

Probing the dark universe with gravitational lensing

by

Mansour Karami

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2018

© Mansour Karami 2018

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Patrick Hall
Professor, Dept. of Physics and Astronomy, York University

Supervisor(s): Niayesh Afshordi
Associate Professor, Dept. of Physics and Astronomy, University of Waterloo
Avery Broderick
Associate Professor, Dept. of Physics and Astronomy, University of Waterloo

Internal Member(s): Brian McNamara
Professor, Dept. of Physics and Astronomy, University of Waterloo
James Taylor
Associate Professor, Dept. of Physics and Astronomy, University of Waterloo

Internal-External Member: Francis Poulin
Associate Professor, Dept. of Applied Mathematics, University of Waterloo

This thesis consists of material all of which I authored or co-authored (see the statement of contributions in the following page). This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of contributions

This dissertation is partially the product of co-authored publications:

- Chapter 2 is based on:

Karami, M.; Afshordi, N.; Zavala, J., “Forward modelling of quasar light curves and the cosmological matter power spectrum on milliparsec scales”, eprint arXiv:1805.06984, Submitted to Physical Review D

I contributed to writing the code, running the analysis and part of the calculations

- Chapter 3 is based on:

Karami, M.; Broderick, A. E.; Rahvar, S.; Reid, M., “Resolving Microlensing Events with Triggered VLBI”, The Astrophysical Journal, Volume 833, Issue 2, article id. 169, 18 pp. (2016)

I contributed to writing the code to calculate the event statistics and running the analysis

- Chapter 4 is based on:

Moniez, M.; Sajadian, S.; **Karami, M.**; Rahvar, S.; Ansari, R., “Understanding EROS2 observations toward the spiral arms within a classical Galactic model framework”, Astronomy & Astrophysics, Volume 604, id.A124, 16 pp. (2017)

I contributed to writing the Monte-Carlo simulation and running the analysis

- Chapter 5 is based on:

Broderick, A. E.; Gold, R.; **Karami, M.**; Preciado-Lopez, J. A.; Tiede, P.; Pu, H.; “Themis: A Parameter Estimation Framework for the Event Horizon Telescope”, To be submitted to Astrophysical Journal

I contributed to writing the code, mostly in the Bayesian inference parts and running some of the analysis

Abstract

Since its early success as an experimental test of the theory of general relativity in 1919, gravitational lensing has come a long way and is firmly established as an indispensable element for many astrophysical applications. In this thesis, we explore novel applications of gravitational lensing that further our understanding of the dark sectors of the cosmos and other astrophysical objects, namely dark matter nanostructure, black holes and the Galactic disk. We pay particular attention to developing concrete and optimal statistical methodologies and numerical implementations for these novel probes.

We start by developing a statistical framework to measure the dark matter power spectrum in the deep nonlinear regime, using *transient weak lensing*, and simultaneously measure the time delays for strongly lensed quasars. We then outline how observations of *microlensing* in optical and radio can unravel the structure, dynamics, and content of the Galactic disk, and in particular, be used to detect stellar mass black holes. Lastly, using the shadow images of the super-massive black holes caused by extreme lensing effect, we can learn about the structure of space-time, accretion flows and astrophysical jets. We present a Bayesian framework for analyzing the data from the Event Horizon Telescope Collaboration.

Acknowledgements

I would like to thank everyone who made this thesis possible. In particular I want to thank my advisors Niayesh Afshordi and Avery Broderick from whom I learned much and to whom I will always be indebted.

Dedication

To my parents

Table of Contents

List of Tables	xiii
List of Figures	xiv
1 Introduction	1
1.1 Bayesian inference	2
1.1.1 The Bayes' theorem	2
1.1.2 Bayesian parameter estimation	3
1.1.3 Markov Chain Monte Carlo	4
1.1.4 Improving the efficiency of MCMC methods	4
1.1.5 Parallel tempering	5
1.1.6 Model selection and criticism	6
1.2 Gravitational lensing	7
1.2.1 The formalism	7
1.2.2 Special cases	9
1.2.3 Gravitational lensing and time delay	10
2 Forward modelling of quasar light curves and the cosmological matter power spectrum on milliparsec scales	12
2.1 Introduction	13
2.2 Datasets	14

2.3	Method description	14
2.4	MCMC Sampling	16
2.5	Finite Size Effect	18
2.6	Recovering time delays and the time-delay challenge	20
2.7	Results	24
2.7.1	Time delays	24
2.7.2	Quasar temporal power spectrum	27
2.7.3	Dark matter power spectrum	28
2.8	Conclusion and Future Prospects	29
3	Resolving microlensing events with triggered VLBI	33
3.1	Introduction	33
3.2	Stellar Remnants in Simulated Microlensing Surveys	37
3.2.1	Galactic Distribution Model and Stellar Lens/Source Population Model	37
3.2.2	Remnant Population Model	39
3.2.3	Mock Observation Statistics	42
3.3	Radio-Bright Microlensing Event Rates	44
3.3.1	Miras as Microlensing Sources	47
3.3.2	Bulge Miras Microlensing Rates	49
3.3.3	Current and Near-future Optical Survey Microlensing Rates	52
3.3.4	Far-future Optical Microlensing Surveys	53
3.3.5	A Future Radio-Continuum Microlensing Survey	54
3.4	Radio VLBI Event Reconstruction	55
3.4.1	Models of Resolved Masers	55
3.4.2	Generating Mock Images	56
3.4.3	Example Lensed Maser Images	58
3.4.4	Parameter Estimation	58
3.4.5	Binaries	67

3.5	Discussion	67
3.5.1	Massive Star Evolution	68
3.5.2	Gravitational Wave Experiments	69
3.5.3	Isolated Neutron Stars	69
3.6	Conclusions	71
4	Understanding EROS2 observations toward the spiral arms within a classical Galactic model framework	73
4.1	Introduction	74
4.2	Microlensing effect	76
4.2.1	Microlensing event characteristics	77
4.2.2	Observables: optical depth, event rate, and t_E distribution	77
4.3	EROS data toward the Galactic spiral arms	78
4.3.1	EROS color-magnitude diagrams	78
4.3.2	Microlensing results	83
4.4	How to synthesize an EROS-like color-magnitude diagram	83
4.4.1	Producing a CMD from the local HIPPARCOS catalog	85
4.4.2	Mass density distributions	86
4.4.3	Extinction	89
4.5	Comparing the EROS observations with simulated populations and microlensing expectations	91
4.5.1	Simulation of the CMDs	92
4.5.2	Simulation of microlensing	92
4.6	Fitting procedure	93
4.6.1	Fit and tuning of the simple model	93
4.6.2	Besançon model: tuning the extinctions	99
4.7	Discussion	99
4.7.1	Comparison with previous results and robustness	99

4.7.2	Lens and source populations	101
4.7.3	Constraining the Galactic model: The specific contribution of microlensing data	102
4.7.4	Limitations of this study	104
4.8	Conclusions and perspectives	105
5	Themis: A Parameter estimation framework for the Event Horizon Telescope	106
5.1	Introduction	107
5.2	Summary of THEMIS	109
5.2.1	Structure	109
5.2.2	Implementation	110
5.3	THEMIS Data Structures	111
5.3.1	Visibility Amplitudes	114
5.3.2	Closure Phases	115
5.3.3	Closure Amplitudes	116
5.3.4	Interferometric Polarization Fractions	116
5.3.5	Flux Measurements	117
5.4	THEMIS Models	117
5.4.1	Image-based Models	118
5.4.2	Phenomenological Geometric Models	118
5.4.3	Interstellar Scattering Models	120
5.4.4	Native Physical Models	122
5.4.5	External Physical Models	125
5.5	Likelihoods, Priors and Samplers	126
5.5.1	Likelihoods	126
5.5.2	Priors	131
5.5.3	Samplers	132
5.6	Validation Tests	135

5.6.1	Validation of the samplers	135
5.6.2	Self-tests with simulated data	137
5.6.3	Reproducing Previous Results	145
5.7	New Results	152
5.7.1	Crescent model	152
5.7.2	Extended RIAF model	152
5.8	Code Performance	154
5.9	Summary	156
6	Conclusion	160
	Bibliography	189
	APPENDICES	189
A	Breaking degeneracy with combined observation	190
A.1	Lens Equation for Single Lens	190
A.2	Parallax Effect in Microlensing Light Curves	192
A.3	Parallax Effect in Microlensing Images	193
B	Producing a local debiased CMD from the HIPPARCOS catalog	195
B.1	Extrapolating the local HIPPARCOS CMD	199
B.2	Comparison with the stellar density expected from the mass function: A coherence check	201
B.3	Parameters of the Besançon Galactic model	202
C	Prediction Accuracy Requirements	205
C.1	Error Distributions of Quantities Associated with Visibility Amplitudes	207
C.2	Visibility Amplitudes	207
C.2.1	Visibility Amplitude Products	209
C.3	Polarization Fractions — Visibility Amplitude Quotients	211
C.4	Closure Amplitudes	212

List of Tables

2.1	Recovered values vs actual time delay values for randomly selected TDC light curves	21
2.2	Time delay values obtained using our method compared to the previous works . .	25
3.1	Parameters of the galactic model	38
4.1	Data and results toward the 4 regions monitored in the EROS spiral arms program	84
4.2	Best fit results on the observables toward the 4 regions monitored in the EROS spiral arms program	98
5.1	Published EHT Data	114
B.1	Age, local mass density $\rho(r_{\odot}, z_{\odot})$, disk axis ratio ϵ , and IMF of the different stellar components of the disks in the Besançon model	203
C.1	Accuracy of Visibility Amplitude Error Distribution Approximations	208
C.2	Accuracy of Product Error Distribution Approximations	210
C.3	Accuracy of Polarization Fraction Error Distribution Approximations	213
C.4	Accuracy of Closure Amplitude Error Distribution Approximations	219

List of Figures

1.1	Schematic view of a gravitational lensing event	8
2.1	Ratio of temporal power spectrum for an extended source to that of a point source.	19
2.2	Synthetic light curves for the two quasar images	21
2.3	Recovered lensing power spectrum for the synthetic quasar light curve	22
2.4	Recovered intrinsic power spectrum for the synthetic quasar light curve	23
2.5	Posterior distributions of the time delays for COSMOGRIAL quasars	26
2.6	Measurements of the temporal power spectrum for intrinsic quasar variability	27
2.7	Measurements of the lensing power spectrum	29
2.8	The light curves for the doubly lensed Quasar system J1001+5027	30
2.9	Upper limits on the dimensionless dark matter power spectrum	31
3.1	Assumed mass functions of stars, white dwarfs and black holes in the disk and the bulge	40
3.2	Assumed mass functions of stars, white dwarfs and black holes in the Spheroid	41
3.3	Distribution of microlensing events in Einstein crossing time from Monte-Carlo simulation	45
3.4	Distribution of lens positions in galaxy for various lens populations	46
3.5	The OGLE III Color-Magnitude diagram	48
3.6	Number of Miras in OGLEII and OGLEIII fields vs COBE/DIRBE $2.2\mu\text{m}$ flux in the same regions	51
3.7	Microlensed images for a $10 M_{\odot}$ black hole	59

3.8	Microlensed images for a $1 M_{\odot}$ lens	60
3.9	light curve for the microlensing event shown in Figure 3.7	64
3.10	Evolution of the image separations for the microlensing event shown in Figure 3.7	65
3.11	Microlensing event images for a black hole-black hole binary lens	66
4.1	Four directions toward the Galactic spiral arms monitored by EROS	75
4.2	Time sampling toward the 4 monitored targets in the Galactic spiral arms	79
4.3	Relative color-magnitude diagrams $n(I, V - I)$ of the EROS catalogs toward the 4 directions toward the Galactic spiral arms	80
4.4	Star detection probability in EROS vs. the relative magnitude	81
4.5	Photometric point-to-point precision along the EROS light-curves vs. $R_{EROS} = I$ and B_{EROS}	82
4.6	Microlensing detection efficiency of the EROS survey toward the Galactic spiral arms	83
4.7	The $V - I$ observed and the simulated distributions	90
4.8	Different mass functions considered in this chapter	96
4.9	Mass-density along the line of sight of γSct from the various Galactic structures	97
4.10	Simulated CMDs toward the 4 monitored directions with the magnitude and color projections for the stars brighter than $I = 18.4$	100
4.11	Simulated optical depths toward the 4 monitored directions	101
4.12	Expected normalized distributions of the distances for the lensed sources	103
4.13	Einstein duration t_E distribution of the microlensing events	104
5.1	Inheritance diagram for the <code>model_image</code> object within THEMIS generated via Doxygen	112
5.2	Inheritance diagram for the <code>likelihood</code> object within THEMIS generated via Doxygen	113
5.3	Image produced by the radiative transfer module VRT ² showing the resulting best-fit RIAF model	123
5.4	Two dimensional Gaussian likelihood sampled by the affine invariant MCMC sampler, the differential evolution MCMC sampler and the grid search sampler	136

5.5	Validation of the sampler with a five dimensional “eggbox”-likelihood	138
5.6	Validation of the sampler with a two dimensional test involving a 16-Gaussians-likelihood	139
5.7	Relative posterior probability mass for the 16 gaussian peaks recovered via MCMC sampling	140
5.8	Posterior distribution of the size of the symmetric Gaussian reconstructed from simulated closure amplitude data	141
5.9	Joint parameter distributions from the crescent model analysis of the simulated visibility amplitude and closure phase data	143
5.10	Joint parameter distributions for the SED-fitted RIAF model analysis of the simulated visibility amplitude and closure phase data	144
5.11	Gaussian size distribution from the analysis of the symmetric Gaussian model to the 2007 and 2009 visibility amplitude data of Sgr A*	146
5.12	Joint parameter distributions from the analysis of the asymmetric Gaussian model to the 2007 and 2009 visibility amplitude data of Sgr A*	147
5.13	Joint parameter distributions from the analysis of the crescent model to the 2007 and 2009 visibility amplitude data of Sgr A*	149
5.14	Joint parameter distributions from the analysis of the SED-fitted RIAF model to the 2007 and 2009 visibility amplitude data of Sgr A*	150
5.15	Joint posterior parameter distributions for the SED-fitted RIAF model implied by the combined visibility amplitude and closure phase data sets bolded in Table 5.1	151
5.16	Joint posterior parameter distributions for the crescent model implied by the combined visibility amplitude and closure phase data sets bolded in Table 5.1	153
5.17	Joint posterior parameter distributions for the SED-fitted RIAF model implied by the combined flux, visibility amplitude and closure phase data sets bolded in Table 5.1	155
5.18	Scaling plot illustrating the multilayered parallelization strategy in THEMIS	157
B.1	HIPPARCOS absolute magnitudes vs. distance distributions (up= M_V , down= M_I)	196
B.2	HIPPARCOS absolute color-magnitude diagram in M_{I_C} vs. $(V - I)_J$	197
B.3	Two-dimensional and 3D distributions of the HIPPARCOS objects	198

B.4	Raw and debiased distributions of color and magnitude for HIPPARCOS stars . .	200
C.1	Comparison of the Rice and Gaussian distributions at various SNRs together with the distribution of products of visibility amplitudes and their comparison to a single Rice distribution	208
C.2	Comparison of the polarization fraction distribution	210
C.3	Comparison of the closure amplitude distributions	214

Chapter 1

Introduction

In this thesis, we apply gravitational lensing to indirectly detect the properties of astrophysical objects on different scales. In most cases, we are targeting objects that are not luminous and hence not directly observable; this includes dark matter structure and black holes.

There has been many observational constraints on the dark matter power spectrum in the last few decades including the CMB observations, large scale structure surveys, Lyman- α forest and more. These provide strong limits on the dark matter power spectrum on cosmological scales down to galactic scales. On smaller scales however, there are only weak upper bounds coming from a handful of methods such as the limits from ultra compact minihalos and primordial black holes [37]. In Chapter 2 we propose a novel way to constrain the dark matter power spectrum in extremely small scales comparable to the size of the solar system.

Stellar-mass black holes, as the final stage of stellar evolution for massive stars present a unique way to study the evolution of high mass stars including the core-collapse supernova physics as well constraining the high-mass end of the stellar initial mass function [110]. In addition having a black hole mass function will allow us to more accurately estimate the merger event rates in gravitational wave experiments. Currently only a couple dozens of stellar-mass black holes with measured masses exist [52] and they belong to a very specific subsets of the stellar mass black holes namely X-ray binaries detected in X-ray observations and merging black holes detected via gravitational wave experiments. The current low statistics coupled with nontrivial biases in the observed populations make it impossible to construct a black hole mass function. Chapter 3 proposes a new way to detect stellar mass black holes at a much higher rate than possible today using gravitational microlensing observed in radio images of radio bright sources. This method allows for breaking the degeneracies present in stellar microlensing and measuring all the lensing parameters such as mass of the black hole, its distance and velocity.

Gravitational microlensing provides a unique opportunity to study the structure of the Milky Way. The microlensing observables such as the optical depth and the distribution of the lensing time scales depend to the density profile along the line of sight, velocity distribution of the sources and the lenses and the mass distribution of the lenses [190]. Combined with other data products from microlensing surveys such as the color-magnitude distribution of the monitored stars this can be used to constraint the structure and dynamics of the Galaxy. In Chapter 4 we use the gravitational microlensing data to study the structure and dynamics of the Milky Way’s disk .

The Event Horizon Telescope is the highest resolution instrument capable of achieving a resolutions of $13\mu\text{as}$. It is a collection of radio telescopes all around the globe and is aimed at resolving event-horizon scale structures around supermassive black holes [79, 82]. The main targets of observations are the supermassive black holes at the center of the Galaxy (Sgr A*) and M87. Chapter 5 introduces Themis which is a parameter estimation framework for the Event Horizon Telescope (EHT). Themis is a Bayesian parameter estimation framework and allows the use of different EHT data products, incorporating various physical models. The code is designed to be easily extensible and to exploit the high performance computing machines efficiently. In what follows, we start with an introduction to Bayesian inference and gravitational lensing as the common themes in this thesis, setting the stage for the following chapters.

1.1 Bayesian inference

Statistical inference as an extension of logical inference is the tool that is used in all branches of science for building and evolving the mathematical models. In particular, Bayesian inference has been widely used in physics over the past decade. Bayesian method allows iterative improvements to the models by combining the new data with the results of the last updated model (used as prior information) in a coherent way. This in essence is the gradual application of the scientific method as the new data becomes available. In this section, we overview the concepts of Bayesian inference and how they are applied to the process of model building.

1.1.1 The Bayes’ theorem

The Bayes’ theorem is derived using the symmetries of Boolean logic and is expressed as:

$$P(\theta|M, D) = \frac{P(D|M, \theta)P(\theta|M)}{P(D|M)}, \tag{1.1}$$

where M represents the mathematical model, θ represents the set of parameters within the model M and D stands for the observed data. The term $P(\theta|M)$ is the prior and is the probability of each set of parameters within model M prior to observing the data D . This would generally come from the data perviously observed and represents the state of knowledge before the new data becomes available. $P(D|M, \theta)$ is called the likelihood and is the probability of the observed data D being observed in the context of model M and parameter set θ . This term represents the predictions of the model at its current state. The term in the denominator $P(D|M)$ is called the Bayesian evidence and is the probability of data D being observed given the model M . This is a normalization factor that can be ignored if one is only considering a single model M . On the other hand, when comparing two different models M and M' the bayesian evidence is used to chose which model is favoured by the data. This is the subject of section 1.1.6. The term on the left hand side $P(\theta|M, D)$ is called the posterior probability and is the probability of the parameter set θ given the observed data D and in the context of model M . The posterior represents the updated plausibility of each θ after taking into account the new observations. One can look at the Bayes' theorem as the application of scientific method used to iteratively improve the mathematical models with new observations.

1.1.2 Bayesian parameter estimation

Oftentimes one needs to estimate parameters within a given model, be it a phenomenological model or built upon a physical theory. In such cases the underlying model M is fixed and the Bayes' theorem can be simplified as:

$$P(\theta|D) \propto P(D|\theta)P(\theta), \quad (1.2)$$

where we have taken out the Bayesian evidence term $P(D)$ since it's a constant. One is generally interested in calculating expectation values of physical quantities over the posterior probability, (these include finding the means and variances)

$$E[f] = \int f(\theta)P(\theta|D)d\theta, \quad (1.3)$$

where $E[f]$ is the expectation value of the function $f(\theta)$ over the posterior probability distribution. However in many cases the model contains many parameters and numerical estimation of the integral poses a challenge. In particular the number of evaluations of the integrand grows exponentially with the number of dimensions. The de facto way of approximating the posterior probability mass in high dimensions is to use the Markov Chain Monte Carlo (MCMC) methods due to their excellent scaling with the dimension of the parameter space θ .

1.1.3 Markov Chain Monte Carlo

The origins of Markov Chain Monte Carlo methods go back to almost the same time as the Monte Carlo methods themselves (i.e 1940s). The first paper on an MCMC algorithm was written in 1953 by Metropolis et al [182] on what is known as the Metropolis method today but it wasn't until 1990s that these methods gained widespread application.

Markov Chain Monte Carlo methods constructs a sequence of random points from the parameter space θ such that the sequence has the Markov property and the stationary distribution of the points converges to the target density (e.g. the posterior distribution). Markov property means that each point in the sequence depends only on the previous point. It has to be noted although this is generally true there are some MCMC methods that don't strictly satisfy this condition such as the adaptive MCMC methods [14]. Once the MCMC chain is converged the density of the sequence of points is proportional to the mass density of the target distribution which is often the posterior distribution [183].

As an example of a simple MCMC algorithm we briefly present the Metropolis-Hastings algorithm. To construct the series of samples from the posterior distribution in Metropolis-Hastings method one proposes the new point (θ_{n+1}) according to:

$$\theta_{n+1} = \theta_n + \mathcal{N}(0, \Sigma). \quad (1.4)$$

In this equation $\mathcal{N}(0, \Sigma)$ is a vector of random gaussian numbers with mean 0 and covariance matrix Σ . The proposed point θ_{n+1} is then accepted or rejected with a probability given by

$$\min\left(1, \frac{P(D|\theta_{n+1})P(\theta_{n+1})}{P(D|\theta_n)P(\theta_n)}\right). \quad (1.5)$$

If the proposed point (θ_{n+1}) is accepted then it's added to the series of points; otherwise the current point (θ_n) is added to the series. The procedure is then repeated until the series is converged to a stationary distribution which is the posterior ($P(\theta|D)$) probability mass.

1.1.4 Improving the efficiency of MCMC methods

Although MCMC methods can be very efficient, if not tuned properly they can suffer from poor convergence. As an example the efficiency of Metropolis-Hastings algorithm depends on the choice of the covariance matrix Σ used to propose the new steps. The optimal choice of the Σ depends on the shape of the likelihood surface which is not known a priori. There are a host of

methods called the adaptive methods that aim at optimizing the stepping probability $\mathcal{N}(0, \Sigma)$ by iteratively changing Σ using the information from the previous steps [14].

There is a class of MCMC methods called the ensemble sampling methods that use many MCMC chains (called the walkers) that work in conjunction to propose new steps and sample the target distribution. The advantage of these methods is that they generally require minimal tuning and can perform efficiently on a large class of target distributions. Examples of such methods are the Affine Invariant method of Goodman et al. [127] and the Differential Evolution method of [36]. Additionally these algorithms can be parallelized and using enough computational resources can run as fast as a single MCMC chain.

There are many other improved MCMC methods devised through the years and it's outside the scope of this thesis to cover all. Here we mention one notable example called the Hamiltonian MCMC [24, 25]. Hamiltonian MCMC was first introduced for numerical simulations of lattice field theory [85] and is particularly suited for very high dimensional problems where the conventional MCMC methods would become inefficient. Intuitively the Hamiltonian method uses "trajectories" to effectively move in the typical set of the target distribution such that subsequent steps can explore points in the typical set that are far away whereas the conventional MCMC methods have to "diffuse" to explore the same space that yields to slow exploration particularly in higher dimensions.

1.1.5 Parallel tempering

One of the major problems in sampling a probability distribution is the risk of getting stuck in a local maximum of a likelihood surface that has several maxima. In particular, if a multi-modal distribution has modes that are sharp and well separated, any MCMC method can struggle to explore the entire distribution. One solution to this problem is to use parallel tempering [87].

Parallel tempering works by introducing a temperature parameter ($T \geq 1$) that is used to make the target probability distribution more smooth through the relation

$$\pi_T(\theta) \propto \pi(\theta)^{\frac{1}{T}}, \tag{1.6}$$

where $\pi_T(\theta)$ is the tempered target distribution at temperature T and $\pi(\theta)$ is the original target probability distribution. As the temperature increases the target distribution becomes more flat and hence more easy to sample (For a flat target distribution MCMC would be doing a random walk).

Parallel tempering normally uses a Temperature ladder ($T_0 = 1 < T_1 < \dots < T_n$) each with it's associated tempered distribution. Each tempered distribution would be sampled with an MCMC

method running in parallel. This allows for the higher temperature MCMC chains to easily explore the entirety of the target distribution while the MCMC routine at $T = 1$ will be sampling the original target distribution. The MCMC chains in adjacent tempering levels are then allowed to swap their position vectors θ_n and θ_{n+1} once in a while with a probability given by:

$$\min \left[1, \left(\frac{\pi_{T_{n+1}}(\theta_{n+1})}{\pi_{T_n}(\theta_n)} \right)^{(1/T_n - 1/T_{n+1})} \right]. \quad (1.7)$$

This will allow for the MCMC chains at $T = 1$ to efficiently explore the target distribution while ensuring the sampled distribution is not biased. A practical aspect of parallel tempering is the choice of the temperature ladder since it can greatly affect the effectiveness of the algorithm. The optimal choice of the temperatures depend on the structure of the target distribution and hence not known a priori. However, the temperatures can be iteratively adjusted to achieve optimal performance [279].

1.1.6 Model selection and criticism

So far the focus of this section has been Bayesian parameter estimation within the context of a mathematical model and often that's all one needs. However, sometimes there is a need to compare various models with different degrees of complexity and assess their plausibility given the observed data. Here we briefly cover the Bayesian solution to this problem.

It can easily be shown that:

$$\frac{P(M|D)}{P(M'|D)} = \frac{P(D|M)P(M)}{P(D|M')P(M')}. \quad (1.8)$$

This means that the relative probability of models M and M' is proportional to the ratio of their Bayesian evidence known as the Bayes factor or the odds ratio. Usually one can simplify this and assume that $P(M) = P(M')$, that is the prior probability of the models are the same. In this case evaluating the Bayesian evidence is all that is needed to compare the plausibility of various models.

There are various methods to calculate the Bayesian evidence. The main methods are thermodynamic integration [164], nested sampling [254] and approximate methods such as Laplace approximation or various information criteria.

Finally, let us discuss the last step which is known as the “model criticism”. Once a model is selected amongst various models and the model parameters are estimated, one has to “criticize”

the model based on the observed data; this is the extension of the traditional goodness of fit tests such as the reduced Chi-squared test. Based on the finding, the models might have to be revised and the process is repeated until the models can pass the criticism step. In the Bayesian context, predictive checks are usually used for this purpose [35, 121]. Posterior predictive checks works by first choosing a sample θ from the posterior distribution. Then a replica data set D' is generated using the likelihood function $P(D'|\theta)$. Finally the discrepancy between the real data and the model $\mathcal{D}(D|\theta)$ and between the replicated data and the model $\mathcal{D}(D'|\theta)$ are calculated using a discrepancy measure and compared against each other. The discrepancy measure \mathcal{D} can be the Chi-squared measure or any number of other measures to assess various aspects of the model.

1.2 Gravitational lensing

The history of the gravitational lensing in the context of General Relativity goes back to 1911 when Einstein calculated the deflection of light by the Sun [91]. He calculated the deflection angle to be 0.83 arcsec which was wrong by a factor of two due to his use of Newtonian mechanics. He later corrected this result, a prediction which was observationally confirmed in 1919 [71]. During 1920's Eddington [90] and Chowlson [61] entertained the possibility of producing multiple images in a gravitational lensing event. In 1930's Einstein suggested that stellar mass lenses are unlikely to be detected due to small angular separation [92] while Zwicky showed that galaxies can produce large enough image separation to be detected [300, 301]. The next major development in the field was the detection of a strong lensing event of a quasar in 1979 [281] and the reintroduction of the idea of lensing by stellar mass objects by Paczynski [205]. With the advancements in the instrumentation and the use of dedicated surveys, the number of both stellar and extragalactic lensing events has greatly increased in the past few decades which has revolutionized the field of gravitational lensing.

1.2.1 The formalism

Here, we shall present a simple description of gravitational lensing in (asymptotically) flat spacetime, but it can be easily extended to cosmological spacetime.

Figure 1.1 shows the geometry of a gravitational lensing event. As the light rays from the source follow the null geodesics the wave-fronts get distorted and deflected. The angular position of the source with respect to the optical axis is denoted as θ_s , the angular position of the lensed image is named θ_l and α denotes the deflection angle. It could be shown that the two dimensional




Figure 1.1: Schematic view of a gravitational lensing event. The source and the Lens are located at a distance of D_S and D_L from the observer respectively. The source is located at an angle θ_S from the optical axis connecting the lens to the observer. The apparent position of the source is represented as θ_L . α is the deflection angle caused by the lensing effect and b is the impact parameter. Image courtesy of Wikimedia [3].

angular position of the source and the two dimensional angular position of the lens can be mapped to each other using the lens equation given by

$$\vec{\theta}_S = \vec{\theta}_L + \vec{\alpha}(\vec{\theta}_L) = \vec{\theta}_L + \nabla\psi. \quad (1.9)$$

ψ is known as the lensing potential which depends on the mass distribution of the lens and the geometry of the lensing [247] and can be written as:

$$\psi(\vec{\theta}_L) = \int \kappa(\vec{\theta}) \ln|\vec{\theta}_L - \vec{\theta}| d^2\theta, \quad (1.10)$$

where κ is the normalized surface density defined as:

$$\kappa(\vec{\theta}) = \frac{\Sigma(D_L\vec{\theta})}{\Sigma_{critical}}, \quad \Sigma_{critical} = \frac{c^2 D_S}{4\pi G D_L D_{LS}}. \quad (1.11)$$

D_L , D_S and D_{LS} represent the distance of the observer to the lens, the distance of the observer to the source and the distance between the lens and the source as shown in figure 1.1. The Jacobian of the lensing transformation is of particular interest since it provides information on how the areas change under lensing transformation and together with conservation of surface brightness yields the magnification caused by the lensing. The Jacobian matrix $A(\theta_L)$ can be written as:

$$A(\theta_L) = \frac{d\vec{\theta}_S}{d\vec{\theta}_L} \quad (1.12)$$

The Jacobian matrix A is usually expressed in terms of a convergence κ and shear γ_1, γ_2 :

$$A = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix} \quad (1.13)$$

The convergence is a measure of change in the size of the image due to gravitational lensing while shear represents the distortion in the shape of the image. The magnification in the flux μ is given by:

$$\mu = \frac{1}{\det(A)} = \frac{1}{(1 - \kappa)^2 - \gamma^2}, \quad \gamma = \sqrt{\gamma_1^2 + \gamma_2^2}. \quad (1.14)$$

When $\det(A)$ is equal to zero the magnification μ diverges. The locus of all points θ_L in the lens plane that satisfy the condition $A(\theta_L) = 0$ are called the critical curves. The corresponding points θ_S in the source plane form the caustics.

1.2.2 Special cases

Now we look at the special case where the lens is a point source. In this case the lens equation can be explicitly written as

$$\theta_S = \theta_L - \frac{\theta_E^2}{\theta_L}, \quad \theta_E = \sqrt{\frac{4GM D_{LS}}{c^2 D_L D_S}}, \quad (1.15)$$

where θ_E is the angular Einstein radius and M is the mass of the lens. If the lens and the source are perfectly aligned due to symmetry the image is going to be in the form of a ring located at θ_E called the Einstein ring. The lens equation 1.15 has two solutions which correspond to two lensed images located at

$$\theta_{L\pm} = \frac{\theta_S \pm \sqrt{\theta_S^2 + 4\theta_E^2}}{2}. \quad (1.16)$$

One of the images is located within the Einstein ring and the other one is on the outside, separated by $\Delta\theta_{L\pm} = \sqrt{\theta_S^2 + 4\theta_E^2}$. The two images have magnifications given by:

$$\mu_{\pm} = \frac{1}{1 - (\theta_E/\theta_{L\pm})^4}. \quad (1.17)$$

If in a lensing event there are visible images or arcs formed it would be what is known as an strong lensing event. This usually happens for high mass lenses and well aligned lens-source systems. If the deflections are too small and there are no visible images but the effect of lensing can be inferred through statistically modifying the orientation of the background sources the lensing system is in the weak gravitational lensing regime. In some cases there are multiple images formed but due to small angular separation between the images they cannot be distinguished and only the variation in the total flux can be observed. In this case the lensing event is called microlensing. This situation typically arises for stellar mass Galactic lenses. This is evident from the fact that the Einstein radius θ_E is approximately given by:

$$\frac{\theta_E}{10^{-3}\text{arcsec}} \approx \left(\frac{M}{M_{\odot}} \frac{10\text{kpc}}{D} \right)^{1/2}, \quad D = \frac{D_S D_L}{D_{LS}}. \quad (1.18)$$

As can be readily seen the typical image separation for Galactic lensing events caused by stellar mass lenses is of order of a milli arcsecond which is too small to be resolved with optical telescopes. Since stars are moving around and the lensing magnification depends on the source-lens alignment, the total flux received from the microlensed images varies with time. The timescale for this variation is called the Einstein time t_E and is roughly of order of a few months. The change in the brightness of the background source as a function of time due to the microlensing effect is called a microlensing light curve.

1.2.3 Gravitational lensing and time delay

Finally, we shall briefly discuss the time delay caused by the gravitational lensing effect. This is in part due to the geometric effect of bending of the light rays and having to take a different optical path. This effect is proportional to $(\theta_L - \theta_S)^2$. The second effect is the gravitational time delay which is proportional to the lensing potential $\psi(\theta_L)$. One can write the total time delay as

$$\tau(\theta_L) \propto \frac{1}{2}(\theta_L - \theta_S)^2 - \psi(\theta_L). \quad (1.19)$$

As a result there is a time delay between the light reaching an observer from the different images in a strong lensing image. The proportionality constant is sensitive to cosmological parameters such as the Hubble constant H_0 and hence measuring the time delay is of particular interest.

Chapter 2

Forward modelling of quasar light curves and the cosmological matter power spectrum on milliparsec scales

In this chapter we devise an optimal method to measure the temporal power spectrum of the lensing and intrinsic fluctuations of multiply-imaged strongly lensed quasar light curves, along with the associated time delays. The method is based on a Monte-Carlo Markov Chain (MCMC) sampling of a putative gaussian likelihood, and accurately recovers the input properties of simulated light curves, as well as the “Time Delay Challenge”. We apply this method to constrain the dimensionless cosmological (non-linear) matter power spectrum on milliparsec scales (comparable to the size of the solar system), to $\Delta_{\text{NL}}^2 < 4 \times 10^7$ at $k_{\text{NL}} \sim 10^3 \text{pc}^{-1}$. Using a semi-analytic nonlinear clustering model which is calibrated to simulations, the corresponding constraint on the primordial (linear) scalar power spectrum is $\mathcal{P}_{\mathcal{R}} < 3 \times 10^{-9}$ at $k_{\text{L}} \sim 3 \text{pc}^{-1}$. This is the strongest constraint on primordial power spectrum at these scales, and is within an order of magnitude from the standard Λ CDM prediction. We also report measurements of temporal spectra for intrinsic variabilities of quasar light curves, which can be used to constrain the size of the emitting region in accretion disks. Future cadenced optical imaging surveys, such as LSST, should increase the number of observed strongly lensed quasars by 3 orders of magnitude and significantly improve these measurements, even though improvements in modelling quasar accretion and stellar microlensing are necessary.

2.1 Introduction

As a light beam travels through the matter distribution in space to reach us it gets both sheared and focused. If the distortions in the wavefront are large enough, there are going to be multiple images due to strong gravitational lensing. This effect was first observed by [280] for the doubly-imaged quasar (Q0957+561). Being bright compact light sources, quasars can be observed up to high redshifts and are excellent candidates as light sources in strong lensing systems. Many such systems have been since observed and are of particular interest in cosmology. They can be both used to study the properties of dark matter and the cosmological parameters such as the Hubble constant [67, 265].

The time delay between different lensed images is the only parameter which depends on the cosmological length scales and hence can be used to measure the Hubble parameter. This was first proposed by [232] and led to dedicated monitoring programs such as Cosmological Monitoring of Gravitational Lenses (COSMOGRAIL). There have been many different approaches to measure the time delay from long term brightness measurements of the images in strongly lensed systems. A comparison of some of these methods against synthetic data generated by the Time Delay Challenge can be found in [166].

Another area where strong lensing can be particularly useful is in detecting (or constraining) the distribution and abundance of dark matter. For example, it can be employed to study the amount of dark matter in the lens system (e.g. [30]), or the sizes of its substructures (e.g. [69, 142]).

Less well-known is how time variability of the images can also be used to measure the statistics of CDM nanostructures (or microhaloes). This is known as the transient weak lensing effect, and is induced by the moving dark matter microhaloes that cross the lines of sight towards multiply-imaged quasars [227]. In this work, we search for the transient weak lensing signal in strongly lensed quasar systems, leading to constraints on the (linear and nonlinear) dark matter power spectrum. As a by-product of our analysis we also measure the strong lensing time delay, as well as the temporal power spectrum of quasar accretion flow.

The chapter is structured as follows: We describe the data used in this work in Section 2.2. Next, our method is described in detail in section 2.3 and the details of the parameter estimation techniques are discussed in Section 2.4. The limitations arising from the finite size of the light emitting region in source quasars is discussed in section 2.5. Section 2.6 presents the results of applying our method to the Time Delay Challenge data. The results for internal and lensing power spectra, as well as constraints on the Λ CDM linear spectrum are presented in Section 2.7, which are followed by the Conclusions.

2.2 Datasets

Strongly lensed quasar systems have been monitored by several different groups. This includes radio observations such as [241] and optical measurements. Optical measurements include dedicated campaigns such as COSMOGRAIL as well as observations by other groups such as the OGLE gravitational microlensing group [161]. We use the data made publicly available by COSMOGRAIL to demonstrate our method.

COSMOGRAIL is a project aimed at constraining the cosmological parameters by monitoring strongly lensed quasars. It has monitored the light curves of a few well known lensed quasars over the course of a decade in an attempt to measure the time delay between different images. The data consists of R-band light curves for each lensed image. There are six publicly available lensed systems, namely HE 04351223, SDSS J1001+5027, RX J11311231, SDSS J1206+4332, HS 2209+1914 and DES J0408-5354 [1].

To validate our method we have used the Time Delay Challenge (TDC) dataset. It includes thousands of light curves that are made to represent different data quality and observational strategies, as well as many realistic features present in real data such as periods of missing data and the effect of gravitational microlensing by stars in the lensed galaxy. As such, the TDC dataset serve as an independent test to measure the performance of our method.

2.3 Method description

The Λ CDM model predicts a hierarchy of dark matter haloes on different length scales. While baryonic matter can cool and form galaxies in the potential wells of larger dark matter haloes, the cooling time is too long in the smaller halos, which are then non-luminous. As the light coming from a distant quasar travels towards the observer it encounters several dark matter haloes of various sizes, each one inducing an additional weak lensing effect. Since the haloes are moving across the line of sight, the lensing effect is time variable. In [227] a relation between the dimensionless matter power spectrum and the temporal power spectrum of the lensing amplification was derived. In this section, we describe how the temporal lensing power spectrum and the time delay are constrained by strongly lensed quasar light curves. Section 2.5 describes how the constraints on temporal lensing power spectrum are related to the limits on the dark matter power spectrum.

In the following, we describe how the likelihood function for a set of observations depends on the free parameters of the model. The data is the observed magnitude of the quasar images in a strongly-lensed system over a period of time. The model consists of time delays between differ-

ent images and two power spectra: the temporal power spectrum of weak lensing amplification and the temporal power spectrum of intrinsic quasar magnitude variations.

A given image in the multiply-imaged quasar system is labeled with the subscript a . Each measurement of the apparent magnitude for image a is decomposed into three parts, an intrinsic part, $m(t_{ia} + T_a)$, a part caused by the gravitational lensing effect, $L_a(t_{ia})$, and the measurement error, n_{ia} :

$$I_{ia} = m(t_{ia} + T_a) + L_a(t_{ia}) + n_{ia} \quad (2.1)$$

where t_{ia} is the time at which the i th measurement is done for image a . T_a represents the gravitational time delay for image labeled by a (with $T_1 = 0$). The vector I contains the light curve data for lensed images of the same quasar stacked together.

The subscript i runs over different time steps from 1 up to the total number of time steps N_T . The subscript a takes N_I different values where N_I is the number of images. These two indices can be combined into a single index (represented by greek letters) and defined as $\mu = i + N_T \times (a - 1)$. Using this convention, Equation 2.1 can be re-written as:

$$I_\mu = m(t_\mu + T_a) + L_a(t_\mu) + n_\mu. \quad (2.2)$$

The covariance matrix for apparent magnitude measurements can be expanded in terms of power spectra:

$$C_{\mu\nu} = \langle I_\mu I_\nu \rangle = \delta_{\mu\nu} \sigma_\mu^2 + \int \frac{d\omega}{2\pi} e^{i\omega(t_\mu - t_\nu)} [\delta_{ab} P_L(\omega) + e^{i\omega(T_a - T_b)} P_m(\omega)]. \quad (2.3)$$

where $P_m(\omega)$ is the intrinsic temporal power spectrum and $P_L(\omega)$ is the lensing temporal power spectrum. $\delta_{\mu\nu}$ is the Dirac delta function and $\sigma_\mu^2 = \langle n_\mu^2 \rangle$. Note that we have assumed that the lensing effect is uncorrelated across different images, which is the key property we use to distinguish the intrinsic from the lensing temporal power spectra.

Now we divide the relevant part of frequency space into N_F frequency bins and approximate the power spectra using the Heaviside step functions as:

$$\begin{aligned} P_L(\omega) &= \sum_{l=1}^{N_F} p_l K_l(\omega), \\ P_m(\omega) &= \sum_{l=1}^{N_F} p_{l+N_F} K_l(\omega), \\ K_l(\omega) &= \Theta(\omega - \omega_l + \Delta\omega_l) \Theta(\omega_l + \Delta\omega_l - \omega). \end{aligned} \quad (2.4)$$

We assume that $P_L(\omega)$ and $P_m(\omega)$ are even functions while p_l 's are unknown weights to be estimated using the data. The covariance matrix can then be rewritten as:

$$C_{\mu\nu} = \delta_{\mu\nu}\sigma_\mu^2 + \sum_{l=1}^{2N_F} p_l \widetilde{K}_{\mu\nu}^l \quad (2.5)$$

Where $\widetilde{K}_{\mu\nu}^l$ for $1 \leq l \leq N_F$ is defined as:

$$\begin{aligned} \widetilde{K}_{\mu\nu}^l &= \delta_{ab} F^l(t_\mu - t_\nu) \\ \widetilde{K}_{\mu\nu}^{l+N_F} &= F^l(t_\mu - t_\nu + T_a - T_b) \\ F^l(\Delta t) &= \int \frac{d\omega}{2\pi} e^{i\omega\Delta t} K^l(|\omega|) \end{aligned} \quad (2.6)$$

where $\Delta T_{ab} = T_a - T_b$ is the time delay between different images. For each unknown parameter set p_l and ΔT_{ab} , the chi-squared can be written in terms of the data vector I and the covariance matrix C as:

$$\chi^2(p_l, \Delta T_{ab}) = I^t C I + \log[\det(C)] \quad (2.7)$$

Having the likelihood function $\equiv \exp(-\chi^2/2)$, we use a Markov Chain Monte Carlo (MCMC) to explore the parameter space and find the best fit values together with their uncertainties.

The power spectra functions are considered between a minimum and a maximum frequency corresponding to a minimum time scale T_{min} equal to one third of the median of the time difference between data points and a maximum time scale T_{max} which is equal to three times the time difference between the first and the last data points. These choices select the frequency range where our data has the highest constraining power. Since the results are not sensitive to these bounds we adopt these value throughout this chapter. N_F has been set to 9 while the width of frequency bins has been chosen such that they are equality spaced in logarithmic scale. In addition there is a bin at very low frequency to take out the very long scale variations.

2.4 MCMC Sampling

In this section we describe the method used to explore the model parameter space and find their posterior probability density function.

First, we need to define a likelihood function and choose priors on model parameters. In this work, we choose the following form for the likelihood function:

$$\mathcal{L}(p_l, T_a) = \exp\left(-\frac{\chi^2}{2}\right) \quad (2.8)$$

We further choose a flat prior over a reasonably wide range for all the parameters in the model. The details on the ranges are presented with the results below.

Having a large number of parameters, MCMC methods would be a natural choice. We tried a range of MCMC algorithms such as Metropolis-Hastings and Gibbs sampling with adaptive step size tuning but they generally struggled to yield reliable answers and suffered from convergence issues. A combination of Affine-Invariant MCMC [127] and parallel tempering proved to give reliable estimates. Here is a brief description of the algorithm used.

Affine invariant MCMC is a particular form of ensemble sampling that performs equally well on a parameter space mapped by any Affine transformation. In particular it can sample highly skewed distributions with linear correlations very efficiently. It’s also straightforward to parallelize and hence take advantage of the available high performance computing facilities. These methods only have a few hyperparameters and can be efficiently used on a large number of problems with minimal need for tuning. Having a highly irregular and spiky likelihood surface, the Affine-invariant ensemble sampler would spend a long time in local extrema and would suffer from slow convergence. To circumvent this problem the ensemble sampler was combined with a parallel tempering scheme [87].

Parallel tempering makes many copies of the likelihood function modified by a “temperature” parameter:

$$\mathcal{L} \propto \exp\left(-\frac{\chi^2}{2T}\right), \quad (2.9)$$

where $1 \leq T \leq T_{max}$ is the temperature parameter. In our runs generally around 5 tempering levels were used. generally At $T = 1$ we have the original likelihood that we wish to sample. We run an independent ensemble sampler at each temperature and let the chains at different temperatures swap their positions in the parameter space after many Monte Carlo steps. This happens with a probability given by:

$$\min\left(1, \frac{\mathcal{L}(\vec{x}_2, T_2)^{1/T_1 - 1/T_2}}{\mathcal{L}(\vec{x}_1, T_1)}\right), \quad (2.10)$$

where $T_2 > T_1$. \vec{x}_1 and \vec{x}_2 are the positions of the two chains in the parameter space. In this way the high-temperature chains easily move in the parameter space and visit places that would have been difficult for the low-temperature physical chains to visit. By performing position swaps, the physical chains can sample the region allowed by the priors effectively even for hard to sample multimodal distributions. The posterior probability distribution function is then given by the density of the lowest temperature chains ($T = 1$) only.

The choice of temperatures has an important effect on the performance of the sampler. Firstly, the maximum temperature should be high enough to allow the chains to effectively move everywhere within the region permitted by the priors. Secondly, the temperature difference between

adjacent temperatures should be small enough to allow position swaps to happen often. The choice of temperature ladder is not clear *a priori*. We used a method to adaptively tune the temperatures so that we get uniform swapping acceptance rate between the adjacent temperatures. This avoids having a bottleneck in propagation of positions visited in the parameter space by the highest temperature chain to the lowest temperature physical chains [279]. The parallel tempering method could also be trivially parallelized which is very important in our case since likelihood calculations are computationally expensive and the runtimes can be otherwise very long.

2.5 Finite Size Effect

Assuming the quasar is a point source, the lensing temporal power spectrum can be calculated and is given by [227]:

$$\omega P_L(\omega) = 18\pi^2 H_0^4 \Omega_m^{(0)2} \int_0^{\chi_s} \left(1 - \frac{\chi'}{\chi_s}\right) \chi'^2 d\chi' \int_0^\infty dv e^{-v^2/\sigma^2} \left(\frac{v}{\sigma}\right)^2 \frac{\Delta^2}{\omega} (1+z_{\chi'})^3, \quad (2.11)$$

where χ_s is the comoving distance to the quasar, z_{χ} is the redshift at comoving distance χ , σ is the velocity dispersion of dark matter halos and Δ^2 is the dimensionless matter power spectrum. For this work, we adopt $\sigma \simeq 500$ km/s, which is dominated by the cosmological bulk flows on large scales (~ 30 Mpc) [227].

We will recalculate this to take into account the finite size of the quasars' emitting region and generalize this result to include the effect of the finite size of the source. Assuming a radial surface brightness profile given by a function $f(r)$ we find the following formula for the lensing temporal power spectrum:

$$\omega P_L(\omega) = 18\pi^2 H_0^4 \Omega_m^{(0)2} \int_0^{\chi_s} \left(1 - \frac{\chi'}{\chi_s}\right)^2 \chi'^2 d\chi' \int_0^\infty \frac{dv}{\sigma^2} e^{-v^2/2\sigma^2} \int_{\omega/v(1+z)}^\infty \frac{dk_\perp}{2\pi} \frac{\omega(1+z_{\chi'})\Delta^2 F_0^2(k_\perp)}{k_\perp^2 \sqrt{k_\perp^2 - \omega^2/v^2(1+z_{\chi'})^2}}, \quad (2.12)$$




Figure 2.1: Ratio of temporal power spectrum for an extended source of size 0.1 light – day to that of a point source. The horizontal axis is $T \equiv 2\pi/\omega$.

k_{\perp} is the transverse wavenumber and $F_0(k_{\perp})$ is the normalized Hankel transform of surface brightness $f(r)$ which is given by:

$$F_0(k_{\perp}) = \frac{\int_0^{\infty} f(r) J_0[rk_{\perp}(1+z)] r dr}{\int_0^{\infty} f(r) r dr}. \quad (2.13)$$

Figure 2.1 shows the ratio of the lensing temporal power spectrum for an extended source of size 0.1 light – day compared to a point source quasar. The extended source is assumed to be a Shakura-Sunyaev disk [250] radiating as a black body. As can be seen here the finite size of the source can drastically suppress the power on short time scales.

Another factor to consider is the effect of strong lensing on amplifying the fluctuations from the transient weak lensing as well as the intrinsic fluctuations[227]. The effect is an enhancement in the power by a factor of:

$$\left(\frac{2(\kappa^2 + \gamma^2)^{1/2}}{|(1 - \kappa)^2 - \gamma^2|} \right)^2, \quad (2.14)$$

where κ and γ are the convergence and shear.

The values of κ and γ are generally not known for the source quasars so the enhancement cannot be calculated. If $\gamma^2 \ll 1$ the enhancement factor can be approximated as

$$4(-1 + \sqrt{\mu})^2 \mu, \quad (2.15)$$

where $\mu = 1/|(1 - \kappa)^2 - \gamma^2|$ is the lensing magnification [266]. We used the estimated magnifications, calculated through lens modelling for three quasars, namely, HE-04351223 [285], RX-J11311231 [27] and DES-J0408-5354 [6] to find the enhancement factor. For all the other quasars we used the average enhancement for the known quasars which is ≈ 400 .

2.6 Recovering time delays and the time-delay challenge

As discussed in Section 2.3, our model consists of the intrinsic and lensing power spectra and the time delay between the light-curves. Therefore, we can recover the time delay for the quasar images. In this section, we test the ability of our pipeline to recover the correct time delays.

The first test involves generating synthetic light curves and using them in our pipeline. The lightcurves are generated using predefined lensing and intrinsic power spectra and time delays. We then compare the recovered values to the actual input values. The data is generated to mimic the observational strategy adopted by COSMOGRAIL. The time sampling is randomized and the time shift due to strong gravitational lensing is included. The light curves include observational errors and missing data intervals corresponding to non-observing seasons. Figure 2.2 shows an example of such a light curve for two images of a strongly lensed quasar.

For each run several tests are performed to ensure the MCMC chains have converged. For this example light curve the true time delay $\Delta T = 23$ days was recovered as $\Delta T = 25.1 \pm 3.7$ days. In addition the lensing and intrinsic power spectra were recovered as shown in Figures 2.3 and 2.4.

The next test involves using TDC light curves [166]. There are thousands of generated light curves separated into different classes with different data quality and observational strategies. These are called different rungs and the differences include different cadence, total observational timespan and dispersion in the cadence. We chose light curves from all 5 available rungs and compared the recovered time delays to the true values. Table 2.1 shows the results of recovered time delays for randomly selected TDC light curves. The results correspond to a randomly selected light curve from each rung.




Figure 2.2: Synthetic light curves for the two quasar images.

rung	N_I	Time delay[days]	
		Actual	Recovered
0	2	5.78	$5.34^{+0.60}_{-0.54}$
1	2	14.23	$14.06^{+0.09}_{-0.08}$
2	2	28.44	$29.16^{+0.96}_{-0.71}$
3	2	57.53	$56.71^{+0.91}_{-1.29}$
4	2	27.2	$26.91^{+0.29}_{-0.29}$

Table 2.1: Recovered values vs actual time delay values for randomly selected TDC light curves. N_I is the number of images and error bars are 1σ significance level.

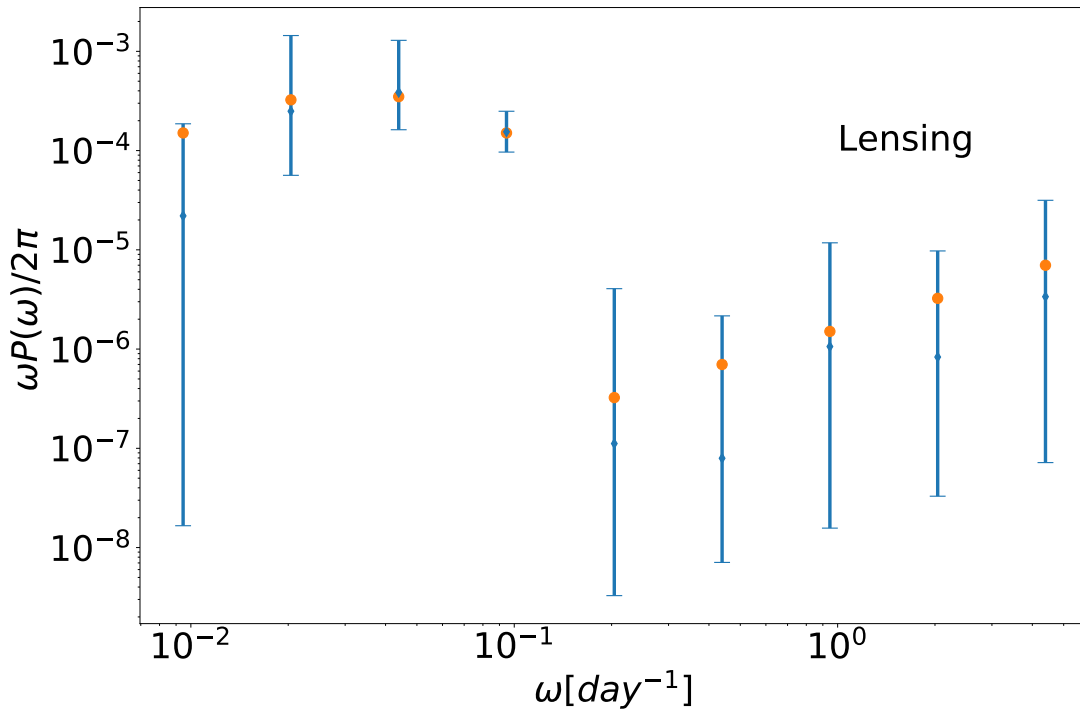


Figure 2.3: Lensing power spectrum for the synthetic quasar light curve measured by our method. The orange dots are the actual values used to generate the light curve. The errorbars show the three sigma uncertainty region.




Figure 2.4: Intrinsic power spectrum for the synthetic quasar light curve measured by our method. the orange dots are the actual values used to generate the light curve. The errorbars show the three sigma uncertainty region.

2.7 Results

In this section, we discuss the results of applying our method to COSMOGRAIL light curves. The results are divided into three separate sections discussing the time-delays, the limits on the dark matter power spectrum and the power spectrum for quasar variability.

Once the limits on the nonlinear dark matter power spectrum are found, one can use the stable clustering hypothesis to transform these into limits on the linear dark matter power spectrum ([209, 257, 297]). Here, we shall use the particle phase space average density (P²SAD) modelling, provided in [297], which is calibrated against numerical N-body simulations [259]. The model is inspired by the stable clustering hypothesis in phase space and supported by the remarkable universality of the clustering of dark matter in phase space as measured by P²SAD across simulated haloes of different masses, environments, and redshifts. On small scales and for primordial power spectra which are reasonably similar to Λ CDM, we can fit the P²SAD predictions for $\Delta_{\text{NL}}^2(k_{\text{NL}})$ by power laws in terms of the linearly extrapolated power spectrum $\Delta_{\text{L}}^2(k_{\text{L}})$:

$$\Delta_{\text{NL}}^2(k_{\text{NL}}) \approx a \left[\Delta_{\text{L}}^2(k_{\text{L}}) \right]^\alpha, \quad k_{\text{NL}} \approx b \left[\Delta_{\text{NL}}^2(k_{\text{NL}}) \right]^{1/3} \times k_{\text{L}}, \quad (2.16)$$

where $a = 0.24$, $b = 1.12$ and $\alpha = 3.05$.

The inferred linear power spectrum $\Delta_{\text{L}}^2(k_{\text{L}})$ can then be converted to the primordial power spectrum $\mathcal{P}_{\mathcal{R}}(k_{\text{L}})$, using the Λ CDM linear transfer and growth functions. For Planck 2015 cosmology the conversion factor is:

$$\Delta_{\text{L}}^2(3 \text{ pc}^{-1}) = 1.5 \times 10^{11} \mathcal{P}_{\mathcal{R}}(3 \text{ pc}^{-1}) \quad (2.17)$$

2.7.1 Time delays

In this section, we present the recovered time-delay values and compare them to the results obtained by the COSMOGRAIL collaboration. Table 2.2 summarizes the results. It shows the time delay values obtained in this work and the time delay values obtained in previous works. It should be noted that in cases where multiple previous estimates existed only one is quoted in the table. The last column in the table provides the references for the quoted time delay values. The fourth column shows the name of the images used for calculating the time delays. The name designations follows that of the corresponding reference given in the last column. There are two sets of results for DESJ0408-5354 in the table. The first result is when only two of the three available light curves were used. The next two lines show the result when all three light curves were fitted simultaneously. The corresponding posterior distributions of the time delays are shown in figure 2.5

Name	Time delay[days]		Images	Ref.
	This work	Previous works		
HE0435-1223	$8.56^{+0.05}_{-0.06}$	8.4 ± 2.1	BA	[31]
RXJ1131-1231	$0.45^{+0.05}_{-0.90}$	0.7 ± 1.0	BA	[270]
HS2209+1914	$22.11^{+2.95}_{-3.33}$	20.0 ± 5.0	BA	[96]
J1206+4332	$109.31^{+2.27}_{-2.34}$	111.3 ± 3.0	AB	[96]
J1001+5027	$116.11^{+2.11}_{-2.62}$	119.3 ± 3.3	BA	[230]
DESJ0408-5354	$113.91^{+39.34}_{-1.46}$	112.1 ± 2.1	BA	[66]
DESJ0408-5354	$113.93^{+26.92}_{-11.58}$	112.1 ± 2.1	BA	[66]
	$151.44^{+47.01}_{-17.57}$	155.5 ± 12.8	DA	[66]

Table 2.2: Time delay values obtained using our method compared to the previous works. The first column is the name of the quasar system. The second and third column present the values found in this work and the values found in previous works. The fourth column lists the images used to calculate the time delays. The image designations follow the conventions in the corresponding reference given in the last column. As we can see in figure 2.5 the asymmetric and/or large errors compared to previous results can be attributed to the extended tail and/or secondary peaks in the posterior. This will not be captured in a Gaussian approximation.

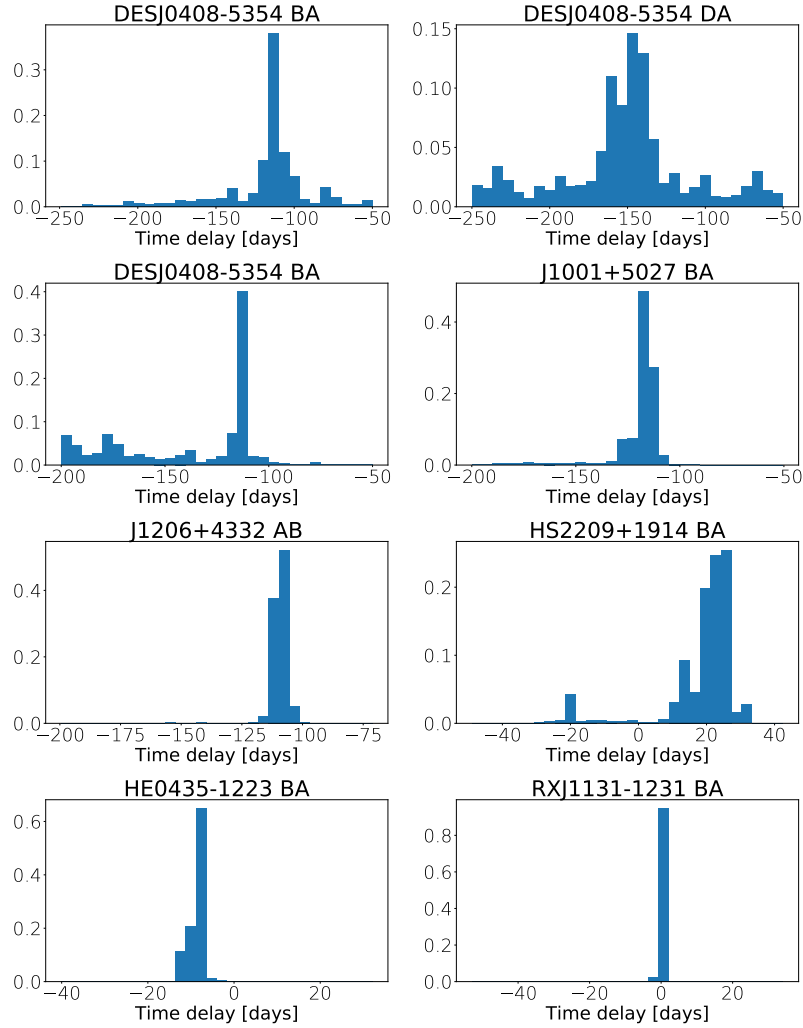


Figure 2.5: Posterior distributions of the time delays from MCMC chains corresponding to the values reported in table 2.2. The title of each panel shows the quasars and image pairs according to table 2.2. The top row shows the two time delays obtained when simultaneously fitting three light curves. The other panels are obtained by only fitting the pair of light curves noted in the title.

2.7.2 Quasar temporal power spectrum

One of the output products of our pipeline is the temporal power spectrum of intrinsic quasar variability. Figure 2.6 shows this measurement for the COSMOGRAIL quasars. The low frequency break in the power spectrum can be used to constrain the size of the accretion disk ([171], [138]). The detailed analysis of disk size will be the subject of a separate paper.




Figure 2.6: Measurements of the temporal power spectrum for quasar intrinsic variability. Error bars are the two sigma confidence intervals. Different colors represent different quasars. The color red is DESJ0408-5354 using two light curves, magenta is the same quasar using three light curves, black is J1001+5027, yellow is J1206+4332, green is HS2209+1914, blue is HE0435-1223 and cyan is RXJ1131-1231.

2.7.3 Dark matter power spectrum

The last output from our analysis is the lensing power spectrum. Figure 2.7 shows the results for COSMOGRAIL quasars. The resulting lensing power spectra are then converted into dimensionless matter power spectra. It should be noted that the lensing power recovered by our pipeline contains both the transient weak lensing signal from dark matter structures and the gravitational microlensing signal from the stars within the lens galaxy and thus should be interpreted as upper limits. These limits depend upon the size of the quasar disk via the finite size effect. We report four sets of limits assuming different disk sizes. The first set is assuming the light emitting region in the quasar disk size is almost a point source at 10^{-5} light – day. This may represent the situation where most of the light comes from a compact hot spot on the accretion disk. The size of the light emitting region in the quasar disk is estimated to be in the range of $\approx 0.1 - 10$ light – day. We plotted four sets of limits for disk sizes of 10^{-5} , 0.1, 1 and 10 light – day. Figure 2.9 shows these limits relative to the Λ CDM predictions using the P²SAD modelling discussed above [297]. The best upper limits are given by the doubly lensed Quasar system J1206+4332 (Black lines in figure 2.9). Figure 2.8 shows the light curves for the two images in this lensed system.




Figure 2.7: Measurements of the lensing power spectrum. Error bars are the two sigma confidence intervals. Different colors represent different quasars. The color red is DESJ0408-5354 using two light curves, magenta is the same quasar using three light curves, black is J1001+5027, yellow is J1206+4332, green is HS2209+1914, blue is HE0435-1223 and cyan is RXJ1131-1231.

2.8 Conclusion and Future Prospects

We presented a novel method to simultaneously fit for the time delays of strongly lensed quasars, as well as the power spectra of their intrinsic variability and the temporal power spectrum of the gravitational lensing, caused by stellar microlensing and dark matter haloes. The recovered time delays are consistent with the previous methods and, depending on the light curve quality, can even yield sub percent level accuracy.

We have presented upper limits on the dimensionless dark matter power spectrum over the

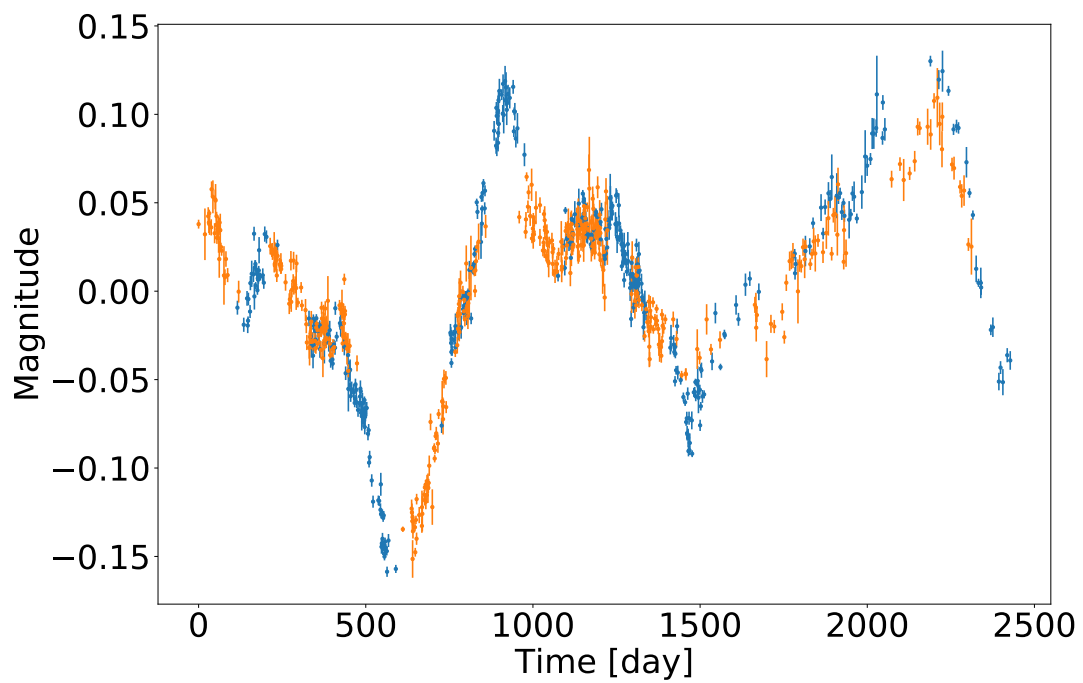


Figure 2.8: The mean subtracted light curves for the two lensed images in doubly lensed Quasar system J1001+5027. The light curves are shifted by the best fit time delay value.




Figure 2.9: Upper limits on dimensionless dark matter power spectrum at the three sigma level. Different sub-plots correspond to different quasar disk sizes given in light days. The lower (upper) grey band shows the prediction for the nonlinear power spectrum using the P²SAD modelling of [297], assuming Λ CDM ($5 \times \Lambda$ CDM) linear power spectrum (the band thickness shows 0.5 dex uncertainty in the substructure volume fraction $0.1 \lesssim f_{\text{subs}} \lesssim 0.4$). Different colors represent different quasars. The color red is DESJ0408-5354 using two light curves, magenta is the same quasar using three light curves, black is J1001+5027, yellow is J1206+4332, green is HS2209+1914, blue is HE0435-1223 and cyan is RXJ1131-1231.

$10^{-7} - 10^{-3}$ pc scales which remain consistent with the Λ CDM predictions, despite the dependence on the size of the emission region in quasar accretion disks. Our strongest limit on the (non-linear) matter power spectrum is on milliparsec scales, and is given by $\Delta_{\text{NL}}^2 < 4 \times 10^7$ for $k_{\text{NL}} \sim 10^3 \text{ pc}^{-1}$. Using the P²SAD model for nonlinear clustering of CDM nanostructure [297], which is calibrated against high resolution N-body simulations, we can translate this limit to a limit on the linear power spectrum. The corresponding constraint on the primordial (linear) scalar power spectrum is given by $\mathcal{P}_{\mathcal{R}} < 3 \times 10^{-9}$ on $k_L \sim 3 \text{ pc}^{-1}$, which is the strongest constraint on these scales, and is within an order of magnitude of the Λ CDM prediction, assuming a power law power spectrum down to these scales. Finally we were able to measure the power spectrum for the intrinsic quasar variability which could help us study the nature of quasar variability and accretion processes.

Future cadence optical imaging surveys, most notably the Large Synoptic Survey Telescope (LSST), are expected to improve the size of the sample of strongly lensed quasars by ~ 3 orders of magnitude, dramatically reducing our statistical errors [200]. However, it is clear that further theoretical modelling in the structure of the emission region in quasar accretion disks, as well as a clean separation of microlensing and transient weak lensing effects (e.g., via the gaussianity of the noise [227]) are necessary to lower the upper limits and/or turn them into a detection.

Chapter 3

Resolving microlensing events with triggered VLBI

Microlensing events provide a unique capacity to study the stellar remnant population of the Galaxy. Optical microlensing suffers from a near complete degeneracy between the mass, the velocity and the distance. However, a subpopulation of lensed stars, Mira variable stars, are also radio bright, exhibiting strong SiO masers. These are sufficiently bright and compact to permit direct imaging using existing very long baseline interferometers such as the Very Long Baseline Array (VLBA). We show that these events are relatively common, occurring at a rate of $\approx 2 \text{ yr}^{-1}$ of which 0.1 yr^{-1} are associated with Galactic black holes. Features in the associated images, e.g., the Einstein ring, are sufficiently well resolved to fully reconstruct the lens properties, enabling the measurement of mass, distance, and tangential velocity of the lensing object to a precision better than 15%. Future radio microlensing surveys conducted with upcoming radio telescopes combined with modest improvements in the VLBA could increase the rate of Galactic black hole events to roughly 10 yr^{-1} , sufficient to double the number of known stellar mass black holes in a couple of years, and permitting the construction of distribution functions of stellar mass black hole properties.

3.1 Introduction

Gravitational lensing presents one of a handful of current observational windows on the dark universe. The distortion of background galaxies by galaxy clusters has provided a striking demonstration of general relativity [21]. More importantly, it has enabled the unique reconstruction

of the projected mass density of the cluster dark matter haloes [205, 231]. While the dynamics of stars and gas probe the dark matter halo within 10^2 kpc of the central galaxy [213], gravitational lensing currently provides the only means to detect and study the halo on larger scales. By combining such measurements, it has become possible to map the dark matter distribution in the local universe [251].

The detailed structure of a gravitationally lensed image may be separated into the intrinsic structure of the lensed object and a distortion imposed by the lens that depends upon the lens location and mass distribution. A typical scale is set by the Einstein angle, this is exactly the angular scale of the Einstein ring cast by a point lens,

$$\theta_E = \sqrt{\frac{4GM(D_S - D_L)}{c^2 D_S D_L}}, \quad (3.1)$$

where M is the total source mass and D_S and D_L are the source and lens distance, respectively. For galaxy clusters this can span many arcminutes.

Gravitational lensing within the Galaxy also occurs when the angular separation of a background star and a foreground object becomes comparable to θ_E . This differs from the gravitational lensing of clusters and galaxies in two important respects. Firstly, the angular size of the lensed image has a typical scale of $\theta_E \approx 1(M/M_\odot)^{1/2}$ mas, well below the resolution of existing optical telescopes. As a result, direct optical imaging has not been feasible. Nevertheless, the conservation of brightness, a consequence of Liouville's theorem, implies substantial magnifications of the total flux associated with the magnified image, typically increasing by more than an order of magnitude. Secondly, the peculiar motions of objects within the Galaxy cause the source-lens system to evolve on time-scales comparable to

$$t_E \equiv \frac{\theta_E}{|\boldsymbol{\mu}_S - \boldsymbol{\mu}_L|}, \quad (3.2)$$

where $\boldsymbol{\mu}_S$ and $\boldsymbol{\mu}_L$ are the source and lens apparent angular velocities, respectively. For typical values in the Galaxy t_E ranges from days to months, depending on the source mass and distance of lens and source from the observer. The relationship between t_E and the lens structure implies a characteristic light curve, referred to as a microlensing event.

The rarity of serendipitous lens-source alignments require the monitoring of large numbers of potential sources to identify candidate microlensing events. For this purpose optical surveys of dense star fields have been undertaken by a number of groups, both to constrain the contribution of stellar remnants to the Galactic dark matter budget [188] and more recently to find extrasolar planets which results in short time-scale features in the microlensing light curve [118]. These

have successfully excluded remnants with masses between $10^{-7}M_{\odot}$ – $10M_{\odot}$ as a candidate for dark matter within the Galactic halo [12, 165].

Motivated by the potential to detect extra-solar planets, recent years have seen the initiation of high-cadence, large-area surveys. These have been enabled by the development of large CCD arrays for telescopes with wide fields of view. As a result, the OGLE¹ (Optical Gravitational Lensing Experiment) collaboration recorded $\approx 2 \times 10^3$ events in 2013, using the Early Warning System [277] to alert the follow-up telescopes for monitoring microlensing events with better time-coverage. The Korean Microlensing Telescope Network (KMTNet) will increase the number of events to 6000 per year with a 10 minute cadence operating continuously [217].

Reconstruction of microlensing events usually only determines the t_E . As a direct consequence, the mass of the lens, and the distances and velocities of the lens and source suffer from a fundamental degeneracy. Methods for breaking the mass-distance degeneracy typically requires higher-order effects. If the microlensing parallax can be measured – an asymmetry in the light curve induced by the non-uniform motion of the Earth – the lens distance and the mass of the lens can be constrained [128, 229]. Parallax effect can also be employed by monitoring microlensing events from earth and space simultaneously. Spitzer has been used to this end, measuring microlensing parallax and thus providing an additional constrain [260]. Alternatively the source finite-size effect could be used to break the degeneracy. If the angular impact parameter is smaller than the source angular size the point source approximation is no longer valid and the finite-size effect becomes important. By analyzing the light curve an additional relation between angular source size and angular Einstein radius could be found. In both cases, the lens and the source apparent motions must be subsequently measured after the event to remove these from consideration [see, e.g., 132], which is only possible for stellar lenses and explicitly excludes those associated with stellar remnants.

Directly resolving the image would greatly simplify the reconstruction of the lens parameters. The typical scales of microlensing events are well matched to those achievable by very long baseline interferometry (VLBI) at centimetre wavelengths, e.g., with the Very Long Baseline Array (VLBA). Unfortunately, imaging microlenses with the VLBA requires compact radio bright sources with typical brightness temperatures in excess of 10^{10} K, well above the typical stellar temperatures. However, the unprecedented rate at which microlensing events are being identified by optical surveys makes it possible to leverage rare source properties for this purpose.

Mira variables provide a natural target for optically triggered radio imaging microlensing experiments.² Miras are asymptotic giant branch (AGB) stars with month time-scale pulsations

¹<http://ogle.astrouw.edu.pl/ogle4/ews/ews.html>

²[141] explore a number of other potential radio-bright microlensing targets, including the continuum emission from a wide class of giants. Of these, only the SiO masers from Miras and active galactic nuclei have sufficiently

largely determined by convection in their envelopes [288]. While the optical luminosities of Miras can vary over three orders of magnitude [233], the typical peak luminosities are on the order of $10^3 L_\odot$, ensuring that they are among the brightest objects in microlensing survey samples [258]. With typical masses similar to $1 M_\odot$ Miras are relatively numerous despite having AGB lifetimes of roughly a million years. Given the number of stellar targets, optical depth of lenses, and typical lensing timescales the rate of microlensing events with hundred percent detection efficiency is given by

$$\frac{N_{events}}{N_{stars} T_{obs}} = \frac{2}{\pi} \frac{\tau}{\langle t_E \rangle}. \quad (3.3)$$

Assuming there are 4×10^4 Mira stars in the bulge (see Section 3.3.2), the optical depth of microlensing events toward the Galactic bulge is $\tau \simeq 2.35 \times 10^{-6}$ [261], and a typical time scale of one month for the microlensing events (i.e. $\langle t_E \rangle \simeq 1$ month) we therefore expect to observe roughly 1 Mira-source microlensing events per year.

Of particular importance here is the presence of SiO masers in the extended atmospheres of Miras. Typical sizes of the masing regions are 1 AU, corresponding to angular sizes of 0.1 mas for distances characteristic of bulge stars [236]. Because of the nonthermal nature of the maser process, the spots can have brightness temperatures as high as 5×10^{10} K, well above the thresholds for imaging with the VLBA [236]. Within the Galactic center, the inner parsec of the Galactic bulge, the SiO masers associated with evolved stars have been imaged with the VLBA at angular resolutions of 0.7 mas, limited by the interstellar electron scattering in that region [235, 236]. Outside of the Galactic center, it should be possible to image these with resolutions approaching 0.3 mas.

Here we explore the rates of suitable microlensing events anticipated by current and future surveys and the precision with which the associated lens parameters can be reconstructed. We propose using optical microlensing surveys to trigger follow-up VLBA observations based on source location in the color-magnitude diagram, anticipated event duration, and the lack of a bright lens counterpart. A subset of these will be Miras, previously identified via the observation of long-time intrinsic variability. Massive lenses are more heavily represented in long-duration events, and those without obvious stellar counterparts are more likely to be associated with massive compact objects.

In section 3.2 we describe a Monte Carlo simulation of microlensing events, discuss the relative rates at which Miras are expected to be lensed by various objects, including stellar remnants, and how event selection can be optimized for imaging based on optical properties. Section 3.3 discusses the absolute rates of radio-bright black hole microlensing events of a variety of potential surveys. Section 3.4 describes the resulting images for a simple maser model and estimates

high brightness temperatures to be amenable to VLBI imaging.

the accuracy with which the lens properties can be reconstructed. The implications for constructing mass and distribution functions for the Galactic remnant population is discussed in 3.5. Finally, conclusions are collected in 3.6.

3.2 Stellar Remnants in Simulated Microlensing Surveys

The rates of microlensing events associated with compact remnants depends on the distribution of sources and lenses within the Galaxy and the detection efficiency of current and upcoming microlensing surveys. Here we describe simulations of an OGLE-like microlensing survey and estimate the number of microlensing events we anticipate to be associated with the various Galactic remnant populations. To do this, we perform mock surveys assuming the monitored stellar field is located in the Galactic bulge.

Necessary inputs are the mass, velocity, color, and magnitude distributions of potential sources and the mass and velocity distributions of potential lenses. For simulation of microlensing events, we adapt the detection efficiency in terms of the Einstein crossing time of OGLE survey sources [290]. We describe the details of the model here.

3.2.1 Galactic Distribution Model and Stellar Lens/Source Population Model

To generate mock microlensing events we require the spatial and velocity distributions of the stellar targets and lenses.

For the structure of the Galaxy, we use thin disk model and standard bulge model from [26]. The density distribution in disk is modeled in cylindrical coordinates by a double exponential function,

$$\rho_D(R, z) = \frac{\Sigma}{2H} \exp\left(\frac{-(R - R_\odot)}{R_d}\right) \exp\left(\frac{-|z|}{H}\right), \quad (3.4)$$

where Σ is the column density of the disk at the Sun position, H the height scale and R_d the length scale of the disk. The distribution of the lens transverse velocity with respect to the line of sight is established from solar motion and the local lens velocity distributions (see Table 3.1).

The bar is described in a Cartesian frame positioned at the galactic center with the major axis x tilted by $\phi = 45^\circ$ with respect to the Galactic center-Sun line. The bar density is given by

$$\rho_B = \frac{M_B}{6.57\pi abc} e^{-r_s^2/2} \quad (3.5)$$

$$r_s^4 = \left[\left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 \right]^2 + \frac{z^4}{c^4},$$

Structure	Parameter	Value	
Disk	Σ ($M_\odot \text{ pc}^{-2}$)	50	
	H (kpc)	0.325	
	R_d (kpc)	3.5	
	velocity	σ_r (km s^{-1})	34.
	disper-	σ_θ (km s^{-1})	28.
	sions	σ_z (km s^{-1})	20.
Bar	M_B (M_\odot)	1.7×10^{10}	
	a (kpc)	1.49	
	b (kpc)	0.58	
	c (kpc)	0.40	
	velocity	σ (km s^{-1})	110
Spheroid	ρ_0 ($M_\odot \text{ pc}^{-3}$)	0.932×10^{-5}	
	a_c (kpc)	0.5	
	velocity	σ (km s^{-1})	120
Halo	$\rho_{H\odot}$ ($M_\odot \text{ pc}^{-3}$)	0.008	
	R_c (kpc)	5.0	
	M in 60 kpc ($10^{10} M_\odot$)	51	
	velocity	σ (km s^{-1})	200

Table 3.1: Assumed parameters of the galactic model, The model is partly adopted from [229]

where M_B is the bulge mass, and a , b and c are the scale length factors. The assumed spherically symmetric dispersion velocity of bulge stars is 110 km s^{-1} . For the density of the spheroid structure we take the following function [240],

$$\rho_{spher} = \begin{cases} \rho_0(a_c/8.5\text{kpc})^{-2.44} & a \leq a_c \\ \rho_0(a/8.5\text{kpc})^{-2.44} & a \geq a_c \end{cases} \quad (3.6)$$

$a^2 = R^2 + z^2/(0.76)^2$, a_c , ρ_c and assumed spherically symmetric spheroid velocity dispersion are given in table (3.1).

Combining the density distribution of Galaxy, kinematics, mass function and stellar population of stars, we can generate microlensing events in terms of t_E , M_l , D_l , D_s , θ_E and π , where π is the parallax parameter that is described in the Appendix A. We use the detection efficiency of OGLE survey in terms of Einstein crossing time (i.e. $\epsilon(t_E)$) to select the observed microlensing events in our simulation [290]. A general overview about the simulation of microlensing events can be found in [226].

3.2.2 Remnant Population Model

In addition to the known stellar populations, gravitational lens candidates include the Galactic population of compact objects. In principle, these may be comprised of the remnants of stellar evolution and primordial objects. We make the conservative assumption that the latter are absent, and consider only the remnants of massive stars.

The end point of stellar evolution depends primarily on the mass of the progenitor. We consider two classes of remnants, white dwarfs and black holes. We ignore the intermediate neutron stars, which we expect to contribute marginally to the high-mass tail of the white dwarf population and otherwise be indistinguishable from them using microlensing observations alone. Hence in practice the remnants are effectively white dwarfs/neutron stars and black holes.

Encoded in the remnant population is the Galactic high-mass star formation history. Once the relationship between the zero-age main sequence progenitor mass (M_{MS}) and the final remnant mass (M_R , where R may be WD or BH) is specified, the mass function³ ($\phi(M_R)$) is given solely in terms of the star formation rate ($\Lambda(t)$) and initial mass function of main sequence stars ($\Phi(M_{\text{MS}})$):

$$\phi(M_R) = \int dt \frac{\Lambda(t - T_{\text{MS}})\Phi(M_{\text{MS}})}{d \log_{10} M_R / d \log_{10} M_{\text{MS}}}, \quad (3.7)$$

where $T_{\text{MS}} \approx 10(M_{\text{MS}}/M_{\odot})^{-1.5}$ Gyr is the lifetime of a main sequence star with mass M_{MS} associated with the remnant mass M_R , and the integration is over the entire Galactic history. Effectively, this is simply the number of remnants formed over the history of Galactic star formation, excluding those stars that remain on the main sequence. For this we adopt the star formation rate from [137], which peaks roughly 10 Gyr ago, and the ‘‘Disk and Young Clusters’’ stellar initial mass function (IMF(M)) from [54], though employing the ‘‘Universal’’ IMF from [163] makes no discernible difference.

We relate M_{WD} and M_{MS} via the piecewise continuous expression presented in [244], obtained for white dwarfs in open clusters by comparing the white dwarf ages inferred by cooling models and cluster age from isochrone fitting:

$$M_{\text{WD}}/M_{\odot} = \begin{cases} 0 & m_{\text{MS}} \leq 0.5588 \\ 0.01m_{\text{MS}} + 0.5418 & 0.5588 < m_{\text{MS}} \leq 1.7 \\ 0.134m_{\text{MS}} + 0.331 & 1.7 < m_{\text{MS}} \leq 4 \\ 0.047m_{\text{MS}} + 0.679 & 4 \leq m_{\text{MS}} < 8 \\ 0 & 8 \leq m_{\text{MS}}, \end{cases} \quad (3.8)$$

³We define the mass function to be the number of objects per logarithmic decade, i.e., $dN/d \log_{10} M$.




Figure 3.1: Assumed mass functions of brown dwarfs and stars (solid), white dwarfs (dashed), and black holes (dotted) in the disk and bulge. The spike features at low (white dwarf) and high (black hole) masses correspond to regions in which the remnant mass is nearly independent of that of the progenitor. These do not substantially impact the results here.

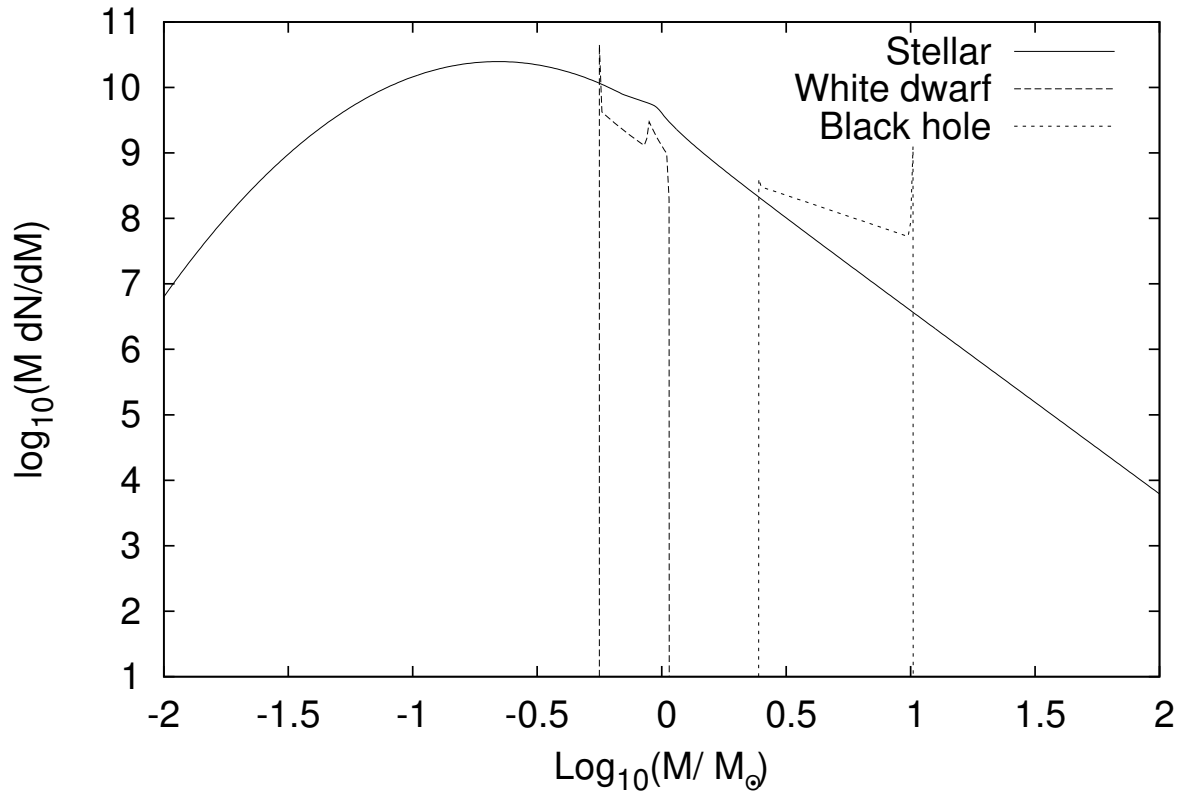


Figure 3.2: Assumed mass function of brown dwarfs and stars (solid), white dwarfs (dashed), and black holes (dotted) in the Spheroid. The spike features at low (white dwarf) and high (black hole) masses correspond to regions in which the remnant mass is nearly independent of that of the progenitor. These do not substantially impact the results here.

where $m_{\text{MS}} \equiv M_{\text{MS}}/M_{\odot}$.

The relationship between M_{BH} and M_{MS} is highly dependent upon the uncertain evolution of massive stars and the details of core collapse supernovae. As a result, above $30M_{\odot}$ it is unclear what the slope of the $M_{\text{BH}}-M_{\text{MS}}$ relationship is [see Figure 9 of 111]. Here we make a moderately conservative assumption that the resulting black hole mass is independent of M_{MS} , setting

$$M_{\text{BH}}/M_{\odot} = \begin{cases} 0 & m_{\text{MS}} \leq 8 \\ 0.3m_{\text{MS}} & 8 < m_{\text{MS}} \leq 33 \\ 10 & 33 < m_{\text{MS}}, \end{cases} \quad (3.9)$$

which lies in the middle of the permitted range presented in [111] and is roughly consistent with narrow black hole mass range inferred from X-ray binaries by [203].

The resulting stellar and remnant mass functions for the disk, bulge and spheroid are shown in Figures 3.1, 3.2. The spikes in the low and high mass ends of the white dwarf and black hole mass functions are associated with regions in which the remnant mass are nearly independent of the progenitor mass, and contain a finite number of objects. For both the disk and the bulge there is a break in the stellar mass function where stars produced during the star formation peak are leaving the main sequence, roughly $1M_{\odot}$.

Because neither white dwarfs nor black holes are expected to experience strong kicks at formation, we assume the same velocity dispersion for the remnants as the surrounding stellar population. This is notably not the case for neutron stars, which often experience sufficient kicks to be launched into the halo, diluting their density in the disk and therefore their representation in the lens sample for current and planned microlensing surveys. The observation of pulsars with ages less than 3 Gyr, indicates a pulsar birth velocity dispersion of about 400 km s^{-1} [144].

3.2.3 Mock Observation Statistics

In order to simulate the microlensing events, we choose the source stars from a Hipparcos-like color-magnitude distribution [212], distributed proportional to the density of matter. After selecting the source star with a given absolute color and magnitude, the position of star in the color-magnitude diagram changes due to the distance modulus and interstellar extinction. Those stars brighter than the limiting magnitude of the microlensing survey (which is taken about 19 in I-band in this simulation) are selected for the observation. It should be noted that the results are not sensitive to the limiting magnitude and could be applied to microlensing surveys with different limiting magnitudes.

We note that not all the microlensing events generated in this simulation are observable in reality. Depending on duration of observation, cadence of photometric data and non-observing nights due to bad weather and technical failures, only a fraction of events as a function of Einstein crossing time, t_E , can be observed. We adapt the detection efficiency function (i.e. $\epsilon(t_E)$) for selecting the observed events from OGLE published function [226, 290].

Each event is constructed by first choosing a random source with the appropriate properties, subsequently a random lens with the appropriate properties, and finally applying the OGLE detection efficiency. That is,

1. Source star selection:

- (a) A source star position is selected with a probability distribution given by the stellar density and geometric factors, i.e.,

$$\frac{dP}{dD_S} \propto \rho(D_S) D_S^2, \quad (3.10)$$

- (b) The source star stellar type is randomly chosen according to Hipparcos Color-Magnitude distribution.
- (c) The stellar component (bulge, disk, spheroid) is selected with weights equal to their density at D_S .
- (d) The source star velocity is the sum of the bulk Galactic motion for the chosen stellar component at the chosen source location and a random contribution pulled from a Gaussian distribution with the appropriate dispersion (listed in Table 3.1).

2. Lens selection:

- (a) A lens position is selected from the stellar density, modified by the lensing cross section, i.e.,

$$\frac{dP}{dD_L} \propto \rho(D_L) \sqrt{\frac{D_L(D_S - D_L)}{D_S^2}}. \quad (3.11)$$

- (b) A lens component (bulge, disk, spheroid) is randomly chosen according to the local densities of each at D_L .
- (c) A lens component type (star, white dwarf, black hole) is randomly chosen according to the local fraction of each for the given component.
- (d) A lens mass is selected according to the total mass function of the chosen component modified by the approximate lensing cross section; for a fixed observing time the probability of detecting a lens is $\propto \theta_E \propto M^{1/2}$.

- (e) A lens velocity is the sum of the bulk Galactic motion for the chosen component and a random contribution pulled from a Gaussian distribution with the appropriate dispersion (listed in Table 3.1).
3. Detection efficiency cut: Given the lens mass, source and lens velocities, and source and lens positions t_E is computed, and the event is accepted with a probability set by the OGLE detection efficiency.

The result is a set of lensing events with known intrinsic properties (source and lens masses, distances, velocities, and types) selected using a realistic microlensing survey biases.

Figure 3.3 shows the distribution in Einstein crossing time of simulated events. Since in the simulation we flag lenses that are either main sequence or remnants and, we can classify lenses based on lens type. Here we compare the overall microlensing events with those produced by black holes specifically; the fraction of events with black hole lenses is 0.035. This fraction is comparable to the estimate of [131], which finds 0.01 assuming a 100% detection efficiency, and thus neglecting the higher detection efficiency of Mira events.

Another important point is the average $\log(t_E)$ for the overall events and black hole events are 1.5 and 2.0, respectively. This is unsurprising since the duration of events scales with lens mass as $M^{1/2}$, and is an immediate consequence of the fact that the average black hole mass is roughly an order of magnitude larger than that of the overall sample. Hence, black hole lensing events are typically more than three times longer than the typical microlensing event, with important consequences for identification and parameter estimation.

Figure 3.4 shows the distribution of lens locations in our simulated survey for each of the Galactic components and lens types. Events are dominated by lenses within the disk (65%) with the remainder from the bulge (35%); the spheroid contributes negligibly. Of note is that the lens distribution within the disk is nearly evenly distributed between the bulge and a heliocentric distance of 2 kpc, implying that microlensing surveys necessarily probe the lens population throughout the intervening disk. This remains true independent of the lens type.

3.3 Radio-Bright Microlensing Event Rates

Key to the radio imaging of microlensing events is the identification of radio-bright lensed sources. In the near term this requires the identification of candidate events based on their optical properties alone. Here we discuss the rates implied by the previous section employing surveys that exploit optical counterparts of radio bright sources (i.e., Mira variables) and the potential




Figure 3.3: Distribution of microlensing events in Einstein crossing time from Monte-Carlo simulation. The detection efficiency of OGLE survey is applied during event selection. The fraction of black hole events to the overall events is 0.035 and weighted heavily towards long t_E .

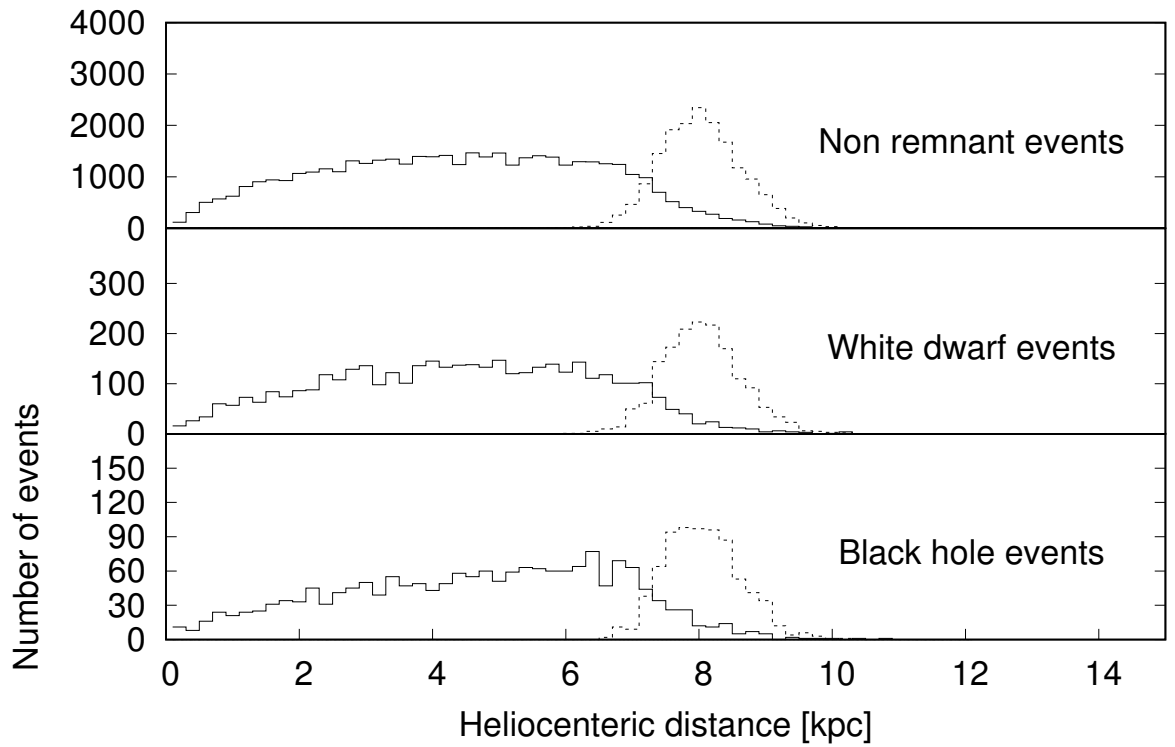


Figure 3.4: Distribution of lens positions in galaxy for various lens populations. Contributions of the disk (solid) and bulge (dashed) are shown separately; the spheroid contribution is negligible.

rates from future radio microlensing surveys. This is ultimately dependent on the population and distribution of candidate sources.

3.3.1 Miras as Microlensing Sources

Miras provide a natural class of optically bright sources with strong radio emission in the form of circumstellar SiO masers. In most microlensing surveys Miras have already been identified due to their variable nature. Where they have not been, they may be crudely identified with in the optical color magnitude diagram (CMD), shown for the OGLE III sample in Figure 3.5. In particular, Miras produce a bright, red cluster that only marginally overlaps with the remainder of the survey targets. From direct inspection of the OGLE III fields we estimate that as many as 5% of the objects near the center of the Mira cluster in Figure 3.5 are identified Miras.⁴ In Section 3.5 we show that there should be nearly 6 times more Miras in the bulge than have been detected thus far, suggesting that it might be possible to increase the fraction of Miras in this region to up to 10%-15%, after accounting for the field of view of the OGLE III survey. Given their high radio brightness, confirmation that a given candidate source is a Mira can then be quickly obtained by direct radio observation with existing large radio telescopes. Equally important is that there is little chance for lensing-induced source confusion – blending of source and lens stars does not move objects appreciably within the region populated by Miras.

Rapid intrinsic stellar variability provides an obvious impediment to the identification and characterization of microlensing events. For this reason variable stars have been generally identified and excluded from past and ongoing surveys. Thus, while there are microlensing events reported by OGLE III that penetrate the Mira cluster (see Figure 3.5), these are almost certainly not associated with Miras. Nevertheless, the presence of such events makes clear that apart from variability there are no intrinsic barriers to including stars lying in the Mira-cluster-region of the CMD in microlensing surveys.

The variable nature of Miras can be mitigated in a number of ways. Typical Miras vary with periods of ≈ 1 yr, I-band variability amplitudes between 1.4 mag and 3 mag [258], and V-band variability amplitudes typically two–three times larger [157]. However, black hole microlensing events can be easily distinguished from the intrinsic variability as a result of three important differences. All of these may be exploited in part due to the high luminosity of Miras, and therefore the high photometric accuracy with which they may be monitored.

First, microlensing light curves have a well-understood structure set by the lens-source geometry and nature of gravitational lensing, characterized by a divergent rise as the source ap-

⁴For this purpose we fit the Mira density distribution with a Gaussian and define “near center” to be within a single standard deviation from the mean.




Figure 3.5: OGLE III Color-Magnitude diagram. Grey points are a subsample of the bulge stars monitored by OGLE III. The red and blue points are the bulge Mira variables and Bulge lensing events observed by OGLE III respectively. The small black interval inside the box shows how much a MIII giant will move in HR diagram as a result of blending with a solar type (G2V) lens star located half way between the source and observer.

proaches the lens caustics (moderated for the largest magnification events by the finite source size; see Section 3.4.2 and Appendix A.1). This differs from the intrinsic variations of Miras. Thus careful light curve fitting should be able to rapidly distinguish between microlensing and intrinsic variability.

Second, the amplitude of the flux variation during microlensing events is typically larger than that due to intrinsic variability. Typical magnifications are of order 200%, while high-magnification events can exceed an 1000%. These correspond to magnitude variations of 0.75 mag and 2.5 mag, comparable to that caused by intrinsic variations in I-band. Hence, the microlensing signal is an order unity modification of the underlying source variability in I-band. At longer wavelengths the intrinsic variability decreases further, on average variations in the magnitude at $1.25 \mu\text{m}$ are 20% of those at optical wavelengths [255]⁵.

Third, the flux variations due to microlensing are achromatic. Again this is in stark contrast to the intrinsic variability of Miras, which varies dramatically with wavelength. Thus, multi-band monitoring should be readily able to disentangle the two variability components.

We make no further attempt to assess the efficiency with which Mira microlensing events can be identified, i.e., we presume that the intrinsic variability of Miras does not result in a significant inefficiency. As a result, our rate estimates may be considered to be optimistic.

3.3.2 Bulge Miras Microlensing Rates

The OGLE III survey monitored approximately 150 million stars in Galactic bulge from 2001 to 2009 in search for gravitational microlensing events and provided us with the largest catalogue of microlensing events available. OGLE used the Early Warning System (EWS) to detect ongoing microlensing events [277]. EWS uses Difference Image Analysis photometry and by analyzing all the stars in the field flags the potential microlensing events [9]. In doing so it currently identifies and filters out variable stars. The flagged stars go through several tests including a visual inspection of the light curve and after satisfying all the criteria they are announced as a microlensing candidate. Figure 3.5 shows the CMD for OGLE III microlensing events [290]. The lensing events have not gone through blending correction for consistency reasons since it cannot be done for non lensed stars.

Miras are a few orders of magnitude brighter than M-dwarfs, increasing their representation in microlensing surveys. An advantage of using bright stars is that blending by background and

⁵There are instances for which near-infrared variability is nevertheless sizable, e.g., R Cas, where the amplitude of variations at $1.25 \mu\text{m}$ are as large as 1.42 mag. However, the variability continues to drop rapidly with increasing wavelength.

lens stars do not affect the brightness of the source star significantly (Figure 3.5). Moreover, from the lens equation, and estimating the mass of the lens star, we can calculate the contribution of blending and correct the position of the source star in the CMD [186, 196].

Within the OGLEIII catalogue there are currently 6528 Miras [258]. The larger number of Miras inferred in the bulge suggest that this can be readily enlarged by surveying the entire bulge region. A rough estimate of the number of Miras in the bulge can be made via the infrared surface brightness.

We follow the approach used in [178]. The method is based on using the empirical relation between Mira number density and infrared surface brightness to infer the total number of Miras in the bulge.

We employ the infrared brightness measurements from Diffuse Infrared Background Experiment (DIRBE) aboard the Cosmic Background Explorer (COBE) satellite. DIRBE made full sky brightness maps in ten infrared bands ranging from $1.25\mu\text{m}$ to $240\mu\text{m}$. For our purpose we make use of the zodiacal-light-subtracted infrared mission-average map at $2.2\mu\text{m}$. The maps are available online at http://lambda.gsfc.nasa.gov/product/cobe/dirbe_prod_table.cfm.

Before using the DIRBE brightness measurements it must be corrected for dust extinction and the contribution from the Galactic disk removed. Following [178] we correct for dust using an empirical relationship between the K-band extinction A_K and infrared colour:

$$A_K = 0.73 \times (-2.5 \log(I_{1.25}/I_{2.2}) + 0.14), \quad (3.12)$$

where $I_{1.25}$ is the $1.25\mu\text{m}$ zodiacal-light-subtracted mission-average DIRBE flux.

To remove the Galactic disk contribution we make use of fits to the infrared brightness maps outside the bulge. [178] adopted an exponential function in Galactic longitude to model the disk

$$I(l, b) = I(0, b) \exp(-|l|/l_0(b)), \quad (3.13)$$

where the scale-height in Galactic longitude $l_0(b)$ is a function of Galactic latitude. This function was fit to the $2.2\mu\text{m}$ DIRBE maps at high Galactic latitudes, $10^\circ < |l| < 45^\circ$. Motivated by the exponential vertical structure of the Galactic disk, we also estimated the disk contribution by fitting an exponential function in Galactic *latitude*. The resulting bulge brightness estimates, and thus the number of bulge Miras, is insensitive to which fitting function is used.

As shown in Figure 3.6, the number of Miras in a given field is strongly correlated with the corrected infrared surface brightness. Like [178] we find a high-quality linear relationship between the Mira number counts in OGLEII fields and the corrected DIRBE $2.2\mu\text{m}$. This remains true when Mira number counts in OGLEIII fields are added. Linear fits to the latter gives an




Figure 3.6: Number of Miras in OGLEII (blue circles) and OGLEIII (red circles) fields vs COBE/DIRBE $2.2\mu m$ flux in the same regions. The best linear fit to OGLEII data is shown by the blue line, consistent with [178].

updated ratio of Mira number density to $2.2\mu\text{m}$ surface brightness of 4.3Sr/MJy , similar to that found by [178]. With a total $2.2\mu\text{m}$ flux of 0.9MJy in the bulge ($-10^\circ < l < 10^\circ$ and $-10^\circ < b < 10^\circ$), this implies that there are around 4×10^4 bulge Miras.⁶

3.3.3 Current and Near-future Optical Survey Microlensing Rates

At least two improved microlensing surveys are underway, exceeding OGLE III in three ways: higher cadence, deeper magnitude limits, and increased sky coverage. OGLE IV, in operation since 2010, has a bulge field of $\approx 330 \text{ deg}^2$, 3.3 times larger than that of OGLE III. The resulting OGLE IV EWS event rate is three times larger than that of OGLE III; OGLE III EWS detected about 650 events in 2008 whereas OGLE IV EWS detected almost 2000 events in 2013 and over 2000 in 2014. The forthcoming The Korean Microlensing Telescope Network (KMTNet) consists of three dedicated 1.6 m telescopes located in Australia, South Africa, and South America, each with 4 deg^2 field of view. Thus, KMTNet anticipates detecting ≈ 2300 microlensing events per year with constant cadence of almost 10 min [140].

The high luminosity of Miras and the long time-scales of black hole microlensing events imply that the first two, higher cadence and deeper magnitude limits, will at best produce modest improvements in the Mira microlensing rate. High cadences will vastly over-sample black hole microlensing light curves, and thus while important for planet searches are unlikely to be helpful for the study of Galactic remnants.

Even within the extincted region near the Galactic center Miras remain visible due to their high intrinsic luminosity. Since optically dim stars do not produce strong radio emission (Stellar photospheres of dwarf stars are essentially undetectable with current radio telescopes), expanding the survey sample to less luminous objects provides little value for the study of black holes. Nevertheless, the deeper magnitude limits of future surveys will permit the inclusion of regions of high extinction, i.e., dark areas of the bulge, in survey fields [249], marginally increasing the number of Miras that can be monitored. This may be ameliorated by operating microlensing surveys in the infrared as well.

Much more important is the expanded sky coverage. We have already assumed that the entire Galactic bulge will be monitored in our previous rate estimate. Further increasing the number of monitored Miras substantially in an optical survey necessarily requires expanding the surveys beyond the bulge. Moderate modifications to the observing strategies of existing optical surveys

⁶This number is insensitive to alternative fitting functions for the Mira number– $2.2\mu\text{m}$ surface brightness relation. We tried linear fits with non-zero offsets, corresponding to an unsubtracted background component, and non-linear fits, none of which substantially changed this number.

present an obvious way in which radio-bright microlensing event rates can be substantially increased. Because Miras are intrinsically luminous, they lie well above the detection threshold of all ongoing microlensing surveys located anywhere within the Galaxy. Therefore, the relative paucity of Miras does suggest an alternative strategy: broad and shallow instead of narrow and deep.

Including the Galactic disk increases the number of stellar targets by a factor of $\approx 6-7$. Assuming the stellar population does not vary widely between the bulge and the disk this would produce an commensurate growth in the Mira sample. Hence, we expect $\approx 2 - 3 \times 10^5$ Miras in total within Milky way. Unfortunately, the microlensing optical depth in the disk is roughly a third of that of the bulge as a result of the former's smaller stellar density. Therefore, after taking into account the disparity in optical depth, the net increase of Mira microlensing events is reduced to a factor of 2-3. As a result, we estimate that a survey that monitors the entire Galactic plane to an I-band magnitude limit of 15 would capture radio-bright microlensing events at a rate of 2 yr^{-1} , corresponding to a radio-bright black hole microlensing events at a rate of 0.1 yr^{-1} .

Importantly, these survey strategies need not be exclusive. Since the objects of primary interest for radio imaging, black holes, result in long-duration microlensing events, even sparse time sampling is sufficient (on the order of once per week). The OGLE IV survey has already adopted such a strategy, monitoring the entire Galactic disk with a cadence of 1-2 days and the bulge with a cadence $\lesssim 1 \text{ hr}$ [278].

3.3.4 Far-future Optical Microlensing Surveys

The LSST (Large Synoptic Survey Telescope) is a 8.4 meter telescope equipped with a wide field camera with 9.6 square degrees field of view. It will provide a deep survey of southern sky (over 20000 square degrees) giving 1000 exposures of each patch during 10 years of observations. The optimal cadence of the survey for different areas of the sky is yet to be determined. The single exposure depth in r-filter will be $\approx 24.5 \text{ mag}$ and it can yet go deeper ($\approx 26.5 \text{ mag}$) by co-adding the images. Being a ground base survey the angular resolution is limited by seeing ($0''.7$) and ultimately by its pixel size ($0''.2$) [147].

The mean cadence of 3 – 4 days is much shorter than the typical time-scales of the long duration black hole lensing events. Furthermore the deep survey combined with the fact that Mira variables are bright stars guarantees that most of the Mira lensing events in the Galaxy, LMC, and SMC will be detected by LSST. Thus, for finding black hole lensing events amenable to radio imaging, the LSST presents a similar capability to current generation microlensing surveys operated with modified observation strategy.

3.3.5 A Future Radio-Continuum Microlensing Survey

Were a microlensing survey performed in the radio directly the need to identify optically-bright radio counterparts would be eliminated altogether, resulting in the finding of events that are amenable to radio imaging with a near 100% efficiency. Moreover such a radio survey would permit monitoring the much more numerous compact continuum radio sources. With the advent of a number of rapid-survey radio telescopes (e.g., the Square Kilometre Array [268], Canadian Hydrogen Intensity Mapping Experiment [19], etc.), enabling rapid nearly all-sky surveys, high-cadence radio transient searches will become common. A radio microlensing survey is a natural byproduct of these.

The stringent requirements on source size ($\lesssim 1$ mas) restricts the potential radio microlensing survey source targets to stars (e.g., Miras) or primarily extragalactic objects. The details of the lens parameter estimation are only weakly dependent on the distance to the radio source; placing sources at extragalactic distances decreases the Einstein angle by at most 30%.

More important is the over-all number of sources, which directly translates into the predicted rate enhancement. These, in turn, depend on the flux limit and wavelength of interest, both set by the VLBA. For a $10 M_{\odot}$ black hole lens multiple image components are easily visible by the VLBA for frequencies above 3 GHz, providing a natural upper limit of 10 cm on the observation wavelength. For 10 min integration times with 500 MHz bandwidths the VLBA flux limit should be near $20 \mu\text{Jy}$ and $40 \mu\text{Jy}$ at 22 GHz and 43 GHz, respectively; where required we adopt a flux limits of $30 \mu\text{Jy}$.⁷ Large-amplification events will reduce the effective flux limit further. As we will see below, either improvements in the sensitivity of the VLBA or new radio-bright source classes will be necessary to leverage a radio microlensing survey.

Radio-bright active galactic nuclei (AGN) provide a natural class of sources, being both numerous, and more importantly compact. Conveniently, they are also typically flat-spectrum radio sources, meaning that neither the source identification nor the subsequent microlensing survey need be performed at the same wavelengths at which the events are ultimately imaged. Thus, to estimate the number of potential sources above the VLBA flux limit we employ the NRAO VLA Sky Survey at 1.4 GHz reported in [64] [though see also 63, 65] finding roughly 1.6×10^7 and 1.5×10^8 radio-bright AGN in ellipticals and spirals, respectively, with fluxes above $30 \mu\text{Jy}$.

These are distributed isotropically, and thus a large fraction of AGN will be visible at high Galactic latitude with correspondingly smaller lensing optical depth. As a result, the average lensing optical depth for an AGN survey to that for the optical surveys that monitor the Galactic

⁷We assume system-equivalent flux densities are roughly 500 Jy and 1000 Jy at 22 GHz and 43 GHz, respectively.

bulge alone is approximately 0.01. Thus, despite having nearly 290 times the number of targets, a radio microlensing survey of objects above $30 \mu\text{Jy}$ would produce a black hole microlensing event rate 1 yr^{-1} , about the same as that from Mira-based optical surveys.

To obtain an event rate of 10 yr^{-1} , doubling the number of known black holes in two years, requires a flux limit of $3 \mu\text{Jy}$. This would require corresponding improvements in the VLBA. However, in principle this may be achieved via an extended integration time combined with an expanded bandwidth. Currently, 4 GHz bandwidths supported by 16 Gbps recorders are planned for stations participating in millimetre-wavelength VLBI observations [283], and could be deployed to VLBA stations for use at centimetre wavelengths. Combined with 2 hr integration times, these reach the $3 \mu\text{Jy}$ flux limit needed. Below $3 \mu\text{Jy}$ the number of sources scale approximately as $N_{>S} \propto S^{-0.5}$, and thus further growth in the number of sources is a slow function of flux limit.

3.4 Radio VLBI Event Reconstruction

Motivated by the prospect of a substantial number of optical microlensing events amenable to radio imaging, we now present illustrative examples of what these may look like for typical event parameters, and discuss the precision with which the event parameters may be reconstructed.

3.4.1 Models of Resolved Masers

SiO masers at $\approx 43 \text{ GHz}$ ($\approx 7 \text{ mm}$) have routinely been detected around late-type giants in the central parsec of the Galaxy [235]. Located within the extended atmospheres of Miras, these are typically within 8 AU of their parent, corresponding to an astrometric offset of 1 mas in the Galactic center. Note that this is comparable to the Einstein angle, θ_E , and therefore the lensing of the maser emission is distinct from that of optical emission of the star due to the finite size effect of the source star [286].

The masing spots are resolved both because of their intrinsic structure and as a result of an interstellar scattering screen that scatter-broadens images of the Galactic center [17, 33, 134]. Because of the latter effect, even point sources exhibit an extended source structure with a full-width half-max (FWHM) of $\approx 0.7 \text{ mas}$, effectively limiting the longest baselines that can be employed by the VLBA. However, this scattering is highly localized, occurring only for sources in the immediate vicinity of the Galactic center, and unlikely to limit the resolution attainable by efforts to image Miras throughout the Galactic bulge.

Typical intrinsic sizes for nearby SiO individual maser spots are 1 AU, corresponding to $\approx 0.1 \text{ mas}$ at $\approx 8 \text{ kpc}$ distance, and therefore also unlikely to significantly limit the resolution of

VLBA observations. We ignore the impact of intervening scatter broadening and model the intrinsic emission from an individual spot by a Gaussian intensity profile with a FWHM of 0.1 mas:

$$I_{\text{int}}(\boldsymbol{\beta}) = I_0 e^{-|\boldsymbol{\beta}|^2/2\sigma_{\text{spot}}^2}, \quad (3.14)$$

where $\sigma_{\text{spot}} = 0.1 \text{ mas} / \sqrt{8 \ln 2} = 0.042 \text{ mas}$.

In practice, SiO masers from Mira variables often form arc-like structures that are dominated by a handful of individual masing spots. For black hole lenses the angular sizes of these structures are typically smaller than θ_E , and thus multiple spots are likely to be strongly lensed simultaneously. For less massive lenses the ring angular scale and θ_E may be more similar. Nevertheless, because the emission between maser spots remains incoherent⁸, the resulting lensed image is a linear superposition of a number of individual spots. Thus, here we consider the simpler problem of a single masing spot to assess the size of the constraints that can be placed on the lens in principle.

Net velocity offsets between the star and masing spots are systematically incorporated into the radio source velocities and thus do not present an additional systematic uncertainty. However, the masing spots can also evolve in size and luminosity. Both are unlikely to produce substantial complications since they don't impact the separation of the lensed images (see Section 3.4.4). While the latter can complicate the determination of the radio light curve, this may be ameliorated via the optical light curve.

3.4.2 Generating Mock Images

Time sequences of mock images are generated via a multi-step process: beginning with the specification of the positions of the source, lens, and observer, the mapping of the source to the image plane via the thin-lens equation, and convolution with a realistic beam.

Over the duration of a lensing event, the position of the source and lens are assumed to evolve with a fixed velocity:

$$\mathbf{x}_S = \mathbf{x}_{S,0} + \mathbf{v}_S t \quad \text{and} \quad \mathbf{x}_L = \mathbf{x}_{L,0} + \mathbf{v}_L t, \quad (3.15)$$

respectively. In practice only the transverse motion is important. In contrast, the observer position, i.e., that of the Earth, is orbital, and thus includes the orbital acceleration:

$$\mathbf{x}_{\oplus} = \mathbf{x}_{\oplus,0} + \int dt \boldsymbol{\Omega}_{\oplus} \times \mathbf{r}_{\oplus}. \quad (3.16)$$

⁸This is distinct from, e.g., for pulsars, which produce many *coherent* spots due to scintillation within the interstellar medium, and therefore require careful consideration of wave-optics effects [211].

The lens geometry is then fully defined by the distances

$$D_L \equiv |\mathbf{x}_L - \mathbf{x}_\oplus| \quad \text{and} \quad D_S \equiv |\mathbf{x}_S - \mathbf{x}_\oplus|, \quad (3.17)$$

and the transverse angular displacement

$$\boldsymbol{\beta} \equiv \frac{(\mathbf{x}_S - \mathbf{x}_\oplus)}{D_S} - \frac{(\mathbf{x}_L - \mathbf{x}_\oplus)}{D_L} = \frac{(\mathbf{x}_L - \mathbf{x}_\oplus)(\mathbf{x}_L - \mathbf{x}_\oplus)}{D_L^2}. \quad (3.18)$$

Neglecting the acceleration in the Earth's motion and the line-of-sight motion of the source and lens permits a simplification of $\boldsymbol{\beta}$ to a linear function of time, though here we will make use of the more general expression in equation (3.18).

For a point-mass lens the observed position on the sky, $\boldsymbol{\theta}$ is given by the thin-lens equation. Usually this is expressed as a potentially multi-valued function of the projected source position. However, we need only to identify the projected source position in terms of the observed position $\boldsymbol{\theta}$, given by

$$\boldsymbol{\beta} = \boldsymbol{\theta} \left(1 - \frac{\theta_E^2}{|\boldsymbol{\theta}|^2} \right). \quad (3.19)$$

Since surface brightness is conserved, this then immediately defines the lensed image in terms of the intrinsic

$$I_{\text{lens}}(\boldsymbol{\theta}) = I_{\text{int}}[\boldsymbol{\beta}(\boldsymbol{\theta})]. \quad (3.20)$$

Finally, we convolve the lensed image with a realistic, if not pessimistic, beam. For this we assume a 43 GHz beam, $B(\boldsymbol{\theta})$ is an anisotropic Gaussian, consistent with the beam in [168]: semi-minor axis of 0.5 mas and semi-major axis of 1.4 mas, oriented 12° East of North. The large aspect ratio is a result of a combination of the low declination of the Galactic center and the North American location of the VLBA antennae. In practice, this resolution is lower than may be achieved by the VLBA at 43 GHz; in [235] the interstellar scatter broadening limited the size of the array that could be effectively employed. Outside of the Galactic center it should be possible to increase the resolution by a factor of 2–3, though we adopt the [168] beam. The resulting observed intensity distribution is then

$$I_{\text{obs}}(\boldsymbol{\theta}) = \int d^2\theta' B(\boldsymbol{\theta} - \boldsymbol{\theta}') I_{\text{lens}}(\boldsymbol{\theta}'), \quad (3.21)$$

where prior to the convolution the angular positions are converted into equatorial coordinates.

All that remains is to specify the relative positions, velocities, and mass of the source and lens. We do this for a handful of illustrative examples in the following sections.

3.4.3 Example Lensed Maser Images

As fiducial source parameters we assume $D_S = 8$ kpc and a proper velocity of 120 km s^{-1} , consistent with sources in Galactic bulge. The fiducial lens parameters are $D_L = 4$ kpc and a proper velocity of 30 km s^{-1} , consistent with the velocity dispersion in the Galactic disk. The relative positions of the source and lens are chosen so that minimum impact parameter is 0.3 mas oriented in declination. These result in typical images that are illustrative of the general situation. All that remains is to specify the mass of the lens.

The first case we consider is that of a $10M_\odot$ black hole, for which $\theta_E = 3.2$ mas. A time-sequence of images for our fiducial case is shown in Figure 3.7, covering nearly a year. Since the minimum impact parameter is small in comparison to θ_E , multiple images are present. Because θ_E is also considerably larger than the intrinsic source and beam sizes, these are well resolved, suggesting that detecting and interpreting the features of strong lensing by black holes will be straightforward. In particular, a direct measurement of θ_E is possible simply by fitting the spot separations, though we defer to how well this may be done in practice to Section 3.4.4.

The same qualitative conclusions hold for a $1M_\odot$ lens, indicative of a white dwarf, shown in Figure 3.8.⁹ Despite the smaller Einstein angle ($\theta_E = 1.0$ mas), multiple images are resolved, again implying that θ_E can be directly measured. Where these differ most dramatically is near peak magnification (center panel in Figures 3.7 and 3.8); the more massive lens produces correspondingly larger distortions in the image as a result of the better resolution of the Einstein ring. The microlensing event is also significantly shorter.

The effect of the accelerated observer frame, due to the Earth's orbital motion, is evident in the asymmetric entrance and exit from the microlensing event. However, this effect may be confused with features arising from the impact of the asymmetric beam and arbitrary event orientation. Thus in Figures 3.7 and 3.8 we have also shown by the red contours the images with the Earth's orbital motion neglected. From these it is clear that asymmetry in the temporal evolution of the image structure at early and late times is a robust indicator of parallax.

3.4.4 Parameter Estimation

The primary parameters of interest are those of the lens: mass, velocity, distance. Given only the magnification light curve these are degenerate, constrained only by the event time-scale and peak

⁹Recall that a solar-type star would not modify the color of the combined lens-source system sufficiently to place it outside of the allowed region in the CMD.




Figure 3.7: Example $10 M_{\odot}$ black hole microlensing event. Impact parameter of 0.3 mas, relative lens-source velocity of 150 km s^{-1} , 1 AU FWHM maser, and radio beam typical of Galactic center observations with the VLBA (semi-minor/major axes of 0.5 mas and 1.4 mas, respectively, with a position angle of 12° east of north). Red contours show images ignoring the orbital motion of the Earth, while black contours include the parallax. For reference, the blue dotted line shows the Einstein ring.




Figure 3.8: Example $1 M_{\odot}$ stellar microlensing event. Impact parameter of 0.3 mas, relative lens-source velocity of 150 km s^{-1} , 1 AU FWHM maser, and radio beam typical of Galactic center observations with the VLBA (semi-minor/major axes of 0.5 mas and 1.4 mas, respectively, with a position angle of 12° east of north). Red contours show images ignoring the orbital motion of the Earth, while black contours include the parallax. For reference, the blue dotted line shows the Einstein ring.

magnification.¹⁰ As a consequence, all that can be recovered even in principle are the projected impact parameter in unites of θ_E , t_E , and the time of maximum magnification. However, images like those in Figures 3.7 and 3.8 introduce at least two key additional observables: θ_E and a lensing parallax.

The peak magnification occurs at the projected closest approach, for which the projected impact parameter is β_{\min} and the magnification is

$$A_{\max} = \frac{\beta_{\min}^2 + 2\theta_E^2}{\beta_{\min} \sqrt{\beta_{\min}^2 + 4\theta_E^2}}. \quad (3.22)$$

The degree to which A_{\max} may be measured depends upon the quality of the photometry that can be performed, and thus is likely to be limited by systematic errors in the radio flux calibration (flux calibration accuracy for the VLBA and the VLA at 43 GHz would be roughly 10%). At this time the separation between the centroid of the primary and secondary images is

$$\Delta\theta = \sqrt{\beta_{\min}^2 + 4\theta_E^2}, \quad (3.23)$$

which may be measured directly from the images to a precision that exceeds the interferometric beam width by up to an order of magnitude for high signal-to-noise detections [235]. Thus, together, these yield a high-precision estimate of θ_E :

$$\begin{aligned} \theta_E &= \frac{\Delta\theta}{\sqrt{2}} \left[A_{\max} \sqrt{A_{\max}^2 - 1} - (A_{\max}^2 - 1) \right]^{1/2} \\ &\approx \frac{\Delta\theta}{2} \left(1 - \frac{1}{8A_{\max}^2} \right), \end{aligned} \quad (3.24)$$

where the latter expression is for large A_{\max} , and a fractional precision of

$$\frac{\sigma_{\theta_E}^2}{\theta_E^2} \approx \frac{\sigma_{\Delta\theta}^2}{\Delta\theta^2} + \frac{\sigma_{A_{\max}}^2}{4A_{\max}^6}. \quad (3.25)$$

When the uncertainty is dominated by the astrometric uncertainty, this implies that for a $10M_{\odot}$ black hole θ_E can be measured to within $\approx 1\%$.

Even without detailed event modeling, θ_E immediately enables the determination of the relative proper motion from equation (3.2):

$$\boldsymbol{\mu}_L - \boldsymbol{\mu}_S = \frac{\theta_E}{t_E} \hat{\mathbf{u}}, \quad (3.26)$$

¹⁰This may be improved substantially through the exploitation of the impact of parallax on the light curve [128, 229].

where $\hat{\mathbf{u}}$ is the asymptotic direction of the vector between the first and second images. The source proper motions can be directly measured by follow-up observations; the angular velocities of individual maser spots have been measured with an accuracy better than 1 mas yr^{-1} . As with the source positions, this is limited by the electron scattering [235], and thus for sources beyond the central pc it should be possible to measure this nearly an order of magnitude better, giving a typical accuracy of 0.1 mas yr^{-1} , corresponding to a physical velocity of 4 km s^{-1} at the distance of the Galactic center. With a typical lens apparent velocity of 30 km s^{-1} , comparable to the velocity dispersion within the disk, this induces to a roughly 10% uncertainty in the inferred ω_L .

The remaining degeneracy with D_L may be broken via the measurement of a parallax from the microlensing event itself. The sources of most interest also are expected to have the longest t_E , and thus encompass a substantial fraction of a year. As made explicit in Figures 3.7 and 3.8, long t_E permit measurements of the impact of the Earth’s orbital motion, i.e., a parallax. This comes in two forms.

First, the light curve itself is asymmetric as a result of the modified evolution of the Earth-lens-source alignment [256]. Figure 3.9 shows the effect of parallax on light curve of a $10 M_\odot$ black hole lensing event. Since the Earth accelerates substantially during this period, the additional component is distinguishable from the otherwise unknown but essentially fixed velocities of the source and lens, resulting in an asymmetry in the magnification light curve. While the maser emission and optical stellar emission are not spatially coincident, the offset in the distance is much smaller than D_S , and thus D_L can be reconstructed from either the optical or infrared light curves. Typical uncertainties of %10 in radio flux calibration may preclude the use of radio light curve for this purpose. Since Miras are among the brightest stars in the microlensing survey fields, they may be good candidates for obtaining D_L in this way, assuming the underlying variability can be adequately modeled.

Second, as described in the previous section, the underlying impact on the lensed images themselves is directly measurable. Note that this is essentially the same effect, the asymmetry in the magnification light curve arises from the asymmetry in the evolving image structure. However, unlike the parallax effect in the light curve, the image structure is insensitive to the Mira variability. While it is possible to attempt a full fit to the sequence of images, most of the information is contained in the locations of the multiple images. Thus, shown in Figure 3.10 is the angular separation of image components as a function of time for the sequence of images in Figure 3.7.

The mock data shown in figure 3.10 provide a convenient way to estimate the accuracy with which the parallax parameter can be reconstructed. We compute the centroid positions of the multiple image components by performing a maximum-likelihood fit of the positions of beam-convolved point sources. The uncertainty in the centroid is constructed assuming a flux limit

of 10% of the maximum brightness of an unlensed maser spot, and introduce corresponding Gaussian fluctuations in the mock data (see in Fig. 3.10). At late times the uncertainties grow as a result of the dimming of one of the two images far from peak magnification.

As a simple model we write the impact parameter as a linear combination of impact parameters with and without parallax effect included:

$$\beta = \lambda\beta_p + (1 - \lambda)\beta_{np}, \quad (3.27)$$

where β_{np} and β_p are the angular separations when the orbital acceleration of the Earth is ignored and included assuming $D_L = D_S/2$. As shown in Appendix A.3, the interpolation parameter is related to the lens distance by $\lambda = (D_S - D_L)/D_L$, and hence a measurement of λ corresponds to a measurement of the lens position. That is, keeping all angular measurements fixed, fitting the evolving image separation directly probes the lens distance.

A maximum-likelihood fit recovers $\lambda = 0.84 \pm 0.20$, where the 1σ errors are indicated. That is, for typical parameters the impact of parallax can be detected within the radio images at more than 4σ . For a given λ the corresponding estimates of the lens distance and its uncertainty are

$$D_L = \frac{D_S}{1 + \lambda} \quad \text{and} \quad \frac{\sigma_{D_L}^2}{D_L^2} = \frac{\sigma_\lambda^2}{(1 + \lambda)^2}, \quad (3.28)$$

and thus, we recover $D_L = 4.3 \pm 0.5$ kpc. That is, typically, D_L can be reconstructed with an accuracy of roughly 10%.

With an estimate of the lens distance it is possible to reconstruct the lens mass and transverse velocity from θ_E and ω_L . The mass may be estimated given measurements of θ_E and λ via

$$M = \frac{c^2}{4G} \theta_E^2 \frac{D_S}{\lambda}. \quad (3.29)$$

Assuming the uncertainty in θ_E is negligible and in D_S is of order 10%, the fractional uncertainty in M is given by that in λ , and thus the mass can be estimated with an accuracy of roughly 14%. The lens transverse velocity is given by

$$\mathbf{v}_L = D_L \boldsymbol{\omega}_L. \quad (3.30)$$

Assuming the typical uncertainties of %10 in both quantities, the transverse velocity can be typically reconstructed to 14% as well.




Figure 3.9: Example light curve for the microlensing event shown in Figure 3.7. The blue dashed and solid red lines correspond to when the effect of parallax is neglected and included, respectively. On ± 50 day time-scales the impact of parallax is clearly visible, resulting in roughly 30% asymmetries in the light curve.

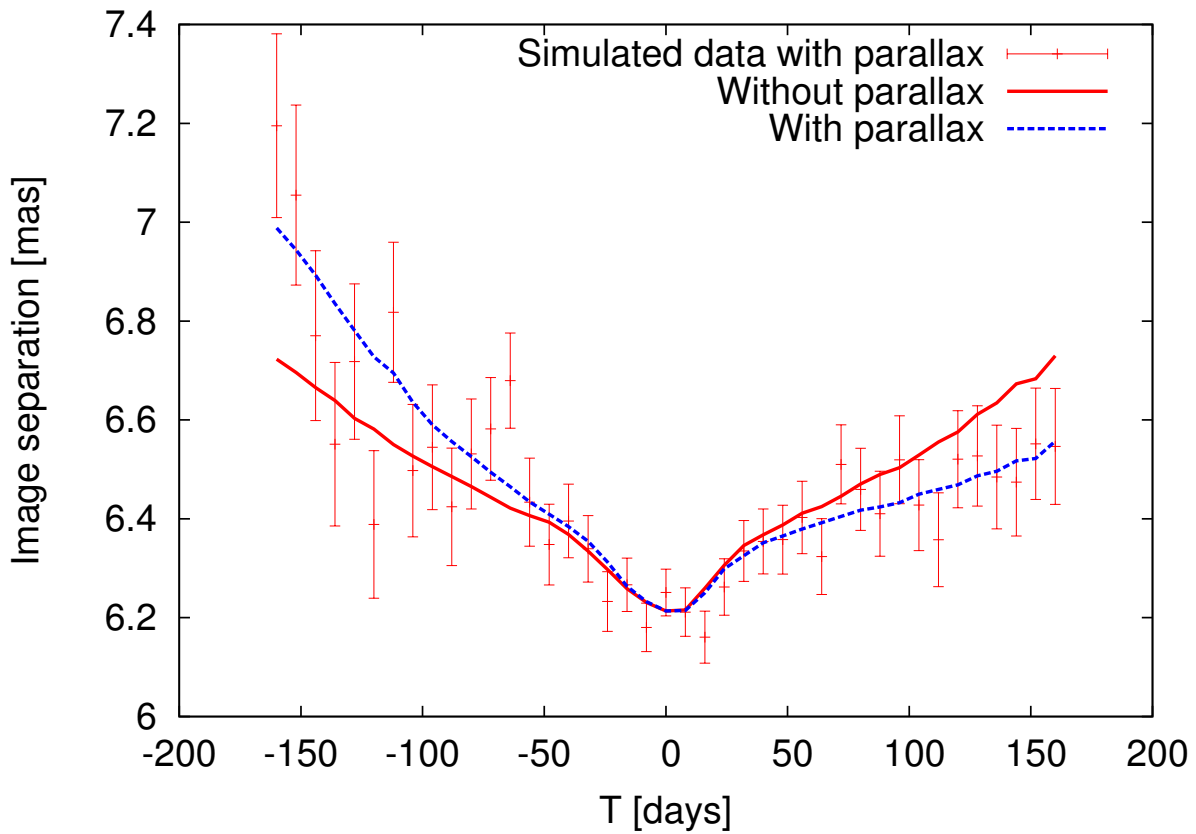


Figure 3.10: Evolution of the image separations for the microlensing event shown in Figure 3.7. Centroid errors are constructed assuming a flux limit of 10% the maximum flux. The dashed and dotted lines correspond to the expected image separations when the effect of parallax is neglected and included, respectively. As with the light curve in Figure 3.9, the largest impact of parallax occurs at early and late times, causing deviations of roughly 50% of a beam width on year time-scales.




Figure 3.11: Microlensing event example for a black hole-black hole binary system of equal mass ($10 M_{\odot}$ each) with orbital period of 30 yr. Impact parameter of 0.3 mas, relative lens-source velocity of 150 km s^{-1} , 1 AU FWHM maser, and radio beam typical of Galactic center observations with the VLBA (semi-minor/major axes of 0.5 mas and 1.4 mas, respectively, with a position angle of 12° east of north). Red contours show images ignoring the orbital motion of the Earth, while black contours include the parallax. For reference, the blue dotted line shows the Einstein ring.

3.4.5 Binaries

Resolving the lensed images provides a natural way to distinguish between single and double lenses, i.e., isolated and binary black holes. The latter, black hole-black hole binary systems are of particular interest as a potential gravitational wave source [5]. At the same time they have the potential to both inform and leverage the observations of X-ray binaries [see, e.g., 237]. Figure 3.11 shows a typical binary image, which may be immediately distinguished from the Figure 3.7 by both the complicated lensed-image morphology and, more directly, the presence of a third image.

The efficiency with which binaries can be detected depends strongly on their projected angular separation. Sufficiently compact binaries will appear as a single lens, while the components of sufficiently wide binaries will produce independent lensing events. After conducting a suite of numerical experiments with binaries in different orientations, peculiar velocities, and impact parameters, we have found that binaries with projected angular separations spanning an order of magnitude about the Einstein angle can be readily identified. The case shown in Figure 3.11 has a binary angular separation comparable to the individual Einstein radii of the black holes, 3.2 mas.

As with the isolated case, the motion of the Earth induces shifts in the image locations due to parallax, and thus the lens distance can in principle be extracted. Unlike the isolated case, typically the image morphology during maximum magnification is complicated. This is mitigated by the fact that the parallax signal is largest when the lens-source separation is largest, and thus where the binary will appear most point-like. Nevertheless, we leave binary parameter estimation for future work.

3.5 Discussion

With moderate changes in strategy existing and future optical/infrared microlensing surveys should achieve black hole event rates of roughly 0.1 yr^{-1} . Dedicated radio microlensing surveys could reach event rates an order of magnitude higher. As a result, radio imaging of microlensing events can provide the novel ability to produce a large sample of stellar remnants with known masses, positions (including distances), and velocities. Importantly, this sample would not suffer from the standard biases that plague studies based on binaries; the selection effects associated with microlensing surveys are well understood and can be modeled and removed. This enables a variety of probes of astrophysical phenomena, which we discuss here.

3.5.1 Massive Star Evolution

The large event rates enables the systematic construction of a black hole mass function over a decade. For radio microlensing surveys (see below) this rate can be an order of magnitude larger. Unlike X-ray binaries, for microlensing the observing biases are well understood and can be addressed via direct modeling, resulting in an accurate representation of the mass distribution of stellar mass black holes in the Milky Way.

The mass distribution of black holes would be immediately diagnostic of the dynamics of core collapse supernovae, the cataclysmic events surrounding their formation. Typically, accretion of fallback material produces a mass gap between neutron stars and black holes, a region of relative paucity in the black hole mass function extending to masses well above $2M_{\odot}$ [272, 289, 298]. In all cases, key uncertainties in the stellar evolution (e.g., wind loss rates, supernovae energetics, etc.) impact the final mass distributions [115]. There have been studies on the effect of variations in neutrino mechanism for core collapse supernovae on the theoretical black hole mass function [210] and using the observed black hole mass function there already have been constraints on core collapse supernova [162]. For compact binary systems, the location and depth of the mass gap is further sensitive to the binary evolution history [112]. While tentative evidence for a mass gap has already been reported based on the inferred masses of black holes in X-ray binaries [203], this is necessarily subject to the variety of strong selection effects and substantial uncertainties mentioned above. Moreover, studies of X-ray binaries are fundamentally limited by the small number of binaries currently known.

In contrast, radio-imaged microlensing events provide a means to systematically accrue large, debiased samples of Galactic black holes, limited only by survey duration. Thus, they offer both a method to assess the poorly known biases inherent in X-ray binaries and ultimately to directly access key elements of massive star evolution.

Furthermore, radio-imaged microlensing events will produce estimates of the lens positions and tangential velocities. Thus, together with the mass function it is possible in principle to construct a Galactic black hole distribution function. The velocity distribution will possibly be strongly impacted by the dynamics of core collapse supernovae via supernova kicks [284], as well as the subsequent evolution of the Galactic remnant population.¹¹

Galactic black holes also serve as a fossil record of massive star formation. Therefore, detailed comparisons of the local¹² densities of stellar mass black holes, white dwarfs, and sub-solar stars yield an estimate of the potentially temporally and spatially varying stellar initial mass func-

¹¹While the dynamical relaxation time of the Galaxy is of order 10^4 Gyr, black holes are much more massive than the typical Galactic object, and thus can have relaxed substantially.

¹²Necessarily accounting for potentially significant differences in relaxation rates.

tion. In particular, this provides a unique ability to constrain the history of the high-mass end of the initial mass function.

3.5.2 Gravitational Wave Experiments

The measurement of the mass of a single isolated stellar-mass black hole would provide an important calibration for population synthesis computations, performed primarily to produce rate estimates for gravitational wave detectors [22]. This is the consequence of the typical nature of lenses in contrast to the particular evolutionary history of X-ray binaries, currently used to calibrate the variety of uncertainties in the rate estimates.

More generally, the binarity rate and source distribution are key inputs into black hole binary gravitational wave source population estimates. High-mass X-ray binaries, the objects likely progenitors of the black hole-black hole binaries visible by experiments like LIGO [4], are necessarily young by virtue of the large mass of the secondary, providing a strong bias towards populations associated with the recent star formation history of the Milky Way. In stark contrast, binary mergers are expected to occur very long times after their formation, and are therefore indicative of the integrated star formation history of the Milky Way. The detection of a black hole-black hole binary presents a means to study the binary black hole population over the entirety of the Galactic history. Unfortunately, the lower limit on the projected binary separation for which radio-imaged microlensing events can efficiently detect binarity (roughly 1 mas at 8 kpc) corresponds to orbital periods of years, and thus will not present candidate gravitational wave sources. Nevertheless, it would provide a significant calibration of binary formation rates, and therefore black hole-black hole binary populations.

3.5.3 Isolated Neutron Stars

Thus far we have focused primarily on black hole lenses. Nevertheless, it is possible to directly detect solar mass objects as well. While we have included a discussion of the white dwarf lensing rates we have largely ignored lensing events from isolated neutron stars. An estimate of the neutron star lensing rate is nevertheless possible using the black hole lensing rate. The two key assumptions are that the massive stellar progenitors of black holes and neutron stars follow a known initial mass function (here we assume Salpeter) and that the neutron stars receive kicks during formations.

From the first we estimate the ratio of the total number of Galactic neutron stars to the total

number of Galactic black holes:

$$\frac{N_{\text{NS}}}{N_{\text{BH}}} \approx \left(\frac{M_{\text{ZNS}}}{M_{\text{ZBH}}} \right)^{-1.35} \approx 4, \quad (3.31)$$

where $M_{\text{ZNS}} \approx 8 M_{\odot}$ and $M_{\text{ZBH}} \approx 21 M_{\odot}$ are the minimum zero-age main-sequence masses of stars that form neutron stars and black holes, respectively [289].

The second assumption modifies the volume occupied by neutron stars. Stars formed at an initial radius r_{init} with a natal kick (v_{kick}) larger than the circular velocity (v_{circ}) will isotropize within a radius of

$$r_{\text{max}} \approx r_{\text{init}} e^{v_{\text{kick}}^2 / 2v_{\text{circ}}^2} \approx 7r_{\text{init}}, \quad (3.32)$$

where we have used typical values $v_{\text{kick}} \approx 400 \text{ km s}^{-1}$ and $v_{\text{circ}} \approx 200 \text{ km s}^{-1}$. Thus, neutron stars formed in the Galactic bulge will be distributed across a volume nearly 400 times larger than similarly formed black holes, and thus exhibit a number density 400 times smaller. The net result is that for typical numbers the rate of neutron star microlensing events is expected to be roughly 1% of that for the black holes, justifying our neglect of neutron stars.

However, we caution that this conclusion is extremely sensitive to the typical kick velocities. If the typical kick velocity is 300 km s^{-1} the neutron star lensing rate rises to 10% of the black hole lensing rate. More importantly, if the neutron star kick velocity distribution contains a low-velocity tail or is bimodal, as suggested by the large number of known Galactic neutron stars [see, e.g., 108, 214], the bulge may retain a large fraction of neutron stars originally formed within it. In this cases the neutron star and black hole lensing rates can be comparable. Thus, the relative frequency of neutron star and black hole events provides an additional probe of the distribution of neutron star formation kicks.

If large numbers of neutron star lensing events are observed, radio-imaged microlensing provides a novel way in which to directly measure the neutron star mass function, independent of the biases inherent in the study of neutron star binaries. Unfortunately, given their large typical age ($\gtrsim 1 \text{ Gyr}$) these are all likely to be exceedingly dim, with luminosities of $\approx 10^{-11}$ - $10^{-9} L_{\odot}$, and therefore not amenable to direct size measurements, complicating any effort to directly constrain the high-density nuclear equation of state.

Nominally, these objects would also fall below the death line for typical pulsar magnetic field strengths (10^{12} G) and periods (1 s). However, due to magnetic field decay, implicated by the low surface fields in recycled millisecond pulsars, these objects could still exhibit observable magnetospheric emission; typical surface fields after a Gyr would then be expected to lie near 10^{11} G . Nevertheless, many would still live near the death line, and therefore would at best be transient radio sources. Hence, in this case the radio emission of isolated neutron stars found by radio-imaged microlensing would provide an unbiased probe of neutron star magnetization and its evolution.

3.6 Conclusions

VLBI observations of radio-bright microlensing events affords the ability to resolve the multi-component structure of microlensing events at radio wavelengths. Such images would break the degeneracies between distance, velocity, and source mass inherent in studies of the light curve alone by providing a direct measurement of the size of the Einstein ring and the presence of signatures of parallax in the evolving image structure and/or asymmetric light curves. Thus, radio imaging enables the reconstruction of the lens parameters to better than 15% accuracy. The dominant source of this remaining uncertainty is due to the unknown distance and velocity of the lensed source, and thus these may be improved by roughly an order of magnitude by follow-up observations that constrain the source's proper motion and distance. Furthermore, imaging provides an immediate method to detect binary lenses, albeit within a rather narrow range of orbital separations (corresponding to periods of roughly 2 yr–20 yr). It is noteworthy that this remains the case when the lens belongs to the otherwise unobservable Galactic remnant population, e.g., neutron stars and black holes.

Radio-bright microlensing events can be identified using existing optical/infrared microlensing surveys. Mira variables exhibit SiO masers in their envelopes and therefore provide a natural radio-luminous target. Moderate modifications to existing survey strategies, some of which are already pursued by OGLE-IV, should produce events amenable to radio imaging at a rate of $\approx 2 \text{ yr}^{-1}$. Of these, assuming the Galactic black hole population arises solely via the evolution of massive stars, roughly 0.1 will be due to black holes. This rate will increase substantially if black holes compose a significant fraction of the Galactic dark matter budget.

Detections of Galactic black holes via microlensing provides an unbiased sample of massive star remnants. The existing sample of stellar mass black holes, obtained via observations of X-ray binaries, necessarily suffers from strong, uncertain biases associated with the formation and evolution of tight massive-star binaries. As a consequence, the detection and characterization of even a handful of black holes using radio-imaged microlensing events will inform the late-stage evolution of massive stars, the energetics of their subsequent supernovae, and the event rates of current and future gravitational wave experiments.

Currently, the rate of radio-bright microlensing events is limited by the number of compact radio sources above the detection limits of existing VLBI facilities, and in particular the VLBA. Nevertheless, an order of magnitude increase in event rates can be obtained by improvements in the flux limits of VLBI observations, achieved, e.g., through increased collecting area (as could be provided with long baselines in the next-generation VLA), expanded bandwidths and longer integration times, and dedicated radio-continuum microlensing surveys. Therefore, it is possible that in the near future it will be possible to double the number of known Galactic black holes with

a survey extending two years, and begin the statistical study of the Galactic black hole properties within a decade.

Chapter 4

Understanding EROS2 observations toward the spiral arms within a classical Galactic model framework

EROS (Expérience de Recherche d'Objets Sombres) has searched for microlensing toward four directions in the Galactic plane away from the Galactic center. The interpretation of the catalog optical depth is complicated by the spread of the source distance distribution. We compare the EROS microlensing observations with Galactic models (including the Besançon model), tuned to fit the EROS source catalogs, and take into account all observational data such as the microlensing optical depth, the Einstein crossing durations, and the color and magnitude distributions of the catalogued stars.

We simulated EROS-like source catalogs using the HIPPARCOS (HIgh-Precision PARallax COLlecting Satellite) database, the Galactic mass distribution, and an interstellar extinction table. Taking into account the EROS star detection efficiency, we were able to produce simulated color-magnitude diagrams that fit the observed diagrams. This allows us to estimate average microlensing optical depths and event durations that are directly comparable with the measured values.

Both the Besançon model and our Galactic model allow us to fully understand the EROS color-magnitude data. The average optical depths and mean event durations calculated from these models are in reasonable agreement with the observations. Varying the Galactic structure parameters through simulation, we were also able to deduce constraints on the kinematics of the disk, the disk stellar mass function (at a few *kpc* distance from the Sun), and the maximum contribution of a thick disk of compact objects in the Galactic plane ($M_{thick} < 5 - 7 \times 10^{10} M_{\odot}$ at 95% confidence, depending on the model). We also show that the microlensing data toward one

of our monitored directions are significantly sensitive to the Galactic bar parameters, although much larger statistics are needed to provide competitive constraints.

Our simulation gives a better understanding of the lens and source spatial distributions in the microlensing events. The goodness of a global fit taking into account all the observables (from the color-magnitude diagrams and microlensing observations) shows the validity of the Galactic models. Our tests with the parameter excursions show the unique sensitivity of the microlensing data to the kinematical parameters and stellar initial mass function (IMF).

4.1 Introduction

Following Paczyński’s seminal publication ([205]), several groups initiated survey programs beginning in 1989 to search for compact halo objects within the Galactic halo. The challenge for the EROS (Expérience de Recherche d’Objets Sombres) and MACHO (MASSive Compact Halo Objects) teams was to clarify the status of the missing baryons in our own Galaxy. In September 1993, the three teams, EROS ([15]), MACHO ([13]), and OGLE (Optical Gravitational Lensing Experiment, [276]), discovered the first microlensing events in the directions of the Large Magellanic Cloud and the Galactic center (GC). Since these first discoveries, thousands of microlensing effects have been detected in the direction of the GC together with a handful of events toward the Galactic spiral arms (GSA) and the Magellanic Clouds.

Microlensing has proven to be a powerful probe of the Milky Way structure. Searches for microlensing toward the Magellanic Clouds (LMC, SMC) and M31 (survey MEGA ; [68] and survey AGAPE; [199]) provide optical depths through the Galactic halo, allowing one to study dark matter in the form of massive compact objects. Searches toward the Galactic plane (GC and Galactic spiral arms) allow one to measure the microlensing optical depth of ordinary stars in the Galactic disk and bar. Kinematical models and mass functions can also be constrained through the event duration distributions.

Several teams have published results about the Galactic structure, through microlensing searches in the Galactic plane, such as MACHO ([215]), EROS ([136]), OGLE ([264]), and MOA (Microlensing Observations in Astrophysics, [16]). The EROS team is the only group that have searched for microlensing toward the Galactic spiral arms, away from the Galactic center. As a matter of fact, the EROS team have measured the microlensing optical depth toward four directions of the Galactic plane (Fig. 4.1), *i.e.*,

- γ Sct ($\bar{b} = -2.1^\circ, \bar{l} = 18.5^\circ$),
- γ Nor ($-2.4^\circ, 331.1^\circ$),

- β Sct ($-2.2^\circ, 26.6^\circ$),
- θ Mus ($-1.5^\circ, 306.6^\circ$),

as far as 55 degrees in longitude away from the Galactic center ([224]). The distinguishing

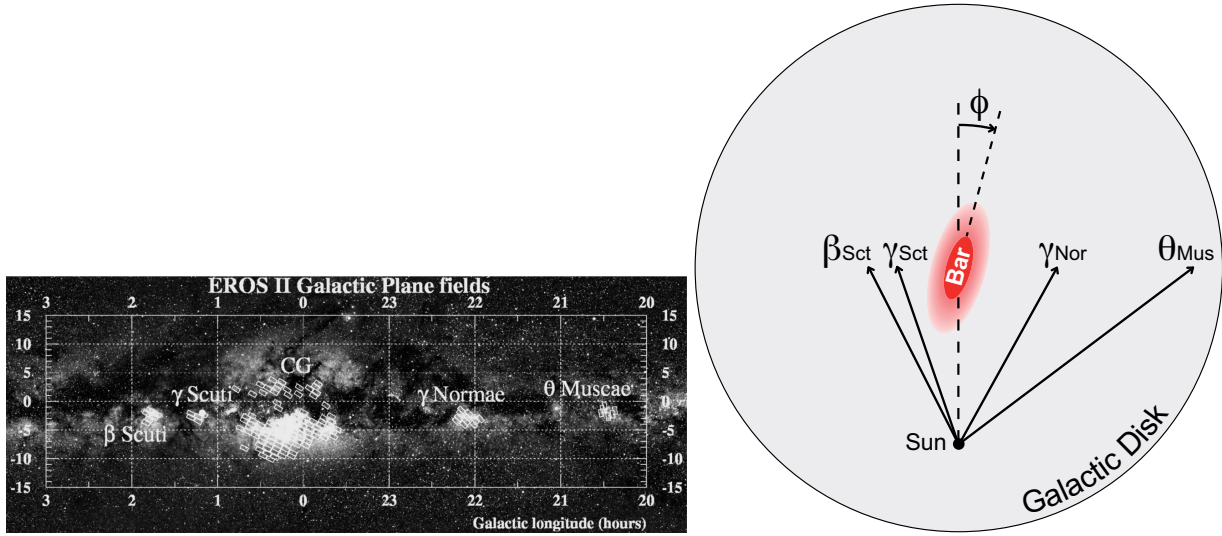


Figure 4.1: Four directions toward the Galactic spiral arms monitored by EROS.

feature of these measurements with respect to other targets like SMC or LMC is the widespread distribution of the distances of the monitored sources. The distances to the sources could not be individually measured and both their average and dispersion are poorly estimated. The concept of “catalog optical depth” was introduced in [224], and in this chapter we describe a complete procedure to compare measured optical depth with model predictions. After the introduction of the microlensing concepts (Sect. 4.2) and the presentation of the EROS data (Sect. 4.3), in Sect. 4.4 we describe the technique to produce synthetic color-magnitude diagrams (CMDs), via the HIPPARCOS catalog (HIGH-Precision PARallax COLlecting Satellite [95], [273]), the spatial distribution of mass from Galactic models, and the absorptions tabulated in a 3D map obtained with infrared observations ([177]). We cross-checked the obtained local stellar number densities with the expectations from the stellar initial mass function (IMF). In Sect. 4.5, we describe the full simulation of the EROS program, in terms of CMDs taking into account the stellar detection efficiency of EROS, and in terms of the microlensing events. Our fitting procedure is described in Sect. 4.6, where we derive constraints on our simple Galactic model and test the Besançon model

([240]); the fit takes into account the observed CMDs as well as the data from the microlensing (optical depths and mean event durations) toward the four observed lines of sight; we use the fit to estimate the allowed range of our simple Galactic model parameters. In the final discussion (Sect. 4.7), we extract from the best fit the distance distributions of the sources and lenses. Finally, we discuss the sensitivity of microlensing observations toward the Galactic arms to the dark thick disk, central bar inclination, stellar mass function and disk kinematics.

4.2 Microlensing effect

The gravitational microlensing effect occurs when a massive compact object passes close enough to the line of sight of a star to produce a temporary magnification of the source. A general overview of the microlensing formalism can be found in [248] and [226]. In the approximation of a single point-like lens deflecting the light from a single point-like source, the total magnification of the source luminosity at time t is given by ([205])

$$A(t) = \frac{u(t)^2 + 2}{u(t) \sqrt{u(t)^2 + 4}}, \quad (4.1)$$

where $u(t)$ is the distance of the deflecting object to the undeflected line of sight, expressed in units of the Einstein radius R_E given by:

$$\begin{aligned} R_E &= \sqrt{\frac{4GM}{c^2} D_S x(1-x)} \\ &\simeq 4.54 \text{ A.U.} \times \left[\frac{M}{M_\odot} \right]^{\frac{1}{2}} \left[\frac{D_S}{10 \text{ kpc}} \right]^{\frac{1}{2}} \frac{[x(1-x)]^{\frac{1}{2}}}{0.5}. \end{aligned} \quad (4.2)$$

Here G is the Newtonian gravitational constant, D_S is the distance of the observer to the source, and $x D_S = D_L$ is its distance to the deflector of mass M . Assuming a deflector moving at a constant relative transverse speed v_T , reaching its minimum distance u_0 (impact parameter) to the undeflected line of sight at time t_0 , $u(t)$ is given by

$$u(t) = \sqrt{u_0^2 + \left(\frac{t-t_0}{t_E} \right)^2}, \quad (4.3)$$

where $t_E = R_E/v_T$, the lensing timescale, is the only measurable parameter bringing useful information regarding the lens parameters in the approximation of simple microlensing,

$$t_E \sim 79 \text{ days} \times \left[\frac{v_T}{100 \text{ km/s}} \right]^{-1} \left[\frac{M}{M_\odot} \right]^{\frac{1}{2}} \left[\frac{D_S}{10 \text{ kpc}} \right]^{\frac{1}{2}} \frac{[x(1-x)]^{\frac{1}{2}}}{0.5}. \quad (4.4)$$

4.2.1 Microlensing event characteristics

The so-called simple microlensing effect (point-like source and point-like lens with uniform relative motion with respect to the line of sight) has some characteristic features that allow one to discriminate it from any known intrinsic stellar variability. These features are as follows: given the low probability for source detector alignment within R_E , the event should be singular in the history of the source (as well as of the deflector); the magnification is independent of the color; the magnification is a simple function of time, depending on (u_0, t_0, t_E) , with a symmetrical shape; as the geometric configuration of the source-deflector system is random, the impact parameters of the events must be uniformly distributed; the passive role of the lensed stars implies that their population should be representative of the monitored sample at any given source distance, particularly with respect to the observed color and magnitude distributions.

This simple microlensing description can be complicated in many different ways: for example, multiple lens and source systems ([174]), extended sources ([292]), and parallax effects ([129]); these complications will not be discussed here.

4.2.2 Observables: optical depth, event rate, and t_E distribution

The optical depth up to a given source distance, D_S , is defined as the instantaneous probability for the line of sight of a target source to intercept a deflector's Einstein disk, which corresponds to a magnification $A > 1.34$. Assuming that the distribution of the deflector masses is described by a density function $\rho(D_L)$ and a normalized mass function $dn_L(D_L, M)/dM$, this probability is

$$\tau(D_S) = \int_0^{D_S} \int_{M=0}^{\infty} \frac{\pi\theta_E^2}{4\pi} \times \frac{\rho(D_L)}{M} \frac{dn_L(D_L, M)}{dM} dM 4\pi D_L^2 dD_L, \quad (4.5)$$

where $\theta_E = R_E/D_L$ is the angular Einstein radius of a lens of mass M located at D_L . The second term of the integral is the differential number of these lenses per mass unit. As the solid angle of the Einstein disk is proportional to the deflectors' mass M , this probability is found to be independent of the deflectors' mass function

$$\tau(D_S) = \frac{4\pi G D_S^2}{c^2} \int_0^1 x(1-x)\rho(x)dx, \quad (4.6)$$

where $\rho(x)$ is the mass density of deflectors located at a distance $x D_S$. This expression is used when the distance to the monitored source population is known (for example, toward the LMC and SMC).

When the monitored population is spread over a wide distance distribution, as is the case toward the Galactic plane, we have to consider the concept of “catalog optical depth” as introduced in [224]; the mean optical depth toward a given population defined by a distance distribution $dn_S(D_S)/dD_S$ of target stars is defined as ([189])

$$\langle \tau \rangle = \frac{\int_0^\infty \frac{dn_S(D_S)}{dD_S} \tau(D_S) D_S^2 dD_S}{\int_0^\infty \frac{dn_S(D_S)}{dD_S} D_S^2 dD_S}. \quad (4.7)$$

Again, this optical depth does not depend on the deflectors’ mass function. On the other hand, for a given optical depth, the microlensing event rate depends on the deflectors’ mass distribution as well as on the velocity and spatial distributions.

Contrary to the optical depth, the microlensing event durations t_E and consequently the event rate (deduced from the optical depth and durations) depend on the deflectors mass distribution as well as on the velocity and spatial distributions. The statistical properties of the durations and event rates can therefore provide global information on the dynamics of the Galaxy and on the mass distribution, which complement other observational techniques based on direct velocity and luminosity measurements.

In this chapter, the optical depth together with the observed event rate and more precisely the duration distributions are compared with simulations to constrain the mass, shape and kinematics of the lensing structures.

4.3 EROS data toward the Galactic spiral arms

In this section, we recall and summarize the EROS2 CCD observations and microlensing results toward the Galactic spiral arms, and describe the efficiencies and uncertainties needed to allow comparisons with simulations. Fig. 4.2 shows the observation time span with the average weekly sampling toward the four targets discussed here. We only provide the information on the data that is relevant for our simulation; more details on the original data can be found in ([224]).

4.3.1 EROS color-magnitude diagrams

The stars detected in EROS are statistically described by their color-magnitude diagrams given in Fig. 4.3 in the (I_C, V_J) photometric system, hereafter simply noted (I, V) . The published EROS-CMDs provide for each catalog, labeled (C) , the observed stellar density $n_C(I, V - I)$ per square degree, magnitude, and color index, as a function of I and $V - I$, sampled in 0.3×0.2 cells ([225]). When using these CMDs, one has to take into account the following uncertainties:

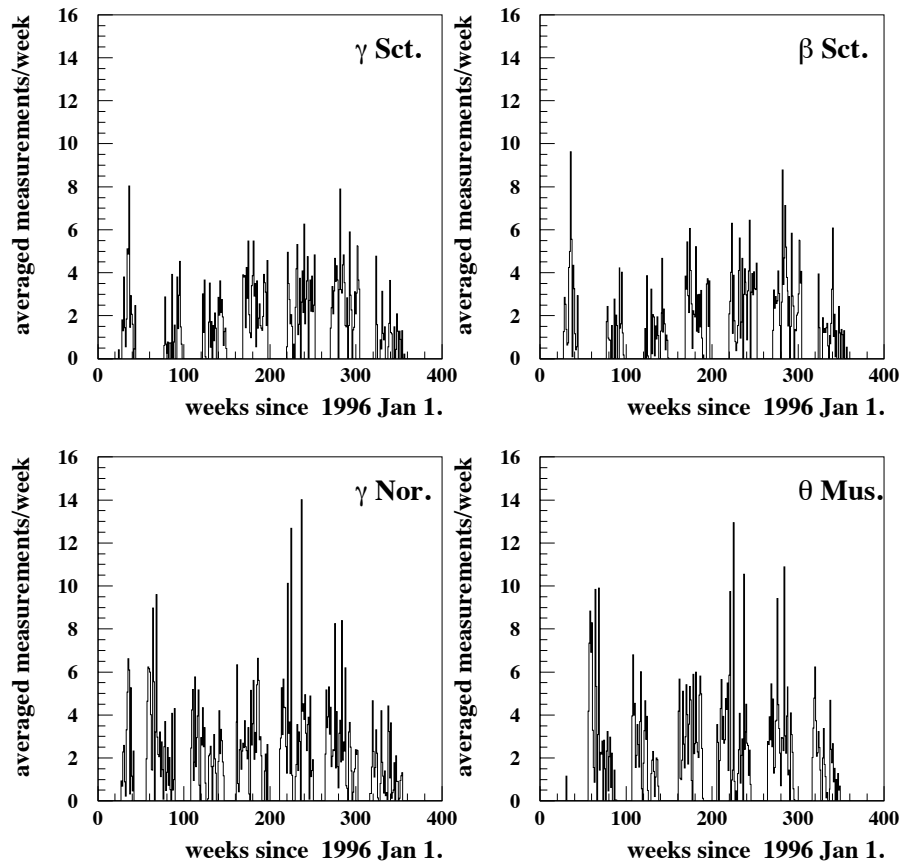


Figure 4.2: Time sampling toward the 4 monitored targets in the Galactic spiral arms: average number of measurements per star and per week.




Figure 4.3: Relative color-magnitude diagrams $n(I, V - I)$ of the EROS catalogs toward the 4 directions toward the Galactic spiral arms. The gray scale gives the number density of stars per square degree, unit of magnitude, and unit of color index.

- Each stellar number density $n_C(I, V-I)$ value is affected by a statistical uncertainty coming from the propagation of the Poissonian noise in the original EROS catalogs, as explained in the header of the published EROS-CMD ([225]).
- Each $n_C(I, V-I)$ value is affected by a systematic uncertainty of $\sim 5.3\%$, owing to the uncertainty on the size of the effective EROS field; this uncertainty is common to all catalogs.
- Another systematic uncertainty is due to the residual 0.07 magnitude EROS calibration uncertainty ([28]), which affects the attribution of a star to a given $[I, (V-I)]$ cell. It has to be taken into account for each EROS color, and therefore induces a systematic uncertainty of $[0.07, 0.16]mag.$ in the $[I, V-I] \equiv [R_{EROS}, (B_{EROS} - R_{EROS})/0.6]$ system.

To generate an ‘‘EROS-like’’ catalog from a model for comparison purposes, one needs to use the efficiency of EROS to detect stars and the photometric uncertainties, both defined in the EROS photometric system $[R_{EROS}, B_{EROS}] \equiv [I, I + 0.6(V-I)]$. The EROS stellar detection efficiency has been studied in ([224]), by comparing EROS data with HST data ([2]). Since we found that an object detected in B_{EROS} is systematically detected in R_{EROS} , the EROS stellar detection efficiency can be parametrized as a function of the relative magnitude B_{EROS} only (Fig. 4.4). The EROS photometric errors on the magnitudes and colors are parametrized as

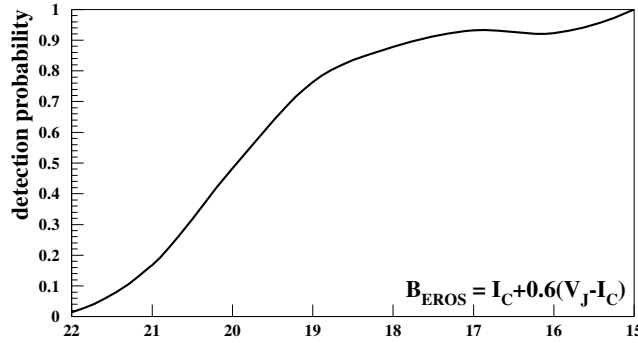


Figure 4.4: Star detection probability in EROS vs. the relative magnitude $B_{EROS} = I + 0.6(V - I)$.

$$\delta I = \sqrt{0.1^2 + \left[\frac{2.5}{\ln 10} \right]^2 \left[\frac{\sigma_\Phi}{\Phi} \right]_{R_{EROS}}^2 \frac{1}{N_{meas.}}}, \quad (4.8)$$

$$\delta(V - I) = \sqrt{0.1^2 + \left[\frac{1}{0.6} \frac{2.5}{\ln 10} \right]^2 \left(\left[\frac{\sigma_\Phi}{\Phi} \right]_{R_{EROS}}^2 + \left[\frac{\sigma_\Phi}{\Phi} \right]_{B_{EROS}}^2 \right) \frac{1}{N_{meas.}}},$$

where the 0.1 constant term (dominant for stars brighter than ~ 18) is a residual uncertainty, as estimated from EROS calibration studies using DENIS catalog data ([94])¹, $[\sigma_{\Phi}/\Phi]$ is the relative image-to-image dispersion of the successive flux measurements given by Fig. 4.5, and N_{meas} is the number of observations (exposures) used to estimate the mean flux of a star, *i.e.*, 268 toward β Sct, 277 toward γ Sct, 454 toward γ Nor and 375 toward θ Mus.




Figure 4.5: Photometric point-to-point precision along the EROS light-curves vs. $R_{EROS} = I$ (upper) and B_{EROS} (lower). Vertical bars in I show the dispersion of this precision in the EROS catalog. The histograms show the magnitude distribution of the full EROS spiral arm catalog (all directions).

Table 4.1 summarizes some of the key numbers regarding the color-magnitude statistical data. When comparing the data with simulations, we focus on the stars brighter than $I = 18.4$, the most reliable part of the EROS-CMD, with the highest and best controlled stellar detection efficiency.

¹This irreducible uncertainty is attributed to the variability of the stellar spectra within the very wide EROS passbands.

4.3.2 Microlensing results

Table 4.1 provides the microlensing results from EROS ([224]). The σ_{t_E} values differ from the values published in table 3 from ([224]) because they were biased, since we assumed large statistics for their estimates. To properly take into account the statistical fluctuations on small numbers, we therefore re-estimated σ_{t_E} from the expression,

$$\sigma_{t_E}^2 = \frac{1}{N_{events} - 1} \sum_{events} (t_E - \bar{t_E})^2, \quad (4.9)$$

where N_{events} is the number of microlensing events toward the target.

The average microlensing detection efficiency of the EROS survey was estimated in [224]; it is defined as the ratio of events satisfying the EROS selection cuts to the theoretical number of events with an impact parameter $u_0 < 1$, and was found to be almost independent of the target, since the time samplings were very similar. Figure 4.6 shows this efficiency as a function of the Einstein duration of the events t_E .




Figure 4.6: Microlensing detection efficiency of the EROS survey toward the Galactic spiral arms, as a function of the event characteristic duration t_E .

4.4 How to synthesize an EROS-like color-magnitude diagram

We now compare the data with realistic simulations. In this section we describe how our modeling takes into account all the known observational constraints and discuss how to handle the

Target	θ Mus	γ Nor	γ Sct	β Sct
$\langle \alpha^\circ \rangle$	200	245	278	281
$\langle \delta^\circ \rangle$	-64	-52	-13	-6
$\langle b^\circ \rangle$	-1.46	-2.42	-2.09	-2.15
$\langle l^\circ \rangle$	306.56	331.09	18.51	26.60
field (deg ²)	3.8	8.4	3.6	4.3
$N_{stars}^{I < 18.4}$	$2.28 \cdot 10^6$	$5.24 \cdot 10^6$	$2.38 \cdot 10^6$	$3.0 \cdot 10^6$
$\rho_* \times 10^{-6} \text{deg}^{-2}$	0.60	0.62	0.66	0.70
$\rho_*^{I < 18.4} \times 10^{-6} \text{deg}^{-2}$	0.245	0.23	0.28	0.34
N_{event} with $u_0 < 0.7$	3	10	6	3
$\tau \times 10^6$	$.67^{+.63}_{-.52}$	$.49^{+.21}_{-.18}$	$.72^{+.41}_{-.28}$	$.30^{+.23}_{-.20}$
$\overline{t_E}$ (days)	97 ± 75	57 ± 10	47 ± 6	59 ± 9
σ_{t_E} (days)	98	31	14	12

Table 4.1: Data and results toward the 4 regions monitored in the EROS spiral arms program. Average coordinates, field extensions, numbers of bright stars ($I < 18.4$), surface densities of all stars, of the bright stars, and the measured microlensing optical depth and duration parameters are provided for each target.

specific difficulties of this kind of analysis.

We generated apparent color-magnitude diagrams based on the following hypotheses and ingredients from direct observations:

- The HIPPARCOS catalog ([95], [273]) provides the magnitudes and colors of 118218 local stars. We assume that the local population is representative of the entire Galactic disk stellar population. This hypothesis is certainly justified for the disk stars. The central bar stellar population is redder, but the EROS observations we are considering here do not point toward its center.
- A random magnitude shift is induced to take into account observational limitations, such as blending and uncertainties, from the HIPPARCOS and EROS data.
- The spatial mass density distribution results from the addition of the contributions of thin and thick disks and of the bar modeled according to [26] and [86] or to the Besançon model ([240]).
- The light propagation is affected by Galactic extinction in I and reddening in $V - I$, obtained from a 3D table of K_S extinctions kindly provided by ([78]).

4.4.1 Producing a CMD from the local HIPPARCOS catalog

We present in Appendix A our procedure to obtain a debiased CMD in the (I, V) color system within the domain $0 < M_V < 8$ from the HIPPARCOS catalog. This debiased catalog is described by the distribution $n(\mathbf{M})$, where \mathbf{M} represents the absolute magnitude and color “vector” of a given stellar type. We established in Appendix A that the numerical contribution of stars brighter than $M_V = 0$ is negligible in a deep Galactic image. In our case, given our limiting magnitude, we can also neglect the contribution of stars fainter than $M_V = 8$.

Assuming that the stellar composition is constant along the line of sight, stars of any given type are distributed along the line proportionally to the total mass density ρ . The number of stars expected per square degree ($\Omega(1^\circ \times 1^\circ) = 3.046 \times 10^{-4} sr$) in the EROS catalog is then the integral along the line of sight

$$n_{EROS}(\mathbf{m}) = \int_0^\infty \frac{\rho(D)}{\rho_\odot} n(\mathbf{m} - \delta\mathbf{m} - \mu(D) - \mathbf{A}(D)) \epsilon_{EROS}(\mathbf{m}) \Omega(1^\circ \times 1^\circ) D^2 dD, \quad (4.10)$$

where

- D is the distance to the star along the line of sight,
- $\mu(D)$ the corresponding distance modulus (independent of the color),
- $\mathbf{A}(D)$ is the interstellar extinction vector (one component per filter)
- $\delta\mathbf{m}$ is a random shift of \mathbf{m} that takes into account blending (see Sect. 4.4.3) and uncertainties from HIPPARCOS parallax and EROS photometry; HIPPARCOS stellar absolute I magnitudes are randomly shifted according to a Gaussian distribution of dispersion

$$\epsilon_I = \sqrt{\left[5 \log e \times \frac{\delta\pi}{\pi}\right]^2 + (\delta I)^2}, \quad (4.11)$$

where π and $\delta\pi$ are the HIPPARCOS parallax and associated error, and δI is the estimated EROS photometric uncertainty from expression (4.8). The colors $V - I$ are similarly randomly shifted with the dispersion

$$\epsilon_{V-I} = \sqrt{\delta(V-I)_H^2 + \delta(V-I)^2}, \quad (4.12)$$

where $\delta(V-I)_H$ is the uncertainty on the color from the HIPPARCOS catalog and $\delta(V-I)$ is given by Eq. (4.8).

- $\epsilon_{EROS}(\mathbf{m})$ is the probability to detect a star with apparent magnitudes \mathbf{m} in the EROS catalog (see Fig. 4.4). Here this probability is a function of B_{EROS} only, which is related to the absolute magnitudes and to the distance D as follows:

$$\begin{aligned} B_{EROS} &= V - 0.4(V - I) \\ &= \mu(D) + M_V + A_V(D) - 0.4(M_V + A_V(D) - M_I - A_I(D)). \end{aligned} \quad (4.13)$$

4.4.2 Mass density distributions

In this section, we describe two mass distribution models used to scale the local densities of lenses and sources along the line of sight. We note the different status of the thick disk: it is considered hypothetical within the framework of the first model (so-called simple) since it is a pure hidden matter contribution; on the other hand, it is considered as one of the components within the framework of the second model (Besançon).

Simple tunable Galactic model

In this model, which is slightly modified (updated) from the so-called model1 we used in [224], the mass density of the Galaxy is described with a thin disk and a central bar structure. The disk is modeled by a double exponential density in galactocentric cylindrical coordinates

$$\rho_D(r, z) = \frac{\Sigma}{2H} \exp\left(\frac{-(r - R_\odot)}{R}\right) \exp\left(\frac{-|z|}{H}\right), \quad (4.14)$$

where $\Sigma = 50M_\odot\text{pc}^{-2}$ is the column density of the disk at the solar radial position $R_\odot = 8.3\text{kpc}$ ([51]), $H = 0.325\text{kpc}$ is the height scale, and $R = 3.5\text{kpc}$ is the radial length scale of the disk. The position of the Sun with respect to the symmetry plane of the disk is $z_\odot = 26\text{pc} \pm 3\text{pc}$ ([173]). The bar is described in a Cartesian frame in the plane of the Galaxy, with its origin at the Galactic center with the major axis X tilted by $\Phi = 13^\circ$ ([239]) with respect to the Galactic center-Sun line, *i.e.*,

$$\rho_B = \frac{M_B}{6.57\pi abc} e^{-r^2/2}, \quad r^4 = \left[\left(\frac{X}{a}\right)^2 + \left(\frac{Y}{b}\right)^2 \right]^2 + \frac{Z^4}{c^4}, \quad (4.15)$$

where $M_B = 1.7 \times 10^{10}M_\odot$ is the bar mass, and $a = 1.49\text{kpc}$, $b = 0.58\text{kpc}$, and $c = 0.40\text{kpc}$ are the scale length factors.

There has been some controversy about the bar orientation Φ ; in particular, the EROS collaboration ([136]) published an erroneously high value ($\Phi = 49^\circ \pm 8^\circ$) deduced from the variation

of the mean distance to the red giant stars with the Galactic longitude. This mean distance was confused with the distance to the bar major axis, but this view is only correct for a zero width bar. As a consequence, the value of Φ was strongly overestimated, since as soon as the bar is elliptic, the barycenters of the stars along the line of sight do not coincide with the bar main axis ([167]). Moreover, this difference between the barycenter line and the main axis increases when Φ decreases and when the width of the bar increases. Correcting this wrong view, we checked that the EROS red giant clump distance measurements are in fact compatible with the low values of Φ recently published ([239], [282]), as discussed in the following sections.

The hypothetical thick disk is also considered in our model, and we fit its fractional contribution f_{thick} to the Galactic structure ($f_{thick} = 1$ would correspond to fully baryonic Galactic hidden matter). This disk is modeled as the thin disk (Eq. (4.14)), with $\Sigma_{thick} = 35M_{\odot}\text{pc}^{-2}$, $H_{thick} = 1.0\text{kpc}$, and $R_{thick} = 3.5\text{kpc}$.

The IMF of the stellar population is taken from [55] (Eq. (B.9)). We already mentioned that we expect the microlensing duration to be especially sensitive to the low-mass side of the IMF of the lens population. We therefore define a tunable function for the low-mass side IMF ($m \leq M_{\odot}$), by introducing a parameter m_0 (with value $m_0 = 0.2M_{\odot}$ for the regular Chabrier IMF):

$$\xi(\log m/M_{\odot}) = 0.093 \times \exp\left[\frac{-(\log m/m_0)^2}{2 \times (0.55)^2}\right], \text{ for } m \leq M_{\odot} \quad (4.16)$$

and we fit this parameter to our microlensing duration data in Section 4.6.

We use the following kinematical parameters:

- The radial (axis pointing toward the Galactic center), tangential and perpendicular solar motions with respect to the disk are taken from ([51]),

$$v_{\odot r} = 11.1_{-0.75}^{+0.69}, \quad v_{\odot \theta} = 12.24_{-0.47}^{+0.47}, \quad v_{\odot z} = 7.25_{-0.36}^{+0.37} \quad (\text{km/s}). \quad (4.17)$$

We found that the microlensing duration distribution obtained in our simulation is almost insensitive to the exact values of these parameters.

- The global rotation of the disk is given as a function of the galactocentric distance by

$$V_{rot}(r) = V_{rot,\odot} \times \left[1.00767 \left(\frac{r}{R_{\odot}} \right)^{0.0394} + 0.00712 \right], \quad (4.18)$$

where r is the projected radius (cylindrical coordinates) and $V_{rot,\odot} = 239 \pm 7 \text{ km/s}$ ([51]).

- The peculiar velocity of the (thin or thick) disk stars is described by an anisotropic Gaussian distribution with the following radial, tangential, and perpendicular velocity dispersions ([208] and [207]):

$$\begin{aligned}
\sigma_r^{thin} &= 27.4 \pm 1.1 \text{ km/s} & \sigma_r^{thick} &= 56.1 \pm 3.8 \text{ km/s} \\
\sigma_\theta^{thin} &= 20.8 \pm 1.2 \text{ km/s} & \sigma_\theta^{thick} &= 46.1 \pm 6.7 \text{ km/s} \\
\sigma_z^{thin} &= 16.3 \pm 2.2 \text{ km/s} & \sigma_z^{thick} &= 35.1 \pm 3.4 \text{ km/s}.
\end{aligned} \tag{4.19}$$

We also found that the microlensing duration distribution is insensitive to the exact values of these parameters.

- The velocity distribution of the bar stars is given by the combination of a global rotation ([114], [216])

$$\Omega_{bar} = 39 \text{ km} \pm 3.5 \text{ s}^{-1} \text{ kpc}^{-1} \tag{4.20}$$

with a Gaussian isotropic velocity dispersion distribution characterized by $\sigma_{bar} \sim 110 \text{ km/s}$. We found that the mean duration of microlensing events toward γ Sct, which is the only line of sight crossing the bar, is almost insensitive to Ω_{bar} , mainly because the global rotation velocity is almost tangent to this line of sight.

Besaçon Galactic model

In this model ([240], with updated parameters from [239]), the distribution of the matter in the Galaxy is described by the superposition of eight thin disk structures with different ages, a thick disk component, and a central (old) bar structure made of two components ([239]). We considered the updated model from ([239]) that appears to be specifically adapted to the Galactic plane, and chose the fitted parameters associated with a two ellipsoid bar (Freundenreich (S) plus exponential (E) shapes). All the parameters from this model can be found in the Appendix B, to enable any useful comparison with our simple model.

From the local CMD and mass density to the stellar distribution

The mass densities are then converted into stellar number densities and distributed according to our debiased HIPPARCOS-CMD (Section 4.4.1). The number density of stars scales with the stellar mass density, such that the total number density of stars within $0 < M_V < 8$ equals the total mass density within the corresponding mass interval $[0.65, 2.8]M_\odot$, divided by the mean stellar mass in this interval, as computed from the IMF. We finally take into account the fact that $\sim 2:3$ of

those stars are in binary systems, as discussed in Section B.2. This 2:3 poorly known factor and the exact mass to stellar number ratio can both be absorbed in a global renormalization factor, and our simulated catalog has been tuned to precisely reproduce the local (debiased) observed HIPPARCOS-CMD.

We have now in hand the full description of stellar number densities according to the mass densities and the debiased HIPPARCOS-CMD, which is our initial ingredient to simulate EROS-like CMDs.

4.4.3 Extinction

We now have to consider the absorption model to simulate the effects of distance and reddening of the sources in expressions (4.10) and (4.13).

After generating the position and type of a star, we estimate the extinction due to dust along the line of sight using the table provided by ([78]). This 3D table provides A_K , the extinction in K_S in the $(b, l) = (\pm 10^\circ, \pm 100^\circ)$ domain, up to ~ 15 kpc, with 0.1° angular resolution and 0.1 kpc distance resolution. We use the following relations to transpose the A_K into I and V passbands

$$A_V = 8.55 \times A_K, A_I = 4.70 \times A_K, A_{V-I} = 3.85 \times A_K. \quad (4.21)$$

We compared the extinctions from this table with the 2D table of ([246]) (through extrapolation at infinite distance), which is notoriously imprecise toward the Galactic plane, and with the calculator of ([245])². We found that up to ~ 5 kpc, the extinctions in I from [245] are compatible with the Marshall table, although systematically lower. At larger distances, the estimates depart from each other, and extrapolations at large distance from Marshall table are much larger than estimates from both [246] and [245]. Nevertheless, as discussed in section 4.5.1, we found it necessary to correct the extinctions (A_K) of the Marshall table for systematic and statistical uncertainties, to get synthetic CMDs of $I < 18.4$ stars that correctly match the observed CMDs (compare Fig. 4.7 with Fig. 4.10); indeed, because of the large multiplicative factor relating A_V and A_I to A_K , a small error on A_K has a very significant impact on the apparent position of a star in our CMD. Fig. 4.12 shows the average extinctions in V along the lines of sights as a function of the distance to the source, after tuning the model parameters according to our fitting procedure.

²https://ned.ipac.caltech.edu/help/extinction_law_calc.html




Figure 4.7: The $V-I$ observed (gray histograms) and the simulated distributions (simple model in black, Besançon model in red) for the bright stars (with $I < 18.4$), using the Marshall table without systematic/statistical uncertainties (to be compared with Fig. 4.10, bottom).

Blending

We know from the comparison of the EROS images with the HST images ([224]) that $\sim 60\%$ of the $I < 16$ objects and $\sim 70\%$ of the $I > 16$ objects detected by EROS are blends. This blending effect is different than the binary blend mentioned at the end of Appendix A. This effect, due to the EROS low separation power, is accounted for by randomly decreasing the magnitudes of 60% of the faint stars (resp. 70% of the bright) according to a Gaussian distribution centered on -0.07 , with $\sigma = 0.25$ (resp. 0.13), truncated at zero.

In principle, blending also contributes to reduce the number of detected objects with respect to the predictions based on the HIPPARCOS catalog. As for the binary blend, this effect can be absorbed in a global renormalization factor.

4.5 Comparing the EROS observations with simulated populations and microlensing expectations

Our aim is now to tune and compare the Galactic models with the observations toward the four Galactic disk lines of sight (characterized by the corresponding EROS catalogs noted C). We use all the available observables for this purpose as follows:

- The four color-magnitude distributions (CMD) of stars brighter than $I = 18.4$, which is the most reliable part of the EROS-CMD. The observable variables we consider are derived from the projected magnitude and color distributions: the total stellar densities ρ_* ³, and the first moments $\overline{V-I}$ and σ_{V-I} of the $V-I$ distribution⁴.
- The measured optical depths $\tau(C)$ ([224]) toward the four catalogs C (Table 4.1).
- The measured means $\overline{t_E}(C)$ ([224]) (Table 4.1). The poor available statistics convinced us not to use the σ_{t_E} parameter in our fitting procedure, since it is affected by such a large uncertainty that it is essentially not constraining.

For quantitative statistical comparisons based on χ^2 studies, we need good control of the uncertainties on these observables. The τ and $\overline{t_E}$ uncertainties are provided in [224]. Table 4.1 summarizes the numerical data toward the EROS monitored populations that we use for the comparison with a simulation (apart from σ_{t_E}).

³After noting that the slopes of the magnitude distributions seem universal, we concluded that the integrated stellar number density ρ_* carries all the information on this distribution.

⁴These variables have the advantage that they do not depend on an arbitrary binning.

4.5.1 Simulation of the CMDs

All the relevant information is already given in Sect. 4.4. Here, we briefly summarize the different stages to simulate EROS CMDs from various models or parameters.

The stellar absolute magnitudes and colors are first randomly chosen according to the HIPPARCOS unbiased color-magnitude density diagram of stars with $0 < M_V < 8$ (Fig. B.4 bottom). Generated magnitudes are then shifted to take into account the blending described in Sect. 4.4.3, as well as the HIPPARCOS parallax uncertainties and EROS photometric uncertainties (Eqs. (4.11) and (4.12)).

To estimate the integral in expression (4.10), we generate the distance distributions of stars according to the mass density distributions of each Galactic structure (bar, thin disk, or thick disk). The EROS stellar apparent magnitudes and colors are estimated from the absolute magnitudes, the distances, and take into account the absorptions tabulated (in K_S) at the position randomly chosen within the EROS fields (see Sect. 4.4.3). After this stage, we obtain the apparent color-magnitude distribution of the stars before detection. Finally, the contribution of each generated star is weighted by the EROS stellar detection efficiency $\epsilon_{EROS}(\mathbf{m})$, which is parametrized as a function of B_{EROS} (figure 4.4).

As mentioned in Sect. 4.4.3, to successfully fit the CMDs we had to introduce the hypothesis of a systematic uncertainty on K_S changing with the catalog, $\Delta A_K(C)$, and a random uncertainty with constant width ϵ_{A_K} , within the tabulated data. Since the table does not provide uncertainties, we used this hypothesis as the simplest way to make our simulation compatible with the observations (A.C. Robin, priv. comm.). Then the $\Delta A_K(C)$ and ϵ_{A_K} parameters were tuned together with the Galactic parameters to obtain synthetic CMDs that fit the observed CMDs (see below).

4.5.2 Simulation of microlensing

The previous procedure, based on the synthesis of the color-magnitude diagrams, allows us to simulate the EROS catalogs of sources. To simulate the microlensing process for these catalogs, we also need to synthesize the population of lenses, containing all massive objects regardless of their visibility. The local lens density population is therefore simulated with the appropriate IMF (depending on the Galactic structure and on the model) scaled with the local mass density. The transverse velocity distribution needed to simulate the microlensing event durations is obtained from the combination of the velocity distributions from the disk(s) and the bar, according to their respective local mass contributions. Finally, we take into account the impact of the time sampling by simulating the microlensing detection efficiency according to Fig. 4.6.

4.6 Fitting procedure

Our simulation program allows us to produce the CMDs and microlensing distributions toward our 4 catalogs labeled (C), with any choice of Galactic parameters. We detail below the procedure developed for our simple tunable model, which we also used to probe the Besançon model (with no tuned parameter other than the systematic uncertainties of the interstellar absorptions).

4.6.1 Fit and tuning of the simple model

We examined the following 16 observables (4 per target C) $\rho_*(C)$, $\overline{V-I}(C)$, $\tau(C)$, and $\overline{t_E}(C)$ as a function of the following parameters, around their nominal values: ϵ_{A_K} , the random uncertainty on the extinctions A_K provided by the table from [177] for each generated stellar position; $\Delta A_K(C)$, the systematic uncertainty on $A_K(C)$, depending on the catalog(C); the Galactic bar orientation Φ (nominal value $\Phi = 13^\circ$); and the (hypothetical) thick disk contribution, which is parametrized by the fraction f_{thick} of the thick disk considered in ([224]). This contribution is modeled like the thin disk (see Eq. (4.14)), with $\Sigma_{thick} = f_{thick} \times 35 \text{M}_\odot \text{pc}^{-2}$, $H_{thick} = 1.0 \text{kpc}$, $R_{thick} = 3.5 \text{kpc}$, and velocity dispersions given by Eq. (4.19).

To benefit from the exclusive time information $\overline{t_E}(C)$ provided by the microlensing data, we also considered some specific parameters that are expected to impact the microlensing optical durations. First, the low-mass part of the IMF, which we generalized from [55] through parameter m_0 (nominal value $m_0 = 0.2$) (Sect. 4.4.2). Second, we explored the sensitivity to the peculiar velocities of the microlensing actors through a scaling of the velocity dispersions reported in expression (4.19). We found that our simulation is insensitive to such a scaling, therefore confirming that orbital velocities dominate the relative transverse motions. Third, for completeness, we also tested the sensitivity of $\overline{t_E}$ with the global rotation of the bar (Eq. 4.20) and found almost no sensitivity; this is mainly because the bar rotation is almost tangent to the line of sight of γ Sct, which is the only line of sight that crosses the bar structure.

Sensitivity of the observables with respect to the Galactic parameters

We used our simulation to establish the sensitivity of the observables with the variations of the different parameters, and we made the following observations.

We find that only the simulated observables from the low longitude fields (β Sct and γ Sct) are sensitive to the variations of Φ , when we test for very large changes, but they are insensitive to few degree variations around the nominal value $\Phi = 13^\circ$. As a consequence, we exclude Φ from our fit.

At first order, the absorption random shift dispersion ϵ_{A_K} , with respect to the tabulated values, is assumed to be the same for all fields, and the widths of the four color distributions $\sigma_{V-I}(C)$ are found to be disconnected from the other observables and parameters. We therefore directly fit ϵ_{A_K} by minimizing the differences between $(\sigma_{V-I}^{obs}(C))^2$ and the width combination $(\sigma_{V-I}^{sim.(0)}(C))^2 + (3.85 * \epsilon_{A_K})^2$, where the 3.85 factor comes from the relation $A_{V-I} = 3.85 \times A_K$ (see Sect. 4.4.3), and the $\sigma_{V-I}^{sim.(0)}(C)$ values are obtained with a simulation that assumes $\epsilon_{A_K} = 0$. The value that minimizes the sum on (C) is $\epsilon_{A_K} = 0.085$, which we assume to be independent of the catalog C . We use this value in the subsequent simulations.

The Chabrier-like IMF parameter m_0 and the observables $\overline{t_E}(C)$ are also disconnected from the other observables and parameters. We therefore make a separate (sub-)fit for these parameters, by minimizing

$$\chi_{t_E}^2 = \sum_C^{catalogs} \frac{(\overline{t_E}^{sim}(C) - \overline{t_E}^{obs}(C))^2}{\sigma_{t_E}^2(C)} \quad (4.22)$$

with respect to m_0 , where the suffixes *sim* and *obs* refer to the simulated and observed catalogs.

The observables $\rho_*(C)$, $\overline{V-I}(C)$, and the microlensing optical depths $\tau(C)$ (12 observables) depend only on f_{thick} and on the systematics $\Delta A_K(C)$ (5 parameters). We performed a combined fit by minimizing the sum of $\chi_{\rho_*}^2$, $\chi_{\overline{V-I}}^2$ and χ_{τ}^2 , which is defined similar to $\chi_{t_E}^2$, but since we have to take into account common systematics, some of the covariant matrices are not diagonal.

In our minimization procedure, we used the first order development of the observables as functions of the parameters to be fitted, from the derivatives computed with our simulation. This allowed us to perform the fit with acceptable computing time, considering the very long runs needed for each model configuration.

Systematic and statistical uncertainties

We have carefully established the budget error for each observable as follows.

For the $\rho_*(C)$ budget error, we have to take into account the uncertainty of $\sim 5.3\%$ on the size of the effective EROS field and the consequences of the 0.07 magnitude EROS calibration uncertainty. The impact of this calibration uncertainty on $\rho_*(C)$ has been estimated from the published EROS-CMD tables, by changing the position of the $I < 18.4$ magnitude cut by the 0.07 systematics. We found that the uncertainty on $\rho_*(C)$ due to this calibration error is $\sim 5\%$. The final systematics results from the quadratic addition of both uncertainties (7.3%) and since it is a multiplicative systematics, it has to be considered as an uncertainty on a global normalization α ; we therefore use a standard procedure to include the extra parameter α and fit the product

$\alpha \times \rho_*^{sim}$ with ρ_*^{obs} . We adopt 15% as the statistical uncertainty on $\rho_*(C)$, which is dominated by residual uncertainties from the absorption model and blending effects.

For $\overline{V-I}(C)$, we have to account for the systematics due to calibration uncertainties on both R_{EROS} and B_{EROS} , thus giving a global systematics of 0.16 mag.. In the covariance matrix associated with the fit minimization, this additive systematics, which is common to the four directions, contributes as a full matrix, to be added to the usual diagonal matrix built from the residual statistical uncertainty that is estimated to be 0.15 mag..

Statistical uncertainties from the EROS-CMD Poissonian fluctuation propagation are estimated as explained in the header of the published EROS-CMD ([225]). Considering the large statistics available in the EROS database, we can neglect the uncertainties due to the Poissonian fluctuations of the number of stars in the original EROS histogram used to produce the CMDs.

As a conclusion, the uncertainties on $\rho_*(C)$, $\overline{V-I}(C)$ and $\sigma_{V-I}(C)$ are dominated by the impact of the calibration uncertainties and the residual uncertainties from blending and absorption effects discussed above. The values used for the fit are summarized in Table 4.6.1.

Results from the fit

We note that the fit is done with the best value for the random uncertainty on the tabulated absorptions A_K : $\epsilon_{A_K} = 0.085$. The best fit is obtained with the following parameters.

First, regarding absorption systematics, we find $\Delta A_K(\beta S ct) = 0.09$ mag, $\Delta A_K(\gamma S ct) = 0.04$ mag, $\Delta A_K(\gamma Nor) = 0.11$ mag, and $\Delta A_K(\theta Mus) = -0.01$ mag.

Second, regarding the fraction of the thick disk, we find $f_{thick} = 0.05 \pm 0.6$. This result does not differ from zero, showing that there is no need for an additional baryonic contribution to the thin disk within the framework of our simple model. We also tested the option of a non-luminous thick disk (made of compact unseen objects), assuming no contribution to the CMD (therefore only impacting the optical depths); we found $f_{invisiblethick} = 0.5 \pm 0.9$, which is again not significantly different than zero. From this estimate, we can conclude that the total mass of an invisible thick disk is smaller than $7 \times 10^{10} M_\odot$ at 95%CL.

Third, regarding the IMF, we find $m_0 = 0.51 \pm 0.25 M_\odot$, which is 1.24 sigmas away from the 0.2 nominal value of the local Chabrier IMF ([55]). Our observations are therefore significantly sensitive to the low-mass side of the lens IMF. This sensitivity belongs to a non-local IMF, since it concerns only the lenses and not the solar neighborhood. Fig. 4.8 shows both IMFs (the local and best fitted lens-IMF).

For this global fit of the CMDs, optical depths and microlensing durations, we find $\chi^2 = 6.5$

for 10 degrees of freedom with a fair repartition between the different types of observables (ρ_* , $\overline{V-I}$, τ and t_E).

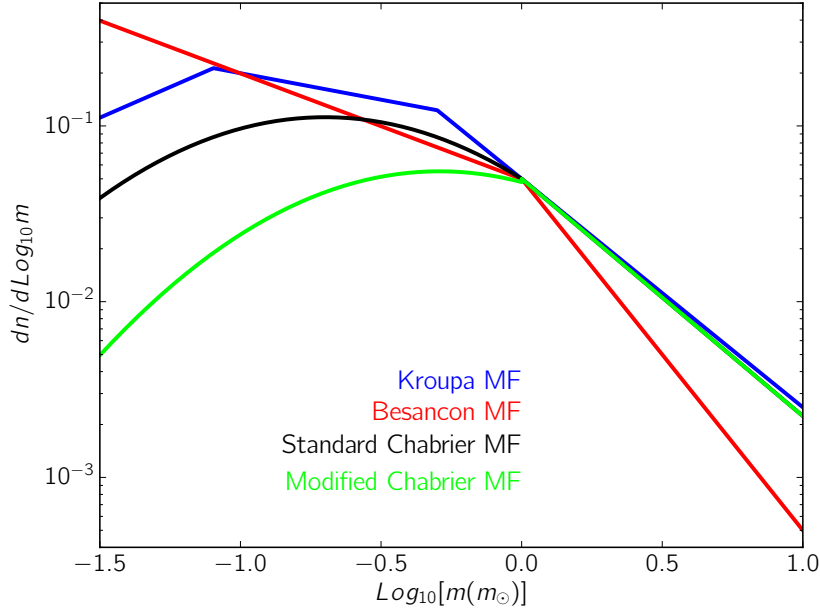


Figure 4.8: Different mass functions considered in this chapter: Standard Chabrier (black) corresponds to the local regular Chabrier IMF (Eq. (B.9) with $m_0 = 0.2M_\odot$); the modified Chabrier ($m_0 = 0.51M_\odot$, in green) gives the best fit for the lens IMF from our simple model.

We exchanged in our simple model the Chabrier IMF for the Kroupa IMF ([163]). The only consequence to this exchange was a significant decrease in the $\overline{t_E}$ values, as expected from the larger contribution of low-mass objects (see Fig. 4.8 and Table 4.6.1). This degrades the fit by $\Delta\chi^2 = 7.4$, showing that the Kroupa IMF is strongly disfavored by our data.

It is clear that a larger statistics of microlensing events toward the spiral arms would have the capability to better constrain the thick disk component and the lens-IMF.

Table 4.6.1 summarizes the best fit results for our simple model compared with previous simulations (model 1) considered in [224], differing mainly through the extinction description. Fig. 4.9 shows the mass density along the line of sight of γSct resulting from our simple fitted model.




Figure 4.9: Mass-density along the line of sight of γSct from the various Galactic structures (disks and bar), as a function of the distance from the Sun for our nominal simple model (thin black lines) and the Besançon model (thick blue lines). The total densities are shown with dashed lines.

	Target	θ Mus	γ Nor	γ Sct	β Sct
$\rho_*^{I < 18.4} \times 10^6$	measured	$0.25 \pm .037$	$0.23 \pm .035$	$0.28 \pm .042$	$0.34 \pm .051$
		$\pm 7.3\%$ common systematics			
	simple model	0.22	0.26	0.28	0.32
	Besaçon	0.23	0.26	0.30	0.33
$\overline{V-I}$	measured	$1.95 \pm .15$	$1.86 \pm .15$	$2.36 \pm .15$	$2.20 \pm .15$
		± 0.16 common systematics			
	simple model	1.83	2.02	2.35	2.13
	Besaçon	1.94	2.11	2.52	2.22
σ_{V-I}	measured	0.71	0.78	0.71	0.75
	simple model	0.72	0.73	0.83	0.74
	Besaçon	0.73	0.74	0.81	0.73
$N_{event}(u_0 < .7)$	observed	3	10	6	3
$\overline{N}_{event}(u_0 < .7)$	model 1	2.8	9.9	7.1	6.3
	simple model	4.0	8.6	3.6	2.2
	Besaçon	4.0	9.9	3.5	2.4
$\tau \times 10^6$	measured	$.67^{+.63}_{-.52}$	$.49^{+.21}_{-.18}$	$.72^{+.41}_{-.28}$	$.30^{+.23}_{-.20}$
	model 1	0.42	0.52	0.71	0.57
	simple model	0.23	0.38	0.43	0.45
	Besaçon	0.22	0.34	0.44	0.40
\overline{t}_E (days)	measured	97 ± 75	57 ± 10	47 ± 6	59 ± 9
	model 1	73.8	67.9	37.9	60.2
	simple model	79.4	54.4	49.1	53.8
	with Kroupa IMF	64	43	38	42
	Besaçon	68.5	51.9	43.0	49.3

Table 4.2: Best fit results on the observables toward the 4 regions monitored in the EROS spiral arms program, compared with previous simulations (model 1) and observations published in [224]: Surface density (per square degree) of stars brighter than $I = 18.4$, mean and width of CMD color distribution, number of microlensing events, optical depth, and mean duration.

4.6.2 Besançon model: tuning the extinctions

In this section, our purpose is to test the agreement of the Besançon model with the EROS microlensing results. We used almost the same procedure as above, but fitting only the uncertainties on the K extinctions. The best fit is obtained for $\epsilon_{A_K} = 0.10$, $\Delta A_K(\beta Sct) = 0.14 \text{ mag}$, $\Delta A_K(\gamma Sct) = 0.13 \text{ mag}$, $\Delta A_K(\gamma Nor) = 0.15 \text{ mag}$, and $\Delta A_K(\theta Mus) = 0.04 \text{ mag}$. The global fit has a $\chi^2 = 8.2$ for 12 d.o.f, with specific contributions of $\chi_{\rho_*}^2 = 1.2$, $\chi_{V-I}^2 = 2.2$, $\chi_{\tau}^2 = 2.8$, and $\chi_{IE}^2 = 2.0$.

Not surprisingly, the values of $\chi_{\tau}^2 = 2.8$ and χ_{IE}^2 are worse than those of our simple model, since no parameters are fitted for the thick disk and the IMF, but the fit is globally satisfying (see Table 4.6.1 for the summary of the fitted parameters and observables). Fig. 4.9 shows the mass density along the line of sight of γSct from the Galactic structures of the Besançon model (in blue), resulting from the best fitted extinction.

As for the previous simple model, we also tested the hypothesis of an invisible extra contribution to the thick disk for this model; we find that the best fitted value for such a thick disk favors an added contribution of 2.5 ± 4.7 times the modeled thick disk ($\chi^2 = 8.0 \text{ per } 11 \text{ d.o.f}$). Again, there is no significant indication of the need for such an invisible contribution and the upper limit of a Besançon-like thick disk (somewhat thinner than in our simple model) is $\sim 5 \times 10^{10} M_{\odot}$ at 95%CL.

4.7 Discussion

As a preliminary to the discussion, we recall here some of the hypotheses used throughout this chapter: First, we assume the disk to have the same CMD as around the sun; then we rely on the extrapolation of the extinction map obtained in K band to I and V bands, and assume reasonable systematic uncertainties on this map.

4.7.1 Comparison with previous results and robustness

Figure 4.10 (to be compared with Fig. 4.3) shows that our best fitted models are able to reproduce satisfactorily the observed CMDs of the ($I < 18.4$) stars. Table 4.6.1 shows that the model we used previously (model 1) was also satisfactory. We tested the robustness of our results by changing some of the uncertainties (systematics and statistics) with insignificant variations of the best fitted numbers.

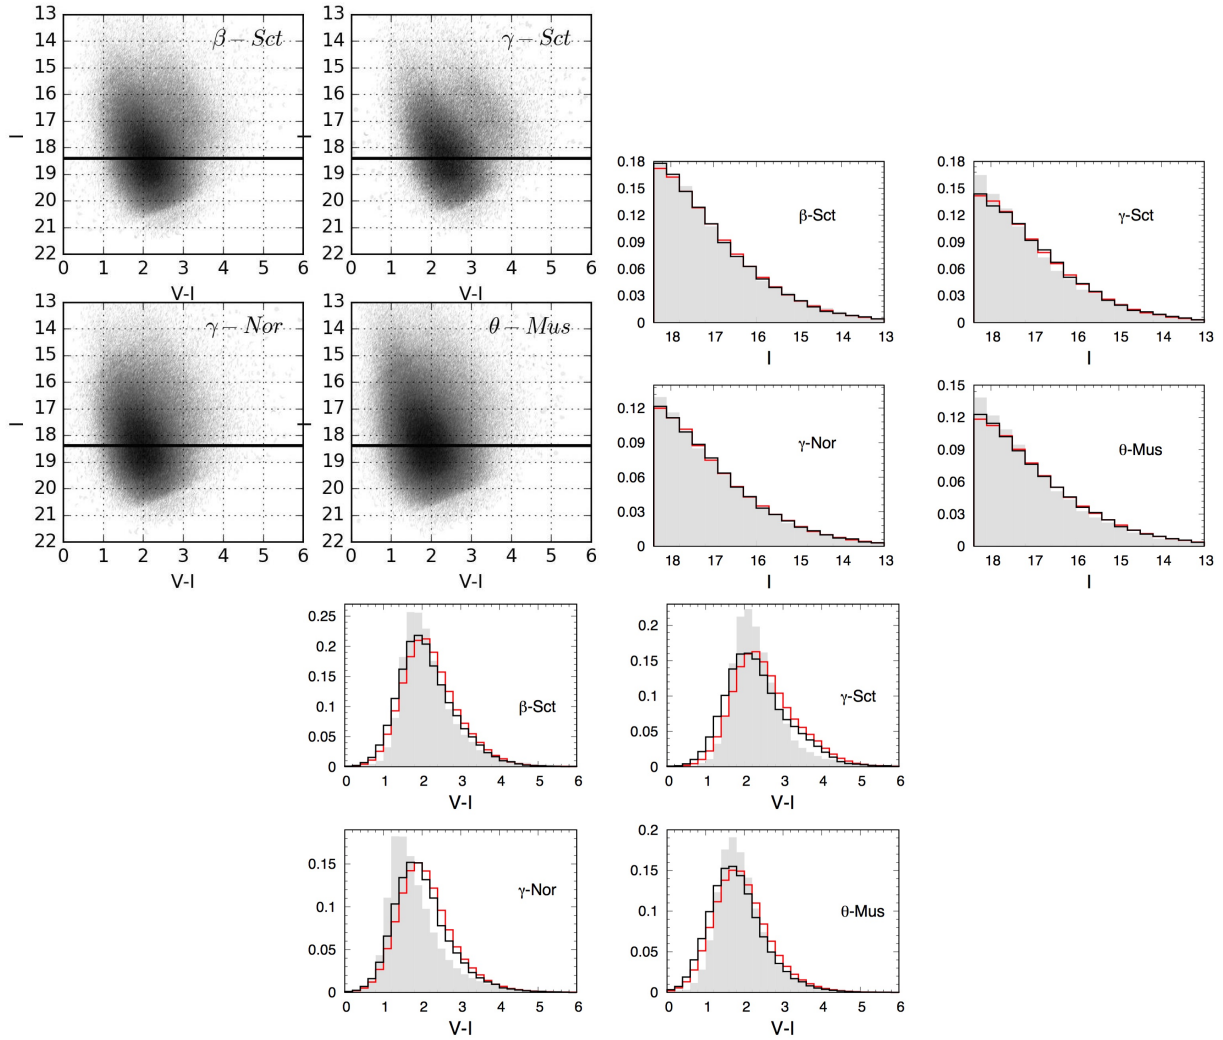


Figure 4.10: Simulated CMDs toward the 4 monitored directions (top) with the magnitude (middle) and color (bottom) projections for the stars brighter than $I = 18.4$, expressed in million of stars per square degree per magnitude. Results from our simple model are plotted with black lines and results from the Besançon model with red lines; the distributions of the EROS observed populations of bright stars ($I < 18.4$) are superimposed on the projections as light gray histograms.

Our model now incorporates enough details to allow one to use the CMD as an observable to be fitted. As a consequence, the main impact of this type of study, apart from constraining the parameters f_{thick} (for our simple model) and m_0 , is to extract information on the underlying stellar populations of sources and lenses.

4.7.2 Lens and source populations

Figure 4.11 shows the fast variation of the simulated optical depth along the line of sight with the distance for the four studied directions and for both models considered in this chapter. This




Figure 4.11: Simulated optical depths toward the 4 monitored directions, as a function of the source distance, for our nominal simple model (thin lines) and the Besançon model (thick lines).

fast variation of the optical depth with the distance shows that the notion of catalog optical depth is crucial when dealing with sources distributed along a line of sight. This notion is not relevant when considering well-defined distance targets such as LMC, SMC, and M31; when considering only bright sources toward the Galactic Center, it is estimated that the relative uncertainty on the bright sources positions is less than 10% ([206]) and it is still possible to ignore the spread of the sources and to use the classical concept of optical depth up to a given distance for the whole catalog. Previous studies concerning the Galactic spiral arms ([75], [74]) performed a simplified analysis, by assuming all sources to be at $7Kpc$ to compare the observed optical depth with simple models, but [224] started to draw attention to the impact of the source distance spread. Now it is clear that precise studies in the Galactic plane are needed to know the distance distribution of the monitored catalog. Figure 4.12 shows the expected distance distributions of the lenses and sources in the EROS microlensing events obtained from our simulation (taking into

account the EROS efficiencies). Again, the source distance distribution illustrates the relevance of the concept of optical depth toward a population in contrast with the optical depth up to a given distance.

4.7.3 Constraining the Galactic model: The specific contribution of microlensing data

The good agreement of our Galactic models with the data shows that there is no need for other or more ingredients. The Besançon model predicts relatively small optical depths, and this observation is in agreement with the deficit of optical depth toward the inner bulge directions noticed by MOA-II ([16]), even if this is not very significant from our reduced statistics.

We also used our simulation to measure the domain of Galactic parameters that is compatible with our observations. We focused on parameters that are expected to impact the microlensing optical depths or durations, *i.e.*, the bar inclination Φ (nominal value $\Phi = 13^\circ$); the thick disk contribution, parametrized by the fraction f_{thick} , either visible (for the simple model) or invisible (for both models); the disk kinematics for which we explored our sensitivity through the scaling of the velocity dispersions (in expression (4.19)); and the IMF parameter m_0 , as defined in Sect. 4.6.

The impact of the Galactic bar is illustrated in Fig. 4.12, where it is clearly visible that it mainly intercepts the γ Sct line of sight; in the present case, owing to low statistics, our data can only distinguish between a small or a large bar angle, but cannot refine its current estimate. Nevertheless, it is clear that systematic microlensing study at relatively small Galactic longitude is a promising technique to precisely measure the bar inclination.

We show that there is no significant need for an extra thick disk component (visible or invisible); otherwise, from our data alone there are not enough constraints to exclude its existence and only a 95%CL upper limit on its total mass could be inferred ($\sim 7 \times 10^{10} M_\odot$ for the simple-model thick disk, $\sim 5 \times 10^{10} M_\odot$ for an extra invisible component in the Besançon-model thick disk).

We found that we cannot constrain the velocity dispersion ellipsoids of the microlensing actors, since the transverse velocities involved in the microlensing durations are dominated by the orbital velocities.

Interestingly, we show that microlensing durations can constrain the low-mass end of the mass function (see Fig. 4.13), and more importantly, they can provide such constraints for non-local stellar populations (the disk lens population); this is in contrast with the other techniques, which can only measure the mass function around the Sun.

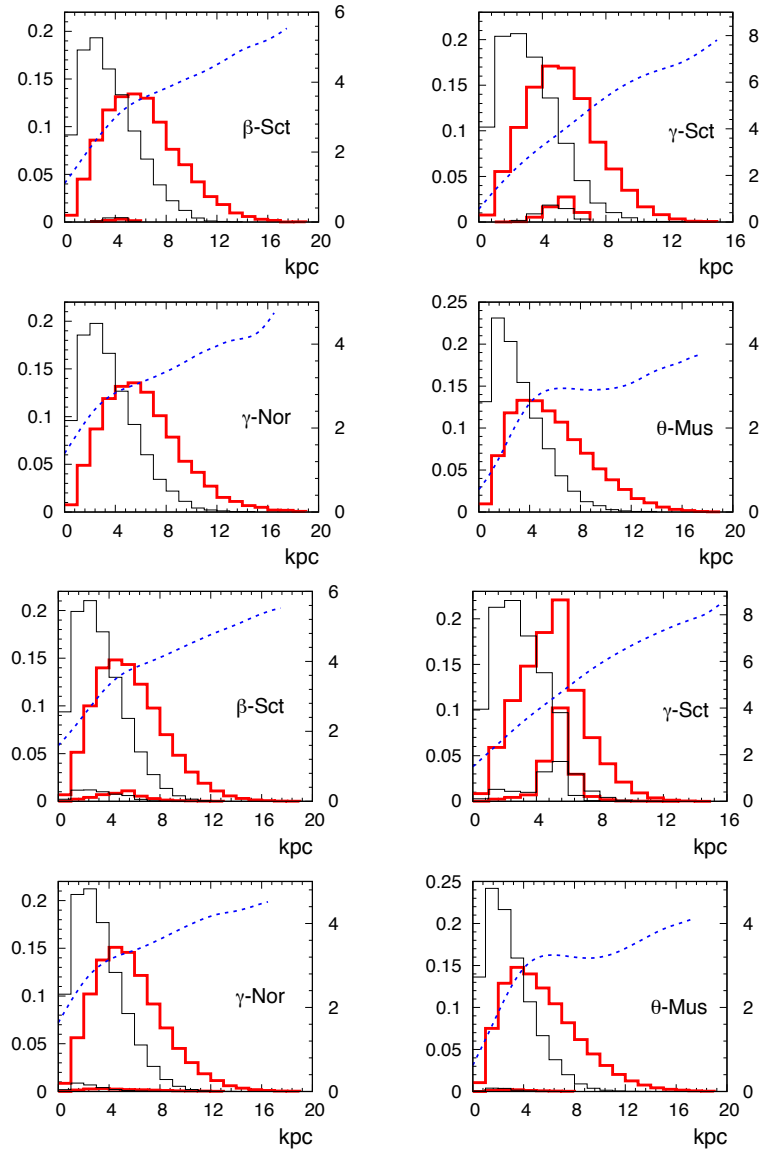


Figure 4.12: Expected normalized distributions of the distances for the lensed sources —when taking into account the EROS microlensing detection efficiencies— (thin lines) and of the lenses (thick lines) from the simulation of our simple model (upper) and the Besançon model (lower). The sparsely-populated distributions around 4 kpc (for β Sct and γ Sct) correspond to the contribution of the bar objects. The dashed curves show, as a function of the distance, the average extinctions of the stars in the simulated EROS-like catalog (in V magnitude, on the right scale). It is strongly biased in favor of small extinctions mainly due to the magnitude selection $I < 18.4$.




Figure 4.13: Einstein duration t_E distribution of the microlensing events expected by assuming 4 different IMFs: the standard Chabrier (black), the Besançon model (red), the modified Chabrier (with $m_0 = 0.57$, green), and the Kroupa IMF (blue).

The best fitted value we obtain for our parametrized Chabrier-type IMF of the lens population of the disk is $m_0 = 0.51 \pm 0.25$, which is in relative disagreement (by one standard deviation) with the parameter of the local mass function ($m_0 = 0.2$) of the Chabrier model. This discrepancy originates in the longer mean durations of the observed events compared with the simulation based on the local IMF. The microlensing technique seems to be significantly sensitive to the IMF low-mass end.

We find that the Kroupa IMF does not correctly reproduce the mean durations of our microlensing events, because of the higher contribution of low-mass objects, inducing a deficit of predicted long duration events.

4.7.4 Limitations of this study

We have made a considerable effort to understand the CMDs and the microlensing data toward directions that have not been examined by other teams. For this reason, we note the limits we encountered during this study to avoid any misinterpretation. Knowledge of the absorption map was one of the most important limitations. Its precision and resolution within the studied fields are parameters that impact the CMD so strongly that we found it necessary to assume (reasonable) systematic and statistical dispersions to understand the observed densities of bright stars.

The blending and the 2:3 estimated fraction of binary stars are also other sources of limitation for understanding the CMDs. All of these elements have fortunately a somewhat degenerated impact on the predicted stellar densities; without any correction to the extinctions, we found that the simulated CMDs had too many stars and were bluer than the data, which could be solved with a systematic extinction increase. These limitations impacts mainly the CMDs; the specific observables from microlensing (optical depth and durations) are mainly impacted through the distance distribution of the lenses.

4.8 Conclusions and perspectives

We have performed a complete simulation of the Galactic structure and the EROS acceptance, which is able to reproduce all the EROS exclusive observations toward the Galactic arms. In this view, we produced a debiased color-magnitude diagram from the HIPPARCOS catalog to feed our simulation with a realistic stellar population. This population was spatially distributed according to the Besançon Galactic model, and to a simple Galactic mass model including a thin disk and a central bar, with an adjustable thick disk contribution and IMF. Every simulated object was then considered as a potential gravitational lens as well as a potential source to gravitational lensing. Taking into account the dust extinction and EROS detection efficiencies, the observed color-magnitude diagrams and the microlensing optical depths and durations are correctly fitted with both our simple Galactic model (with no thick disk) and the Besançon model. We then used the simulation as a tool to obtain information on the configuration space of the microlensing actors (lens and source distance distributions). The large width found in this way for the source distance distribution validates the concept of “catalog optical depth” by contrast with the usual optical depth to a given distance. This concept is to be used as soon as the sources are widely distributed in distance. Finally, even with the small statistics of microlensing events, we were able to extract interesting constraints on the Galactic parameters – *i.e.*, bar inclination confirmation, disk kinematics, mass function, and hidden matter– that have an impact on the microlensing distributions.

The running VISTA Variables in the Via Lactea (VVV) survey [185], which is monitoring stars within the Galactic plane in infrared, is well suited to enlarge the field of view within the Galactic plane, by searching for microlensing in dusty regions. This survey should be able to better constrain the parameters mentioned above, with promising perspectives such as measuring the mass function in areas other than the solar neighborhood. The Large Synoptic Survey Telescope (LSST) will also have the capability to monitor a wide domain of the Galactic plane for microlensing, but only limited to the clear windows, free from large dust column densities.

Chapter 5

Themis: A Parameter estimation framework for the Event Horizon Telescope

The Event Horizon Telescope (EHT) provides the unprecedented ability to directly resolve the structure and dynamics of black hole emission regions on scales smaller than their horizons. This has the potential to probe critically the mechanisms by which black holes accrete and launch outflows, and the structure of supermassive black hole spacetimes. However, accessing this information is a formidable analysis challenge for two reasons. First, the EHT natively produces a variety of data types that encode information about the image structure in non-trivial ways; these are subject to a variety of systematic effects associated with very-long baseline interferometry, and are supplemented by a wide variety of auxiliary data on the primary EHT targets from decades of other observations. Second, models of the emission regions and their interaction with the black hole are complex, highly uncertain, and computationally expensive to construct. As a result, the scientific utilization of EHT observations requires a flexible, extensible, and powerful analysis framework. We present such a framework, `THEMIS`, which defines a set of interfaces between models, data, and sampling algorithms that facilitates future development. We describe the design and currently existing components of `THEMIS`, how `THEMIS` has been validated thus far, and present additional analyses made possible by `THEMIS` that illustrate its capabilities. Importantly, we demonstrate that `THEMIS` is able to reproduce prior EHT analyses, extend these, and do so in a computationally efficient manner that can efficiently exploit modern high-performance computing facilities. We expect that `THEMIS` will play an important role in extracting scientific conclusions from recent and future EHT observations.

5.1 Introduction

The Event Horizon Telescope (EHT), a global array of millimeter and sub-millimeter radio telescopes, has resolved the horizons of at least two black holes [79–82]. This provides a unique window on the high-energy astrophysical processes responsible for the substantial growth and inordinate influence of supermassive black holes [98, 139], the dynamics and thermodynamics of material in the strong gravity regime [197, 295], and the fundamental nature of black holes [42, 220]. However, efficiently and accurately extracting this information from the observational data presents numerous challenges, requiring the development of novel analysis tools tailored to the EHT data products, EHT-target properties, and auxiliary information.

The EHT achieves an extraordinary resolution of $13 \mu\text{as}$, making it the highest resolution imaging instrument in the history of astronomy. It does this via very long baseline interferometry (VLBI), in which information from pairs of individual stations separated by Earth-sized distances are combined to measure small-scale structure on the sky. The resulting data takes the form of complex visibilities, directly related to the Fourier transform of the image. This can be performed in all four Stokes parameters, yielding complete information about the resolved polarization structures (e.g., [153]). In the near future, this will be extended to multiple wavelengths (1.3 mm and 0.87 mm) [99]. Millimeter-VLBI observations of the primary EHT targets have already been carried out at multiple epochs, covering times ranging from 10 seconds to 10 years [82, 104, 105, 153].

Difficulties in the phase calibration, and lesser — though still significant — complications in the amplitude calibration of these visibilities, has motivated the construction of a set of VLBI observables (e.g., visibility amplitudes, closure phases [148], closure amplitudes [274, 275], visibility polarization fractions [153], etc.) that probe the underlying image structure in nonintuitive ways. These have traditionally been interpreted within the context of a simple set of phenomenological models, e.g., multi-component Gaussians. However, the substantial structure anticipated on horizon scales exhibited by the primary EHT targets has given rise to a broader modeling effort, which includes a variety of physical processes [40, 41, 45, 56, 57, 59, 77, 84, 101, 126, 193, 195, 197, 253, 294, 295].

This modeling effort is further motivated by the large amount of ancillary data that exists for EHT targets. All EHT targets are necessarily bright radio sources, and thus have been the object of substantial astronomical scrutiny. Both the Galactic center (Sgr A*) and M87 have been studied across the electromagnetic spectrum, from decameter wavelengths [72] to very-high energy gamma rays ($> 1 \text{ TeV}$) [88]. Moreover, due to the close proximity to their central black holes, both are empirically highly variable, providing statistical information about the dynamics within the mm-wavelength emission regions and creating opportunities to probe these dynamics

directly at multiple wavelengths [89]. Physical modeling of these sources provides the unique ability to synthesize all of these observations, which when combined with EHT data, can provide a detailed description of the conditions and dynamics of material near black hole horizons.

There are substantial challenges to such a broad modeling effort. First and foremost, models of the near horizon region are necessarily complicated, invoking multiple emission components (non-thermal and thermal emission regions with uncertain and potentially distinct locations [56, 60, 70, 202, 238]), a variety of dynamical processes (orbital motion, winds, jets, explosive events, etc., [45, 84, 149, 181, 221]), strong lensing in a potentially uncertain spacetime (Kerr or beyond, [42, 150, 151, 187]), polarization transfer effects (e.g., Faraday rotation and conversion [145, 252]), and propagation effects (e.g., interstellar scattering [154]). To add to the complexity of this comparison only Fourier modes along specific tracks in the two-dimensional Fourier domain are probed via Earth-Aperture synthesis on time scales that can be comparable to intrinsic source variability [170]. Thus, any tools constructed to make comparisons between physical models of EHT targets and the collection of EHT and auxiliary data must be extremely flexible.

Second, there are clear emission-model independent features in many images that arise from the structure of the underlying spacetime. These include the black hole shadow, the image of the event horizon at infinity, first described by [20]. This is bounded by the photon ring, a bright ring arising from the stacking of multiple images, in which the gross features of the spacetime are encoded. Thus, there is substantial motivation to directly extract these generic features from the EHT data alone. Again, this is complicated by the indirect relationship between the VLBI observables and the image, resulting in frequently counter-intuitive conclusions. Hence, ideally, any tools for assessing the presence and properties of image structures should be able to extend to phenomenological models as well.

Third, the nature of EHT data has evolved rapidly over the past decade, growing as the sensitivity and baseline coverage improved. It is far from clear that any particular set of EHT data types are optimal for a given astrophysical or gravitational question. In some cases new data types have been developed based on both instrumental and observational limitations (e.g., visibility polarization fractions). Given the broad range of EHT and ancillary data types, any model comparison effort must maintain substantial flexibility in the kinds of information that it can utilize.

Finally, in many cases the construction of physically realistic models is computationally expensive, requiring ray tracing (relatively cheap) and radiative transfer (often expensive) through model structures. This difficulty is compounded by the often multimodal nature of the reconstructed posterior parameter distributions (see, e.g., [41]). As a result, any analysis tools must be both computationally efficient and be able to exploit the large investment in high-performance computing resources.

It is to address these challenges that we have begun the development of an analysis framework for EHT and ancillary data: THEMIS (means "order" and is the name of the Greek goddess of divine law and order). THEMIS is designed to be modular, extensible, and highly parallel, enabling the extraction of increasingly detailed information from EHT observing campaigns, both individually and in aggregate. Here we present the underlying design philosophy, structure, and validation tests of THEMIS, including the reproduction of a variety of published analyses. We then demonstrate the ability of THEMIS to trivially extend these, presenting new analyses of phenomenological models that include the full set of published EHT observations.

In Section 5.2, we summarize the algorithms, components and implementation details of THEMIS. Individual features are described in Sections 5.3-5.5. Various tests used to validate THEMIS features are presented in Section 5.6. A handful of novel results enabled by THEMIS are collected in Section 5.7. The computational performance of THEMIS and its key components, including the implications for high-performance computing (HPC) systems, is addressed in Section 5.8. Finally, conclusions are summarized in Section 5.9.

5.2 Summary of THEMIS

5.2.1 Structure

The primary goal for THEMIS is to provide an extensible framework for unifying existing and developing future analyses of EHT and auxiliary data. Thus, a key element of THEMIS is the partitioning of the problem, defining a set of independent components in an extensible fashion, ensuring that each may be independently developed. Thereby THEMIS may be continually and effectively developed by the EHT community. Importantly, the practical bar to do so is substantially reduced, requiring would-be developers to understand only the elements of the interface. In the presence of a rapidly evolving data type and modeling effort, this is critical to leveraging the substantial preceding efforts.

THEMIS consists of three distinct collections of components, each of which is designed to be interchangeable:

Data Structures Management and standardization of observational data throughout THEMIS. These facilitate the rapid introduction of new data products, expand the capability of existing data products, and define the objects for which predictions are ultimately made.

Models Any algorithm that produces a prediction for some data object given a list of parameters. Models may be physically motivated or purely phenomenological. They are directly tied to underlying data structures via the declaration of those for which predictions can be made.

Likelihoods, Priors and Samplers Likelihoods provide a method for directly comparing model and data objects. Note that in many cases elements of the underlying model may be subsumed into a likelihood (e.g., nuisance parameters that can be analytically marginalized over). Priors and samplers provide methods for efficiently exploring the model parameter space, providing information about the model parameters.

In practice, there is some overlap between component classes (e.g., Likelihoods and Models), which may be implemented in more than one way. Nevertheless, this has proven sufficiently modular to enable rapid and significant model development already.

All THEMIS-based analyses are structured in the following way:

1. Generate the desired data objects, e.g., by reading in existing data sets.
2. Create an appropriate model object, i.e., declare a model capable of making predictions for the data selected.
3. Specify prior probability distributions for each model parameter.
4. Construct the relevant likelihood objects, combining data sets as desired.
5. Execute a sampler, reporting sampler-specific parameter information (e.g., generate chains for MCMC samplers).

In this way the execution of the analysis is conceptually modularized, enabling variations in each stage to be made trivially.

5.2.2 Implementation

The `main` function is kept concise and is the only element of THEMIS a user that is simply running THEMIS needs to modify. The user may choose interchangeably different EHT data set(s), theoretical model(s), likelihoods, priors, and samplers to employ. Conceptually, this function is organized in a fashion that closely follows the analysis pipeline listed at the end of the previous section to improve usability.

THEMIS also allows users to add wholly new functionality, such as additional models, which can be included easily into a clear and well-established structure. An object-oriented programming framework, along with inheritance, permits a clear and concise definition of component

interfaces. Examples of how these are propagated through various THEMIS components are explicitly illustrated in the inheritance diagrams shown in Figures 5.1 and 5.2. Importantly, in the former, the various predictions enabled by a particular model type (in this case, image-based models, see Section 5.4.1) are shown; for more details see Section 5.4.

THEMIS is under version control provided by `git` with a modern, state-of-the-art branching strategy including master, development and feature branches. Users are encouraged to generate new code branches, develop and contribute to the code in the form of a pull request that will be reviewed by the THEMIS core development team.

A suite of tests is run regularly via a script in an effort to identify bugs or regressions as early as possible. The script performs these tests and sends a report to the THEMIS core development team. These include short tests using EHT data, and range all the way to less frequent and slower, to full scale parameter estimation validation tests similar to the ones presented in Section 5.6.

The code is written in C++ making it maximally portable, and has been tested on a variety of systems. THEMIS is designed with minimal dependency on external libraries to avoid installation conflicts; currently, the only required external libraries are FFTW [107] and the Message Passing Interface (MPI)¹. Up-to-date documentation is critical in a rapid development environment. To meet this challenge, THEMIS has integrated documentation comments which may be optionally rendered via Doxygen² to produce a comprehensive, cross-linked html and/or PDF document.

We now turn to describing each component collection independently.

5.3 THEMIS Data Structures

Within THEMIS, observational data are collected in type-specific data structures. Each has a singular data element defined (a datum object) and an associated plural data structure (a data object) that provide additional input/output facilities and element access functions. At a minimum, these provide access to the values and their uncertainties. Typically, they include a variety of additional “accoutrements”, information necessary or useful in modeling the data. Importantly, these accoutrements are both data-type specific and extensible: information that only becomes useful in subsequent observations or analyses can be added without modifying the data-model interface. For example, observed fluxes may initially include frequency as an accoutrement and later expand to include time, observation facility, etc.

¹Information on the MPI 3.1 standard can be found at www.mpi-forum.org.

²Information on Doxygen features, directives, and on how to obtain and install it, may be found at www.doxygen.org.

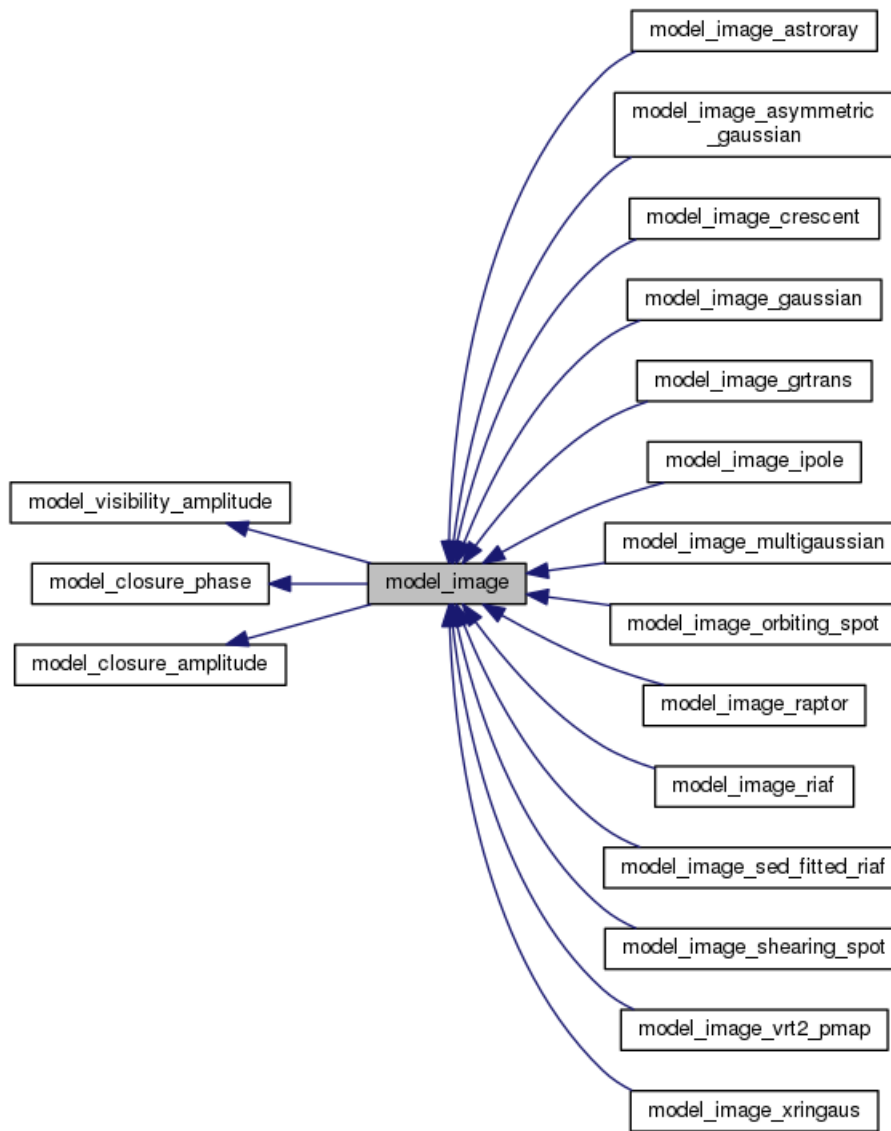


Figure 5.1: Inheritance diagram for the `model_image` object within THEMIS generated via Doxygen. These are models whose primary output is a raster image. Note that a number of models that are either analytically tractable, or extend beyond a single, raster image are not shown. A full listing of THEMIS models can be found in the online documentation.

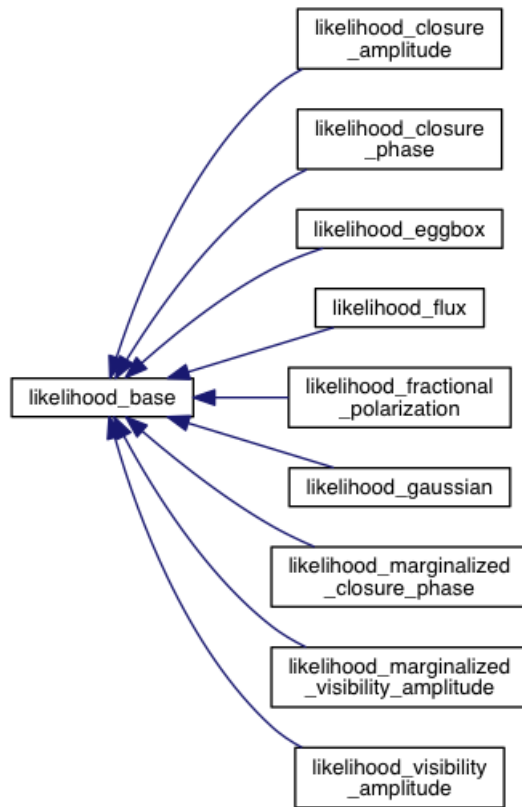


Figure 5.2: Inheritance diagram for the likelihood object within THEMIS generated via Doxygen.

Target	Type	Obs. Yr	Campaign Day(s)	N	Reference
Sgr A*	F	1998-2006		11	[100, 176, 296, 299]
–	VA	2007	100-101	19	[82]
–	VA	2009	95-97	51	[104]
–	CP	2009	93,96-97	24	[105]
–	CP	2011	88,90-91,94	31	[105]
–	CP	2012	81	25	[105]
–	CP	2013	80-82,85-86	101	[105]
–	VA	2013	80-82,85-86	128	[153]
–	LP	2013	80-82,85-86	662	[153]
–	VA	2013	80-82,85-86	861	[169]
–	CP	2013	80-82,85-86	267	[169]
M87	VA	2009	95-97	104	[81]
–	F	2012	81	8	[8]
–	CP	2012	81	17	[8]

Table 5.1: Published EHT Data. Data types include visibility amplitudes (VA), closure phases (CP), interferometric linear polarization fraction (LP), and fluxes (F) and N stands for the Number of data points, including detections only.

Organizing data this way within THEMIS both permits, evolution in how data is employed in model comparisons, and presents a simple way in which to include additional types of data that are currently unforeseen. This is especially important given the wide variety of auxiliary data that exists for EHT targets, most of which has yet to be fully utilized. This has already been implemented for a number of existing data types, including all for which EHT data has already been reported (see Table 5.1). We summarize each of these below.

5.3.1 Visibility Amplitudes

The primary product of VLBI observations are complex visibilities, corresponding to the Fourier modes of the image on the sky at spatial frequencies given by the projected baseline presented by pairs of VLBI stations. Specifically, in the absence of confounding effects, the complex visibility is given by

$$V_{ij} = \int d\alpha d\beta I(\alpha, \beta) e^{-2\pi i(\alpha u + \beta v)}, \quad (5.1)$$

where (u, v) is the two-dimensional projected baseline length between the i th and j th stations expressed in units of the observed wavelength, and $I(\alpha, \beta)$ is the spatial intensity distribution at angular position (α, β) [for a comprehensive introduction to radio interferometry, see, e.g., 271].

In practice, these are modified by a variety of observational complications, chief among which are atmospheric absorption and phase delays at individual stations, which impact the amplitude and complex phase of V_{ij} . Of these, the latter are especially problematic, resulting in phase shifts of the V_{ij} by many times 2π , effectively randomizing the phase on every baseline. As a result, often the magnitudes of the visibilities, $|V_{ij}|$, are employed, which are subject only to a comparably modest uncertainty, 1%-20% depending on station and atmospheric conditions [see, e.g., 153, 169], albeit containing less information on the structure of the image. The number of visibility amplitudes generated by an interferometer grows quadratically with the number of stations, N , scaling as $\propto N(N-1)/2$. Throughout an observing campaign, the rotation of the Earth produces a large number of independent measurements at different projected baselines.

Already a large number of EHT visibility amplitudes have been published for the primary EHT targets, beginning in 2007, and extending through 2013 [82, 104, 153, 169]. We list these in Table 5.1.

5.3.2 Closure Phases

While atmospheric phase delays typically preclude the reconstruction of the phase of the complex visibilities³, it is possible, nevertheless, to obtain some information about these phases via the closure phase,

$$\Phi_{ijk} = \arg(V_{ij}V_{jk}V_{ki}), \quad (5.2)$$

i.e., the sum of the phases of a triplet of visibilities measured on the baselines between some triplet of stations⁴. Because the baselines “close”, i.e., $(u, v)_{ij} + (u, v)_{jk} + (u, v)_{ki} = 0$, all station-specific phase errors vanish identically, leaving a quantity that depends solely on the image structure. Of particular importance, closure phases are also insensitive to the image blurring induced by the diffractive component of the interstellar scattering. Closure phases are not unique — for an array with N stations only $(N-1)(N-2)/2$ are independent — a result that is presaged by their independence of the phase delays.

Closure phases have been reported by the EHT for Sgr A* for a number of years in [104] and [169], and summarized in Table 5.1.

³This is not true if a phase reference is used, typically an extragalactic background source (see, e.g., [40]), or if multiple wavelengths are simultaneously observed, permitting one to be phase referenced to the other [184].

⁴This is the argument of the bispectrum.

5.3.3 Closure Amplitudes

Station-specific amplitude calibration errors can also be mitigated by combining visibilities measured on multiple baselines. The closure amplitude is constructed from combinations of visibilities measured on four stations,

$$\mathcal{V}_{ijklm} = \frac{|V_{ij}||V_{km}|}{|V_{ik}||V_{jm}|}, \quad (5.3)$$

and is insensitive to variations in the flux calibration and phase delays. Again, closure amplitudes are also insensitive to the image blurring induced by the diffractive component of the interstellar scattering. As with the closure phase, this comes at the price of uniqueness; there are only $N(N-3)/2$ independent closure amplitudes.

Closure amplitudes constructed from EHT data have not yet been published, primarily due to the limited number of stations participating in early observations. However, recent observations have generated a number of trivial closure amplitudes, i.e., amplitudes for which one baseline is very short, as part of the calibration process [see, e.g., 153]. Beginning with the April, 2017 observations, many non-trivial closure amplitudes can be anticipated.

5.3.4 Interferometric Polarization Fractions

The EHT observes in all four Stokes parameters, (I, Q, U, V). Independently, these can be used to construct visibility amplitudes, closure phases, and closure amplitudes. However, additional information may be obtained by combining observations made in different Stokes parameters. The interferometric polarization fraction,

$$\check{m}_{ij} = \frac{\sqrt{|V_{ij}^Q|^2 + |V_{ij}^U|^2}}{|V_{ij}|}, \quad (5.4)$$

where $V_{ij}^{Q,U}$ are the visibilities associated with Stokes Q and U , and V_{ij} is the visibility defined in Equation (5.1), is the extension of the familiar polarization fraction to the individual Fourier modes of the image. \check{m} is not to be mistaken with the Fourier transform of the linear polarization fraction as measured in the image domain. Unlike the standard polarization fraction, \check{m}_{ij} may be larger than unity, and can exhibit counter-intuitive pathologies for even simple source models [see the discussion surrounding Figure S6 in the supplemental material in 153]. Like closure amplitudes, the interferometric polarization fractions are insensitive to station-specific flux calibration uncertainties and the diffractive component of the interstellar scattering.

Interferometric polarization fractions have been reported for Sgr A*, are typically quite large and indicate the presence of ordered horizon-scale polarization structures [153]. We summarize these in Table 5.1.

5.3.5 Flux Measurements

A key auxiliary set of observations are the spectral energy density distributions (SED) for primary EHT targets, which typically place strong limits on the uncertain emitting particle distributions. In addition, multi-wavelength light curves are a key probe of the nature and origin of variability in the emission regions of the source. Both empirical constraints are intrinsically encoded in measurements of the unresolved source flux, F_ν , effectively equivalent to the visibility amplitudes measured at “zero-baseline”, i.e. neighboring antennas. The distinction between these arises in the accoutrements associated with the data, e.g., the origin of the observation, wavelength, time, etc.

Multiple sets of flux measurement data for Sgr A* and M87 exist. For Sgr A*, one set is summarized in Table 5.1.

5.4 THEMIS Models

Within THEMIS, a model is any algorithm capable of generating a prediction for any THEMIS data type. Thus, THEMIS models are closely aligned with THEMIS data structures — for each data type there is a corresponding base model type. Models can encompass multiple base model types, i.e., they can be capable of generating predictions for more than one type of data. This enables a broad, easily extensible, and backwards compatible framework for defining models that permits the incrementally increasing sophistication. Importantly, it provides a uniform interface for both phenomenological models, which are designed to make predictions for a handful of data types, and physically motivated models, which can simultaneously make predictions for a wide variety of data types.

The manner in which predictions are made is not prescribed. That is, where analytical expressions for the relevant data type exist (e.g., visibilities from simple geometric models), models are capable of employing these. For more complex models, numerical computations are often required. In anticipation of numerically produced predictions, THEMIS permits the passing of an accuracy parameter for each value that specifies the accuracy with which these must be generated; typically setting this to 25% of the measurement uncertainty is sufficient to generate accurate parameter estimates (see Appendix C).

5.4.1 Image-based Models

Because the EHT directly probes the structure of horizon-scale images, THEMIS contains an image-based model type. This provides a set of utilities for generating and manipulating visibility-based data from models that primarily generate images.

Because image generation is frequently computationally intensive, the image-based model introduces an additional position angle parameter, permitting the specific model implementations to dispense with trivial image rotations, leading to a substantial potential reduction in the time required to sample a broad range of parameters.

Once generated, images are padded with zeros by a factor of 8 by default to effect sinc-interpolation in the numerically computed complex visibilities. The complex visibilities are computed on a two-dimensional grid of (u, v) values via a two-dimensional Fast Fourier Transform using the FFTW library [107]. There are no restrictions on the image dimensions, though it is expected that the image is computed on a rectilinear grid with uniform pixel size; dimensions that factor into small primes will be marginally faster.

Complex visibilities are then estimated at any arbitrary (u, v) via interpolation. By default THEMIS employs bicubic interpolation, though a user may specify bicubic spline interpolation if desired. From these, the closure phases are constructed via Equation (5.2). While visibility magnitudes may also be constructed from the interpolated complex visibilities, it is considerably more accurate to interpolate the visibility magnitudes directly.⁵ These are then used directly or to compute closure amplitudes via Equation (5.3).

5.4.2 Phenomenological Geometric Models

Within THEMIS, a number of phenomenological geometric models have been implemented. These are models for which no underlying physical emission mechanism is identified for the origin of the image structures. However, such models are capable of extracting signatures of geometric features associated with underlying physical processes of interest, e.g., black hole shadows. Currently implemented phenomenological models include the following.

⁵The magnitude of the gradient of the complex visibility and the gradient of the visibility amplitude are related via $|\nabla V|^2 = (\nabla|V|)^2 + |V|^2(\nabla\phi)^2 \geq (\nabla|V|)^2$, and thus the former is generally smaller than the latter. As a result, the errors in interpolation at any order are typically smaller when interpolating visibility amplitudes directly. Alternatively, this permits considerably smaller image sizes when only amplitudes are required.

Symmetric Gaussian

Historically, the first shadow size estimates from mm-VLBI observations of Sgr A* and M87 arose from fitting symmetric Gaussians to visibility amplitude measurements [82]. Therefore, we have implemented within THEMIS a model consisting of a single symmetric Gaussian component, characterized by a size, σ , and an amplitude, V_0 . This makes predictions for visibility amplitudes, closure phases (trivially zero), and closure amplitudes.

Asymmetric Gaussian

The introduction of asymmetry in mm-VLBI images was initially characterized by an asymmetric Gaussian. Within THEMIS we have implemented such a Gaussian model parameterized as in [40], and characterized by a size, σ , an asymmetry parameter, A , the amplitude, V_0 , and the position angle, ξ .

Multiple Symmetric Gaussian

THEMIS also includes a model consisting of an arbitrary number of symmetric Gaussian components, each characterized by a size, σ_j , location, (x_j, y_j) , and amplitude, V_j .

Crescent Model

THEMIS includes an implementation of the crescent model described in [155], for which the image is obtained by subtracting two non-concentric discs, with the smaller disc lying completely inside the larger one. The complex visibilities for this model can be obtained analytically and are given by Equation (3) of [155]. As in [155], we reparameterize this in terms of an amplitude, V_0 , overall size, R , relative thickness, ψ , degree of symmetry, τ , and the position angle, ξ . Both ψ and τ are defined on the unit interval.

The “xringaus” model

THEMIS also contains an implementation of the nine-parameter *xringaus* model proposed in [23]. This model was constructed in an effort to mimic a more realistic black hole accretion image like the ones commonly obtained from physically motivated models. The xringaus image is the combination of an eccentric slashed ring and an elliptical Gaussian located in the brighter side of the ring.

This model is then described by a tuple of nine parameters: the zero-spacing flux, V_0 , the external radius, R_{ex} , the internal radius, R_{in} , the distance between centers of the circles, d , the “fading” parameter controlling the minimum brightness, the Gaussian axes sizes, a and b , the fraction of the total flux in the Gaussian, g_q , and the position angle, ξ . The complex visibilities for this model, in terms of these parameters, can be also obtained analytically. The reader is referred to Section 2 of [23] for a more detailed description.

Visual Binary

THEMIS also features a model of two Radio emitting Gaussian components in orbit around each other. The model is characterized by a tuple of 13 parameters including the total flux F_i , size σ_i , and spectral index α_i of each component, the total mass of the system M , the binary mass ratio, $q \leq 1$, the orbital separation, R , the source distance d , the phase offset Φ_0 , the cosine of the inclination angle, $\cos(i)$ of the orbital angular momentum vector and the position angle in the sky, ξ . This model includes (and therefore also takes advantage of) relativistic effects such as Doppler boost and relativistic aberration. It is explicitly time-dependent while being fully analytic and thus fast to evaluate.

This model is to be compared to long time scale monitoring campaigns of sources such as OJ 287 or other binary candidates. Details will be published in a separate paper which focusses on this topic.

5.4.3 Interstellar Scattering Models

Interstellar scattering modifies the intrinsic images of Sgr A* by both blurring the image (diffractive component) and adding small scale structures associated with a random realization of refractive modes that vary slowly throughout the night [refractive component; see, e.g., 154]. These significantly modify visibilities on long baselines, and must be included in analyses of EHT observations of Sgr A*.

When many realizations of the scattering screen are averaged over, e.g., after many observing nights, i.e., when the scattering may be treated in the ensemble average limit, only the diffractive component is present. This appears as an image smoothing via convolution with a Gaussian kernel whose parameters depend on the details of the intervening scattering screen(s). THEMIS has implemented two models for addressing interstellar scattering, both in the ensemble average limit, which we list below. In both, the impact of scattering is imposed directly on the visibilities, for which the convolution in image space reduces to a multiplicative factor. Within THEMIS, each

are implemented as a model that modifies an existing intrinsic model, with the latter introducing additional parameters. Hence, scattering provides an explicit example of how the modular structure of THEMIS enables the rapid construction of new models.

Default Diffractive Screen

Multi-wavelength observations have produced a model for the scattering kernel that is asymmetric and wavelength dependent, consistent with that anticipated by models of the scattering screen that invoke Kolomogorov turbulence within a plasma sheet [34]. The associated semi-major and semi-minor axis sizes are given by

$$\begin{aligned}\sigma_{\text{maj}} &= 9.39 \left(\frac{\lambda}{1.3 \text{ mm}} \right)^2 \mu\text{as} \\ \sigma_{\text{min}} &= 4.59 \left(\frac{\lambda}{1.3 \text{ mm}} \right)^2 \mu\text{as},\end{aligned}\tag{5.5}$$

and are oriented such that the major axis lies 78° East of North.

Parameterized Diffractive Screen

Recently, it has been shown that even for thin scattering screens, the wavelength dependence of anisotropic scattering screens may be substantially more complicated [218]. The main uncertainty is the inner-scale of the turbulence within the screen, corresponding to the dissipative scale within the sheet. For some plausible values, the wavelength dependence could depart from that found in [34] near 1.3 mm. As a result, a second scattering model has been implemented in which σ_{maj} , σ_{min} , and the position angle are all parameterized as power laws of wavelength with unknown coefficients and powers. That is,

$$\begin{aligned}\sigma_{\text{maj}} &= \sigma_A \left(\frac{\lambda}{\lambda_p} \right)^\alpha, & \sigma_{\text{min}} &= \sigma_B \left(\frac{\lambda}{\lambda_p} \right)^\beta, \\ \text{and } \xi &= \xi_0 + \xi_1 \left[\left(\frac{\lambda}{\lambda_p} \right)^\gamma - 1 \right],\end{aligned}\tag{5.6}$$

where the 7 parameters, σ_A , σ_B , ξ_0 , ξ_1 , α , β , and γ may be varied. The pivot wavelength, λ_p , is set by the user.

5.4.4 Native Physical Models

The past two decades have seen the development of a number of physically motivated models which employ ray tracing and radiative transfer in black hole spacetimes. These have two main components: the construction of photon trajectories within the spacetime under consideration, and the radiative transfer through some emitting plasma distribution. Both elements are directly affected by variations in the spacetime structure, with the emission also depending on a number of astrophysical considerations.

While this class of models is substantially more complicated than geometric models, their physical origin presents a number of significant advantages. First, they are capable of making predictions for a wide range of observations, making it possible to bring far more empirical data to bear upon them. For example, they necessarily make simultaneous, self-consistent predictions for images, fluxes, variability, and polarization features of the EHT and auxiliary data [44, 48, 56, 59, 77, 126, 146, 193, 221]. Hence, physical modeling enables a concordance fitting effort that promises far more power to constrain the nature of the emission region [41, 59, 126]. Second, the spacetime structure impacts the image in many ways beyond gravitational lensing. The dynamics of the material in the emission region modifies its optical depth, and therefore appearance [38, 39, 43, 45, 46, 49, 58, 77, 126, 149, 194, 195, 221]. Thus, in principle, modeling the brightness distribution offers additional probes of gravity [42, 45, 150, 151, 187]. Third, it provides direct information about the high-energy astrophysical processes responsible for the growth of black holes and the launching of jets [46, 47, 77, 126, 143, 193].

Within THEMIS, two general relativistic ray tracing and radiative transfer packages are provided. The first of these is the vacuum ray tracing and radiative transfer package VRT² (Vacuum Ray Tracing and Radiative Transfer). VRT² is based on the plasma radiative transfer package described in [38, 39] and provides a modular framework for adding novel plasma distributions, radiative transfer mechanisms, and spacetime structures. It was the basis for the images generated in e.g. [40] and used in the analysis of [41]. It also natively interfaces with THEMIS, having been written in the same programming language (C++), in a similar style. Models based on VRT² within THEMIS include those listed below.

In addition, the vacuum ray tracing and radiative transfer package Odyssey described in [222] has also been incorporated within THEMIS. Based on the ray tracing algorithm in [113] and the radiative transfer formula presented in [293], Odyssey can exploit graphics processing unit (GPU) cards to realize substantial speed gains for models that employ it. It requires the Compute Unified Device Architecture (CUDA)-enabled GPU cards and the CUDA compiler *nvcc*. Again, like THEMIS, Odyssey is implemented in C/C++ and CUDA in C/C++, making its integration straightforward.




Figure 5.3: Image produced by the radiative transfer module VRT² showing the resulting best-fit RIAF model after sampling the full parameter space. This reproduces the result in [41]. The X and Y axis show image coordinates in units of the gravitational radius.

SED-fitted RIAF

This is an image at a single wavelength associated with the radiatively inefficient accretion flow (RIAF) models described in [44] and refined in [40]. This model employs a tabulated set of accretion flow parameters, obtained at different black hole spins and inclinations, that reproduce the observed SED of Sgr A*. The model parameters are the dimensionless spin magnitude, a (in the range $[0,1]$), the cosine of the inclination, $\cos\theta$, ($[-1,1]$), and the position angle, ξ ($[-180^\circ, 180^\circ]$, as part of a model image). The intensity normalization may be included in via the likelihood (see Section 5.5.1). An example image from the THEMIS-integrated VRT² package is shown in Figure 5.3.

Extended RIAF

This is an extension of the SED-fitted RIAF model that permits a wide range of structural parameters in the RIAF model to vary. This consists of two populations of synchrotron emitting electrons, orbiting a Kerr black hole in the presence of a toroidal magnetic field. Specifically, the proper number density and temperature of a thermal population of electrons are given by

$$n_{\text{th}} = n_{e,t} r^{\eta_t} e^{-z^2/2h_t^2 R^2}, \quad T_{\text{th}} = T_e r^{\tau_t}, \quad (5.7)$$

where $z = r \cos \theta$ and $R = r \sin \theta$ where r is the standard Boyer-Lindquist radius (measured in GM/c^2) and θ is the Boyer-Lindquist polar angle. Similarly, the proper number density of the nonthermal electrons is given by

$$n_{\text{nth}} = n_{e,nt} r^{\eta_{nt}} e^{-z^2/2h_{nt}^2 R^2}, \quad (5.8)$$

and has a power-law distribution in microscopic Lorentz factor above γ_{min} with a power law corresponding to an optically thin spectral index of α (i.e., $2\alpha - 1$). These are emitting within a toroidal magnetic field with comoving strength

$$\frac{B^2}{8\pi} = \beta^{-1} \frac{m_p c^2}{6r}, \quad (5.9)$$

and orbiting with a four-velocity outside of the innermost stable circular orbit (ISCO) given by

$$u_\mu = u_t(1, 0, 0, \kappa \ell_K) \quad (5.10)$$

where ℓ_K is the specific Keplerian angular momentum and u_t is determined by the standard normalization condition on u^μ ; inside of the ISCO the material plunges on ballistic orbits. Thus, there are 15 parameters: black hole spin, a , cosine of the black hole spin inclination, $\cos \theta$, black hole spin position angle, ξ ; thermal electron density normalization, $n_{e,t}$, radial power-law, η_t , and scale-height, h_t , electron temperature normalization, T_e , and radial power law, τ_t ; nonthermal electron density normalization, $n_{e,nt}$, radial power-law, η_{nt} , scale-height h_{nt} , minimum microscopic Lorentz factor, γ_{min} , and spectral index, α ; plasma beta, β and sub-Keplerian fraction κ . Note that subsets of these may be held fixed or varied simultaneously via the definition of a wrapper model.

Orbiting Hot Spots

Major dissipative events within the accretion flow, such as magnetic reconnection events and shocks, can generate initially compact, orbiting, synchrotron emitting hot spots. These may

increase the emission of Sgr A* by orders of magnitude before inducing any dynamical effects. Therefore, they may be roughly modeled as orbiting, Gaussian non-thermal particle overdensities that subsequently synchrotron emit in the radio and infrared, restricted to the equatorial plane [43, 45]. To model the velocity profile of the spot we use a two-parameter, $(\alpha_r, \kappa) \in [0, 1]$, four-velocity given by

$$u^\mu = (u^t, u_{\alpha_r}^r, 0, u^t \Omega_\kappa). \quad (5.11)$$

Here $u_{\alpha_r}^r = u_K^r + \alpha_r(u_{ff}^r - u_K^r)$ and $\Omega_\kappa = \Omega_{ff} + \kappa(\Omega_K - \Omega_{ff})$, where K, ff subscripts denote Keplerian and free fall motion respectively, and $\Omega_i = u_i^\phi / u_i^t$ [see also 221]. Thus, there are 10 parameters needed for this model: black hole spin, a , cosine of the black hole spin inclination $\cos\theta$, black hole spin position angle ξ ; central spot non-thermal electron density $n_{e,spot}$, spot radial size R_s ; initial spot location in time, t_0 , radius, r_0 , and azimuthal angle ϕ_0 ; the sub-Keplerian parameter κ , and the radial infall parameter α_r .

Shearing Hot Spots

In practice, hot spots will subsequently shear and cool. Thus, THEMIS also includes a shearing hot spot model [149], that incorporates the expansion of the hot spots within a background accretion flow. The parameters of this model are identical to the orbiting spot model above.

5.4.5 External Physical Models

There is no intrinsic bar to including additional ray tracing and radiative transfer packages within THEMIS. Doing so offers a number of benefits, including the ability to rapidly generate new models within THEMIS itself, efficient parallelization and improved portability. However, native integration is not necessary. It is often initially faster, and occasionally necessary, to externally include modeling software. For THEMIS this has been done for a number of existing packages:

GRTRANS A publicly available general relativistic, polarized radiative transfer code written in *FORTRAN*, see [76, 77]. *GRTRANS* and by extension also THEMIS is coupled to the HARM3D GRMHD code [77, 117, 180].

ASTRORAY A significantly extended version of the general relativistic polarized radiative transfer code written in C/C++ based on [252] and substantially extended in [126]. *ASTRORAY* and by extension THEMIS is coupled to HARM3D [117, 179, 180].

iPOLE A publicly available general relativistic, polarized radiative transfer code [194] based on the covariant formulation presented in [116] and written in standard C. *iPOLE* and by extension `THEMIS` are coupled to `HARM3D` [83, 117, 195].

RAPTOR A publicly available general relativistic radiative transfer code, see [49] written in standard C. *RAPTOR* and by extension `THEMIS` is coupled to the *BHAC* [201] GRMHD code, `HARM3D` [117, 195] and is GPU capable.

Note that many of these are directly coupled to a variety of existing GRMHD simulation codes such as *HARM3D* and *BHAC*. As of now, `THEMIS` has successfully interfaced, in at least a limited form, with the vast majority of the image generation tools employed by the EHT collaboration.

5.5 Likelihoods, Priors and Samplers

Models and data are systematically compared via likelihoods, which express the probability that the data was obtained from the model. These are then explored by samplers, which explore the dependence of the likelihood on the model parameters, incorporating any priors on the parameter values. Here we describe the various elements of each as implemented in `THEMIS`.

5.5.1 Likelihoods

Within `THEMIS` likelihood is any method for taking a parameter vector, \mathbf{p} , and construct a log-likelihood, \mathcal{L} . When this is generated using a `THEMIS` data object (consisting of a number of individual values) and a `THEMIS` model object the log-likelihood is the probability of obtaining the data given the model. Likelihoods can be combined with user-supplied weights, enabling the combination of various data sets. However, when doing so it is assumed that the model parameters are unchanged, i.e., the same set of model parameters are to be supplied to each likelihood being used. All likelihoods expect a matching data type and model type, e.g., visibility amplitude data and a model that generates visibility amplitude predictions.

The likelihood generally requires information about the underlying error distribution of the data, which is typically provided via an error estimate. All currently implemented likelihoods in `THEMIS` assume Gaussian errors, though this is not required — likelihood classes that assume alternative error distributions (e.g., Rice distributions, etc.) are possible. Similarly, all currently implemented likelihoods assume the data values are independent — this too may be relaxed in

principle. An obvious example of both that is of considerable interest is the covariance induced by the refractive modes in the scattering screen.

Likelihoods also can incorporate model features. In many instances, a subset of model parameters may be analytically marginalized over, and in the process subsumed into the likelihood itself. We have implemented a number of examples of such “marginalized” likelihoods, i.e., likelihoods in which sets of nuisance parameters have been treated analytically.

The likelihoods currently implemented in THEMIS include the following.

Test Cases

To facilitate testing samplers, THEMIS includes two artificial likelihoods with given distributions. The first is a multi-dimensional Gaussian, with user-specified mean and size. The second, the Egg Box, is considerably more complicated, producing a highly multimodal likelihood function in 5 dimensions:

$$\mathbb{L}(\mathbf{p}) = \left[2 + \prod_{i=0}^5 \cos(p_i) \right]^5. \quad (5.12)$$

The number of peaks can be set by the range over which the priors permit the parameters, \mathbf{p} , to vary. This is typically used to assess the ability of a sampler to accurately find widely separated, high-likelihood regions.

Visibility Amplitudes

THEMIS includes a log-likelihood that assumes Gaussian errors for visibility amplitudes:

$$\mathbb{L}(\mathbf{p}) = - \sum_j \frac{[|V|_j - |\hat{V}|_j(\mathbf{p})]^2}{2\sigma_j^2}, \quad (5.13)$$

where $|V|_j$ and σ_j are the observed visibility amplitudes and their errors, and $|\hat{V}|_j(\mathbf{p})$ are the model visibility amplitudes given parameters \mathbf{p} . Note that the true visibility amplitude error distribution is given by the Rice distribution, and for low signal-to-noise ratio (SNR) is both biased and non-Gaussian [271]. However, when data are selected such that $SNR \geq 2$ and approximately debiased via $|V|_j \rightarrow \sqrt{|V|_j^2 - \sigma_j^2}$, the visibility amplitude error distribution is within 8% of an unbiased Gaussian distribution at all $|V|$, reproduces the mode to better than 1% and the 68% and 95% cumulative widths to better than 6% (see Appendix C.1). Currently, the user is required to independently implement the debiasing procedure in the generation of the data tables prior to reading them in THEMIS.

Closure Phases

Similarly, THEMIS includes a log-likelihood that assumes Gaussian errors for closure phases:

$$\mathfrak{L}(\mathbf{p}) = - \sum_j \frac{[\Phi_j - \hat{\Phi}_j(\mathbf{p})]^2}{2\sigma_j^2}, \quad (5.14)$$

where Φ_j and σ_j are the observed closure phases and their errors, and $\hat{\Phi}_j(\mathbf{p})$ is the model closure phase given parameters \mathbf{p} . This is similar to, but distinct from the visibility amplitudes likelihood in that the difference, $[\Phi_j - \hat{\Phi}_j(\mathbf{p})]$, is chosen on $[-180^\circ, 180^\circ)$, selecting the branch that minimizes the angular difference.

Closure Amplitudes

Closure amplitudes provide an example of a non-Gaussian likelihood within THEMIS. Because closure amplitudes are constructed via taking ratios of visibility amplitudes, the likelihood of a single value exhibits a significant asymmetry and extended tail towards large values, characteristic of quotient distributions (see Appendix C.1). For $\text{SNR} \geq 4$, this is well approximated by a Gaussian quotient distribution, given in Equation (C.24). In principle, this can make use of ancillary information in the form of station system equivalent flux densities (SEFDs), though this is left for future development. Thus, at present, we assume that the parameter ρ , defined in Equation (C.21), is fixed to unity, for which the Gaussian quotient approximation is accurate at all \mathcal{V} to better than 13% for $\text{SNR} \geq 4$. The associated log-likelihood is

$$\begin{aligned} \mathfrak{L}(\mathbf{p}) = & - \sum_j \left\{ \frac{(\mathcal{V}_j - \hat{\mathcal{V}}_j)^2}{2\Sigma_j^2} - \log\left(\frac{\Sigma_j}{\sigma_j}\right) \right. \\ & \left. + \log \left[\Delta_j \text{erf}\left(\frac{\Delta_j}{\sqrt{2}\Omega_j}\right) - \frac{2\Omega_j}{\sqrt{2\pi}} e^{-\Delta_j^2/2\Omega_j^2} \right] \right\}, \end{aligned} \quad (5.15)$$

where $\Sigma_j^2 = \sigma_j^2 \frac{1 + \mathcal{V}_j^2}{1 + \hat{\mathcal{V}}_j^2}$, $\Delta_j = \frac{1 + \mathcal{V}_j \hat{\mathcal{V}}_j}{1 + \mathcal{V}_j^2}$,

and $\Omega_j^2 = \frac{\sigma_j^2}{(1 + \mathcal{V}_j^2)(1 + \hat{\mathcal{V}}_j^2)}$,

where the \mathcal{V}_j and σ_j are the observed closure amplitudes, the $\hat{\mathcal{V}}_j(\mathbf{p})$ are the model visibility amplitudes given parameters \mathbf{p} (with the functional dependence suppressed for clarity), and $\text{erf}(x)$

is the error function. This differs from Equation (C.24) by constant normalization factors. In the limit of $\sigma_j/\hat{\mathcal{V}}_j \rightarrow 0$ the third term vanishes. However, for \mathcal{V}_j of order unity, Equation (5.15) does not reduce to a Gaussian distribution in any SNR limit.

Finally, we note that this approximation is significantly better when $\hat{\mathcal{V}}_j$ is small. Generally, the closure amplitudes can be constructed such that $\mathcal{V}_j < 1$, approximating this requirement. Currently, the user is expected to independently define the set of closure amplitudes such that this is true prior to reading them into THEMIS.

Interferometric Polarization Fractions

The interferometric polarization fraction provides a second example of a non-Gaussian likelihood available in THEMIS. As with the closure amplitude, the source of the non-Gaussianity is the presence of the ratio in their definition. This leads to an asymmetric likelihood with an extended tail towards large \check{m} that is also well approximated by a Gaussian quotient distribution for $\text{SNR} \geq 2$, given in Equation (C.17). That \check{m} is defined by the ratio of visibilities constructed simultaneously on the same baseline places an additional constraint on the likelihood, permitting it to accurately be described by a single noise parameter (Appendix C.1): for $\text{SNR} \geq 2$, the Gaussian quotient distribution is accurate at all \check{m} to 13% for $\text{SNR}=2$ and 6% for $\text{SNR} \geq 4$. The associated log-likelihood is identical in form to Equation (5.15):

$$\begin{aligned} \mathbb{L}(\mathbf{p}) = & - \sum_j \left\{ \frac{(\check{m}_j - \hat{m}_j)^2}{2\Sigma_j^2} - \log\left(\frac{\Sigma_j}{\sigma_j}\right) \right. \\ & \left. + \log\left[\Delta_j \text{erf}\left(\frac{\Delta_j}{\sqrt{2}\Omega_j}\right) - \frac{2\Omega_j}{\sqrt{2\pi}} e^{-\Delta_j^2/2\Omega_j^2} \right] \right\}, \\ \text{where } \Sigma_j^2 = & \sigma_j^2 \frac{1 + \check{m}_j^2}{1 + \hat{m}_j^2}, \quad \Delta_j = \frac{1 + \check{m}_j \hat{m}_j}{1 + \check{m}_j^2}, \\ \text{and } \Omega_j^2 = & \frac{\sigma_j^2}{(1 + \check{m}_j^2)(1 + \hat{m}_j^2)}, \end{aligned} \tag{5.16}$$

were \check{m}_j and σ_j are the observed polarization fraction and its uncertainty, $\hat{m}_j(\mathbf{p})$ are the model polarization fractions associated with parameters \mathbf{p} (with the functional dependence suppressed for clarity). This differs from Equation (C.17) by constant normalization factors. As with the closure amplitudes, this is non-Gaussian even in the limit of $\sigma_j/\check{m}_j \rightarrow 0$.

Norm-Marginalized Visibility Amplitudes

Variations in the total source flux can be directly incorporated into the likelihood. Assuming Gaussian errors for visibility amplitudes, it is possible to introduce and analytically marginalize over an over-all normalization, V_{00} , presuming a flat prior [42]. This provides both the maximum log-likelihood:

$$\mathcal{L}_{\max} = -\frac{(\sum_j V_j^2/\sigma_j^2)(\sum_j \hat{V}_j^2/\sigma_j^2) - (\sum_j V_j \hat{V}_j/\sigma_j^2)^2}{2 \sum_j \hat{V}_j^2/\sigma_j^2}, \quad (5.17)$$

which may be identified with the minimum χ^2 , occurring at

$$V_{00,\max} = \frac{\sum_j V_j \hat{V}_j/\sigma_j^2}{\sum_j \hat{V}_j^2/\sigma_j^2}. \quad (5.18)$$

More relevant for sampling is the marginalized log-likelihood:

$$\bar{\mathcal{L}} = \mathcal{L}_{\max} + \frac{1}{2} \log \left(\frac{2\pi V_{00,\max}^2}{\sum_j \hat{V}_j^2/\sigma_j^2} \right) \quad (5.19)$$

with the corresponding marginalized normalization, $V_{00,\text{marg}} = V_{00,\max}$.

By breaking the visibility amplitude data into epochs with similar visibility normalizations, corresponding, e.g., to a variable accretion rate, this can substantially increase the efficiency of sampling the remaining parameter space.

Shift-Marginalized Closure Phases

At lowest order, refractive scattering induces shifts in the closure phase. These, again, may be incorporated into an appropriately constructed likelihood. Assuming Gaussian errors for closure phases, it is possible to analytically marginalize over an over-all shift, ϕ , presuming a user-supplied Gaussian prior [41]. This provides both the maximum log-likelihood

$$\mathcal{L}_{\max} = -\frac{\phi_{\max}^2}{2\Sigma^2} - \sum_j \frac{(\Phi_j - \hat{\Phi}_j)^2}{2\sigma_j^2}, \quad (5.20)$$

where the most likely phase offset is,

$$\phi_{\max} = \Sigma^2 \sum_j \frac{\Phi_j - \hat{\Phi}_j}{\sigma_j^2} \quad \text{where} \quad \frac{1}{\Sigma^2} \equiv \sum_j \frac{1}{\sigma_j^2}, \quad (5.21)$$

The maximum log-likelihood is trivially related to the χ^2 and is relevant for fit quality assessment. More relevant for parameter estimation is the marginalized likelihood, for which the log-likelihood is given by

$$\bar{\mathcal{L}} = \mathcal{L}_{\max} - \frac{\phi_{\max}^2}{2(\sigma_{\Phi}^2 + \Sigma^2)} + \log \left(\frac{\Sigma}{\sqrt{\sigma_{\Phi}^2 + \Sigma^2}} \right), \quad (5.22)$$

with an associated marginalized value of the closure phase shift of

$$\bar{\phi} = \frac{\sigma_{\Phi}^2}{\sigma_{\Phi}^2 + \Sigma^2} \phi^M. \quad (5.23)$$

Here σ_{Φ} is the width of the Gaussian prior on ϕ ; it is indicative of the amplitude of the refraction or turbulence responsible for the inter-epoch closure phase fluctuations. This marginalized log-likelihood is appropriate for sampling the remaining parameters. By breaking the closure phase data into epochs with similar visibility normalizations, corresponding, e.g., to a variable accretion rate, this can substantially increase the efficiency of sampling the remaining parameter space.

5.5.2 Priors

THEMIS provides a number of potential priors for individual parameters. These may be imposed in two distinct ways: as “priors” that modify the likelihood and “transforms” that modify the parameter values. Within THEMIS, “priors” add a term associated with a given prior distribution. These are trivially implemented and easy to understand. However, they can be inefficient, assuming that the sampler will efficiently incorporate the modified likelihood. In contrast, “transforms” impose priors indirectly by mapping the variable being sampled into the desired prior via a coordinate transformation. These are more complicated to implement, typically requiring the integration of the desired prior probability distribution. However, they are optimally efficient, permitting the sampler to apply a more natural distribution. Note that “transforms” may be implemented intrinsically within models by choosing a convenient set of parameters.

Likelihood evaluation is short-circuited on the evaluation of priors, i.e., where the prior has zero probability (e.g., outside the limits of a linear range), the likelihood is not evaluated but rather returns the appropriate vanishing value. This achieves two goals: first, THEMIS is made marginally more efficient by avoiding unnecessary computation, and second, permits priors to be used to avoid unphysical parameter combinations, where models may return nonsensical results, e.g., negative densities passed to a radiative transfer code or black hole spin outside the range permitted by General Relativity.

Currently, THEMIS has only implemented priors and transforms of a single variable. This is sufficient for most situations. However, there are situations which may benefit from priors that depend on many parameters, e.g., enforcing an ordering among the intensities of multiple Gaussian components, thereby eliminating the trivial degeneracy associated with swapping components. Nevertheless, there is no reason that such a prior cannot be implemented within THEMIS.

Implemented priors include:

- **None:** a flat prior without boundary.
- **Linear:** a flat prior given two bounding values.
- **Logarithmic:** a logarithmic prior given two bounding values.
- **Gaussian** a Gaussian prior given a mean and standard deviation.

And implemented transforms include:

- **None:** no transformation (default).
- **Fixed:** returns a single, user-defined value.
- **Logarithmic:** effectively imposes a logarithmic prior.

5.5.3 Samplers

The process of sampling is conceptually separated from the definitions of data and models through the standardization of the likelihood objects. Thus, within THEMIS, a sampler is any method for exploring the values of a likelihood for various choices of the parameter vector. There is no standard output or input for a sampler, which may even vary qualitatively depending on the goal of the sampling process. However, all samplers interface with data and model objects solely through the use of likelihood objects, and thereby permit analyses of a wide variety of combinations of data and models. Implemented samplers include the following.

Grid Search

The conceptually simplest but least efficient is a simple grid search where the parameter space is probed in predetermined fixed steps in each dimension. While limited in computational efficiency this scheme is often used to cross-check results obtained by other samplers for smaller parameter spaces that both schemes can handle.

Parallel-Tempered, Affine-Invariant Markov Chain Monte Carlo

The natural choice for high dimensional models is to use MCMC. Having scalability in mind we chose to implement ensemble sampling methods in which many MCMC chains sample the parameter space in parallel. The chains interact and use the information from their spatial distribution to effectively adjust their next jump proposals. This has the added benefit of being able to sample the unknown likelihood surfaces efficiently and with minimal user input. We have implemented two different ensemble sampling methods, namely, an affine-invariant method and a differential evolution method; we discuss the latter in the next subsection.

The affine-invariant method can sample likelihood functions that are related by affine transformations with the same efficiency [127]. This means it is very efficient in sampling highly stretched likelihood distributions as long as the nonlinear correlations among parameters are sufficiently weak.

MCMC algorithms are generally not very efficient on highly multimodal distributions. In order to overcome this problem we have implemented parallel tempering for each MCMC sampler. Parallel tempering makes copies of the log-likelihood (\mathcal{L}) function that are made smoother through the introduction of a temperature parameter, the higher the temperature the smoother the likelihood surface:

$$\mathcal{L}_i \propto \frac{\mathcal{L}}{T_i} \tag{5.24}$$

The different temperatures are chosen from a temperature ladder such that $1 \leq T_i \leq T_{max}$. Then we run a copy of our MCMC sampler for each tempered likelihood copy in parallel. The highest temperature chains can freely move in the parameter space, while the low temperature chains can be trapped in local likelihood maxima. By allowing the different temperature chains to exchange their positions with some prescription we let the low temperature chains to get out of the local maxima and explore the entire parameter space. In the end the lowest temperature chain, which samples the original untempered likelihood, yields the posterior probability distribution.

In order to get an efficient parallel tempering algorithm the temperature ladder has to be chosen carefully. There are two main factors to consider. First, The highest temperature used should be large enough to let the chains move freely within the likelihood surface. Furthermore the temperatures should not be too widely spaced as that could hinder efficient swaps between chains from adjacent temperatures and lead to inefficient tempering. The choice of an efficient temperature ladder depends on the likelihood surface and could be difficult to guess. To mitigate this problem we have implemented a method to adaptively change the temperatures in order to get near optimal efficiency. Our method follows that of [279].

We have implemented parallelism in different levels of the MCMC sampler. Tempering is parallelized, MCMC chains at each tempering level run in parallel and the likelihood can itself

use multiple threads to run. This allows for the effective use of large high performance computing machines.

It is a well-known feature of MCMC schemes that there is an initial so-called “burn-in” phase when the sampling exhibits comparatively large changes in parameter predictions, followed by a phase where the MCMC chains settle down to smaller and more consistent changes, before ultimately converging to a final answer. As is customary in this approach, we exclude a certain number of MCMC steps corresponding to the “burn-in” process at the beginning of the “chains”, i.e. the history of an MCMC walker.

Parallel-Tempered, Differential-Evolution Markov Chain Monte Carlo

The second MCMC algorithm implemented in THEMIS is the parallel-tempered, differential-evolution algorithm [36]. The differential evolution method adjusts the collective move of its chains in a way to achieve an optimal acceptance rate during Monte Carlo steps. It has the added benefit of being able to jump between modes in multimodal problems even without any tempering, thus representing a better option for multimodal distributions. Our implementation follows that of ([198]) and it also makes use of the same parallel-tempering algorithm as the affine invariant method.

Bayesian Evidence

The MCMC sampling described above provides the posterior probability distribution on parameters within the context of a given model. This allows us to calculate expectation values for any quantity of interest and to assess the goodness of fit for a given model to the data. However, if we need to compare the plausibility of different models given the same data set, running conventional MCMC is not enough.

There are different ways of performing the model comparison, this includes the reversible jump MCMC, calculating the Bayesian evidence (via thermodynamic integration, nested sampling or Laplace approximation), and information criteria [160].

In Bayesian probability theory the relative probability of two models given the same data set is related to the ratio of the Bayesian evidences for the two models which is known as the Bayes factor or the odds ratio. The relative posterior probability of the two models can be written as:

$$\frac{P(M_1|D)}{P(M_2|D)} = \frac{P(D|M_1) P(M_1)}{P(D|M_2) P(M_2)} \quad (5.25)$$

In this equation M_1 and M_2 are the two models we wish to compare and D represents the data used to make the comparison. $P(D|M_1)$ and $P(D|M_2)$ are the Bayesian evidence for the two models. One can often assume that the prior probability of the two models, $P(M_1)$ and $P(M_2)$, are equal, and hence the ratio of the Bayesian evidence or the Bayes factor is all one needs to calculate.

Themis implements the thermodynamics integration method to calculate the Bayesian evidence [164]. In order to do this many MCMC chains are run in parallel on tempered versions of the likelihood. The temperature ladder for this purpose is provided by the user. Given the posterior distributions and the values of the likelihood at these points, the Bayesian evidence (Z) is obtained from

$$\ln Z = \int_0^1 E_\beta(\mathcal{L}) d\beta, \quad (5.26)$$

Where \mathcal{L} is the log-likelihood, $\beta = 1/T$, and $E_\beta(\mathcal{L})$ is the expectation value of the log-likelihood calculated using the posterior probability distribution of chains at a tempered level corresponding to $T = 1/\beta$.

5.6 Validation Tests

We now turn to validating THEMIS. For this we focus on the sampling methods, for which the implementations are novel, and reproducing prior analyses of EHT observations of Sgr A*. The variety in algorithmic improvements present in THEMIS result in a considerable speed-up and simplicity in implementation in comparison to the previous work to which we compare — often analyses that took many months are now executed in days. Moreover, all of these tests have been integrated into THEMIS, both as validation tools and tutorials for future users.

5.6.1 Validation of the samplers

Here we test the sampling part of the code thoroughly. In particular, we demonstrate the ability of the affine sampler to reliably probe non-trivial parameter spaces. That is, we anticipate that the complex models ultimately of most interest in the context of EHT analyses will produce multimodal probability distributions in high-dimensional parameter spaces. It will be necessary, therefore, to consistently identify all of the high-likelihood islands and determine accurately their relative posterior probabilities.




Figure 5.4: Two dimensional Gaussian likelihood sampled by the affine invariant MCMC sampler (blue), the differential evolution MCMC sampler (red) and the grid search sampler (black).

Two dimensional Gaussian likelihood

If our model has a small number of parameters the grid search sampler can be efficient in sampling the parameter space. In this test the grid search sampler was used to sample a two dimensional symmetric Gaussian likelihood. Figure 5.4 shows the log likelihood recovered using the grid search sampler as well as the marginalized posterior distributions for the same likelihood sampled using MCMC methods.

Egg box test

In this test a 5 dimensional parameter space with a highly multimodal egg box like distribution is sampled. The likelihood is described in Section 5.5.1 and contains $5^5 = 3125$ sharp peaks within the prior range: $p_i \in [-8, 8]$ for all i .

This presents a significant challenge to most sampling schemes. The narrowness of the peaks and the dimensionality of the parameter space precludes a grid search, which would require more than 3×10^{12} samples to robustly detect all of them. The large dynamic range in the likelihood, i.e., the very low likelihoods between peaks, precludes typical MCMC schemes, which are unable to efficiently explore the full parameter space. Therefore, it provides a strong test of the ability of the parallel tempered, affine-invariant and differential-evolution parallel MCMC samplers to efficiently find and reconstruct the various high-likelihood regions. As seen in Figure 5.5, both capture all of the features of the Egg Box likelihood. This run was executed employing the differential evolution sampler with 5 tempering levels and took only 1 min on a typical laptop. There were 100 walkers running for 8000 steps and the burn-in period was less than 1000 steps.

16-Gaussian test

Here we show that the sampling scheme can correctly probe a two dimensional parameter space with a likelihood consisting of 16 Gaussians, and accurately reconstruct the relative posterior probabilities of different peaks. To do this, the likelihood of the Gaussian located at $(x_0, x_1) = (20, 10)$, is chosen to be nine times higher than the others. This test was run using the affine invariant sampler with 4 tempering levels. The number of MCMC steps in this case was 4000 steps and the sampler used 100 walkers.

As shown in Figure 5.6, the sampler finds all of the Gaussian components. In addition, it recovers the non-uniformity of the likelihood surface, accurately reconstructing the posterior weight of the appropriate component. Figure 5.7 shows the relative probability mass correctly recovered for these 16 gaussian peaks.

5.6.2 Self-tests with simulated data

Here, we demonstrate the ability of THEMIS to accurately reconstruct model parameters. This presents a simultaneous test of many of the components of THEMIS, including the data structures, models, likelihoods, and samplers. We generate simulated images using THEMIS' native model classes, from which the appropriate simulated data is constructed. Thermal noise is then included, producing a data set similar in character to that associated with a single night of the 2017 EHT campaign. The simulated data is then analyzed with THEMIS using the corresponding model. Note that this does not address model discrimination. For this purpose, we considered three models: the symmetric gaussian and the geometric crescent models, for which visibilities can be computed analytically, and the SED-fitted RIAF model, which incorporates the ray-tracing components of VRT² and numerical data generation of THEMIS.




Figure 5.5: Validation of the sampler with a five dimensional “eggbox”-likelihood test with the five artificial parameters $x_0, x_1, x_2, x_3,$ and x_4 . All peaks in the likelihood surface are successfully recovered.

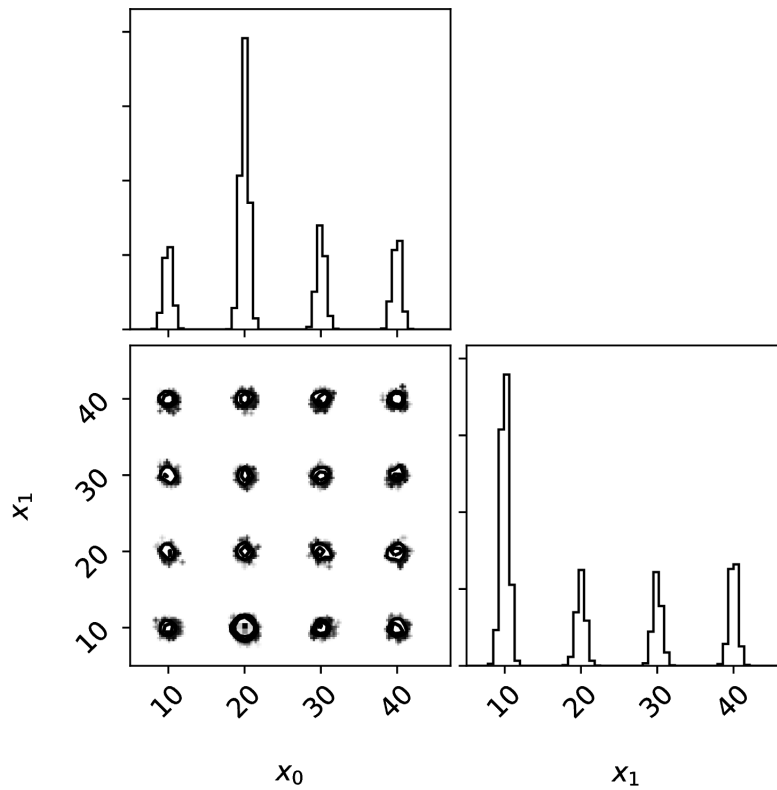


Figure 5.6: Validation of the sampler with a two dimensional test involving a 16-Gaussians-likelihood and two artificial parameters x_0 and x_1 . Note that the sampler not only probes the likelihood surfaces comprehensively, but it also correctly retrieves the Gaussian with the higher likelihood at $x_0 = 20, x_1 = 10$.




Figure 5.7: Relative posterior probability mass for the 16 gaussian peaks recovered via MCMC sampling.

Gaussian Model

We generated simulated closure amplitude data from a compact, symmetric Gaussian with $V_0 = 2.5$ Jy and $\sigma = 5 \mu\text{as}$. We adopted a Gaussian for simplicity. The very compact size was selected to ensure high SNR detections on even the longest baselines of the 2017 EHT campaign; such high SNRs are typical of more complex models. We analyze this data with THEMIS’ symmetric Gaussian model (Section 5.4.2) to assess potential biases associated with the non-Gaussian nature of the closure amplitude error distribution. We imposed an SNR minimum on the simulated closure amplitude of 4.

For this analysis we used the closure amplitude likelihood described in Section 5.5.1. We sampled the posterior distribution with the parallel-tempered affine-invariant MCMC sampler, adopting linear priors on each model parameter. The analysis converged using 5 tempering levels with 128 walkers communicating every 50 MCMC steps, and taking 100 samples per walker.

As expected, the total intensity is not constrained by closure amplitudes, recovering our prior distribution. The resulting posterior distribution for the size of the Gaussian is shown in Figure 5.8. The reconstructed size is $\sigma = 5.0004 \pm 0.0004 \mu\text{as}$ ⁶, consistent with the input value.

⁶Note that for all of our analysis we report the marginalized values for the parameters of each model instead of




Figure 5.8: Posterior distribution of the size of the symmetric Gaussian reconstructed from simulated closure amplitude data with $\sigma_0 = 5 \mu\text{as}$. The expected value is indicated by the vertical blue line.

Repeating the analysis with different realizations of the simulated data produces qualitatively similar results, though they do exhibit 2σ fluctuations marginally more often than anticipated. No experiment produced a deviation larger than 3σ . As a result, we conclude that the likelihood in Equation (5.15) does not fully eliminate the bias inherent in the closure amplitude error distribution, though does so at the 2σ level. Decreasing the SNR minimum increases this bias substantially, suggesting that additional development is required to fully exploit low-SNR data.

While the closure amplitude likelihood in Equation (5.15) is not Gaussian, and thus does not admit a well-defined χ^2 , we do construct an approximate expression via $\chi^2 = -2\mathcal{L}$. In the limit of small closure amplitudes, this identification is well-motivated. The associated reduced- χ^2 is 0.97 with 1656 degrees of freedom, suggesting that this statistic will be informative of fit quality.

Crescent Model

We generated simulated visibility amplitude and closure phase data from a diffractively-scattered crescent image with $V_0 = 2.24 \text{ Jy}$, $R = 28 \mu\text{as}$, $\psi = 0.14$, $\tau = 0.07$, and $\xi = 0^\circ$, and added thermal noise to it. We then analyze this data with THEMIS' crescent model (Section 5.4.2), demonstrating that THEMIS properly recovers the parameters of the original image.

the maximum values, which for complicated likelihood distributions may differ significantly.

For the analysis we used the standard visibility amplitude and closure phase likelihoods described in Sections 5.5.1 and 5.5.1, and modeled the effects of diffractive scattering with the default scattering model implemented in THEMIS, and described in Section 5.4.3.

In this case we sampled the posterior distribution with the parallel-tempered differential evolution MCMC sampler adopting linear priors on each parameter of the model. The analysis converged using 4 tempering levels with 16 walkers per level communicating every 50 MCMC steps, and taking 10000 samples per walker. The resulting posterior distributions for the parameters of this model are shown in Figure 5.9 where the blue lines represent the true parameter values of the original image.

Our analysis shows that the marginalized values for the parameters of the model are $V_0 = 2.2399 \pm 0.0001$ Jy for the total flux, $R = 28.0064 \pm 0.0054$ μ as for the overall size of the crescent, with a relative thickness $\psi = 0.1404 \pm 0.0003$, an asymmetry parameter $\tau = 0.0691 \pm 0.0005$, and a position angle $\xi = 0.040^\circ \pm 0.023^\circ$. Individually, these are consistent at the 2σ -level with the true values of the original crescent image. The model gives a satisfactory fit to the data as confirmed by the reduced- χ^2 of 0.9813 with 1670 degrees of freedom, which implies that high-quality fits exist.

RIAF Model

We generated visibility amplitude and closure phase data from a diffractively-scattered RIAF image with $(a, \theta, \xi) = (0.10, 60^\circ, 0^\circ)$. We added thermal noise to the simulated data and then analyze it with THEMIS' SED-fitted RIAF model (Section 5.4.4) to show that THEMIS can properly recover the parameters of the original image.

For the analysis the standard visibility amplitude and closure phase likelihoods (Sections 5.5.1 and 5.5.1) were used, and the effects of diffractive scattering were modeled using THEMIS' default scattering model (Section 5.4.3). We used the parallel-tempered differential evolution MCMC sampler with 3 tempering levels, 14 walkers per level communicating every 50 MCMC steps, and took 5000 samples per walker.

The posterior distributions for the parameters of the model are shown in Figure 5.10. We find that the marginalized values for the black hole spin parameters are $a = 0.0997^{+0.0006}_{-0.0007}$, $\theta = 59.9983^{+0.0155^\circ}_{-0.0167^\circ}$, and $\xi = 0.0017^{+0.0223^\circ}_{-0.0199^\circ}$. These parameter estimates are consistent at the 1σ -level with the true values of the original RIAF image. In this case we find a reduced χ^2 of 0.9868 with 1664 degrees of freedom, indicating that high-quality fits were found.

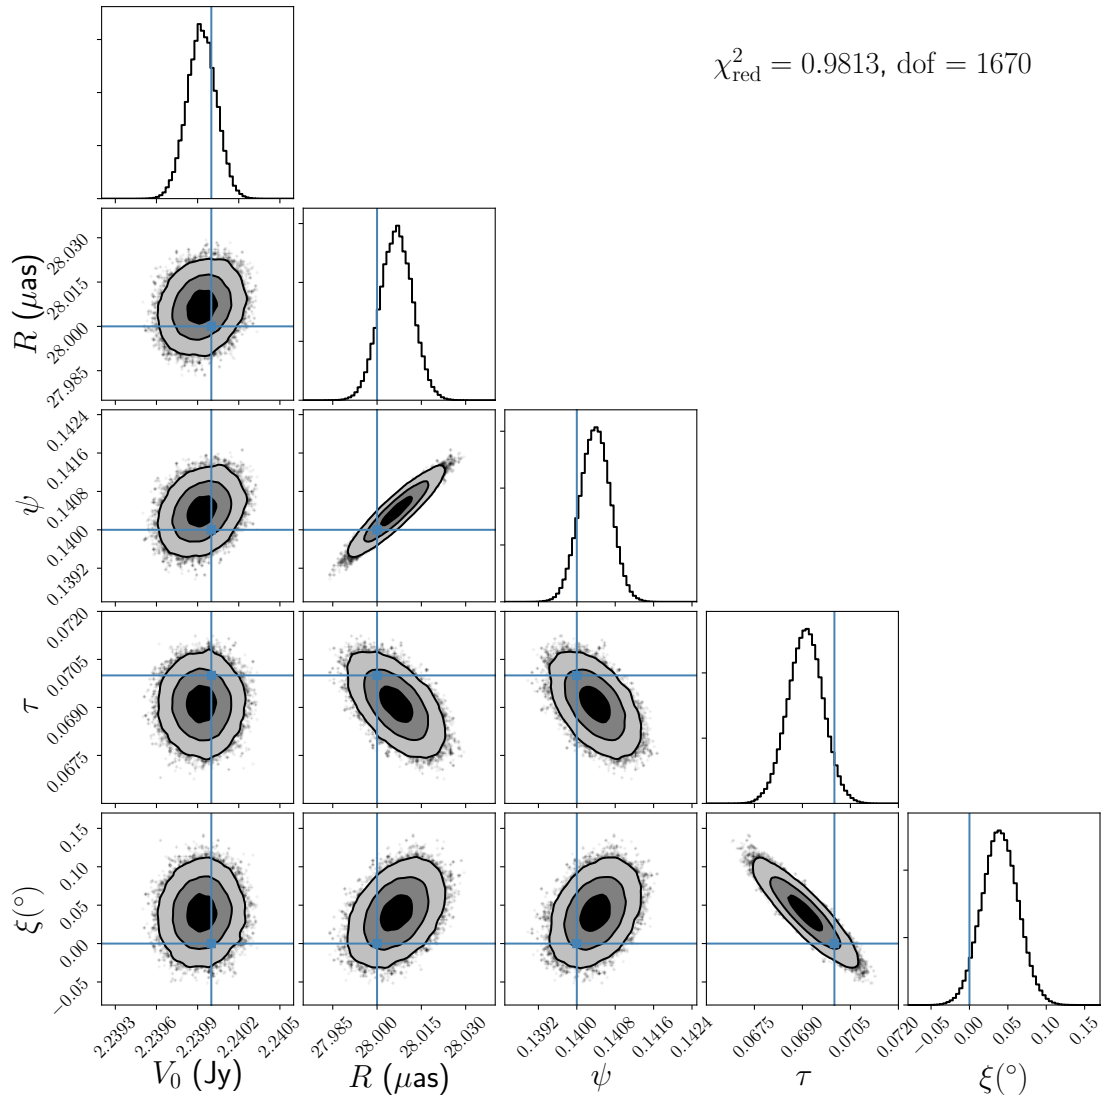


Figure 5.9: Joint parameter distributions from the crescent model analysis of the simulated visibility amplitude and closure phase data generated from a crescent image with $V_0 = 2.24$ Jy, $R = 28 \mu\text{as}$, $\psi = 0.07$, $\tau = 0.14$, and $\xi = 0$ rad. Here the true parameter values are represented by the blue lines, and the contours show the 1, 2 and 3 σ levels of the sampled posterior distribution.




Figure 5.10: Joint parameter distributions for the SED-fitted RIAF model analysis of the simulated visibility amplitude and closure phase data generated for a RIAF image with $(a, \cos\theta, \xi) = (0.1, 0.5, 0.0)$. Here the true parameter values are represented by the blue lines, and the contours show the 1, 2 and 3- σ of the sampled posterior distribution.

5.6.3 Reproducing Previous Results

The variety of published analyses of EHT observations of Sgr A* provides a natural validation test of THEMIS, as well as a demonstration of its flexibility. These include comparisons of purely phenomenological and physically motivated models of the image structure. In constructing these, we make use of the published EHT data sets listed in bold in Table 5.1, consisting of visibility amplitudes measured in 2007 and 2009 and closure phases measured between 2009 and 2013, inclusively.

Symmetric Gaussian

We analyze the visibility amplitude data from 2007 and 2009 using the symmetric Gaussian model described in Section 5.4.2 in order to show that THEMIS can reproduce previous model fitting studies made to estimate the source size of Sgr A*.

For the analysis we employed the norm-marginalized visibility amplitude likelihood described in Section 5.5.1 to account for variations in the total flux of Sgr A* between observation nights. The effects of diffractive scattering were modeled using THEMIS' default scattering model (Section 5.4.3).

We employed the parallel-tempered affine Invariant sampler with 4 tempering levels with 32 walkers per level, and adopted linear priors on each parameter of this model. The posterior distribution for the intrinsic size of Sgr A* after 10000 MCMC iterations is shown in Figure 5.11. The reconstructed size is $\sigma = 15.73 \pm 0.25 \mu\text{as}$ (FWHM = $37.05 \pm 0.60 \mu\text{as}$). The model gives a satisfactory fit to the data with an associated reduced- $\chi^2 = 1.15$ with 65 degrees of freedom. These results are in good agreement with the best fits found in [40] when all epochs are combined, and the inferred sizes for each night reported in [82] and [104].

Asymmetric Gaussian

We also analyze the visibility amplitude data from 2007 and 2009 with the asymmetric Gaussian model described in Section 5.4.2 to show that THEMIS can reproduce previous studies made to probe the asymmetry of the emitting region of Sgr A*.

For this analysis we also analytically marginalize variations in the total flux of Sgr A* between observation nights using the norm-marginalized visibility amplitude likelihood described in Section 5.5.1, and modeled the effects of diffractive scattering using THEMIS' default scattering model (Section 5.4.3).




Figure 5.11: Gaussian size distribution from the analysis of the symmetric Gaussian model to the 2007 and 2009 visibility amplitude data of Sgr A*.

We employed the Affine Invariant sampler and adopted linear priors on each parameter of this model. The results converge using 4 tempering levels, with 32 walkers per level, and taking 20000 samples per walker. The posterior distributions for the different parameters of the model are shown in Figure 5.12. Our analysis shows that the marginalized values for the parameters of the model are $\sigma = 19.07^{+1.07}_{-2.51} \mu\text{as}$, $A = 0.54^{+0.13}_{-0.18}$, and $\xi = -64.3^{\circ+17.0^{\circ}}_{-4.7^{\circ}}$, and $\xi = 114.7^{\circ+19.2^{\circ}}_{-5.1^{\circ}}$.

This model also gives a satisfactory fit to the data with an associated reduced- $\chi^2 = 0.75$ with 63 degrees of freedom. These results are in good agreement with the best fits found in [40] when all epochs are combined.

Crescent Model

We analyze the visibility amplitude data from 2007 and 2009 with the crescent model outlined in Section 5.4.2 in order to show that THEMIS can reproduce the earlier findings reported by [155].

We proceeded in a similar fashion to the analysis performed with the Gaussian models employing the norm-marginalized visibility amplitude likelihood described in Section 5.5.1 to account for variations in the total flux of Sgr A* between days and modeling the effects of diffractive scattering with the default scattering model implemented in THEMIS (Section 5.4.3).

In this case we sampled posterior distribution with the differential evolution MCMC sampler adopting linear priors on each parameter of the model. We used the reported values in Table 1

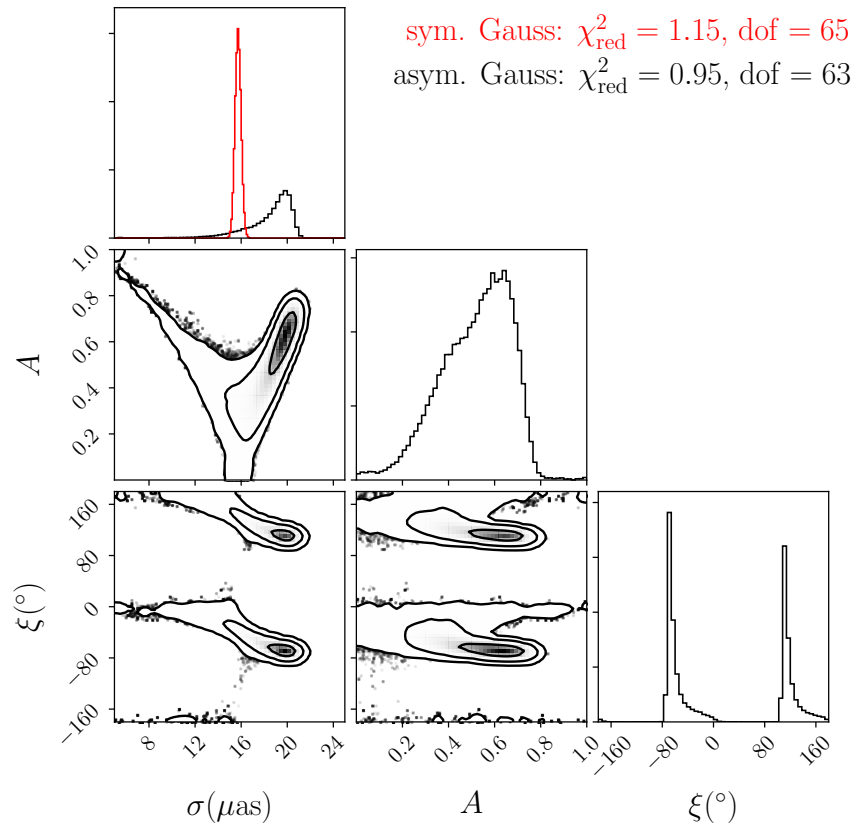


Figure 5.12: Joint parameter distributions from the analysis of the asymmetric Gaussian model to the 2007 and 2009 visibility amplitude data of Sgr A*. The gray contours show the 1σ , 2σ , and 3σ confidence regions for the size, the asymmetry parameter and the position angle. For reference, the symmetric Gaussian size distribution from Figure 5.11 is shown in red in the top left panel. These parameter distributions are consistent with the results of [40].

of [155] as initial guesses for the values of the parameters of this model. The analysis converged using 4 tempering levels with 32 walkers per level, and taking 20,000 samples per walker. The resulting posterior distributions for the parameters of this model are shown in Figure 5.13. The marginalized values for the parameters of the model are $R = 29.8_{-3.2}^{+4.8} \mu\text{as}$, $\psi = 0.28_{-0.17}^{+0.12}$, and $\tau = 0.20_{-0.13}^{+0.15}$. In this case the analysis finds two values for the position angle $\xi = 60.5_{-12.4}^{+8.4} \circ$, $\xi = -119.4_{-13.7}^{+8.7} \circ$, and a minimum reduced- $\chi^2 = 0.76$ with 62 degrees of freedom

RIAF Model: Visibility Amplitude Analysis

We now turn to the first example of a physical model. First, we demonstrate THEMIS' ability to reproduce the analysis published in [40]. For that purpose, we analyze the visibility amplitude data of Sgr A* from 2007 and 2009 with the SED-fitted RIAF model described in Section 5.4.4, using a tabulated set of accretion flow parameters obtained at different black hole spins and inclinations — and distributed with THEMIS— that reproduce the observed SED of Sgr A*.

For this analysis we employed a set of linear priors for each parameter of the model and the norm-marginalized visibility amplitude likelihood described in Section 5.5.1 to account for variations in the total flux of Sgr A* between observation nights. The effects of diffractive scattering were modeled using THEMIS' default scattering model (Section 5.4.3).

We used the parallel-tempered differential evolution MCMC sampler with 5 tempering levels and 16 walkers per level communicating every 50 MCMC steps. The test completed 5000 MCMC iterations and the posterior distributions for the black hole spin parameters is shown in Figure 5.14. We find a spin $a = 0.25_{-0.19}^{+0.35}$, while the inclination angle has two values located at $\theta = -61.06_{-10.02}^{+8.77} \circ$, and $\theta = 62.31_{-9.33}^{+10.54} \circ$. The model has a minimum reduced- $\chi^2 = 0.80$ with 63 degrees of freedom. These results are in good agreement with the best fits found in [40] when all epochs are combined.

RIAF Model: Visibility Amplitude and Closure Phase Analysis

We analyze the visibility amplitude and closure phase data sets that are bolded in Table 5.1 with the SED-fitted RIAF model described in Section 5.4.4 using 128x128 pixel RIAF images, to show that THEMIS successfully reproduces the results of the analysis published by [41].

For this analysis we employed the norm-marginalized visibility amplitude likelihood described in Section 5.5.1 to account for variations in the total flux of Sgr A* between observation nights. We also used the shift-marginalized closure phase likelihood 5.5.1 to model the effects

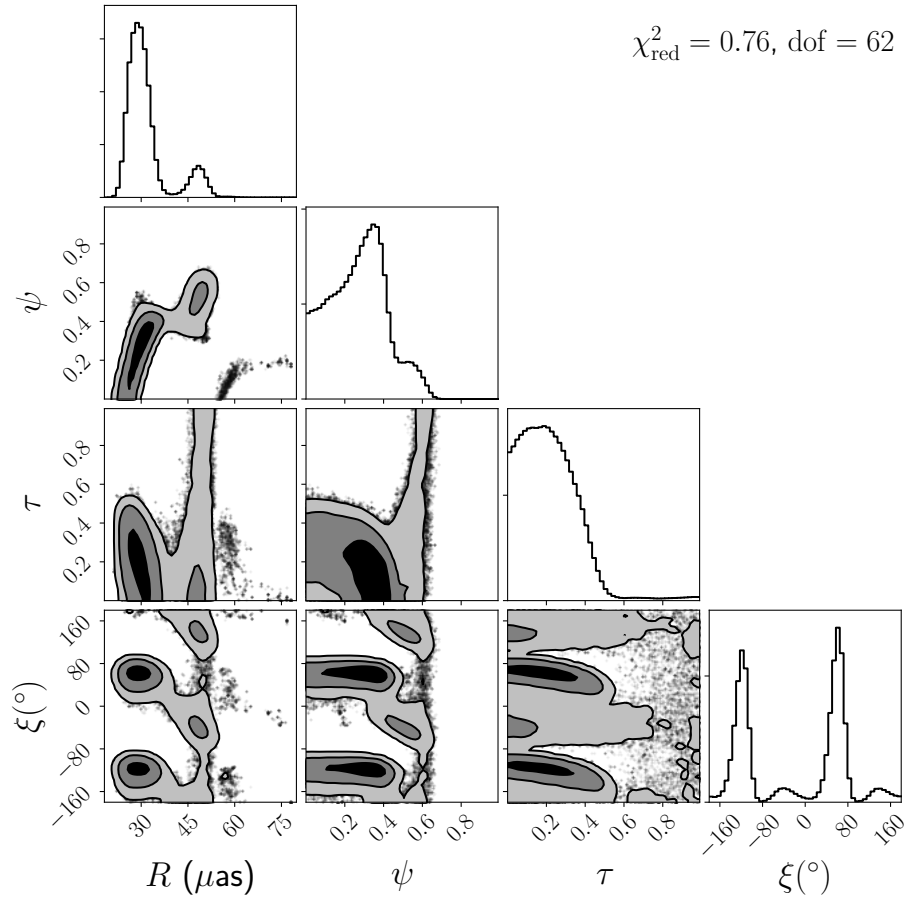


Figure 5.13: Joint parameter distributions from the analysis of the crescent model to the 2007 and 2009 visibility amplitude data of Sgr A*. The gray contours show the 1σ , 2σ , and 3σ confidence regions for the overall radius, the relative thickness, the degree of symmetry and the position angle of the crescent. These parameter distributions are consistent with the results of [155].




Figure 5.14: Joint parameter distributions from the analysis of the SED-fitted RIAF model to the 2007 and 2009 visibility amplitude data of Sgr A*. The gray contours show the 1σ , 2σ , and 3σ confidence regions for the spin magnitude, a , the cosine of the inclination, $\cos\theta$, and the position angle, ξ . These parameter constraints are consistent with the results of [40].

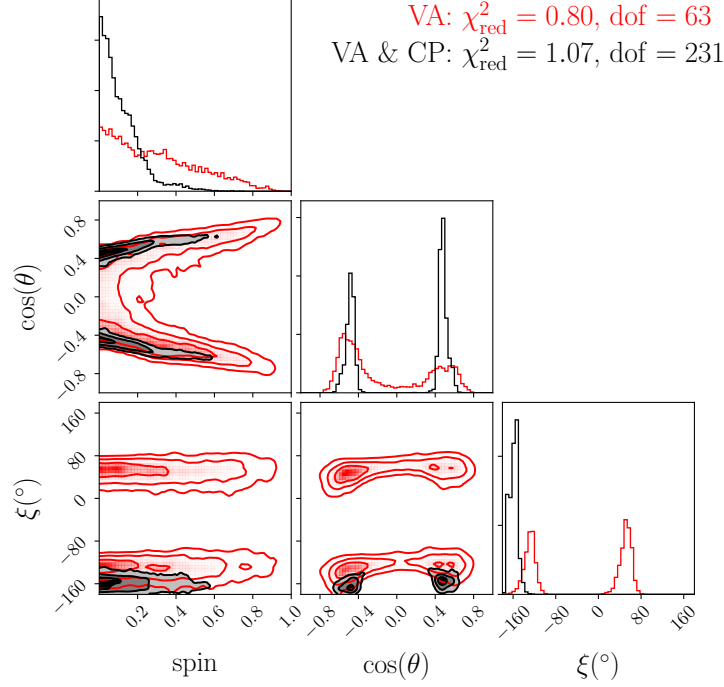


Figure 5.15: Joint posterior parameter distributions for the SED-fitted RIAF model implied by the combined visibility amplitude and closure phase data sets bolded in Table 5.1. For reference, the posteriors implied by the visibility amplitude data alone are shown in red. The contours show the 1σ , 2σ , and 3σ confidence regions for the spin magnitude, a , the cosine of the inclination, $\cos\theta$, and the position angle, ξ . All parameter constraints are consistent with the results of [41].

of refractive scattering, while the effects of diffractive scattering were modeled using THEMIS’ default scattering model (Section 5.4.3).

We used the parallel-tempered differential evolution MCMC sampler with 5 tempering levels and 16 walkers per level communicating every 50 MCMC steps. The MCMC chain was run for 8000 steps and the resulting posterior distributions for the parameters of this model are shown in Figure 5.15 in comparison to the results of analysis with visibility amplitude data only discussed in the previous section. We find that the black hole spin parameters are similarly constrained after the inclusion of the closure phase data, with $a = 0.09^{+0.11}_{-0.07}$, $\theta = -61.00^{+2.74}_{-2.12}$, and $\theta = 61.55^{+1.97}_{-3.25}$, and $\xi = -165.13^{+7.42}_{-4.62}$. In this case we find a reduced χ^2 of 1.07 with 231 degrees of freedom, indicating that high-quality fits were found. These results are consistent with best fit parameters reported by [41].

5.7 New Results

In this section we present first novel results obtained with THEMIS. These make use of the ability of THEMIS to rapidly generate new data comparisons and/or combine data sets in a uniform manner. Additional results obtained by combining a more complete combination of the data sets in Table 5.1 and applying additional model features will be reported elsewhere.

5.7.1 Crescent model

The extensible nature of THEMIS makes the extension of the analysis of the [155] crescent model to include additional data trivial. We demonstrate this by including the closure phase data sets that are bolded in Table 5.1. To account for refractive scattering, we employ the shift-marginalized closure phase likelihood (Section 5.5.1) when including the contribution to the total likelihood from closure phases. In all other respects, the analysis is similar to that presented in Section 5.6.3.

The inclusion of closure phase data places strong new constraints on the crescent structure in a number of respects. The resulting posterior distributions for the parameters of this model are shown in Figure 5.16. The constraints on all of the crescent parameters are substantially improved quantitatively, often settling ambiguities in the previous analysis. The crescent overall size has been restricted to $R = 46.3_{-1.5}^{+1.4} \mu\text{as}$; the relative thickness parameter is now $\psi = 0.41_{-0.04}^{+0.07}$; and asymmetry parameter is $\tau = 0.23_{-0.14}^{+0.21}$. Individually, these are consistent at the 2σ -level with the expectation based on visibility amplitudes alone.

Similarly, the position angle is also strongly constrained, with the prior degeneracy eliminated, finding $\xi = 179.4_{-9.1^\circ}^{+19.2^\circ}$. Unlike the other parameters, this is inconsistent with the estimates from the visibility amplitudes alone at the 2σ -level. This is apparent in the bottom panels of Figure 5.16. This is modestly disconcerting given the qualitatively distinct natures of the closure phases and visibility amplitudes. Nevertheless, the reduced- $\chi^2 = 1.01$, implies that high-quality fits exist.

5.7.2 Extended RIAF model

The SED-fitted RIAF model treats the comparisons to the EHT data and flux measurements differently, which utilize a set of prior set of SED fits. Again, the extensibility of THEMIS enables relaxing this procedure, and comparing simultaneously to the flux and mm-VLBI measurements.

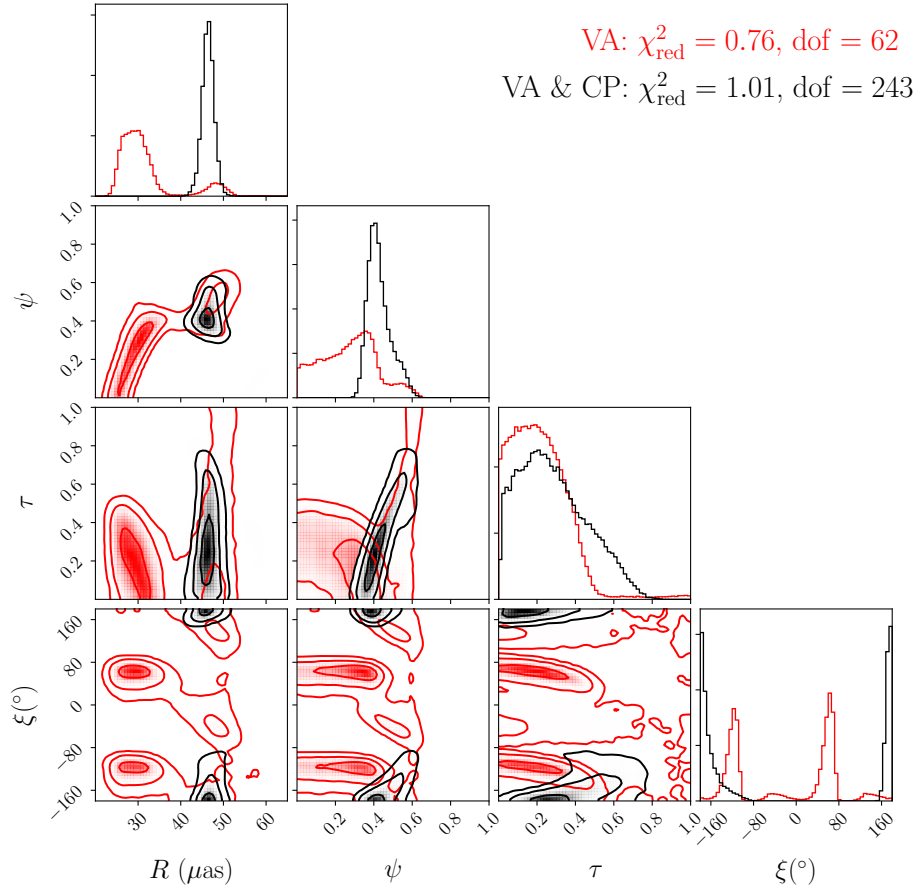


Figure 5.16: Joint posterior parameter distributions for the crescent model implied by the combined visibility amplitude and closure phase data sets bolded in Table 5.1. For reference, the posterior distributions implied by the visibility amplitude data alone from Figure 5.13 are shown in red. The gray contours show the 1σ , 2σ , and 3σ confidence regions for the overall radius, the relative thickness, the degree of symmetry and the position angle of the crescent.

In principle, this may broaden the black hole parameter estimates, trading worse SED fits for better structural fits. To explore this we performed a new analysis, similar in spirit to that presented in Section 5.6.3, in which we analyze both data sets concurrently.

We performed a new analysis using the extended RIAF model described in Section 5.4.4, which generates flux measurements in addition to the mm-VLBI observations in a fashion identical to the SED-fitted RIAF model. In addition to the parameters describing the black hole spin (magnitude, inclination, and position angle), three additional parameters were introduced, describing the normalizations of the densities ($n_{e,t}$, $n_{e,nt}$) and temperature (T_e) of the emitting electron population; all remaining parameters were held fixed at the values employed in the SED-fitted RIAF model: $\eta_t = -1.1$, $\eta_{nt} = -2.02$, $\tau_t = -0.84$, $h_t = h_{nt} = 1.0$, $\alpha = 1.25$, $\gamma_{\min} = 100$, $\beta = 10$, and $\kappa = 0$. This model was compared to the flux and mm-VLBI data bolded in Table 5.1. For this run the affine invariant sampler with 4 tempering levels was used. There were 120 walkers used and the MCMC chain was run for 3300 steps.

In Figure 5.17, the resulting set of parameter constraints are presented in comparison to the prior analyses described in Sections 5.6.3 and 5.6.3. In all cases, the spacetime parameters are consistent with those found previously. Including the flux data produces a marginally stronger constraint on the black hole spin, $a = 0.1^{+0.19}_{-0.08}$, arising from the systematic decrease in the quality of the SED fits at higher spins (that was ignored in the prior analyses). Nevertheless, as anticipated, the inclination constraints are broadened, permitting $\theta = 62.2^{+5.3}_{-4.6}$ and $\theta = -62.8^{+5.5}_{-6.2}$.

5.8 Code Performance

As seen in many of the validation tests and example analyses presented in Sections 5.6 and 5.7, even for models with modest numbers of parameters, it is typical for the posterior probability distributions to be multimodal. As the models increase in sophistication, introducing additional physical freedoms and addressing various systematic uncertainties, this problem will be compounded by the need to explore high-dimensional parameters spaces. This is further complicated by the computational expense of numerically generating images of realistic astrophysical models. As a result, THEMIS has been designed to exploit the proliferation of modern HPC systems. Here we discuss the ways in which this has been, and may be, implemented, along with a description of THEMIS’s scaling efficiency, demonstrating that it can run efficiently on very large machines.

THEMIS explicitly supports parallelization via MPI, and implicitly via OpenMP and CUDA. MPI parallelization has been implemented at a number of levels, including the samplers, likelihood evaluation, and model generation, permitting users maximum flexibility in distributing the computational workload of an analysis.

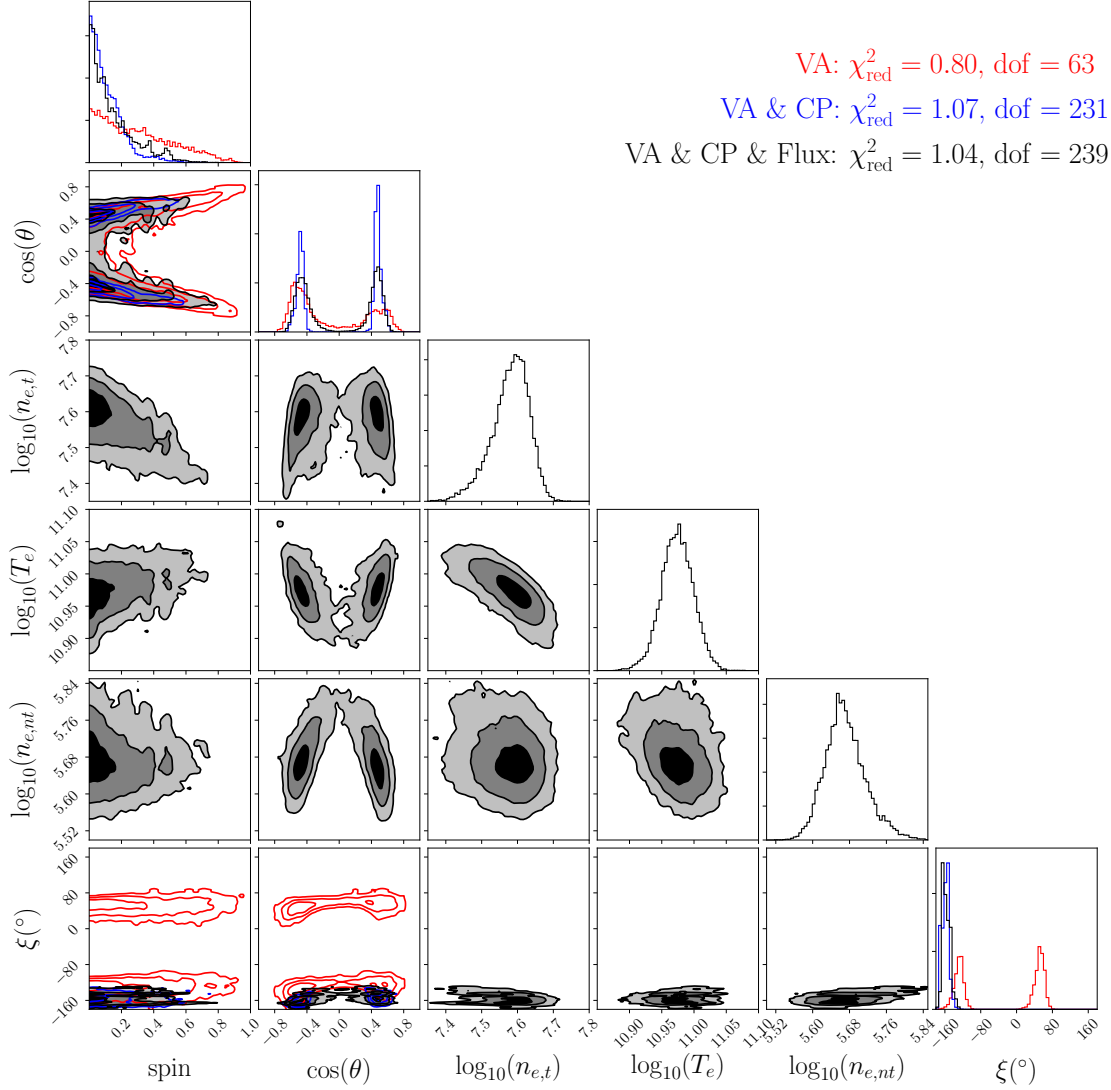


Figure 5.17: Joint posterior parameter distributions for the SED-fitted RIAF model implied by the combined flux, visibility amplitude and closure phase data sets bolded in Table 5.1. For reference, the posteriors implied by the visibility amplitude data alone and by the combined visibility amplitude and closure phase data sets from Figures 5.14 and 5.15 are shown in red and blue respectively. The contours show the 1σ , 2σ , and 3σ confidence regions for the spin magnitude, a , the cosine of the inclination, $\cos\theta$, and the position angle, ξ . All parameter constraints are consistent with the results of [41].

Both, the Parallel Tempered Affine Invariant and Differential Evolution MCMC sampling algorithms are designed to exploit parallelization in two levels. First, the use of parallel tempering levels may be further parallelized by assigning separate tempering levels to different collections of processors. Second, the use of ensemble methods may be trivially parallelized among the individual walkers. Our implementation of the ensemble sampler evolves half of the walkers simultaneously while using the other non-evolving half to determine the next proposed jump. Each walker in the “active” set can be evolved on a separate CPU. Upon completion, the “active” and “passive” sets swap, and the process is repeated. The result is a set of samplers that can immediately utilize $N_T N_W / 2$ processors, where N_T is the number of tempering levels and N_W is the number of walkers, typically many times the number of parameters.

Image generation is an intrinsically parallelizable task. The VRT² library already natively supports MPI parallelization and vectorization via OpenMP. On modern Xeon based systems, VRT² can efficiently use $N_L = 32$ cores to produce 128x128 pixel images before ancillary memory and communication costs become significant. Odyssey employs GPUs via CUDA, and provides an example of mixed MPI/GPU support within THEMIS. The performance of mixed MPI/GPU computation depends mainly on the number and specifications of the GPU cards, and is less sensitive to the number of MPI cores. Users implementing new THEMIS models are provided an MPI communicator and are only responsible for determining if and how parallelization should be implemented in their instance; they will be able to trivially exploit parallelization at the other levels.

Figure 5.18 shows the scaling of THEMIS on a representative sample problem with $N_T = 4$ tempering levels, $N_W = 16$ MCMC walkers, and $N_L = 32$ processors per likelihood evaluation. For this case, THEMIS scales with 94% efficient to 32, 88% efficient at 512 cores, and 84% at 1024 cores. Note that even modest increases in problem complexity involving larger images or higher-dimensional parameter spaces, require a larger set of walkers, tempering levels, and allow more processors per likelihood evaluation. Thus, the scaling efficiency of THEMIS will *improve* with problem size. Already, THEMIS can run efficiently on several thousand cores.

5.9 Summary

THEMIS provides a powerful new framework in which to develop and implement analyses of EHT observations. By focusing on the construction of interfaces, THEMIS enforces a modularity that facilitates rapid future development, ensuring flexibility and permitting extensibility. This flexibility is illustrated by the existing set of current THEMIS components, which span a wide variety of types of data, models, and sampling techniques. The clear definition of component

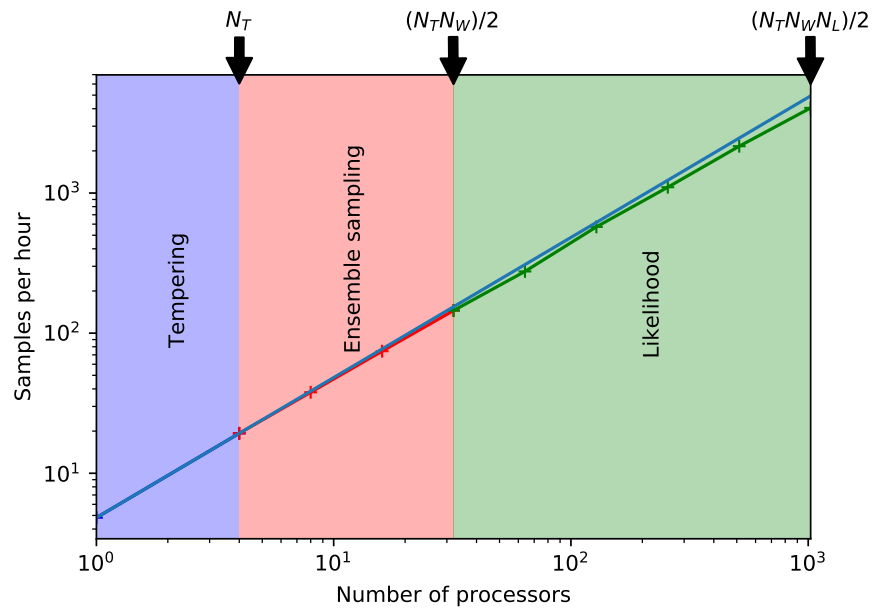


Figure 5.18: Scaling plot illustrating the multilayered parallelization strategy in THEMIS: The parallelization in tempering levels (perfect by construction) is shown in blue, the parallelization in different walkers is shown in red, the parallelization in likelihood evaluations (e.g. images here via calling VRT² is shown in green. The blue, red and green data points show the scaling of the code. The blue line acts as a reference for perfect linear scaling.

inputs and outputs enables future developers to rapidly contribute additional components (e.g., image models) without the need for a global understanding of the internal structure of the code.

Implemented data types include both mm-VLBI observables (visibility amplitudes, closure phases, closure amplitudes, polarization fractions) and ancillary data (fluxes). The ability to easily add accoutrements to these data objects, e.g., time stamps, observing stations, atmospheric conditions, observation resolution, etc., significantly increases their flexibility and the potential sophistication of subsequent analyses.

The generic nature of the model interface produces a correspondingly broad array of acceptable models, solving a key difficulty with unifying prior EHT analyses. As a result, THEMIS can construct analyses of phenomenological (e.g., gaussians) and physically motivated models (e.g., polarized images of synchrotron-emitting GRMHD simulations). It also naturally allows the inclusion of optional, additional, independent model features (e.g., interstellar scattering) in a uniform way. In principle, it can also facilitate in non-parametric modeling, e.g., image inversion, though this has yet to be implemented.

A number of likelihoods have been implemented, including likelihoods that analytically address nuisance parameters. This will become increasingly important as additional EHT systematics are considered, e.g., telescope gain corrections, refractive scattering, and intrinsic source variability. Similarly, a number of samplers have been implemented, including samplers that efficiently explore high-dimensional, multimodal likelihood surfaces.

A key feature of THEMIS is the ability to mix and match the above, constructing new analyses via minor changes in the model used, data included, and sampler used. This will be critical to evaluating the robustness of features, teasing apart subtle interactions in aspects of complex models, and systematically assessing the impact of additional types of data. At the same time, this permits rapid, distributed development: as features are added in the service of one analysis, e.g., a new sampler or a new scattering model, they may be rapidly deployed to others.

In anticipation of increasingly complex, physically motivated emission models for EHT targets, THEMIS enables the implementation of parallelization at multiple levels via multiple schemes. At present, this is implemented in a number of samplers via MPI and models via MPI, OpenMP, and CUDA. As a result, for typical analyses, THEMIS scales efficiently to thousands of cores, depending on problem complexity, and can effectively exploit modern HPC systems. For the implemented samplers, this parallel-performance scaling *improves* with problem complexity (i.e., number of parameters), partially mitigating the introduction of additional physical features.

Both the individual components of THEMIS and their integration have been extensively tested. THEMIS is able to accurately and consistently explore high-dimensional multimodal posterior probability distributions. It is able to recover the parameters of models used to construct realistic simulated EHT data for both geometric and physically motivated RIAF models. It has accurately

reproduced previous analyses of published EHT data. In the case of the RIAF models, it has done so in an order of magnitude less user time.

The extensibility of THEMIS is evident in the extension of these prior analyses. The [155] crescent model has been reassessed in light of the EHT closure phase measurements of Sgr A* published in [105]. The weak degeneracy in the size of the crescent is now broken, selecting $R = 46.3_{-1.5}^{+1.4} \mu\text{as}$. This is considerably larger than the size implied by the 1σ region obtained when only visibility amplitude are considered, though still consistent at 2σ . Nevertheless, high-quality fits of the combined closure phase and prior visibility amplitude data sets do exist. Note that this implies a crescent diameter that is nearly twice as large as the $55 \mu\text{as}$ anticipated for Sgr A* by identifying the crescent with the gravitationally lensed image of a geometrically thick accretion flow.

Where prior RIAF analyses have separated the fitting the SED and EHT data for Sgr A*, THEMIS now simplifies the process of fitting both simultaneously. While this may yield weaker parameter constraints in principle, in practice the black hole spin parameters are similarly constrained, with $a = 0.1_{-0.08}^{+0.19}$, $\theta = 62.2_{-4.6}^{+5.3}$, and $\theta = -62.8_{-6.2}^{+5.5}$; and $\xi = -158.1_{-10.4}^{+11.5^\circ}$. Future analyses that will systematically explore the relaxation of assumptions about the structure of the inner accretion flow will be published elsewhere.

THEMIS is meant to facilitate continuous, vigorous development. Already, plans are underway to implement schemes to correct individual station gains, address refractive scattering in the interstellar medium, model stochastic variability in the intrinsic emission region, introduce jet models, exploit GRMHD simulations, and perform non-parametric analyses. Future data type development will include polarized fluxes, Faraday rotation measurements, circular polarization, visibility variances. As a result, THEMIS is prepared to play fundamental role in the scientific exploitation of the new window on black hole physics being opened by the EHT.

Chapter 6

Conclusion

In this thesis we explored the power of gravitational lensing to measure the properties of various hard-to-detect astrophysical objects. We showed that these methods can substantially change our understanding of these objects. We explored how using the transient weak lensing effect of dark matter halos can be used to put the strongest observational limit on the dark matter power spectrum in the deep nonlinear regime. We found the constraint on the primordial (linear) scalar power spectrum to be $\mathcal{P}_{\mathcal{R}} < 3 \times 10^{-9}$ at $k_L \sim 3 \text{ pc}^{-1}$. Additionally our method was able to measure time delays between light curves of strongly lensed quasars which can be used to measure the Hubble constant. We then showed how gravitational microlensing can be used to understand the structure and dynamics of the Galaxy. Concretely, we studied the stellar mass function, contribution of the thick disk to the Galactic disk density, and the disk kinematics. The same gravitational microlensing effect can be used to detect stellar mass black holes at a much higher rate than what was possible before and transform our knowledge of these objects. Taking advantage of the very long baseline interferometers, we showed that the degeneracy present in the optical microlensing can be overcome and one can find the mass and velocity of the lens, as well as its distance. With modest improvements to the VLBA we can detect stellar mass black holes at a rate of roughly 10 events per year. Finally, the images of super-massive black hole shadows can be used to learn about the structure of space-time, accretion flows and astrophysical jets. We introduced a Bayesian framework for analyzing the data from the Event Horizon Telescope Collaboration. We then outlined the architecture of this framework, the tests performed to validate it, and its first applications to real data.

Bibliography

- [1] Cosmograil homepage. <http://www.cosmograil.org>.
- [2] HST archive. <https://archive.stsci.edu/>.
- [3] Wikimedia under creative commons licence. https://commons.wikimedia.org/wiki/File:Gravitational_lens_geometry.svg, 2012 (accessed July 2, 2018).
- [4] J. Abadie et al. TOPICAL REVIEW: Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27(17):173001, September 2010.
- [5] J. Abadie et al. Search for gravitational waves from low mass compact binary coalescence in ligo’s sixth science run and virgo’s science runs 2 and 3. *Phys. Rev. D*, 85:082002, Apr 2012.
- [6] A. Agnello, H. Lin, L. Buckley-Geer, T. Treu, V. Bonvin, F. Courbin, C. Lemon, T. Morishita, A. Amara, M. W. Auger, S. Birrer, J. Chan, T. Collett, A. More, C. D. Fassnacht, J. Frieman, P. J. Marshall, R. G. McMahon, G. Meylan, S. H. Suyu, F. Castander, D. Finley, A. Howell, C. Kochanek, M. Makler, P. Martini, N. Morgan, B. Nord, F. Ostrovski, P. Schechter, D. Tucker, R. Wechsler, T. M. C. Abbott, F. B. Abdalla, S. Allam, A. Benoit-Lévy, E. Bertin, D. Brooks, D. L. Burke, A. C. Rosell, M. C. Kind, J. Carretero, M. Croce, C. E. Cunha, C. B. D’Andrea, L. N. da Costa, S. Desai, J. P. Dietrich, T. F. Eifler, B. Flaugher, P. Fosalba, J. García-Bellido, E. Gaztanaga, M. S. Gill, D. A. Goldstein, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, K. Honscheid, D. J. James, K. Kuehn, N. Kuropatkin, T. S. Li, M. Lima, M. A. G. Maia, M. March, J. L. Marshall, P. Melchior, F. Menanteau, R. Miquel, R. L. C. Ogando, A. A. Plazas, A. K. Romer, E. Sanchez, R. Schindler, M. Schubnell, I. Sevilla-Noarbe, M. Smith, R. C. Smith, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, D. Thomas, and A. R. Walker. Models of the strongly lensed quasar DES J0408-5354. *MNRAS*, 472:4038–4050, December 2017.

- [7] K. Akiyama, K. Kuramochi, S. Ikeda, V. L. Fish, F. Tazaki, M. Honma, S. S. Doeleman, A. E. Broderick, J. Dexter, M. Mościbrodzka, K. L. Bouman, A. A. Chael, and M. Zazizen. Imaging the Schwarzschild-radius-scale Structure of M87 with the Event Horizon Telescope Using Sparse Modeling. *The Astrophysical Journal*, 838:1, March 2017.
- [8] K. Akiyama, R.-S. Lu, V. L. Fish, S. S. Doeleman, A. E. Broderick, J. Dexter, K. Hada, M. Kino, H. Nagai, M. Honma, M. D. Johnson, J. C. Algaba, K. Asada, C. Brinkerink, R. Blundell, G. C. Bower, R. Cappallo, G. B. Crew, M. Dexter, S. A. Dzib, R. Freund, P. Friberg, M. Gurwell, P. T. P. Ho, M. Inoue, T. P. Krichbaum, L. Loinard, D. MacMahon, D. P. Marrone, J. M. Moran, M. Nakamura, N. M. Nagar, G. Ortiz-Leon, R. Plambeck, N. Pradel, R. A. Primiani, A. E. E. Rogers, A. L. Roy, J. SooHoo, J.-L. Tavares, R. P. J. Tilanus, M. Titus, J. Wagner, J. Weintroub, P. Yamaguchi, K. H. Young, A. Zensus, and L. M. Ziurys. 230 GHz VLBI Observations of M87: Event-horizon-scale Structure during an Enhanced Very-high-energy γ -Ray State in 2012. *The Astrophysical Journal*, 807:150, July 2015.
- [9] C. Alard and R. H. Lupton. A Method for Optimal Image Subtraction. *ApJ*, 503:325–331, August 1998.
- [10] C. Alcock et al. First Observation of Parallax in a Gravitational Microlensing Event. *ApJL*, 454:L125, December 1995.
- [11] C. Alcock et al. The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations. *ApJ*, 542:281–307, October 2000.
- [12] C. Alcock et al. MACHO Project Limits on Black Hole Dark Matter in the 1-30 M_{solar} Range. *ApJL*, 550:L169–L172, April 2001.
- [13] C.h. Alcock, C. W. Akerlof, R.A. Allsman, T.S. Axelrod, D.P. Bennett, S. Chan, K.H. Cook, K.C. Freeman, K. Griest, S.L. Marshall, et al. Possible gravitational microlensing of a star in the large magellanic cloud. *Nature*, 365(6447):621, 1993.
- [14] C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Statistics and computing*, 18(4):343–373, 2008.
- [15] E. Aubourg, P. Bareyre, S. Brehin, M. Gros, M. Lachieze-Rey, B. Laurent, E. Lesquoy, C.h. Magneville, A. Milsztajn, L. Moscoso, et al. Evidence for gravitational microlensing by dark objects in the galactic halo. *Nature*, 365(6447):623, 1993.

- [16] S. Awiphan, E. Kerins, and A.C. Robin. Besançon galactic model analysis of moa-ii microlensing: evidence for a mass deficit in the inner bulge. *Monthly Notices of the Royal Astronomical Society*, 456(2):1666–1680, 2015.
- [17] D. C. Backer. Scattering of radio emission from the compact object in Sagittarius A. *ApJL*, 222:L9–L12, May 1978.
- [18] E. A. Baltz. Microlensing Surveys of M31 in the Wide-Field Imaging Era. *ApJ*, 624:168–177, May 2005.
- [19] K. Bandura et al. Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 9145 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 22, July 2014.
- [20] J. M. Bardeen. In B. S. DeWitt and C. DeWitt, editors, *Black holes (Les astres occlus)*, page 215. Gordon and Breach, New York, 1973.
- [21] M. Bartelmann. TOPICAL REVIEW Gravitational lensing. *Classical and Quantum Gravity*, 27(23):233001, December 2010.
- [22] K. Belczynski, V. Kalogera, and T. Bulik. A Comprehensive Study of Binary Compact Objects as Gravitational Wave Sources: Evolutionary Channels, Rates, and Physical Properties. *ApJ*, 572:407–431, June 2002.
- [23] L. Benkevitch, K. Akiyama, R. Lu, S. Doeleman, and V. Fish. Reconstruction of Static Black Hole Images Using Simple Geometric Forms. *ArXiv e-prints 1609.00055*, August 2016.
- [24] M. Betancourt. A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*, 2017.
- [25] M. Betancourt, S. Byrne, S. Livingstone, M. Girolami, et al. The geometric foundations of hamiltonian monte carlo. *Bernoulli*, 23(4A):2257–2298, 2017.
- [26] J. Binney and S. Tremaine. *Galactic dynamics*. 1987.
- [27] J. A. Blackburne, D. Pooley, S. Rappaport, and P. L. Schechter. Sizes and Temperature Profiles of Quasar Accretion Disks from Chromatic Microlensing. *ApJ*, 729:34, March 2011.

- [28] G. Blanc, C. Afonso, C. Alard, J.N. Albert, G. Aldering, A. Amadon, J. Andersen, R. Ansari, É. Aubourg, C. Balland, et al. Type Ia supernova rate at a redshift of ~ 0.1 . *Astronomy & Astrophysics*, 423(3):881–894, 2004.
- [29] A. Boehle, A. M. Ghez, R. Schödel, L. Meyer, S. Yelda, S. Albers, G. D. Martinez, E. E. Becklin, T. Do, J. R. Lu, K. Matthews, M. R. Morris, B. Sitarski, and G. Witzel. An Improved Distance and Mass Estimate for Sgr A* from a Multistar Orbit Analysis. *The Astrophysical Journal*, 830:17, October 2016.
- [30] A. S. Bolton, S. Burles, L. V. E. Koopmans, T. Treu, R. Gavazzi, L. A. Moustakas, R. Wayth, and D. J. Schlegel. The Sloan Lens ACS Survey. V. The Full ACS Strong-Lens Sample. *ApJ*, 682:964–984, August 2008.
- [31] V. Bonvin, F. Courbin, S. H. Suyu, P. J. Marshall, C. E. Rusu, D. Sluse, M. Tewes, K. C. Wong, T. Collett, C. D. Fassnacht, T. Treu, M. W. Auger, S. Hilbert, L. V. E. Koopmans, G. Meylan, N. Rumbaugh, A. Sonnenfeld, and C. Spiniello. HOLiCOW - V. New COSMOGRAIL time delays of HE 0435-1223: H_0 to 3.8 per cent precision from strong lensing in a flat Λ CDM model. *MNRAS*, 465:4914–4930, March 2017.
- [32] G. C. Bower, A. Deller, P. Demorest, A. Brunthaler, R. Eatough, H. Falcke, M. Kramer, K. J. Lee, and L. Spitler. The Angular Broadening of the Galactic Center Pulsar SGR J1745-29: A New Constraint on the Scattering Medium. *Astrophysical Journal, Letters*, 780:L2, January 2014.
- [33] G. C. Bower, W. M. Goss, H. Falcke, D. C. Backer, and Y. Lithwick. The Intrinsic Size of Sagittarius A* from 0.35 to 6 cm. *ApJL*, 648:L127–L130, September 2006.
- [34] G. C. Bower, W. M. Goss, H. Falcke, D. C. Backer, and Y. Lithwick. The Intrinsic Size of Sagittarius A* from 0.35 to 6 cm. *Astrophysical Journal, Letters*, 648:L127–L130, September 2006.
- [35] G. E.P. Box. Sampling and bayes' inference in scientific modelling and robustness. *Journal of the Royal Statistical Society. Series A (General)*, pages 383–430, 1980.
- [36] Cajo J. F. Ter Braak. A markov chain monte carlo version of the genetic algorithm differential evolution: easy bayesian computing for real parameter spaces. *Statistics and Computing*, 16(3):239–249, Sep 2006.
- [37] T. Bringmann, P. Scott, and Y. Akrami. Improved constraints on the primordial power spectrum at small scales from ultracompact minihalos. *Physical Review D*, 85(12):125027, 2012.

- [38] A. Broderick and R. Blandford. Covariant magnetoionic theory - I. Ray propagation. *Monthly Notices of the Royal Astronomical Society*, 342:1280–1290, July 2003.
- [39] A. Broderick and R. Blandford. Covariant magnetoionic theory - II. Radiative transfer. *Monthly Notices of the Royal Astronomical Society*, 349:994–1008, April 2004.
- [40] A. E. Broderick, V. L. Fish, S. S. Doeleman, and A. Loeb. Evidence for low black hole spin and physically motivated accretion models from millimeter-vlbi observations of sagittarius a*. *The Astrophysical Journal*, 735(2):110, 2011.
- [41] A. E. Broderick, V. L. Fish, M. D. Johnson, K. Rosenfeld, C. Wang, S. S. Doeleman, K. Akiyama, T. Johannsen, and A. L. Roy. Modeling seven years of event horizon telescope observations with radiatively inefficient accretion flow models. *The Astrophysical Journal*, 820(2):137, 2016.
- [42] A. E. Broderick, T. Johannsen, A. Loeb, and D. Psaltis. Testing the no-hair theorem with event horizon telescope observations of sagittarius a*. *The Astrophysical Journal*, 784(1):7, 2014.
- [43] A. E. Broderick and A. Loeb. Imaging bright-spots in the accretion flow near the black hole horizon of Sgr A*. *Monthly Notices of the Royal Astronomical Society*, 363:353–362, October 2005.
- [44] A. E. Broderick and A. Loeb. Frequency-dependent Shift in the Image Centroid of the Black Hole at the Galactic Center as a Test of General Relativity. *Astrophysical Journal, Letters*, 636:L109–L112, January 2006.
- [45] A. E. Broderick and A. Loeb. Imaging optically-thin hotspots near the black hole horizon of Sgr A* at radio and near-infrared wavelengths. *Monthly Notices of the Royal Astronomical Society*, 367:905–916, Apr 2006.
- [46] A. E. Broderick and A. Loeb. Signatures of Relativistic Helical Motion in the Rotation Measures of Active Galactic Nucleus Jets. *Astrophysical Journal, Letters*, 703:L104–L108, October 2009.
- [47] A. E. Broderick and A. Tchekhovskoy. Horizon-scale Lepton Acceleration in Jets: Explaining the Compact Radio Emission in M87. *The Astrophysical Journal*, 809:97, August 2015.
- [48] B. C. Bromley, F. Melia, and S. Liu. Polarimetric Imaging of the Massive Black Hole at the Galactic Center. *Astrophysical Journal, Letters*, 555:L83–L86, July 2001.

- [49] T. Bronzwaer, J. Davelaar, Z. Younsi, M. Moscibrodzka, H. Falcke, M. Kramer, and L. Rezzolla. Raptor. i. time-dependent radiative transfer in arbitrary spacetimes. *Astronomy and Astrophysics*, 2018.
- [50] A. Brunthaler, M. J. Reid, and H. Falcke. Atmosphere-Corrected Phase-Referencing. In J. Romney and M. Reid, editors, *Future Directions in High Resolution Astronomy*, volume 340 of *Astronomical Society of the Pacific Conference Series*, page 455, December 2005.
- [51] A. Brunthaler, M. J. Reid, K. M. Menten, X-W. Zheng, A. Bartkiewicz, Y. K. Choi, T. Dame, K. Hachisuka, K. Immer, G. Moellenbrock, et al. The bar and spiral structure legacy (bessel) survey: mapping the Milky Way with vlbi astrometry. *Astronomische Nachrichten*, 332(5):461–466, 2011.
- [52] J. Casares, P. G. Jonker, and G. Israelian. *X-Ray Binaries*, page 1499. 2017.
- [53] G. Chabrier. The galactic disk mass function: reconciliation of the hubble space telescope and nearby determinations. *The Astrophysical Journal Letters*, 586(2):L133, 2003.
- [54] G. Chabrier. Galactic Stellar and Substellar Initial Mass Function. *PASP*, 115:763–795, July 2003.
- [55] G. Chabrier. The initial mass function: from salpeter 1955 to 2005. In *The Initial Mass Function 50 Years Later*, pages 41–50. Springer, 2005.
- [56] A. Chael, M. Rowan, R. Narayan, M. Johnson, and L. Sironi. The role of electron heating physics in images and variability of the Galactic Centre black hole Sagittarius A*. *Monthly Notices of the Royal Astronomical Society*, May 2018.
- [57] C.-K. Chan, S. Liu, C. L. Fryer, D. Psaltis, F. Özel, G. Rockefeller, and F. Melia. MHD Simulations of Accretion onto Sgr A*: Quiescent Fluctuations, Outbursts, and Quasiperiodicity. *The Astrophysical Journal*, 701:521–534, August 2009.
- [58] C.-k. Chan, L. Medeiros, F. Ozel, and D. Psaltis. GRay2: A General Purpose Geodesic Integrator for Kerr Spacetimes. *ArXiv e-prints*, June 2017.
- [59] C.-K. Chan, D. Psaltis, F. Özel, R. Narayan, and A. Sądowski. The Power of Imaging: Constraining the Plasma Properties of GRMHD Simulations using EHT Observations of Sgr A*. *The Astrophysical Journal*, 799:1, January 2015.
- [60] M. Chandra, C. F. Gammie, F. Foucart, and E. Quataert. An Extended Magnetohydrodynamics Model for Relativistic Weakly Collisional Plasmas. *The Astrophysical Journal*, 810:162, September 2015.

- [61] O. Chwolson. Über eine mögliche form fiktiver doppelsterne. *Astron. Nachr.*, 221:329, 1924.
- [62] A Claret. A new non-linear limb-darkening law for lte stellar atmosphere models. calculations for $-5.0 \leq \log [m/h] \leq +1$, $2000 \leq k \leq 50000$ k at several surface gravities. *Astronomy and Astrophysics*, 363:1081–1190, 2000.
- [63] J. J. Condon. Cosmological evolution of radio sources. *ApJ*, 287:461–474, December 1984.
- [64] J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick. The NRAO VLA Sky Survey. *AJ*, 115:1693–1716, May 1998.
- [65] J. J. Condon et al. Resolving the Radio Source Background: Deeper Understanding through Confusion. *ApJ*, 758:23, October 2012.
- [66] F. Courbin, V. Bonvin, E. Buckley-Geer, C. D. Fassnacht, J. Frieman, H. Lin, P. J. Marshall, S. H. Suyu, T. Treu, T. Anguita, V. Motta, G. Meylan, E. Paic, M. Tewes, A. Agnello, D. C.-Y. Chao, M. Chijani, D. Gilman, K. Rojas, P. Williams, A. Hempel, S. Kim, R. Lachaume, M. Rabus, T. M. C. Abbott, S. Allam, J. Annis, M. Banerji, K. Bechtol, A. Benoit-Lévy, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, C. B. D’Andrea, L. N. da Costa, C. Davis, D. L. DePoy, S. Desai, B. Flaugher, P. Fosalba, J. Garcia-Bellido, E. Gaztanaga, D. A. Goldstein, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, K. Honscheid, D. J. James, K. Kuehn, S. Kuhlmann, N. Kuropatkin, O. Lahav, M. Lima, M. A. G. Maia, M. March, J. L. Marshall, R. G. McMahon, F. Menanteau, R. Miquel, B. Nord, A. A. Plazas, E. Sanchez, V. Scarpine, R. Schindler, M. Schubnell, I. Sevilla-Noarbe, M. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, G. Tarle, D. L. Tucker, A. R. Walker, and W. Wester. COSMOGRAIL XVI: Time delays for the quadruply imaged quasar DES J0408-5354 with high-cadence photometric monitoring. *ArXiv e-prints*, June 2017.
- [67] F. Courbin, P. Saha, and P. L. Schechter. Quasar Lensing. In F. Courbin and D. Minniti, editors, *Gravitational Lensing: An Astrophysical Tool*, volume 608 of *Lecture Notes in Physics*, Berlin Springer Verlag, page 1, 2002.
- [68] Arlin PS Crotts and Austin B Tomaney. Results from a survey of gravitational microlensing toward m31. *The Astrophysical Journal Letters*, 473(2):L87, 1996.
- [69] N. Dalal and C. S. Kochanek. Direct detection of CDM substructure. *Astrophys. J.*, 572:25–33, 2002.

- [70] J. Davelaar, M. Moscibrodzka, T. Bronzwaer, and H. Falcke. General relativistic magnetohydrodynamical κ -jet models for Sgr A*. *ArXiv e-prints*, December 2017.
- [71] Mr C Davidson. IX. a determination of the deflection of light by the sun's gravitational field, from observations made at the total eclipse of may 29, 1919. *Phil. Trans. R. Soc. Lond. A*, 220(571-581):291–333, 1920.
- [72] F. de Gasperin, E. Orrú, M. Murgia, A. Merloni, H. Falcke, R. Beck, R. Beswick, L. Birzan, A. Bonafede, M. Brüggen, G. Brunetti, K. Chyży, J. Conway, J. H. Croston, T. Enßlin, C. Ferrari, G. Heald, S. Heidenreich, N. Jackson, G. Macario, J. McKean, G. Miley, R. Morganti, A. Offringa, R. Pizzo, D. Rafferty, H. Röttgering, A. Shulevski, M. Steinmetz, C. Tasse, S. van der Tol, W. van Driel, R. J. van Weeren, J. E. van Zwieten, A. Alexov, J. Anderson, A. Asgekar, M. Avruch, M. Bell, M. R. Bell, M. Bentum, G. Bernardi, P. Best, F. Breitling, J. W. Broderick, A. Butcher, B. Ciardi, R. J. Dettmar, J. Eisloffel, W. Frieswijk, H. Gankema, M. Garrett, M. Gerbers, J. M. Griessmeier, A. W. Gunst, T. E. Hassall, J. Hessels, M. Hoeft, A. Horneffer, A. Karastergiou, J. Köhler, Y. Koopman, M. Kuniyoshi, G. Kuper, P. Maat, G. Mann, M. Mevius, D. D. Mulcahy, H. Munk, R. Nijboer, J. Noordam, H. Paas, M. Pandey, V. N. Pandey, A. Polatidis, W. Reich, A. P. Schoenmakers, J. Sluman, O. Smirnov, C. Sobey, B. Stappers, J. Swinbank, M. Tagger, Y. Tang, I. van Bemmelen, W. van Cappellen, A. P. van Duin, M. van Haarlem, J. van Leeuwen, R. Vermeulen, C. Vocks, S. White, M. Wise, O. Wucknitz, and P. Zarka. M 87 at metre wavelengths: the LOFAR picture. *Astronomy and Astrophysics*, 547:A56, November 2012.
- [73] X. Delfosse, T. Forveille, D. Ségransan, J-L. Beuzit, S. Udry, C. Perrier, and M. Mayor. Accurate masses of very low mass stars: Iv improved mass-luminosity relations. *arXiv preprint astro-ph/0010586*, 2000.
- [74] F Derue, C Afonso, C Alard, J-N Albert, J Andersen, R Ansari, É Aubourg, P Bareyre, F Bauer, J-P Beaulieu, et al. Observation of microlensing toward the galactic spiral arms. EROS II 3 year survey. *Astronomy & Astrophysics*, 373(1):126–138, 2001.
- [75] F Derue, C Afonso, C Alard, and JN Albert. Observation of microlensing towards the galactic spiral arms. EROS II 2 year survey. *arXiv preprint astro-ph/9903209*, 1999.
- [76] J. Dexter. A public code for general relativistic, polarised radiative transfer around spinning black holes. *Monthly Notices of the Royal Astronomical Society*, 462(1):115–136, 2016.

- [77] J. Dexter, E. Agol, and P. C. Fragile. Millimeter Flares and VLBI Visibilities from Relativistic Simulations of Magnetized Accretion Onto the Galactic Center Black Hole. *Astrophysical Journal, Letters*, 703:L142–L146, Oct 2009.
- [78] Marshall D.J. Priv. com. 2015.
- [79] S. Doeleman. Building an event horizon telescope: (sub)mm VLBI in the ALMA era. In *10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the New Generation of Radio Arrays*, page 53, 2010.
- [80] S. Doeleman, E. Agol, D. Backer, F. Baganoff, G. C. Bower, A. Broderick, A. Fabian, V. Fish, C. Gammie, P. Ho, M. Honman, T. Krichbaum, A. Loeb, D. Marrone, M. Reid, A. Rogers, I. Shapiro, P. Strittmatter, R. Tilanus, J. Weintraub, A. Whitney, M. Wright, and L. Ziurys. Imaging an Event Horizon: submm-VLBI of a Super Massive Black Hole. In *astro2010: The Astronomy and Astrophysics Decadal Survey*, volume 2010 of *ArXiv Astrophysics e-prints*, page 68, 2009.
- [81] S. S. Doeleman, V. L. Fish, D. E. Schenck, C. Beaudoin, R. Blundell, G. C. Bower, A. E. Broderick, R. Chamberlin, R. Freund, P. Friberg, M. A. Gurwell, P. T. P. Ho, M. Honma, M. Inoue, T. P. Krichbaum, J. Lamb, A. Loeb, C. Lonsdale, D. P. Marrone, J. M. Moran, T. Oyama, R. Plambeck, R. A. Primiani, A. E. E. Rogers, D. L. Smythe, J. SooHoo, P. Strittmatter, R. P. J. Tilanus, M. Titus, J. Weintraub, M. Wright, K. H. Young, and L. M. Ziurys. Jet-Launching Structure Resolved Near the Supermassive Black Hole in M87. *Science*, 338:355–358, Oct 2012.
- [82] S. S. Doeleman, J. Weintraub, A. E. E. Rogers, et al. Event-horizon-scale structure in the supermassive black hole candidate at the Galactic Centre. *Nature*, 455:78–80, 2008.
- [83] J. C. Dolence et al. grmonty: A Monte Carlo Code for Relativistic Radiative Transport. *Astrophysical Journal, Supplement*, 184:387–397, October 2009.
- [84] J. C. Dolence, C. F. Gammie, H. Shiokawa, and S. C. Noble. Near-infrared and X-Ray Quasi-periodic Oscillations in Numerical Models of Sgr A*. *Astrophysical Journal, Letters*, 746:L10, February 2012.
- [85] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.
- [86] E Dwek, RG Arendt, MG Hauser, T Kelsall, CM Lisse, SH Moseley, RF Silverberg, TJ Sodroski, and JL Weiland. Morphology, near-infrared luminosity, and mass of the

- galactic bulge from coBE observations. *The Astrophysical Journal*, 445:716–730, 1995.
- [87] D. J. Earl and M. W. Deem. Parallel tempering: Theory, applications, and new perspectives. *Physical Chemistry Chemical Physics*, 7(23):3910–3916, 2005.
- [88] A. Eckart, F. K. Baganoff, R. Schödel, M. Morris, R. Genzel, G. C. Bower, D. Marrone, J. M. Moran, T. Viehmann, M. W. Bautz, W. N. Brandt, G. P. Garmire, T. Ott, S. Trippe, G. R. Ricker, C. Straubmeier, D. A. Roberts, F. Yusef-Zadeh, J. H. Zhao, and R. Rao. The flare activity of Sagittarius A*. New coordinated mm to X-ray observations. *Astronomy and Astrophysics*, 450:535–555, May 2006.
- [89] A. Eckart, R. Schödel, M. García-Marín, G. Witzel, A. Weiss, F. K. Baganoff, M. R. Morris, T. Bertram, M. Dovčiak, W. J. Duschl, V. Karas, S. König, T. P. Krichbaum, M. Krips, D. Kunneriath, R.-S. Lu, S. Markoff, J. Mauerhan, L. Meyer, J. Moutaka, K. Mužić, F. Najarro, J.-U. Pott, K. F. Schuster, L. O. Sjouwerman, C. Straubmeier, C. Thum, S. N. Vogel, H. Wiesemeyer, M. Zamaninasab, and J. A. Zensus. Simultaneous NIR/sub-mm observation of flare emission from Sagittarius A*. *Astronomy and Astrophysics*, 492:337–344, December 2008.
- [90] A.S. Eddington. Space, time and gravitation. *Camb. Univ. Press*, 1920.
- [91] A. Einstein. On the influence of gravitation on the propagation of light. *Annalen der Physik*, 35(898-908):906, 1911.
- [92] A. Einstein. Lens-like action of a star by the deviation of light in the gravitational field. *Science*, 84(2188):506–507, 1936.
- [93] A. Einstein. Lens-Like Action of a Star by the Deviation of Light in the Gravitational Field. *Science*, 84:506–507, December 1936.
- [94] N. Epchtein, E. Deul, S. Derriere, J. Borsenberger, D. Egret, G. Simon, C. Alard, L.G. Balázs, B. De Batz, M-R. Cioni, et al. A preliminary database of denis point sources. *Astronomy and Astrophysics*, 349:236–242, 1999.
- [95] ESA, editor. *The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission*, volume 1200 of *ESA Special Publication*, 1997.

- [96] E. Eulaers, M. Tewes, P. Magain, F. Courbin, I. Asfandiyarov, Sh. Ehgamberdiev, S. R. Kumar, C.S. Stalin, T.P. Prabhu, G. Meylan, et al. Cosmograil: the cosmological monitoring of gravitational lenses-xii. time delays of the doubly lensed quasars sdss j1206+ 4332 and hs 2209+ 1914. *Astronomy & Astrophysics*, 553:A121, 2013.
- [97] N. W. Evans. The power-law galaxies. *MNRAS*, 267:333–360, March 1994.
- [98] A. C. Fabian. Observational Evidence of Active Galactic Nuclei Feedback. *Annual Review of Astron and Astrophys*, 50:455–489, September 2012.
- [99] H. Falcke. Imaging black holes: past, present and future. In *Journal of Physics Conference Series*, volume 942 of *Journal of Physics Conference Series*, page 012001, December 2017.
- [100] H. Falcke, W. M. Goss, H. Matsuo, P. Teuben, J.-H. Zhao, and R. Zylka. The Simultaneous Spectrum of Sagittarius A* from 20 Centimeters to 1 Millimeter and the Nature of the Millimeter Excess. *The Astrophysical Journal*, 499:731–734, May 1998.
- [101] H. Falcke and S. Markoff. The jet model for Sgr A*: Radio and X-ray spectrum. *Astronomy and Astrophysics*, 362:113–118, October 2000.
- [102] H. Falcke and S. B. Markoff. Toward the event horizon—the supermassive black hole in the Galactic Center. *Classical and Quantum Gravity*, 30(24):244003, December 2013.
- [103] Heino Falcke, Fulvio Melia, and Eric Agol. Viewing the shadow of the black hole at the galactic center. *Astrophysical Journal, Letters*, 528(1):L13–L16, 2000.
- [104] V. L. Fish, S. S. Doeleman, C. Beaudoin, R. Blundell, D. E. Bolin, G. C. Bower, R. Chamberlin, R. Freund, P. Friberg, M. A. Gurwell, M. Honma, M. Inoue, T. P. Krichbaum, J. Lamb, D. P. Marrone, J. M. Moran, T. Oyama, R. Plambeck, R. Primiani, A. E. E. R., D. L. Smythe, J. SooHoo, P. Strittmatter, R. P. J. Tilanus, M. Titus, J. Weintroub, M. Wright, D. Woody, K. H. Young, and L. M. Ziurys. 1.3mm wavelength vlbi of sagittarius a*: Detection of time-variable emission on event horizon scales. *Astrophysical Journal, Letters*, 727(2):L36, February 2011.
- [105] V. L. Fish, M. D. Johnson, S. S. Doeleman, A. E. Broderick, D. Psaltis, R.-S. Lu, K. Akiyama, W. Alef, J. C. Algaba, K. Asada, C. Beaudoin, A. Bertarini, L. Blackburn, R. Blundell, G. C. Bower, C. Brinkerink, R. Cappallo, A. A. Chael, R. Chamberlin, C.-K. Chan, G. B. Crew, J. Dexter, M. Dexter, S. A. Dzib, H. Falcke, R. Freund, P. Friberg, C. H. Greer, M. A. Gurwell, P. T. P. Ho, M. Honma, M. Inoue, T. Johannsen, J. Kim, T. P. Krichbaum, J. Lamb, J. León-Tavares, A. Loeb, L. Loinard, D. MacMahon, D. P.

- Marrone, J. M. Moran, M. Mościbrodzka, G. N. Ortiz-León, T. Oyama, F. Özel, R. L. Plambeck, N. Pradel, R. A. Primiani, A. E. E. Rogers, K. Rosenfeld, H. Rottmann, A. L. Roy, C. Ruszczyk, D. L. Smythe, J. SooHoo, J. Spilker, J. Stone, P. Strittmatter, R. P. J. Tilanus, M. Titus, L. Vertatschitsch, J. Wagner, J. F. C. Wardle, J. Weintroub, D. Woody, M. Wright, P. Yamaguchi, A. Young, K. H. Young, J. A. Zensus, and L. M. Ziurys. Persistent Asymmetric Structure of Sagittarius A* on Event Horizon Scales. *The Astrophysical Journal*, 820(2):90, April 2016.
- [106] E. Fomalont, S. Kopeikin, G. Lanyi, and J. Benson. Progress in Measurements of the Gravitational Bending of Radio Waves Using the VLBA. *ApJ*, 699:1395–1402, July 2009.
- [107] M. Frigo and S.G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. Special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [108] C. Fryer, A. Burrows, and W. Benz. Population Syntheses for Neutron Star Systems with Intrinsic Kicks. *ApJ*, 496:333–351, March 1998.
- [109] C. L. Fryer. Mass Limits For Black Hole Formation. *ApJ*, 522:413–418, September 1999.
- [110] C. L. Fryer. Compact object formation and the supernova explosion engine. *Classical and Quantum Gravity*, 30(24):244002, December 2013.
- [111] C. L. Fryer, K. Belczynski, G. Wiktorowicz, M. Dominik, V. Kalogera, and D. E. Holz. Compact Remnant Mass Function: Dependence on the Explosion Mechanism and Metallicity. *ApJ*, 749:91, April 2012.
- [112] C. L. Fryer, A. Heger, N. Langer, and S. Wellstein. The Limiting Stellar Initial Mass for Black Hole Formation in Close Binary Systems. *ApJ*, 578:335–347, October 2002.
- [113] S. V. Fuerst and K. Wu. Radiation transfer of emission lines in curved space-time. *Astronomy and Astrophysics*, 424:733–746, September 2004.
- [114] R. Fux. 3D self-consistent N-body barred models of the Milky Way: II. Gas dynamics. *arXiv preprint astro-ph/9903154*, 1999.
- [115] C. Gallart, M. Zoccali, and A. Aparicio. The adequacy of stellar evolution models for the interpretation of the color-magnitude diagrams of resolved stellar populations. *Annu. Rev. Astron. Astrophys.*, 43:387–434, 2005.
- [116] C. F. Gammie and P. K. Leung. A Formalism for Covariant Polarized Radiative Transport by Ray Tracing. *The Astrophysical Journal*, 752:123, June 2012.

- [117] C. F. Gammie, J. C. McKinney, and G. Tóth. HARM: A Numerical Scheme for General Relativistic Magnetohydrodynamics. *The Astrophysical Journal*, 589:444–457, May 2003.
- [118] B. S. Gaudi. Microlensing Surveys for Exoplanets. *ARA&A*, 50:411–453, September 2012.
- [119] B. S. Gaudi. Microlensing Surveys for Exoplanets. *ARA&A*, 50:411–453, September 2012.
- [120] B. S. Gaudi and A. Gould. Spectrophotometric Resolution of Stellar Surfaces with Microlensing. *ApJ*, 513:619–625, March 1999.
- [121] A. Gelman, X-L. Meng, and H. Stern. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica sinica*, pages 733–760, 1996.
- [122] R. Genzel, F. Eisenhauer, and S. Gillessen. The Galactic Center massive black hole and nuclear star cluster. *Rev. Mod. Phys.*, 82:3121–3195, Dec 2010.
- [123] A. M. Ghez, S. Salim, N. N. Weinberg, J. R. Lu, T. Do, J. K. Dunn, K. Matthews, M. R. Morris, S. Yelda, E. E. Becklin, T. Kremenek, M. Milosavljevic, and J. Naiman. Measuring distance and properties of the Milky Way’s central supermassive black hole with stellar orbits. *The Astrophysical Journal*, 689(2):1044–1062, 2008.
- [124] A. Gianninas, B. D. Strickland, M. Kilic, and P. Bergeron. Limb-darkening Coefficients for Eclipsing White Dwarfs. *ApJ*, 766:3, March 2013.
- [125] I. S. Glass. *Handbook of infrared astronomy*, volume 1. Cambridge University Press, 1999.
- [126] R. Gold, J. C. McKinney, M. D. Johnson, and S. S. Doeleman. Probing the magnetic field structure in sgr a* on black hole horizon scales with polarized radiative transfer simulations. *The Astrophysical Journal*, 837(2):180, Mar 2017.
- [127] J. Goodman and J. Weare. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80, 2010.
- [128] A. Gould. Extending the MACHO search to about 10×10^6 solar masses. *ApJ*, 392:442–451, June 1992.
- [129] A. Gould. Extending the macho search to about 10×10^6 solar masses. *The Astrophysical Journal*, 392:442–451, 1992.

- [130] A. Gould. Theory of Pixel Lensing. *ApJ*, 470:201, October 1996.
- [131] A. Gould. Measuring the Remnant Mass Function of the Galactic Bulge. *ApJ*, 535:928–931, June 2000.
- [132] A. Gould. Microlens Masses from 1-D Parallaxes and Heliocentric Proper Motions. *Journal of Korean Astronomical Society*, 47:215–218, December 2014.
- [133] A. Gould and A. Loeb. Discovering planetary systems through gravitational microlenses. *ApJ*, 396:104–114, September 1992.
- [134] C. R. Gwinn, Y. Y. Kovalev, M. D. Johnson, and V. A. Soglasnov. Discovery of substructure in the scatter-broadened image of sgr a*. *ApJL*, 794(1):L14, 2014.
- [135] Jahreiss H. and Wielen R. Esa symp. 402, hipparcos-venice. *ESA Publications*, page 675, 1997.
- [136] C Hamadache, L Le Guillou, Patrick Tisserand, C Afonso, JN Albert, J Andersen, R Ansari, É Aubourg, P Bareyre, JP Beaulieu, et al. Galactic bulge microlensing optical depth from eros-2. *Astronomy & Astrophysics*, 454(1):185–199, 2006.
- [137] F. D. A. Hartwick. Cosmic Star Formation History from Local Observations and an Outline for Galaxy Formation and Evolution. *ApJ*, 603:108–118, March 2004.
- [138] M. R. S. Hawkins. Timescale of variation and the size of the accretion disc in active galactic nuclei. *Astronomy & Astrophysics*, 462:581–589, February 2007.
- [139] T. M. Heckman and P. N. Best. The Coevolution of Galaxies and Supermassive Black Holes: Insights from Surveys of the Contemporary Universe. *Annual Review of Astron and Astrophys*, 52:589–660, August 2014.
- [140] C. B. Henderson, B. S. Gaudi, C. Han, J. Skowron, M. T. Penny, D. Nataf, and A. P. Gould. Optimal Survey Strategies and Predicted Planet Yields for the Korean Microlensing Telescope Network. *ApJ*, 794:52, October 2014.
- [141] J. S. Heyl. Diffractive microlensing - II. Substellar disc and halo objects. *MNRAS*, 411:1780–1786, March 2011.
- [142] Yashar D. Hezaveh et al. Detection of lensing substructure using ALMA observations of the dusty galaxy SDP.81. *Astrophys. J.*, 823(1):37, 2016.

- [143] K. Hirotani and H.-Y. Pu. Energetic Gamma Radiation from Rapidly Rotating Black Holes. *The Astrophysical Journal*, 818:50, February 2016.
- [144] G. Hobbs, D. R. Lorimer, A. G. Lyne, and M. Kramer. A statistical study of 233 pulsar proper motions. *MNRAS*, 360:974–992, July 2005.
- [145] L. Huang and R. V. Shcherbakov. Faraday conversion and rotation in uniformly magnetized relativistic plasmas. *Monthly Notices of the Royal Astronomical Society*, 416:2574–2592, October 2011.
- [146] L. Huang, R. Takahashi, and Z.-Q. Shen. Testing the Accretion Flow with Plasma Wave Heating Mechanism for Sagittarius A* by the 1.3 mm VLBI Measurements. *The Astrophysical Journal*, 706:960–969, December 2009.
- [147] Z. Ivezić et al. LSST: from Science Drivers to Reference Design and Anticipated Data Products. *arXiv:0805.2366*, May 2008.
- [148] R. C. Jennison. A phase sensitive interferometer technique for the measurement of the Fourier transforms of spatial brightness distributions of small angular extent. *Monthly Notices of the Royal Astronomical Society*, 118:276, 1958.
- [149] B. Jeter, A. E. Broderick, and R. Gold. Differentiating Disk and Black Hole Driven Jets with EHT Images of Variability in M87. *ArXiv e-prints*, April 2018.
- [150] T. Johannsen. Systematic study of event horizons and pathologies of parametrically deformed Kerr spacetimes. *Physical Review D*, 87(12):124017, June 2013.
- [151] T. Johannsen, C. Wang, A. E. Broderick, S. S. Doeleman, V. L. Fish, A. Loeb, and D. Psaltis. Testing General Relativity with Accretion-Flow Imaging of Sgr A*. *Physical Review Letters*, 117(9):091101, August 2016.
- [152] M. D. Johnson, K. L. Bouman, L. Blackburn, A. A. Chael, J. Rosen, H. Shiokawa, F. Roelofs, K. Akiyama, V. L. Fish, and S. S. Doeleman. Dynamical Imaging with Interferometry. *The Astrophysical Journal*, 850:172, December 2017.
- [153] M. D. Johnson, V. L. Fish, S. S. Doeleman, D. P. Marrone, R. L. Plambeck, J. F. C. Wardle, K. Akiyama, K. Asada, C. Beaudoin, L. Blackburn, R. Blundell, G. C. Bower, C. Brinkerink, A. E. Broderick, R. Cappallo, A. A. Chael, G. B. Crew, J. Dexter, M. Dexter, R. Freund, P. Friberg, R. Gold, M. A. Gurwell, P. T. P. Ho, M. Honma, M. Inoue, M. Kosowsky, T. P. Krichbaum, J. Lamb, A. Loeb, R.-S. Lu, D. MacMahon, J. C. McKinney, J. M. Moran, R. Narayan, R. A. Primiani, D. Psaltis, A. E. E. Rogers, K. Rosenfeld,

- J. SooHoo, R. P. J. Tilanus, M. Titus, L. Vertatschitsch, J. Weintraub, M. Wright, K. H. Young, J. A. Zensus, and L. M. Ziurys. Resolved magnetic-field structure and variability near the event horizon of sagittarius a*. *Science*, 350(6265):1242–1245, Dec 2015.
- [154] M. D. Johnson and C. R. Gwinn. Theory and Simulations of Refractive Substructure in Resolved Scatter-broadened Images. *The Astrophysical Journal*, 805:180, June 2015.
- [155] A. B. Kamruddin and J. Dexter. A geometric crescent model for black hole images. *Monthly Notices of the Royal Astronomical Society*, 434(1):765–771, 2013.
- [156] E. Kerins, A. C. Robin, and D. J. Marshall. Synthetic microlensing maps of the Galactic bulge. *MNRAS*, 396:1202–1210, June 2009.
- [157] P. N. Kholopov. *General catalogue of variable stars. Vol.4: Reference tables*. 1990.
- [158] J. Kim, D. P. Marrone, C.-K. Chan, L. Medeiros, F. Özel, and D. Psaltis. Bayesian Techniques for Comparing Time-dependent GRMHD Simulations to Variable Event Horizon Telescope Observations. *The Astrophysical Journal*, 832:156, December 2016.
- [159] S-L. Kim et al. Technical specifications of the kmtnet observation system. In *SPIE Astronomical Telescopes+ Instrumentation*, pages 77333F–77333F. International Society for Optics and Photonics, 2010.
- [160] K. H. Knuth, M. Habeck, N. K. Malakar, A. M. Mubeen, and B. Placek. Bayesian evidence and model selection. *Digital Signal Processing*, 47:50–67, 2015.
- [161] C. S. Kochanek. Quantitative Interpretation of Quasar Microlensing Light Curves. *ApJ*, 605:58–77, April 2004.
- [162] C. S. Kochanek. Constraints on core collapse from the black hole mass function. *MNRAS*, 446:1213–1222, January 2015.
- [163] P. Kroupa. On the variation of the initial mass function. *Monthly Notices of the Royal Astronomical Society*, 322(2):231–246, 2001.
- [164] N. Lartillot and H. Philippe. Computing bayes factors using thermodynamic integration. *Systematic biology*, 55(2):195–207, 2006.
- [165] T. Lasserre et al. Not enough stellar mass Machos in the Galactic halo. *A&A*, 355:L39–L42, March 2000.

- [166] K. Liao, T. Treu, P. Marshall, C. D. Fassnacht, N. Rumbaugh, G. Dobler, A. Aghamousa, V. Bonvin, F. Courbin, A. Hojjati, N. Jackson, V. Kashyap, S. Rathna Kumar, E. Linder, K. Mandel, X.-L. Meng, G. Meylan, L. A. Moustakas, T. P. Prabhu, A. Romero-Wolf, A. Shafieloo, A. Siemiginowska, C. S. Stalin, H. Tak, M. Tewes, and D. van Dyk. Strong Lens Time Delay Challenge. II. Results of TDC1. *ApJ*, 800:11, February 2015.
- [167] M. López-Corredoira, A. Cabrera-Lavers, T.J. Mahoney, P.L. Hammersley, F. Garzón, and C. González-Fernández. The long bar in the Milky Way: Corroboration of an old hypothesis. *The Astronomical Journal*, 133(1):154, 2006.
- [168] R.-S. Lu, T. P. Krichbaum, A. Eckart, S. König, D. Kunneriath, G. Witzel, A. Witzel, and J. A. Zensus. Multiwavelength VLBI observations of Sagittarius A*. *A&A*, 525:A76, January 2011.
- [169] R.-S. Lu, T. P. Krichbaum, A. L. Roy, V. L. Fish, S. S. Doeleman, M. D. Johnson, K. Akiyama, D. Psaltis, W. Alef, K. Asada, C. Beaudoin, A. Bertarini, L. Blackburn, R. Blundell, G. C. Bower, C. Brinkerink, A. E. Broderick, R. Cappallo, G. B. Crew, J. Dexter, M. Dexter, H. Falcke, R. Freund, P. Friberg, C. H. Greer, M. A. Gurwell, P. T. P. Ho, M. Honma, M. Inoue, J. Kim, J. Lamb, M. Lindqvist, D. Macmahon, D. P. Marrone, I. Martí-Vidal, K. M. Menten, J. M. Moran, N. M. Nagar, R. L. Plambeck, R. A. Primiani, A. E. E. Rogers, E. Ros, H. Rottmann, J. SooHoo, J. Spilker, J. Stone, P. Strittmatter, R. P. J. Tilanus, M. Titus, L. Vertatschitsch, J. Wagner, J. Weintraub, M. Wright, K. H. Young, J. A. Zensus, and L. M. Ziurys. Detection of Intrinsic Source Structure at ~ 3 Schwarzschild Radii with Millimeter-VLBI Observations of SAGITTARIUS A*. *The Astrophysical Journal*, 859(1):60, May 2018.
- [170] R.-S. Lu, F. Roelofs, V. L. Fish, H. Shiokawa, S. S. Doeleman, C. F. Gammie, H. Falcke, T. P. Krichbaum, and J. A. Zensus. Imaging an Event Horizon: Mitigation of Source Variability of Sagittarius A*. *The Astrophysical Journal*, 817:173, February 2016.
- [171] Y. E. Lyubarskii. Flicker noise in accretion discs. *MNRAS*, 292:679, December 1997.
- [172] A. Maeder. Light Curves of the Gravitational Lens-like Action for Binaries with Degenerate Stars. *A&A*, 26:215, July 1973.
- [173] D. J. Majaess, D. G. Turner, and D. J. Lane. Characteristics of the galaxy according to cepheids. *Monthly Notices of the Royal Astronomical Society*, 398(1):263–270, 2009.
- [174] S Mao and R Di Stefano. Interpretation of gravitational microlensing by binary systems. *The Astrophysical Journal*, 440:22–27, 1995.

- [175] S. Mao and B. Paczynski. Gravitational microlensing by double stars and planetary systems. *ApJL*, 374:L37–L40, June 1991.
- [176] D. P. Marrone. *Submillimeter properties of Sagittarius A*: The polarization and spectrum from 230 to 690 GHz and the submillimeter array polarimeter*. PhD thesis, Harvard University, 2006.
- [177] D. J. Marshall, A. C. Robin, C. Reylé, M. Schultheis, and S. Picaud. Modelling the Galactic interstellar extinction distribution in three dimensions. *A&A*, 453:635–651, July 2006.
- [178] N. Matsunaga, H. Fukushi, and Y. Nakada. Mira variables in the Galactic bulge with OGLE-II data. *MNRAS*, 364:117–125, November 2005.
- [179] J. C. McKinney, A. Tchekhovskoy, and R. D. Blandford. General relativistic magnetohydrodynamic simulations of magnetically choked accretion flows around black holes. *Monthly Notices of the Royal Astronomical Society*, 423:3083–3117, July 2012.
- [180] J. C. McKinney, A. Tchekhovskoy, A. Sadowski, and R. Narayan. Three-dimensional general relativistic radiation magnetohydrodynamical simulation of super-Eddington accretion, using a new code HARMRAD with M1 closure. *Monthly Notices of the Royal Astronomical Society*, 441:3177–3208, July 2014.
- [181] L. Medeiros, C.-k. Chan, F. Özel, D. Psaltis, J. Kim, D. P. Marrone, and A. Sądowski. GRMHD Simulations of Visibility Amplitude Variability for Event Horizon Telescope Images of Sgr A*. *The Astrophysical Journal*, 856:163, April 2018.
- [182] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- [183] S. P. Meyn and R. L. Tweedie. Markov chains and stochastic stability. communication and control engineering series. *Springer-Verlag London Ltd., London*, 1:993, 1993.
- [184] E. Middelberg, A. L. Roy, R. C. Walker, and H. Falcke. VLBI observations of weak sources using fast frequency switching. *Astronomy and Astrophysics*, 433:897–909, April 2005.
- [185] D. Minniti, P.W. Lucas, J.P. Emerson, R.K. Saito, M. Hempel, P. Pietrukowicz, A.V. Ahumada, M.V. Alonso, J. Alonso-Garcia, J. I. Arias, et al. Vista variables in the via lactea (vvv): The public eso near-ir variability survey of the milky way. *New Astronomy*, 15(5):433–443, 2010.

- [186] N. Miyake et al. A Sub-Saturn Mass Planet, MOA-2009-BLG-319Lb. *ApJ*, 728:120, February 2011.
- [187] Y. Mizuno, Z. Younsi, C. M. Fromm, O. Porth, M. De Laurentis, H. Olivares, H. Falcke, M. Kramer, and L. Rezzolla. The current ability to test theories of gravity with black hole shadows. *Nature Astronomy*, April 2018.
- [188] M. Moniez. Microlensing as a probe of the Galactic structure: 20 years of microlensing optical depth studies. *General Relativity and Gravitation*, 42:2047–2074, September 2010.
- [189] M. Moniez. Microlensing as a probe of the galactic structure: 20 years of microlensing optical depth studies. *General Relativity and Gravitation*, 42(9):2047–2074, 2010.
- [190] M. Moniez, S. Sajadian, M. Karami, S. Rahvar, and R. Ansari. Understanding EROS2 observations toward the spiral arms within a classical Galactic model framework. *Astronomy and Astrophysics*, 604:A124, August 2017.
- [191] C. W. Morgan, C. S. Kochanek, N. D. Morgan, and E. E. Falco. The Quasar Accretion Disk Size-Black Hole Mass Relation. *The Astrophysical Journal*, 712:1129–1136, April 2010.
- [192] M. Mościbrodzka and H. Falcke. Coupled jet-disk model for Sagittarius A*: explaining the flat-spectrum radio core with GRMHD simulations of jets. *Astronomy and Astrophysics*, 559:L3, Nov 2013.
- [193] M. Mościbrodzka, H. Falcke, H. Shiokawa, and C. F. Gammie. Observational appearance of inefficient accretion flows and jets in 3D GRMHD simulations: Application to Sagittarius A*. *Astronomy and Astrophysics*, 570:A7, October 2014.
- [194] M. Moscibrodzka and C. F. Gammie. ipole - semianalytic scheme for relativistic polarized radiative transport. to be published, 2017.
- [195] M. Mościbrodzka, C. F. Gammie, J. C. Dolence, H. Shiokawa, and P. K. Leung. Radiative Models of SGR A* from GRMHD Simulations. *The Astrophysical Journal*, 706:497–507, November 2009.
- [196] Y. Muraki et al. Discovery and Mass Measurements of a Cold, 10 Earth Mass Planet and Its Host Star. *ApJ*, 741:22, November 2011.

- [197] R. Narayan, R. Mahadevan, J. E. Grindlay, R. G. Popham, and C. Gammie. Advection-dominated accretion model of Sagittarius A*: evidence for a black hole at the Galactic center. *The Astrophysical Journal*, 492:554–568, January 1998.
- [198] B. Nelson, E.B. Ford, and M. J. Payne. Run dmc: An efficient, parallel code for analyzing radial velocity observations using n-body integrations and differential evolution markov chain monte carlo. *Astrophysical Journal, Supplement*, 210(1):11, Jan 2014.
- [199] S Calchi Novati, V Bozza, I Bruni, M Dall’Ora, F De Paolis, M Dominik, R Gualandi, G Ingrassio, Ph Jetzer, L Mancini, et al. The m31 pixel lensing plan campaign: macho lensing and self-lensing signals. *The Astrophysical Journal*, 783(2):86, 2014.
- [200] M. Oguri and P. J. Marshall. Gravitationally lensed quasars and supernovae in future wide-field optical imaging surveys. *MNRAS*, 405:2579–2593, July 2010.
- [201] H. Olivares, O. Porth, and Y. Mizuno. The Black Hole Accretion Code: adaptive mesh refinement and constrained transport. *ArXiv e-prints*, February 2018.
- [202] F. Özel, D. Psaltis, and R. Narayan. Hybrid Thermal-Nonthermal Synchrotron Emission from Hot Accretion Flows. *The Astrophysical Journal*, 541:234–249, September 2000.
- [203] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock. The Black Hole Mass Distribution in the Galaxy. *ApJ*, 725:1918–1927, December 2010.
- [204] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock. The Black Hole Mass Distribution in the Galaxy. *ApJ*, 725:1918–1927, December 2010.
- [205] B. Paczynski. Gravitational microlensing by the galactic halo. *The Astrophysical Journal*, 304:1–5, May 1986.
- [206] B. Paczyński and K. Z. Stanek. Galactocentric Distance with the Optical Gravitational Lensing Experiment and HIPPARCOS Red Clump Stars. *The Astrophysical Journal*, 494:L219–L222, February 1998.
- [207] S. Pasetto, E. K. Grebel, T. Zwitter, C. Chiosi, G. Bertelli, O. Bienayme, G. Seabroke, J. Bland-Hawthorn, C. Boeche, B. K. Gibson, et al. Thick disk kinematics from rave and the solar motion. *Astronomy & Astrophysics*, 547:A70, 2012.
- [208] S. Pasetto, E. K. Grebel, T. Zwitter, C. Chiosi, G. Bertelli, O. Bienayme, G. Seabroke, J. Bland-Hawthorn, C. Boeche, B. K. Gibson, et al. Thin disk kinematics from rave and the solar motion. *Astronomy & Astrophysics*, 547:A71, 2012.

- [209] J. A. Peacock and S. J. Dodds. Non-linear evolution of cosmological power spectra. *MNRAS*, 280(3):L19–L26, 1996.
- [210] O. Pejcha and T. A. Thompson. The Landscape of the Neutrino Mechanism of Core-collapse Supernovae: Neutron Star and Black Hole Mass Functions, Explosion Energies, and Nickel Yields. *ApJ*, 801:90, March 2015.
- [211] U.-L. Pen, J.-P. Macquart, A. T. Deller, and W. Brisken. 50 picoarcsec astrometry of pulsar emission. *MNRAS*, 440:L36–L40, May 2014.
- [212] M. A. C. Perryman et al. The HIPPARCOS Catalogue. *Astronomy and Astrophysics*, 323, July 1997.
- [213] M. Persic, P. Salucci, and F. Stel. The universal rotation curve of spiral galaxies - I. The dark matter connection. *MNRAS*, 281:27–47, July 1996.
- [214] P. Podsiadlowski, E. Pfahl, and S. Rappaport. Neutron-Star Birth Kicks. In F. A. Rasio and I. H. Stairs, editors, *Binary Radio Pulsars*, volume 328 of *Astronomical Society of the Pacific Conference Series*, page 327, July 2005.
- [215] P. Popowski, K. Griest, C.L. Thomas, K. H. Cook, D.P. Bennett, A. C. Becker, D. R. Alves, D. Minniti, A. J. Drake, C. Alcock, et al. Microlensing optical depth toward the galactic bulge using clump giants from the macho survey. *The Astrophysical Journal*, 631(2):879, 2005.
- [216] M. Portail, O. Gerhard, C. Wegg, and M. Ness. Dynamical modelling of the galactic bulge and bar: The Milky Ways pattern speed, stellar and dark matter mass distribution. *Monthly Notices of the Royal Astronomical Society*, page stw2819, 2016.
- [217] W. M. Poteet, H. K. Cauthen, N. Kappler, L. G. Kappler, B.-G. Park, C.-U. Lee, S.-L. Kim, and S.-M. Cha. Design and fabrication of three 1.6-meter telescopes for the Korea Microlensing Telescope Network (KMTNet). In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8444 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, September 2012.
- [218] D. Psaltis, M. Johnson, R. Narayan, L. Medeiros, L. Blackburn, and G. Bower. A Model for Anisotropic Interstellar Scattering and its Application to Sgr A*. *ArXiv e-prints*, May 2018.
- [219] D. Psaltis, F. Özel, C.-K. Chan, and D. P. Marrone. A General Relativistic Null Hypothesis Test with Event Horizon Telescope Observations of the Black Hole Shadow in Sgr A*. *The Astrophysical Journal*, 814:115, December 2015.

- [220] D. Psaltis, N. Wex, and M. Kramer. A Quantitative Test of the No-hair Theorem with Sgr A* Using Stars, Pulsars, and the Event Horizon Telescope. *The Astrophysical Journal*, 818:121, February 2016.
- [221] H.-Y. Pu, K. Akiyama, and K. Asada. The Effects of Accretion Flow Dynamics on the Black Hole Shadow of Sagittarius A*. *The Astrophysical Journal*, 831:4, November 2016.
- [222] H.-Y. Pu, K. Yun, Z. Younsi, and S.-J. Yoon. Odyssey: A Public GPU-based Code for General Relativistic Radiative Transfer in Kerr Spacetime. *The Astrophysical Journal*, 820:105, April 2016.
- [223] Y. R. Rahal et al. The EROS2 search for microlensing events towards the spiral arms: the complete seven season results. *A&A*, 500:1027–1044, June 2009.
- [224] Y.R. Rahal, C. Afonso, J.N. Albert, J. Andersen, R. Ansari, E. Aubourg, P. Bareyre, J.P. Beaulieu, X. Charlot, F. Couchot, et al. The eros2 search for microlensing events towards the spiral arms: the complete seven season results. *Astronomy & Astrophysics*, 500(3):1027–1044, 2009.
- [225] Y.R. Rahal, C. Afonso, J.N. Albert, J. Andersen, R. Ansari, E. Aubourg, P. Bareyre, J.P. Beaulieu, X. Charlot, F. Couchot, et al. Vizier online data catalog: Eros cmd towards 4 galactic spiral arms fields (rahal+, 2009). *VizieR Online Data Catalog*, 350, 2009.
- [226] S. Rahvar. Gravitational microlensing I: A unique astrophysical tool. *International Journal of Modern Physics D*, 24:30020, April 2015.
- [227] S. Rahvar, S. Baghran, and N. Afshordi. Transient weak lensing by cosmological dark matter microhaloes. *PRD*, 89(6):063001, March 2014.
- [228] S. Rahvar and M. Dominik. Planetary microlensing signals from the orbital motion of the source star around the common barycentre. *MNRAS*, 392:1193–1204, January 2009.
- [229] S. Rahvar, M. Moniez, R. Ansari, and O. Perdureau. Study of a strategy for parallax microlensing detection towards the Magellanic Clouds. *A&A*, 412:81–90, December 2003.
- [230] S. Rathna Kumar, M. Tewes, C. S. Stalin, F. Courbin, I. Asfandiyarov, G. Meylan, E. Eu-laers, T. P. Prabhu, P. Magain, H. Van Winckel, and S. Ehgamberdiev. COSMOGRAIL: the COSmological MONitoring of GRAVItational Lenses. XIV. Time delay of the doubly lensed quasar SDSS J1001+5027. *Astronomy & Astrophysics*, 557:A44, September 2013.
- [231] A. Refregier. Weak Gravitational Lensing by Large-Scale Structure. *ARA&A*, 41:645–668, 2003.

- [232] S. Refsdal. On the possibility of determining Hubble's parameter and the masses of galaxies from the gravitational lens effect. *MNRAS*, 128:307, 1964.
- [233] M. J. Reid and J. E. Goldston. How Mira Variables Change Visual Light by a Thousandfold. *ApJ*, 568:931–938, April 2002.
- [234] M. J. Reid and K. M. Menten. Imaging the radio photospheres of mira variables. *ApJ*, 671(2):2068, 2007.
- [235] M. J. Reid, K. M. Menten, R. Genzel, T. Ott, R. Schödel, and A. Eckart. The Position of Sagittarius A*. II. Accurate Positions and Proper Motions of Stellar SiO Masers near the Galactic Center. *ApJ*, 587:208–220, April 2003.
- [236] M. J. Reid, K. M. Menten, S. Trippe, T. Ott, and R. Genzel. The Position of Sagittarius A*. III. Motion of the Stellar Cusp. *ApJ*, 659:378–388, April 2007.
- [237] R. A. Remillard and J. E. McClintock. X-Ray Properties of Black-Hole Binaries. *Annual Review of Astronomy and Astrophysics*, 44:49–92, September 2006.
- [238] S. M. Ressler, A. Tchekhovskoy, E. Quataert, M. Chandra, and C. F. Gammie. Electron thermodynamics in GRMHD simulations of low-luminosity black hole accretion. *Monthly Notices of the Royal Astronomical Society*, 454(2):1848–1870, December 2015.
- [239] A. C. Robin, D. J. Marshall, M. Schultheis, and C. Reylé. Stellar populations in the Milky Way bulge region: towards solving the galactic bulge and bar shapes using 2mass data. *Astronomy & Astrophysics*, 538:A106, 2012.
- [240] A. C. Robin, C. Reylé, S. Derrière, and S. Picaud. A synthetic view on structure and evolution of the Milky Way. *A&A*, 409:523–540, October 2003.
- [241] N. Rumbaugh, C. D. Fassnacht, J. P. McKean, L. V. E. Koopmans, M. W. Auger, and S. H. Suyu. Radio monitoring campaigns of six strongly lensed quasars. *MNRAS*, 450:1042–1056, June 2015.
- [242] S. Sajadian and S. Rahvar. Illuminating hot Jupiters in caustic crossing. *MNRAS*, 407:373–380, September 2010.
- [243] S. Sajadian and S. Rahvar. Photometric, Astrometric and Polarimetric observations of gravitational microlensing events. *ArXiv e-prints*, September 2014.

- [244] M. Salaris, A. Serenelli, A. Weiss, and M. Miller Bertolami. Semi-empirical White Dwarf Initial-Final Mass Relationships: A Thorough Analysis of Systematic Uncertainties Due to Stellar Evolution Models. *ApJ*, 692:1013–1032, February 2009.
- [245] E. F. Schlafly and D. P. Finkbeiner. Measuring reddening with sloan digital sky survey stellar spectra and recalibrating sfd. *The Astrophysical Journal*, 737(2):103, 2011.
- [246] D.J. Schlegel, D. P. Finkbeiner, and M. Davis. Maps of dust infrared emission for use in estimation of reddening and cosmic microwave background radiation foregrounds. *The Astrophysical Journal*, 500(2):525, 1998.
- [247] P Schneider, J Ehlers, and EE Falco. Gravitational lenses Springer-Verlag. *Berlin Inc., New York*, 1992.
- [248] P. Schneider, C. Kochanek, and J. Wambsganss. *Gravitational lensing: strong, weak and micro: Saas-Fee advanced course 33*, volume 33. Springer Science & Business Media, 2006.
- [249] M. Schultheis, S. Ramírez, K. Sellgren, S. Stolovy, and S. Ganesh. Interstellar extinction in the Galactic Center and its impact on the study of AGB stars. In *The Evolving ISM in the Milky Way and Nearby Galaxies*, page 59, January 2009.
- [250] N. I. Shakura and R. A. Sunyaev. Black holes in binary systems. Observational appearance. *AAP*, 24:337–355, 1973.
- [251] H. Y. Shan et al. Weak lensing mass map and peak statistics in Canada-France-Hawaii Telescope Stripe 82 survey. *MNRAS*, 442:2534–2542, August 2014.
- [252] R. V. Shcherbakov. ASTRORAY: General relativistic polarized radiative transfer code. Astrophysics Source Code Library, jul 2014.
- [253] H. Shiohawa, C. F. Gammie, and S. S. Doeleman. Time Domain Filtering of Resolved Images of Sgr A*. *The Astrophysical Journal*, 846:29, September 2017.
- [254] John Skilling et al. Nested sampling for general bayesian computation. *Bayesian analysis*, 1(4):833–859, 2006.
- [255] B. J. Smith, D. Leisawitz, M. W. Castelaz, and D. Luttermoser. Infrared Light Curves of Mira Variable Stars from COBE DIRBE Data. *AJ*, 123:948–964, February 2002.
- [256] M. C. Smith, S. Mao, and B. Paczyński. Acceleration and parallax effects in gravitational microlensing. *MNRAS*, 339:925–936, March 2003.

- [257] R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman. Stable clustering, the halo model and non-linear cosmological power spectra. *MNRAS*, 341:1311–1332, June 2003.
- [258] I. Soszyński et al. The Optical Gravitational Lensing Experiment. The OGLE-III Catalog of Variable Stars. XV. Long-Period Variables in the Galactic Bulge. *AcA*, 63:21–36, March 2013.
- [259] V. Springel, J. Wang, M. Vogelsberger, A. Ludlow, A. Jenkins, A. Helmi, J. F. Navarro, C. S. Frenk, and S. D. M. White. The Aquarius Project: the subhaloes of galactic haloes. *MNRAS*, 391:1685–1711, December 2008.
- [260] R. A. Street et al. Spitzer Parallax of OGLE-2015-BLG-0966: A Cold Neptune in the Galactic Disk. *ApJ*, 819(2):93, 2016.
- [261] T. Sumi et al. Microlensing Optical Depth toward the Galactic Bulge from Microlensing Observations in Astrophysics Group Observations during 2000 with Difference Image Analysis. *ApJ*, 591:204–227, July 2003.
- [262] T. Sumi et al. Microlensing Optical Depth toward the Galactic Bulge Using Bright Sources from OGLE-II. *ApJ*, 636:240–260, January 2006.
- [263] T. Sumi et al. The Microlensing Event Rate and Optical Depth Toward the Galactic Bulge from MOA-II. *ArXiv:1305.0186S*, may 2013.
- [264] T. Sumi, P.R. Woźniak, A. Udalski, M. Szymański, M. Kubiak, G. Pietrzyński, I. Soszyński, K. Żebruń, O. Szewczyk, B. Paczyński, et al. Microlensing optical depth toward the galactic bulge using bright sources from ogle-ii. *The Astrophysical Journal*, 636(1):240, 2006.
- [265] S. H. Suyu, V. Bonvin, F. Courbin, C. D. Fassnacht, C. E. Rusu, D. Sluse, T. Treu, K. C. Wong, M. W. Auger, X. Ding, S. Hilbert, P. J. Marshall, N. Rumbaugh, A. Sonnenfeld, M. Tewes, O. Tihhonova, A. Agnello, R. D. Blandford, G. C.-F. Chen, T. Collett, L. V. E. Koopmans, K. Liao, G. Meylan, and C. Spiniello. H0LiCOW - I. H₀ Lenses in COSMOGRAIL's Wellspring: program overview. *MNRAS*, 468:2590–2604, July 2017.
- [266] R. Takahashi, M. Oguri, M. Sato, and T. Hamana. Probability Distribution Functions of Cosmological Lensing: Convergence, Shear, and Magnification. *ApJ*, 742:15, November 2011.

- [267] A. Tamm, E. Tempel, P. Tenjes, O. Tihhonova, and T. Tuvikene. Stellar mass map and dark matter distribution in M 31. *A&A*, 546:A4, October 2012.
- [268] A.R. Taylor. The square kilometre array. *Proceedings of the International Astronomical Union*, 3(S248):164–169, 2007.
- [269] M. Tewes, F. Courbin, G. Meylan, C. S. Kochanek, E. Eulaers, N. Cantale, A. M. Mosquera, I. Asfandiyarov, P. Magain, H. van Winckel, D. Sluse, R. K. S. Keerthi, C. S. Stalin, T. P. Prabhu, P. Saha, and S. Dye. COSMOGRAIL: Measuring Time Delays of Gravitationally Lensed Quasars to Constrain Cosmology. *The Messenger*, 150:49–52, December 2012.
- [270] M. Tewes, F. Courbin, G. Meylan, C. S. Kochanek, E. Eulaers, N. Cantale, A. M. Mosquera, P. Magain, H. Van Winckel, D. Sluse, G. Cataldi, D. Vörös, and S. Dye. COSMOGRAIL: the COSmological MONitoring of GRAVItational Lenses. XIII. Time delays and 9-yr optical monitoring of the lensed quasar RX J1131-1231. *Astronomy & Astrophysics*, 556:A22, August 2013.
- [271] A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr. *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*. 2017.
- [272] F. X. Timmes, S. E. Woosley, and T. A. Weaver. The Neutron Star and Black Hole Initial Mass Function. *ApJ*, 457:834, February 1996.
- [273] C. Turon, Y. Requieme, M. Grenon, A. Gomez, D. Morin, F. Crifo, F. Arenou, M. Froeschle, F. Mignard, M.A.C. Perryman, et al. Properties of the hipparcos input catalogue. *Astronomy and Astrophysics*, 304:82, 1995.
- [274] R. Q. Twiss, A. W. L. Carter, and A. G. Little. Brightness distribution over some strong radio sources at 1427 Mc/s. *The Observatory*, 80:153–159, August 1960.
- [275] R. Q. Twiss, A. W. L. Carter, and A. G. Little. Radio Source Interferometry at 1427 Mc/s. *Australian Journal of Physics*, 15:378, September 1962.
- [276] A Udalski, M Szymanski, J Kaluzny, M Kubiak, W Krzeminski, M Mateo, GW Preston, and B Paczynski. The optical gravitational lensing experiment. discovery of the first candidate microlensing event in the direction of the galactic bulge. *Acta Astronomica*, 43:289–294, 1993.
- [277] A. Udalski, M. Szymanski, J. Kaluzny, M. Kubiak, M. Mateo, W. Krzeminski, and B. Paczynski. The Optical Gravitational Lensing Experiment. The Early Warning System: Real Time Microlensing. *AcA*, 44:227–234, July 1994.

- [278] A. Udalski, M. K. Szymański, and G. Szymański. OGLE-IV: Fourth Phase of the Optical Gravitational Lensing Experiment. *AcA*, 65:1–38, March 2015.
- [279] W. D. Vausden, W. M. Farr, and I. Mandel. Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations. *MNRAS*, 455:1919–1937, January 2016.
- [280] D. Walsh, R. F. Carswell, and R. J. Weymann. 0957 + 561 A, B - Twin quasistellar objects or gravitational lens. *Nature*, 279:381–384, May 1979.
- [281] Dennis Walsh, Robert F Carswell, and Ray J Weymann. 0957+ 561 a, b: twin quasistellar objects or gravitational lens? *Nature*, 279(5712):381, 1979.
- [282] C. Wegg, O. Gerhard, and M. Portail. The structure of the Milky Way’s bar outside the bulge. *Monthly Notices of the Royal Astronomical Society*, 450(4):4050–4069, 2015.
- [283] A. R. Whitney et al. Demonstration of a broadband-RF VLBI system at 16 Gbps data rate per station. *arXiv:1210.5961*, October 2012.
- [284] B. Willems, M. Henninger, T. Levin, N. Ivanova, V. Kalogera, K. McGhee, F. X. Timmes, and C. L. Fryer. Understanding Compact Object Formation and Natal Kicks. I. Calculation Methods and the Case of GRO J1655-40. *ApJ*, 625:324–346, May 2005.
- [285] L. Wisotzki, P. L. Schechter, H. V. Bradt, J. Heinmüller, and D. Reimers. HE 0435-1223: A wide separation quadruple QSO and gravitational lens. *AAP*, 395:17–23, November 2002.
- [286] H. J. Witt and S. Mao. Can lensed stars be regarded as pointlike for microlensing by MACHOs? *ApJ*, 430:505–510, August 1994.
- [287] H. J. Witt and S. Mao. Can lensed stars be regarded as pointlike for microlensing by MACHOs? *ApJ*, 430:505–510, August 1994.
- [288] P.R. Wood. Convection and pulsation in red giant stars. *Proceedings of the International Astronomical Union*, S239, 2(S239):343–348, 2007.
- [289] S. E. Woosley, A. Heger, and T. A. Weaver. The evolution and explosion of massive stars. *Reviews of Modern Physics*, 74:1015–1071, November 2002.
- [290] Ł. Wyrzykowski et al. OGLE-III Microlensing Events and the Structure of the Galactic Bulge. *ApJS*, 216:12, January 2015.

- [291] Ł. Wyrzykowski, A. E. Rynkiewicz, J. Skowron, S. Kozłowski, A. Udalski, M. K. Szymański, M. Kubiak, I. Soszyński, G. Pietrzyński, R. Poleski, et al. OGLE-III microlensing events and the structure of the galactic bulge based on observations obtained with the 1.3 m Warsaw Telescope at the Las Campanas Observatory of the Carnegie Institution for Science. *The Astrophysical Journal Supplement Series*, 216(1):12, 2015.
- [292] J. Yoo, D.L. DePoy, A. Gal-Yam, B.S. Gaudi, A. Gould, C. Han, Y. Lipkin, D. Maoz, E.O. Ofek, B-G Park, et al. OGLE-2003-BLG-262: Finite-source effects from a point-mass lens. *The Astrophysical Journal*, 603(1):139, 2004.
- [293] Z. Younsi, K. Wu, and S. V. Fuerst. General relativistic radiative transfer: formulation and emission from structured tori around black holes. *Astronomy and Astrophysics*, 545:A13, September 2012.
- [294] F. Yuan, S. Markoff, and H. Falcke. A Jet-ADAF model for Sgr A*. *Astronomy and Astrophysics*, 383(3):854–863, Mar 2002.
- [295] F. Yuan and R. Narayan. Hot Accretion Flows Around Black Holes. *Annual Review of Astron and Astrophys*, 52:529–588, August 2014.
- [296] F. Yuan, E. Quataert, and R. Narayan. On the Nature of the Variable Infrared Emission from Sagittarius A*. *The Astrophysical Journal*, 606:894–899, May 2004.
- [297] J. Zavala and N. Afshordi. Universal clustering of dark matter in phase space. *MNRAS*, 457:986–992, March 2016.
- [298] W. Zhang, S. E. Woosley, and A. Heger. Fallback and Black Hole Production in Massive Stars. *ApJ*, 679:639–654, May 2008.
- [299] J.-H. Zhao, K. H. Young, R. M. Herrnstein, P. T. P. Ho, T. Tsutsumi, K. Y. Lo, W. M. Goss, and G. C. Bower. Variability of Sagittarius A*: Flares at 1 Millimeter. *Astrophysical Journal, Letters*, 586:L29–L32, March 2003.
- [300] F. Zwicky. Nebulae as gravitational lenses. *Physical Review*, 51(4):290, 1937.
- [301] F. Zwicky. On the probability of detecting nebulae which act as gravitational lenses. *Physical Review*, 51(8):679, 1937.

APPENDICES

Appendix A

Breaking degeneracy with combined observation

In this section we introduce the mathematical basis of gravitational microlensing, introduce the photometric and astrometric observations of microlensing events and the way we can break the degeneracy between the lens parameters.

A.1 Lens Equation for Single Lens

The single lens equation is given by geometrical relation between the angular position of source β , angular position of the image θ and Einstein angle θ_E as follows:

$$\theta^2 - \theta\beta - \theta_E^2 = 0, \quad (\text{A.1})$$

where β and θ are given with respect the center of coordinate system where lens is located. In the gravitational microlensing the angular position of the source β moves along a straight line, varying as

$$\beta^2 = \beta_0^2 + \theta_E^2 \left(\frac{t - t_0}{t_E} \right)^2, \quad (\text{A.2})$$

where β_0 is the minimum impact factor and t_0 is the time of closest angular approach of the lens and source. From the equation (A.1), we get two solutions for the position of two images at two sides of the lens as follows:

$$\theta^\pm = \frac{\beta \pm \sqrt{\beta^2 + 4\theta_E^2}}{2}. \quad (\text{A.3})$$

Substituting equation (A.2) in equation(A.3) results in the radial dynamics of the images as a function of time.

The angular distance between the two images ($\Delta\theta = \theta^+ - \theta^-$), regardless of the position of source, can be written in terms of the impact parameter and the Einstein angle:

$$\Delta\theta = \sqrt{\beta^2 + 4\theta_E^2} \quad (\text{A.4})$$

At the minimum impact parameter where $\beta = \beta_0$, two images get nearest approach to each other and in this case the trajectory of the source is perpendicular to the line connecting these two images

$$\beta_0 = \sqrt{\Delta\theta^2 - 4\theta_E^2}. \quad (\text{A.5})$$

We note that on the right hand side the term in the square-root should be positive which implies $\Delta\theta \geq 2\theta_E$. So the separation between the two images is always larger than the diameter of the Einstein ring while one image is inside and the other one is outside the Einstein ring.

The other observable parameter which can be measured from the photometry is the magnification of the source during the lensing. Since in the gravitational lensing the flux of light is conserved, the apparent change of area for lensed images would lead to a change in the energy we receive from each image. The area of images compared to the source size is given by the inverse of determinant of Jacobian. Denoting the area for the two images by A^+ and A^- the differential element of area of an image to that of the source is given by:

$$A^\pm = \frac{\theta^\pm}{\beta} \frac{\partial\theta^\pm}{\partial\beta}. \quad (\text{A.6})$$

Substituting equation (A.3) in equation(A.6) the area for each images could be written as:

$$A^\pm = \frac{1}{2} \left(\frac{\beta^2 + 2\theta_E^2}{\beta \sqrt{\beta^2 + 4\theta_E^2}} \pm 1 \right) \quad (\text{A.7})$$

The overall magnification is the sum of these two terms,

$$A = |A^+| + |A^-| = \frac{\beta^2 + 2\theta_E^2}{\beta \sqrt{\beta^2 + 4\theta_E^2}}. \quad (\text{A.8})$$

Substituting equation for (A.2) the magnification is independent of the Einstein angle and only depends on t_0 , t_E and $u_0 = \beta_0/\theta_E$.

At the time of minimum approach the magnification would have its maximum as a function of β_0 and θ_E . Substituting β_0 from equation (A.5) in (A.8), maximum magnification can be obtained as a function of Einstein angle and angular separation between the images as follows:

$$A_{\max} = \frac{\Delta\theta^2 - 2\theta_E^2}{\Delta\theta \sqrt{\Delta\theta^2 - 4\theta_E^2}} \quad (\text{A.9})$$

As a result the Einstein angle can be directly extracted from A_{\max} as described in equation (3.24).

If the distance of the source from the earth is known (which is the case for microlensing events in Galactic Center) measuring θ_E can put constraints on the mass and the lens distance from observer. On the other hand parallax measurements can impose another constraint on the lens parameters and in theory can resolve the lens parameters.

A.2 Parallax Effect in Microlensing Light Curves

The microlensing parallax effect results from the accelerating motion of the earth around sun. It manifests itself as the relative motion of the source with respect to lens deviates from a straight line. Assuming the earth moves around the Sun in a circular orbit, the relative angular motion of the lens and source with respect to the observer in the case of stationary source and lens is given by

$$\pi_x = \pi \cos \xi(t), \quad (\text{A.10})$$

$$\pi_y = \pi \cos \beta \sin \xi(t), \quad (\text{A.11})$$

where

$$\pi = \frac{1}{\theta_E} \left(\frac{1 \text{ a.u.}}{D_l} - \frac{1 \text{ a.u.}}{D_s} \right), \quad (\text{A.12})$$

β is the angular deviation of the orbital plane with respect to our line of sight, $\xi(t) = \omega t + \xi_0$, ω is the orbital angular frequency and ξ_0 is the initial phase of the earth orbit which can be set to zero at the equinox. On the other hand relative displacement of the source with respect to position of the sun on the lens plane is given by

$$\beta_x^S = \beta(t) \cos \alpha - \beta_0 (\cos \alpha + \sin \alpha), \quad (\text{A.13})$$

$$\beta_y^S = \beta(t) \sin \alpha + \beta_0 (\cos \alpha - \sin \alpha), \quad (\text{A.14})$$

where the superscript represents relative motion of the source with respect to the sun and β_0 is the closest approach of the source to the sun and α is the angle between the trajectory of source

and semi-major axis of the earth orbit. Now combining these two motions the relative motions of the source with respect to the earth is given by

$$\vec{\beta}(t) = \vec{\beta}^{\delta}(t) - \vec{\pi}(t). \quad (\text{A.15})$$

The second term corresponds to a perturbation compared to the simple microlensing events. The photometric effect of parallax can be measured as an asymmetry of the microlensing light curve [10]. Also the astrometric effect of parallax is observable, if the lensed images are resolved. By combining equation (3.24) and (A.12) the distance of the lens from the observer (D_l) can be extracted. Substituting in the definition of θ_E , the mass of lens can be obtained. In fact the lens mass can be obtained only given θ_E and π using the relation:

$$\frac{M}{M_{\odot}} = \frac{1}{8\pi} \left(\frac{\theta_E}{1 \text{ mas}} \right). \quad (\text{A.16})$$

From photometric observations and microlensing light curves the physical parameter that is obtained by fitting to the light curve is the Einstein crossing time. Having measured this parameter the transverse velocity of lens can be derived as $v_t = D_l \theta_E / t_E$.

A.3 Parallax Effect in Microlensing Images

The acceleration of the Earth impacts the image separations for the same reason it affects the light curve. Here we estimate the size of this effect, explicitly determining the way in which the lens distance enters.

We begin by defining averaged and perturbed lens and source positions,

$$\mathbf{s} \equiv \mathbf{x}_S - \mathbf{x}_{\oplus} = \bar{\mathbf{s}} - \boldsymbol{\delta} \quad \text{and} \quad \boldsymbol{\ell} \equiv \mathbf{x}_L - \mathbf{x}_{\oplus} = \bar{\boldsymbol{\ell}} - \boldsymbol{\delta}, \quad (\text{A.17})$$

where $\bar{\mathbf{s}}$ and $\bar{\boldsymbol{\ell}}$ include the linear lens and source motions and the linear component of the Earth's motion. The acceleration of the Earth is contained in $\boldsymbol{\delta}$ which is by construction quadratic in time to lowest order. We will presume this to be small in comparison to $\bar{\mathbf{s}}$ and $\bar{\boldsymbol{\ell}}$.

Inserting these into equation (3.18) and linearizing in δ and ignoring terms of order β^2 yields

$$\begin{aligned}
\boldsymbol{\beta} &\equiv \left(\mathbf{1} - \frac{\boldsymbol{\ell}\boldsymbol{\ell}}{D_L^2} \right) \cdot \frac{\mathbf{s}}{D_S} \\
&\approx \left(\mathbf{1} - \frac{\bar{\boldsymbol{\ell}}\bar{\boldsymbol{\ell}}}{D_L^2} \right) \cdot \frac{\bar{\mathbf{s}}}{D_S} \\
&\quad + \left(\frac{D_S}{D_L} - 1 \right) \left(\mathbf{1} - \frac{\bar{\boldsymbol{\ell}}\bar{\boldsymbol{\ell}}}{D_L^2} \right) \cdot \frac{\boldsymbol{\delta}}{D_S} \\
&\quad + \frac{\bar{\boldsymbol{\ell}}}{D_L} \frac{\bar{\mathbf{s}}}{D_L} \cdot \left(\mathbf{1} - \frac{\bar{\boldsymbol{\ell}}\bar{\boldsymbol{\ell}}}{D_L^2} \right) \cdot \frac{\boldsymbol{\delta}}{D_S}.
\end{aligned} \tag{A.18}$$

Therefore, noting that $\boldsymbol{\beta} \cdot \boldsymbol{\ell} = 0$,

$$\beta \approx \beta_{np} + \left(\frac{D_S - D_L}{D_L} \right) \frac{\boldsymbol{\beta}_{np}}{\beta_{np}} \cdot \frac{\boldsymbol{\delta}}{D_S}, \tag{A.19}$$

where $\boldsymbol{\beta}_{np} \equiv \left(\mathbf{1} - \bar{\boldsymbol{\ell}}\bar{\boldsymbol{\ell}}/D_L^2 \right) \cdot \bar{\mathbf{s}}/D_S$ is the angular position of the lens ignoring the Earth's acceleration. The corresponding angular separation between the multiple images is then given by equation (A.4)

$$\Delta\theta \approx \Delta\theta_{np} + \left(\frac{D_S - D_L}{D_L} \right) \frac{\boldsymbol{\beta}_{np}}{\Delta\theta_{np}} \cdot \frac{\boldsymbol{\delta}}{D_S}. \tag{A.20}$$

Note that equation (A.19) interpolates between β_{np} and β as the lens distance changes. With the definition

$$\beta_p \equiv \beta_{np} + \frac{\boldsymbol{\beta}_{np}}{\beta_{np}} \cdot \frac{\boldsymbol{\delta}}{D_S}, \tag{A.21}$$

i.e., when $D_L = D_S/2$, we recover equation (3.27),

$$\beta = \lambda\beta_p + (1 - \lambda)\beta_{np}, \tag{3.27}$$

where

$$\lambda \equiv \frac{D_S - D_L}{D_L}. \tag{A.22}$$

Appendix B

Producing a local debiased CMD from the HIPPARCOS catalog

The HIPPARCOS catalog provides equatorial coordinates (α, δ) , apparent magnitudes $V_J (= V)$, color indexes $(B - V)_J$, $(V - I)$, and parallaxes π . To produce a local color-magnitude diagram, we calculate the absolute magnitudes M from the relative magnitudes and from the parallax, neglecting the local absorption. Fig. B.1 shows the distribution of these absolute magnitudes M_V and M_I as a function of the distance for the catalogued stars. It has been established ([135]) that the HIPPARCOS catalog is complete until apparent visual magnitude $V = 7.5$, *i.e.*, above the red curves of fig. B.1. This means that for a given absolute magnitude M_V , the catalog is complete up to the distance $d_c(M_V)$ associated with the distance modulus $\mu_c = 7.5 - M_V$; for example, within 50 pc the catalog is complete up to $M_V = 4.0$, which corresponds approximately to $M_I = 3.1$. Since we want to estimate the *local* CMD, we considered only those objects closer than 50 pc to avoid bias due to the very fast density variations with the distance to the Galactic plane. Fig. B.2 shows the full HIPPARCOS-Tycho M_I versus $V - I$ distribution and the distribution limited to stars within 50 pc (in red). It is clear that the full catalog is strongly biased in favor of bright (remote) objects.

To benefit from the whole statistics without suffering from selection bias, we calculate the differential volumic density of stars as a function of the absolute magnitude $0 < M_V < 6$ (interval chosen for statistical reasons, see next section) from the numbers of stars found within the corresponding completion distance

$$d_c(M_V) = 10\text{pc} \times 10^{\frac{\mu_c}{5}} = 10\text{pc} \times 10^{\frac{7.5 - M_V}{5}} \simeq 50\text{pc} \times 10^{\frac{4.0 - M_V}{5}}, \quad (\text{B.1})$$

divided by the corresponding completion volume $4\pi/3 \times d_c(M_V)^3$. Those stars that we accounted for lie above the completion (red) curve and between the two horizontal full lines in Fig. B.1. As




Figure B.1: HIPPARCOS absolute magnitudes vs. distance distributions (up= M_V , down= M_I). The red curves indicate the absolute magnitude completeness limit as a function of the distance. The vertical line shows our distance limit to get the local stellar population. The horizontal full lines at $M_V = 0$ and $M_V = 6$ correspond to the domain that contains enough stars from the HIPPARCOS catalog to enable our debiasing procedure.




Figure B.2: HIPPARCOS absolute color-magnitude diagram in M_{I_C} vs. $(V-I)_J$. The black squares correspond to the full catalog (statistically biased). The red squares correspond to the subsample of stars closer than 50pc; this subsample is statistically unbiased only for absolute magnitude $M_V < 4.0$ (corresponding to $M_I < 3.1$). The size scales are different between the red and black squares for readability.

we need a diagram that is representative of the solar neighborhood, we also consider only those stars that are inside a sphere of radius 50pc (left of the vertical line in Fig. B.1) to avoid depleted regions away from the Galactic median plane; indeed, as shown in Fig. B.3, the spatial 2D and 3D distributions of stars within 50pc distance of the catalog do not show global anisotropies. With all these constraints, a total of 2307 stars from the HIPPARCOS catalog are used to build our debiased local CMD. The upper panels of Fig. B.4 show the absolute magnitude and color




Figure B.3: Two-dimensional and 3D distributions of the HIPPARCOS objects within 50pc. The excess toward ($\alpha = 67^\circ$ $\delta = 16^\circ$) corresponds to the Hyades open cluster.

distributions of all the HIPPARCOS stars within 50pc (full lines) and of the stars that are closer than $\min(d_c(M_V), 50\text{pc})$, where $d_c(M_V)$ is the completion distance defined in Eq. (B.1) (dashed red lines).

We represent the HIPPARCOS catalog as a multi-dimensional distribution function defined by

$$f(\mathbf{x}, \mathbf{M}) = \sum_{catalog} \delta(\mathbf{M} - \mathbf{M}_i) \delta^3(\mathbf{x} - \mathbf{x}_i), \quad (\text{B.2})$$

where \mathbf{x}_i is the position of star i and \mathbf{M}_i represents its absolute magnitude and color “vector” (*i.e.*, its type). As explained above, to extract the unbiased local density for a given stellar type characterized by the vector \mathbf{M} (here (M_I, M_V)), we only account for the objects that are both

within the completion volume ($d < d_c(M_V)$) and closer than 50 pc, *i.e.*,

$$n(\mathbf{M}) = \frac{3}{4\pi \cdot \min[d_c(M_V), 50\text{pc}]^3} \int_{d < \min[d_c(M_V), 50\text{pc}]} f(\mathbf{x}, \mathbf{M}) k(d) d^3 x, \quad (\text{B.3})$$

where $k(d)$ is a correction factor that takes into account the variation of the density within the completion volume (this correction varies from 1 to 1.09).

B.1 Extrapolating the local HIPPARCOS CMD

The number of usable HIPPARCOS objects (closer than $\min(d_c(M_V), 50\text{pc})$) is statistically limited in the faint ($M_V > 6$) and bright ($M_V < 0$) ends, as can be seen in Fig. B.4 (upper left, dashed line). Moreover, there is no star with $M_V > 9$ within its corresponding completion distance $d_c(9) \simeq 5\text{pc}$ (*i.e.*, above the red curve of Fig. B.1), because the volume is too small; there is also no local star (within 50pc) brighter than $M_V = -3$.

Therefore, when building a debiased density color-magnitude diagram, we need to examine specifically the contribution of the stars with absolute M_V magnitudes out of $[0, 6]$ range to avoid statistical limitations or biases:

- First, we can neglect the contribution of the brightest stars; indeed, the HIPPARCOS catalog contains only 35 stars brighter than $M_V < 0$ within 50pc (complete sample). This corresponds to a maximum contribution of

$$35 \times \left[\frac{10\text{kpc}}{50\text{pc}} \right]^3 \times \frac{\Omega(1^\circ \times 1^\circ)}{4\pi} \sim 6800 \text{ stars/sq. deg.} \quad (\text{B.4})$$

of $M_V < 0$ stars within 10kpc distance (typically less than 2 – 3% toward the directions studied in this paper). This contribution will be neglected in the following discussions¹.

- Stars fainter than $M_V = 6$ have a minor, but not negligible contribution to a deep Galactic exposure. Instead of debiasing the statistically limited subsample of the HIPPARCOS catalog, we choose to linearly extrapolate the local stellar density of these faint stars as (see Fig. B.4 middle, left):

$$\frac{dn}{dM_V} = \text{Const.} + 4.6 \times 10^{-4} M_V \text{ (pc}^{-3} \text{ mag}^{-1}\text{)}. \quad (\text{B.5})$$

¹For this very conservative estimate, we assume a constant density along the line of sight, we neglect the absorption, and we assume a 100% detection efficiency.

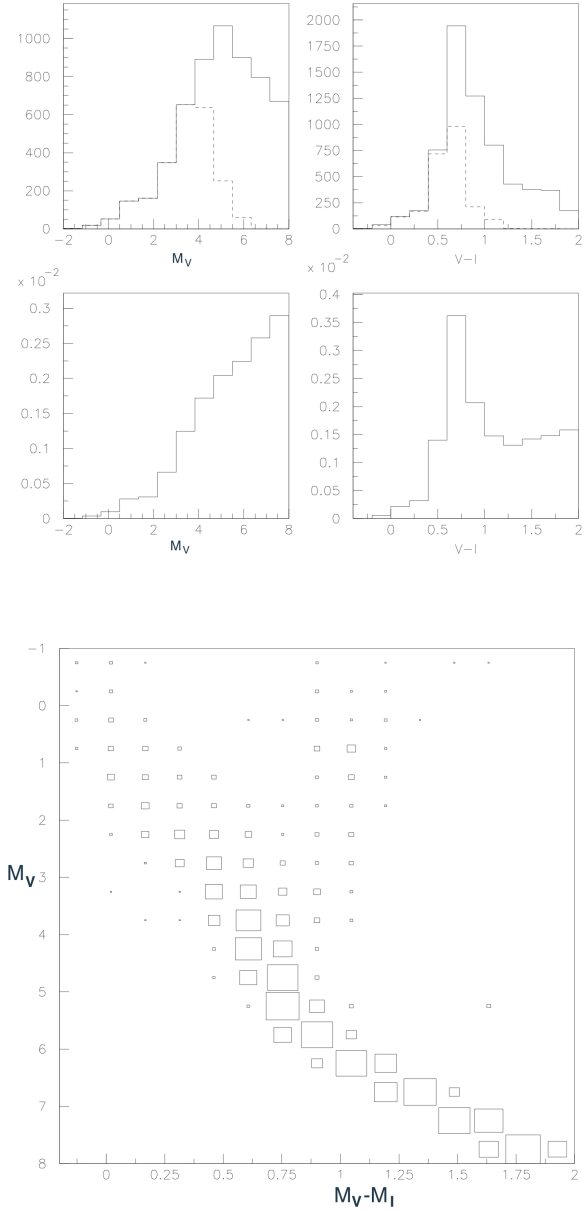


Figure B.4: Top: Raw distributions of M_V and $(V - I)$ of all HIPPARCOS stars within 50pc (6911 objects). The dashed lines show the numbers of stars within the completion volume corresponding to their magnitude (see text). Middle: Local debiased volumic density of stars (per magnitude unit, in pc^{-3}) estimated from the ratio of stars within the completion volume and extrapolated beyond $M_V = 6$. Bottom: Debiased M_V vs. $V - I$ stellar density of stars closer than 200

Since we deduce from Fig. B.2 (lower right branch) that

$$V - I \sim 0.47 \times (M_I - 3.97), \quad (\text{B.6})$$

or equivalently

$$V - I \sim 0.33 \times (M_V - 4.0) \quad (\text{B.7})$$

the type of these faint stars is also completely extrapolated.

Fig. B.4 (bottom) shows the local CMD obtained following our complete procedure using the HIPPARCOS stars with $0 < M_V < 6$ within $d_c(M_V)$ and our extrapolated distribution for $6 < M_V < 8$.

B.2 Comparison with the stellar density expected from the mass function: A coherence check

We can crosscheck the stellar number density found from the HIPPARCOS catalog and the density expected from the mass function as follows: stars with $0 < M_V < 6$ belong to the mass domain defined by $0.85M_\odot < m < 2.8M_\odot$ ([73]). The local number density of objects within this mass range is given by

$$n(0.85M_\odot < m < 2.8M_\odot) = \int_{0.85}^{2.8} \frac{dn}{dm} dm, \quad (\text{B.8})$$

where $\frac{dn}{dm}$ is the stellar mass function in the solar neighborhood. We use the mass function $\xi(\log m/M_\odot) = \frac{dn}{d \log m/M_\odot}$ of [53], revised in [55],

$$\begin{aligned} \xi(\log m/M_\odot) &= 0.093 \times \exp \left[\frac{-(\log m/0.2M_\odot)^2}{2 \times (0.55)^2} \right], m \leq M_\odot \\ &= 0.041(m/M_\odot)^{-1.35}, m > M_\odot \end{aligned} \quad (\text{B.9})$$

(see fig 4.8). We find that the mean density of disk stars with $0 < M_V < 6$ in a sphere of 50pc centered on the sun (located at 26 pc from the disk plane ([173])) is 0.012pc^{-3} . This is compatible with the estimates from the integral of the M_V debiased distribution of the volumic density of stars plotted in Fig. B.4, $n_{\text{HIPPARCOS}} = 0.0076 \text{pc}^{-3}$, when taking into account the fact that $\sim 2/3$ of the stars are in binary systems ([55]) not deblended in the HIPPARCOS observations.

B.3 Parameters of the Besançon Galactic model

The Sun is located at $R_{\odot} = 8.0\text{kpc}$ and $z_{\odot} = 15\text{ pc}$, which is different than in our simple model. The thin disk structures are parametrized in cylindrical galactocentric coordinates (r, z) , and for various ranges of age, as follows:

$$\begin{aligned} \rho_D(r, z)_{age} &\propto \left[\exp\left(-\frac{a^2}{R_d^2}\right) - \exp\left(-\frac{a^2}{R_h^2}\right) \right] \text{ if age} < 0.15 \text{ Gyr,} \\ &\propto \left[\exp\left(-\sqrt{0.25 + \frac{a^2}{R_d^2}}\right) - \exp\left(-\sqrt{0.25 + \frac{a^2}{R_h^2}}\right) \right], \\ &\text{if age} > 0.15 \text{ Gyr,} \end{aligned} \quad (\text{B.10})$$

where

- $R_d = 5.0\text{ kpc}$ and $R_h = 3.0\text{ kpc}$ if age $< 0.15\text{ Gyr}$,
- $R_d = 2.17\text{ kpc}$ and $R_h = 1.33\text{ kpc}$ if age $> 0.15\text{ Gyr}$,
- $a^2 = r^2 + (z/\epsilon_{age})^2$;
- ϵ_{age} and the local mass densities corresponding to $\rho_D(r_{\odot}, z_{\odot})_{age}$ values are given in table B.3 for the different ranges of stellar age, together with the IMFs.

The thick disk contribution is expressed by

$$\begin{aligned} \rho_D^{thick}(r, z) &= \rho_D^{thick}(r_{\odot}, z_{\odot}) \\ &\times \left(1 - \frac{z^2}{x_l(2h_z + x_l)}\right) \exp\left[-\frac{r - R_{\odot}}{R_{thick}}\right] \text{ if } |z| < x_l, \\ &\times \frac{\exp(x_l/h_z)}{1 + x_l/2h_z} \exp\left[-\frac{|z|}{h_z}\right] \exp\left[-\frac{r - R_{\odot}}{R_{thick}}\right] \text{ if } |z| > x_l, \end{aligned} \quad (\text{B.11})$$

where $x_l = 400\text{pc}$, $h_z = 800\text{pc}$ and $R_{thick} = 2.5\text{Kpc}$. Table B.3 also gives the total local density $\rho_D^{thick}(r_{\odot}, z_{\odot})$ for the thick disk together with the IMF.

The two components of the bar are described in a Cartesian frame positioned at the Galactic center with the major axis X tilted by $\Phi = 12.8$ degree with respect to the Galactic center-Sun direction. The mass density for each component of the bar is given by ([239])

$$\rho_{bar1}(X, Y, Z) = \rho_0 \text{sech}^2(-R_s) \times f_c(X, Y) \quad (\text{B.12})$$

$$\rho_{bar2}(X, Y, Z) = \rho_0 \exp(-R_s) \times f_c(X, Y), \quad (\text{B.13})$$

where

$$R_s^{C_p} = \left[\left| \frac{X}{a} \right|^{C_n} + \left| \frac{Y}{b} \right|^{C_n} \right]^{\frac{C_p}{C_n}} + \left| \frac{Z}{c} \right|^{C_p}, \quad (\text{B.14})$$

and f_c is a cutoff function

$$\begin{aligned} f_c(X, Y) &= 1. \quad \text{if } X^2 + Y^2 < R_c^2, \\ &= \exp \left[-\frac{(\sqrt{X^2 + Y^2} - R_c)^2}{0.25 \text{ kpc}^2} \right] \quad \text{if } X^2 + Y^2 > R_c^2. \end{aligned} \quad (\text{B.15})$$

The parameters for bar1 are $\rho_0 = 9.21 \text{ M}_\odot \cdot \text{pc}^{-3}$ ⁽²⁾, $a = 1.46 \text{ kpc}$, $b = 0.49 \text{ kpc}$, $c = 0.39 \text{ kpc}$, $R_c = 3.43 \text{ kpc}$ are the scale length factors and $C_p = 3.007$, $C_n = 3.329$. The total mass of this bar is $35.45 \times 10^9 \text{ M}_\odot$

The parameters for bar2 are $\rho_0 = 0.026 \text{ M}_\odot \cdot \text{pc}^{-3}$, $a = 4.44 \text{ kpc}$, $b = 1.31 \text{ kpc}$, $c = 0.80 \text{ kpc}$, $R_c = 6.83 \text{ kpc}$ are the scale length factors, and $C_p = 2.786$, $C_n = 3.917$. The total mass of this bar is $2.27 \times 10^9 \text{ M}_\odot$

The IMF for these two bars is $dn/dm \propto (m/M_\odot)^{-2.35}$.

	Age (Gyr)	$\rho(r_\odot, z_\odot)$ ($\text{M}_\odot \text{pc}^{-3}$)	ϵ	IMF
disk	0-0.15	4.0×10^{-3}	0.0140	
	0.15-1	7.9×10^{-3}	0.0268	
	1-2	6.2×10^{-3}	0.0375	$dn/dm \propto (m/M_\odot)^{-\alpha}$
	2-3	4.0×10^{-3}	0.0551	$\alpha = 1.6$ for $m < 1 \text{ M}_\odot$
	3-5	5.8×10^{-3}	0.0696	$\alpha = 3.0$ for $m > 1 \text{ M}_\odot$
	5-7	4.9×10^{-3}	0.0785	
	7-10	6.6×10^{-3}	0.0791	
	WD	3.96×10^{-3}		
Thick disk	all	1.64×10^{-3}		$dn/dm \propto (m/M_\odot)^{-0.5}$

Table B.1: Age, local mass density $\rho(r_\odot, z_\odot)$, disk axis ratio ϵ , and IMF of the different stellar components of the disks in the Besançon model. WD represents the white dwarfs.

²Not to be confused with the local density of the bar $\rho_{bar}(r_\odot, z_\odot)$

As far as kinematics is concerned, we use the ellipsoids of velocity dispersions provided for each structure and age in table 4 of ([240]).

Appendix C

Prediction Accuracy Requirements

When predictions are made numerically frequently the computational expense is strongly dependent on the accuracy of the theoretical estimate required. Thus, significant efficiencies can be realized by understanding and limiting the accuracy requested where possible. Generally, comparisons with data with large uncertainties require far less accurate theoretical estimates than with data that has small uncertainties. Here we determine the relationship with parameter estimation uncertainty and the accuracy of the theoretical estimates, thereby estimating the accuracy required by THEMIS.

We begin by assuming that the measurement errors are Gaussian. We further assume that the posterior parameter probability is also nearly Gaussian, and thus adopt a Fisher matrix approach to the estimation of the uncertainty of the parameter estimates. Finally, we assume that errors in the predicted values are Gaussian and uncorrelated. That is, we set the log-likelihood to

$$\mathcal{L} = - \sum_j \frac{(f_j + \delta_j - y_j)^2}{2\sigma_j^2}, \quad (\text{C.1})$$

where the predicted value is f_j , δ_j is the error in the predicted value, y_j are the data, and σ_j are the observational uncertainties.

The assumption that the prediction errors are Gaussian corresponds to assuming that the δ_j are Gaussian random variables. This is not true in an absolute sense: each time a prediction is made for the same independent variables the δ_j does not change. However, in a statistical sense we are assuming that at different independent variable values and for different parameter values the δ_j are well approximated by a random variable. It will be useful henceforth to characterize the size of the distribution of the δ_j in terms of σ_j , i.e., we set the variance of the prediction error

in terms of the measurement uncertainty as $\langle \delta_j^2 \rangle = \Delta^2 \sigma_j^2$. This implies that higher prediction accuracy is possible for more accurately measured quantities.

The δ_j modify the minimum \mathbb{L} (and thus χ^2) expected: averaging over realizations of the data and the prediction errors,

$$\langle \mathbb{L} \rangle = - \sum_j \left\langle \frac{(f_j - y_j)^2 + \delta_j^2}{2\sigma_j^2} \right\rangle \approx -\frac{N}{2} (1 + \Delta^2), \quad (\text{C.2})$$

where we have further assumed the number of degrees of freedom is close to the number of data points, N . For a sufficiently large number of degrees of freedom and a large enough Δ^2 , the deviation will be statistically noticeable in the reduced χ^2 when $\Delta \gtrsim (8/N)^{1/4}$. This expression grows slowly with N , though, and is therefore not a fundamental limit as far as parameter estimation is concerned.

The uncertainty in the estimate of the parameters is set by the inverse of the covariance matrix, given by

$$C_{ab} = - \frac{\partial^2 \mathbb{L}}{\partial p_a \partial p_b} = \sum_j \frac{1}{\sigma_j^2} \left[\frac{\partial f_j}{\partial p_a} \frac{\partial f_j}{\partial p_b} + (f_j - y_j + \delta_j) \frac{\partial^2 f_j}{\partial p_a \partial p_b} \right]. \quad (\text{C.3})$$

Again, averaging over realizations of the data and the prediction errors gives

$$\langle C_{ab} \rangle = \sum_j \frac{1}{\sigma_j^2} \frac{\partial f_j}{\partial p_a} \frac{\partial f_j}{\partial p_b}. \quad (\text{C.4})$$

However, a typical value will be modified by the presence of the linear term. That is, the variance in the inverse covariance is

$$\langle (C_{ab})^2 \rangle - \langle C_{ab} \rangle^2 = \sum_j (1 + \Delta^2) \left(\frac{\partial^2 f_j}{\partial p_a \partial p_b} \right)^2. \quad (\text{C.5})$$

The covariance matrix, whose eigenvalues indicate the magnitude of the uncertainties, is given by

$$C_{ab}^{-1} = [\langle C_{ab} \rangle + \delta C_{ab}]^{-1} \approx \langle C_{ab} \rangle^{-1} \left[1 - \langle C_{ab} \rangle^{-1} \delta C_{ab} \right] \quad (\text{C.6})$$

where

$$\delta C_{ab} = \mu \sqrt{\sum_j (1 + \Delta^2) \left(\frac{\partial^2 f_j}{\partial p_a \partial p_b} \right)^2} \approx \mu \left(1 + \frac{\Delta^2}{2} \right) \sqrt{\sum_j \left(\frac{\partial^2 f_j}{\partial p_a \partial p_b} \right)^2} \quad (\text{C.7})$$

in which μ is a Gaussian random variable with unit variance. This error term is suppressed by approximately a factor of $N^{-1/2}$ relative to the mean variance, and thus in the limit of large N becomes insignificant.

The error term contains two elements, associated with the measurement and prediction errors, respectively. The ratio of the latter to the former is $\Delta^2/2$. When Δ^2 is small and N is large, the uncertainty on the parameter estimates then grows by a multiplicative factor of $\Delta^2/4$. This is unconditionally small when Δ is small.

In **THEMIS**, we typically set $\Delta = 0.25$, for which $\Delta^2/4 = 0.016$, which broadens the posterior parameter distributions by 1.6%. This does complicate the interpretation of fit quality for $N \gtrsim 2048$, however.

C.1 Error Distributions of Quantities Associated with Visibility Amplitudes

THEMIS has three data types associated with visibility amplitudes: visibility amplitudes themselves, interferometric polarization fractions, and closure amplitudes. The underlying error distributions of none of these is Gaussian, and the latter two poorly approximated by Gaussians. Here, we summarize what the relevant error distributions are and quantify how well they are approximated in **THEMIS**. In all cases, we will assume that the complex visibilities are well described by a Gaussian random variable with non-zero mean.

C.2 Visibility Amplitudes

The probability distribution of the magnitude of a complex Gaussian random variable, V , with mean V_0 and standard deviation σ is given by the Rice distribution [see, e.g., 271]:

$$p_r(|V|; |V_0|, \sigma) = \frac{|V|}{\sigma^2} e^{-(|V|^2 + |V_0|^2)/2\sigma^2} I_0\left(\frac{|V||V_0|}{\sigma^2}\right). \quad (\text{C.8})$$

At high SNR (defined here by $|V_0|/\sigma$), a Gaussian with mean $\sqrt{|V_0|^2 + \sigma^2}$ and standard deviation σ becomes an increasingly good proxy for the Rice distribution, with the quality of this approximation increasing with SNR. When $\text{SNR} \geq 2$, the biased Gaussian is within 8% of the maximum probability of the Rice distribution for all values of $|V|$. These are compared in the left panel of Figure C.1 for various choices of SNR. We provide a set of quantitative estimates of the accuracy of the of Gaussian approximation in Table C.1 for various SNRs.




Figure C.1: Left: Comparison of the Rice and Gaussian distributions at various SNRs. Right: Distribution of products of visibility amplitudes and their comparison to a single Rice distribution. The SNR indicates $\sigma/W = \sqrt{\sigma_A^2/|V_{A,0}|^2 + \sigma_B^2/|V_{B,0}|^2}$, with the range for each SNR value corresponding to that obtained from different ways of apportioning the errors between $\sigma_A/|V_{A,0}|$ and $\sigma_B/|V_{B,0}|$. Summaries of the accuracy of the comparisons shown can be found in Tables C.1 and C.2.

Approx.	Notes	Error	SNR				
			1	2	4	8	32
Gauss	—	δ_{\max}	29%	8%	2%	0.4%	0.02%
—	—	δ_{mode}	7%	1%	< 0.1%	< 0.1%	< 0.1%
—	—	δ_w	15%/17%	5%/6%	2%/2%	< 0.5%/0.5%	< 2%/ < 1%

Table C.1: Accuracy of Visibility Amplitude Error Distribution Approximations. δ_{\max} is the maximum absolute difference, measured relative to probability maximum. δ_{mode} represents the fractional error in the location of the mode and δ_w is the Fractional error in the width of the region containing 68%/95% of the cumulative probability.

C.2.1 Visibility Amplitude Products

Before discussing the data quantities of interest, we begin by considering the distribution of the product of visibility amplitudes, i.e., $W = |V_A||V_B|$. We do this both to illustrate the procedure by which we construct exact probability distributions for combinations of products and quotients of visibility amplitudes and to show explicitly that these are typically well-approximated by a single Rice distribution.

We begin by exploiting the non-negative behavior of $|V|$ to define $v = \log(|V|)$. This simplifies the construction of the product by reducing it to a sum, i.e., in terms of v_A and v_B , $W = e^{v_A+v_B} \equiv e^w$. The probability distribution of v is given in terms of the Rice distribution by

$$q_r(v; |V_0|, \sigma) = e^v p_r(e^v; |V_0|, \sigma), \quad (\text{C.9})$$

which we will call the logarithmic Rice distribution. Its characteristic function is

$$\phi_r(k; |V_0|, \sigma) = \int_{-\infty}^{\infty} e^{-ikv} q_r(v; |V_0|, \sigma) dv. \quad (\text{C.10})$$

In practice, this may be computed efficiently via FFT. In terms of ϕ_r , the characteristic function of the probability distribution of w is, $\phi_p(k; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B) = \phi_r(k; |V_{A,0}|, \sigma_A) \phi_r(k; |V_{B,0}|, \sigma_B)$, and thus, the probability distribution of w is

$$q_p(w; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikw} \phi_r(k; |V_{A,0}|, \sigma_A) \phi_r(k; |V_{B,0}|, \sigma_B) dk, \quad (\text{C.11})$$

which again may be computed efficiently via FFT. Finally, the desired probability distribution of W is then

$$p_p(W; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B) = W^{-1} q_p(\log W; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B). \quad (\text{C.12})$$

These are shown in the right panel of Figure C.1.

While formally, the distribution of W is characterized by four parameters, in practice it is well approximated by a single Rice distribution with $W_0 = |V_{A,0}||V_{B,0}|$ and $\sigma = \sqrt{\sigma_A^2 + \sigma_B^2}$, differing by 17% of the maximum probability for $\text{SNR} \geq 2$. This comparison is also shown in the right panel of Figure C.1, and estimates of the accuracy of the approximation are tabulated in Table C.2 for various SNRs. Note that this also implies that the product distribution is well fit by a Gaussian for sufficiently high SNR.

Approx.	Notes	Error	SNR				
			1	2	4	8	32
Rice	SNR ratio=1	δ_{\max}	33%	17%	9%	4%	1%
-	-	δ_{mode}	36%	14%	4%	1%	< 0.1%
-	-	δ_w	16%/21%	7%/10%	2%/2%	0.6%/0.6%	< 2%/ < 1%
Rice	SNR ratio=4	δ_{\max}	11%	5%	2%	1%	0.2%
-	-	δ_{mode}	13%	4%	1.0%	0.2%	< 0.1%
-	-	δ_w	4%/9%	1%/3%	0.4%/0.7%	< 0.5%/ < 0.2%	< 2%/ < 1%

Table C.2: Accuracy of Product Error Distribution Approximations. SNR ratio is the Ratio of $V_{A,0}/\sigma_A$ to $V_{B,0}/\sigma_B$ in the construction of the product distribution. δ_{\max} is the maximum absolute difference, measured relative to probability maximum. δ_{mode} represents the fractional error in the location of the mode and δ_w is the fractional error in the width of the region containing 68%/95% of the cumulative probability.

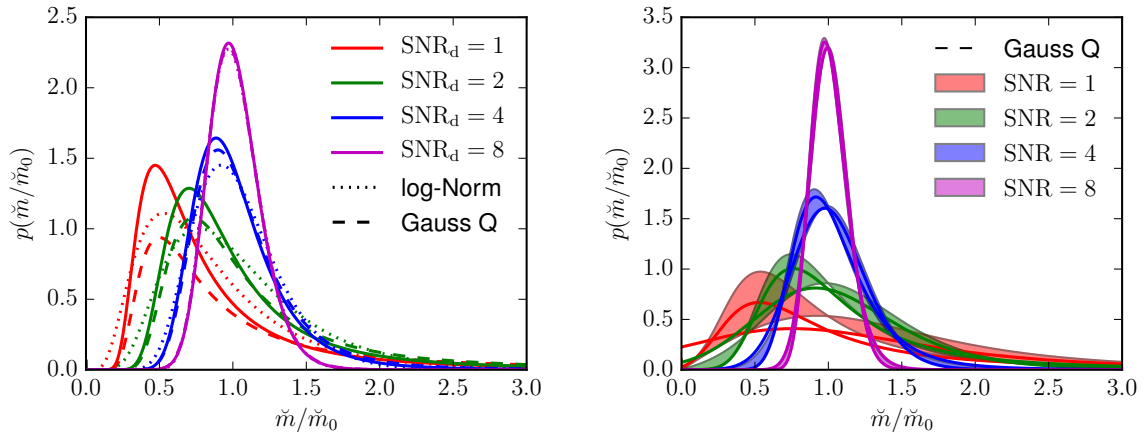


Figure C.2: Left: Comparison of the polarization fraction distribution (visibility amplitude quotient distribution) with the Gaussian quotient distribution approximate (Gauss Q) for various input values of the denominator SNR, SNR_d . In all cases the numerator SNR was set to 8. For comparison, a log-Normal distribution is also shown. Right: Comparison of the polarization fraction distribution (boundaries of the shaded region), the Gauss Q approximation in Equation (C.17), and the log-Normal model, for various total SNR (i.e., $\check{m}/\sigma_{\check{m}}$) at $\check{m} = 0.5$ and 2.0 . Quantitative estimates of the accuracy of the various approximations can be found in Table C.3.

C.3 Polarization Fractions — Visibility Amplitude Quotients

We follow a similar procedure to that in the previous section to compute the distribution of the quotient of visibility amplitudes, i.e., $Q = |V_A|/|V_B|$. Unfortunately, we find that the remarkable simplicity of the distribution of visibility products does not extend to quotients. The characteristic function of the logarithmic quotient distribution is $\phi_q(k; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B) = \phi_r(k; |V_{A,0}|, \sigma_A) \phi_r^*(k; |V_{B,0}|, \sigma_B)$, where the $*$ denotes complex conjugation. The resulting logarithmic quotient probability distributions is

$$q_q(q; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikw} \phi_r(k; |V_{A,0}|, \sigma_A) \phi_r^*(k; |V_{B,0}|, \sigma_B), \quad (\text{C.13})$$

with corresponding quotient probability distribution

$$p_q(Q; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B) = Q^{-1} q_q(\log Q; |V_{A,0}|, \sigma_A, |V_{B,0}|, \sigma_B). \quad (\text{C.14})$$

This is directly applicable to the polarization fraction, for which this distribution is shown in the left panel of Figure C.2.

The polarization fraction distribution clearly deviates from the Gaussian and Rice distributions in two key respects. First, even at high SNR, the distributions are asymmetric, with the probability maximum lying below $Q_0 = |V_{A,0}|/|V_{B,0}|$. Second, there is a significant tail extending to high values of Q , containing sufficient weight to move the average Q above Q_0 for $\text{SNR} \geq 2$. More accurate are fitted (i.e., same mean and standard deviation) log-normal approximations, shown in the left panel of Figure C.2, which recover the asymmetry, though still exhibit deviations for $\text{SNR} \leq 4$.

Combined with the accuracy of the Gaussian approximation to the Rice distribution, this motivates an exploration of better approximates to the quotient distribution of visibility amplitudes. For two Gaussian variables, with non-zero means, it is possible to analytically construct the quotient distribution analytically:

$$p_{\text{GaussQ}}(r) = \int dx dy \frac{e^{-(x-\bar{x})/2\sigma_x^2}}{\sqrt{2\pi}\sigma_x} \frac{e^{-(y-\bar{y})^2/2\sigma_y^2}}{\sqrt{2\pi}\sigma_y} \delta\left(r - \frac{x}{y}\right) = \frac{e^{-(r-r_0)^2/2\Sigma^2}}{\sqrt{2\pi}\Sigma} \left[\Delta \text{erf}\left(\frac{\Delta}{\sqrt{2}\Omega}\right) - \frac{2\Omega}{\sqrt{2\pi}} e^{-\Delta^2/2\Omega^2} \right]$$

$$\text{where } r_0 \equiv \frac{\bar{x}}{\bar{y}}, \Sigma^2 \equiv r_0^2 \frac{\sigma_x^2}{\bar{x}^2} + r^2 \frac{\sigma_y^2}{\bar{y}^2}, \Delta \equiv \Sigma^{-2} \left(\frac{\sigma_x^2}{\bar{x}^2} + \frac{r}{r_0} \frac{\sigma_y^2}{\bar{y}^2} \right), \Omega \equiv \frac{\sigma_x \sigma_y}{\Sigma \bar{x} \bar{y}}, \quad (\text{C.15})$$

and $\text{erf}(x)$ is the standard error function. This is also shown in the left panel of Figure C.2. While deviations from p_q continue to exist, the Gaussian quotient (Gauss Q) distribution accurately reproduces the large high- Q tail.

As with $p_p(W)$, $p_q(Q)$ depends on the SNR of both the numerator and denominator independently. Unlike the product distribution, the quotient distribution is not symmetric in this dependency, with the properties of the denominator controlling the asymmetry and tail. Therefore, characterizing this distribution by a single pair of numbers — a central value and width — will result in a substantial uncertainty in the resulting quotient distribution. This is simplified for polarization fractions by the fact that the stations used to construct the visibility amplitudes in the numerator and denominator are the same, and thus both quantities have similar noise profiles in principle, i.e., $\sigma_A \approx \sigma_B$. This implies that

$$\frac{\sigma_A}{|V_{A,0}|} \approx \frac{1}{\sqrt{1+\check{m}_0^2}} \frac{\sigma_{\check{m}}}{\check{m}_0} \quad \text{and} \quad \frac{\sigma_B}{|V_{B,0}|} \approx \frac{\check{m}_0}{\sqrt{1+\check{m}_0^2}} \frac{\sigma_{\check{m}}}{\check{m}_0}, \quad (\text{C.16})$$

where $\check{m}_0 = |V_{A,0}|/|V_{B,0}|$ and $\sigma_{\check{m}}$ is the uncertainty obtained by the standard error propagation formula. Note that these are the only two quantities required to fully specify the Gauss Q and quotient distributions. Combining this with Equation (C.15), we obtain for the polarization fraction,

$$p_{\check{m}}(\check{m}; \check{m}_0, \sigma_{\check{m}}) = \frac{e^{-(\check{m}-\check{m}_0)^2/2\Sigma_{\check{m}}^2}}{\sqrt{2\pi}\Sigma_{\check{m}}} \left[\Delta_{\check{m}} \operatorname{erf}\left(\frac{\Delta_{\check{m}}}{\sqrt{2}\Omega_{\check{m}}}\right) - \frac{2\Omega_{\check{m}}}{\sqrt{2\pi}} e^{-\Delta_{\check{m}}^2/2\Omega_{\check{m}}^2} \right] \quad (\text{C.17})$$

where $\Sigma_{\check{m}}^2 = \sigma_{\check{m}}^2 \frac{1+\check{m}^2}{1+\check{m}_0^2}$, $\Delta_{\check{m}} = \frac{1+\check{m}\check{m}_0}{1+\check{m}^2}$, $\Omega_{\check{m}}^2 = \frac{\sigma_{\check{m}}^2}{(1+\check{m}^2)(1+\check{m}_0^2)}$.

For various $\check{m}_0/\sigma_{\check{m}}$, the right panel of Figure C.2 shows comparisons of the exact quotient and Gauss Q distributions assuming the errors in Equation (C.16) for $\check{m}_0 = 0.5$ and 2.0. The Gauss Q distribution is within 11% and 13% of the maximum probability of the exact quotient distribution, respectively, at all values of \check{m} when $\text{SNR} \geq 2$. The accuracy of the Gauss Q approximation for the polarization fraction distribution is tabulated for different SNRs and \check{m} in Table C.3.

C.4 Closure Amplitudes

We now turn to the problem of constructing the error distribution for closure amplitudes generally. Again, the exact expression can be constructed using the characteristic functions of the logarithmic Rice distributions:

$$\phi_a(k; |V_A|, \sigma_A, |V_B|, \sigma_B, |V_C|, \sigma_C, |V_D|, \sigma_D) = \phi_r(k; |V_{A,0}|, \sigma_A) \phi_r(k; |V_{B,0}|, \sigma_B) \phi_r^*(k; |V_{C,0}|, \sigma_C) \phi_r^*(k; |V_{D,0}|, \sigma_D), \quad (\text{C.18})$$

Approx.	Notes	Error	SNR				
			1	2	4	8	32
Gauss Q	$\check{m}_0 = 0.5$	δ_{\max}	41%	11%	5%	2%	0.6%
–	–	δ_{mode}	22%	7%	2%	0.5%	< 0.1%
–	–	δ_w	3%/13%	5%/4%	2%/2%	< 0.5%/0.4%	< 2%/ < 1%
Gauss Q	$\check{m}_0 = 2.0$	δ_{\max}	31%	13%	6%	3%	0.6%
–	–	δ_{mode}	2%	2%	1%	0.4%	< 0.1%
–	–	δ_w	17%/20%	13%/18%	5%/6%	1%/1%	< 2%/ < 1%
log-Norm	$\check{m}_0 = 0.5$	δ_{\max}	41%	17%	10%	5%	1%
–	–	δ_{mode}	26%	0.2%	0.4%	< 0.1%	< 0.1%
–	–	δ_w	46%/27%	21%/29%	5%/6%	1%/2%	< 2%/ < 1%
log-Norm	$\check{m}_0 = 2.0$	δ_{\max}	68%	48%	27%	13%	3%
–	–	δ_{mode}	27%	23%	9%	3%	0.2%
–	–	δ_w	> 100%/63%	51%/23%	12%/3%	3%/0.6%	< 2%/ < 1%

Table C.3: Accuracy of Polarization Fraction Error Distribution Approximations. δ_{\max} is the maximum absolute difference, measured relative to probability maximum. δ_{mode} represents the fractional error in the location of the mode and δ_w is the fractional error in the width of the region containing 68%/95% of the cumulative probability.

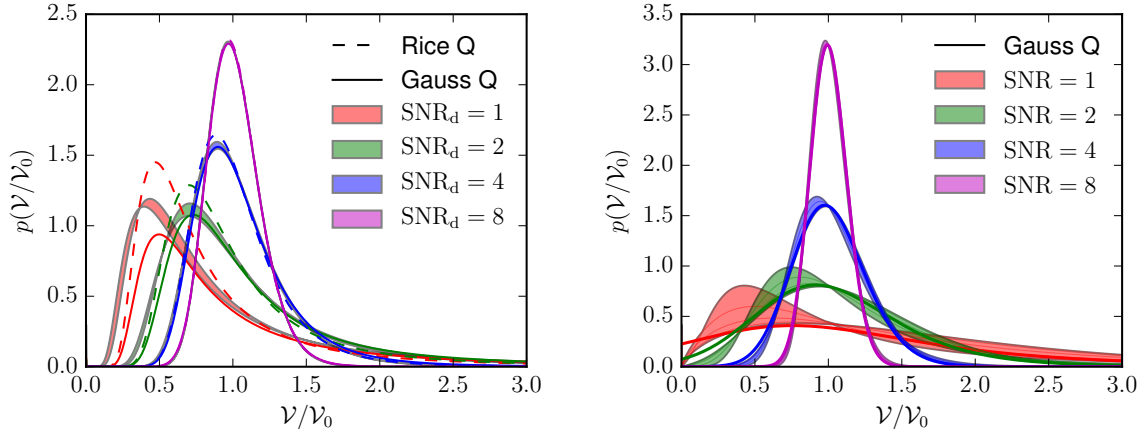


Figure C.3: Left: Comparison of the closure amplitude distribution with the Gauss Q approximation in Equation (C.24) for various choices of the denominator SNR, SNR_d . In all cases the numerator SNR was set to 8, divided equally among the visibility amplitudes in the numerator. For comparison, the visibility amplitude quotient (Rice Q) distribution is also shown. The range of the filled bands indicate the uncertainty associated with various choices of how SNR_d is apportioned between the two visibility amplitudes in the denominator. In all cases, the value of ρ was set to the proper value in all models. Right: Comparison of the closure amplitude distribution and the Gauss Q approximation with $\rho = 1$, for $V_0 = 0.5$. The colored bands indicate the range of closure amplitude distributions when the true ρ is varied within the permissible range for the 2017 EHT campaign, $[0.3, 0.33]$. The SNR within the numerator and denominator is distributed uniformly. Quantitative estimates of the accuracy of the various approximations for illustrative cases can be found in Table C.4.

from which we obtain

$$q_a(a; |V_{A,0}, \sigma_A, \dots) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ika} \phi_a(k; |V_{A,0}, \sigma_A, \dots), \quad (\text{C.19})$$

and

$$p_{\mathcal{V}}(\mathcal{V}; |V_{A,0}, \sigma_A, \dots) = \mathcal{V}^{-1} q_a(\log \mathcal{V}; |V_{A,0}, \sigma_A, \dots). \quad (\text{C.20})$$

These are shown in the left panel of Figure C.3. Similar to the polarization fractions, they are clearly asymmetric and exhibit large tails to high values, typical of quotient distributions.

Formally, this requires knowledge of eight values to define. However, again, it is possible to accurately estimate $p_{\mathcal{V}}$ with only a handful of combinations of these value. Due to the similarity between the amplitude product distribution (Section C.2.1) and the Rice distribution, both the numerator and the denominator can be effectively described by only two parameters each. As a result, the closure amplitude distribution is similar to the amplitude quotient distribution described in Section C.3, shown by the dashed lines in the left panel of Figure C.3.

If the SNRs of the denominator and numerator are independently known, this is well approximated by the Gauss Q distribution in Equation (C.15). These may be reconstructed with knowledge of the total SNR ($\mathcal{V}/\sigma_{\mathcal{V}}$), and the ratio of the thermal uncertainties in the numerator and denominator, i.e., $\sigma_n = \sqrt{\sigma_A^2 + \sigma_B^2}$ and $\sigma_d = \sqrt{\sigma_C^2 + \sigma_D^2}$; also shown in the left panel of Figure C.3. The latter are not independent, related by the repeated presence of each of the four stations required to produce a closure amplitude in the numerator and denominator. That is, identifying the baselines A , B , C , and D , with stations 1 and 2, 3 and 4, 1 and 4, and 2 and 3, respectively, assuming identical bandwidths and scan duration,

$$\rho^2 \equiv \frac{\sigma_n^2}{\sigma_d^2} = \frac{S_1 S_2 + S_3 S_4}{S_1 S_4 + S_2 S_3} = \frac{1 + (S_3/S_1)(S_4/S_2)}{(S_4/S_2) + (S_3/S_1)}, \quad (\text{C.21})$$

where the S_j is are station specific SEFDs. Despite the appearance of four SEFDs, this is a function of only two variables: the ratios S_3/S_1 and S_4/S_2 . Where both of these ratios are of order unity, i.e., for a homogeneous array, $\rho \approx 1$. For highly heterogeneous arrays, in which *more than one* station is much more sensitive than the others, this can deviate from unity substantially, by an amount that depends on the second lowest and second highest SEFDs, $S_{2\text{nd min}}$ and $S_{2\text{nd max}}$, respectively:

$$\sqrt{\frac{S_{2\text{nd min}}}{S_{2\text{nd max}}}} \lesssim \rho \lesssim \sqrt{\frac{S_{2\text{nd max}}}{S_{2\text{nd min}}}}. \quad (\text{C.22})$$

For the 2017 EHT campaign, the station SEFDs ranged from 90 Jy to 6000 Jy, with most being near 5000 Jy. The phased Atacama Large Millimeter/submillimeter Array (ALMA) is

an extremely low-noise outlier at 90 Jy, followed by the Large Millimeter Telescope at 600 Jy (CITE). As a result, ρ ranges from 0.3 to 3.3, with most of the potential closure amplitude squares having ρ within 11% of unity. Therefore, even without prior knowledge about the visibilities that comprise the numerator and denominator of the closure amplitude, a similar procedure to that used for the closure amplitudes, where $\rho \approx 1$ is assumed, is well motivated.

Given a value of ρ , either from the station SEFDs or setting it to unity,

$$\frac{\sigma_n}{n} \approx \frac{\rho}{\sqrt{\rho^2 + \mathcal{V}_0^2}} \frac{\sigma_{\mathcal{V}}}{\mathcal{V}_0} \quad \text{and} \quad \frac{\sigma_d}{d} \approx \frac{\mathcal{V}_0}{\sqrt{\rho^2 + \mathcal{V}_0^2}} \frac{\sigma_{\mathcal{V}}}{\mathcal{V}_0}, \quad (\text{C.23})$$

where n and d are the mean values of $|V_A||V_B|$ and $|V_C||V_D|$, and $\mathcal{V}_0 = n/d$. As a result, an approximation of the closure amplitude distribution based on the Gauss Q distribution can be constructed via

$$p_{\mathcal{V}}(\mathcal{V}; \mathcal{V}_0, \sigma_{\mathcal{V}}) = \frac{e^{-(\mathcal{V}-\mathcal{V}_0)^2/2\Sigma_{\mathcal{V}}^2}}{\sqrt{2\pi}\Sigma_{\mathcal{V}}} \left[\Delta_{\mathcal{V}} \text{erf}\left(\frac{\Delta_{\mathcal{V}}}{\sqrt{2}\Omega_{\mathcal{V}}}\right) - \frac{2\Omega_{\mathcal{V}}}{\sqrt{2\pi}} e^{-\Delta_{\mathcal{V}}^2/2\Omega_{\mathcal{V}}^2} \right] \quad (\text{C.24})$$

where $\Sigma_{\mathcal{V}}^2 = \sigma_{\mathcal{V}}^2 \frac{\rho^2 + \mathcal{V}^2}{\rho^2 + \mathcal{V}_0^2}$, $\Delta_{\mathcal{V}} = \frac{\rho^2 + \mathcal{V}\mathcal{V}_0}{\rho^2 + \mathcal{V}^2}$, $\Omega_{\mathcal{V}}^2 = \frac{\rho^2 \sigma_{\mathcal{V}}^2}{(\rho^2 + \mathcal{V}^2)(\rho^2 + \mathcal{V}_0^2)}$.

Quantitative assessments of the performance of this approximation for a number of variations in the distribution of SNRs among the various components and for fixed and known ρ are listed in Table C.4. For comparison, we also provide accuracy for estimates for a log-Normal distribution with mean $\log(\mathcal{V}_0 + \sigma_{\mathcal{V}}^2/\mathcal{V}_0)$ and standard deviation $\sigma_{\mathcal{V}}/\mathcal{V}_0$. In the Gauss Q cases with fixed ρ and for the log-Normal distribution, we permit the true value of ρ to range from 0.3 to 3.3, reporting the maximum deviation for each measure independently.

Knowledge of ρ significantly improves the quality of the approximation, which for $\text{SNR} \geq 2$ is accurate to 13% for $\mathcal{V} < 1$. The asymmetric impact of noise in the denominator and numerator of the closure amplitude is responsible for the worsening performance of the approximation when the $\mathcal{V} > 1$; generally, closure amplitudes can be constructed so that $\mathcal{V} \lesssim 1$. The performance of the approximation in reconstructing the mode and width of the distribution is very good in this limit, better than 1%. When ρ is not known a priori, setting it to unity introduces an additional error in the approximation of the closure amplitude distribution. Nevertheless, even with excursions of a factor of 3, by $\text{SNR} \geq 4$, the Gauss Q approximation is accurate to 13% at all \mathcal{V} .

In practice, the primary difficulty with applying Equation (C.24) is the accuracy with which σ_n/n and σ_d/d can be reconstructed, which depends how close \mathcal{V} is to \mathcal{V}_0 . At low SNR, this can lead to a significant error in the estimation of the likelihood. Where the estimate of σ_n/n or σ_d/d

are higher than their true values, this makes little difference. However, where they are much lower than their true values, this can result in a distribution that is considerably more narrowly concentrated about \mathcal{V}_0 than the true distribution, biasing any resulting parameter estimates. This ceases to be a significant bias for $\text{SNR} > 4$ for the 2017 campaign.

Approx.	Notes	Error	SNR				
			1	2	4	8	32
Gauss Q	$\mathcal{V}_0 = 0.5$, SNR _{n,d} ratios=0.5, $\rho \in [0.3, 3.3]$	δ_{\max} δ_{mode} δ_w	54% 38% 15%/14%	28% 17% 14%/11%	13% 6% 6%/8%	6% 2% 2%/4%	2% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 0.5$, SNR _{n,d} ratios=2, $\rho \in [0.3, 3.3]$	δ_{\max} δ_{mode} δ_w	54% 38% 15%/14%	28% 17% 14%/11%	13% 6% 6%/8%	6% 2% 2%/4%	2% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 2.0$, SNR _{n,d} ratios=0.5, $\rho \in [0.3, 3.3]$	δ_{\max} δ_{mode} δ_w	47% 25% 15%/5%	25% 17% 5%/26%	13% 6% 4%/15%	6% 2% 2%/5%	2% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 2.0$, SNR _{n,d} ratios=2, $\rho \in [0.3, 3.3]$	δ_{\max} δ_{mode} δ_w	47% 25% 15%/5%	25% 17% 5%/26%	13% 6% 4%/15%	6% 2% 2%/5%	2% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 0.5$, SNR _{n,d} ratios=0.5, $\rho = 3.3$	δ_{\max} δ_{mode} δ_w	53% 5% 9%/0.6%	13% 0.6% 0.4%/5%	6% <0.1% 0.5%/0.1%	6% <0.1% <0.5%/<0.2%	0.6% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 0.5$, SNR _{n,d} ratios=2, $\rho = 3.3$	δ_{\max} δ_{mode} δ_w	53% 5% 9%/0.6%	13% 0.6% 0.4%/5%	6% <0.1% 0.5%/0.1%	6% <0.1% <0.5%/<0.2%	0.6% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 2.0$, SNR _{n,d} ratios=0.5, $\rho = 0.3$	δ_{\max} δ_{mode} δ_w	46% 15% 5%/6%	20% 3% 5%/16%	8% 0.4% 3%/10%	8% 0.4% 0.9%/2%	0.6% <0.1% <2%/<1%
Gauss Q	$\mathcal{V}_0 = 2.0$, SNR _{n,d} ratios=2, $\rho = 0.3$	δ_{\max} δ_{mode} δ_w	46% 15% 5%/6%	20% 3% 5%/16%	8% 0.4% 3%/10%	8% 0.4% 0.9%/2%	0.6% <0.1% <2%/<1%
log-Norm	$\mathcal{V} = 0.5$, SNR _{n,d} ratios=0.5, $\rho \in [0.3, 3.3]$	δ_{\max} δ_{mode} δ_w	64% 36% 62%/28%	43% 22% 32%/38%	23% 8% 10%/13%	11% 2% 3%/5%	3% <0.1% <2%/<1%
log-Norm	$\mathcal{V} = 0.5$, SNR _{n,d} ratios=2, $\rho = 0.3$	δ_{\max} δ_{mode} δ_w	64% 36% 62%/28%	43% 22% 32%/38%	23% 8% 10%/13%	11% 2% 3%/5%	3% <0.1% <2%/<1%

–	SNR _{n,d} ratios=2, ρ ∈ [0.3, 3.3]	δ _{mode}	36%	22%	8%	2%	< 0.1%
–		δ _w	62%/28%	32%/38%	10%/13%	3%/5%	< 2%/ < 1%
log-Norm	ℳ = 2.0,	δ _{max}	76%	52%	29%	14%	3%
–	SNR _{n,d} ratios=0.5, ρ ∈ [0.3, 3.3]	δ _{mode}	42%	27%	10%	3%	0.2%
–		δ _w	> 100%/40%	41%/32%	12%/11%	4%/5%	< 2%/ < 1%
log-Norm	ℳ = 2.0,	δ _{max}	76%	52%	29%	14%	3%
–	SNR _{n,d} ratios=2, ρ ∈ [0.3, 3.3]	δ _{mode}	42%	27%	10%	3%	0.2%
–		δ _w	> 100%/40%	41%/32%	12%/11%	4%/5%	< 2%/ < 1%

Table C.4: Accuracy of Closure Amplitude Error Distribution Approximations. δ_{\max} is the maximum absolute difference, measured relative to probability maximum. SNR_{n,d} ratios stand for the ratio of $V_{A,0}/\sigma_A$ to $V_{B,0}/\sigma_B$ and $V_{C,0}/\sigma_C$ to $V_{D,0}/\sigma_D$ in the construction of the closure amplitude distribution. δ_{mode} represents the fractional error in the location of the mode and δ_w is the fractional error in the width of the region containing 68%/95% of the cumulative probability.