

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Predicting Microlensing Rates and Properties in Wide-Field Surveys

Permalink

<https://escholarship.org/uc/item/7k9461vh>

Author

Fournier, Amanda Patrice

Publication Date

2015

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Predicting Microlensing Rates and Properties in Wide-Field Surveys

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Amanda Patrice Fournier

Committee in charge:

Professor Lars Bidsten, Chair
Dr. Rachel Street
Professor Carl Gwinn
Professor Tommaso Treu

March 2015

The Dissertation of Amanda Patrice Fournier is approved.

Dr. Rachel Street

Professor Carl Gwinn

Professor Tommaso Treu

Professor Lars Bidsten, Committee Chair

March 2015

Predicting Microlensing Rates and Properties in Wide-Field Surveys

Copyright © 2015

by

Amanda Patrice Fournier

To the teachers

Acknowledgements

I have had a difficult journey through grad school, and could not have made it without the support of a great many wonderful people.

My mother, Vikki Fournier, has taught me since before I could speak, and I have never stopped learning from her. Her love was always ready to remind me that there are good things in the world. Her example of strength and perseverance is the best anyone could aspire to.

My boyfriend, Christopher Takacs, has supported me through illness and uncertainty. He has always had my back, and has always encouraged me to take the next step, even when I falter. With him I have become a better person, and learned to better forgive my own shortcomings.

My advisors Rachel Street and Lars Bildsten took in a misplaced student and gave me a place to grow. They gave me the safety and resources I needed to develop as a scientist. Their patience and kindness, their excitement about science, their guidance and teaching was exactly what I needed.

Adrian Price-Whelan and Marcel Agüeros of Columbia University were my first scientific collaborators. Working with them was one of the best experiences of my academic career; they spurred me to do better.

Tina Panteleakos and Turi Honegger are two of the many doctors that helped me recover from devastating illness, but I owe them a particular debt for their generous time, friendship, and care.

I will always be indebted to my Uncle Tony and Aunt Patty, who put up with my adolescent angst for a summer when I quite desperately needed it, and took care of me the best they could.

My sisters and brother - Briana, Christina, and Austin - whenever I am with you, I

am home.

I'd like to thank Las Cumbres Observatory Global Telescope Network and its people for taking a chance on a student and employing me through much of my graduate career in a warm and welcoming atmosphere.

I thank the RoboNet collaboration for all they taught me and for the wonderful experience of working with them.

I thank every author of a scientific paper who has gone out of their way to ensure that others can build on their work. In particular, I would like to thank David Schlegel, not only for his meticulous work but for his responsiveness to queries.

And finally, I thank the teachers. There have been too many to count.

Curriculum Vitæ

Amanda Patrice Fournier

Education

- 2015 Ph.D. in Physics (Expected), University of California, Santa Barbara.
- 2013 M.A. in Physics, University of California, Santa Barbara.
- 2007 B.S. in Physics, University of Southern California

Publications

- 2014 Statistical Searches for Microlensing Events in Large, Non-uniformly Sampled Time-Domain Surveys: A Test Using Palomar Transient Factory Data
- 2010 The Size and Origin of Metal-enriched Regions in the Intergalactic Medium from Spectra of Binary Quasars

Fields of study

- 2010-2015 Time-domain data trend removal, microlensing observation, and microlensing simulation under Dr. Rachel Street and Prof. Lars Bildsten
- 2007-2010 Diffuse interstellar media and quasar spectroscopy under Prof. Crystal Martin

Abstract

Predicting Microlensing Rates and Properties in Wide-Field Surveys

by

Amanda Patrice Fournier

The rates and properties of out-of-plane microlensing events have been understudied in the past. We seek to remedy this by building a simulation of galactic stars and lensing events, drawing upon numerous up-to-date sources. The resulting code is a well-verified software tool which can be adapted to simulate a wide range of potential survey strategies and parameters. It will be a useful tool for the community to optimize the design of the deep, wide-angle surveys coming on line in the next decades, such as ATLAS, Evryscope, and LSST. In the text, we provide baseline all-sky lensing properties for a deep survey.

Contents

Curriculum Vitae	vii
Abstract	viii
1 Introduction	1
1.1 A brief introduction to microlensing	1
1.2 Microlensing and dim, massive objects	3
1.3 Microlensing detection of planets	6
1.4 Field microlensing	12
1.5 Our project: Predictions for observers	14
2 Introduction to Microlensing	16
2.1 Introduction	16
2.2 Deflection angle	17
2.3 The lens equation in the point source, point lens case	21
2.4 Magnification	25
2.5 Binary mass lensing	33
2.6 Microlensing rate	36
2.7 Second-order microlensing effects	37
2.8 Numeric values of microlensing quantities	40
3 Introduction to Numerical Integration	44
3.1 Numerical integration methods	44
3.2 Computational efficiency in multidimensional problems	54
3.3 Monte Carlo Integration	63
4 Astrophysical inputs to the code	72
4.1 Summary of inputs	72
4.2 Comparison to earlier work	94
5 Code design	98
5.1 Summary of simulation design	98

6	Code verification	115
6.1	Coded persistent tests	115
6.2	Comparison to models and observations	117
7	Results	133
7.1	Results	133
7.2	Discussion	145
7.3	Final words	149
	Bibliography	150

Chapter 1

Introduction

1.1 A brief introduction to microlensing

A key result of general relativity is that gravity is not so much a force as a bending of time and space. Because of this, gravity can influence the path of even massless particles, like photons [Einstein, 1916]. Thus it is possible for massive, gravitating objects, like stars and galaxies, to act as “lenses” for bright objects, like other stars and galaxies, turning more of their light towards an observer on earth [Einstein, 1936]. This phenomenon is known as lensing, and when the lens mass is a star or smaller object, it is called microlensing.

When the lensing effect was first proposed, it was thought to be unobservable [Einstein, 1936]. However, improvements in telescope technology, especially the introduction of the charge coupled device (CCD) camera and the resulting integration of telescopes with computers, have transformed the observation of microlensing from nigh impossible to an achievable endeavor, albeit one requiring daunting amounts of resources and coordination [Paczynski, 1986].

The first effort at deliberate microlensing observation, the MACHO project, required

enormous efforts to overcome the technical obstacles of microlensing, ultimately observing less than two dozen events over half a dozen years [Alcock et al., 2000]. Detection, distinction from other celestial events, and characterization were all nontrivial, but these obstacles were ultimately overcome, laying the groundwork for later surveys.

A combination of improved technique, clever use of resources, and heroic efforts ultimately brought the Optical Gravitational Lensing Experiment (OGLE) [Udalski et al., 1992, Udalski et al., 1997, Udalski, 2003, Udalski et al., 2008, Wyrzykowski et al., 2014] to the forefront of MACHO's successors, presently uncovering over a thousand lensing events per year, followed by the Microlensing Observations in Astrophysics (MOA) [Muraki et al., 1999, Yanagisawa et al., 2000, Hearnshaw et al., 2006, Sako et al., 2008]. OGLE and other modern surveys use a two-part strategy for microlensing observations; while the OGLE survey proper continually scans the sky to detect events, it relies on outside collaborators for the detailed, high-fidelity observations and continuous coverage needed to actually characterize lensing events. This strategy, carried out by a dedicated survey team and their community of supporters, has met with great success; OGLE is now uncovering over a thousand lensing events per year.

It has been proposed that the Large Synoptic Survey Telescope (LSST) [LSST Science Collaboration et al., 2009] could represent a next major step in microlensing surveys, although their effectiveness will depend on the details of their observing strategy [Gould, 2013]. The LSST will survey most of the visible sky multiple times a night, potentially identifying many microlensing events in the process [LSST Science Collaboration et al., 2009, p. 268]. We will present extensive analysis of likely outcomes of such an effort, noting the effects of possible changes in the still-developing survey plan.

1.2 Microlensing and dim, massive objects

Lensing has its roots in finding and characterizing hard-to-see celestial objects. The very first microlensing study, the MACHO Project [Hart et al., 1996, Alcock et al., 2000], sought MAssive Compact Halo Objects such as black holes or brown dwarfs that were thought to be a major component of the galaxy’s dark matter. While traditional astronomical methods wouldn’t do - the hypothesized MACHOs did not give off much light, thus qualifying as dark matter - any large enough mass is perfectly capable of acting as the lens mass in a microlensing event.

The MACHO Project suffered through the teething problems that come with being first. They had to deal with the difficulties of ground-based time-domain astronomy; not only is data discontinuous as the daytime sky blocks out the stars, but bad weather disrupts observations, sometimes halting them completely, sometimes lowering their quality. Sometimes the disturbances are detected; sometimes not. They had to distinguish different kinds of transient events not using perfect data, but incomplete, noisy data. Although some false signals were relatively easy to be rid of, they had to develop their own means of filtering out some variable stars [Alcock et al., 2000, p. 290], and still could not be fully certain of every interesting event. Even so, they were able to come away fairly certain of more than a dozen microlensing events, and to provide a first estimate of the galactic dark matter’s MACHO population.

The microlensing field has matured a great deal since those early efforts. What was once done with difficulty is now done routinely; MACHO took years to find a scant dozen events, but OGLE detects more than a thousand lensing events per year, and MOA detects hundreds. Large portions of the data pipeline have been given over to automation, but just as importantly, the community has grown a population of expert lensing event analysts. The field is now more than capable of studying a wide variety of

other dim objects, such as brown dwarfs, rogue planets, and stellar remnants.

Brown dwarfs are failed stars that never accumulated enough mass to start a self-sustaining fusion reaction. Any star-like object smaller than about seven or eight percent of our sun's mass will not be able to ignite and sustain a hydrogen-to-helium fusion reaction, the main source of power for the vast majority of stars. First theorized in the 1960s [Kumar, 1963, Hayashi and Nakano, 1963], brown dwarfs were not detected until the 80s [Becklin and Zuckerman, 1988] or confirmed until the 90s [Golimowski et al., 1995, Pavlenko, 1997]. They emit a faint light as they slowly cool, mostly in the infrared, but are very dim. Brown dwarfs are difficult to detect, and their detection by conventional means is virtually impossible from the ground; our atmosphere blocks most of their radiation. They can be found either in orbit about a larger star, using many of the same techniques as planet-finding, or solitary, using space-based infrared surveys. Brown dwarfs are in many ways the missing link between low-mass stars and high-mass planets; learning more about them will inform studies of stars and planets alike [Beichman et al., 2014].

Infrared surveys have been successful in detecting lone brown dwarfs, but only in the very close neighborhood of our sun. The Deep Near Infrared Survey (DENIS) [Delfosse et al., 1997], the Two Micron All Sky Survey (2MASS) [Kleinmann et al., 1994], the Wide-field Infrared Survey Explorer (WISE) [Wright et al., 2010, Mainzer et al., 2011] and others have discovered hundreds of brown dwarfs [Phan-Bao et al., 2001, Phan-Bao et al., 2003, Kirkpatrick et al., 2000, Kirkpatrick et al., 2011, Mace et al., 2013], but all very close by, within about a hundred parsecs. Microlensing, by contrast, has discovered and taken good measurements of brown dwarfs without stellar companions at distances of nearly two thousand parsecs [Choi et al., 2013].

At low masses, brown dwarfs come to closely resemble gas giant planets like Jupiter and Saturn. There is some debate about where to draw the line between the two classes

of objects. For our purposes, the mass at which deuterium ignites is a useful boundary; objects larger than about 13 Jupiter masses or 1.2 percent of a solar mass will briefly ignite and fuse their deuterium into helium, but smaller objects will not. Anything below this mass, we will refer to as a planet.

“Free-floating” or “rogue” planets, i.e. planet-sized masses not bound to any home star, are even more difficult to detect than brown dwarfs. Their presence cannot be inferred using the most widespread planet-finding techniques because those methods require the planet to orbit a visible star. They emit very little light of their own, and are not near enough to a star to shine even by reflection; no telescope on or around earth can hope to see any but the nearest, youngest, and largest. The population of free planets could provide a wealth of information about how planets form and how well they fare over long time scales, but they are terribly difficult to detect at all.

Nonetheless, rogue planets have been detected. The MOA collaboration announced in 2011 that they had detected a population of low-mass objects with no sign of stellar hosts, at least some of which were very likely rogue planets [Sumi et al., 2011]; others may be brown dwarfs or planets in very wide orbits. Using this data, they were able to provide an estimate that there are roughly twice as many free-floating planets as there are stars in our galaxy. A handful of rogue planet candidates have been found by other methods, some of them very plausible but uncommonly young and close to our sun [Delorme et al., 2012]. Only microlensing can detect enough rogue planets to seriously attempt population statistics.

The same principles apply to any class of dim, dense, massive objects; if it is in our galaxy, microlensing has the potential to detect it. As with brown dwarfs and rogue planets, microlensing is capable of providing mass measurements for stellar remnants such as black holes and neutron stars regardless of whether they are bright enough to see directly, and can provide an independent mass measurement for those that can already

be detected and characterized.

1.3 Microlensing detection of planets

Much of microlensing’s current popular appeal comes from its use as a planet detection technique. During a microlensing event in which a star magnifies a background source, it is possible for a planet orbiting the lens star to cause additional magnification from which the planet’s presence and properties can be deduced. This effect was first described by [Mao and Paczynski, 1991] as a possible by-product of the MACHO survey, and elaborated upon by [Gould and Loeb, 1992].

This suggestion came at a time of great excitement in the astrophysics community, as the first confirmed planets outside our solar system (“exoplanets”) had just been discovered. Wolszczan & Frail observed a variation in the timing of a nearby pulsar and demonstrated that this variation was consistent with not one, but two planets in orbit about the pulsar [Wolszczan and Frail, 1992].

This discovery was an enormous landmark for astronomy. The discovery of exoplanets has long been a Holy Grail of astronomy. Speculations about “other worlds” are perhaps as old as humanity, but speculations about exoplanets specifically were recorded shortly after the advent of the Copernican model in the sixteenth century by Giordano Bruno, and mistaken discoveries published as early as the nineteenth century [Jacob, 1855]. All prior claims had been discredited, however, either by demonstrations that phenomena were better explained otherwise, or by a failure to reproduce observations that should be repeatable.

Not only was the discovery of exoplanets a great technical feat, it represented an important philosophical point. It is one of the most important rebuffs to anthropocentrism, world views in which humanity and earth are unique, special, and in extreme cases, liter-

ally thought to be the center of the universe. Copernicus was one of the first to propose otherwise, postulating that the earth orbited the sun, not vice versa; this was borne out by the observations of Brahe and Kepler. Hubble's observations of so-called spiral nebulae demonstrated that they were much further away than previously thought [Hubble, 1925], leading to the conclusion that they were in fact spiral galaxies; our galaxy is not the only one. Astronomers have long thought that it is unlikely that our star was the only one to have planets and thus potentially be home to life, but while the view was certainly defensible, it had no concrete proof. The existence of exoplanets is the latest - but we hope not the last - in a series of discoveries showing that the universe is a much bigger, more interesting place than humanity alone.

Wolszczan & Frail's discovery started something akin to a gold rush, as projects to search for planets were suddenly the subject of much interest and old proposals for planet-finding methods received a wealth of new attention. A great many planets have been discovered since using a wide variety of methods, but a handful deserve special mention and will be covered here.

As a planet orbits a star, the star also moves, counterbalancing the planet. The star's movement is much smaller than the planet's, making it difficult to detect by looking for stars that "wobble" from side to side; the first such detection by tracking a star's position was achieved in 2010 [Muterspaugh et al., 2010a-2010e]. However, a closely related effect was the first detection method to achieve great success.

Radial velocity (RV) measurement relies on the fact that a star's wobble towards and away from the earth can be detected because it has a Doppler effect on the light from the star [Struve, 1952]; its light shifts to higher, bluer frequencies as the star moves towards us, and lower, redder frequencies as it moves away. Precision spectroscopy can detect the shift in color and track it over time.

Radial velocity measurements were the first major source of planets. From 1995




Figure 1.1: Planet detections per year by detection method. Radial velocity measurements dominated until 2009, when Kepler Observatory began reporting planetary candidates. Image courtesy of the NASA Exoplanet Archive.

[Mayor and Queloz, 1995, Marcy et al., 1997] to about 2010, RV was the reigning champion of planet-hunting; to date it has racked up more than 500 planet detections. However, it suffered from being difficult to mass produce; each and every detection was the product of exacting, time-consuming work by astronomers. The spectroscopes used in RV detection are some of the best in the world, developed specifically for the task by dedicated teams, and not easily replicated or replaced. Furthermore, to reach the precision required for planet searches, a star must be very bright by astronomers' standards; such stars are rare, and it is uncommon to see several close together on the sky, where one measurement might capture all at once. Each detection requires an individual allotment of telescope time with a world-class instrument. This particular weakness is one of the main reasons RV was ultimately upstaged by an unlikely-seeming contender.

Transit detection [Charbonneau et al., 2000] relies on a very special coincidence; it requires that a planet's orbit pass directly between its parent star and earth. The slight dimming as the planet blocks out some of its star and the return to full brightness as the planet moves away [Struve, 1952] is what transit surveys watch for. At least three transits need to be observed before the existence of a planet can be confidently deduced from transit data, and unlike RV detection, transit detection requires an almost unblinking watch; if you're not observing at the time a transit happens, you'll have no hint of the planet's presence.

Planet hunters dealt with the rarity of transiting planets by searching many, many stars, and the timing of transit events by keeping nightly watch for months at a time. Surveys such as the Wide Angle Search for Planets (WASP) [Pollacco et al., 2006, Butters et al., 2010, Smith and WASP Consortium, 2014] and the Hungarian Automated Telescope Network (HATnet) [Bakos et al., 2004, Bakos et al., 2013] both had to handle massive quantities of data at a very fast pace, much faster than a human team could possibly keep up with; they both had to rely on automation for a large portion of their analysis needs. Although some RV data was still required to prove that the star did not have a stellar companion that could explain the variations in its light, confirming or denying a likely candidate whose properties were already well-constrained was much quicker and easier than proving the presence of a planet through RV data alone, while searching without any foreknowledge of which stars would prove interesting. This strategy paid off; each survey netted dozens of new discoveries, and ultimately paved the way for the fantastically successful Kepler survey.

The Kepler Observatory is a telescope in orbit about earth, designed from the ground up for the purpose of finding transit events. Its placement in space has important advantages; it never needs to close because of the day-night cycle, and has no atmosphere blurring its view of the stars. Its sensitivity lets it detect transits of planet masses as

small as earth; its continuous operation of five years running lets it confirm planetary orbits as long as an earth year. Kepler's enormous success - over one thousand confirmed planets, including some with very earth like properties - transformed transit detection into the dominant form of planet-hunting in the space of a few years.

Microensing, like transits, relies on a coincidence: one star passing in front of another, close enough to act as a magnifying lens. To detect a planet via microensing, an observer must find microensing events as they happen, searching millions of stars at once. They must then record the microensing event in detail, especially the brightest point in the event, which contains a lot of information about the objects involved. Most events monitored by scientists are long - there may be a month or more between the time when an event begins, and the time when it peaks - but not all, and short events may prove very interesting.

At first glance, microensing seems to share the worst difficulties of both radial velocity detection and transit searches. Like RV, each microensing event requires a fair amount of individual attention and skilled labor to characterize in detail. Like transits, microensing events are rare and give no advance warning. While lensing events do generally give an observer more opportunity to prepare than a previously unknown transit, they are strictly one-shot events that will never recur, and cannot be observed in more detail at a later date.

What makes microensing uniquely useful is the type of planets it detects. Both transits and radial velocity searches are more likely to pick up on a planet if it is closer to its star, in scalding hot orbits like Mercury's, or warmer; microensing tends to detect planets near Earth's orbit or cooler. RV favors massive planets like Jupiter and Saturn, and transits likewise favor gas giants very strongly over small, rocky planets like Earth; microensing also prefers large planets, but does much better at reaching down to earth-like masses. RV works best when a planet's orbit is near to edge-on, and a transit requires

that the orbit be almost perfectly edge-on; microlensing has a chance to detect planets in any orbit about their star. For nearly every aspect of a planet’s properties, microlensing favors a different set of characteristics than the currently dominant methods, and it is generally much more sensitive to planets earth-like in nature than RV or transit searches.




Figure 1.2: Planet properties by detection method. Microlensing probes lower masses and longer periods, hence cooler orbits, than other detection methods. Image courtesy of the NASA Exoplanet Archive.

The one truly unique characteristic of microlensing as a planet-finding technique, though, is that we need not be able to see the star a planet orbits. Stars smaller than our sun are plentiful, long-living, and relatively stable, in many ways excellent candidates to host earth-like life, but dim. The brightest stars, by contrast, are large; some massive and short-lived, others at the end of their lifetime, neither type particularly hospitable. As a technique that doesn’t depend on an object’s own visible light, microlensing has a truly unique level of access to the small, dim, stable stars of our galaxy, able to detect

them at much greater distances.

The net effect of these differences is that microlensing serves to fill in an area of parameter space that other methods cannot access. Microlensing is able to detect earth-mass planets in earth-like orbits, from the ground, using relatively modest telescopes [Abe et al., 2013]. This claim is absolutely unique amongst planet-finding methods.

1.4 Field microlensing

To date, microlensing projects have focused on dense stellar fields. MACHO, searching for massive dark objects in the galactic halo, pointed towards the Magellanic clouds. OGLE and MOA focus primarily of the galactic bulge, where microlensing rates are highest. These projects are excellently designed to fill their intended purposes, but leave untouched the more sparsely starred areas of the sky, the “field”.

1.4.1 The Tago event

In October 2006 A. Tago, an amateur astronomer unaffiliated with any microlensing programs, reported what he believed to be a supernova of an already bright star. Fortunately for all involved, his report was noticed by scientists and recognized for what it was: microlensing.

A team of Japanese microlensing astronomers began an observing campaign the following night [Fukui et al., 2007]. They concluded, and later follow-up analysis agreed [Gaudi et al., 2008b], that this was indeed a microlensing event, the lensing of a “field” star, located out of the galactic disk rather than towards the galactic bulge. This event was the first of its kind, and to the day of this publication, remains the only field lensing event observed.

Partly due to its unique nature, there were strong efforts to characterize the Tago

event as well as possible. However, this event yielded much tighter constraints on the properties of the stars involved than a typical bulge lensing event subjected to the same level of scrutiny. The stars were closer, and as a result had brighter apparent magnitudes, larger proper motions, and a shorter separation time than the stars in typical bulge lensing events; they also occupied an uncrowded area of the sky, where it was possible to resolve them from uninvolved stars. The benefits of the stars being near us were important to their successful study, but the clear sky behind them was even more so.

1.4.2 Wide-field surveys and possibilities for the future

It has long been known that the rate of lensing will be much lower in regions with fewer stars, but the Tago event inspired a renewed interest in the possibilities of field lensing. About a year later, Cheongho Han published a paper intended to estimate the microlensing rate in the field, in the sparsely-starred areas above and below the galactic disk [Han, 2008]. That paper forms a major part of the inspiration for this thesis, which aims to make a similar estimation. Han ultimately concluded that there would be roughly two dozen events per year caused by the lensing of field stars visible to a survey that could see stars with a brightness of eighteen visual magnitudes.

It is not possible to justify the building of a suitable sky survey for a mere two dozen events each year, not with thousands of events already being detected each year by other projects. The survey to do this would need to be high-cadence (it should photograph the sky once every few nights, at a minimum), wide-angle (it should cover a significant fraction of the visible sky), and deep (able to see dim stars); these three characteristics together describe a very advanced and expensive system, unlike any survey now operating. However, it is a long-standing tradition of astronomy to allow multiple teams access to a data set, each studying a different type of object or event. If a suitable survey existed,

there is an excellent chance that microlensers could gain access for the study of lensing events.

Enter the Large Synoptic Survey Telescope, or LSST [LSST Science Collaboration et al., 2009]. LSST will survey most of its visible sky, a space of 10,000 square degrees, about one-quarter of the total sphere about earth; it satisfies the wide angle requirement. LSST will return to each position every three to four days, satisfying the high cadence requirement. LSST will see stars as dim as 24.5 magnitudes in the red band, by far surpassing the requirement for depth. It has enormous potential as the detection stage of a lensing detection-followup collaboration similar to MOA or OGLE.

The LSST group is alert to the possibility of studying lensing events [LSST Science Collaboration et al., 2009], but does not have many microlensers on team to advocate for such a program. Some aspects of their proposed program may need adjustment to reach LSST's full potential in this area [Gould, 2013]. This is at the core of the problem we hope to address in this thesis.

1.5 Our project: Predictions for observers

As we will explain in chapter 2, translating between the directly observable microlensing light curve and the physical properties of the objects involved is often a complex task, and careful analysis often calls for good models of stellar populations. These models are used to provide prior probabilities about the stars involved; if a type of star is known to be rare, it is less likely to be involved in a given event.

In this thesis I provide a detailed, working model of microlensing events in our galaxy, with extensive comparison to known quantities wherever possible. The model can take common observational restrictions into account, such as field of view, survey cadence, and limiting magnitudes, returning predictions about the properties of events such an

observation would detect. It is intended that this should be a useful tool to observers seeking to characterize events.

I also hope to inform the design of microlensing observations. Upcoming large field of view surveys like LSST may benefit greatly from being able to determine in advance their likely yield of microlensing events and how it depends on survey design decisions.

The equations governing microlensing are derived in full in chapter 2, and technical aspects of the discipline expanded upon. Numerical integration techniques, special challenges of multidimensional integration, and Monte Carlo integration methods are covered in chapter 3. Chapter 4 contains the astrophysical inputs that we combine to form a model of galactic stellar populations. Chapter 5 outlines the code's design and notable features. Chapter 6 describes the verification of the code, including software tests and comparison to known observable quantities. Chapter 7 showcases our results, discussions, and conclusions.

Chapter 2

Introduction to Microlensing

2.1 Introduction

Albert Einstein published on gravitational lensing twenty-four years after he first published on the gravitational bending of light [Einstein, 1916, Einstein, 1936], but did not think it was important. He could see that for any given star, the odds of being lensed were minuscule, and that a lensed star would be nigh impossible to distinguish from an unlensed one using still-frame photometry. Lensing was, to all appearances, an unobservable phenomenon. In the end, the result was published only because a Czech engineer, R. W. Mandl, pleaded with him to do so.

Einstein was missing two crucial pieces of information, though. First is that the stars are all in motion, not just orbiting the galactic center, but each deviating from that average by some amount. We need not rely on their appearance at any one moment to learn whether they are lensed, but rather can observe how their appearance changes over time as they pass in and out of lensing alignments. Time-dependent photometry *can* distinguish a lensing event, as we will describe in this chapter.

Second is that our photographic and computing technology would improve by orders

of magnitude in the century following his work. What was far too dim to see in his day can now be observed, and the digitization of photography means that computers can search millions of signals for one that matches the microlensing light curve.

In this chapter we will derive useful equations to describe microlensing properties and touch on phenomenology to help the reader gain some intuition about the observation of microlensing events, as well as the historical development of microlensing observation. §2.2 derives the deflection angle of light under general relativity. §2.3 finds the resulting position of lens images. §2.4 derives the magnification caused by a lensing event, and §2.4.1 discusses the light curve that results. Binary mass lensing, as with lens stars orbited by planets, is covered in §2.5, and some important points on planet detection in §2.5.1. The microlensing rate is derived in §2.6. Important second-order effects are described in §2.7. §2.8 gives estimates of the numeric values of lensing quantities.

2.2 Deflection angle

Here we derive the angle by which light is deflected when passing a massive, gravitating object. This derivation references without proof certain results of general relativity which are beyond the scope of this thesis. Einstein's original result can be found in [Einstein, 1916] and a more detailed derivation in [Wald, 1984].

Assumptions:

The gravitating object - the "lens" - is a single object, is to an excellent approximation spherical, is not rotating at relativistic speeds, and has a radius significantly smaller than the impact parameter of the passing photon.

The angle by which our test photon is deflected is very small.

Definitions:

r is the photon's impact parameter, i.e. the distance at which it passes closest to the

lens star.

M is the mass of the lens star.

ϵ is the angle of deflection.

Derivation:

The curvature of space-time outside a spherically symmetric gravitating object, such as a star, is given by the Schwarzschild metric [Schwarzschild, 1916]:

$$ds^2 = - \left(1 - \frac{2GM}{c^2 r}\right) (cdt)^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} dr^2 + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) \quad (2.1)$$

We consider the path a photon takes in this geometry. Light travels along null metrics ($ds^2 = 0$), and we can choose our coordinates such that throughout the photon's travel $d\phi = 0$.

$$0 = - \left(1 - \frac{2GM}{c^2 r}\right) (cdt)^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} dr^2 + r^2 d\theta^2 \quad (2.2)$$

The static Killing field gives us a conserved momentum-like quantity, E .

$$E = \left(1 - \frac{2GM}{c^2 r}\right) ct \quad (2.3)$$

The rotational Killing field gives us a conserved angular momentum-like quantity, L .

$$L = r^2 \dot{\phi} \quad (2.4)$$

Substituting L and E into the metric and rearranging produces

$$\dot{r}^2 = E^2 - \left(1 - \frac{2GM}{c^2 r}\right) \frac{L^2}{r^2} \quad (2.5)$$

We can now write an equation for the change of ϕ with respect to r by combining

this latest equation with the equation for L .

$$\frac{d\phi}{dr} = \frac{L}{r^2} \left(E^2 - \left(1 - \frac{2GM}{c^2 r} \right) \frac{L^2}{r^2} \right)^{-1/2} \quad (2.6)$$

Recognizing that a conserved angular momentum and a conserved momentum imply an impact parameter, we write

$$L/E = b \quad (2.7)$$

where b is the impact parameter the photon would have if the metric were flat and/or M were zero. We can find b 's relation to the actual distance of closest approach, r_0 , by finding the value of r for which $\dot{r} = 0$.

$$0 = E^2 - \left(1 - \frac{2GM}{c^2 r_0} \right) \frac{E^2 b^2}{r_0^2} \quad (2.8)$$

$$b^2 = \frac{r_0^3}{\left(r_0 - \frac{2GM}{c^2} \right)} \quad (2.9)$$

Note in the flat-space limit $M = 0$, we find that $b = r_0$.

We simplify $d\phi/dr$ using $L/E = b$.

$$\frac{d\phi}{dr} = \frac{1}{r^2} \left(b^{-2} - \left(1 - \frac{2GM}{c^2 r} \right) \frac{1}{r^2} \right)^{-1/2} \quad (2.10)$$

The total change in ϕ over the photon's trajectory will be given by

$$\Delta\phi = 2 \int_{r_0}^{\infty} \frac{d\phi}{dr} dr \quad (2.11)$$

This integral becomes easier to handle if we use a substitution $\rho \equiv 1/r$.

$$\Delta\phi = 2 \int_0^{1/r_0} \frac{d\rho}{\sqrt{b^{-2} - \rho^2 + \frac{2GM}{c^2}\rho^3}} \quad (2.12)$$

This is as simple as we can make the integral, but it is still not trivial to evaluate in full. Instead, we will evaluate $\Delta\phi$ at $M = 0$ and the derivative with respect to M at the same point, creating a first-order expansion in M . Recall that for $M = 0$, $b = r_0$.

$$\Delta\phi|_{M=0} = 2 \arcsin(b/r_0) = \pi \quad (2.13)$$

For the first derivative, we write b in terms of r_0 and hold r_0 constant.

$$\left. \frac{\partial\phi}{\partial M} \right|_{M=0} = 2\partial_M \int_0^{1/r_0} \frac{d\rho}{\sqrt{r_0^{-2} - \frac{2GM}{c^2}r_0^{-3} - \rho^2 + \frac{2GM}{c^2}\rho^3}} \Big|_{M=0} \quad (2.14)$$

$$= 2 \int_0^{1/r_0} \frac{(r_0^{-3} - \rho^{-3}) d\rho}{(r_0^{-2} - \frac{2GM}{c^2}r_0^{-3} - \rho^2 + \frac{2GM}{c^2}\rho^3)^{3/2}} \Big|_{M=0} \quad (2.15)$$

$$= 2 \int_0^{1/r_0} \frac{\frac{2G}{c^2} (r_0^{-3} - \rho^{-3}) d\rho}{(r_0^{-2} - \rho^2)^{3/2}} \quad (2.16)$$

$$= -\frac{4G}{c^2 r_0} \quad (2.17)$$

Therefore the bending angle $\Delta\phi$ (henceforth referred to as ϵ) is, to first-order approximation,

$$\epsilon = \frac{4GM}{r_0 c^2} \quad (2.18)$$

2.3 The lens equation in the point source, point lens case

Here we derive position of “images” produced when a background source is lensed, i.e. where it appears to be to an observer. Original work on this subject can be found in [Liebes, 1964, Refsdal, 1964].

Assumptions:

The gravitating object is spherical, is not rotating at relativistic speeds, and has a radius significantly smaller than the impact parameter, and hence can be treated as a point mass under relativistic mechanics.

The angle by which light is deflected is small.

The radius of the source star is much, much smaller than the distance from observer to source star, and hence it can be treated as a point source of light.

The fraction of the distance where the light is being significantly influenced by the lens’ gravity is much smaller than the total distance traveled, and hence we can approximate the light’s path well by a sudden bending as it passes through the lens plane.

These assumptions taken together are often referred to as the “point source, point lens” or PSPL case for a single lens.

Definitions (see fig. 2.1):

M is the mass of the lens star whose gravity deflects light from the source star.

$r_{1,2}$ is the impact parameter i.e. distance of closest passage between the photon’s trajectory and the gravitating mass. As we will see, this is a double-valued quantity for the PSPL case.

$\alpha_{1,2}$ is the angular separation (from observer’s point of view) between the source and the lens.

$\beta_{1,2}$ are the angles between source true position and image locations.




Figure 2.1: Diagram of the bending of light by a lens mass. In the case described here, light can and does take two paths from the source to the observer. Quantities are labeled accordingly. Note that angles are greatly exaggerated for the purpose of illustration and the two paths are nearly parallel.

$\theta_{1,2}$ are the angles between lens and image locations in the lens plane; $r_{1,2} = D_L \theta_{1,2}$.

$\epsilon_{1,2}$ are bending angles of light passing near a massive object, $\frac{4GM}{r_{1,2}c^2}$. Again, a solution for each route light can take to the observer.

D_S is the distance from observer to source star, D_L is the distance from observer to lens star, and D_{LS} is the distance between source star and lens star.

Derivation:

We use two geometric identities involving the unnamed third angle in the β , β' triangle.

$$180 - \beta_{1,2} - \beta'_{1,2} = 180 - \epsilon_{1,2} \quad (2.19)$$

$$\beta_{1,2} + \beta'_{1,2} = \epsilon_{1,2} \quad (2.20)$$

Use small-angle approximation for source to source image distance in the lens plane, then substitute in previous relation.

$$\beta_{1,2} D_L = \beta'_{1,2} D_{LS} \quad (2.21)$$

$$\beta_{1,2} (1 + D_L/D_{LS}) = \epsilon_{1,2} \quad (2.22)$$

Define $(1 + D_L/D_{LS}) \equiv \mu$.

$$\beta_{1,2} = \epsilon_{1,2}/\mu \quad (2.23)$$

Recall value of ϵ derived in §2.2; relate to θ and β .

$$\epsilon_{1,2} = 4GM/r_{1,2}c^2 \quad (2.24)$$

$$\theta_{1,2}D_L = r_{1,2} \quad (2.25)$$

$$\beta_{1,2} = 4GM/\theta_{1,2}D_L\mu c^2 \quad (2.26)$$

Relate β , α , and θ . Here, the ‘ $\pm\alpha$ ’ term takes the value of $+\alpha$ for the 1-subscript case and $-\alpha$ for the 2-subscript case. Substitute for β .

$$\beta_{1,2} \pm \alpha = \theta_{1,2} \quad (2.27)$$

$$4GM/\theta_{1,2}D_L\mu c^2 \pm \alpha = \theta_{1,2} \quad (2.28)$$

Define $4GM/D_L\mu c^2 \equiv \theta_E^2$ and solve the quadratic equation for $\theta_{1,2}$ with the restriction that the value of θ must be real and positive.

$$\theta_{1,2} = \pm \frac{\alpha}{2} + \sqrt{(\alpha/2)^2 + \theta_E^2} \quad (2.29)$$

This equality is referred to as the lens equation.

Note that α changes as the source and lens move relative one another, but θ_E is a constant composed of the physical quantities involved.

θ_E , the angular Einstein radius, or its equivalent distance in the lens plane $r_E \equiv D_L\theta_E$, the Einstein radius, is useful to keep in mind as the basic physical scale of a lensing problem. We will be using units normalized by these quantities in much of this chapter.

2.4 Magnification

Here we derive the solid-angle size of the images produced by microlensing, and the resulting change in the source's apparent brightness. Original work on this subject can be found in [Liebes, 1964, Refsdal, 1964].

Assumptions:

We once again take the PSPL assumptions enumerated in previous sections, but we now consider the fact that while the source is very, very small and appears pointlike to the observer, it is not a point; it has some finite radius and its image spans some finite albeit miniscule angle.

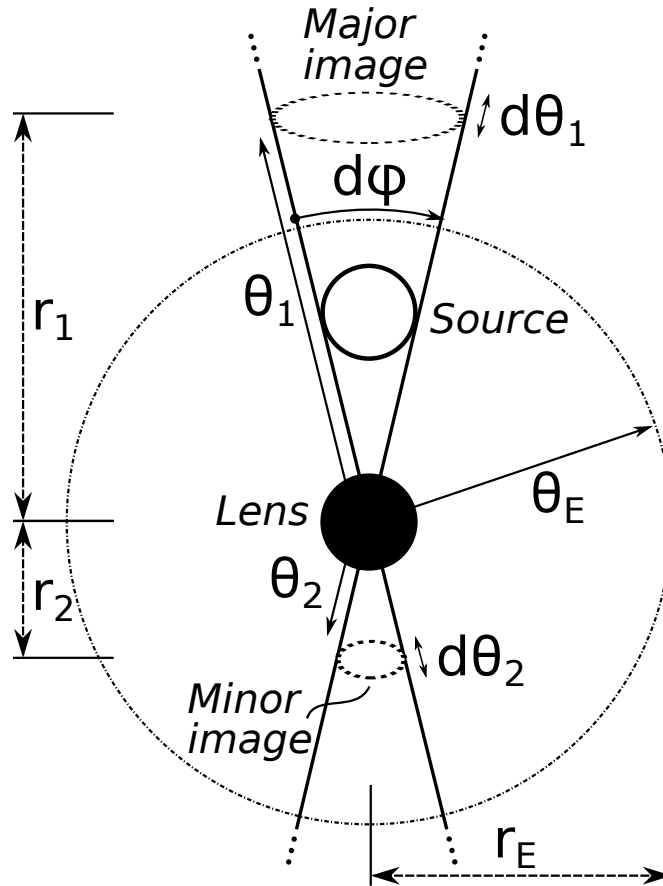


Figure 2.2: Diagram of lensing from the observer's point of view. Distances r are measured in lens plane. Proportions are greatly exaggerated for the purpose of illustration.

Definitions:

The coordinates we have been using so far tacitly acknowledge the lens mass as the pole of the coordinate system. α and θ are both polar angles from the observer-lens line; we now introduce ϕ , an azimuthal angle about the observer-lens axis; see fig. 2.2.

Derivation:

While the process of lensing changes an object's apparent polar angle, it has no effect on its azimuthal angle. Hence the object's "width" in azimuthal angle, $d\phi$, remains constant, but the spatial extent of that apparent width varies. The true location of the source is $\sin(\alpha)d\phi$ wide, but the source's image is $\sin(\theta)d\phi$ wide. The width therefore changes by a factor of $\sin\theta/\sin\alpha$. Because we are in the small angle limit, this can be approximated as θ/α , or

$$\frac{\theta_{1,2}}{\alpha} = \pm\frac{1}{2} + \sqrt{\left(\frac{1}{2}\right)^2 + \left(\frac{\theta_E}{\alpha}\right)^2} \quad (2.30)$$

To find the image's "thickness" in the θ direction, we consider how a change in α affects θ . We find this by taking the derivative of the lens equation found in section 2.3.

$$d\theta_{1,2} = d\alpha \left(\pm\frac{1}{2} + \frac{\alpha/4}{\sqrt{(\alpha/2)^2 + \theta_E^2}} \right) \quad (2.31)$$

We now introduce $u \equiv \alpha/\theta_E$; this quantity will be useful to us. We re-write the expressions for θ/α and $d\theta/d\alpha$.

$$\frac{\theta_{1,2}}{\alpha} = \left(\pm\frac{1}{2} + \frac{1}{2u} \sqrt{u^2 + 4} \right) \quad (2.32)$$

$$\frac{d\theta_{1,2}}{d\alpha} = \left(\pm\frac{1}{2} + \frac{u}{2\sqrt{u^2 + 4}} \right) \quad (2.33)$$

Multiplying together the change in the object's apparent width and the change in its apparent thickness, we find the multiplicative change in image size, A :

$$A_{1,2} = \left(\frac{\theta}{\alpha} \frac{d\theta}{d\alpha} \right)_{1,2} = \frac{1}{2} \pm \frac{u^2 + 2}{2u\sqrt{u^2 + 4}} \quad (2.34)$$

There are several things to note in this equation. First is that u is defined as positive and so image 1, the major image, is always larger than image 2, the minor image. Second, note that in the small u limit, i.e. when the source passes very close to the lens, the term on the right is approximately $1/u$. Hence for small u , the value of A_2 is negative; this reflects the fact that the second image created by lensing, the minor image, is reversed. We are interested in the size of the image rather than its orientation, so in finding the full magnification, we add the absolute value of the two images' size.

$$|A_1| + |A_2| = A = \frac{u^2 + 2}{u\sqrt{u^2 + 4}} \quad (2.35)$$

This quantity, A , is the magnification. It is the change in size of the source's image as seen by an observer. It is *also* the change in the amount of light seen by an observer, as we will demonstrate.

Consider figure 2.3, a visual representation of the finite spread of paths taken by light to the observer. If unlensed, the source subtends a solid angle $d\alpha d\phi \equiv d\Omega_S$ of the observer's view, but its appearance has been split into two images with sizes $d\theta_{1,2} d\phi \equiv d\Omega_{1,2}$ with total size $d\Omega$. Our last calculation showed that $d\Omega/d\Omega_S = (u^2 + 2)/(u\sqrt{u^2 + 4}) = A$. The physical extent of the source's image projected in the lens plane occupies an area equal to $D_L^2 d\Omega$, whereas its unlensed image would have size $D_L^2 d\Omega_S$.

For a moment, humor the notion (approximately correct, enumerated in our PSPL assumptions) that light deflects instantaneously at an angle of ϵ as it passes through the




Figure 2.3: As in figure 2.1, with the finite spread of light made explicit. Proportions are greatly exaggerated for the purpose of illustration.

lens plane. Whatever light is radiated by the star and passes through the image in the lens plane is redirected to the observer. Likewise we can imagine the unlensed case as light passing through an image in the lens plane, being deflected by an angle of zero, and continuing on to the observer. If this model holds true, then the amount of light sent to the observer by the image is proportional to the amount of light received by the image from the source star, and this holds true for any set of possible images in the lens plane, so long as the small-angle approximation holds and we can consider all lens-plane images to be equidistant from the source star. Hence the ratio of the lensed image size to the unlensed image size is also the ratio of light received by the observer in the lensed case to light received in the unlensed case; both quantities are the magnification A .

Please note that while it was necessary to assume the source had finite size for this derivation, the term “finite source effects” is used in microlensing to refer to a much more extreme case in which the source subtends a larger angle. This phenomenon will be briefly discussed in §2.7.

2.4.1 Phenomenology of microlensing light curves

We have derived the magnification A as a function of lens-source separation u ; we will now discuss the light curves that result from this. Original work on this subject can be found in [Paczynski, 1986].

We can derive the basic microlensing light curve from our $A(u)$ by imagining a lens passing a source on a linear trajectory. Define (x, y) to be the source's apparent position from the observer's view point projected onto the lens plane; define coordinates so that the source's y position is constant as it moves in x . Choosing $t = 0$ so that it is the time of closest passage and assigning $y = y_0$ and $x = \omega t$, we obtain

$$u = \frac{1}{r_E} \sqrt{(\omega t)^2 + y_0^2} \quad (2.36)$$

Moving the Einstein distance inside the square root, defining the “crossing time” $t_E \equiv r_E/\omega$, and defining $u_0 \equiv y_0/r_E$, this can be rewritten as

$$u(t) = \sqrt{\left(\frac{t}{t_E}\right)^2 + u_0^2} \quad (2.37)$$

In this form, it is easy to see that the minimum value of u is u_0 .

Choosing a particular value of u_0 and feeding the resulting $u(t)$ curve into $A(u)$ gives us the magnification as a function of time, $A(u)$. This function is plotted in figure 2.4 for a value of $u_0 = 0.6$.

The smaller the value of u_0 , the greater the corresponding peak magnification A_0 . We demonstrate this in figure 2.5, showing the light curves for a range of values of u_0 .

This iconic microlensing light curve, also known as the Paczynski curve, has several properties worth noting.

First, the light curve is symmetrical about the time of closest passage t_0 , chosen

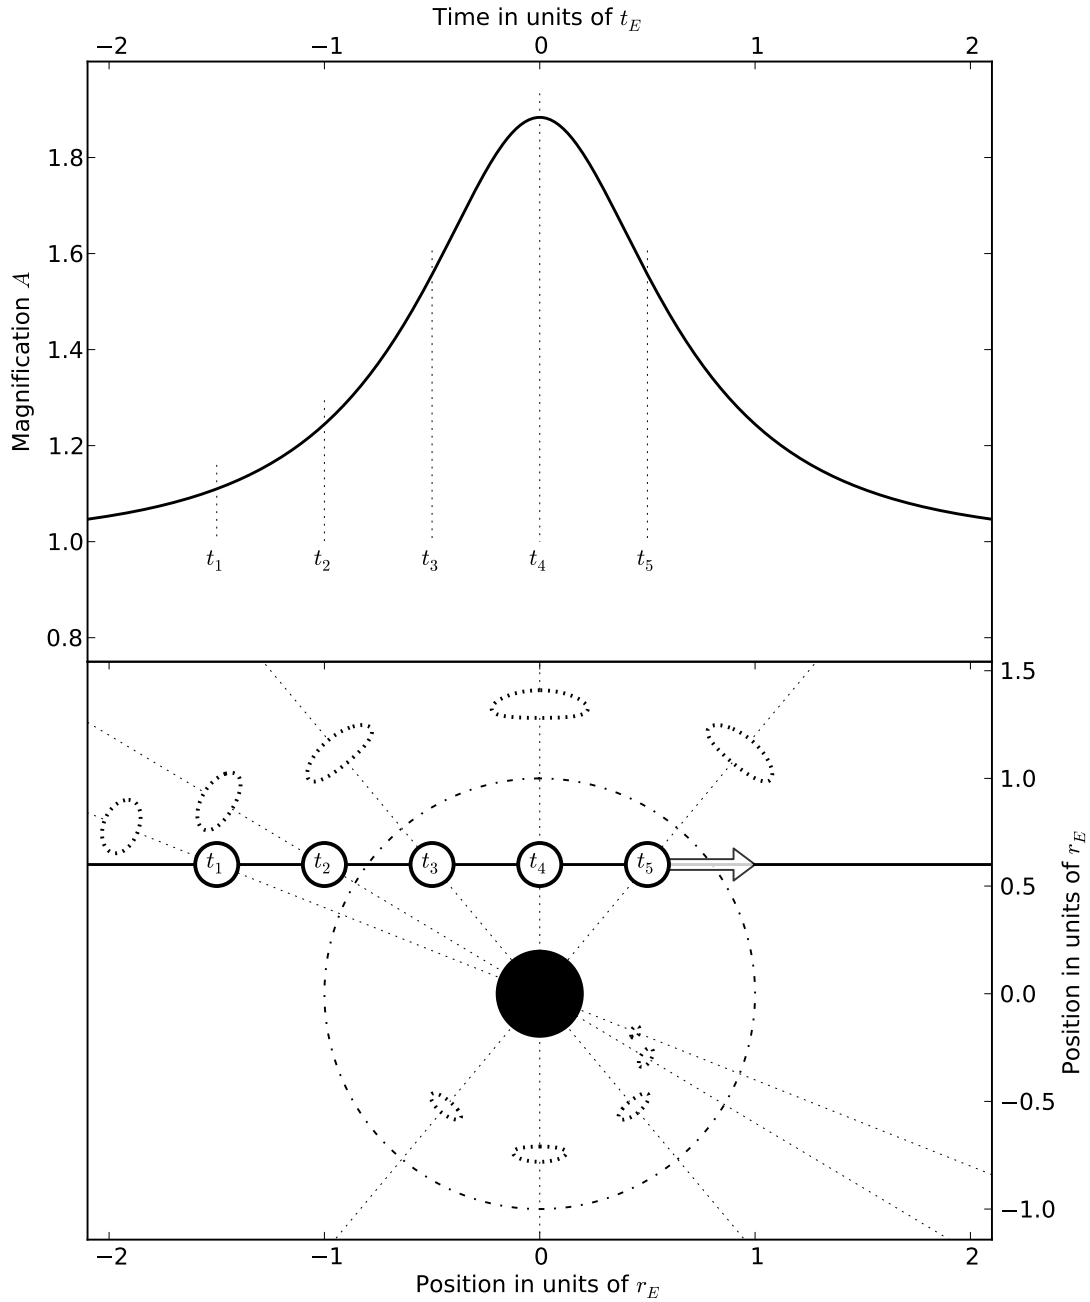


Figure 2.4: A path with $u_0 = 0.6$, showing images at selected times (bottom) and corresponding light curve (top). Images shown here are actual size, shape, and position for the portrayed source-lens geometry; dotted lines connect associated source positions and image positions.




Figure 2.5: Paths with different values of u_0 (bottom) and corresponding light curves (top).

here to be $t = 0$. Asymmetry in an observed light curve may indicate a parallax effect, described in §2.7, but more often simply indicates that the event being observed is not microlensing.

Second, there is no wavelength dependency in the bending angle and magnification equations; microlensing magnifies the source star without changing its color. This property is useful in distinguishing microlensing events from other transients; it is also important to disentangling microlensing magnification from any other unresolved stars in the vicinity, including the lens star. Taking at least some multi-band photometry is standard in microlensing observation.

Finally, there are three quantities that are directly determined from this light curve: the time of closest approach t_0 , a parameter that tells us nothing about the properties of the stars, the peak magnification A_0 , and the time scale or crossing time t_E . While they do provide information, these parameters do not directly determine the value of any of the physical properties of the system; each depends on multiple quantities. Writing these out explicitly, and invoking $D_L + D_{LS} = D_S$ to simplify θ_E :

$$A_0 = A(u_0) = \frac{u_0^2 + 2}{u_0 \sqrt{u_0^2 + 4}} \quad (2.38)$$

$$u_0 = \alpha_0 / \theta_E = \alpha_0 \left[\frac{4GM}{c^2} \frac{D_{LS}}{D_L D_S} \right]^{-1/2} \quad (2.39)$$

$$t_E = \theta_E / \omega = \frac{1}{\omega} \left[\frac{4GM}{c^2} \frac{D_{LS}}{D_L D_S} \right]^{1/2} \quad (2.40)$$

These quantities mix information about the apparent position and motion of the source (α_0, ω), the lens mass M , and the distances to the lens and source (D_L, D_{LS}). The light curve of an archetypal lensing event does not, by itself, provide enough information for

the study of the lens. For this, we will generally need additional observations, or a special case that provides additional information; we will describe several such cases in §2.7.

2.5 Binary mass lensing

Although not strictly necessary to understand the rest of this thesis, this section will help the reader understand how planets are detected using microlensing. The search for planets is a driving motivation in many current microlensing projects, and a basic understanding of binary lensing will be needed to read many papers in the field. Original work on binary lenses can be found in [Mao and Paczynski, 1991]; the planetary case described here can be found in [Gould and Loeb, 1992].

We take the PSPL assumptions and add the presence of a secondary mass or planet, $m \ll M$, near the primary lens. We require the second mass to be close enough to the lens that we can approximate it as being in the lens plane. Because we have assumed $m \ll M$, we will treat the effects of the second lens as a perturbation on the effects of the first lens. These assumptions together we will refer to as the planetary approximation.

We imagine the source's true position being lensed by the primary lens, and then the resulting major and minor images being lensed by the secondary lens. In most cases, only one image passes close enough to the planet to be significantly affected; exceptions are most likely to occur for very small u_0 , where both major and minor image sit close to the Einstein radius. We can approximate this using the same tools as in the PSPL case, but they grow increasingly inaccurate and the "source" being lensed is no longer so simple in shape. More typically in scientific work, the actual image shapes are computed and their area integrated to find the resulting total image area. An application of the planetary approximation is demonstrated in figure 2.6.

Although a double application of PSPL transformations is not accurate enough for




Figure 2.6: Diagram of planetary lensing from the observer’s point of view. On the left, we show lensing in the approximation that only the primary lens mass is considered. On the right, we show lensing in the approximation that one image is affected by the secondary mass. Image size, shape, and positions are accurate for the shown geometry and assumptions, but proportions are greatly exaggerated for the purpose of illustration; shown here, $r_{E,m} = 0.25 r_{E,M}$, in which case the planetary approximation would not ordinarily be applied because it is insufficiently accurate.

most scientific work, it can be used to gain some basic insight into how the resulting light curve will appear. Recall that the physical scale of lensing $r_E \propto M^{1/2}$; likewise $t_E \propto M^{1/2}$. A planetary lensing curve can look much like a microlensing light curve with another, shorter microlensing event occurring in the midst of the event; Gould and Loeb [Gould and Loeb, 1992] demonstrate a theoretical example where this approximation is close to reality. By comparing the timescales of the two microlensing peaks, we can obtain an estimate of their relative masses:

$$t_{E,m}/t_{E,M} \approx (m/M)^{1/2} \quad (2.41)$$

2.5.1 Planet detection

A planet in orbit about a lens star will affect the light curve if one of the images passes near it in the lens plane. Because the minor image goes to the lens star's position at large t , this can include planets in close orbits about the lens; because the major image goes to the source's true position at large t , this can include planets in distant orbits about the lens star. However, the orbits best covered are those near the Einstein radius. For most of a high-magnification lensing event, the major image is just outside the Einstein radius, and the minor image is just inside.

The Einstein radius r_E is typically of order a few astronomical units (see §2.8), significantly larger than the typical orbital radius of a planet discovered by radial velocity or transit (REF). Unlike RV and transits, microlensing places no requirement on the orbital orientation of a planet; any orbital orientation has a chance to be detected, and a face-on orbit near the Einstein radius has a high detection probability indeed. Hence, lensing detection is most likely in a region of parameter space that is very difficult to reach with other methods.

More importantly, though, microlensing is the *only* planet-finding technique that does not require the observer to directly observe the host star to detect the planets orbiting it. This allows us to find planets about stars smaller and further away than other techniques are sensitive to.

More extensive discussion of planetary detection probability and how it depends on the properties of a microlensing event or of the lens system can be found in [Gould and Loeb, 1992, Peale, 2001, Han and Kang, 2003, Han, 2011, Jung et al., 2014].

2.6 Microlensing rate

We will now derive the microlensing rate, which will be a major focus of our thesis. A discussion of lensing rate can be found in [Han, 2008].

Lensing events can be conceptualized as a two-dimensional collision rate problem. If a source star (particle type 1) passes within a certain angular distance (cross section σ) of a lens star (particle type 2), a microlensing event occurs (a collision). If all quantities were constant and spatially uniform, the resulting rate could be calculated in seconds using a simple collision rate equation.

$$\Gamma = n_1 n_2 \sigma v \quad (2.42)$$

The space in which these collisions occur is not three-dimensional; it is the two-dimensional canvas of the sky, whose angular coordinates from the observer's point of view we will call (l, b) . Where $n(D, l, b)$ is the volume density of stars at a distance D from earth, the contribution $d\eta$ to the angular density of stars η from a volume element $dV = D^2 dD d\Omega$ is given by

$$d\eta d\Omega = n D^2 dD d\Omega \quad (2.43)$$

There is no angular distance where an alignment of stars goes abruptly from lensed to not-lensed, but convention is to use $\alpha = \theta_E$ as the dividing line. As we will see in §2.8, this corresponds to a magnification of $A \approx 1.34$. Our collisional cross section, then, is $\sigma = 2\theta_E$.

As for v , the first step is to find each object's velocity relative an observer on earth; we subtract away earth's velocity, \vec{v}_\oplus . We do not care about the portion of the velocity that is parallel to our line of sight (\vec{v}_\parallel , in the \hat{D} direction), only the portion that is perpendicular (\vec{v}_\perp , in the \hat{l} and \hat{b} directions). Angular velocity ω is given by v_\perp/D , and

the relative angular velocity of our lens star and source star in two-dimensional angular space is

$$|\vec{\omega}| = \left| \frac{(\vec{v}_L - \vec{v}_\oplus)_\perp}{D_L} - \frac{(\vec{v}_S - \vec{v}_\oplus)_\perp}{D_S} \right| \quad (2.44)$$

We return to the rate equation with these values in hand:

$$\Gamma(l, b) = \int 2 n_S n_L \theta_E |\vec{\omega}| D_S^2 D_L^2 dD_S dD_L \quad (2.45)$$

An astute reader will note that this expression does not explicitly integrate over a number of variable quantities that we will need to account for in a detailed calculation: θ_E depends on M as well as D_S and D_L , the velocities \vec{v} of stars are by no means uniform, and a full calculation of Γ will require integration over l and b . It is also important to note that the integrand is not separable.

2.7 Second-order microlensing effects

There are a great many other variations on microlensing that do not directly pertain to this thesis, but will be mentioned here to give the reader an idea of the effects that can alter the iconic microlensing curve. Many of these effects would be novelties to a planet hunter or stellar mass measurer if it were not for their value in breaking degeneracies in microlensing curve fits, allowing for more precise mass measurements. We describe them here in brief.

2.7.1 Parallax

As we will explain in §2.6, any combination of straight-line motions for observer, lens, and source can be reduced to a single apparent straight-line motion, like the one we used in §2.4.1 to produce the iconic microlensing light curve. However, straight-line motion

is never quite an exact description; the observer, sitting on earth's surface, is rotating about earth's center, and earth is orbiting the sun. When the observer's deviation from a straight-line motion is great enough to cause observable changes in the light curve, it is known as parallax.

There are several ways to achieve enough deviation from a straight-line trajectory to cause observable parallax; all fall under the common theme of observation events not on a Newtonian inertial trajectory. In long timescale events, with t_E of order a year, the orbit of the earth about the sun can cause parallax. In high magnification events, where a small change in u can make a significant change in A , smaller deviations like the rotation of the earth can be enough to produce measurable parallax. Likewise, for events with small r_E and θ_E , a small change in observer position can cause a significant change in u .

Because we know the motion of earth-based and satellite observatories to high precision, we can translate an observed parallax into a measure of the angular Einstein radius θ_E [Gould and Yee, 2014]. D_S can generally be measured well using the color of the source star to infer its magnitude [Yoo et al., 2004], but there is still a D_L - M degeneracy in θ_E . Follow-up astrometry in which the direction [Gould and Yee, 2014] or net speed [Yoo et al., 2004] of source-lens separation is measured can provide the final piece of information needed to fully disentangle the physical parameters of a lensing event.

Long- t_E light curves with parallax are often noticed due to a slight asymmetry, as was the case with the first parallax detection [Alcock et al., 1995]. In cases of high magnification or small r_E , observers from different locations on or around earth may see significantly different light curves that become consistent when modeled as a parallax event. The further from the earth-average trajectory, in these cases, the better the sensitivity to parallax; satellites in orbit about earth, for instance, can give greatly distanced vantage points from earth-based observatories. Parallaxes that could not be

detected from ground-based data alone have been observed by incorporating data from satellite observatories such as the Spitzer Space Telescope [Yee et al., 2014], and it has been proposed that most microlensing events will have measurable parallax if space-based observations are combined with earth-based data [Gould and Yee, 2014].

2.7.2 Xallarap

Xallarap is the opposite of parallax. It occurs when the observer’s acceleration is not significant to the light curve, but the source or lens undergoes significant acceleration, typically due to being a member of a binary star system.

Note that this effect appears in systems very different from the planetary binaries described in §2.5. The planetary case requires $M \gg m$, and that their separation be of order r_E . Xallarap, by contrast, is strongest when the lens star accelerates most, when the binary companion is massive and/or close. Detectable xallarap may occur when a massive planet is in a close orbit a small fraction of r_E from the star [Rahvar and Dominik, 2008], or when the companion is of comparable mass at a more distant orbit [Miyake et al., 2012].

Xallarap can serve much the same function as parallax in breaking degeneracies if observers are able to determine the orbits of the stars involved.

2.7.3 Multiple lens masses

The $M \gg m$ binary lens case, typical of planetary observations, can be further complicated by additional planets in the sensitive region. As the number of lens masses goes up, so does the potential complication of the light curve. Multiple local maxima, abrupt local maxima, plateaus, and wildly asymmetric light curves are all common results of multi-lens systems; the phenomenology of these systems is very rich and far beyond

the scope of this thesis. Theoretical triple-lens curves are discussed in [Song et al., 2014], and an observed multi-lens system is analyzed in [Gaudi et al., 2008a, Bennett et al., 2010].

2.7.4 Finite source effects

When $d\alpha$, the angular size of the source star, is not small enough to allow the approximations used in §2.4, the result is known as “finite source effects” in the microlensing community. While one part of the source is strongly magnified, the remainder is only slightly magnified. The net effect is a light curve that plateaus near a constant value while the source star crosses the lens. As the source approaches the lens, its magnification increases, but once its near edge reaches the peak-magnification position, the total magnification will remain almost constant until the far edge passes that position; as the image leaves the lens, the magnification goes back down to resting values. A more detailed discussion can be found in [Gould and Gaucherel, 1996].

Finite source effects are often observed with red giant source stars, due to their large radius, but can occur for a source star of any size if it passes close enough to a magnification peak. The width of the plateau region allows an observer to determine the ratio of the source size $d\alpha$ to angular Einstein radius θ_E [Gould and Gaucherel, 1996].

2.8 Numeric values of microlensing quantities

It is useful to have a sense of just how large the quantities involved in microlensing are, so we include this section giving worked examples of typical values. Note also that we prefer to express quantities in units that render them a small number, between one and twenty or so, as it is helpful to human intuition to work in small numbers; thus this section will reveal our motives for many of our favorite units throughout this thesis.

The masses and radii of stars we will measure by comparing them to our sun. $1M_{\odot}$ is the mass of the sun, and $1R_{\odot}$ is the sun's radius.

The deflection of light by a star ϵ that we calculated in §2.2, a first order approximation, has been experimentally verified as accurate to a few parts in 10^4 [Bodenner and Will, 2003]. We express it here in arcseconds, notated with $''$.

$$\epsilon = 1.7504'' \cdot \left[\frac{M}{M_{\odot}} \cdot \frac{R_{\odot}}{r} \right] \quad (2.46)$$

The bending of background star light around our sun was the first major experimental proof of general relativity.

Rearranging the equation for the Einstein radius r_E , using the identity $D_L + D_{LS} = D_S$, we can write

$$r_E = \sqrt{\frac{4GM}{c^2} \frac{D_L D_{LS}}{D_S}} \quad (2.47)$$

Using a substitution $D_L = xD_S$ and $D_{LS} = (1 - x)D_S$, where x can take values from zero to one, gives us an illuminating form for this equation:

$$r_E = \sqrt{\frac{4GMD_S}{c^2} x(1 - x)} \quad (2.48)$$

If we hold D_S constant and take the derivative with respect to x , we will find that r_E is greatest for $D_L = 0.5D_S$. This provides a useful rule of thumb for us; all other things being constant, the lens is most likely to be halfway between the observer and the source.

We find r_E 's value in astronomical units, or AU:

$$r_E = 4.0346 \text{ AU} \cdot \left[\frac{M}{M_{\odot}} \cdot \frac{D_L D_{LS}}{D_S \cdot 2 \text{ kpc}} \right]^{1/2} \quad (2.49)$$

For most microlensing observations, done towards the galactic bulge, the source is in the

bulge itself and the lens is about halfway between us and the source; $D_S \approx 8$ kpc, $D_L \approx 4$ kpc, and $D_{LS} \approx 4$ kpc; we will measure most lens and source distances in kiloparsecs. Hence a typical value for a sun-like star towards the bulge would be:

$$r_E(1M_\odot) \approx 4 \text{ AU} \quad (2.50)$$

For this same sun-like star towards the bulge, we can calculate the angular Einstein radius $\theta_E = r_E/D_L$. Arcseconds are far too large; we'll express this in milliarcseconds, abbreviated mas.

$$\theta_E \approx 1 \text{ mas} \quad (2.51)$$

Note that the resolution of ground-based telescopes are generally of order a few arcseconds, and world-class telescopes in favorable weather can reach resolutions of about 0.5". A typical bulge microlensing event, therefore, is not resolvable; the two images of the source star and any light that may be coming from the lens will be unresolved.

In §2.6, we use $u_0 \leq 1$ as the minimum threshold for a microlensing event. The corresponding peak magnification is $A_0 \geq 1.3416$.

At very small u , we have previously noted, the value of A approaches $1/u$; the first-order Taylor expansion is

$$\lim_{u \rightarrow 0} A(u) = \frac{1}{u} \left(1 + \frac{3}{8}u^2 \right) \quad (2.52)$$

Note, however, that this expression of $A(u)$ is only accurate as long as the PSPL approximations hold, and the smaller the value of u , the greater the likelihood that those assumptions will fail.

We can obtain a very naive estimate of microlensing rates towards the bulge using the equations found in §2.6. We assign “typical” values for a bulge-ward microlensing event, taking n_S near the center of the bulge using a model from [Jurić et al., 2008], n_L

from halfway to the bulge, θ_E for a solar mass lens, and choose a value for ω using a model from [Bond et al., 2010]. This gives $n_L \approx 0.16 \text{ pc}^{-3}$, $n_S \approx 4.1 \text{ pc}^{-3}$, $\theta_E \approx 1 \text{ mas}$, and $\omega \approx 2.3 \text{ mas/yr}$; we treat these quantities as constants with this typical value. This gives a very rough estimate of

$$\Gamma \approx 2 n_S n_L \theta_E \omega \frac{1}{3} D_S^3 \frac{1}{3} D_L^3 \approx 80 \text{ year}^{-1} \text{ deg}^{-2} \quad (2.53)$$

towards the bulge.

Chapter 3

Introduction to Numerical Integration

The term “Monte Carlo method” can technically be applied to almost any method of computation or solution-finding that involves repeated random selection of variable values. In this thesis, we use Monte Carlo integration (henceforth MC integration), one of the simpler MC techniques, as a major component of our simulation. This chapter will help the reader understand both our choice of integration technique and the math behind it. We will describe the basics of numerical integration, compare the effectiveness of various techniques when applied to different integrands, and demonstrate the inherent difficulties of multidimensional numeric integration. We will also describe techniques for increasing the efficiency, and therefore the accuracy, of our numerical integration.

3.1 Numerical integration methods

When it is not possible to analytically solve an integral, the alternative is numerical integration. I will cover one-dimensional integration in §3.1.1, then describe the

generalization to multiple dimensions in §3.1.2.

3.1.1 One-dimensional integration techniques

The simplest of all numerical integration methods is integration on a grid. The method used is to divide the range into N small, equal intervals, use a polynomial of order k to approximate the function in each interval, and analytically compute the integral over that polynomial approximation. A more detailed description of these methods can be found in [Press et al., 1992, p. 130-136].

In practice, this boils down to a very simple algorithm, where each interval is sampled at evenly spaced points, the values at those points are multiplied by constants, and the lot are added together.

$$F = \int_{x_1}^{x_2} f(x)dx \approx \sum_i^N \sum_j^h \Delta x_{ij} c_j f(x_{ij}) = \frac{x_2 - x_1}{N} \sum_i^N \sum_j^h c_j f(x_{ij}) \quad (3.1)$$

where

$$\Delta x_i = (x_2 - x_1)/N$$

$$\Delta x_{ij} = \Delta x_i/h$$

$$x_i = x_1 + i\Delta x_i$$

$$x_{ij} = x_i + j\Delta x_{ij}$$

See figs. 3.1 and 3.2 for illustration.

In general, a method of order k requires $k + 1$ terms within each interval Δx_i , i.e. $h = k + 1$. In some cases, the symmetry of the polynomial being used to approximate the integrand allows some terms to cancel, allowing for a integral of order k to be performed with fewer than $k + 1$ terms. Notably, an integral of order 3 requires only 3 terms, and




Figure 3.1: A function which would require numerical integration, divided into equal intervals for integration on a grid. For a 0th or 1st order approximation, the values at the marked locations are the only samples needed.

an integral of order 5 requires only 5 terms.

This method is exact for integrands that are polynomials of order k or less; for other functions, their error is of order $\Delta x^{k+2} f(x)^{(k+1)}$, where $f(x)^{(n)}$ is the n th derivative of $f(x)$.

$$\Delta F_i \approx \Delta x_i^{k+2} f(x_i)^{(k+1)} \quad (3.2)$$

$$\Delta x_i = \text{const.}$$

Adaptive integration improves on the accuracy of grid integration by sampling more densely in the regions where the integrand diverges heavily from the polynomial approximation, i.e. where the $(k+1)$ th derivative $f(x)^{(k+1)}$ is largest. Once again, the integration




Figure 3.2: The same function as fig. 3.1 with the same number of intervals, but a higher-order approximation applied. Each interval is sampled at three evenly spaced locations; the value of the function at these locations is used to approximate the function locally as a 3rd order polynomial, which in turn is analytically integrated. This approximation is generally more accurate than a 1st-order approximation based on the same number of samples.

is given by

$$F = \int_{x_1}^{x_2} f(x)dx \approx \sum_i^N \sum_j^h \Delta x_{ij} c_j f(x_{ij}) = \sum_i^N \frac{\Delta x_i}{h} \sum_j^h c_j f(x_{ij}) \quad (3.3)$$

where

$$\Delta x_{ij} = \Delta x_i / h$$

$$x_{ij} = x_i + j\Delta x_{ij}$$

but we can no longer make the simplifications that depend on the Δx_i being equal.

By adjusting the step size Δx_i for each step, the error estimate for each step $\Delta F_i \approx$

$\Delta x_i^{k+2} f(x)^{(k+1)}$ is kept within some acceptable range; a predetermined standard of accuracy is maintained, and excessively fine steps over well-approximated regions are avoided. The effectiveness of this approach ultimately stems from minimizing the summed squared errors of the individual steps, $\sum \Delta F_i^2$, relative to the number of steps taken, by bringing the ΔF_i to a near-constant value. The resulting error per step is the same as the one given for grid integration, but the step size is chosen so the error estimates each interval are nearly equal.

$$\Delta F_i \approx \Delta x_i^{k+2} f(x_i)^{(k+1)} \approx \text{const.} \quad (3.4)$$

A demonstration of this method, and one of its potential pitfalls, is given in fig. 3.3.

Most methods of adaptive integration will therefore require the integrand to be differentiable $k + 1$ times so that an error estimate can be obtained. This is generally not a problem, although it can become one when integration over data is required. A more comprehensive description of an adaptive integration algorithm can be found in [Press et al., 1992, p. 707-722].

More importantly to the general problem of integration, the sampling locations are not known beforehand and are chosen based on local conditions. Two one-dimensional paths through a multi-dimensional integrand may choose very different samplings, which will make generalizing to an n -dimensional integration difficult. This issue presents a major barrier to attempts to generalize one-dimensional adaptive integration.

Simple Monte Carlo integration, by contrast, chooses its samples without any reference to the locations of already-existing samples. In the most basic possible implementation, sample locations x_i are drawn randomly from a flat probability distribution between x_1 and x_2 . The integral is then given by

$$F = \int_{x_1}^{x_2} f(x) dx \approx \frac{x_2 - x_1}{N} \sum_i^N f(x_i) = (x_2 - x_1) \langle f \rangle \quad (3.5)$$




Figure 3.3: Simple adaptive sampling of an integrand which is to be evaluated using a first-order approximation. At each step, the second derivative is computed and used to estimate an appropriate step size. Practical adaptive integration schemes will include safeguards against very large steps, such as the one near the minimum on the left, caused by a sample falling near an inflection point in the integrand, where the curvature changes from positive to negative. Samples do not fall very close to the other two inflection points on each side of the peak, and so the sampling there is smooth. If a sample were to fall directly on an inflection point where the curvature is equal to zero, this naive algorithm would choose an infinitely large step size. Avoiding this problem is the primary reason most adaptive algorithms do not force the error estimate for each step to be exactly equal.

with an error estimate of

$$\Delta F \approx \frac{x_2 - x_1}{N} \left(\sum_i^N (f(x_i) - \langle f \rangle)^2 \right)^{1/2} = (x_2 - x_1) \left(\frac{\langle f^2 \rangle - \langle f \rangle^2}{N} \right)^{1/2} \quad (3.6)$$

Note that, unlike previous equations which gave the one-step error estimate ΔF_i , this equation gives an error estimate for the entire Monte Carlo integral.

This simple MC integration is analogous to the $k = 0$ case of grid or adaptive inte-

gration. It has similar accuracy to 0th-order grid integration, but generally less accuracy than 0th-order adaptive integration. A demonstration of this sampling is given in fig. 3.4.




Figure 3.4: A simple Monte Carlo sampling of an integrand using the same total number of samples as fig. 3.3. The same sampling algorithm, if run again, would choose a different set of points for its sampling locations, and more intelligent randomized sampling is possible. In one-dimensional integration, simple MC integration is outperformed by methods that make better use of our knowledge of the integrand.

In one-dimensional integration tasks, MC integration is usually unnecessary (see §3.2.3 for exceptions) and not particularly accurate. It uses a poor-man's sampling algorithm (one that requires no record of past samples) combined with a poor-man's approximation (the lowest possible polynomial approximation). The relative simplicity that comes with such a low-effort algorithm, however, will become its greatest strength when confronting more complex problems.

3.1.2 Multidimensional integration

Generalizing one-dimensional grid or adaptive integration techniques to a multi-dimensional integration problem is not trivial [Press et al., 1992, p. 161-164]. The number N of function evaluations $f(x_i)$ needed for an n -dimensional integration scales as the n th power; an integrand that can be evaluated to sufficient accuracy with 100 samples in one dimension will require 10^6 samples for a three-dimensional evaluation of comparable fineness. We will elaborate on the need for more samples and discuss some further practical concerns in §3.1.2.

In addition to the increased number of required samples, boundaries of integration tend to become more complex as the dimensionality of the integrand increases, which leads to a drop in computational efficiency. We will discuss this at length in §3.2.1.

Adaptive integration like techniques can be applied to multidimensional integration, but the technique outlined in §3.1.1 cannot be applied practically to a repeated one-dimensional integration. One practical alternative is Markov chain Monte Carlo integration, which can be applied with relative safety to many-dimensional problems. Our own approach will be discussed in §3.2.2 and §3.3.2.

It is possible to reduce a multi-dimensional integration to a series of one-dimensional integrations, but generally inadvisable for an integrand of more than three dimensions due to the $N(n)$ scaling mentioned above. The organizational costs of the calculation also become increasingly overwhelming, for reasons explained in §3.1.2. The problem addressed in this thesis has of order ten dimensions to integrate over, and will need another method.

If the integrand is flat in any dimension, separable, or even reasonably simple and uncoupled from other dimensions, a canny investigator will search out ways to reduce the number of dimensions that must be simultaneously numerically integrated. The count of

approximately ten dimensions given above was made after every possible simplification.

Despite all of these issues, which present significant barriers to grid and adaptive integration in the style of §3.1.1, multidimensional MC integration remains almost as simple as one dimensional integration to implement. It shares the need for large samples counts N , but has minimal additional organizational costs. Where V is the n -dimensional volume integrated over, the Monte Carlo integral F is given by

$$F = \int_V f(\vec{x})dV \approx V \langle f \rangle \quad (3.7)$$

with an error estimate

$$\Delta F \approx V \left(\frac{\langle f^2 \rangle - \langle f \rangle^2}{N} \right)^{1/2} \quad (3.8)$$

where

$$\langle f \rangle = \frac{1}{N} \sum_i^N f(\vec{x}_i) \quad (3.9)$$

$$\langle f^2 \rangle = \frac{1}{N} \sum_i^N (f(\vec{x}_i) - \langle f \rangle)^2 \quad (3.10)$$

Computational overhead

Both grid integration and adaptive integration become computationally difficult in a space with more than a couple of dimensions. The reason is simple: you cannot reduce your data below a certain complexity if you need to keep track of where you've sampled the integrand and take that into account. For some problems - not all, but many-dimensional problems are especially prone to this, for reasons I'll explain below - it is entirely possible for this complexity to overrun the RAM of even modern computing clusters.

Suppose that you have an equation to integrate over in n dimensions, and a way of

checking how closely your answer mimics the truth, even though you do not know exact value of the integral. (This case is broadly analogous to the problem of simulating the galactic lensing rate; we cannot expect that calculation to be accurate unless certain other, better-known galactic properties are reproduced accurately.) You naively start out with a number of samples N that seems reasonable to run on the computing power you have, and check your metrics to see whether there are clear signs that the answer you got is far from the truth.

However, you find this falls short. Your metrics indicate undersampling even in lower-dimensional projections of the problem, ones you understand well. You need more samples.

Perhaps you simply need to double the sampling density in each dimension. In a one-dimensional problem, this is a very modest increase in density, but for our n -dimensional problem you must take 2^n times as many samples, for a new total of N' . With a value of $n = 10$, close to the value in this problem, this translates to three orders of magnitude more samples.

If you want to record all N' samples, you need 2^n times as much storage space. If your calculation needs to compare each sample to some fraction of the other samples, this will take at least 2^{2n} times as many operations. Common sorting algorithms have scaling efficiencies running from $\propto N \log_2 N$ to $\propto N^2$ for large N [Press et al., 1992, p. 329-338], so again you may get an increased cost of $n2^n$ to 2^{2n} .

And these changes in computation time assume that you haven't overrun the available computing resources entirely. At one extreme, you may just crash your computer trying to run the new large sample size all in one batch; you'll need to add in a lot of code structure to allow you to divide the job into smaller batches. At another extreme, your computer operating system or software can divide the task without further interference, but now needs to read and write data between hard disk and memory mid-calculation at

a rate much slower than it can read from RAM. While a good enough programmer can avoid the former case, ultimately the latter will come into play when your calculation gets large enough. It becomes unavoidable that at some sufficiently large N , the calculation's a in $t_{calc} \propto N^a$ grows.

In short, a very modest increase in sampling (such as doubling the sampling rate in each dimension) can lead to very dramatic, horribly expensive outcomes (like the calculation taking longer than writing the code did).

3.2 Computational efficiency in multidimensional problems

As the number of dimensions n of an integrand increases, the fraction of the enclosed n -volume that contributes significantly to the integral tends to decrease. There are two reasons that this tends to be true, and although neither is universal, both are common in physical problems, and both are relevant to our project. We will cover these in §3.2.1 and §3.2.2. We will discuss the special cases in which Monte Carlo integration methods are vastly superior in §3.2.3.

3.2.1 Complex boundary conditions

When nonlinear expressions of multiple variables, piecewise expressions, or worse, actual data dictate part of the integration boundaries of a problem, the “region of interest” that you want to integrate over can become quite complicated. An otherwise manageable integrand can be rendered non-analytic by such a choice of boundaries. When this happens, one way to deal with it is to choose a region of integration that entirely encloses the region of interest, and define the integrand to be zero outside the region of interest. We

define the fractional relevant volume (FRV) as the size of the region of interest divided by the size of the region of integration. A two-dimensional example is demonstrated in fig. 3.5.




Figure 3.5: A complicated two-dimensional region of interest (solid black) and simple regions of integration (dashed lines), a rectangle and an ellipse, that might be used to encompass it. The fractional relevant volume in this case is of order 0.5.

I'll give three simplistic examples to demonstrate how scaling dimensionality can affect the fractional relevant volume (FRV). First, a uniform function, completely filling the volume of integration; second, an n -dimensional top hat function that does not necessarily fill the volume of integration; third, an n -ball, i.e. an n -dimensional analogue of a spherical volume. All three of these examples are simple and handled analytically; they are intended to be illustrative of FRV, not of complex boundary conditions.

In all cases, the volume of integration will be bounded by uniform values in each

dimension, i.e., the integral will be over an n -dimensional analogue to a rectangular prism. These values will be labeled $a_1, b_1, a_2, b_2, \dots, a_{n-1}, b_{n-1}, a_n, b_n$, and the total volume of the integrated region will be

$$V = \int_V \prod_{i=1}^n dx_i = \int_V dV = \prod_{i=1}^n (b_i - a_i). \quad (3.11)$$

Example 1: Uniform function

Our integrand in this case is simply a uniform function $f(\vec{x}) = f_0$; integrating over it returns the volume of integration times a constant.

$$\int_V f_0 dV = f_0 V_n \quad (3.12)$$

The total volume we have integrated over is V_n , and the volume we needed to integrate over to find the final value is V_n . The fractional relevant volume is 1, the theoretical maximum.

Example 2: Top hat function

We define the function $f(\vec{x})$ as the product of n one-dimensional top hat functions $f_i(x_i)$, i.e.:

$$f_i(x_i) = \begin{cases} 0 & : x_i \leq a'_i \\ 1 & : x_i \leq b'_i \\ 0 & : x_i > b'_i \end{cases}$$

$$f(\vec{x}) = \prod_{i=1}^n f_i(x_i) \quad (3.13)$$

The integral over this function is

$$\int_V f(\vec{x}) dV = \prod_{i=1}^n (b'_i - a'_i) = V \prod_{i=1}^n \frac{(b'_i - a'_i)}{(b_i - a_i)} \quad (3.14)$$

Defining $c_i \equiv (b'_i - a'_i)/(b_i - a_i)$, we can write this as

$$\int_V f(\vec{x}) dV = V \prod_{i=1}^n c_i \quad (3.15)$$

The fractional relevant volume is

$$FRV = \prod_{i=1}^n c_i \quad (3.16)$$

We have not yet said anything about what the values of c_i are, but they represent fractional coverage in a given dimension. They can range in value from 0 to 1. If all values of c_i are one, the FRV is one; if even one value of c_i is zero, the value of the FRV is zero. We can gain some insight, however, by manipulating the equation for FRV, and by proposing distributions of values for c_i .

$$\log(FRV) = \log\left(\prod_{i=1}^n c_i\right) = \sum_{i=1}^n \log(c_i) \quad (3.17)$$

If the distribution of values of c_i is such that there is a well-defined average value of $\log(c_i)$, we can write

$$\langle \log(FRV) \rangle = n \langle \log(c_i) \rangle \quad (3.18)$$

and therefore a typical value of the FRV would be

$$FRV \approx \exp(\langle \log(c_i) \rangle)^n \quad (3.19)$$

where, we note, $\exp(\langle \log(c_i) \rangle)$ must be a number between one and zero.

If all values of c_i are the same number c between one and zero, the value of FRV is

$$FRV = c^n \quad (3.20)$$

While the actual value of FRV depends on the detailed choice of c_i , these two equations both follow the form of $FRV = a^n$, where a is a number between one and zero. We can safely guess that this is the general form of $FRV(n)$ in this case.

Example 3: N-ball

The general formula for the volume of an n-ball with radius R is

$$V_n(R) = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2} + 1)} R^n \quad (3.21)$$

where Γ is Euler's gamma function, a continuous analogue to the factorial, given by

$$\Gamma(t) = \int_0^{\infty} x^{t-1} e^{-x} dx \quad (3.22)$$

For even and odd integers n , respectively, $V_n(R)$ takes the form of

$$V_{2k}(R) = \frac{\pi^k}{k!} R^{2k} \quad (3.23)$$

$$V_{2k+1}(R) = \frac{2(k!)(4\pi)^k}{(2k+1)!} R^{2k+1} \quad (3.24)$$

We divide these volumes by the rectangular prism like volume V of space we must integrate over to contain them, $(2R)^n$.

$$FRV_{2k} = \frac{\pi^k}{k!} 2^{-2k} \quad (3.25)$$

$$FRV_{2k+1} = \frac{2(k!)(4\pi)^k}{(2k+1)!} 2^{-(2k+1)} \quad (3.26)$$

These relations are plotted in fig. 3.6.




Figure 3.6: Relationship between fractional relevant volume and number of dimensions for spherical analogue regions of interest enclosed by cubic analogue regions of integration. Note the logarithmic scaling. Odd integer n are marked with open circles, even n with closed circles; these are computed using the exact relations in eqns. 3.25 and 3.26. The trace in the background is the continuous approximation provided by eqn. 3.28.

To get a better idea of how these expressions scale with dimensionality, we will use Stirling's approximation for large factorials.

$$n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n \quad (3.27)$$

We will also substitute back in the number of dimensions, $n = 2k$ in the even case,

$n = 2k + 1$ in the odd case.

$$FRV_{even\ n} \approx \frac{1}{\sqrt{n\pi}} \left(\frac{e\pi}{2n}\right)^{n/2} \quad (3.28)$$

$$FRV_{odd\ n} \approx \sqrt{\frac{e}{n\pi}} \left(\frac{e\pi(n-1)}{2n^2}\right)^{n/2} \quad (3.29)$$

From these formulas, we conclude that an n -sphere's FRV goes by

$$FRV_n \approx ab^n n^{(n-1)/2} \quad (3.30)$$

where a and b are constants. Unlike the previous example, b here is a number greater than one; however, the exponential b^n and power law $n^{-1/2}$ are overwhelmed by the superexponential $n^{-n/2}$ term. This is visible in fig. 3.6; in logarithmic space, the relation $FRV(n)$ has negative curvature.

Putting it together

The uniform function and the top hat function initially appear to represent two conceptual extremes. One's FRV is unaffected by dimensional scaling, the other exponentially affected. The n -ball, however, reveals that the scaling of $FRV(n)$ can be even more punishing than a mere exponential relation, for even the simplest of geometric constructions.

While for one special case, when the region of interest exactly matches the region of integration, there is no affect of dimensional scaling, that is untrue of any case where the boundaries of the two regions do not exactly match. For all other cases, an increase in dimensionality implies a decrease in the FRV. As a rule of thumb,

$$FRV \propto a^n \quad (3.31)$$

where a is a number less than one and n is the number of dimensions.

Although we chose our boundaries of integration the “hard way” in these examples, there are of course easy ways to bound both the top-hat problem and the n-sphere problem so that the FRV is one. When dealing with multidimensional integrals, it is important to be on the lookout for such opportunities. We will, of course, make every effort to maximize the FRV of the problem described in this thesis by choosing our boundaries and sampling strategies carefully. Some of these methods will be described in §3.3.

3.2.2 High sensitivity regions

Recall from §3.1.1 that an integral technique of order k has an accuracy of order $\Delta F_i \approx \Delta x_i^{k+2} f(x_i)^{(k+1)}$ over a single step, where $f(x)^{(n)}$ is the n th derivative of $f(x)$. The one-dimensional adaptive integration technique we described allows a user to get higher accuracy for the same number of samples by varying the individual step sizes Δx_i , taking smaller steps and hence sampling more frequently where $f(x_i)^{(k+1)}$ is larger.

The same idea can be applied in Monte Carlo and multidimensional integration, although the implementation we will need to use is different. Broadly speaking, we will make it more probable that a sample will be chosen where samples are most needed, and we will decrease the statistical weighting of samples in that area. Mathematical details of this process will be given in §3.3. Our one-dimensional adaptive algorithm performed both tasks at once by decreasing Δx_i .

For all the same reasons that the fractional relevant volume FRV tends to decrease as the number of dimensions n increases, the fraction of the region of interest that contributes *significantly* to the integral also tends to decrease as n increases. Choosing our sample locations accordingly will give us the same kind of advantage adaptive integration has over grid integration in one dimension, but the payout is even higher.

3.2.3 Special cases

There are a few cases where Monte Carlo integration has a special advantage over other methods of numerical integration.

The first and most notable of these is when the integrand has a periodic component. Periodically sampling a periodically varying integrand can lead to drastic misinterpretations known as “aliasing”, as demonstrated in fig. 3.7. Monte Carlo sampling, even in its simplest implementation, is much more robust against aliasing than grid integration. Likewise, one-dimensional adaptive integration techniques need to be formulated with special care to avoid inappropriate step sizes.

MC integration is more easily adapted to problems with complex boundaries than other numeric methods. While any method using a polynomial approximation above order zero assumes some degree of smoothness and differentiability of the integrand, MC has no such requirement; hence choosing a larger volume of integration, and defining the integrand to be zero outside the volume of interest, has no negative impact on the functioning of MC integration. Sudden changes in value are not an issue.

Simple MC integration does not tend to work well for functions with strongly localized peaks; multidimensional numerical integration in general is often stymied by strongly localized peaks. This is because, as we explained in §3.2.1 and 3.2.2, the fraction of the integration volume taken up by that peak has a tendency to plummet as n increases, and sampling frequently enough near the peak is generally important to obtaining an accurate integral. Although $k = 0$ grid integration suffers from the same loss of accuracy, and neither technique can be adapted without some knowledge of the peak locations, we will describe ways to rephrase the integral and diminish the impact of an integrable peak on total integral accuracy in §3.3.2.

The integrand we will treat in this thesis is “peaky”, but only in a few of the ten




Figure 3.7: Possible effects of aliasing. Each graph shows an integrand in solid black line, tick marks representing sample locations, and a dashed line indicating a lower-frequency function that would produce the same sample values. Whenever the sampling does not have at least twice the frequency of the integrand’s variability, as is the case in the top graph, the integrand will appear to have lower frequency than it truly does, and usually a different phase as well. If the sampling is close to an integer multiple of the integrand’s variability, as in the bottom graph, the integrand may appear to have little or no variability, and the odds of badly misestimating the integral are high.

or so dimensions, and that peakiness can be reduced using the methods described in §3.3; this goes a long way towards lessening the primary objection to using Monte Carlo integration in our project.

3.3 Monte Carlo Integration

We will now cover important topics that often go unmentioned in descriptions of Monte Carlo integration: producing arbitrary probability distributions using a random number generator, and changing variables in MC integration. Both topics are extremely

important to coding an efficient MC integration. When computing resources are limited, such efficiency can make a considerable difference to the net accuracy of an integration.

Source: [Press et al., 1992, p. 316-328]

3.3.1 Mapping an arbitrary probability distribution to linear space

Most existing random number generation algorithms, including those included in the NumPy suite of mathematical tools that we used to code our simulation, are made to produce numbers drawn from a flat probability distribution, i.e. equally likely to be at any point, between zero and one. Tools elaborating on this theme are common - NumPy, for instance, has functions that accept arbitrary limits, or generate values at discrete intervals - but almost all can be mapped back to this central tool, a flat probability distribution.

Although it is common to need other probability distributions for practical purposes, the flat probability distribution can be transformed into almost any other sensible probability distribution with minimal difficulty. The only requirement is that the desired probability distribution be integrable in the region of interest. We will show how to do this here.

w is a random number between zero and one, the output of a basic random number generator. x is the parameter for which we are trying to produce a value, and $P(x)$ the probability distribution we wish to sample x with. x_1 and x_2 are the lower and upper limits we allow for values of x . We stipulate that the fraction of probability space between w and $w + dw$ be the same as the fraction of probability space between x and $x + dx$.

$$\frac{P(w)dw}{\int_{w_1}^{w_2} P(w')dw'} = \frac{P(x)dx}{\int_{x_1}^{x_2} P(x')dx'} \quad (3.32)$$

The integrals in the denominator of each side of this equation are equal to one. We simplify on the left, but leave the right in its current form, which will allow us to avoid errors due to misplaced normalizing constants. We then integrate over each side, over a range of $w_1 = 0$ to w on the left, and over the corresponding x_1 to x on the right.

$$w = \frac{\int_{x_1}^x P(x') dx'}{\int_{x_1}^{x_2} P(x') dx'} \quad (3.33)$$

The right side of this equation is also known as the cumulative probability distribution. For practical purposes, using this expression generally means integrating over $P(x)$, evaluating, and solving for x in terms of x_1 , x_2 , and w .

For instance, sampling a power law distribution with a probability of the form $P(x)dx = ax^b dx$, we can evaluate and solve to obtain

$$w = \frac{(x^{b+1} - x_1^{b+1})}{(x_2^{b+1} - x_1^{b+1})} \quad (3.34)$$

$$x = (w \cdot (x_2^{b+1} - x_1^{b+1}) + x_1^{b+1})^{1/(b+1)} \quad (3.35)$$

Note that the probability distribution need not even be analytically integrable, so long as it is integrable in the sense of encompassing a finite region of parameter space. One useful example of this is a gaussian distribution. A flat distribution can be converted to a gaussian distribution by means of an error function, the integral over a gaussian. To do so oneself is more trouble than it is worth, since the programmer would need to numerically integrate the gaussian and then numerically interpret to the variable of interest, all with sufficient precision for whatever task was at hand; the result, coming from a person not formally trained in computational methods, would likely be inefficient as well. Fortunately for us, a gaussian distribution generator is a very common supplemental tool

in random number generation suites, including NumPy.

We will make extensive use of this probability distribution mapping technique throughout this thesis and the corresponding code.

3.3.2 Variable substitution in MC integration

Suppose you wish to integrate some function $f(\vec{x})$ over a spherical volume with radius r_0 centered on coordinate zero.

$$\int_V f(\vec{x}) dV = \int_{-r_0}^{r_0} dx \int_{-\sqrt{r_0^2-x^2}}^{\sqrt{r_0^2-x^2}} dy \int_{-\sqrt{r_0^2-x^2-y^2}}^{\sqrt{r_0^2-x^2-y^2}} dz f(x, y, z) \quad (3.36)$$

As we've noted previously, you can simplify the boundaries in this coordinate system by choosing a simple region of integration that contains the region of interest, e.g.

$$\int_V f(\vec{x}) dV = \int_{-r_0}^{r_0} dx \int_{-r_0}^{r_0} dy \int_{-r_0}^{r_0} dz \delta(\vec{x}) f(\vec{x}) \quad (3.37)$$

where $\delta(\vec{x}) = 1$ if \vec{x} is inside the volume of interest, and 0 if \vec{x} is outside the volume of interest. The Monte Carlo equivalent of this integral can be written

$$\int_V f(\vec{x}) dV \approx V \frac{\sum_i^N \delta_i(\vec{x}_i) f(\vec{x}_i)}{\sum_i^N \delta_i(\vec{x}_i)} \quad (3.38)$$

where the values of \vec{x}_i are obtained by choosing randomly a value of x between $-r_0$ and r_0 , y between $-r_0$ and r_0 , and z between $-r_0$ and r_0 .

We can rewrite the sums to count over i' , indicating only those samples that fall within the volume of interest, up to N' , the total number of samples that fall in the volume of interest. Although the exact value of N' is left to chance, we know it will be

approximately $FRV \cdot N$.

$$\int_V f(\vec{x})dV \approx \frac{V}{N'} \sum_{i'}^{N'} f(\vec{x}_{i'}) \approx \frac{V}{FRV \cdot N} \sum_{i'}^{N'} f(\vec{x}_{i'}) \quad (3.39)$$

This is somewhat inefficient. This scheme gives us a fractional relevant volume of $\pi/6 \approx 0.523$, meaning that nearly half of our samples fall outside the volume of interest and are effectively wasted.

We can improve on this scheme by choosing a coordinate system in which our spherical volume of interest is much simpler to express. We re-express the integral in polar coordinates.

$$\int_V f(\vec{x})dV = \int_0^{r_0} dr \int_{-\pi/2}^{\pi/2} d\theta \int_0^{2\pi} d\phi r^2 \cos \theta f(r, \theta, \phi) \quad (3.40)$$

Changing coordinates is not always a trivial business when one is trying to analytically solve an integral. An expression that is compact and elegant in one coordinate system can become sprawling and cumbersome in another. Not only is such a change aesthetically offensive, it can move the expression outside of the mathematician's repertoire, becoming something that can barely be comprehended, much less solved.

In numeric integration, however, this is not a concern. The expression being integrated doesn't have to be simple; no human mind needs to hold a working model of it. All that is needed is that the expression be evaluable by a computer, transcribable in code. Traditional grid and adaptive integration also ask that the integrand be differentiable in order to obtain an error estimate, but even that much is unnecessary when using MC integration.

We can rewrite our polar coordinates expression in MC terms as

$$\int_V f(\vec{x})dV \approx \frac{V}{N} \sum_i^N r_i^2 \cos \theta_i f(r_i, \theta_i, \phi_i) \quad (3.41)$$

where r_i is chosen randomly between 0 and r_0 , θ_i is chosen between $-\pi/2$ and $\pi/2$, and ϕ_i between 0 and 2π . Our FRV is now one. None of our samples are wasted. However, a disturbing pattern remains.

Our samples are more dense in integration volumes near $r = 0$ than in volumes with large r . A quick calculation will show that half of the samples (on average) are between 0 and $r_0/2$, a volume of $\pi r_0^3/6$, and the other half in the volume between $r_0/2$ and r_0 , a volume of $7\pi r_0^3/6$. The sampling in θ has a similar problem; samples are more dense in volumes near the poles, $\theta = \pi/2$ and $\theta = -\pi/2$, than near the equator, $\theta = 0$. Sampling in ϕ is still uniform.

This does not invalidate the calculation. The factor of $r_i^2 \cos \theta_i$ that has appeared next to the function $f(r_i, \theta_i, \phi_i)$ is counterbalancing it. Effectively, this term increases the weight of samples at high r above those at low r , increases the weight of samples near the equator above those near the poles. The integral will still converge to the right answer as N increases, for the exact same reasons that the analytic version of the integral, eqn. 3.40, is still correct.

The pattern remains disturbing. Suppose we know that the integrand does *not* have higher sensitivity at the poles than at the equator, does *not* have higher sensitivity at low r than high r . Our integral accuracy will be driven by the region with the highest sensitivity per sampling density; if sensitivity is similar throughout, we want sampling to be similar throughout. In the general case, where we have no special knowledge of the integrand's sensitivity, we would like the sampling to be uniform in V .

Fortunately, this isn't hard to achieve! We can do this simply by drawing our coor-

dinate values from probability distributions given by

$$P(r) \propto \frac{\partial V}{\partial r} \propto r^2 \quad (3.42)$$

$$P(\theta) \propto \frac{\partial V}{\partial \theta} \propto \cos \theta \quad (3.43)$$

$$P(\phi) \propto \frac{\partial V}{\partial \phi} \propto \text{const.} \quad (3.44)$$

where the relation between the probability on the left and the expression on the right becomes exact with the application of a normalizing constant such that the integral of P over the permitted range of values is equal to one.

These probability distributions we have found are, up to a constant, identical to the volume element prefix that appeared before the function $f(\vec{x})$ when we converted to polar coordinates. This is no coincidence. This change in our sampling is a reflection of the volume element.

For practical purposes, we will need to convert a flat, zero-to-one random number w to these probability distributions. We described how to do so in §3.3.1. The result in this case is

$$r = r_0 \cdot w_r^{1/3} \quad (3.45)$$

$$\theta = \arcsin(1 - 2w_\theta) \quad (3.46)$$

$$\phi = 2\pi w_\phi \quad (3.47)$$

When we are choosing our \vec{x}_i according to this scheme, the MC integral reverts to a

simpler form, no longer requiring a volume element prefix.

$$\int_V f(\vec{x})dV \approx \frac{V}{N} \sum_i^N f(r_i, \theta_i, \phi_i) \quad (3.48)$$

This integral is now as efficient as it can possibly become without additional information about the integrand. Although our work on improving this expression stops here, we have one more lesson to learn from this final re-phrasing.

We made this last change in the expression of the integral motivated by a physical phenomenon we could easily understand, the idea of a volume element. However, the same sort of operations can be applied regardless of motive, without compromising the convergence of the integral, as follows.

$$F = \int_V f(x_1, \dots x_n)dV \approx \frac{V}{N} \sum_i^N f(x_{1,i}, \dots x_{n,i}) \quad (3.49)$$

where flat sampling of the variables of integration is used

$$x_{1,i} = (x_{1,max} - x_{1,min}) \cdot w_{1,i} + x_{1,min} \quad (3.50)$$

...

$$x_{n,i} = (x_{n,max} - x_{n,min}) \cdot w_{n,i} + x_{n,min} \quad (3.51)$$

Choose some probability function, $P(x_1)$, you wish to use for sampling x_1 , and define

$$h(x_1, \dots x_n) \equiv f(x_1, \dots x_n)/P(x_1) \quad (3.52)$$

Redefine your sampling accordingly.

$$\int_{x_{1,min}}^{x_{1,i}} P(x_1) dx_1 = w_{1,i} \int_{x_{1,min}}^{x_{1,max}} P(x_1) dx_1 \quad (3.53)$$

$$x_{2,i} = (x_{2,max} - x_{2,min}) \cdot w_{2,i} + x_{2,min} \quad (3.54)$$

...

$$x_{n,i} = (x_{n,max} - x_{n,min}) \cdot w_{n,i} + x_{n,min} \quad (3.55)$$

And re-write the integral using this sampling as

$$F = \int_V f(x_1, \dots, x_n) dV \approx \frac{V}{N} \sum_i^N h(x_{1,i}, \dots, x_{n,i}) \quad (3.56)$$

This operation can be done for any integrable, nonzero $P(x)$ over the range of possible values of x . This operation does not change the value that F converges to with large N , but it does change the speed and accuracy with which the integral converges. A wise investigator will use these transformations to bring the integrand as close to a constant as possible, allowing the integral to converge as quickly and as accurately as possible.

Chapter 4

Astrophysical inputs to the code

4.1 Summary of inputs

The starting point and inspiration for our galactic model is the one used by [Han, 2008] to estimate field microlensing rates, but we strove to update and improve nearly every aspect of that model. We sought out up-to-date and data-driven models for all major components of the simulation.

We focus on the galactic disk and halo, as they are the most important components in the sparse fields we are considering.

Our model does not currently include a bulge component. It is not intended for and should not be used to determine lensing rates towards the bulge. Our science focus is the areas of the sky away from the bulge, and preferably, away from the galactic plane; our choices throughout the development of the code reflect this. In choosing among the available models of the galaxy, we gave precedence to those best supported by observations towards high galactic latitudes.

We could not find a single galactic model that included descriptions of every property important to microlensing. Properties essential to estimating lensing rates include the

spatial distribution of stars, the velocity distribution, stellar mass, luminosity, stellar color, and color-dependent extinction by dust. Other details such as stellar radius can provide additional information, allowing us to determine whether effects such as finite source extent will be observable.

In the absence of a coherent pre-made model, we found and combined multiple up-to-date sources describing various galactic properties and did our best to bring them into agreement with one another, changing as little as possible, favoring the stronger observational case where conflicts arose.

The coordinate systems we use and how they relate to one another are described in §4.1.1. We draw the physical locations of stars from the SDSS tomography papers, described in §4.1.2. Our model of velocity combines SDSS tomography with Sofue’s measurements over a wide range of radii R in §4.1.3. §4.1.4 is devoted to stellar masses, ages, and absolute magnitudes, assembled from a variety of sources. Schlegel et al.’s map of galactic dust forms the core of our dust model, described in §4.1.5 along with the associated extinction. We compare our model to previous work on the same subject in §4.2.

We hope to make the full code, including reference data, available to the scientific community.

4.1.1 Coordinate systems

In order to describe galactic properties in the simplest way possible, and to unify information from multiple sources, we employ several coordinate systems. One group of closely related coordinate systems has its origin on the galactic center and its $z = 0$ plane through the galactic disk; we describe a galactocentric cartesian, cylindrical, and polar system. Another group is centered on our sun; we describe solarcentric cartesian

and polar systems. We demonstrate the relationship between our solarcentric cartesian and galactocentric cartesian coordinate systems in figs. 4.1, 4.2, and 4.3.

Our coordinate choices throughout this thesis are motivated by a desire to express physical quantities as simply as possible. Galactic properties are easiest described in galactocentric cylindrical coordinates, due to the galaxy's approximate cylindrical symmetry. Observations from earth, by contrast, are best described using solarcentric polar coordinates, as are descriptions of lensing properties such as the Einstein radius. Galactocentric and solarcentric cartesian coordinates are useful for comparing absolute velocities of stars, and form a stepping stone between galactocentric cylindrical and solarcentric polar coordinates.

This profusion of coordinate systems would be a considerable obstacle in an analytically evaluated problem; in our numerically evaluated, machine-computed problem, it is only a minor inconvenience.

Galactocentric coordinates

From our vantage point about the sun, the galactic center is roughly eight kiloparsecs away in the direction of Sagittarius. Sagittarius A* (Sgr A*), a bright radio source and supermassive black hole [Ghez et al., 2008], is generally acknowledged as the best observable landmark for the center of our galaxy.

Most papers, however, do not carefully describe their coordinate systems, and some coordinate systems predate our knowledge of Sgr A*. In this thesis, our coordinates do not necessarily center precisely on Sgr A*; we set our coordinate zero at the hypothetical galactic center of mass, presumably very close to Sgr A* itself. We define galactocentric $(0, 0, 0)$ at this point.

The galactic bulge and halo each form rough spheroidal clouds of stars about the galactic center. Around this, a thin, fuzzy plane of stars known as the galactic disk

rotates, with density perturbations giving the appearance of spiral arms. We define galactic $z = 0$ to be in the plane of this disk.

We choose positive z pointing towards the conventional north galactic pole, and positive x from the galactic center towards the point in the $z = 0$ plane closest to our sun. We choose y so that the coordinate system is right-handed.

We note that in these coordinates, the galactic disk has approximate mirror symmetry about the xy plane, and rotates about the z axis. Galactic rotation is in the $-y$ direction at the location of the sun.

Our galactic cylindrical coordinates (R, ϕ, z) are easily expressed in terms of galactic x , y , and z . The z coordinate is identical in both systems.

$$R = \sqrt{x^2 + y^2} \quad (4.1)$$

$$\phi = \arctan(y/x) \quad (4.2)$$

$$z = z \quad (4.3)$$

Likewise, polar coordinates (r, θ, ϕ) are simply expressed, with ϕ identical to cylindrical ϕ .

$$r = \sqrt{x^2 + y^2 + z^2} \quad (4.4)$$

$$\theta = \arctan(z/R) \quad (4.5)$$

$$\phi = \arctan(y/x) \quad (4.6)$$

Solarcentric coordinates

Our sun orbits the galactic center at a distance $R_{\odot} \approx 8$ kiloparsecs, and is currently $z_{\odot} \approx 25$ parsecs above the galactic plane [Jurić et al., 2008]. This location,

$(x, y, z) = (8000, 0, 25)$ parsecs relative to the galactic origin, will be the coordinate zero for our solarcentric coordinate systems. To distinguish them from galactocentric cartesian coordinates, we name our solarcentric cartesian coordinates $(x_{sol}, y_{sol}, z_{sol})$.

One of the solarcentric systems we will describe is commonly called “galactic coordinates” by astronomers. It is chosen so that the galactic center lies at angular $(0, 0)$; the azimuthal angle l is called the “galactic longitude” and the polar angle b the “galactic latitude”. This system of angular coordinates is in wide use among astronomers and we will employ it extensively ourselves.

Unfortunately, it does not match up well with the axes we chose to simply describe the galaxy itself. Our sun is slightly out of the galactic plane. For an axis to run between the sun and the galactic center, our solar axes must be slightly tilted relative to our galactic axes. The difference between our galactocentric and solarcentric coordinates, then, is not just a translation, but a small rotation as well. The relative tilt between the two I will call Δb .

$$x_{sol} = (x - R_{\odot}) \cos(\Delta b) \quad (4.7)$$

$$+(z - z_{\odot}) \sin(\Delta b)$$

$$y_{sol} = y \quad (4.8)$$

$$z_{sol} = (z - z_{\odot}) \cos(\Delta b) \quad (4.9)$$

$$-(x - R_{\odot}) \sin(\Delta b)$$

where

$$\Delta b \equiv \arctan(z_{\odot}/R_{\odot}) \quad (4.10)$$

is a small angle, about a fifth of a degree. Figures 4.1, 4.2, and 4.3 demonstrate the relationship between our two cartesian systems graphically; table 4.1 demonstrates by

giving the same locations in different systems.




Figure 4.1: Relationship between our galactocentric and solarcentric coordinates in the galactic xy plane. The spiral symbol represents the galactic center, and the circle represents our sun.




Figure 4.2: Relationship between our galactocentric and solarcentric coordinates in the galactic xz plane. The spiral symbol represents the galactic center, and the circle represents our sun. The size of Δb is exaggerated for legibility.

We then define polar coordinates; distance from our sun is D , l is galactic longitude, and b is galactic latitude.

$$D = \sqrt{x_{sol}^2 + y_{sol}^2 + z_{sol}^2} \quad (4.11)$$

$$l = \arctan(y_{sol}/x_{sol}) - \pi \quad (4.12)$$

$$b = \arcsin(z_{sol}/D) \quad (4.13)$$




Figure 4.3: Relationship between our galactocentric and solarcentric coordinates in the solar xz plane. The spiral symbol represents the galactic center, and the circle represents our sun. The size of Δb is exaggerated for legibility.

Note that the azimuthal angle l does not have its zero along the positive x_{sol} axis, but rather the negative axis.

As our coordinate systems are generally not identical to the systems used in our various source papers, and on occasion the source papers do not explicitly describe their coordinates, we take care throughout this thesis to give values and equations in our own coordinate systems in the most unambiguous way possible.

Table 4.1 gives the coordinates of three notable locations within each of four coordinate systems we use.

	Galactic center	Solar location	Point on galactic plane nearest sun
(x, y, z) , [kpc, pc, pc]	(0, 0, 0)	(8, 0, 25)	(8, 0, 0)
(R, ϕ, z) , [kpc, deg, pc]	(0, \emptyset , 0)	(8, 0, 25)	(8, 0, 0)
$(x_{sol}, y_{sol}, z_{sol})$, [kpc, pc, pc]	(-8, 0, 0)	(0, 0, 0)	(0, 0, -25)
(D, l, b) , [pc, deg, deg]	(8e3, 0, 0)	(0, \emptyset , \emptyset)	(25, 0, $\Delta b - 90$)

Table 4.1: Coordinates of three locations in each of four coordinate systems discussed in §4.1.1.

Although astronomers also commonly use right ascension (RA) and declination (DEC) to indicate the location of stars, that earth-centric system is not useful for describing galactic positions intuitively. Our program does not use RA and DEC at all, but instead

indicates coordinates in galactic latitude and longitude.

4.1.2 Stellar number density model

Our model of the spatial distribution of stars comes primarily from [Jurić et al., 2008], which drew upon an SDSS catalog of 47.7 million stars over 6500 square degrees of sky and up to 20 kpc from the sun. The number of objects catalogued, the large volume probed, and the pointing out of the galactic plane align well with our scientific goal of estimating lensing rates away from the galactic center.

In addition, the completeness of the study was enormously appealing. While this paper, by itself, described only the physical distribution of stars in the galaxy, companion papers [Ivezić et al., 2008] and [Bond et al., 2010] described stellar metallicity and stellar velocities, respectively. We use Bond et al. 2010 extensively in our model of stellar velocities (see §4.1.3).

Jurić et al.’s model of the galaxy includes a thin disk, a thick disk, and a halo. They neglected a bulge component because their sample gives no insight on areas near the bulge; none of their sampled stars are within five kiloparsecs of the galactic center.

The number density of stars belonging to the thin and thick disks are each represented by double exponentials in z and R . The exponential length scale in R is notated L , and the height scale in z is called H ; L and H have different values in the two disks. The absolute normalization of disk density is provided by $n(R_{\odot}, z_{\odot})$ and f , where $n(R_{\odot}, z_{\odot})$ is the thin disk’s stellar number density in the vicinity of the sun, and f is the ratio of thick disk stars to thin disk stars in the solar neighborhood.

$$n_{thin}(R, z) = n(R_{\odot}, z_{\odot}) \cdot \exp\left(-\frac{R - R_{\odot}}{L_{thin}} - \frac{|z| - z_{\odot}}{H_{thin}}\right) \quad (4.14)$$

$$n_{thick}(R, z) = n(R_{\odot}, z_{\odot}) \cdot f \quad (4.15)$$

$$\cdot \exp\left(-\frac{R - R_{\odot}}{L_{thick}} - \frac{|z| - z_{\odot}}{H_{thick}}\right)$$

The number density of halo stars is represented by an ellipsoid power law, slightly flattened by the overall rotation of the galaxy. The ratio of halo height to halo width is given by q_H , the power law scaling by n_H , and f_H is the ratio of halo stars to thin disk stars in the solar neighborhood.

$$n_{halo}(R, z) = n(R_{\odot}, z_{\odot}) \cdot f_H \quad (4.16)$$

$$\cdot \left(\frac{R^2 + (z/q_H)^2}{R_{\odot}^2 + (z_{\odot}/q_H)^2}\right)^{n_H}$$

Table 4.2 holds the values of these parameters determined by Jurić et al. and used in our simulation to describe the galactic distribution of stars.

R_{\odot}	8000 pc		
z_{\odot}	25 pc	\pm	20%
L_{thin}	2600 pc	\pm	20%
H_{thin}	300 pc	\pm	20%
f	0.12	\pm	10%
L_{thick}	3600 pc	\pm	20%
H_{thick}	900 pc	\pm	20%
f_H	0.0051	\pm	25%
q_H	0.64	\pm	$\lesssim 0.1$
n_H	2.77	\pm	$\lesssim 0.2$

Table 4.2: Parameter estimates from Jurić et al. 2008, used in our simulation, and Jurić et al.’s estimated errors on those parameter values. Their value of R_{\odot} was not measured but taken from literature, and no error estimate was given.

In the future, we do hope to add a bulge component to our model. We consider this outside the scope of the present paper, though we do realize it would provide us additional credibility to reproduce known bulge lensing event properties. The primary barrier to this was the lack of empirically rigorous studies that included both the bulge and the

out-of-plane regions of the galaxy; every study we found focused on either the bulge or high latitudes, always to the exclusion of the other. This is unsurprising; studying the bulge and studying stars out-of-plane present vastly different technical challenges.

4.1.3 Velocity model

The primary source for our stellar velocity model was [Bond et al., 2010], companion paper to Jurić et al. 2008, the source of our stellar number density model. We drew additional, complementary information from [Sofue, 2012]. Bond breaks their model into two components, providing separate velocity distributions for disk and halo.

Bond et al. find that their data are consistent with non-rotating halo that has spatially uniform velocity dispersion, where the probability distribution in each dimension can be written as a normalized gaussian G with variance σ_i :

$$G(x, \sigma, x_0) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x - x_0)^2}{2\sigma^2}\right) \quad (4.17)$$

$$P(v_i)dv_i = G(v_i, \sigma_i, 0)dv_i \quad (4.18)$$

We use the values for the dispersions σ_i given in cylindrical coordinates rather than those given in spherical coordinates in part because the former are most directly derived from the data, and in part because the oblate spheroid of our halo density model has cylindrical symmetry but not spherical symmetry. These values are tabulated in table 4.3.

The stellar disk's velocity distribution is slightly more complicated. While the radial and vertical components are simply gaussian distributions with zero mean as in eqn. 4.18, the rotational speed requires a more complex description. Bond describes this as a pair

$\overline{v_{\phi,H}}$	$\lesssim 10 \text{ km s}^{-1}$
σ_R	135 km s^{-1}
σ_ϕ	85 km s^{-1}
σ_z	85 km s^{-1}

Table 4.3: Parameter estimates for the halo velocity model from Bond et al. 2010, used in our simulation. Their measurement of a net halo rotation component was uncertain and is neglected in our model.

of offset gaussians

$$P(v_\phi)dv_\phi = [0.75 G(\bar{v}_\phi, \sigma_{\phi,1}, v_1) + 0.25 G(\bar{v}_\phi, \sigma_{\phi,2}, v_2)] dv_\phi \quad (4.19)$$

where

$$v_2 = v_1 - 34 \text{ km s}^{-1} \quad (4.20)$$

$$0.75v_1 + 0.25v_2 = \bar{v}_\phi \quad (4.21)$$

This distribution is distinctly asymmetrical about the mean rotational velocity.

Additionally, Bond et al. describe a z -dependence of each quantity in the disk velocity distribution out to a distance of 5 kpc from the galactic plane of the form

$$f(z) = a + b |z|^c \quad (4.22)$$

Bond cautions, however, that the model is not reliable beyond a distance of 5 kpc from the galactic plane. With this warning in mind, we hold the properties of the velocity distribution constant at the 5 kpc values for heights higher than that. We list parameters for the disk velocity model, including z -dependence, in table 4.4.

Bond's sample has poor coverage in R near the galactic plane and hence no R -dependence, so we supplement their model with Sofue 2012, which deals solely with

	a	b	c
$\overline{v_\phi}$	-194	19.2	1.25
$\sigma_{\phi,1}$	12	1.8	2.0
$\sigma_{\phi,2}$	34	1.2	2.0
σ_R	40	5.0	1.5
σ_z	25	4.0	1.5

Table 4.4: Parameter estimates for the disk velocity model of Bond et al. 2010, used in our simulation. Bond gives uncertainty estimates of $\sim 10 \text{ km s}^{-1}$ for a , $\sim 20\%$ for b , and 0.1-0.2 for c . The dependence of each parameter on z is given by $f(z) = a + b|z|^c$.

the dependence of rotational velocity on R . We give the Bond paper precedence due to its larger sample size and due to its sample’s better coincidence with our target region, the area out of the galactic plane.

Sofue assembles data from many sources in order to determine the R -dependence of the galactic rotation over a wide dynamic range, and normalizes each source datum to assume the same solar radius and velocity.

The Sofue paper ultimately focuses on deriving a mass distribution for the galaxy rather than on the rotation velocity curve for its own sake, and as such does not explicitly provide a formula for the rotation curve. However, it does provide the data to which the rotation curve was fit, a set of 128 radius and velocity pairs. These are not raw measurements of individual objects but rather an intentionally smoothed curve formed by averaging over sets of objects that are close together in galactocentric radius, with the exception of the eight data points within 2 kpc of the galactic center; these were left in their original form in order to preserve a higher spatial resolution near the bulge.

In order to construct a rotation curve for our code, we connect the data forming the rotation curve using cubic spline interpolation. In addition, we add one point at the galactic center, chosen by extrapolating linearly between the two innermost points of Sofue’s data. We show the resulting rotation curve in fig. 4.4.

Bond models only the z -dependence of rotation speed, while Sofue models only




Figure 4.4: Two subfigures showing the dependence of rotation speed on radius R . Top: Linear radius scaling from zero to twenty kiloparsecs, the portion of Sofue’s velocity curve we will actually use. Bottom: Logarithmic scaling out to the maximum radius allowed by the data, ≈ 1 Mpc, far more than we need. Both: Data points from Sofue are shown as dots with the error bars representing the uncertainties he lists, and our interpolated curve is shown as a solid line. The point we added at $R = 0$ is also shown in the first subfigure as a triangle symbol sans error bars.

the R -dependence. Both papers use a value of eight kiloparsecs for R_\odot . Both the z - and R -dependencies could be applied without conflict if they assumed the same speeds for the local disk rotation and the same velocities for our sun. They do not, and so we give precedence to Bond's model, using Bond's local disk rotation speed and Bond's solar velocity, $(v_x, v_y, v_z)_\odot = (-10, -225, 7)$ km/s. We incorporate Sofue's R -dependence into Bond's model by multiplying Sofue's entire rotation curve by $v_{rot,Bond}(R_\odot)/v_{rot,Sofue}(R_\odot)$, and treating Bond's z -dependence and Sofue's R -dependence as multiplicative.

4.1.4 Stellar properties

Both the lens and the source in a microlensing event, in general, will be stars; by and large, they will be drawn from the same population. A microlensing simulation, then, should have a coherent, consistent description of both stellar mass and stellar luminosity. Unifying the two considerations will require us to make some provision for stellar age, as well; the largest and brightest stars are also the most short-lived, and the remains of spent stars litter the galaxy, no longer shining, but still gravitating.

Because we wish to unify several concerns - mass, luminosity, and age - into a single model, and we wish to do so over the entire range of the main sequence from the coolest hydrogen-burning dwarf (a likely lens star) to the hottest blue giant (a likely source star), we will need to assemble several sources of information into a coherent whole. Some of these sources take very simple forms; their lack of precision is a necessary price to pay for their generality and compatibility.

Initial mass

One of the first and simplest pieces we will need is an initial mass function (IMF), a relation describing how often stars of a given mass are formed. We draw this relation

from [Kroupa et al., 1993], who gives a three-part broken power law

$$\xi(m) = \begin{cases} 0.035m^{-1.3} & \text{if } 0.08 \leq m < 0.5 \\ 0.019m^{-2.2} & \text{if } 0.5 \leq m < 1.0 \\ 0.019m^{-2.7} & \text{if } 1.0 \leq m < \infty \end{cases}$$

where m is stellar mass in units of M_{\odot} , and ξ is proportional to the probability that a star is formed with mass m ; $\xi(m)dm \propto P(m)dm$.

We extend this mass function down to $0.07 M_{\odot}$ and cap it at $117.5 M_{\odot}$, the limits over which our luminosity relation will extend in §4.1.4.

This relation, as given, has no dependence on metallicity, location in the galaxy, or population membership.

Main-sequence luminosity

In keeping with our use of Kroupa’s initial mass function, we use the mass-luminosity relation he employed to construct the IMF, given in the appendices of [Kroupa et al., 1993]. However, Kroupa’s listed luminosities trail off quite rapidly above $1.6 M_{\odot}$, fitting with his emphasis on low-mass stars. We need a relation that treats the entire main sequence, and so we must supplement Kroupa’s relation with another.

We find the necessary complement in [Zombeck, 2007]. Zombeck’s mass-luminosity relation is, in fact, just a broken power law, but it will do admirably for our purposes. Where Kroupa’s more detailed relation is available, we use Kroupa’s; where Kroupa’s relation grows sparse, we transition to Zombeck’s. The result can be seen in fig. 4.5. The two relations line up well at the transition, and no further adjustment is necessary.

This mass-luminosity relation applies only to the V band magnitude. In order to examine the effect of observing bandpass on event detection, we must also know the

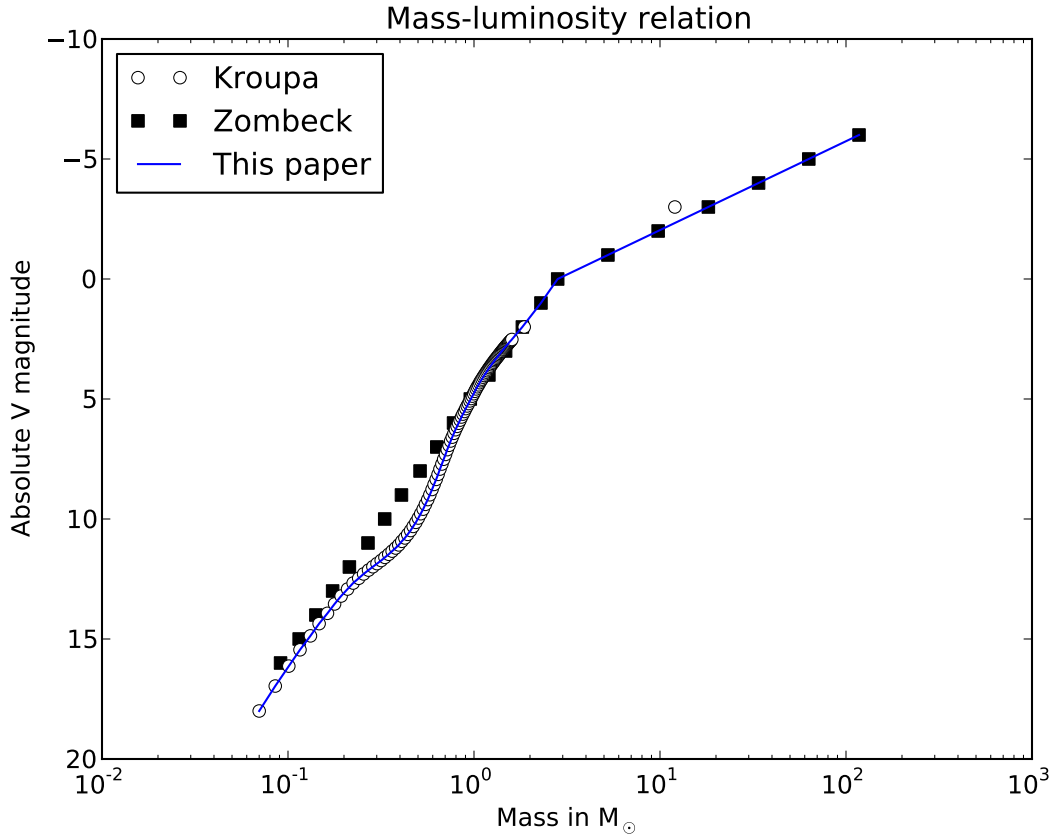


Figure 4.5: Zombeck gives a broken power law mass-luminosity relation, marked with black squares. Kroupa’s mass-luminosity relation provides superior detail at low masses, but becomes sparse above a few M_{\odot} . We use Kroupa’s relation where it is well-sampled, and Zombeck’s relation at high masses; our relation is marked with a solid blue line.

magnitude in other bands. We obtain colors for stars between -6 and 12 visual magnitudes from [Zombeck, 2007], and extrapolate to obtain colors for stars as dim as 18 V magnitudes. We use these colors to calculate U , B , R , and I band magnitudes. Our color-magnitude relation is illustrated in fig. 4.6.

Population age

Our initial treatment of the galactic model showed us that applying dust extinction to a zero-age population could not reproduce known star counts with sufficient accuracy.




Figure 4.6: Zombeck provides colors for stars with V magnitudes between -6 and 12 . We interpolate linearly in magnitude-color space to form a continuous relation, and extrapolate to find colors for stars between 12 and 18 magnitudes.

We needed an additional qualifier on the stellar population; we chose to apply ages to each star.

For this, we referred to the highly respected Besançon model, described in [Robin et al., 2003]. Their primary claim to fame is an excellent reproduction of known galactic properties using a synthetic model, i.e. one in which stars are generated and allowed to evolve.

The Besançon model treats the galactic halo as consisting entirely of stars formed 14 billion years ago, the thick disk as populated by 11 billion-year-old stars, and the thin

disk as having been steadily forming stars at a uniform rate for the past 10 gigayears. We adopt this same scheme for our own stellar population ages.

Lifetime

We use the relationship

$$t_{nuc} \approx 10^{10} \left(\frac{M}{M_{\odot}} \right) \left(\frac{L_{\odot}}{L} \right) \quad (4.23)$$

drawn from Hansen, Kawaler, and Trimble [Hansen et al., 2004] to estimate stellar main-sequence lifetimes. This relation should be very generally applicable.

We estimate bolometric luminosity as a function of mass using our mass-V magnitude relation, plus the definition of absolute magnitudes

$$M_{bol} - M_{bol,\odot} = -2.5 \log_{10}(L/L_{\odot}) \quad (4.24)$$

We plot the result in fig. 4.7.

This approximation does clearly overestimate the lifetimes of high-mass stars. Stars with mass of order one hundred solar masses have lifetimes of order a million years; this approximation estimates their lifetime to be of order one hundred million years. This approximation has given us the marked improvements in matching stellar populations we needed, and so stands as it is for the time being. A more accurate lifetime estimate could be obtained by using bolometric magnitudes instead of V magnitudes.

Stellar evolution

A full description of stellar evolution is far beyond the scope of this thesis, but some consideration of stellar aging is necessary for our purposes. We therefore treat




Figure 4.7: Estimated main-sequence lifetime as a function of mass.

all main-sequence stars as being freshly formed, on the zero-age main sequence. Post-main-sequence stars are reduced to zero luminosity and so cannot act as sources; their mass remains the same, allowing them to still act as lenses. This approximation appears to be sufficient for our purposes.

Present day population

As stated in §4.1.4, we model the halo and thick disk each as population of uniform age. Combined with our approach to stellar evolution, this corresponds to a sharp turn-off at a mass slightly below one solar mass; all stars smaller than the given age-dependent

limit are still on the main sequence, and all stars larger than the limit have burnt out.

The thin disk, by contrast, is modeled as having undergone continuous star formation for the past ten gigayears. A star of arbitrarily high mass has a chance to be still on the main sequence, if it was formed recently enough. We must use a more sophisticated description for the thin-disk population.

Out of ten stars with lifetimes one gigayear, formed at random times within the last ten gigayears, on average one will still be burning; the “main sequence fraction” is one-tenth. Using the calculated lifetimes from §4.1.4, we determine what fraction of stars of a given mass were born recently enough to still be on the main sequence. The result is shown in fig. 4.8. Stars with masses below about one solar mass are guaranteed to still be burning, and have a main sequence fraction of one.

Multiplying this main sequence fraction by the initial mass function produces a present day mass function (PDMF) of main sequence stars, shown in fig. 4.9. This PDMF is the distribution from which we draw source stars, which must still be shining. Lens stars, by contrast, are drawn from the IMF; they may or may not be on the main sequence.

4.1.5 Extinction model

The high-resolution dust maps produced by [Schlegel et al., 1998] form the core of our extinction model. They map the infrared emission throughout the sky, correcting for instrumental artifacts, point sources, and other systematic effects to produce a map solely of the dust emission. They then relate dust quantity to extinction at large distances in the V band.

The V band extinction, in turn, is related to extinction in other bandpasses. Schlegel et al. give both a table of many common bandpasses and an analytic approximation of




Figure 4.8: Fraction of stars on the main sequence as a function of mass in a population that has experienced ten gigayears of constant star formation.

extinction as a function of wavelength. We show the latter relation in fig. 4.10.

This falls slightly short of what we need, however; these maps describe the variance of extinction in only two dimensions, galactic latitude and longitude. They provide the extinction due to dust upon a star beyond all of that dust, potentially implying a stellar distance of kiloparsecs or more. We must use some scheme by which to apply a distance-dependent extinction effect, so that nearby stars are not unduly dimmed.

The traditional solution is to assume that dust is organized with falling-exponential density about the galactic plane. Our source for stellar distribution in space, Jurić et al. 2008, concluded that most of the dust was within 70 parsecs of the galactic plane based

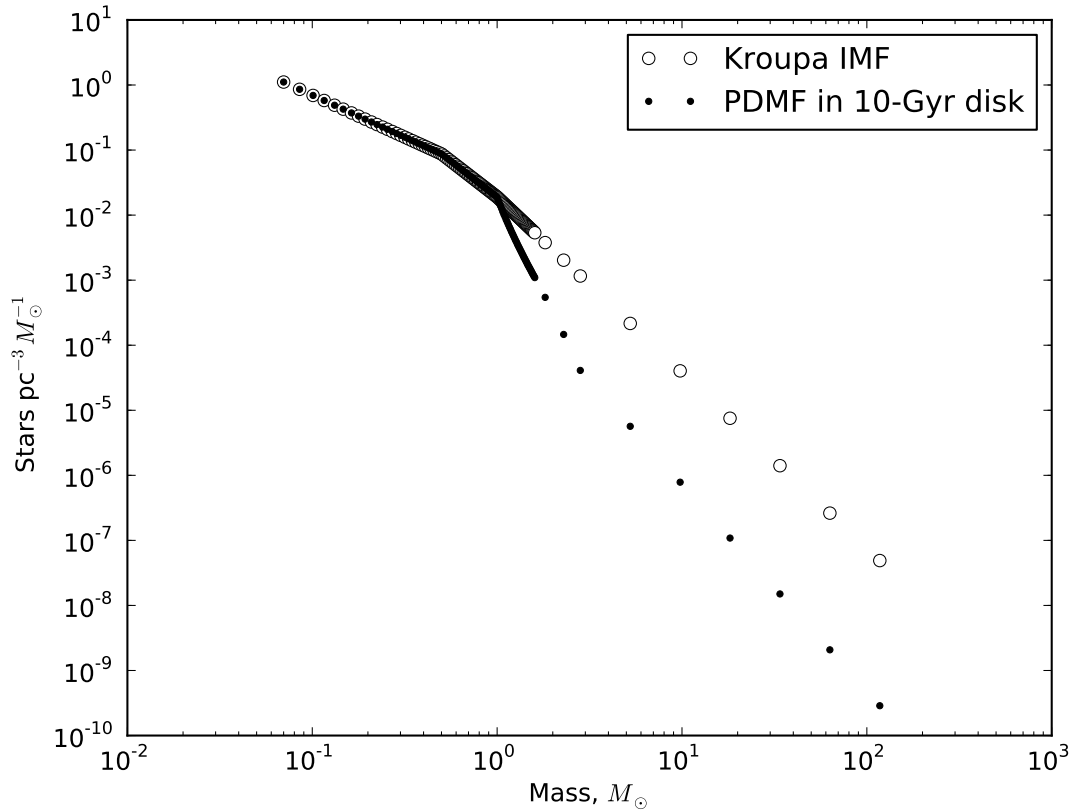


Figure 4.9: Total stars (white) and main-sequence stars (black) in a population that has experienced ten gigayears of constant star formation.

on their analysis of stellar colors. Other papers focused on the dust itself, however, have published a variety of contradicting opinions about the dust scale height, most of which seem to indicate larger scale heights for the dust. [Marshall et al., 2006], for instance, estimates the dust scale height to be 134_{-11}^{+44} pc. We use a scale height of 100 pc in this paper, and in an attempt to alleviate the asymmetry of dust about the galactic center implied by this relation, we introduce an exponential scale length in R of 2140 parsecs. We integrate over this simple relation to find what fraction of the total dust is between us and a star along a given line of sight.

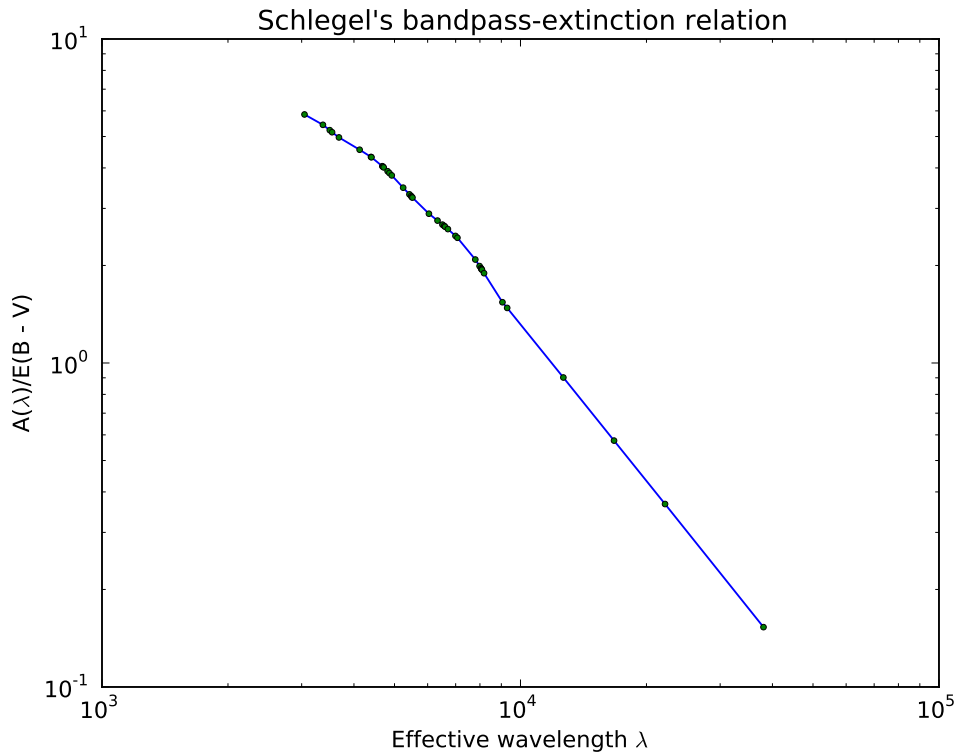


Figure 4.10: Dependence of extinction on wavelength from Schlegel 1998. We interpolate in log-log space to form a continuous relation.

4.2 Comparison to earlier work

A great many papers have been written regarding the lensing rate towards the galactic bulge; what makes our project unusual is the focus on fields away from the bulge. This subject has been very rarely treated. We will cover the few papers we are aware of on field microlensing rates and properties in this section.

The inspiration for our work, and our guide in the basic process of assembling the necessary model components, is [Han, 2008]. We have tried to improve on his work in almost every respect.

Han modeled the stellar density distribution as a double exponential disk of the form

$$n(R, z) \propto \exp\left(-\frac{R}{L} - \frac{z}{h}\right) \quad (4.25)$$

using a constant value for L , but varying h for different stellar populations separated by absolute magnitude, with values ranging from $h = 90$ pc for V magnitudes of less than three to a scale height of 600 pc for stars with V magnitudes of seven or greater. This was meant to approximate the general pattern that bright young stars are generally near to the plane, while older and dimmer ones may be further. Han’s stellar density model does not include a halo component, and includes no solar offset from the galactic plane. After building this model, he applies an unspecified normalization to the stellar number density in order to match observed star counts.

Han entirely neglects non-main-sequence stars, using the same mass and spatial distribution for source and lens stars. Stellar masses are drawn from the [Kroupa et al., 1993] initial mass function.

Han models the velocity of stars as a spatially independent three-dimensional gaussian dispersion with widths $(\sigma_R, \sigma_\phi, \sigma_z) = (38, 25, 20)$ km s⁻¹ about a bulk velocity of $(\overline{v_R}, \overline{v_\phi}, \overline{v_z}) = (0, 220, 0)$ km s⁻¹. Where we have separately recorded latitudinal and longitudinal angular velocity components for use in parallax measurements, Han recorded only total angular velocity.

Han describes dust distribution using a falling exponential in z with a dust scale height of 150 pc. He applies no R -dependence to the location of dust, and does not describe the relation between dust density and extinction. Han does not record color dependence, but considers only V band magnitudes of events.

Han’s primary concession to event observability is in the matter of source star magnitude; he does not discuss the effects of event time scale upon observability, for instance, nor does he make predictions for any telescope or survey’s observations.

Han gives event property distributions of timescale, lens mass, and lens and source distances. However, these properties are presented neither as a normalized distribution nor as an absolute distribution with definite quantity.

Han neglects to include a halo component in the stellar populations, and neglects our offset from the galactic plane. Han does not give normalizations for stellar density or dust density. Han does not use dust maps, but simply a mathematical relation, to describe extinction by dust, and provides no description of color. Han's model of velocity is spatially uniform, without dependence on R or z . Han does not allow for post-main-sequence stars to act as lenses. Han's paper, while good work, is somewhat lacking in detailed explanation of that work. In all of these respects we have improved upon Han's 2008 paper.

[Gould, 2013] is the only other paper we are aware of on the same topic. His choices were in many respects similar to Han's. Like Han, Gould used an exponential in R and z to describe the spatial distribution of stars; he also acknowledged a solar displacement from the galactic plane. Gould gives a luminosity function of stars which does not extend above two solar masses. Gould's velocity model was very similar to Han's, using slightly different numbers and adding a solar proper velocity. Where Han used an exponential in z only to describe the placement of dust, Gould used an exponential in R and z . Where Han did make predictions on distributions of some observable quantities, Gould reveals only two very simple metrics: total lensing rate, and lensing rate as a function of pointing.

Gould 2013's primary shortcoming, however, is a shortage of checks against reality. Where Han made an attempt to compare his model galaxy against known star counts before applying an arbitrary and unspecified normalization to make them consistent, Gould does not record even this much effort to verify his results. In this respect, too, we have made a great effort to build a superior code, demonstrably consistent with known

galactic properties.

Finally, though it is not a microlensing prediction tool, we wish to remark on the Besançon model, described in [Robin et al., 2003]. While their model is extraordinarily successful in replicating known galactic properties and takes an admirably self-consistent description of stellar populations, their published description of that model is not complete enough that we could reproduce their work in the context of a microlensing code.

Chapter 5

Code design

5.1 Summary of simulation design

To calculate microlensing rates and properties, I wrote a Monte Carlo integration code in Python. Its accuracy can be improved by running it longer, and it is easy to parallelize.

I made my code as modular as possible, constructing the code so that most elements were capable of working independently of one another, and of being tested independently. I built in safeguards to ensure that obvious mistakes were not made. I strove for efficiency, as much as I could, in the choice of Monte Carlo variable distributions. I took steps to make the code more human-readable, recording outputs in clear formats and providing commentary within the code itself.

Our model inputs focus on the galactic disk and halo, as they are the most important components in the sparse fields we are considering. Our model inputs do not currently include a bulge component; our code is not intended for and should not be used to determine lensing rates towards the bulge. The astrophysical inputs to our code are described at length in chapter 4.

An description of how events are generated in my code is found in §5.1.1. §5.1.2 briefly explains the dimensions of integration we must handle, and §5.1.3 describes how we transform this problem into a Monte Carlo integral. §5.1.4 explains observational cuts. Our method of parallelizing code for large runs is described in §5.1.5. §5.1.6 describes an important component of our code, the units-handling library I created.

We hope to make the full code and instructions on its use available to the scientific community.

5.1.1 Outline of information flow

My code can be conceptually divided into three major sections. The first section is the galactic model; it generates the properties, extrinsic and intrinsic, of every star in the simulation, one lens and one source for each event. The second section is the observational constraints; it determines whether a given event can be observed by a given observing program, configured by the user. The third section is output; it sorts and sums the events into meaningful quantities and useful formats. A schematic visual representation of information flow in my code can be found in fig. 5.1.

Description of event generation

Source star positions are chosen within the allowed volume, V_S . The number densities of thin disk, thick disk, and halo components are calculated at that location. Component membership is chosen for the source star. A location for the lens star is chosen; its local density and component membership are also determined.

The source star is assigned a mass from its component's present day mass function of main-sequence stars; the lens mass is chosen from the initial mass function. Ages are assigned; halo stars and thick disk stars all have uniform ages, but thin disk lens stars




Figure 5.1: A schematic representation of information flow in my code.

may have formed any time within the past ten gigayears, and thin disk sources must have formed at some time less than their main sequence lifetime in the past.

Velocities are chosen for both source and lens stars according to their component membership and position in the galaxy. From this velocity and their positions, their proper motion and relative angular velocity are calculated.

The source stars' absolute magnitudes and colors are calculated based on their mass.

The extinction by dust applying to a source star is calculated based on its position and color, the distance modulus is calculated based on its position, and apparent magnitude is calculated.

5.1.2 Dimensions of integration

As we explained at length in chapter 3, making a multidimensional integration efficient is paramount if your computing resources have some practical limit. We claimed in that chapter that our integration problem has of order ten dimensions; we will justify that claim now.

We derived an expression for a local lensing rate in chapter 2:

$$\Gamma(l, b) = \int 2 n_S n_L \theta_E |\vec{\omega}| D_S^2 D_L^2 dD_S dD_L \quad (5.1)$$

This expression explicitly integrates over D_L and D_S , and a rate for a solid-angle patch of sky would require integration over l and b (four spatial dimensions). However, there are additional variables hiding in this problem, and the boundaries of integration have not yet been described.

The relative angular velocity, $\vec{\omega}$, depends on the source star's velocity (three dimensions) and the lens star's velocity (three dimensions), as well as D_L and D_S . The angular cross section $2\theta_E$ depends on a less-than-straightforward function of D_L and D_S , plus

the mass of the lens (one dimension).

Other dependencies still will creep in when we consider the boundaries of integration, but the dimensions we have named here alone add up to a count of eleven dimensions. Arguably two dimensions of velocity could be neglected because motion towards and away from the viewer is far too small to have an effect on the lensing event, but that only brings the count down to nine dimensions.

5.1.3 Calculating lensing rate

Bearing the full eleven dimensions in mind, we transform the analytic expression of Γ to a Monte Carlo integral.

Explicitly integrating over angles l and b in eqn. 5.1, we have

$$\Gamma = \int 2 n_S n_L \theta_E |\vec{\omega}| D_L^2 D_S^2 dD_S \cos b db dl dD_L \quad (5.2)$$

Note that $D_S^2 dD_S \cos b db dl \equiv dV_S$ is the volume element of physical space in which a source star is found. We begin our transformation to a proper MC integral by indicating some volume V_S in which a source may be found, and choosing the source position \vec{x}_S with constant probability within that volume.

$$\Gamma = \frac{V_S}{N} \sum_i^N 2 n_S \int n_L \theta_E |\vec{\omega}| D_L^2 dD_L \quad (5.3)$$

where

$$P(\vec{x}_S) dV_S = \frac{1}{V_S} dV_S \quad (5.4)$$

The upper limit for D_L is always D_S . As this is not a constant, it cannot be moved

outside the integral sum, but rather remains inside.

$$\Gamma = \frac{V_S}{N} \sum_i^N D_S \cdot 2 n_{S n_L} D_L^2 \int \theta_E |\vec{\omega}| \quad (5.5)$$

where

$$P(D_L)dD_L = \frac{1}{D_S} dD_L \quad (5.6)$$

and eqn. 5.4 apply.

From chapter 4, we use the velocity distribution - already in the form of a normalized probability distribution - as the probability distribution from which we draw velocities.

$$\Gamma = \frac{V_S}{N} \sum_i^N 2 n_{S n_L} D_S D_L^2 |\vec{\omega}| \int \theta_E \quad (5.7)$$

where $P(\vec{v}_S)$ and $P(\vec{v}_L)$ are each given by the velocity distributions described in chapter 4, and eqns. 5.4 and 5.9 apply.

Finally, we draw our lens masses from a normalized initial mass function, completing the conversion to Monte Carlo.

$$\Gamma = \frac{V_S}{N} \sum_i^N D_S \cdot 2 n_{S n_L} D_L^2 |\vec{\omega}| \theta_E \quad (5.8)$$

where

$$P(M)dM = \frac{\text{IMF}(M)dM}{\int \text{IMF}(M)dM} \quad (5.9)$$

and all previous constraints apply.

5.1.4 Limits of integration and observational constraints

A small subset of the boundaries of integration we use are what might be termed artificial boundaries, boundaries that do not directly relate to the observability of an

event. Foremost among these is the volume from which we draw source locations, V_S . We wished to allow source stars from anywhere in the galaxy, in a meaningful sense; however, the galaxy does not have neatly demarcated boundaries. We settled for choosing a large spheroid extending to 20 kpc in R , and to $q_H \cdot 20$ kpc in z , so that the volume of integration mirrors the symmetry of the halo component described in ch. 4.

Our other boundaries of integration are natural, directly related to event observability. We impose boundaries based on event position in the sky, which determines whether a given telescope can be pointed towards it. We cut out source stars that are too dim, such that a given telescope cannot detect them to reasonable precision. We eliminate events with short timescales; a survey that scans once a day may entirely miss a 24-hour-long event. We eliminate events with long timescales; a pointing that can only be observed six months out of the year will not allow complete coverage of a year-long event.

All of these natural limits are adjustable by the user. Cutoffs in different passbands, field of view, and minimum and maximum timescale can all be set by the user. We designed the code this way for use as an observation-planning tool.

5.1.5 Parallelization

Our code is parallelized in the simplest way possible; it is designed to run multiple times using the same input parameters, and then combine the results.

In each run of the code, event sources and lenses are generated by the process described in §5.1.1. They are then pared down to observable events only, using the cuts described in §5.1.4. The results are then coadded to get the complete lensing rate, and binned in various ways to produce plottable distributions, many of which are shown in chapters 6 and 7. The ultimate data product stored by each run of the code is a python dictionary of interesting quantities and arrays of plot data.

After enough runs have been executed and enough data points have been generated, the code can be ordered to combine the full set of outputs generated from a given set of inputs.

5.1.6 Units-handling module

Mistakes in unit conversion are a scourge upon engineers, scientists, and programmers everywhere. A high-profile demonstration of the severity of the problem is provided by the Mars Climate Orbiter, which on September 23, 1999, went into an unplanned orbital insertion and subsequent crash landing. The problem was ultimately traced back to a software glitch in which thrust was expressed in pounds-force instead of Newtons.

I do not wish to highlight the costliness of this mistake, but rather the rarified environment in which it occurred. The National Aeronautics and Space Administration (NASA) has planned and executed countless space missions, the grand majority of which concluded successfully despite the extraordinary precision required, despite the continual exploration of new technological territory, despite the inevitability of human error, due to their extremely thorough and meticulous organization.

At NASA, everything that can be tested, is tested. Hypothetical failure scenarios are constructed, and procedures are developed to prevent those failures, and to recover if they occur. Processes are in place for filing reports of any anomalous spacecraft behavior. Despite these many precautions, an error in unit conversion slipped through the cracks, and evaded all the tests designed to find such errors. The failure investigation board ultimately identified eight points of technical and managerial failure that allowed the initial error to reach the catastrophic conclusion it did [NASA-JPL, 1999].

There came a point in debugging my code when I realized that the precautions I had taken up to that point were not enough, and a unit conversion error had slipped through,

due to the fact that my variables did not have explicit units. Rather than spend a truly excessive amount of time and energy combing over my code for similar mistakes, I built a tool to catch and forestall all such errors.

To avoid errors due to unit ambiguities and to simplify unit conversion tasks, I wrote a handful of modules in Python for the purpose of tracking and calculating units of scalars and arrays. These modules are not only friendly to but require NumPy, a numerical calculation library in common use amongst the Python programmers of the scientific world.

The resulting library is fairly easy to learn and use, reasonably efficient, and dependent only on NumPy and the Python standard library. While there are certainly ways it can be improved, I hope that this library, by itself, might be of use to other scientists.

It is worth noting that a similar project, the AstroPy.units module, has attempted to address the same concerns for the python-using astronomer community as a whole, with functions including notation and conversion of units [AstroPy Developers, 2015]. AstroPy.units has a similar but not identical feature set to my own units module; my module was built specifically for my needs in this code. However, my own solution to the problem was written before AstroPy.units was publicly available, has all the features needed for my code, and includes a complete library of the units I need for my own calculations.

Core features

These are the basic tasks my units module was designed to perform.

1. If you add or subtract quantities of the same type but with different units, it converts the arguments to a shared unit.

E.g.: 1 meter + 50 centimeters = 1.5 meters.

2. If you multiply, divide, or raise to a power a quantity with units, it works out what the units of the resulting quantity will be.

E.g.: 1 meter · 1 Newton = 1 Joule.

3. If you need a quantity converted from one unit to another, it can do so.

E.g.: 1 pc = 2.06×10^5 AU.

4. If you feed it a quantity expressed in unnecessarily convoluted units, it can simplify.

E.g.: 1 ft² · 1 in = 2.359×10^{-3} cubic meters.

5. It adapts many standard NumPy functions so that it is not possible to overlook their requirements on the unit of the argument.

E.g.: $\sin(45 \text{ degrees}) = 1/\sqrt{2}$, but $\sin(45 \text{ grams})$ throws an error.

Method

I will explain how the module works using the examples listed in §5.1.6, explaining the steps the code takes to handle inputs, and how it reaches the results.

Example 1: Addition with conversion

We begin by opening an interactive Python interpreter. In line [1], we import the units module, calling it `u` for short.

```
In [1]: import MCL_units as u
```

We create a variable `a` that represents a one meter length in line [2]. In the code, this object is a member of the `hasUnit` object class.

```
In [2]: a = u.hasUnit(1., 'm')
```

The code recognizes ‘m’ as an abbreviation for meter. It uses the full name of the unit internally. It determines that meters are a measure of distance. We confirm this by asking it to repeat `a` back to us, and querying `a`’s unit type.

```
In [3]: a
Out[3]: hasUnit(1.0, meter)
```

```
In [4]: a.unitType
Out[4]: 'distance'
```

If we multiply by a plain number, the result still has appropriate units.

```
In [5]: 2*a
Out[5]: hasUnit(2.0, meter)
```

We make a fifty centimeter length and add it with `a`.

```
In [6]: b = u.hasUnit(50., 'cm')
```

```
In [7]: a + b
Out[7]: hasUnit(1.5, meter)
```

In line [6], the code created a fifty-centimeter length. In line [7], it added one meter with fifty centimeters. Several things happened behind the scenes in the process.

The code first asked whether centimeters and meters can be added. Because they are both units of distance, it allowed the addition to go forward.

Had both units been the same, it would have added without any further manipulation, but in this case it needs to do a unit conversion. To keep that process as simple as possible, it converts each addend to MKS units, in which the basic units of distance, mass, and time are meters, kilograms, and seconds. `a` is still represented as one meter, but `b` is now represented as 0.5 meters.

Finally, the code adds, creating a new, 1.5-meter object.

Example 2: Multiplication and powers

To understand how multiplication is done in the code, you must first understand how units are classified and stored. We create a unit mass, distance, and time.

```
In [8]: x = u.hasUnit(1., 'kg')
```

```
In [9]: y = u.hasUnit(1., 'm')
```

```
In [10]: z = u.hasUnit(1., 's')
```

The human-readable versions of the unit types listed are as we would expect:

```
In [11]: x.unitType
Out[11]: 'mass'
```

```
In [12]: y.unitType
Out[12]: 'distance'
```

```
In [13]: z.unitType
Out[13]: 'time'
```

These are mostly for the sake of the user, however. The code's machinery mostly deals in a different expression, `unitUnit`. The `unitUnit` variable is a series of four numbers. The first represents powers of mass; the second represents powers of distance; the third is powers of time.

```
In [14]: x.unitUnit
Out[14]: array([ 1.,  0.,  0.,  0.])
```

```
In [15]: y.unitUnit
Out[15]: array([ 0.,  1.,  0.,  0.])
```

```
In [16]: z.unitUnit
Out[16]: array([ 0.,  0.,  1.,  0.])
```

The fourth digit, powers of temperature, was added as a quick way to use kelvins without accounting for degrees of freedom and other fiddly details. This feature should be considered incomplete.

When two units-bearing quantities are multiplied, their `unitUnits` add.

```
In [17]: r = x*y
```

```
In [18]: r.unitUnit
Out[18]: array([ 1.,  1.,  0.,  0.])
```


The code creates a human-readable version for the user based on the value of `unitUnit`.

```
In [19]: r.unitType
Out[19]: 'mass distance'
```

If the unit type has a common name, it will give that instead of a stilted description.

```
In [20]: v = y/z

In [21]: v.unitUnit
Out[21]: array([ 0.,  1., -1.,  0.])

In [22]: v.unitType
Out[22]: 'velocity'
```

If a quantity is raised to a power, the `unitUnit` is multiplied by that power.

```
In [23]: q = y**2

In [24]: q.unitUnit
Out[24]: array([ 0.,  2.,  0.,  0.])

In [25]: q.unitType
Out[25]: 'area'
```

The code can even handle non-integer powers, which are occasionally necessary, though rarely physically meaningful alone.

```
In [26]: p = z**0.5

In [27]: p.unitUnit
Out[27]: array([ 0. ,  0. ,  0.5,  0. ])

In [28]: p.unitType
Out[28]: 'time**0.5'
```

These operations form the basics of unit manipulation. We can now return to the problem of multiplying two quantities.

We once again create some quantities with units, and query them.

```
In [29]: a = u.hasUnit(1., 'm')
```

```
In [30]: b = u.hasUnit(1., 'n')
```

```
In [31]: a.unitUnit
```

```
Out[31]: array([ 0.,  1.,  0.,  0.])
```

```
In [32]: b.unitUnit
```

```
Out[32]: array([ 1.,  1., -2.,  0.])
```

The product of a force with `unitUnit = [1., 1., -2., 0.]` and a distance with `unitUnit = [0., 1., 0., 0.]` will have `unitUnit = [1., 2., -2., 0.]` and will be an energy.

```
In [33]: c = a*b
```

```
In [34]: c.unitUnit
```

```
Out[34]: array([ 1.,  2., -2.,  0.])
```

```
In [35]: c.unitType
```

```
Out[35]: 'energy'
```

To reach the final expression for `c`, the code converted both `a` and `b` to MKS units (a trivial operation, in this case), and multiplied those MKS units, knowing that in metric systems, one unit of force times one unit of distance equals one unit of mass.

The program's final description of `c` expresses it as

```
In [36]: c
```

```
Out[36]: hasUnit(1.0, joule)
```

Example 3: Unit conversion

Suppose we want to convert from one non-metric unit, like parsecs (abbreviated 'pc'), to another, such as astronomical units (abbreviated 'AU'). We create a one-parsec distance, and ask the code to convert to AU.

```
In [37]: a = u.hasUnit(1., 'pc')
```

```
In [38]: a.convert('au')
```

The code does not store any information directly relating the size of a parsec to the size of an AU. Instead, it stores information relating each unit to its metric equivalent. The code can convert any distance quantity to meters, and can convert meters to any other distance quantity.

To convert parsecs to AU, the code first converts parsecs to meters, then converts meters to AU.

```
In [39]: a
Out[39]: hasUnit(206264.805734, astronomicalunit)
```

Any unit conversion in my code uses the same scheme, although the details can become more complicated. For instance, converting kilometers per hour to feet per minute requires a total of four conversion operations.

```
In [40]: b = u.hasUnit(1., 'km hour**-1')

In [41]: b.convert('foot min**-1')

In [42]: b
Out[42]: hasUnit(54.6806649169, foot minute**-1)
```

In this example, kilometers is converted to meters, and then to feet, while per-hour is converted to per-second, and then to per-minute.

Example 4: Unit simplification

When given a problem in absurd units, my code will generally simplify by converting everything to its favorite metric units.

```
In [43]: a = u.hasUnit(1., 'foot**2')

In [44]: b = u.hasUnit(1., 'inch')

In [45]: a*b
Out[45]: hasUnit(0.002359737216, meter**3)
```

If a more complicated unit expression is needed, it can always be restored by use of the `convert` command. In the meantime, the actual value of the quantity is preserved in a clear and unambiguous format.

Example 5: Implied unit checking

Many mathematical functions have an implied requirement that their arguments be unitless. The expression `arctan(2 grams)` does not have any clear meaning, for instance. Because of this, we have built in checks on our version of many of the standard NumPy functions and operations.

Angles are a special case. `sin(45 degrees)` *does* have meaning. Our treatment of angles in our code ultimately centers around the fact that angles are technically unitless. Two meters divided by one meter might be two radians, or it might be simply two, depending on the physical relationship between those two lengths.

As the saying goes, “Never send a computer to do a human’s job.” The rules governing whether something is an angle or just a number depend on information that the computer most likely will not have. That decision is left in the hands of the user, and tools are in place for the user to specify which case they are invoking.

In the absence of an obvious metric equivalent, we use radians as the basic measure of angle in our code. This preserves compatibility with the NumPy and SciPy packages, which also assume by default that the argument of a trigonometric function is given in radians. Conversions to other measures of angle, such as degrees and arcseconds, work as usual.

Other features

Advanced users can choose which system they prefer to have quantities expressed in. Whenever the code has to simplify an expression, it converts to the user-defined preferred system. For instance if the user chooses MKS as the preferred system of units, one dyne

is re-expressed as 10^{-5} newtons.

Several common systems of units (MKS, CGS) are described in the code, as well as a few less-common ones and a number of physical constants. Many astronomical units such as solar masses, parsecs, and arcseconds are included because of their usefulness for my calculations. Although the set of units and systems described in my library at present is not comprehensive, it is extensible, and advanced users can easily define additional units, unit systems, and physical constants with relative ease.

It is useful to note that my `hasUnit` objects can contain NumPy arrays as the data at their heart, and that while units-handling operations are handled by my own code wrapped around the data, the multiplication and other arithmetic operations of those data arrays still falls back on NumPy's efficiently coded methods. There should be no significant slowdown of code re-written for use with my units package, as the core computational operations are still done by the reliable and efficient NumPy.

Chapter 6

Code verification

The final code written for this project was thousands of lines long. Although I made every effort to modularize the code, fool-proof the calculations, and generally construct as clean and simple a design as possible, software development does not stop when the code has been written. Code *must* be tested.

Relying solely on my own good sense to debug such a large, interconnected program would be foolhardy in the extreme. My strategy has been to take two major avenues of attack. I have included coded tests to check elements of code functionality every time the program runs, described in §6.1. I have also performed extensive tests to verify that the known physical properties of the galaxy are being accurately reproduced; these are described in §6.2.

6.1 Coded persistent tests

These tests are computationally cheap, and can reasonably be run every time the code is executed.

Their purpose is not only in the value of the initial check they provided when first

performed; these tests stand guard against meaningful changes, either in my own code or in the Python installation being used, as other aspects of development continue.

6.1.1 Module tests

The development of the units-handling code library greatly increased how robust the code overall was to errors, but more testing and verification was needed. To help find and pinpoint errors, to prove that the code was working as intended, and to ensure that once-corrected errors did not recur, I built a set of functionality tests permanently into the code.

The first of these tests were simply problems with known answers, verifying that the output given by the code was correct. Many of these early tests were performed with several input values, to eliminate the possibility that they functioned correctly for certain special cases but not in general. In modules with random elements, expectation values and expected variances were computed, and the answers were compared to see if they fell reasonably close. In others tests, I set the code problems whose randomized answers had to fall within certain limits, and ran the problem many times to be sure that it always obeyed those limits.

Later tests compared the properties of the Monte Carlo generated stars to the constraints posed by the model used and the known properties of the real galaxy. However, for the bulk of our verification of the galactic-model components, we rely on comparing them to known galactic properties, demonstrated in §6.2.

The modules I consider foundational - modules such as the units-handling library, the binning of data points, and the calculation of errors - I tested with particular thoroughness, creating of order a dozen tests each that run every time the code runs.

6.2 Comparison to models and observations

The outputs of my code all take the form of histograms due to the discrete nature of the data involved, inherent to Monte Carlo methods. Interpreting the data products as continuous inevitably lowers the precision.

Where there is more than one data point in a particular histogram bin, an error estimate is calculated based on the variance of the data in that bin. Note that this is a one-sigma error estimate, and it will not always be a particularly accurate estimate, but it is the best estimate available given the data, and is generally decent. See the chapter on Monte Carlo integration for further details on how this estimate is obtained. This error estimate is visually represented on the graph by a vertical line segment through the histogram bin.

Bins that contain only one data point cannot have an error estimate derived by the same means. These bins are capped by a circular marker to designate their lack of a proper error estimate. These bins should be viewed as highly suspect; their values should not be relied upon. They represent under-sampled regions of parameter space.

Bins that have x-marks are bins in which no samples fell. This does not mean that the value of these bins is zero, nor does it necessarily mean that the value is low; it means that these bins were very under-sampled. In some cases this is perfectly acceptable, e.g. unpopulated and uninteresting areas of parameter space. It is not surprising, for instance, to have empty bins five standard deviations away from the average value of a normally distributed quantity.

Finally, note that a high value in a bin does *not* necessarily mean that it is well-sampled, nor that it necessarily contains more data points than other, lower-value bins. The contribution of a data point to a histogram bin depends on the value of the quantity portrayed (e.g., lensing rate) at that data point.

The following plots were made using a dataset of 5.1×10^7 generated lens-source pairs. Such high sampling is, in most cases, overkill where these metrics are concerned, and hence error bars are very small in most bins. In some cases error bars may be too small to see, but they have been plotted nonetheless.

This high sampling will become necessary in the next chapter, when we present our results; only about one out of a hundred of these hypothetical events is actually observable. Please note for graphs throughout this chapter that the error bars are of very reasonable magnitudes. Where there is a definite model to compare to, the one-sigma error bars usually encompass that model. Even in cases where they do not, it is common to see larger error bars on wilder data bins.

The nature of Monte Carlo methods means that wild values far from the expected value will occur, but with sufficient sampling, they usually show correspondingly large error estimates.

6.2.1 Distribution of stars in space

Figure 6.1 shows the density of stars in space as a function of height z out of the galactic plane. Each bin is averaged over the volume of integration in R and ϕ . Overplotted are the upper model limit, i.e. the density at zero R where disk and halo are thickest, and the lower model limit, the density at large R where stars are few and far between.

Figure 6.2 shows the density of stars in space as a function of galactocentric cylindrical distance R . Each bin is averaged over the volume of integration in z and ϕ . Overplotted are the upper model limit, i.e. the density at zero z where the disk dominates, and the lower model limit, the density at large z where scarce outer halo stars are the majority.

The upturn at large R in figure 6.2 is an effect of the spheroidal integration volume




Figure 6.1: The distribution of stars in z , the direction perpendicular to the galactic plane, averaged over the volume of integration. Model maximum and minimum shown.




Figure 6.2: The distribution of stars in R , galactocentric cylindrical radius, averaged over the volume of integration. Model maximum and minimum shown. The upward curve at large R is an edge effect illustrated in fig. 6.3.

chosen. Figure 6.3 illustrates.




Figure 6.3: Represented are the volume of integration (large ellipse), galactic disk (horizontal bar), halo (small ellipse), and three bins from plot 6.2 (labeled a, b, c). **a**: At small R , the density is dominated by the power-law halo. **b**: At moderate R , the density is dominated by the disk, which occupies only a small fraction of the total space. **c**: At large R , the total density is still dominated by the disk, but the disk occupies a much larger fraction of the volume of integration. As a result, the average density in c is higher than the average density in b.

Figure 6.4 shows the density of stars in space as a function of galactocentric spherical radius r . Each bin is averaged over the volume of integration in θ and ϕ . Overplotted are the upper model limit, i.e. the density at zero z where the disk dominates, and the lower model limit, the density at large z where scarce outer halo stars are the majority.

The upturn at high r , also noted in fig. 6.2, has the same cause, illustrated in fig. 6.3.

Figure 6.5 shows the density of stars in space as a function of galactocentric azimuthal angle ϕ , averaged over z and R . In other words, if you stood at the galactic center and rotated about the pole, this is what you would see; the graph should be essentially a constant, with variation about the mean galactic density. The heavy line on this plot is




Figure 6.4: The distribution of stars in r , galactocentric spherical radius, averaged over the volume of integration. Model maximum and minimum shown. The upward curve at large r is an edge effect illustrated in fig. 6.3.

the mean density of source stars, calculated using

$$\overline{n_S} = \frac{1}{N} \sum_i^N n_{S,i} \quad (6.1)$$

Figure 6.5 has higher variance than the others in this section because each bin encompasses values from the very densest parts of the galactic center out to the furthest tail of the disk and halo, meaning that the data averaged within that bin have orders of magnitude of intrinsic variation. No limiting upper bound can be plotted because our expression for the halo component diverges at the galactic center; the lower bound is of order $5e-5$, and would not be visible with linear scaling.

6.2.2 Velocity distribution

Figures 6.6, 6.7, and 6.8 display the distributions of stellar velocity in the radial (R), vertical (z), and rotational (ϕ) directions respectively.

These distributions are averaged over the entire volume of integration, and so are a mishmash of halo and disk star velocity distributions. This is not very apparent in plots 6.6 and 6.7, where all populations are gaussians symmetric about zero; however, in figure 6.8, the rotational velocity, the distinction becomes much more obvious.

Here, the majority of stars are offset some 200 km s^{-1} from zero due to the bulk rotation of the disk. A long tail towards zero is composed at least partially of halo stars, which we modeled as having no net rotation, but is also contributed to by disk stars far from the galactic plane. Our final relation very closely resembles the original data of [Bond et al., 2010].

On each velocity plot, we plot three traces representing three of the many velocity distributions that contribute to the whole. One represents the halo's velocity distribution; one the local velocity distribution of disk stars; the third the distribution of stars




Figure 6.5: The distribution of star location as a function of galactic azimuthal angle, averaged over the volume of integration. This histogram should be noise about a constant value, represented by the solid horizontal line.




Figure 6.6: Velocity distribution in the R direction, averaged over the integration volume. The three traces represent normalized probability distributions from three different populations; the whole should be close to some combination of these three traces. The three populations share a mean in this dimension, and so their influences are difficult to distinguish.




Figure 6.7: Velocity distribution in the z direction, averaged over the integration volume. The three traces represent normalized probability distributions from three different populations; the whole should be close to some combination of these three traces. The three populations share a mean in this dimension, and so their influences are difficult to distinguish.




Figure 6.8: Velocity distribution in the ϕ direction, averaged over the integration volume. The three traces represent normalized probability distributions from three different populations; the whole should be close to some combination of these three traces. These three populations have very distinct rotational behaviors and mean values, leading to a very asymmetrical distribution.

five kiloparsecs out of the plane from our location. Each of these distributions is independently normalized to $\overline{n_S}$ as given in eqn. 6.1. Although these three distributions are not the sole contributors to the total velocity distribution, it is easily visible that some combination of these three could describe the whole reasonably well.

6.2.3 Stellar populations

Figure 6.9 gives the mass distribution of lens stars. It follows the initial mass function described in ch. 4 so closely that it is difficult to distinguish the histogram from the comparison trace.

Figure 6.10, the mass distribution of source stars, falls off from the initial mass function above one solar mass, as expected. The total distribution closely resembles our PDMF given in ch. 4, but becomes somewhat undersampled at high masses, where living main-sequence stars are rare.

Figure 6.11 represents the last metric still reporting badly in our code. After all our refinements, the star counts predicted by our model still show a significant mismatch with the known galactic star counts, roughly an order of magnitude from the observed counts we used.

We are not yet certain where, exactly, the fault lies in this disagreement. We note that [Han, 2008] also had difficulty reproducing known star counts when using a similar model, resorting to applying an unspecified scale factor. We note that the source of our model is backed by very recent, extremely rigorous observations and modeling efforts applied to a population of nearly fifty million stars. We note that our code has demonstrated sensible results in a very great number of other metrics, failing only here. And finally, we note that the source of our comparison star counts is incredibly antiquated.

We drew these star counts from [Cox, 2000], better known as Allen’s Astrophysical




Figure 6.9: The distribution of lens masses, which may be of any age, compared to the initial mass function.




Figure 6.10: The distribution of source star masses, which we require still be on the main sequence. The initial mass function is plotted for comparison. The final distribution closely resembles the mass function of main sequence stars after applying the aging effects described in ch. 4.




Figure 6.11: This plot shows stellar counts ultimately derived from [Searles et al., 1925, Searles and Joyner, 1928, van Rhijn, 1929], plotted as symbols, compared to star counts generated by our own simulation, shown as solid lines. The lines are should match up with the symbols of the same color. From top to bottom, the trace/symbol combinations represent 20, 18, 16, 14, and 12 V magnitude populations, respectively. The dip at equatorial latitudes is caused by dust extinction not balanced out by a bulge component.

Quantities, a text respected and widely used amongst astronomers, recently updated to a fourth edition. However, investigating the sources used ultimately led us to a set of papers published in the 1920s: [Seares et al., 1925, Seares and Joyner, 1928, van Rhijn, 1929]. The original work by which these numbers were determined was done in the era of photographic plate astronomy.

Age alone does not, of course, disqualify any scientific work. Photographic plates would doubtless introduce minor errors, but we would ordinarily expect those to be small, less than an order of magnitude. Star counts determined a hundred years ago should be just as valid today. However, we find ourselves somewhat skeptical nonetheless. We intend to look further into this problem at a later date, looking for more up to date star counts, and making additional checks for errors in the code.

One possible explanation for the discrepancy is found in the fact that old photographic plates were more sensitive in the blue than the red. Old observations were also conducted almost entirely from the northern hemisphere, where the bulge is obscured by a high airmass that preferentially scatters blue light. The two effects could combine to cause a major underestimate of star count at equatorial latitudes.

We intend, ultimately, to find and use more modern sources of star counts from sources such as SDSS and Gaia, and bring our model in line with them.

Chapter 7

Results

7.1 Results

7.1.1 Whole-sky lensing rates and properties

The following plots are made from the same set of 5.1×10^7 datapoints as the graphs in ch. 6, but have been subjected to several observational cuts. Events with timescales shorter than one day or longer than one hundred days are rejected as unobservable. Events with source quiescent magnitudes dimmer than 23.9 in U band, 25.0 in V band, 24.7 in R band, or 24.0 in I band are discarded, approximating the requirement that an event source be visible in four of the filters used by LSST [LSST Science Collaboration et al., 2009]. Events from the whole sky are kept and displayed.

This set of restrictions leaves roughly one event in one hundred observable. As a result, these plots are not as well-sampled as the corresponding plots of ch. 6, but they are sufficiently sampled to provide an extremely illuminating look at lensing event properties.

Note that the comments in ch. 5 on my histogram plotting format apply, and are important to understanding the plots in this chapter. I would strongly recommend that

the reader familiarize themselves with that material before continuing.

Distribution of events in space

Figure 7.1 shows the distribution of observable events in b , summed over D and l .

The dip in event rate at the equator is not surprising in light of the dip in star counts at the equator; however, it is suspect as long as our star counts remain suspect. Further investigation on that front will hopefully reveal whether this effect is real and observable. Additionally, note that any adjustment to the dust model will likewise change the width and depth of the excluded region about the equator.

We *were* surprised to see the asymmetry of event rate about the equator. This is almost certainly an effect of our sun's offset from the galactic plane. As we are slightly north (positive galactic z direction) of the galactic plane, more stars are south of us than north.


We included our z offset from the plane more out of a sense of duty than of clever suspicion; we did not think an offset of 25 pc and a tilt of $\Delta b \approx 0.17$ degrees were likely to have a noticeable effect.

Figure 7.2 displays the distribution of events in galactic longitude l , summed over D and b .

This distribution also appears to have an asymmetry in the form of a dip around -20 to -40 degrees. We believe this to be caused by dust lanes in our dust maps from [Schlegel et al., 1998].

Figure 7.3 shows the distribution of event source distances, summed over l and b .

The peak of this distribution at about eight kpc is entirely unsurprising; it is consistent with stars near the dense galactic center being the most common source stars. The near-symmetry of the distribution is interesting; it implies that the decreasing availability of sufficiently bright sources counterbalances the increasing availability of lenses.



88

Figure 7.1: Distribution of events in galactic latitude b . Event rate in the disk drops almost to zero at the equator, due to the heavier dust extinction there and corresponding decrease in visible stars.




Figure 7.2: Distribution of events in galactic longitude l . Event rate spikes towards the galactic center due to the higher stellar population there.




Figure 7.3: Distribution of event source distances D_S . Note that these continue nearly to the maximum possible distance.

Figure 7.4 shows the distribution of event lens distances, summed over l and b .

As with event source stars, the lens distance distribution peaks around eight kiloparsecs, the distance to the galactic center. It does, however, appear to be convex instead of concave in the lead up to that peak, and falls off more dramatically afterwards.

Angular velocity distribution of events

Figure 7.5 shows the relative angular velocity of stars in the b direction.

This distribution appears to be symmetrical, as expected. The shape, however, is odd; the distribution is sharply peaked and distinctly non-gaussian.

Figure 7.6 shows the relative angular velocity of stars in the l direction.

As with the rotational velocity distribution shown in ch. 6, this distribution is asymmetrical and has a definite nonzero mean. This is unsurprising, although not particularly expected.

Figure 7.7 shows the total relative angular velocity of event stars. This holds no particular surprises; we expect it to look similar to a Maxwell distribution constructed for two dimensions instead of the usual three, and indeed it does. We learn from this plot that the typical lens-source pair separate at a rate of 0 to 20 milliarcseconds per year.

Stellar populations

Figure 7.8 shows the distribution of lens masses. We see that the distribution is dominated by low-mass stars. Although small stars have smaller cross-sections to lensing, the numeric advantage of low mass stars causes them to overwhelmingly outweigh high mass stars in the lens mass distribution.

Figure 7.9 shows the distribution of source masses. A given massive star, being brighter, is more likely to be visible to act as a source star, but bright stars are fewer in




Figure 7.4: Distribution of event lens distances D_L . Note that these average smaller than D_S .




Figure 7.5: Relative angular velocity of source and lens along the longitudinal direction, \hat{b} . These average about zero, as expected.




Figure 7.6: Relative angular velocity of source and lens along the latitudinal direction, \hat{l} . Like the rotational velocity distribution, this distribution is asymmetrical with nonzero mean.




Figure 7.7: Total angular velocity. Appearance as expected.




Figure 7.8: Distribution of event lens masses. Dominated by low-mass stars.




Figure 7.9: Distribution of event source masses. Peaks at about 0.9 solar masses.

number. These two effects create a peak at about 0.9 solar masses, with power law like tails to each side.

7.1.2 Other results of interest

The most exclusive observational cut we made was the exclusion of events with source stars too dim to be seen by our hypothetical telescope. This was expected; a great many stars exist in our galaxy that are too dim to see.

We did not expect the effect of our timescale cuts, however. Requiring that no events with timescales shorter than one day be detected excluded about one percent of events. This is an interesting result in itself because it speaks to the rarity of short-duration star-star lensing events, but more interesting is the effect of the long-event cutoff.

Excluding events with timescales longer than one hundred days eliminated approximately one half of the events generated. Although the exact cutoff for what constitutes a “too long” event is uncertain, 100 days is order of magnitude correct, erring high, for most surveys. Bulge-facing ground-based surveys, for instance, can only run effectively for about half the year; an event with timescale 100 days is not likely to fit entirely into that half-year window. Long events may be observed by combining multiyear data, and OGLE presently does so for many large- t_E events.

These long-timescale events represent an population of events likely to be neglected by standard survey strategies. They must be accounted for in any survey seeking to capture lensing events. We intend to investigate this further before publication.

7.2 Discussion

Our results so far indicate that LSST will indeed observe microlensing events at a rate of a few tens per year, but the yield can be optimized with a suitable observing

strategy. My software provides a tool to test these observing strategies before first light.

The results we showed here are for all-sky observation, and even very wide surveys like LSST will not cover the entire sky; they will be able to survey some 50% to 70% at any given time, weather permitting, and the portion of the sky they have access to will change throughout the year. Our final estimates for LSST lensing event rates will take this into account.

Additionally, we uncovered unexpected trends in event timescale, and an unexpectedly strong effect of our small displacement from the galactic plane.

7.2.1 Experiments

There are a few easily accessible and interesting experiments we can do with our code, and are likely to perform before publication.

Foremost among these is to create plots, like those in this chapter, describing lensing event time scale distributions. The unexpected result of many long-timescale events is very much deserving of additional investigation on the effects of observing strategies on the capture of these events.

We can use the tools we have already built to apply constraints comparable to the survey observing plans of the upcoming LSST, ATLAS, and Evryscope surveys, estimating the lensing rates and event property distributions they can expect to find.

A number of small but potentially interesting experiments can be performed by varying our sun's modeled displacement from the galactic plane and the shape of the dust distribution, and relating these to the effect on lensing event properties.

7.2.2 Outstanding problems

As discussed in chapter 6, our generated star counts do not presently match the antiquated observed star counts we compared them against. We are actively investigating this, and have identified possible causes of the effect, as well as identified modern sources to use as a comparison.

7.2.3 Potential future directions

There is always more to investigate, and this section covers a few of the major avenues of investigation we have identified as interesting.

Galactic model improvements

We are very interested in expanding the range of lens types considered in our code. We would very much like to extend the mass function of lens objects down to brown dwarf and planetary masses, although we note the existing estimates of these populations are very much ongoing work. If sufficient constraints on their population should become available, we would like to incorporate them. Alternatively, we could make predictions of the relationship between observed rate and population size.

Similarly, we would like to properly handle mass loss at end of main sequence life. Presently our stars do not undergo a change in mass at any point in their evolution; this is not an accurate reflection of stellar evolution, and we would like to amend it.

We would also be pleased to add a bulge component to our model for comparison against known bulge lensing event properties.

Units code improvements

While my units-handling code is something I am very proud of, I can see a number of avenues by which it could be improved.

My code relies on NumPy, but as an unfortunate side effect of otherwise useful behavior, multiplying a NumPy array by a hasUnit array can cause a major error. While workarounds exist, I would not expect a casual user to employ them. Finding a solution to this, perhaps in the form of creating array subclasses that NumPy itself will recognize, would be a major step in improving the usefulness of my code to others.

Adding units and unit systems such as imperial units is easy to do, and an efficient use of my time when it comes to making the code useful for other users.

My code does not currently have any description of electricity and magnetism, as it was not necessary for my purposes. This seems like a major candidate for improvement. Likewise, putting in proper handling of temperature seems like it would broaden the applicability.

Currently, every unit in my code that can be converted is “zeroed” and “linear”. Zeroed means that a value of unit zero indicates an absolute zero (kelvins are zeroed, but degrees fahrenheit are not). Linear means that doubling a physical quantity doubles the unit measurement that describes it (lumens are linear, but absolute magnitudes are not). This allowed some important simplifications in the baseline code, but ultimately it might be useful to include proper descriptions of conversions to and from nonlinear, nonzeroed units.

I have some interest in adding provisions for the unit equivalencies used by some physicists. Relativists, for example, often equate mass to distance via the Schwarzschild radius, and equate distance to time by means of the speed of light. I do not yet consider my understanding of this concept deep enough to determine whether implementing such

a system is feasible, much less how to do it, but I do consider it a potential long-term goal for development.

I also wish to note that the basic method I used to apply units to arrays can be used to apply error estimates that transform properly under arithmetic operations. This, as with the units code, would be a broadly useful tool, but is less urgently needed, and likewise represents a possible long-term development goal.

I intend to publish all of my simulation code, but the units-handling library in particular takes high priority for independent publication, due to its extremely broad applicability.

7.3 Final words

My code is a well-verified software tool which can be adapted to simulate a wide range of potential survey strategies and parameters. As a result, it will be a useful tool for the community to optimize the design of the deep, wide-angle surveys coming on line in the next decades, such as ATLAS, Evryscope, and LSST.

Bibliography

- [Abe et al., 2013] Abe, F., Airey, C., Barnard, E., Baudry, J., Botzler, C., Douchin, D., Freeman, M., Larsen, P., Niemiec, A., Perrott, Y., Philpott, L., Rattenbury, N., and Yock, P. (2013). Extending the planetary mass function to Earth mass by microlensing at moderately high magnification. *Monthly Notices of the Royal Astronomical Society*, 431:2975–2985.
- [Alcock et al., 1995] Alcock, C., Allsman, R. A., Alves, D., Axelrod, T. S., Bennett, D. P., Cook, K. H., Freeman, K. C., Griest, K., Guern, J., Lehner, M. J., Marshall, S. L., Peterson, B. A., Pratt, M. R., Quinn, P. J., Rodgers, A. W., Stubbs, C. W., and Sutherland, W. (1995). First Observation of Parallax in a Gravitational Microlensing Event. *Astrophysical Journal Letters*, 454:L125.
- [Alcock et al., 2000] Alcock, C., Allsman, R. A., Alves, D. R., Axelrod, T. S., Becker, A. C., Bennett, D. P., Cook, K. H., Dalal, N., Drake, A. J., Freeman, K. C., Geha, M., Griest, K., Lehner, M. J., Marshall, S. L., Minniti, D., Nelson, C. A., Peterson, B. A., Popowski, P., Pratt, M. R., Quinn, P. J., Stubbs, C. W., Sutherland, W., Tomaney, A. B., Vandehei, T., and Welch, D. (2000). The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations. *The Astrophysical Journal*, 542:281–307.
- [AstroPy Developers, 2015] AstroPy Developers (2015). Units and quantities (astropy.units). Online documentation. <http://docs.astropy.org/en/stable/units/index.html>, accessed Mar. 9 2015.
- [Bakos et al., 2004] Bakos, G., Noyes, R. W., Kovács, G., Stanek, K. Z., Sasselov, D. D., and Domsa, I. (2004). Wide-Field Millimagnitude Photometry with the HAT: A Tool for Extrasolar Planet Detection. *The Publications of the Astronomical Society of the Pacific*, 116:266–277.
- [Bakos et al., 2013] Bakos, G. Á., Csubry, Z., Penev, K., Bayliss, D., Jordán, A., Afonso, C., Hartman, J. D., Henning, T., Kovács, G., Noyes, R. W., Béky, B., Suc, V., Csák, B., Rabus, M., Lázár, J., Papp, I., Sári, P., Conroy, P., Zhou, G., Sackett, P. D., Schmidt, B., Mancini, L., Sasselov, D. D., and Ueltzhoeffer, K. (2013). HATSouth: A Global Network of Fully Automated Identical Wide-Field Telescopes. *The Publications of the Astronomical Society of the Pacific*, 125:154–182.

- [Becklin and Zuckerman, 1988] Becklin, E. E. and Zuckerman, B. (1988). A low-temperature companion to a white dwarf star. *Nature*, 336:656–658.
- [Beichman et al., 2014] Beichman, C., Gelino, C. R., Kirkpatrick, J. D., Cushing, M. C., Dodson-Robinson, S., Marley, M. S., Morley, C. V., and Wright, E. L. (2014). WISE Y Dwarfs as Probes of the Brown Dwarf-Exoplanet Connection. *The Astrophysical Journal*, 783:68.
- [Bennett et al., 2010] Bennett, D. P., Rhie, S. H., Nikolaev, S., Gaudi, B. S., Udalski, A., Gould, A., Christie, G. W., Maoz, D., Dong, S., McCormick, J., Szymański, M. K., Tristram, P. J., Macintosh, B., Cook, K. H., Kubiak, M., Pietrzyński, G., Soszyński, I., Szewczyk, O., Ulaczyk, K., Wyrzykowski, Ł., OGLE Collaboration, DePoy, D. L., Han, C., Kaspi, S., Lee, C.-U., Mallia, F., Natusch, T., Park, B.-G., Pogge, R. W., Polishook, D., μ FUN Collaboration, Abe, F., Bond, I. A., Botzler, C. S., Fukui, A., Hearnshaw, J. B., Itow, Y., Kamiya, K., Korpela, A. V., Kilmartin, P. M., Lin, W., Ling, J., Masuda, K., Matsubara, Y., Motomura, M., Muraki, Y., Nakamura, S., Okumura, T., Ohnishi, K., Perrott, Y. C., Rattenbury, N. J., Sako, T., Saito, T., Sato, S., Skuljan, L., Sullivan, D. J., Sumi, T., Sweatman, W. L., Yock, P. C. M., MOA Collaboration, Albrow, M., Allan, A., Beaulieu, J.-P., Bramich, D. M., Burgdorf, M. J., Coutures, C., Dominik, M., Dieters, S., Fouqué, P., Greenhill, J., Horne, K., Snodgrass, C., Steele, I., Tsapras, Y., PLANET, F. t., RoboNet Collaborations, Chaboyer, B., Crocker, A., and Frank, S. (2010). Masses and Orbital Constraints for the OGLE-2006-BLG-109Lb,c Jupiter/Saturn Analog Planetary System. *The Astrophysical Journal*, 713:837–855.
- [Bodenner and Will, 2003] Bodenner, J. and Will, C. M. (2003). Deflection of light to second order: A tool for illustrating principles of general relativity. *American Journal of Physics*, 71(8):770–773.
- [Bond et al., 2010] Bond, N. A., Ivezić, Ž., Sesar, B., Jurić, M., Munn, J. A., Kowalski, A., Loebman, S., Roškar, R., Beers, T. C., Dalcanton, J., Rockosi, C. M., Yanny, B., Newberg, H. J., Allende Prieto, C., Wilhelm, R., Lee, Y. S., Sivarani, T., Majewski, S. R., Norris, J. E., Bailer-Jones, C. A. L., Re Fiorentin, P., Schlegel, D., Uomoto, A., Lupton, R. H., Knapp, G. R., Gunn, J. E., Covey, K. R., Allyn Smith, J., Miknaitis, G., Doi, M., Tanaka, M., Fukugita, M., Kent, S., Finkbeiner, D., Quinn, T. R., Hawley, S., Anderson, S., Kiuchi, F., Chen, A., Bushong, J., Sohi, H., Haggard, D., Kimball, A., McGurk, R., Barentine, J., Brewington, H., Harvanek, M., Kleinman, S., Krzesinski, J., Long, D., Nitta, A., Snedden, S., Lee, B., Pier, J. R., Harris, H., Brinkmann, J., and Schneider, D. P. (2010). The Milky Way Tomography with SDSS. III. Stellar Kinematics. *The Astrophysical Journal*, 716:1–29.
- [Butters et al., 2010] Butters, O. W., West, R. G., Anderson, D. R., Collier Cameron, A., Clarkson, W. I., Enoch, B., Haswell, C. A., Hellier, C., Horne, K., Joshi, Y., Kane, S. R., Lister, T. A., Maxted, P. F. L., Parley, N., Pollacco, D., Smalley, B., Street,

- R. A., Todd, I., Wheatley, P. J., and Wilson, D. M. (2010). The first WASP public data release. *Astronomy and Astrophysics*, 520:L10.
- [Charbonneau et al., 2000] Charbonneau, D., Brown, T. M., Latham, D. W., and Mayor, M. (2000). Detection of Planetary Transits Across a Sun-like Star. *The Astrophysical Journal Letters*, 529:L45–L48.
- [Choi et al., 2013] Choi, J.-Y., Han, C., Udalski, A., Sumi, T., Gaudi, B. S., Gould, A., Bennett, D. P., Dominik, M., Beaulieu, J.-P., Tsapras, Y., Bozza, V., Abe, F., Bond, I. A., Botzler, C. S., Chote, P., Freeman, M., Fukui, A., Furusawa, K., Itow, Y., Ling, C. H., Masuda, K., Matsubara, Y., Miyake, N., Muraki, Y., Ohnishi, K., Rattenbury, N. J., Saito, T., Sullivan, D. J., Suzuki, K., Sweatman, W. L., Suzuki, D., Takino, S., Tristram, P. J., Wada, K., Yock, P. C. M., MOA Collaboration, Szymański, M. K., Kubiak, M., Pietrzyński, G., Soszyński, I., Skowron, J., Kozłowski, S., Poleski, R., Ulaczyk, K., Wyrzykowski, L., Pietrukowicz, P., OGLE Collaboration, Almeida, L. A., DePoy, D. L., Dong, S., Gorbikov, E., Jablonski, F., Henderson, C. B., Hwang, K.-H., Janczak, J., Jung, Y.-K., Kaspi, S., Lee, C.-U., Malamud, U., Maoz, D., McGregor, D., Muñoz, J. A., Park, B.-G., Park, H., Pogge, R. W., Shvartzvald, Y., Shin, I.-G., Yee, J. C., The μ FUN Collaboration, Alsubai, K. A., Browne, P., Burgdorf, M. J., Calchi Novati, S., Dodds, P., Fang, X.-S., Finet, F., Glittrup, M., Grundahl, F., Gu, S.-H., Hardis, S., Harpsøe, K., Hinse, T. C., Hornstrup, A., Hundertmark, M., Jessen-Hansen, J., Jørgensen, U. G., Kains, N., Kerins, E., Liebig, C., Lund, M. N., Lundkvist, M., Maier, G., Mancini, L., Mathiasen, M., Penny, M. T., Rahvar, S., Ricci, D., Scarpetta, G., Skottfelt, J., Snodgrass, C., Southworth, J., Surdej, J., Tregloan-Reed, J., Wambsganss, J., Wertz, O., Zimmer, F., MiNDSTeP Consortium, T., Albrow, M. D., Bachelet, E., Batista, V., Brilliant, S., Cassan, A., Cole, A. A., Coutures, C., Dieters, S., Dominis Prester, D., Donatowicz, J., Fouqué, P., Greenhill, J., Kubas, D., Marquette, J.-B., Menzies, J. W., Sahu, K. C., Zub, M., The PLANET Collaboration, Bramich, D. M., Horne, K., Steele, I. A., Street, R. A., and The RoboNet Collaboration (2013). Microlensing Discovery of a Population of Very Tight, Very Low Mass Binary Brown Dwarfs. *The Astrophysical Journal*, 768:129.
- [Cox, 2000] Cox, A. N., editor (2000). *Allen’s astrophysical quantities*. Springer-Verlag.
- [Delfosse et al., 1997] Delfosse, X., Tinney, C. G., Forveille, T., Epchtein, N., Bertin, E., Borsenberger, J., Copet, E., de Batz, B., Fouque, P., Kimeswenger, S., Le Bertre, T., Lacombe, F., Rouan, D., and Tiphene, D. (1997). Field brown dwarfs found by DENIS. *Astronomy and Astrophysics*, 327:L25–L28.
- [Delorme et al., 2012] Delorme, P., Gagné, J., Malo, L., Reylé, C., Artigau, E., Albert, L., Forveille, T., Delfosse, X., Allard, F., and Homeier, D. (2012). CFBDSIR2149-0403: a 4-7 Jupiter-mass free-floating planet in the young moving group AB Doradus? *Astronomy and Astrophysics*, 548:A26.

- [Einstein, 1916] Einstein, A. (1916). Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354:769–822.
- [Einstein, 1936] Einstein, A. (1936). Lens-Like Action of a Star by the Deviation of Light in the Gravitational Field. *Science*, 84:506–507.
- [Fukui et al., 2007] Fukui, A., Abe, F., Ayani, K., Fujii, M., Iizuka, R., Itow, Y., Kabumoto, K., Kamiya, K., Kawabata, T., Kawanomoto, S., Kinugasa, K., Koff, R. A., Krajci, T., Naito, H., Nogami, D., Narusawa, S., Ohishi, N., Ohnishi, K., Sumi, T., and Tsumuraya, F. (2007). Observation of the First Gravitational Microlensing Event in a Sparse Stellar Field: The Tago Event. *The Astrophysical Journal*, 670:423–427.
- [Gaudi et al., 2008a] Gaudi, B. S., Bennett, D. P., Udalski, A., Gould, A., Christie, G. W., Maoz, D., Dong, S., McCormick, J., Szymański, M. K., Tristram, P. J., Nikolaev, S., Paczyński, B., Kubiak, M., Pietrzyński, G., Soszyński, I., Szewczyk, O., Ulaczyk, K., Wyrzykowski, Ł., OGLE Collaboration, DePoy, D. L., Han, C., Kaspi, S., Lee, C.-U., Mallia, F., Natusch, T., Pogge, R. W., Park, B.-G., μ -Fun Collaboration, Abe, F., Bond, I. A., Botzler, C. S., Fukui, A., Hearnshaw, J. B., Itow, Y., Kamiya, K., Korpela, A. V., Kilmartin, P. M., Lin, W., Masuda, K., Matsubara, Y., Motomura, M., Muraki, Y., Nakamura, S., Okumura, T., Ohnishi, K., Rattenbury, N. J., Sako, T., Saito, T., Sato, S., Skuljan, L., Sullivan, D. J., Sumi, T., Sweatman, W. L., Yock, P. C. M., MOA Collaboration, Albrow, M. D., Allan, A., Beaulieu, J.-P., Burgdorf, M. J., Cook, K. H., Coutures, C., Dominik, M., Dieters, S., Fouqué, P., Greenhill, J., Horne, K., Steele, I., Tsapras, Y., Planet Collaboration, RoboNet Collaborations, Chaboyer, B., Crocker, A., Frank, S., and Macintosh, B. (2008a). Discovery of a Jupiter/Saturn Analog with Gravitational Microlensing. *Science*, 319:927–.
- [Gaudi et al., 2008b] Gaudi, B. S., Patterson, J., Spiegel, D. S., Krajci, T., Koff, R., Pojmański, G., Dong, S., Gould, A., Prieto, J. L., Blake, C. H., Roming, P. W. A., Bennett, D. P., Bloom, J. S., Boyd, D., Eyler, M. E., de Ponthière, P., Mirabal, N., Morgan, C. W., Remillard, R. R., Vanmunster, T., Wagner, R. M., and Watson, L. C. (2008b). Discovery of a Very Bright, Nearby Gravitational Microlensing Event. *The Astrophysical Journal*, 677:1268–1277.
- [Ghez et al., 2008] Ghez, A. M., Salim, S., Weinberg, N. N., Lu, J. R., Do, T., Dunn, J. K., Matthews, K., Morris, M. R., Yelda, S., Becklin, E. E., Kremenek, T., Milosavljević, M., and Naiman, J. (2008). Measuring Distance and Properties of the Milky Way’s Central Supermassive Black Hole with Stellar Orbits. *The Astrophysical Journal*, 689:1044–1062.
- [Golimowski et al., 1995] Golimowski, D. A., Nakajima, T., Kulkarni, S. R., and Oppenheimer, B. R. (1995). Detection of a very low mass companion to the astrometric binary Gliese 105A. *The Astrophysical Journal Letters*, 444:L101–L104.

- [Gould, 2013] Gould, A. (2013). LSST’s DC Bias Against Planets and Galactic-Plane Science. *ArXiv e-prints*.
- [Gould and Gaucherel, 1996] Gould, A. and Gaucherel, C. (1996). Finite Source Effects in Microlensing Events. *ArXiv Astrophysics e-prints*.
- [Gould and Loeb, 1992] Gould, A. and Loeb, A. (1992). Discovering planetary systems through gravitational microlenses. *The Astrophysical Journal*, 396:104–114.
- [Gould and Yee, 2014] Gould, A. and Yee, J. C. (2014). Microlens Masses from Astrometry and Parallax in Space-based Surveys: From Planets to Black Holes. *The Astrophysical Journal*, 784:64.
- [Han, 2008] Han, C. (2008). Near-Field Microlensing from Wide-Field Surveys. *The Astrophysical Journal*, 681:806–813.
- [Han, 2011] Han, C. (2011). Dependence of the Sensitivity to Planets on the Properties of High-Magnification Gravitational Microlensing Events. *Journal of Korean Astronomical Society*, 44:109–113.
- [Han and Kang, 2003] Han, C. and Kang, Y. W. (2003). Probing the Spatial Distribution of Extrasolar Planets with Gravitational Microlensing. *The Astrophysical Journal*, 596:1320–1326.
- [Hansen et al., 2004] Hansen, C. J., Kawaler, S. D., and Trimble, V. (2004). *Stellar interiors : physical principles, structure, and evolution*. Springer Science+Business Media.
- [Hart et al., 1996] Hart, J., van Hermelen, J., Hovey, G., Freeman, K. C., Peterson, B. A., Axelrod, T. S., Quinn, P. J., Rodgers, A. W., Allsman, R. A., Alcock, C., Bennett, D. P., Cook, K. H., Griest, K., Marshall, S. L., Pratt, M. R., Stubbs, C. W., and Sutherland, W. (1996). The Telescope System of the MACHO Program. *Publications of the Astronomical Society of the Pacific*, 108:220.
- [Hayashi and Nakano, 1963] Hayashi, C. and Nakano, T. (1963). Evolution of stars of small masses in the pre-main-sequence stages. *Progress of Theoretical Physics*, 30(4):460–474.
- [Hearnshaw et al., 2006] Hearnshaw, J. B., Abe, F., Bond, I. A., Gilmore, A. C., Itow, Y., Kamiya, K., Masuda, K., Matsubara, Y., Muraki, Y., Okada, C., Rattenbury, N. J., Sako, T., Sasaki, M., Sullivan, D. J., and Yock, P. C. M. (2006). The MOA 1.8-metre alt-az Wide-field Survey Telescope and the MOA Project. In Sutantyo, W., Premadi, P. W., Mahasena, P., Hidayat, T., and Mineshige, S., editors, *The 9th Asian-Pacific Regional IAU Meeting*, page 272.

- [Hubble, 1925] Hubble, E. P. (1925). Cepheids in spiral nebulae. *The Observatory*, 48:139–142.
- [Ivezić et al., 2008] Ivezić, Ž., Sesar, B., Jurić, M., Bond, N., Dalcanton, J., Rockosi, C. M., Yanny, B., Newberg, H. J., Beers, T. C., Allende Prieto, C., Wilhelm, R., Lee, Y. S., Sivarani, T., Norris, J. E., Bailer-Jones, C. A. L., Re Fiorentin, P., Schlegel, D., Uomoto, A., Lupton, R. H., Knapp, G. R., Gunn, J. E., Covey, K. R., Smith, J. A., Miknaitis, G., Doi, M., Tanaka, M., Fukugita, M., Kent, S., Finkbeiner, D., Munn, J. A., Pier, J. R., Quinn, T., Hawley, S., Anderson, S., Kiuchi, F., Chen, A., Bushong, J., Sohi, H., Haggard, D., Kimball, A., Barentine, J., Brewington, H., Harvanek, M., Kleinman, S., Krzesinski, J., Long, D., Nitta, A., Snedden, S., Lee, B., Harris, H., Brinkmann, J., Schneider, D. P., and York, D. G. (2008). The Milky Way Tomography with SDSS. II. Stellar Metallicity. *The Astrophysical Journal*, 684:287–325.
- [Jacob, 1855] Jacob, W. S. (1855). On certain Anomalies presented by the Binary Star 70 Ophiuchi. *Monthly Notices of the Royal Astronomical Society*, 15:228.
- [Jung et al., 2014] Jung, Y. K., Park, H., Han, C., Hwang, K.-H., Shin, I.-G., and Choi, J.-Y. (2014). Reevaluating the Feasibility of Ground-based Earth-mass Microlensing Planet Detections. *The Astronomical Journal*, 786:85.
- [Jurić et al., 2008] Jurić, M., Ivezić, Ž., Brooks, A., Lupton, R. H., Schlegel, D., Finkbeiner, D., Padmanabhan, N., Bond, N., Sesar, B., Rockosi, C. M., Knapp, G. R., Gunn, J. E., Sumi, T., Schneider, D. P., Barentine, J. C., Brewington, H. J., Brinkmann, J., Fukugita, M., Harvanek, M., Kleinman, S. J., Krzesinski, J., Long, D., Neilsen, Jr., E. H., Nitta, A., Snedden, S. A., and York, D. G. (2008). The Milky Way Tomography with SDSS. I. Stellar Number Density Distribution. *The Astrophysical Journal*, 673:864–914.
- [Kirkpatrick et al., 2011] Kirkpatrick, J. D., Cushing, M. C., Gelino, C. R., Griffith, R. L., Skrutskie, M. F., Marsh, K. A., Wright, E. L., Mainzer, A., Eisenhardt, P. R., McLean, I. S., Thompson, M. A., Bauer, J. M., Benford, D. J., Bridge, C. R., Lake, S. E., Petty, S. M., Stanford, S. A., Tsai, C.-W., Bailey, V., Beichman, C. A., Bloom, J. S., Bochanski, J. J., Burgasser, A. J., Capak, P. L., Cruz, K. L., Hinz, P. M., Kartaltepe, J. S., Knox, R. P., Manohar, S., Masters, D., Morales-Calderón, M., Prato, L. A., Rodigas, T. J., Salvato, M., Schurr, S. D., Scoville, N. Z., Simcoe, R. A., Stapelfeldt, K. R., Stern, D., Stock, N. D., and Vacca, W. D. (2011). The First Hundred Brown Dwarfs Discovered by the Wide-field Infrared Survey Explorer (WISE). *The Astrophysical Journal Supplement*, 197:19.
- [Kirkpatrick et al., 2000] Kirkpatrick, J. D., Reid, I. N., Liebert, J., Gizis, J. E., Burgasser, A. J., Monet, D. G., Dahn, C. C., Nelson, B., and Williams, R. J. (2000). 67 Additional L Dwarfs Discovered by the Two Micron All Sky Survey. *The Astronomical Journal*, 120:447–472.

- [Kleinmann et al., 1994] Kleinmann, S. G., Lysaght, M. G., Pughe, W. L., Schneider, S. E., Skrutskie, M. F., Weinberg, M. D., Price, S. D., Matthews, K., Soifer, B. T., and Huchra, J. P. (1994). The Two Micron All Sky Survey. *Astrophysics and Space Science*, 217:11–17.
- [Kroupa et al., 1993] Kroupa, P., Tout, C. A., and Gilmore, G. (1993). The distribution of low-mass stars in the Galactic disc. *Monthly Notices of the Royal Astronomical Society*, 262:545–587.
- [Kumar, 1963] Kumar, S. S. (1963). The Structure of Stars of Very Low Mass. *The Astrophysical Journal*, 137:1121.
- [Liebes, 1964] Liebes, S. (1964). Gravitational Lenses. *Physical Review*, 133:835–844.
- [LSST Science Collaboration et al., 2009] LSST Science Collaboration, Abell, P. A., Allison, J., Anderson, S. F., Andrew, J. R., Angel, J. R. P., Armus, L., Arnett, D., Asztalos, S. J., Axelrod, T. S., and et al. (2009). LSST Science Book, Version 2.0. *ArXiv e-prints*.
- [Mace et al., 2013] Mace, G. N., Kirkpatrick, J. D., Cushing, M. C., Gelino, C. R., Griffith, R. L., Skrutskie, M. F., Marsh, K. A., Wright, E. L., Eisenhardt, P. R., McLean, I. S., Thompson, M. A., Mix, K., Bailey, V., Beichman, C. A., Bloom, J. S., Burgasser, A. J., Fortney, J. J., Hinz, P. M., Knox, R. P., Lowrance, P. J., Marley, M. S., Morley, C. V., Rodigas, T. J., Saumon, D., Sheppard, S. S., and Stock, N. D. (2013). A Study of the Diverse T Dwarf Population Revealed by WISE. *The Astrophysical Journal Supplement*, 205:6.
- [Mainzer et al., 2011] Mainzer, A., Bauer, J., Grav, T., Masiero, J., Cutri, R. M., Dailey, J., Eisenhardt, P., McMillan, R. S., Wright, E., Walker, R., Jedicke, R., Spahr, T., Tholen, D., Alles, R., Beck, R., Brandenburg, H., Conrow, T., Evans, T., Fowler, J., Jarrett, T., Marsh, K., Masci, F., McCallon, H., Wheelock, S., Wittman, M., Wyatt, P., DeBaun, E., Elliott, G., Elsbury, D., Gautier, IV, T., Gomillion, S., Leisawitz, D., Maleszewski, C., Micheli, M., and Wilkins, A. (2011). Preliminary Results from NEOWISE: An Enhancement to the Wide-field Infrared Survey Explorer for Solar System Science. *The Astrophysical Journal*, 731:53.
- [Mao and Paczynski, 1991] Mao, S. and Paczynski, B. (1991). Gravitational microlensing by double stars and planetary systems. *The Astrophysical Journal*, 374:L37–L40.
- [Marcy et al., 1997] Marcy, G. W., Butler, R. P., Williams, E., Bildsten, L., Graham, J. R., Ghez, A. M., and Jernigan, J. G. (1997). The Planet around 51 Pegasi. *The Astrophysical Journal*, 481:926–935.
- [Marshall et al., 2006] Marshall, D. J., Robin, A. C., Reylé, C., Schultheis, M., and Picaud, S. (2006). Modelling the Galactic interstellar extinction distribution in three dimensions. *Astronomy and Astrophysics*, 453:635–651.

- [Martin et al., 2010] Martin, C. L., Scannapieco, E., Ellison, S. L., Hennawi, J. F., Djorgovski, S. G., and Fournier, A. P. (2010). The Size and Origin of Metal-enriched Regions in the Intergalactic Medium from Spectra of Binary Quasars. *The Astrophysical Journal*, 721:174–192.
- [Mayor and Queloz, 1995] Mayor, M. and Queloz, D. (1995). A Jupiter-mass companion to a solar-type star. *Nature*, 378:355–359.
- [Miyake et al., 2012] Miyake, N., Udalski, A., Sumi, T., Bennett, D. P., Dong, S., Street, R. A., Greenhill, J., Bond, I. A., Gould, A., Kubiak, M., Szymański, M. K., Pietrzyński, G., Soszyński, I., Ulaczyk, K., Wyrzykowski, L., OGLE Collaboration, Abe, F., Fukui, A., Furusawa, K., Holderness, S., Itow, Y., Korpela, A., Ling, C. H., Masuda, K., Matsubara, Y., Muraki, Y., Nagayama, T., Ohnishi, K., Rattenbury, N., Saito, T., Sako, T., Sullivan, D. J., Sweatman, W. L., Tristram, P. J., Yock, P. C. M., MOA Collaboration, Allen, W., Christie, G. W., DePoy, D. L., Gaudi, B. S., Han, C., Lee, C.-U., McCormick, J., Monard, B., Natusch, T., Park, B.-G., Pogge, R. W., μ FUN Collaboration, Allan, A., Bode, M., Bramich, D. M., Clay, N., Dominik, M., Horne, K. D., Kains, N., Mottram, C., Snodgrass, C., Steele, I., Tsapras, Y., RoboNet Collaboration, Albrow, M. D., Batista, V., Beaulieu, J. P., Brilliant, S., Burgdorf, M., Caldwell, J. A. R., Cassan, A., Cole, A., Cook, K. H., Coutures, C., Dieters, S., Dominis Prester, D., Donatowicz, J., Fouqué, P., Jorgensen, U. G., Kane, S., Kubas, D., Marquette, J. B., Martin, R., Menzies, J., Pollard, K. R., Sahu, K. C., Wambsganss, J., Williams, A., Zub, M., and PLANET Collaboration (2012). A Possible Binary System of a Stellar Remnant in the High-magnification Gravitational Microlensing Event OGLE-2007-BLG-514. *The Astrophysical Journal*, 752:82.
- [Muraki et al., 1999] Muraki, Y., Sumi, T., Abe, F., Bond, I., Carter, B., Dodd, R., Fujimoto, M., Hearnshaw, J., Honda, M., Jugaku, J., Kabe, S., Kato, Y., Kobayashi, M., Koribalski, B., Kilmartin, P., Masuda, K., Matsubara, Y., Nakamura, T., Noda, S., Pennycook, G., Rattenbury, N., Reid, M., Saito, T., Sato, H., Sato, S., Sekiguchi, M., Sullivan, D., Takeuti, M., Watase, Y., Yanagisawa, T., Yock, P., and Yoshizawa, M. (1999). Search for Machos by the MOA Collaboration. *Progress of Theoretical Physics Supplement*, 133:233–246.
- [Muterspaugh et al., 2010a] Muterspaugh, M. W., Fekel, F. C., Lane, B. F., Hartkopf, W. I., Kulkarni, S. R., Konacki, M., Burke, B. F., Colavita, M. M., Shao, M., and Williamson, M. (2010a). The Phases Differential Astrometry Data Archive. IV. The Triple Star Systems 63 Gem A and HR 2896. *The Astronomical Journal*, 140:1646–1656.
- [Muterspaugh et al., 2010b] Muterspaugh, M. W., Hartkopf, W. I., Lane, B. F., O’Connell, J., Williamson, M., Kulkarni, S. R., Konacki, M., Burke, B. F., Colavita, M. M., Shao, M., and Wiktorowicz, S. J. (2010b). The Phases Differential Astrometry

- Data Archive. II. Updated Binary Star Orbits and a Long Period Eclipsing Binary. *The Astronomical Journal*, 140:1623–1630.
- [Muterspaugh et al., 2010c] Muterspaugh, M. W., Lane, B. F., Kulkarni, S. R., Konacki, M., Burke, B. F., Colavita, M. M., and Shao, M. (2010c). The Phases Differential Astrometry Data Archive. III. Limits to Tertiary Companions. *The Astronomical Journal*, 140:1631–1645.
- [Muterspaugh et al., 2010d] Muterspaugh, M. W., Lane, B. F., Kulkarni, S. R., Konacki, M., Burke, B. F., Colavita, M. M., Shao, M., Hartkopf, W. I., Boss, A. P., and Williamson, M. (2010d). The Phases Differential Astrometry Data Archive. V. Candidate Substellar Companions to Binary Systems. *The Astronomical Journal*, 140:1657–1671.
- [Muterspaugh et al., 2010e] Muterspaugh, M. W., Lane, B. F., Kulkarni, S. R., Konacki, M., Burke, B. F., Colavita, M. M., Shao, M., Wiktorowicz, S. J., and O’Connell, J. (2010e). The Phases Differential Astrometry Data Archive. I. Measurements and Description. *The Astronomical Journal*, 140:1579–1622.
- [NASA-JPL, 1999] NASA-JPL (1999). Mars climate orbiter failure board releases report, numerous nasa actions underway in response. Press release. <http://mars.jpl.nasa.gov/msp98/news/mco991110.html>, accessed Mar. 9 2015.
- [Paczynski, 1986] Paczynski, B. (1986). Gravitational microlensing by the galactic halo. *The Astrophysical Journal*, 304:1–5.
- [Pavlenko, 1997] Pavlenko, Y. V. (1997). Analysis of the Spectra of Two Pleiades Brown Dwarfs: Teide 1 and Calar 3. *Astrophysics and Space Science*, 253:43–53.
- [Peale, 2001] Peale, S. J. (2001). Probability of Detecting a Planetary Companion during a Microlensing Event. *The Astrophysical Journal*, 552:889–911.
- [Phan-Bao et al., 2003] Phan-Bao, N., Crifo, F., Delfosse, X., Forveille, T., Guibert, J., Borsenberger, J., Epchtein, N., Fouqué, P., Simon, G., and Vetois, J. (2003). New neighbours. V. 35 DENIS late-M dwarfs between 10 and 30 parsecs. *Astronomy and Astrophysics*, 401:959–974.
- [Phan-Bao et al., 2001] Phan-Bao, N., Guibert, J., Crifo, F., Delfosse, X., Forveille, T., Borsenberger, J., Epchtein, N., Fouqué, P., and Simon, G. (2001). New neighbours: IV. 30 DENIS late-M dwarfs between 15 and 30 parsecs. *Astronomy and Astrophysics*, 380:590–598.
- [Pollacco et al., 2006] Pollacco, D. L., Skillen, I., Collier Cameron, A., Christian, D. J., Hellier, C., Irwin, J., Lister, T. A., Street, R. A., West, R. G., Anderson, D. R., Clarkson, W. I., Deeg, H., Enoch, B., Evans, A., Fitzsimmons, A., Haswell, C. A.,

- Hodgkin, S., Horne, K., Kane, S. R., Keenan, F. P., Maxted, P. F. L., Norton, A. J., Osborne, J., Parley, N. R., Ryans, R. S. I., Smalley, B., Wheatley, P. J., and Wilson, D. M. (2006). The WASP Project and the SuperWASP Cameras. *The Publications of the Astronomical Society of the Pacific*, 118:1407–1418.
- [Press et al., 1992] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical recipes in C. The art of scientific computing. Second Edition*. Cambridge University Press.
- [Price-Whelan et al., 2014] Price-Whelan, A. M., Agüeros, M. A., Fournier, A. P., Street, R., Ofek, E. O., Covey, K. R., Levitan, D., Laher, R. R., Sesar, B., and Surace, J. (2014). Statistical Searches for Microlensing Events in Large, Non-uniformly Sampled Time-Domain Surveys: A Test Using Palomar Transient Factory Data. *The Astrophysical Journal*, 781:35.
- [Rahvar and Dominik, 2008] Rahvar, S. and Dominik, M. (2008). Detecting exoplanets with the xallarap microlensing effect. In *Manchester Microlensing Conference*, page 37.
- [Refsdal, 1964] Refsdal, S. (1964). The gravitational lens effect. *Monthly Notices of the Royal Astronomical Society*, 128:295.
- [Robin et al., 2003] Robin, A. C., Reylé, C., Derrière, S., and Picaud, S. (2003). A synthetic view on structure and evolution of the Milky Way. *Astronomy and Astrophysics*, 409:523–540.
- [Sako et al., 2008] Sako, T., Sekiguchi, T., Sasaki, M., Okajima, K., Abe, F., Bond, I. A., Hearnshaw, J. B., Itow, Y., Kamiya, K., Kilmartin, P. M., Masuda, K., Matsubara, Y., Muraki, Y., Rattenbury, N. J., Sullivan, D. J., Sumi, T., Tristram, P., Yanagisawa, T., and Yock, P. C. M. (2008). MOA-cam3: a wide-field mosaic CCD camera for a gravitational microlensing survey in New Zealand. *Experimental Astronomy*, 22:51–66.
- [Schlegel et al., 1998] Schlegel, D. J., Finkbeiner, D. P., and Davis, M. (1998). Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds. *The Astrophysical Journal*, 500:525–553.
- [Schwarzschild, 1916] Schwarzschild, K. (1916). Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, 1916, Seite 189-196, pages 189–196.
- [Seares and Joyner, 1928] Seares, F. H. and Joyner, M. C. (1928). Systematic Deviations from the Mean Stellar Distribution. *Astrophysical Journal*, 67:24.
- [Seares et al., 1925] Seares, F. H., van Rhijn, P. J., Joyner, M. C., and Richmond, M. L. (1925). Mean Distribution of Stars According to Apparent Magnitude and Galactic Latitude. *Astrophysical Journal*, 62:320.

- [Smith and WASP Consortium, 2014] Smith, A. M. S. and WASP Consortium (2014). The SuperWASP exoplanet transit survey. *Contributions of the Astronomical Observatory Skalnaté Pleso*, 43:500–512.
- [Sofue, 2012] Sofue, Y. (2012). Grand Rotation Curve and Dark Matter Halo in the Milky Way Galaxy. *Publications of the Astronomical Society of Japan*, 64:75.
- [Song et al., 2014] Song, Y.-Y., Mao, S., and An, J. H. (2014). Degeneracies in triple gravitational microlensing. *Monthly Notices of the Royal Astronomical Society*, 437:4006–4018.
- [Struve, 1952] Struve, O. (1952). Proposal for a project of high-precision stellar radial velocity work. *The Observatory*, 72:199–200.
- [Sumi et al., 2011] Sumi, T., Kamiya, K., Bennett, D. P., Bond, I. A., Abe, F., Botzler, C. S., Fukui, A., Furusawa, K., Hearnshaw, J. B., Itow, Y., Kilmartin, P. M., Korpela, A., Lin, W., Ling, C. H., Masuda, K., Matsubara, Y., Miyake, N., Motomura, M., Muraki, Y., Nagaya, M., Nakamura, S., Ohnishi, K., Okumura, T., Perrott, Y. C., Rattenbury, N., Saito, T., Sako, T., Sullivan, D. J., Sweatman, W. L., Tristram, P. J., Udalski, A., Szymański, M. K., Kubiak, M., Pietrzyński, G., Poleski, R., Soszyński, I., Wyrzykowski, Ł., Ulaczyk, K., and Microlensing Observations in Astrophysics (MOA) Collaboration (2011). Unbound or distant planetary mass population detected by gravitational microlensing. *Nature*, 473:349–352.
- [Udalski, 2003] Udalski, A. (2003). The Optical Gravitational Lensing Experiment. Real Time Data Analysis Systems in the OGLE-III Survey. *Acta Astronomica*, 53:291–305.
- [Udalski et al., 1997] Udalski, A., Kubiak, M., and Szymanski, M. (1997). Optical Gravitational Lensing Experiment. OGLE-2 – the Second Phase of the OGLE Project. *Acta Astronomica*, 47:319–344.
- [Udalski et al., 1992] Udalski, A., Szymanski, M., Kaluzny, J., Kubiak, M., and Mateo, M. (1992). The Optical Gravitational Lensing Experiment. *Acta Astronomica*, 42:253–284.
- [Udalski et al., 2008] Udalski, A., Szymanski, M. K., Soszynski, I., and Poleski, R. (2008). The Optical Gravitational Lensing Experiment. Final Reductions of the OGLE-III Data. *Acta Astronomica*, 58:69–87.
- [van Rhijn, 1929] van Rhijn, P. J. (1929). Distribution of Stars According to Apparent Magnitude, Galactic Latitude and Galactic Longitude. *Publications of the Kapteyn Astronomical Laboratory Groningen*, 43:1–104.
- [Wald, 1984] Wald, R. M. (1984). *General relativity*. University of Chicago Press.

- [Wolszczan and Frail, 1992] Wolszczan, A. and Frail, D. A. (1992). A planetary system around the millisecond pulsar PSR1257 + 12. *Nature*, 355:145–147.
- [Wright et al., 2010] Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., Ressler, M. E., Cutri, R. M., Jarrett, T., Kirkpatrick, J. D., Padgett, D., McMillan, R. S., Skrutskie, M., Stanford, S. A., Cohen, M., Walker, R. G., Mather, J. C., Leisawitz, D., Gautier, III, T. N., McLean, I., Benford, D., Lonsdale, C. J., Blain, A., Mendez, B., Irace, W. R., Duval, V., Liu, F., Royer, D., Heinrichsen, I., Howard, J., Shannon, M., Kendall, M., Walsh, A. L., Larsen, M., Cardon, J. G., Schick, S., Schwalm, M., Abid, M., Fabinsky, B., Naes, L., and Tsai, C.-W. (2010). The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance. *The Astronomical Journal*, 140:1868–1881.
- [Wyrzykowski et al., 2014] Wyrzykowski, Ł., Kostrzewa-Rutkowska, Z., Kozłowski, S., Udalski, A., Poleski, R., Skowron, J., Błagorodnova, N., Kubiak, M., Szymański, M. K., Pietrzyński, G., Soszyński, I., Ulaczyk, K., Pietrukowicz, P., and Mróz, P. (2014). OGLE-IV Real-Time Transient Search. *Acta Astronomica*, 64:197–232.
- [Yanagisawa et al., 2000] Yanagisawa, T., Muraki, Y., Matsubara, Y., Abe, F., Masuda, K., Noda, S., Sumi, T., Kato, Y., Fujimoto, M., Sato, S., Bond, I., Rattenbury, N., Yock, P., Kilmartin, P., Hearnshaw, J., Reid, M., Sullivan, D., Carter, B., Dodd, R., Nankivell, G., Rumsey, N., Honda, M., Sekiguchi, M., Yoshizawa, M., Nakamura, T., Sato, H., Kabe, S., Kobayashi, M., Watase, Y., Jugaku, J., Saito, T., and Koribalsky, B. (2000). Wide-Field Camera for Gravitational Microlensing Survey: MOA-cam2. *Experimental Astronomy*, 10:519–535.
- [Yee et al., 2014] Yee, J. C., Udalski, A., Calchi Novati, S., Gould, A., Carey, S., Poleski, R., Gaudi, B. S., Pogge, R. W., Skowron, J., Kozłowski, S., Mróz, P., Pietrukowicz, P., Pietrzyński, G., Szymański, M. K., Soszyński, I., Ulaczyk, K., and Wyrzykowski, Ł. (2014). First Space-based Microlens Parallax Measurement of an Isolated Star: Spitzer Observations of OGLE-2014-BLG-0939. *ArXiv e-prints*.
- [Yoo et al., 2004] Yoo, J., DePoy, D. L., Gal-Yam, A., Gaudi, B. S., Gould, A., Han, C., Lipkin, Y., Maoz, D., Ofek, E. O., Park, B.-G., Pogge, R. W., Mu-Fun Collaboration, Udalski, A., Soszyński, I., Wyrzykowski, Ł., Kubiak, M., Szymański, M., Pietrzyński, G., Szewczyk, O., Żebruń, K., and OGLE Collaboration (2004). OGLE-2003-BLG-262: Finite-Source Effects from a Point-Mass Lens. *The Astrophysical Journal*, 603:139–151.
- [Zombeck, 2007] Zombeck, M. (2007). *Handbook of Space Astronomy and Astrophysics: Third Edition*. Cambridge University Press.