
C. Power  <https://orcid.org/0000-0002-4003-0904>

References

- Baldry, I. K., Driver, S. P., Loveday, J., et al. 2012, *MNRAS*, **421**, 621
Behroozi, P. S., Conroy, C., & Wechsler, R. H. 2010, *ApJ*, **717**, 379
Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013a, *ApJ*, **770**, 57
Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013b, *ApJ*, **762**, 109
Berlind, A. A., Weinberg, D. H., Benson, A. J., et al. 2003, *ApJ*, **593**, 1
Bocquet, S., Saro, A., Dolag, K., & Mohr, J. J. 2016, *MNRAS*, **456**, 2361
Cooray, A., & Sheth, R. 2002, *PhR*, **327**, 1
de Haan, T., Benson, B. A., Bleem, L. E., et al. 2016, *ApJ*, **832**, 95
Driver, S. P., Davies, L. J., Meyer, M., et al. 2016, in *Astrophysics and Space Science Proc. 42, The Universe of Digital Sky Surveys*, ed. N. R. Napolitano (Cham: Springer), 205
Driver, S. P., Hill, D. T., Kelvin, L. S., et al. 2011, *MNRAS*, **413**, 971
Durkalec, A., Le Fèvre, O., de la Torre, S., et al. 2015, *A&A*, **576**, L7
Eke, V. R., Baugh, C. M., Cole, S., Frenk, C. S., & Navarro, J. F. 2006, *MNRAS*, **370**, 1147
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
Gelman, A., & Rubin, D. B. 1992, *StaSc*, **7**, 457
Hofstad, R. V. D., & Keane, M. 2008, *American Mathematical Monthly*, 753
Hogg, D. W. 1999, arXiv:[astro-ph/9905116](https://arxiv.org/abs/astro-ph/9905116)
Hogg, D. W., & Baldry, I. K. 2002, arXiv:[astro-ph/0210394](https://arxiv.org/abs/astro-ph/0210394)
Ishiyama, T., Enoki, M., Kobayashi, M. A. R., et al. 2015, *PASJ*, **67**, 61
Jenkins, A., Frenk, C. S., White, S. D. M., et al. 2001, *MNRAS*, **321**, 372
Knebe, A., Knollmann, S. R., Muldrew, S. I., et al. 2011, *MNRAS*, **415**, 2293
Lagos Álvarez, B., Ferreira, G., & Valenzuela Hube, M. 2011, *Proyecciones (Antofagasta)*, **30**, 415
Lawless, J. F. 1982, *Canadian Journal of Statistics*, **10**, 316
Loveday, J., Norberg, P., Baldry, I. K., et al. 2012, *MNRAS*, **420**, 1239
McConnachie, A., Babusiaux, C., Balogh, M., et al. 2016, arXiv:[1606.00043](https://arxiv.org/abs/1606.00043)
Moster, B. P., Somerville, R. S., Maulbetsch, C., et al. 2010, *ApJ*, **710**, 903
Murray, S. G., Power, C., & Robotham, A. S. G. 2013, *A&C*, **3-4**, 23
Mutch, S. J., Croton, D. J., & Poole, G. B. 2013, *MNRAS*, **435**, 2445
Obreschkow, D., Murray, S. G., Robotham, A. S. G., & Westmeier, T. 2018, *MNRAS*, **474**, 5500
Pacaud, F., Clerc, N., Giles, P. A., et al. 2016, *A&A*, **592**, A2
Planck Collaboration 2014, *A&A*, **571**, A29
Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A24
Press, W. H., & Schechter, P. 1974, *ApJ*, **187**, 425
Robotham, A. S. G., Norberg, P., Driver, S. P., et al. 2011, *MNRAS*, **416**, 2640
Schechter, P. 1976, *ApJ*, **203**, 297
Schmidt, M., & Lipson, H. 2009, *Sci*, **324**, 81
Sheth, R. K., Mo, H. J., & Tormen, G. 2001, *MNRAS*, **323**, 1
Tinker, J. L., Kravtsov, A. V., Klypin, A., et al. 2008, *ApJ*, **688**, 709
Warren, M. S., Abazajian, K. N., Holz, D. E., & Teodoro, L. 2006, *ApJ*, **646**, 881
Watson, W. A., Iliev, I. T., D'Aloisio, A., et al. 2013, *MNRAS*, **433**, 1230
White, S. D. M., & Rees, M. J. 1978, *MNRAS*, **183**, 341
Zehavi, I., Zheng, Z., Weinberg, D. H., et al. 2005, *ApJ*, **630**, 1
Zehavi, I., Zheng, Z., Weinberg, D. H., et al. 2011, *ApJ*, **736**, 59