

INVERTING COLOR–MAGNITUDE DIAGRAMS TO ACCESS PRECISE STAR CLUSTER PARAMETERS: A NEW WHITE DWARF AGE FOR THE HYADES

STEVEN DEGENNARO¹, TED VON HIPPEL^{1,2}, WILLIAM H. JEFFERYS^{1,3}, NATHAN STEIN¹, DAVID VAN DYK⁴,
AND ELIZABETH JEFFERY¹

¹ Department of Astronomy, University of Texas at Austin, 1 University Station C1400, Austin, TX 78712-0259, USA; deg@astro.as.utexas.edu

² Visiting Scientist, Department of Physics, Florida International University, Miami, FL, USA

³ Department of Mathematics and Statistics, University of Vermont, Burlington, VT, USA

⁴ Department of Statistics, University of California, Irvine, USA

Received 2008 June 5; accepted 2009 January 7; published 2009 April 13

ABSTRACT

We have extended our Bayesian modeling of stellar clusters—which uses main-sequence stellar evolution models, a mapping between initial masses and white dwarf (WD) masses, WD cooling models, and WD atmospheres—to include binary stars, field stars, and two additional main-sequence stellar evolution models. As a critical test of our Bayesian modeling technique, we apply it to Hyades *UBV* photometry, with membership priors based on proper motions and radial velocities, where available. Under the assumption of a particular set of WD cooling models and atmosphere models, we estimate the age of the Hyades based on cooling WDs to be 648 ± 45 Myr, consistent with the best prior analysis of the cluster main-sequence turnoff (MSTO) age by Perryman et al. Since the faintest WDs have most likely evaporated from the Hyades, prior work provided only a lower limit to the cluster’s WD age. Our result demonstrates the power of the bright WD technique for deriving ages and further demonstrates complete age consistency between WD cooling and MSTO ages for seven out of seven clusters analyzed to date, ranging from 150 Myr to 4 Gyr.

Key words: open clusters and associations: general – stars: evolution – white dwarfs

Online-only material: color figures

1. INTRODUCTION

Ages are fundamental in understanding astrophysical processes from the formation of planets to the formation of the Universe. Yet, at present, we have precise ages for only the Solar System (4.566 ± 0.002 Gyr; Allègre et al. 1995) and the Universe as a whole (13.7 ± 0.2 Gyr; Spergel et al. 2003, 2007). For the ages of the Milky Way and its components, we rely on two techniques that typically yield substantially less ($\gtrsim 20\%$) age precision, even with excellent data sets. These two techniques, based on the luminosity and/or color of the main-sequence turnoff (MSTO) and the luminosity of white dwarfs (WDs), are based upon mature theories, though considerable technical difficulties remain in both theory and observation.

Our goal is to improve the age precision of both the MSTO and WD techniques to $\sim 5\%$. Many investigators have collected high-quality data sets, yet this 5% age precision is generally beyond reach. Until the next generation of space-based trigonometric parallaxes from satellites such as SIM and GAIA, we expect no qualitative advances in precision of absolute photometry, stellar abundances, or cluster distances. In our judgment, the greatest gains we can currently make in age precision will come from improved modeling techniques (see also Tosi et al. 1991, 2007; Hernandez & Valls-Gabaud 2008). Any such modeling technique should both fully leverage the data we can collect today and provide a pathway to fully exploit the higher quality data we expect in the future.

We introduced our modeling technique in von Hippel et al. (2006, hereafter Paper I). Briefly, in paper I we developed a Bayesian technique that objectively incorporates our prior knowledge of stellar evolution, star cluster properties, and data quality estimates to derive posterior probability distributions for

a cluster’s age, metallicity, distance, and line-of-sight absorption. We simulated artificial data with a set of oft-used stellar evolution models and realistic photometric error, then recovered the posterior probability distributions of the cluster parameters as well as the masses for each star. We found that our technique yielded high precision for even modest numbers of cluster stars. For clusters with 50 to 400 members and one to a few dozen WDs, we found typical internal errors of $\sigma([\text{Fe}/\text{H}]) \leq 0.03$ dex, $\sigma((m - M_V)) \leq 0.02$ mag, and $\sigma(A_V) \leq 0.01$ mag. We derived cluster WD ages with internal errors of typically only 0.04 dex (10%) for clusters with only three WDs and almost always ≤ 0.02 dex ($\leq 5\%$) with ten WDs.

In Jeffery et al. (2007, hereafter Paper II), we demonstrated the theoretical feasibility of determining cluster WD ages from just the bright WDs, when the coolest WDs are not observed. This technique, discussed below, exploits the slope and position of the WD cooling sequence relative to the MS, extracting age information from the brightest cluster WDs, rather than from the coolest WDs (as in the traditional method). With the assumption of a single-valued initial (MS) - final (WD) mass relation (IFMR), we achieved age precision of 10% with $S/N \geq 30$ when observations limit the WD sample to $M_V < 12$.

In this paper, we extend our model ingredients and Bayesian approach to include two additional models of main-sequence stellar evolution (Yi et al. 2001; Dotter et al. 2008), binary stars, and field stars. We then apply our updated model to the Hyades star cluster, for which more is known than perhaps any other star cluster. The Hyades is our benchmark for determining the precision in the cluster parameters our model can recover, for shaking out subtleties with the current limits of stellar evolution theory, for understanding the complexities introduced by binary and field stars, and for beginning the process of calibrating our bright WD technique, placing it on an absolute age scale.

2. STATISTICAL METHOD

The principles of Bayesian analysis that lie at the heart of our statistical method are more thoroughly discussed in Paper I and van Dyk et al. (2009). Briefly, the goal of our technique is to use information from the data and from our prior knowledge to obtain posterior probability distributions on the parameters of our model. Our prior knowledge is encoded in prior probability distributions on the model parameters, which include cluster parameters such as age and metallicity and individual stellar parameters such as mass, cluster membership, and the masses of any unresolved binary companions. These parameters are the inputs to our cluster evolution model, which we use to derive predicted photometric magnitudes. The likelihood function then compares the predicted magnitudes with the observed data.

Bayes' theorem relates the posterior distribution to the prior distribution and the likelihood function. If $\mathbf{M} = (M_1, M_2, \dots, M_N)$ is a vector of masses of all stars in the cluster and $\Theta = (T, [\text{Fe}/\text{H}], A_V, (m - M_V))$ is a vector of cluster parameters, then we can treat the cluster evolution model as a function $G(\mathbf{M}, \Theta)$ that maps every reasonable choice of (\mathbf{M}, Θ) to a resultant set of expected photometric magnitudes. To obtain the likelihood, we assume that the errors in our measurements are independently distributed and Gaussian with known variance. Suppose there are N stars in the cluster and we have observed them through n different filters. Then the observed data form an $n \times N$ matrix \mathbf{X} with typical element x_{ij} representing the magnitude in the i th filter of the j th star. By assumption, each observed magnitude is normally distributed,

$$x_{ij} \sim N(\mu_{ij}, \sigma_{ij}^2), \quad (1)$$

where μ_{ij} and σ_{ij}^2 are the mean and variance of the modeled photometry through filter i of star j . The means and variances also form $n \times N$ matrices, which we call $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. The full likelihood is then

$$L(\boldsymbol{\mu}, \boldsymbol{\Sigma} | \mathbf{X}) \propto p(\mathbf{X} | \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \prod_{j=1}^N \left(\prod_{i=1}^n \left[\frac{1}{\sqrt{2\pi\sigma_{ij}^2}} \exp\left(-\frac{(x_{ij} - \mu_{ij})^2}{2\sigma_{ij}^2}\right) \right] \right). \quad (2)$$

The variances $\boldsymbol{\Sigma}$ come from our knowledge of the precision of our observations. The means $\boldsymbol{\mu}$ are the predicted photometric magnitudes that we obtain from the cluster evolution model $\boldsymbol{\mu} = G(\mathbf{M}, \Theta)$.

In a more generic notation, where y represents the observed data (e.g., \mathbf{X}) and θ represents the model parameters (e.g., \mathbf{M} and Θ), Bayes' theorem states that the posterior density $p(\theta | y)$ on model parameters θ given data y is

$$p(\theta | y) = \frac{p(y | \theta)p(\theta)}{p(y)}, \quad (3)$$

where $p(y | \theta)$ is the likelihood, $p(\theta)$ is the prior density on the model parameters, θ , and $p(y) = \int p(y | \theta)p(\theta)d\theta$ is a normalizing constant.

From a Bayesian perspective, the posterior distribution is a complete summary of what is known about the model parameters. We can compute means and intervals of this distribution as parameter estimates and error bars. Because the distribution is complex and high dimensional, we use Monte Carlo integration to compute these and other summaries. In particular, we use Markov chain Monte Carlo (MCMC) to generate a sample

from the posterior distribution (Casella & George 1992; Chib & Greenberg 1995). MCMC constructs a Markov chain that upon convergence delivers simulated values that are distributed according to the posterior distribution. The history of the chain can be regarded as a correlated random sample from the posterior distribution. We can thus obtain quantities of interest, such as sample means and quantiles, without having to analytically integrate the normalized posterior distribution. These sample quantities serve as numerical approximations of the corresponding quantities of the posterior distribution.

We use the Metropolis–Hastings algorithm (Chib & Greenberg 1995) to construct our MCMC sampler. In particular, we sample one parameter at a time, conditioning on the current values of all other parameters. For a given single parameter θ , at iteration t , the sampled parameter is generated from a density $q(\theta^* | \theta^{(t)})$, where θ^* is a proposed new value that is accepted with probability, α equal to

$$\alpha = \min \left[\frac{p(\theta^* | y)q(\theta^{(t)} | \theta^*)}{p(\theta^{(t)} | y)q(\theta^* | \theta^{(t)})}, 1 \right] = \min \left[\frac{p(y | \theta^*)p(\theta^*)q(\theta^{(t)} | \theta^*)}{p(y | \theta^{(t)})p(\theta^{(t)})q(\theta^* | \theta^{(t)})}, 1 \right]. \quad (4)$$

If the proposal is accepted, we set $\theta^{(t+1)} = \theta^*$ and otherwise, set $\theta^{(t+1)} = \theta^{(t)}$. Our sample is the parameter sequence $(\theta^{(n)}, \theta^{(n+1)}, \dots, \theta^{(N)})$, where N is the total number of iterations, and $n - 1$ is the number of iterations before the chain converges, which are referred to as the burn-in.

Many of the cluster and stellar parameters are highly correlated (e.g., an increase in the age of the cluster requires a decrease in the mass of a WD to keep its modeled photometry near the observed photometry; a similar correlation exists between MS masses and metallicity). Correlations such as these do not bias our results, but they tend to make the MCMC algorithm inefficient, as it wanders very slowly through parameter space. Fortunately, these correlations are nearly linear and can be removed using dynamic linear transformations computed via linear regression during the burn-in.

To remove the WD age–mass correlation, for each star we introduce a new parameter, U_j , and a new constant, β_j , defined by

$$M_j^{(k)} = \beta_j(T^{(k)} - \bar{T}) + U_j^{(k)}, \quad (5)$$

where $M_j^{(k)}$, $U_j^{(k)}$, and $T^{(k)}$ are the mass, decorrelated mass parameter, and logarithm of the cluster age of the j th star at the k th iteration, respectively, and \bar{T} is the mean log cluster age. Then, rather than directly sampling mass, we sample on U_j for each star. The MCMC algorithm computes the mass at each iteration from Equation (5). The new parameter, U_j , is then decorrelated from distance modulus and metallicity in a similar manner. Finally, the distance modulus and metallicity are decorrelated from each other and then from reddening. Further mathematical details can be found in Paper I and van Dyk et al. (2009).

The most significant additions to our methods since Paper I are the ability to handle unresolved binaries and the ability to handle field star contamination. We have also implemented additional stellar evolution models.

2.1. New Stellar Evolution Models

In Paper I, we used exclusively the following model ingredients: main sequence and giant branch stellar evolution time




Figure 1. Comparison of isochrones for three different main-sequence model sets at the nominal age, distance, and metallicity of the Hyades. Solid (purple) lines = Girardi models, dotted (red) lines = Yale-Yonsei models, dashed (blue) lines = DSED models. (A color version of this figure is available in the online journal.)




Figure 2. Close-up of the WD region of Figure 1. The crosses are the Hyades WDs, the other symbols show the positions of individual theoretical WDs of a given mass along the WD tracks. The solid (purple) line = Girardi models, the dotted (red) line = Yale-Yonsei models, the dashed (blue) line = DSED models. (A color version of this figure is available in the online journal.)

scales of Girardi et al. (2000), the initial-final mass relation of Weidemann (2000), WD cooling time scales of Wood (1992), and WD atmosphere colors of Bergeron et al. (1995). We have since incorporated two new sets of MS models—Yale-Yonsei (Yi et al. 2001), and a finer grid of models from the Dartmouth Stellar Evolution Database (DSED; Dotter et al. 2008) created specifically for this work in a range of ages and metallicities appropriate for the Hyades—as well as updated versions of Bergeron WD atmosphere models.⁵

Figure 1 shows the differences among the main-sequence models at the nominal age, metallicity, distance, and reddening of the Hyades (see Section 3). There are some differences in the position of the turnoff, as well as some important differences

in the shapes of the main sequence. The WD tracks are nearly coincident for all three sets of MS models. A closer look at the WD sequences (Figure 2) shows some differences, however. Here, individual WDs for different zero-age main-sequence (ZAMS) masses are plotted. Although the tracks are nearly coincident, individual WDs of the same mass fall in different places along the tracks. Near the top of the cooling sequence, where the progenitor lifetimes are a significant fraction of the WDs' ages, small differences in the MS timescales have a larger impact on the exact position of the WD along the sequence. For the more massive WDs, which have been cooling much longer, the MS lifetime is a smaller fraction of the WD's total age, and by about $5 M_{\odot}$, the different MS models yield nearly identical positions. Note that these discrepancies in the upper part of

⁵ <http://www.astro.umontreal.ca/~bergeron/CoolingModels/>.

the sequence will result in different mass determinations for individual stars, but will not meaningfully alter the derived age for the cluster (see Section 2.4).

Our goals in this work are to provide precise relative ages, and to use the well-studied Hyades cluster as a step toward calibrating the absolute accuracy of our method. To this end, we restrict ourselves to single sets of WD cooling and atmosphere models. In the case of the Hyades, where the WDs are still relatively warm, the physics of WD cooling is well-understood (Wood 1990; Fontaine et al. 2001). More astrophysically complicated phenomena such as crystallization, phase separation, and the onset of convective coupling do not occur until cooler temperatures than the Hyades WDs have had time to reach. Still our results and error bars should be understood as being model dependent.

We restrict our WD atmosphere models to hydrogen atmosphere (DA) WDs. The ratio of DAs to non-DAs for the field WDs (i.e., those not in open clusters) is between five and six to one, depending on temperature (Williams et al. 2006; DeGennaro et al. 2008). In the so-called DB gap, $25000 \text{ K} \lesssim T_{\text{eff}} \lesssim 45000 \text{ K}$, the ratio of DAs to non-DAs jumps to 12.5:1. The exact corresponding ratios for open clusters are still a matter of debate, but they are almost certainly no smaller, and are likely to be larger (Williams et al. 2006; Kalirai et al. 2008). Moreover, the WDs included in our analysis of the Hyades are all spectroscopically confirmed DAs (Reid 1996).

2.2. Binaries

In our original model, each stellar system (i.e., each point of photometry in the CMD to be modeled) had a single mass. This mass, together with the cluster parameters of age, metallicity, distance, and reddening, allowed our stellar evolution model to derive a predicted photometry, which we compared to the observed photometry to calculate the likelihood function. The stellar evolution model has since been modified to accommodate possible unresolved binary stars.

We now treat each stellar system as if it were a binary, and parameterize it in terms of the larger of the two masses, which we call the primary mass, M_j , and the ratio of the smaller mass to the larger mass, q_j , by definition confined to the interval $[0, 1]$. These two parameters allow us to calculate the mass of a possible unresolved secondary component. The two stars are evolved independently and their fluxes are added to determine their combined photometry. The form of the likelihood probability distribution function (Equation (2)) remains unchanged.

In the Bayesian framework, all fitted parameters require a prior. We choose to place priors on the physically meaningful variables of primary and secondary mass, rather than the sampled variables of decorrelated mass parameter and mass ratio. The prior on the primary mass, as in Paper I, is the Miller–Scalo IMF (Miller & Scalo 1979) normalized from $\log_{10}(M/M_{\odot}) = -1$ to the maximum mass for WDs. The prior on the secondary mass is taken to be flat between 0 and the primary mass.

The decorrelated mass parameter and the mass ratio are sampled on and accepted or rejected in series independently for each stellar system. We use a uniform symmetrical step sampler for the mass ratio, centered on the current value with a step size determined dynamically by the code during the burn-in period. Proposals are reflected if necessary at the boundaries $q_j = 0$ and $q_j = 1$. Note that a star *without* a secondary companion is equivalent to a mass ratio $q_j = 0$, and can be adequately modeled by our cluster evolution model.

As with several other parameters in this multi-dimensional problem, the primary mass and mass ratio for a given MS star are often correlated, meaning that a change in a star’s primary mass can be at least partially compensated for by a change in the mass ratio. Fortunately, this correlation is essentially linear over the sampled range of values, so again for reasons of computational efficiency, we remove it using the same procedure outlined in Section 2.

In order to create secondary stars below the mass limits of the main-sequence stellar evolution models, we extrapolate from the lowest two mass entries. In the future, we plan to incorporate improved models for low-mass stars. For all but the lowest MS primaries, a secondary companion below the mass limits (typically $M < 0.4 M_{\odot}$) makes little or no difference to the photometry of the system. Our extension exists only to provide the evolution model with a means to traverse the distance between the smallest mass in the input models and 0. This serves adequately to differentiate between binaries and single stars, and does not affect the fundamental cluster parameters, which are the target of this study.

Since low-mass MS companions do, however, have a measurable impact on the photometry of the much fainter WDs, we have chosen to restrict the binary models to MS–MS binaries only. While there may theoretically be some age information in WD–WD or WD–MS binaries, in practice these types of systems, particularly when they are too close to be resolved, have often undergone a much more astrophysically complicated evolutionary history due to mass exchange, etc. Modeling such systems would often, if not always, introduce a greater level of uncertainty than whatever we might be able to gain by including them in the analysis.

Moreover, in the case of the Hyades, the WDs have all been studied spectroscopically (Reid 1996), ruling out any significant binary companions for all but HZ9, which we have excluded from our analysis. More distant clusters could potentially contain unresolved white dwarf–M dwarf (WD+M) binaries. However, the vast majority of these stars reside in the broad space between the MS and the WD region in the CMD. At present, our code would reject such a star as a field star, so it would have no effect on our results.

One aspect of MS–MS binary sampling that does affect the cluster parameters is a degeneracy between secondary companions and cluster distance. A given point of photometry on the main sequence can equally well correspond to a single star or an equal mass binary somewhat farther away. Although these two scenarios are indistinguishable from the data for a single-star system, clusters have both a WD sequence and actual MS–MS binaries, which create a secondary MS above the true, single-star sequence. Still, in our analysis of simulated clusters, we observe a slight bias in the distance modulus, especially when we include the redder bands of photometry (I through K), which are more sensitive to low-mass companions. Our analysis of the Hyades, which uses only U, B, V photometry, contains many binary stars, and has a clearly defined WD sequence, is essentially unaffected by this bias.

For simplicity, from this point onward we use the word “star” to mean each individual point of photometry, regardless of whether that point is indeed a single star or an unresolved binary.

2.3. Field Stars

To deal with potential field star contamination in the CMD, we introduce an additional variable for each star, Z_j , which exclusively takes on a value of 0 or 1. If $Z_j = 1$, then the

star is considered to be a cluster member (for that iteration). It is evolved as described above and contributes to the posterior probability distribution as outlined in Section 2. If $Z_j = 0$, then the star is considered to be a field star, and its contribution to the posterior probability is calculated in a different manner described below.

The cluster star model relies on the ability to pool information across stars to leverage the data to determine cluster-wide parameters. In other words, we assume that all of the stars in the cluster share the same age, metallicity, distance, and reddening. These parameters, along with the individual stellar masses and mass ratios, allow us to calculate a predicted magnitude to compare with the observed magnitude in the likelihood. For the field star model, we do not have enough information about an individual star to determine a predicted magnitude, so the likelihood function for a star in the field star model must depend only on its observed magnitude.

Currently this likelihood is taken to be uniform across the entire observed CMD and normalized for each band of photometry between a minimum and maximum determined from the photometry data. In principle, a probability map could be created from, say, an adjacent field or from some generalized map of field stars at a specific Galactic latitude, provided that (1) the map is properly normalized, and (2) the map is not created from the data to be analyzed. We have plans to incorporate such features into our model in the future, but our testing has so far indicated that even our very rough approximation (i.e., uniform across the CMD) is enough for our statistical model to arrive at reasonable answers for posterior distributions on each star's membership status.

The cluster star model has two parameters for each star, one that is related to the star's primary mass and another that is equal to the ratio of secondary mass to primary mass. In the field star model, on the other hand, these two parameters are insufficient, in the absence of any other information about that star (e.g., age, metallicity, distance, absorption—none of which we know for a field star, nor are they of direct interest) to allow our model to predict where the star should lie in the CMD. Since the likelihood in the field star model is not dependent on the value of the two mass variables, their prior distributions alone inform where they are allowed to wander when a star is classified as a field star for multiple consecutive iterations. If we were to leave these priors unchanged in the field star model, the variables would soon wander to regions where a proposed jump back to the cluster star model would be very unlikely to be accepted, and the sampler might jump between models so rarely that the model space would not be adequately sampled.

Fortunately, since these variables have no physical meaning in the field star model, we can place whatever priors we choose on them, including distributions that force them to remain in areas of high probability in the cluster star model. We accomplish this by using a section of the burn-in period to calculate distributions for the decorrelated mass parameter and the mass ratio in the cluster star model. We assume that these distributions can be approximated by Gaussians, and we calculate a mean and variance for each. This is essentially the final step in the burn-in period so that we can be confident that the cluster parameters have converged. For the actual priors, we use distributions with wider tails than Gaussian (specifically, student T distributions with 6 degrees of freedom), so that if our means and variances are somewhat off, the samplers will still be able to find their way to areas of higher probability.




Figure 3. Results of a single MCMC run with main-sequence stars down to $V = 12.5$ included. The points with error bars (including the filled circles, which often obscure the error bars) are the observed data. The (red) X's are the mean positions of the modeled photometry for all iterations of the run after the burn-in period. Each star's posterior probability of being a cluster member is plotted according to the key in the lower right. Note how even our simple field star model serves to distinguish the cluster members from the field stars.

(A color version of this figure is available in the online journal.)

At each iteration, for each star, we sample Z_j , potentially proposing a jump from field star to cluster star or vice versa. We have chosen proposal distributions to be the same as the star's prior probability of membership. That is, if the star has a 70% prior probability of being a member of the cluster (as input by the user), then 70% of the time we propose it to be a cluster member, and 30% of the time we propose it to be a field star. This proposal is independent of the value of Z_j in the current iteration. If the proposed status is the same as the current status, nothing changes and no further calculation is necessary. If the proposed status is different (i.e., a jump), then we use the fully normalized posterior probabilities and proposal distributions of each state to construct the Metropolis–Hastings factor and accept or reject the new state in the standard MCMC manner. Note that we use the prior distributions to inform our proposal distributions solely because we have found this to lead to efficient sampling. Our final results should not—and do not—depend on the choice of proposal distribution.

In the final analysis, we can determine the probability of each star's membership in the cluster by dividing the number of iterations during which it was a cluster member by the total number of iterations in the run. More importantly, the posterior distributions on all of the cluster parameters are now fully marginalized over every possible combination of the other cluster parameters, the individual stellar masses, and each star's individual posterior probability of being a cluster member. A star that spends very few iterations in the cluster star state will contribute minimally to the final cluster parameter distributions. A star that spends every iteration as a cluster member will contribute fully.




Figure 4. WD portion of the color–magnitude diagram showing the subtle differences in slope and position of the WD cooling sequences relative to the fixed MS (not plotted) for clusters of different ages. The isochrones are plotted in intervals of 0.3 dex in $\log(\text{age})$, with the center isochrone (red) at the age we derive for the Hyades in this work.

(A color version of this figure is available in the online journal.)

Figure 3 shows the results of one MCMC run, including main-sequence stars down to $V = 12.5$. The filled circles indicate the posterior probabilities of cluster membership, with black points representing stars that were rejected completely, naked error bars representing stars that were determined to be cluster members on more than 90% of the iterations, and gray points scaling between according to the key in the lower right. Our naive field star likelihood does an adequate job of distinguishing field stars from cluster members. If we were interested in the precise values of cluster membership posterior probabilities on a star-by-star basis, we could employ a more sophisticated model. For the current analysis, we are not modeling field stars realistically to learn anything about the stars, we are simply modeling them adequately to remove them.

The (red) X’s in the diagram show the mean position of the modeled photometry for each star. That is, on each iteration, the model uses the current values of the cluster parameters, each star’s individual mass and mass ratio, and the cluster evolution model to derive a modeled magnitude in each band, which it then compares to the observed photometry in the likelihood. The means of these values for each star across every iteration where the star was considered a cluster member are shown as the X’s, and indicate where the most likely fit for that particular star is located. In the case of the WDs, it is easy to see the correspondence between the six observed points and their respective modeled photometry. The X’s essentially trace out the best-fit WD isochrone.

2.4. The Bright White Dwarf Technique

In Paper II, we demonstrated the theoretical feasibility of determining WD ages from the brighter WDs alone. Briefly, because stars evolve off the main sequence in order of descending mass, the bright WDs in young clusters come from higher mass progenitors than the bright WDs in older clusters. If the initial-final mass relation is universal and single-valued, then the bright WDs in young clusters will themselves also have




Figure 5. Similar to Figure 4 but in magnitude–magnitude space. This is a more accurate representation of what our Bayesian MCMC method sees when it analyzes the WDs in a cluster.

(A color version of this figure is available in the online journal.)

higher masses than their counterparts in older clusters. Through the WD mass–radius relation and the Stefan–Boltzmann law, the mass of a WD affects its position in the CMD relative to the MS.

Figure 4 shows the subtle differences in slope and position of the WD cooling sequences relative to the MS for clusters of different ages. It is easier to see why the method chooses one isochrone over another in a magnitude–magnitude plot, as in Figure 5. In both figures, the isochrones are plotted in ~ 0.3 dex intervals in $\log(\text{age})$, with the center isochrone (in red) at the age we derive for the Hyades in this work.

We noted in Paper II that the *absolute* ages derived by this technique are only as convincing as the IFMR used in the analysis. However, under the generally accepted assumptions that the IFMR is single-valued and the same from cluster to cluster (Weidemann 2000; Williams 2007; Kalirai et al. 2008), the technique yields precise *relative* ages. Indeed, with enough data in hand on many open clusters, we could not only test these assumptions by inter-comparing bright WD and MSTO ages, we could, in essence *recover* the IFMR. The technique requires extensive calibration before this age indicator can be used as an absolute chronometer. This work is a step in such a calibration.

3. THE HYADES

The Hyades is one of the most well-studied open clusters in the sky. The most reliable estimates for the age and the distance to the center of the cluster ($T = 625 \pm 50$ Myr and $m - M = 3.33 \pm 0.01$) come from Perryman et al. (1998), who used Hipparcos astrometry to derive precise trigonometric distances to individual cluster members. Taylor & Joner (2005) find $[\text{Fe}/\text{H}] = +0.103 \pm 0.008$, based on their re-analyses of Paulson et al. (2003), Taylor (1998), and Taylor (1994). Their formal error, however, includes only internal uncertainties. We use a more conservative estimate of error, $\sigma([\text{Fe}/\text{H}]) = 0.05$, to account for systematic uncertainties. For interstellar absorption toward the Hyades, we adopt the result of Taylor (1980), $E(B - V) = 0.003 \pm 0.002$, or $A_V = 0.009 \pm 0.006$. With the exception of age, we use the above values and their stated errors as Gaussian

Table 1
List of WDs in the Hyades, with Cross-References

GCPD identifier	Reid (1996)	EGGR ^a	Reid (1992)	McCook & Sion (1999)
HG7041	HZ4	26	...	0352+096
vA292	VR7	36	192	0421+162
vA490	VR16	37	265	0425+168
vA722	HZ7	39	330	0431+125
4003	HZ14	42	408	0438+108
...	LB227	29	81	0406+169
vA673	HZ9	38	308	0429+176
HG7126	LP 414–120	...	102	0410+188

Note. ^a Eggen & Greenstein (1965).

priors on the cluster parameters. The age prior is flat in $\log(\text{age})$ between the limits of our models.

The data for the Hyades come from the General Catalogue of Photometric Data (GCPD; <http://obswww.unige.ch/gcpd/gcpd.html>; Mermilliod et al. 1997). For each Hyades star in the database, they calculate weighted mean and dispersions in the V band of photometry and the $B - V$ and $U - B$ color indices. Their two-step process, outlined in Mermilliod & Mermilliod (1994, p. 1387), combines data from diverse sources with the first step assigning weights based on the number of independent measurements reported in the source, and the second step slightly shifting the weights to give lower weighting to discrepant measurements. We use the most recent reported values in the database as of 2008 January, and use only those stars for which $U - B$ values are reported.

The errors they report are the dispersions between sources. For stars with three or more sources, we use the reported dispersions to calculate the errors needed by our code. For stars with two or fewer sources, the reported dispersions can often be anomalously low or non-existent. Therefore, we adopt minimum dispersions using the estimates of the average error reported in Mermilliod & Mermilliod (1994, p. 1387), namely: $\epsilon_V = 0.01$, $\epsilon_{B-V} = 0.0075$, $\epsilon_{U-B} = 0.011$. Since our method needs errors in each band of photometry, we take the error in the B band to be the sum of the errors in V and $B - V$ (i.e., the quadrature sum of the variances), and similarly the error in U to be the sum of the errors in B and $U - B$. We have eliminated four stars with anomalously high dispersions ($\sigma_B > 0.1$).

We match what stars we can with the stars in Perryman et al. (1998), and determine a prior probability on each star's membership in the cluster on the basis of their reported χ^2 value (column (w) in their Table 2) and the number of degrees of freedom (3 for stars with radial velocity measurements, 2 for those with proper motions only). Stars without reported values of membership probability are assigned a (somewhat arbitrary) 0.5 prior probability.

Reid (1996) lists eight WDs as members of the Hyades. Table 1 shows these WDs, along with cross-identifications to other major works on Hyades WDs. Of these, vA673 is a known WD–MS binary (Reid 1996), and LP 414–120 does not have available U -band photometry. We eliminate both of these stars from our analysis. The remaining six WDs have individual mass determinations (Weidemann 2000, and references therein). Two of these, LB227 and HZ4, are used as photometric standards in Landolt & Uomoto (2007), and we use their photometry rather than the (less precise) photometry from GCPD.

The precision of the error bars on the photometry for these two objects in particular highlights a problem one encounters with very few open clusters: differential distance. The Hyades

is one of the only clusters for which differential distances pose a serious problem. Perryman et al. (1998) quote a tidal radius for the Hyades of order 10 pc, with stars as far away as 20 pc from the center of mass. A star that is 20 pc nearer to us than the roughly 50 pc distance to the center of the cluster will appear ~ 1.2 mag too bright. A star 20 pc too distant will appear ~ 0.8 mag too faint. For the bright stars, which have individual parallaxes from Hipparcos, this differential distance can be accounted for, though we do not do so in the present work. The WDs in any case are too faint to have Hipparcos parallaxes.

For the purposes of the present problem, we note that an uncertainty in the distance to an individual WD essentially translates to an error in the apparent magnitude. For the two stars in question, we assume an unknown systematic error of 0.1 mag in the V band, and use the errors in color as quoted in Landolt & Uomoto (2007). These errors are plotted as the red error bars in Figure 2. We remind the reader as well that our MCMC analysis does not directly analyze the color–magnitude plot, but rather each band of photometry independently. What our Bayesian method sees looks more like Figure 5. Notice that because the errors in U and B are the sums of the errors in the V band and the appropriate color term(s), the 0.1 mag systematic added to the V band translates to similar errors in the other two magnitudes.

The starting masses and mass ratios for each star are chosen so as to minimize the time necessary during the initial wandering period for each run. In theory, the posterior should not depend on the choice of initial values. In practice, this is usually the case, provided that most of the starting values are within a few tenths of a solar mass for the MS stars, and large enough for the WDs that the code does not try to put the star on the MS or giant branch on the first iteration. If the latter happens, the star will usually be rejected as a field star. Consequently, we choose the starting masses for any objects that lie significantly below and to the left of the main sequence to be $> 3.0 M_{\odot}$.

Since we are interested in a comparison between ages derived from traditional main-sequence turnoff fitting and our technique to determine WD ages, we remove the MS turnoff and giant stars from the data so that our code cannot derive any age information from these stars. We cut off any stars with $V < 4.5$. There is still some age information in the MS insofar as the lack of a turnoff fainter than $V = 4.5$ sets an upper limit on the age. Our bright WD technique in part exploits this phenomenon.

A closer examination of Figure 1 reveals a long-standing problem in stellar astronomy (for recent discussions see von Hippel et al. 2002; Grocholski & Sarajedini 2003). As we move down the main sequence, the models and the data begin to diverge. There is a considerable difference among the models with regard to the shape of the MS. The Yale models fit best overall, particularly in the range of $7 < V < 10$. All three sets of models considerably underestimate the V magnitude for a given $B - V$ (or vice versa) at the bottom of the MS, with the Girardi models dropping off first and most dramatically at $V \simeq 10.5$, the Yale models following shortly thereafter, though not as dramatically, and the DSED models showing reasonable agreement until dropping off precipitously at $V \simeq 12.5$. The DSED models also show considerable diversion from the data in the $U - B$ versus B color–magnitude diagram, particularly on the bright end. We expect, therefore, that the Yale models, showing the best overall agreement with the data, will yield the most accurate results for the cluster parameters.

The proper functioning of our code depends on the models' being accurate. We are currently exploring methods of dealing




Figure 6. Our derived $\log(\text{age})$ for the Hyades as a function of the faintest main-sequence magnitude star included in the analysis. The (blue) squares are the DSED models, the (red) circles are two different runs of the Yale–Yonsei models (to test sensitivity to starting values), and the (purple) triangles are the Girardi models. The horizontal lines are the mean and $\pm 1\sigma$ deviations of the most reliable estimate for the age of the Hyades as determined by MS turnoff fitting (Perryman et al. 1998). It should be noted that for the case of age, this information was *not* used as a prior to inform our analysis.

(A color version of this figure is available in the online journal.)




Figure 7. Our derived $[\text{Fe}/\text{H}]$ for the Hyades as a function of the faintest main-sequence magnitude star included in the analysis. The symbols have the same meaning as in Figure 6. The metallicity value comes from Taylor & Joner (2005) and represents the prior information used in the analysis.

(A color version of this figure is available in the online journal.)

with this discrepancy, and may derive an empirical correction for the lower main sequence for a future paper on MS binary masses. However, our goal in this paper is to derive cluster ages, in particular WD ages, on which the lower MS has a very limited impact. The primary function of the MS in our code is to help pin down the cluster parameters of distance, metallicity, and absorption, yielding more precise ages from the WDs. However, as shown below, we can derive reasonable cluster parameters without using the main-sequence stars at all. Ultimately, the lack of MS model agreement points to a limitation inherent in the models themselves, not our method.

4. RESULTS

To test the sensitivity of our results and errors on the shape of the inaccurately modeled main sequence, we cut off the faint end of the MS at half magnitude intervals and run each set of data separately through the code. Each data set thus contains all of the WDs and the main sequence between $V = 4.5$ and




Figure 8. Our derived distance modulus for the Hyades as a function of the faintest main-sequence magnitude star included in the analysis. The symbols have the same meaning as in Figure 6. The distance value comes from Perryman et al. (1998) and represents the prior information used in the analysis.

(A color version of this figure is available in the online journal.)




Figure 9. Our derived A_V for the Hyades as a function of the faintest main-sequence magnitude star included in the analysis. The symbols have the same meaning as in Figure 6. The A_V value comes from Taylor (1980) and represents the prior information used in the analysis.

(A color version of this figure is available in the online journal.)

some low cutoff, ultimately looking something like the points with error bars in Figure 3. Each data set is run through the code using each set of MS models.

Figures 6–9 show the results of these runs. At each MS cutoff, the three different MS models are slightly offset in X for clarity, with the (purple) triangles representing the Girardi models, the (blue) squares the DSED models, and the (red) circles two different runs with the Yale–Yonsei models (to test the sensitivity of our results to starting values). The horizontal lines are the most recently accepted values for the parameters and their standard errors, as discussed in Section 3. The first cluster of points in each figure represents information obtained from the WDs alone, with no main-sequence stars used in the analysis.

All three models give very reasonable answers for all of the parameters, and the internal precision of the method is comparable to, and in some cases better than, the most precise values measured to date using other methods. This is particularly surprising in the case of distance. We will show below that our choice of prior does have an effect on this result, particularly on the size of the error bars, but that we obtain consistent results even with less restrictive priors. We also remind the reader




Figure 10. Our derived $\log(\text{age})$ for the Hyades when we use less restrictive priors on distance, metallicity, and absorption, similar to what we might know for a less well-studied cluster (black squares). For clarity, only a single set of models (Yale–Yonsei) is plotted with the (red) circles, the same as in Figure 6 plotted for comparison.

(A color version of this figure is available in the online journal.)




Figure 11. Similar to Figure 10, but for metallicity.

(A color version of this figure is available in the online journal.)

that while the horizontal lines in the metallicity, distance, and absorption figures represent the means and standard deviations of the prior probability distributions used in the analysis, the prior on $\log(\text{age})$ is uniform between the model limits, and the horizontal lines in the $\log(\text{age})$ diagram merely represent the most widely accepted turnoff age for the Hyades (Perryman et al. 1998). Thus, the excellent agreement between our WD age and the MS turnoff age is *not* a function of prior information but inherent in the data.

The metallicity (Figure 7) is particularly influenced by the shape and slope of the main sequence. When the WDs are run through the code by themselves (i.e., an MS cutoff of $V = 4.5$ —the first point in each figure), the code essentially has no information on metallicity except for the prior. Hence, the first few points, with no or minimal main sequence information, agree in value and error with our prior. As we add more main sequence, its shape and slope contribute more and more to the likelihood, and thus the posterior. However, since the agreement in shape and slope of the MS models and the data becomes progressively worse as we include the fainter MS stars, the metallicity values determined by our method also tend to become worse.

There are clear breaks in the $[\text{Fe}/\text{H}]$ determinations for each of the MS model sets, between 8.5 and 9.0 for the Yale and




Figure 12. Similar to Figure 10, but for distance modulus.

(A color version of this figure is available in the online journal.)




Figure 13. Similar to Figure 10, but for absorption.

(A color version of this figure is available in the online journal.)

DSED models, and between 6.5 and 7.0 for the Girardi models. Consistent with the visual inspection of the model fits, the Yale models show the best agreement with previous results, as well as the smallest variation between runs. The DSED models fit the poorest overall, particularly in the areas where the models and data diverge visually in the CMD ($V \geq 6.5$ and $V \leq 9.0$).

Figures 10–13 compare the results of a single set of MS models (Yale–Yonsei) using a less restrictive prior, similar to what we might know for a typical cluster. Specifically, the values we used for the standard deviations of the Gaussian priors were 0.3, 0.2, and 0.1 dex for the metallicity, distance modulus, and absorption, respectively. The black squares are the results using the less restrictive priors. The age derived for the cluster is largely insensitive to the choice of prior. The other parameters, as one would expect, are farther from the accepted (i.e., prior) value when the prior distributions are wider. Still, the results are consistent with the accepted distance to the Hyades in almost all cases. Metallicity shows the largest discrepancy, in part because of the aforementioned problems with the models, and in part because metallicity is the parameter that the photometric data constrain the least.

Figure 14 shows the individual WD mass determinations for the various runs, with the plotting symbols the same as in Figures 6–9. In this case, the horizontal lines represent mass determinations from various authors as quoted in Weidemann (2000). For the most part, our mass determinations lie well within the range

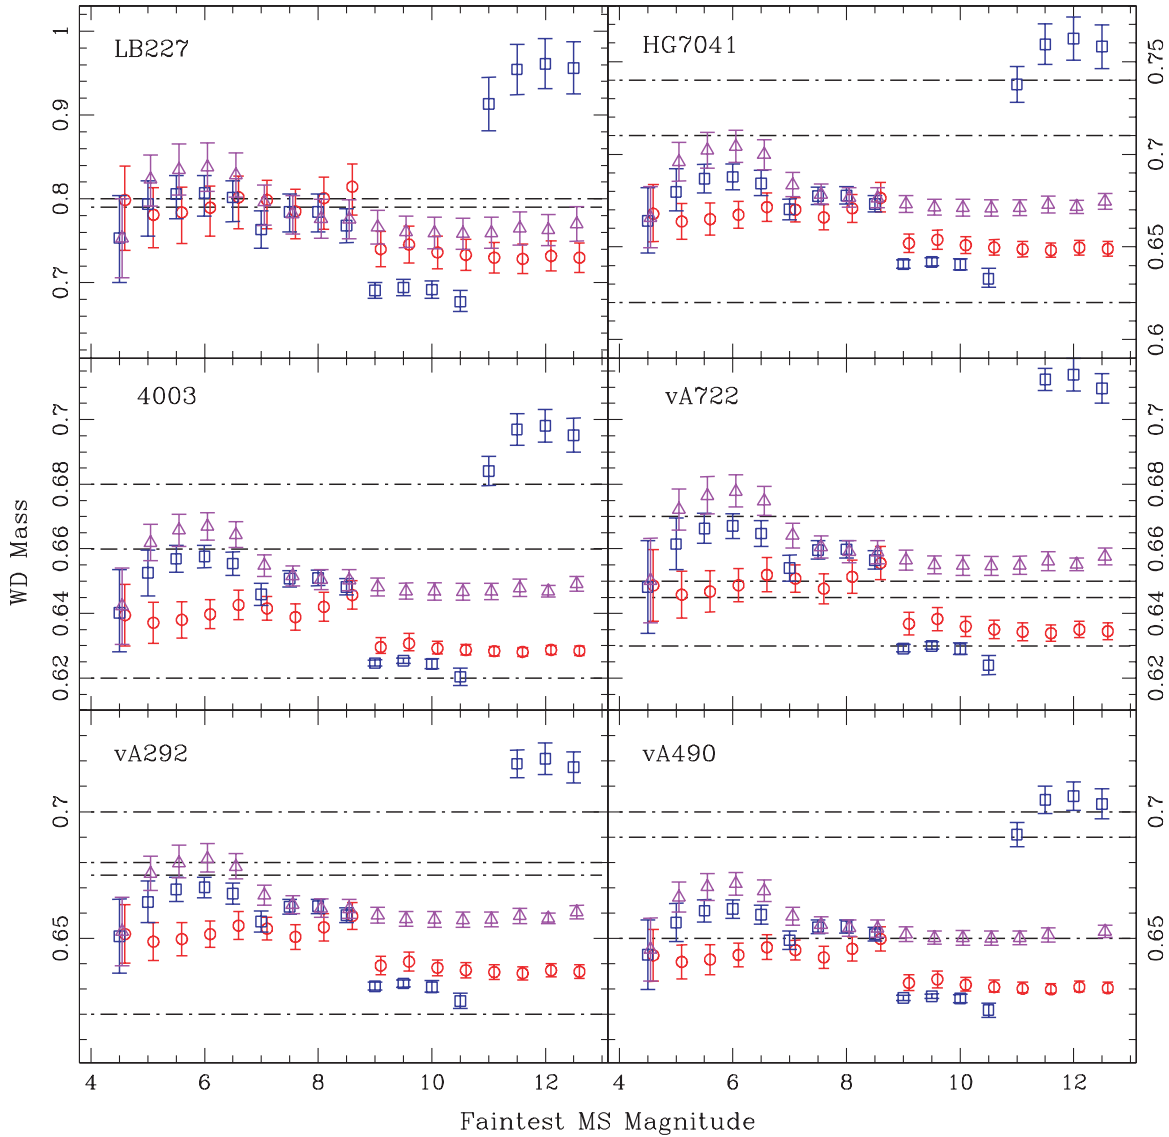


Figure 14. Our mass determinations for the Hyades WDs as a function of the faintest MS magnitude included in the analysis for three sets of models: Yale–Yonsei (red circles), Girardi (purple triangles), and DSED (blue squares). The horizontal lines are previous mass determinations as compiled by Weidemann (2000). (A color version of this figure is available in the online journal.)

of possible masses for the WDs, as determined by other methods, including gravitational redshifts and spectroscopic gravities.

Figures 6–14 show graphically that the choice of model and the amount of main sequence included have a significant effect on the results we obtain. However, there is a clear difference between the results obtained for $V \leq 8.5$ and $V > 8.5$ in nearly all of the models in each parameter. This is not entirely surprising, given that stellar evolution theorists have traditionally concentrated on the upper main sequence, which is easier to model physically and more critical for studies of traditional cluster ages via MS turnoff methods.

Assuming that the models are accurate down to $V = 8.5$, we take a weighted average of the three models for each parameter for the runs which include MS stars down to $V = 8.5$. In the interest of providing a conservative estimate of our errors, we add the standard deviation of the three determinations in quadrature to the average error of the individual determinations. This estimate of the error bars takes into account the internal precision of the method and the systematic errors caused by

differences among the MS models. It does not, however, include other known sources of systematic error, including errors in the WD cooling or atmosphere models or in the imprecisely known IFMR.

The overall systematic uncertainty in the WD age of the Hyades is dominated by uncertainties in the WD cooling models. We have not yet explored these model systematics. Salaris et al. (2008) provide a detailed analysis of systematic uncertainties in WD cooling for various parameters (e.g., core composition and surface layer masses), and we refer the interested reader to their paper. For clusters, rather than individual WDs, the possible effect of systematic errors in the models is complicated since different WDs in the same cluster can pull the age solution in opposite directions.

We derive a WD age for the Hyades of 648 ± 45 Myr ($\log[\text{age}] = 8.81 \pm 0.03$), an $[\text{Fe}/\text{H}]$ of 0.078 ± 0.065 , a distance modulus of 3.35 ± 0.02 (corresponding to a distance of 46.75 ± 0.50 pc), and an absorption in the visual band of 0.014 ± 0.007 . Of these, only the age really provides much new




Figure 15. MSTO vs. WD ages for seven clusters, adapted and updated from von Hippel (2005). The age we derive from the WDs in the Hyades using our “bright white dwarf” technique brings the WD age of the Hyades into agreement with the main-sequence turnoff age for the first time (the solid triangle). The solid line shows a one-to-one correspondence between WD and MSTO ages, and the gray point shows the lower limit of the most reliable WD age for the Hyades prior to this work (Weidemann et al. 1992).

(A color version of this figure is available in the online journal.)

information of astrophysical interest, as the other quantities (for the Hyades, at least) have been measured elsewhere by more accurate methods.

4.1. The Bright White Dwarf Age of the Hyades

Previous studies to determine the WD age of the Hyades cluster have produced a result (300 Myr; Weidemann et al. 1992) that is about half the measured MSTO age (625 Myr; Perryman et al. 1998). Weidemann et al. (1992) suggested that this discrepancy is due to the dynamical evaporation of stars from this cluster; the coolest WDs have been ejected. In the absence of any data on these missing faint WDs, traditional techniques to determine WD ages can provide at best a lower limit to the WD age.

As summarized in Section 2.4, in Paper II we demonstrated the possibility of determining cluster WD ages from just the bright WDs, when the coolest WDs are not observed. Because the coolest WDs are missing from the Hyades, we require the bright WD technique to measure the true WD age (rather than a lower limit, as was done previously). Our Hyades results provide empirical evidence that the bright WD technique yields reasonable and precise ages for real date, as well as providing an important step in calibrating the technique.

Figure 15 is an updated version of Figure 1 from von Hippel (2005), plotting WD age versus MSTO age for open clusters up to 4 Gyr. Our results, a measure of the bright WD age of the Hyades, bring the WD age of this cluster into agreement with the MSTO age for the first time. The solid line shows a one-to-one correspondence between WD and MSTO ages and the gray point shows the most reliable WD age for the Hyades prior to this work (Weidemann et al. 1992).

5. CONCLUSION

We have extended our Bayesian approach from Paper I to include binary stars, field stars, and two additional main-sequence stellar evolution models. As a critical test of our technique, we apply it to Hyades *UBV* photometry, with membership

priors based on proper motions and radial velocities, where known. We find complications arising from the poor fits of most main-sequence isochrones to low mass main-sequence stars, but bypass these difficulties by applying our Bayesian technique to the better-modeled portion of the Hyades main sequence. We further remove the MSTO portion of the data in order to derive WD ages without any further age information. Although we employ only one set of WD cooling (Wood 1992) and atmosphere models (Bergeron et al. 1995), WD cooling in this age range is thought to be well-understood (Wood 1990; Fontaine et al. 2001). We find only small differences in WD-age fits based on each of the three MS stellar evolution models (Girardi et al. 2000; Yi et al. 2001; Dotter et al. 2008). We use the small differences between the results for each of these models as well as the scatter within each result to determine a conservative age estimate for the Hyades based on cooling WDs. This age estimate is 648 ± 45 Myr, consistent with the best prior analysis of the cluster MSTO age by Perryman et al. (1998). The faintest WDs have most likely evaporated from the Hyades, limiting prior work on the WD age of this cluster to report only a lower limit of 300 Myr (Weidemann et al. 1992). Our result demonstrates the power of the bright WD technique for deriving ages, first presented in Paper II. Our WD age for the Hyades also demonstrates complete age consistency between WD and MSTO ages for seven out of seven clusters analyzed, ranging from 150 Myr to 4 Gyr.

We thank Niv Drory for help with computational techniques. This material is based upon work supported by the National Aeronautics and Space Administration under Grant No. NAG5-13070 issued through the Office of Space Science. This material is also based on work supported by the National Science Foundation under grant AST 06-07480.

REFERENCES

- Allègre, C. J., Manhès, G., & Göpel, C. 1995, *Geochim. Cosmochim. Acta*, **59**, 1445
- Bergeron, P., Wesemael, F., & Beauchamp, A. 1995, *PASP*, **107**, 1047
- Casella, G., & George, E. I. 1992, *Am. Stat.*, **46**, 167
- Chib, S., & Greenberg, E. 1995, *Am. Stat.*, **49**, 327
- DeGennaro, S., von Hippel, T., Winget, D. E., Kepler, S. O., Nitta, A., Koester, D., & Althaus, L. 2008, *AJ*, **135**, 1
- Dotter, A., Chaboyer, B., Jevremovic, D., Kostov, V., Baron, E., & Ferguson, J. W. 2008, *ApJS*, **178**, 89D
- Eggen, O. J., & Greenstein, J. L. 1965, *ApJ*, **141**, 83
- Fontaine, G., Brassard, P., & Bergeron, P. 2001, *PASP*, **113**, 409
- Girardi, L., Bressan, A., Bertelli, G., & Chiosi, C. 2000, *A&AS*, **141**, 371
- Grocholski, A. J., & Sarajedini, A. 2003, *MNRAS*, **345**, 1015
- Hernandez, X., & Valls-Gabaud, D. 2008, *MNRAS*, **383**, 1603
- Jeffery, E. J., von Hippel, T., Jefferys, W. H., Winget, D. E., Stein, N., & DeGennaro, S. 2007, *ApJ*, **658**, 391
- Kalirai, J. S., Hansen, B. M. S., Kelson, D. D., Reitzel, D. B., Rich, R. M., & Richer, H. B. 2008, *ApJ*, **676**, 594
- Landolt, A. U., & Uomoto, A. K. 2007, *AJ*, **133**, 768
- McCook, G. P., & Sion, E. M. 1999, *ApJS*, **121**, 1
- Mermilliod, J.-C., & Mermilliod, M. 1994, *Catalogue of Mean UB_V Data on Stars*, VI (Berlin: Springer)
- Mermilliod, J.-C., Mermilliod, M., & Hauck, B. 1997, *A&AS*, **124**, 349
- Miller, G. E., & Scalo, J. M. 1979, *ApJS*, **41**, 513
- Paulson, D. B., Sneden, C., & Cochran, W. D. 2003, *AJ*, **125**, 3185
- Perryman, M. A. C., et al. 1998, *A&A*, **331**, 81
- Reid, I. N. 1996, *AJ*, **111**, 2000
- Reid, N. 1992, *MNRAS*, **257**, 257
- Salaris, M., Serenelli, A., Weiss, A., & Miller Bertolami, M. 2009, *ApJ*, **692**, 1013
- Spergel, D. N., et al. 2003, *ApJS*, **148**, 175
- Spergel, D. N., et al. 2007, *ApJS*, **170**, 377
- Taylor, B. J. 1980, *AJ*, **85**, 242

- Taylor, B. J. 1994, [PASP](#), **106**, 600
Taylor, B. J. 1998, [PASP](#), **110**, 708
Taylor, B. J., & Jonev, M. D. 2005, [ApJS](#), **159**, 100
Tosi, M., Bragaglia, A., & Cignoni, M. 2007, [MNRAS](#), **378**, 730
Tosi, M., Greggio, L., Marconi, G., & Focardi, P. 1991, [AJ](#), **102**, 951
van Dyk, D. A., De Gennaro, S., Stein, N., Jefferys, W. H., & von Hippel, T. 2009, *Ann. of Appl. Stat.*, in press
von Hippel, T. 2005, [ApJ](#), **622**, 565
von Hippel, T., Jefferys, W. H., Scott, J., Stein, N., Winget, D. E., DeGennaro, S., Dam, A., & Jeffery, E. 2006, [ApJ](#), **645**, 1436
von Hippel, T., Steinhauer, A., Sarajedini, A., & Deliyannis, C. P. 2002, [AJ](#), **124**, 1555
Weidemann, V. 2000, [A&A](#), **363**, 647
Weidemann, V., Jordan, S., Iben, I. J., & Casertano, S. 1992, [AJ](#), **104**, 1876
Williams, K. A. 2007, in ASP Conf. Ser. 372, 15th European Workshop on White Dwarfs ed. R. Napiwotzki & M. R. Burleigh (San Francisco, CA: ASP), 85
Williams, K. A., Liebert, J., Bolte, M., & Hanson, R. B. 2006, [ApJ](#), **643**, L127
Wood, M. A. 1990, PhD thesis, Univ. of Texas, Austin
Wood, M. A. 1992, [ApJ](#), **386**, 539
Yi, S., Demarque, P., Kim, Y.-C., Lee, Y.-W., Ree, C. H., Lejeune, T., & Barnes, S. 2001, [ApJS](#), **136**, 417