

Dynamical Inference of Black Hole Populations in Globular Clusters

by

Peter J. Smith

A Thesis Submitted to Saint Mary's University, Halifax, Nova Scotia in Partial Fulfillment
of the Requirements for the Degree of MSc in Astronomy
(Department of Astronomy and Physics)

2024, Halifax, Nova Scotia

© Peter Smith, 2024

Approved: _____

Dr. Vincent Hénault-Brunet

Approved: _____

Dr. Robert Thacker

Approved: _____

Dr. Marcin Sawicki

Date: August 14, 2024.

Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Vincent Hénault-Brunet, for his guidance, support, and patience throughout the many years that we have worked together. Having started some of this work before the pandemic as an undergraduate and finishing it more than four years later, I have been very fortunate to work with someone who has taught me so much about being a scientist and has been an invaluable mentor. I would also like to thank my committee members, Dr. Robert Thacker and Dr. Marcin Sawicki, for their time and feedback on this work and for their support throughout my time at Saint Mary's.

I would like to thank my fellow graduate students, for their support, advice and camaraderie. In particular, I would like to thank Nolan Dickson with whom I have worked closely with for several years and who has been a great source of support and technical advice throughout our many projects together.

Finally, I would like to thank my family and friends for their support and encouragement throughout this process. The past few years have been challenging for everyone, and I am grateful to have had such a strong support network to help me through it and to celebrate my successes with me.

Contents

1	Introduction	1
1.1	Black Holes in Globular Clusters	1
1.1.1	Numerical and theoretical predictions	2
1.1.2	Gravitational waves	6
1.2	Direct Detections	8
1.2.1	Radio and X-ray detections	9
1.2.2	Radial velocity detections	13
1.2.2.1	The uncertain extrapolation from direct detections	15
1.2.3	Microlensing	16
1.3	Dynamical Modelling of Globular Clusters	17
1.3.1	Dynamical signatures of BH populations	18
1.4	Summary and Goals of this Thesis	21
2	Multimass modelling of globular clusters and their black hole populations: method validation with mock data	23
2.1	Background	24
2.2	Validation of BH population inference.	26
2.2.1	Mock observations	27

2.2.1.1	Number density profiles	27
2.2.1.2	Proper motion dispersion profiles	28
2.2.1.3	Line-of-sight velocity dispersion profiles	29
2.2.1.4	Stellar mass functions	29
2.2.2	Validation results	30
2.2.3	Appendix: Validation results	37
3	Discussion of other methods for BH population inference	39
3.1	Modelling of individual clusters	39
3.1.1	Evolutionary models	40
3.1.2	Equilibrium models	42
3.2	Scaling relations	46
3.2.1	MOCCA	47
3.2.2	CMC	51
3.3	Comparison of inferred black hole populations	54
4	Probing populations of dark stellar remnants in the globular clusters 47 Tuc and Terzan 5 using pulsar timing: Introduction, Data and Methods	58
4.1	Introduction	59
4.2	Data	62
4.2.1	Kinematics and density profiles	63
4.2.1.1	Proper motion dispersion profiles	63
4.2.1.2	Line-of-sight velocity dispersion profiles	63
4.2.1.3	Number density profiles	64
4.2.2	Stellar mass functions	65
4.2.3	Pulsar data	66

4.3	Methods	67
4.3.1	Models	67
4.3.2	Fitting	69
4.3.3	Likelihoods	70
4.3.3.1	Pulsars	71
4.3.4	Stellar mass functions	80
4.3.5	Stellar populations and mass bins	81
5	Probing populations of dark stellar remnants in the globular clusters 47 Tuc and Terzan 5 using pulsar timing: Results, Discussion and Conclusions	83
5.1	Results	84
5.1.1	47 Tuc	86
5.1.2	Terzan 5	90
5.2	Discussion	95
5.2.1	Constraints from pulsars	95
5.2.2	Mass of Terzan 5	98
5.2.3	Comparison of mass in BHs to literature results	100
5.2.4	An IMBH in 47 Tuc?	101
5.2.5	Central velocity dispersion of Terzan 5	104
5.2.6	Terzan 5: a comparison with ω Cen	105
5.3	Conclusions	106
5.4	Supplementary Material	108
5.4.1	Incorporating dispersion measures	108
5.4.2	Pulsar Data	109
5.4.3	Supplementary Figures	113

6	Conclusions	120
6.1	Summary of indirect methods for detecting black holes	121
6.2	This thesis	122
6.2.1	Tests of equilibrium models with mock data	122
6.2.2	Pulsar timing constraints for dynamical modelling	123
6.3	Future work	126

List of Figures

1.1	Primary mass distribution of BH-BH mergers detected by LIGO	8
1.2	Radio/X-ray ratios for X-ray binaries and stellar mass BHs.	10
1.3	The fitted radial velocity curve of the companion of the candidate BH in NGC 3201.	14
1.4	Surface brightness profile of NGC 3201 compared to several model profiles.	21
2.1	Example fit to mock observations of a CMC snapshot matched to NGC 6256.	31
2.2	Example fit to mock mass function data of a CMC snapshot matched to NGC 6256.	32
2.3	Comparison of true vs inferred mass and half-mass radius.	33
2.4	Selection of problematic snapshots in $\delta - N_{\text{relax}}$ plane and true vs inferred BH fraction for all snapshots.	33
2.5	True vs inferred values of f_{BH} from mock observations of CMC snapshots.	36
3.1	The density of the BHS as a function of the average surface luminosity from MOCCA models.	48
3.2	The density of the BHS as a function of the number of BHs in the BHS from the CMC catalogue.	50
3.3	The number of BHs in various CMC snapshots as a function of the two mass segregation parameters.	53

3.4	Comparison of the inferred BH populations for a selection of Milky Way GCs from three modelling studies discussed in this chapter.	56
3.5	Error-weighted residuals between the inferred value of f_{BH} from Dickson et al. (2024) and the values from two literature sources, Askar et al. (2018) and Weatherford et al. (2020).	56
4.1	The $P-\dot{P}$ plane for the field pulsars with a Gaussian KDE.	73
4.2	Construction of $P(a_{\text{cl},z} R_i)$ distribution from model quantities.	76
4.3	Construction of \dot{P}/P likelihood function.	78
5.1	Best-fitting model for the 47Tuc-AllData fit compared to the data.	87
5.2	Continuation of Figure 5.1 showing the fit to the stellar mass function data. . .	88
5.3	Likelihood functions corresponding to the best-fitting model for the observed \dot{P}/P and \dot{P}_b/P_b for one pulsar in each cluster.	89
5.4	Best-fitting model for the 47Tuc-NoKin fit compared to the data	90
5.5	Posterior probability distribution of the mass in BHs in 47 Tuc for each of the fits summarized in Table 5.1.	91
5.6	Best-fitting model for the Ter5-AllData fit compared to the data.	92
5.7	Continuation of Figure 5.6 showing the best-fitting model for the Ter5-AllData fit compared to the stellar mass function data.	94
5.8	Posterior probability distribution of the mass in BHs in Terzan 5 for each of the fits summarized in Table 5.1.	95
5.9	Minimum and maximum values of \dot{P}/P allowed from the acceleration in the cluster potential for the best-fitting models of our 47Tuc-AllData (top) and Ter5-AllData (bottom) fits.	97

5.10	The inferred cumulative mass profile of Terzan 5 from our <code>Ter5-AllData</code> fit with literature values also plotted.	99
5.11	Likelihood functions corresponding to the best-fitting model (<code>47Tuc-AllData</code>) for the observed \dot{P}/P for each pulsar in 47 Tuc.	114
5.12	Continuation of Figure 5.11.	115
5.13	Likelihood functions corresponding to the best-fitting model (<code>47Tuc-AllData</code>) for the observed \dot{P}_b/P_b for each pulsar in 47 Tuc with an orbital timing solution.	116
5.14	Likelihood functions corresponding to the best-fitting model (<code>Ter5-AllData</code>) for the observed \dot{P}/P for each pulsar in Terzan 5.	117
5.15	Continuation of Figure 5.14.	118
5.16	Likelihood functions corresponding to the best-fitting model (<code>Ter5-AllData</code>) for the observed \dot{P}_b/P_b for each pulsar in Terzan 5 with an orbital timing solution.	119

List of Tables

2.1	Results of the validation fits.	38
4.1	Model parameters and their priors.	69
5.1	Summary of model fits for 47 Tuc and Terzan 5	84
5.2	Results for each of the six fits we perform of 47 Tuc and Terzan 5.	85
5.3	Reported masses in BHs (or upper limits) in 47 Tuc and Terzan 5 from dynamical studies in the literature and from this work.	101
5.4	Pulsar timing data used in this work for 47 Tuc.	110
5.5	Pulsar timing data used in this work for Terzan 5.	111
5.6	Pulsar timing data used in this work for Terzan 5.	112

ABSTRACT

Peter J. Smith

Dynamical Inference of Black Hole Populations in Globular Clusters

August 14, 2024

The present-day black hole (BH) populations of globular clusters (GCs) have far-reaching implications for the evolution of these systems and the production of gravitational wave sources. Despite this, our ability to constrain these populations remains unclear. Using synthetic observations from CMC models, we test one dynamical method, finding it to be effective but likely underestimating uncertainties. Additionally, we introduce a new technique for combining dynamical models with pulsar timing, and apply it to two Milky Way GCs. This reveals a modest BH population in 47 Tuc (with a total mass of $446^{+75}_{-72} M_{\odot}$), suggesting no need for a large central BH, and sets a significant upper limit for the mass in BHs in Terzan 5 of $\sim 4000 M_{\odot}$. This technique, not reliant on stellar kinematics, offers insights into the dynamics of even heavily obscured GCs and will benefit from future radio telescope observations.

Chapter 1

Introduction

1.1 Black Holes in Globular Clusters

Globular Clusters (GCs) are old, dense, spherical collections of stars that are found in most galaxies. These clusters typically contain $10^4 - 10^6$ stars, all of which formed from the same initial cloud of gas (e.g. Carroll & Ostlie 2014). These systems are of great interest to astronomers for a number of reasons, including as pristine examples of stellar evolution, as relics of the earliest epochs of galaxy formation, and as ideal systems in which to study stellar interactions (e.g. Binney & Merrifield 1998; Heggie & Hut 2003; Forbes et al. 2018).

GCs are some of the oldest objects in the Universe with ages typically in the range of 10 – 13 Gyr, meaning that they formed during the earliest epochs of galaxy formation (Forbes et al. 2018). The old ages of GCs also means that all stars more massive than $\sim 1 M_{\odot}$ have evolved off the main sequence, leaving behind a population of white dwarfs, neutron stars, and stellar-mass BHs (e.g. Dickson et al. 2023). The dense environments of GCs mean that rare stellar interactions are much more common than in the field, allowing us to study exotic stellar systems like blue stragglers, X-ray binaries, and millisecond pulsars (e.g. Ye et al. 2022).

The present-day stellar-mass black hole (BH) populations of globular clusters are of great

interest to astronomers for a number of reasons. As the most massive objects in GCs, BHs are expected to have a significant effect on the dynamical evolution of their host cluster, and as such their present-day populations are important to constrain if we wish to understand the evolution of these systems (e.g MacKey et al. 2008). The dense environments of GCs also make them ideal sites for the formation of BH-BH binaries (Portegies Zwart & McMillan 1999), which are expected to be the progenitors of a large portion of the gravitational wave sources that have been detected by the LIGO/Virgo/KAGRA collaboration (The LIGO Scientific Collaboration et al. 2023). Despite the great interest in these systems, the present-day populations of BHs in GCs are still poorly constrained. This stems from the inherent difficulty in observing BHs combined with the difficulty of modelling these dense stellar systems. In this chapter, we will discuss the current state of knowledge of the present-day BH populations of GCs, and the methods that have been used to constrain these populations. We will examine some of the theoretical and numerical predictions that have been made regarding these populations before discussing the observational methods that can be used to detect these objects. Finally, we will look at the dynamical effects of BHs in GCs and the implications of these effects for our ability to infer the presence of BHs in these systems through dynamical modelling.

1.1.1 Numerical and theoretical predictions

The presence of BHs is expected to have a significant impact on the dynamics of GCs, especially in the central regions where the BHs are expected to be concentrated due to mass segregation (e.g. Kremer et al. 2019). These centrally concentrated black hole systems (BHS) have been the subject of a number of theoretical and numerical studies.

One of the first papers to look at the retention and dynamical ejection of BHs in GCs was Spitzer (1969). In this work, the author performs an analytic exploration of the effects of two-

body encounters and the tendency for a cluster to move towards kinetic energy equipartition, specifically in the case where there is a population of much more massive objects. The key finding of this study was the prediction of the so-called ‘‘Spitzer stability criterion’’ that governs the long-term evolution of the BHS within a cluster:

$$M_2/M_1 < \beta (m_1/m_2)^{3/2}, \quad (1.1)$$

where M_1 and M_2 are the total masses of the two populations (M_1 being the population with the larger total mass and lower mean stellar mass, representing the typical star in a GC, while M_2 is the population with the smaller total mass and higher mean stellar mass, representing the BHs), m_1 and m_2 are the mean stellar masses of the two populations, and β is a numerical constant that depends on the densities and half-mass radii of the two populations. For systems that satisfy this criterion the BHS will be dynamically stable and will trend towards kinetic energy equipartition with the rest of the cluster. For systems that do not satisfy this criterion, the BHS will not be able to reach energy equipartition with the rest of the cluster and will instead form a dynamically decoupled subsystem in the core of the cluster which will rapidly evaporate via strong dynamical interactions within the BHS and ejections. By inspecting Equation 1.1 we can see that systems where the BHs make up a large fraction of the total mass of the cluster, or where the BHs are much more massive than the typical star in the cluster, will be Spitzer unstable. This work further predicted that most open clusters, globular clusters and nuclear star clusters would be unstable and thus would not be expected to ultimately retain a significant population of BHs (though with different timescales of BH depletion for each type of system).

The consequences of this result were explored by Sigurdsson & Hernquist (1993) who highlighted the role of BH-BH binaries which tend to form in the cores of clusters on short timescales due to the large gravitational cross-sections of BHs. These binaries are firmly in the ‘‘hard binary’’

regime, meaning that they will tend to get tighter over time through three-body interactions with other objects in the cluster (Heggie 1975). The binding energy of these binaries is the primary energy source by which other BHs are ejected. Overall, Sigurdsson & Hernquist (1993) predict just a few (0-4) BHs to be retained in the core of a typical cluster, along with a few in the outer regions of the cluster due to incomplete ejections where the ejection velocity is insufficient to fully eject the BH from the cluster.

These early theoretical studies are in stark contrast to more recent theoretical and numerical works that generally predict much larger populations of BHs to be retained in GCs by the present day. A more recent study by Breen & Heggie (2013) used idealized two-component clusters to study the evolution of systems with two mass components. They make assumptions about the two components that are expected to be reasonable for GCs, namely that the ratio of the mean mass of the two components is larger than 10 and the more massive component (the BHs) is about 1% of the total cluster mass. They first find that the ratio of the half-mass radii of the two components approximately scales with the ratio of the mean masses for both Spitzer stable and unstable clusters (though with a different dependence in each case). Breen & Heggie (2013) then go on to examine the evaporation and ejection rates of the BHS, first finding that evaporation from two-body encounters within the BHS is a fairly insignificant source of mass loss for the BHS. Ejections on the other hand, which are typically driven by three-body interactions between BH-BH binaries and single BHs are found to be largely independent of the properties of the BHS (M_2/M_1 and m_2/m_1), instead depending on the half-mass relaxation time of the whole cluster. This dependence on the half-mass relaxation time of the whole cluster is due to the dependence of the half-mass relaxation time on the central density of the cluster, which sets the interaction rate between BHs in the core of the cluster.

Due to the negligible contribution of evaporation, this implies that the evolution of the BHS

as a whole is independent of the BHS properties and is instead coupled to the properties of the cluster as a whole. By comparing their theoretical results to numerical N -body models of these idealized systems, the mass loss rate of the BHS (primarily due to ejections) is found to occur on much slower timescales than previously predicted, with the BHS taking between 2 and 10 initial half-mass relaxation times ($t_{rh} \approx 1$ Gyr) to lose 90% of its mass. Overall, this study lays the theoretical groundwork for the idea that the BHS in GCs are not as short-lived as previously thought, and that the BHS can be expected to survive for a significant fraction of the cluster's lifetime.

On the numerical side, Morscher et al. (2013) and Morscher et al. (2015) performed Monte Carlo simulations (we discuss the use of Monte Carlo models in greater depth in Chapter 3) of GCs, this time with a full mass spectrum instead of only two mass components. Their simulations adopt a reasonable range of initial conditions and despite not being tuned to any particular Milky Way cluster are still able to qualitatively reproduce the observed properties of real GCs (e.g. mass and half-light radius). These numerical studies confirm a number of the results of the theoretical results of Breen & Heggie (2013), namely that the evolution of the BHS is indeed coupled to the evolution of the whole cluster, and that the BHS can be expected to survive for a significant portion of the cluster's lifetime and up to the present day in many cases. A particularly interesting result of these simulations is the behaviour of the BHS and whether it ever reaches the "Spitzer unstable" regime. The authors find in their simulations that the BHS does form very rapidly in the core of the cluster (within the first Gyr) but that three-body interactions in the BHS are a critical driver of the overall evolution of the BHS. Through these interactions, BHs are frequently ejected from the BHS (though retained in the cluster) meaning that the BHs are able to mix with the rest of the cluster and exchange energy. This means that the BHS as a whole is able to reach energy equipartition with the rest of the cluster. In their simulations they find that less than half of the BHs are ejected from the clusters within 12 Gyr, leaving anywhere from a few hundred to

thousands of BHs per cluster by the present day.

The results of both Breen & Hoggie (2013) and Morscher et al. (2013, 2015) are in stark contrast to the earlier theoretical results of Spitzer (1969) and Sigurdsson & Hernquist (1993) which predicted very few remaining BHs left by the present day. These studies highlight the importance of three-body interactions in the evolution of the BH populations in GCs, and the importance of numerical simulations in understanding the behaviour of these complex systems.

1.1.2 Gravitational waves

An additional motivation for studying the present-day BH populations of GCs is the possibility that dense clusters may be the source of a significant fraction of the gravitational wave (GW) events that have started to be detected over the past decade. Some of the early predictions of GW sources in GCs came from Portegies Zwart & McMillan (1999) who looked at dynamically formed BH-BH binaries in GCs. Through the dynamical processes we discussed in the previous section, BHs rapidly sink to cores of clusters due to mass segregation. In the cores of clusters these BHs frequently form binaries through two-body interactions (e.g. Sigurdsson & Hernquist 1993). These binaries are hardened through three-body interactions, and in the N -body simulations of Portegies Zwart & McMillan (1999) about 30% of the BHs are ejected as intact binaries. These ejected binaries have relatively short orbital periods, and are predicted to undergo mergers within a few Gyr, which we can detect as gravitational waves.

Since the first detection of a GW event by LIGO in 2015 (Abbott et al. 2016), we have slowly built up a picture of the population of GW sources in the Universe. Of particular interest to studies of GCs are the sources whose progenitor BHs seem to fall within the mass gap in the $40 - 60 M_{\odot}$ range predicted by stellar evolution models (e.g. Spera & Mapelli 2017). These sources can be explained by hierarchical mergers of BHs in dense stellar environments like GCs

(e.g. Rodriguez et al. 2016; Antonini et al. 2023). These hierarchical mergers are cases where a BH-BH binary merges within the dense core of a GC, and the resulting BH is then retained in the cluster where it undergoes yet another merger, filling in the predicted mass gap. Another property of some GW signals that is of interest to studies of GCs is the misalignment of the spins of the BHs in the binary. Misaligned spins are expected to be common in GCs due to the chaotic nature of the BH dynamics in these systems and are expected to be much harder to produce with isolated binary evolution (e.g. Fishbach & Fragione 2023).

Simulations of the GW sources in GCs performed by Antonini et al. (2023) show that GCs are able to account for the mass-gap sources detected by LIGO, only in the case where hierarchical mergers are considered. These simulations use semi-analytic models of GCs, tuned to N -body models to follow the evolution of GCs along with their BH populations without having to perform computationally expensive N -body simulations. Figure 1.1 shows the distribution of BH-BH binary primary masses detected by LIGO (green) compared to the distribution of primary masses in GCs with hierarchical mergers (blue) and without (black). The hierarchical merger model is able to reproduce the observed distribution of primary masses in the mass gap, while the model without hierarchical mergers clearly displays the upper limit in BH mass predicted by stellar evolution models.

Studying the present-day populations of BHs in GCs provides a unique opportunity to study the progenitors of these GW sources. In particular, models of cluster evolution that purport to explain the observed population of GW signals must also predict present-day populations of BHs that are consistent with the populations that we can infer from observations and modelling. This means that the present-day populations of BHs in GCs can act as boundary conditions for models of GC evolution, and that the inferred populations of BHs in GCs can be used to constrain the contribution of GCs to the overall population of GW signals (e.g. Fishbach & Fragione 2023).




Figure 1.1: Panel 1 of Figure 1 of Antonini et al. (2023). The primary mass distribution of BH-BH mergers detected by LIGO (green) compared to the distribution of primary masses in GCs with hierarchical mergers (blue) and without (black). The mass gap is clearly visible in the model without hierarchical mergers around $50 M_{\odot}$.

As we have discussed throughout this introduction, the present-day populations of BHs in GCs are relevant to many areas of astrophysics. The dynamical evolution of GCs is intimately tied to the presence of BHs, and the present-day populations of BHs in GCs can be used to constrain models of GC evolution. The detection of GW signals from BH-BH mergers, particularly those with massive progenitors or misaligned spins have provided further motivation for studying the present-day populations of BHs in GCs. In the following sections we will discuss the observational methods that have been used to directly detect BHs in GCs and the dynamical signatures of BH populations that have been used to indirectly detect the presence of these populations.

1.2 Direct Detections

The conceptually simplest way for astronomers to constrain the BH population in a GC would be to simply individually detect each BH and count them. This is of course impossible in practice

as isolated BHs emit no light of their own and can only be detected through their gravitational influence on other objects. This means that the only way to directly detect BHs with electromagnetic radiation is through their interactions with other matter, typically in the form of accretion disks. This section will discuss the various methods for detecting BHs in GCs, paying particular attention to the handful of confirmed detections and how they exemplify the difficulties of this approach.

1.2.1 Radio and X-ray detections

Among the various methods for detecting BHs in GCs, direct detections of accreting BHs in the radio and X-ray bands are perhaps the most obvious method, though as we will see, the confident identification of a source as a BH is by no means straightforward. These detections rely on the emission from the super-heated accretion disks that surround accreting compact objects. Stellar-mass BHs are observed to emit in the radio and X-ray bands due to synchrotron and blackbody processes in the jet and corona (e.g. Gallo & Fender 2005). These emission processes are however present for most other accreting compact objects, including cataclysmic variables (CVs), neutron stars (NSs), white dwarfs (WDs), and various types of pulsars (e.g. Chomiuk et al. 2013). Observers must also consider the possibility that the emission is not associated with the GC at all and is instead from a background active galaxy, meaning the confident identification of a BH is challenging. In the absence of spectroscopic observations, the main diagnostic for identifying the nature of a source is by examining the relation between the radio and X-ray luminosities (e.g. Gallo et al. 2006, 2012). Figure 1.2 shows the relation between the radio and X-ray luminosities for a variety of accreting compact objects that have been confidently classified with observed relations for BHs and NSs plotted as dashed lines. The use of this relation to identify objects is further complicated by the nature of the objects themselves. The luminosities of accreting objects are often inherently variable (e.g. Longair 2011) meaning that

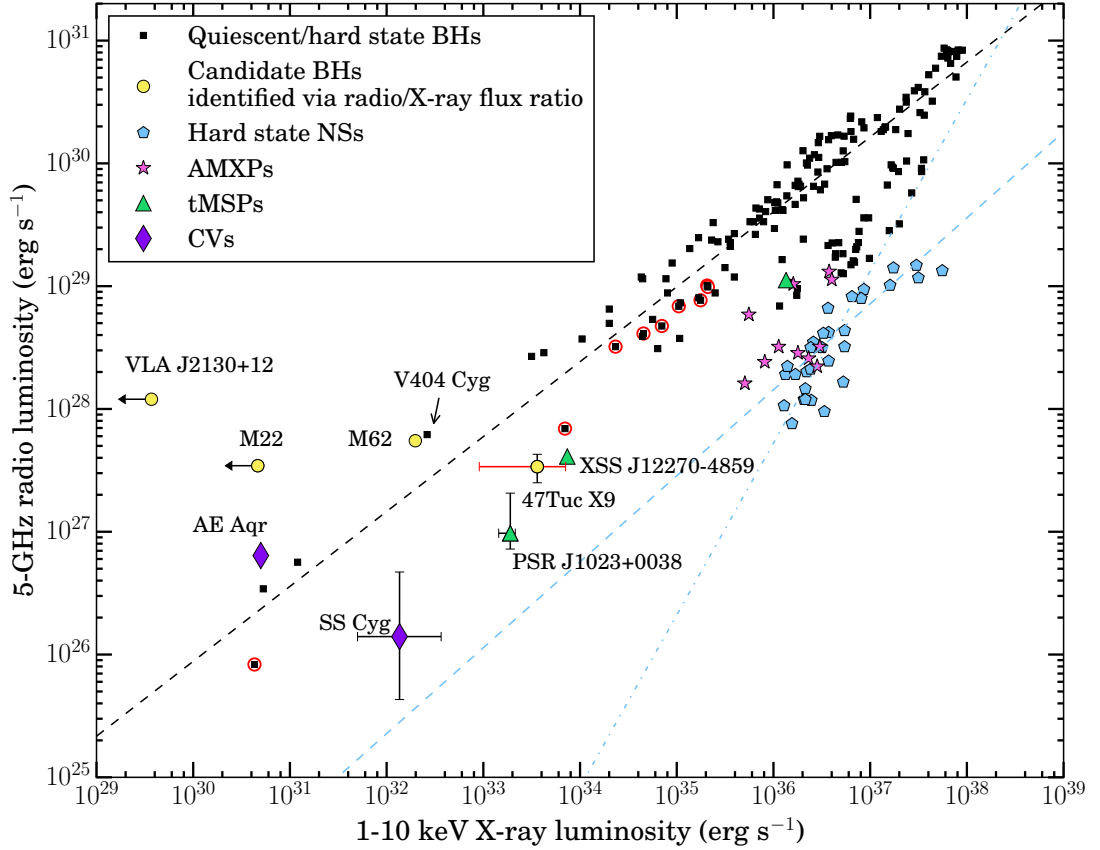


Figure 1.2: Figure 11 from Bahramian et al. (2017). The radio/X-ray ratios for X-ray binaries and stellar mass BHs (black points) are plotted along with the observed relation for BHs (black dashed line). Various other types of accreting compact objects are in blue (NSs), pink (accreting X-ray millisecond pulsars), green (transitional millisecond pulsars) and purple (cataclysmic variables.). The two dashed blue lines are two different observed relations for NSs. The candidate BHs are the yellow points and mostly fall into the region where this relation is poorly constrained. The black squares circled in red show the track of one individual BH candidate, XTE J1118+480.

not only are simultaneous radio and X-ray observations required to confidently place an object on the relation, there remains the possibility that the object is undergoing a transient flaring phase that is not representative of its typical state. Finally, due to the small number of stellar-mass BHs known in the Milky Way, the fainter end of the radio to X-ray relation is poorly constrained and appears to have multiple classes of objects overlapping with the observed relation for BHs, making the confident classification of sources even more difficult.

Among the first reported detections of a BH candidate in a GC was the detection of a radio source in M22 by Strader et al. (2012). This work reported the detection of two potential BHs in M22 using radio observations made with the Very Large Array (VLA). The authors classified these objects as actively accreting based on their radio emission, noting that the flat nature of their spectra indicated relatively low accretion rates, about 2 – 3% of the Eddington rate. The authors make an argument based on mass segregation and energy equipartition to estimate the masses of the BHs based on their projected radial distances. In brief, the authors apply a relation between the mean mass ratio of BHs to stars and core-radius to BH half-mass radius from Kulkarni et al. (1993):

$$m_{\text{BH}}/m_{\star} = (r_c/r_{\text{BH}})^2$$

where m_{BH} and m_{\star} are the masses of the BH and typical star in the cluster, and r_c and r_{BH} are the core radius and half-mass radius of the BH population. By assuming that the cluster is in energy equipartition and that the BHs are located at roughly the half-mass radius of the BH population, the authors estimate the masses of the BHs, finding masses of roughly $15 M_{\odot}$ which are indeed consistent with the systems being binary systems hosting a BH. The circumstantial nature of this argument, combined with the lack of X-ray data to place the objects on the radio to X-ray luminosity relation, means that the identification of these objects as BHs needs further confirmation.

The first candidate BH in a Milky Way GC to be identified using *both* radio and X-ray data was reported by Chomiuk et al. (2013). The authors again used radio data from the VLA to identify a source as a potential BH candidate, this time in M62. What makes this candidate more compelling than the previous candidates is the existence of archival X-ray data of the same region from Chandra. While the X-ray data and radio data are separated by more than a decade, the X-ray data nonetheless allows the authors to place the source on the radio to

X-ray luminosity relation, where it falls squarely on the expected relation for BHs (with the caveat that there is possibly variability in the X-ray and radio luminosity of the source that they are not accounting for). The authors further analyze the X-ray spectrum of the source, finding that it is also consistent with the spectrum expected for an accreting NS, meaning that the radio to X-ray ratio remains the key diagnostic tool for discerning the nature of the object. The authors enumerate the various alternative explanations for the source, and conclude that the next most likely explanation is that the source is a flaring NS. This would imply that the NS was in a quiescent state during the X-ray observations but was flaring during the radio observations, causing the radio to X-ray ratio to be roughly two orders of magnitude higher than expected for a NS, again highlighting the importance of simultaneous radio and X-ray observations.

Perhaps the most promising X-ray/radio detection of a BH in a GC to date is the detection of a BH in 47 Tuc first reported by Miller-Jones et al. (2015) and later followed-up upon by Bahramian et al. (2017). While this source has been known to astronomers since the launch of the Einstein observatory in the 1980s (Hertz & Grindlay 1983), it was not until more recent radio observations were obtained that the source was identified as a potential BH using the same radio to X-ray luminosity relation as the previous studies. UV spectra of the source were also obtained, showing a double-peaked C IV emission line interpreted as a clear sign of accretion, further adding to the novelty of this source. What makes this source particularly interesting is the follow-up observations made by Bahramian et al. (2017) who obtained simultaneous radio and X-ray observations, finally allowing for the confident placement of the source on the radio to X-ray luminosity relation. While the authors note that the X-ray spectrum of the source is unusual for a BH, they also note that the spectrum is not indicative of any other class of objects, and that we simply do not have a large enough sample of these objects to make any definitive statements about specific spectral features.

The current state-of-the-art search for BHs in GCs is represented by the MAVERIC survey (Shishkovsky et al. 2020; Tudor et al. 2022), designed to obtain deep radio observations of Milky Way GCs in order to identify actively accreting BHs and other exotic binary systems. The MAVERIC survey has already identified a new candidate BH in NGC 6539 (Bahramian et al. 2020), though this source is subject to the same caveats as previous candidates, with a transitional MSP being a possible alternative explanation. The MAVERIC survey has also looked at NGC 3201 (Paduano et al. 2022), a GC with known BHs confirmed through radial velocity measurements (Giesers et al. 2018, 2019) which we will discuss in the next section. These BHs in NGC 3201 were undetected in the radio and X-ray bands due to the detached nature of the BH-hosting binaries, again demonstrating the limitations of radio and X-ray searches for BHs in GCs.

While these detections are very exciting, it is worth noting that detections of individual systems do not necessarily tell us about the size of the whole BH population. In the next section we discuss radial velocity detections of BHs in GCs and, crucially, how the number of BHs detected in binaries is not necessarily correlated with the total number of BHs in a cluster.

1.2.2 Radial velocity detections

With the installation of multi-object spectrograph instruments like MUSE (Bacon et al. 2010) on major observatories, large radial velocity surveys have become much more feasible. One of the few methods for detecting non-accreting BHs is through analyzing the radial velocity curves of binary systems. With well-sampled data it is possible to constrain the minimum mass of an unresolved or dark binary companion, making it possible to detect BHs in binary systems (e.g. Giesers et al. 2018, 2019).




Figure 1.3: Figure 5 from Giesers et al. (2018). The fitted radial velocity curve of the companion of the candidate BH in NGC 3201. The middle and lower panels show the residuals of the fit and the radial velocity curve of a reference star respectively.

The first detection of a BH in a GC with radial velocity measurements was presented by Giesers et al. (2018) who reported a stellar-mass BH found in NGC 3201. This candidate was found in a survey of 25 Milky Way clusters with MUSE (Kamann et al. 2018) which obtained spectra for hundreds to tens of thousands of stars per cluster. After the spectra are processed to obtain radial velocities at multiple epochs, the orbital parameters of systems identified as binaries can be fit to obtain a lower limit on the mass of any unseen companion. The particular candidate reported by Giesers et al. (2018) is a main-sequence turn-off (MSTO) star that shows radial velocity variations on the scale of 100 km s^{-1} (see Figure 1.3 for the phase-folded radial velocity curve). No other components were detected in the spectrum, making a dark companion highly likely. The authors found a minimum companion mass of $4.36 M_{\odot}$ which allowed them to confidently identify it as a BH. The authors do note this could in theory be a triple system with a compact NS binary instead of a BH, but this would require two, very massive NSs just to reach the minimum companion mass, and such a system has not been observed (Giesers et al. 2018). In the same NGC 3201 dataset, Giesers et al. (2019) reported two additional BH candidates, one

of which requires more data to confirm due the current RV curve being well-fit by systems with a minimum companion mass of either $4.40 M_{\odot}$ or $1.10 M_{\odot}$. The other candidate is only explained by a companion with a minimum mass of $7.68 M_{\odot}$ and is therefore highly likely to be a BH.

While only these three BH candidates were detected in the NGC 3201 dataset, it is important to note that these easily detected binaries are expected to be fairly rare relative to the total number of BHs in a cluster. Kremer et al. (2018b) did a modelling-based analysis of the BH population in NGC 3201, partly motivated by the results of Giesers et al. (2018), and found that the models that best replicate the observables of NGC 3201 (see Section 3.1.1 for more details) are those with a significant retained BH population ($\sim 200 - 420$ BHs at the present day). Within these models they found only a handful of BHs in detached binaries like those reported in Giesers et al. (2018) and Giesers et al. (2019), highlighting the rarity of these kinds of systems.

1.2.2.1 The uncertain extrapolation from direct detections

Given the relatively rare nature of BHs in either a mass-transferring binary or a detached binary with a detectable companion, it is tempting to assume that by observing just a few of these systems in a given cluster, we could infer the presence of a much larger, unseen population of BHs in the cluster. The results of both Kremer et al. (2018a) and Kremer et al. (2018b) however would caution against applying this line of reasoning in general. They find that the number of BH-non-BH binaries is largely independent of the total number of BHs in the cluster. The BH-non-BH binaries that exist at the present day in their models are all formed dynamically; no primordial binaries survive as an intact pair within the cluster. The dynamical interactions that form these systems are predominately three-body interactions between BHs and stellar binaries, meaning the number of BH-non-BH binaries depends strongly on the number density of stellar binaries near the core of the cluster. Adding more BHs to a cluster heats the core and pushes

these stellar binaries further out, reducing the chances of an exchange interaction between the BH and the binary. Similarly, reducing the number of BHs in a cluster, while increasing the number of stellar binaries available in the core, reduces the number of BHs available to form these systems.

In addition to the effect mentioned above, Kremer et al. (2018b) reports that the median lifetimes of BH-non-BH binaries in their models are only about 400 Myr, even in a relatively low-density cluster like NGC 3201. This short lifetime is due to the frequent dynamical interactions that these systems are subject to in the core of the cluster. Even if several BH-non-BH binaries are formed in a cluster, these systems will quickly sink to the core of the cluster where they will inevitably be disrupted by dynamical interactions with other objects in the core. Increasing the number of these binaries only makes the disruption of these systems more likely, providing another channel through which the number of BH-non-BH binaries is self-regulating.

While we should not extrapolate the number of BHs in a cluster from the number of BH-non-BH binaries, the detection of a few of these systems can still be useful for indicating the presence of BHs in a cluster. The targeted modelling of NGC 3201 by Kremer et al. (2018b) indicates that the small number of detached BH binaries detected in NGC 3201 is still consistent with a much larger population of BHs in the cluster, showing the power of the modelling based approaches we will examine in Chapters 2 and 3 which rely on the dynamical signatures we discuss in Section 1.3.1.

1.2.3 Microlensing

The final method we will discuss for directly detecting BHs in GCs is microlensing. Microlensing occurs when the light from a background source is gravitationally lensed by a foreground object. In the case of microlensing, the lensing object is much smaller than a galaxy cluster and typically only distorts the light from the background source by a small amount (e.g. Paczynski

1986). These distortions manifest as both astrometric and photometric effects, where the lensed object is either projected onto a new position, magnified such that it appears brighter, or some combination of the two.

The idea of detecting BHs in GCs through microlensing has been known for some time (e.g. Paczynski 1994), but has not yet been fruitful due to the difficulties in detecting small astrometric or photometric perturbations that may only last for a few days (Kains et al. 2018; Zaris et al. 2020). In globular clusters the lensing objects that we would be most interested in would be the stellar-mass BHs near the centre of the cluster (lensing by e.g. WDs would be much more common, but harder to detect, and we are mainly interested in lensing by BHs) and the lensed objects would be either background stars or stars within the cluster itself. Numerical simulations indicate that for clusters with large background populations like 47 Tuc (the SMC) and M22 (the galactic bulge) lensing of background sources occurs at much higher rates, making these clusters and others like them the most promising targets for microlensing searches (Kiroğlu et al. 2022). Predicted rates of microlensing events from BHs in these clusters are on the order of 0.003 - 0.02 events per year within the half-light radii of the clusters (Kiroğlu et al. 2022), meaning that even for the most promising targets, the detection of significant populations of BHs through microlensing is unlikely.

1.3 Dynamical Modelling of Globular Clusters

Having now discussed the various methods for directly detecting BHs in globular clusters, we now turn to the modelling-based approaches for constraining the BH populations that will be the focus of this thesis. These methods make use of dynamical models of globular clusters to infer the BH population from observables. The methods can be broadly split into two categories: targeted modelling of individual clusters (e.g. Kremer et al. 2018b; Hénault-Brunet et al. 2020)

and studies that use populations of models to derive scaling relations between observables and the BH populations that hold for a wide range of clusters (e.g. Askar et al. 2018; Weatherford et al. 2020). Within these two categories, there is an assortment of models that can be used, including both evolutionary and equilibrium dynamical models.

1.3.1 Dynamical signatures of BH populations

Before we discuss the methods for indirectly detecting populations of BHs in GCs (Chapters 2, 3, 4 and 5), it is worth discussing the dynamical signatures of BHs in GCs that underpin these methods. As mentioned in Section 1.1.1, BHs will rapidly sink to the centre of GCs, forming a BHS in the core of the cluster. Critically, this BHS will not become fully dynamically decoupled from the rest of the cluster (e.g. Breen & Heggie 2013), meaning that the BHS will influence the evolution of the cluster as a whole. The key concept to understanding the dynamical signatures of a BH population is the idea of the BHS as an energy source in the core of the cluster. As stars in the core of the cluster interact with BHs, and BH-BH binaries in particular, they gain kinetic energy. This process has come to be known as “binary burning” and is the principal mechanism responsible for halting core collapse in GCs (Chatterjee et al. 2013; Kremer et al. 2018b). It is because of this process that we can be fairly certain that core-collapsed clusters do not contain BHs, as the presence of BHs in the core would prevent the runaway contraction of the cluster.

Aside from direct interactions with stars, the BHS can also influence the evolution of the cluster through internal interactions that kick BHs out of the core of the cluster. This process was explored in detail by MacKey et al. (2008) who analyzed N -body simulations designed to match the massive clusters found in the Small and Large Magellanic Clouds. Of particular note is their comparison of the evolution of two models that were identical in all aspects except for their BH content. These two simulations evolve very similarly until about 650 Myr where the model with BHs begin to show significant expansion of its core radius. Before this point, the

clusters undergo the expected process of core contraction both in visible stars and in the BH subsystem. After this point, the density of the BHS is approximately 80 times higher than the central stellar density, reaching the point when stable BH-BH binaries can form through three-body interactions. After the formation of these binaries begins, core-collapse is halted and the core radius of the cluster instead begins to expand. As mentioned earlier, this expansion is driven by the energy transfer from the BHS to the rest of the cluster which, in these simulations, occurs through two main channels. The first channel is the kick-out of single BHs into the cluster halo through three-body interactions with binary BHs. One of the three BHs in the interaction is kicked out of the core of the cluster into halo of the cluster, where it will rapidly sink back to the core of the cluster through two-body relaxation. This process of two-body relaxation deposits the kinetic energy of the BH into the stars in the halo of the cluster, heating the halo and causing it to expand. The second channel is the full ejection of BHs and BH-BH binaries from the cluster. This occurs when a three-body interaction occurs between a BH and a very tight binary BH. In this case, the recoil velocity of the BH or BH-BH binary is high enough that it can be fully ejected from the cluster, again depositing its kinetic energy into the halo of the cluster as it escapes the gravitational potential.

These processes of energy transfer from the BHS to the rest of the cluster are responsible for two of the most commonly discussed signatures of a BHS in GCs: a heightened central velocity dispersion and a shallower surface brightness profile. The physical intuition for both of these effects is the same, namely that the presence of BHs in the core of the cluster prevents the visible stars from becoming highly concentrated in the core by kicking them out or by heating halo stars before they would have entered the core. This means that the stars in and around the core would have higher than expected velocity dispersions and would be less concentrated than expected, leading to the observable heightened central velocity dispersion and shallower surface brightness profiles (e.g. Askar et al. 2018). This effect also manifests itself as a reduced degree

of mass segregation (MacKey et al. 2008; Weatherford et al. 2018), as the most massive stars are the most likely to interact with BHs and thus are the most likely to be either kicked out of the core or to gain kinetic energy from interactions with BHs, effectively stopping their inward migration.

It is this dynamical coupling between the BHS and the rest of the cluster that we can use in dynamical models to constrain the BH population in a given cluster. The presence of a BHS in a cluster has a significant impact on the phase-space distribution of the stars and thus with a sufficiently accurate dynamical model and sufficient observational data, we can in theory infer the presence of and the properties of a BHS in a cluster.

One illustrative example of a study that uses evolutionary models to constrain the BH population in a globular cluster that we have already discussed is Kremer et al. (2018b). In this study, the authors use the CMC code (Rodriguez et al. 2022) to simulate the evolution of the cluster NGC 3201 from a grid of models with different initial conditions specifically chosen to reproduce the observed properties of NGC 3201. The key parameter varied in this grid is the scaling factor applied to the natal kicks of the BHs, effectively setting the BH retention for each model. By comparing the evolution of just the core radius and half-light radius the authors are already able to show that the models with more BHs retained by the present day are better able to match the observed properties of the real cluster. Kremer et al. (2018b) then compare the surface brightness profiles of the model snapshots from each model in the grid to the observed profile to show that only the models with more than 200 BHs retained by the present day are able to match the data. We show this result in Figure 1.4 to illustrate the power of this technique. This is a result of the dynamical effects we have discussed above, namely that the presence of BHs in the core of the cluster prevents the visible stars from concentrating in the core, leading to a shallower SBP.




Figure 1.4: Figure 2 of Kremer et al. (2018b). Plotted in yellow is the surface brightness profile of NGC 3201 as measured by Trager et al. (1995). The coloured lines show the surface brightness profiles of the models from the grid of Kremer et al. (2018b). The models with more than 200 BHs are plotted in black while the models with fewer than 10 BHs are plotted in blue.

1.4 Summary and Goals of this Thesis

In this chapter we have examined the various methods for searching for and modelling BHs in GCs. We have seen how the theoretical predictions have evolved such that we now expect many GCs to harbour non-zero populations of BHs. On the observational side, we have looked at the challenges of confidently identifying BHs through their X-ray and radio emission and the challenges of finding binary systems containing BHs through radial velocity measurements. Critically, we have emphasized the fact that the number of observable BHs in a cluster is not necessarily representative of the total number of BHs in a cluster, making inference about the total number of BHs in a cluster very challenging. On the modelling side, we have looked at how the presence of BHs affects the overall dynamics of a cluster and how we can use these dynamical signatures to infer the presence of BHs and the size of the BH population in a cluster, and we will examine many of these methods in Chapter 3.

The goals of this thesis lie in improving the state of the modelling side of the field. In Chapter

2, we will test the ability of multimass distribution function models to infer the BH content of clusters using mock data and identify areas where these models might struggle. We will discuss the other methods used to dynamically infer the BH content of clusters in Chapter 3, focusing on their strengths and weaknesses. In Chapters 4 and 5 we will present a new method to incorporate pulsar-timing based constraints for these models and apply it to two Milky Way clusters with large pulsar populations. Through the work presented in this thesis we will show that the BH population of Milky Way GCs *can* be constrained through careful dynamical modelling when combined with a variety of observational constraints. We summarize our findings and briefly discuss the future directions of these methods in Chapter 6.

Chapter 2

Multimass modelling of globular clusters and their black hole populations: method validation with mock data

Section 2.2 below is reproduced from Section 3 of the following paper:

Multimass modelling of Milky Way globular clusters – II. present-day black hole populations

N. Dickson, **P. J. Smith**, V. Hénault-Brunet, M. Gieles, H. Baumgardt

Originally published in Monthly Notices of the Royal Astronomical Society, Volume 529, Issue 1, Pages 331-347.

I led and performed the analysis and wrote the entirety of the text for this section of the original publication. In this version we additionally include Figures 2.1, 2.2 and 2.3 which were not included in the original publication for brevity. Section 2.1 is a supplementary background discussion to provide context for the rest of the chapter.

2.1 Background

As part of the work presented in Dickson et al. (2024), I led the validation of the methodology used to infer the BH populations of Milky Way GCs. This involved applying the method to synthetic observations extracted from Monte Carlo models of GCs with known BH populations. The results of this validation are contained in Section 3 of Dickson et al. (2024) and reproduced in this chapter as Section 2.2.

In this study, multimass LIMEPY models are fit to a variety of datasets (number density profiles, line-of-sight velocity dispersion profiles, radial and tangential proper motion dispersion profiles as well as stellar mass function data) from a collection of Milky Way GCs. Most studies that fit evolutionary models to data (e.g. Rui et al. 2021b) typically only look at a few datasets, usually the surface brightness profile (SBP) and a combined (both line-of-sight and proper motions) velocity dispersion profile (VDP). Because the LIMEPY models are fast to compute, we are able to vary the (initial) mass function of the models, meaning they not only fit to the SBPs and VDPs, but also the radially varying stellar mass functions (Baumgardt et al. 2023) of the clusters. This allows us to constrain the present-day stellar mass function for each cluster individually, rather than assuming a canonical universal IMF, evolved to the present day. We then use these models to infer posterior probability distributions for all model parameters as well as derived quantities like the BH (and other stellar remnant) content for the clusters in the sample. The fact that this method is able to recover posterior probability distributions for the BH populations is an advantage over most studies that use evolutionary models (discussed in more detail in Section 3.1.1) which typically only produce point estimates for the BH population or a rough estimate of the uncertainty using the handful of closest matching models.

The LIMEPY models used in this work are a set of distribution function-based equilibrium models that are isothermal near the cluster centre and described by polytropes in the outer re-

gions near the escape energy . The models have been extensively tested against N -body models (Zocchi et al. 2016; Peuten et al. 2017) and their multimass version is able to effectively reproduce the effects of mass segregation. We further describe the models and their parameters in Section 4.3. These models work by solving the Poisson equation for a given distribution function, effectively giving the mass profiles and velocity dispersion profiles of each mass component. This allows for straightforward comparisons to a variety of observables, including observables that trace different mass components. The typical process of inferring the BH population of a cluster using these models relies on the fact that the total mass of a cluster is well constrained by the kinematics of the cluster and the visible mass in turn can be constrained by observations of the stellar mass function and surface brightness profile. This means that with sufficient kinematic and stellar mass function data it is possible to accurately constrain the amount of dark mass in a cluster. The inference of the composition of the dark mass relies on the ability of these models to freely vary the BH retention, varying it such that it best reproduces the observed properties of the cluster. Because of the dynamical coupling of the BHs in the core of the cluster with the rest of the cluster, there is limited freedom to arrange the dark mass within the cluster while still reproducing properties like the degree of mass segregation among visible stars, allowing us to infer how much of the dark mass is in the form of BHs.

These models have typically been coupled with the `SSPTOOLS` package (first presented by Balbinot & Gieles 2018, updated in Dickson et al. 2023) in order to produce realistic present-day mass functions, populated with an accurate mix of WDs, NSs and BHs (see Dickson et al. 2023). The `evolve_mf` algorithm of `SSPTOOLS` approximates the evolution of a stellar population using analytic fits to full stellar evolution models, allowing for the fast computation of the present-day mass function from a given initial mass function. This present-day mass function is used as the input mass spectrum for the `LIMEPY` models, allowing for the models to match the present-day mass function of individual clusters rather than adopting a single canonical IMF.

2.2 Validation of BH population inference.

In order to test the reliability of our method in inferring BH populations, we first apply it to simulated observations from Monte Carlo models with known BH populations.

In order to explore a number of models with similar properties as real Milky Way clusters, we select snapshots from the existing grid of Cluster Monte Carlo (CMC; Rodriguez et al. 2022) models presented in Kremer et al. (2020)¹. We select the snapshots using the same methodology as Rui et al. (2021b) which we briefly summarize here.

The selections are based on the SBPs of Trager et al. (1995) and the velocity dispersion profiles (VDP) compiled by Baumgardt & Sollima (2017) and Baumgardt & Hilker (2018)². We search for snapshots that are a good match to any clusters from the Harris (1996, 2010 edition) catalogue present in both the VDP and SBP compilations, leaving us with about 100 clusters to match to snapshots. We first use the metallicities from Harris (1996, 2010 edition) and the present-day galactocentric radii from Baumgardt et al. (2019a) to select the subset of models which are closest to the true values for each cluster.

From this subset, we then search every model for snapshots that match suitably well to a given cluster’s observed SBP and VDP simultaneously³. In order to select a snapshot, we adopt a threshold of $s \equiv \max(\tilde{\chi}_{\text{SBP}}^2, \tilde{\chi}_{\text{VDP}}^2) < 10$ for the “fitting heuristic” s of Rui et al. (2021b), which describes the goodness-of-fit of a snapshot based on the $\tilde{\chi}^2$ statistic between the observations and the interpolated model profiles. We have found that a threshold of $s < 10$ provides an acceptable balance between the number and fit quality of the retained snapshots. While a number of snapshots passing this filter have an apparently poor overall fit to one or both of the observational profiles, we opt to still include these snapshots in the sample, as our goal

¹Available at <https://cmc.ciera.northwestern.edu/home>

²Available at <https://people.smp.uq.edu.au/HolgerBaumgardt/globular>

³See Rui et al. (2021b) for details on how the model profiles are extracted from the CMC snapshots for comparison with the observations.

is not to select only snapshots which are perfect matches to specific real clusters but instead to build a sample of snapshots that are qualitatively similar to the Milky Way clusters examined in this work. For each cluster covered by the observational datasets we select the single best-fitting snapshot, where one exists.

2.2.1 Mock observations

The search described above results in a sample of 53 CMC model snapshots, representative of Milky Way GCs. From these we next extract synthetic observations designed to emulate the real observational data used to constrain the models examined in this work.

We place each cluster at its respective heliocentric distance as reported by Baumgardt & Vasiliev (2021) and then use the CMCTOOLKIT library (Rui et al. 2021b) to calculate projected positions and velocities as well as simulated photometry for objects in each snapshot.

2.2.1.1 Number density profiles

We extract projected number density profiles from the models, designed to emulate those of de Boer et al. (2019). We select all stars brighter than *Gaia* $G = 20$, sort them into 50 radial bins, with equal numbers of stars, and calculate the number density in each radial bin. All densities are assigned a Poisson counting error. The de Boer et al. (2019) profiles are combinations of *Gaia* star counts in the outer regions and *HST* and archival SBPs in the inner regions, where crowding becomes an issue, however we find that using the same $G < 20$ cut over the entire radial extent of the cluster results in well-sampled profiles which cover a similar radial extent and have similar uncertainties to the de Boer et al. (2019) profiles.

2.2.1.2 Proper motion dispersion profiles

We extract two sets of PM dispersion profiles for each snapshot, in order to represent the performance of the two different sources of PM observations used.

In the inner regions, we seek to emulate the performance of the *HST* based PM dispersion profiles of Libralato et al. (2022). We select stars within the central $100''$ of the cluster to mimic the footprint of an *HST* ACS image, and limit our selection to stars within $15 < V < 18$. We split the stars into radial bins containing at least 120 stars each, up to a maximum of five bins. This provides sufficient radial coverage of the cluster while still allowing us to construct profiles for distant clusters, where limited numbers of stars pass the magnitude cut. We assume a typical uncertainty of 0.1 mas yr^{-1} on all stars. Within each bin we compute the mean velocity and velocity dispersion along with their associated uncertainties, assuming the velocities are drawn from a Gaussian error distribution, using MCMC. This is repeated for both the radial and tangential components of PM. The median and 1σ values of the dispersion in each bin are used going forward.

In the outer regions, we seek to emulate the *Gaia* DR3 based profiles of Vasiliev & Baumgardt (2021). We base our magnitude cuts on their profiles, selecting all stars in the $13 < G < 19$ range outside of the innermost $100''$, to avoid overlapping with the *HST* profiles. We assign each star an uncertainty in proper motion based on its *G* band magnitude using the calibrations provided in Table 4 of Lindegren et al. (2021), allowing us to replicate the performance of the *Gaia* DR3 catalogue. We again bin the stars using the same conditions as in the inner *HST* profiles, and calculate the velocity dispersion in each bin using the same method, again for both radial and tangential components.

2.2.1.3 Line-of-sight velocity dispersion profiles

In addition to the PM dispersion profiles, we also extract LOS velocity dispersion profiles, designed to emulate those presented by Baumgardt & Sollima (2017) and Baumgardt & Hilker (2018). As the compilation of velocity dispersions that make up these profiles consist of several different inhomogeneous datasets, with varying precisions, we adopt the simplifying assumption of a typical uncertainty of 1 km s^{-1} on all observed stars. We limit this dataset to only giants brighter than $V = 17$, which is typical of the datasets used in the observed compilations. We again sort the stars into several radial bins, requiring at least 70 stars per bin, up to a maximum of 10 bins, and compute the velocity dispersion for each in the same way as for the PM profiles.

2.2.1.4 Stellar mass functions

We extract stellar mass function data for each of our snapshots, designed to emulate the datasets presented in Baumgardt et al. (2023). These datasets consist of star counts, binned by stellar mass, extracted from archival *HST* observations. While the real *HST* fields are distributed somewhat randomly around each GC according to the various goals of each proposal for which they were originally observed, in general there is typically at least one exposure centred on the cluster centre and a number of fields placed outside of the central region. For simplicity, we opt to place one field over the centre of the cluster, covering a range of $0' - 1.6'$ in projected angular separation from the centre, and two outer annuli at radial distances of $2.5'$ and $5'$, each sized such that they cover the same area as the central field. The central field is split into 4 annuli, giving us a total of 6 annular fields covering the central, intermediate and outer regions of the cluster. In many of the more well-studied clusters in our sample, such as NGC 104 and ω Cen, the large number of observed *HST* fields actually provides much better coverage of the clusters than our fields simulated here, and thus the results of this section may actually be conservative. In each of these fields we extract stellar counts separately, in bins of stellar mass with widths of $0.1 M_{\odot}$.

In the real datasets, the faintest stars (lowest stellar masses) for which stellar counts can be extracted with reasonable completeness (> 90 per cent) is a function of crowding. To replicate this effect in our synthetic mass functions we construct an empirical relation between the surface number density and the lowest observable mass within a field. We use NGC 104 as the basis for this calibration because it covers a wide range of number densities from its core to the outskirts and has a large number of *HST* fields for which the mass function was extracted. We extract star counts in each field, down to a lower mass limit calculated by the above relation and up to the main sequence turn-off, replicating the performance of the observed *HST* stellar counts.

We assign Poisson counting errors to our stellar counts, though to reflect the scatter we typically see in the real data we also inflate these errors by a factor of $F = 3$ and re-sample each point within the errors, resulting in mass functions that are very similar to those of Baumgardt et al. (2023) (see Section 4.1.1.3 of Dickson et al. (2023) for a description of this F parameter and its motivation).

2.2.2 Validation results

After extracting the synthetic datasets, we then directly apply our model fitting method (as described in Section 2 of Dickson et al. (2024) and Dickson et al. (2023)) and compare the resulting inferred BH mass fractions to the known BH population in the CMC models from which the mock data was extracted. We show an example fit to the mock data in Figures 2.1 and 2.2.

As in Dickson et al. (2023), we discard any obviously poor fits. We also discard any snapshots which have much smaller datasets (mostly consisting of snapshots matched to very distant clusters). This leaves us with 44 final snapshots with datasets of similar quality to the Milky Way clusters we study in this work, and with model fits which satisfyingly reproduce the mock observations and recover the various cluster parameters well, such as the total mass, which we generally recover within ~ 10 per cent and the half-mass radius, which we generally recover




Figure 2.1: Example fit to mock observations of a CMC snapshot matched to NGC 6256. *Top Left*: The number density profile of the snapshot (blue points) and the best-fitting model (orange line). *Top Right*: The line-of-sight velocity dispersion profile of the snapshot (blue points) and the best-fitting model (orange line). *Bottom Left*: The radial proper motion dispersion profile of the snapshot (blue points) and the best-fitting model (orange line). *Bottom Right*: The tangential proper motion dispersion profile of the snapshot (blue points) and the best-fitting model (orange line). The shaded regions represent the 1σ uncertainties on the model profiles. This figure was not included in the original publication.

within uncertainties (we show the mass and radius recovery in Figure 2.3).

Our inferred values of f_{BH} , compared to the true values for our collection of snapshots, are shown in Figures 2.4 (right panel), 2.5, and Table 2.1.

While in general our fits satisfyingly recover the mass fraction in BHs, we find a number of snapshots for which we underpredict the black hole mass fraction. A common feature of these problematic snapshots is their dynamical age. All of the mock clusters for which we significantly underpredict f_{BH} have an age that is less than 3 times their present-day half-mass relaxation time. In other words, they are dynamically very young. There are also, however, many dynamically young clusters for which we correctly recover f_{BH} .

To understand this behaviour, we plot in the left panel of Figure 2.4 all of the mock clusters used in the validation of our models on the $\delta - N_{\text{relax}}$ plane (where N_{relax} is the ratio of cluster age to present-day half-mass relaxation time), coloured by the distance of their median inferred




Figure 2.2: Example fit to mock mass function data of a CMC snapshot matched to NGC 6256. Each panel is a different radial annulus of the cluster. The blue points are the simulated observations and the orange line is the best-fitting model. The shaded regions represent the 1σ and 2σ uncertainties on the model mass functions, respectively. This figure was not included in the original publication.




Figure 2.3: *Left*: The true vs inferred mass for each snapshot in our sample. *Right*: The true vs inferred half-mass radius for each snapshot in our sample. In both figures, the dashed line represents the one-to-one relation. The error bars represent the 1σ statistical uncertainties on the inferred values. This figure was not included in the original publication.

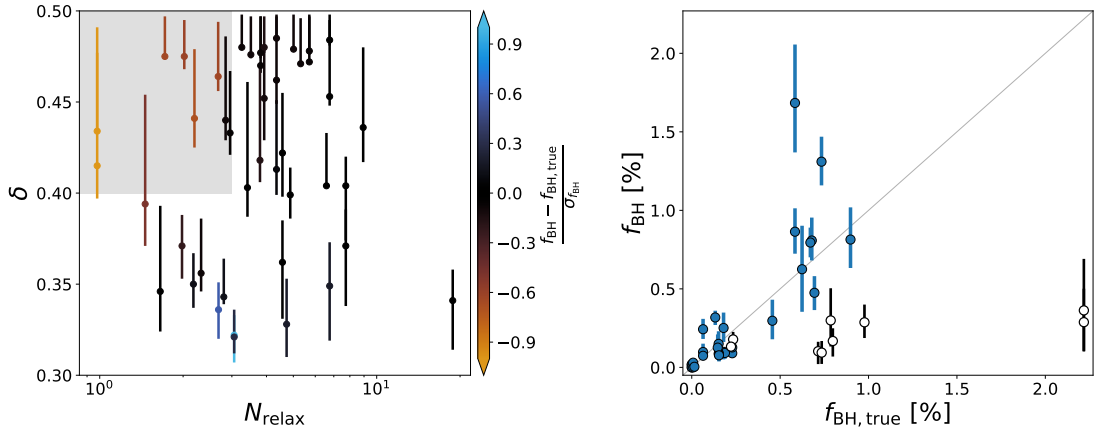


Figure 2.4: Left panel: Validation snapshots plotted in the $\delta - N_{\text{relax}}$ plane. A selection of dynamically young clusters with high inferred values of δ is shown by the grey shaded region ($N_{\text{relax}} < 3$, $\delta > 0.4$). The points are coloured based on the number of $\sigma_{f_{\text{BH}}}$ the median inferred f_{BH} is away from the true value (where $\sigma_{f_{\text{BH}}}$ is the width of the posterior for f_{BH}). Under this colour scheme, snapshots for which we underpredict (overpredict) f_{BH} are increasingly orange (blue). Right panel: The f_{BH} values inferred based on the multimass model fits to mock observations extracted from CMC models, against the true values in those validation snapshots ($f_{\text{BH,true}}$). Snapshots which fall within the shaded region in the left panel are plotted with unfilled markers in the right panel, demonstrating that they also largely represent the snapshots for which we significantly underpredict f_{BH} .

f_{BH} value from the true value. As δ is a measure of the degree of energy equipartition in our multimass models (with typical values of $\delta \approx 0.5$ for mass-segregated, dynamically evolved clusters), we should normally expect dynamically young clusters ($\lesssim 3 N_{\text{relax}}$) to have lower values of δ . The mock clusters for which we significantly underpredict f_{BH} are clearly concentrated in the upper-left corner of the $\delta - N_{\text{relax}}$ plane, highlighted with the grey-shaded rectangle in the left panel of Figure 2.4. Their value of δ is inconsistent with our expectations for these systems.

Previous comparisons of the multimass LIMEPY models used here with N -body models (Peuten et al. 2017) have confirmed this intuition that dynamically young clusters should have lower values of δ (with $\delta \sim 0.35$) and have also shown that models with any significant black hole population should have similarly low values of δ until they eject most of their black holes (after several relaxation times). Therefore, we should not expect to find any clusters in the upper-left or lower-right sections of the $\delta - N_{\text{relax}}$ plane and any model fits which fall into these regions are worthy of some suspicion, even though it is not immediately clear why the data for some of these dynamically young clusters would prefer models with values of δ close to 0.5. Simple test fits with δ forced to a lower value ($\delta \leq 0.35$) have shown that a lower value of δ can result in an increase in the inferred f_{BH} for these snapshots, as generally expected.

Within our validation sample there are a handful of problematic clusters in the upper-left, suspect region of the $\delta - N_{\text{relax}}$ plane. In the right panel of Figure 2.4, we show our inferred values of f_{BH} compared to the true values for our collection of snapshots, with the snapshots that fall into this problematic region plotted with unfilled markers. It can be clearly seen that, if these clusters are ignored, the correlation between the inferred and true black hole mass fraction is even stronger. In Figure 2.5, we plot the inferred values of f_{BH} compared to the true values but now ignoring these problematic snapshots and showing the 1σ statistical error bars (in blue) for the retained mock clusters.

With these caveats in mind we can turn to the sample of real clusters we investigate in this

work to look for clusters that fall into this suspect region. In the real sample only three clusters out of the 34 fall into this region of relatively high δ and low N_{relax} (compared to 9 of the 44 validation snapshots). The relative lack of problematic fits in the real sample compared to the mock sample highlights the fact that our mock sample is not a perfect imitation of the population of MW clusters we investigate in this work and is limited by the constraints of the CMC grid. One of these problematic clusters is NGC 3201, which is discussed further below (Section 5.3.4 of Dickson et al. 2024). The other two are NGC 104 (47 Tuc) and NGC 5139 (ω Cen), two of the most massive clusters in our sample, which still possess notable BH populations despite their high δ . Our results for these two well-studied clusters concurs well with other literature values (see the discussion of each in Section 5.3 of Dickson et al. 2024). One real cluster in our sample falls into the lower-right portion of $\delta - N_{\text{relax}}$. This is NGC 6624, a core-collapsed cluster which is discussed further in Section 4.1 of Dickson et al. (2024), and for which we argue that we incorrectly infer a non-zero number of black holes.

As may be expected, the statistical uncertainties derived solely from the fitting procedure seem to underestimate the real uncertainties slightly. Our fitting procedure operates under the assumption that our models are a perfect representation of the data, and as such, in reality, may underestimate the true errors. It has been shown that multimass DF models, such as those used here, may underestimate the uncertainties when compared to more flexible models, such as Jeans models (Hénault-Brunet et al. 2019), which could be indicative of systematic errors not captured in the statistical uncertainties and limitations in the ability of these models to perfectly reproduce the data.

In an attempt to quantify this underestimation, we search for the factor by which the statistical uncertainties on our inferred values of f_{BH} need to be inflated to make them fully consistent with a one-to-one relation with the true model values. We define a nuisance parameter ϵ by which we multiply the statistical errors $\sigma_{f_{\text{BH}}}$ to determine the total, inflated error Δ on each




Figure 2.5: The f_{BH} values inferred based on the mock observations extracted from CMC models, against the true values in those models ($f_{\text{BH,true}}$). The one-to-one line is shown in grey, representing perfect agreement. The median and 1σ values, based solely on the statistical uncertainties from the fit, are shown in blue. The inflated errors, based on the nuisance parameter ϵ ($\Delta = \epsilon \sigma_{f_{\text{BH}}}$), are shown in orange.

inferred f_{BH} . We then fit a fixed, one-to-one (i.e. slope of 1, intercept of 0) line through the points in Figure 2.5, assuming a Gaussian likelihood and allowing ϵ to vary freely and inflate the uncertainties on the inferred f_{BH} . This fit results in a value of $\epsilon = 2.5^{+0.3}_{-0.3}$. We also show the inflated 1σ error bars on the inferred f_{BH} in Figure 2.5 (orange), demonstrating the additional uncertainty needed to make our inferred values consistent with the true values and illustrating the typical uncertainties expected when applying our method to real data.

Overall, this comparison with mock observations extracted from dynamical simulations lends confidence in the ability of our methods to correctly recover the mass fraction in BHs in GCs, with the important note that the statistical uncertainties on these inferred values may underestimate the true uncertainties by up to a factor of ~ 2.5 , and with a reminder that dynamically young clusters with inferred values of $\delta \sim 0.5$ should be treated with caution.

2.2.3 Appendix: Validation results

CMC Model	Snapshot	Closest Milky Way Analogue	$f_{\text{BH,true}} [\%]$	$f_{\text{BH}} [\%]$
N4e5-rv1-rg2-Z0.0002	248	NGC 6558	0.0	$0.000^{+0.003}_{-0.000}$
N4e5-rv1-rg8-Z0.0002	217	NGC 5946	0.008	0^{+0}_{-0}
N4e5-rv1-rg8-Z0.0002	217	NGC 5986	0.008	$0.03^{+0.02}_{-0.02}$
N4e5-rv1-rg8-Z0.0002	252	NGC 6544	0.015	$0.004^{+0.007}_{-0.004}$
N4e5-rv2-rg2-Z0.002	165	NGC 6352	0.152	$0.15^{+0.08}_{-0.07}$
N4e5-rv2-rg8-Z0.0002*	160	NGC 6584	0.222	$0.13^{+0.06}_{-0.05}$
N4e5-rv2-rg8-Z0.0002	165	NGC 6981	0.147	$0.13^{+0.08}_{-0.07}$
N4e5-rv4-rg8-Z0.0002*	128	NGC 5897	0.786	$0.3^{+0.2}_{-0.2}$
N4e5-rv4-rg8-Z0.0002	134	NGC 288	0.624	$0.6^{+0.3}_{-0.3}$
N8e5-rv0.5-rg2-Z0.0002	608	NGC 6355	0.0	$0.019^{+0.011}_{-0.008}$
N8e5-rv0.5-rg2-Z0.0002	608	NGC 6681	0.0	$0.02^{+0.01}_{-0.01}$
N8e5-rv0.5-rg2-Z0.002	516	NGC 6624	0.0	0^{+0}_{-0}
N8e5-rv1-rg2-Z0.0002	375	NGC 6293	0.064	$0.10^{+0.05}_{-0.03}$
N8e5-rv1-rg2-Z0.0002	375	NGC 6453	0.064	$0.07^{+0.01}_{-0.02}$
N8e5-rv1-rg2-Z0.0002	375	NGC 6522	0.064	$0.24^{+0.06}_{-0.06}$
N8e5-rv1-rg2-Z0.002	340	NGC 6316	0.229	$0.09^{+0.02}_{-0.03}$
N8e5-rv1-rg2-Z0.002	340	NGC 6539	0.229	$0.13^{+0.03}_{-0.03}$
N8e5-rv1-rg2-Z0.002	345	NGC 6569	0.18	$0.09^{+0.03}_{-0.04}$
N8e5-rv1-rg2-Z0.002	345	NGC 6637	0.18	$0.10^{+0.03}_{-0.03}$
N8e5-rv1-rg2-Z0.002	345	NGC 6642	0.18	$0.25^{+0.10}_{-0.09}$
N8e5-rv1-rg2-Z0.002	354	NGC 6638	0.179	$0.09^{+0.02}_{-0.03}$
N8e5-rv1-rg2-Z0.002	358	NGC 6401	0.175	$0.09^{+0.02}_{-0.02}$
N8e5-rv1-rg2-Z0.002	362	NGC 6256	0.155	$0.09^{+0.03}_{-0.02}$
N8e5-rv1-rg2-Z0.002	362	NGC 6304	0.155	$0.09^{+0.02}_{-0.02}$
N8e5-rv1-rg20-Z0.0002	394	NGC 1904	0.154	$0.08^{+0.02}_{-0.04}$
N8e5-rv1-rg8-Z0.0002	379	NGC 6284	0.188	$0.09^{+0.03}_{-0.04}$
N8e5-rv1-rg8-Z0.0002	386	NGC 6779	0.153	$0.08^{+0.03}_{-0.04}$
N8e5-rv1-rg8-Z0.0002	386	NGC 6934	0.153	$0.09^{+0.03}_{-0.04}$
N8e5-rv1-rg8-Z0.002*	366	NGC 6760	0.234	$0.18^{+0.05}_{-0.04}$
N8e5-rv2-rg2-Z0.0002	315	NGC 6218	0.456	$0.3^{+0.1}_{-0.1}$
N8e5-rv2-rg2-Z0.002*	273	NGC 6723	0.977	$0.3^{+0.1}_{-0.1}$
N8e5-rv2-rg2-Z0.002*	286	NGC 6712	0.798	$0.17^{+0.08}_{-0.10}$
N8e5-rv2-rg20-Z0.0002*	296	NGC 1261	0.735	$0.09^{+0.07}_{-0.07}$
N8e5-rv2-rg8-Z0.0002	299	NGC 3201	0.694	$0.5^{+0.1}_{-0.1}$
N8e5-rv4-rg8-Z0.0002*	194	NGC 4372	2.218	$0.3^{+0.2}_{-0.2}$
N8e5-rv4-rg8-Z0.0002*	194	NGC 6101	2.218	$0.4^{+0.3}_{-0.3}$
N1.6e6-rv0.5-rg2-Z0.002	875	NGC 6440	0.133	$0.32^{+0.04}_{-0.05}$
N1.6e6-rv1-rg2-Z0.0002	694	NGC 6402	0.898	$0.8^{+0.2}_{-0.2}$
N1.6e6-rv1-rg2-Z0.0002	727	NGC 6656	0.734	$1.3^{+0.2}_{-0.2}$
N1.6e6-rv1-rg2-Z0.0002	737	NGC 6541	0.679	$0.8^{+0.1}_{-0.1}$
N1.6e6-rv1-rg2-Z0.0002	755	NGC 6333	0.584	$0.9^{+0.1}_{-0.1}$
N1.6e6-rv1-rg2-Z0.0002	755	NGC 6342	0.584	$1.7^{+0.4}_{-0.3}$
N1.6e6-rv1-rg8-Z0.002*	698	NGC 6356	0.715	$0.10^{+0.06}_{-0.08}$

Table 2.1: Results of the validation fits. The “True” model values ($f_{\text{BH,true}}$) and median and 1σ credibility intervals inferred by the fits to mock observations ($f_{\text{BH,infer}}$) are shown. The model names (as given in the CMC cluster catalogue) indicate the initial number of particles, virial radii, galactocentric radii and metallicity of each CMC model. The Milky Way cluster to which each snapshot was matched is also given. All dynamically young snapshots which fall into the suspect region are denoted by an asterisk.

Chapter 3

Discussion of other methods for BH population inference

Having now discussed in detail the use of multimass distribution function models to infer the presence of BHs in clusters, it is worth discussing the wide array of other methods that have been used to study the BH populations of GCs. This is a discussion chapter in which we will not present any new results, but rather place the results of studies that use other methods in the context of the results presented in Chapter 2 and focus on their respective strengths and weaknesses.

3.1 Modelling of individual clusters

The most straightforward way to use dynamical models to constrain the BH population in a given GC (and indeed the method used in Chapter 2) is to construct an accurate model of the cluster using the available data (e.g. kinematics, mass functions, and number/surface brightness profiles) and then use the model to infer the BH population. This can be accomplished by either directly fitting an equilibrium model to data or by computing a grid of evolutionary models and

then selecting the model from the grid that best matches the data. While the former approach allows for the use of Bayesian inference schemes, it is only feasible for models that are fast to compute, meaning that the grid-based approach is often the only feasible option for evolutionary models.

3.1.1 Evolutionary models

Evolutionary models are models that can simulate the dynamical evolution of a GC from initial conditions to the present day. These models typically incorporate stellar evolution codes such that each object in the cluster is evolved according to its mass, eventually forming a stellar remnant (for more massive stars). We have already discussed the work of Kremer et al. (2018b) which employs this methodology to study the BH population of NGC 3201.

Similar work was done by Giersz & Heggie (2011), where the authors use a Monte Carlo code to simulate the evolution of the cluster 47 Tuc from a grid of models with varying initial conditions. The authors vary several parameters in their grid, including the initial mass function and binary content of the cluster, attempting to match the observed properties of 47 Tuc. Unlike Kremer et al. (2018b), where the models were only compared to the SBP and VDP of the cluster, Giersz & Heggie (2011) also compare to measurements of the stellar luminosity function (a proxy for the stellar mass function) as well as confirming that accelerations of pulsars derived from pulsar timing are consistent with the maximum allowed acceleration in their models. The authors found that simultaneously fitting the surface brightness profile, and pulsar accelerations was the most important constraint on the models, with the final best-fitting model having just 19 BHs retained by the present day. We will further investigate the potential that pulsar timing measurements have for dynamical modelling of GCs in Chapters 4 and 5.

The primary limitation of this approach is the computational cost of running the models. The days-long runtimes of these models mean that it is not feasible to run a grid with more than a few parameters varied and even then the grid must be relatively sparse. This typically means that the stellar initial mass function (IMF) of the cluster is held constant, greatly reducing the freedom of the models to vary the population of heavy remnants (WDs, NSs, BHs). Recent results (Baumgardt et al. 2023; Dickson et al. 2023) have shown that the IMF of GCs is likely to be more bottom-light (i.e. deficient in low-mass stars) than the canonical Kroupa (2001) IMF that is typically assumed in these models. For a cluster of a given total mass, a Kroupa IMF would produce fewer high-mass stars and thus fewer BHs than a bottom-light IMF. This illustrates the importance of the choices made in the initial conditions of these models and highlights the utility of fast-to-compute equilibrium models we tested in Chapter 2 where the remnant content (including BHs) can be varied with more flexibility.

A more recent example of a study that used a very similar methodology is Ye et al. (2022). In this study, the authors again use the CMC code to simulate the evolution of an individual cluster (47 Tuc) from a grid of initial conditions, though they also present a new binary formation mechanism that they implement in the code. This binary formation mechanism implements a more realistic treatment of tidal forces during close encounters and allows for more realistic simulations of dynamical binary formation. The authors use the surface brightness profile and the velocity dispersion profile to match the grid of models to the data much like Kremer et al. (2018b) did, however they also use the large population of exotic objects (millisecond pulsars, cataclysmic variables, low-mass X-ray binaries) in 47 Tuc as a further constraint on the models showing that the new binary formation mechanism is critical for reproducing the neutron star binaries that are responsible for most of the exotic objects in 47 Tuc. Like the work of Kremer et al. (2018b), this method is still limited by the fact that it only produces a single point-estimate for the BH population of the cluster.

A related work that we discussed in Chapter 2 is that of Rui et al. (2021b), who provided a formal description for the snapshot fitting method used by Kremer et al. (2018b) and also formed the basis for the fitting performed by Ye et al. (2022). The authors use this method in combination with a large grid of pre-computed CMC models (Kremer et al. 2020) to take a slightly different approach to the problem of fitting models to data. Rather than focusing on a single cluster, the authors instead took the grid of ~ 150 models and search for *any cluster* for which the SBP and VDP match well with an existing snapshot. This method makes use of the existing grid of models computed for other work (we will examine the method of Weatherford et al. 2018, 2020, another work which uses this pre-computed grid of CMC models in Section 3.2) to greatly reduce the computational cost, however it is reliant on the grid of models being sufficiently dense in parameter space to contain good representations of individual Milky Way clusters. Once the best-fitting snapshots for a given Milky Way cluster are identified, the authors then calculate the median BH number from all snapshots that pass their goodness-of-fit criterion. This does allow the authors to report a range of possible values for the BH population, however the method is still limited by the low number of well-fitting snapshots that exist for any given cluster. Unlike the method we described and tested in Chapter 2, this method does not attempt to match the stellar mass function of individual clusters, meaning that a snapshot identified as a good fit may not actually be a good fit to all available data.

3.1.2 Equilibrium models

In contrast to evolutionary models, equilibrium dynamical models are models that capture the phase-space distribution of stars and stellar remnants in a GC at a single point in time (e.g. Gieles & Zocchi 2015). We introduced and applied these models in Chapter 2 and discuss them in the context of other modelling techniques here. Equilibrium models are typically much faster to compute than evolutionary models meaning that it is possible to directly fit the model to

the data using Bayesian inference techniques (techniques like MCMC and Nested Sampling require anywhere from a few thousands of model evaluations to a few million). The drawback of these models in comparison to evolutionary models is that they do not model star-by-star dynamics, and they do not model the evolution of the cluster over time. This means that they often neglect the effects of more complicated objects like binary systems and can encounter issues at very small scales in the cores of clusters where continuous mass profiles are no longer a good approximation of the discrete mass distribution of centrally concentrated heavy remnants (Hénault-Brunet et al. 2019). State-of-the-art equilibrium models incorporate a full mass spectrum and some, like multimass distribution function based models, are able to fully reproduce the effects of mass segregation. Many of the recent studies (including the work presented in this thesis) that have used equilibrium models to constrain the BH population in clusters have used the (multimass) LIMEPY models (Gieles & Zocchi 2015).

Multimass LIMEPY models have also been used in several studies that sought to investigate the presence of an intermediate-mass BH (IMBH) in several clusters, a slightly different but related problem to the one we are focusing on in this thesis (discussed in more detail in Section 5.2). The first of the studies we will discuss is the work of Zocchi et al. (2019) which used multimass LIMEPY models to investigate a reported IMBH in a Milky Way GC. This particular study focuses on ω Cen, the most massive GC in the Milky Way and a cluster that is still too large to be properly modelled with star-by-star N -body simulations. Zocchi et al. (2019) use a simplified two-component mass spectrum, with one component representing the stars and low-mass remnants and the other representing the BHs. Even with this simplified mass spectrum the authors find that their models are able to reproduce the observed kinematics of the cluster without the need for an IMBH. The authors additionally show that velocity anisotropy is a key ingredient when modelling GCs, with anisotropic models being able to better reproduce the

observed kinematics than isotropic models. This echoes the results of Zocchi et al. (2017) which found that radial anisotropy can produce some of the same kinematic signatures that are often associated with the presence of an IMBH like a higher velocity dispersion near the centre of the cluster.

The work of Gieles et al. (2018) was also motivated by a reported IMBH in a Milky Way GC, this time NGC 6624. In this study the authors use a full mass spectrum in their modelling. Because NGC 6624 is a core-collapsed cluster, the authors fix the BH content of their models to zero. They show that the observed kinematics of the cluster can be fully explained with their models, with no need for an IMBH or any significant population of stellar-mass BHs. The collection of datasets used in this work goes beyond the usual surface brightness profiles and velocity dispersion profiles and includes stellar mass function data and pulsar timing measurements.

The methodology of Dickson et al. (2023, 2024) is essentially a refinement of the methodology presented in Hénault-Brunet et al. (2020), applied to a large sample of clusters. To summarize the methodology (see Chapter 2 for more details), these studies fit multimass LIMEPY models to the VDPs, number density profiles and stellar mass functions of 34 Milky Way clusters. These models, populated with stellar mass functions computed with the `evolve_mf` algorithm of `SSPTOOLS`, allow for the inference of not just the BH population, but also the present-day mass function of each cluster. This ability to vary and infer the present-day mass function of each cluster individually is one of the biggest advantages of this method over the evolutionary models discussed in the previous section.

Another flavour of equilibrium models that we should mention are the “Jeans Models”, which are based on the Jeans equations that describe the motion of a collection of stars in a gravitational field (e.g. Watkins et al. 2013; Sollima et al. 2016; Kamann et al. 2016; Vitral & Mamon 2021; Vitral et al. 2023). These models are typically fit to the density and velocity dispersion profiles of

a tracer population and vary the mass profile of the model to match the observed kinematics. They make no assumptions about the shape of the underlying mass profile, allowing for very flexible though sometimes unphysical solutions for the mass and mass-to-light ratio profiles (e.g. Hénault-Brunet et al. 2019).

A recent example of a study that uses these models to investigate the dark remnant population of a GC is Vitral & Mamon (2021), who use the MAMPOSST-PM Jeans model to investigate the dark remnant population of NGC 6397. They fit a few variations of their models to the observed surface brightness profile and velocity dispersion profile: one with no dark central component, one with a so-called “cluster of unseen objects” (CUO), one with a central IMBH and one with both an IMBH and a CUO. They find that the models with either an IMBH or a CUO are best able to match the data, with the model with no dark central component clearly underpredicting the central velocity dispersion. These authors report that the models with a diffuse CUO produce fits with higher likelihoods and Bayesian evidence than the models with an IMBH, concluding that the data is most consistent with a CUO of $1000 - 2000 M_{\odot}$. The authors initially interpret this as evidence for a large population of stellar-mass BHs in the cluster, making up the majority of the CUO. As was pointed out by Rui et al. (2021a), this interpretation is extremely unlikely due to the core-collapsed nature of NGC 6397 (see also discussion in Section 1.3.1), arguing instead that the CUO is likely made up of primarily WDs.

The ambiguity surrounding the nature of the CUO in NGC 6397 is a good example of the limitations of Jeans models. These models are very flexible and can fit a wide range of mass profiles to the data, however distinguishing between different types of dark objects is challenging. In follow-up work, Vitral et al. (2022) and Vitral et al. (2023) use their Jeans models in combination with CMC models to first detect and then characterize the dark mass in NGC 3201, NGC 6397 and M4. In these studies, the data are first fit with the Jeans models to detect the presence of a dark component before the best-matching CMC snapshots are examined to identify the make-up

of the dark component. For NGC 3201 and NGC 6397, they are able to show that the dark mass is likely made up of a population of BHs for NGC 3201 and a population of WDs for NGC 6397. For M4, the authors find that CUO inferred from the Jeans models is more massive than the small population of BHs that are present in the best-matching CMC models but more compact than a cluster of WDs, leading them to speculate that this cluster may host an unusually compact sub-cluster of BHs or even an IMBH. These studies highlight the fact that while Jeans models are excellent tools for detecting the presence of dark mass in a cluster, they are less well-suited to characterizing the nature of the dark mass. Indeed, the final results of these works could be in part reproduced without the use of Jeans models, by simply analyzing the properties of the best-matching CMC snapshots directly.

3.2 Scaling relations

The other main way that dynamical models are used to constrain BH populations in GCs is by using large populations of models to find scaling relations between observable quantities like the central velocity dispersion or degree of mass segregation of a cluster and its BH content. This method is typically applied to evolutionary models, for which individual modelling of many clusters would be too expensive. In this section we will look at two sets of studies that each use large grids of Monte Carlo models to derive relations between observables and the BH population of a cluster. The first set of studies is that of Arca Sedda et al. (2018) and Askar et al. (2018), which use the MOnte Carlo Cluster simulAtor (MOCCA) code (e.g. Hypki & Giersz 2013; Giersz et al. 2013) and associated MOCCA database (Askar et al. 2017) to identify and then apply relations between observables and the BH content of a cluster. The second set of studies is that of Weatherford et al. (2018) and Weatherford et al. (2020). These studies use the grid of CMC models presented in Kremer et al. (2020) and follow the same pattern of first iden-

tifying relations between observables and the BH population and then applying these relations to a large sample of clusters. We will examine both of these sets of studies in detail focusing on their respective strengths and areas for improvement.

3.2.1 MOCCA

It is useful to first give a brief overview of the MOCCA database and the models it contains. The entire MOCCA database contains more than 2000 models, in which the metallicity, number of stars, central concentration, tidal radius, half-mass radius, binary fraction and galactocentric radius are all varied across the initial conditions. Of all the models in the database, almost 1300 survive to 12 Gyr and of these models, 172 have at least 10 BHs remaining at 12 Gyr. It is these 172 models that the authors use to derive their scaling relations.

This particular method focuses specifically on the BH subsystem (BHS), the centrally concentrated grouping of BHs that forms in the core of a cluster due to mass segregation (e.g. Breen & Heggie 2013). There are essentially two parts to the correlations that they observe in their models. The first is that the various properties of the BHS are correlated with each other, with relations between the mass of the BHS and the radius of the BHS, the number of BHs and the mass of the BHS, the mean mass of the BHs and the radius of the BHS and finally the density of the BHS and its radius all showing relatively tight groupings. What these relations mean is that if we can find a relation between just *one* property of the BHS and an observable quantity, we should in theory be able to infer the rest of the properties of the BHS, allowing us to constrain the BH content of a given cluster.

The second part of the correlations that they observe allows them to do exactly this. They find correlations of varying strength between the density of the BHS and the following observational properties: observational half-mass radius ($r_{h,obs}^2$), average surface luminosity, central velocity




Figure 3.1: Bottom panel of figure 10 of Arca Sedda et al. (2018). The density of the BHS as a function of the average surface luminosity. The points are coloured by the mean mass of the BHs in the BHS.

dispersion and total luminosity (L). While individually each of these relations is not particularly strong, the authors are able to combine the various observational quantities into a tighter relation by defining the “average surface luminosity” as $L/r_{h,obs}^2$ which is tightly correlated with the density of the BHS. We show this final relation in Figure 3.1. It is this relation that the authors ultimately apply in Askar et al. (2018) to infer the BH content of 29 MW GCs.

Askar et al. (2018) begin their study by selecting 29 Milky Way clusters that match the properties of the MOCCA models the relations were derived from. They have four criteria: they require the galactocentric radius to be less than 17 kpc, the V-band absolute magnitude to be brighter than $M_V = -6.5$, the average surface luminosity must be less than $10^4 L_{\odot}/pc^2$ and finally the half-mass relaxation time must be longer than 0.9 Gyr. These criteria ensure that they do not extrapolate their relations beyond what the underlying models cover.

After selecting appropriate clusters, they then straightforwardly apply their scaling relations to first infer the density of the BHS from the average surface luminosity and then the size and mass of the BHS by applying the further relations. This method ultimately allows the authors to infer the density, size, total mass and number of BHs in the BHS for each cluster from easily observable quantities, namely the total luminosity and the half-light radius.

Having outlined the methodology, we now turn to the shortcomings of this study. The first is simply an artifact of the fact that the relations were derived from models that all contained at least some BHs. The authors themselves point out in Askar et al. (2018) that models with very few or no BHs occupy the same region of the average surface luminosity - total luminosity plane as models with many BHs. These models are not included in the relations, and are largely resultant from models where there is no fallback during BH formation, resulting in very high natal kicks and very low BH retention. This means that if a Milky Way cluster has a very small BH population, it very well might be placed in the same region of the average surface luminosity - total luminosity plane as a model with a much higher BH population. While the current state of SNe modelling suggests that BHs do in fact experience fallback and thus lowered natal kicks (e.g Fryer et al. 2012), this does still add uncertainty for models with very low BH populations which likely are very similar to models that lost all of their BHs to natal kicks. This is especially concerning given that Askar et al. (2018) state that roughly 20% of the low-BH models that they exclude are models with the normal fallback and natal kick prescriptions, highlighting the fact that this relation is not as robust as it may seem.

The other issue with this particular methodology was pointed out in Weatherford et al. (2020) and concerns the formulation of the scaling relations themselves. Arca Sedda et al. (2018) use linear relations in log-log space to describe their various observed correlations. Weatherford et al. (2020) point out that the assumption of linear relations will inherently lead to biases, and this

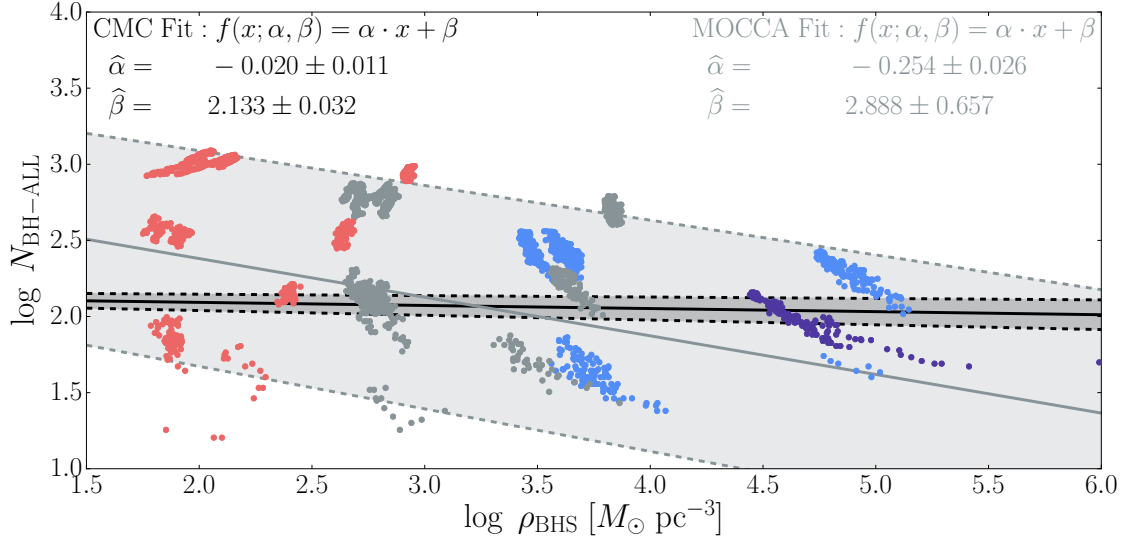


Figure 3.2: Panel e of Figure 8 of Weatherford et al. (2020). The density of the BHS as a function of the number of BHs in the BHS. The points are coloured by the initial virial radius of the CMC models. The light grey line and wide grey band are the linear fit and 1σ scatter of the chained relations from Arca Sedda et al. (2018). The dark grey line and narrow grey band are the resulting fit when bypassing the intermediate steps and directly fitting the number of BHs to the density of the BHS.

problem is only compounded when multiple linear relations are combined. They illustrate this issue by showing that while each of the individual relations are recovered with similar values in their grid of CMC models, skipping the intermediate steps and directly plotting the density of the BHS against the number of BHs in the BHS shows that even in the MOCCA grid, the correlation is quite weak. The range in the number of BHs in the BHS for a given density spans two orders of magnitude while in the CMC grid there appears to be no real relation of any significance. We show this in Figure 3.2.

This lack of a significant correlation between the density of the BHS and the number of BHs in the BHS draws into question the accuracy of the inferred BH populations derived in Askar et al. (2018) but also serves as motivation for the works of Weatherford et al. (2018) and Weatherford et al. (2020) which we discuss next.

3.2.2 CMC

The works of Weatherford et al. (2018) and Weatherford et al. (2020) are conceptually very similar to the works of Arca Sedda et al. (2018) and Askar et al. (2018). The main differences between the two sets of studies are the models used and the ways in which the relations are formulated. Firstly, the works of Weatherford et al. (2018) and Weatherford et al. (2020) use the CMC code and its associated grid of models instead of the MOCCA database. We have already discussed the CMC grid briefly in our discussion of targeted modelling of individual clusters however it is worth going over some parameters of the grid. While the MOCCA grid varies seven parameters when generating initial conditions, the CMC grid varies only four: the number of stars, the metallicity, and the virial and galactocentric radii. This means that, in theory, the CMC grid covers less of the possible parameter space than the MOCCA grid, however the CMC grid has still been shown to be able to reproduce the properties of a wide range of Milky Way clusters (Kremer et al. 2020).

Weatherford et al. (2018) and Weatherford et al. (2020) use the degree of mass segregation as their observable, building on the theoretical work that suggests that mass segregation is suppressed by the formation of a BH subsystem in the core of a cluster (e.g. MacKey et al. 2008). This is a similar process to the one we discussed when examining the effect of a BHS on the SBP of a cluster. The BHS causes the centrally concentrated stars to be ejected back out into the outer regions of the cluster, effectively suppressing the process of mass segregation. This means that a suitably designed measure of mass segregation should in theory correlate with the number of BHs in a cluster. Weatherford et al. (2018) formulate an observable that is specifically designed to be easy to extract from observations and to tightly correlate with the number of BHs in a cluster. They define two quantities, Δ_A and Δ_{r50} , which are two different ways to quantify the differences between the cumulative radial distribution functions (CDFs) of two populations of

stars in a cluster. The first, Δ_A , is the area between the two CDFs while the second, Δ_{r50} , is the difference between the projected distances of the medians of the two CDFs. The two populations of stars are chosen such that they are both easy to distinguish and different in mass.

Weatherford et al. (2020) tweak the method for better performance before applying it to 50 Milky Way GCs. The primary change they make is to switch from using two populations to using three populations made up solely of main sequence stars. This change was motivated by the fact that post main sequence stars are comparatively much rarer than main sequence stars and thus have noisier CDFs. The three main sequence populations are then used in pairings to calculate both Δ_A and Δ_{r50} , resulting in 6 measurements of the degree of mass segregation (3 sets of pairing for each of the two measures). The key statistical improvement the authors make is to use a non-parametric kernel density estimator (KDE) to map the degree of mass segregation to the number of BHs in the cluster. This avoids the need to make any assumptions about the functional form of the relation between the degree of mass segregation and the number of BHs, avoiding any potential biases. When constructing the KDE, the authors restrict themselves to the models that are closest in metallicity to the cluster of interest (the CMC grid has 3 different values of metallicity). They then construct a multidimensional KDE relating the Δ parameters to the number of BHs in a cluster. In order to calculate the probability density function (PDF) of the number of BHs in a cluster, they then simply slice the KDE along the value of the degree of mass segregation for the cluster of interest. We show the relation between the two measures of mass segregation and the number of BHs in Figure 3.3.

Weatherford et al. (2020) apply this method to 50 Milky Way GCs, deriving PDFs for the number of BHs in each cluster. They compare their results to the results of Askar et al. (2018) and find that they systematically predict smaller BH populations than Askar et al. (2018). Weatherford et al. (2020) attribute this to the bias towards larger BH populations in the MOCCA relations

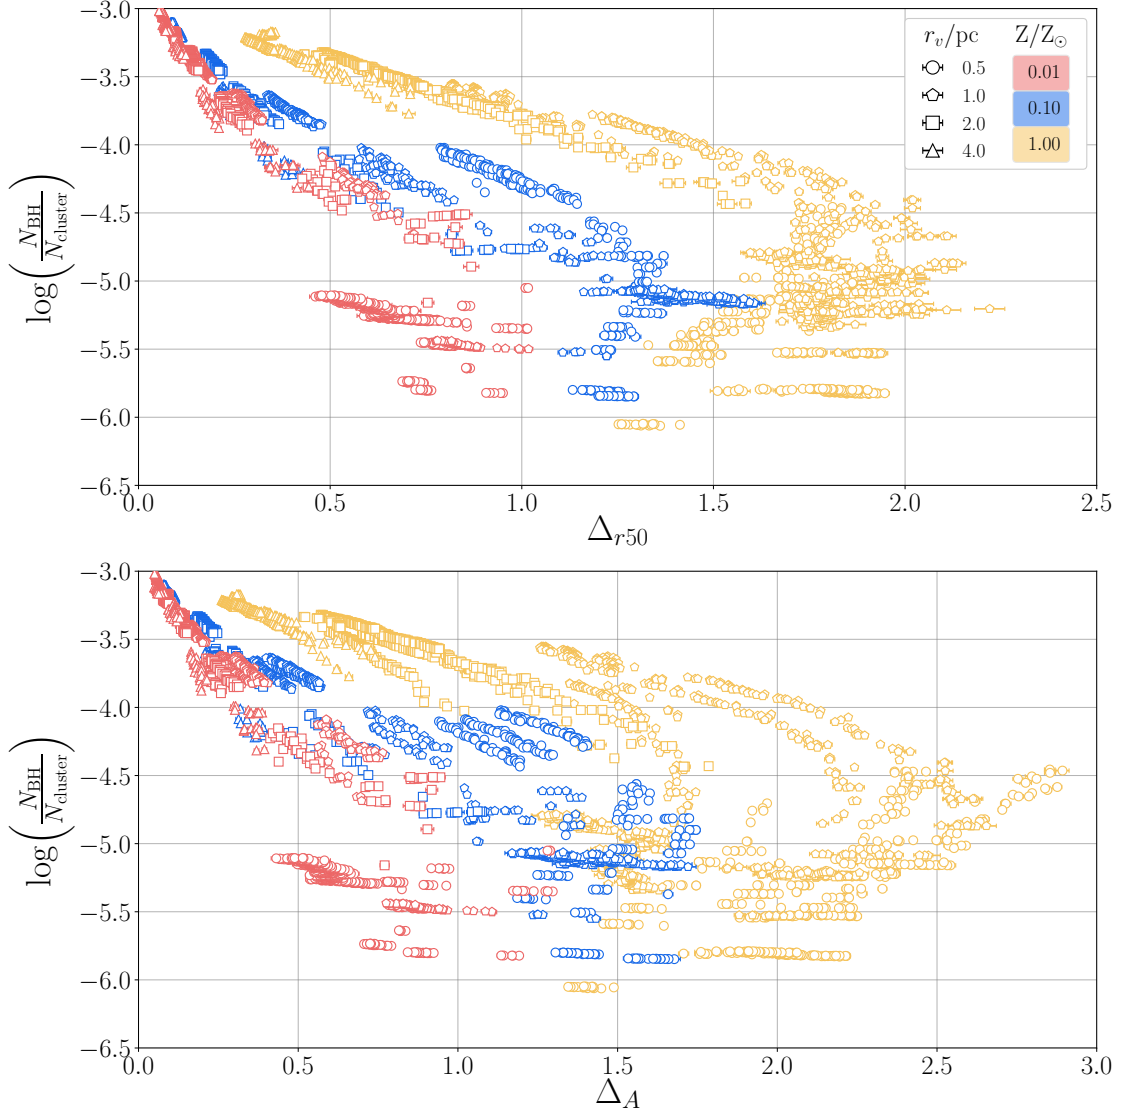


Figure 3.3: Figure 2 of Weatherford et al. (2020). The number of BHs in various CMC snapshots as a function of the two mass segregation parameters. The top panel shows the number of BHs as a function of Δ_{r50} while the bottom panel shows the number of BHs as a function of Δ_A . The points are coloured by the metallicity of the snapshots and the markers show the initial virial radii of the snapshots.

that we discussed earlier. They find that 35 of their 50 clusters likely contain more than 20 BHs, with 8 of these clusters likely containing more than 80 BHs at the present day, adding further support to the idea that many Milky Way GCs have retained a large population of BHs by the present day.

The primary limitation of this method is the fact that it is still reliant on a grid of models that is relatively sparse. The limitation of only having three different metallicities is particularly apparent when examining Figure 3.3, where gaps between the metallicity tracks are clearly visible. The other major limitation is the fact that this method is not particularly precise, with the PDFs for the number of BHs in a cluster typically spanning a very large range of possible values. Finally, the fact that this method relies on a single observable, and does not take into account the kinematics, density profile or stellar mass function of the cluster means that it is not making full use of the available data, potentially leading to lower precision or accuracy in the inferred BH populations. Overall, the work of Weatherford et al. (2020) represents a meaningful improvement over the work of Askar et al. (2018) while still being conceptually quite similar. The study presented the largest sample of BHS estimates to date and also provided a robust methodology for deriving these BHS estimates.

3.3 Comparison of inferred black hole populations

Having demonstrated in Chapter 2 that multimass LIMEPY models are able to successfully infer the BH populations of Milky Way GCs (with some caveats), and having discussed the strengths and weaknesses of the other methods used to infer BH populations, it is worth discussing the discrepancies between these methods. In Figure 3.4 we show the inferred BH populations for a selection of Milky Way GCs from Dickson et al. (2024) as well as the results of Weatherford et al. (2020), Askar et al. (2018), and Rui et al. (2021b). Comparing the results of these studies,

we see that there is good agreement about the BH content of Milky Way GCs on a population level, with all studies in agreement that most clusters seem to have between 0 and 1% of their mass in the form of BHs by the present day. Comparing individual clusters however, we see that there is often a significant discrepancy between the inferred BH populations of the different studies. We show these discrepancies in a different form in Figure 3.5, which shows the error-weighted residuals between the inferred values of f_{BH} from Dickson et al. (2024) and the values from Askar et al. (2018) and Weatherford et al. (2020). In this representation, points falling within the ± 1 range indicate that the two studies are in agreement within their 1σ uncertainties. Examining this plot in detail, not only do we see that the various studies disagree by more than 1σ for many clusters, we also see a trend where the inferred values from Dickson et al. (2024) are lower than the values from the literature for clusters where they infer a low value of f_{BH} and higher for clusters where they infer a high value of f_{BH} . This suggests that the disagreement between the studies is not simply due to random errors but rather due to systematic differences in the methodologies used to infer the BH populations of the clusters.

Given that we have shown in Chapter 2 that the inferred BH populations from multimass LIMEPY models do not seem to be systematically biased as a function of f_{BH} , these differences suggest that the other methods may be underpredicting the variance in f_{BH} across the Milky Way GC population.

As we discussed previously, a test of the methodology used by Askar et al. (2018) was presented by Weatherford et al. (2020) where they applied the methodology of Askar et al. (2018) to a grid of CMC models. The testing of both the methodology used in Askar et al. (2018) by Weatherford et al. (2018) and the methodology we use in Dickson et al. (2024) presented in Chapter 2 has in both cases revealed important caveats and sources of error that were not immediately apparent from simply examining the results of each study. In particular, the exclusion of models with very low BH populations in Askar et al. (2018) when deriving their scaling rela-

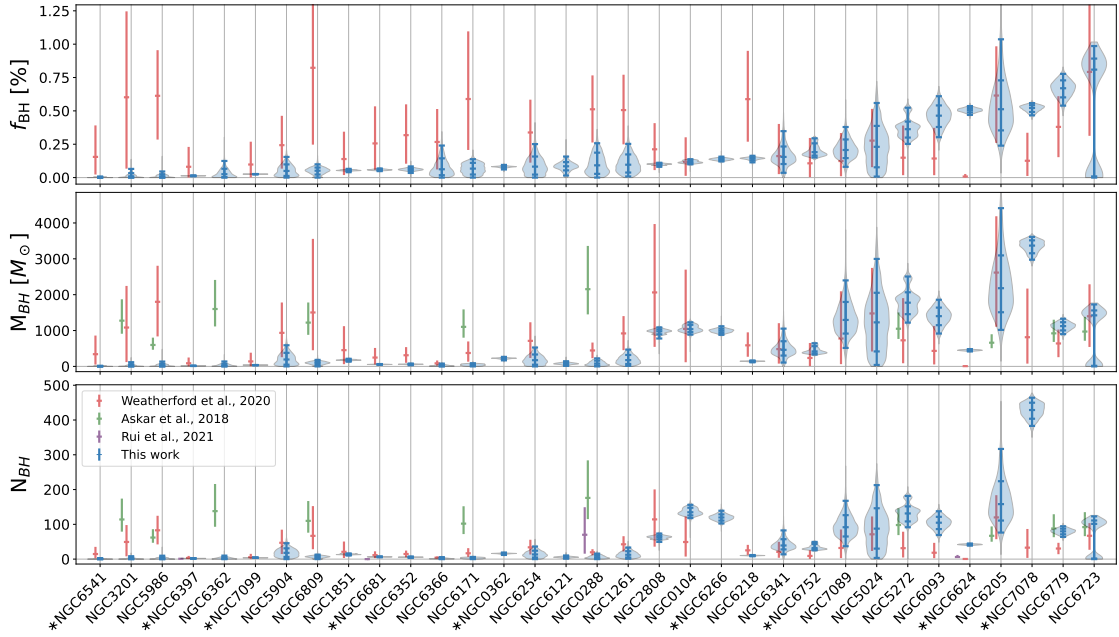


Figure 3.4: Figure 3 of Dickson et al. (2024). Comparison of the inferred BH populations for a selection of Milky Way GCs from three modelling studies discussed in this chapter. The red points are the results of Weatherford et al. (2020), the green points are the results of Askar et al. (2018), the light blue points are the results of Rui et al. (2021b) while the violin plots are the results from Dickson et al. (2024). The top panel is BH mass fraction, the middle is total mass in BHs and the bottom panel is the total number of BHs.

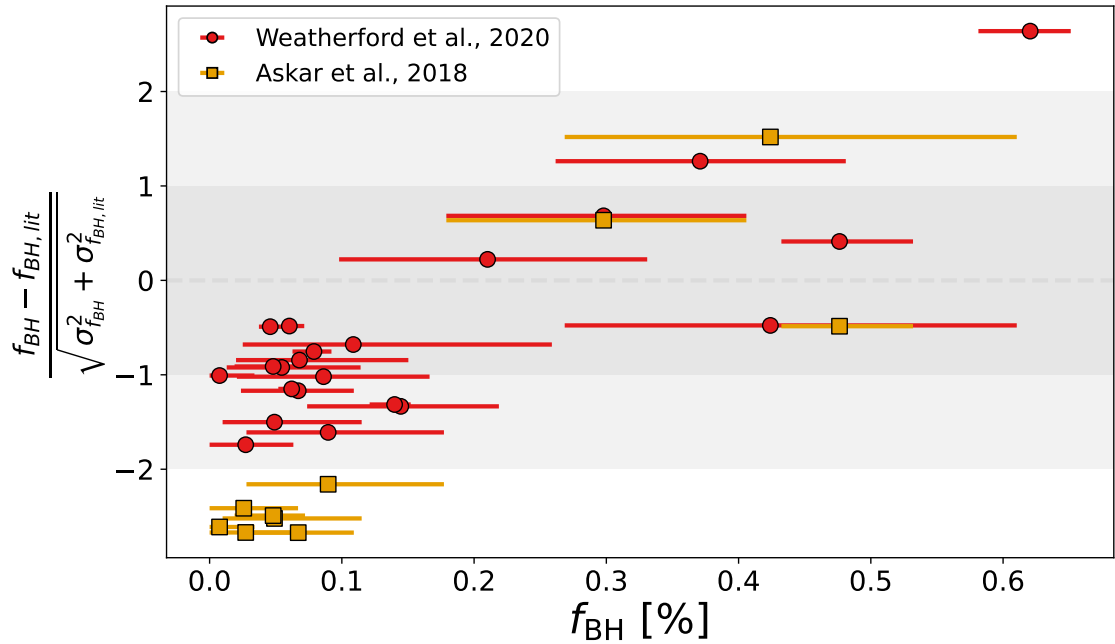


Figure 3.5: Figure 8 from Dickson et al. (2024). Error-weighted residuals between the inferred value of f_{BH} from Dickson et al. (2024) and the values from two literature sources, Askar et al. (2018) and Weatherford et al. (2020). In general, we see a trend where the inferred values from Dickson et al. (2024) are lower than the values from the literature for clusters where they infer a low value of f_{BH} and higher for clusters where they infer a high value of f_{BH} .

tions would bias their predictions towards higher BH populations, which is one way to explain the systematic differences seen in Figure 3.5. Looking specifically at the results of Weatherford et al. (2020), we see that they generally report larger uncertainties on their inferred BH populations than the other studies. When considering the likely underestimation of the uncertainties on the results of Dickson et al. (2024) (see discussion in Section 2.2), it is possible that these two studies are actually in agreement for most clusters and the discrepancies we see are simply due to the poorly understood systematic sources of error in both studies. Given the results of these tests, it is critical that future studies that seek to infer the BH populations of Milky Way GCs do similar testing to ensure that their results are robust and reliable, ideally validating their methods using mock data generated from models unrelated to those that they are testing.

Chapter 4

Probing populations of dark stellar remnants in the globular clusters 47

Tuc and Terzan 5 using pulsar timing:

Introduction, Data and Methods

This chapter contains the Introduction, Data and Methods sections (sections 1, 2 and 3) of the following manuscript submitted to the Astrophysical Journal for publication:

Probing populations of dark stellar remnants in the globular clusters 47 Tuc and Terzan 5 using pulsar timing

P. J. Smith, V. Hénault-Brunet, N. Dickson, M. Gieles, H. Baumgardt

I am the lead author of this manuscript and led the analysis and writing of the paper. I produced all figures and tables and wrote the text.

Abstract

We present a new method to combine multimass equilibrium dynamical models and pulsar timing data to constrain the mass distribution and remnant populations of Milky Way globular clusters (GCs). We first apply this method to 47 Tuc, a cluster for which there exists an abundance of stellar kinematic data and which is also host to a large population of millisecond pulsars. We demonstrate that the pulsar timing data allow us to place strong constraints on the overall mass distribution and remnant populations even without fitting on stellar kinematics. Our models favor a small population of stellar-mass BHs in this cluster (with a total mass of $446^{+75}_{-72} M_{\odot}$), arguing against the need for a large ($> 2000 M_{\odot}$) central intermediate-mass black hole. We then apply the method to Terzan 5, a heavily obscured bulge cluster which hosts the largest population of millisecond pulsars of any Milky Way GC and for which the collection of conventional stellar kinematic data is very limited. We improve existing constraints on the mass distribution and structural parameters of this cluster and place stringent constraints on its black hole content, finding an upper limit on the mass in BHs of $\sim 4000 M_{\odot}$. This method allows us to probe the central dynamics of GCs even in the absence of stellar kinematic data and can be easily applied to other GCs with pulsar timing data, for which datasets will continue to grow with the next generation of radio telescopes.

4.1 Introduction

Pulsars have a long history of being used to investigate the mass distribution of globular clusters (GCs). Early work from Phinney (1992, 1993) examined the effect of a pulsar’s surroundings on its measured spin (\dot{P}) and orbital period derivatives (\dot{P}_b ; for pulsars in binary systems), including quantifying the effects of the cluster potential, Galactic potential, proper motion and intrinsic effects like magnetic braking. Recently, several works have presented detailed analyses of pul-

sar data for probing the gravitational potential of GCs (see e.g., Prager et al. 2017 for Terzan 5, Freire et al. 2017 and Abbate et al. 2018, 2019a for 47 Tuc, Gieles et al. 2018 for NGC 6624, Abbate et al. 2019b for M62, Corongiu et al. (2024) for NGC 6752, and Bañares-Hernández et al. 2024 for NGC 5139).

Pulsars indeed present a unique opportunity to probe the central dynamics of these systems, especially when crowding and extinction make it challenging to obtain detailed stellar kinematic data in their central regions. Because of their extremely stable periods (spin and orbital), measured changes in the periods of pulsars beyond their (unknown) intrinsic spin down due to magnetic braking can be almost entirely attributed to external factors. By performing timing measurements over long time scales and precisely measuring the changes in their periods, we can learn about the host potential of the pulsars. In particular, the observed period derivatives due to the changing ‘Doppler shift’ from the line-of-sight gravitational acceleration felt by pulsars in GCs allow us to constrain the gravitational potential and mass distribution of GCs hosting pulsars.

Most of the studies mentioned above used single-mass dynamical models, without a mass spectrum¹. Therefore, they cannot capture the effect of mass segregation, which is affected by the presence/absence of central black holes (e.g. Merritt et al. 2004; MacKey et al. 2008; Gill et al. 2008; Peuten et al. 2016, 2017; Weatherford et al. 2018). These studies also usually focus on fitting models to the pulsar data, ignoring the velocity dispersion profiles, surface density profiles and stellar mass functions which are typically used to constrain mass models of GCs². As an example, Freire et al. (2017) presented a single-mass King model of 47 Tuc which was compared to the pulsar acceleration data as well as the measured “jerk” (the time derivative of

¹With the exception of Gieles et al. (2018), who compared the observed period derivatives of pulsars in NGC 6624 to predictions from multimass models but did not directly fit these models to the pulsar data.

²Although see Bañares-Hernández et al. (2024) for an application of joint modeling of pulsar timing data and stellar kinematics in a Jeans analysis of NGC 5139. Their work, submitted shortly after ours, is based on a similar methodology in many aspects, but also adopts different models and assumptions while targeting a different cluster, and is therefore complementary.

the acceleration) of the pulsars. Their model does not contain an intermediate-mass black hole (IMBH) but is still able to account for all of the pulsars' period derivatives as well as the inferred jerks for all the pulsars within the cluster core. Corongiu et al. (2024) presented similar work for NGC 6752, using the first and second period derivatives of five pulsars in this cluster to infer the presence of non-luminous matter in its core and investigate the possibility of that mass residing in an IMBH.

In this work, we present new self-consistent multimass models of 47 Tuc and Terzan 5 that are fitted both to traditional observables (velocity dispersion profiles, number density profile, local stellar mass function observations) and to the variety of pulsar timing data available for these clusters. The direct inclusion of the pulsar data in the likelihood function allows us to revisit and address the previous claim from Kızıltan et al. (2017a) that the pulsar accelerations favor a large central mass in 47 Tuc in the form of an IMBH. We use updated stellar mass function data where available, and adopt the latest constraints on the distance to our clusters from *Gaia* (Baumgardt & Vasiliev 2021) as a prior on the distance parameter in our models. To properly model the effect and constrain the size of a possible population of stellar-mass BHs, our multimass models use stellar evolution recipes based on recent prescriptions for the masses and natal kicks of BHs, and can therefore include realistic BH populations. While our models do not explicitly allow for an IMBH, given that the effects of a central IMBH and a large population of centrally concentrated stellar-mass BHs are expected to be similar beyond the sphere of influence of the IMBH, especially when the mass fraction of the cluster in BHs is small (e.g. Aros & Vesperini 2023), a best-fitting model without a large population of BHs would provide evidence that an IMBH is not required to explain the dynamics of the cluster.

Since we are interested in accurately modeling the present-day mass distribution within these clusters, equilibrium distribution function-based models are a good choice. Compared to the more computationally expensive evolutionary models like Monte-Carlo or N -body models

(which are limited to relatively small grids), our multimass models offer much-increased flexibility to vary the cluster’s structural properties, stellar mass function, and population of dark remnants (including BHs) at a small fraction of the computational cost.

Both 47 Tuc and Terzan 5 are host to large populations of pulsars. 47 Tuc is a well-studied cluster with a great deal of stellar kinematic data, making it an ideal candidate to test our method. Terzan 5 is a heavily obscured bulge cluster which is host to the largest population of millisecond pulsars of any Milky Way GC and which is located in the bulge, making the collection of conventional stellar kinematic data extremely challenging. The combination of the large pulsar population and the lack of traditional stellar kinematic data means that our method is particularly well-suited to studying the central dynamics of this cluster. This large population of pulsars may be partially explained by the cluster’s collision rate, which is the highest of any Milky Way GC (e.g. Lanzoni et al. 2010). The central dynamics of this cluster are therefore of great interest.

The remainder of this paper is structured as follows. In Section 4.2 we describe the data to which we fit our models. Section 4.3 describes the models, fitting procedure and individual likelihoods. In Section 5.1 we present our fits and discuss our results. In Section 5.2 we discuss the implications of our results and compare them to other studies. We finally summarize our findings in Section 5.3.

4.2 Data

For this study, we use the same data that was used by Dickson et al. (2023, 2024), with the addition of the pulsar data described below and some additional datasets for Terzan 5. We summarize the data below.

4.2.1 Kinematics and density profiles

4.2.1.1 Proper motion dispersion profiles

We use both *Hubble Space Telescope* (*HST*) proper motion data and *Gaia* DR3 proper motions to constrain the kinematics of the clusters. For 47 Tuc, there are proper motion measurements from *HST* that cover the inner regions of the cluster. These data were presented in Libralato et al. (2022) and are split into radial and tangential dispersion profiles, allowing some leverage on the velocity anisotropy of the cluster.

For both clusters, we also use *Gaia* DR3 proper motion dispersion profiles which are based on the membership catalogs presented in Vasiliev & Baumgardt (2021). These profiles are split into radial and tangential components for 47 Tuc where there is an abundance of high-probability members (the radial and tangential profiles were derived in Dickson et al. 2023), but left as total proper motion ($\mu_{\text{tot}}^2 = \mu_{\alpha^*}^2 + \mu_{\delta}^2$) for Terzan 5 where isolating high-probability members is more difficult due to bulge contamination.

4.2.1.2 Line-of-sight velocity dispersion profiles

We use the line-of-sight velocity dispersion profiles from Baumgardt & Hilker (2018) to further constrain the kinematics of the clusters. These dispersion profiles are based on archival spectra obtained at the European Southern Observatory’s (ESO) Very Large Telescope (VLT) and the Keck observatory, supplemented with published radial velocity data from the literature from Baumgardt (2017).

For 47 Tuc, we additionally use the line-of-sight dispersion profile presented by Kamann et al. (2018), who used the MUSE spectrograph (Bacon et al. 2010) to collect data for 22 GCs.

As these radial velocity samples are dominated by bright stars, we assume that these velocity dispersion profiles trace the kinematics of upper main-sequence and evolved stars (which we

assume trace the kinematics of giants) in our models.

4.2.1.3 Number density profiles

We use the number density profiles from de Boer et al. (2019) and Lanzoni et al. (2010) to constrain the size and structural parameters of the clusters. The de Boer et al. (2019) profile is made up of a combination of the number density profile of cluster members based on *Gaia* DR2 data in the outer regions and a surface brightness profile from Trager et al. (1995) in the central regions, which is matched to the *Gaia* number density profile in the region where the two profiles overlap.

Terzan 5 is not included in the compilation of de Boer et al. (2019) due its position in the bulge, where crowding and extinction are significant issues. Because of these challenges, we use the number density profile from Lanzoni et al. (2010) which is based on a combination of data from the *HST*, the Multi-conjugate Adaptive optics Demonstrator (MAD) on VLT and the Two Micron All Sky Survey (2MASS) which combine to cover the entire radial extent of the cluster.

The *Gaia* data only includes bright stars ($m > 0.6 M_{\odot}$) and the *HST* and ground-based data are likewise dominated by bright stars, so we assume that these number density profiles trace the distribution of upper main-sequence and evolved stars in our models.

Note that there are two density profiles we could have chosen from for Terzan 5: the surface brightness profile of Trager et al. (1995) or the number density profile of Lanzoni et al. (2010). In their region of overlap, these profiles do not match very well with each other, even when scaled vertically. The Lanzoni et al. (2010) profile decreases faster in the outer regions compared to the Trager et al. (1995) profile. We opted to use the Lanzoni et al. (2010) profile because it is based on *HST* data and modern ground-based data which should help to more reliably subtract bulge contamination, and also because it has well-defined uncertainties. We also found that tests with the surface brightness profile of Trager et al. (1995) resulted in best-fitting models where

the surface brightness profile and the other datasets and profiles could not be simultaneously reproduced as well as when using the Lanzoni et al. (2010) number density profile.

4.2.2 Stellar mass functions

As a constraint on the global present-day stellar mass function of 47 Tuc, we use the completeness-corrected stellar mass function data that was derived from archival *HST* photometry by Baumgardt et al. (2023). These are based on various archival *HST* images (20 different pointings for 47 Tuc, proposal IDs shown in Figure 5.2) from which stellar number counts were derived as a function of magnitude and projected distance from the cluster center and were then converted into stellar mass functions through isochrone fits. For 47 Tuc, there are extensive observations which cover stars within a mass range of $\sim 0.1 - 0.8 M_{\odot}$ as well as a radial range of 0-40 arcminutes from the cluster center. The large span of radii and stellar masses allows us to constrain the varying local stellar mass function as a function of distance from the cluster center, and therefore the degree of mass segregation in the cluster.

The stellar mass function for Terzan 5 was derived in the same way as for 47 Tuc, but was not included in the compilation of Baumgardt et al. (2023) due to limitations with the data resulting from the cluster's position in the bulge. For this cluster, we have a single mass function field, in the infrared (filters F110W and F160W of *HST* proposal 12933, PI: Ferraro) which covers a region from 0.6-1.6 arcminutes from the cluster center and a mass range of $\sim 0.6 - 0.9 M_{\odot}$. While this dataset provides weaker constraints on the model than the mass function data for 47 Tuc, it is still useful for constraining the amount of visible mass in the cluster around its half-mass radius and verifying that the assumed global stellar mass function in our model of Terzan 5 is reasonable.

4.2.3 Pulsar data

For 47 Tuc, we use timing solutions from Freire et al. (2017), Ridolfi et al. (2016) and Freire & Ridolfi (2018) which include both the spin and orbital periods (the latter when applicable, for pulsars in binaries) and their time derivatives. For Terzan 5, we use the timing solutions presented in Lyne et al. (2000), Ransom et al. (2005), Prager et al. (2017), Cadelano et al. (2018), Andersen & Ransom (2018), Ridolfi et al. (2021) and Padmanabh et al. (2024), again including both the spin and orbital periods and their derivatives, where available.

Pulsars with non-degenerate companions are classified as either ‘black-widow’ or ‘redback’ systems, where black widows have companions with masses less than $\sim 0.1 M_{\odot}$ and redbacks have more massive companions (Roberts 2012). All redbacks and some black widows (e.g. Shaifullah et al. 2016) display changes in their observed orbital periods that are likely due to Roche lobe overflow of the companion or interactions between the companion and the pulsar wind (e.g. Thongmeearkom et al. 2024). This orbital variability can also affect the measured spin period derivatives, and is harder to correct for in redbacks where the orbital variability is ubiquitous and complex. The observed changes in the spin periods of redbacks could therefore be incorrectly interpreted as effects from the cluster potential, so we follow Prager et al. (2017) and exclude redback systems from our analysis. Among the well-timed pulsars, this means we exclude pulsar *W* from 47 Tuc and pulsars *A*, *P*, *ad* and *ar* from Terzan 5. The pulsar data is summarized in Tables 5.4 and 5.5 in the Appendix.

Finally, we make use of the Australia Telescope National Facility’s pulsar database³ presented by Manchester et al. (2005) in order to build a representative population of Galactic MSPs which we use to estimate the probability distribution for the intrinsic spin-down of the pulsars as a function of their spin period (see Section 4.3.3.1).

³<http://www.atnf.csiro.au/research/pulsar/psrcat>

4.3 Methods

4.3.1 Models

To model the dynamics and mass distribution of 47 Tuc and Terzan 5, we use the GCFIT⁴ package, recently presented by Dickson et al. (2023, 2024). This package couples a fast mass evolution algorithm with the LIMEPY⁵ family of models presented by Gieles & Zocchi (2015). We refer readers to these papers for a detailed description of the models, and provide a brief summary here. The LIMEPY models are a set of distribution function-based equilibrium models that are isothermal for the most bound stars near the cluster center and described by polytropes in the outer regions near the escape energy. The models have been extensively tested against N -body models (Zocchi et al. 2016; Peuten et al. 2017) and their multimass version is able to effectively reproduce the effects of mass segregation. Their suitability for mass modeling of GCs has been tested on mock data (Hénault-Brunet et al. 2019), and they have recently been applied to real datasets as well (for example, Gieles et al. 2018; Hénault-Brunet et al. 2020; Dickson et al. 2023, 2024).

The input parameters needed to compute our models include the dimensionless central potential ϕ_0 , the truncation parameter g^6 , the anisotropy radius r_a which determines the degree of radial anisotropy in the models, δ which sets the mass dependence of the velocity scale and thus governs the degree of mass segregation, and finally the specific mass bins to use as defined by the mean stellar mass (m_j) and total mass (M_j) of each bin, which together specify the stellar mass function. In order to scale the model units to physical units, the total mass of the cluster M and a size scale (the half-mass radius of the cluster r_h) are provided as well. Finally, we provide the distance to the cluster (D) which is used in converting between angular and linear quantities.

⁴<https://github.com/nmdickson/GCfit/>

⁵<https://github.com/mgieles/limepy/>

⁶Several well-known classes of models are reproduced by specific values of g : Woolley models (Woolley 1954) have $g = 0$, King models (King 1966) $g = 1$, and (non-rotating) Wilson models (Wilson 1975) $g = 2$.

In order to generate the input mass bins (the m_j and M_j sets of input values) for the multi-mass LIMEPY models, the GCFIT models use the `evolve_mf` algorithm, originally presented by Balbinot & Gieles (2018) and updated in Dickson et al. (2023). This algorithm combines pre-computed grids of stellar evolution models, isochrones and initial-final mass relations to model the evolution of a given initial mass function (IMF), including the effects of stellar evolution as well as (optionally) mass loss due to escaping stars and dynamical ejections. The algorithm returns a binned mass function at a requested evolutionary time, for specified metallicity, ideal for use as an input in the LIMEPY models.

We parameterize the mass function as a three-segment broken power law with break points at $0.5 M_{\odot}$ and $1.0 M_{\odot}$. We provide to `evolve_mf` the IMF slopes⁷ (α_1 , α_2 and α_3) and break points, the cluster age, metallicity and initial escape velocity. We adopt the same methodology as Dickson et al. (2023) to determine the initial escape velocities of our clusters. Briefly, we run an initial fit with an initial guess of the escape velocity and use the present-day escape velocity of this preliminary fit to set the initial escape velocity of the cluster. We use double the present-day value as our estimate for the initial escape velocity, which accounts for adiabatic expansion of the cluster after mass-loss due to stellar evolution. We note that after the initial fit, changing the initial escape velocity by 20 km s^{-1} in either direction has no discernible effect on the final model. We additionally specify parameters which control the mass loss (if any) due to escaping stars and the specific binning to be used when returning the final discrete mass-function bins. We finally provide the black hole retention fraction (BH_{ret}) which controls the percentage of the mass in black holes initially created from the initial mass function that is retained to the present day after natal kicks and dynamical ejections. We first eject primarily low-mass BHs through natal kicks and then eject the rest of the required mass by ejecting the most massive BHs first, capturing the effect of dynamical ejections (see Dickson et al. 2023 for details). For this study

⁷Throughout this work we adopt the convention that $\xi(m) \equiv dN/dm \propto m^{-\alpha}$ such that a positive value of α gives a decreasing power-law slope.

Table 4.1: Model parameters and their priors. Most priors are uniform and are chosen to bound the parameters around a reasonable range of values for both clusters. For the mass function slopes, we add the additional constraint that α_2 must be steeper than α_1 and α_3 steeper than α_2 . For the distance, we use a Gaussian prior with the distance measurement from Baumgardt & Vasiliev (2021) and its uncertainty providing the mean μ and standard deviation σ of the prior. These distances are $D = 4.521 \pm 0.031$ kpc and $D = 6.62 \pm 0.15$ kpc for 47 Tuc and Terzan 5 respectively.

Parameter	Prior Form	Value
ϕ_0	Uniform	[0.1, 15]
$M[10^6 M_\odot]$	Uniform	[0.01, 3]
r_h [pc]	Uniform	[0.5, 15]
$\log_{10}(r_a/\text{pc})$	Uniform	[0, 8]
g	Uniform	[0, 3.5]
δ	Uniform	[0.3, 0.5]
s^2	Uniform	[0, 20]
F	Uniform	[1, 7.5]
α_1	Uniform	[-1, 2.35]
α_2	Uniform	[-1, 2.35] and $\geq \alpha_1$
α_3	Uniform	[1.6, 4] and $\geq \alpha_2$
BH _{ret} [%]	Uniform	[0, 100]
D [kpc]	Gaussian	BV21

we do not model the mass loss due to escaping stars, so we set this mass loss due to escaping stars to be zero, and we are effectively specifying the present-day mass function for low-mass stars, not their IMF.

4.3.2 Fitting

The GCFIT package provides a uniform interface for fitting the coupled LIMEPY and `evolve_mf` models to a variety of observables using either a Markov Chain Monte Carlo (MCMC) or nested sampling algorithm. For this work, we use the nested sampling algorithm, which is implemented using the `dynesty` package (Speagle 2020; Kozlov et al. 2023). For the majority of our parameters we adopt wide, uniform priors, with the exception of the distance where we adopt the measurements of Baumgardt & Vasiliev (2021) as Gaussian priors and the mass function power-law slopes where we adopt the physically motivated priors described in Dickson et al. (2023). We list the priors in Table 4.1.

4.3.3 Likelihoods

The majority of the likelihood functions we use for different datasets are Gaussian likelihoods. Provided below is the log-likelihood for velocity dispersion profile data as an example, but all other likelihoods are of a similar form⁸:

$$\ln(\mathcal{L}) = -\frac{1}{2} \sum_{i=1}^{N_p} \left\{ \frac{[\sigma_{\text{obs}}(r_i) - \sigma_{\text{model}}(r_i)]^2}{\delta\sigma_{\text{obs}}^2(r_i)} - \ln[\delta\sigma_{\text{obs}}^2(r_i)] \right\}, \quad (4.1)$$

where \mathcal{L} is the likelihood, N_p is the number of data points, σ_{obs} is the measured velocity dispersion, σ_{model} is the model velocity dispersion at the corresponding radius, r is the projected distance from the cluster center, and $\delta\sigma_{\text{obs}}$ is the uncertainty in the velocity dispersion. The likelihoods for other observables are formulated in the same way, and the specifics are discussed in Dickson et al. (2023) as well as in the GCFIT documentation⁹. The total log-likelihood is the sum of all the log-likelihoods for each set of observations.

For the mass function and number density profile likelihoods, we include additional nuisance parameters and scaling terms. For the number density data, we introduce a parameter s^2 which is added in quadrature to the existing measurement uncertainties. This parameter allows us to add a constant contribution to all values in the dataset, effectively lowering the weight of the data located farthest from the cluster center where the number density is lowest. This allows us to account for limitations in the models such as the effects of potential escapers near the cluster tidal boundary that the LIMEPY models do not account for (see Claydon et al. 2019 for a discussion of potential escapers in equilibrium models).

Finally, for the number density profile data, we make an additional modification to equation (4.1) and introduce a scaling factor K which allows us to fit only on the shape of the

⁸We note that in Dickson et al. (2023) the leading minus sign in the log-likelihood functions was missing from the text.

⁹[gcfits.readthedocs.io](https://github.com/gcfits/gcfits.readthedocs.io)

number density profile instead of the absolute values. K is defined as follows (see section 3.3 of Hénault-Brunet et al. 2020 for a complete explanation) and the model data points ($\Sigma_{\text{model},i}$) are multiplied by this value before they are compared to the data:

$$K = \frac{\sum_{i=1}^{N_p} \Sigma(r)_{\text{obs},i} \Sigma(r)_{\text{model},i} / \delta \Sigma(r)_{\text{model},i}^2}{\sum_{i=1}^{N_p} (\Sigma(r)_{\text{model},i})^2 / \delta \Sigma(r)_{\text{obs},i}^2}, \quad (4.2)$$

where N_p is the number of data points, $\Sigma(r)_{\text{obs},i}$ are the number density measurements, $\Sigma(r)_{\text{model},i}$ are the model number densities at the corresponding radii and $\delta \Sigma(r)_{\text{obs},i}$ are the uncertainties on the number density measurements where r is the projected radius of a given measurement of $\Sigma(r)_{\text{obs},i}$.

We discuss the likelihoods for the pulsar timing data and the stellar mass function data separately in the subsections below.

4.3.3.1 Pulsars

Pulsar period derivatives, as measured by an observer, are made up of several distinct components, with contributions from the cluster's gravitational potential, the gravitational potential of the Milky Way, the pulsar's proper motion, intrinsic effects like magnetic braking, and the changing dispersion measure between the pulsar and the observer. The effects of the cluster's proper motion and the Galactic potential are fairly well constrained based on the pulsar's position and motion in the Galaxy but, the effects of processes like magnetic braking which are intrinsic to the pulsar itself require more careful consideration. The breakdown of the measured period derivative $(\dot{P}/P)_{\text{obs}}$ into separate components is:

$$\left(\frac{\dot{P}}{P}\right)_{\text{obs}} = \left(\frac{\dot{P}}{P}\right)_{\text{int}} + \frac{a_{\text{cl},z}}{c} + \frac{a_G}{c} + \frac{a_S}{c} + \frac{a_{\text{DM}}}{c}, \quad (4.3)$$

where $(\dot{P}/P)_{\text{int}}$ is any change in period due to the effects intrinsic to the pulsar like magnetic breaking, c is the speed of light, $a_{\text{cl},z}$ the line-of-sight acceleration of the pulsar due to the cluster’s gravitational potential and is the quantity we are most interested in, a_{G} is the acceleration of the pulsar along the line of sight due to the Galaxy’s gravitational potential, a_{S} is the ‘Shklovskii’ effect (Shklovskii 1970), an apparent acceleration due to the proper motion of the pulsar and a_{DM} is the effect of the changing dispersion measure between the pulsar and the observer.

For each of the components in equation (4.3), we explain below how we calculate either a point estimate or a probability distribution for the quantities of interest, and how we combine these to obtain a probability distribution for the measured period derivative of a pulsar, given a model (i.e. the likelihood).

All pulsars have some intrinsic spin-up or spin-down caused by processes like magnetic breaking or active accretion. We exclude any redback pulsars (pulsars with massive, non-degenerate companions) from this work and therefore assume that none of the pulsars are actively accreting and that any intrinsic effects are purely in the spin-down direction (positive $(\dot{P}/P)_{\text{int}}$ term in Equation 4.3). To estimate the probability distribution for the intrinsic spin-down distribution, we assume that the intrinsic spin-down of cluster pulsars follows the same distribution as the Galactic field pulsars, and that it is dependent only on their period. We use the ATNF pulsar catalog¹⁰ (Manchester et al. 2005) to build a distribution of possible \dot{P} values for a given value of P using the Galactic field pulsars as a reference (for which the period derivative can be directly linked to the intrinsic spin-down after correcting for Galactic and proper motion contributions due to there being no cluster acceleration for these pulsars). We compute a Gaussian kernel density estimator (KDE) in the field P - \dot{P} space, which is sliced along each cluster pulsar’s period to extract a distribution of intrinsic values. We show the field pulsars in the P - \dot{P} plane, the KDE

¹⁰<https://www.atnf.csiro.au/research/pulsar/psrcat/>

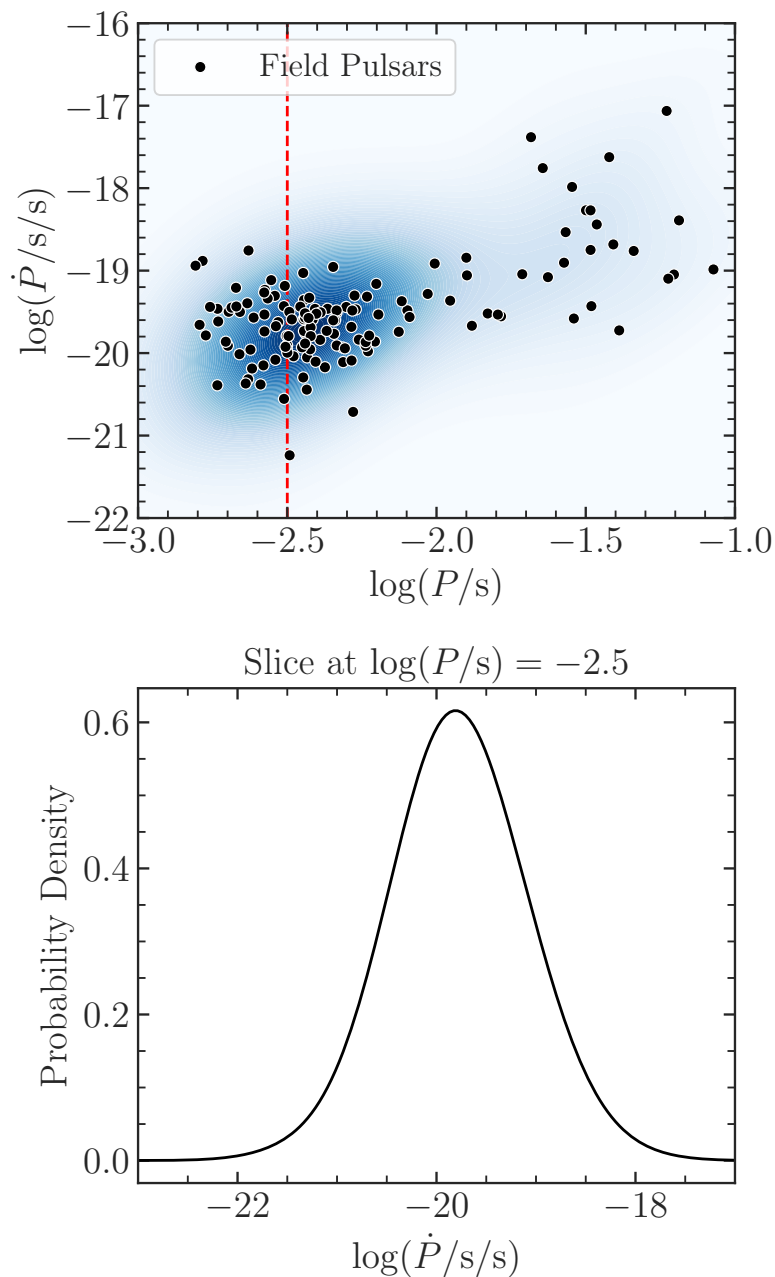


Figure 4.1: *Top*: The $P-\dot{P}$ plane for the field pulsars with the Gaussian KDE of the $P-\dot{P}$ distribution shown in blue. *Bottom*: an example of a slice from this KDE. The slice is taken at $\log P = -2.5$ s and shows the distribution of $\log \dot{P}$ values for pulsars with this period.

and an example of a slice from this KDE in Figure 4.1.

The next two components, $a_{\text{cl},z}/c$ and a_{G}/c , are fundamentally similar in that they are both manifestations of the Doppler effect. In the typical case, we infer a star’s radial velocity by measuring the frequency shift of some known spectral feature, but in the case of pulsars, we instead measure the acceleration of the pulsar along the line of sight by measuring the rate of change of the pulsar’s period.

We will first look at a_{G}/c , the effect of the Galaxy’s gravitational potential on the pulsar’s period derivative. The acceleration due to the Galaxy’s potential is a function of the pulsar’s position in the Galaxy and the Galaxy’s mass distribution. We use the GALA package (Price-Whelan 2017) to calculate the acceleration due to the Galactic potential at each cluster’s position. We adopt the MILKYWAYPOTENTIAL2022 potential from GALA as well as the cluster positions measured by Vasiliev & Baumgardt (2021) who used *Gaia* EDR3 data to measure the positions (including distances) and kinematics of Milky Way GCs. After projecting this acceleration along the line of sight, the effect on the period derivative is the following:

$$\dot{P}_{\text{G}} = \frac{a_{\text{G}}P}{c}, \quad (4.4)$$

where \dot{P}_{G} is the contribution to the period derivative due to the Galactic potential.

The effect of the cluster’s gravitational potential on the pulsar’s period derivative (the $a_{\text{cl},z}/c$ term in equation 4.3) is the effect we are most interested in as it helps us to constrain the internal mass distribution of the cluster. For a pulsar with a known 3-D position within the cluster, the period derivative due to the cluster’s gravitational potential (\dot{P}_{cl}) will simply be:

$$\dot{P}_{\text{cl}} = \frac{a_{\text{cl},z}P}{c}. \quad (4.5)$$

Because the pulsar position along the line of sight is unknown, we need to generate a probabil-

ity distribution over possible line-of-sight accelerations for a given pulsar at a given projected radius for a given model, based on the enclosed mass over the full range of possible line-of-sight positions. We can then weight this distribution by the probability of a pulsar being at a given position z along the line of sight to generate a probability distribution for \dot{P}_{cl} given a projected radius for a given model.

In order to generate the probability distribution of a pulsar being at a given position along the line of sight we insert a tracer mass bin into the LIMEPY models. This tracer mass bin is a single mass bin with a mean mass of $1.6 M_{\odot}$ ¹¹ and a negligible total mass. This allows us to calculate the (line-of-sight) density profile of pulsar-mass objects in our model at a given projected radius. This profile is proportional to the probability that a pulsar has a given line-of-sight acceleration, given the model parameters, which indirectly provides constraints on the position of the pulsar along the line-of-sight. Combining the line-of-sight density profile with the line-of-sight acceleration profile we then obtain a probability distribution over the range of period derivatives for a given projected radius for a given model.

We show the expression from which we can calculate this probability distribution in equation (4.6):

$$P(a_{\text{cl},z} | R_i) \propto \frac{dm}{da_{\text{cl},z}} = \frac{dm}{dz} \left| \frac{dz}{da_{\text{cl},z}} \right| = \frac{\rho(z)}{|da_{\text{cl},z}/dz|}, \quad (4.6)$$

where $P(a_{\text{cl},z} | R_i)$ is the probability of a given line-of-sight acceleration measurement ($a_{\text{cl},z}$) for a projected radius R_i , m is mass column density of pulsar-mass objects along the line of sight at projected radius R_i , $a_{\text{cl},z}$ is the line-of-sight acceleration for a given line-of-sight position and $\rho(z)$ is the mass density of pulsar-mass objects at a given line-of-sight position. We note that, as stated on the right-hand side of Equation 4.6, this probability distribution is not

¹¹We use $1.6 M_{\odot}$ as the tracer mass because most of the pulsars (which we assume to have masses of $1.4 M_{\odot}$) have binary companions with typical masses of $0.2 M_{\odot}$ (see e.g. Freire et al. 2017).




Figure 4.2: An example of the combination of the line-of-sight density and acceleration profiles leading to the probability distribution of line-of-sight accelerations. *Left panel:* The normalized line-of-sight density profile for pulsar-mass objects for a model fit to Terzan 5 at the projected radius of pulsar aa . *Middle panel:* The line-of-sight acceleration profile for the same model and projected radius. The derivative of this profile, used in the calculation of the probability distribution of line-of-sight accelerations is shown with a red, dashed line. *Right panel:* The probability distribution for the line-of-sight acceleration for the same model, at the same projected radius. This distribution is a result of combining the density profile with the derivative of the acceleration profile.

normalized. After constructing this distribution, we explicitly normalize it such that it behaves like a probability density function. Each of the quantities on the right-hand side of Equation (4.6) ($\rho(z)$, $da_{cl,z}/dz$) are calculated for each LIMEPY model. We show the combination of these distributions in Figure 4.2.

The Shklovskii effect (a_S/c) is the effect of the pulsar’s proper motion on its observed period derivative. Any transverse motion of a pulsar acts to increase the distance to the pulsar, regardless of the direction of motion. A constant transverse motion results in a non-linear increase in the distance, manifesting as an apparent line-of-sight acceleration (e.g. Verbiest et al. 2008). This effect is calculated as:

$$\dot{P}_S = \frac{a_S P}{c} = \frac{\mu^2 D P}{c}, \quad (4.7)$$

where \dot{P}_S is the rate of change of the period due to the Shklovskii effect and μ is the proper motion. We use the cluster’s bulk proper motion to calculate this effect and again adopt the measurements of Vasiliev & Baumgardt (2021)¹². Finally, D is the distance to the cluster, one

¹² $\mu_{\alpha^*} = 5.253 \pm 0.008$ mas yr⁻¹ and $\mu_\delta = -2.557 \pm 0.008$ mas yr⁻¹ for 47 Tuc and $\mu_{\alpha^*} = -1.864 \pm$

of the parameters which we allow to vary in our fitting. This effect is of order $a_S \sim 10^{-11} - 10^{-10} \text{ m s}^{-2}$ which is negligible compared the acceleration due to the cluster potential which is of order $10^{-9} - 10^{-8} \text{ m s}^{-2}$ (e.g. Figure 4.2).

The final component, a_{DM}/c , is the effect of the changing dispersion measure between a pulsar and the observer. The total amount and distribution of the ionized gas along our line of sight is not necessarily constant over the full time-span over which these observations were performed and small changes in the total dispersion measure between us and the cluster can cause small variations in the observed period derivatives (see, for example, Prager et al. 2017). This effect is stochastic, meaning it is unlikely to bias the timing solution in one direction or the other. Furthermore, the magnitude of this effect is expected to be very small, on the order of $10^{-13} \text{ m s}^{-2}$ (Prager et al. 2017), several orders of magnitude smaller than the typical acceleration from the cluster potential, therefore we do not consider it in our analysis.

One potential contribution to the observed values of \dot{P}/P that we do not model is the acceleration and its higher-order derivatives caused by nearby stars in the dense core of the cluster. For the acceleration of the pulsars in particular, this effect has been shown to be typically ~ 2 orders of magnitude smaller than the mean-field acceleration from the cluster potential as a whole (Phinney 1993; Prager et al. 2017). This effect is however relevant for the higher order derivatives, where nearby stars contribute at a similar level to the bulk cluster potential (Blandford et al. 1987; Gieles et al. 2018). It is for this reason that while higher-order period derivatives are measured for many of the pulsars we use in this work, we chose not to consider these measurements in our determination of the mass distributions of our clusters.

To combine these various effects into a likelihood function for $(\dot{P}/P)_{\text{meas}}$ given a model, we start with the distribution of \dot{P}_{cl}/P from the cluster's gravitational potential and convolve with it the intrinsic distribution of $(\dot{P}/P)_{\text{int}}$ values for the period of a given pulsar. We additionally 0.030 mas yr⁻¹ and $\mu_\delta = -5.108 \pm 0.027 \text{ mas yr}^{-1}$ for Terzan 5.




Figure 4.3: *Left panel:* The probability distribution of \dot{P}_{int} (intrinsic spin-down) for pulsar *aa* in Terzan 5. Note that this distribution is of the same form as the one shown in the bottom panel of Figure 4.1 but here shown on a linear scale. *Middle panel:* The probability distribution of \dot{P}_{cl} (due to the cluster’s potential at the projected radius of pulsar *aa* in a model fit to Terzan 5). *Right panel:* The convolution of the intrinsic and cluster \dot{P} distributions, then convolved with a Gaussian distribution representing the uncertainty on the measured period derivative, transformed to \dot{P}/P , with the observed value for pulsar *aa* shown as a vertical dashed line.

convolve the distribution with a Gaussian distribution, centered at zero with a width equal to the uncertainty of $(\dot{P}/P)_{\text{meas}}$ in order to fully incorporate the uncertainty of the period derivative measurement. We then shift this distribution by the point estimates for the contributions of the Galactic potential and the Shklovskii effect. This results in a probability distribution for a measured period derivative which fully incorporates the physical effects within the cluster, which depend on our model parameters as well as the effects of the Galactic gravitational potential and the effects of the pulsars’ proper motions. We use this probability distribution to compute the likelihood of each measured period derivative. We show an example of the combination of these various distributions into the final likelihood in Figure 4.3.

Many pulsars in GCs are in binary systems, and for systems with well-determined timing solutions, the orbital period derivatives of these systems can be measured. The orbital period solutions are useful because the orbital periods of these systems are of the order of days, while the intrinsic orbital decay of these systems acts over millions of years (e.g. Binney & Tremaine 2008). This means that, in the cases where the orbital period derivative can be measured, the changes in the orbital period can be entirely attributed to the acceleration from the cluster and

the well-constrained effects of the Galactic potential and pulsar's proper motion. We note that some black widow pulsars, like pulsars *J* and *O* in 47 Tuc, show orbital variability that is likely due to interactions with their companion (e.g. Shaifullah et al. 2016; Freire et al. 2017) which could be incorrectly interpreted as changes in the period derivative due to the cluster potential. The orbital period derivatives of these systems cannot be measured, so these problematic systems are already excluded from our sample.

Due to the longer timescales and the difficulties associated with determining the orbital periods of these systems (see Ridolfi et al. 2016 for details), the relative uncertainties on the orbital period derivatives are much larger than those of the spin-period derivatives. These larger uncertainties mean that the likelihood functions for the observed orbital period derivatives are wider and provide weaker constraints on the mass distribution of the cluster, but we nonetheless use the orbital period derivatives of these systems (pulsars *E*, *H*, *I*, *Q*, *R*, *S*, *T*, *U*, *X* and *Y* in 47 Tuc and pulsars *ao*, *ap*, *au*, *av*, *aw* and *ax* in Terzan 5¹³) as an additional constraint on the cluster potential, independent of any intrinsic effects on the period derivatives. We construct these likelihood functions in an identical way to the spin period likelihoods but we neglect any effects intrinsic to the binary systems.

One avenue for future improvement of this methodology lies in including the dispersion measures of the pulsars in the analysis. The dispersion measure of a pulsar provides a measure of the amount of ionized gas between the pulsar and an observer, with a higher column density of free electrons producing a larger dispersion measure. Given an estimate of the average dispersion measure between an observer and a cluster and a model for the internal gas distribution within a cluster, the dispersion measures provide an estimate of the line-of-sight position of each pulsar

¹³There are many additional pulsars in binary systems in Terzan 5 (e.g. Ransom et al. 2005) however the timing solutions for these systems lack reported uncertainties which are required for our method. The spin-period timing solutions from Ransom et al. (2005) are similarly lacking reported uncertainties however, in practice, uncertainties on spin-period solutions are so small that our method is insensitive to their value. In these cases we adopt a single value for the uncertainty on the spin period derivatives, taking 5×10^{-21} s/s as a conservative estimate for these pulsars though we stress that our results are not sensitive to the adopted value.

within the cluster.

For 47 Tuc, Abbate et al. (2018) used the pulsars within the cluster to infer the internal gas distribution. These authors found that the pulsar data preferred a uniform gas distribution within the cluster rather than a distribution that follows the stellar density (see also Pancino et al. 2024), finding a gas density n_g of $0.23 \pm 0.05 \text{ cm}^{-3}$. This measurement, combined with the average cluster dispersion measure DM_{cl} of $24.38 \pm 0.02 \text{ pc cm}^{-3}$ allows us to infer the 3-D position of each pulsar within 47 Tuc, independent of our modeling. We describe the necessary modifications to equation 4.6 in Appendix 5.4.1.

We implemented this alternative formulation and applied it to 47 Tuc to test if the dispersion measures would enable the pulsar timing data to place stronger constraints on the mass distribution within our models. We found that while the dispersion measures did in some cases provide stronger constraints from individual pulsars, the uncertainties on the dispersion-measure-based line-of-sight positions are such that the overall constraints are ultimately very similar to those provided by the density-based calculation described in equation (4.6). Because the required internal gas models do not yet exist for Terzan 5 and because the dispersion measures provided little to no improvement for 47 Tuc, we opt to simply use the density-based calculation for both clusters for the remainder of this paper.

4.3.4 Stellar mass functions

The stellar mass function likelihoods are also Gaussian likelihoods, however, care must be taken when extracting model values due to the non-trivial footprint of the observed *HST* fields from which the mass functions were extracted. To ensure that we are extracting mass functions from the same corresponding regions in the models, we employ a Monte Carlo integration method which allows us to handle the irregular overlapping *HST* fields. This process is described in detail by Dickson et al. (2023) and is implemented in the GCFIT package.

The only uncertainty formally included with the stellar mass function data is the Poisson counting error. We introduce a nuisance parameter F which scales up the uncertainties on the absolute counts by a constant factor, leading to larger relative errors in regions with lower counts. This error encapsulates additional sources of error that may not have been accounted for such as the error associated with the conversion from luminosity to mass with an isochrone and the fact that the mass function is being approximated as a broken power law, a functional form which may not be a perfectly accurate representation of the true mass function of the cluster.

4.3.5 Stellar populations and mass bins

To generate the input mass bins for the LIMEPY models the `evolve_mf` algorithm requires as inputs the mass function power-law slopes, as well as the age and metallicity of the stellar population. For 47 Tuc we adopt an age of 11.75 Gyr (VandenBerg et al. 2013) and a metallicity of $[\text{Fe}/\text{H}] = -0.72$ (Harris 1996,2010 edition). Depending on the fit, as detailed below, we either allow the mass function power-law slopes to vary or fix the mass function.

Terzan 5 has been found to have at least three distinct stellar populations (Ferraro et al. 2009; Origlia et al. 2013; Ferraro et al. 2016). One of these populations, the most metal-poor population ($[\text{Fe}/\text{H}] = -0.8$), only makes up a small fraction of the cluster and can be neglected in our analysis. The other two population are a young (4.5 Gyr) super-solar ($[\text{Fe}/\text{H}] = 0.2$) population, making up about 40% of the cluster (by mass) and an old (12 Gyr) population with $[\text{Fe}/\text{H}] = -0.2$ making up the other 60% of the cluster. With the young population making up a significant fraction of the cluster, we cannot simply assume a single, old stellar population for Terzan 5 as we do for 47 Tuc as we want our remnant populations to be as realistic as possible.

Tests with our stellar evolution algorithm show that this 40/60 mixture of the young and old stellar populations results in a remnant mass fraction around 35% at the present day, made up of mostly white dwarfs. We can achieve a similar remnant fraction ($\sim 34\%$) with a single metal-

poor ($[\text{Fe}/\text{H}] = -0.2$), intermediate-age population of about 8 Gyr. Using this intermediate-age stellar population also makes the main-sequence turnoff mass consistent with the maximum mass of main-sequence stars of $\sim 0.9 M_{\odot}$ in our stellar mass function data, avoiding possible issues when comparing the observed mass function and model predictions.

While populations of different metallicities are expected to produce different remnants from similarly massive progenitors, this is a minor effect for our modeling. Our primary goal is to produce a realistic mix of remnants that together make up the correct fraction of the total mass of the cluster, a goal for which an intermediate-age population is a useful simplification.

Finally, the stellar mass function data available for Terzan 5 does not cover a wide enough range of masses or distances from the cluster center to leave the mass function power-law slopes free while fitting as they are derived from a single field and only extend down to $\sim 0.6 M_{\odot}$. As such, we chose to fix the mass function slopes for Terzan 5 while still including the mass function data as a constraint on the visible stellar mass in three radial bins between 0.67 and 1.67 arcminutes from the center.

We chose to adopt for the present-day mass function of Terzan 5 the bottom-light IMF of Baumgardt et al. (2023), which was measured from star clusters in the Milky Way and Magellanic Clouds and represents the best estimate for the IMF of massive star clusters. We show in Section 5.1 that this mass function provides a satisfactory match to the available stellar mass function data. This mass function has slopes of $\alpha_1 = 0.3$, $\alpha_2 = 1.65$, $\alpha_3 = 2.3$ and we again place our breakpoints at 0.5 and $1 M_{\odot}$.

Chapter 5

Probing populations of dark stellar remnants in the globular clusters 47

Tuc and Terzan 5 using pulsar timing:

Results, Discussion and Conclusions

This chapter contains the Results, Discussion and Conclusions sections (sections 4, 5 and 6) of the following manuscript submitted to the *Astrophysical Journal* for publication:

Probing populations of dark stellar remnants in the globular clusters 47 Tuc and Terzan 5 using pulsar timing

P. J. Smith, V. Hénault-Brunet, N. Dickson, M. Gieles, H. Baumgardt

I am the lead author of this manuscript and led the analysis and writing of the paper. I produced all figures and tables and wrote the text.

Table 5.1: Summary of the different model fits for 47 Tuc and Terzan 5, showing which datasets are included or held out in each case. The columns indicate if the models are fit to the number density profile (NDP), line-of-sight velocity dispersion profile (LOS), proper motion dispersion profile (PM), stellar mass function data (MF) and the pulsar timing data.

Fit	NDP	LOS	PM	MF	Pulsars
47Tuc-AllData	✓	✓	✓	✓	✓
47Tuc-NoPulsars	✓	✓	✓	✓	
47Tuc-NoKin	✓			✓(one field)	✓
Ter5-AllData	✓	✓	✓	✓	✓
Ter5-NoPulsars	✓	✓	✓	✓	
Ter5-NoKinNoMF	✓				✓

5.1 Results

To test the performance of our method, we run several fits for each cluster with different subsets of the data introduced in Section 4.2.

For 47 Tuc, we perform three fits: (1) a fit to all available data for this cluster (47Tuc-AllData), (2) a fit with the pulsar timing data held out (47Tuc-NoPulsars), and (3) a fit to the number density profile, pulsar timing data and a single field of mass function data¹ (47Tuc-NoKin), designed to emulate the data available for Terzan 5.

For Terzan 5 we also run three fits: (1) a fit with all of the available data for this cluster (Ter5-AllData), (2) a fit with the pulsar data held out (Ter5-NoPulsars), and finally (3) a fit on just the number density profile and pulsar data, with the kinematic data and mass function data held out (Ter5-NoKinNoMF), to test the reliability of the limited stellar kinematic data available for Terzan 5.

These fits are summarized in Table 5.1. We present the median of the posterior probability distribution and 1σ credibility intervals for the parameters of each of these fits in Table 5.2².

¹The field from 5.0 – 8.33 arcmin, *HST* proposal ID 11677. This field was chosen to roughly probe a similar radial region as the single field available for Terzan 5.

²We have made the plots and sampler outputs for all six of our fits available in an online repository: [10.5281/zenodo.12004419](https://zenodo.org/record/12004419)

Table 5.2: Medians and 1σ uncertainties of each model parameter for each of our fits. Entries without uncertainties indicate parameters that have been held fixed during fitting. We note that as discussed in Dickson et al. (2023), the statistical uncertainties listed here likely underestimate the true uncertainties on each parameter and, in particular, our uncertainties on the cluster mass are likely closer to 10% (Dickson et al. 2024).

Cluster	47Tuc-AllData	47Tuc-NoPulsars	47Tuc-NoKin	Ter5-AllData	Ter5-NoPulsars	Ter5-NoKinNoMF
$\hat{\phi}_0$	$6.08^{+0.08}_{-0.08}$	$6.05^{+0.07}_{-0.06}$	$6.00^{+0.10}_{-0.06}$	$5.9^{+0.3}_{-0.3}$	$6.0^{+1.3}_{-0.3}$	$5.9^{+0.4}_{-0.3}$
M [$10^6 M_\odot$]	$0.899^{+0.006}_{-0.006}$	$0.907^{+0.006}_{-0.005}$	$0.96^{+0.02}_{-0.01}$	$0.67^{+0.06}_{-0.04}$	$0.79^{+0.06}_{-0.07}$	$0.70^{+0.06}_{-0.07}$
r_h [pc]	$6.68^{+0.04}_{-0.04}$	$6.70^{+0.04}_{-0.04}$	$7.03^{+0.06}_{-0.06}$	$2.1^{+0.3}_{-0.2}$	$2.0^{+0.4}_{-0.3}$	$2.3^{+0.4}_{-0.3}$
$\log_{10}(\hat{r}_a)$	$1.73^{+0.05}_{-0.03}$	$1.78^{+0.05}_{-0.04}$	$4.88^{+2.01}_{-1.88}$	$4.75^{+2.07}_{-2.27}$	$5.39^{+1.77}_{-2.50}$	$4.49^{+2.10}_{-2.18}$
g	$1.50^{+0.03}_{-0.03}$	$1.54^{+0.02}_{-0.02}$	$1.50^{+0.02}_{-0.03}$	$1.3^{+0.5}_{-0.6}$	$2.0^{+0.2}_{-0.5}$	$1.6^{+0.4}_{-0.6}$
δ	$0.47^{+0.01}_{-0.01}$	$0.48^{+0.01}_{-0.01}$	$0.489^{+0.008}_{-0.013}$	$0.38^{+0.06}_{-0.05}$	$0.34^{+0.06}_{-0.03}$	$0.41^{+0.06}_{-0.06}$
s^2 [arcmin $^{-4}$]	$0.0006^{+0.0003}_{-0.0002}$	$0.0011^{+0.0028}_{-0.0006}$	$0.0005^{+0.0002}_{-0.0002}$	$7.84^{+4.87}_{-5.18}$	$6.37^{+5.47}_{-4.33}$	$7.33^{+5.13}_{-4.80}$
F	$2.6^{+0.1}_{-0.1}$	$2.59^{+0.09}_{-0.10}$	$5.8^{+0.8}_{-0.7}$	$1.9^{+0.3}_{-0.2}$	$1.7^{+0.3}_{-0.2}$	$1.9^{+0.3}_{-0.2}$
α_1	$0.38^{+0.03}_{-0.02}$	$0.38^{+0.02}_{-0.02}$	0.3	0.3	0.3	0.3
α_2	$1.31^{+0.04}_{-0.04}$	$1.31^{+0.04}_{-0.04}$	1.65	1.65	1.65	1.65
α_3	$2.23^{+0.03}_{-0.03}$	$2.24^{+0.02}_{-0.02}$	2.3	2.3	2.3	2.3
BH _{ret} [%]	$0.28^{+0.05}_{-0.04}$	$0.31^{+0.06}_{-0.05}$	$0.31^{+0.21}_{-0.10}$	$1.77^{+2.07}_{-1.24}$	$5.35^{+2.90}_{-2.96}$	$1.84^{+1.60}_{-1.19}$
d [kpc]	$4.41^{+0.02}_{-0.01}$	$4.43^{+0.01}_{-0.01}$	$4.45^{+0.03}_{-0.03}$	$6.7^{+0.1}_{-0.1}$	$6.7^{+0.1}_{-0.1}$	$6.7^{+0.1}_{-0.1}$

5.1.1 47 Tuc

The `47Tuc-NoPulsars` fit is very similar to the fit presented in Dickson et al. (2024), and we use this fit as a baseline to evaluate the additional leverage provided by the pulsar data. The `47Tuc-AllData` fit to all of the available data is shown in Figure 5.1 and Figure 5.2, along with examples of the likelihood functions for the measured pulsar period derivatives for the best-fitting models shown in Figure 5.3 (top panels) and all pulsars in Figures 5.11, 5.12 and 5.13.

As an additional check on our fits, we compare the velocity dispersion of the pulsars to the prediction from our models. For 47 Tuc, we limit this comparison to pulsars within $1'$ from the center, which corresponds to the isothermal (for pulsars) portion of our model. This leaves us with 22 pulsars for which we calculate a total proper motion dispersion of $0.37 \pm 0.10 \text{ mas yr}^{-1}$. The prediction from our `47Tuc-AllData` model is $0.430 \pm 0.004 \text{ mas yr}^{-1}$, in good agreement with the measured value.

An initial comparison of the `47Tuc-AllData` and `47Tuc-NoPulsars` fits reveals no significant differences, either in model parameters, fit quality or derived quantities like the black hole mass fraction. We take this agreement as an indication that the `47Tuc-NoPulsars` fit already provides a very good description of the underlying mass distribution and dynamics of the cluster, as probed by and fully consistent with the pulsar data.

Given the agreement between the `47Tuc-AllData` and `47Tuc-NoPulsars` fits, we turn to the third case in order to evaluate the leverage provided by the pulsar data. In the `47Tuc-NoKin` fit we seek to emulate the data that we have for Terzan 5. For this fit, we fix the mass function to the bottom-light IMF of Baumgardt et al. (2023) discussed in Section 4.3.5. This mass function is a reasonable approximation for 47 Tuc and is similar to the best-fitting mass function we infer when the mass function is allowed to vary (see Table 5.2). We show part

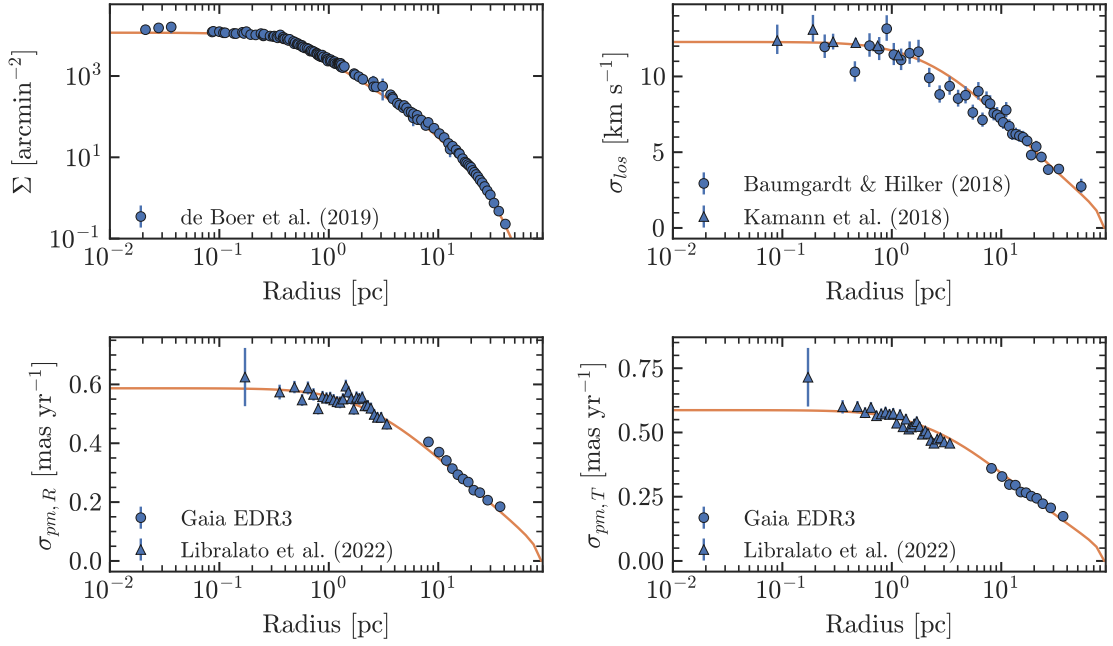


Figure 5.1: Best-fitting model (in orange) for the `47Tuc-AllData` fit compared to different datasets (in blue). *Top left*: Model fit to the projected number density profile. *Top right*: Model fit to the projected line-of-sight velocity dispersion profile. *Bottom panels*: Model fit to the projected proper motion dispersion profile, separated into radial (left panel) and tangential (right panel) components. The shaded regions represent the 1σ and 2σ credible intervals.

of the results for this third fit in Figure 5.4, where the best-fitting model is plotted along with the stellar kinematic data even though this data is excluded from the fit. This model is in excellent agreement with the data, similar to the models that are directly fit on the stellar kinematics and the best fit parameters for this fit are similar to the previous two fits (see Table 5.2). A comparison of the enclosed mass profiles of the `47Tuc-AllData` and `47Tuc-NoKin` fits reveal that the mass profiles vary by less than $\sim 5\%$ within the innermost 1 pc (where the `47Tuc-NoKin` fit contains less mass) and the total mass varies only by $\sim 5\%$ with the `47Tuc-NoKin` fit favoring a slightly higher mass.

We show our inferred posterior probability distribution for the cluster mass in BHs for each of our three fits of 47 Tuc in Figure 5.5. While the results for the `47Tuc-AllData` and `47Tuc-NoPulsars` fits are quite similar and both resemble the results of Dickson et al. (2024), the `47Tuc-NoKin` fit is worth discussing in more detail. The most obvious feature of

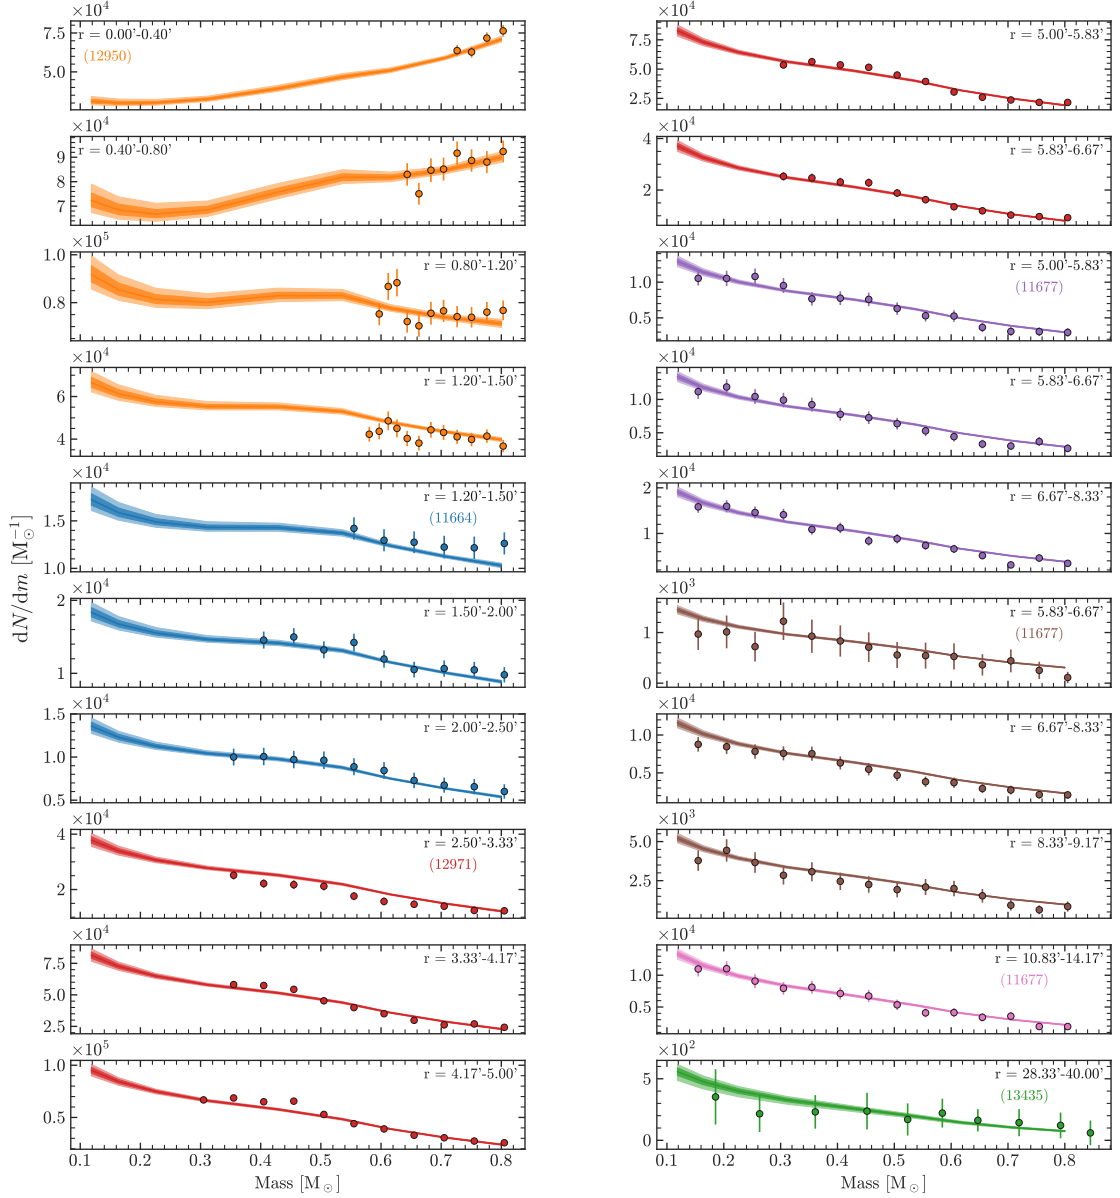


Figure 5.2: Continuation of Figure 5.1 showing the best-fitting model for the 47Tuc-AllData fit compared to the stellar mass function data of 47 Tuc. Each colour corresponds to a different *HST* field and each panel to a different radial region within a field. We display the *HST* proposal ID for in the first panel of each field.




Figure 5.3: Likelihood functions corresponding to the best-fitting model for the observed \dot{P}/P and \dot{P}_b/P_b (recall that P and P_b are the spin and orbital periods of the pulsars) for one pulsar in each cluster. *Top left*: The \dot{P}/P likelihood for pulsar S in 47 Tuc from the 47Tuc-AllData fit. *Top right*: The \dot{P}_b/P_b likelihood for pulsar S in 47 Tuc from the 47Tuc-AllData fit. *Bottom left*: The \dot{P}/P likelihood for pulsar ap in Terzan 5 from the Ter5-AllData fit. *Bottom right*: The \dot{P}_b/P_b likelihood for pulsar ap in Terzan 5 from the Ter5-AllData fit. In each panel, we show the observed period derivative as a vertical orange line. We show similar plots for all of the pulsars in Figures 5.11 through 5.16.

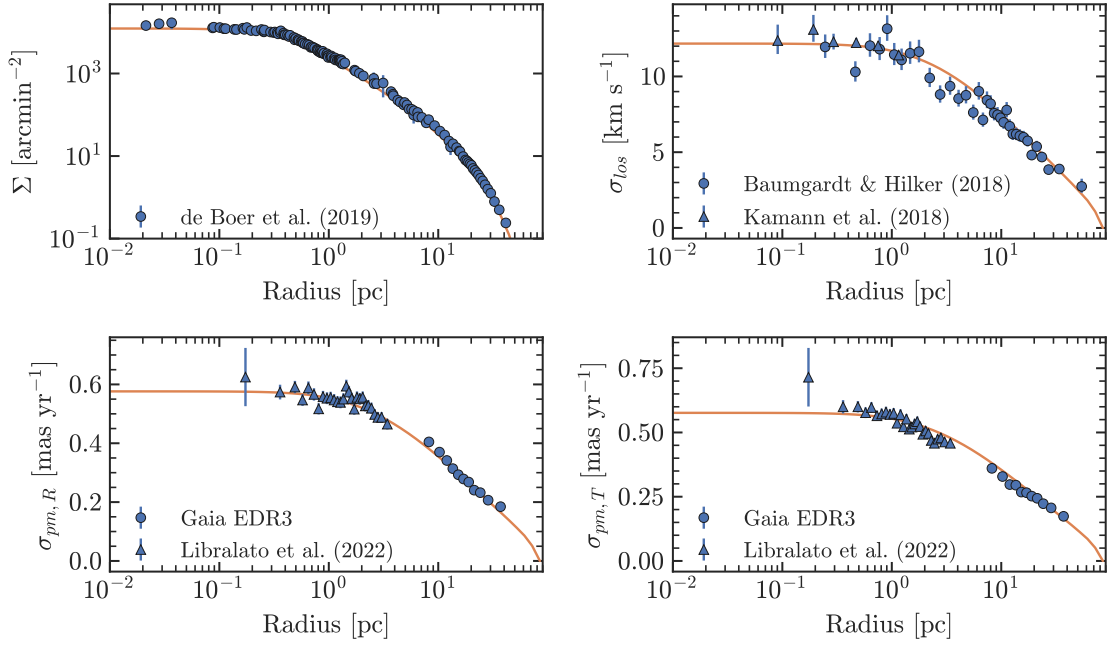


Figure 5.4: Same as Figure 5.1 but for the 47Tuc-NoKin fit. The best-fitting model is plotted along with the stellar kinematic data even though this data is excluded from the fit.

this fit is that the posterior distribution of the mass in BHs is much broader than the cases with abundant stellar kinematic and mass function data, and this posterior is also not uni-modal. Despite this, we can still place a very stringent upper limit (99th percentile) on the mass in BHs, limiting this mass to less than $\sim 0.1\%$ of the total cluster mass, even in the 47Tuc-NoKin fit. We further note that even though we have adopted a fixed IMF for the 47Tuc-NoKin fit, the inferred BH content is consistent with the fits where we allow the IMF to vary.

These tests demonstrate that the pulsar timing data can provide constraints on the central dynamics, BH content, and mass distribution of the cluster very similar to the stellar kinematics, and in cases like Terzan 5 where the stellar kinematic data are lacking, pulsar timing data may provide an excellent substitute. With this in mind, we discuss the case of Terzan 5 next.

5.1.2 Terzan 5

We show the Ter5-AllData fit in Figures 5.6 and 5.7, along with examples of the likelihood functions for the measured period derivatives for the best-fitting models shown in Figure 5.3




Figure 5.5: Posterior probability distribution of the mass in BHs in 47 Tuc for each of the fits summarized in Table 5.1.

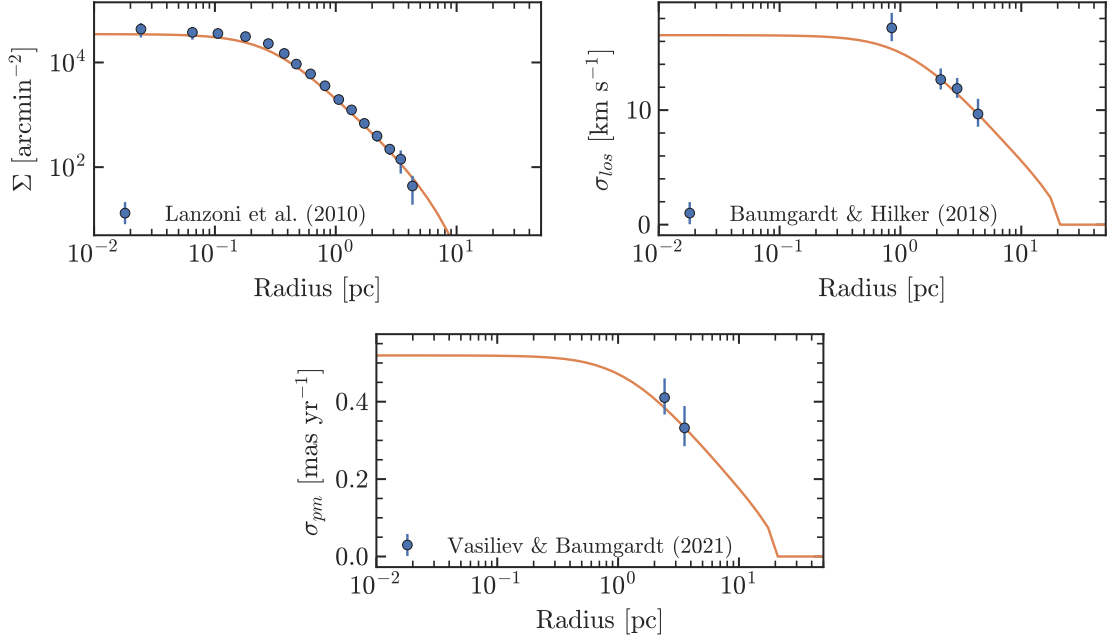


Figure 5.6: Best-fitting model (in orange) for the `Ter5-AllData` fit compared to different datasets (in blue). *Top left*: Model fit to the projected number density profile. *Top right*: Model fit to the projected line-of-sight velocity dispersion profile. *Bottom*: Model fit to the projected proper motion dispersion profile. The shaded regions represent the 1σ and 2σ credible intervals of the model fits.

(bottom panels) and all the pulsars in Figures 5.14, 5.15 and 5.16. Our best-fitting model for the `Ter5-AllData` fit is in excellent agreement with the data and is able to fully reproduce all of the observables, including the pulsar data. The comparative lack of data for Terzan 5 means that our inferred model parameters for Terzan 5 generally have larger uncertainties compared to our fits of 47 Tuc. We present the median and 1σ uncertainties on all model parameters in Table 5.2.

For Terzan 5, most pulsars do not have reported proper motions, and fewer still have both components of the proper motions reported, meaning that we have insufficient data to calculate a velocity dispersion from the pulsar data to compare to our model prediction.

When we do not fit on the pulsar data (`Ter5-NoPulsars`), the uncertainties on model parameters and related quantities generally become larger than was the case for the `Ter5-AllData` fit. This comparison suggests that the pulsar data can play a more dominant role in constraining the models in the case of Terzan 5 (compared to 47 Tuc), given the lack of stellar kinematic data

for this bulge cluster.

The third fit, `Ter5-NoKinNoMF`, was done to test the reliability of the stellar kinematics data given the challenge of membership determination in the Galactic bulge. As we find for 47 Tuc, even when we exclude the stellar kinematic data from the fit, the pulsar timing data provides enough constraints that the resulting best-fitting model is relatively similar and generally in good agreement with the held-out data. With the insight from 47 Tuc (see Section 5.1.1) that pulsars can provide strong constraints on the mass distribution of a cluster, we interpret this agreement as a sign that the existing stellar kinematics for Terzan 5, while sparse, are likely not suffering from significant contamination or other systematic effects.

As mentioned previously, the stellar mass function data available for Terzan 5 does not cover a wide enough range of masses or radii to a sufficient completeness level to allow us to leave the mass function slopes as free parameters when fitting models. We adopted the IMF of Baumgardt et al. (2023) for each of our three fits and we show in Figure 5.7 that this mass function is in excellent agreement with the available data.

We show our inferred cluster masses in BHs for each of our fits of Terzan 5 in Figure 5.8. Comparing the fits, it is again obvious that the pulsar data is providing most of the constraining power, with the `Ter5-NoPulsars` fit resulting in an upper bound on the mass in BHs roughly double that of the two fits that do include the pulsar data. Even with the inclusion of the pulsar data we are only able to place an upper limit on the mass in BHs in this cluster, though we do significantly improve on existing estimates, lowering the range of allowed masses by a factor of ~ 10 (see discussion in Section 5.2.3). With these results, we cannot rule out the possibility that Terzan 5 contains zero BHs at the present day and indeed our posterior distribution of mass in BHs is peaked towards zero for both fits that include the pulsar data.




Figure 5.7: Continuation of Figure 5.6 showing the best-fitting model for the Ter5-AllData fit compared to the stellar mass function data of Terzan 5. Each panel corresponds to a different radial region within the *HST* field.

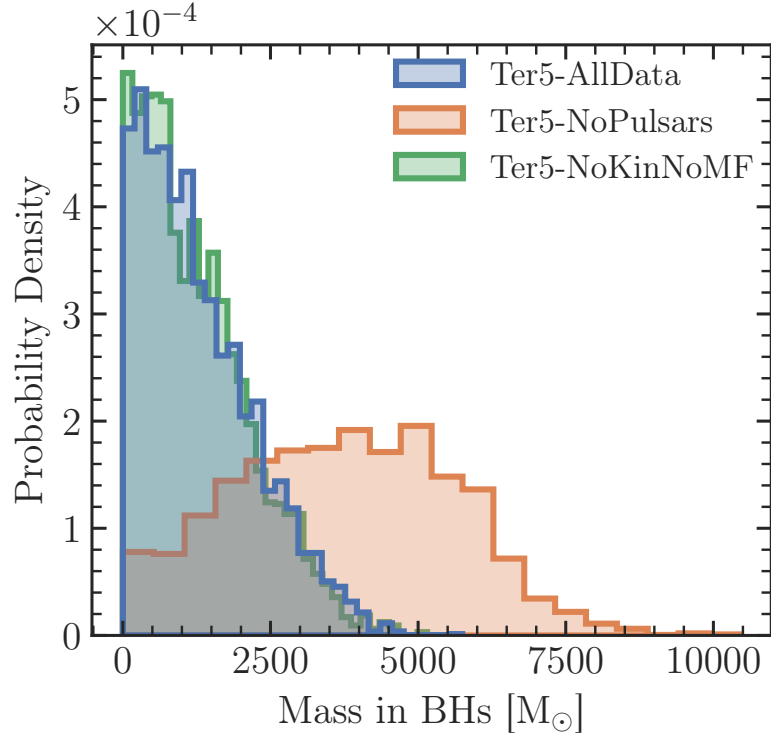


Figure 5.8: Posterior probability distribution of the mass in BHs in Terzan 5 for each of the fits summarized in Table 5.1.

5.2 Discussion

5.2.1 Constraints from pulsars

The specific constraints provided by the pulsar data are on the local acceleration of each pulsar, and thus on the enclosed mass profile of the cluster. The most stringent constraints on the cluster gravitational field come from pulsars with large *negative* observed values of \dot{P}/P . This can be seen most easily in the rightmost panel of Figure 4.3 where the positive side of the \dot{P}/P distribution has a long tail due to the convolution with the intrinsic spin-down distribution. By contrast, the negative side of the \dot{P}/P distribution truncates sharply to zero at the \dot{P}/P corresponding to the maximum possible acceleration at a given projected radius. In terms of specific pulsars, this means that pulsar *S* and *aa* in 47 Tuc provide the strictest limits on the enclosed mass while pulsar *ae* provides the strictest limits in Terzan 5. When inferring the dark remnant content of a cluster (in particular BHs), the ideal pulsars would be very close to the center of

the cluster and would have large negative observed values of \dot{P}/P^3 , as these pulsars place the strongest constraints on the central mass distribution of their host cluster, where the mass density is dominated by heavy remnants due to mass segregation. We show in Figure 5.9 the minimum and maximum allowed values of \dot{P}/P and \dot{P}_b/P_b due to the cluster potential, where pulsars providing these stronger constraints fall along the bottom contour of the allowed values of \dot{P}/P .

Our finding of pulsar *S* being the most constraining pulsar in 47 Tuc is not new and was discussed by Giersz & Heggie (2011). These authors reported that when attempting to find a suitable Monte Carlo model of 47 Tuc, the tension between the central surface brightness and the large negative acceleration of pulsar *S* was the single most impactful factor. Similarly, Figure 2 (extended) of Kızıltan et al. (2017a) shows that pulsar *S* is just barely compatible with their model and could potentially be a driving factor in requiring more mass in the cluster center, therefore favoring models with an IMBH. Pulsars 47Tuc-*S* and Ter5-*ae* do have a binary companions (Freire et al. 2017; Prager et al. 2017), however, due to the fact that the companions are low-mass WDs any spin-up effects from accretion which would shift the observed spin period derivatives towards more negative values are very unlikely. Pulsar 47Tuc-*aa* appears to be an isolated pulsar, leaving little possibility for accretion-induced spin-up. Because these pulsars are not likely to experience any spin period change from accretion, the constraints that their spin period derivatives place on the mass distribution of the cluster are likely trustworthy.

While pulsars with large negative observed values of \dot{P}/P provide the strongest constraints, the majority of pulsars have values that are much closer to zero or even on the positive extreme of the distribution. These pulsars nonetheless provide useful constraints on the mass distribution of the cluster, particularly when it comes to constraining the total mass of the cluster. These pulsars have observed values of \dot{P}/P that would be technically compatible with any model containing

³These pulsars would fall on the far side of the cluster along the line of sight. Pulsars on the near side of the cluster would have positive line-of-sight accelerations due to the cluster potential which would shift the observed \dot{P}/P to positive values where our method is less constraining because intrinsic spin-down also shifts the observed \dot{P}/P towards positive values.

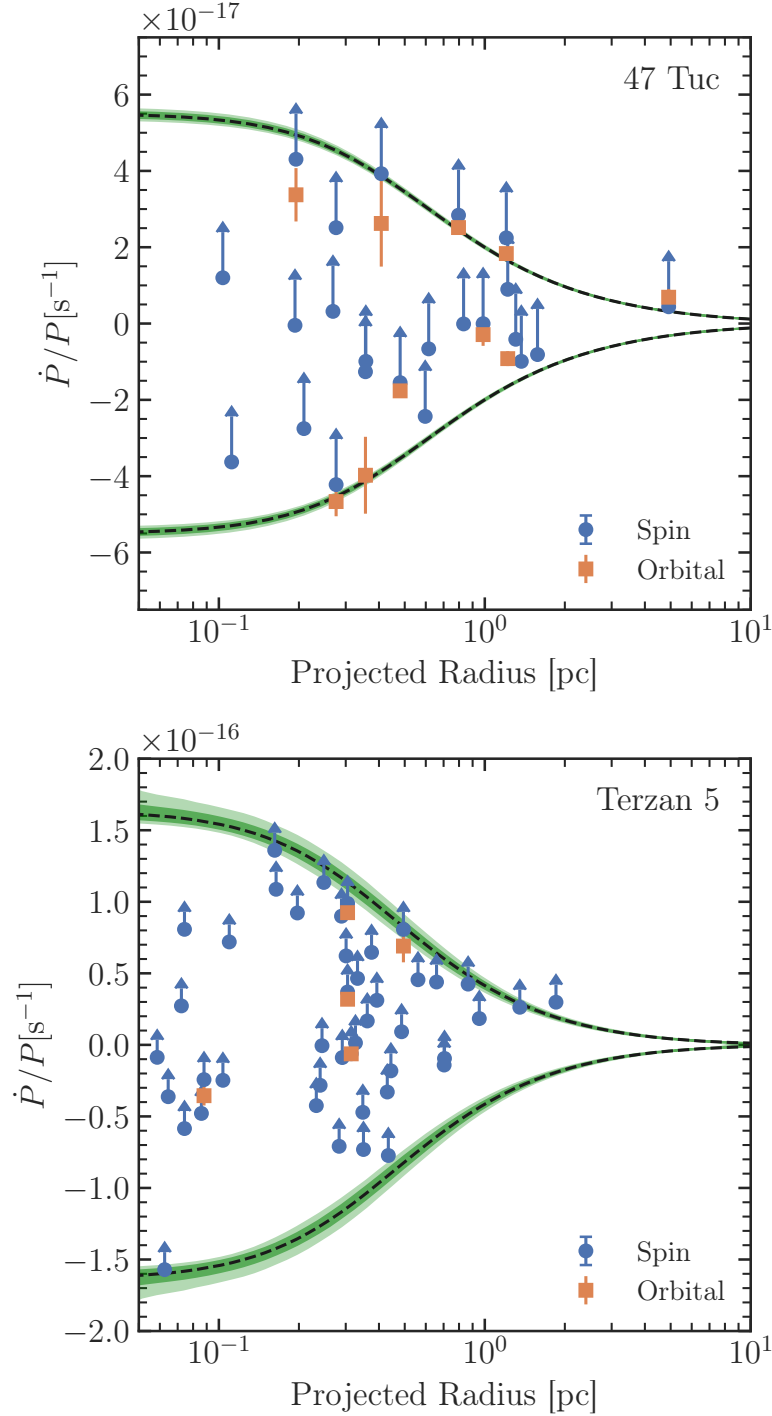


Figure 5.9: Minimum and maximum values of \dot{P}/P allowed from the acceleration in the cluster potential for the best-fitting models of our 47Tuc-AllData (top) and Ter5-AllData (bottom) fits. We show the median values as a dashed line while the shaded regions represent the 1σ and 2σ credible intervals of the model fits. The observed \dot{P}/P of the pulsars are shown with blue circles and the observed \dot{P}_b/P_b are shown with orange squares. Uncertainties on the \dot{P}_b/P_b points represent measurement errors while the upward-facing error bars on the \dot{P}/P points show the typical width of the \dot{P}_{int} distribution at each pulsar’s location in the P - \dot{P} plane shown in Figure 4.1. Note that pulsars above the maximum contour are not disallowed by our models due to the long tails on the spin-down side of the likelihood function shown in Figure 4.3. We note that we have converted the projected radius of each pulsar from angular to linear units using the median value of our inferred distances to each cluster from the AllData fits (see Table 5.2).

some minimum mass at their radius. As the mass of the model grows however, the range of possible \dot{P}/P values grows, lowering the probability of observing these less extreme values. In this way the pulsars provide not just a minimum enclosed mass at their projected radius, but also some leverage on the exact value of the enclosed mass. We can see this when we compare our fits of Terzan 5 with different subsets of the data, as the best-fitting models found when including the pulsar data in the fits (`Ter5-AllData` and `Ter5-NoKinNoMF`) are less massive than the fit that excludes the pulsar data (`Ter5-NoPulsars`).

5.2.2 Mass of Terzan 5

The total mass of Terzan 5 is somewhat uncertain in the current literature, with mass estimates based on photometry (e.g. Lanzoni et al. 2010) a factor of a few higher than those based on kinematics and dynamical modeling (e.g. Baumgardt & Hilker 2018; Prager et al. 2017). We show in Figure 5.10 our inferred cumulative mass profile along with several literature values for the cluster mass. Our profile is in good agreement with the estimate of the enclosed mass at 1 pc from Prager et al. (2017) and our total mass estimate of $0.67_{-0.04}^{+0.06} \times 10^6 M_{\odot}$ is in good agreement with the total masses of Baumgardt & Hilker (2018) and Baumgardt et al. (2019b) given that the true uncertainty on our inferred mass from multimass modeling is likely closer to 10% (see Section 3 in Dickson et al. 2024). The masses inferred from kinematics and dynamical modeling, including our own value, are a factor of 2 – 3 times smaller than the mass inferred by Lanzoni et al. (2010) from photometry. A lower present-day mass for Terzan 5 potentially has many important implications, in particular for studies that seek to model the star-formation and chemical enrichment histories of this system (e.g. Romano et al. 2023).




Figure 5.10: The inferred cumulative mass profile of Terzan 5 from our `Ter5-AllData` fit. We show the median values as a solid line (blue) while the shaded regions represent the 1σ and 2σ credible intervals of the model fit. The enclosed mass at 1 pc measured by Prager et al. (2017) is shown with an orange square. We show our inferred total mass and the values from Lanzoni et al. (2010), Baumgardt & Hilker (2018) and Baumgardt et al. (2019b) as horizontal dashed lines where the shaded regions represent the 1σ credible intervals.

5.2.3 Comparison of mass in BHs to literature results

Our models allow us to place strong constraints on the mass in BHs in both 47 Tuc and Terzan 5. We compare our results with other studies that investigate the BH population in these clusters in Table 5.3.

47 Tuc is a well-studied cluster with many previous works investigating its BH content, using a variety of methods. In general, we see that recent works which employ Monte Carlo models (Weatherford et al. 2020; Ye et al. 2022) generally infer larger BH populations (with larger uncertainties) despite taking very different approaches while works employing equilibrium models favor somewhat smaller populations of BHs. Our models place an upper limit (99th percentile) on the total mass in BHs in 47 Tuc of $649 M_{\odot}$, very similar the results of Dickson et al. (2024) who did not consider pulsar timing data. Our results are also in very good agreement with the upper limit of $578 M_{\odot}$ on the mass of a central IMBH reported by Della Croce et al. (2024). These authors employ action-based distribution function models to derive an upper limit on the mass of a putative IMBH in 47 Tuc, which as discussed earlier (see discussion in Section 4.1), is expected to have similar dynamical effects to a compact central cluster of BHs. Our inferred upper limit on the mass in BHs corresponds to an upper limit of 0.07% on the BH mass fraction f_{BH} , slightly lower than the upper limit of 0.09% found by Della Croce et al. (2024)⁴. These upper limits significantly limit the room for a more massive potential IMBH in 47 Tuc, which we discuss in the following section.

Terzan 5 on the other hand, has only a single estimate of its BH content reported in the literature, by Prager et al. (2017). These authors, also fitting on pulsar data, test for the presence of an IMBH in the center of 47 Tuc, reporting marginal evidence of a $\sim 500 M_{\odot}$ IMBH and an upper limit of $30\,000 M_{\odot}$. Our posterior for the mass in BHs in Terzan 5 is peaked near zero

⁴Our larger mass in BHs corresponds to a smaller f_{BH} because we infer a larger total mass for 47 Tuc with our multimass models than Della Croce et al. (2024) with their single-mass models, in line with the tendency for single-mass equilibrium models to underestimate the total cluster mass (see Hénault-Brunet et al. 2019).

Table 5.3: Reported masses in BHs (or upper limits) in 47 Tuc and Terzan 5 from dynamical studies in the literature and from this work. Note that we have scaled the BH mass fraction reported by Weatherford et al. (2020) to our inferred total mass to facilitate comparison. We note that the uncertainties on the black hole content reported by studies that employ LIMEPY models (Hénault-Brunet et al. 2020; Dickson et al. 2024 and this work) are likely underestimated by a factor of 2 – 3 (see discussion in Dickson et al. 2024). Listed uncertainties correspond to the 1σ uncertainties.

†Note that the the posterior of Hénault-Brunet et al. (2020) is peaked towards zero.

*These are upper limits on the mass of a putative IMBH, not the total mass in BHs.

Study	47 Tuc	Terzan 5
Hénault-Brunet et al. (2020)†	$430^{+386}_{-301} M_{\odot}$	
Weatherford et al. (2020)	$1037^{+1640}_{-922} M_{\odot}$	
Ye et al. (2022)	$\sim 2375 M_{\odot}$	
Della Croce et al. (2024)*	$< 578 M_{\odot}$	
Dickson et al. (2024)	$420^{+150}_{-80} M_{\odot}$	
This work	$446^{+75}_{-72} M_{\odot}$	$< 3860 M_{\odot}$
Prager et al. (2017)*		$< 30\,000 M_{\odot}$

and the 99th percentile upper limit is $3860 M_{\odot}$ ($f_{\text{BH}} < 0.6\%$), representing an improvement of nearly one order of magnitude on the existing constraints on the mass in BHs.

5.2.4 An IMBH in 47 Tuc?

GCs have long been suggested to host IMBHs. For proposed detections in multiple clusters using various techniques, see e.g. Gerssen et al. (2002), Noyola et al. (2008), Jalali et al. (2012), Lützgendorf et al. (2015), Kamann et al. (2016), Baumgardt (2017), Kızıltan et al. (2017a), Paduano et al. (2024), and Häberle et al. (2024). Many of these possible detections have however been contested. Typically, follow-up studies find that the dynamical effects of a central IMBH and other ingredients like a large population of centrally concentrated dark stellar remnants (including stellar-mass black holes - BHs) are highly degenerate, making the detection of a non-accreting IMBH very difficult (see e.g. McNamara et al. 2003; Van Der Marel & Anderson 2010; Freire et al. 2017; Zocchi et al. 2017, 2019; Gieles et al. 2018; Baumgardt et al. 2019c; Mann et al. 2019, 2020; Hénault-Brunet et al. 2020). Ruling out the presence of an IMBH is also very

difficult, with most works placing only an upper limit on the dark mass in the core of a cluster (e.g. Mann et al. 2019; Häberle et al. 2021; Della Croce et al. 2024). Häberle et al. (2024) find seven stars with velocities above the inferred escape velocity in the inner 0.08 pc of ω Cen, and they conclude that this implies the presence of an IMBH with a mass $\gtrsim 8200 M_{\odot}$. We note that only two of those seven stars are above the escape velocity of the mass model of Dickson et al. (2023), but they are the two closest to the cluster center and remain difficult to reconcile with explanations other than an IMBH, making this the strongest case for an IMBH in a Galactic GC so far.

Particularly relevant to this work are the results from Kızıltan et al. (2017a,b) and Paduano et al. (2024) who each report evidence for an IMBH in 47 Tuc. Kızıltan et al. (2017a) reported evidence for a $2300 M_{\odot}$ IMBH in the center of 47 Tuc on the basis of pulsar timing measurements. In order to constrain the mass of a central IMBH, they compared snapshots from a grid of N -body models (with and without an IMBH) to surface density and velocity dispersion profiles. The best-fitting models with and without an IMBH were then compared to measurements of pulsar accelerations (based on period derivatives) due to the cluster’s gravitational potential, and the set of models with a central IMBH was found to be more consistent with these measurements.

Critically, several follow-up studies identified limitations and possible issues with the analysis of Kızıltan et al. (2017a) (see Freire et al. 2017; Mann et al. 2019, 2020; Hénault-Brunet et al. 2020). Among the issues raised was the assumption of a short cluster distance of 4 kpc⁵, the lack of primordial binaries in their N -body models, and the use of a grid of isolated N -body models that have a much steeper present-day mass function than the bottom-light mass function that is observed for 47 Tuc, all of which could affect the inferred amount of dark mass in the cluster core.

We note that models of 47 Tuc that do not include a central IMBH have been shown to

⁵The latest estimates of the distance to 47 Tuc based on *Gaia* data place it 4.52 ± 0.03 kpc away (Baumgardt & Vasiliev 2021).

satisfyingly reproduce its velocity dispersion profile, number density profile, and stellar mass function data (Baumgardt et al. 2019a; Hénault-Brunet et al. 2020; Dickson et al. 2023, 2024). These studies, however, either did not consider the pulsar data, or they checked that the models were consistent with the pulsars’ maximum accelerations and projected radial distribution but did not directly incorporate the pulsar timing data in the fitting process.

Our results allow us to address the claims of Kızıltan et al. (2017a) since we are directly fitting our models to the pulsar timing data. As mentioned previously, our fit to all the available 47 Tuc data (47Tuc-AllData) yields an upper limit of $649 M_{\odot}$ in BHs ($\sim 0.07\%$ of the total cluster mass). This argues against the conclusions of Kızıltan et al. (2017a) that the pulsar data requires a central IMBH of $2300 M_{\odot}$ to explain the observed values of \dot{P}/P . We performed a fit of 47 Tuc with the distance set to 4 kpc, resulting in a noticeably worse quality of fit but also a larger ($\sim 2400 M_{\odot}$) population of BHs, perhaps indicating that the low distance assumed by Kızıltan et al. (2017a) contributes to the differences in our inferred dark mass in 47 Tuc.

The potential IMBH detection reported by Paduano et al. (2024) has a much wider range of possible masses, with the nominal 1σ uncertainty range spanning $56 - 6000 M_{\odot}$. The results of Della Croce et al. (2024), who infer an upper limit on the mass of an IMBH in 47 Tuc of $578 M_{\odot}$ reduce this range substantially. The fact that our inferred upper limit on the mass in BHs is consistent with the upper limit on the mass of an IMBH from Della Croce et al. (2024) highlights the fact there is very little room for a significant dark mass in the center of 47 Tuc.

While the effects of an IMBH and a centrally concentrated population of BHs are expected to be similar, they are not identical because the dark mass is more centrally concentrated if it is in an IMBH. This means that an upper limit on the mass in BHs does not necessarily correspond directly to the upper limit on the mass of an IMBH. Despite this, we expect that in our models, these effects are likely degenerate due to the spatial regions probed by our datasets. The half-mass radius of the BH population in our 47Tuc-AllData fit is ~ 0.1 pc, which is more

centrally concentrated than the majority of the pulsars in this cluster. While we do have a few pulsars inside this radius, the pulsars with large negative period derivatives (those that provide the most stringent constraints on the enclosed mass) are located outside of this radius meaning we are largely insensitive to the spatial extent and concentration of the mass in BHs in the central regions of the cluster.

We note that given the central line-of-sight velocity dispersion of 47 Tuc, the radius of influence of an IMBH of $\sim 600 M_{\odot}$ is $\lesssim 0.02$ pc which at the distance of 47 Tuc corresponds to $\lesssim 0.8''$, a much smaller region than is probed by either the pulsar data or the stellar kinematic data. This means that any future work seeking to use stellar kinematics to understand the nature of the dark mass in the core of this cluster would need to precisely measure the velocities of individual stars in the central arcsecond of 47 Tuc.

5.2.5 Central velocity dispersion of Terzan 5

The existing stellar kinematic data for Terzan 5 does not probe the central regions of the cluster, but the pulsar timing data allows us to independently predict the central velocity distribution of bright stars, without directly relying on any stellar kinematics. This also allows us to validate the existing observed stellar kinematics for a cluster like Terzan 5 where bulge contamination makes membership selection especially difficult. Our predicted central dispersion for Terzan 5 based on fit `Ter5-NoKinNoMF` is $15.7^{+0.5}_{-0.4}$ km s⁻¹. Future work seeking to constrain the central kinematics of this cluster can use this prediction as a benchmark to which new measurements can be compared. For example, a cusp in the central velocity dispersion within the radius of influence of an IMBH should manifest as a velocity dispersion larger than this predicted central velocity dispersion given that our models do not contain an IMBH.

5.2.6 Terzan 5: a comparison with ω Cen

The formation of Terzan 5 and its multiple populations is a topic of much debate in the literature, with suggested formation channels ranging from the stripped core of an accreted dwarf galaxy (Ferraro et al. 2009, 2016) to a surviving fragment of the proto-bulge (Ferraro et al. 2009, 2016; Taylor et al. 2022) to the product of a collision between a typical GC and a giant molecular cloud or young massive cluster (McKenzie & Bekki 2018; Pfeffer et al. 2021; Bastian & Pfeffer 2022).

While it is generally accepted that the metal-rich population in Terzan 5 is too enriched to have formed in a low-mass dwarf galaxy (e.g. Bastian & Pfeffer 2022), we can further investigate the dynamical evolution of this object by comparing Terzan 5 to ω Cen, a well-studied GC which is frequently suggested to be the core of an accreted dwarf galaxy (e.g. Pfeffer et al. 2021). ω Cen is the most massive GC in the Milky Way (e.g. Harris 1996, 2010 edition) and hosts multiple stellar populations with a large spread in iron abundance (e.g. Johnson & Pilachowski 2010; Bellini et al. 2017). Several studies have presented dynamical models of ω Cen which suggest that the cluster is host to a very large population of black holes (e.g. Zocchi et al. 2019; Baumgardt et al. 2019c; Dickson et al. 2024) typically making up about 5% of the total cluster mass (consistent with having retained almost all its BHs). The fact that we infer a much smaller population of BHs and indeed rule out a population larger than 0.6% of the total cluster mass in Terzan 5 suggests a different evolutionary history for this object. We note that ω Cen is more massive than Terzan 5 by a factor of $\sim 3 - 5$, and has a factor of ~ 5 larger r_h , which both contribute to a higher retention of BHs by the present day for ω Cen through an order of magnitude longer relaxation time (diminishing the importance of dynamical ejections of BHs throughout the evolution of the cluster). Additionally, the higher metallicity of Terzan 5 compared to ω Cen is expected to produce lower-mass BHs from the same initial progenitor masses (e.g. Fryer et al. 2012; Banerjee et al. 2020), increasing the number of BHs lost to natal kicks and reducing the

fraction of the cluster mass in BHs resulting from a given IMF.

While the origin of Terzan 5 remains uncertain, the discussion above highlights that further studies of the evolution of Terzan 5 along with its BH population could help to shed light on the initial conditions and formation of this peculiar system.

5.3 Conclusions

In this work, we presented a method to directly fit multimass dynamical models of GCs to pulsar timing data, allowing to infer the mass distribution of clusters. We applied our method to 47 Tuc, a well-studied cluster with a wealth of conventional stellar kinematic and local stellar mass function data as well as a large population of pulsars. We use this cluster as a benchmark by which we evaluate the performance of our method. We then applied our method to Terzan 5, a bulge cluster host to the largest population of pulsars of any Milky Way GC and lacking in conventional stellar kinematic data. Our main conclusions are as follows:

1. For clusters like 47 Tuc and Terzan 5 that are host to large populations of pulsars, the timing solutions of these pulsars can place similar constraints on the mass distribution and dynamics of their host cluster as conventional stellar kinematics. We demonstrate that the pulsar timing data allows us to accurately predict held-out stellar kinematic data and place strong constraints on the BH content of clusters.
2. We infer new and improved values for the mass and structural parameters of Terzan 5, finding a total mass of $0.67_{-0.04}^{+0.06} \times 10^6 M_{\odot}$, and a (3D) half-mass radius of $2.1_{-0.2}^{+0.3}$ pc. This mass is consistent with other dynamical estimates but is smaller by a factor of 2 – 3 than the estimate derived from photometry.
3. We refine existing constraints on the BH content of 47 Tuc and lower the existing upper limit on the mass in BHs of Terzan 5 by an order of magnitude. We infer the presence of

$446_{-72}^{+75} M_{\odot}$ in BHs in 47 Tuc and place an upper limit on the mass in BHs in Terzan 5 of $3860 M_{\odot}$.

4. Our results do not support the $\sim 2300 M_{\odot}$ IMBH reported by Kızıltan et al. (2017a) in the center of 47 Tuc on the basis of pulsar timing data, as we instead infer an upper limit of $649 M_{\odot}$ in BHs in this cluster, representing $\sim 0.07\%$ of the total cluster mass. This adds to several follow-up studies that refuted this original claim, but it is the first time that pulsar timing data, a crucial component of the Kızıltan et al. (2017a) study, is revisited as a direct constraint on the dynamical models.
5. We predict the central velocity dispersion of Terzan 5, independently of any stellar kinematic data, finding a dispersion of $15.7_{-0.4}^{+0.5} \text{ km s}^{-1}$. This prediction provides a baseline to which future work seeking to measure the central kinematics of this cluster can be compared.

The next generation of radio telescopes is expected to dramatically increase the number of detected pulsars in GCs (e.g. Hessels et al. 2015; Ridolfi et al. 2021; Chen et al. 2023; Berteaud et al. 2024), allowing us to apply our methodology to clusters that are not currently known to host large numbers of pulsars.

Perhaps the most promising candidate to host a large number of undiscovered pulsars is Liller 1, a bulge cluster that is qualitatively very similar to Terzan 5 (e.g. Ferraro et al. 2021). Liller 1 is a massive cluster ($\sim 1 \times 10^6 M_{\odot}$, Baumgardt et al. 2019b) that is similarly compact to Terzan 5 and has been found to have a similarly high stellar encounter rate (Saracino et al. 2015). This high stellar encounter rate, combined with strong gamma-ray emission detected from this cluster, suggests that it may be host to hundreds of undiscovered MSPs (Tam et al. 2011).

Like Terzan 5, Liller 1 is difficult to observe due to its location, with bulge contamination and strong differential reddening (e.g. Pallanca et al. 2021) making the collection of conventional

stellar kinematic data very difficult. As we have shown for Terzan 5, pulsars present a unique opportunity to investigate the internal dynamics of even heavily obscured clusters.

Acknowledgments

We thank the anonymous referee for helpful comments and suggestions that improved the quality of this work.

We thank Barbara Lanzoni for kindly sharing with us the number density profile of Terzan 5 presented in Lanzoni et al. (2010).

VHB acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) through grant RGPIN-2020-05990, and a New Faculty Grant from the Faculty of Graduate Studies and Research of Saint Mary's University. ND is grateful for the support of the Durland Scholarship in Graduate Research. MG acknowledges financial support from the grants PID2021-125485NB-C22, EUR2020-112157, CEX2019-000918-M funded by MCIN/AEI/10.13039/501100011033 (State Agency for Research of the Spanish Ministry of Science and Innovation) and SGR-2021-01069 (AGAUR).

This work made extensive use of Paulo Freire's database of pulsars in GCs (<https://www3.mpifr-bonn.mpg.de/staff/pfreire/GCpsr.html>).

This research was enabled in part by support provided by ACENET (www.ace-net.ca) and the Digital Research Alliance of Canada (<https://alliancecan.ca>).

5.4 Supplementary Material

5.4.1 Incorporating dispersion measures

Given a description of the gas distribution within a cluster (e.g. Abbate et al. 2018) Equation 4.6 can be modified to use the dispersion measure of a pulsar to form the line-of-sight probability

distribution instead of using the model line-of-sight density profile:

$$P(a_{\text{cl},z} | R_i) \propto \frac{dm}{da_{\text{cl},z}} = \frac{dm}{dz} \left| \frac{dz}{da_{\text{cl},z}} \right| = \frac{P(z | z_{\text{DM},i}, \sigma_{z,\text{DM},i})}{\left| \frac{da_{\text{cl},z}}{dz} \right|}, \quad (5.1)$$

where we have replaced the $\frac{dm}{dz}$ term with $P(z | z_{\text{DM}}, \sigma_{z,\text{DM}})$, the probability distribution of line-of-sight positions z given the predicted line-of-sight position from the dispersion measure z_{DM} and its uncertainty $\sigma_{z,\text{DM}}$. The values of $z_{\text{DM},i}$ and $\sigma_{z,\text{DM}}$ can be calculated from the an individual dispersion measure DM_i , n_g and DM_c as $z_{\text{DM},i} = (\text{DM}_i - \text{DM}_c)/n_g$ (under the assumption of a uniform gas distribution) with $\sigma_{z,\text{DM}}$ following from Gaussian error propagation. $P(z | z_{\text{DM},i}, \sigma_{z,\text{DM},i})$ is then a Gaussian probability density function, centered at $z_{\text{DM},i}$ with a dispersion of $\sigma_{z,\text{DM}}$:

$$P(z | z_{\text{DM},i}, \sigma_{z,\text{DM},i}) = \frac{1}{\sigma_{z,\text{DM},i} \sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{z - z_{\text{DM},i}}{\sigma_{z,\text{DM},i}}\right)^2\right). \quad (5.2)$$

5.4.2 Pulsar Data

Table 5.4: Pulsar timing data used in this work for 47 Tuc. The columns indicate the pulsar ID, projected radius, spin period, spin-period derivative, the uncertainty on the spin-period derivative, the orbital period, the orbital period derivative, the uncertainty on the orbital period derivative, and the proper motion in R.A. and Dec. with associated uncertainties. In the rightmost column we indicate the work in which the timing solution was derived. Timing References: F17: Freire et al. (2017), R16: Ridolfi et al. (2016), FR18: Freire & Ridolfi (2018).

Pulsar ID	R [arcmin]	P [ms]	\dot{P} [s/s]	$\Delta\dot{P}$ [s/s]	P_b [days]	\dot{P}_b [s/s]	$\Delta\dot{P}_b$ [s/s]	μ_{α^*} [mas yr ⁻¹]	μ_{δ} [mas yr ⁻¹]	Ref.
J0024-7204C	1.2298	5.75678	-4.99e-20	2e-24	–	–	–	5.2 ± 0.1	–3.1 ± 0.1	F17
J0024-7204D	0.6483	5.35757	-3.42e-21	9e-25	–	–	–	4.24 ± 0.07	–2.24 ± 0.05	F17
J0024-7204E	0.6205	3.53633	9.85e-20	5e-25	2.2568483	4.8e-12	2e-13	6.15 ± 0.03	–2.35 ± 0.06	F17
J0024-7204F	0.2149	2.62358	6.45e-20	7e-25	–	–	–	4.52 ± 0.08	–2.50 ± 0.05	F17
J0024-7204G	0.2781	4.04038	-4.22e-20	2e-24	–	–	–	4.5 ± 0.1	–2.9 ± 0.1	F17
J0024-7204H	0.7677	3.21034	-1.83e-21	1e-24	2.357696895	-7e-13	6e-13	5.1 ± 0.2	–2.8 ± 0.2	F17
J0024-7204I	0.2772	3.48499	-4.59e-20	2e-24	0.2297922489	-8e-13	2e-13	5.0 ± 0.2	–2.1 ± 0.2	F17
J0024-7204J	1.0185	2.10063	-9.79e-21	9e-25	–	–	–	5.27 ± 0.06	–3.59 ± 0.09	F17
J0024-7204L	0.1627	4.34617	-1.22e-19	1e-24	–	–	–	4.4 ± 0.2	–2.4 ± 0.2	F17
J0024-7204M	1.0688	3.67664	-3.84e-20	5e-24	–	–	–	5.0 ± 0.3	–2.0 ± 0.4	F17
J0024-7204N	0.4793	3.05395	-2.19e-20	2e-24	–	–	–	6.3 ± 0.2	–2.8 ± 0.2	F17
J0024-7204O	0.0806	2.64334	3.03e-20	6e-25	–	–	–	5.01 ± 0.05	–2.58 ± 0.08	F17
J0024-7204Q	0.9502	4.03318	3.4e-20	6e-25	1.1890840496	-1e-12	2e-13	5.2 ± 0.1	–2.6 ± 0.1	F17
J0024-7204R	0.1519	3.48046	1.48e-19	3e-24	0.06623147751	1.9e-13	4e-14	4.8 ± 0.1	–3.3 ± 0.2	F17
J0024-7204S	0.215	2.83041	-1.21e-19	1e-24	1.2017242354	-4.9e-12	4e-13	4.5 ± 0.1	–2.5 ± 0.1	F17
J0024-7204T	0.3179	7.58848	2.94e-19	1e-23	1.126176771	2.5e-12	1.1e-12	5.1 ± 0.6	–2.6 ± 0.7	F17
J0024-7204U	0.9386	4.34283	9.52e-20	2e-24	0.42910568324	6.6e-13	5e-14	4.6 ± 0.2	–3.8 ± 0.1	F17
J0024-7204W	0.087	2.35234	-8.66e-20	1e-24	–	–	–	6.1 ± 0.5	–2.6 ± 0.3	R16
J0024-7204X	3.828	4.77152	1.84e-20	7e-25	10.921183545	6e-12	2e-12	5.8 ± 0.1	–3.3 ± 0.2	R16
J0024-7204Y	0.3743	2.19666	-3.52e-20	8e-25	0.5219386107	-8.2e-13	7e-14	4.4 ± 0.1	–3.4 ± 0.1	F17
J0024-7204Z	0.1506	4.55445	-4.56e-21	1e-22	–	–	–	4 ± 2	1 ± 2	F17
J0024-7204aa	0.465	1.84538	-4.59e-20	1.5e-23	–	–	–	4.6 ± 0.8	–4.6 ± 1.3	FR18
J0024-7204ab	0.2092	3.70464	9.82e-21	8e-24	–	–	–	4.2 ± 0.6	–2.9 ± 0.5	F17

Table 5.5: Same as Table 5.4 but for Terzan 5. Timing References: L00: Lyne et al. (2000), R05: Ransom et al. (2005), P17: Prager et al. (2017), C18: Cadelano et al. (2018), A18: Andersen & Ransom (2018), R21: Ridolfi et al. (2021), P24: Padmanabh et al. (2024).

Pulsar ID	R [arcmin]	P [ms]	\dot{P} [s/s]	$\Delta\dot{P}$ [s/s]	P_b [days]	\dot{P}_b [s/s]	$\Delta\dot{P}_b$ [s/s]	Ref.
J1748-2446C	0.179	8.4361	-6.06e-19	4e-21	–	–	–	L00
J1748-2446D	0.693	4.71398	1.3e-19	–	–	–	–	R05
J1748-2446E	0.361	2.1978	-1.8e-20	–	–	–	–	R05
J1748-2446F	0.125	5.54014	4e-21	–	–	–	–	R05
J1748-2446G	0.185	21.6719	3.9e-19	–	–	–	–	R05
J1748-2446H	0.227	4.92589	-8.3e-20	–	–	–	–	R05
J1748-2446I	0.03	9.57019	-7.1e-20	–	–	–	–	R05
J1748-2446J	0.948	80.3379	2.5e-18	–	–	–	–	R05
J1748-2446K	0.22	2.96965	-9.4e-20	–	–	–	–	R05
J1748-2446L	0.149	2.2447	-1.7e-20	–	–	–	–	R05
J1748-2446M	0.083	3.56957	4.9e-19	–	–	–	–	R05
J1748-2446N	0.154	8.6669	5.5e-19	–	–	–	–	R05
J1748-2446O	0.119	1.67663	-6.9e-20	–	–	–	–	R05
J1748-2446Q	0.36	2.812	-3.6e-20	–	–	–	–	R05
J1748-2446R	0.101	5.02854	4.7e-19	–	–	–	–	R05
J1748-2446S	0.249	6.11664	6.4e-20	–	–	–	–	R05
J1748-2446T	0.443	7.08491	3.1e-19	–	–	–	–	R05
J1748-2446U	0.148	3.28914	3e-19	–	–	–	–	R05
J1748-2446V	0.178	2.07251	-9.5e-20	–	–	–	–	R05
J1748-2446W	0.037	4.20518	1.2e-19	–	–	–	–	R05
J1748-2446X	0.488	2.99926	5.9e-20	–	–	–	–	R05
J1748-2446Y	0.056	2.04816	1.5e-19	–	–	–	–	R05
J1748-2446Z	0.033	2.46259	-8.6e-20	–	–	–	–	P17

Table 5.6: Same as Table 5.4 but for Terzan 5. Timing References: L00: Lyne et al. (2000), R05: Ransom et al. (2005), P17: Prager et al. (2017), C18: Cadelano et al. (2018), A18: Andersen & Ransom (2018), R21: Ridolfi et al. (2021), P24: Padmanabh et al. (2024).

Pulsar ID	R [arcmin]	P [ms]	\dot{P} [s/s]	$\Delta\dot{P}$ [s/s]	P_b [days]	\dot{P}_b [s/s]	$\Delta\dot{P}_b$ [s/s]	Ref.
J1748-2446aa	0.222	5.78804	-4.4e-19	–	–	–	–	P17
J1748-2446ab	0.038	5.11971	4.2e-19	–	–	–	–	P17
J1748-2446ac	0.337	5.08691	2.3e-19	–	–	–	–	P17
J1748-2446ae	0.032	3.65859	-5.7e-19	–	–	–	–	P17
J1748-2446af	0.145	3.30434	-2.3e-19	–	–	–	–	P17
J1748-2446ag	0.167	4.44803	1.2e-20	–	–	–	–	P17
J1748-2446ah	0.127	4.96515	5.7e-19	–	–	–	–	P17
J1748-2446ai	0.192	21.22838	1.4e-18	–	–	–	–	P17
J1748-2446aj	0.17	2.95891	1.41232e-19	6e-24	–	–	–	C18
J1748-2446ak	0.287	1.8901	8.8495e-20	6-24	–	–	–	C18
J1748-2446am	0.044	2.93382	-1.368e-19	3e-23	–	–	–	A18
J1748-2446an	0.201	4.802	1.55746e-19	6-24	–	–	–	R21
J1748-2446ao	0.156	2.27438	8.6979e-20	1e-24	57.5556	1.65e-10	9e-12	P24
J1748-2446ap	0.253	3.74469	3.07e-19	1e-24	21.3882	1.3e-10	2.1e-11	P24
J1748-2446aq	0.038	12.52194	-7.16198e-19	6e-24	–	–	–	P24
J1748-2446as	0.084	2.32646	2.559829e-19	6e-25	–	–	–	P24
J1748-2446at	0.123	2.18819	-5.89966e-20	4e-25	–	–	–	P24
J1748-2446au	0.053	4.54822	-1.06797e-19	2e-24	–	–	–	P24
J1748-2446av	0.045	1.84945	-4.25047e-20	2e-25	3.38166	-1e-11	2e-12	P24
J1748-2446aw	0.156	13.04908	1.306465e-18	3e-24	0.73138	5.92e-12	2e-14	P24
J1748-2446ax	0.161	1.9435	-9.5495e-21	7e-25	30.2088	-1.3e-11	9e-12	P24

5.4.3 Supplementary Figures

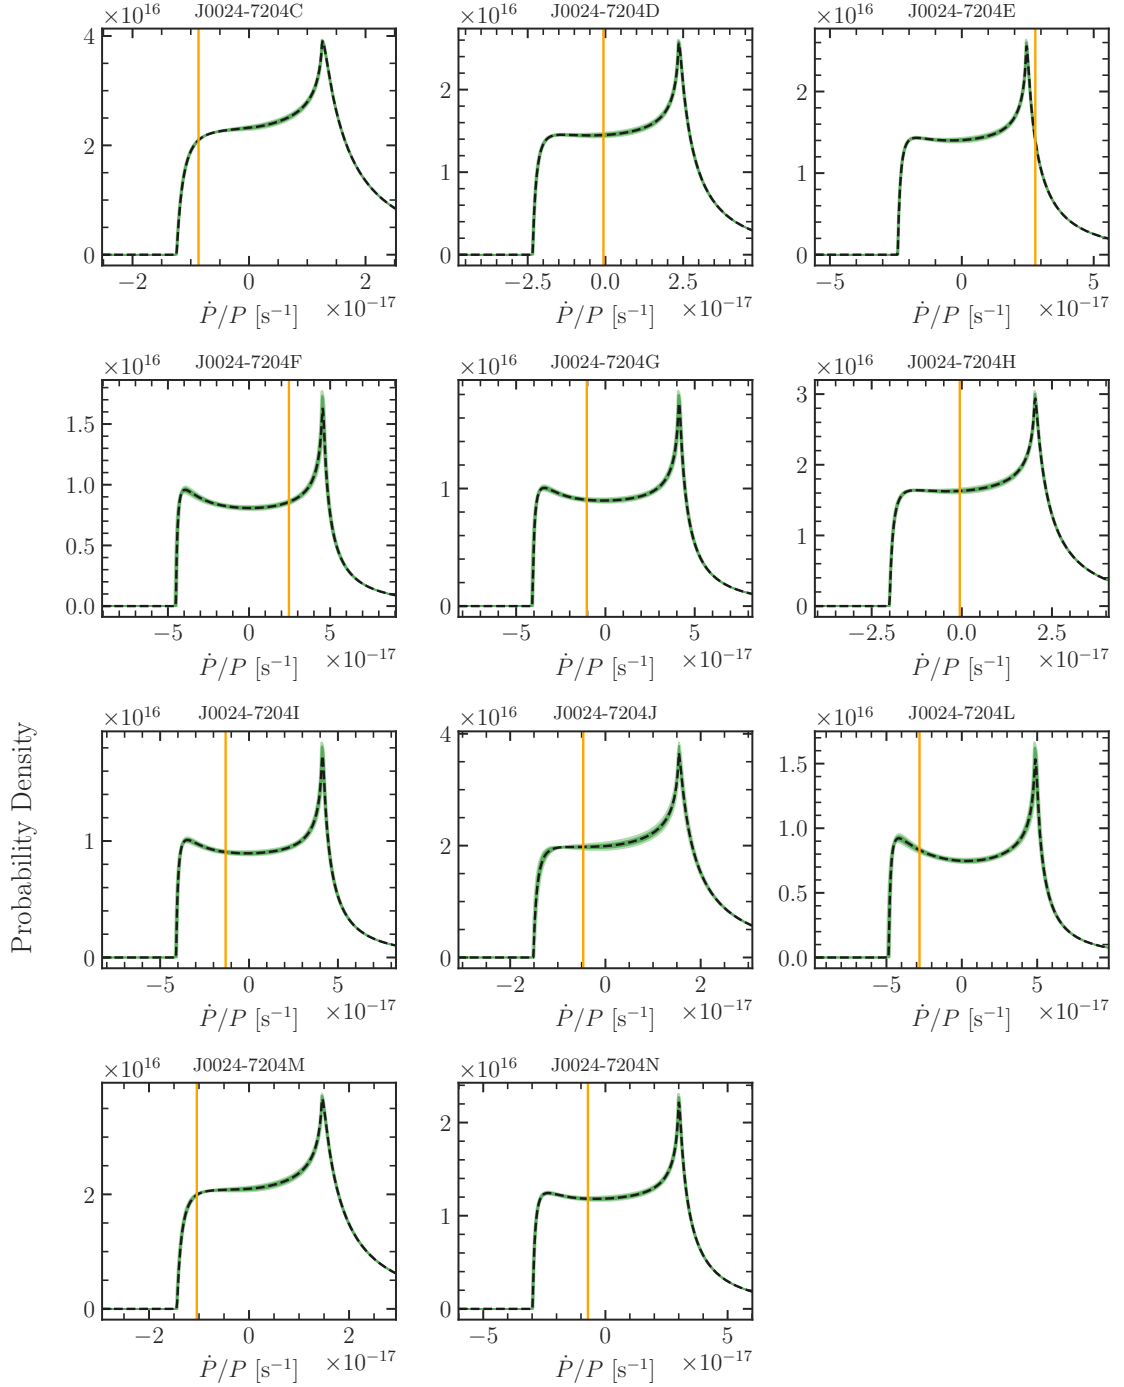


Figure 5.11: Likelihood functions corresponding to the best-fitting model (47Tuc-AllData) for the observed \dot{P}/P for each pulsar in 47 Tuc. In each panel, we show the observed period derivative as a vertical orange line.

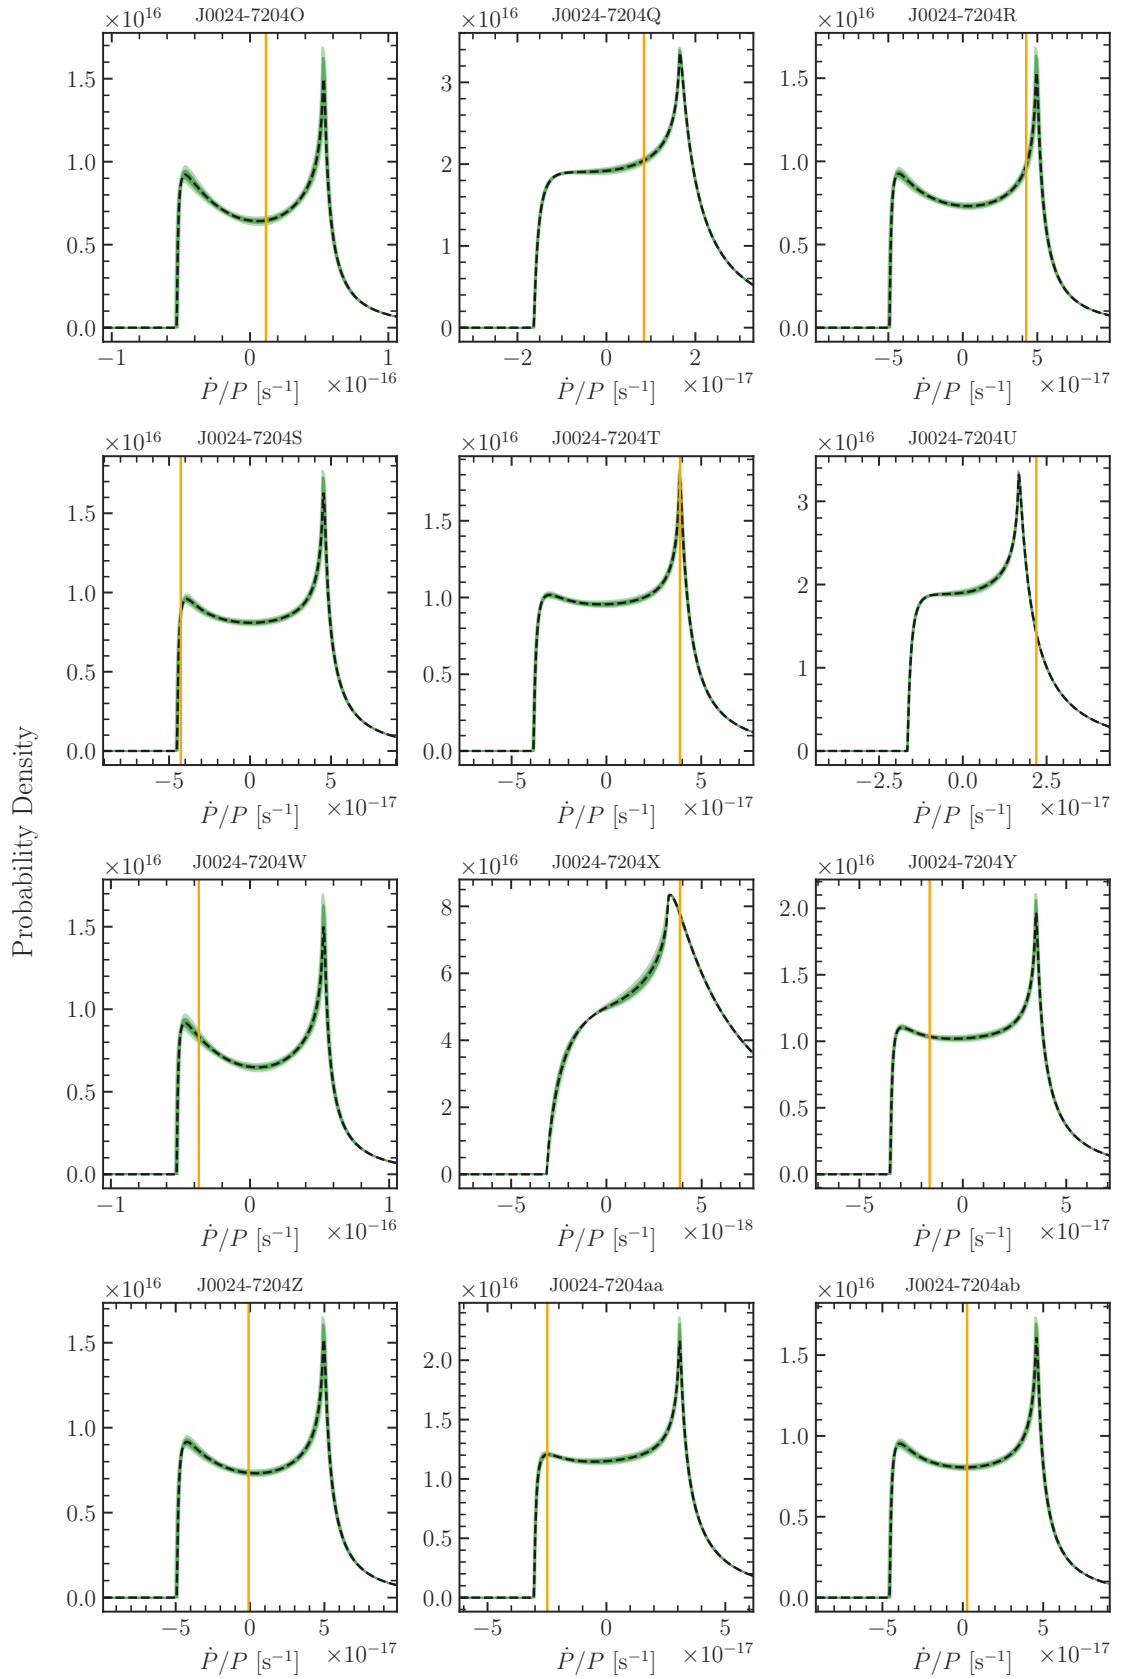


Figure 5.12: Continuation of Figure 5.11.

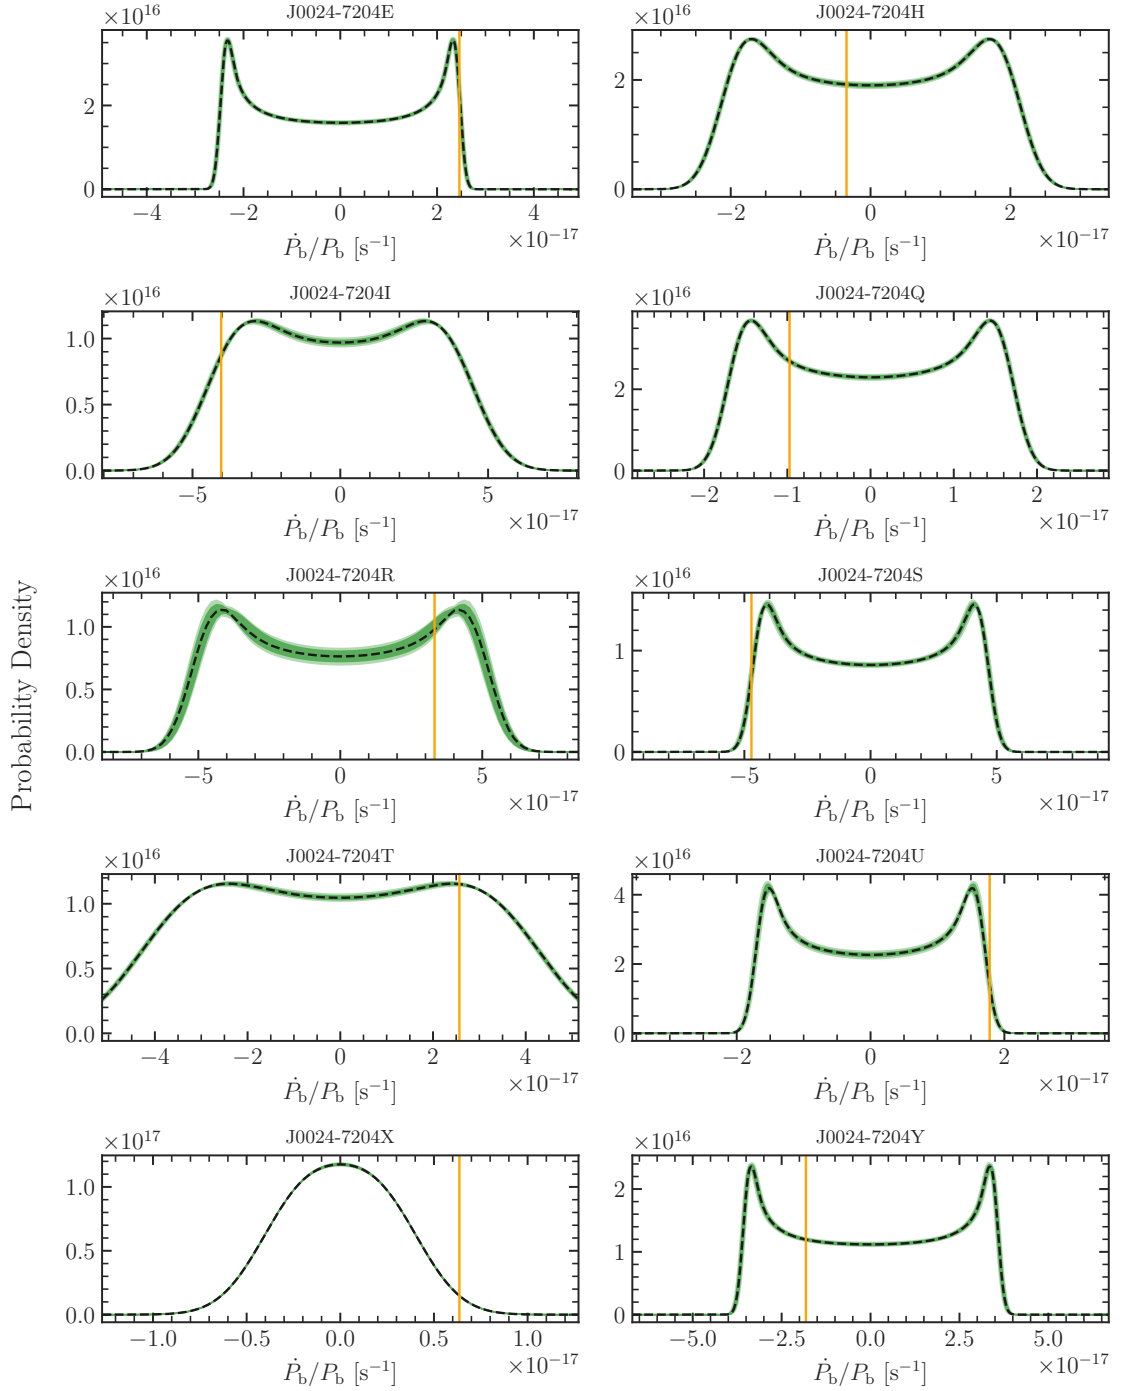


Figure 5.13: Likelihood functions corresponding to the best-fitting model (47Tuc-AllData) for the observed \dot{P}_b/P_b for each pulsar in 47 Tuc with an orbital timing solution. In each panel, we show the observed period derivative as a vertical orange line.

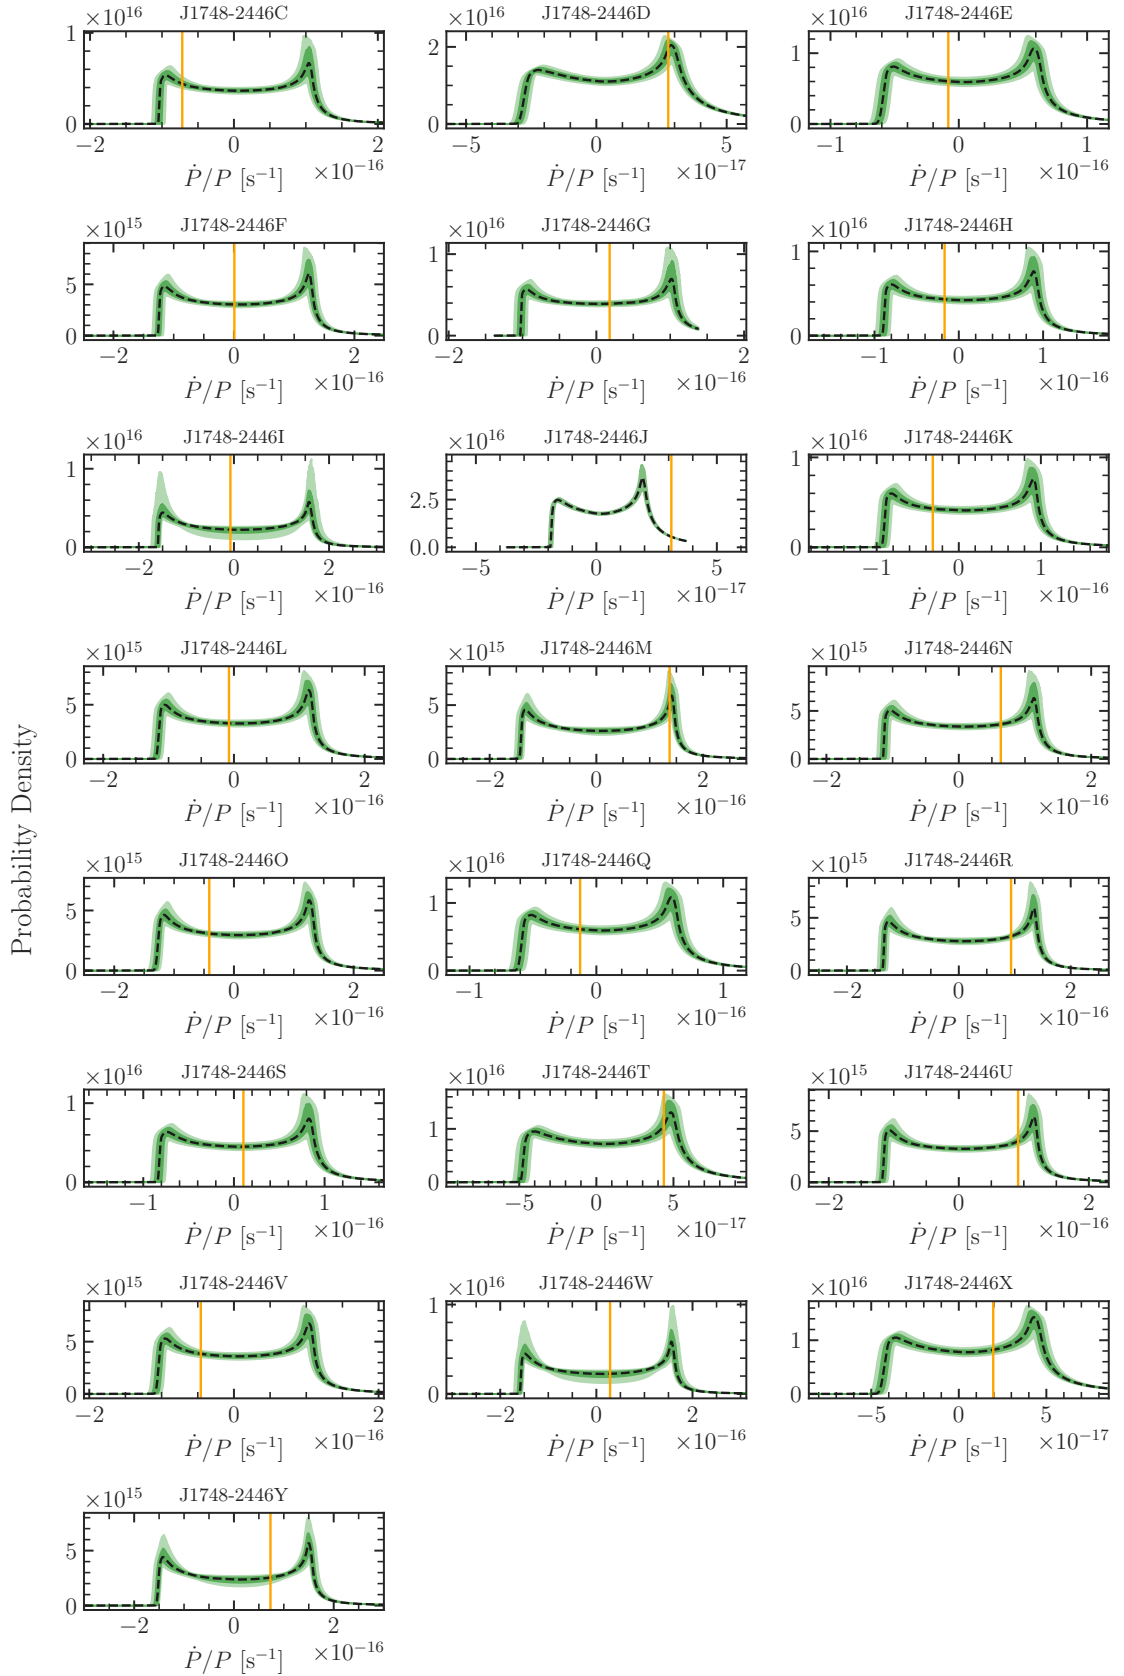


Figure 5.14: Likelihood functions corresponding to the best-fitting model (Ter5-AllData) for the observed \dot{P}/P for each pulsar in Terzan 5. In each panel, we show the observed period derivative as a vertical orange line.

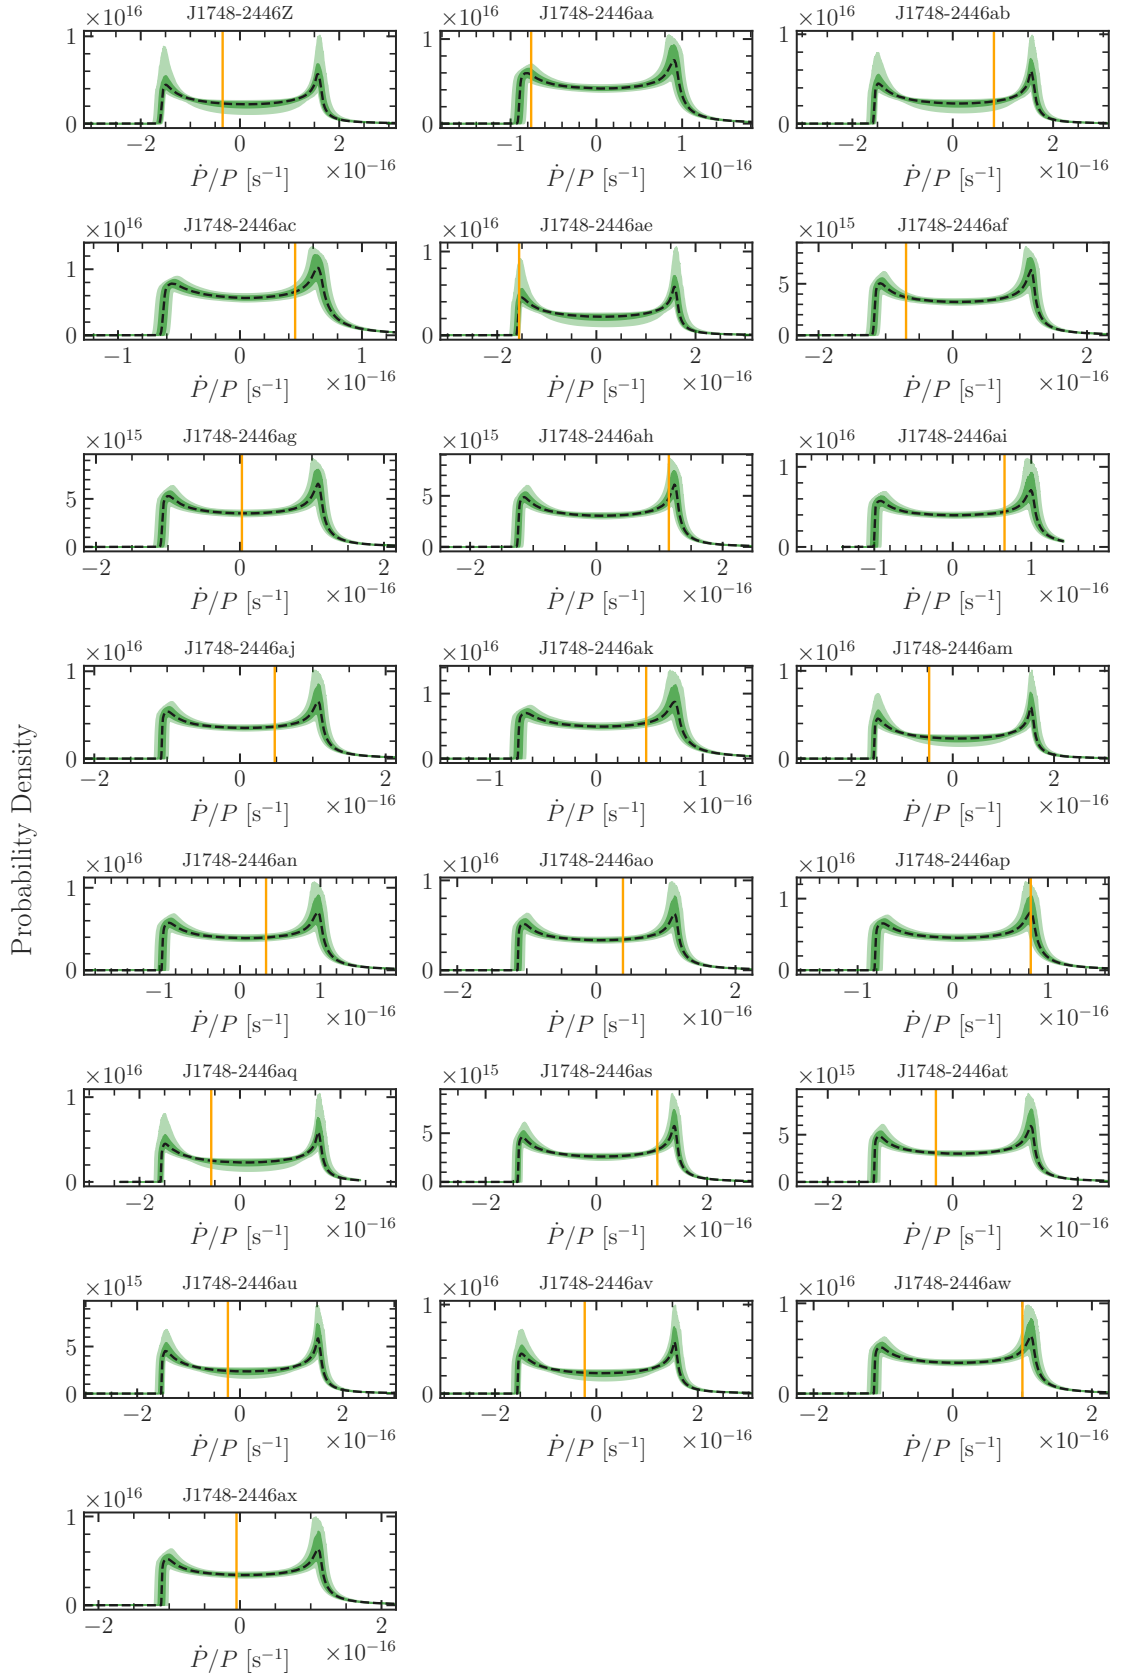


Figure 5.15: Continuation of Figure 5.14.




Figure 5.16: Likelihood functions corresponding to the best-fitting model (Ter5-AllData) for the observed \dot{P}_b/P_b for each pulsar in Terzan 5 with an orbital timing solution. In each panel, we show the observed period derivative as a vertical orange line.

Chapter 6

Conclusions

In this thesis we have examined the present-day BH populations in Galactic GCs. These BHs are expected to form through the evolution of massive stars in clusters, and the properties of the present-day populations have wide-ranging implications for many areas of astronomy (Section 1.1).

Among the array of possible methods for directly detecting BHs in GCs, we have briefly discussed X-ray and radio detections of accreting BHs (Section 1.2.1, e.g. Strader et al. 2012), radial velocity detections of BHs in binary systems (Section 1.2.2, e.g. Giesers et al. 2018), and the potential for gravitational lensing detections with future observatories (Section 1.2.3, e.g. Kiroğlu et al. 2022).

Of these methods, those that are expected to find the most BHs are X-ray and radio detections and radial velocity detections. Common to both of these methods is the significant caveat that the number of BHs we expect to detect via their accretion or binary companions is not a good indicator of the total number of BHs in a cluster (Kremer et al. 2018b,a). This is a significant issue as the total number of BHs in a cluster is the most important quantity for understanding the role of BHs in the dynamical evolution of the cluster, and for understanding the broader astrophysical implications of BHs in GCs.

6.1 Summary of indirect methods for detecting black holes

Given the limitations of direct detections, we have focused on indirect methods for detecting BHs in GCs. These methods rely on the dynamical effects of the BHs on the cluster as a whole, and are therefore sensitive to the total number of BHs in the cluster. Underlying each of the indirect methods we have discussed is the fact that the presence of BHs in a cluster has a significant effect on the dynamics of the cluster as a whole. This means that by observing the dynamics of the stars in a cluster, we can infer the presence of BHs in the cluster through the use of dynamical models (Section 1.3).

These dynamical signatures manifest themselves in a number of observable ways, each a consequence of the BHs acting as a heat source in the core of the cluster. In brief, BHs and especially BH binaries transfer energy to surrounding stars through dynamical interactions, leading to an expansion of the core of the cluster, an increase in the velocity dispersion of the core, and a decrease in the degree of mass segregation in a cluster (Section 1.3.1). Each of these effects can be used to infer the presence of BHs in a cluster, and to estimate the number of BHs in a cluster, when coupled with sufficiently detailed dynamical models.

There are several approaches to using these dynamical signatures to infer the presence of BHs in a cluster. The most direct approach is to compute a dynamical model of the cluster, accurately reproducing the observed properties of the cluster, and then inspecting the model to infer the BH population (Section 3.1, e.g. Kremer et al. 2018b; Ye et al. 2022; Dickson et al. 2024). This approach has been used in a number of studies, and we have shown in this thesis that it can be successful in inferring the presence of BHs in a cluster, given a sufficient array of observational data. This approach can be conducted with either evolutionary models or equilibrium models, with the former being more computationally expensive but providing the evolutionary history of the cluster, and the latter omitting the evolutionary history in favour of being fast

enough to be used in a Bayesian analysis.

The other approach to inferring the presence of BHs in a cluster is to compute a large grid of models (typically evolutionary models), designed to approximate the population of GCs, and then to look for correlations between observable properties of the models and their BH content (Section 3.2, Askar et al. 2018; Weatherford et al. 2020). This approach has the advantage of being applicable to a large number of clusters so long as they fall within the parameter space of the grid of models, without the need to compute dedicated models for individual clusters. This approach relies on the grid of models being sufficiently large and dense to accurately represent the population of Milky Way GCs, and on the ability to construct a robust relationship between the observable properties of a cluster and its BH content. We presented a discussion of these methods as well as their strengths and limitations in Chapter 3.

6.2 This thesis

6.2.1 Tests of equilibrium models with mock data

In this thesis we have focused on the use of multimass distribution function (LIMEPY) models to infer the presence of BHs in GCs. Through the use of mock data extracted from a large grid of Monte Carlo models, we have shown that these models can be used to accurately infer the presence of BHs in a cluster (Chapter 2, see also Dickson et al. 2024). We matched snapshots of the Monte Carlo models to real Milky Way clusters and extracted mock observations, designed to mimic the observational data available for real clusters. We then fit the LIMEPY models to the mock data using the same procedure as we do for real clusters, and compared the inferred BH populations to the true BH populations of the models. We examined the median BH mass fraction inferred by the models, and the associated uncertainties. Our main conclusions from this first part can be summarized as follows:

- When coupled with a mass-evolution algorithm, the LIMEPY models are able to accurately infer the BH content of a cluster from observational data, with the inferred BH mass fraction, f_{BH} falling along the one-to-one line with the true f_{BH} for most snapshots. We find no evidence of systematic biases in the inferred BH populations as a function of the true BH content of the cluster.
- The models have a tendency to significantly underestimate the BH content of some dynamically young clusters. We developed a criterion to efficiently identify these clusters and exclude them from the analysis.
- The uncertainties on the BH mass fraction are found to be underestimated by a factor of ~ 2.5 . Under the assumption that the estimates are not biased as a function of f_{BH} , we developed a simple correction to the uncertainties that reflects the true spread in the estimates of the LIMEPY models around the true values.

6.2.2 Pulsar timing constraints for dynamical modelling

In addition to the tests of the LIMEPY models with mock data, we have also developed a new method to include timing data from millisecond pulsars as additional constraints on dynamical models (Chapters 4 and 5). Pulsar timing provides one of the few opportunities in astrophysics to directly measure accelerations of individual objects, allowing us to directly probe the gravitational potential of the cluster.

While modern pulsar timing data gives extremely precise measurements of periods and period derivatives (both spin and orbital) for individual millisecond pulsars (Section 4.2), the observed period derivatives does not by itself provide a direct measurement of the acceleration of the pulsar. Instead, the observed period derivative is a combination of the intrinsic spin-down of the pulsar with several real or apparent accelerations that the pulsar is subject to. Because we

are only concerned with the acceleration due to the cluster potential, we need to account for the accelerations due to galactic potential, the apparent acceleration due to the proper motion of the pulsar, and the intrinsic spin-down of the pulsar (Section 4.3). Among these effects, the galactic acceleration and the apparent acceleration due to the proper motion of the pulsar are straightforward to calculate and account for while the intrinsic spin-down of the pulsar requires a more careful treatment. To marginalize over the range of possible intrinsic spin-downs, we take the field population of pulsars, which are only subject to the galactic acceleration and the apparent acceleration due to the proper motion of the pulsar, as a reference population. We then construct a Gaussian KDE over the observed period-period derivative space, allowing us to estimate the distribution of intrinsic spin-downs for a pulsar of a given period. Having corrected for the above-mentioned effects, we can then calculate a likelihood for the observed period derivative and fit dynamical models of the cluster to the pulsar timing data.

For this study, we build on the work we presented in Dickson et al. (2023, 2024), using the same models and fitting procedure, but with the addition of the pulsar timing data. We focus on two clusters, Terzan 5 and 47 Tuc, the clusters which host the two largest populations of millisecond pulsars in the Milky Way.

In general, we treat 47 Tuc as a test case for the method, as the cluster is well-studied and has an abundance of conventional stellar kinematic data. We run three fits to the cluster, one with all data included (`47Tuc-AllData`), one without the pulsar timing data (`47Tuc-NoPulsars`), and one with the pulsar timing data included, but the stellar kinematics excluded (`47Tuc-NoKin`). We find that the pulsar timing data is able to provide very strong constraints on the mass distribution and dynamics of the cluster, with the `47Tuc-NoKin` fit very accurately predicting the held-out stellar kinematic data. Looking at the inferred BH content, we see a similar agreement between the fits, with the `47Tuc-NoKin` fit resulting in a slightly wider posterior distribution for the BH fraction, but recovering a similar median value to the `47Tuc-AllData` fit.

In contrast to 47 Tuc, Terzan 5 is a cluster with very little conventional stellar kinematic data, owing primarily to its location in the Galactic bulge. The fact that this cluster is host to a large population of millisecond pulsars gives us a unique opportunity to probe the dynamics of a cluster that would otherwise be extremely challenging to model. We again run three fits to the cluster, one with all data included (`Ter5-AllData`), one without the pulsar timing data (`Ter5-NoPulsars`), and one with all stellar kinematics and stellar mass function data excluded (`Ter5-NoKinNoMF`). Comparing the three fits, we find that the fit without the pulsar timing data is unable to place strong constraints on the BH content, resulting in a wide posterior distribution for the BH fraction spanning several thousand solar masses.

Our conclusions can be summarized as follows:

- Pulsar timing data can provide strong constraints on the dynamics of GCs, and can be used to infer the presence of BHs in a cluster even when conventional stellar kinematic data is lacking.
- We are able to place a firm upper limit on the mass in BHs in Terzan 5 of $\sim 4000 M_{\odot}$, an order of magnitude lower than the previous estimate from the literature (Prager et al. 2017).
- We refine existing constraints on the BH content of 47 Tuc, find that a population of stellar-mass BHs with a total mass of $446^{+75}_{-72} M_{\odot}$ is preferred by the data. The presence of an IMBH in 47 Tuc of several thousand solar masses is not supported by our analysis, but a less massive IMBH of $\lesssim 600 M_{\odot}$ is still possible.
- We present a new measurement of the total mass of Terzan 5, finding a value of $0.67^{+0.06}_{-0.04} \times 10^6 M_{\odot}$, consistent with previous dynamical measurements (Prager et al. 2017; Baumgardt & Hilker 2018; Baumgardt et al. 2019b) but a factor of 2 – 3 lower than the photometric measurement from the literature (Lanzoni et al. 2010).

- We predict the central velocity dispersion of Terzan 5 to be $15.7^{+0.5}_{-0.4}$ km s⁻¹, providing a baseline for future efforts to measure the central kinematics of this cluster, which are most sensitive to the presence of BHs or a potential IMBH.

As the next generation of radio telescopes come online, the power of this method will only increase. The SKA is expected to provide a significant increase in the number of pulsars detected, many of which will be in GCs (Hessels et al. 2015). Adding to the census of pulsars in clusters will allow us to apply the method we have developed here to a larger number of clusters, providing a more complete picture of the dynamics of GCs and the populations of BHs within them. As we have seen in Section 5.2, all pulsars do not contribute equally to the dynamical constraints, and the method is most effective when we have pulsars with large negative period derivatives located close to the cluster centre. This means that discoveries of individual pulsars that match these criteria can by themselves be extremely valuable, even if the total number of pulsars in a cluster is not significant.

6.3 Future work

Much of the work presented in this thesis has the potential to be expanded upon in the future. The most obvious extension is to simply wait for the discovery and timing of new pulsars, and applying the methodology we have presented here to new clusters or new pulsars. Likewise, the framework we have developed for testing inferred BH populations with mock data can be applied to other methodologies, such as those discussed in Section 1.3. Beyond these immediate extensions, there are several areas of theoretical work that could be pursued, for both of the main topics of this thesis.

When testing inferred BH populations with mock data, we currently use the pre-existing public grid of CMC models presented in Kremer et al. (2020). The models included in this grid

cover a wide range of initial conditions, but the limitations of the grid are especially apparent towards the high-mass end of the cluster population. Additional parameters like the primordial binary fraction and the initial mass function of the cluster could be also varied to provide a more complete picture of the parameter space of Milky Way GCs. Beyond changes to the initial conditions of the models, the models themselves have the inherent limitation of spherical symmetry and do not include cluster rotation (a property shared by the LIMEPY models). While most clusters are not strongly rotating (Sollima et al. 2019; Vasiliev & Baumgardt 2021), the effects of modelling a rotating cluster with a non-rotating model could be explored. In order to do this, we would likely need to employ direct N -body models, which can include rotation.

For the pulsar timing work, one possible extension is to jointly consider the spin and orbital timing solutions of individual pulsars. In this thesis, we treat the spin and orbital solutions as separate data but, in reality, their period derivatives should be affected by the same accelerations. By comparing the two period derivatives, it is possible to constrain the true intrinsic spin-down of the pulsar (e.g. Padmanabh et al. 2024), which would allow us to forgo the current KDE-based method of marginalizing over the intrinsic spin-downs of the pulsars. This would allow us to more accurately infer the accelerations of individual pulsars, and would mean that pulsars with large positive period derivatives could be just as informative as those with large negative period derivatives. While the relative uncertainties on the orbital period timing solutions are larger than the spin period solutions, longer timing baselines should improve these measurements, allowing us to directly measure line-of-sight accelerations for individual pulsars and directly measure the gravitational potential within clusters for the first time.

Bibliography

- Abbate F., Possenti A., Ridolfi A., Freire P. C. C., Camilo F., Manchester R. N., D'Amico N., 2018, *Monthly Notices of the Royal Astronomical Society*, 481, 627
- Abbate F., Spera M., Colpi M., 2019a, *Monthly Notices of the Royal Astronomical Society*, 487, 769
- Abbate F., Possenti A., Colpi M., Spera M., 2019b, *The Astrophysical Journal*, 884, L9
- Abbott B. P., et al., 2016, *The Astrophysical Journal*, 818, L22
- Andersen B. C., Ransom S. M., 2018, *The Astrophysical Journal*, 863, L13
- Antonini F., Gieles M., Dosopoulou F., Chattopadhyay D., 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 466
- Arca Sedda M., Askar A., Giersz M., 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 4652
- Aros F. I., Vesperini E., 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 3136
- Askar A., Szkudlarek M., Gondek-Rosińska D., Giersz M., Bulik T., 2017, *Monthly Notices of the Royal Astronomical Society: Letters*, 464, L36
- Askar A., Arca Sedda M., Giersz M., 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 1844
- Bañares-Hernández A., Calore F., Camalich J. M., Read J. I., 2024, *New Constraints on the Central Mass Contents of Omega Centauri from Combined Stellar Kinematics and Pulsar Timing* ([arXiv:2408.00939](https://arxiv.org/abs/2408.00939))
- Bacon R., et al., 2010, *SPIE Proceedings*, 7735, 773508
- Bahramian A., et al., 2017, *Monthly Notices of the Royal Astronomical Society*, 467, 2199
- Bahramian A., et al., 2020, *The Astrophysical Journal*, 901, 57
- Balbinot E., Gieles M., 2018, *Monthly Notices of the Royal Astronomical Society*, 474, 2479
- Banerjee S., Belczynski K., Fryer C. L., Berczik P., Hurley J. R., Spurzem R., Wang L., 2020, *Astronomy & Astrophysics*, 639, A41
- Bastian N., Pfeffer J., 2022, *Monthly Notices of the Royal Astronomical Society*, 509, 614
- Baumgardt H., 2017, *Monthly Notices of the Royal Astronomical Society*, 464, 2174
- Baumgardt H., Hilker M., 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 1520

- Baumgardt H., Sollima S., 2017, *Monthly Notices of the Royal Astronomical Society*, 472, 744
- Baumgardt H., Vasiliev E., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 5957
- Baumgardt H., Sollima A., Hilker M., 2019a, *Proceedings of the International Astronomical Union*, 14, 400
- Baumgardt H., Hilker M., Sollima A., Bellini A., 2019b, *Monthly Notices of the Royal Astronomical Society*, 482, 5138
- Baumgardt H., et al., 2019c, *Monthly Notices of the Royal Astronomical Society*, 488, 5340
- Baumgardt H., Hénault-Brunet V., Dickson N., Sollima A., 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 3991
- Bellini A., Milone A. P., Anderson J., Marino A. F., Piotto G., van der Marel R. P., Bedin L. R., King I. R., 2017, *The Astrophysical Journal*, 844, 164
- Berteaud J., Eckner C., Calore F., Clavel M., Haggard D., 2024, *Simulation-Based Inference of Radio Millisecond Pulsars in Globular Clusters* ([arXiv:2405.15691](https://arxiv.org/abs/2405.15691))
- Binney J., Merrifield M., 1998, *Galactic Astronomy*. Princeton Series in Astrophysics, Princeton University Press, Princeton, NJ
- Binney J., Tremaine S., 2008, *Galactic Dynamics*. Princeton University Press
- Blandford R. D., Romani R. W., Applegate J. H., 1987, *Monthly Notices of the Royal Astronomical Society*, 225, 51p
- Breen P. G., Heggie D. C., 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 2779
- Cadelano M., Ransom S. M., Freire P. C. C., Ferraro F. R., Hessels J. W. T., Lanzoni B., Pallanca C., Stairs I. H., 2018, *The Astrophysical Journal*, 855, 125
- Carroll B. W., Ostlie D. A., 2014, *An Introduction to Modern Astrophysics*, 2. ed., pearson new internat. ed edn. Pearson, Harlow
- Chatterjee S., Umbreit S., Fregeau J. M., Rasio F. A., 2013, *Monthly Notices of the Royal Astronomical Society*, 429, 2881
- Chen W., et al., 2023, *Monthly Notices of the Royal Astronomical Society*, 520, 3847
- Chomiuk L., Strader J., Maccarone T. J., Miller-Jones J. C. A., Heinke C., Noyola E., Seth A. C., Ransom S., 2013, *The Astrophysical Journal*, 777, 69
- Claydon I., Gieles M., Varri A. L., Heggie D. C., Zocchi A., 2019, *Monthly Notices of the Royal Astronomical Society*, 487, 147
- Corongiu A., et al., 2024, *Timing of Millisecond Pulsars in NGC\,6752 – III. On the Presence of Non-Luminous Matter in the Cluster’s Core* ([arXiv:2407.03271](https://arxiv.org/abs/2407.03271))
- Della Croce A., Pascale R., Giunchi E., Nipoti C., Cignoni M., Dalessandro E., 2024, *Astronomy and Astrophysics*, 682, A22
- Dickson N., Hénault-Brunet V., Baumgardt H., Gieles M., Smith P. J., 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 5320

- Dickson N., Smith P. J., Hénault-Brunet V., Gieles M., Baumgardt H., 2024, *Monthly Notices of the Royal Astronomical Society*, 529, 331
- Ferraro F. R., et al., 2009, *Nature*, 462, 483
- Ferraro F. R., Massari D., Dalessandro E., Lanzoni B., Origlia L., Rich R. M., Mucciarelli A., - 2016, *The Astrophysical Journal*, 828, 75
- Ferraro F. R., et al., 2021, *Nature Astronomy*, 5, 311
- Fishbach M., Fragione G., 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 5546
- Forbes D. A., et al., 2018, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474, 20170616
- Freire P. C., Ridolfi A., 2018, *Monthly Notices of the Royal Astronomical Society*, 476, 4794
- Freire P. C., et al., 2017, *Monthly Notices of the Royal Astronomical Society*, 471, 857
- Fryer C. L., Belczynski K., Wiktorowicz G., Dominik M., Kalogera V., Holz D. E., 2012, *Astrophysical Journal*, 749, 91
- Gallo E., Fender R. P., 2005, *Memorie della Societa Astronomica Italiana*, 76, 600
- Gallo E., Fender R. P., Miller-Jones J. C. A., Merloni A., Jonker P. G., Heinz S., Maccarone T. J., Van Der Klis M., 2006, *Monthly Notices of the Royal Astronomical Society*, 370, 1351
- Gallo E., Miller B. P., Fender R., 2012, *Monthly Notices of the Royal Astronomical Society*, 423, 590
- Gerssen J., van der Marel R. P., Gebhardt K., Guhathakurta P., Peterson R., Pryor C., 2002, *The Astronomical Journal*, 124, 3270
- Gieles M., Zocchi A., 2015, *Monthly Notices of the Royal Astronomical Society*, 454, 576
- Gieles M., Balbinot E., Yaaqib R. I. S. M., Hénault-Brunet V., Zocchi A., Peuten M., Jonker P. G., 2018, *Monthly Notices of the Royal Astronomical Society*, 473, 4832
- Giersz M., Heggie D. C., 2011, *Monthly Notices of the Royal Astronomical Society*, 410, 2698
- Giersz M., Heggie D. C., Hurley J. R., Hypki A., 2013, *Monthly Notices of the Royal Astronomical Society*, 431, 2184
- Giesers B., et al., 2018, *Monthly Notices of the Royal Astronomical Society: Letters*, 475, L15
- Giesers B., et al., 2019, *Astronomy & Astrophysics*, 632, A3
- Gill M., Trenti M., Miller M. C., van der Marel R., Hamilton D., Stiavelli M., 2008, *The Astrophysical Journal*, 686, 303
- Häberle M., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 503, 1490
- Häberle M., et al., 2024, *Nature*, 631, 285
- Harris W. E., 1996, *The Astronomical Journal*, 112, 1487
- Heggie D. C., 1975, *Monthly Notices of the Royal Astronomical Society*, 173, 729

- Heggie D., Hut P., 2003, *The Gravitational Million–Body Problem*. Cambridge University Press, doi:10.1017/CBO9781139164535
- Hénault-Brunet V., Gieles M., Sollima A., Watkins L. L., Zocchi A., Claydon I., Pancino E., Baumgardt H., 2019, *Monthly Notices of the Royal Astronomical Society*, 483, 1400
- Hénault-Brunet V., Gieles M., Strader J., Peuten M., Balbinot E., Douglas K. E., 2020, *Monthly Notices of the Royal Astronomical Society*, 491, 113
- Hertz P., Grindlay J. E., 1983, *The Astrophysical Journal*, 275, 105
- Hessels J., et al., 2015, in *Proceedings of Advancing Astrophysics with the Square Kilometre Array — PoS(AASKA14)*. Sissa Medialab, Giardini Naxos, Italy, p. 047, doi:10.22323/1.215.0047
- Hypki A., Giersz M., 2013, *Monthly Notices of the Royal Astronomical Society*, 429, 1221
- Jalali B., Baumgardt H., Kissler-Patig M., Gebhardt K., Noyola E., Lützgendorf N., De Zeeuw P. T., 2012, *Astronomy and Astrophysics*, 538, 1
- Johnson C. I., Pilachowski C. A., 2010, *The Astrophysical Journal*, 722, 1373
- Kains N., Calamida A., Sahu K. C., Anderson J., Casertano S., Bramich D. M., 2018, *The Astrophysical Journal*, 867, 37
- Kamann S., et al., 2016, *Astronomy and Astrophysics*, 588, A149
- Kamann S., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 473, 5591
- King I. R., 1966, *The Astronomical Journal*, 71, 64
- Kıroğlu F., Weatherford N. C., Kremer K., Ye C. S., Fragione G., Rasio F. A., 2022, *The Astrophysical Journal*, 928, 181
- Kızıltan B., Baumgardt H., Loeb A., 2017a, *Nature*, 542, 203
- Kızıltan B., Baumgardt H., Loeb A., 2017b, *Nature*, 545, 510
- Koposov S., et al., 2023, *Joshspeagle/Dynesty: V2.1.3*, Zenodo, doi:10.5281/zenodo.8408702
- Kremer K., Chatterjee S., Rodriguez C. L., Rasio F. A., 2018a, *The Astrophysical Journal*, 852, 29
- Kremer K., Ye C. S., Chatterjee S., Rodriguez C. L., Rasio F. A., 2018b, *The Astrophysical Journal*, 855, L15
- Kremer K., Chatterjee S., Ye C. S., Rodriguez C. L., Rasio F. A., 2019, *The Astrophysical Journal*, 871, 38
- Kremer K., et al., 2020, *The Astrophysical Journal Supplement Series*, 247, 48
- Kroupa P., 2001, *Monthly Notices of the Royal Astronomical Society*, 322, 231
- Kulkarni S. R., Hut P., McMillan S. J., 1993, *Nature*, 364, 421
- Lanzoni B., et al., 2010, *The Astrophysical Journal*, 717, 653

- Libralato M., et al., 2022, *The Astrophysical Journal*, 934, 150
- Lindegren L., et al., 2021, *Astronomy and Astrophysics*, 649, A2
- Longair M. S., 2011, *High Energy Astrophysics*, 3rd edn. Cambridge University Press, Cambridge ; New York
- Lützgendorf N., Gebhardt K., Baumgardt H., Noyola E., Neumayer N., Kissler-Patig M., De Zeeuw T., 2015, *Astronomy and Astrophysics*, 581
- Lyne A. G., Mankelov S. H., Bell J. F., Manchester R. N., 2000, *Monthly Notices of the Royal Astronomical Society*, 316, 491
- MacKey A. D., Wilkinson M. I., Davies M. B., Gilmore G. F., 2008, *Monthly Notices of the Royal Astronomical Society*, 386, 65
- Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, *The Astronomical Journal*, 129, 1993
- Mann C. R., et al., 2019, *The Astrophysical Journal*, 875, 1
- Mann C. R., et al., 2020, *The Astrophysical Journal*, 893, 86
- McKenzie M., Bekki K., 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 3126
- McNamara B. J., Harrison T. E., Anderson J., 2003, *The Astrophysical Journal*, 595, 187
- Merritt D., Piatek S., Zwart S. P., Hemsendorf M., 2004, *The Astrophysical Journal*, 608, L25
- Miller-Jones J. C. A., et al., 2015, *Monthly Notices of the Royal Astronomical Society*, 453, 3918
- Morscher M., Umbreit S., Farr W. M., Rasio F. A., 2013, *The Astrophysical Journal Letters*, 763, L15
- Morscher M., Pattabiraman B., Rodriguez C., Rasio F. A., Umbreit S., 2015, *Astrophysical Journal Letters*, 800
- Noyola E., Gebhardt K., Bergmann M., 2008, *The Astrophysical Journal*, 676, 1008
- Origlia L., Massari D., Rich R. M., Mucciarelli A., Ferraro F. R., Dalessandro E., Lanzoni B., 2013, *The Astrophysical Journal*, 779, L5
- Paczynski B., 1986, *The Astrophysical Journal*, 304, 1
- Paczynski B., 1994, *Acta Astronomica*, 44, 235
- Padmanabh P. V., et al., 2024, *Astronomy and Astrophysics*, 686, A166
- Paduano A., et al., 2022, 16, 1
- Paduano A., et al., 2024, *The Astrophysical Journal*, 961, 54
- Pallanca C., et al., 2021, *The Astrophysical Journal*, 917, 92
- Pancino E., et al., 2024, *Astronomy and Astrophysics*, 686, A283
- Peuten M., Zocchi A., Gieles M., Gualandris A., Hénault-Brunet V., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 2333

- Peuten M., Zocchi A., Gieles M., Hénault-Brunet V., 2017, *Monthly Notices of the Royal Astronomical Society*, 470, 2736
- Pfeffer J., Lardo C., Bastian N., Saracino S., Kamann S., 2021, *Monthly Notices of the Royal Astronomical Society*, 500, 2514
- Phinney E. S., 1992, *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 341, 39
- Phinney E. S., 1993, *ASP Conference Series*, 50, 141
- Portegies Zwart S. F., McMillan S. L. W., 1999, *The Astrophysical Journal*, 528, L17
- Prager B. J., Ransom S. M., Freire P. C. C., Hessels J. W. T., Stairs I. H., Arras P., Cadelano M., 2017, *The Astrophysical Journal*, 845, 148
- Price-Whelan A. M., 2017, *Journal of Open Source Software*, 2, 388
- Ransom S. M., Hessels J. W. T., Stairs I. H., Freire P. C. C., Camilo F., Kaspi V. M., Kaplan D. L., 2005, *Science*, 307, 892
- Ridolfi A., et al., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 2918
- Ridolfi A., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 504, 1407
- Roberts M. S. E., 2012, *Proceedings of the International Astronomical Union*, 8, 127
- Rodriguez C. L., Chatterjee S., Rasio F. A., 2016, *Physical Review D*, 93, 084029
- Rodriguez C. L., et al., 2022, *The Astrophysical Journal Supplement Series*, 258, 22
- Romano D., et al., 2023, *The Astrophysical Journal*, 951, 85
- Rui N. Z., Weatherford N. C., Kremer K., Chatterjee S., Fragione G., Rasio F. A., Rodriguez C. L., Ye C. S., 2021a, *Research Notes of the AAS*, 5, 47
- Rui N. Z., Kremer K., Weatherford N. C., Chatterjee S., Rasio F. A., Rodriguez C. L., Ye C. S., 2021b, *The Astrophysical Journal*, 912, 102
- Saracino S., et al., 2015, *The Astrophysical Journal*, 806, 152
- Shaifullah G., et al., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 1029
- Shishkovsky L., et al., 2020, *The Astrophysical Journal*, 903, 73
- Shklovskii I. S., 1970, *Soviet Astronomy*, 13, 562
- Sigurdsson S., Hernquist L., 1993, *Nature*, 364, 423
- Sollima A., et al., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 1937
- Sollima A., Baumgardt H., Hilker M., 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 1460
- Speagle J. S., 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 3132
- Spera M., Mapelli M., 2017, *Monthly Notices of the Royal Astronomical Society*, 470, 4739

- Spitzer Jr. L., 1969, *The Astrophysical Journal*, 158, L139
- Strader J., Chomiuk L., Maccarone T. J., Miller-Jones J. C. A., Seth A. C., 2012, *Nature*, 490, 71
- Tam P. H. T., Kong A. K. H., Hui C. Y., Cheng K. S., Li C., Lu T.-N., 2011, *The Astrophysical Journal*, 729, 90
- Taylor D. J., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 513, 3429
- The LIGO Scientific Collaboration et al., 2023, *Physical Review X*, 13, 041039
- Thongmeearkom T., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 530, 4676
- Trager S. C., King I. R., Djorgovski S., 1995, *The Astronomical Journal*, 109, 218
- Tudor V., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 513, 3818
- Van Der Marel R. P., Anderson J., 2010, *Astrophysical Journal*, 710, 1063
- VandenBerg D. A., Brogaard K., Leaman R., Casagrande L., 2013, *The Astrophysical Journal*, 775, 134
- Vasiliev E., Baumgardt H., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 5978
- Verbiest J. P. W., et al., 2008, *The Astrophysical Journal*, 679, 675
- Vitral E., Mamon G. A., 2021, *Astronomy and Astrophysics*, 646, A63
- Vitral E., Kremer K., Libralato M., Mamon G. A., Bellini A., 2022, *Monthly Notices of the Royal Astronomical Society*, 514, 806
- Vitral E., Libralato M., Kremer K., Mamon G. A., Bellini A., Bedin L. R., Anderson J., 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 5740
- Watkins L. L., van de Ven G., den Brok M., van den Bosch R. C. E., 2013, *Monthly Notices of the Royal Astronomical Society*, 436, 2598
- Weatherford N. C., Chatterjee S., Rodriguez C. L., Rasio F. A., 2018, *The Astrophysical Journal*, 864, 13
- Weatherford N. C., Chatterjee S., Kremer K., Rasio F. A., 2020, *The Astrophysical Journal*, 898, 162
- Wilson C. P., 1975, *The Astronomical Journal*, 80, 175
- Woolley R. V. D. R., 1954, *Monthly Notices of the Royal Astronomical Society*, 114, 191
- Ye C. S., Kremer K., Rodriguez C. L., Rui N. Z., Weatherford N. C., Chatterjee S., Fragione G., Rasio F. A., 2022, *The Astrophysical Journal*, 931, 84
- Zaris J., Veske D., Samsing J., Márka Z., Bartos I., Márka S., 2020, *The Astrophysical Journal*, 894, L9
- Zocchi A., Gieles M., Hénault-Brunet V., Varri A. L., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 696

Zocchi A., Gieles M., Hénault-Brunet V., 2017, Monthly Notices of the Royal Astronomical Society, 468, 4429

Zocchi A., Gieles M., Hénault-Brunet V., 2019, Monthly Notices of the Royal Astronomical Society, 482, 4713

de Boer T. J. L., Gieles M., Balbinot E., Hénault-Brunet V., Sollima A., Watkins L. L., Claydon I., 2019, Monthly Notices of the Royal Astronomical Society, 485, 4906