Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **468,** 4374–4388 (2017) doi:10.1093/mnras/stx765
Advance Access publication 2017 March 29

A hierarchical model for the ages of Galactic halo white dwarfs

Shijing Si,[1]★ David A. van Dyk,[1] Ted von Hippel,[2,3] Elliot Robinson,[4] Aaron Webster[2] and David Stenning[5]

[1]*Statistics Section, Department of Mathematics, Imperial College London, London SW7 2AZ, UK*
[2]*Center for Space and Atmospheric Research, Embry-Riddle Aeronautical University, Daytona Beach, USA*
[3]*Max Planck Institute for Astronomy, Königstuhl 17, D-69117 Heidelberg, Germany*
[4]*Argiope Technical Solutions, Florida, USA*
[5]*Institut d'astrophysique de Paris, Paris, France*

Accepted 2017 March 24. Received 2017 March 24; in original form 2016 August 12

ABSTRACT

In astrophysics, we often aim to estimate one or more parameters for each member object in a population and study the distribution of the fitted parameters across the population. In this paper, we develop novel methods that allow us to take advantage of existing software designed for such case-by-case analyses to simultaneously fit parameters of both the individual objects and the parameters that quantify their distribution across the population. Our methods are based on Bayesian hierarchical modelling that is known to produce parameter estimators for the individual objects that are on average closer to their true values than estimators based on case-by-case analyses. We verify this in the context of estimating ages of Galactic halo white dwarfs (WDs) via a series of simulation studies. Finally, we deploy our new techniques on optical and near-infrared photometry of 10 candidate halo WDs to obtain estimates of their ages along with an estimate of the mean age of Galactic halo WDs of $12.11^{+0.85}_{-0.86}$ Gyr. Although this sample is small, our technique lays the ground work for large-scale studies using data from the *Gaia* mission.

Key words: methods: statistical – white dwarfs – Galaxy: halo.

1 INTRODUCTION

1.1 White dwarfs and the galactic halo age

In the astrophysical hierarchical structure formation model, the present Galactic stellar halo is the remnant of mergers of multiple smaller galaxies (e.g. Freeman & Bland-Hawthorn 2002; Tumlinson et al. 2010; Scannapieco et al. 2011), most of which presumably formed some stars prior to merging, and some of which may have experienced, triggered or enhanced star formation during the merging process. The age distribution of Galactic halo stars encodes this process. Any perceptible age spread for the halo thus provides information on this complex star formation history.

At present, we understand the Galactic stellar halo largely through the properties of its globular clusters. These star clusters are typically grouped into a few categories: (i) those with thick disc kinematics and abundances, (ii) those with classical halo kinematics and abundances, (iii) the most distant population that is a few Gyr younger than the classical halo population and (iv) a few globular clusters such as M54 that are ascribed to known, merging systems, in this case the Sagittarius dwarf galaxy (see Forbes & Bridges 2010; Pawlowski, Pflamm-Altenburg & Kroupa 2012; Leaman, VandenBerg & Mendel 2013). Globular clusters in category (ii) appear consistent with the simple collapse picture of Eggen, Lynden-Bell & Sandage (1962), yet those of categories (iii) and (iv) argue for a more complex precursor plus merging picture. The newly appreciated complexity of multiple populations in many or perhaps all globular clusters (Gratton, Sneden & Carretta 2004) adds richness to this story, and may eventually help us better understand the earliest star formation environments.

Despite the tremendous amount we have learned from globular clusters, they are unlikely to elucidate the full star formation history of the Galactic halo because today's globular clusters represent a ∼1 per cent minority of halo stars. Without studying the age distribution of halo field stars, we do not know whether globular cluster ages are representative of the entire halo population. We do know that globular clusters span a narrower range in abundances than field halo stars (see Roederer et al. 2010; Yamada et al. 2013), so there is every reason to be suspicious that there is more to the story than globular clusters can themselves provide.

In order to determine the age distribution of the Galactic halo, we need to supplement the globular cluster-based story with ages for individual halo stars. This is not practical for the majority of main-sequence or red giant stars because of well-known degeneracies in their observable properties as a function of age. Gyrochronology (see Barnes 2010; Soderblom 2010) does hold some hope for determining the ages of individual stars, but this is unlikely to provide

★ E-mail: ss2913@ic.ac.uk

© 2017 The Authors
Published by Oxford University Press on behalf of the Royal Astronomical Society

precise ages for very old stars even after the technique sees considerably more development. Our best current hope for deriving the Galactic halo distribution is to determine the ages of halo field white dwarfs (WDs).

WDs have the advantages that they are the evolutionary end-state for the vast majority of stars and their physics is relatively well understood (Fontaine, Brassard & Bergeron 2001). A WD's surface temperature, along with its mass and atmospheric type, is intimately coupled to its cooling age, i.e. how long it has been a WD. The mass of a WD, along with an assumed initial–final mass relation (IFMR), provides the initial main-sequence mass of the star, which along with theoretical models provides the lifetime of the precursor star. Pulling all of this information together provides the total age for the WD. The weakest link in this chain is typically the IFMR. Yet fortunately, the uncertainty in the IFMR often has little effect on the *relative* ages of WDs, and thus the precision of any derived age distribution. Additionally, among the higher mass WDs, the uncertainty in the precursor ages can be reduced to a level where the IFMR uncertainties do not dominate uncertainties in the absolute WD ages.

While WDs provide all of these advantages for understanding stellar ages, the oldest are very faint, and thus few are known, with fewer still known with the kinematics of the Galactic halo. The paucity of data for these important objects will shortly become a bounty when *Gaia* both finds currently unknown WDs with halo kinematics and provides highly accurate and precise trigonometric parallaxes, which constrain WD surface areas and thus masses. The number of cool, halo WDs is uncertain by a factor of perhaps 5, and depending on the Galaxy model employed, Carrasco et al. (2014) calculate that *Gaia* will derive parallaxes for ∼60 or ∼350 single halo WDs with $T_{\rm eff} \leq 5000$. *Gaia* will measure parallaxes for more than 200 000 WDs with thick disc and disc kinematics.

We have developed a Bayesian statistical technique to derive the ages of individual WDs (van Dyk et al. 2009; O'Malley, von Hippel & van Dyk 2013) and intend to apply this to each WD for which *Gaia* obtains excellent parallaxes. Yet the number of halo WDs for which we can derive high-quality ages may still be modest, particularly because we also require accurate optical and near-infrared (near-IR) photometry. Because of the importance of the age distribution among halo stars, we have developed a hierarchical modelling technique to pool halo WDs and derive the posterior distributions of their ages.

1.2 Statistical analysis of a population of objects

In statistics, hierarchical models are viewed as the most efficient and principled technique to estimate the parameters of each object in a population (e.g. Gelman 2006b). In astrophysics, we often aim to estimate a set of parameters for each of a number of objects in a population, which motivates the application of hierarchical models. Noticeably, these models have gained popularity in astronomy mainly for two reasons. First, they provide an approach to combining multiple sources of information. For instance, Mandel et al. (2009) employed Bayesian hierarchical models to analyse the properties of Type Ia supernova light curves by using data from Peters Automated Infrared Imaging Telescope and the literature. Secondly, they generally produce estimates with less uncertainty. By combining information from Type Ia supernova light curves, March et al. (2014) and Shariff et al. (2016) illustrate how a hierarchical model can improve the estimates of cosmological parameters. Similarly, Licquia & Newman (2015) obtained improved estimates of several global properties of the Milky Way by using a hierarchical model to combine previous measurements from the literature.

However, fully modelling a population of objects within a hierarchical model requires substantial computational investment and often specialized computer code, especially for complicated problems. In this study, we develop novel methods to conveniently obtain the improved estimates available under a hierarchical model. While taking advantage of the existing code for case-by-case analyses, our methods simultaneously estimate parameters of the individual objects and parameters that describe their population distribution. Our methods are based on Bayesian hierarchical modelling that are known to produce estimators of parameters of the individual objects that are on average closer to their true values than estimators based on case-by-case analyses.

There are many possible applications of hierarchical models in astrophysics. In this paper, we focus on the analysis of a sample of candidate halo WDs. We perform a simulation study to illustrate the advantage of our approach over the commonly used case-by-case analysis in this setting. We find that approximately two thirds of the estimated WD ages are closer to their true age under the hierarchical model. Using optical and near-IR photometry of 10 candidates halo WDs, we simultaneously estimate their individual ages and the mean age of these halo WDs, the latter which we estimate as $12.11^{+0.85}_{-0.86}$ Gyr. Another application to the distance modulus of the Large Magellanic Cloud (LMC) is included in Appendix B as a pedagogical illustration of our methods in a simpler setting.

One of the primary benefits of our approach is that it takes advantages of the existing code that fits one object at a time. We only need to write wrapper code that calls the existing programs, see Si et al. (in preparation) for more details.

This saves substantial human capital that might otherwise be devoted to developing and coding a complex new algorithm. The power of this approach can be conceived of as coming from (i) an informative assumption, which is that all the objects belong to a population with a particular distribution of the parameters of the objects across the population, and (ii) that it otherwise is difficult to come up with a technique that can combine the individual results when they may have asymmetric posterior density functions.

The remainder of this paper is organized into five sections. We introduce hierarchical modelling and its statistical inference methods in Section 2. We present methods for case-by-case and hierarchical analyses of the ages of a group of WDs in Section 3. In Section 4, we use a simulation study to verify the advantages of the hierarchical approach. In Section 5, we apply both the case-by-case and our hierarchical model to 10 Galactic halo WDs, and then interpret the Galactic halo age in the context of known Milky Way ages. Section 6 summarizes the proposed methodology and our results. In Appendix A, we describe the statistical background of hierarchical models and explain why they tend to provide better estimates. We illustrate the application of hierarchical models and their advantageous statistical properties via the LMC example in Appendix B. Appendices C and D outline the computational algorithms we use to efficiently fit the hierarchical models.

2 HIERARCHICAL MODELLING

Suppose we observe a sample of objects from a population of astronomical sources, for example, the photometry of 10 WDs from the Galactic halo, and we wish to estimate a particular parameter or a set of parameters for each object. We refer to these as the *object-level parameters*. By virtue of the population, there is a distribution of these parameters across the population of objects.

This distribution can be described by another set of parameters that we refer to as the *population-level parameters*. Often we aim to estimate both the object-level parameters and population-level parameters. As we shall see however, even if we are only interested in the object-level parameters, they can be better estimated if we also consider their population distribution.

Hierarchical models (e.g. Gelman et al. 2014), also called random effect models, can be used to combine data from multiple objects in a single coherent statistical analysis. Potentially, this can lead to a more comprehensive understanding of the overall population of objects. Hierarchical models are widely used in many fields, spanning the medical, biological, social and physical sciences. Because they leverage a more comprehensive data set when fitting the object-level parameters, they tend to result in estimators that on average exhibit smaller errors (e.g. James & Stein 1961; Efron & Morris 1972; Carlin & Louis 2000; Morris & Lysy 2012). Because a property of these estimators is that they are 'shrunk' towards a common central value relative to those derived from the corresponding case-by-case analyses, they are often called *shrinkage estimators*. More details about shrinkage estimators appear in Appendix A.

A concise hierarchical model is

$$Y_i|\theta_i \sim N(\theta_i, \sigma),\ i = 1, 2, \ldots, n, \tag{1}$$

$$\theta_i \sim N(\gamma, \tau), \tag{2}$$

where $\boldsymbol{Y} = (Y_1, \ldots, Y_n)$ are observations, σ is the standard error of the observations, $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$ are objective-level parameters of interest, γ and τ are the unknown population-level mean and standard deviation parameters of a Gaussian distribution.[1]

Bayesian statistical methods use the conditional probability distribution of the unknown parameters given the observed data to represent uncertainty and generate parameter estimates and error bars. This conditional probability distribution is called the *posterior distribution* and in our notation written as $p(\gamma, \tau, \boldsymbol{\theta}|\boldsymbol{Y})$. To derive the posterior distribution via Bayes theorem requires us to specify a *prior distribution* that summarizes our knowledge about likely values of the unknown parameters having seen the data, see van Dyk et al. (2001) and Park, van Dyk & Siemiginowska (2008) for applications of Bayesian methods in the context of astrophysical analyses. Our prior distribution on $\boldsymbol{\theta}$ is given in equation (2), and we choose the non-informative prior distribution $p(\gamma, \tau) \propto 1$ for γ and τ, which is a standard choice in this setting (Gelman 2006a). Two commonly used Bayesian methods to fit the hierarchical model in equations (1) and (2) are the fully Bayesian (FB) and the empirical Bayes (EB) methods.

FB (e.g. Gelman et al. 2014) fits all of the unknown parameters via their joint posterior distribution

$$p(\gamma, \tau, \boldsymbol{\theta}|\boldsymbol{Y}) \propto p(\gamma, \tau) \prod_{i=1}^{n} p(\theta_i|\gamma, \tau) \prod_{i=1}^{n} p(Y_i|\theta_i). \tag{3}$$

Generally, we employ Markov chain Monte Carlo (MCMC) algorithms to obtain a sample from the posterior distribution, $p(\gamma, \tau, \boldsymbol{\theta}|\boldsymbol{Y})$. The MCMC sample can be used to (i) generate parameter estimates, e.g. by averaging the sampled parameter values, (ii) generate error bars, e.g. by computing percentiles of the sampled parameter values, and (iii) represent uncertainty, e.g. by plotting histograms or scatter plots of the sampled values. For intricate hierarchical models, however it may be computationally challenging to obtain a reasonable MCMC sample.

EB (e.g. Morris 1983; Casella 1985; Efron 1996) uses the data to first fit the parameters of the prior distribution in equation (2) and then given these fitted parameters infer parameters in equation (1) in the standard Bayesian way. Specifically, γ and τ are first estimated as $\hat{\gamma}$ and then $\hat{\tau}$ and the prior distribution $\theta_i \sim N(\hat{\gamma}, \hat{\tau})$ is used in a Bayesian analysis to estimate the θ_i. Thus, EB proceeds in two steps.

Step 1. Find the maximum a posterior (MAP) estimates of γ and τ by maximizing their joint posterior distribution, i.e.

$$(\hat{\gamma}, \hat{\tau}) = \arg\max_{\gamma,\tau} p(\gamma, \tau|\boldsymbol{Y}) = \arg\max_{\gamma,\tau} \int p(\gamma, \tau, \boldsymbol{\theta}|\boldsymbol{Y}) \mathrm{d}\boldsymbol{\theta}. \tag{4}$$

Step 2. Use $N(\hat{\gamma}, \hat{\tau})$ as the prior distribution for θ_i, $i = 1, \ldots, n$ and estimate θ_i in the standard Bayesian way, i.e.

$$p(\theta_i|Y_i, \hat{\gamma}, \hat{\tau}) \propto p(Y_i|\theta_i) p(\theta_i|\hat{\gamma}, \hat{\tau}). \tag{5}$$

When applying the EB to fit a hierarchical model, it is possible that the estimate of the standard deviation τ is equal to 0, which leads to $\theta_1 = \cdots = \theta_n = \hat{\gamma}$. This is generally not a desirable result. We can avoid $\hat{\tau} = 0$ by using the transformations $\xi = \log \tau$ or $\delta = 1/\tau$ (e.g. Park et al. 2008; Gelman et al. 2014). We refer to EB implemented with these transformations as EB-log and EB-inv, respectively. Step 2 of EB-log and EB-inv remains exactly the same as that of EB, but Step 1 changes. Specifically, Step 1 of EB-log is

Step 1. Find the MAP estimates of γ and ξ by maximizing their joint posterior distribution, i.e.

$$(\hat{\gamma}, \hat{\xi}) = \arg\max_{\gamma,\xi} p(\gamma, \exp(\xi)|\boldsymbol{Y}) \exp(\xi)$$

and setting $\hat{\tau} = \exp(\hat{\xi})$, where $p(\cdot|\boldsymbol{Y})$ is the posterior distribution of γ and τ. Thus,

$$(\hat{\gamma}, \hat{\tau}) = \arg\max_{\gamma,\tau} p(\gamma, \tau|\boldsymbol{Y})\tau. \tag{6}$$

Comparing equation (6) with equation (4), the added τ in equation (6) prevents τ from being zero. Step 1 of EB-inv proceeds similarly.

3 ANALYSES FOR FIELD HALO WDS

Our model is based on obtaining photometric magnitudes for n WDs from the Galactic halo. We denote the l-dimensional observed photometric magnitudes for the ith WD by $\boldsymbol{X}_i$ and the known variance-covariance matrix of its measurement errors by $\boldsymbol{\Sigma}_i$. Our goal is to use $\boldsymbol{X}_i$ to estimate the age, distance modulus, metallicity and zero-age main-sequence (ZAMS) mass of the WD. Our WD model is specified in terms of the $\log_{10}$(age), distance modulus, metallicity and ZAMS mass of WDs, and we denote these parameters by A_i, D_i, Z_i and M_i for $i = 1, \ldots, n$, respectively. Because we are primarily interested in WD ages, we group the other stellar parameters into $\boldsymbol{\Theta}_i = (D_i, Z_i, M_i)$. Finally to simplify notation, we write $\boldsymbol{X} = (\boldsymbol{X}_i, \ldots, \boldsymbol{X}_n), A = (A_1, \ldots, A_n)$ and $\boldsymbol{\Theta} = (\boldsymbol{\Theta}_1, \ldots, \boldsymbol{\Theta}_n)$. Here we review a case-by-case analysis method for WDs and develop convenient approaches to obtain the hierarchical modelling fits with improved statistical properties.

3.1 Existing case-by-case analysis

The public-domain Bayesian software suite, Bayesian Analysis of Stellar Evolution with 9 parameters (BASE-9), allows one to precisely estimate cluster parameters based on photometry

[1] In this paper, we parametrize univariate Gaussian distributions in terms of their means and standard deviations. Generally, we write $Y|\theta \sim N(\mu, \sigma)$ to indicate that given θ, Y follows a Gaussian (or normal) distribution with mean μ and standard deviation σ.

Table 1. Distributions used to simulate the parameter values.

Parameters	Distributions
$\log_{10}$(age)	$N(\gamma = 10.09, \tau)$ truncated between 9.7 and 10.17
Distance modulus	$N(4.0, 2.5^2)$ truncated to interval (2.7, 5.7), i.e. 34–138 pc
Metallicity	$N(-1.5, 0.02)$
Mass	$\log_{10}(\text{Mass}) \sim N(-1.02, 0.677\,29)$ truncated to (0.8, 3.0)

(von Hippel et al. 2006; DeGennaro et al. 2009; van Dyk et al. 2009). We have applied BASE-9 to key open clusters (DeGennaro et al. 2009; Jeffery et al. 2011; Hills et al. 2015), extended BASE-9 to study mass-loss from the main sequence through the WD stage, the so-called IFMR (Stein et al. 2013), and have demonstrated that BASE-9 can derive the complex posterior age distributions for individual field WD stars (O'Malley et al. 2013).

In this paper, we focus on the development of BASE-9 for fitting the parameters of individual WD stars. BASE-9 employs a Bayesian approach to fit parameters. The statistical model underlying BASE-9 relates a WD's photometry to its parameters,

$$\boldsymbol{X}_i|A_i, \boldsymbol{\Theta}_i \sim N_l(G(A_i, \boldsymbol{\Theta}_i), \boldsymbol{\Sigma}_i), \tag{7}$$

where N_l represents an l-variate Gaussian distribution and $G(\cdot)$ represents the underlying astrophysical models that predict a star's photometric magnitudes as a function of its parameters. Specifically G combines models for the main sequence through red giant branch (e.g. Dotter et al. 2008) and the subsequent WD evolution (e.g. Bergeron, Wesemael & Beauchamp 1995; Montgomery et al. 1999).

The Bayesian approach employed by BASE-9 requires a joint prior density on $(A_i, \boldsymbol{\Theta}_i)$ for each WD. We assume that this prior can be factored into

$$\begin{aligned} p(A_i, \boldsymbol{\Theta}_i) = {} & p(A_i|\mu_{A_i}, \sigma_{A_i}) p(D_i|\mu_{D_i}, \sigma_{D_i}) \\ & \times p(Z_i|\mu_{Z_i}, \sigma_{Z_i}) p(M_i), \end{aligned} \tag{8}$$

where the individual prior distributions on age, distance modulus and metallicity $p(A_i|\mu_{A_i}, \sigma_{A_i})$, $p(D_i|\mu_{D_i}, \sigma_{D_i})$ and $p(Z_i|\mu_{Z_i}, \sigma_{Z_i})$ are normal densities each with its own prior mean (i.e. μ_{A_i}, μ_{D_i} and μ_{Z_i}) and standard deviation (i.e. σ_{A_i}, σ_{D_i} and σ_{Z_i}). When possible, these prior distributions are specified using external studies. The prior on the mass M_i is specified as the initial mass function (IMF) taken from Miller & Scalo (1979), i.e. $\log_{10}(M_i) \sim N(\mu = -1.02, \sigma = 0.677\,29)$. BASE-9 deploys an MCMC sampler to separately obtain an MCMC sample from each of the WD's joint posterior distributions,

$$\begin{aligned} p(A_i, \boldsymbol{\Theta}_i|\boldsymbol{X}_i) \propto {} & p(\boldsymbol{X}_i|A_i, \boldsymbol{\Theta}_i) p(A_i|\mu_{A_i}, \sigma_{A_i}) \\ & \times p(D_i|\mu_{D_i}, \sigma_{D_i}) p(Z_i|\mu_{Z_i}, \sigma_{Z_i}) p(M_i). \end{aligned} \tag{9}$$

In this manner, we can obtain case-by-case fits of A_i and $\boldsymbol{\Theta}_i$ for each WD using BASE-9.

In this paper for both the case-by-case and the hierarchical analysis, we obtain MCMC samples for most of the parameters. After we obtain a reasonable MCMC sample for $A_i, \boldsymbol{\Theta}_i, i = 1, \ldots, n$, we estimate these quantities and their 1σ error bars using the means and standard deviations of their MCMC samples, respectively. For example, letting $A_i^{(s)}, s = 1, \ldots, S$ be an MCMC sample for A_i of size S, after suitable burn-in (DeGennaro et al. 2009), the posterior mean and standard deviation of A_i are approximated by

$$\hat{A}_i = \sum_{s=1}^{S} A_i^{(s)}/S, \tag{10}$$

$$\hat{\sigma}_{A_i} = \sqrt{\sum_{s=1}^{S} (A_i^{(s)} - \hat{A}_i)^2/(S-1)}. \tag{11}$$

When the posterior distribution of the parameter A_i is highly asymmetric, its posterior mean and 1σ error bar may not be a good representation of the posterior distribution. In this case, we might instead compute the 68.3 per cent posterior interval of A_i as the range between the 15.87 per cent and 84.13 per cent quantiles of the MCMC sample.

3.2 Hierarchical modelling of a group of WDs

In this section, we embed the model in equation (7) into a hierarchical model for a sample of halo WDs,

$$\begin{aligned} \boldsymbol{X}_i|A_i, \boldsymbol{\Theta}_i &\sim N_l(G(A_i, \boldsymbol{\Theta}_i), \boldsymbol{\Sigma}_i), \\ A_i &\sim N(\gamma, \tau). \end{aligned} \tag{12}$$

In this hierarchical model, A_i, D_i, Z_i and M_i are the object-level parameters, while γ and τ are population-level parameters, the mean and standard deviation of the $\log_{10}$(age) of WDs in the Galactic halo. The assumption of a common population incorporating an age constraint is the source of the statistical shrinkage that we illustrate below. For the prior distributions of each $\boldsymbol{\Theta}_i$, we take the same strategy as in the case-by-case analysis in equation (8). For the population-level parameters γ and τ, we again choose the uninformative prior distribution, i.e. $p(\gamma, \tau) \propto 1$. The joint posterior distribution for parameters in the hierarchical model is

$$\begin{aligned} p(\gamma, \tau, A, \boldsymbol{\Theta}|\boldsymbol{X}) \propto {} & p(\gamma, \tau) \prod_{i=1}^{n} p(\boldsymbol{X}_i|A_i, \boldsymbol{\Theta}_i) p(A_i|\gamma, \tau) \\ & \times p(D_i|\mu_{D_i}, \sigma_{D_i}) p(Z_i|\mu_{Z_i}, \sigma_{Z_i}) p(M_i). \end{aligned} \tag{13}$$

3.2.1 *FB method*

The FB approach obtains an MCMC sample from the joint posterior distribution in equation (13). Here we employ a two-stage algorithm

Table 2. Comparing the statistical properties of the various shrinkage and case-by-case estimates under five simulated settings.

	EB		EB-log		EB-inv		FB		Case-by-case	
Simulation settings	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
I: $\tau = 0.05$, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$	0.028	0.038	0.029	0.042	0.033	0.052	0.032	0.047	0.084	0.18
II: $\tau = 0.03$, $\boldsymbol{\Sigma} = 0.8^2\boldsymbol{\Sigma}_0$	0.024	0.035	0.025	0.038	0.027	0.044	0.025	0.041	0.086	0.19
III: $\tau = 0.03$, $\boldsymbol{\Sigma} = 1.2^2\boldsymbol{\Sigma}_0$	0.023	0.031	0.024	0.034	0.026	0.037	0.025	0.035	0.099	0.20
IV: $\tau = 0.06$, $\boldsymbol{\Sigma} = 0.8^2\boldsymbol{\Sigma}_0$	0.038	0.063	0.038	0.070	0.040	0.074	0.034	0.049	0.12	0.26
V: $\tau = 0.06$, $\boldsymbol{\Sigma} = 1.2^2\boldsymbol{\Sigma}_0$	0.033	0.046	0.032	0.044	0.032	0.045	0.032	0.044	0.10	0.22

(Si et al., in preparation) to obtain the FB results. This algorithm takes advantage of the case-by-case samples in Section 3.1 and is easy to implement. A summary of the computational details of FB appears in Appendix C.

3.2.2 EB method

We also illustrate how to fit the hierarchical model in equation (12) with EB. First, the joint posterior distribution for γ and τ is calculated as

$$p(\gamma, \tau|\boldsymbol{X}) = \int \cdots \int p(\gamma, \tau, A, \boldsymbol{\Theta}|\boldsymbol{X}) \mathrm{d}A \mathrm{d}\boldsymbol{\Theta}. \quad (14)$$

The integration in equation (14) is $4 \times n$ dimensional, which is computationally challenging. To tackle this, we use the Monte Carlo expectation-maximization (MCEM) algorithm (see e.g. Dempster, Laird & Rubin 1977; Wei & Tanner 1990) to find the MAP estimates of γ and τ. To avoid estimating τ as zero when its (profile) posterior distribution is highly skewed (e.g. Park et al. 2008), we again implement EB-log ($\xi = \log \tau$) or EB-inv ($\delta = 1/\tau$). For EB-log, the joint posterior distribution of γ and ξ equals

$$p(\gamma, \exp(\xi)|\boldsymbol{X}) \exp(\xi),$$

where $p(\cdot|\boldsymbol{X})$ is the joint posterior distribution of γ and τ. The EB-log method proceeds in two steps.

Step 1. Deploy MCEM to obtain the MAP estimates of γ and ξ, and transform to γ and τ, i.e.

$$(\hat{\gamma}, \hat{\tau}) = \arg\max_{\gamma,\tau} p(\gamma, \tau|\boldsymbol{X})\tau. \quad (15)$$

For details of MCEM in this setting, see Appendix D.

Step 2. For WD $i = 1, \ldots, n$, we obtain an MCMC sample from

$$p(A_i, \boldsymbol{\Theta}_i|\boldsymbol{X}_i, \hat{\gamma}, \hat{\tau}) \propto p(\boldsymbol{X}_i|A_i, \boldsymbol{\Theta}_i) p(\boldsymbol{\Theta}_i) p(A_i|\hat{\gamma}, \hat{\tau})$$

using BASE-9.

EB-inv proceeds in a similar manner, but with equation (15) replaced with

$$(\hat{\gamma}, \hat{\tau}) = \arg\max_{\gamma,\tau} p(\gamma, \tau|\boldsymbol{X})\tau^2,$$

where $p(\cdot|\boldsymbol{X})$ is again the posterior distribution of γ and τ.

4 SIMULATION STUDY

To illustrate the performance of the various estimators of the object-level WD ages and the population-level parameters γ and τ, we perform a set of simulation studies. Because the relative advantage of the shrinkage estimates compared with the case-by-case estimates depends both on the precision of the case-by-case estimates and the degree of heterogeneity of the object-level parameters, we repeat the simulation study under five scenarios, each with different values of observation error matrix $\boldsymbol{\Sigma}$ and population standard deviation of $\log_{10}$(age), i.e. of τ. We simulate the parameters $\{A_i, D_i, T_i, M_i, i = 1, 2, \ldots, N_1\}$ for each group of WDs from the distributions in Table 1, where $\gamma = 10.09$ (12.30 Gyr) is the population mean and τ varies among the simulation settings given in Table 2. For consistency with the data analyses in Section 5, we simulate u, g, r, i, z, J, H, K magnitudes for all WDs. Using BASE-9 for each setting, we simulate $N_2 = 25$ replicate data sets, each composed of $N_1 = 10$ halo WDs. For each WD in every group, we generate its $\log_{10}$(age), distance modulus, metallicity and mass from distributions in Table 1, where τ is given in Table 2. The particular values and truncations in Tables 1 and 2 are chosen because they reflect plausible values for actual halo WDs.

We compute the empirical standard error for each simulated magnitude by averaging the errors from the observed halo WDs in Section 5, and we denote by $\boldsymbol{\Sigma}_0$ the variance matrix of observed magnitudes, i.e. the square of empirical standard errors for all eight magnitudes. Specifically, $\boldsymbol{\Sigma}_0$ is a diagonal matrix with diagonal elements equal to (0.304^2, 0.092^2, 0.027^2, 0.026^2, 0.068^2, 0.062^2, 0.086^2, 0.083^2).

For simplicity in each setting, all stars share the same diagonal observation variance, that is each $\boldsymbol{\Sigma}_i = \boldsymbol{\Sigma}$, $i = 1, 2, \ldots, N_1$. The observation error variances for five simulation settings are described in terms of $\boldsymbol{\Sigma}_0$ in Table 2. In the entire simulation study, we employ the Dotter et al. (2008) WD precursor models, Renedo et al. (2010) WD interior models, Bergeron et al. (1995) WD atmospheres and Williams, Bolte & Koester (2009) IFMR.

Subsequently, we recover parameters with multiple approaches: EB, EB-log, EB-inv, FB and the case-by-case analysis. We specify non-informative broad prior distributions on each $\boldsymbol{\Theta}_i$, namely $D_i \sim N(4.0, 2.4^2)$, $Z_i \sim N(-1.5, 0.5^2)$ and $\log_{10}(M_i) \sim N(-1.02, 0.677\,29)$. The case-by-case analyses require a prior distribution on each A_i and we use $p(A_i) \propto 1$. The hierarchical model in equation (12) requires priors on γ and τ, and we again use $p(\gamma, \tau) \propto 1$. We compare the case-by-case estimates with shrinkage estimates based on the hierarchical model. Results from the case-by-case analyses [obtained by fitting the model in equation (7)] are indicated with a superscript I (for 'individual') and those from hierarchical analyses [obtained by fitting to model in equation (12)] are indicated with an H.

We denote $\log_{10}$(age) of the ith simulated WD in the jth replicate group by A_{ij}. Using both the case-by-case and hierarchical analyses, we obtain MCMC samples of the parameters A_{ij}, $i = 1, 2, \ldots, N_1$, $j = 1, 2, \ldots, N_2$. We estimate A_{ij} by taking the MCMC sample mean as in equation (10) and denote the estimates based on case-by-case and hierarchical analyses by $\hat{A}^I_{ij}$ and $\hat{A}^H_{ij}$, respectively. We compute the absolute value of the error[2] of each estimator $\hat{A}_{ij}$ as

$$\mathrm{error}(\hat{A}_{ij}) = |\hat{A}_{ij} - A_{ij}|.$$

In our simulation study, we are mainly concerned with the difference between absolute errors from shrinkage and case-by-case estimates

$$\begin{aligned}\mathrm{Diff}(A_{ij}) &= \mathrm{error}(\hat{A}^I_{ij}) - \mathrm{error}(\hat{A}^H_{ij}) \\ &= |\hat{A}^I_{ij} - A_{ij}| - |\hat{A}^H_{ij} - A_{ij}|,\end{aligned}$$

which compares the prediction accuracy of the two methods. If $\mathrm{Diff}(A_{ij}) \geq 0$, then the absolute deviation of the case-by-case estimate of star A_{ij} is greater than that of the shrinkage estimate.

Fig. 1 compares the performance of shrinkage estimates under simulated setting I ($\tau = 0.05$, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$). The corresponding summary plots for the other simulation settings are similar and appear in Figs 6–9 (the online supplement). The histograms in Fig. 1 demonstrate that the estimates of γ and τ are generally close to their true values (thick dashed red lines). Under all five settings, however, for some replicate groups of halo WDs, EB produces estimates of τ equal to 0, see the first row, middle panel of Fig. 1. (This phenomenon is fully discussed in Appendix B and Fig. B1.) In these cases, the shrinkage estimates of the age of each WD in these are equal, which potentially leads to large errors. As we mentioned in Section 3.2.2, this highlights a difficulty with EB, and demonstrates

[2] We use this term to refer to the absolute value of the error.

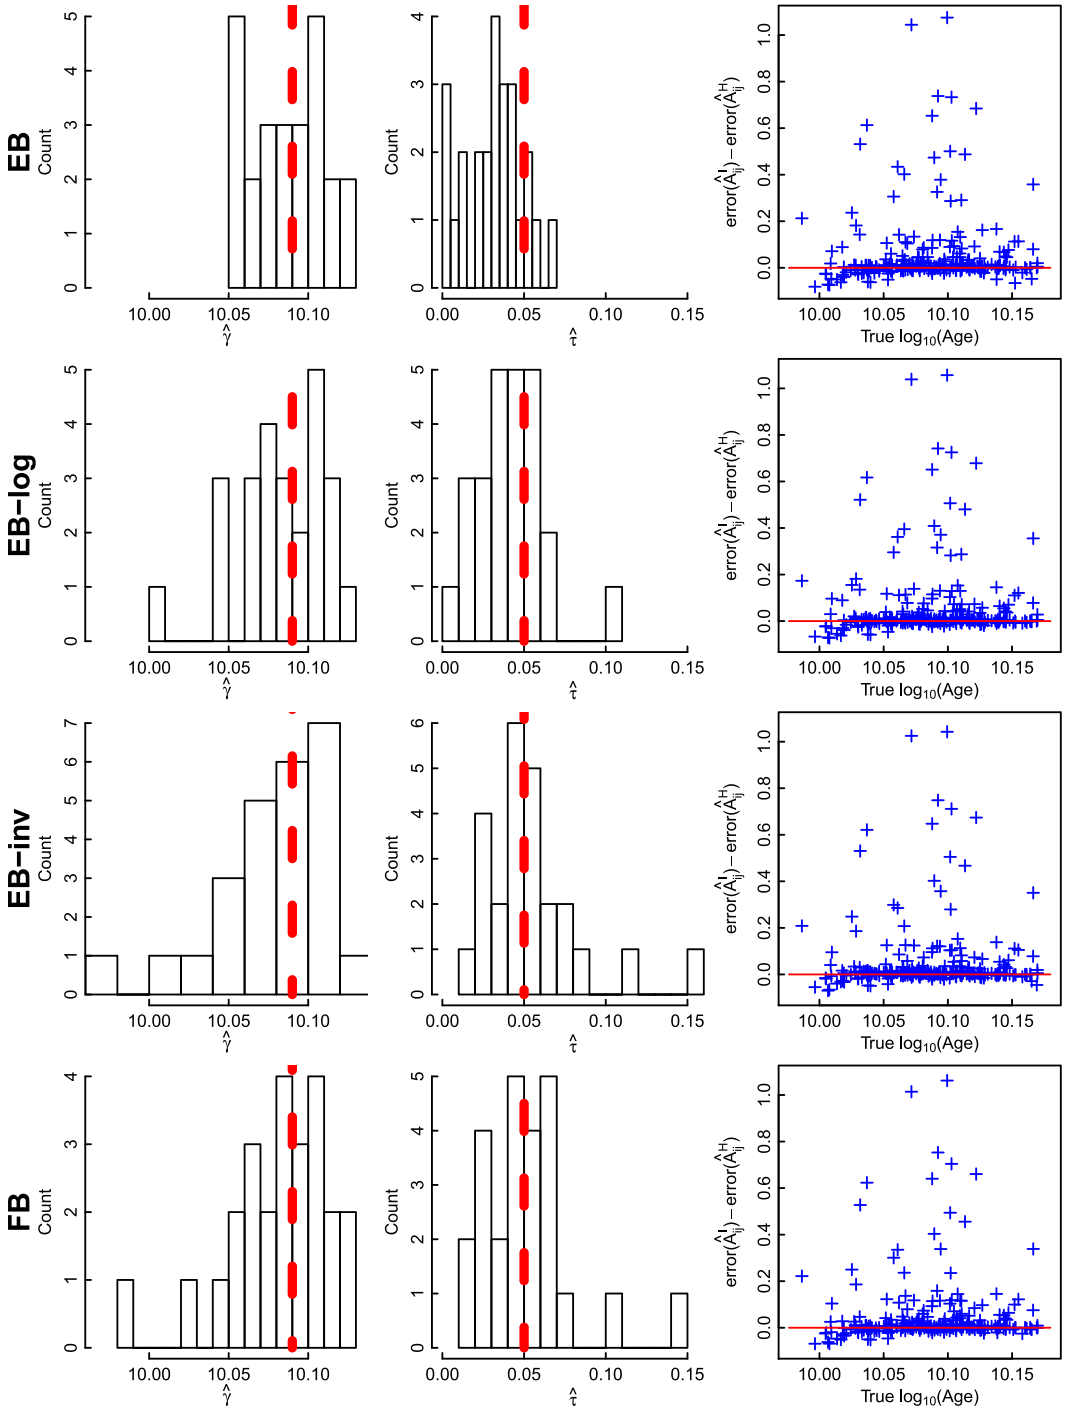


Figure 1. Results for simulated setting I, with $\tau = 0.05$ [standard deviation of $\log_{10}$(age)] and $\Sigma = \Sigma_0$ (photometric error variances). The rows correspond to various fitting methods: EB, EB-log, EB-inv and FB, respectively. Each row includes (i) a histogram of the fitted values of $\hat{\gamma}$ across the $N_2 = 25$ simulation replicates, (ii) a histogram of $\hat{\tau}$ and (iii) a comparison of the difference in absolute errors in the fitted WD ages obtained with the case-by-case analysis versus each of the hierarchical methods. The thick dashed red lines in each histogram indicate the true values of γ or τ. In the rightmost column, positive vertical values indicate larger error for the case-by-case fit and the thin horizontal lines correspond to equal errors for the two methods. Figs 6–9 for simulated settings II–V are shown in the online supplementary material.

the need for the transformed EB-log or EB-inv. Both of these approaches produce similar results to EB, but avoid the possibility of $\hat{\tau} = 0$. The third column in Fig. 1 shows the scatter plot of Diff(A_{ij}) against A_{ij}, $i = 1, 2, \ldots, N_1, j = 1, 2, \ldots, N_2$. Because most of the scatter in these plots is above the solid red zero line, the estimates of $\log_{10}$(age) from the case-by-case analyses tend to be further from the true values than the shrinkage estimates. Approximately two thirds of the $N_1 \times N_2$ simulated stars in each setting are better estimated with the shrinkage method than the case-by-case fit. For stars below the red solid lines, nominally the case-by-case fit is better,

Table 3. The percentage of WDs with improved age estimates.

Simulation settings		EB (per cent)	EB-log (per cent)	EB-inv (per cent)	FB (per cent)
$\tau = 0.05$	$\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$	65.6	69.6	68.8	66.8
$\tau = 0.03$	$\boldsymbol{\Sigma} = 0.8^2\boldsymbol{\Sigma}_0$	74.4	72.8	76.4	73.2
	$\boldsymbol{\Sigma} = 1.2^2\boldsymbol{\Sigma}_0$	67.6	69.6	70.4	72.8
$\tau = 0.06$	$\boldsymbol{\Sigma} = 0.8^2\boldsymbol{\Sigma}_0$	65.6	67.6	68.0	64.4
	$\boldsymbol{\Sigma} = 1.2^2\boldsymbol{\Sigma}_0$	60.4	64.0	66.8	63.2

but the advantage is small. In fact, for almost all simulated stars, $\text{Diff}(A_{ij}) = \text{error}(A_{ij}^I) - \text{error}(A_{ij}^H) > -0.1$, so the shrinkage estimates do not perform much worse than case-by-case estimates for any WD and often perform much better. For some stars, we have $\text{Diff}(A_{ij}) > 0.5$. From the point of view of reliability of the technique, it is comforting that the four hierarchical fits (EB, EB-log, EB-inv, FB) perform similarly, at least when $\hat{\tau} > 0$ for EB.

Table 2 presents a numerical comparison of the shrinkage and case-by-case estimates of $\log_{10}$(age). Specifically, it presents the average mean absolute error (MAE) and the average root of the mean squared error (RMSE) for each method, i.e.

$$\text{MAE}(A) = \frac{1}{N_1 \cdot N_2} \sum_{j=1}^{N_2} \sum_{i=1}^{N_1} |\hat{A}_{ij} - A_{ij}|,$$

$$\text{RMSE}(A) = \sqrt{\frac{1}{N_1 \cdot N_2} \sum_{j=1}^{N_2} \sum_{i=1}^{N_1} (\hat{A}_{ij} - A_{ij})^2}.$$

Both MAE and RMSE measure the distance between the true values and their estimates. Smaller MAE and RMSE indicate that the estimate is more accurate. Table 2 summarizes the performance of different estimators under the five simulated settings. In terms of MAE and RMSE, all four shrinkage estimates (EB, EB-log, EB-inv, FB) are significantly better than the case-by-case estimates, though there are slight differences among the four shrinkage estimates. Table 3 reports the percentage of simulated WDs that are better estimated by shrinkage methods than the case-by-case fits for each of the four statistical approaches and each of the five simulation settings. We conclude that 60 per cent–75 per cent of simulated stars have a more reliable age estimate from the hierarchical analyses than from the case-by-case analyses.

From Tables 2 and 3, we conclude that shrinkage estimates from both EB-type and FB approaches outperform the case-by-case analyses in terms of smaller RMSE and MAE. Under the five simulated settings, all four computational methods, EB, EB-log, EB-inv and FB, behave similarly. Their MAEs and RMSEs are comparable. Also, the percentages of better estimated WDs from these four approaches are consistent.

Simulation setting III ($\tau = 0.03, \boldsymbol{\Sigma} = 1.2^2\boldsymbol{\Sigma}_0$) benefits most from the shrinkage estimates in terms of reduced RMSE. The RMSE from the case-by-case fits under simulation setting III is approximately 0.20, while the RMSE from the EB is around 0.031, less than one sixth of the former. Simulation setting IV ($\tau = 0.06, \boldsymbol{\Sigma} = 0.8^2\boldsymbol{\Sigma}_0$) gains the least from the shrinkage estimates; the RMSE of EB (0.063) is about a quarter of the RMSE of case-by-case (0.26). Generally, when Σ is large and τ is small, the advantage of shrinkage estimates is the greatest. With small Σ and large τ, the advantage of shrinkage estimates over case-by-case estimates decreases. This is consistent with statistical theory (see Gelman et al. 2014, chapter 5). Generally speaking, using EB-log rather than EB to avoid a fitted variance of zero. In terms of computational investment, the FB algorithm is less time-consuming than all of our EB algorithms.

Table 4. Prior distributions for distance moduli and atmosphere composition for halo WD sample

WD	Distance modulus $(V - M_V)$	Comp.	Reference
J0301−0044	$N(4.35, 0.2)^a$	He	Paper III[d]
J0346+246	$N(2.24, 0.33)^a$	H	Paper II[c]
J0822+3903	$N(5.19, 2.4)$	H	Paper IV[e]
J1024+4920	$N(5.54, 2.4)$	He	Paper IV[e]
J1102+4113	$N(2.64, 0.13)^a$	H	Paper II[c]
J1107+4855	$N(3.41, 0.19)^a$	H	Paper III[d]
J1205+5502	$N(5.04, 2.4)$	He	Paper IV[e]
J2137+1050	$N(4.46, 2.4)$	H	Paper I[b]
J2145+1106N	$N(4.19, 2.4)$	H	Paper I[b]
J2145+1106S	$N(4.23, 2.4)$	H	Paper I[b]

Notes. [a]Precise distance moduli are from trigonometric parallax measurements.
[b]Paper I is Kilic et al. (2010).
[c]Paper II is Kilic et al. (2012).
[d]Paper III is Gianninas et al. (2015).
[e]Paper IV is Dame et al. (2016).

5 ANALYSIS OF A GROUP OF CANDIDATE HALO WDS

Now we turn to the hierarchical and case-by-case analysis of the 10 field WDs from the Galactic halo listed in Table 4. In the entire analysis, we employ the Dotter et al. (2008) WD precursor models, Renedo et al. (2010) WD interior models, Bergeron et al. (1995) WD atmospheres and Williams et al. (2009) IFMR.

We acquire prior densities on M_i, D_i, and Z_i, $i = 1, \ldots, 10$ from the literature (Kilic et al. 2010, 2012; Gianninas et al. 2015; Dame et al. 2016). The atmospheric compositions and priors on distance moduli for these 10 stars are listed in Table 4. We use a ZAMS mass prior IMF from Miller & Scalo (1979) on M_i and a diffuse prior on metallicity $Z_i \sim N(-1.50, 0.5)$. In this paper, we do not leave the WD core composition as a free parameter, but instead, we use the WD cooling model derived from the work of Renedo et al. (2010).

For the case-by-case fitting of each WD, we employ a minimally informative flat prior on A_i, specifically $A_i \sim \text{Unif}(8.4, 10.176\,09)$.

5.1 Case-by-case analysis

We derive the joint posterior density for the parameters using Bayes' theorem:

$$p(A_i, D_i, Z_i, M_i|\boldsymbol{X}_i) \propto p(\boldsymbol{X}_i|A_i, D_i, Z_i, M_i)p(A_i)p(D_i)p(Z_i)p(M_i). \quad (16)$$

Before specifying a hierarchical modelling for the 10 WDs, we obtain their case-by-case fits using BASE-9 (as in O'Malley et al. 2013). By using the priors in Table 4 and as described above, we fit each of 10 halo WDs individually with BASE-9. We present results for five typical stars in Fig. 2.

Each column in Fig. 2 corresponds to one WD. The rows provide different two-dimensional projections of the posterior distributions. The asymmetric errors in the fitted parameters, including age, are evident. The first row illustrates that the correlation between the metallicity and age for these five WDs is weak. In the second row, the distance and age of WDs have a strong positive correlation for

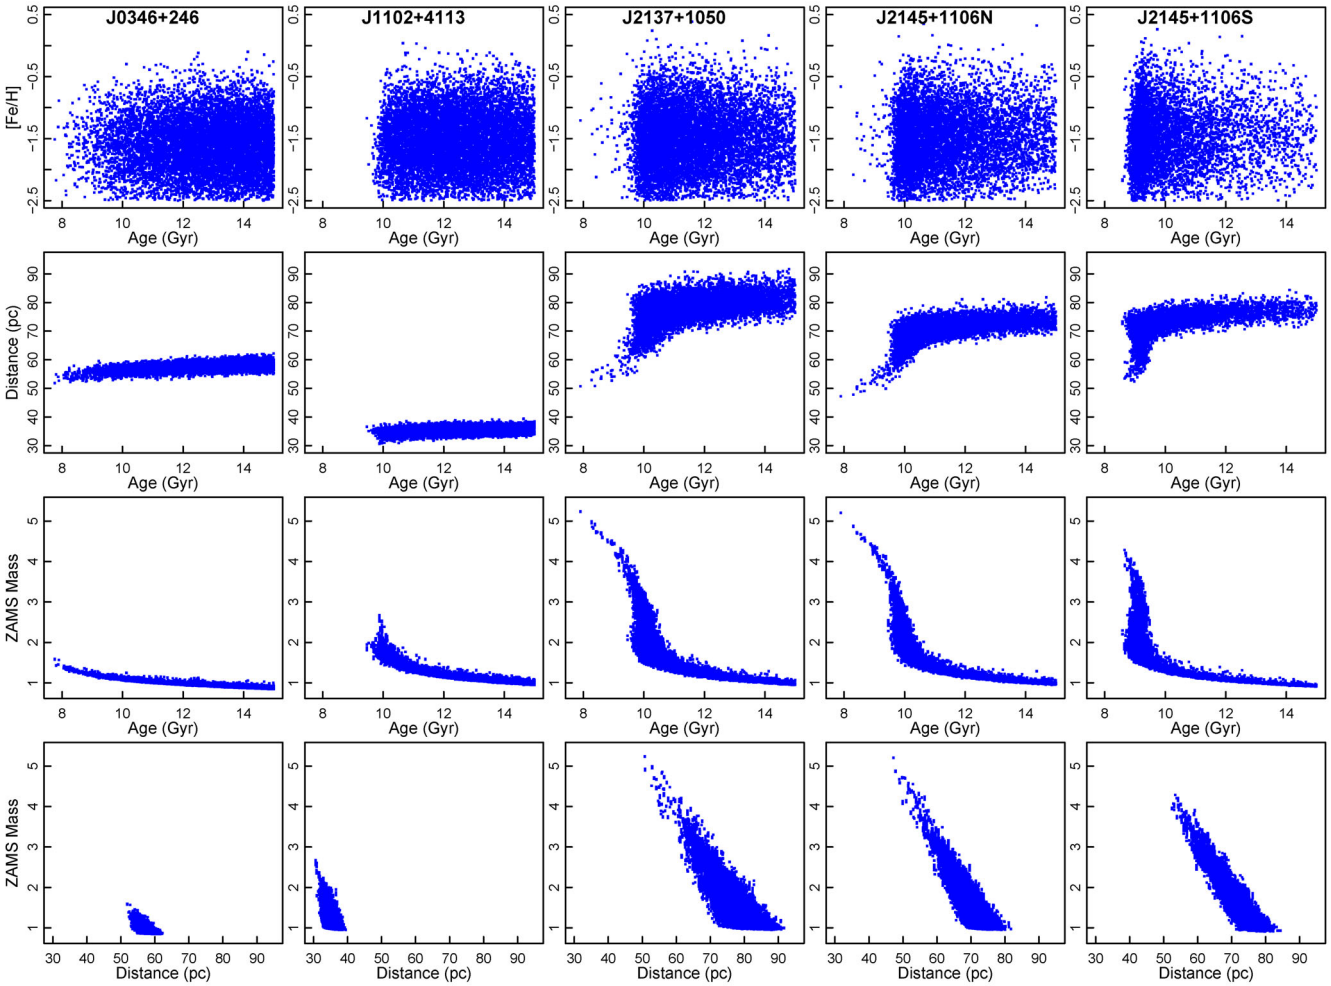


Figure 2. Case-by-case results: projections of the joint posterior distributions on to the two-dimensional planes of (from top to bottom) age–metallicity, age–distance, age–ZAMS mass and distance–ZAMS mass for five of the Galactic halo WDs (columns).

ages less than 10 Gyr. However, this pattern generally disappears for ages greater than 10 Gyr. From the third and fourth rows, the ZAMS mass displays a clear negative correlation with both age and distance.

The plot shows that the range of possible ages for these five stars is large, from 8 to 15 Gyr. Assuming that each of these 10 WDs is a bona fide Galactic halo member, we expect their ages to be similar. However, their masses, distance moduli and metallicities may vary substantially. In this situation, it is sensible to deploy hierarchical modelling on the ages of these 10 WDs, which provides substantial additional information.

5.2 Hierarchical analysis

Here we deploy both EB-log and FB to obtain fits of the hierarchical model in equation (12) based on 10 candidate Galactic halo WDs. In Fig. 3, we compare the posterior density distributions of the age of each WD obtained with the case-by-case method and with that obtained with both EB-log and FB.

Fig. 3 demonstrates that the posterior distributions of the ages under the hierarchical model – both EB-log and FB – peak near a sensible halo age, whereas the case-by-case estimates (solid lines) disperse over a much wider range. Both EB-log and FB estimates are consistent, which we discuss further below.

The photometric errors of these 10 WDs are close to Σ_0 in the simulation study. So the data are similar to simulation setting I ($\tau = 0.07$, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$). Hence, the advantage of the shrinkage estimates over the case-by-case estimates shown in simulation setting I should be predictive, and we expect that the hierarchical fits (dotted lines) in Fig. 3 are better estimates of the true ages of these halo WDs.

Table 5 and Fig. 4 summarize the estimated ages. The 68.3 per cent posterior intervals of ages of WDs from EB-log and FB are generally narrower than those from the case-by-case analyses, which means that shrinkage estimates (FB and EB-log) produce more precise estimates. The fits and errors from EB-log and FB are quite consistent.

In both BASE-9 and the hierarchical model (equation 12), the ages, A_i, of stars are specified on the $\log_{10}$(Year) scale. Given an MCMC sample from the posterior distribution of age on the $\log_{10}$(Year) scale, we can obtain an MCMC sample on the age scale by backwards transforming each value in the sample via

$$\text{age} = 10^{(A_i - 9)}, \tag{17}$$

where the units for age and $\log_{10}$(age) are Gyr and $\log_{10}$(Year), respectively. For the population-level parameters γ and τ, however, the transformation from the $\log_{10}$(Year) scale to the Gyr scale is more complicated. Again, starting with the MCMC sample of γ and τ, for each sampled pair, we (i) generate a Monte Carlo sample of A_i,

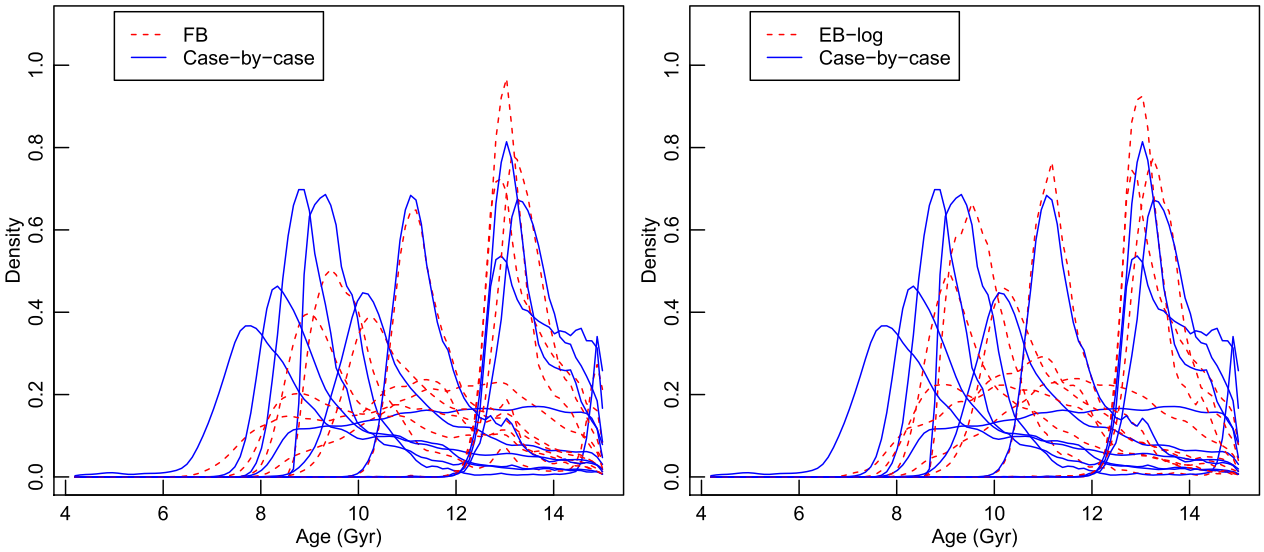


Figure 3. Comparison of the posterior distributions of ages of WDs obtained with the hierarchical analyses and with the case-by-case analysis. The left-hand panel uses the FB approach to fit the hierarchical model, whereas the right-hand panel uses EB-log. Both hierarchical approaches give age estimates that are more consistent with each other than with the case-by-case estimates.

Table 5. Estimates of the ages (in Gyr) of 10 candidate halo WDs.

WD name	FB		EB-log		Case-by-case	
	Mean	Posterior interval	Mean	Posterior interval	Mean	Posterior interval
J0301−0044	11.97	(10.26, 13.65)	11.85	(10.24, 13.5)	11.84	(9.6, 14)
J0346+246	10.68	(9.24, 12.97)	9.92	(9.18, 10.56)	10.23	(9.07, 10.88)
J0822+3903	11.33	(10, 12.91)	11.12	(9.97, 12.43)	11.1	(9.81, 12.74)
J1024+4920	10.65	(8.57, 12.5)	10.55	(8.98, 11.9)	8.94	(7.43, 10.9)
J1102+4113	13.41	(12.74, 14.23)	13.41	(12.74, 14.22)	13.64	(12.84, 14.53)
J1107+4855	10.4	(8.82, 12.48)	10.07	(8.83, 11.51)	9.46	(8.53, 10.35)
J1205+5502	10.89	(8.78, 13.01)	10.67	(8.88, 12.45)	9.77	(8.22, 11.83)
J2137+1050	13.47	(12.93, 14.08)	13.46	(12.91, 14.06)	13.63	(13.03, 14.35)
J2145+1106N	13.25	(12.74, 13.84)	13.24	(12.73, 13.81)	13.4	(12.8, 14.15)
J2145+1106S	11.68	(10.87, 12.71)	11.54	(10.83, 12.37)	11.65	(10.84, 12.64)

(ii) transform this sample to the Gyr scale as in equation (17), and (iii) compute the mean and standard deviation of the transformed Monte Carlo sample. Histograms of the resulting sample from the posterior distribution of the mean and standard deviation of the age on the Gyr scale appear in Fig. 5.

We present estimates of the population distribution of the age of Galactic halo WDs in Table 6, on both the Gyr and $\log_{10}$(Year) scales. In the first two rows, we report the 68.3 per cent posterior intervals for the mean (γ) and standard deviation (τ) of the distribution of ages of halo WDs, $12.11^{+0.85}_{-0.86}$ Gyr and $1.18^{+0.57}_{-0.62}$ Gyr, respectively. The point estimates of the population mean (11.43 Gyr) and standard deviation (1.86 Gyr) from EB-log are quite consistent with results of FB. EB-log does not directly provide error estimates for the population mean and standard deviation, but bootstrap techniques (Efron 1979) could be used. We do not pursue this here, because it is computationally expensive and uncertainties are provided by FB.

In Table 6, we also report the 68.3 per cent predictive intervals of the age distribution from FB and EB-log, which summarizes the underlying distribution of halo WD ages. These are our estimates of an interval that contains the ages of 68.3 per cent of halo WDs. From FB, the 68.3 per cent predictive interval for the distribution of halo WD ages is $12.11^{+1.40}_{-1.53}$ Gyr. In other words, we predict that 68.3 per cent of WDs in the Galactic halo have ages between 10.58 and 13.51 Gyr. The 68.3 per cent predictive interval from EB-log is 11.43 ± 1.86 Gyr, slightly broader than that from the FB.

In summary, our hierarchical method finds that the Galactic halo has a mean age of $12.11^{+0.85}_{-0.86}$ Gyr. Furthermore, the halo appears to have a measurable age spread with standard deviation $1.18^{+0.57}_{-0.62}$ Gyr. This result is preliminary as we await *Gaia* parallaxes to tightly constrain distances, which constrains both ages and stellar space motions. If one or a few of these WDs have anomalous atmospheres, are unresolved binaries or are not true halo members, including them in this hierarchical analysis could artificially increase the estimated halo age spread.

Our mean halo age estimate is consistent with other WD-based age measurements for the Galactic halo. For halo field WDs, these estimates are 11.4 ± 0.7 Gyr (Kalirai 2012), 11–11.5 Gyr (Kilic et al. 2012) and $10.5^{+2.0}_{-1.5}$ (Kilic et al. 2017). Although broadly consistent, these studies all use somewhat different techniques. The study of Kalirai (2012) relies on spectroscopic determinations of field and globular cluster WDs. The Kilic et al. (2012) analysis depends on photometry and trigonometric parallaxes, as does our

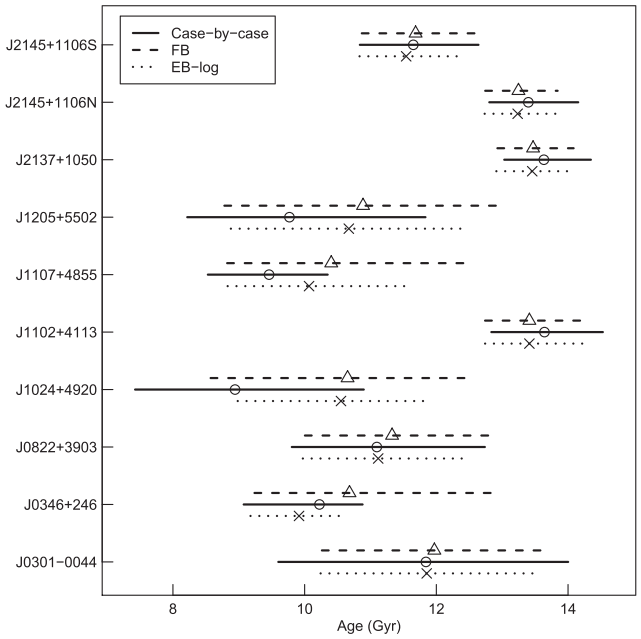


Figure 4. The EB-log, FB and case-by-case estimates of the ages of the 10 candidate halo WDs. The interval estimates are central 68.3 per cent posterior intervals of the posterior distributions. The dashed, solid and dotted intervals are computed from FB, case-by-case and EB-log, respectively. The triangle, circle and cross signs are the posterior means of the ages for each approach. The hierarchical age estimates of J1205+5502 and J1024+4920 in particular are shrunk towards the age estimates of the other WDs.

work, yet at that point only two halo WDs were available for their study. The Kilic et al. (2017) study is based on the halo WD luminosity function isolated by Munn et al. (2017). Although this sample contains 135 likely halo WDs, there are as yet no trigonometric parallaxes or spectroscopy to independently constrain their masses. Thus, all of these samples suffer some defects, and it is comforting to see that different approaches to these different halo WD data sets yield consistent halo ages.

Another comparison to the field WD halo age is the WD age of those globular clusters that have halo properties. Three globular clusters have been observed to sufficient depth to obtain their WD ages, and two of these (M4 and NGC 6397) exhibit halo kinematics and abundances. The WD age of M4 is 11.6 ± 2 Gyr (Bedin et al. 2009) and that age for NGC 6397 is 11.7 ± 0.3 Gyr (Hansen et al. 2013). These halo ages for globular cluster stars are almost identical to those for the halo field. If there is any problem with these ages, it may only be that they are too young, at ∼2 Gyr younger than the age of the Universe. At this point, we lack sufficient data to determine whether this is a simple statistical error, with most techniques having uncertainties in the range of 1 Gyr, or whether there is a systematic error with the WD models or IFMR for these stars, or whether these WD studies have simply failed to find the oldest Galactic stars. Alternatively, as mentioned above, the field halo age dispersion may really be of the order of ±2 Gyr, in which case the halo field is sufficiently old, yet these globular clusters may not be. We look forward to future results from *Gaia* and LSST that should reduce the observational errors in WD studies substantially while dramatically increasing sample size. This will allow us to precisely measure the age distribution of the Galactic halo and place the globular cluster ages into this context.

6 CONCLUSION

We propose the use of hierarchical modelling, fitted via EB and FB to obtain shrinkage estimates of the object-level parameters of a population of objects. We have developed novel computational algorithms to fit hierarchical models even when the likelihood function is complicated. Our new algorithms are able to take advantage of available case-by-case code, with substantial savings in software development effort.

By applying hierarchical modelling to a group of 10 Galactic halo WDs, we estimate that 68.3 per cent of Galactic halo WDs have ages between 10.58 and 13.51 Gyr. This tight age constraint from the photometry of only 10 halo WDs demonstrates the power of our Bayesian hierarchical analysis. In the near future, we expect not only better calibrated photometry for many more WDs, but also to incorporate highly informative priors on distance and population

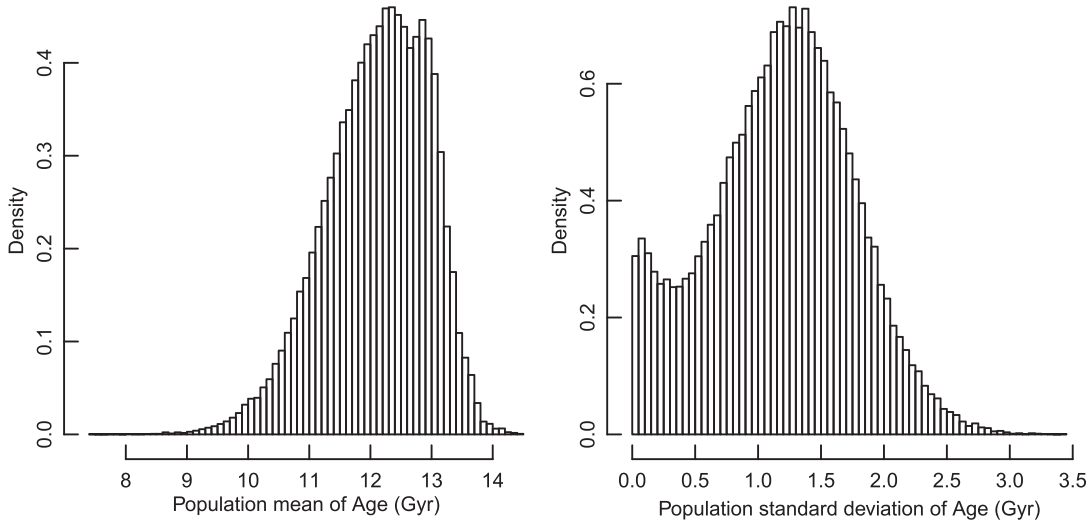


Figure 5. Histograms of MCMC samples from the posterior distribution of the mean and standard deviation of the population of halo WD ages on the Gyr scale, obtained using FB fit to 10 candidate Galactic halo WDs.

Table 6. The 68.3 per cent posterior/predictive intervals for the age of Galactic halo WDs

Items	Units	FB	EB-log
Posterior interval for the mean age of halo WDs	$\log_{10}$(Year)	$10.07^{+0.03}_{-0.04}$	n.a.
	Gyr	$12.11^{+0.85}_{-0.86}$	n.a.
Posterior interval for the standard deviation of ages of halo WDs	$\log_{10}$(Year)	0.05 ± 0.03	n.a.
	Gyr	$1.18^{+0.57}_{-0.62}$	n.a.
Predictive interval for the age distribution of halo WDs[a]	$\log_{10}$(Year)	10.08 ± 0.05	10.06 ± 0.07
	Gyr	$12.11^{+1.40}_{-1.53}$	11.43 ± 1.86

Note. [a]The 68.3 per cent predictive interval for the ages of WDs in the Galactic is an estimate of an interval that contains 68.3 per cent of halo WD ages, taking account of uncertainties of both population-level parameters, γ and τ, and of the variability in the ages of halo WDs.

membership from the *Gaia* satellite's exquisite astrometry. We look forward to using these WDs to fit our hierarchical model in order to derive an accurate and precise Galactic halo age distribution.

ACKNOWLEDGEMENTS

We would like to thank the Imperial College High Performance Computing support team for their kind help, which greatly accelerates the simulation study in this project. SS thanks the Department of Mathematics in Imperial College London for a Roth studentship, which supports his research. DvD acknowledges support from a Wolfson Research Merit Award (WM110023) provided by the British Royal Society and from Marie-Curie Career Integration (FP7-PEOPLE-2012-CIG-321865) and Marie-Skodowska-Curie RISE (H2020-MSCA-RISE-2015-691164) Grants both provided by the European Commission.

REFERENCES

Barnes S. A., 2010, ApJ, 722, 222
Bedin L. R., Salaris M., Piotto G., Anderson J., King I. R., Cassisi S., 2009, ApJ, 697, 965
Bergeron P., Wesemael F., Beauchamp A., 1995, PASP, 1047
Carlin B. P., Louis T. A., 2000, J. Am. Stat. Assoc., 95, 1286
Carrasco J., Catalan S., Jordi C., Tremblay P.-E., Napiwotzki R., Luri X., Robin A. C., Kowalski P. M., 2014, A&A, 565, A11
Casella G., 1985, Am. Stat., 39, 83
Clementini G., Gratton R., Bragaglia A., Carretta E., Di Fabrizio L., Maio M., 2003, AJ, 125, 1309
Dame K., Gianninas A., Kilic M., Munn J. A., Brown W. R., Williams K. A., von Hippel T., Harris H. C., 2016, MNRAS, 463, 2453
DeGennaro S., von Hippel T., Jefferys W. H., Stein N., van Dyk D., Jeffery E., 2009, ApJ, 696, 12
Dempster A. P., Laird N. M., Rubin D. B., 1977, J. R. Stat. Soc. B, 39, 1
Di Benedetto G., 1997, ApJ, 486, 60
Dotter A., Chaboyer B., Jevremović D., Kostov V., Baron E., Ferguson J. W., 2008, ApJS, 178, 89
Efron B., 1979, Ann. Stat., 7, 1
Efron B., 1996, J. Am. Stat. Assoc., 91, 538
Efron B., Morris C., 1972, J. Am. Stat. Assoc., 67, 130
Eggen O. J., Lynden-Bell D., Sandage A. R., 1962, ApJ, 136, 748
Feast M., 2000, preprint (arXiv:astro-ph/0010590)
Feast M. W., Catchpole R. M., 1997, MNRAS, 286, L1
Fitzpatrick E. L., Ribas I., Guinan E. F., DeWarf L. E., Maloney F. P., Massa D., 2002, ApJ, 564, 260
Fontaine G., Brassard P., Bergeron P., 2001, PASP, 113, 409
Forbes D. A., Bridges T., 2010, MNRAS, 404, 1203
Freeman K., Bland-Hawthorn J., 2002, ARAA, 40, 487
Gelman A., 2006a, Bayesian Anal., 1, 515
Gelman A., 2006b, Technometrics, 48, 241
Gelman A., Carlin J. B., Stern H. S., Rubin D. B., 2014, Bayesian Data Analysis, Vol. 2. Taylor and Francis, London
Gianninas A., Curd B., Thorstensen J. R., Kilic M., Bergeron P., Andrews J. J., Canton P., Agüeros M. A., 2015, MNRAS, 449, 3966
Gieren W. P., Fouqué P., Gómez M., 1998, ApJ, 496, 17
Gratton R., Sneden C., Carretta E., 2004, ARA&A, 42, 385
Groenewegen M. A., 2000, A&A, 363, 901
Hansen B. et al., 2013, Nature, 500, 51
Hills S., von Hippel T., Courteau S., Geller A. M., 2015, AJ, 149, 94
James W., Stein C., 1961, in Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Vol. 1, p. 361
Jeffery E. J., von Hippel T., DeGennaro S., van Dyk D. A., Stein N., Jefferys W. H., 2011, ApJ, 730, 35
Kalirai J. S., 2012, Nature, 486, 90
Kilic M. et al., 2010, ApJ, 715, L21
Kilic M., Thorstensen J. R., Kowalski P., Andrews J., 2012, MNRAS, 423, L132
Kilic M. et al., 2017, AJ, in press
Laney C., Stobie R., 1994, MNRAS, 266, 441
Leaman R., VandenBerg D. A., Mendel J. T., 2013, MNRAS, 436, 122
Licquia T. C., Newman J. A., 2015, ApJ, 806, 96
Mandel K. S., Wood-Vasey W. M., Friedman A. S., Kirshner R. P., 2009, ApJ, 704, 629
March M., Karpenka N. V., Feroz F., Hobson M., 2014, MNRAS, 437, 3298
Miller G. E., Scalo J. M., 1979, ApJS, 41, 513
Montgomery M., Klumpe E., Winget D., Wood M., 1999, ApJ, 525, 482
Morris C. N., 1983, J. Am. Stat. Assoc., 78, 47
Morris C. N., Lysy M., 2012, Stat. Sci., 27, 115
Munn J. A. et al., 2017, in press
O'Malley E. M., von Hippel T., van Dyk D. A., 2013, ApJ, 775, 1
Panagia N., 1998, Mem. Soc. Astron. Ital., 69, 225
Park T., van Dyk D. A., Siemiginowska A., 2008, ApJ, 688, 807
Pawlowski M., Pflamm-Altenburg J., Kroupa P., 2012, MNRAS, 423, 1109
Pietrzyński G., Gieren W., 2002, AJ, 124, 2633
Renedo I., Althaus L. G., Bertolami M. M., Romero A. D., Córsico A. H., Rohrmann R. D., García-Berro E., 2010, ApJ, 717, 183
Roederer I. U., Sneden C., Thompson I. B., Preston G. W., Shectman S. A., 2010, ApJ, 711, 573
Romaniello M., Salaris M., Cassisi S., Panagia N., 2000, ApJ, 530, 738
Sarajedini A., Grocholski A. J., Levine J., Lada E., 2002, AJ, 124, 2625
Scannapieco C., White S. D., Springel V., Tissera P. B., 2011, MNRAS, 417, 154
Shariff H., Dhawan S., Jiao X., Leibundgut B., Trotta R., van Dyk D. A., 2016, MNRAS, 463, 4311
Soderblom D. R., 2010, ARA&A, 48, 581
Stein N. M. et al., 2013, Stat. Anal. Data Min.: ASA Data Sci. J., 6, 34
Tumlinson J. et al., 2010, ApJ, 718, L156
van Dyk D. A., Meng X.-L., 2010, Stat. Sci., 25, 429
van Dyk D. A., Connors A., Kashyap V. L., Siemiginowska A., 2001, ApJ, 548, 224
van Dyk D. A., Degennaro S., Stein N., Jefferys W. H., von Hippel T., 2009, Ann. Appl. Stat., 3, 117
van Leeuwen F., Feast M., Whitelock P., Yudin B., 1997, MNRAS, 287, 955

von Hippel T., Jefferys W. H., Scott J., Stein N., Winget D. E., De Gennaro S., Dam A., Jeffery E., 2006, ApJ, 645, 1436
Wei G. C., Tanner M. A., 1990, J. Am. Stat. Assoc., 85, 699
Williams K. A., Bolte M., Koester D., 2009, ApJ, 693, 355
Yamada S., Suda T., Komiya Y., Aoki W., Fujimoto M. Y., 2013, MNRAS, 436, 1362

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

Figure 6. Results for Simulated Setting II, with $\tau = 0.03$ (standard deviation of $\log_{10}$(age)) and $\Sigma = 0.8^2\Sigma_0$ (photometric error variances).
Figure 7. Similar to Fig. 6, but for Simulated Setting III ($\tau = 0.03$, $\Sigma = 1.2^2\Sigma_0$).
Figure 8. Similar to Fig. 6, but for Simulated Setting IV ($\tau = 0.06$, $\Sigma = 0.8^2\Sigma_0$).
Figure 9. Similar to Fig. 6, but for Simulated Setting V ($\tau = 0.06$, $\Sigma = 1.2^2\Sigma_0$).

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: SHRINKAGE ESTIMATES

In this appendix, we discuss shrinkage estimates and their advantages. Consider, for example, a Gaussian model,

$$Y_i|\theta_i \sim N(\theta_i, \sigma), i = 1, 2, \ldots, n, \tag{A1}$$

where $\boldsymbol{Y} = (Y_1, \ldots, Y_n)$ is a vector of independent observations of each of n objects, $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$ is the vector of object-level parameters of interest and σ is the known measurement error. A simple technique to fit equation (A1) is to estimate each θ_i individually, $\hat{\theta}_i^{\rm ind} = Y_i$, i.e. $\hat{\boldsymbol{\theta}}^{\rm ind} = \boldsymbol{Y}$, using only its corresponding data. If a population is believed to be homogeneous, however, we might suppose, in the extreme case, that all of the objects have the same parameter, $\theta_1 = \theta_2 = \cdots = \theta_n$. For example, one might suppose stars in a cluster all have the same age. Under this assumption, the pooled estimate, $\hat{\theta}_i^{\rm pool} = \bar{Y} = \frac{1}{n}\sum Y_i$, i.e. $\hat{\boldsymbol{\theta}}^{\rm pool} = (\bar{Y}, \ldots, \bar{Y})$, is appropriate.

The mean squared error (MSE) is a statistical quantity that can be used to evaluate the quality of an estimator. As its name implies, it measures the average of the squared deviation between the estimator and true parameter value. Thus, the MSE of $\hat{\boldsymbol{\theta}}^{\rm ind}$ is

$$\mathrm{MSE}(\hat{\boldsymbol{\theta}}^{\rm ind}) = E\left[\sum_{i=1}^{n}(\hat{\theta}_i^{\rm ind} - \theta_i)^2|\boldsymbol{\theta}\right] = n\sigma^2,$$

where $E(\cdot|\boldsymbol{\theta})$ is the conditional expectation function that assumes that $\boldsymbol{\theta}$ is fixed and here that $\hat{\boldsymbol{\theta}}^{\rm ind} = \boldsymbol{Y}$ varies according to the model in equation (A1). It is well known in the statistics literature that the individual estimators $\hat{\boldsymbol{\theta}}^{\rm ind}$ are *inadmissible* if $n > 3$. This means that there is another estimator that has smaller MSE regardless of the true values of $\boldsymbol{\theta}$ or σ^2. In particular, the James–Stein estimator of $\boldsymbol{\theta}$, $\hat{\boldsymbol{\theta}}^{\rm JS} = (1 - \hat{B})\hat{\boldsymbol{\theta}}^{\rm ind} + \hat{B}\hat{\boldsymbol{\theta}}^{\rm pool}$, where $S^2 = \sum(Y_i - \bar{Y})^2/(n-1)$ and $\hat{B} = (n-3)\sigma^2/(n-1)S^2$, is known to have smaller MSE than $\hat{\boldsymbol{\theta}}^{\rm ind}$ if $n > 3$ (James & Stein 1961; Efron & Morris 1972; Morris 1983).[3] When $n > 3$, $\hat{B} > 0$ and the James–Stein estimator of θ_i is a weighted average of the individual estimates, $\hat{\theta}_i^{\rm ind} = Y_i$, and the pooled estimates, $\hat{\theta}_i^{\rm pool} = \bar{Y}$. The James–Stein estimator is an example of *shrinkage estimators*, which are estimates of a set of object-level parameters that are 'shrunk' towards a common central value relative to those derived from the corresponding case-by-case analyses.

The population-level parameters that describe the distribution of $(\theta_1, \ldots, \theta_n)$ are often also of interest. Suppose we model the population by assuming that θ_i follows a common normal distribution, i.e. we extend the model in equation (A1) to

$$Y_i|\theta_i \sim N(\theta_i, \sigma), i = 1, 2, \ldots, n; \tag{A2}$$

$$\theta_i \sim N(\gamma, \tau^2), \tag{A3}$$

where γ and τ are unknown population-level parameters. The model in equations (A2) and (A3) is a hierarchical model and can be fitted using EB (e.g. Morris 1983; Efron 1996). We choose the non-informative, $p(\gamma, \tau) \propto 1$, which is a standard choice in this setting (e.g. Gelman 2006a). The EB approach is Bayesian in that it views equation (A3) as a prior distribution and is empirical in that the parameters of this prior are fitted to the data. Specifically, EB proceeds by first deriving the marginal posterior distribution of γ and τ^2,

$$p(\gamma, \tau^2|\boldsymbol{Y}) = p(\gamma, \tau^2)\prod_{i=1}^{n}\int p(Y_i|\theta_i)p(\theta_i|\gamma, \tau^2)\mathrm{d}\theta_i, \tag{A4}$$

and then estimating γ and τ^2 with the values that maximize equation (A4). These estimates are $\hat{\gamma} = \bar{Y}$ and $\hat{\tau}^2 = \max\left\{\frac{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}{n+1} - \sigma^2, 0\right\}$. [Even with the normal assumptions in equations (A2) and (A3), closed form estimates of γ and τ^2 are available only under the simplifying assumption that the measurement errors for each Y_i are the same, i.e. $\sigma_1^2 = \sigma_2^2 = \cdots = \sigma_n^2$.] Finally, the posterior distribution of θ_i can be expressed as

$$p(\theta_i|Y_i, \hat{\gamma}, \hat{\tau}^2) \propto p(Y_i|\theta_i)p(\theta_i|\hat{\gamma}, \hat{\tau}^2). \tag{A5}$$

Each θ_i can be estimated with its posterior mean under equation (A5). Under certain conditions, EB is consistent with James–Stein estimators (e.g. Morris 1983). EB can produce estimators having the same advantages as James–Stein, and it is readily able to handle more complicated problems whereas James–Stein would require model-specific derivation of MSE-reducing estimators.

APPENDIX B: LARGE MAGELLANIC CLOUD

We illustrate the construction and fitting of hierarchical models and the advantages of shrinkage estimates through an illustrative application to data used to estimate the distance to the LMC. The LMC

[3] It can be shown that the MSE of $\hat{\boldsymbol{\theta}}^{\rm JS}$ is

$$E\left[\sum_{i=1}^{n}(\hat{\theta}_i^{\rm JS} - \theta_i)^2|\boldsymbol{\theta}\right] = n\sigma^2 - \sigma^2(n-3)E(\hat{B}) < n\sigma^2 = E\left[\sum_{i=1}^{n}(\hat{\theta}_i^{\rm ind} - \theta_i)^2|\boldsymbol{\theta}\right],$$

which shows the advantage of James–Stein estimator in terms of MSE over the individual estimator when $n > 3$.

Table B1. Population I distance indicators.

Method	Reported distance modulus	References
Cepheids: trig. paral.	18.70 ± 0.16	Feast & Catchpole (1997)
Cepheids: MS fitting	18.55 ± 0.06	Laney & Stobie (1994)
Cepheids: B-W	18.55 ± 0.10	Gieren, Fouqué & Gómez (1998); Di Benedetto (1997)
Cepheids: P/L relation	18.575 ± 0.2	Groenewegen (2000)
Eclipsing binaries	18.4 ± 0.1	Fitzpatrick et al. (2002)
Clump	18.42 ± 0.07	Clementini et al. (2003)
Clump	18.45 ± 0.07	Clementini et al. (2003)
Clump	18.59 ± 0.09	Romaniello et al. (2000)
Clump	18.471 ± 0.12	Pietrzyński & Gieren (2002)
Clump	18.54 ± 0.10	Sarajedini et al. (2002)
Miras	18.54 ± 0.18	van Leeuwen et al. (1997)
Miras	18.54 ± 0.14	Feast (2000)
SN 1987a	18.54 ± 0.05	Panagia (1998)

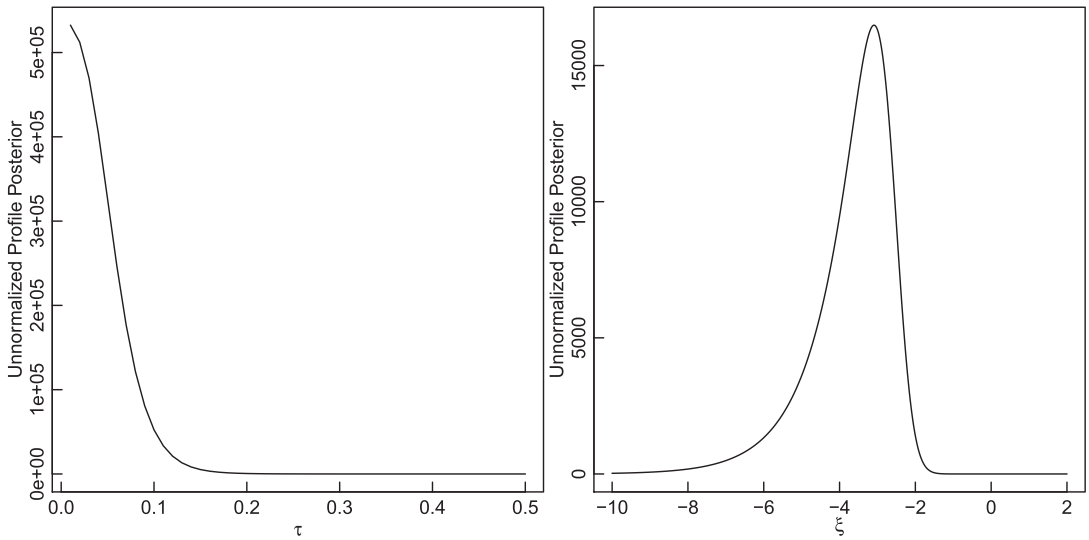


Figure B1. The profile posterior distribution of τ (left-hand panel) and of $\xi = \ln \tau$ (right-hand panel). The modal estimate of τ is zero and is less representative of the distribution than is the modal estimate of ξ.

is a satellite galaxy of the Milky Way. Numerous estimates based on various data sources have been made of the distance modulus to the LMC. The population of stars used affects the estimated distance modulus. Estimates based on Population I tend to be larger than those based on Population II stars. We use a set of estimates based on Population I stars, and formulate a hierarchical model for these estimates in order to develop a comprehensive estimate. We use the data in Table B1, which was compiled by Clementini et al. (2003).

Besides statistical errors, the various distance estimates may be subject to systematic errors. We aim to estimate the magnitude of these systematic errors. If we further assume that the systematic errors tend to average out among the various estimators, we can obtain a better comprehensive estimator of the distance modulus. Let μ_i be the best estimate of the distance modulus that could be obtained with method i, i.e. with an arbitrarily large data set. Because of systematic errors, μ_i does not equal the true distance modulus, but is free of statistical error.

Consider the statistical model,

$$D_i \sim N(\mu_i, \sigma_i), i = 1, \ldots, 13, \tag{B1}$$

$$\mu_i \sim N(\gamma, \tau), \tag{B2}$$

where D_i is the actual estimated distance modulus based on the method/data set i including statistical error, σ_i is the known standard deviation of the statistical error, γ is the true distance modulus of the LMC and τ is the standard deviation of the systematic errors of the various estimates. Equation (B2) specifies our assumption that the systematic errors tend to average out. We denote $\boldsymbol{D} = (D_1, \ldots, D_{13})$ and $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_{13})$.

We take an EB approach to fitting the hierarchical model in equations (B1) and (B2). This involves first estimating the population-level parameters γ and τ and then plugging these estimates in equation (B2) and using it as the prior distribution for each μ_i. Finally, the individual μ_i are estimated with their posterior expectations, $E(\mu_i|\boldsymbol{D}, \hat{\gamma}, \hat{\tau})$, and their posterior standard deviations, $\mathrm{SD}(\mu_i|\boldsymbol{D}, \hat{\gamma}, \hat{\tau})$, are used as 1σ uncertainties. Our EB approach requires a prior distribution for γ and τ. We choose the standard non-informative prior, $p(\gamma, \tau) \propto 1$, in this setting.

We estimate γ and τ by maximizing their joint posterior density,

$$p(\gamma, \tau|\boldsymbol{D}) \propto \int p(\tau, \gamma, \boldsymbol{\mu}|\boldsymbol{D})\mathrm{d}\boldsymbol{\mu}$$
$$= p(\gamma, \tau)\prod_{i=1}^{13}\int p(D_i|\mu_i)p(\mu_i|\gamma, \tau)\mathrm{d}\mu_i. \tag{B3}$$

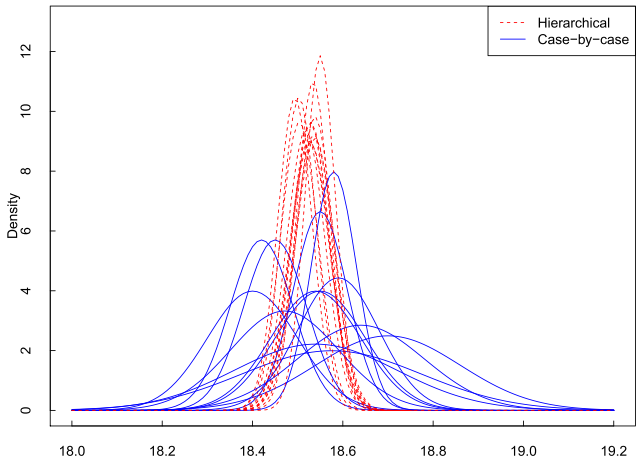


Figure B2. Posterior density distributions of the LMC distance modulus based on hierarchical fitting (red dashed lines) and based on the case-by-case analysis (blue solid lines). Those based on hierarchical fitting shrink towards the centre relative to those based on the case-by-case analysis.

The values of γ and τ that maximize equation (B3) are known as maximum a posterior (MAP) estimates. For any τ, equation (B3) is maximized with respect to γ by

$$\hat{\gamma}(\tau) = \frac{\sum_{i=1}^{13} D_i/(\tau^2+\sigma_i^2)}{\sum_{i=1}^{13} 1/(\tau^2+\sigma_i^2)}, \tag{B4}$$

where $\hat{\gamma}(\tau)$ is a function of τ. The profile posterior density of τ is obtained by evaluating equation (B3) at $\hat{\gamma}(\tau)$ and τ, i.e. $p(\hat{\gamma}(\tau), \tau|\boldsymbol{D})$. The global maximizer of the profile posterior distribution is the MAP estimate of τ and must be obtained numerically.

As shown in the left-hand panel of Fig. B1, the profile posterior density of τ monotonically decreases from its peak at 0, which means that the MAP estimator of τ is 0, a poor summary of the profile posterior. This is because 0 is the lower boundary of the possible values of τ. A better estimate can be obtained using a transformation of the population standard deviation, specifically, $\xi = \ln\tau$. The joint posterior of γ and ξ can be expressed as

$$p(\gamma, \xi|\boldsymbol{D}) = p(\gamma, \exp(\xi)|\boldsymbol{D})\exp(\xi),$$

where $p(\cdot|\boldsymbol{D})$ is the posterior distribution of γ and τ. The profile posterior of ξ is plotted in the right-hand panel of Fig. B1, is more symmetric and is better summarized by its mode. After having estimated ξ with its MAP estimate, we compute $\hat{\tau} = \exp(\hat{\xi})$ and $\hat{\gamma} = \gamma(\hat{\tau})$. See Park et al. (2008) for a discussion of transforming parameters to achieve approximate symmetry in the case of mode-based estimates.

Plugging $\hat{\gamma}$ and $\hat{\tau}$ into the prior for μ_i given in equation (B2), we can compute the posterior distribution of each μ_i as

$$p(\boldsymbol{\mu}|\boldsymbol{D}, \hat{\gamma}, \hat{\tau}^2) \propto \prod_{i=1}^{13} p(D_i|\mu_i)p(\mu_i|\hat{\gamma}, \hat{\tau}).$$

Fig. B2 shows the hierarchical and case-by-case posterior distributions of the individual estimates, μ_i. The hierarchical results (dashed red lines) are shrunk towards the centre relative to the case-by-case results (blue solid lines). The case-by-case density functions of μ_i range from 18.0 to 19.2, whereas the hierarchical posterior density functions are more precise, ranging from 18.3 to 18.7. This is an example of the shrinkage of the case-by-case fits towards their average that occurs when fitting a hierarchical model. We can also see this effect in the posterior means,

$$E(\mu_i|\tau, \boldsymbol{D}) = \frac{\hat{\gamma}(\tau)/\tau^2 + D_i/\sigma_i^2}{1/\tau^2 + 1/\sigma_i^2}, \tag{B5}$$

which are weighted averages of the case-by-case estimates, D_i, and the combined (MAP) estimate of the distance modulus, given in equation (B4). The MAP estimate of the distance modulus is $\hat{\gamma} = 18.525$ and the standard deviation of the systematic errors is $\hat{\tau} = 0.045$ and the distance modulus is $\hat{\gamma} = \gamma(\hat{\tau}) = 18.525$. We compute

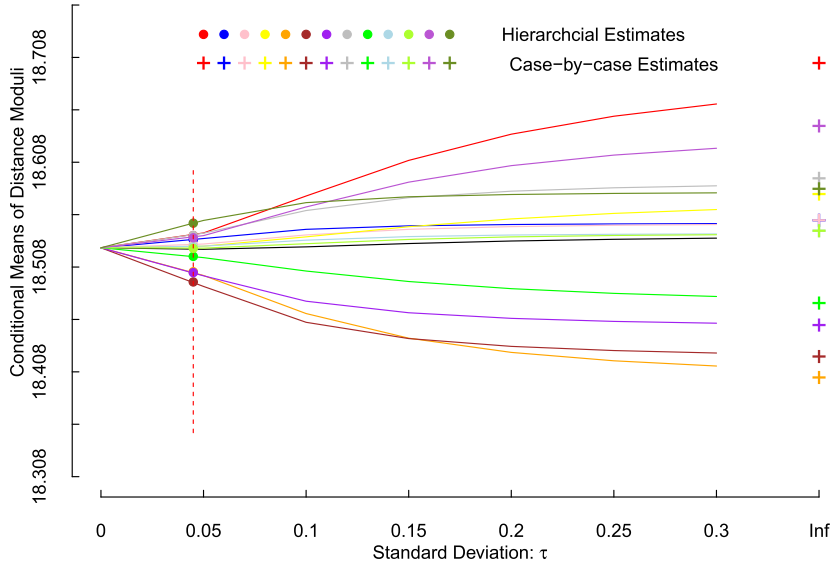


Figure B3. Conditional posterior means of distance modulus $E(\mu_i|\tau, \boldsymbol{D})$, as a function of the standard deviation of the systematic errors τ. The red dashed line indicates our estimate $\hat{\tau} = 0.045$. The 13 solid round dots of different colours show the hierarchical estimates of the distance modulus to the LMC using various methods, whereas the plus signs are estimates from the case-by-case analyses. (Plus signs and round dots of the same colour correspond to the same published result; see Table B1.)

$\hat{\tau}$ via the MAP estimate of ξ as described above. It measures the extent of heterogeneity between 13 different published results. To compute the uncertainty of $\hat{\gamma}$, we generate 200 bootstrap samples (Efron 1979) of $\boldsymbol{D}$, and for each we compute the MAP estimate for γ, resulting in 200 bootstrap estimates of γ with standard deviation 0.024. Thus, our estimate of γ, the distance modulus of LMC, is 18.53 ± 0.024, that is, 50.72 ± 0.56 kpc.

For illustration, in Fig. B3, we plot the posterior expectation $E(\mu_i|\tau, \boldsymbol{D})$ of each of the best estimates of the distance moduli from each method as a function of τ as 13 coloured lines. The black solid line is the MAP estimate $\hat{\gamma}(\tau)$ plotted as a function of τ. When τ is close to zero, the conditional posterior means of each μ_i shrink towards the overall weighted mean $\hat{\gamma}(0) = \frac{\sum D_i/\sigma_i^2}{1/\sigma_i^2}$. The $\tau = 0$ case corresponds to no systematic error and relatively large statistical error. As τ becomes larger, the conditional posterior means approach the case-by-case estimators of the distance moduli marked by plus signs at the far right in Fig. B3. The red dashed vertical line indicates our estimate of τ and intersects the coloured curves at the hierarchical estimates of each μ_i. Fig. B3 shows how the hierarchical fit reduces to the case-by-case analyses as the variance of the systematic errors goes to infinity. We include Fig. B3 to illustrate the 'shrinkage' of the estimates produced with hierarchical models, but such a plot is not needed to obtain the final fit.

APPENDIX C: THE TWO-STAGE ALGORITHM FOR FB

In this appendix, we illustrate how to fit the hierarchical model (in equation 12) via our two-stage algorithm. For more details about this algorithm, see Si et al. (in preparation).

Step 0a. For each WD, run BASE-9 to obtain an MCMC sample of $p(A_i, \boldsymbol{\Theta}_i|\boldsymbol{X}_i)$ under the case-by-case analysis. Thin each chain to obtain an essentially independent MCMC sample and label it $\{A_1^{(t)}, \boldsymbol{\Theta}_1^{(t)}, \ldots, A_n^{(t)}, \boldsymbol{\Theta}_n^{(t)}, t = 1, 2, \ldots, t_{\rm MC}\}$.

Step 0b. Initialize each WD age at $\tilde{A}_i^{(1)} = A_i^{(1)}$ and the other parameters at $\tilde{\boldsymbol{\Theta}}_i^{(1)} = \boldsymbol{\Theta}_i^{(1)}$.

For $s = 1, 2, \ldots,$ we run Step 1 and Step 2.

Step 1. Sample $\tilde{\gamma}^{(s)}$ and $\tilde{\tau}^{(s)}$ from $p(\gamma, \tau|\tilde{A}_1^{(s)}, \ldots, \tilde{A}_n^{(s)})$.

Step 2. Randomly generate n integers between 1 and $t_{\rm MC}$, and denote them $r_1, \ldots, r_n$. For each i, set $A_i^* = A_i^{(r_i)}$, $\boldsymbol{\Theta}_i^* = \boldsymbol{\Theta}_i^{(r_i)}$ as the new proposal and set $\tilde{A}^{(s+1)} = A_i^*$, $\boldsymbol{\Theta}^{(s+1)} = \boldsymbol{\Theta}_i^*$ with probability $\alpha = \min\{1, \frac{p(A_i^*|\tilde{\gamma}^{(s)}, \tilde{\tau}^{(s)})/p(A_i^*|\mu_{A_i}, \sigma_{A_i})}{p(\tilde{A}_i^{(s)}|\tilde{\gamma}^{(s)}, \tilde{\tau}^{(s)})/p(\tilde{A}_i^{(s)}|\mu_{A_i}, \sigma_{A_i})}\}$. Otherwise, set $\tilde{A}^{(s+1)} = \tilde{A}^{(s)}$, $\tilde{\boldsymbol{\Theta}}_i^{(s+1)} = \tilde{\boldsymbol{\Theta}}_i^{(s)}$.

Steps 1 and 2 are iterated until a sufficiently large MCMC sample is obtained. If a good sample from the case-by-case analysis is available, this two-stage sampler only takes a few minutes to obtain an MCMC sample from the FB posterior distribution for the hierarchical model in equation (12).

APPENDIX D: MCEM-TYPE ALGORITHM

In this appendix, we present our algorithm to optimize population-level parameters in Step 1 of EB-type methods (EB, EB-log and EB-inv). We employ an MCEM algorithm with importance sampling for our EB-type methods. MCEM is a Monte Carlo implementation of expectation-maximization (EM) algorithm. See van Dyk & Meng (2010) for more details on EM and MCEM, and an illustration of their application in astrophysics. To apply EM, we treat the object-level parameters, namely $A_1, M_1, D_1, T_1, \ldots, A_n, M_n, D_n, T_n$ as latent variables. Due to the complex structure of this astrophysical model, it is impossible to obtain the expectation step (E-step) of the ordinary EM algorithm in closed form. MCEM avoids this via a Monte Carlo approximation to the E-step. We employ two algorithms to compute the MAP estimate of (γ, τ): Approach 1 is MCEM and Approach 2 uses importance sampling to evaluate the integral in the expectation step instead of drawing samples from the conditional density of the latent variables.

Using Approach 1 to update γ and $\xi = \ln\tau$ requires invoking BASE-9 once for each WD at each iteration of MCEM. This is computationally expensive and motivates Approach 2. We suggest interleaving Approaches 1 and 2 to construct a more efficient algorithm for computing the MAP estimates of γ and τ.

Approach 1: MCEM

Step 0. Initialize $\gamma = \gamma^{(1)}$, $\xi = \xi^{(1)}$, $d_1 = 1$ and $\tau = \exp(\xi^{(1)})$.

Repeat for $t = 1, 2, \ldots,$ until an appropriate convergence criterion is satisfied.

Step 1. For star $i = 1, \ldots, n$, sample $A_i^{[s,t]}$, $\boldsymbol{\Theta}_i^{[s,t]}$, $s = 1, \ldots, S_t$ from their joint posterior distribution

$$p(A_i, \boldsymbol{\Theta}_i|\boldsymbol{X}_i, \gamma^{(t)}, \tau^{(t)}) \propto p(\boldsymbol{X}_i|A_i, \boldsymbol{\Theta}_i)p(A_i|\gamma^{(t)}, \tau^{(t)})p(\boldsymbol{\Theta}_i),$$

where S_t is the MCMC sample size at the tth iteration and should be an increasing function of t (we take $S_t = 1000 + 500t$).

Step 2. Set

$$\gamma^{(t+1)} = \frac{1}{S_t \cdot n}\sum_{i=1}^{I}\sum_{s=1}^{S_t} A_i^{[s,t]},$$

$$\xi^{(t+1)} = \log\left(\frac{1}{S_t \cdot (I-1)}\sum_{i=1}^{n}\sum_{s=1}^{S_t}(A_i^{[s,t]} - \gamma^{(t+1)})^2\right)/2;$$

$$\tau^{(t+1)} = \exp(\xi^{(t+1)}).$$

Approach 2: EM with importance sampling

Suppose we have a sample at the $*$th iteration, $(A_i^{[*,s]}, \boldsymbol{\Theta}_i^{[*,s]})$, $i = 1, \ldots, n$, $s = 1, \ldots, S_*$ from the joint posterior distribution $p(A_i, \boldsymbol{\Theta}_i|\boldsymbol{X}_i, \gamma^*, \tau^*)$ given $\gamma = \gamma*$, $\tau = \tau*$. Set

$$w_i^{[t,s]} = \frac{\frac{\phi(A_i^{[*,s]}|\gamma^{(t)}, \tau^{(t)})}{\phi(A_i^{[*,s]}|\gamma^*, \tau^*)}}{\sum_{s=1}^{S_t}\frac{\phi(A_i^{[*,s]}|\gamma^{(t)}, \tau^{(t)})}{\phi(A_i^{[*,s]}|\gamma^*, \tau^*)}};$$

$$\gamma^{(t+1)} = \frac{1}{n}\sum_{i=1}^{n}\sum_{s=1}^{S_t} A_i^{[*,s]} w_i^{[t,s]};$$

$$\xi^{(t+1)} = \log\left(\frac{1}{(n-1)}\sum_{i=1}^{n}\sum_{s=1}^{S_t}[A_i^{[*,s]} - \gamma^{(t+1)}]^2\right)/2;$$

$$\tau^{(t+1)} = \exp(\xi^{(t+1)});$$

where $\phi(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp(-\frac{(x-\mu)^2}{2\sigma^2})$.

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