

Copyright
by
Steven Andrew De Gennaro
2009

The Dissertation Committee for Steven Andrew De Gennaro
certifies that this is the approved version of the following dissertation:

White Dwarfs and the Ages of Stellar Populations

Committee:

Ted von Hippel, Supervisor

Don Winget, Supervisor

Edward Robinson

Volker Bromm

David van Dyk

White Dwarfs and the Ages of Stellar Populations

by

Steven Andrew De Gennaro, B.A., M.A.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

December 2009

For the Queen of the Dorks, from the Viscount of Dweebs.

Acknowledgments

My sincerest appreciation to the following people: Ted von Hippel and Don Winget for their patient support and advice. The members of my dissertation and master's committees. The members of our collaboration—Elizabeth Jeffery, Bill Jefferys, David van Dyk, Nathan Stein, and Albert Dam. The members of the white dwarf group at UT, both current and former—Mike Montgomery, Kurtis Williams, Ross Falcon, J.J. Hermes, Fergal Mullally, Agnes Kim, Mukremin Kilic, and Barbara Castanheira. The co-authors on my papers—in addition to many of the above, S.O. Kepler, Detlev Koester, Atsuko Nitta, and Leandro Althaus. I also thank my parents, who many, many long years ago began laying the foundation for this work.

White Dwarfs and the Ages of Stellar Populations

Publication No. _____

Steven Andrew De Gennaro, Ph.D.
The University of Texas at Austin, 2009

Supervisors: Ted von Hippel
Don Winget

Our group has developed a Bayesian modeling technique to determine the ages of stellar populations (in particular, open and globular clusters) using white dwarf (WD) cooling physics. As the theory of WD cooling is both simpler than, and essentially independent of, main sequence evolutionary theory, white dwarfs provide an independent measure of the ages of Galactic populations.

We have 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, as well as the individual stellar parameters of mass, mass ratio (for unresolved binaries) and cluster membership probability. The key advantage of our Bayesian method is that we can calculate probability distributions for cluster and stellar parameters

with reference only to known, quantifiable, objective, and repeatable quantities. In doing so, we also have more sensitivity to subtle changes in cluster isochrones than traditional “chi-by-eye” cluster fitting methods.

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. We use secular parallaxes derived from Hipparcos proper motions via the moving cluster method to put all members of the Hyades at a common distance. Under the assumption of a particular set of WD cooling and atmosphere models, we estimate the age of the Hyades based on cooling white dwarfs to be 610 ± 110 Myr, consistent with the best prior analysis of the cluster main-sequence turn-off age (Perryman et al., 1998). Since the faintest white dwarfs have most likely evaporated from the Hyades, prior work provided only a lower limit to the cluster’s white dwarf age. Our result demonstrates the power of the bright white dwarf technique for deriving ages (Jeffery et al., 2007) and further demonstrates complete age consistency between white dwarf cooling and main-sequence turn-off ages for seven out of seven clusters analyzed to date, ranging from 150 Myr to 4 Gyr.

We then turn our attention to the white dwarf luminosity function. We use Sloan Digital Sky Survey (SDSS) data to create a white dwarf luminosity function with nearly an order of magnitude (3,358) more spectroscopically confirmed white dwarfs than any previous work. We determine the completeness of the SDSS spectroscopic white dwarf sample by comparing a proper-motion selected sample of WDs from SDSS imaging data with a large catalog of spec-

troscopically determined WDs. We derive a selection probability as a function of a single color ($g - i$) and apparent magnitude (g) that covers the range $-1.0 < g - i < 0.2$ and $15 < g < 19.5$. We address the observed upturn in $\log g$ for white dwarfs with $T_{\text{eff}} \lesssim 12,000\text{K}$ and offer arguments that the problem is limited to the line profiles and is not present in the continuum. We offer an empirical method of removing the upturn, recovering a reasonable mass function for white dwarfs with $T_{\text{eff}} < 12,000\text{K}$.

Finally, we outline several other current and future applications of our method and our code to determine not only ages of Galactic stellar populations, but helium abundances of clusters, ages of individual field WDs, and the initial (main sequence) to final (WD) mass relation.

Table of Contents

Acknowledgments	v
Abstract	vi
List of Tables	xii
List of Figures	xiii
Chapter 1. Introduction	1
1.1 A Brief History Of The White Dwarfs	1
1.2 White Dwarf Cooling And The Ages Of Galactic Populations .	4
1.3 Limitations Of Current Methods	4
1.4 Prior Work	9
1.4.1 von Hippel, et al. (2006)	9
1.4.2 Jeffery, et al. (2007)	9
1.5 Dissertation Outline	11
Chapter 2. Statistical Method	14
2.1 Bayes Theorem	14
2.2 Stellar Evolution Models	16
2.3 Statistical Model	18
2.3.1 Basic Likelihood	18
2.3.2 Binary Stars	23
2.3.3 Field Star Contamination	25
2.3.4 Prior Distributions	26
2.4 Statistical Computations	28
2.4.1 Markov Chain Monte Carlo	29
2.4.2 Sampling and Correlation	31

2.4.2.1	Removing Linear Autocorrelations Between Variables	32
2.4.2.2	Field Star Calculations	37
2.4.3	Starting Values For Parameters	39
2.4.4	Initial Run Calculations	40
Chapter 3.	A New WD Age For The Hyades	43
3.1	The Data	44
3.1.1	Correction For Differential Distance	49
3.1.2	Problems with the Main-Sequence Models	54
3.2	Results	57
3.2.1	Cluster-wide Parameter Determinations	64
3.2.2	The Bright White Dwarf Age Of The Hyades	66
3.2.3	White Dwarf Mass Determinations	68
3.2.4	Main Sequence Mass Determinations	68
3.2.5	Effect of Less Restrictive Priors	72
Chapter 4.	White Dwarf Luminosity and Mass Functions from Sloan Digital Sky Survey Spectra	78
4.1	Introduction	78
4.2	The Data	82
4.2.1	Issues With Model Spectra	85
4.3	Constructing The Luminosity And Mass Functions	90
4.4	Completeness Corrections	93
4.5	Luminosity Functions And Discussion	103
4.5.1	White Dwarf Mass Functions	103
4.5.2	White Dwarf Luminosity Functions	107
4.6	Conclusions	110
Chapter 5.	Current and Future Applications of the Bayesian Method	112
5.1	Analysis of Open (and Globular) Clusters	112
5.2	Initial-Final Mass Relation	114
5.3	Analysis of Single Field White Dwarfs	117

5.4	MCMC Analysis of Binary Stars	122
5.4.1	Binaries in the Field	122
5.4.2	Binaries in Open Clusters	122
5.5	Helium Abundance of NGC 6791	123
	Bibliography	126
	Vita	134

List of Tables

2.1	Stellar and cluster parameters in our Bayesian method.	19
2.2	Periods in the initial run used to compute correlation-reducing transformations.	42
3.1	A list of WDs in the Hyades, with cross-references.	46
3.2	Proper motions for the Hyades white dwarfs	51
4.1	The fraction of stars in Eisenstein et al. (2006) listed as DA or DA_auto.	109

List of Figures

1.1	Fits of isochrones to the same set of data using different values for cluster parameters (Taken from VandenBerg & Stetson (2004)).	6
1.2	A histogram representing a posterior probability distribution for $\log(\text{age})$ for a particular cluster analyzed with our method.	7
1.3	The WD portion of the Hyades 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.	10
2.1	Similar to Figure 1.3 but in magnitude-magnitude space.	21
2.2	A comparison of isochrones for three different main-sequence model sets.	22
2.3	The correlation between ZAMS mass and age for a particular white dwarf.	33
2.4	The Markov chain for a particular cluster before and after the process of decorrelation.	34
2.5	The correlation between primary mass and mass ratio for a particular star.	36
3.1	A comparison of isochrones for three different main-sequence model sets at the nominal age, distance, and metallicity of the Hyades.	47
3.2	The WD Portion of the Hyades CMD, Uncorrected for Distance.	48
3.3	Color-magnitude diagram of the Hyades in absolute magnitude space, after correction for differential distances.	54
3.4	A close up of the white dwarf region of Figure 3.3	55
3.5	Our derived $\log(\text{age})$ for the Hyades as a function of the faintest magnitude main-sequence star included in the analysis.	58
3.6	Our derived $[\text{Fe}/\text{H}]$ for the Hyades as a function of the faintest magnitude main-sequence star included in the analysis.	59

3.7	Our derived distance modulus for the Hyades as a function of the faintest magnitude main-sequence star included in the analysis.	60
3.8	Our derived A_V for the Hyades as a function of the faintest magnitude main-sequence star included in the analysis.	61
3.9	Color-magnitude diagram of the Hyades in absolute magnitude space, overlaid with the isochrones ultimately fit by our method.	63
3.10	MSTO versus WD ages for seven clusters, adapted and updated from von Hippel (2005).	67
3.11	Our mass determinations for the Hyades white dwarfs as a function of the faintest MS magnitude included in the analysis for three sets of models.	69
3.12	A comparison of our results for the masses of twelve MS-MS binaries with mass determinations from Bender & Simon (2008).	70
3.13	Similar to Figure 3.12 for the star vB022.	73
3.14	Similar to Figure 3.5, with errors on the prior on the cluster parameters increased.	74
3.15	Similar to Figure 3.6, with errors on the priors on the cluster parameters increased.	75
3.16	Similar to Figure 3.7, with errors on the priors on the cluster parameters increased.	76
3.17	Similar to Figure 3.8, with errors on the priors on the cluster parameters increased.	77
4.1	$\log g$ v. $\log T_{\text{eff}}$ for white dwarfs in the Sloan Digital Sky Survey	86
4.2	Theoretical vs. observed colors for SDSS white dwarfs	87
4.3	$\log g$ v. $\log T_{\text{eff}}$ with the upturn in $\log g$ removed.	89
4.4	Color-color plot of the white dwarfs used to derive our completeness correction	96
4.5	A map of our completeness correction	97
4.6	A histogram of the completeness values in Figure 4.5	99
4.7	Luminosity functions for three different limiting magnitudes.	101
4.8	Mass functions for three different limiting magnitudes.	102
4.9	The white dwarf mass function for all WDs with $T_{\text{eff}} > 12,000\text{K}$ and $g < 19.0$	104
4.10	White dwarf mass functions for WDs with $T_{\text{eff}} > 8,000\text{K}$ and $g < 19.0$	105

4.11	White dwarf luminosity functions derived in this paper.	106
4.12	Spectroscopically vs. photometrically derived effective temperatures of the SDSS white dwarfs.	108
5.1	An updated version of Figure 3.10 including new and reanalyzed clusters from Jeffery (2009).	113
5.2	The empirical initial-final mass relation.	115
5.3	Mass-age-distance-metallicity correlations for a single (simulated) white dwarf of ~ 8 Gyr in the field.	119
5.4	Same as Figure 5.3 for a simulated star of lower mass.	120
5.5	Similar to Figure 5.3 for the real white dwarf WD0243+026.	121
5.6	Color-magnitude diagram showing the effect of helium abundance on the shape of the main-sequence turn-off region.	124

Chapter 1

Introduction

1.1 A Brief History Of The White Dwarfs

A star like the Sun spends most of its 10 billion year life slowly and stably fusing hydrogen into helium and radiating the energy thus produced out into space. During this phase of its life, the energy produced in the center of the star nearly exactly balances the energy it radiates away. The production of energy—and thus heat—in the core also keeps the star from collapsing under its own immense weight. The pressure of the hot gas pushing outward and the pull of gravity inward are perfectly balanced at each layer in the star. Thus, while the star continues to stably burn hydrogen on what is known as the “main sequence”(MS), it undergoes remarkably little change.

A star’s supply of nuclear fuel, however, is finite, and after its allotted time on the main-sequence (10 billion years in the case of the Sun, shorter for more massive stars—which burn their fuel more quickly—longer for less massive stars), its central fire goes out. As the star cools, its core begins to collapse.

What happens next depends mostly on the mass of the star. If the star is massive enough, then the contracting core will eventually become hot

enough and dense enough to begin burning helium into heavier elements. As it does so, the star's envelope expands greatly and the outer layers of the star cool—the star becomes a red giant. This phase lasts only a small fraction of the main-sequence lifetime, since helium burns at a much faster pace than hydrogen, and there is far less of it. Eventually, as with hydrogen, the helium fuel is exhausted and again the core begins to collapse.

Very massive stars will then begin to burn heavier and heavier elements in a similar manner, until eventually exploding violently in a supernova and leaving behind a black hole or a neutron star. Stars of low or intermediate mass (up to about 8 or 9 times the mass of the sun), will be unable to generate enough heat and pressure in their cores to burn the carbon and oxygen created by the earlier burning of helium.

Through processes still not entirely understood, the outer layers of the star are shed, often forming a planetary nebula. The ultra-hot, ultra-dense core continues to collapse until the particles of which it is composed (specifically the electrons) become so tightly packed that they cannot move any closer together. The laws of quantum mechanics are such that no two particles of the same type can occupy the same space with the same energy. In the collapsed core of a low mass star, the matter is packed so tightly that the electrons begin to fill up all of the lowest possible energy states and are said to be “degenerate”. The fact that the electrons are barred by the laws of physics from moving any closer together is now the only thing holding the star up against its own gravity.

What remains is a stellar ember a little more than half the mass of the Sun, with a core composed of some combination of carbon and oxygen surrounded by thin layers of helium and hydrogen. This is what we call a white dwarf (WD).

Because a white dwarf is no longer replenishing the energy it radiates away, its luminosity decreases monotonically with time. If we know the rate at which it cools, we can determine how long it has been cooling, and by extension, how old it is, just as we could determine from the remains of a campfire approximately at what time the fire went out by the temperature of the embers. Thus, WD cooling can provide a valuable “cosmic clock” to determine the ages of many Galactic populations, including the disk (Winget et al. 1987; Liebert, Dahn & Monet 1988), and open and globular clusters (von Hippel, Gilmore, & Jones 1995; von Hippel et al. 2006; Jeffery et al. 2006). With more accurate models of the cooling physics of white dwarfs, heavily constrained by empirical evidence, it will be possible to determine absolute ages with greater precision than using main-sequence evolution theory.

In addition to astronomical considerations, white dwarfs also provide our only means of studying the physics of compact degenerate matter. The temperatures and densities that occur inside white dwarfs are not only completely impossible to duplicate on Earth, but they will likely remain so for a very long time.

1.2 White Dwarf Cooling And The Ages Of Galactic Populations

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 ($\lesssim 20\%$) age precision, even with excellent datasets. These two techniques, based on the luminosity and/or color of the main-sequence turn-off (MSTO) and the luminosity of white dwarfs, are based upon mature theories, though considerable technical difficulties remain in both theory and observation.

1.3 Limitations Of Current Methods

The methods employed to determine cluster ages cover a wide range in their statistical sophistication. On one end of the spectrum lie methods that first fix the distance, metallicity, and reddening at previously determined values, then estimate the magnitude or color (or, in the “advanced” version, both) of the MSTO in the data, and finally pick an isochrone that has those same value(s), with no reference at all to the shape of the sequence. On the other end are methods that divide the CMD into regions and use counting statistics and Monte Carlo simulations to determine marginalized probability distributions for cluster parameters (e.g., Hansen et al., 2004).

Traditionally, however, most methods for determining ages from either the MSTO or the WD cooling sequence involve some version of “chi-by-eye”. The parameters of a model are tweaked until the authors create an isochrone that *looks* like it fits the data. But as Figure 1.1 shows, a single set of data can often be fit by several (or many) different values of the parameters. The “best fit” choice is often made in a very inexact and subjective manner, and the many subtle trade-offs that must be made in order to fit the data as a whole are ignored or treated very roughly. Often degeneracies make it is necessary to assume several of the cluster parameters (e.g., distance and $[\text{Fe}/\text{H}]$) in order to derive others, and the method provides only the most crude determinations of error bars.

This lack of precision and repeatability motivated us to search for a more objective method. Fortunately, advances in computing power have opened up possibilities for new statistical methods that were impossible just a few decades ago.

We have 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, as well as the individual stellar parameters of mass, mass ratio (for unresolved binaries) and cluster membership probability.

The key advantage of our Bayesian method is that we can calculate probability distributions for cluster and stellar parameters with reference only




Figure 1.1: Fits of isochrones to the same set of data using different values for cluster parameters, demonstrating the inherent limits of the “chi-by-eye” method. Despite the wide range of ages, metallicities, etc., used in the fits, it is nearly impossible to choose any one fit over the others based on a visual inspection alone. [Taken from Vandenberg & Stetson (2004).]




Figure 1.2: A histogram representing a posterior probability distribution for $\log(\text{age})$ for a particular cluster analyzed with our method. By simulating draws from the posterior distributions, rather than simply calculating a single number with errors, we are able to more accurately determine the true distribution of values for the parameters, even in cases, like the above, where the distributions are non-Gaussian and/or multi-modal.

to known, quantifiable, objective, and repeatable quantities. Any biases in our analysis, including our choices for how to include prior information on cluster parameters, are completely transparent.

Our method also has more sensitivity to subtle changes in cluster isochrones than traditional “chi-by-eye” cluster fitting methods. The method

not only allows us to determine cluster parameters and error estimates with a precision not available to the human eye, it also provides highly accurate estimates of the complete marginal probability distributions for each of the parameters of our model. In other words, the outputs of our method are not simply values with some sort of half-assedly propagated error bars, but entire probability distributions for each parameter, as in Figure 1.2. When these distributions are non-Gaussian or multi-modal (and in a problem in many hundreds or thousands of dimensions, this is often the case), we can *see it*. In any case, we can use the probability distributions to calculate quantities of interest, such as confidence intervals, moments, etc..

Our eventual goal is to improve the age precision of both the MSTO and WD techniques to $\sim 5\%$. Many investigators have collected high quality datasets, 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.

1.4 Prior Work

1.4.1 von Hippel, et al. (2006)

We introduced our modeling technique in von Hippel et al. (2006). Briefly, in that paper, 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 white dwarfs, 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 white dwarf ages with internal errors of typically only 0.04 dex (10%) for clusters with only three white dwarfs and almost always ≤ 0.02 dex ($\leq 5\%$) with ten white dwarfs.

1.4.2 Jeffery, et al. (2007)

In Jeffery et al. (2007), we demonstrated the theoretical feasibility of determining white dwarf ages from the brighter white dwarfs 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 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 1.3 shows the subtle differences in slope and position of the




Figure 1.3: The WD portion of the Hyades 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 for the particular model set shown.

WD cooling sequences relative to the MS for clusters of different ages. 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 for the particular set of models shown.

We noted in Jeffery et al. (2007) 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; Williams et al., 2009), 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 dissertation (in particular Chapter 3) is a pivotal step in such a calibration.

1.5 Dissertation Outline

This dissertation is organized as follows:

Chapter 2 outlines the astronomical and statistical models that underly the Bayesian analysis of open clusters and the numerical and computational issues encountered in the implementation of the method. Portions of this material have been published previously in DeGennaro et al. (2009). Other portions are partially adapted from van Dyk et al. (2009).

Chapter 3 describes the application of our method to determine a new WD age for the Hyades open star cluster. The Hyades, one of the most well-studied clusters in the sky, 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, and for beginning the process of calibrating our bright white dwarf technique, placing it on an absolute age scale.

We apply our technique to Hyades *UBV* photometry, with membership priors based on proper motions and radial velocities, where available. We estimate the age of the Hyades based on cooling white dwarfs to be 610 ± 110 Myr, consistent with the best prior analysis of the cluster main-sequence turn-off age by Perryman et al. (1998). Since the faintest white dwarfs have most likely evaporated from the Hyades, prior work provided only a lower limit to the cluster’s white dwarf age. Our result demonstrates the power of the bright white dwarf technique for deriving ages and further demonstrates complete age consistency between white dwarf cooling and main-sequence turn-off ages for seven out of seven clusters analyzed to date, ranging from 150 Myr to 4 Gyr. Portions of this chapter have been published previously in DeGennaro et al. (2009).

In Chapter 4, we turn our attention to the white-dwarf luminosity function. Using data from the Sloan Digital Sky Survey (SDSS), we present a white dwarf luminosity function with nearly an order of magnitude (3,358) more spectroscopically confirmed white dwarfs than any previous work. To

accomplish this, we determine the completeness of the SDSS spectroscopic white dwarf sample by comparing a proper-motion selected sample of WDs from SDSS imaging data with a large catalog of spectroscopically determined WDs. We derive a selection probability as a function of a single color ($g - i$) and apparent magnitude (g) that covers the range $-1.0 < g - i < 0.2$ and $15 < g < 19.5$. We address the observed upturn in $\log g$ for white dwarfs with $T_{\text{eff}} \lesssim 12,000\text{K}$ and offer arguments that the problem is limited to the line profiles and is not present in the continuum. We offer an empirical method of removing the upturn, recovering a reasonable mass function for white dwarfs with $T_{\text{eff}} < 12,000\text{K}$. Most of this chapter was published previously in DeGennaro et al. (2009) and DeGennaro (2007).

The final chapter outlines further current and future applications of the method and code that we have developed. These include: using *JHK* photometry to determine the helium abundance of the open cluster NGC 6791 ; Elizabeth Jeffery’s work on the ages of open clusters; simultaneous fitting of cluster ages with the initial-final mass relation; fitting of single field white dwarfs; and determining masses of binary stars in open clusters and the field.

Chapter 2

Statistical Method

2.1 Bayes Theorem

Bayes' theorem relates the conditional probabilities of two events. In the context of most scientific endeavors (in particular, the subject of the current work), those two events are a particular hypothesis and a set of data that potentially supports or refutes the hypothesis.

If y represents a set of observed phenomena (e.g. cluster photometry) and θ represents a hypothesis (e.g. the values of the parameters of our stellar models), Bayes' theorem states that the posterior density $p(\theta|y)$ on model parameters θ given data y is

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)} \propto p(y|\theta)p(\theta). \quad (2.1)$$

Which is to say, the probability of the truth of hypothesis θ , given observed phenomena y is proportional to the probability that one would observe phenomena y , *if* hypothesis θ were true, times the prior probability of the truth of hypothesis θ . The denominator on the right side of (2.1) is the normalizing constant:

$$p(y) = \int p(y|\theta)p(\theta)d\theta. \quad (2.2)$$

The advantages of the Bayesian method are many. First, by inverting the probability in this way, we can calculate the probability of our hypothesis with reference only to other known, quantifiable, repeatable quantities: the likelihood, $p(y|\theta)$, and our prior degree of belief in the hypothesis, $p(\theta)$. In essence, the method allows us to update our belief in hypothesis θ in the face of new evidence y . While discussion and dissent may exist regarding the best form of the likelihood function and our prior degree of belief, the assumptions underlying any particular choice are transparent and quantifiable.

Second, the method allows us to objectively incorporate our prior degree of belief in a systematic and repeatable way. Each new piece of evidence or set of data adds to our picture of a phenomenon and does not merely supplant all previous determinations. The prior probability provides a way to incorporate (or choose *not* to incorporate) prior knowledge objectively.

The disadvantage of the Bayesian method has historically been the inability to calculate the normalizing integral for all but the simplest of functions. However, due to recent advances in computing power and computational methods, calculating this quantity explicitly is generally not necessary.

The goal of our technique is to use information from photometry 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 as well as individual stellar parameters such as mass, cluster membership, and the masses of any unresolved binary compan-

ions. 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.

2.2 Stellar Evolution Models

The Bayesian method relies on our ability to predict the photometry of an individual star (or unresolved binary system) given a set of values for model parameters. To accomplish this, we employ a set of stellar evolutionary models that allows us to map any reasonable choice of cluster-wide parameters (e.g., age, metallicity, distance) and individual stellar masses to a predicted set of photometry for any (relevant) stage of stellar evolution.

For MS stellar evolution, we currently employ three sets of models: Girardi et al. (2000), Yale-Yonsei (Yi et al., 2001), and Dartmouth Stellar Evolution Database (DSED; Dotter et al., 2008). For a given age (T), metallicity ($[\text{Fe}/\text{H}]$), and—in the case of some of the DSED models—helium abundance (Y), these models allow us to calculate an absolute magnitude for any star on the main-sequence or giant branch.

Stars end their lives in one of three states¹: white dwarf, neutron star, or black hole. The exact mass limit for the formation of WDs is unknown, but Williams et al. (2009) assert that it lies between 6 and 9.5 M_{\odot} with 95%

¹There is the possibility that some ultra-massive stars end their lives in an explosion so catastrophic that it leaves no remnant (Woosley et al., 2002). This outcome, if it happens at all, is exceedingly rare, poorly understood, and unlikely to affect the principal results of our work.

confidence. We currently use a mass limit of $8.0M_{\odot}$. Stars with zero-age main sequence (ZAMS) masses above this limit are assumed to form neutron stars and black holes.

For the stars with masses below this limit that have already evolved off of the asymptotic giant branch (AGB) for a particular set of cluster parameters and become WDs, we use the main-sequence + giant branch models to calculate a main-sequence lifetime for the star. Subtracting this lifetime from the age of the cluster gives us a cooling time for the WD formed.

We then employ an initial (MS) - final (WD) mass relation (IFMR) to determine the mass of the resulting WD. Currently, we use the IFMR of Weidemann (2000).

Once we know the mass and cooling age of the WD, we use a set of WD cooling models to determine the effective temperature and radius of the WD. Currently, we employ the cooling models of Wood (1992). We then use the mass, radius, and effective temperature to determine the absolute magnitudes of the WD via a set of WD atmosphere models (Bergeron et al., 1995b).²

Once we have absolute magnitudes for each star in the cluster, we use the distance modulus and reddening to determine apparent magnitudes. These magnitudes are our “predicted” magnitudes, which are compared with the observed magnitudes in the likelihood.

In the implementation of our method, each model element is fully in-

²<http://www.astro.umontreal.ca/~bergeron/CoolingModels>

dependent, meaning that an end user can readily substitute his or her own choice of MS model, IFMR, WD cooling model, and WD atmosphere model based on the specific nature of the problem at hand or individual preference for one set of models over another.

2.3 Statistical Model

2.3.1 Basic Likelihood

In more mathematical terms³, given a vector, Θ , of cluster-wide parameters:

$$\Theta = (\theta_T, \theta_{[\text{Fe}/\text{H}]}, \theta_{m-M_V}, \theta_Y, \theta_{A_V}), \quad (2.3)$$

(see Table 2.1 for explanation of symbols), and a vector, \mathbf{M} , of ZAMS masses for each star in the cluster, we can define a function, $\mathbf{G}(\mathbf{M}, \Theta)$ whose output is a magnitude in each photometric band of interest for each star.

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 of star i through filter j . By assumption, each observed magnitude is normally distributed:

$$x_{ij} \sim N(\mu_{ij}, \sigma_{ij}^2). \quad (2.4)$$

³To maintain consistency with previous works, I have largely adopted the mathematical notation of van Dyk et al. (2009).

Table 2.1: Stellar and cluster parameters

Parameter	Value
θ_T	\log_{10} age in \log_{10} years
$\theta_{[\text{Fe}/\text{H}]}$	\log_{10} metallicity relative to solar ^a .
θ_Y	helium abundance
θ_{m-M_V}	distance modulus
θ_{A_V}	line-of-sight absorption in the V band
M_{i1}	mass of the more massive star in binary-star system i
M_{i2}	mass of the less massive star in binary-star system i

^aWe use $[\text{Fe}/\text{H}]$ as a proxy for metallicity (Z), with MS model sets each having a standard scaling of heavy elements

For each star i through filter j , the variance σ_{ij}^2 comes from our knowledge of the precision of our observations, and the mean μ_{ij} is the predicted photometric magnitude that we obtain from the cluster evolution model. The means and variances also form $n \times N$ matrices, which we call $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. Each row of $\boldsymbol{\mu}$ is the complete set of modeled photometry in all bands:

$$\boldsymbol{\mu}_i = \boldsymbol{G}(M_i, \boldsymbol{\Theta}). \quad (2.5)$$

The full likelihood is then

$$L_p(\boldsymbol{M}_1, \boldsymbol{\Theta} \mid \boldsymbol{X}, \boldsymbol{\Sigma}) = \prod_{i=1}^N \left(\prod_{j=1}^n \left[\frac{1}{\sqrt{2\pi\sigma_{ij}^2}} \exp \left(-\frac{(x_{ij} - G_j(M_{i1}, \boldsymbol{\Theta}))^2}{2\sigma_{ij}^2} \right) \right] \right). \quad (2.6)$$

This is the preliminary likelihood we used in von Hippel et al. (2006) and takes neither field stars nor binaries into account.

Note that our likelihood is constructed in terms of individual bands of photometry and not colors. As such, the method does not operate directly on the CMD itself (as in Figure 1.3). Rather, MCMC “sees” something more akin to Figure 2.1. The reason for this choice has to do with the correlation of errors in the CMD. In modern CCD astronomy, it is much more common to measure the different bands of photometry independently rather than differentially. The likelihood in (2.6) assumes that the errors in the various colors are independent. In practice, this assumption often holds, but certain systematics may correlate the errors. For older data sets taken via differential photometry, we must be more careful.

This likelihood also assumes that the variances, Σ , are solely a function of the measurement errors, and that the stellar models contain no errors of their own. In reality, this turns out to be a poor assumption, as Figure 2.2 demonstrates. The figure shows the differences among the three different MS models for the same cluster parameters. Clearly there are systematic errors caused by uncertainties in the underlying physics of stellar evolution and stellar atmospheres.

Failure to take these errors into account will at best result in an underestimation of the errors on the final parameters. More often, though, we have discovered that for certain areas of parameter space, this discrepancy between MS models (and in particular, the discrepancies between models and data) presents critical difficulties to the numerical method we implement to determine posterior probabilities.




Figure 2.1: Similar to Figure 1.3 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. 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 for the particular model set shown.

We note, however, that these discrepancies will affect *any* method that attempts to use these models. Fortunately, our method provides a way to take




Figure 2.2: A 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.

these errors in the models into account. There is no particular reason why the variances, Σ , cannot include multiple sources of error, including errors in the models themselves. Taking account of the errors in the models can, in the simplest sense, be accomplished by increasing the components of Σ . How *much* we increase them is the tricky part, as it depends on many factors—the parameters of the particular cluster we are fitting, the masses and/or colors and magnitudes of the individual stars. Certain parts of the models fit better

than others, and a complete numerical description needs to take this into account.

We are currently exploring the implications of these model discrepancies. In Section 3.1, we discuss their impact specifically on our determination of the age of the Hyades.

2.3.2 Binary Stars

Approximately one third to one half of all stars are actually binary or multi-star systems. When the components of the binary are far enough apart to be resolved independently, we can model each individual star. However, for binaries too close to be resolved, the measured magnitude in a given band will contain contributions from both components. In such a case, the measured luminosity of the system will be the sum of the individual luminosities of the two stars. The added luminosity of a close binary companion tends to shift stars upwards and to the right in the CMD, often forming a secondary binary sequence above the main sequence. Failure to take this effect into account would introduce a bias into our posterior probabilities.

We therefore model each point of photometry as if it were a binary. We redefine \mathbf{M} as an $N \times 2$ matrix with a typical row $\mathbf{M}_i = (M_{i1}, M_{i2})$ representing the masses of the two star system. We can then evolve each star independently and combine the magnitudes by transforming to luminosity space, summing, and transforming back to magnitudes. The mean values used

in (2.5) and (2.6) become

$$\mu_i = -2.5 \log_{10} \left[10^{-G(M_{i1}, \Theta)/2.5} + 10^{-G(M_{i2}, \Theta)/2.5} \right]. \quad (2.7)$$

For stars with no binary companion, we simply set $M_{i2} = 0$, and ensure that $G(0, \Theta) \equiv \infty$, in which case (2.7) reduces to (2.5).

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 $M < 0.4M_{\odot}$) 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, common envelope evolution, 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.

2.3.3 Field Star Contamination

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, helium abundance, 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:

$$p_{\text{field}}(\mathbf{X}_i) = c_{ij}, \text{ for } \min_j \leq x_{ij} \leq \max_j \text{ for } j = 1, \dots, n, \quad (2.8)$$

and zero elsewhere, where \mathbf{X}_i is the row of \mathbf{X} containing the observed magnitudes for star i , (\min_j, \max_j) is the range of values for magnitude j and $c_{ij} = \left[\prod_{j=1}^n (\max_j - \min_j) \right]^{-1}$ is a normalization constant.

In principle, a probability map could be created from e.g., an adjacent field or a generalized map of field stars at a specific Galactic latitude, provided

that a) the map is properly normalized, and b) 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.

We construct our final likelihood by assuming that the cluster is a mixture of cluster stars and field stars. We introduce a new vector, $\mathbf{Z} = (Z_1, \dots, Z_N)$, where $Z_i = 1$ if the star is a member of the cluster, and $Z_i = 0$ if it is not. Our final likelihood is:

$$L(\mathbf{M}, \mathbf{\Theta}, \mathbf{Z} \mid \mathbf{X}, \Sigma) = \prod_{i=1}^N \prod_{j=1}^n [(Z_i)p_{\text{cluster}}(\mathbf{X}_i) + (1 - Z_i)p_{\text{field}}(\mathbf{X}_i)] \quad (2.9)$$

with $p_{\text{field}}(\mathbf{X}_i)$ as defined in (2.8) and

$$p_{\text{cluster}}(\mathbf{X}_i) = \frac{1}{\sqrt{2\pi}\sigma_{ij}^2} \exp \left(-\frac{\left\{ x_{ij} + 2.5 \log_{10} \left[10^{-G_j(M_{i1}, \mathbf{\Theta})/2.5} + 10^{-G_j(M_{i2}, \mathbf{\Theta})/2.5} \right] \right\}^2}{2\sigma_{ij}^2} \right). \quad (2.10)$$

2.3.4 Prior Distributions

The Bayesian approach allows us to quantify our degree of prior knowledge in a systematic, objective, and repeatable way. In the Bayesian method, each variable (e.g. the components of $\mathbf{\Theta}$ and \mathbf{M}) requires a prior probability distribution.

For the prior on the primary mass, we choose to employ a version of the Miller & Scalo (1979) initial mass function (IMF):

$$p(\log_{10}(M_{i1})) \propto \exp\left(-\frac{1}{2}\left(\frac{\log_{10}(M_{i1}) + 1.02}{0.677}\right)^2\right), \quad (2.11)$$

truncated and normalized to the range $0.1M_{\odot}$ and $8.0M_{\odot}$. As stated in Section 2.2, we have chosen $8.0M_{\odot}$ as the largest ZAMS mass that will form a white dwarf. Anything more massive than this is not of interest. Stars with masses less than $\sim 0.1M_{\odot}$ will not undergo hydrogen fusion in their cores. The prior on the secondary mass is flat between 0 and the primary mass.

For clusters with parallaxes, proper motions, and/or radial velocities of individual stars, we can place individual priors on cluster membership for each star. For less well-studied clusters, we use a single cluster-wide prior based on an estimate of the total ratio of cluster members to field stars. In practice, cluster membership posterior probabilities are often insensitive to the choice of prior.

With the exception of age, the priors on the components of Θ are Gaussian distributions with parameters specific to the cluster being studied. Often the means and variances of the distributions will be chosen based on the most recent and/or most reliable determinations in the literature. In practice, as with cluster membership priors, we have found that narrow prior distributions are often unnecessary for precise results. Generally, the more stars in the cluster, the more the posterior is driven by the likelihood rather than the priors.

Since one goal of this project is to compare ages derived from the MSTO with those derived from the WDs, and to compare our method objectively with other methods of age determination, we have chosen to place a more generalized prior on age than on the other cluster parameters. Specifically, we have chosen a uniform prior on θ_T (the \log_{10} of the age) between a minimum and maximum value determined by the specific combination of model sets being used. This is equivalent to a power law prior on age with an exponent of -1, reflecting the fact that young clusters are more prevalent than older clusters. In practice, we have found the difference between this prior and a uniform (flat) prior on age to be negligible.

2.4 Statistical Computations

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.

However, given the incredibly complex nature of stellar evolution, there is little hope for a closed-form expression for \mathbf{G} . For our problem, the normalization in (2.2) is an intractable integral in $3N + 4$ dimensions (or $3N + 5$, if we include helium abundance). We must turn to a numerical method for help in evaluating the posterior probability.

2.4.1 Markov Chain Monte Carlo

We use Markov chain Monte Carlo (MCMC— Casella & George, 1992; Chib & Greenberg, 1995) to generate samples from the posterior distributions of each parameter. 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, quantiles, etc. 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:

$$\alpha = \min \left[\frac{p(\theta^*|y)q(\theta^{(t)}|\theta^*)}{p(\theta^{(t)}|y)q(\theta^*|\theta^{(t)})}, 1 \right] \quad (2.12)$$

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 $(\theta^{(s)}, \theta^{(s+1)}, \dots, \theta^{(S)})$ where S is the total number of iterations and $s - 1$ is the number of iterations before the chain converges, which is referred to as the burn-in.

In plainer language, we start with values for each of the $3N + 4$ (or $3N + 5$) parameters of our model (i.e., the cluster parameters, each stellar primary and secondary mass, and values for each star’s cluster membership variable, Z_i). For each of these parameters in series, we propose a new value for the parameter, holding all the other parameters constant (the distributions, $q(\theta^*|\theta^{(t)})$, from which we draw the proposed values are discussed further in Section 2.4.2). We calculate the posterior probabilities for the current value of the parameter and for the proposed new value for the parameter. We then construct the Metropolis-Hastings factor in (2.12), compare it to a random variable on the interval $[0,1]$, accept or reject the proposed value for the parameter, and then move on to the next parameter. We discard the first $s - 1$ iterations and treat the remaining iterations as random draws from the posterior distributions for each parameter, from which we can calculate statistics of interest.

The computational advantage of Metropolis-Hastings algorithm derives from the fact that it allows us to generate random samples from a distribution known only up to a normalization constant. In the case of a Bayesian calculation, this obviates the need to calculate the intractable integral in (2.2), since in the Bayesian method, the posterior distribution is proportional to the product of the likelihood and the prior distribution, as in (2.1).

Replacing the posterior probabilities, $p(\theta^*|y)$ and $p(\theta^{(t)}|y)$ in 2.12 via Bayes’ equation (2.1), and supressing the iteration superscripts, the Metropolis-

Hastings factor becomes

$$\begin{aligned}\alpha &= \min \left[\frac{p(\theta^*|y) q(\theta|\theta^*)}{p(\theta|y) q(\theta^*|\theta)}, 1 \right] \\ &= \min \left[\frac{p(y|\theta^*) p(\theta^*)}{\int p(y|\theta^*) p(\theta^*) d\theta^*} \frac{\int p(y|\theta) p(\theta) d\theta}{p(y|\theta) p(\theta)} \frac{q(\theta|\theta^*)}{q(\theta^*|\theta)}, 1 \right].\end{aligned}\quad (2.13)$$

Since $\int p(y|\theta^*) p(\theta^*) d\theta^* = \int p(y|\theta) p(\theta) d\theta$, (2.13) reduces to

$$\alpha = \min \left[\frac{p(y|\theta^*) p(\theta^*) q(\theta|\theta^*)}{p(y|\theta) p(\theta) q(\theta^*|\theta)}, 1 \right] \quad (2.14)$$

where $p(y|\theta)$ and $p(\theta)$ are the likelihood and prior densities, respectively.

2.4.2 Sampling and Correlation

By definition, a Markov chain has the property

$$p(\theta^{(t+1)}|\theta^{(0)}, \theta^{(1)}, \dots, \theta^{(t)}) = p(\theta^{(t+1)}|\theta^{(t)}). \quad (2.15)$$

That is, the distribution of the current link of the chain depends only upon the value of the link immediately preceding it, and not on any other previous value. The efficiency of the sampler depends heavily on the choice of proposal distribution $q(\theta^*|\theta^{(t)})$.

For most of the parameters in our problem, we use a “random walk” sampler, i.e., a symmetrical proposal distribution centered on the current value of the parameter. The advantages of this sampler include its simplicity and ease of implementation. For a symmetric step, $q(\theta^*|\theta) = q(\theta|\theta^*)$, and the proposal distributions cancel in (2.14).

The size and shape of the steps is of critical importance, however. If the step size is too small, the chain may take too long to adequately traverse

parameter space. If the step size is too large, the sampler may propose values of the parameter in areas of low probability, leading to excessive rejection of proposed values.

For the cluster-wide parameters except age, we sample from a T-distribution with six degrees of freedom. For age we use a uniform distribution except during the first 100 samples, when we use a much wider Gaussian distribution to ensure that if the initial age is grossly in error, the sampler can rapidly find its way to an area of higher probability.

For the individual stellar parameters of primary mass and mass ratio (a more detailed explanation of the latter follows in Section 2.4.2.1), we use uniform distributions. The mass ratio, being confined to the interval $[0,1]$, is allowed to reflect at the boundaries. Proposal distributions for a star’s cluster membership status are more complicated and are described below in Section 2.4.2.2.

2.4.2.1 Removing Linear Autocorrelations Between Variables

An additional problem arises when two (or more) parameters are highly correlated. For example, in our problem, an increase in the age of the cluster requires a decrease in the mass of a white dwarf to keep its modeled photometry near the observed photometry. The two-dimensional probability cloud looks something like Figure 2.3. A naive “random walk” sampler can take steps in age only about as wide as the blue line before stepping into an area of low probability. Since the parameters are updated in series, the chain must take a

series of small steps in mass and age (e.g., the red arrows) in order to traverse the entire range of both parameters. The result is a highly correlated chain that looks something like the first half of Figure 2.4.




Figure 2.3: The correlation between ZAMS mass and age for a particular white dwarf. Taking steps in the x and y directions (red arrows) results in an inefficient sampler. If instead we take steps along the direction of the green arrows, we can potentially traverse the entire distributions of both age and mass in any given step.

The solution to the problem involves reparameterizing so that instead of sampling in the X–Y direction, we sample *along* the axis of correlation (i.e., the green arrows in Figure 2.3). For linear correlations, such as that in




Figure 2.4: The Markov chain for a particular cluster before and after the process of decorrelation outlined in this section. While the chain in the first half of the figure is drawn from the same distribution as the chain in the second half, the latter will cover the entire range of the distribution with considerably less computational time than the former.

Figure 2.3, this transformation is relatively simple. For the case of the WD age-mass correlation, we introduce a new parameter, U_i , and a new constant, $\beta_{T,i}$, defined by:

$$M_i^{(t)} = \beta_{T,i} \left(\theta_T^{(t)} - \hat{\theta}_T \right) + U_i^{(t)}, \quad (2.16)$$

where $M_i^{(t)}$, $U_i^{(t)}$, and $\theta_T^{(t)}$ are the mass, decorrelated mass parameter, and logarithm of the cluster age of the i th star at the t th iteration, respectively,

and $\hat{\theta}_T$ is the approximate mean log cluster age. Then, rather than directly sampling mass, we sample on U_i for each star. The MCMC algorithm computes the mass at each iteration from (2.16). The results after this decorrelation process look more like the second half of Figure 2.4, where each iteration is largely independent of the previous iteration, and the chain is able to traverse parameter space more freely, quickly, and completely.

Many other such autocorrelations exist. Fortunately, these correlations are usually close enough to linear that we can remove them with successive versions of the same process.

One exception is the correlation between the primary and secondary masses of each star. To remove this correlation, we first define the mass ratio, $R_i = M_{i2}/M_{i1}$, and sample on R_i for each star, rather than M_{i2} .⁴ Figure 2.5 shows the correlation between primary mass and mass ratio for one particular star. Due to the tightness and shape of the correlation, a linear fit is insufficient. Instead, we fit the correlation with a power law:

$$M_{i1} = U_i + \beta_{R,i} R_i^{\alpha_i}. \quad (2.17)$$

The full set of transformations, worked out using a little bit of astrophysics, a modicum of statistical intuition, and a lot of trial and error, are

$$M_{i1} = U_i + \beta_{R,i} R_i^{\alpha_i} + \beta_{T,i} (\theta_T - \hat{\theta}_T) + \beta_{[\text{Fe}/\text{H}],i} (\theta_{[\text{Fe}/\text{H}]} - \hat{\theta}_{[\text{Fe}/\text{H}]})$$

⁴The reasons for this reparameterization have as much to do with the historical development of the code than for any concrete statistical or astrophysical reason.




Figure 2.5: The correlation between primary mass and mass ratio for a particular star. Since a linear fit (red line) is inadequate, we use a power law to effectively sample along the blue curve.

$$+ \beta_{m-M_V, i} \left(\theta_{m-M_V} - \hat{\theta}_{m-M_V} \right) \quad (2.18)$$

$$\theta_{A_V} = V + \gamma_{[\text{Fe}/\text{H}]} \left(\theta_{[\text{Fe}/\text{H}]} - \hat{\theta}_{[\text{Fe}/\text{H}]} \right) + \gamma_{m-M_V} \left(\theta_{m-M_V} - \hat{\theta}_{m-M_V} \right) \quad (2.19)$$

where hats denote approximate means. These, as well as the components of β and γ are calculated dynamically via linear regression during a series of periods in an initial run before the main run. Ultimately, then, the actual variables that MCMC samples on are $\{(U_1, R_1), \dots, (U_N, R_N), \theta_T, \theta_{[\text{Fe}/\text{H}]}, \theta_{m-M_V}, \theta_Y, V\}$.

It is important to note that we remove these correlations solely for

reasons of computational efficiency. The chains in the two halves of Figure 2.4 are drawing from the exact same posterior distribution. The difference is that the chain in the first half will require considerably more samples before we can be confident that the entire distribution is represented. Were we blessed with an infinitely fast processor chip (or an infinite amount of time), the results of both runs would be identical.

2.4.2.2 Field Star Calculations

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 the values they are allowed to take on 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 two mass 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 initial run 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 initial run 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.

At each iteration then, 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 ac-

cept 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 then 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 as a cluster member will contribute minimally to the final cluster parameter distributions. A star that spends every iteration as a cluster member will contribute fully.

This ability to marginalize over a large number of binary variables is one of the abilities for which MCMC is uniquely suited. Its applications extend beyond cluster membership. Likely extensions of the principle in our work will include WD parameters such as hydrogen (DA) vs. helium (DB) atmospheres and carbon-oxygen vs. oxygen-neon cores.

2.4.3 Starting Values For Parameters

In theory, the posterior should not depend on the choice of initial values. In practice, this is usually the case, provided that the starting values do not

put the posterior in an area of such low probability that MCMC is unable to find its way out. However, the closer the starting values are to the eventual posterior means, the fewer iterations we need to run to get the samplers to “burn in”.

For all of the cluster-wide parameters except age, we start at the means of the prior distributions. A best guess for the starting age is input by the user at run time.

Starting values for masses can either be input by the user or calculated by the code before the start of the run. The latter is accomplished by creating an isochrone at the starting values for the cluster parameters and using the observed magnitude of each star in a given band to derive the masses of the stars from the isochrone. Starting mass ratios are trickier, as any universal choice for all stars will either over- or underestimate the position of the MS and binary sequence. Currently, the user must input a starting mass ratio for each star, regardless of whether the initial primary mass is input or calculated by the code. The starting values of each star’s cluster membership status are also input by the user.

2.4.4 Initial Run Calculations

We start with an initial wandering period (called the “burn-in”) that allows the parameters to find their way to an area of high posterior probability. The length needed for this period depends highly on the choice of starting values for the parameters. The closer they are to their eventual posterior

mean values, the less time the sampler will need to burn in.

This initial wandering period is followed by seven (usually shorter) periods used to calculate the various correlation factors described in Sections 2.4.2.1 and 2.4.2.2. Details on these periods are contained in Table 2.4.4

During all of these periods, the step size for proposal distributions for each parameter that uses a symmetric step distribution (everything except for cluster membership status) is dynamically adjusted. If the acceptance rate for the parameter drops below 20%, or rises above 30% during the previous 200 samples, the step size is decreased or increased respectively by a factor that scales according to the acceptance rate.

At the end of the completed initial run, all of the parameters of interest are output to a separate file, allowing the user to re-run the cluster without the computational overhead of running the initial run over again.

Table 2.2: Periods in the initial run used to compute correlation-reducing transformations.

In the initial burn-in period and in the first 6 periods of the initial run each star’s cluster membership status is held constant. That is, the Z_i ’s are not updated from the starting values input by the user.

0. Burn-in period.
1. Compute each $\beta_{R,i}$ by regressing each $U_i^{(0)}$ on R_i .
2. Compute each $\beta_{T,i}$ by regressing each $U_i^{(1)}$ on θ_T . During this period $\theta_{[\text{Fe}/\text{H}]}$, θ_Y , θ_{m-M_V} , and θ_{A_V} are fixed at our best estimate of their posterior means (as determined dynamically during period (1)).
3. Compute each $\beta_{m-M_V,i}$ by regressing each $U_i^{(2)}$ on θ_{m-M_V} . Compute γ_{m-M_V} by regressing $V^{(0)}$ on θ_{m-M_V} .
4. Compute each $\beta_{[\text{Fe}/\text{H}],i}$ by regressing each $U_i^{(3)}$ on $\theta_{[\text{Fe}/\text{H}]}$. Compute $\gamma_{[\text{Fe}/\text{H}]}$ by regressing $V^{(1)}$ on $\theta_{[\text{Fe}/\text{H}]}$.
5. Approximate the posterior mean and variance of $U_i^{(3)}$ and R_i to construct the alternative prior distributions on the masses for field stars.
6. Fine tune step sizes used in the Metropolis proposals to optimize acceptance rates.

In a second set of 7 periods, the above runs are repeated (including a second burn-in period), but this time the cluster memberships are sampled.

Step sizes for all parameters are adjusted continuously throughout all of the initial runs.

Chapter 3

A New WD Age For The Hyades

More is known about the Hyades open cluster than perhaps any other star cluster in the sky. As such, it is a natural choice to be 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, and for beginning the process of calibrating our bright white dwarf technique, placing it on an absolute age scale.

Moreover, prior to the publication of DeGennaro et al. (2009), the most reliable estimate for the age of the Hyades derived primarily from WD cooling physics (300 Myr; Weidemann et al., 1992) was a factor of two too small when compared with the 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. In DeGennaro et al. (2009), our “bright white dwarf” technique allowed us to publish a WD age for the Hyades (648 ± 45 Myr) that was, for the first time, consistent with the MSTO age.

Since the publication of that paper, we have expanded on our results. The Hyades, because of its nearness to us, not only subtends a considerable angle on the sky, but also spans a depth ($\sim 10\text{-}20$ pc) that is comparable to its distance from us (~ 50 pc). As such, the nearest and furthest stars in the cluster show a difference in distance modulus of nearly 2 magnitudes.

In this chapter we update the results of DeGennaro et al. (2009) by using secular parallaxes derived from Hipparcos proper motions to correct for the different distances to individual members of the Hyades. This not only considerably tightens the main sequence, but also puts to rest one long-standing mystery regarding several of the white dwarfs, which have the kinematics of cluster members and yet appear too bright or too faint to lie on the proper WD sequence.

We derive a new WD age for the Hyades (610 ± 110 Myr) which is both more accurate and more reliable than our previous result. Though the formal errors on this determination are larger than those we have previously published, we believe that they more accurately reflect the true internal uncertainty in our age determination.

3.1 The Data

The most reliable estimates for the age and the distance to the center of the Hyades ($T = 625 \pm 50$ 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 adopt 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 priors on the cluster parameters. The age prior is flat in $\log(\text{age})$ between the limits of our models.

The photometry data for the Hyades come from The General Catalogue of Photometric Data (GCPD¹ — Mermilliod et al., 1997). For each Hyades star in the database, they calculate weighted means 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), 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 January 2008, 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 method. For stars with two or fewer sources, the re-

¹<http://obswww.unige.ch/gcpd/gcpd.html>

Table 3.1: A list of WDs in the Hyades, with cross-references.

GCPD ID	Reid (1996)	EGGR ¹	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

¹Eggen & Greenstein (1965)

ported 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), 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 4 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 8 WDs as members of the Hyades. Table 3.1 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 determi-




Figure 3.1: A 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.

nations (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 complete data set is plotted in Figure 3.1, along with isochrones for three different MS model sets.




Figure 3.2: An expanded view of the WD portion of Figure 3.1. The two objects in red are LB227 and HZ4, both of which are standard stars in Landolt & Uomoto (2007), and thus have unusually small error bars. That they fail to lie along the WD cooling sequence, despite having kinematics that strongly suggest their membership in the Hyades, is accounted for by the fact that they are at significantly different distances from our vantage point on Earth.

3.1.1 Correction For Differential Distance

The precision of the error bars on the photometry for LB227 and HZ4 (shown as the red points in Figure 3.2) highlights a problem one encounters with very few open clusters. The Hyades is close enough to the Earth that the tidal radius of the cluster (~ 10 pc) is a significant fraction of its distance from us (~ 45 pc). Perryman et al. (1998) find members as far as 20 pc from the center of mass of the cluster. 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 magnitudes too bright. A star 20 pc too distant will appear ~ 0.8 mags too faint. In this case, our assumption that all of the stars in the cluster are at a single distance breaks down, and we must take into account the differential distances to each individual star.

In DeGennaro et al. (2009), we chose to address this issue by noting 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 assumed an unknown systematic error of 0.1 mag in the V band, and used the errors in color as quoted in Landolt & Uomoto (2007). 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 translated to similar errors in the other two magnitudes.

However, in presenting our results at a conference, an alternative method

was suggested to us by Eric Mamajek of the University of Rochester².

The nearness of the Hyades itself provides a solution to the problem, as it is also one of the only clusters close enough to the Earth to have measured geometric distances to most individual members. The Hipparcos and Tycho catalogs (Perryman & ESA, 1997) contain individual trigonometric parallaxes for most of the brightest Hyades members. The white dwarfs and lower main sequence stars in the Hyades are too faint for Hipparcos trigonometric parallaxes, but do have measured proper motions.

de Bruijne et al. (2001) used the moving cluster method to derive the parameters for the true space velocity (and velocity dispersion) of the Hyades from Hipparcos proper motions. With this information in hand, distances to individual stars can be worked out from their individual proper motions. In the simplest case, assuming negligible velocity dispersion,

$$\pi_{sec} = \frac{A|\mu|}{|\mathbf{v}|\sin\lambda}, \quad (3.1)$$

where \mathbf{v} is the space motion of the cluster, μ is the proper motion vector of the star, λ is the angle on the sky between the star and the cluster convergent point, and A is a constant equal to the ratio of one astronomical unit in kilometers to the number of seconds in a Julian year ($=4.740470446 \text{ km yr s}^{-1}$).

de Bruijne et al. (2001) performed a more complicated analysis, taking the velocity dispersion of the cluster into account and simultaneously fitting

²I am grateful to Dr. Mamajek not only for suggesting this line of inquiry, but also for helping to locate the relevant papers in the literature and performing the initial calculations of distances to the Hyades WDs.

Table 3.2: Proper motions for the Hyades white dwarfs from various sources. Adapted from a table compiled by E. Mamajek (2008—private communication). The proper motion determinations come from the USNO-B1.0 Catalog (Monet et al., 2003), the PM2000 Bordeaux Proper Motion Catalogue (Ducourant et al., 2006), the Second U.S. Naval Observatory CCD Astrograph Catalog (Zacharias et al., 2004), and the Lick NPM2 Catalog (Hanson et al., 2004). The symbols ν and τ refer, respectively, to motion in the direction of, and perpendicular to, the cluster convergent point. V_{pec} is the velocity perpendicular to the motion of the cluster. All of these stars show kinematics in strong agreement with the cluster as a whole.

WD Name	Catalog index	R.A. (deg)	Dec (deg)	PM ν (mas/yr)	PM τ (mas/yr)	$\sigma_{PM\nu}$ (mas/yr)	$\sigma_{PM\tau}$ (mas/yr)	D (pc)	σ_D (pc)	π (mas)	σ_π (mas)	V_{pec} (km/s)	σ_V (km/s)
VR16	PM2000_171118	67.1641354	16.9700192	100.37	-4.55	5.90	5.90	50.4	3.0	19.86	1.18	-1.1	1.4
	USNO-B1.0_1069-0	67.1639528	16.9700361	81.22	-8.86	8.72	3.74	62.2	6.7	16.07	1.75	-2.6	1.1
VR7	PM2000_167294	65.9820137	16.3542164	119.49	1.02	5.80	5.80	43.4	2.1	23.02	1.13	0.2	1.2
	USNO-B1.0_1063-0	65.9820083	16.3542444	113.03	0.49	3.35	6.84	45.9	1.4	21.77	0.66	0.1	1.5
HZ14	PM2000_181740	70.2570179	10.9944556	89.68	-0.95	2.20	2.30	49.5	1.3	20.19	0.51	-0.2	0.5
	UCAC2_35516984	70.2572331	10.9944406	91.06	-0.69	11.23	2.04	48.8	6.0	20.50	2.57	-0.2	0.5
HZ4	USNO-B1.0_0997-0	58.8416000	9.7883806	167.86	7.95	4.00	4.00	36.0	0.9	27.79	0.69	1.4	0.7
HZ7	NPM2_+12.0517	68.4375600	12.7113300	115.76	5.34	6.00	6.00	41.0	2.1	24.41	1.28	1.0	1.2
	PM2000_175308	68.4361821	12.7114225	104.47	0.54	0.70	0.70	45.4	0.4	22.02	0.21	0.1	0.2
	UCAC2_36236653	68.4373558	12.7112153	99.14	2.23	2.70	2.60	47.8	1.3	20.90	0.59	0.5	0.6
LB227	PM2000_157579	62.3703196	17.1317700	110.33	-5.19	5.30	5.30	51.6	2.5	19.38	0.94	-1.3	1.3
	USNO-B1.0_1071-0	62.3703778	17.1318222	106.73	-0.97	2.06	2.96	53.3	1.1	18.74	0.38	-0.2	0.7

the various parameters of the velocity field and convergent point. In the end, they provide a table of likely Hyades members on the main sequence with individual secular parallaxes they claim are ~ 3 times more precise than the Hipparcos trigonometric parallaxes.

Table 3.2 shows the individual proper motion determinations for the WDs in the Hyades, where the proper motions come from the USNO-B1.0 Catalog (Monet et al., 2003), the PM2000 Bordeaux Proper Motion Catalogue (Ducourant et al., 2006), the Second U.S. Naval Observatory CCD Astrograph Catalog (Zacharias et al., 2004), and the Lick NPM2 Catalog (Hanson et al., 2004). For stars with multiple measurements, we take a weighted average of the individual proper motion determinations for each star and use (3.1), together with the cluster velocity and convergent point determined by de Bruijne et al. (2001) to calculate the final secular parallax for each WD.

We match these stars up with the photometry from the GCPD and Landolt & Uomoto (2007), and use the parallax to determine absolute magnitudes for each star. We then propagate the errors in the distance determinations and the apparent magnitudes to determine the total error in the absolute magnitude in each band for each star.

Theoretically, correcting the photometry in this manner puts every star at a distance of exactly 10 pc, with the errors in the distance determinations incorporated into the errors in the photometry. However, the “individual” distance determinations are not, in fact, independent at all, but rely on a common set of parameters determined by fitting the cluster as a whole (e.g., the

components of the velocity field and the cluster convergent point). Correcting for the individual distances, even accounting for the errors in the proper-motion determinations, leaves an untreated systematic error in the distance.

de Bruijne et al. (2001) find an upper limit to the systematic uncertainty in their secular parallaxes to be ~ 0.3 mas. At the mean distance to the Hyades, this translates to an error in the distance modulus of ~ 0.03 . To account for this systematic error, we use this value as the uncertainty on the prior on the cluster distance. This is probably an overly conservative estimate, as the 0.3 mas figure cited is the *maximum* systematic error that could hide in their analysis without detection. Elsewhere, they show that their secular parallaxes, as a set, agree with Hipparcos trigonometric parallaxes to within 0.1 mas. Our own calculations comparing their fits to the cluster velocity and convergent point with five previous fits in the literature, indicates a total systematic of only ~ 0.2 mas. We choose to use the more conservative value of 0.3 mas in our analysis, and discuss the effect of this choice in Section 3.2.1.

Figure 3.3 shows the CMD of the Hyades in absolute magnitude space after correcting for differential distances. The main sequence is considerably tighter, with the remaining dispersion accounted for almost entirely by binaries and photometric scatter. There is essentially no field star contamination.

Figure 3.4 shows an expanded view of the WD portion of Figure 3.3. Accounting for differential distances has moved the WDs in the expected directions, with both LB227 and HZ4 (as well as the other WDs) now lying in a much tighter sequence. While the isochrones fit the $B - V$ data well, the




Figure 3.3: Color-magnitude diagram of the Hyades in absolute magnitude space, after correction for differential distances. Three different sets of MS models are overplotted. Solid (purple) line=Girardi models, dotted (red) line=Yale-Yonsei models, dashed (blue) line=DSED models.

$U - B$ isochrones clearly lie blue-ward of the actual data, which may indicate a problem with the U band in either the data or the WD model atmospheres.

3.1.2 Problems with the Main-Sequence Models

Figure 3.3 also shows the differences among the main-sequence models at the nominal age, metallicity, distance, and reddening of the Hyades. There are some differences in the position of the turnoff, as well as some important




Figure 3.4: A close up of the white dwarf region of Figure 3.3. The crosses are the Hyades WDs, the other symbols show the positions of individual theoretical WDs of a given mass along the WD tracks. Solid (purple) line=Girardi models, dotted (red) line=Yale-Yonsei models, dashed (blue) line=DSED models.

differences in the shapes of the main sequence. These differences represent a long-standing problem in stellar astrophysics (for recent discussions see von Hippel et al., 2002; Grocholski & Sarajedini, 2003; Jeffery, 2009).

The proper functioning of our code depends on the models' being accurate. However, our goal in this paper is to derive cluster ages, in particular WD ages, on which the MS has a limited impact. The primary function of the MS

in our method is to help pin down the cluster parameters of distance, metallicity, and absorption, yielding more precise ages from the WDs. However, as we will shortly show, 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.

We note also that the differences between the various main-sequence model predictions are a measure of the uncertainty in the theoretical underpinnings of the models. As such, results obtained under different sets of models will represent the spread in cluster parameters that result from these underlying theoretical uncertainties.

The WD tracks are nearly coincident for all three sets of MS models. A closer look at the WD sequences (Figure 3.4) 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 white dwarfs' 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 $5M_{\odot}$, the different MS models yield nearly identical positions. Note that these discrepancies in the upper part of the sequence will result in different mass determinations for individual stars, but will not meaningfully alter the derived age for the cluster (see Section 1.4.2).

3.2 Results

Since we are interested in a comparison between ages derived from traditional main-sequence turn-off 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 $M_V < 1.5$ (shown in gray in Figure 3.3). There is still some age information in the MS insofar as the lack of a turnoff fainter than $M_V = 1.5$ sets an upper limit on the age. Our bright WD technique in part exploits this phenomenon.

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 $M_V = 1.5$ and some low cutoff.

Figures 3.5 through 3.8 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 3.1. 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 under the assumption of a specific set of models is comparable to, and in some cases better than, the most precise values measured to date using other methods. This is particularly surprising




Figure 3.5: Our derived $\log(\text{age})$ for the Hyades as a function of the faintest magnitude main-sequence 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 turn-off 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.

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 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




Figure 3.6: 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 3.5. The metallicity value comes from Taylor & Joner (2005) and represents the prior information used in the analysis.

uniform between the model limits, and the horizontal lines in the $\log(\text{age})$ diagram merely represent the most widely accepted turn-off age for the Hyades (Perryman et al., 1998). Thus, the excellent agreement between our WD age and the MS turn-off age is *not* a function of prior information but inherent in the data.




Figure 3.7: Our derived distance modulus for the Hyades as a function of the faintest magnitude main-sequence star included in the analysis. The symbols have the same meaning as in Figure 3.5. The distance value is an estimate of the systematic uncertainty in the distances derived via Hipparcos secular parallaxes in de Bruijne et al. (2001) and represents the prior information used in the analysis.

The metallicity (Figure 3.6) is particularly influenced by the shape and slope of the main sequence. When the WDs are run through the code by themselves (i.e., a MS cutoff of $M_V = 1.5$ — 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




Figure 3.8: Our derived A_V for the Hyades as a function of the faintest magnitude main-sequence star included in the analysis. The symbols have the same meaning as in Figure 3.5. The A_V value comes from Taylor (1980) and represents the prior information used in the analysis.

slope contribute more and more to the likelihood, and thus the posterior.

In all of the parameters, the values show more variability in the left half of the figures, but begin to stabilize as we include main sequence stars fainter than $M_V \simeq 6.5$. With a shorter main sequence, the method has more wiggle-room in fitting the parameters, but as we include the lower part of the MS, we constrain the possible values that the parameters can take on and still adequately fit the entire sequence.

The discrepancy in the distance for the DSED models (blue) can be at least partly understood by a closer examination of Figure 3.3, where we can see that the DSED models tend to overestimate the B luminosity for a given $U - B$ color on the main sequence (the same is true to a lesser extent in the right half of the figure as well). This discrepancy does not affect the distance determination when we only include stars down to $M_V \sim 5.5$ because the MS in the B vs. $U - B$ figure is nearly vertical. But as we include the lower part of the main sequence, the code compensates for this problem in the models by moving the modeled cluster further away to better match the observations. The underestimation of the distance in the case of the Yale-Yonsei models (red) can be understood in the same way.

The slight offset in the WD sequence in the B vs. $U - B$ diagram—though it is probably a color effect—can also be reduced by adjusting the cluster distance, in this case by moving the theoretical cluster closer (i.e., upward in the CMD). The first few points for all three model sets in Figure 3.7—where we include little or no MS information—show precisely this trend.




Figure 3.9: Color-magnitude diagram of the Hyades in absolute magnitude space, after correction for differential distances. Three different sets of MS models are overplotted, this time at the values for the cluster parameters determined by our method. Solid (purple) line=Girardi models, dotted (red) line=Yale-Yonsei models, dashed (blue) line=DSED models.

Figure 3.9 shows the theoretical isochrones for the three different models sets at the values determined by our method when we include the entire MS (i.e., , the rightmost points in Figures 3.5 through 3.8).

These isochrones fit better overall than the isochrones created with the cluster parameters fixed at their prior values. Some of the trade-offs required to fit the stars as a whole are apparent. In particular, the lower part of the WD cooling sequence has shifted varying amounts for the three model sets in order to fit the coolest two WDs in both the B vs. $U-B$ diagram and the V vs. $B-V$ diagram. The shape of the upper MS for the DSED models (blue) has adjusted red-ward in the $U-B$ diagram, causing a slight shift in the opposite direction in the top of the MS in the $B-V$ diagram. These subtle trade-offs highlight the advantage our method has over traditional “chi-by-eye” fitting, which has considerable difficulty handling such nuanced distinctions in a consistent and repeatable way. At the same time, they also highlight the limitations in the stellar evolution models, while also providing valuable feedback to those who create such models.

3.2.1 Cluster-wide Parameter Determinations

Figures 3.5 through 3.8 show graphically that the choice of model and the amount of main sequence included have a significant effect on the results we obtain. However, as pointed out in Section 3.2 the values begin to stabilize as we include main sequence stars fainter than $M_V \simeq 6.5$, essentially flattening out well before $M_V = 10.0$.

We have chosen to use the final point in each figure for each model set to determine the values for the cluster-wide parameters for the Hyades. We take a weighted average of the three models for each parameter for the runs which include the entire MS. 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 (including photometric errors) and the systematic errors caused by differences among the MS models. It does not, however, include other known sources of systematic error, most notably errors in the WD cooling or atmosphere models or in the imprecisely known IFMR.

In fact, the overall systematic uncertainty in the WD age of the Hyades may be dominated by uncertainties in the WD cooling models. We have not yet explored these model systematics. Salaris et al. (2009) 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 610 ± 110 Myr ($\log[\text{age}] = 8.78 \pm 0.07$), an $[\text{Fe}/\text{H}]$ of 0.061 ± 0.051 , and an absorption in the visual band of 0.014 ± 0.007 , all of which are consistent with the best previous results. Of these, only the age really provides much new information of astrophysical

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

The systematic offset in the posterior distribution for the distance modulus (0.03 ± 0.08), is consistent with the claims of de Bruijne et al. (2001) that the distances derived via secular parallax can be considered absolute.

The age we derive here is closer to the accepted value of the MSTO age derived by Perryman et al. (1998) than the value (648 ± 45 Myr) we previously published in DeGennaro et al. (2009), though with a larger formal error. This larger error is due at least in part to the slightly larger variance on the prior on the cluster distance. In fact, when we repeat the analysis with a less conservative, but probably more realistic systematic distance error of 0.01 dex, we derive an age of 602 ± 75 Myr.

3.2.2 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 1.4.2, in Jeffery et al. (2007) 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




Figure 3.10: MSTO versus WD ages for seven clusters, adapted and updated from von Hippel (2005). The age we derive from the white dwarfs in the Hyades using our “bright white dwarf” technique brings the WD age of the Hyades into agreement with the main-sequence turn-off age for the first time (solid red 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 (Weidemann et al., 1992) prior to our results in DeGennaro et al. (2009).

ages for real data, as well as providing an important step in calibrating the technique.

Figure 3.10 is an updated version of Figure 1 from von Hippel (2005), plotting WD age vs. 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 results of the most reliable WD age for the Hyades (Weidemann et al., 1992) prior to our results in DeGennaro et al. (2009).

3.2.3 White Dwarf Mass Determinations

Figure 3.11 shows the individual WD mass determinations for the various runs, with the plotting symbols the same as in Figures 3.5 through 3.8. 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 of possible masses for the WDs, as determined by other methods, including gravitational redshifts and spectroscopic gravities.

3.2.4 Main Sequence Mass Determinations

Bender & Simon (2008) undertook a survey of Hyades stars to look for low-mass binary companions via infrared and optical spectroscopy. Using spectra, they derived dynamical masses for the binaries, then used Hipparcos distances, 2MASS photometry, and a set of main-sequence models to derive




Figure 3.11: Our mass determinations for the Hyades white dwarfs 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).




Figure 3.12: A comparison of our results for the masses of twelve MS-MS binaries with mass determinations from Bender & Simon (2008). The colored points are the individual MCMC samples for one run of our code (which included the entire MS down to $M_V = 10.0$) for each of the three MS model sets. The black points with error bars are the mass determinations of Bender & Simon (2008) based on infrared and optical spectroscopy combined with 2MASS photometry, Hipparcos parallaxes, and an assumed mass-luminosity relationship from theoretical isochrones.

individual masses for the two components of the binary. Twelve of the stars in their sample overlap with stars we fit in our analysis.

Figure 3.12, shows a comparison of our results with theirs for these twelve stars. The colored points are the individual MCMC samples for one run of our code (which included the entire MS down to $M_V = 10.0$) for each of the three MS model sets. The black points with error bars are the primary and secondary masses determined by Bender & Simon (2008).

The shape of the correlation between primary and secondary mass for most of these stars in our results is due to the insensitivity of the photometry to binaries with mass ratios less than ~ 0.5 . Because of the steep relationship between mass and luminosity on the main sequence, companions with less than half the mass of their primaries tend to get lost in the glare of their brighter hosts, especially in the bluer bands used in our analysis of the Hyades (U , B , and V). As such, our method is essentially unable to distinguish a star without a companion from a star with a low-mass companion. This accounts for the essentially vertical portion of the MCMC “curves”. As the mass of the secondary becomes comparable to that of the primary, the mass of the primary must decrease to keep the modeled photometry at the same color and luminosity as the observed photometry, which accounts for the hook-shaped portion of the correlation structure.

Because Bender & Simon (2008) were specifically searching for low mass companions via infrared spectroscopy, most of their results lie in the region where our method is insensitive to the secondary masses. The primary masses

we derive, however, are each individually consistent with the results of Bender & Simon (2008). In aggregate, they show a small systematic offset from the spectroscopic masses, with our masses tending to be smaller. However, spectroscopy can only provide masses up to an inclination angle. Bender & Simon (2008) couple this information with main sequence models, trigonometric parallaxes, and 2MASS photometry to derive the masses of the two components. In doing so, they rely on a theoretical mass-luminosity relation derived from a single set of main-sequence models (different than any of the models sets we used), at a single age and metallicity. These model-dependent factors may account for the small systematic difference we see between their masses and ours.

Peterson & Solensky (1988) also used radial velocity measurements to derive dynamical mass ratios for several Hyades binaries. Of these, one system (vB022) is an eclipsing binary, and using the inclination angle derived by Schiller & Milone (1987), they were able to derive essentially model-independent primary and secondary masses for the system. Figure 3.13 compares their results with ours. Our secondary mass is consistent with their result, but the primary mass is $\sim 5\%$ lower than the more reliable external estimate.

3.2.5 Effect of Less Restrictive Priors

In order to quantify the effects of the prior variances on our derived values for the cluster parameters, we re-analyzed our data using less restric-




Figure 3.13: Similar to Figure 3.12 for the star vB022. In this case, the black point is the mass determined dynamically from radial velocity measurements combined with an inclination angle derived from eclipse light curves (Peterson & Solensky, 1988).

tive priors on distance, metallicity, and reddening. The values we used for the standard deviations of the Gaussian priors—0.3, 0.2, and 0.1 dex for the metallicity, distance modulus, and absorption, respectively—are similar to what we might know for a more typical open cluster.

Figures 3.14 through 3.17 show the results of these runs. While all of the values now lie further away from the prior mean value, two of the model

sets remain well within the errors of the new, less restrictive priors (gray horizontal lines), indicating that for these models the posterior is dominated by the likelihood and not the prior. The third model set (DSED, in blue), which fit the data most poorly on visual inspection, may be relying slightly more on the prior to constrain the fit.




Figure 3.14: Similar to Figure 3.5. For these runs, the variances on the prior probability distributions for the cluster-wide parameters except $\log(\text{age})$ have been increased to values that are more in line with what we might be able to obtain for a typical cluster. We note again that the prior on $\log(\text{age})$ is flat. The horizontal lines in this figure merely represent the most reliable age for the Hyades based on the MSTO (Perryman et al., 1998).

The age we derive with these less restrictive priors is 560 ± 105 Myr ($\log(\text{age}) = 8.75 \pm 0.08$), which is consistent with both the MSTO age and our own WD age from Section 3.2.1. Interestingly, increasing the variances on the priors on distance, metallicity, and reddening appears to have had no effect on the uncertainty in the age.

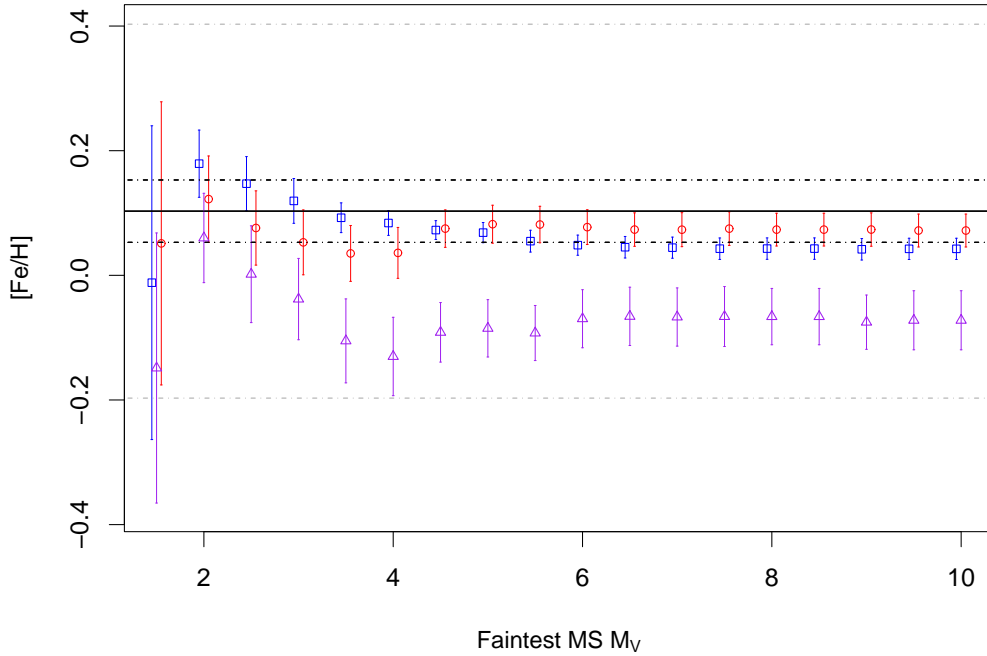


Figure 3.15: Similar to Figure 3.6, except that the errors on the prior probability distributions for the cluster-wide parameters (except $\log(\text{age})$) have been increased to values that are more in line with what we might be able to obtain for a typical cluster. The dark horizontal lines are the same as in Figure 3.6. The gray horizontal lines are the $\pm 1\sigma$ deviations of the new, less restrictive prior.

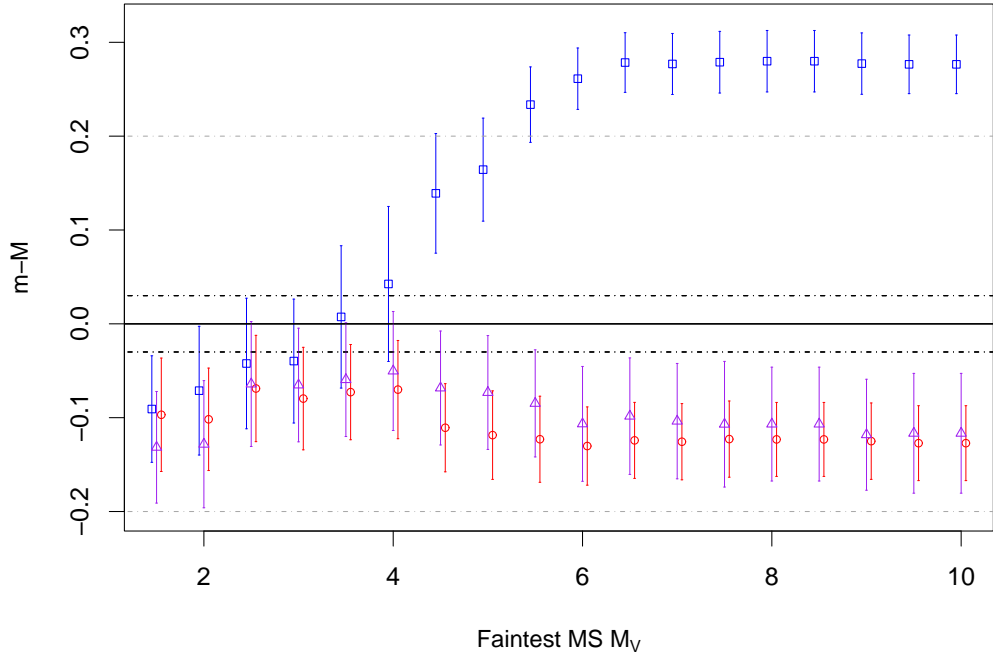


Figure 3.16: Similar to Figure 3.7, except that the errors on the prior probability distributions for the cluster-wide parameters (except $\log(\text{age})$) have been increased to values that are more in line with what we might be able to obtain for a typical cluster. The dark horizontal lines are the same as in Figure 3.7. The gray horizontal lines are the $\pm 1\sigma$ deviations of the new, less restrictive prior.

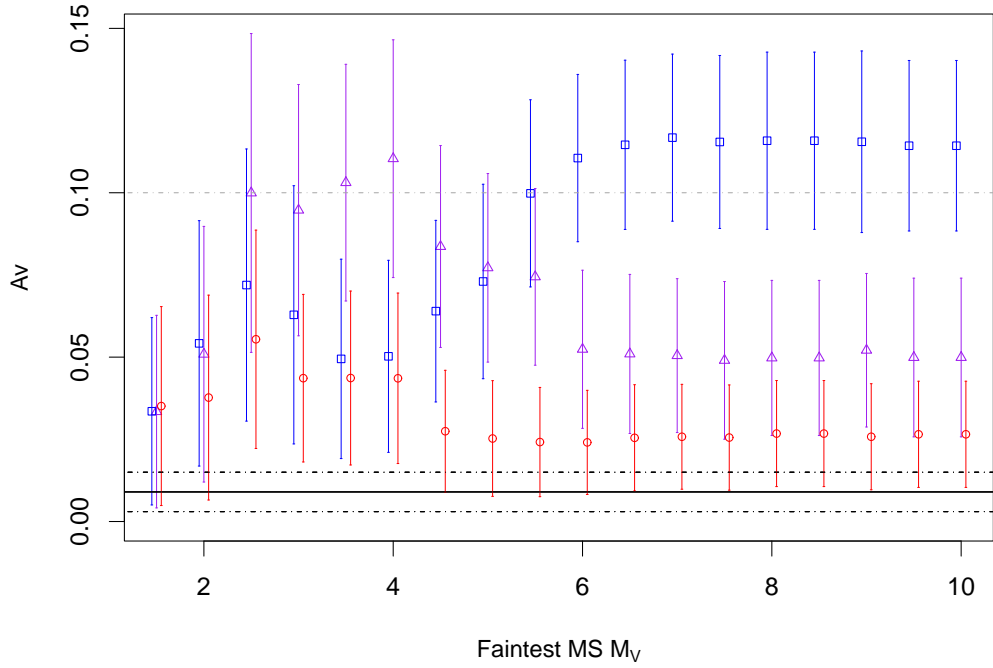


Figure 3.17: Similar to Figure 3.8, except that the errors on the prior probability distributions for the cluster-wide parameters (except $\log(\text{age})$) have been increased to values that are more in line with what we might be able to obtain for a typical cluster. The dark horizontal lines are the same as in Figure 3.8. The gray horizontal lines are the $\pm 1\sigma$ deviations of the new, less restrictive prior.

Chapter 4

White Dwarf Luminosity and Mass Functions from Sloan Digital Sky Survey Spectra

4.1 Introduction

In addition to the open and globular clusters, white dwarfs have long been used as age indicators for other Galactic populations. The same cooling physics applies equally to stars in the thin disk, thick disk, and halo of the Galaxy, and while the specific methods employed to determine the ages of these populations vary, one fact remains constant: the ages we determine are ultimately only as accurate as the WD cooling models themselves, which in turn are only as good as the physics underlying those models.

The history of the development of WD models, as with much in astrophysics, involves a sort of leap-frog process: theoretical considerations yield testable predictions; observations, in turn, illuminate limitations in theory and suggest new avenues of theoretical exploration. At the heart of any astrophysical model is a set of self-consistent physical laws, but given the complexity of the universe, new or more precise observations almost invariably lead to new physics, or at least new treatments of old physics.

In the case of the white dwarfs, one of the oldest and most useful ob-

servational tools for constraining the theoretical models is the white dwarf luminosity function (WDLF). Attempts at an empirical luminosity function (LF) for white dwarfs date as far back as Luyten (1958) and Weidemann (1967). The low-luminosity shortfall, discovered by Liebert et al. (1979), and attributed by Winget et al. (1987) to the finite age of the Galactic disk, was confirmed and explored more fully when a greater volume of reliable data on low-luminosity WDs became available (Liebert et al., 1988; Wood, 1992). More recently, photometric data from the Sloan Digital Sky Survey (SDSS—York et al., 2000) have been used to provide a much more detailed luminosity function with more than an order of magnitude more white dwarfs than previously attempted (Harris et al., 2006), as well as a new LF of a large sample of spectroscopically confirmed WDs (Hu et al., 2007)¹.

To date, however, empirical WD luminosity functions, especially those derived from stars with spectra, have been hampered by a limited volume of reliable data. This has forced a trade-off between the number of stars included in a sample and their homogeneity; either a broad range of temperatures, masses, and spectral types must be used, or else the sample population of stars would be so small as to render reliable conclusions difficult.

In 2004, at the commencement of my Master’s project, our hope was that this situation had changed. Data from the Fourth Data Release (DR4) of SDSS had yielded nearly 10,000 white dwarf spectra. All of these spectra had

¹Hu et al. (2007), based on spectra from the First Data Release of SDSS, suffered from methodological flaws serious enough to render their conclusions virtually meaningless.

been fitted with model atmospheres to determine their effective temperatures and surface gravities (Kleinman et al., 2004; Krzesiński et al., 2004; Eisenstein et al., 2006; Hügelmeyer et al., 2006; Kepler et al., 2007).

With this massive quantity of data in hand, we hoped that for the first time, we could begin to disentangle the many disparate factors that make up the WD LF by focusing solely on a single type of white dwarf (hydrogen atmosphere, or DA), and moreover, by creating for the first time separate well-populated WD luminosity functions for two or more different ranges of mass. Theoretical work had been done in this area (Wood, 1992; Fontaine et al., 2001), but to that date (and, indeed, to this one), attempts at creating an empirical LF to explore the effects of mass have relied on limited sample sizes (Liebert et al., 2005; Catalán et al., 2008).

In addition to helping unlock the physics of white dwarfs, creating luminosity functions for several mass ranges would also help to disentangle the effects of changes in cooling rates from changes in star formation rates. A burst or dip in star formation at a given instant in Galactic history should be recorded in all of the luminosity functions, regardless of mass, and could be confirmed by its position across the various mass bins. For example, a short burst of increased star formation would be seen as a bump in each luminosity function, occurring at cooler temperatures in the higher mass LF (these stars, with shorter MS lifetimes, have had longer to cool). On the other hand, features intrinsic to the cooling physics of the white dwarfs themselves should be seen in places that correspond with the underlying physics, which

may be earlier, later, or nearly concurrent across mass bins. These effects include neutrino cooling, crystallization, phase separation, the onset of convective coupling (Fontaine et al., 2001), and Debye cooling (Althaus et al., 2007).

Ultimately, we discovered that the data available to us was of insufficient quality to allow us to complete this ambitious project in full. Despite these limitations, however, in DeGennaro et al. (2008) we published the first WDLF for spectroscopically determined WDs in the SDSS and the first white dwarf mass function corrected for selection effects. Most of the remainder of this chapter is derived and updated from that paper.

In Section 4.2, we introduce the data, examining the methods used to classify spectra and derive quantities of interest (dominant atmospheric element, T_{eff} , and $\log g$). We also address the observed upturn in $\log g$ for DAs below $T_{\text{eff}} \sim 12,000\text{K}$. We present several lines of reasoning that the upturn is an artifact of the line fitting procedure, and propose an empirical method for correcting the problem. Section 4.3 outlines the methods used to construct the luminosity and mass function and determine error bars.

In Section 4.4, we present an analysis of the completeness of our data sample. We use a well-defined sample of proper-motion selected, photometrically determined white dwarfs in SDSS (Harris et al., 2006) to determine our completeness and derive a correction as a function of $g - i$ color and g magnitude. Finally, in Section 4.5, we present our best luminosity and mass functions for the entire DA spectroscopic sample and discuss the impact of

both our empirical $\log g$ correction and our completeness correction.

4.2 The Data

Our white dwarf data come mainly from Eisenstein et al. (2006), a catalog of spectroscopically identified white dwarfs from the Fourth Data Release (DR4) of the Sloan Digital Sky Survey (York et al., 2000). The SDSS is a survey of $\sim 8,000$ square degrees of sky at high Galactic latitudes. It is, first and foremost, a redshift survey of galaxies and quasars. Large “stripes” of sky are imaged in 5 bands (u,g,r,i,z) and objects are selected, on the basis of color and morphology, to be followed up with spectroscopy, accomplished by means of twin fiber-fed spectrographs, each with separate red and blue channels with a combined wavelength coverage of about 3800 to 9200Å and a resolution of 1800. Objects are assigned fibers based on their priority in accomplishing SDSS science objectives, with high redshift galaxies, “bright red galaxies” and quasars receiving the highest priority. Stars are assigned fibers for spectrophotometric calibration, and other classes of objects are only assigned fibers that are left over on each plate. More detailed descriptions of the target selection and tiling algorithms can be found in Stoughton et al. (2002) and Blanton et al. (2003).

Though white dwarfs are given their own (low priority) category in the spectroscopic selection algorithms, very few white dwarfs are targeted in this way. Rather, most of the white dwarfs in SDSS obtain spectra only through the “back door,” most often when the imaging pipeline mistakes them for quasars.

Kleinman et al. (2004) list the various algorithms that target objects ultimately determined to be white dwarfs in DR1 (their Table 1). White dwarfs are most commonly targeted by the QSO and SERENDIPITY_BLUE algorithms, with significant contributions also from HOT_STANDARD (standard stars targeted for spectrophotometric calibration) and SERENDIPITY_DISTANT. Of the 5 significant contributors, the STAR_WHITE_DWARF category contributes the least to the population of WD spectra.

The SDSS Data Release 4 contains nearly 850,000 spectra. Several groups have already attempted to sort through them to find white dwarfs: Harris et al. (2003) for the Early Data Release, Kleinman et al. (2004) for Data Release 1 (DR1), and most recently, Eisenstein et al. (2006) for the DR4, from which the majority of our data sample derives, though a handful of stars from DR1 omitted by Eisenstein have been re-included from Kleinman et al. (2004). Most recently, Kepler et al. (2007) have refit the DA and DB stars from Eisenstein et al. (2006) with an expanded grid of models. A complete analysis of the methods by which candidate objects are chosen, spectra fitted, and quantities of interest are calculated can be found in Kleinman et al. (2004), Eisenstein et al. (2006), and Kepler et al. (2007). We put forth a brief outline here, with special attention paid to those aspects important to our own analysis.

Objects in the SDSS spectroscopic database were put through several cuts in color designed to separate the WDs from the main stellar locus. Figure 1 in Eisenstein et al. (2006) shows the location of these cuts. The chief failing

of their particular choices of cuts, as noted by the authors, is that WDs with temperatures below $\sim 8,000\text{K}$ begin to overlap in color-color space with the far more numerous A and F stars. The SDSS spectroscopic pipeline calculates a redshift for each object by looking for prominent lines in the spectrum. Objects with redshifts higher than $z=0.003$ were eliminated, unless the object had a proper motion from USNO-A greater than $0.3''$ per year. Because the spectroscopic pipeline is fully automated, occasionally DC white dwarfs show weak noise features that can be misinterpreted as low-confidence redshifts. Other types of WD, particularly magnetic WDs, can fool the pipeline as well. In the present work we are concerned chiefly with DA white dwarfs, so this incompleteness is of importance only insofar as we use the entire set of white dwarf spectral types to derive our completeness correction, as outlined in Section 4.4. We explore the implications of this more fully in that section.

Eisenstein et al. (2006) then used a χ^2 minimization technique to fit the spectra and photometry of the candidate objects with separate model atmospheres of pure hydrogen and pure helium (Finley et al., 1997; Koester et al., 2001) to determine the dominant element, effective temperature, surface gravity, and associated errors. As their Figure 2 demonstrates, they recovered a remarkably complete and uncontaminated sample of the candidate stars. They believe that they recovered nearly all of the DA white dwarfs hotter than $10,000\text{K}$ with SDSS spectra.

These stars form the core of our data sample. Their final table lists data on 10,088 white dwarfs. Of these, 7,755 are classified as single, non-

magnetic DAs. Kepler et al. (2007) re-fit the spectra for these stars using the same autofit method and Koester model atmospheres, but with a denser grid which also included models up to $\log g$ of 10.0. Where they differ from Eisenstein’s, we use these newer fits in our analysis. Of these 7,755 entries, ~ 600 are actually duplicate spectra of the same star. For our analysis we take an average of the values derived from each individual spectra weighted by the quoted errors. Our final sample contains 7,128 single, non-magnetic DA white dwarfs.

4.2.1 Issues With Model Spectra

As noted by Kleinman et al. (2004) and others, the surface gravities determined from Sloan spectra show a suspicious upturn below temperatures of about 12,000K which increases at cooler temperatures, as shown in Figure 4.1.

A number of separate pieces of evidence argue that this upturn in $\log g$ —and thus mass—is an artifact of the models and not a real effect. Not least among these is that no one has yet provided any satisfactory mechanism by which WDs could gain enough mass or shrink enough in radius as they cool to account for the magnitude of the effect. We do expect a slight increase in mass at cooler temperatures because in a galaxy of finite age, the cooler white dwarfs must come from higher mass progenitors. This is the reason for the upward slope of the blue dashed line in Figure 4.1. However, this effect is far smaller and more gradual than the upturn observed in the actual data.




Figure 4.1: $\log g$ v. $\log T_{\text{eff}}$ for the white dwarfs in our sample. At temperatures below $\sim 12,500\text{K}$, the $\log g$ values begin to rise to an extent unexplained by current theory. The solid line is a function empirically fit to the real data. The dashed line is the modest rise predicted by theory. The excess at a given T_{eff} is subtracted from the measured $\log g$ value for some of our luminosity functions.

Furthermore, Engelbrecht & Koester (2007), and Kepler et al. (2007) demonstrated that the masses derived solely from the colors do not show an increase in mass for cooler stars, which indicates that the problem is not physical, but a result of either the line fitting procedure or the line profiles themselves.

Figure 4.2 further illustrates the above point. The upper panel shows




Figure 4.2: A comparison of the theoretical colors of the SDSS WDs, derived from the atmospheric fits (black triangles), with the observed colors, as measured by the SDSS photometry (open blue circles). In the upper panel, the colors of the model atmospheres do not agree with the observed colors at low temperatures, indicating a problem with the line fitting for stars cooler than $\sim 12,500\text{K}$. In the lower panel, where the excess $\log g$ has been removed, the colors agree much better.

the colors derived from the synthetic spectra at the values of T_{eff} and $\log g$ quoted by Kepler et al. (2007) (i.e., the values in Figure 4.1), overlaid on the actual SDSS photometry for the same objects. Contrast this with the lower panel, which instead shows the colors derived from the synthetic spectra when the excess $\log g$ has been removed (in a manner described below; the resulting values are shown in Figure 4.3). The colors in the latter figure agree much better with the measured colors of the objects.

Furthermore, Kepler et al. (2007) found a similar increase in mean mass for the SDSS DB white dwarfs below $T_{\text{eff}} \sim 16,000\text{K}$. They conclude that since a) the problem only shows up in the line profiles and not the continuum, and b) the onset of the effect in both hydrogen (DA) and helium (DB) atmosphere WDs occurs at just the effective temperature where the neutral species of the atmospheric element begins to dominate, then the problem lies in the treatment of line broadening by neutral particles. This is supported further by the fact that as the species continues to become more neutral (i.e., as the temperature drops), the problem grows worse.

However, more recent model calculations indicate that neutral broadening is not important in the DA white dwarfs at temperatures down to at least 8,500K. Koester et al. (2009) consider a number of other possible mechanisms, drawing no hard conclusions but favoring the hypothesis that the observed effect is caused by a flawed or incomplete treatment of convection, leading to errors in the temperature structure of the outer layers of the WD models. They specifically reject the conclusions of Bergeron et al. (2007) (see also

Bergeron et al., 1990, 1995a), that the convective mixing of helium from a lower layer in the atmosphere leads to a higher pressure, mimicking the effect of a higher surface gravity, as this would require a hydrogen layer much thinner than any seismologically determined in a DA so far (Bradley, 1998, 2001, 2006; Castanheira & Kepler, 2008).




Figure 4.3: $\log g$ v. $\log T_{\text{eff}}$ with the upturn in $\log g$ removed.

The lack of resolution to the problem of the model atmospheres is one of the chief stumbling blocks to our analysis. Until the problem is resolved, the best we can do is to empirically remove the $\log g$ upturn. For a given T_{eff} , we subtract the excess in the measured mean value (as fit by the red solid lines

in Figure 4.1) over the theoretically expected mean (blue dashed line). Figure 4.3 shows the resulting values used. In fitting out the upturn this way, we make two implicit assumptions. First, that the excess $\log g$ is a function only of T_{eff} ; for most of the candidate solutions to the problem, we would expect only a weak dependence on $\log g$. Second, we assume that the problem affects only the $\log g$ determination and not T_{eff} . This latter assumption is unlikely to be true, as the two parameters are correlated. In Section 4.5 we explore more fully the impact of this fitting procedure on the luminosity and mass functions.

4.3 Constructing The Luminosity And Mass Functions

Since we are dealing with a magnitude-limited sample, the most luminous stars in our sample can be seen to much further distances than the intrinsically fainter stars. We thus expect more of them, proportionally, than we would in a purely volume-limited sample, and must make a correction for the different observing volumes. As shown by Wood & Oswalt (1998) and Geijo et al. (2006), the $1/V_{\text{max}}$ method of Schmidt (1968) (described more fully in, e.g. Green, 1980; Fleming et al., 1986) provides an unbiased and reliable characterization of the WDLF.

In the $1/V_{\text{max}}$ method, each star’s contribution to the total space density is weighted in inverse proportion to the total volume over which it would still be included in the magnitude limited sample. Since the stars are not spherically distributed, but lie preferentially in the plane of the Galaxy, an additional

correction for the scale height of the Galactic disk must be included. For the purposes of comparison with previous work, we adopt a scale height of 250 pc.

To determine the absolute magnitude of each WD, we use the effective temperatures and $\log g$ values provided by Kepler et al. (2007)—as corrected in Section 4.2—and fit each WD with an evolutionary model to determine the mass and radius. For $7.0 < \log g < 9.0$, we use the mixed C/O models of Wood (1995) and Fontaine et al. (2001), as calculated by Bergeron et al. (1995b). For $9.0 < \log g < 10.0$, we use the models of Althaus et al. (2005) with O/Ne cores, including additional sequences for masses larger than $1.3 M_{\odot}$ calculated specifically for Kepler et al. (2007). Once we know the radius, we can calculate the absolute magnitude in each Sloan band by convolving the synthetic WD atmospheres of Koester (Finley et al., 1997; Koester et al., 2001) with the Sloan filter curves. We apply bolometric corrections from Bergeron et al. (1995b) to determine the bolometric magnitude. For the handful of stars (~ 80 -100) with $\log g$ values outside the range covered by Bergeron’s tables, we use a simple linear extrapolation.

We then determine photometric distances to each star from the observed SDSS g magnitude. SDSS, being concerned mostly with extragalactic objects, reports the total interstellar absorption along each line of sight from the reddening maps of Schlegel et al. (1998). Since the objects in our sample lie within the Galaxy, and most of them within a few hundred parsecs, they are affected by only a fraction of this reddening. Following Harris et al. (2006), we therefore assume: 1) that objects within 100 pc are not affected

by reddening, 2) objects with Galactic height $|z| > 250$ pc are reddened by the full amount, and 3) that the reddening varies linearly between these two values. The distances and reddening are then fit iteratively from the observed and calculated absolute g magnitude. In practice, the reddening correction makes very little difference to the final LF (typical A_g values range from 0.01 to 0.05).

We calculate error bars on the luminosity function using a Monte-Carlo simulation, drawing random deviates in T_{eff} , $\log g$, and each band of photometry from Gaussian distributions centered around the measured value. The standard deviations in T_{eff} and $\log g$ we use for this scattering are 1.2 times the formal errors quoted in Eisenstein et al. (2006) (their own analysis, based on repeated autofit measurements on duplicate spectra of the same stars, suggests that the formal errors derived by their method are $\sim 20\%$ too small). The photometry errors come directly from the SDSS database. After scattering the parameters in this way, we recalculate the LF. We then add in quadrature the standard deviation of each LF bin after 200 iterations and the counting error for each bin (the errors for each individual star—taken to be of the order of the star’s $1/V_{\text{max}}$ statistical weight—summed in quadrature).

At a S/N of 16—the mean for the stars in our sample brighter than $g = 19.5$ —formal errors in T_{eff} and $\log g$ are of order 1.5%. When propagated through our code, the mean errors in M_{bol} and mass are 0.35 dex and 9% ($0.05 M_{\odot}$) respectively. For the stars brighter than $g = 19.0$ used to compile our mass functions the average S/N is 19.5, leading to errors in M_{bol} and mass of

0.35 dex and 7% ($0.04 M_{\odot}$).

4.4 Completeness Corrections

Aside from the problems with the model spectra, the chief difficulty we have encountered in deriving our luminosity functions is unraveling the complicated way in which SDSS objects are assigned spectral fibers. SDSS is foremost a survey of extragalactic objects and rarely targets white dwarfs for follow-up spectroscopy explicitly. Most of the objects in our sample are targeted by some other algorithm. In particular, there is considerable overlap in color between the white dwarfs and many QSOs.

A completeness correction could, in theory, be built from “first principles.” For each object in the SDSS spectroscopic database, we know by which algorithm(s) it was targeted (or rejected) for spectroscopy, and by which algorithm it was ultimately assigned a fiber. And for each algorithm, we know which objects were targeted, which were ultimately assigned a fiber, and which, of the targeted objects, turned out to be WDs. However, the selection process is a multi-variate function of 5 apparent magnitudes, and colors in spaces of as many as 4 dimensions (which vary based on the algorithm), as well as the complex tiling algorithm. Furthermore, these algorithms are also time-dependent. The QSO-targeting algorithm, for instance, “improved” as SDSS progressed, rejecting more non-QSO objects (e.g., white dwarfs) in the later stages of the survey.

We believe unraveling this complicated process to be unnecessary. In-

stead we have chosen to compare our sample with the stars used to derive the WDLF of Harris et al. (2006). Given certain assumptions about completeness and contamination in both data sets, we derive a completeness correction as a function of a single color index ($g - i$) and g magnitude.

The Harris et al. (2006) sample comes from photometric data in the SDSS Data Release 3. They selected objects by using the reduced proper motion diagram to separate WDs from more luminous subdwarfs of the same color. Briefly, they used color and proper motion (from USNO-B Munn et al., 2004) to determine WD candidates from SDSS imaging data. They then fit candidates with WD model atmosphere colors to determine temperatures and absolute magnitudes, from which they derived photometric distances and—together with proper motion—tangential velocities. In order to minimize contamination, they adopted a tangential velocity cutoff of 30 km/s and rejected all stars below this limit. The remaining 6,000 objects are, with a high and well-defined degree of certainty ($\sim 98 - 99\%$), likely to be white dwarfs.

If the database of SDSS spectra were complete, all of these objects would (eventually) have spectra, and all but the 1-2% of contaminating objects would be confirmed to be WDs. Furthermore, all of the WDs that did *not* make it into the Harris et al. sample—because they were either missing from the Munn et al. (2004) proper motion catalog, or had a tangential velocity below 30 km/s—would also all have spectra. In such a perfect world, of course, no completeness correction would be necessary. However, since SDSS does not obtain a spectrum of every object in its photometric database, a significant

percentage of the objects in Harris et al. will not have spectra, or else will be dropped at some later point by Eisenstein et al. and thus not make it into our spectroscopic sample. Our goal, then, is to look at all of the WDs in the Harris et al. sample that potentially *could* have made it into our sample, and determine which ones in fact did. If we assume that the WDs *not* in Harris et al. follow the same distribution (an assumption we discuss more fully below), then we can take this as a measure of the overall detection probability and invert it to get a completeness correction.

The imaging area of the DR3, from which Harris et al. derive their sample, is not the same as the spectroscopic area in the DR4. Therefore, for the purposes of this comparison, we removed all stars not found in the area of sky common to the two data sets from their respective samples. This left 5,340 objects classified as white dwarfs by Harris et al. that could potentially have been recovered by Eisenstein et al. Of these, 2,572 were assigned spectral fibers in DR4, and 2,346 were ultimately confirmed by Eisenstein et al. to be white dwarfs.

Since we wish to restrict our analysis to single (i.e., non-binary) DA white dwarfs, we removed all stars classified as DA+M stars in either catalog. Unfortunately, given that the Harris catalog contains no further information as to the type of WD, we were unable to remove the non DA stars and simply compare what remains with the Eisenstein sample. Instead, we compute the completeness for all of the WDs, under the assumption—explored more fully below—that DAs, as the largest component of the WD population, dominate




Figure 4.4: Color-color plot of the white dwarfs in the two samples used to derive our completeness correction. Open symbols are WDs from the Harris et al. (2006) sample that a) were in the area of sky covered by spectroscopy in DR4, b) had $V_{tan} \geq 30$ km/s, and c) were not determined by i - and z -band excess to be WD + main-sequence binaries. The filled circles are the stars for which SDSS obtained spectra and Eisenstein et al. (2006) confirmed to be WDs. The dashed box shows a two-dimensional projection of the QSO targeting algorithm's exclusion region. The open gray squares are the WDs from Harris et al. that lie outside Eisenstein et al.'s color-color cuts. For clarity, only half of the points have been plotted.

the selection function.

Figure 4.4 shows a comparison of the two samples. The open symbols are the complete Harris et al. sample (excluding those, as mentioned above,




Figure 4.5: A map of our completeness correction. Darker areas indicate more complete regions of the figure, with black being 100% complete. The overall completeness is of order $\sim 50\%$.

with $V_{tan} < 30$ km/s, those not in the region of sky covered by spectroscopy, and the DA+M stars). The gray squares lie outside the cuts in color-color space imposed by Eisenstein et al. They may have spectra in SDSS, but they were not fit by Eisenstein et al., and therefore will not have made it into our sample. The filled circles are the stars that *are* in Eisenstein et al. In other words, if the SDSS spectral coverage of WDs were complete, and Eisenstein et al. recovered every WD spectra in SDSS, then all of the open circles would

be filled. The inside of the dashed box is the exclusion region for SDSS's QSO targeting algorithm (Richards et al., 2002), specifically implemented to eliminate WDs from their sample. Note that our sample is more complete for the stars outside this region.

Figure 4.5 shows the discovery probability as a function of $g - i$ color and g magnitude. Darker areas mean a higher probability of discovery, with black indicating that all the WDs in the Harris et al. sample in that area of color-magnitude space made it into our sample. We have performed a box smoothing to eliminate small scale fluctuations.

There is a drop off in discovery probability for stars bluer than $g - i \sim -0.2$ at all apparent magnitudes. This corresponds to the red edge of the exclusion region of the QSO targeting algorithm, as noted above. The QSO algorithm is also itself a function of apparent magnitude, which accounts for the general decrease at fainter magnitudes in the red half of the diagram, and the much steeper drop off between $g \simeq 19$ and $g \simeq 19.5$. The bluer stars ($g - i \lesssim -0.2$), most of which are targeted by the HOT_STANDARD or SERENDIPITY_BLUE algorithms, show the opposite: a slight increase at fainter magnitudes.

To give a better sense of the order of magnitude of our completeness, Figure 4.6 shows a histogram of the values in Figure 4.5. For most of the cells that end up in the bins for 0, 1, and 0.5, the Harris et al. sample contains only one or two stars. The mean completeness for the whole sample is $\sim 51\%$.




Figure 4.6: A histogram of the completeness values in Figure 4.5. Most of the 0, 1, and 0.5 values come from color-magnitude regions in which there are only one or two stars in the Harris et al. sample available for comparison.

To derive our final completeness correction, we must further consider the incompleteness and contamination in the Harris et al. sample itself. Assuming that the SDSS photometric database is essentially complete down to $g = 19.5$, then the incompleteness in Harris et al. comes mainly from two sources: 1) the incompleteness in the Munn et al. (2004) proper motion catalog, and 2) the tangential velocity limit of 30 km/s imposed, which results in some low tangential velocity WDs being dropped from the sample. How-

ever, with one negligible exception, none of the criteria used to target objects for spectroscopy in SDSS, nor those used by Eisenstein et al. to select white dwarf candidates, depends explicitly on proper motion or tangential velocity. Thus we assume that the low-velocity stars—dropped from the Harris et al. sample—will be recovered by Eisenstein with the same probability as the high-velocity stars—i.e., the stars in Figure 4.4.

Contamination poses a bit more challenging problem. At first glance, it would seem that the reverse of the above process could be applied, whereby those objects in Harris et al. which did get spectral fibers—but were ultimately rejected as WDs by Eisenstein et al.—could be removed from the sample, and those that did *not* get spectra could be assumed to follow the same distribution. This latter assumption, however, is unlikely to be true. SDSS gives very low priority to targeting white dwarfs specifically, and we would thus expect a larger fraction of the objects that get spectral fibers to turn out to be contaminating objects (in particular QSOs, of which we found 13 in the Harris et al. sample) than if the fibers were assigned purely randomly. Furthermore, many of the 225 objects which have spectra in DR4 but are not included in the Eisenstein catalog may actually be white dwarfs which Eisenstein’s algorithms dropped for some other reason, e.g. they lie outside the color and magnitude ranges used for initial candidate selection, or there is a problem (low S/N, bad pixels) with the spectrum. Approximately 100 appear to be DC white dwarfs to which the SDSS spectroscopic pipeline assigned erroneous redshifts on the basis of weak noise features. Ultimately, we have chosen to




Figure 4.7: Luminosity functions for three different limiting magnitudes. We take the good agreement between the curves to indicate that our completeness correction (and the $1/V_{\text{max}}$ correction) are working properly.

adopt the contamination fraction of Harris et al. (2%) for the whole sample, and have reduced our final completeness correction accordingly. This choice has a negligible effect on the small scale structure of the WDLF in which we are interested.

Finally, we note that the Harris et al. sample has an apparent magnitude limit of $g = 19.5$, whereas the spectroscopic sample contains objects down to $g \simeq 20.5$. Given that the SDSS targeting algorithms are themselves




Figure 4.8: Mass functions for three different limiting magnitudes. Because of the essentially fixed integration time for SDSS spectra, objects with fainter apparent magnitudes generally have lower signal to noise, which translates directly into larger uncertainties in the derived parameters (T_{eff} , $\log g$, and mass). Hence, as we include stars with fainter apparent magnitudes, more stars scatter out of the peak, broadening the mass function.

functions of apparent magnitude, our completeness correction is as well. An extrapolation of our discovery probability is problematic in this area, though, because this is just the apparent magnitude where the QSO targeting algorithm drops off rapidly. We have decided to impose a magnitude cutoff of $g = 19.5$ in our sample. This reduces our sample by nearly a half, with a corresponding increase in counting error. However, because SDSS spectra have a small range

of exposure times (45-60min), fainter apparent magnitude translates directly into lower S/N and larger errors in derived parameters.

Figure 4.7 shows the luminosity functions we derive for different choices of limiting magnitude. We take the generally good agreement between the curves to indicate that our completeness correction is doing its job correctly in the g magnitude direction.

Figure 4.8 similarly shows the mass functions we derive for different choices of limiting magnitude. In the case of the mass function, the S/N of the spectra becomes a much bigger factor. As a consequence of the essentially constant exposure times of SDSS spectra, the parameters (T_{eff} and $\log g$) determined from the spectra of fainter objects have larger errors, which causes a larger error in mass. Thus, the MF is broadened when stars with $g > 19.0$ are included. For this reason, Kepler et al. (2007) limited their mass functions to stars with $g \leq 19.0$, and we follow their lead for the remaining MFs in the current paper.

4.5 Luminosity Functions And Discussion

4.5.1 White Dwarf Mass Functions

Figure 4.9 shows the WD mass function we derive for all stars with $T_{\text{eff}} > 12,000K$ and $g < 19.0$. The red dashed line is the MF corrected only by $1/V_{\text{max}}$ —i.e., before we apply our completeness correction. It generally shows good agreement with the MF derived in Kepler et al. (2007) (blue points), not surprising considering we use nearly the same data set and very similar WD

models. The small differences are due to our use of slightly different sets of data and models, as well as differing treatment of duplicate spectra, and can largely be considered statistical fluctuations. We refer the interested reader to their paper for a more in-depth analysis of the WDMF.




Figure 4.9: The white dwarf mass function for all WDs with $T_{\text{eff}} > 12,000\text{K}$ and $g < 19.0$. The dashed line in the upper panel is the MF corrected only for $1/V_{\text{max}}$, without our completeness correction applied. It agrees very well with Kepler et al. (2007—dots). The solid line is with our completeness correction applied, and represents the true local space density of white dwarfs. The bottom panel shows the ratio of our two mass functions—i.e., the cumulative completeness correction for each bin. The small variation indicates that the completeness correction, while changing the overall normalization by roughly a factor of 2.2, has little effect on the shape of the MF.

The solid black line in the upper panel shows our MF after correcting for the completeness of the spectroscopic sample. This curve represents the true local space density of WDs per cubic parsec per M_{\odot} interval. The bottom panel shows the total weight of each bin above the uncorrected MF—essentially the final completeness correction for each bin. There is little small scale variation from bin to bin, and our completeness correction mainly has the effect of raising the normalization of the whole MF by a factor of ~ 2.2 . In other words, the shape of the MF is not strongly affected by the completeness correction.




Figure 4.10: White dwarf mass functions for WDs with $T_{\text{eff}} > 8,000\text{K}$ and $g < 19.0$ both with and without the upturn in $\log g$ for cooler stars removed. The solid line is the MF from Figure 4.9 renormalized for comparison purposes.

Figure 4.10 is the WDMF for all stars down to 8,000K. The dashed red line is for the data as reported by Kepler et al. (2007), the dotted blue line is after our correction for the upturn in $\log g$. The solid black line is the WDMF for only those stars above 12,000K (i.e., the same as Figure 4.9) renormalized to the same scale for comparison purposes. There are more high mass stars in general, and one spuriously large bin, but on the whole, our $\log g$ correction recovers a reasonable mass distribution for stars cooler than 12,000K.




Figure 4.11: LFs derived in this paper. Removing the $\log g$ upturn makes each affected star less massive, and therefore larger and brighter, pushing it to a more leftward M_{bol} bin. The results of Harris et al. (2006) and Liebert et al. (2005) are shown for comparison.

4.5.2 White Dwarf Luminosity Functions

Figure 4.11 shows the luminosity function we derive for all of the DA stars in our sample down to 7,000K for all stars with $g < 19.5$. In red is the LF for the data as reported; in black is the LF for the data with the increase in $\log g$ at low temperature removed. The process of removing the excess $\log g$ pushes stars to lower masses, making them larger and therefore brighter for the same T_{eff} . In the range plotted, the black curve contains a total of 3,358 WDs, while the red contains 2,940.

The lack of agreement between our best LF (black) and the Harris et al. (2006) luminosity function (blue) can be attributed, at least in part, to the differing assumptions used in creating the two LFs. Harris et al. derived their temperatures by fitting Bergeron models to the photometry assuming a $\log g$ of 8.0 for every star, a poor assumption for more than 30% of WDs (Liebert et al., 2005; Kepler et al., 2007). The temperatures they derive are systematically different from the spectroscopic temperatures; Figure 4.12 shows the fractional difference between the spectroscopically and photometrically derived effective temperatures. When we use the photometrically derived temperatures and set $\log g = 8.0$, we recover the Harris et al. LF fairly well.

It should also be noted that the Harris et al. luminosity function is for WDs of *all* types, whereas ours is comprised only of the DAs. For each bin in the Harris et al. LF, we have used the full Eisenstein et al. (2006) catalog to determine a rough DA fraction, and reduced the LF of Harris et al. accordingly. This DA fraction—shown in Table 4.5.2—is in generally good agreement with




Figure 4.12: A comparison of the spectroscopically and photometrically derived temperatures for the WDs common to the Harris et al. (2006) and Eisenstein et al. (2006) samples.

previous works (Fleming et al., 1986), but we have made no attempt to address selection biases in the Eisenstein et al. catalog.

One other source of the discrepancy between our results and Harris et al. is due to our assumption that whatever causes the observed upturn in $\log g$ in the cooler stars affects only the $\log g$ determination and does not alter the spectroscopically derived T_{eff} . As the effects of the two parameters on the line profiles are interdependent, this assumption is probably not valid.

Table 4.1: The fraction of stars in Eisenstein et al. (2006) listed as DA or DA_auto. Though they generally agree with previous results, they should be used with much caution, as they were calculated crudely and we have taken no care to correct for biases in the sample. We have employed them here simply to compare our DA-only luminosity function to previous work.

M_{bol}	DA Fraction
7.25	0.9338
7.75	0.9243
8.25	0.9246
8.75	0.8980
9.25	0.8433
9.75	0.8146
10.25	0.7958
10.75	0.8158
11.25	0.7957
11.75	0.7721
12.25	0.7985
12.75	0.7976
13.25	0.8173
13.75	0.8009

The curves in Figure 4.11 suggest that in addition to the excess $\log g$, the temperatures determined by line fitting for the cooler stars are probably too high. Ultimately, this area of the spectroscopic WDLF will remain uncertain until the problems with the model atmospheres have been resolved.

The LF of Liebert et al. (2005) shown in green in figure 4.11 was compiled from a small dataset (348 DA white dwarfs) based on a survey done on photographic plates over 20 years ago on a 0.5m telescope. In addition to low number statistics, the dataset suffers from a very difficult-to-quantify incompleteness at the faint end, which is probably responsible for the lack of agreement below $M_{bol} \sim 9.5$.

4.6 Conclusions

Our initial goal with this phase of the project was to take advantage of the tremendous number of WDs spectroscopically observed by SDSS and studied by Eisenstein et al. (2006) and others to create separate WD luminosity functions for two or more different ranges of mass. Should we one day be successful in this endeavor, we will effectively add a third dimension, currently unexplored, to observational WD luminosity functions.

In order to carry out this analysis, we needed to more fully understand the manner in which white dwarfs were selected to receive spectra in SDSS. By comparing the proper-motion selected sample of Harris et al. (2006) with the spectroscopically determined WDs of Kleinman et al. (2004) and Eisenstein et al. (2006), we derived a WD selection probability over a range of parameters

that included nearly the entire useful range of $g - i$ color ($-1.0 < g - i < 0.2$) and apparent g magnitude ($15 < g < 19.5$).

We also presented additional arguments that the observed upturn in $\log g$ is an artifact of the model atmosphere line-fitting procedure, or—more likely—a problem with the line profiles themselves. Since it may be some time before this problem is fully understood and addressed, we implemented a procedure to remove the excess $\log g$ empirically and showed that the mass function recovered for the stars cooler than 12,000K reasonably agrees with the MF for the hotter stars, which in turn agrees well with previous work.

Finally, we presented the first WDLF for spectroscopically determined WDs in the Fourth Data Release of the SDSS. In addition to addressing the issues of completeness and the observed $\log g$ upturn in a more systematic manner than previously attempted, our LF contained the largest sample of spectroscopically determined WDs to date (3,358), more than six times the 531 presented in Hu et al. (2007), and more than an order of magnitude more than the 298 stars included in the LF of Liebert et al. (2005).

Ultimately, however, we have determined that the quality of the data is insufficient to attain our goal of separate high- and low-mass luminosity functions. In particular, the problems with the model spectra render any discussion of masses for WDs with $T_{\text{eff}} < 12,000\text{K}$ highly suspect. The WDs that remain cover an insufficient temperature range to adequately study the features in the WDLF that would allow us to explore new physics or confirm features long-predicted on theoretical grounds.

Chapter 5

Current and Future Applications of the Bayesian Method

As of this writing, our group has several projects in various states of completion that utilize the methods we have developed in this dissertation. Some of these applications are the natural extension of our work here—using the method to find white dwarf (and main sequence) ages for open clusters. But other applications of the method are meant to extend its capabilities into new frontiers.

5.1 Analysis of Open (and Globular) Clusters

For her doctoral dissertation at University of Texas in 2009, Elizabeth Jeffery used our method to analyze data from four new open clusters (NGC 2360, NGC 2660, NGC 3960, NGC 188) and analyzed new data for two that had previous age determinations (NGC 2420, NGC 2477). Using ground-based photometry for the main-sequence, she determined MSTO ages for these clusters in the absence of WDs. Then, using a combination of the ground-based data and HST data, she re-analyzed the same clusters using the white dwarfs and main sequence *without* the turn-off.




Figure 5.1: An updated version of Figure 3.10 including new and reanalyzed clusters from Jeffery (2009). The solid points show new results. The line represents the one-to-one correspondence between the two ages. [Adapted from Jeffery (2009).]

Figure 5.1 is an updated version of Dr. Jeffery’s Figure 7.1, showing all of the clusters to date with both MSTO and WD ages. This version of the figure includes the new WD age for the Hyades derived in this work.

The MSTO and WD ages show considerable agreement across the entire range of ages considered. For the only major outlier, NGC 2660, the WD data did not extend to the terminus of the sequence, and the S/N of the brighter white dwarfs was too low to determine a reliable bright WD age. The value

plotted is the lower limit estimated by Dr. Jeffery based on her completeness on the faint end of the WD sequence. Dr. Jeffery also considers her determination for WD age of the cluster NGC 188 to be a lower limit as well, though in this case, the age she derives does come from the bright white dwarfs via our Bayesian method and may therefore be reliable.

These results give us confidence that the ages we obtain using our Bayesian method show very little systematic error, at least in this age range. One of the fundamental goals of Dr. Jeffery’s ongoing project is to extend the reliability of WD ages (and in particular, the bright WD technique) to older ages. The WD sequences of several globular clusters have already been imaged with HST (Hansen et al., 2004, 2007), and these rich datasets are ripe for analysis with our methods. The successful extension of the bright WD technique to older clusters will allow us to find ages for more distant clusters, even in the absence of photometry data down to the terminus of the WD sequence.

5.2 Initial-Final Mass Relation

Considerable work has gone into refining the initial-final mass relation (IFMR) in recent years (see, e.g., Kalirai et al., 2008; Williams et al., 2009). The method most often employed to determine the IFMR involves in some sense the reverse of what we do in our Bayesian calculation. First, main-sequence turn-off fitting gives a cluster age. Spectroscopy of the WDs in the cluster give T_{eff} and $\log g$. WD cooling models turn those parameters

into a cooling age for the white dwarf. Subtracting the cooling age from the cluster age gives a precursor lifetime, from which one can determine (via MS evolutionary models) a precursor mass. Repeat as needed.




Figure 5.2: The empirical initial-final mass relation. Different colors represents points from different clusters. [Taken from Ferrario et al. (2005)]

Care must be taken, therefore, to avoid circularity in inverting these same model sets to then determine the WD age of the cluster. Partly this circularity is alleviated by the use of many (and ideally *different*) clusters to determine the IFMR. A single cluster, at any rate, rarely has white dwarfs that

span enough of a mass range to meaningfully comment on the entire IFMR. Only by analyzing many clusters do we obtain a broadly applicable IFMR.

Our Bayesian MCMC method assumes an initial-final mass relation in order to derive predicted photometry to compare with the observed photometry. To date, we have exclusively used the IFMR of Weidemann (2000). However, there is no reason, in principle, that we cannot parameterize the IFMR and fit those parameters *simultaneously* with the cluster age, distance, etc., and the ZAMS masses of the WDs.

Fitting both parameters simultaneously has several benefits. First, we avoid the potential circularity in using the ages and masses of the WDs to derive the IFMR, then using the IFMR to derive the ages and masses of the WDs. Second, we obtain the parameters of the IFMR marginalized over the possible ages of the cluster, and vice versa. As the IFMR is currently one of our largest sources systematic uncertainty in our WD age determinations, the errors on those determinations (as well as on all of the other parameters of the model, including WD masses) will more accurately reflect the true uncertainty in the parameters. Furthermore, as Williams et al. (2009) showed, systematic errors in the turn-off age of the cluster can have a significant effect on the parameters of the IFMR. By marginalizing over these uncertainties, we get a better picture of the true probability distributions of *all* the parameters of our model.

In general, we believe we can accomplish this analysis with a simple two parameter model of the IFMR. Though there is some evidence of non-linearity

in the IFMR, particularly on the low-mass end, a quick glance at Figure 5.2 shows that the WDs in any individual cluster span a small enough mass range for the IFMR to be considered linear. By fitting multiple clusters of different ages in this manner, we can eventually build a more detailed and accurate picture of the true initial-final mass relation for the open clusters, and derive age fits for the clusters that are less model-dependent.

5.3 Analysis of Single Field White Dwarfs

Currently, the only reliable method for determining WD ages for non-cluster Galactic populations (i.e., the thin and thick disks and halo) is via the white dwarf luminosity function. Members of our collaboration are currently working to increase the number of stars in the coolest bins of the disk luminosity function to provide a better age, but the method is extremely observation intensive and time consuming. Luminosity functions for the less well-populated thick disk and halo will need at least one order of magnitude (and perhaps several) more stars in the observational sample to determine a reliable age.

Dating these populations accurately will have a profound effect not only on the study of WDs, but also on our understanding of the age and evolution of the Galaxy. If we could apply our code reliably to single field white dwarfs, it would open an entirely new avenue for determining the ages of the Galactic thin disk, thick disk, and halo.

The traditional stumbling block to this type of analysis is the multi-

way degeneracy between age, distance, and WD mass. Because our method provides complete marginalized probability distributions for these parameters, we can begin to analyze single WDs and explore the true effects of this degeneracy in a more sophisticated manner than simply assuming a distance and then determining a distance-dependent age. At the same time, the Bayesian nature of our method allows us to include prior information on an object. If we can pin down the distance to an individual WD via parallax (or some other method), we can place a prior on the distance that reflects our confidence in that distance determination, rather than simply assuming a single value and attempting to propagate the errors through in a way that assumes our distributions are Gaussian when they clearly are not.

One member of our collaboration (Ted von Hippel) has already begun work in this area. Figures 5.3 through 5.5 show some of his early results and highlight some of the difficulties we have yet to overcome.

Figures 5.3 and 5.4 are runs for simulated WDs of different masses at an age of ~ 8 Gyr. The blue points are the individual draws from the posterior distributions in MCMC with standard errors on the prior on distance of ± 10 pc (10%), consistent with what we might be able to obtain via parallax on a nearby WD. The derived parameters are consistent with the “true” values (red circles) used to create the simulated star.

The complicated, multi-dimensional nature of the age-mass-distance correlation is apparent in the shapes of the 2-D projections presented here. The two branches in the age-mass correlation in the lower left panel of Fig-




Figure 5.3: Mass-age-distance-metallicity correlations for a single (simulated) white dwarf of ~ 8 Gyr in the field. The blue points are the individual samples from the Markov chain with a Gaussian prior on the distance with a standard error of ± 10 pc (10%). All of the parameters are consistent with the “true” values (red circles) used to create this simulated star. [Taken from von Hippel (2009—private communication).]

ure 5.3 correspond to the same two branches seen in the upper right panel. The same is true of the banana-shaped correlations seen in the corresponding panels in Figure 5.4. Here, having access to the entire posterior probability distribution is essential to understanding the complex interaction between these three parameters.




Figure 5.4: Same as Figure 5.3 for a simulated star of lower mass. [Taken from von Hippel (2009—private communication).]

Figure 5.5 shows the results for the real white dwarf WD0243+026. Again, the bottom two panels show the bifurcated correlations among mass, age, and distance. Here, the green points are for a run which included only *BVI* photometry in the analysis. The red points are for a run which included *BVRIJHK*.

For the simulated runs of white dwarfs with ages ~ 8 Gyr, we currently achieve an accuracy of order 10-20%. Our goal is to move that number closer to the 5-10% range.




Figure 5.5: Similar to Figure 5.3 for the real white dwarf WD0243+026. Again, the bottom two panels show the bifurcated correlations among mass, age, and distance. Here, the green points are for a run which included only *BVI* photometry in the analysis. The red points are for a run with included *BVR_IJHK*. [Taken from von Hippel (2009—private communication).]

5.4 MCMC Analysis of Binary Stars

5.4.1 Binaries in the Field

One possible way to achieve errors in the 5-10% range is to look at WDs that are part of wide binary systems in the field. In essence, a non-interacting binary system is an open cluster containing exactly two stars. The main-sequence companion can, through spectroscopy, help to pin down the metallicity of the system. It can also, via precise photometry, constrain the distance to the system. Most important, though, is the possibility of using radial velocity and/or astrometric data to determine dynamical masses for both stars. As Figures 5.3 through 5.5 show, because of the correlation between mass and age, constraining the mass of the WD will yield a more precise age.

5.4.2 Binaries in Open Clusters

Figures 3.12 and 3.13 in Chapter 3 show the results of individual mass determinations for unresolved binaries in the Hyades. As we pointed out in that chapter, there appears to be a systematic offset between our mass determinations and those derived by other methods. This is likely to indicate a problem somewhere in the main-sequence models, either in the translation from T_{eff} and L to broadband colors (i.e., the model atmospheres), or in the underlying stellar evolution models.

Ted von Hippel and the WIYN Open Cluster Study (WOCS— von Hippel & Sarajedini, 1998; Sarajedini et al., 1999) currently have plans to use our method to determine masses of individual binary systems in the open cluster

NGC 188 and other open clusters, and compare these results with dynamical masses obtained via radial velocities. The goals of the project are three-fold: first, a large volume of data on stars with both photometric and dynamical masses will help refine stellar evolution and/or atmosphere models considerably. Second, this data will provide a means of determining the distribution of binaries and mass ratios in open clusters, which has implications in such diverse studies as planet formation and dynamical simulations. Finally, there is a possibility that this work will uncover binaries in unusual and scientifically interesting states.

5.5 Helium Abundance of NGC 6791

NGC 6791, a rare well-populated, old (~ 8 Gyr, Chaboyer et al., 1999), and metal-rich ($[\text{Fe}/\text{H}] = +0.4$, Peterson & Green, 1998) open cluster, is a natural object of study to learn about chemical enrichment processes in the Galaxy. In particular, this cluster provides a valuable tool for determining the relationship between $[\text{Fe}/\text{H}]$ and helium abundance. Evidence suggests that this relationship is linear (Pagel & Portinari, 1998; Chaboyer et al., 1999), but the slope of the correlation is not well constrained. A better determination of this relationship will not only enable us to improve stellar evolution models, but also has a significant impact on several broader areas of astronomy. These include a more complete picture of the chemical enrichment of the Galaxy (and by extension, the Universe), more accurate pictures of star formation and death processes, and age determinations of galactic bulges and elliptical




Figure 5.6: Color-magnitude diagram showing the effect of helium abundance on the shape of the main-sequence turn-off region. The data were simulated using values for cluster parameters similar to NGC 6791 and scattered photometrically. The curves show the effect of helium abundance on the shape and position of the isochrones for $Y = 0.27$ to $Y = 0.39$ in intervals of $\Delta Y = 0.03$, with the red curve representing the isochrone used to simulate the data.

galaxies (e.g., Charlot et al., 1996).

Unfortunately, due to its advanced age, NGC 6791 no longer contains any stars hot enough to show helium lines. Thus, spectroscopy is unable to help us. Recently obtained JK photometry of NGC 6791 may provide a solution. Figure 5.6 shows the effect of changes in helium abundance on the shape and position of the main sequence and MSTO. Our Bayesian method is a natural fit for the analysis of this data, and with minimal extension of our models to include sampling on the helium abundance of the cluster (a feature that we have already added to the code, but have not extensively tested), we should be able to determine the helium abundance of the cluster from the photometry alone.

Bibliography

- Allègre, C. J., Manhès, G., & Göpel, C. 1995, *Geochim. Cosmochim. Acta*, 59, 1445
- Althaus, L. G., García-Berro, E., Isern, J., & Córscico, A. H. 2005, *A&A*, 441, 689
- Althaus, L. G., García-Berro, E., Isern, J., Córscico, A. H., & Rohrmann, R. D. 2007, *A&A*, 465, 249
- Bender, C. F. & Simon, M. 2008, *ApJ*, 689, 416
- Bergeron, P., Gianninas, A., & Boudreault, S. 2007, in *Astronomical Society of the Pacific Conference Series*, Vol. 372, 15th European Workshop on White Dwarfs, ed. R. Napiwotzki & M. R. Burleigh, 29–+
- Bergeron, P., Liebert, J., & Fulbright, M. S. 1995a, *ApJ*, 444, 810
- Bergeron, P., Wesemael, F., & Beauchamp, A. 1995b, *PASP*, 107, 1047
- Bergeron, P., Wesemael, F., Fontaine, G., & Liebert, J. 1990, *ApJ*, 351, L21
- Blanton, M. R., Lin, H., Lupton, R. H., Maley, F. M., Young, N., Zehavi, I., & Loveday, J. 2003, *AJ*, 125, 2276
- Bradley, P. A. 1998, *ApJS*, 116, 307

- . 2001, *ApJ*, 552, 326
- . 2006, *Memorie della Societa Astronomica Italiana*, 77, 437
- Casella, G. & George, E. I. 1992, *The American Statistician*, 46, 167
- Castanheira, B. G. & Kepler, S. O. 2008, *MNRAS*, 385, 430
- Catalán, S., Isern, J., García-Berro, E., & Ribas, I. 2008, *MNRAS*, 387, 1693
- Chaboyer, B., Green, E. M., & Liebert, J. 1999, *AJ*, 117, 1360
- Charlot, S., Worthey, G., & Bressan, A. 1996, *ApJ*, 457, 625
- Chib, S. & Greenberg, E. 1995, *The American Statistician*, 49, 327
- de Bruijne, J. H. J., Hoogerwerf, R., & de Zeeuw, P. T. 2001, *A&A*, 367, 111
- DeGennaro, S., von Hippel, T., Jefferys, W. H., Stein, N., van Dyk, D., & Jeffery, E. 2009, *ApJ*, 696, 12
- DeGennaro, S., von Hippel, T., Winget, D. E., Kepler, S. O., Nitta, A., Koester, D., & Althaus, L. 2008, *AJ*, 135, 1
- DeGennaro, S. A. 2007, Master's thesis, AA(THE UNIVERSITY OF TEXAS AT AUSTIN.)
- Dotter, A., Chaboyer, B., Jevremovic, D., Kostov, V., Baron, E., & Ferguson, J. W. 2008, *ApJS*, in press, arXiv:0804.4473
- Ducourant, C. et al. 2006, *A&A*, 448, 1235

- Eggen, O. J. & Greenstein, J. L. 1965, *ApJ*, 141, 83
- Eisenstein, D. J. et al. 2006, *ApJS*, 167, 40
- Engelbrecht, A. & Koester, D. 2007, in press, *Proceedings of the 15th European Workshop on White Dwarfs, Leicester 2006*
- Ferrario, L., Wickramasinghe, D., Liebert, J., & Williams, K. A. 2005, *MNRAS*, 361, 1131
- Finley, D. S., Koester, D., & Basri, G. 1997, *ApJ*, 488, 375
- Fleming, T. A., Liebert, J., & Green, R. F. 1986, *ApJ*, 308, 176
- Fontaine, G., Brassard, P., & Bergeron, P. 2001, *PASP*, 113, 409
- Geijo, E. M., Torres, S., Isern, J., & García-Berro, E. 2006, *MNRAS*, 369, 1654
- Girardi, L., Bressan, A., Bertelli, G., & Chiosi, C. 2000, *A&AS*, 141, 371
- Green, R. F. 1980, *ApJ*, 238, 685
- Grocholski, A. J. & Sarajedini, A. 2003, *MNRAS*, 345, 1015
- Hansen, B. M. S. et al. 2004, *ApJS*, 155, 551
- . 2007, *ArXiv Astrophysics e-prints*
- Hanson, R. B., Klemola, A. R., Jones, B. F., & Monet, D. G. 2004, *AJ*, 128, 1430

- Harris, H. C. et al. 2003, *AJ*, 126, 1023
- . 2006, *AJ*, 131, 571
- Hernandez, X. & Valls-Gabaud, D. 2008, *MNRAS*, 383, 1603
- Hu, Q., Wu, C., & Wu, X.-B. 2007, *A&A*, 466, 627
- Hügelmeier, S. D., Dreizler, S., Homeier, D., Krzesiński, J., Werner, K., Nitta, A., & Kleinman, S. J. 2006, *A&A*, 454, 617
- Jeffery, E. J. 2009, PhD thesis, AA(Texas Univ., Austin.)
- 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
- Kepler, S. O., Kleinman, S. J., Nitta, A., Koester, D., Castanheira, B. G., Giovannini, O., Costa, A. F. M., & Althaus, L. 2007, *MNRAS*, 375, 1315
- Kleinman, S. J. et al. 2004, *ApJ*, 607, 426
- Koester, D., Kepler, S. O., Kleinman, S. J., & Nitta, A. 2009, *Journal of Physics Conference Series*, 172, 012006
- Koester, D. et al. 2001, *A&A*, 378, 556
- Krzesiński, J., Nitta, A., Kleinman, S. J., Harris, H. C., Liebert, J., Schmidt, G., Lamb, D. Q., & Brinkmann, J. 2004, *A&A*, 417, 1093

- Landolt, A. U. & Uomoto, A. K. 2007, AJ, 133, 768
- Liebert, J., Bergeron, P., & Holberg, J. B. 2005, ApJS, 156, 47
- Liebert, J., Dahn, C. C., Gresham, M., & Strittmatter, P. A. 1979, ApJ, 233, 226
- Liebert, J., Dahn, C. C., & Monet, D. G. 1988, ApJ, 332, 891
- Luyten, W. J. 1958, *On the Frequency of White Dwarfs in Space* (Minneapolis: University of Minnesota Observatory)
- McCook, G. P. & Sion, E. M. 1999, ApJS, 121, 1
- Mermilliod, J.-C. & Mermilliod, M. 1994, Catalogue of Mean UBV Data on Stars (Catalogue of Mean UBV Data on Stars, VI, 1387 pp.. Springer-Verlag Berlin Heidelberg New York)
- Mermilliod, J.-C., Mermilliod, M., & Hauck, B. 1997, A&AS, 124, 349
- Miller, G. E. & Scalo, J. M. 1979, ApJS, 41, 513
- Monet, D. G. et al. 2003, AJ, 125, 984
- Munn, J. A. et al. 2004, AJ, 127, 3034
- Pagel, B. E. J. & Portinari, L. 1998, MNRAS, 298, 747
- Paulson, D. B., Sneden, C., & Cochran, W. D. 2003, AJ, 125, 3185

- Perryman, M. A. C. & ESA, eds. 1997, ESA Special Publication, Vol. 1200,
The HIPPARCOS and TYCHO catalogues. Astrometric and photometric
star catalogues derived from the ESA HIPPARCOS Space Astrometry Mis-
sion
- Perryman, M. A. C. et al. 1998, A&A, 331, 81
- Peterson, D. M. & Solensky, R. 1988, ApJ, 333, 256
- Peterson, R. C. & Green, E. M. 1998, ApJ, 502, L39+
- Reid, I. N. 1996, AJ, 111, 2000
- Reid, N. 1992, MNRAS, 257, 257
- Richards, G. T. et al. 2002, AJ, 123, 2945
- Salaris, M., Serenelli, A., Weiss, A., & Miller Bertolami, M. 2009, ApJ, 692,
1013
- Sarajedini, A., von Hippel, T., Kozhurina-Platais, V., & Demarque, P. 1999,
AJ, 118, 2894
- Schiller, S. J. & Milone, E. F. 1987, AJ, 93, 1471
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
- Schmidt, M. 1968, ApJ, 151, 393
- Spergel, D. N. et al. 2003, ApJS, 148, 175

- . 2007, *ApJS*, 170, 377
- Stoughton, C. et al. 2002, *AJ*, 123, 485
- Taylor, B. J. 1980, *AJ*, 85, 242
- . 1994, *PASP*, 106, 600
- . 1998, *PASP*, 110, 708
- Taylor, B. J. & Joner, 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., DeGennaro, S., Stein, N., Jefferys, W. H., & von Hippel, T. 2009, *The Annals of Applied Statistics*, 3, 117
- VandenBerg, D. A. & Stetson, P. B. 2004, *PASP*, 116, 997
- 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. & Sarajedini, A. 1998, *AJ*, 116, 1789
- von Hippel, T., Steinhauer, A., Sarajedini, A., & Deliyannis, C. P. 2002, *AJ*, 124, 1555
- Weidemann, V. 1967, *Zeitschrift fur Astrophysik*, 67, 286

- . 2000, *A&A*, 363, 647
- Weidemann, V., Jordan, S., Iben, I. J., & Casertano, S. 1992, *AJ*, 104, 1876
- Williams, K. A. 2007, in *Astronomical Society of the Pacific Conference Series*, Vol. 372, 15th European Workshop on White Dwarfs, ed. R. Napiwotzki & M. R. Burleigh, 85–+
- Williams, K. A., Bolte, M., & Koester, D. 2009, *ApJ*, 693, 355
- Winget, D. E., Hansen, C. J., Liebert, J., van Horn, H. M., Fontaine, G., Nather, R. E., Kepler, S. O., & Lamb, D. Q. 1987, *ApJ*, 315, L77
- Wood, M. A. 1992, *ApJ*, 386, 539
- Wood, M. A. 1995, in *Lecture Notes in Physics*, Berlin Springer Verlag, Vol. 443, White Dwarfs, ed. D. Koester & K. Werner, 41
- Wood, M. A. & Oswalt, T. D. 1998, *ApJ*, 497, 870
- Woosley, S. E., Heger, A., & Weaver, T. A. 2002, *Reviews of Modern Physics*, 74, 1015
- Yi, S., Demarque, P., Kim, Y.-C., Lee, Y.-W., Ree, C. H., Lejeune, T., & Barnes, S. 2001, *ApJS*, 136, 417
- York, D. G. et al. 2000, *AJ*, 120, 1579
- Zacharias, N. et al. 2004, *AJ*, 127, 3043

Vita

Steven DeGennaro is the 17th of 128 children born to a couple of really poor immigrants from the planet Gleep-Glop B. He was raised in the wild by a pack of ravenous wolves until discovered in the early 30's by a group of Zoroastrian missionaries, who adopted him and taught him how to read, write, and bathe himself. After murdering the missionaries, burning down the mission, and collecting the insurance money, he set out on his own to travel the world, ultimately ending up in Burma (now Myanmar), where he quickly rose through the ranks of the ruling military junta. Unfortunately, a mistimed burst of flatulence led to a diplomatic nightmare at the UN and he was sentenced to death by firing squad. Having had the presence of mind to wear his Kevlar vest on the day the sentence was carried out, he was spared the worst of it. He took his really, really sore ribs and his internet poker winnings and escaped the country on a raft made from the bones of dead revolutionaries bound together with VHS tape from his formerly extensive collection of homemade pornography. He emigrated to the United States, was granted asylum (after a hefty bribe which included the last of *both* his internet poker winnings *and* his increasingly less-extensive collection of homemade pornography), and enrolled at the University of Texas at Austin in the doctoral program in Astronomy. These days, he can often be found drinking from the turtle pond on campus.

Permanent address: 3805 Bailey Ln.
Austin, Texas 78756

This dissertation was typeset with \LaTeX^\dagger by the author.

[†] \LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's \TeX Program.