

Made-to-measure modelling of observed galaxy dynamics

Jo Bovy,^{1,2★†} Daisuke Kawata³ and Jason A. S. Hunt⁴

¹Department of Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

²Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, New York, NY 10010, USA

³Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey RH5 6NT, UK

⁴Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, Ontario, M5S 3H4, Canada

Accepted 2017 September 14. Received 2017 September 6; in original form 2017 April 9

ABSTRACT

Amongst dynamical modelling techniques, the made-to-measure (M2M) method for modelling steady-state systems is amongst the most flexible, allowing non-parametric distribution functions in complex gravitational potentials to be modelled efficiently using N -body particles. Here, we propose and test various improvements to the standard M2M method for modelling observed data, illustrated using the simple set-up of a one-dimensional harmonic oscillator. We demonstrate that nuisance parameters describing the modelled system’s orientation with respect to the observer – e.g. an external galaxy’s inclination or the Sun’s position in the Milky Way – as well as the parameters of an external gravitational field can be optimized simultaneously with the particle weights. We develop a method for sampling from the high-dimensional uncertainty distribution of the particle weights. We combine this in a Gibbs sampler with samplers for the nuisance and potential parameters to explore the uncertainty distribution of the full set of parameters. We illustrate our M2M improvements by modelling the vertical density and kinematics of F-type stars in *Gaia* DR1. The novel M2M method proposed here allows full probabilistic modelling of steady-state dynamical systems, allowing uncertainties on the non-parametric distribution function and on nuisance parameters to be taken into account when constraining the dark and baryonic masses of stellar systems.

Key words: Galaxy: kinematics and dynamics – solar neighbourhood – galaxies: fundamental parameters – galaxies: general – galaxies: kinematics and dynamics – galaxies: structure.

1 INTRODUCTION

Constraining the orbital structure and mass distribution of astrophysical systems through dynamical modelling is one of the fundamental ways to learn about the dark-matter and baryonic distribution in external galaxies (e.g. Rix et al. 1997; Cappellari et al. 2012), supermassive black holes at the centres of galaxies (e.g. Magorrian et al. 1998), and the mass distribution of the Milky Way (e.g. Bovy & Rix 2013), to name but a few. Of particular interest are systems – such as galaxies or star clusters – that may be assumed to be in a steady state. Many techniques have been proposed to model such systems, typically combining the steady-state assumption with further assumptions about the orbital structure (e.g. the velocity anisotropy) or symmetry (e.g. spherical or axisymmetric) of the system. The simplest amongst these techniques are those based on moments of the collisionless Boltzmann equation, e.g. the Jeans equations, which despite their restrictive assumptions remain a useful tool for interpreting data (e.g. Cappellari et al. 2013). A

second class of techniques directly uses parametrized distribution functions (DFs) that satisfy the collisionless Boltzmann equation by only depending on integrals of the motion. While restricted to gravitational potentials for which such integrals can be computed, this class of models has reached a high level of sophistication, especially in the Milky Way (e.g. Binney 2010; Bovy & Rix 2013; Trick, Bovy & Rix 2016). A third class of methods eschews parametrized DFs, but rather builds a steady-state model in a fixed gravitational potential from a large number of orbit building blocks with weights determined by fitting a set of constraints (Schwarzschild 1979, 1993).

Syer & Tremaine (1996) proposed a method known as made-to-measure (M2M) modelling that is closely related to orbit-based modelling. In M2M, the DF is represented not by entire orbits but instead by a set of N -body particles with positions and velocities $(\mathbf{x}_i, \mathbf{v}_i)$ and weights w_i . They demonstrated that a steady-state solution to a set of constraints on the phase-space distribution (expressed as a χ^2 , the mean-squared difference between the model and the data) can be obtained by slowly adjusting the weights of each particle in the direction of decreasing χ^2 while integrating the orbits of the particles. The advantages of this particle-based technique over orbit-based methods are that only the current snapshot needs to be

* E-mail: bovy@astro.utoronto.ca

† Alfred P. Sloan Fellow.

stored in memory rather than entire orbits, that the N -body particles can contribute a self-consistent part of the gravitational potential, and that one ends up with an actual sampling from the steady-state DF. The latter makes M2M also an ideal technique for initializing N -body simulations (e.g. Dehnen 2009; Malvido & Sellwood 2015).

Since its original description, various improvements have been made to the basic M2M set-up, such as allowing for observational uncertainties and kinematic data in the constraints (De Lorenzi et al. 2007; Long & Mao 2010), integrating particles on individual time-scales for problems with a range of orbital frequencies (Dehnen 2009), improvements in the smoothing applied to the model (Dehnen 2009), and allowing data for individual stars as constraints (Hunt & Kawata 2013; Hunt, Kawata & Martel 2013; Hunt & Kawata 2014). As currently conceived, M2M modelling applies to the particle weights only and any other parameter describing the system is held fixed during the optimization. This includes nuisance parameters describing the modelled system's orientation with respect to the observer, for example, the inclination of an external galaxy or the Sun's distance to the Galactic centre for Milky Way applications, and the parameters of the external gravitational field. Furthermore, as methods for modelling observed data both Schwarzschild and M2M modelling remain problematic in that they are fundamentally optimization algorithms that do not take into account the uncertainties in the DF resulting from the strong degeneracies amongst the large number of orbit or particle weights (Magorrian 2006). For obtaining the best constraints from a given set of observables, a fully probabilistic treatment is warranted that samples from the full uncertainty distribution for the particle weights, nuisance parameters and the parameters describing the potential. In this paper, we extend the basic M2M modelling framework to optimize for nuisance and potential parameters simultaneously with the particle weights and we introduce sampling methods to sample the uncertainty distribution of all parameters.

The outline of this paper is as follows. In Section 2, we describe the simple, one-dimensional set-up that we use as a toy problem: modelling an isothermal population in an external harmonic-oscillator potential. We describe the standard M2M method in Section 3. In Section 4, we discuss how to sample from the uncertainty distribution of the particle weights. We show how one can optimize the value of the nuisance parameters at the same time as the values of the particle weights in Section 5 and give a Markov Chain Monte Carlo (MCMC) algorithm to sample both the particle weights and the nuisance parameters. In Section 6, we discuss how we can also optimize the value of the parameters describing the external gravitational potential simultaneously with the particle weights and the nuisance parameters and present an MCMC algorithm for sampling all parameters. To illustrate how the M2M improvements perform for real data, we apply the new M2M algorithm to data on the density and kinematics of F stars in *Gaia* DR1 in Section 7. We discuss various aspects of this novel M2M method and avenues for future work in Section 8 and present our conclusions in Section 9.

2 HARMONIC-OSCILLATOR M2M: A SIMPLE TESTBED FOR M2M MODELLING

To illustrate and test our modelling extensions of the basic M2M algorithm below, we consider a one-dimensional system with the gravitational potential of a harmonic oscillator (HO). This set-up is chosen for its simplicity; everything that we describe below applies more generally to full, three-dimensional M2M modelling. This

set-up is an ideal testbed for M2M modelling because (a) orbit integration is analytic, (b) the DF corresponding to a given potential and a given density is unique (thus, there is a well-defined unique solution to the M2M problem; e.g. Kuijken & Gilmore 1989), (c) it is easy to write down simple DF models and (d) running the M2M modelling in practice is very fast. While simple, this model is a also semirealistic, approximate representation of the vertical dynamics in the solar neighbourhood close to the mid-plane and thus has some practical applicability (see Section 7). We ignore the self-gravity of the M2M N -body particles and the potential is thus assumed to be external and fixed. In this section, we describe the basic notation, equations and concepts of this model.

We denote the phase-space coordinates as (z, v_z) . The HO potential is

$$\Phi(z; \omega) = \frac{\omega^2 z^2}{2}, \quad (1)$$

specified by a single parameter ω , the oscillator's frequency. Orbit integration in the HO potential is analytic: orbits are given by

$$z_i(t) = A_i \cos(\omega t + \phi_i), \quad (2)$$

$$v_{z,i}(t) = -A_i \omega \sin(\omega t + \phi_i), \quad (3)$$

where

$$A_i = z_{\max} = \frac{\sqrt{2E_i}}{\omega} = \sqrt{z_i^2(0) + \frac{v_{z,i}^2(0)}{\omega^2}}, \quad (4)$$

$$\phi_i = \arctan2(-v_{z,i}(0)/\omega, z_i(0)), \quad (5)$$

in which $(z_i(0), v_{z,i}(0))$ is the initial phase-space position of an orbit indexed by i and $\arctan2$ is the arc-tangent function that chooses the quadrant correctly.

In this HO potential, we attempt to match a population drawn from a DF given by

$$f(z, v_z) \propto e^{-E/\sigma^2}, \quad (6)$$

where $E = \omega^2 z^2/2 + v_z^2/2$ is the energy and σ is the velocity dispersion parameter. This DF is isothermal – it has the same velocity dispersion at all heights $\langle v^2 \rangle = \sigma^2$ – and in a steady state, because it is only a function of the conserved energy E . The density distribution for this distribution is

$$v(z) \propto \exp\left(-\frac{\omega^2 z^2}{2\sigma^2}\right), \quad (7)$$

which is a Gaussian distribution with a standard deviation $\sigma_v = \sigma/\omega$. The velocity distribution at each z is a Gaussian with dispersion σ . Sampling orbits at initial phase-space locations $(z_i(0), v_{z,i}(0))$ from $f(z, v_z) \propto e^{-E/\sigma^2}$ is simple: (i) sample E_i from the exponential distribution and convert it to A_i ; (ii) sample ϕ_i uniformly between 0 and 2π ; (iii) convert (A_i, ϕ_i) to $(z_i, v_{z,i})$.

To fit this DF using M2M below, we start with $(z_i, v_{z,i})$ drawn with uniform weights w_i from an isothermal DF, but with a different σ from the true velocity dispersion: $f(z, v_z) \propto e^{-E/\sigma_{\text{in}}^2}$, with σ_{in} typically 0.2. It is then easy to see that the correct output particle weights for a true velocity-dispersion parameter σ_{target} should be

$$w_i \propto \exp\left(-E_i \left[\frac{1}{\sigma_{\text{target}}^2} - \frac{1}{\sigma_{\text{in}}^2}\right]\right), \quad (8)$$

if the potential remains fixed. If the potential is adiabatically changed from a HO potential with frequency ω_{in} to one with frequency ω_{target} the correct output particle weights are

$$w_i \propto \exp\left(-J_i \left[\frac{\omega_{\text{target}}}{\sigma_{\text{target}}^2} - \frac{\omega_{\text{in}}}{\sigma_{\text{in}}^2} \right]\right), \quad (9)$$

where $J_i = E_i/\omega$ is the action.

3 STANDARD M2M MODELLING

We first describe the standard M2M case. Standard M2M models a steady-state DF as a set of N particles $(z_i, v_{z,i})$ indexed by i orbiting in a fixed potential. Each particle has a weight w_i that is adjusted on-the-fly during orbit integration to fit a set of constraints, like the density in bins, or the velocity dispersion. By only adjusting the particle weights w_i on time-scales \gg the orbital time-scale, an approximate equilibrium distribution is obtained.

In practice, M2M maximizes an objective function F that represents a balance between reproducing the constraints, expressed as χ^2 differences between data and model, and a penalty S that disfavors non-smooth DFs

$$F = S - \frac{1}{2} \sum_j \chi_j^2. \quad (10)$$

Traditionally, the penalty S is implemented through a maximum-entropy constraint by setting

$$S = -\mu \sum_i w_i [\ln(w_i/\hat{w}_i) - 1], \quad (11)$$

where \hat{w}_i is a default set of particle weights. In the absence of constraints, the entropy penalty prefers $w_i = \hat{w}_i$. The parameter μ quantifies the strength of the penalty.

Constraints are expressed as a kernel applied to the DF $f(z, v_z)$:

$$Y_j = \int dz dv_z K_j(z, v_z) f(z, v_z), \quad (12)$$

which for the N -body snapshot is computed as

$$y_j = \sum_i w_i K_j(z_i, v_{z,i}). \quad (13)$$

To illustrate the standard M2M case, we use the density and the density-weighted mean-squared velocity, both observed at a few points indexed by j . The model density is given by

$$\nu(\tilde{z}_j) = \sum_i w_i K^0(|\tilde{z}_j + z_\odot - z_i|; h), \quad (14)$$

where $K^0(r; h)$ is a kernel function with a width parameter h that integrates to one ($\int dr K^0(r; h) = 1$) and we assume that the observations are done as a function of \tilde{z} , which is measured with respect to the observer's position, located at z_\odot from the $z = 0$ mid-plane (we give the specific kernel and \tilde{z} used in this paper in Section 3.2). In what follows, we will abbreviate $K_j^0(z_i; h) \equiv K^0(|\tilde{z}_j + z_\odot - z_i|; h)$. We assume that the density is observed with a Gaussian error distribution characterized by a variance $\sigma_{0,j}^2$ and the contribution $\chi_{j,0}^2$ from the density to χ^2 is then

$$\chi_{j,0}^2 = [\Delta_j^0/\sigma_{0,j}]^2 = (\nu(\tilde{z}_j) - \nu_j^{\text{obs}})^2/\sigma_{0,j}^2, \quad (15)$$

where we have defined $\Delta_j^0 = \nu(\tilde{z}_j) - \nu_j^{\text{obs}}$, with superscript '0' to indicate that this is a zeroth moment quantity.

The model density-weighted mean-squared velocity is given by

$$\nu\langle v_z^2 \rangle(\tilde{z}_j) = \sum_i w_i v_{z,i}^2 K^0(|\tilde{z}_j + z_\odot - z_i|; h), \quad (16)$$

where we have chosen a kernel $K_j^{\text{II}}(z_i, v_{z,i}) = v_{z,i}^2 K_j^0(z_i; h)$. As for the density, we assume that this quantity is observed with a Gaussian error distribution with variance $\sigma_{2,j}^2$ and the contribution $\chi_{j,\text{II}}^2$ to χ^2 is

$$\chi_{j,\text{II}}^2 = [\Delta_j^{\text{II}}/\sigma_{2,j}]^2 = (\nu\langle v^2 \rangle(\tilde{z}_j) - \nu\langle v^2 \rangle_j^{\text{obs}})^2/\sigma_{2,j}^2, \quad (17)$$

where we have defined $\Delta_j^{\text{II}} = \nu\langle v^2 \rangle(\tilde{z}_j) - \nu\langle v^2 \rangle_j^{\text{obs}}$, where the superscript indicates that this is a second-moment quantity (not a square), like for the density difference Δ_j^0 . The reason that we work with the density-weighted mean-squared velocity is that it has a simple form; for applications to data, one might want to use the mean-squared velocity directly (for example in the example application in Section 7), but this requires normalizing by the density and thus leads to more complicated derivatives below (see Appendix A).

The standard M2M *force of change* equation is then given by

$$\begin{aligned} \frac{dw_i}{dt} &= \epsilon w_i \frac{\partial F}{\partial w_i} \\ &= \epsilon w_i \left[\frac{\partial S}{\partial w_i} - \frac{1}{2} \sum_j \frac{\partial \chi_{j,0}^2}{\partial w_i} - \frac{1}{2} \sum_j \frac{\partial \chi_{j,\text{II}}^2}{\partial w_i} \right]. \end{aligned} \quad (18)$$

In this equation, we have that

$$-\frac{1}{2} \frac{\partial \chi_{j,0}^2}{\partial w_i} = -\Delta_j^0 K_j^0(z_i; h)/\sigma_{0,j}^2, \quad (19)$$

$$-\frac{1}{2} \frac{\partial \chi_{j,\text{II}}^2}{\partial w_i} = -\Delta_j^{\text{II}} v_{z,i}^2 K_j^0(z_i; h)/\sigma_{2,j}^2, \quad (20)$$

and

$$\frac{\partial S}{\partial w_i} = -\mu \ln[w_i/\hat{w}_i]. \quad (21)$$

We solve equation (18) using a simple Euler method with a fixed step size, computing the orbital evolution as we go along using equations (2) and (3). Unlike most previous applications of M2M, we do not require $\sum_i w_i = 1$, but instead let the total weight be constrained by the data (see discussion in Section 8.3).

The M2M method for optimizing the objective function can be thought of as a sort of gradient ascent. Gradient-ascent optimization of an objective function does not have a physical time-scale associated with it. However, by writing the gradient-ascent algorithm in the manner of equation (18), we are essentially performing gradient ascent on a clock that runs with time $\tau = \epsilon t$ compared to the orbit integration that runs with time t . If $\Delta t \approx 1$ is the orbital time-scale, substantial changes to the objective function and the particle weights only happen on time-scales $1/\epsilon$. M2M works by adjusting ϵ such that $1/\epsilon \gg 1$, the orbital time-scale, which pushes the particle weights to an equilibrium distribution.

3.1 Previous extensions to the standard M2M algorithm

For the sake of completeness, we discuss some of the previous extensions to the standard M2M method that have been proposed. These are all concerned with how the M2M optimization for the particle weights is run and are thus different from the extensions that we propose in the following sections on how to fit additional parameters beside the particle weights and how to sample from the uncertainty distribution of all parameters.

Syer & Tremaine (1996) propose to lessen the impact of Poisson noise due to the finite number of N -body particles by smoothing the Δ_j^0 and Δ_j^{II} deviations that appear in equation (18) with smoothed

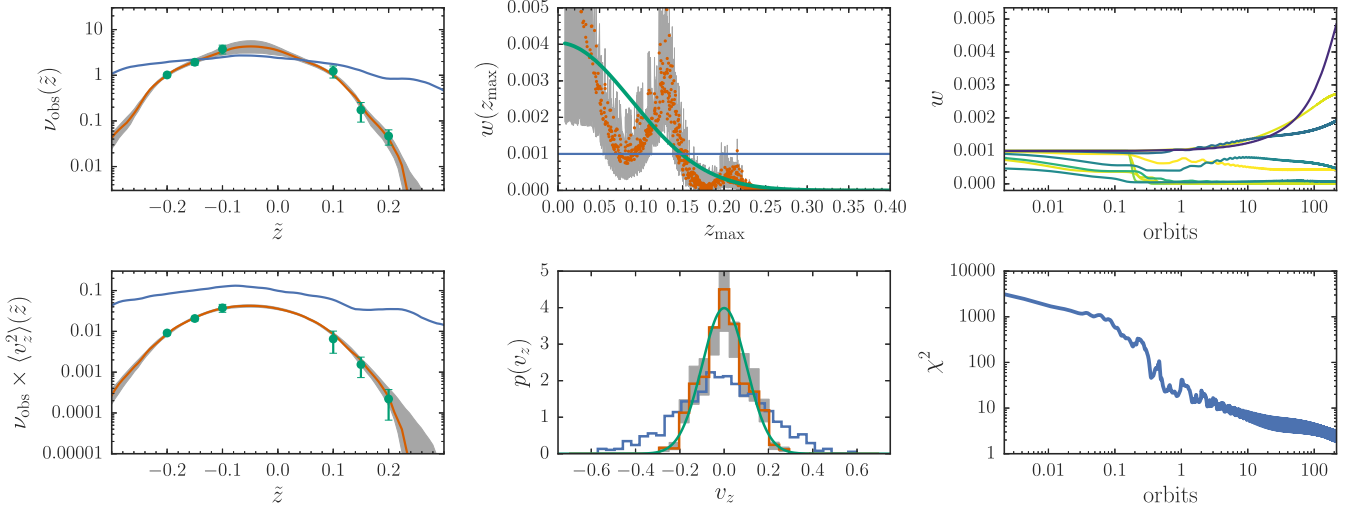


Figure 1. Basic M2M. The left-hand panels display the observed mock data in green: density $\nu(\bar{z})$ (top) and density-weighted mean-squared velocity $\nu\langle v_z^2 \rangle(\bar{z})$ (bottom). The blue curve shows the initial model, while the red curve displays the model for the best-fitting particle weights. The top, middle panel shows the best-fitting particle weights in red, the initial weights in blue, and the true weights in green. The top, middle panel shows the velocity distribution (for all z) for the initial model (blue), the final model (red), and the true, Gaussian distribution (green). The right-hand panels demonstrate how ten randomly-selected particle weights evolve (top) and how the total χ^2 converges in the M2M optimization. The grey band in the left four panels displays the 68 per cent uncertainty region in the fit obtained from 100 samples of the PDF for the particle weights.

versions $\tilde{\Delta}_j^0$ and $\tilde{\Delta}_j^{\text{II}}$. In the end, this leads one to solve for $(\tilde{\Delta}_j^0, \tilde{\Delta}_j^{\text{II}})$ using the differential equations

$$\frac{d\tilde{\Delta}_j^0}{dt} = \alpha (\Delta_j^0 - \tilde{\Delta}_j^0), \quad (22)$$

$$\frac{d\tilde{\Delta}_j^{\text{II}}}{dt} = \alpha (\Delta_j^{\text{II}} - \tilde{\Delta}_j^{\text{II}}), \quad (23)$$

where α is another inverse-time-scale parameter. Because we only want to smooth on shorter time-scales than that over which we substantially change the particle weights, we typically need $\alpha \gtrsim \epsilon$ (see Syer & Tremaine 1996 for a detailed discussion of this constraint). Dehnen (2009) considers a modified version of this procedure in which not the constraint but the objective function itself is smoothed. This leads one to smooth the force-of-change factor $\partial F/\partial w_i$ itself in a similar manner as the Syer & Tremaine (1996) smoothing

$$\frac{d}{dt} \left(\frac{\partial F}{\partial w_i} \right) = \alpha \left(\frac{\partial F}{\partial w_i} - \widetilde{\frac{\partial F}{\partial w_i}} \right). \quad (24)$$

Note that if we discretize the solution of equation (18) with a step-size δt , setting $\alpha = 1/\delta t$ is equivalent to no smoothing and α cannot be set to a larger value. Malvido & Sellwood (2015) argue that for large particle numbers, smoothing is redundant in that the unsmoothed algorithm already leads to final particle weights based on the smoothed objective function. We do not apply any smoothing in any of the examples in this paper.

Dehnen (2009) also introduced a method for solving the M2M optimization where each particle gets integrated on its own (approximate) time-scale. This is a necessary addition when modelling systems with a wide range of orbital time-scales (e.g. Hunt & Kawata 2013) and all of our extensions of the traditional M2M algorithm below apply to this formalism from Dehnen (2009) as well. However, we do not consider it here further, because all orbits in our example problem of the HO have the exact same orbital frequency.

3.2 An example M2M fit

Fig. 1 shows an example of the standard M2M algorithm. We draw 100 000 mock data points from an isothermal DF with $\sigma = 0.1$ and $\sum_i w_i = 1$ in a HO potential with $\omega = 1.3$. We evaluate the density $\nu(\bar{z})$ and the density-weighted mean-squared velocity $\nu\langle v_z^2 \rangle(\bar{z})$ at $\bar{z} = \{\pm 0.10, \pm 0.15, \pm 0.20\}$ for $z_\odot = 0.05$ using the expressions in equations (14) and (16) with a kernel width of $h = 0.025$ for an Epanechnikov kernel

$$K^0(x; h) = \begin{cases} \frac{3}{4h} \left[1 - \left(\frac{x}{h}\right)^2 \right], & 0 \leq x \leq h, \\ 0, & \text{otherwise.} \end{cases} \quad (25)$$

We then assume Gaussian uncertainties $\sigma_{0,j}$ and $\sigma_{2,j}$ and obtain the measurements ν_j^{obs} and $\nu\langle v_z^2 \rangle_j^{\text{obs}}$ displayed in the left-hand panels of Fig. 1. These are the measurements that we use for all of the tests in this paper.

To model these mock data, we draw 1000 M2M particles from the isothermal DF with $\sigma = 0.2$ – twice the true σ – and assign them initial weights $w_i = 1/1000$. We fix z_\odot and ω to their true values. We run the standard M2M optimization algorithm with $\epsilon = 10^{-3.5}$ and solve the M2M evolution with a stepsize of $\pi/3 \times 10^{-2}$ for 10^5 steps or about 217 orbits. We do not apply a roughness penalty ($\mu = 0$) to let the data fully determine the particle weights. We compute observables from these 1000 particles using a kernel with size $h = 0.075$, three times larger than the kernel used to generate the mock data. We chose this larger kernel to demonstrate that the kernel size or even its shape may be different between the data and the model observables, as long as they consistently measure the observable in question. In Section 7, we apply the new methods developed in this paper to *Gaia* data, where to account for the *Gaia* selection function the kernel used must be a set of rectangular bins, while the model observables are computed using an Epanechnikov kernel, because rectangular bins do not have well-behaved derivatives.

The resulting fit is shown in red in Fig. 1. In the left-hand panels the red line is the model’s density and density-weighted mean-squared velocity evaluated at the final snapshot of the particles with

their best-fitting weights. The model is smooth and fits the data well. The top, middle panel displays the best-fitting weights w_i . These oscillate around their true value, indicated by the green curve. The bottom, middle panel shows the velocity distribution (for all z) of the final particle distribution. This velocity distribution is close to a Gaussian with $\sigma = 0.1$, the true distribution displayed in green. The right-hand panels demonstrate how the particle weights (top) and χ^2 (bottom) converge. At the end of the procedure, we have that $\chi^2 \approx 2.5$ and we do not optimize further (the true minimum is $\chi^2 \approx 2$). Some of the weights that are largely unconstrained by the data are still evolving somewhat, without affecting the model fit.

4 UNCERTAINTIES ON THE PARTICLE WEIGHTS

The standard M2M algorithm returns the best-fitting particle weights without any estimate of their uncertainties. Standard algorithms for sampling from the uncertainty distribution for the particle weights, such as MCMC methods of various sorts, could in principle be applied if we interpret the objective function in equation (10) as the logarithm of a posterior PDF. However, these algorithms do not work well for the M2M problem, because this posterior PDF evaluated at any given snapshot is noisy, the weights-space is high-dimensional (dimension 1000 in the test example employed in this paper), and the uncertainties of the particle weights are highly correlated.

The method for obtaining uncertainties on the particle weights that we propose here is based on the following simple observation. Consider a linear model in which the vector of observations \mathbf{Y} is modelled as a function of a parameter vector $\mathbf{W} \in \mathbb{R}^N$ as $\mathbf{Y} = \mathbf{K} \mathbf{W} + \boldsymbol{\delta}$, where $\boldsymbol{\delta} \sim \mathcal{N}(\mathbf{0}, \mathbf{S})$ is Gaussian noise with mean $\mathbf{0}$ and known variance \mathbf{S} (which may include correlations between different components of \mathbf{Y}), and \mathbf{K} is a constant matrix. For this model, the posterior probability distribution function (PDF) under a uniform prior is given by

$$p(\mathbf{W}|\mathbf{Y}, \mathbf{S}) = \mathcal{N}(\mathbf{M} = \mathbf{V}[\mathbf{K}^T \mathbf{S}^{-1} \mathbf{Y}], \mathbf{V}), \quad (26)$$

where the variance \mathbf{V} is given by

$$\mathbf{V} = [\mathbf{K}^T \mathbf{S}^{-1} \mathbf{K}]^{-1} \quad (27)$$

(e.g. Hogg, Bovy & Lang 2010). Rather than computing the mean and variance of this Gaussian posterior PDF, we can sample from the posterior PDF as follows:

$$\tilde{\mathbf{Y}} \sim \mathcal{N}(\mathbf{Y}, \mathbf{S}) \quad (28)$$

$$\tilde{\mathbf{M}} = \mathbf{V}[\mathbf{K}^T \mathbf{S}^{-1} \tilde{\mathbf{Y}}]. \quad (29)$$

That is, we sample new observations $\tilde{\mathbf{Y}}$ from the uncertainty distribution of \mathbf{Y} and compute the ‘best-fitting’ $\tilde{\mathbf{M}}$ for this new set of observations. This $\tilde{\mathbf{M}}$ is a sample from the posterior PDF: (a) the distribution of $\tilde{\mathbf{M}}$ is Gaussian, because $\tilde{\mathbf{M}}$ is a linear transformation of another Gaussian variable $\tilde{\mathbf{Y}}$, (b) the expectation value of $\tilde{\mathbf{M}} = \mathbf{M}$ and (c) the variance $\langle \tilde{\mathbf{M}} \tilde{\mathbf{M}}^T \rangle = \mathbf{V}$; because a Gaussian distribution is fully characterized by its mean and variance, this proves that the distribution of $\tilde{\mathbf{M}}$ is the correct posterior PDF.

In the M2M objective function in equation (10), the observations $\mathbf{Y} = Y_j$ are linearly related to the weight parameters $\mathbf{W} = w_i$ through the kernels $\mathbf{K} = K_j(z_i, v_{z,i})$. The algorithm above is based on the assumption that there is no constraint on the sign of each

Algorithm 1: Particle Weights Monte Carlo Sampling

```

/* To draw  $K$  sets of particle weights
    $\{w_i\}_k$  for data points  $\mathbf{Y}$  with uncertainty
   covariance  $\mathbf{S}$  */
1 for  $k = 1, 2, \dots, K$  do
2    $\tilde{\mathbf{Y}} \sim \mathcal{N}(\mathbf{Y}, \mathbf{S})$ 
3    $\{w_i\}_k \leftarrow$  M2M optimize  $w_i$  for data points  $\tilde{\mathbf{Y}}$ 
4     with uncertainty covariance  $\mathbf{S}$ 
5    $(z_i, v_{z,i}) \leftarrow$  value at the end of M2M
6     optimization
7 end

```

of the weight parameters w_i . Our M2M application, however, requires that all weights be non-negative for the DF to be everywhere non-negative. To deal with this, we therefore force the \tilde{w}_i to remain positive, which is automatically the case when using the M2M optimization algorithm described above; we discuss the effect of this constraint in more detail below. Thus, we sample particle weights from the weights PDF by (a) sampling new observations \tilde{Y}_j from the uncertainty distribution for each Y_j and (b) by computing the best-fitting particle weights \tilde{w}_i using the standard M2M algorithm (which includes the $\tilde{w}_i > 0$ constraint). Each such set \tilde{w}_i is an independent sample from the weights PDF, unlike in a Markov chain. We will refer to this as the ‘data-resampling method for sampling the particle weights PDF’. This method is presented in Algorithm 1. The algorithm, as written down there, draws K samples from the uncertainty distribution for the particle weights; when we use this algorithm as part of a larger Gibbs MCMC chain, we will typically use it to draw just a single sample ($K = 1$ in Algorithm 1).

This method does not properly deal with particle weights for which the penalty term in equation (11) (which becomes a prior when sampling the particle weights) has a significant effect or for weights that, if they were allowed to be negative, have much probability mass at $w_i < 0$. An extreme case of the former are weights of orbits that do not pass through any observed volume. Under the optimization algorithm, these will always return the prior weight with no scatter. If the prior on the particle weights was Gaussian, we could similarly sample new prior means as the first step in the algorithm in equation (28) (because the prior mean \hat{w}_i is in this case equivalent to an ‘observation’ of w_i with an error variance equal to the prior variance). We do not implement this here, but see further discussion of this in Section 8.4. For weights that want $w_i < 0$, the optimization algorithm will effectively associate all probability mass at $w_i < 0$ with $w_i = 0$. While this is not technically correct – it does not sample from the posterior PDF – it is reasonable to set weights to zero that want to be less than zero. Some M2M algorithms remove orbits with small or zero weights and our sampling method effectively samples from the two alternative models for such orbits with the probability of these two alternatives determined by the data: (a) they get removed ($w_i = 0$) and (b) they have non-zero weights ($w_i > 0$). We discuss the issue of particle weights that prefer to be negative in the context of the M2M extensions in the next two sections further in Section 8.2.

An example of the data-resampling method for sampling the particle-weights PDF is shown in Fig. 1. We draw 100 samples from the weight PDF, that is, 100 sets of 1000 particle weights. Each set is optimized using the same optimization settings as in Section 3.2; each set’s initial particle distribution is set to the final snapshot of the previous sample. Because we set $\mu = 0$, we assume a flat,

improper prior $w_i > 0$ for all particle weights. The grey band displays the $\approx 1\sigma$ range spanned by this sample of particle weights. The uncertainty in the particle weights (top, middle panel) and consequent uncertainty in the density and density-weighted mean-squared velocity (left-hand panels) and the velocity distribution (bottom, middle panel) adheres to our physical intuition. For example, orbits with $z_{\max} < 0.05$ are essentially only constrained by the observations at $\tilde{z} = -0.1$, which corresponds to $z = -0.05$ because $z_{\odot} = 0.05$; the uncertainty in the particle weights blows up at $z_{\max} < 0.05$ because of this. The density kernel for an observation at z is dominated by orbits with $z_{\max} \approx z$, while the velocity-squared kernel at all z gets large contributions from stars with large z_{\max} . Therefore, weights at high z_{\max} are strongly constrained by the velocity data. The uncertainty in the density in the left-hand panel is therefore large near $\tilde{z} \approx 0$, while the uncertainty in the velocity is small at the same \tilde{z} . At large \tilde{z} the data allow a more steeply declining density and/or velocity, but not a shallower distribution (which would have too large velocities at low heights). Keep in mind that these strong relations depend on knowing the gravitational potential and keeping it fixed.

5 OPTIMIZING AND SAMPLING NUISANCE PARAMETERS

Dynamical modelling of observed galaxy kinematics often requires the knowledge of parameters separate from those specifying the DF (the particle weights in the M2M case) and those related to the gravitational potential. These are typically related to the observer's perspective: for example, the observer's three-dimensional position and velocity with respect to the centre of the system being modelled (e.g. the Sun's distance from the Galactic centre for Milky Way dynamics) or the observer's viewing angle (e.g. a galaxy's inclination for an external galaxy, the Sun's position with respect to the bar when modelling the central Milky Way). These types of parameters enter into the kernel evaluation in the M2M objective function. The standard method for determining these parameters is to optimize the M2M objective function for the particle weights on a grid of nuisance parameters. Here, we demonstrate that the M2M objective function in equation (10) can be optimized simultaneously for the particle weights and the nuisance parameters.

As an example we consider the Sun's height z_{\odot} above the plane. The Sun's height enters the kernels through $z = \tilde{z} + z_{\odot}$. Similar to the standard M2M algorithm, we can form a force of change equation for z_{\odot} as

$$\begin{aligned} \frac{dz_{\odot}}{dt} &= \epsilon_{\odot} \frac{\partial F}{\partial z_{\odot}} \\ &= \epsilon_{\odot} \left[-\frac{1}{2} \sum_j \frac{\partial \chi_{j,0}^2}{\partial z_{\odot}} - \frac{1}{2} \sum_j \frac{\partial \chi_{j,\text{II}}^2}{\partial z_{\odot}} \right], \end{aligned} \quad (30)$$

where we have allowed the freedom to use a different ϵ_{\odot} from the ϵ parameter used in the force-of-change equation for the particle weights. We have that

$$-\frac{1}{2} \frac{\partial \chi_{j,0}^2}{\partial z_{\odot}} = -\frac{\Delta_j^0}{\sigma_{0,j}^2} \sum_i w_i \frac{dK_j^0(r;h)}{dr} \bigg|_{|\tilde{z}_j + z_{\odot} - z_i|} \text{sign}(\tilde{z}_j + z_{\odot} - z_i), \quad (31)$$

Algorithm 2: MCMC sampling of nuisance parameters

```

/* To draw K MCMC samples  $z_{\odot,k}$ , given a
   set of particle weights  $\{w_i\}$  and a
   gravitational potential, for data
   points Y with uncertainty covariance
   S
   */
// Average objective function for
   current  $z_{\odot}$ :
1  $\tilde{F} \leftarrow 0$ 
2 for  $m = 1, 2, \dots, M$  do
3    $(z_i, v_{z,i}) \leftarrow$  advance orbits by 1 step
4    $\tilde{F} += F(z_i, v_{z,i} | z_{\odot}, w_i, Y, S) / M$ 
5 end
   // MCMC sample using
   Metropolis-Hastings:
6 for  $k = 1, 2, \dots, K$  do
   // Draw proposed  $z'_{\odot}$ :
7    $z'_{\odot} \sim Q(z'_{\odot} | z_{\odot})$ 
8    $(z_i, v_{z,i}) \leftarrow$  rewind orbits by M steps
   // Average objective function for  $z'_{\odot}$ :
9    $\tilde{F}' \leftarrow 0$ 
10  for  $m = 1, 2, \dots, M$  do
11     $(z_i, v_{z,i}) \leftarrow$  advance orbits by 1 step
12     $\tilde{F}' += F(z_i, v_{z,i} | z'_{\odot}, w_i, Y, S) / M$ 
13  end
   // Accept/reject:
14   $q \leftarrow \tilde{F}' - \tilde{F}$ 
15   $r \sim [0, 1]$ 
16  if  $\ln r < q$  then
17     $z_{\odot} \leftarrow z'_{\odot}$ 
18     $\tilde{F} \leftarrow \tilde{F}'$ 
19  else
20     $z_{\odot} \leftarrow z_{\odot}$ 
21  end
22   $z_{\odot,k} \leftarrow z_{\odot}$ 
23 end

```

$$-\frac{1}{2} \frac{\partial \chi_{j,\text{II}}^2}{\partial z_{\odot}} = -\frac{\Delta_j^{\text{II}}}{\sigma_{2,j}^2} \sum_i w_i v_{z,i}^2 \frac{dK_j^0(r;h)}{dr} \bigg|_{|\tilde{z}_j + z_{\odot} - z_i|} \text{sign}(\tilde{z}_j + z_{\odot} - z_i). \quad (32)$$

If one wants to include a prior on z_{\odot} there would be an additional contribution to the force of change from this prior. We then again solve the system of equations (18) and (30) using an Euler method with a fixed step size, computing the orbital evolution as we go along using equations (2) and (3).

An example of this is displayed in Fig. 2, where we fit the same data as in the example described in Section 3.2, but now also fitting z_{\odot} . All of the optimization parameters are kept the same and we set $\epsilon_{\odot} = 10^{-6} \approx \epsilon/300$. We start at an initial guess of $z_{\odot} = -0.05$, far from the true value. We see that z_{\odot} quickly and smoothly converges to $z_{\odot} = 0.053$, close to the true value.

After finding the best-fitting z_{\odot} from the M2M optimization, we can sample the joint posterior PDF for (w_i, z_{\odot}) using a Metropolis-Hastings-within-Gibbs sampler by repeating the following steps

$$(a) \quad w_i \sim p(w_i | z_{\odot}, \text{observations}), \quad [\text{Algorithm 1}] \quad (33)$$

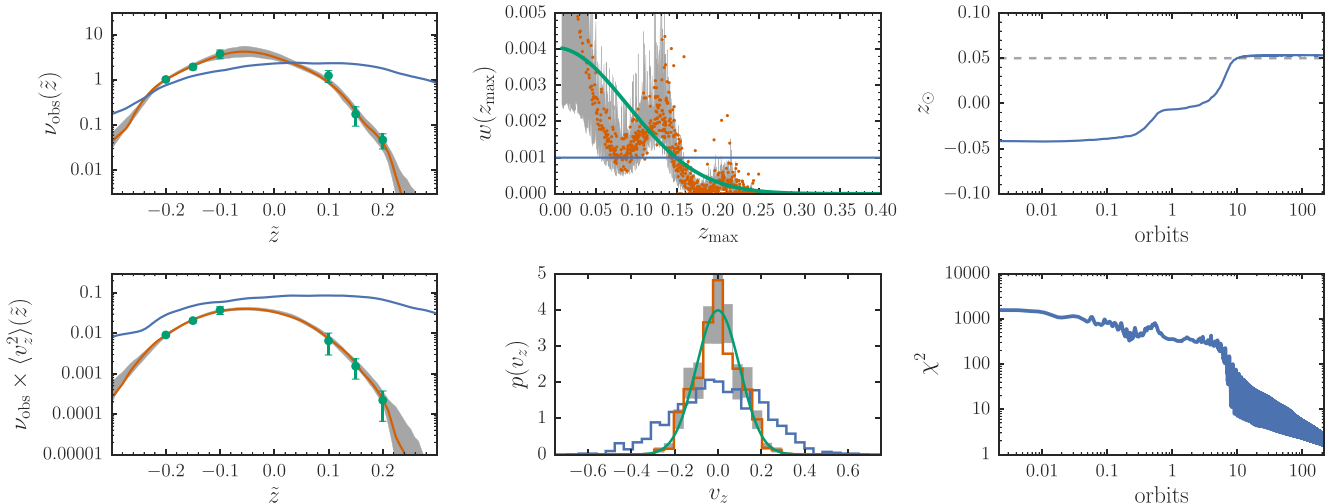


Figure 2. M2M with nuisance parameters. Like Fig. 1, except that we also fit for the Sun’s height above the plane z_{\odot} using the force of change for z_{\odot} and we sample the uncertainty in both the particle weights and z_{\odot} . The top right-hand panel demonstrates how z_{\odot} converges during the joint M2M optimization of the particle weights and z_{\odot} . We find that $z_{\odot} = 0.0527 \pm 0.0042$, in good agreement with its true value of 0.0500, shown as the dashed grey line in the top right-hand panel.

$$(b) \quad z_{\odot} \sim p(z_{\odot} | w_i, \text{observations}), \quad [\text{Algorithm 2}], \quad (34)$$

where we sample particle weights in the (a) step using the data-resampling technique of Section 4 (see Algorithm 1 with $K = 1$ to draw a single particle-weights sample) and sample z_{\odot} using a Metropolis-Hastings (MH) update using the objective function as the log posterior PDF $\ln p(z_{\odot} | w_i, \text{observations})$, in which the weights w_i are held fixed. Step (b) is presented in detail in Algorithm 2. In practice, we average the objective function in step (b) over about one orbital period (lines 2–5 and 10–13 in Algorithm 2) and use the exact same orbital trajectories (thus the rewind step in line 8 of Algorithm 2) to reduce the noise in the objective function. Because the optimization in step (a) typically requires tens to hundreds of orbital periods, step (b) proceeds quickly compared to step (a). We can improve mixing in the MCMC chain by performing multiple MH steps for each weights sample ($K > 1$ in Algorithm 2) and keeping only the final z_{\odot} sample in each step (b); as long as the total number of orbital steps in (b) is much less than that for a single optimization, this does not increase the computational cost significantly.

The result of this procedure for the example problem is shown in Figs 2 and 3. We have drawn 100 samples from the joint PDF of the particle weights and z_{\odot} , using a Gaussian proposal distribution with standard deviation $\sigma_{z_{\odot}} = 0.01$ and performing 10^5 M2M optimization time steps in step (a) and 20 MH steps for each particle-weight sample. The chain is initialized at the best-fitting z_{\odot} from the M2M optimization described above. We average the objective function using $M = 500$ steps or about 1 orbital period. The behaviour of the MCMC chain is displayed in Fig. 3. This figure demonstrates that the chain is well mixed and has a small correlation length (adjacent samples have very different values). The chain for the particle weights demonstrates that weights with similar z_{max} are strongly correlated. The acceptance ratio for the MH steps for z_{\odot} for this chain is 0.30.

The uncertainty in the density and velocity profiles in Fig. 2 now includes the uncertainty in z_{\odot} and this increases the overall uncertainty. We find that $z_{\odot} = 0.0527 \pm 0.0042$. We can compare this to the standard method of constraining z_{\odot} : we optimize the

particle weights for a set of fixed z_{\odot} and record the minimum χ^2 for each z_{\odot} . This gives $z_{\odot} = 0.0534 \pm 0.0046$. We can also compare our M2M-based result to the result if we assume that the DF is isothermal with unknown σ and normalization. In that case, the data constrain $z_{\odot} = 0.0560 \pm 0.0048$, similar to the M2M analyses. Overall, we find that the novel M2M procedure performs well.

6 OPTIMIZING AND SAMPLING THE GRAVITATIONAL POTENTIAL

Traditionally, M2M modelling, much like Schwarzschild modelling, keeps the external gravitational field fixed during the M2M fit. The gravitational potential is optimized by running the fit for different fixed potentials and choosing the potential that provides the best fit. While the overall distance and velocity scale can be optimized by writing down a force of change equation for these (De Lorenzi et al. 2008), this does not apply to other parameters of the potential. However, similar to the force of change for nuisance parameters, we can write down the force of change for parameters describing the potential and adjust these parameters during the fit. Naively, the problem with this procedure is that the instantaneous objective function F does not depend on the potential, because it is only a function of the current phase-space position of the M2M particles. In this section, we discuss how to get around this problem, such that we can fit and MCMC sample the parameters describing the gravitational potential.

Using our HO example, we vary the frequency ω of the HO potential. The force of change equation for ω is

$$\begin{aligned} \frac{d\omega}{dt} &= \epsilon_{\omega} \frac{\partial F}{\partial \omega} \\ &= \epsilon_{\omega} \left[-\frac{1}{2} \sum_j \frac{\partial \chi_{j,0}^2}{\partial \omega} - \frac{1}{2} \sum_j \frac{\partial \chi_{j,\text{II}}^2}{\partial \omega} \right], \\ &= -\epsilon_{\omega} \left[\frac{\Delta_j^0}{\sigma_{0,j}^2} \frac{\partial \Delta_j^0}{\partial \omega} + \frac{\Delta_j^{\text{II}}}{\sigma_{2,j}^2} \frac{\partial \Delta_j^{\text{II}}}{\partial \omega} \right], \end{aligned} \quad (35)$$

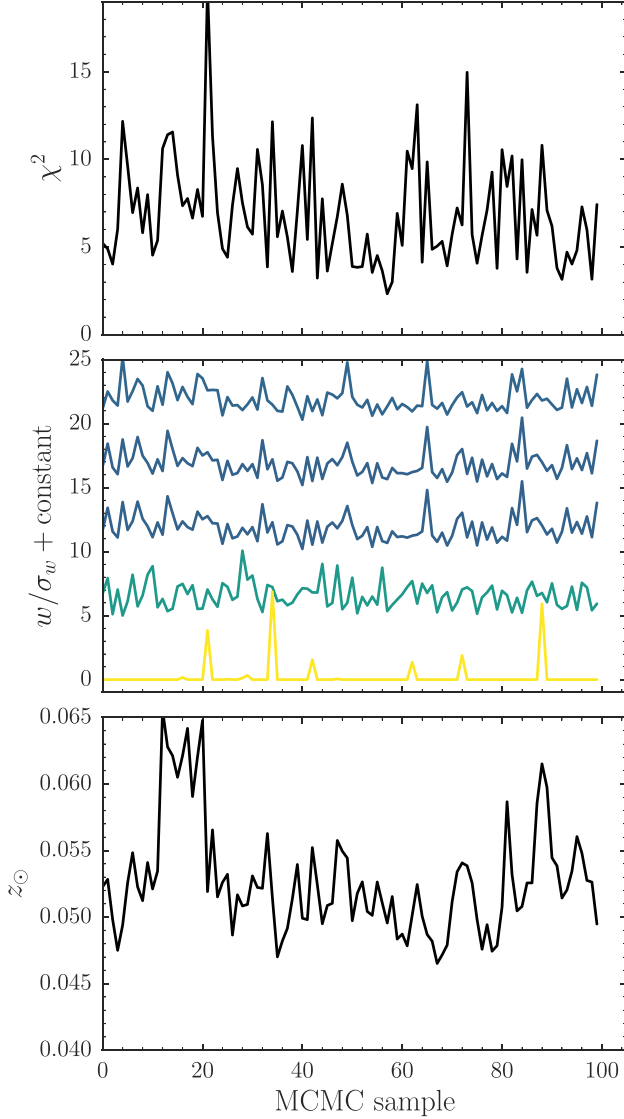


Figure 3. MCMC sampling of the particle weights and z_{\odot} . This figure demonstrates the MCMC chain of 100 samples from the uncertainty distribution of the particle weights and z_{\odot} constrained by the mock data. The top panel displays the behaviour of χ^2 , the middle panel that of five random particle weights (normalized by the standard deviation of their samples, colour-coded by z_{max}) and the bottom panel shows the z_{\odot} samples. The chain has a small correlation length, because we perform 20 MH steps for z_{\odot} for each sample from the weights PDF.

where we have again allowed the freedom to use a different ϵ_{ω} from the ϵ parameter used in the force-of-change equation for the particle weights or for the nuisance parameters. We can directly evaluate $\frac{\partial \Delta_j^0}{\partial \omega}$ and $\frac{\partial \Delta_j^{\text{II}}}{\partial \omega}$ using a finite difference approximation, e.g. $\frac{\partial \Delta_j^0}{\partial \omega} = \frac{\Delta_j^0(\omega + \Delta\omega) - \Delta_j^0(\omega)}{\Delta\omega}$, by integrating the orbit starting from the previous time step in the two potentials characterized by frequencies ω and $\omega + \Delta\omega$ and comparing the $(\Delta_j^0, \Delta_j^{\text{II}})$ at the current time. In practice, we compute these finite differences with a $\Delta\omega$ large enough to give a substantial difference in $(z_i, v_{z,i})$ over the time step Δt . The parameter ϵ_{ω} should be small enough such that substantial changes to ω only happen on many orbital time-scales. In that case, the (non-resonant) orbits change adiabatically and the orbital struc-

Algorithm 3: MCMC sampling of potential parameters

```

/* To draw  $K$  MCMC samples  $\omega_k$ 
   characterizing potentials  $\Phi(z; \omega_k)$ ,
   given a set of particle weights  $\{w_i\}$ 
   and nuisance parameter  $z_{\odot}$ , for data
   points  $Y$  with uncertainty covariance
    $S$  */
// Average objective function for
// current  $\omega$ :
1  $\tilde{F} \leftarrow 0$ 
2 for  $m = 1, 2, \dots, M$  do
3    $(z_i, v_{z,i}) \leftarrow$  advance orbits by 1 step in  $\Phi(z; \omega)$ 
4    $\tilde{F} += F(z_i, v_{z,i} | \omega, z_{\odot}, w_i, Y, S) / M$ 
5 end
// MCMC sample using
// Metropolis-Hastings:
6 for  $k = 1, 2, \dots, K$  do
// Draw proposed  $\omega'$ :
7  $\omega' \sim Q(\omega' | \omega)$ 
// Adiabatically change  $\omega$  to  $\omega'$ :
8  $(z'_i, v'_{z,i}) \leftarrow (z_i, v_{z,i})$ 
9 for  $l = 1, 2, \dots, L$  do
10   $\omega_l \leftarrow \omega + (\omega' - \omega) l / L$ 
11   $(z'_i, v'_{z,i}) \leftarrow$  rewind orbits by 1 step in  $\Phi(z; \omega_l)$ 
12 end
13  $(z'_i, v'_{z,i}) \leftarrow$  advance orbits by  $L - M$  steps in
14    $\Phi(z; \omega')$ 
// Average objective function for  $\omega'$ :
15  $\tilde{F}' \leftarrow 0$ 
16 for  $m = 1, 2, \dots, M$  do
17   $(z'_i, v'_{z,i}) \leftarrow$  advance orbits by 1 step in
18    $\Phi(z; \omega')$ 
19   $\tilde{F}' += F(z'_i, v'_{z,i} | \omega', z_{\odot}, w_i, Y, S) / M$ 
20 end
// Accept/reject
21  $q \leftarrow \tilde{F}' - \tilde{F}$ 
22  $r \sim [0, 1]$ 
23 if  $\ln r < q$  then
24   $\omega \leftarrow \omega'$ 
25   $\tilde{F} \leftarrow \tilde{F}'$ 
26   $(z_i, v_{z,i}) \leftarrow (z'_i, v'_{z,i})$ 
27 else
28   $\omega \leftarrow \omega$ 
29 end
30  $\omega_k \leftarrow \omega$ 
31 end

```

ture corresponding to the M2M particles does not change between potentials. In certain applications, it may also be necessary to adiabatically change the potential to that, in this case, corresponding to $\omega + \Delta\omega$ when computing the finite difference, but we do not find this to be necessary here.

An example of fitting for ω is shown in Fig. 4 where we fit the same data as in the example described in Section 3.2, but now also fitting ω (while keeping z_{\odot} fixed to its true value). We keep the optimization parameters for the particle weights the same as in Section 3.2, but use $\epsilon_{\omega} = 10^{-3}$. We compute the finite difference using equation (35) with $\Delta\omega = 0.3$, and we only update ω every 10 time steps (and we therefore compute the finite difference using a time step $\Delta t = 10$ times the basic stepsize). We start at an initial

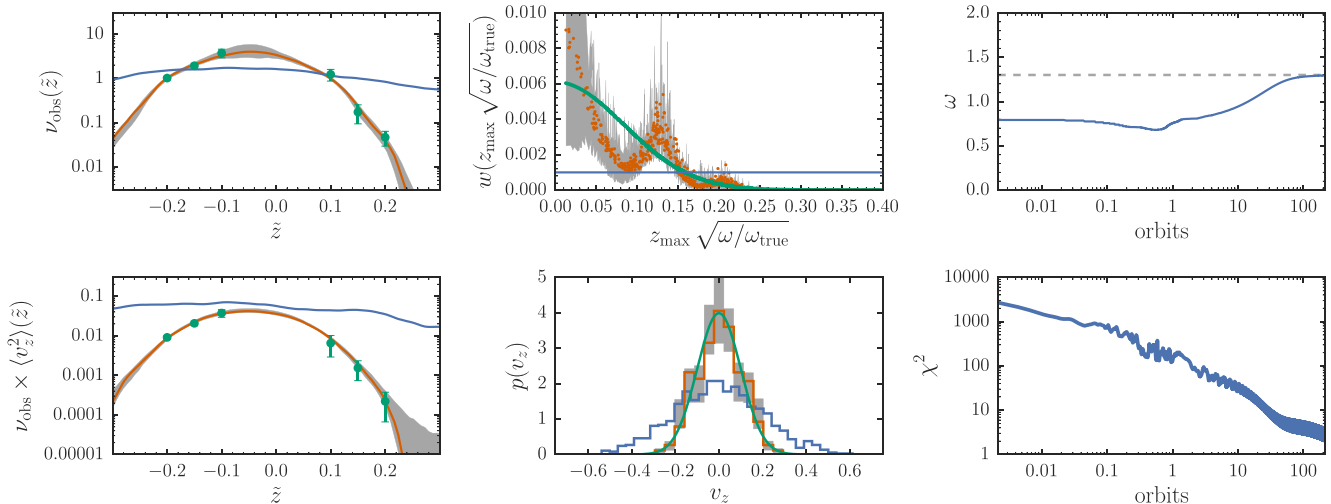


Figure 4. M2M for the parameters of the gravitational potential. Like Fig. 1, except that we also fit for the HO potential’s frequency ω using the force of change for ω and we sample the uncertainty in both the particle weights and ω . Because z_{\max} is not conserved when changing ω , we plot the weights as a function of $z_{\max} \sqrt{\omega}$, which is proportional to the square root of the action, which is approximately conserved during the M2M fit and sampling. The top right-hand panel demonstrates how ω converges during the joint M2M optimization of the particle weights and ω . We find that $\omega = 1.32 \pm 0.08$, in good agreement with its true value of 1.3, indicated by the dashed grey line in the top right-hand panel.

guess $\omega = 0.8$ and the fit converges to $\omega = 1.297$, close to the true value ($\omega = 1.3$).

Like for nuisance parameters, we can sample the joint posterior PDF for the particle weights and the potential parameters, in this case ω , using Metropolis-Hastings-within-Gibbs. The full MCMC sampler including the nuisance parameter z_{\odot} is then given by

$$(a) \quad w_i \sim p(w_i | z_{\odot}, \omega, \text{observations}), \quad [\text{Algorithm 1}], \quad (36)$$

$$(b) \quad z_{\odot} \sim p(z_{\odot} | \omega, w_i, \text{observations}), \quad [\text{Algorithm 2}], \quad (37)$$

$$(c) \quad \omega \sim p(\omega | z_{\odot}, w_i, \text{observations}), \quad [\text{Algorithm 3}], \quad (38)$$

where we again sample particle weights in the (a) step using the data-resampling technique of Section 4 (using $K = 1$ to draw a single particle-weight sample) and sample z_{\odot} and ω in steps (b) and (c) using a MH update using the objective function as the log posterior PDF, presented in detail in Algorithms 2 and 3. Like for the nuisance parameters on their own, we average the objective function in steps (b) and (c) over about one orbital period. Rather than simply changing the potential abruptly from a frequency ω to a proposal ω' for the likelihood evaluation in step (c), we adiabatically change the potential parameter from its current value to its proposed value before evaluating the likelihood (lines 8–12 in Algorithm 3). We perform this adiabatic change by integrating backwards in time and then partially integrating forwards in time, in such a way that the subsequent likelihood evaluation would use the exact same orbital trajectories if ω were not changed (line 13 in Algorithm 3). This reduces the noise from the particle distribution in the likelihood evaluation. We can again improve mixing in the MCMC chain by performing multiple MH steps (b) and (c) for each particle-weight sample ($K > 1$ in Algorithms 2 and 3, not necessarily equal) and keeping only the final ω sample in each step (c). The adiabatic change of the potential is important for maintaining the reversibility of the MCMC chain. If the potential is changed non-adiabatically, orbits differ when revisiting the same potential $\Phi(z; \omega)$ and the likelihood of a given set of particle weights is then different at later times. This does not happen when the potential is changed

adiabatically, because the nature of the orbits represented by the M2M particles does not change.

We apply this MCMC algorithm to sample the uncertainty distribution of the particle weights and ω given the mock data, fixing z_{\odot} to its true value [that is, skipping step (b)]. In step (c), we use a Gaussian proposal with standard deviation $\sigma_{\omega} = 0.2$ and again perform 10^5 M2M optimization time steps in step (a) and 20 MH steps for each step (a). We adiabatically change the potential using $L = 10000$ steps – or about 20 orbital periods – and average the objective function using 1000 orbital time steps. The MCMC chain is started at the best-fitting ω in the M2M optimization described above. The behaviour of the MCMC chain is displayed in Fig. 5. The MH acceptance ratio for the ω steps in the chain is 0.37. The chain is again well mixed and has a short correlation length.

From the MCMC samples, we find that the mock data constrain $\omega = 1.32 \pm 0.08$. We can compare this to the standard M2M procedure, where the PDF for ω is approximated using the best-fitting particle weights for each trial ω . This gives $\omega = 1.31 \pm 0.08$, similar to the MCMC result. If we assume that the DF is isothermal and marginalize over the amplitude and velocity dispersion of this isothermal DF, we find $\omega = 1.19 \pm 0.07$. All of these are consistent with the true value $\omega_{\text{true}} = 1.3$. That the isothermal DF gives a different best-fitting ω than the M2M modelling is unsurprising, because it fits a different functional shape to the density and velocity constraints: the best-fitting M2M DF is close to, but not exactly isothermal.

As a final test problem, we fit the particle weights, nuisance parameter z_{\odot} , and the potential parameter ω simultaneously to the mock data and then perform full MCMC sampling using steps (a) through (c) above. For the optimization part, we use $(\epsilon, \epsilon_{\odot}, \epsilon_{\omega}) = (10^{-3.5}, 10^{-6}, 10^{-3})$ and integrate for 3×10^5 time steps, again updating ω only every 10 time steps. Otherwise the set-up is the same as above. We use the best-fitting (z_{\odot}, ω) as the initial condition for MCMC sampling. In the MCMC sampling, we use 10^5 optimization time steps in step (a) of the algorithm and we average the likelihood using 500 steps for sampling z_{\odot} and using 1000 steps for sampling ω and again adiabatically change the

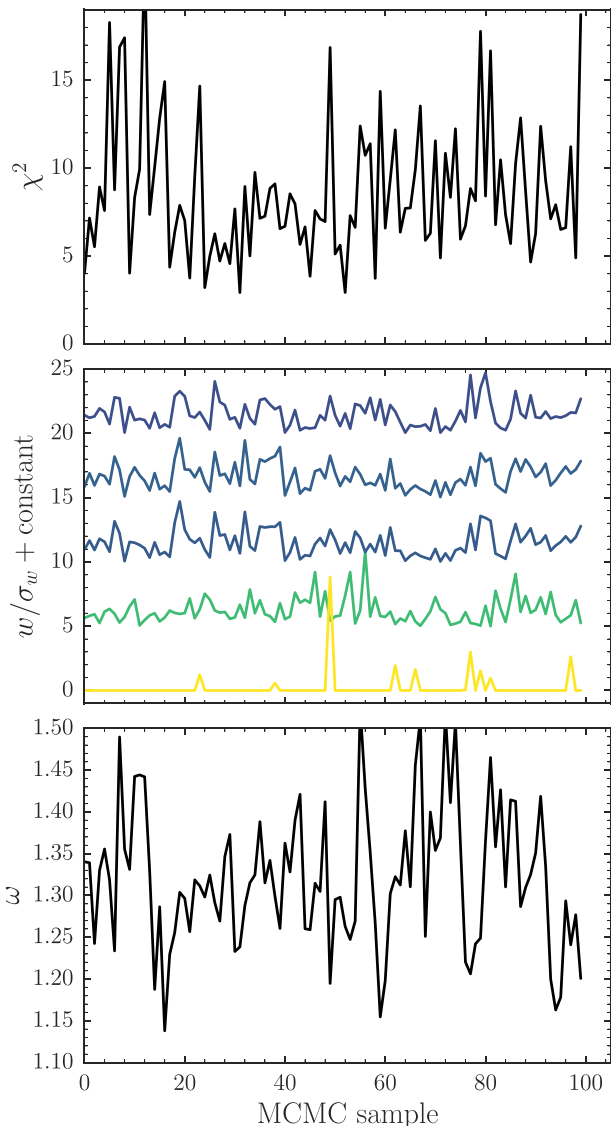


Figure 5. MCMC sampling of the particle weights and ω . This figure demonstrates the MCMC chain of 100 samples from the uncertainty distribution of the particle weights and ω constrained by the mock data. The top panel displays the behaviour of χ^2 , the middle panel that of five random particle weights (normalized by the standard deviation of their samples, colour-coded by $z_{\max}\sqrt{\omega}$) and the bottom panel shows the ω samples. The chain has a small correlation length, because we perform 20 MH steps for ω for each sample from the weights PDF.

frequency in MH steps over 10 000 time steps. The result is shown in Fig. 6. The parameters z_{\odot} and ω converge to best-fitting values of $z_{\odot} = 0.0530$ and $\omega = 1.27$. From the MCMC chain, we find that $z_{\odot} = 0.053 \pm 0.005$ and $\omega = 1.316 \pm 0.085$, similar to the analyses where one of these was kept fixed at its true value.

7 APPLICATION TO GAIA DR1

As a first real-data application of the M2M extensions described in this paper, we model the vertical dynamics of F-type dwarfs using data from the *Gaia* DR1 *Tycho-Gaia Astrometric Solution* (TGAS; Gaia Collaboration et al. 2016; Lindegren et al. 2016). We stress that the point of this application is only to illustrate the performance

of the new M2M method on real data. Because we use the same HO model for the potential, which is not a fully realistic model for the vertical potential near the Sun, the parameter constraints that we derive below cannot be easily translated into a constraint on the local mass distribution and we do not attempt to put any constraint on the local gravitational potential from this modelling.

We define F-type dwarfs as those with near-infrared $J - K_s$ in the range of $0.143 < J - K_s < 0.3$. Bovy (2017) have measured the vertical stellar density profiles for different sub-types of F dwarfs (e.g. F0V) from the TGAS data, correcting for the selection biases inherent in the TGAS data. We use similar measurements of the vertical stellar density of all F-type dwarfs (F0V through F9V), defined as the combination of all of the sub-types considered by Bovy (2017). These density measurements cover the range of $-400 \text{ pc} \leq \tilde{z} \leq 400 \text{ pc}$ in 25 pc wide bins and are shown in the top left-hand panel of Fig. 7; \tilde{z} is the vertical height as measured from the Sun’s position, similar to the toy example above.

We also measure the vertical velocity dispersion as a function of vertical height from the TGAS data. For this, we select 103 603 F-type dwarfs using the same colour and magnitude cuts as in Bovy (2017) and requiring relative parallax uncertainties less than 10 per cent. These data provide us with $(v_{\alpha}, v_{\delta}) = (\mu_{\alpha}\cos\delta/\varpi, \mu_{\delta}/\varpi)$, where ϖ is the parallax and $(\mu_{\alpha}\cos\delta, \mu_{\delta})$ are the proper motion components in right ascension and declination. We obtain the uncertainty covariance for each data point by Monte Carlo sampling 10 001 points from the correlated uncertainty covariance for the parallax and proper motions. We fit the v_z distribution from these data by deconvolving the observed two-dimensional distribution of (v_{α}, v_{δ}) using a mixture-of-Gaussians model of the velocity distribution in rectangular Galactic coordinates $(v_x, v_y, v_z) = (U, V, W)$ using the extreme-deconvolution (XD; Bovy, Hogg & Roweis 2011) algorithm (see Bovy, Hogg & Roweis 2009 for a similar application to *Hipparcos* data). Because we are only interested in the v_z distribution and are not interested in the details of this distribution, we use only two Gaussians. We fit this model in 25 pc bins covering $-200 \text{ pc} < \tilde{z} < 200 \text{ pc}$ and extract σ_z^2 . Outside of this \tilde{z} range, the data are too few and the proper motions constrain v_z too little to provide a useful measurement of σ_z^2 . We obtain uncertainties on these σ_z^2 using 200 bootstrap resamplings. In the context of our modelling, we use these σ_z^2 measurements as a stand-in for $\langle v_z^2 \rangle$, that is, we assume that these have been corrected for the solar motion. In principle, we could marginalize over the solar motion in the same way as we marginalize over the solar position, but for the purpose of this illustration, we will assume that the correction for the solar motion is perfect. These data are shown in the bottom left-hand panel of Fig. 7.

We thus model the density $\nu_{\text{obs}}(\tilde{z})$ and mean-squared velocity $\langle v_z^2 \rangle$. The latter is different from the observable $\nu_{\text{obs}}\langle v_z^2 \rangle(\tilde{z})$ that we have considered so far and requires us to write down the various forces of change for the particle weights, z_{\odot} , and ω for the $\langle v_z^2 \rangle$ observable. These forces of change are similar to the earlier expressions, although they are slightly more complicated because the weights enter into the normalization in the denominator. We give the relevant expressions in the Appendix A. Because the particle weights enter into the denominator of each $\langle v_z^2 \rangle$ measurement, the model is no longer linear in the particle weights and the procedure for sampling the uncertainty distribution of the particle weights is no longer strictly correct. However, for large numbers of particle weights, the normalization factor is only slightly affected by each individual particle and the model is still close to linear in the particle weights. We have run all of the mock tests described in the previous sections for a mock data set consisting of density

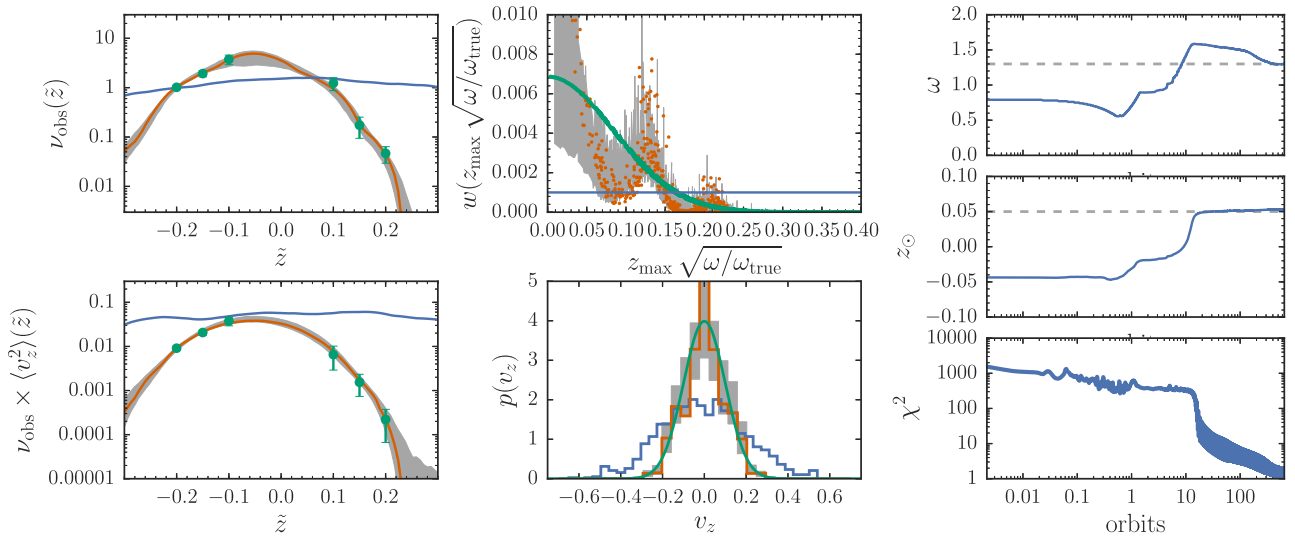


Figure 6. Full probabilistic M2M modelling. Like Fig. 1, except that we also fit for the HO potential’s frequency ω as well as the Sun’s height above the plane z_{\odot} . We MCMC sample from the joint PDF for the particle weights, z_{\odot} , and ω . As in Fig. 4, because z_{max} is not conserved when changing ω , we plot the weights as a function of $z_{\text{max}} \sqrt{\omega}$, which is proportional to the square root of the action. The top right-hand and middle panels demonstrate how ω and z_{\odot} , respectively, converge during the joint M2M optimization of the particle weights, ω , and z_{\odot} . We find that $\omega = 1.316 \pm 0.085$ and $z_{\odot} = 0.053 \pm 0.005$ in good agreement with their true values of $\omega_{\text{true}} = 1.3$ and $z_{\odot, \text{true}} = 0.05$ (dashed lines in the right-hand panels).

and $\langle v_z^2 \rangle$ measurements and find that the method proposed here still works well. We thus apply it as is to the *TGAS* data.

We use 10 000 N -body particles and start from a HO potential with a frequency of $100 \text{ km s}^{-1} \text{ kpc}^{-1}$, an isothermal DF with $\sigma = 12 \text{ km s}^{-1}$, and a solar offset of $z_{\odot} = 25 \text{ pc}$. We use a kernel width of 35 pc. We then optimize the values of the particle weights, z_{\odot} , and the frequency ω using the observed *TGAS* data using 30 000 steps with $\Delta t \approx 0.1 \text{ Myr}$ or a total time $\approx 3 \text{ Gyr}$. We use $\epsilon = 10^{-5.5}$, $\epsilon_{\odot} = 10^{-5}$, and $\epsilon_{\omega} = 100$ and only update ω every 10 steps using $\Delta\omega = 30 \text{ km s}^{-1} \text{ kpc}^{-1}$. We use a flat prior, $\mu = 0$.

The result from the M2M optimization is displayed in Fig. 7. The M2M optimization quickly converges to a well-constrained DF with $z_{\odot} = -0.3 \text{ pc}$ and $\omega = 69.8 \text{ km s}^{-1} \text{ kpc}^{-1}$. We run the MCMC algorithm for sampling the particle weights, z_{\odot} , and ω using 30 000 M2M optimization time steps for sampling the weights and using a proposal distribution for z_{\odot} with width $\sigma_{z_{\odot}} = 7 \text{ pc}$ and a proposal for ω with width $\sigma_{\omega} = 3 \text{ km s}^{-1} \text{ kpc}^{-1}$. We use 500 steps to average the objective function for the MH steps for z_{\odot} and 1000 steps for ω , again changing ω to proposed values using 10 000 steps. We obtain 20 MH samples for z_{\odot} and 10 MH samples for ω for each sample from the uncertainty distribution for the particle weights. The MH acceptance fraction for z_{\odot} and ω was 0.25 and 0.27, respectively.

The resulting uncertainty in the observed density and $\langle v_z^2 \rangle$ as well as that in the inferred DF is shown in Fig. 7. Because the density is so well measured, the uncertainty in the model density is barely visible, but the uncertainty in the kinematics is larger. The DF becomes uncertain at large z_{max} , but is well determined for orbits that stay closer to the plane. Marginalizing over the uncertainty in the DF, we find that $z_{\odot} = -1 \pm 3 \text{ pc}$ and $\omega = 69.1 \pm 1.1 \text{ km s}^{-1} \text{ kpc}^{-1}$.

The HO potential fits the data that we chose to model well. This is surprising, because the local vertical potential should be quite different from a constant density model (the HO model) over the 800 pc range over which we have observed the density. The HO model is able to fit the density data by having a large, high-energy component in the DF, that is, the peak at $z_{\text{max}} \approx 0.3 \text{ kpc}$ in the top

middle panel in Fig. 7. This leads to two observable consequences in other panels of this figure: the velocity dispersion increases for $|z| \gtrsim 150 \text{ pc}$ (bottom left-hand panel) and the local velocity distribution should display a wide, high-velocity tail. An inspection of the *TGAS* F-star kinematics close to the plane where the vertical velocity is approximately given by the vertical component of the proper motion shows that such a high-velocity tail is absent in the observations (see also Holmberg & Flynn 2000). This means that the HO potential is *not* a good model for the local vertical potential. Therefore, we do not compare our constraint on ω to previous determinations of the local gravitational potential (e.g. Holmberg & Flynn 2000) or interpret our measurement of z_{\odot} , which may be affected by the model for the potential. Still, it is promising that the novel M2M algorithm proposed in this paper works reasonably well with the observational data with realistic uncertainties. We defer a more realistic treatment of the vertical potential to future work.

8 DISCUSSION

In the previous sections, we have introduced various extensions of the basic M2M method that are crucial to applying this method to model observational data. Here, we discuss the formal assumptions and underpinnings of the sampling methods in more detail, comment on some aspects of the method further, and describe other extensions and improvements that could be made.

8.1 On interpreting the M2M objective function as a PDF

The algorithm for sampling the uncertainty distribution of the particle weights and the MCMC algorithms for sampling nuisance and potential parameters depend on our assumption that we can interpret the M2M objective function as the logarithm of a PDF for the parameters. However, the M2M objective function, defined in equation (10), is not a static function, but fluctuates as the M2M particles orbit, even when all parameters are held fixed. Thus, the interpretation of the M2M objective function as a PDF is not

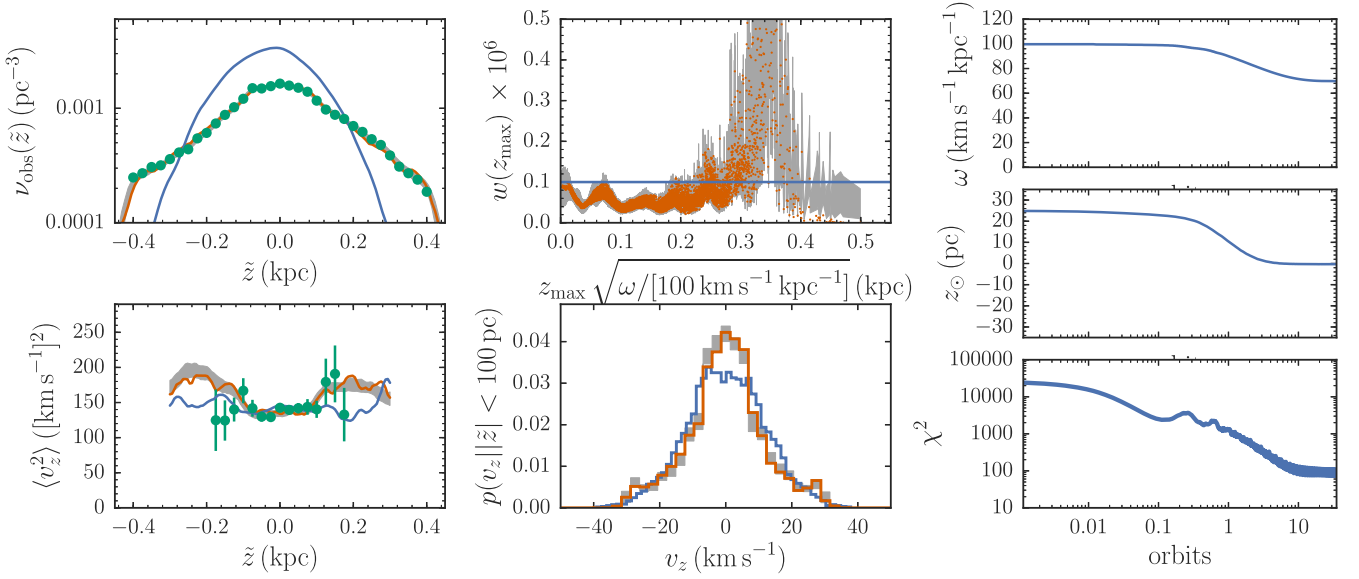


Figure 7. Probabilistic M2M modelling of the vertical dynamics of F-type dwarfs in *Gaia* DR1. Panels are the same as in Fig. 6, except that we model the mean-squared velocity directly rather than the density-weighted mean-squared velocity (bottom left-hand panel) and we display the local velocity distribution in the bottom middle panel. We can successfully model the density and the velocity dispersion of F-type dwarfs in a simple harmonic oscillator potential with $\omega = 69.1 \pm 1.1 \text{ km s}^{-1} \text{ kpc}^{-1}$, but this predicts the existence of an extended tail in the local velocity distribution (bottom middle panel).

obvious. We argue now that when run properly, the M2M procedure optimizes and samples from a well-defined correct PDF.

M2M modelling can be seen as an approximation of Schwarzschild modelling. Schwarzschild modelling uses the same form of the objective function, except that the kernels that in M2M are evaluated for a snapshot are in Schwarzschild modelling averaged in time. The objective function in that case defines the logarithm of a well-defined static PDF. Malvido & Sellwood (2015) have shown that M2M optimization is formally equivalent to Schwarzschild optimization in the limit of large times, large N , and small ϵ . Therefore, the basic M2M optimization procedure in fact optimizes a well-defined static objective function if the optimization is performed sufficiently slowly, that is, over a long enough time and with small ϵ . Moreover, our proposed sampling procedure for the uncertainty in the particle weights also optimizes the same objective function and thus effectively samples the well-defined Schwarzschild PDF. To the extent that the objective function is convex (exactly so for the objective function in our mock example above when no smoothing is applied), there is also no danger of optimizing to a local maximum.

To sample parameters other than the particle weights, we have introduced MH algorithms that use the averaged objective function as the logarithm of the PDF. The correct objective function is once again the equivalent Schwarzschild objective function. The question is in what limit these two are equivalent. For a single observation Y , we can schematically write down the contribution to χ^2 as (ignoring the observational uncertainty in the denominator)

$$\chi_{\text{M2M}}^2 = \left(\sum_i w_i K_i - Y \right)^2, \quad (39)$$

where K_i are the relevant kernel functions. The equivalent Schwarzschild form of this equation is

$$\chi_{\text{Schwarzschild}}^2 = \left(\sum_i w_i \langle K_i \rangle - Y \right)^2, \quad (40)$$

where $\langle K_i \rangle$ denotes the orbit-averaged kernel. Averaging equation (39) gives

$$\begin{aligned} \langle \chi_{\text{M2M}}^2 \rangle &= \sum_{i,j} w_i w_j \langle K_i K_j \rangle - 2Y \sum_i w_i \langle K_i \rangle + Y^2, \\ &= \sum_{i,j} w_i w_j \rho_{ij} \sigma_{K_i} \sigma_{K_j} + \chi_{\text{Schwarzschild}}^2, \end{aligned} \quad (41)$$

where ρ_{ij} is the correlation matrix of the orbital kernels: $\rho_{ij} \sigma_{K_i} \sigma_{K_j} = \langle (K_i - \langle K_i \rangle)(K_j - \langle K_j \rangle) \rangle$ and $\sigma_{K_i} = \sqrt{\langle (K_i - \langle K_i \rangle)^2 \rangle}$. Thus, for the orbit-averaged M2M objective function to be a good approximation to the Schwarzschild objective function, we need

$$\sum_{i,j} w_i w_j \rho_{ij} \sigma_{K_i} \sigma_{K_j} \ll \chi_{\text{Schwarzschild}}^2. \quad (42)$$

Orbits with very different trajectories have $\rho_{ij} \approx 0$, while orbits with similar trajectories have $w_i \approx w_j$ and $K_i \approx K_j$. Therefore, we can simplify the left-hand side of the previous equation to a sum over sets of orbits with similar trajectories

$$\sum_{i,j} w_i w_j \rho_{ij} \sigma_{K_i} \sigma_{K_j} \approx \sum_{\text{sets of orbits } i} w_i^2 \sigma_{K_i}^2 \sum_{j \text{ similar to } i} \rho_{ij}. \quad (43)$$

For a large enough number of M2M particles distributed randomly in orbital phase, $\sum_{j \text{ similar to } i} \rho_{ij} \approx 0$. Thus, if the M2M system consists of a large number N of particles with well-mixed phases, the averaged M2M objective function is a good approximation to the Schwarzschild objective function and can therefore be used as the logarithm of the PDF in a MH update.

8.2 On the approximate data-resampling technique for sampling the particle-weight PDF

As we already stressed in Section 4, the method for sampling the particle-weight PDF by resampling the data and obtaining the best-fitting weights for the resampled data is approximate in the

sense that it technically only applies in the case that the particle weights can take any value, both positive and negative. When the weights are required to be positive, the data-resampling technique will oversample the zero-weight edge of parameter space – oversample it relative to its correct proportion under the PDF. Because a non-parametric method such as M2M can only constrain the gravitational potential by excluding solutions that require particles to have negative weights – without this restriction, the data can always be perfectly fit – this is a matter of concern. Here, we provide some arguments that demonstrate that this is not a major problem.

There are two general statements that we can make about any M2M particle-weight PDF. The first is that for gravitational potentials that are close to the true potential, all orbits should be able to contribute non-negatively to the DF. Some orbital families may not exist and thus be constrained to have weights close to zero, but no orbits should require negative weights. This means in particular that we may assume that the mode of the PDF lies in the volume of \mathbb{R}^N where all weights are non-negative (if the weights are underconstrained, the mode is a trough that for the correct potential will extend into this volume; we will continue to use ‘mode’ below, but this always includes the ‘trough’ case as well). A second general property of the M2M PDF for the weights is that it is log convex if the uncertainties are Gaussian and the predicted observables depend linearly on the particle weights. Combined with the fact that the mode of the PDF has all weights non-negative for gravitational potentials close to the true one, this implies that a significant amount of probability mass has non-negative weights.

We use the data-resampling technique for sampling the particle-weight PDF for two purposes: (i) for sampling the uncertainty in the particle DF when investigating the DF and (ii) to marginalize over the uncertainty in the particle DF when constraining the gravitational potential or any nuisance parameters. When using the data-resampling technique for studying the uncertainty in the orbital DF, presumably this will be done for a reasonable gravitational potential (otherwise one would adjust the potential first). The data-resampling technique will oversample weights exactly equal to zero for those orbits that are required to have negative weights to provide a good fit to the observations. The arguments in the previous paragraph demonstrate that this oversampling will only have a minor effect, because the mode is expected to have all non-negative particle weights.

In the full MCMC method for constraining nuisance parameters and the parameters of an external gravitational potential, the data-resampling technique is used to marginalize over the uncertainty in the particle DF. The MH steps are based on the ratio of the PDF for the current and the proposed set of parameters. That the data-resampling method oversamples weights being exactly equal to zero has a different effect based on whether the mode of the weights PDF has all non-negative weights (the case for well-fitting potentials) or whether it wants to have some negative weights (the case for badly-fitting potentials). In the first case, the likelihood of the model will be biased low by the oversampling, because there is a higher chance than there should be (under the correct PDF) of sampling the zero-weight edge, where the PDF has a lower value than in the $w_i > 0$ region (this follows from the log-convexity of the likelihood). In the second case, the likelihood of the model will be biased high, because the oversampled edge now has higher probability than the $w_i > 0$ region. Thus, the MH sampling is biased towards worse-fitting models and the data-resampling technique will therefore lead to conservative, somewhat inflated uncertainties on the parameters of the model. However, the more severe the oversampling of zero

weights, the worse-fitting the model must be in the first place, because the oversampling fraction – the fraction of samples that land on the zero-weight edge – gets larger the further the mode of the particle-weights PDF is from the non-negative part of parameter space (the oversampling fraction is equal to the integral of the PDF over the part of space where at least one w_i is negative when the w_i are allowed to take on negative values, divided by the integral of the PDF over all space; for this to be large for a Gaussian PDF, the mode needs to be deep in the negative-weight region). Therefore, the bias will be small, because even when evaluated too often with weights set to zero, these models will still have low likelihoods relative to better fitting models. We see no evidence of significantly inflated uncertainties on the nuisance or gravitational-potential parameters from the limitations of the data-resampling technique in any of the experiments in this paper.

8.3 Aspects of the method

Fixing the sum of the particle weights: From when M2M was first proposed, the sum of the particle weights has typically been fixed to a constant, under the assumption that the total mass of the modelled system is known. The standard M2M algorithm does not conserve the sum of the particle weights and the weights are typically simply renormalized after each update step. We have left the sum of the particle weights free to be constrained by the data, which is the appropriate thing to do because the total mass is never exactly known. This completely gets around the issue of the weights renormalization. Nevertheless, when setting up an N -body simulation using the M2M method one may want to constrain the total mass of the system to a specific value. A simple way to do this is to (a) define the particle weights to sum to one, in which case the weights cover the simplex embedded in N -dimensional space and (b) transform the simplex to a $N - 1$ dimensional space that covers all of \mathbb{R}^{N-1} . We discuss how to do this in Appendix B.

The importance of the integration method: We have sidestepped the issue of orbit integration in our example of a HO potential, because orbit integration can be performed analytically in this model. However, in more realistic models, orbits need to be integrated numerically with a small enough time step such that numerical errors are small. While typically not important in galactic dynamics, we recommend use of a symplectic integrator such as leapfrog for the following reason. When performing the entire sampling procedure, orbits can be integrated for thousands of dynamical times or more and small energy errors can accumulate to a significant fraction of the energy. Symplectic integrators with a small time step typically avoid growth of the energy error and faithfully follow the evolution of the dynamical system being modelled over many dynamical times.

Other MCMC samplers: In Algorithms 2 and 3, we have opted to use a simple MH sampler to sample the nuisance and potential parameters. However, in applications with more nuisance parameters or more complicated potential models, we may want to use an MCMC sampler that is less sensitive to the proposal step size or explores the PDF more efficiently. Of particular interest is Hamiltonian Monte Carlo (Duane et al. 1987; Neal 2011), which can make large strides across the PDF by making use of the derivatives of the PDF. For the nuisance parameters, we can straightforwardly compute these derivatives as the average force of change similar to how the average objective function is computed in algorithm 2.

8.4 Directions for future work

Self-consistently generating the potential: One attractive aspect of M2M modelling compared to other dynamical modelling approaches is that it is possible to let the M2M particles generate the gravitational force field or some part of it (e.g. Hunt & Kawata 2013). That is, when modelling the stellar kinematics of, for example, an external galaxy, one can run the M2M optimization while integrating the particles in the gravitational potential that they themselves generate (plus perhaps additional dark matter). If the particle weights are changed slowly enough, the potential changes adiabatically and if the number of particles is large enough, the potential, being the combination of many particles, changes on longer time-scales than the individual particle weights. Therefore, the arguments above that demonstrate that the M2M procedure optimizes a well-defined objective function still hold. The data-resampling method for obtaining uncertainties on the values of the particle weights should therefore still perform well. In the MCMC updates of the nuisance parameters, the particle weights are held fixed and the gravitational force generated by the particles should therefore not change much (it could be held fixed). In the MCMC updates for the parameters of the external gravitational potential, the orbits are changed adiabatically and the potential generated by the particles needs to be updated on the fly as well to preserve the consistency between the M2M particles and the potential.

Dynamical stability: When we do not demand that the M2M particles generate (part of) the gravitational potential, one can end up with a solution or an MCMC sample that is dynamically unstable. M2M modelling, by virtue of using particles, can easily add the constraint of dynamical stability after the fact, by using the set of particle weights for a given MCMC sample to initialize an N -body simulation and determining whether it is dynamically stable or not. Samples that are not stable could be rejected and pruned from the chain.

Priors on the particle weights: We have paid little attention to the penalization term (the prior) in the M2M objective function and set it to zero in all of our examples (corresponding to an improper, flat prior on the weights). While it is clear that we do have definite prior beliefs about the particle weights, these are not well expressed by the standard entropy-like M2M or Schwarzschild penalization terms in the objective function. These standard forms express the prior belief that the particle weights are close to a reference set of weights, but without any correlation between the weights of similar orbits. This is problematic when we want to sample the uncertainty distribution of the particle weights. Interpreting the standard penalization as the logarithm of a prior PDF and sampling from this prior PDF gives sets of particle weights in which similar orbits can have widely different weights. A better prior would express the fact that similar orbits have similar weights without necessarily having strong prior beliefs about the actual value of the weights. This could, for example, be done using a Gaussian process with a kernel function in the space of integrals of the motion. Alternatively, a local smoothing of the current set of particle weights could be substituted for the prior (Morganti & Gerhard 2012). One advantage of using a Gaussian process is that this would allow the prior to be taken into account in the data-resampling technique for sampling the uncertainty in the particle weights: we can ‘resample’ the mean of the prior applied in each optimization sequence similar to how each data point is resampled in this technique and this returns formally correct samples from the posterior PDF for the particle weights (as long as they are positive). For spherical or axisymmetric systems, integrals of the motion are available that can be used to evaluate the similarity of

orbits, but even in general time-independent systems, the energy could be used or one can construct other similarity functions.

Modelling multiple populations: In our mock example, we have assumed that only a single population of stars is being modelled. However, if density and kinematic measurements are available for different populations of stars, one could use the same set of particles with multiple weights associated with each particle, one for each stellar population. That is, suppose that we had modelled both F and G-type dwarfs in *Gaia* DR1 as an example, we could have used N particles with two weights for each particle, one for F-type stars and one for G-type stars. These weights can all be optimized simultaneously. More generally, if we have additional information such as overall metallicity Z , abundance ratios, or ages for stars, we can replace the particle weights w_i associated with each particle with parametrized functions, e.g. $w_i(Z)$, of these additional quantities and fit for the parameters of these functions. One particular attractive way of doing this is to represent these functions in terms of basis functions with free amplitude parameters, e.g. $w_i(Z) = \sum_k \alpha_{ik} \beta_k(Z)$ with $\beta_k(\cdot)$ a set of fixed basis functions. In this case, the observables remain linearly related to the parameters (α_{ik}) and the data-resampling technique for obtaining uncertainties on the particle weights then also applies to the amplitudes of the basis functions.

9 CONCLUSION

M2M modelling is one of the most promising dynamical-modelling methods for fitting observational constraints on relaxed stellar systems without making additional assumptions about the shape of the system’s DF. This generality is a prerequisite to making the most robust inferences regarding the stellar, baryonic and dark masses of stellar systems. M2M has been used successfully to model the dynamics of external galaxies (e.g. De Lorenzi et al. 2008) and of the bar-shaped inner Milky Way region (e.g. Portail et al. 2017). However, so far M2M models (or Schwarzschild models for that matter; Magorrian 2006) have not dealt with the massive degeneracies that necessarily accompany a DF model as flexible as that used in M2M. Because these degeneracies can have a large influence on the inferences about the gravitational potential made using M2M modelling, results obtained without taking the uncertainty in the particle distribution into account should be viewed with suspicion.

We have improved and extended the standard M2M algorithm for fitting observational data in various ways. First, we have shown that all parameters describing the system – particle weights, nuisance parameters, and the parameters of an external gravitational field – can be optimized simultaneously in the M2M optimization. This makes it much easier to fit M2M models to observational data, as only a single M2M run is necessary, no matter how complicated the nuisance parameters or external gravitational potential is.

Secondly, we have introduced algorithms to sample from the full posterior PDF that describes the uncertainty in the particle weights and the nuisance and gravitational-potential parameters. For the particle weights, which can be very numerous, this is done through a technique that resamples the data within its uncertainties and turns the sampling problem into an optimization problem. This technique is formally correct when the model is linear in the parameters and the data uncertainties are Gaussian. This is typically the case for M2M, where the model typically consists of kernels combined using linear weights, but we have also shown that this technique works when the data is the second moment of the velocity distribution. We sample the nuisance parameters and those describing the external gravitational field through a carefully designed MH MCMC

algorithm, where the averaged M2M objective function is used as the logarithm of the PDF and the potential is only ever changed adiabatically. Because of the tight connection between M2M and Schwarzschild modelling demonstrated in Malvido & Sellwood (2015) and in Section 8.1, these new techniques can also be useful for Schwarzschild modelling.

The full M2M method described in this paper allows for large-scale, fully-probabilistic modelling of observational data. It will be useful in future modelling of data on Milky Way stars (e.g. Hunt & Kawata 2014) and on external galaxies. As a first example, we have analysed data on the vertical density and kinematics of F-type dwarfs from *Gaia* DR1 in a simple harmonic-oscillator model for the local gravitational potential. We find that we can fit the data that we have chosen to model, but a more realistic model for the gravitational potential is necessary to make definitive statements about what these data imply about the local mass distribution.

All of the analysis in this paper can be reproduced using the code found at <https://github.com/jobovy/simple-m2m>.

ACKNOWLEDGEMENTS

We thank the anonymous referee for a constructive report. JB received support from the Natural Sciences and Engineering Research Council of Canada. JB also received partial support from an Alfred P. Sloan Fellowship and from the Simons Foundation. DK acknowledges the support of the UK's Science & Technology Facilities Council (STFC Grant ST/K000977/1 and ST/N000811/1). This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<http://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC; <http://www.cosmos.esa.int/web/gaia/dpac/consortium>).

REFERENCES

- Betancourt M., 2012, AIP Conf. Proc. Vol. 1443, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 31st International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. Am. Inst. Phys., New York, p. 157
- Binney J., 2010, MNRAS, 401, 2318
- Bovy J., 2017, MNRAS, 470, 1360
- Bovy J., Rix H.-W., 2013, ApJ, 779, 115
- Bovy J., Hogg D. W., Roweis S. T., 2009, ApJ, 700, 1794
- Bovy J., Hogg D. W., Roweis S. T., 2011, Ann. Appl. Stat., 5, 1657
- Cappellari M. et al., 2012, Nature, 484, 485
- Cappellari M. et al., 2013, MNRAS, 432, 1709
- De Lorenzi F., Debattista V. P., Gerhard O., Sambhus N., 2007, MNRAS, 376, 71
- De Lorenzi F., Gerhard O., Saglia R. P., Sambhus N., Debattista V. P., Pannella M., Méndez R. H., 2008, MNRAS, 385, 1729
- Dehnen W., 2009, MNRAS, 395, 1079
- Duane S., Kennedy A. D., Pendleton B. J., Roweth D., 1987, Phys. Lett. B, 195, 216
- Gaia Collaboration et al., 2016, A&A, 595, A2
- Hogg D. W., Bovy J., Lang D., 2010, ApJS, preprint (arXiv:1008.4686)
- Holmberg J., Flynn C., 2000, MNRAS, 313, 209
- Hunt J. A. S., Kawata D., 2013, MNRAS, 430, 1928
- Hunt J. A. S., Kawata D., 2014, MNRAS, 443, 2112
- Hunt J. A. S., Kawata D., Martel H., 2013, MNRAS, 432, 3062
- Kuijken K., Gilmore G., 1989, MNRAS, 239, 571
- Lindgren L. et al., 2016, A&A, 595, A4
- Long R. J., Mao S., 2010, MNRAS, 405, 301
- Magorrian J., 2006, MNRAS, 373, 425
- Magorrian J. et al., 1998, AJ, 115, 2285
- Malvido J. C., Sellwood J. A., 2015, MNRAS, 449, 2553

- Morganti L., Gerhard O., 2012, MNRAS, 422, 1571
- Neal R. M., 2011, in Brooks S., Gelman A., Jones G., Meng X.-L., eds, Handbook of Markov Chain Monte Carlo. CRC Press, Boca Raton, Florida.
- Portail M., Gerhard O., Wegg C., Ness M., 2017, MNRAS, 465, 1621
- Rix H.-W., de Zeeuw P. T., Cretton N., van der Marel R. P., Carollo C. M., 1997, ApJ, 488, 702
- Schwarzschild M., 1979, ApJ, 232, 236
- Schwarzschild M., 1993, ApJ, 409, 563
- Syer D., Tremaine S., 1996, MNRAS, 282, 223
- Trick W. H., Bovy J., Rix H.-W., 2016, ApJ, 830, 97

APPENDIX A: USING THE MEAN-SQUARED VELOCITY AS THE OBSERVABLE

Instead of the density-weighted mean-squared velocity shown in equation (16), we can use the mean-squared velocity, $\langle v_z^2 \rangle(\tilde{z}_j)$, itself as an observable. This allows us to use the velocity measurements from a subsample of the one used for the density measurement. The model mean-squared velocity is defined as

$$\langle v_z^2 \rangle(\tilde{z}_j) = \sum_i w_i v_{z,i}^2 K^0(|\tilde{z}_j + z_\odot - z_i|; h) / v_{v^2,j}, \quad (\text{A1})$$

where $v_{v^2,j} = \sum_i w_i K_j^0(z_i; h)$, corresponding to a choice of a kernel of $K_j^{v^2}(z_i, v_{z,i}) = v_{z,i}^2 K_j^0(z_i; h) / v_{v^2,j}$. Note that the denominator can be calculated using a different kernel (or kernel width) than the density itself (equation 14) and, therefore, we use $v_{v^2,j}$ which can be different from $v(\tilde{z}_j)$. Assuming that the $\langle v_z^2 \rangle(\tilde{z}_j)$ observations have a Gaussian error distribution with variance $\sigma_{v^2,j}^2$, the contribution to χ^2 from $\langle v_z^2 \rangle(\tilde{z}_j)$ is given by

$$\chi_{j,v^2}^2 = \left[\Delta_j^{v^2} / \sigma_{v^2,j} \right]^2 = \left(\langle v_z^2 \rangle(\tilde{z}_j) - \langle v_z^2 \rangle_j^{\text{obs}} \right)^2 / \sigma_{v^2,j}^2. \quad (\text{A2})$$

In this case, the contribution from $\langle v_z^2 \rangle(\tilde{z}_j)$ to the force of change for the particle weights becomes

$$-\frac{1}{2} \frac{\partial \chi_{j,v^2}^2}{\partial w_i} = -\frac{\Delta_j^{v^2} [v_{z,i}^2 - \langle v_z^2 \rangle(\tilde{z}_j)] K_j^0(z_i; h)}{\sigma_{v^2,j}^2 v_{v^2,j}}. \quad (\text{A3})$$

Similarly, the contribution from $\langle v_z^2 \rangle(\tilde{z}_j)$ to the force of change for z_\odot is

$$-\frac{1}{2} \frac{\partial \chi_{j,v^2}^2}{\partial z_\odot} = -\frac{\Delta_j^{v^2}}{\sigma_{v^2,j}^2 v_{v^2,j}} \times \sum_i w_i [v_{z,i}^2 - \langle v_z^2 \rangle(\tilde{z}_j)] \frac{dK_j^0(r; h)}{dr} \Big|_{|\tilde{z}_j + z_\odot - z_i|} \text{sign}(\tilde{z}_j + z_\odot - z_i). \quad (\text{A4})$$

The force-of-change for ω is again computed using a direct finite difference, similar to equation (35).

APPENDIX B: M2M ON THE SIMPLEX

If one wants to run M2M modelling under a hard constraint on the sum of the particle weights (e.g. if the total mass represented by the M2M particles is exactly known, as in setting up an N -body simulation), the standard M2M force-of-change-based algorithm fails because the update equations for the particle weights do not conserve the sum of the weights. No satisfactory solution of this problem has been proposed in the literature.

If the particle weights must sum to a constant value we can always redefine them such that they sum to one: $\sum_i w_i = 1$. The weights

are then constrained to be positive and to sum to one and they therefore define a $N - 1$ dimensional simplex embedded in \mathbb{R}^N . We can then rewrite the M2M algorithm in terms of a transformed set of variables y_i that cover all of \mathbb{R}^{N-1} and that parametrize the simplex. In this case, the particle weights always exactly sum to one and the algorithm cannot stray from this condition. Generically, such a transformation would require $\mathcal{O}(N^2)$ operations to compute the derivatives with respect to the y_i from those with respect to the w_i . Here, we propose a specific transformation that is simple to implement and for which the derivatives with respect to y_i can be computed in $\mathcal{O}(N)$ time. Transforming to the y_i is then a feasible method even for very large numbers of N -body particles.

The transformation from w_i to y_i is the combination of the following transformations (partially following Betancourt 2012)

$$x_i = 1 - \frac{w_i}{\prod_{k=1}^{i-1} x_k}, \quad (\text{B1})$$

$$y_i = \text{logit}(x_i) - \text{logit}(X_N), \quad (\text{B2})$$

where $\text{logit}(\cdot)$ is the log-odds function $\text{logit}(x) = \ln(x/[1-x])$ with the inverse $\text{logit}^{-1}(x) = 1/[1 + e^{-x}]$. X_N is a $N - 1$ dimensional vector with entries $[\frac{N-1}{N}, \frac{N-2}{N-1}, \dots, \frac{1}{2}]$, which causes the simplex with all particle weights equal to each other, $w_i = 1/N$, to be mapped to the zero vector. The inverse transformation is given by

$$x_i = \text{logit}^{-1}(y_i + \text{logit}(X_N)), \quad (\text{B3})$$

$$w_i = \left(\prod_{k=1}^{i-1} x_k \right) \cdot \begin{cases} 1 - x_i, & i < N \\ 1, & i = N \end{cases}. \quad (\text{B4})$$

This inverse transformation is straightforward to implement using vectorized operations, while the $w_i \rightarrow y_i$ transformation requires a loop to accumulate the product in the first line. The inverse transformation is the one that is relevant for evaluating the objective function during the running of the M2M algorithm. The $w_i \rightarrow y_i$ transformation is only needed at initialization (if the weights are initialized as $w_i = 1/N$, then the initial $y_i = 0$ for all i).

To run the M2M algorithm in terms of the y_i variables, we compute the derivative of the objective function F using the chain rule. The Jacobian $\partial w_k / \partial y_i$ of this transformation is (cf. Betancourt 2012)

$$\frac{\partial w_k}{\partial y_i} = \begin{cases} w_k (1 - x_i), & i < k \\ -w_i x_i, & i = k \\ 0, & i > k \end{cases}. \quad (\text{B5})$$

This is a lower-triangular matrix. The chain rule can then be simplified to

$$\frac{\partial F}{\partial y_i} = -x_i w_i \frac{\partial F}{\partial w_i} + (1 - x_i) \sum_{k=i+1}^N w_k \frac{\partial F}{\partial w_k}. \quad (\text{B6})$$

All $N - 1$ derivatives can be computed together in $\mathcal{O}(N)$ time by accumulating the sum.

If one interprets the objective function as the logarithm of a probability distribution, transforming to a new set of variables requires tracking the determinant of the Jacobian. Because the Jacobian is a lower-triangular matrix, its determinant is given by the product of the diagonal entries

$$\left| \frac{\partial w}{\partial y} \right| = \prod_{k=1}^{N-1} w_k x_k. \quad (\text{B7})$$

The derivative of the logarithm of the Jacobian with respect to y_i is given by

$$\frac{\partial}{\partial y_i} \ln \left| \frac{\partial w}{\partial y} \right| = (N - i)(1 - x_i) - x_i, \quad (\text{B8})$$

for $i = 1, \dots, N - 1$.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.