Orbital Architectures of M Dwarf Companions

by

Eliot Halley Vrijmoet

Under the Direction of Todd J. Henry, Ph.D.

A Dissertation Submitted in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

in the College of Arts and Sciences

Georgia State University

2023

ABSTRACT

The M dwarfs are the most numerous stars in our neighborhood, but also the faintest, leaving open challenges in determining their fundamental properties and the details of their formation histories. Considering that 27% of M dwarfs are in systems of two or more stars, with this dissertation we are surveying the properties of these multiples and how they behave dynamically. Our goal is to establish the orbital parameter distributions for M dwarf companions with orbital periods up to 30 years, forming one picture that encompasses the regimes of several observing techniques. To map this wide parameter space, we are using: (1) astrometry from the 23-year REsearch Consortium On Nearby Stars (RECONS) program at CTIO, (2) speckle interferometry from SOAR HRCam+SAM, and (3) additional imaging and spectroscopic orbits from the literature. This combined data set of nearly 200 orbits reveals three new fundamental results. The first is that the tidal circularization period is between 4 and 7 days for M dwarfs, shorter than for solar-type binaries. The second is that the eccentricity distribution at orbital periods >5 years depends on the mass of the primary star and potentially the mass ratio, as nearly-circular orbits occur only for systems in which at least one component has mass <0.1 $M_{\odot}$. The third result is that the mass ratios of M dwarf multiples are strongly weighted toward unity, with most systems having $q = M_2/M_1 \geq 0.9$. This suggests these low-mass systems' early dynamical evolution is dominated by migration through circumstellar disks. This picture of companions includes brown dwarfs as well as stars across the entire low-mass main sequence, showing that formation and dynamical evolution processes create and mold a continuum of mass all the way down to planetary regimes.

INDEX WORDS: Sample index words, Sample, Sample, Sample

INDEX WORDS: Sample index words, Sample, Sample, Sample

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Introduction: Motivating Orbital Architectures of the Ubiquitous Red Dwarfs

As we carry out our urge to explore, discover, and learn, we measure every new place or experience against our own home and its environments — either to rediscover the familiar, or to identify the most strange.

In the context of modern astronomy, this tendency is most visible in our search for extrasolar planets ("exoplanets"), but our search for the familiar can also be seen in studies of stars and galaxies, as we compare each observed phenomenon to the Sun or the Milky Way. Despite phenomenal success in finding distant worlds, astronomical exploration has yet to discover another multi-planet system exactly like our own. Rather than discouraging us, this difficulty has motivated us toward the longer journey of unraveling how planetary systems form, evolve, and are influenced by their stellar environments and histories. Those environments are, in turn, affected by how those stars form and evolve, and those processes are affected by how the diversity of other stars in their neighborhood form and evolve — and each of these things depends on factors spanning molecular to galactic scales.

It is with this motivation that we embark on this dissertation that examines low-mass stars in systems with multiple stellar components. In this work, each star considered has less than 60% the mass of our Sun and is gravitationally bound to other low-mass stars (or, in a few cases, to brown dwarfs), which admittedly makes them more strange than familiar. On the other hand, these non-solar systems are an inevitable part of the discussion on star and planet formation, as the chain of processes that formed our Sun produces these small

stars just as well; it produces them in such abundance that they comprise three out of every four stars in our neighborhood (Henry et al. 2006, 2018). Predicting the frequency and multiplicity of these stars is thus an essential test of our models of star formation. Predicting the subsequent properties of these multiples — the most common stellar pairs, how widely they are separated, and the orbits they settle into — then tests our understanding of how stars interact with the gas, dust, and other stars in their post-formation environments.

This dissertation addresses that final line of inquiry. We begin by formally asking,

What are the orbital architectures of companions to red dwarf stars?

Our goals are to establish empirically the sizes and shapes of primarily stellar companions' orbits around low-mass stars, and in doing so, to provide constraints for models of star formation and dynamical evolution. The key relationship that will be explored is the comparison of orbital periods (representing the sizes of orbits) to eccentricities (orbital shapes), or the plot of $P_{\rm orb}$ vs. e. The $P_{\rm orb}$ vs. e plot reveals a remarkable amount of information about which formation scenarios apply to various types of stars. These models and empirical results will, in turn, have implications for all phenomena that depend on stellar properties and histories, including planet formation and evolution.

In this first chapter we begin with nomenclature on multiple stars (§1.1), then discuss the current family of paradigms of star formation (§1.2) and how stellar orbits change and evolve (§1.3). We then review what observations have revealed on these topics for a variety of types of main sequence stars (§1.4), including our low-mass "red dwarf" stars of interest.

The remaining chapters of this dissertation describe our investigation of the orbital architectures of low-mass multiples: Chapter 2 defines our stellar regime of focus, Chapter 3 and Chapter 4 explain our observations of several hundred low-mass star systems, Chapter 5 presents the results of each of those observing campaigns, and Chapter 6 synthesizes those results and explores the implications for these star systems' formation and dynamical evolution. Finally, Chapter 7 outlines the conclusions that can be drawn from this work and posits the next steps for future studies.

1.1 Terms and Conditions

Throughout this dissertation, we will frequently refer to vocabulary specific to systems of multiple stars and their orbits. Here we pause to define some of these terms, for clarity.

First, the subjects of this dissertation are the stars on the main sequence with masses of 0.075–0.60 $M_{\odot}$ that have spectral types of M0.0V to L2.5V. We will refer to these objects as "M dwarfs" or "red dwarfs" throughout.

In this work, a *system* is considered to be one or more stars; thus, a single star is considered to be a system. A system consisting two or more stars that are gravitationally bound is known as a multi-star system or *multiple system*. A multiple of exactly two stars is a binary, three stars is a triple, and four is a quadruple. The stars in these systems are sometimes referred to as its *components*, and individually these components are named the primary star (denoted "A"), the secondary star ("B"), the tertiary star ("C"), etc. The latter stars are generally called companions of the primary. These components are named in

order of most massive to least massive, using each star's maximum mass (over its lifetime) as the distinguishing factor. Note that this means in a binary with a white dwarf and M dwarf, the white dwarf is the primary, because earlier in the system's history the white dwarf was the most massive component. In systems with no white dwarfs, the brightest component is assumed to be the most massive — the primary star (component A) is identified to be the one that is brightest at visible wavelengths, first considered in the *Gaia* G photometric band, or in the Johnson system V photometric band if G is not available. Component B is the next brightest, component C the third brightest, etc.

Due to their mutual gravitation, the stars in a multiple orbit their common center of mass. Often we will consider the alternate reference frame centered on the primary star, in which case the companions are said to orbit the primary. Most of the parameters describing an orbit do not change with this change of reference frame (with the exception of the semi-major axis).

The parameters describing the orbit are illustrated in Figure 1.1 and described as follows:

- Orbital period ($P_{\rm orb}$): the time required for each star to complete its orbital path and return to a given point. In a binary system, the period of the primary and secondary around their center of mass are identical. The orbital period of the secondary in the frame where it orbits the primary is the same as in that center of mass frame.

- Eccentricity (e): the shape of the orbit, given as a dimensionless value where a perfectly circular orbit has $e = 0$. For these bound systems, eccentricity is limited as $e < 1$.

- Semi-major axis (a): the size of the orbit, defined as its length along its longest axis. The semi-major axis of the relative orbit is the sum of the semi-major axes of the primary and secondary's orbits around the center of mass: $a = a_1 + a_2$. In this work, this property is typically measured in arcseconds, milliarcseconds (mas), and astronomical units (AU).

- Argument of periastron (ω): angular location along the orbit at which the primary and secondary are closest, measured with respect to the ascending node (see below) and the primary star or center of mass.

- Longitude of the ascending node (Ω): angular location at which the orbit crosses the plane of the sky and the orbiting body is receding from the observer (the other intersection point, at which the star is approaching the observer, is the descending node). This angle is typically measured with respect to the vernal equinox.

- Inclination (i): angle between the plane of the orbit and the plane of the sky, with $i = 90°$ corresponding to an "edge-on" orbit from the observer's perspective.

- Time of periastron passage (T_0): the date at which the stars in a binary cross their orbit's periastron, in the center of mass frame, or the date when the secondary crosses its periastron with respect to the primary. Some authors give alternative reference times in lieu of periastron passage time — for example, eclipsing binary publications refer to the epoch of the first observed primary eclipse.

More detailed descriptions of these parameters are given in Chapter 2 of Hilditch (2001).

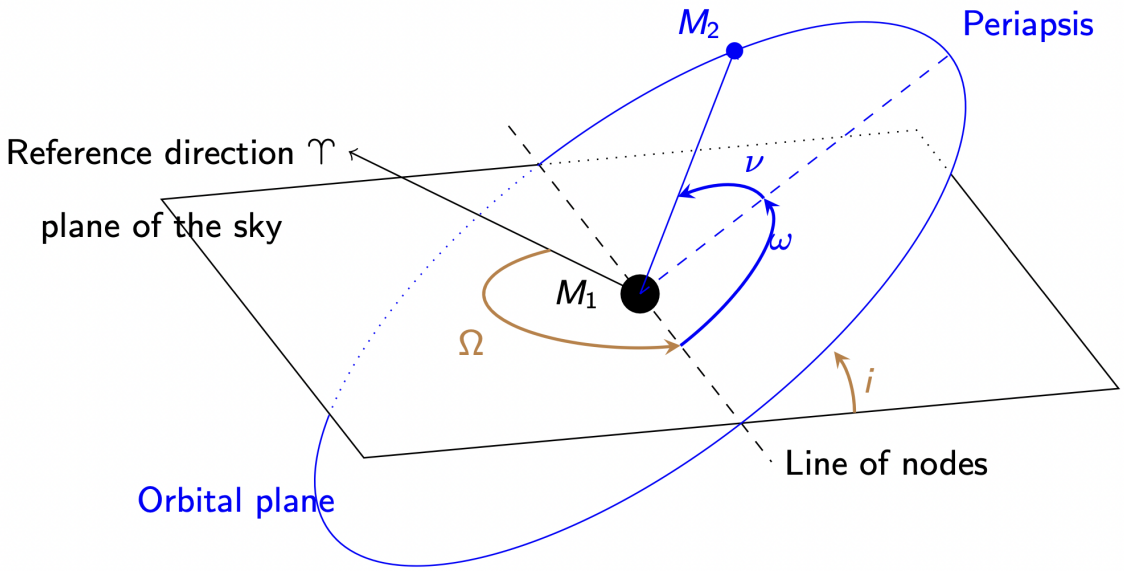

Figure 1.1 Illustration of a secondary star (M_2) orbiting a primary star (M_1) that is fixed in the reference frame, with the angular parameters of the orbit labeled. The primary reference point for the orbit is the vector direction of the vernal equinox, represented by the Aries symbol (♈), in the plane of the sky, represented by the black outlined rectangle. The longitude of the ascending node is described by the angle Ω, which is measured from the reference vector to a point on the line of nodes at which M_2 is moving from below to above the plane of the sky (i.e., away from the observer). The angle ω orients the orbit with an angle that swings from the location of the line of nodes to the periapsis, which is the location of closest approach of M_2 to M_1. The true anomaly ν describes the current location of M_2 with respect to M_1 and the line of nodes. Finally, the inclination i of the orbit is defined to be 90° when the orbit is oriented edge-on to the observer; in this schematic, the inclination is $\sim 30°$.

For any pair of stars in a binary or triple, we will also discuss their *mass ratio* (q), which will always be the ratio of the mass of the secondary or tertiary star to that of the primary star, e.g., $q = M_2/M_1$ or M_3/M_1. In some instances, we will refer to the *mass fraction* (f), which quantifies the amount of mass in one component relative to the total mass in the system, e.g., $f = M_1/(M_1 + M_2)$ or $M_2/(M_1 + M_2)$ or $M_1/(M_1 + M_2 + M_3)$, etc.

1.2 How To Form a System of Multiple Stars

In each multiple system of two (or more) objects, the components are moving along paths that are the cumulative result of how those objects formed and how they have subsequently evolved dynamically. Observing the formation processes from start to finish is not an option because they take longer than the modern astronomical record, but we have gleaned some understanding by matching physics-heavy models to observational snapshots of large sets of systems. In this section we review our current understanding of the different formation paths for systems of low-mass stars and brown dwarfs, as well as the subsequent dynamical evolution that could proceed in each of those scenarios.

In general, stars form from the collapse of localized regions (typically about a parsec in size) within molecular clouds that are relatively more dense than surrounding regions in the clouds. These overdense regions range from spherical to filamentary in shape, and are the natural result of the cloud's balance of turbulence, gas molecule kinetics, gravitational pressure, and magnetic fields. These overdensities are unstable, and prone to collapse when they experience a perturbative event, such as the collision of two molecular clouds, when a cloud passes through galaxy's spiral arm, or when a shock event occurs, such as a supernova. As gas within a localized region falls toward a central point, it quickly forms a disk with a dense, spherical center as angular momentum is conserved. If the central sphere grows enough, the pressure at the center triggers hydrogen to begin fusing into helium in the core, and a star is born. The star's contraction slows to a stop as hydrostatic equilibrium is established, and its leftover material spins around it in the form of a disk that slowly thins

in the breeze of the stellar winds.

Two or more stars may end up gravitationally bound at several points during the above formation process, or even shortly afterward. As detailed in the following subsections, current models group the formation of multiples into three methods: first, one molecular cloud breaking into multiple stellar cores through turbulence (§1.2.1); second, the disk around a stellar core fragmenting into additional cores (§1.2.2); third, fully-formed stars capturing neighboring stars (§1.2.3). Additional dynamical evolution may follow these events due to local gas mechanics (§1.3). Many of these models were constrained by observations of multi-star systems and trends of multiplicity with stellar mass, which are discussed in §1.4.

Figure 5 in Offner et al. (2022) summarizes these scenarios, including their typical length scales and timescales, and is reproduced here as Figure 1.2.

1.2.1 Turbulent Fragmentation of Cores and Filaments

The earliest point at which a star might gain a companion is just as it begins to form, during the collapse of local gas in its molecular cloud. The balance of turbulence vs. rotation keeps the cloud stable, and this turbulence naturally leads to the overdense regions discussed above to be shaped as filamentary structures as well as spherical cores. Each of these structures (and the spectrum of morphologies in between) could form multiples if the conditions are right during their gravitational collapse. In both cases the fragmentation of the structure into collapsing sub-structures is triggered by existing density anomalies within the core or filament exceeding a critical threshold that depends on the region's size, temperature, mass, and — for filaments in particular — magnetic field pressure.

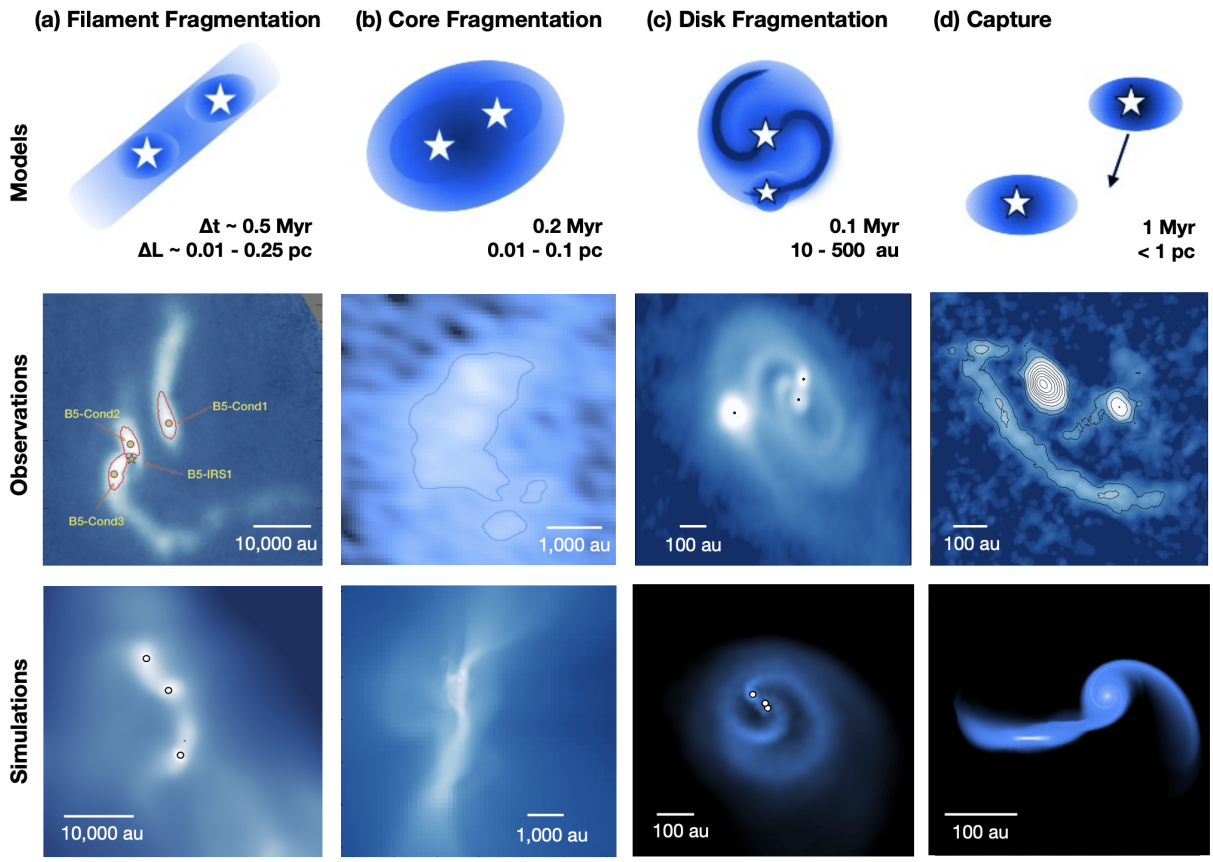

Figure 1.2 Summary of formation scenarios, reproduced here from Offner et al. (2022) Figure 5. *Top row:* schematic of the main formation scenarios discussed here, including typical time- and length scales. *Middle row:* Observational results that are potential examples for each scenario (from Pineda et al. 2015; Kirk et al. 2017; Reynolds et al. 2021; Rodriguez et al. 2018). *Bottom row:* Results from numerical simulations showing examples of each scenario (from Guszejnov et al. 2021; Offner et al. 2016; Bate 2018; Muñoz et al. 2015).

When filaments collapse and fragment they do so along their lengths, creating multiple cores. Such cores would thus be aligned like "beads on a string," and if they are within ∼0.1 pc of each other they may be gravitationally bound. Multiples that form from fragmentation of more spherical overdensities (molecular cloud cores) would be similarly separated but have no particular alignment with one another. Each overdensity collapse suppresses the

propagation and growth of additional perturbations in its vicinity as material is depleted, so each multiple formed is typically restricted to 2–3 members (Offner et al. 2016; Guszejnov et al. 2017).

Measurable attributes of multiple systems formed via turbulent fragmentation include:

Mass ratios — Very low-mass companions ($\lesssim$0.08 $M_\odot$) are expected to be uncommon because the collapsing and fragmenting overdensities are relatively massive and these processes occur while the local gas is plentiful (Fisher 2004), although this is subject to debate. Furthermore, when each protostar forms, its radiation suppresses further smaller-scale fragmentation within a few hundred AU (Offner et al. 2009; Bate 2012). Small mass ratios are thus unlikely.

Initial separations — The initial separations of stars formed in multiples via this method is $\gtrsim$100 AU, regardless of the specific physics that sets the minimum fragmentation scale (e.g., Guszejnov et al. 2017; Lee & Hennebelle 2018). If the stars are connected by gas "bridges" (as suggested by Kuffmeier et al. 2019) and actively accreting, they would migrate inward to $\lesssim$100 AU within a Myr through gas-dynamical friction (Bate & Bonnell 1997; Zhao & Li 2013; Lee et al. 2019).

Additional observational signatures — Due to their initially wide separations and isolated development, the spin axes and early circumstellar disks of the stars in each system should be distributed randomly (Offner et al. 2016; Bate 2018).

A complicating facet of the formation of multiples via turbulent fragmentation is that some properties of the systems may also reflect the initial turbulence of their molecular cloud,

thereby skewing trends of mass ratios vs. orbital periods for binaries (Fisher 2004). These dependencies are not straightforward to untangle, however, for more complicated physics involving magnetic fields.

1.2.2 Disks Broken via Gravitational Instability

As a single protostar is collapsing, it may also gain companions that coalesce from leftover material of its circumstellar disk. At its outer edge, the disk is receiving infalling gas from the collapsing molecular cloud, and at its inner edge its material is falling onto the newly forming star. This growing star heats the disk, providing thermal stabilization, thus the overall stability of the disk is dictated by the balance of accretion onto the star vs. accretion onto the outer disk. The disk will fragment if the cold outer region grows too quickly while its temperature remains constant, or if the central star's luminosity drops (e.g., if accretion slows) and the disk loses temperature. Regardless of the cause, disk fragmentation occurs at roughly the boundary between the protostar-heated region and accretion-fed region, or $\sim$100 AU.

The result is that the fragmented part of the disk quickly collapses to a single mass orbiting the original protostar. Initially the companion may be larger than its own Hill radius (the companion's radius of gravitational influence), which places it in danger of being sheared apart by the competing gravitational influence of the protostar and surrounding material. The companion does not survive unless it becomes small enough to be contained within the safety of its own Hill sphere, thus it must cool efficiently enough and become dense enough to collapse quickly. The companion may migrate inward during this process,

but doing so shrinks its Hill sphere even smaller, as the Hill radius is directly proportional to the semi-major axis of the companion's orbit. Thus, to avoid tidal disruption, the companion must migrate inward slowly enough that its shrinking Hill radius stays larger than its still-collapsing physical radius (Gammie 2001; Kratter & Murray-Clay 2011).

Measurable attributes of multiple systems formed via broken disks include:

Mass ratios — The companion formed from this scenario tends to be similar in mass to the initial protostar due to the conditions that led to the disk fragmentation in the first place — the bountiful infall of gas onto the outer disk. The new companion will also tend to outpace the first star's accretion rate from the disk, making an equal mass ratio even more likely.

Initial separations — The heat from the central star gives thermal support to the disk within tens to hundreds of AU (depending on the star). Thus, fragmentation occurs only in the outer region, and the resulting stars are separated by $\gtrsim$100 AU.

Additional observational signatures — Given that the new companion is born into the remnants of the leftover disk, the properties of the resulting binaries depend substantially on how that leftover disk affects the companion's dynamical evolution. How the companion's migration is affected by disk properties will determine the distributions of companion separations and mass ratios expected from this formation scenario. Gas-driven migration is discussed in detail in §1.3.

1.2.3 Capture of Neighboring Stars

Even if a star remains single through its entire formation, it could gain a bound companion if another star passes near enough to enter its gravitational well and loses so much energy that it cannot escape. There are two main categories of capture scenarios: gasless and gaseous, depending on whether gas is available to assist with the capture.

In a gasless scenario, enough energy can be lost to capture the new companion if the stars pass closely enough to generate tidal forces (Fabian et al. 1975), or if the addition is initially a third companion and the subsequent dynamical evolution ejects one of the stars (usually the least massive member; Valtonen & Mikkola 1991). Neither of these scenarios are likely enough to account for a significant fraction of field binaries (Bate 2015).

For gaseous capture, however, the odds of gaining a companion are improved, as long as there is enough gas in circumstellar disks that two closely passing stars lose energy to that material as they move through it (a process termed "gas-mediated capture"). This mechanism requires close stellar passes and stellar velocities that are relatively slow to avoid simply truncating the disks, making it more likely in smaller, dense clusters (Clarke & Pringle 1991a,b). Because this scenario requires substantial disk material, it is only an option during a star's initial formation phase.

Measurable attributes of multiple systems formed via capture include:

Mass ratios — Dynamical gasless capture leaves systems with mass ratios closer to unity, as the ejected companion tends to be the least massive of the three. Capture via gaseous disks, however, produces a distribution of systems with mass ratios that dramatically rise

toward lower values (McDonald & Clarke 1995), primarily because lower mass stars are more easily captured than massive stars.

Initial separations — Gasless capture usually leaves the resulting binary with a similar separation as the original pair, as ejecting the least massive companion does not dramatically change the total angular momentum and energy (Kratter 2011). Gaseous capture, on the other hand, results in separations necessarily closer than the initial size of the circumstellar disk, which observations suggest are tens of AU in radius for low-mass stars (Pascucci et al. 2016; Burn et al. 2021). Migration to wider orbits (as discussed in §1.3) is not expected in this case because the disk is likely to be truncated (Cuello et al. 2023) or destroyed by the capture, with the latter scenario significantly more likely for higher-mass companions (Olczak et al. 2006).

Additional observational signatures — The capture of another star would not leave an orbit with any particular eccentricity or orientation, although the stars' random initial trajectories make circular orbits less likely unless there is subsequent gas-driven evolution (§1.3). In general, then, there are no observational signatures that allow us to distinguish between multiples formed via gaseous capture and fragmentation of a disk (§1.2.2), or between gasless capture and turbulent core/filament fragmentation (§1.2.1). The only smoking gun would be circumstellar disks in the system that are on clearly colliding paths, especially if their morphologies indicate they collided with each other previously. However, the lifetimes of those disks are relatively short compared to the average age of field stars, so this line of evidence has presumably been erased in samples of multiples found anywhere other than

young clusters.

1.3 Orbital Evolution After a Multiple is Formed

As discussed above (§1.2), models indicate that most multiples form with relatively wide separations (tens to thousands of AU) and no particular orbital eccentricity. The large numbers of closely-separated multiples we observe today at separations $\lesssim$10 AU are thus the result of subsequent orbital evolution. Models currently suggest that that evolution is mainly the result of (1) interactions between the stars and the gas and dust in their environment — the circumstellar disks as well as any surrounding molecular cloud — and in some cases also (2) tidal forces.

1.3.1 Orbital Evolution Due to Gas and Dust

If there is substantial gas left in the surrounding molecular cloud, interactions between the stars and that gas creates torque, which in turn reduces the stars' separations via a process known as "gas-dynamical friction." These interactions include accretion and the effects of the gas distribution that forms a wake trailing behind each star. Angular momentum may also be lost by magnetic braking as the stars move through the gas (Zhao & Li 2013). In less than a Myr, these combined processes can drive stars from thousands of AU apart to tens of AU (Offner et al. 2022). Thus, it appears that substantial gas — either in the leftover molecular cloud or in circumstellar disks — *is required* to create the distributions of companion separations for various types of main sequence stars.

If there are disks around both stars in a binary and both are still accreting from the

surrounding cloud, the balance of the torques from these processes may drive the stars closer together or further apart. Early simulations showed that a single circumbinary disk will consistently induce *only* inward migration (Lubow & Artymowicz 1996). The options changed when star formation models included circumstellar disks and accretion torque, although the specific results depend substantially on the initial parameters of a binary pair, such as its mass ratio, disk morphologies and densities, and orbital eccentricity. For typical disks, pairs with initially low mass ratios will generally migrate inward, whereupon the companion will accrete mass from the disk and drive the mass ratio closer to unity. As the mass ratio increases, the migration slows and potentially reverses direction (Tokovinin & Moe 2020), resulting in a distribution of separations wherein not all companions are at very small distances from their primaries.

1.3.2 Orbital Evolution Due to Kozai-Lidov Cycles

It has been suggested that pairs at separations $\lesssim$10 AU could also be the result of Kozai-Lidov cycles (Kozai 1962; Lidov 1962) in combination with tidal friction (Fabrycky & Tremaine 2007). This mechanism can shrink an orbit if there is a third, wide companion in the system, as angular momentum exchanged between the outer and inner orbits induces oscillations in eccentricity and relative inclination. During the period when the inner pair is eccentric, tidal forces during periastron passage circularize its orbit and generally reduce its semi-major axis.

This process, however, is generally not rapid enough to explain all current $\lesssim$10 AU orbits. Kozai-Lidov oscillations occur on timescales of a Gyr, whereas most binaries evolved to their current $\lesssim$10 AU semi-major axes on Myr timescales, as described by Moe & Kratter (2018).

Gas-driven migration is more efficient and thus more likely.

1.3.3 Orbital Evolution Due to Tides

Under the right circumstances, one of the most efficient mechanisms of orbital evolution is loss of angular momentum to the tidal force arising between two or more bodies. The tidal force is induced when one object passes near to a second object and the gravitational force on its near side is greater than the force on its far side. This differential force deforms the object along that force axis, creating a tidal bulge that lags ahead or behind the force axis (depending on the rotation direction) as the object continues to rotate, inducing a torque. This torque adjusts the rotation of the object until the tidal bulge is aligned with the force axis, synchronizing rotation with orbital motion, and in the process the orbit is circularized.

The torque due to the tidal force is inversely proportional to a^8 (Zahn 1977), where a is the separation between the objects, and the subsequent rate of change of a is proportional to $1/a^5$. Consequently, the torque is ineffective unless the objects are close enough together to mitigate the a^8 or a^5 factor, at which point its proportionalities to the objects' masses, radii, and interior structure constants become significant. The tidal torque then works efficiently such that the orbit is synchronized and circularized on Myr timescales (typically within the pre-main sequence phase), set by the two objects' initial separation, eccentricity, and mass ratio (Zahn & Bouchet 1989).

The tidal dependence on stellar properties implies that for a population of identical binaries of a specific age, every orbit with a semi-major axis smaller than some specific threshold will be completely circularized. That threshold would then move outward as that

population ages and tides less efficiently move the wider binaries closer together; Mathieu & Mazeh (1988) even proposed using this mechanism to age-date binaries. In observational studies, that semi-major axis threshold is discussed in terms of the equivalent orbital period, denoted $P_{\rm circ}$, below which all orbits are circularized. The $P_{\rm circ}$ threshold has been observed to range from $\sim$2 days for high-mass binaries to $\sim$12 days for solar-type binaries (§1.4.1 vs. §1.4.2). The larger radii of high-mass stars improves the efficiency of tides, increasing the $P_{\rm circ}$ at which they are effective, but the stellar lifetimes of the high-mass stars are so short that circularization can only occur for the closest pairs before evolution removes one or both components from the tidal process. In the same vein, Meibom & Mathieu (2005) showed that $P_{\rm trans}$ (analogous to $P_{\rm circ}$, as discussed below) for G-type stars in clusters increases with increasing population age, from $\sim$7 days (Pleiades) to 10 days (Solar neighborhood) to 16 days (halo).

One complication in the discussion of $P_{\rm circ}$ is that the timescale for circularization of a particular binary depends strongly on its mass ratio; for example, Mazeh (2008) pointed out that the mass ratio can change the circularization timescale by a factor of 8, corresponding to a factor of 1.5 in $P_{\rm circ}$. Mass ratio can vary significantly even among relatively uniform samples like "solar-type stars," thus if we have a rich data set we are unlikely to find a $P_{\rm circ}$ for which zero orbits above it are circular (and zero orbits below it are eccentric). Most likely this ambiguity is why techniques to determine $P_{\rm circ}$ for an observed population vary substantially, adding difficulty to comparing results across studies. To address this issue, Meibom & Mathieu (2005) proposed a "transition period" $P_{\rm trans}$ replace the concept of $P_{\rm circ}$:

below $P_{\rm trans}$, the *average* eccentricity of the set is zero, and above it the *average* eccentricity rises exponentially. This approach may be important if the results of this thesis show signs of a transition period rather than $P_{\rm circ}$ cutoff.

1.4 Trends Observed for Stellar Multiples and their Properties

Among main-sequence field stars, observations have established that multiplicity increases with stellar mass. At the lowest masses, ∼30% of M dwarfs are multiples, as found intially in a sample of 27 M dwarfs within 5 pc in (Henry & McCarthy 1990), to the much larger sample of more than 1000 M dwarfs over the entire sky in (Winters et al. 2019). In contrast, at the highest masses of the O and B stars, >90% are multiples, as shown by Mason et al. (1998, 2009) and later (Moe & Di Stefano 2017). Solar-type stars have multiplicity rates of ∼50% (e.g., Raghavan et al. 2010), falling between the rates for the highest and lowest mass stars. This overall trend of multiplicity with mass is summarized in Figure 1 in Offner et al. (2022), reproduced here in Figure 1.3.

Whereas the multiplicity fraction is well-determined for stars along the main sequence, trends among the *properties* of these multiples are not yet evenly established for all stellar masses or spectral types. As discussed in §1.2, the distributions of mass ratios, semi-major axes, orbital periods, and eccentricities are our strongest observational clues to the how these systems formed and evolved dynamically. In this section we review what observations have shown to date about how stellar multiples' properties compare between systems of different masses.

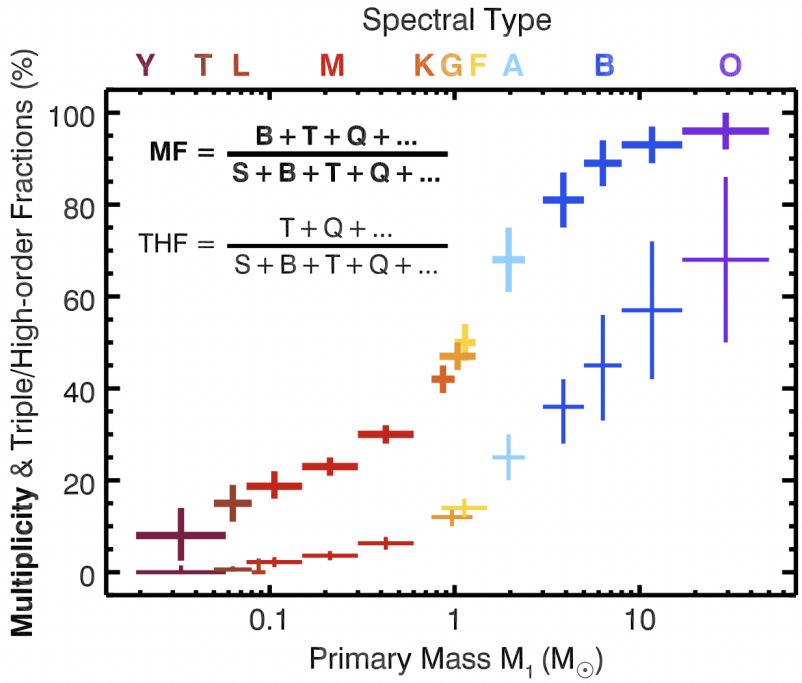

Figure 1.3 Multiplicity (or "multiplicity fraction") of stars as a function of spectral type, reproduced here from Figure 1 of Offner et al. (2022) using a wide accumulation of stellar surveys. The thick lines represent multiplicity of all types, while the thin lines reflect only the fractions of triples/higher-order multiples ("THF;" not discussed here). The letters along the top indicate the approximate spectral type of each stellar mass.

1.4.1 Properties of High-mass Multiples

Multiples among the O, B, and A stars (main sequence stars with masses of at least 1.5 $M_\odot$) have been studied extensively through spectroscopy, long-baseline interferometry, adaptive optics, and eclipsing binaries. These techniques are generally sensitive to companions $\gtrsim$10% the mass of their primary star ($q \gtrsim 0.1$). For the highest-mass stars in this group, the O stars, binary and dynamical evolution quickly and strongly alters the multiples' properties (Hoogerwerf et al. 2001), so multiplicity studies often avoid targets that have undergone evolution by considering only O stars that are members of young star clusters.

Earlier notable surveying of the highest-mass stars was done by Mason et al. (1998), then updated in Mason et al. (2009), who established these systems' remarkable multiplicity and used those results to infer their significant degree of dynamical evolution. Nearly a decade later, Moe & Di Stefano (2017) presented an even more comprehensive survey of these stars by combining data from several earlier studies (e.g., Abt et al. 1990; Shatsky & Tokovinin 2002; Rizzuto et al. 2013). They showed that separations between companions and primary stars are distributed uniformly in $\log a$, known as "Öpik's law" (Öpik 1924). All three spectral types in this group also follow a similar trend for companion mass ratios, keeping in mind the *caveat* that these surveys miss the lowest mass companions: at close separations ($\lesssim$0.5–1 AU) the mass ratio distribution is uniform, at moderate separations (1–100 AU) it peaks at $q \approx 0.3$, and at wide separations ($\gtrsim$100 AU) mass ratios are skewed toward small mass ratios (Offner et al. 2022, and references therein). On top of these trends, however, these systems have an excess of "twin" companions ($q \gtrsim 0.95$) that is strong at the closest separations and small but still significant at moderate separations.

The eccentricities of high-mass multiples in Moe & Di Stefano (2017) generally increased as $P_{\rm orb}$ increased, with this rise somewhat steeper for more massive primary stars ($\gtrsim$5 $M_{\odot}$). We will show this distribution and discuss it in detail in Chapter 6. For systems with $P_{\rm orb}$ $\lesssim$ 15 days this $P_{\rm orb}$ vs. e trend is partly related to their ages: the systems with B and A primaries are significantly older than O stars, and thus have had more time to accumulate the effects of the tidal forces that are felt at those separations. Beyond $P_{\rm orb} \gtrsim$ 15 days, higher-eccentricity orbits for higher-mass systems can only be explained by these systems

being systematically more eccentric at formation. Moe & Di Stefano (2017) proposed two potential explanations:

- The more massive stars have a higher frequency of triples and higher-order multiples, which more efficiently drive their eccentricities upward during the first ∼1 Myr of the pre-main sequence phase.

- All O, B, and A systems are born with high eccentricties, but the less massive stars spend more time contracting during pre-main sequence and thus more time under the influence of tides. This lets them accumulate more eccentricity damping than their O-star cousins. Similarly, circumstellar disks are longer-lived around less massive stars, adding another potential source of eccentricity damping.

The idea that evolution within disks is a factor is also well supported by the observation that twins are more common at tighter separations, as the accretion that comes with interactions between companion and disk drives a system's mass ratio toward unity (§1.3).

1.4.2 Properties of Solar-mass Multiples

The F, G, and K dwarf stars (main sequence stars with masses of 0.6–1.5 $M_{\odot}$, frequently grouped together as "solar-type") are established to have a multiplicity rate of ∼50% (Duquennoy & Mayor 1991; Raghavan et al. 2010).

The distribution of their separations is normal in logarithmic space ("log-normal"), peaking at ∼40 AU (Raghavan et al. 2010; Tokovinin 2014), although certainly increasing to smaller separations in linear space, with more companions found per AU within 10 AU than

at any other distance. The mass ratio distribution over all separations is roughly uniform, but at the closest separations the occurrence of "twin" companions ($q \gtrsim 0.95$) is ∼30% higher than the uniform distribution. This excess decreases but persists at intermediate separations, and even at wide separations there remains a notable twins excess (El-Badry et al. 2019). This trend of twin companions is consistent with these companions forming from disk fragmentation (§1.2.2) and migrating inward or outward (§1.3).

The most complete eccentricity distribution for solar-type multiples is presented by Raghavan et al. (2010), who completed a comprehensive multiplicity study of F, G, and K dwarfs within 25 pc, with companions revealed using multiple observational methods. We will show and discuss this distribution in detail in Chapter 6. To summarize, they showed that the eccentricities of solar-type multiples generally increase as orbital period increases, with long orbits ($P_{\rm orb} \gtrsim 100$ days) being rarely circular, and very short orbits ($P_{\rm orb} \lesssim 12$ days) exclusively circular for systems with no signs of youth. The latter structure is due to tidal forces, which are so efficient those smaller semi-major axes that they have circularized every close orbit in that field-age population (§1.3).

The upper envelope to the Raghavan et al. (2010) $P_{\rm orb}$ vs. e distribution is well below the limit at which two stars would collide at periastron, and is most likely set by early dynamical interactions with neighboring systems. They also noted that the highest-eccentricity systems at each $P_{\rm orb}$ are triples, thus these extreme eccentricities are the result of the triple-specific Kozai-Lidov cycles (§1.3).

1.4.3 Properties of Low-mass Multiples

The low-mass multiples that are the subject of this work include systems with primary star masses of 0.075–0.60 $M_{\odot}$ — these are main sequence dwarfs with M and L spectral types. Their companions have the same spectral types, with the except of the few percent that are the rare substellar brown dwarfs. Pre-2000 efforts targeting these stars struggled with their intrinsic faintness, limiting the current availability of historical and long-term data coverage. As newer instruments made lower-mass stars accessible, efforts focused on them for exoplanet searches, intentionally avoiding targets with stellar companions in the process. The result is that stellar multiples among M and L dwarfs are somewhat less well-studied than their solar-mass neighbors, particularly for the orbits of those multiples.[1]

Low-mass multiples are a particularly critical population to study, however, because they are ubiquitous (Henry et al. 2006, 2018) and encompass a wide range of stellar physics that presents persistent challenges to interior models (e.g., Dieterich et al. 2021; Brandner et al. 2023). Their interiors may be partially or fully convective, depending on their mass (Chabrier & Baraffe 1997; Jao et al. 2018), and these stellar structures are partly to blame for the breakdown of the standard stellar rotation vs. age relations at M dwarf masses (Newton et al. 2016; Douglas et al. 2017; Pass et al. 2022). At the same time, these stars have grown in popularity as targets of searches for habitable exoplanets (e.g., Irwin et al. 2009; Law et al. 2011; Muirhead et al. 2018; Ribas et al. 2023). This combination of persistent challenges

[1]Most surveys targeting low-mass stars are sensitive to brown dwarf companions as well, and it is challenging to exclude them entirely from the results. Because the brightness of the lowest-mass stars drops steeply with mass below $\sim$0.1 $M_{\odot}$ (e.g., Benedict et al. 2016), and because brown dwarfs lose brightness over time as they cool, it is generally not possible to determine whether or not a very low-mass companion is above or below 0.075 $M_{\odot}$ without precisely measuring its dynamical mass.

and the demand for exoplanet host star characterization escalates our need to understand how these stars form and how they affect their immediate environments on Gyr timescales.

Winters et al. (2019) completed the most comprehensive survey of nearby M dwarf multiplicity to date by combining several previous studies with their own imaging observations. They found that 27% of M dwarfs are in systems of multiples, with this fraction smaller for less massive vs. more massive M dwarf primaries. Their separations are log-normally distributed around 20 AU, somewhat closer than the 40 AU peak of solar-type binaries (Raghavan et al. 2010). Limiting the sample to primary stars $\lesssim$0.15 $M_{\odot}$ produced an even closer peak at $\sim$4 AU. As with the solar-type multiples, in linear separation space there are more companions orbiting M dwarfs per AU at the smallest separations than at wider realms. Because comprehensive surveys for the closest companions are still underway (e.g., with radial velocity searches), the peak of the separation distribution will most likely be revised as further inward in future work.

The mass ratios of Winters et al. (2019) were uniformly distributed for M dwarfs $\gtrsim$0.3 $M_{\odot}$ and skewed toward higher q for lower-mass M dwarfs. This trend is not due to formation or evolution, but rather the result of the cutoff in mass at the end of the main sequence, as that exclusion leaves the lowest-mass M dwarfs without the possibility of companions of $q \lesssim 0.5$. After taking this effect into account, no excess of twins is evident in the M dwarf mass ratios, although the distributions are not broken down by separation. This contrasts with the twins excesses seen in different separation regimes around solar-type and massive multiples (§1.4.2 and §1.4.1), potentially indicating a very different formation or dynamical

history for these lower-mass systems.

The Winters et al. (2019) multiples did not have orbits, and thus no eccentricities determined, but a preliminary eccentricity distribution was presented by Udry et al. (2000) using 48 binaries characterized via radial velocities. Their $P_{\rm orb}$ vs. e plot was similar to the solar-type distribution, with $P_{\rm orb} \gtrsim$11 days orbits exclusively eccentric and orbits with $P_{\rm orb} \lesssim$10 days tidally circularized. For the lack of circular orbits they suggest a range of explanations, e.g., circular orbits may be unlikely to form, or star-disk interactions may have increased the systems' eccentricities (modeled by Lubow & Artymowicz 2000). These scenarios are difficult to corroborate further without information about the systems' mass ratios.

This picture of low-mass multiples was supplemented at the very end of the main sequence by Dupuy & Liu (2017), who presented 27 orbits for systems with total masses 0.2 $M_{\odot}$ and less, including brown dwarfs. They showed an eccentricity distribution that peaks strongly at $e \approx 0.1$, significantly lower than the uniform distribution for solar-type stars from Raghavan et al. (2010) and the preliminary M dwarf distribution of Udry et al. (2000). Underscoring that result, their $P_{\rm orb}$ vs. e plot shows $\sim$20% of their very low-mass systems with $P_{\rm orb} \gtrsim 5$ years have $e \lesssim 0.1$ — whereas these orbits were entirely absent among the solar-type systems.

Although the Dupuy & Liu (2017) survey design did not allow them to draw conclusions on the distribution of mass ratios or prevalence of twin binaries, their eccentricity results indicate a clear difference in the formation and dynamical evolution of the lowest-mass stars and brown dwarfs vs. the solar-type stars. The potential explanations mirror those of the analogous situation with the high-mass stars (§1.4.1): either the two classes of objects have

different distributions of orbit shapes from birth, or the distributions are initially similar but evolve differently due to different eccentricity-damping environments. Considering that observational evidence shows circumstellar disks persist for much longer around M dwarfs vs. higher-mass stars such as G dwarfs (e.g., Mamajek 2009; Ribas et al. 2015, and references therein), the latter explanation is perhaps more likely.

Extrapolating the Dupuy & Liu (2017) results to the slightly more massive M dwarfs is hazardous due to the prevalence of brown dwarfs in their sample, which may form from the similar processes as low-mass stars but with significant deviations that affect their final statistics (Whitworth et al. 2007; Ma & Ge 2014). Differences in post-formation evolution of stars vs. brown dwarfs are also inevitable because a star's luminosity is higher and sustained, whereas a brown dwarf starts fainter and diminishes in output forever. This will have significant impacts on the temperatures, viscosities, and ultimately the lifetimes of the disks in those young systems, as photoevaporation is the mechanism that ultimately dissipates the disks (Clarke et al. 2001).

Given our relatively poor understanding of M dwarf orbital architectures, we have undertaken this study of the companions to low-mass stars. Our goals are to establish the distribution of $P_{\rm orb}$ vs. e for these systems and identify any trends (or absence of trends in this distribution with respect to primary stars, secondary stars, and mass ratios. By accumulating many of these orbits and carefully distinguishing the stellar companions from the brown dwarfs, we can explore the extent to which those populations are similar to or different from the outcomes from formation and evolution for more massive stars.

CHAPTER 2

Defining the Sample of Nearby M Dwarfs

To understand how well our results represent the true population of M dwarf multiples, we must carefully define the sample of objects that we will consider in this group. This task requires defining the observational attributes of M dwarf stars, then identifying the list of stars to be targeted and monitored for detecting orbits.

In fact, our samples for orbit monitoring extend slightly beyond the end of the M spectral sequence and into the early L types because we adopt an astrophysical cutoff — the line between stars and brown dwarfs — rather than lettering in a spectral sequence. RECONS member Serge Dieterich determined this cutoff using the minimum in the radius sequence at spectral type L2.0V (Dieterich et al. 2014), and this is the line we will adopt as the end of the stellar main sequence. Thus, our sample is really the "red dwarfs" rather than limited to "M dwarfs," although hereafter we refer to the sample as M dwarfs.

Also note that in this work, we are studying systems in which the primary is a main sequence M dwarf, although we include a few cool subdwarfs and young stars that are of spectral type M.

Ideally, we would define the M dwarf by mass, as the M dwarfs include all stars 62% the mass of the Sun or less (Benedict et al. 2016). But mass is not a directly observable quantity or able to be determined model-independently for most systems, hence we need to use other properties that strongly depend on mass. In this thesis we use absolute magnitude as a proxy for luminosity and thus mass, and if necessary we also consider colors from several

combinations of photometric bandpasses.

In this chapter we first discuss in §2.1 the stellar distances from which each star's absolute magnitude will be derived. The absolute magnitude limits defining an M dwarf are determined in §2.2, which involves a reconciliation of mass-luminosity relations (§2.2.1) and then an extension of those relations to additional bandpasses (§2.2.2). Finally, in §2.3 we review the sample resulting from these definitions. This broad sample will be the set from which we will draw smaller samples for each observing program.

2.1 Setting Distance Limits

In order to derive the absolute magnitudes corresponding to M dwarf masses, we must establish reliable distances to an ensemble of *bona fide* M dwarfs. Although there are several methods to determine distances to astronomical objects, for the most nearby stars the most accurate method is trigonometric parallax.

Trigonometric parallax is a purely geometric technique, requiring no assumptions about the properties of the star of interest. The observer measures the angle between the Earth, nearby star of interest, and the Sun — known as the *parallax angle* — and the distance to the star is proportional to the inverse of that angle following the small-angle approximation. If the angle is expressed in arcseconds and the Earth-Sun distance is 1 AU, then the constant of proportionality is 1.0, giving the formula the convenient form $d = 1/\pi$ (where d is the distance in parsecs (pc) and π is the trigonometric parallax angle in arcseconds).

The strength of this method is its simplicity, but the catch is that the largest parallax

angles in our neighborhood (for the most nearby stars) are tenths to hundredths of an arcsecond in size, which itself is only 1/3600 of a degree on the sky. These tiny angles made parallaxes impossible to measure until the 1830s, when they were measured independently by Friedrich Bessel, Wilhelm Struve, and then Thomas Henderson, after several centuries of astronomers' dedicated efforts toward that goal (Perryman 2012, and references therein). Measuring parallaxes for stars beyond a few tens of parsecs required late-20$^{\rm th}$ century techniques and instrumentation (in particular, charge-coupled devices or CCDs), and achieving precision better than $\sim$5% for stars beyond 100 parsecs was not possible until the space-based astrometric mission *Hipparcos* and, more recently, the space-based *Gaia* mission (Gaia Collaboration et al. 2016). Now that we have moved beyond the instrumental difficulties of the 1800s and have entered an era of highly accurate and precise astrometry, the straightforward geometry underpinning parallax makes it the ideal technique for measuring stellar distances, especially for stars in our immediate neighborhood.

In every sample in this work for which we have set a distance limit, each target is required to have a trigonometric parallax placing its distance within that stated limit. All samples considered here are larger than $0\farcs040$, or 40 milliarcseconds (mas), corresponding to a distance of 25 pc.[1] The two primary sources of parallaxes used here are the RECONS program and the *Gaia* mission due to these efforts' reliability and precision.

The RECONS astrometry program began in 1999 on the 0.9m CTIO/SMARTS[2] telescope, and throughout its tenure it has focused almost exclusively on red dwarfs within 25

[1] A few systems considered here are slightly beyond 25 pc, but these are not used for statistical treatments.

[2] The Small and Moderate Aperture Research Telescope System (SMARTS) at Cerro Tololo Inter-American Observatory (CTIO) near La Serena, Chile.

pc. RECONS parallaxes are measured relative to reference grids of background stars, then corrected to absolute parallaxes using the distances to the reference stars. The program continues regular observations on a volume-complete sample of Southern Hemisphere red dwarfs with parallaxes of at least 60 mas (within 16.7 pc), a distance which provides ∼500 M dwarfs that can each be observed a few times per year. At the thesis sample definition stage in August 2018, RECONS had determined parallaxes for 482 M dwarfs, and a further 155 M dwarfs were expected to have enough data to measure their parallaxes by the thesis conclusion in 2023.

Gaia is a space telescope launched in 2013 for the purposes of measuring parallaxes and other astrometry for stars across the entire sky. Their Data Release 2 (DR2) in April 2018 gave parallaxes for 1.7 billion objects, including for ∼3000 red dwarfs within 25 pc, with a median precision of 0.08 mas. *Gaia* parallaxes are absolute, as they are linked to a reference grid of distant quasars having zero parallax. Comparison between RECONS and *Gaia* results in Vrijmoet et al. (2020) showed that the *Gaia* and RECONS parallaxes were reliably consistent for single stars to within ∼2.5 mas.

To be included in any sample discussed here, at least one of those two programs must have measured the target star's parallax to be within the stated distance limit. Other sources of parallax measurements (e.g., *Hipparcos*) are only considered if there is no RECONS or *Gaia* measurement for that object. A weighted average of the RECONS and *Gaia* parallaxes is not used because, at the time that the observational samples were established for this thesis, the available *Gaia* results showed that parallaxes for unresolved multiples tended to deviate

from the RECONS values by $\gtrsim$2.5 mas, often $\gtrsim$5 mas (Vrijmoet et al. 2020).

These two programs provide nearly all of the M dwarfs in the samples discussed here, with each target having a parallax measured that places it within the stated distance limit. Other sources of parallax measurements (e.g., *Hipparcos*) are only rarely considered if there is no RECONS or *Gaia* measurement for that object. A weighted average of the RECONS and *Gaia* parallaxes is not used because, at the time that we established the observational samples for this thesis, the available *Gaia* results showed that parallaxes for unresolved multiples tended to deviate from the RECONS values by $\gtrsim$2.5 mas, and occasionally were offset by $\gtrsim$5 mas (Vrijmoet et al. 2020).

2.2 Deriving Absolute Magnitude Limits for M Dwarfs

In this section we establish the absolute magnitudes corresponding to the most and least massive M dwarf stars. These photometric limits are then used to select a large sample of M dwarfs to monitor for orbital motion during the observing campaigns detailed in Chapters 3 and 4. In short, the procedure is to use an empirical mass-luminosity relation to estimate masses of a sample of M dwarfs, then use the photometry of those systems in other bandpasses to derive those bandpasses' absolute magnitude limits corresponding to specific masses.

First, we gathered an ensemble of 327 *bona fide* M and L dwarfs within 25 pc that have no known close or unresolved companions. Requiring that these sources be single, without companions, ensured that the photometric limits were representative of individual objects rather than unique combinations of two or more luminous bodies. Although most of these

sources were stars, brown dwarfs were not excluded at this stage because one goal of this exercise was to determine the line dividing stars from brown dwarfs.

Of the 327 objects, 163 were selected from systems on the RECONS astrometry program (discussed in detail in Chapter 3). Selecting from the astrometry program allowed us to exclude targets that showed evidence of an unresolved companion orbiting on multi-year to multi-decade scales. The median baseline of these 163 RECONS systems was 18.55 years, and the shortest-observed system had a baseline of 6.58 years. This list was supplemented by 164 single stars from the RECONS 25 pc Database, which also includes systems well-vetted for companions, bringing the total sample to 327. To ensure that this final *bona fide* single-star sample would be suitable for deriving absolute magnitude cutoffs and would be free from any remaining unseen companions, we imposed the following requirements on all stars:[3]:

- *Gaia* DR2 RUWE $\leq$ 1.4. Reduced unit weight error, or RUWE, is a measure of the deviation between a star's *Gaia* DR2 astrometry and the single-star model fit to that data (Lindegren et al. 2018). For stars with no orbital motion over *Gaia* DR2's 2.4 years of observations, RUWE $\approx$1.

- no flags indicative of youth in the RECONS 25 pc Database or astrometry program

- measured V magnitude available from RECONS or another similar source (Reid et al. 2003; Weis 1991, 1996)

[3] *Gaia* DR2 results were not used in creating this initial sample because stars were not as well-vetted for companions as for the RECONS astrometry program and RECONS 25 pc Database.

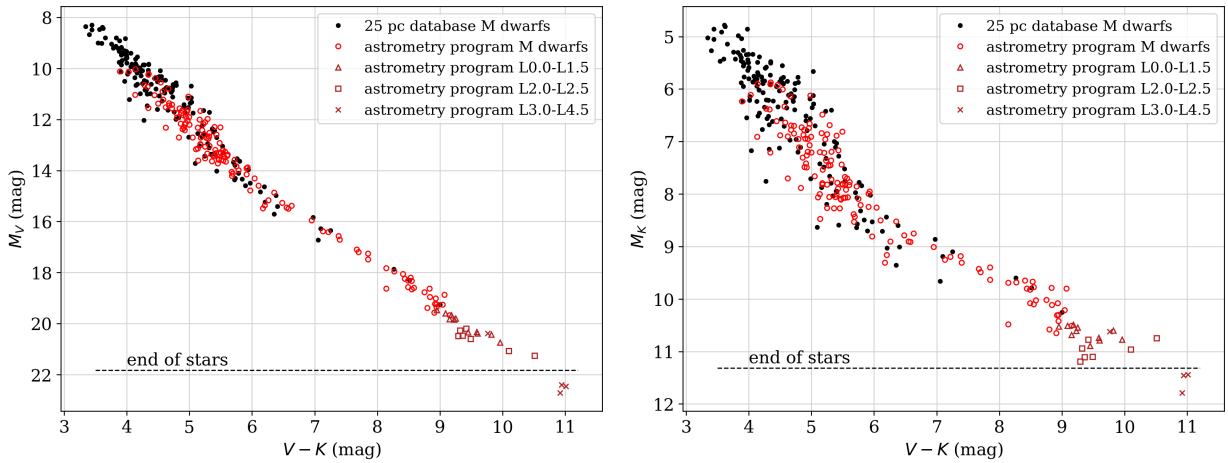

Figure 2.1 Hertzsprung-Russell Diagrams in $V-K$ vs. M_V (*left*) and M_K (*right*), showing stars within 25 pc that have no evidence of multiplicity. All stars have spectral types confirmed using a consistent technique (see text), distances from *Gaia* DR2, *Gaia* DR3 RUWE < 1.4, and no evidence of a perturbation in RECONS astrometry (if RECONS data are available). Light red symbols are M dwarfs from the RECONS astrometry program, those in dark red are L dwarfs from the RECONS astrometry program, and those in black are sources from the RECONS 25 pc Database that are not on the RECONS astrometry program. Dashed horizontal lines indicate the end of the main sequence as defined for this work (§2.2.2).

- measured K magnitude available from 2MASS (Skrutskie et al. 2006; Cutri et al. 2003)

- confirmed spectroscopically to be an M or L dwarf in systematic spectroscopic surveys of low mass dwarfs, including RECONS efforts, Reid et al. (1995), Hawley et al. (1996), and Gray et al. (2003, 2006)

The full sample of 327 *bona fide* single M and L dwarfs is shown in the Hertzsprung-Russell diagrams (HRDs) of Figure 2.1. Red and dark red symbols indicate sources chosen from the RECONS astrometry program, and black symbols are from the RECONS 25 pc database. Additional lines shown in that Figure are discussed later in §2.2.1.

2.2.1 One MLR to Rule Them All

To pinpoint the absolute magnitudes corresponding to the highest and lowest mass M dwarfs, we use the mass-luminosity relations (MLRs) of Benedict et al. (2016), which provides conversions between mass and V and K absolute magnitudes specifically for M dwarf stars. These empirical MLRs are based on dynamical masses of 46 M dwarfs, making them the most comprehensive relations in a multi-decade series of incrementally improving M dwarf MLRs (Popper 1980; Henry & McCarthy 1993; Henry et al. 1999; Delfosse et al. 2000). The similarly precise MLR of Mann et al. (2019) was not used for this sample because it offers a relation only in K band, and because it is based on mass sums rather than individual dynamical masses.

Because Benedict et al. (2016) provides two MLRs (V and K band), estimating the mass of any one target requires choosing between one of these two relations. This choice is not necessarily straightforward: in K band the M dwarfs show less photometric scatter from stellar activity, resulting in a tighter relation between dynamical mass and M_K than for M_V. At V, however, the M dwarfs span twice as many absolute magnitudes, allowing any M_V measurement to provide relatively more leverage in placing an object at a given mass compared to an M_K measurement.

The sample selection discussed below in §2.3 required a single relation to be applied to all targets. To determine the best choice, we estimated the mass of each target in the *bona fide* single M and L dwarf sample using their M_V magnitudes, then we estimated their masses again using their M_K magnitudes. The comparison of those values is shown in the top panel

of Figure 2.2. Targets with magnitudes outside the stated limits of the Benedict et al. (2016) MLR (i.e., with $M_V > 19$ or $M_K > 10$) were excluded from this part of the analysis, which eliminated all L dwarfs and reduced the sample to 289 stars. The comparison (Figure 2.2, top panel) showed significant scatter around the 1:1 line with masses from M_K often larger than masses from M_V. These trends did not depend on whether the target was drawn from the RECONS 25 pc Database or the astrometry program with its extra years of monitoring, hence undetected multiplicity was unlikely to play a role.

Next, we constructed the HRD of these points in terms of $V - K$ and M_V, and fit a 3$^{\rm rd}$-order polynomial to the main sequence using a standard least-squares algorithm (Python's `numpy.polyfit` package, default settings). This HRD is shown in Figure 2.2, middle left panel, with the polynomial fit in black, and the points colored based on their proximity to the main sequence fit — those within 0.3 magnitudes are in blue, and those outside 0.3 magnitudes are pink. The point color thus indicates whether each star is near the center of the main sequence or more scattered.

Finally, in the middle right panel of Figure 2.2 we again plotted mass from M_V vs. mass from M_K with the points colored according to the HRD position/scatter from the Figure 2.2 middle left panel. This second mass vs. mass plot shows that the scatter in mass estimate roughly tracks with scatter on the HRD. Below $\sim$0.25 $M_\odot$, the HRD-scattered stars show consistently higher mass from M_V than the HRD-tight stars; this may reflect the effects of metallicity, especially because for $M_V \gtrsim 13.5$ on the HRD the scattered points are always bluer than the main sequence fit.

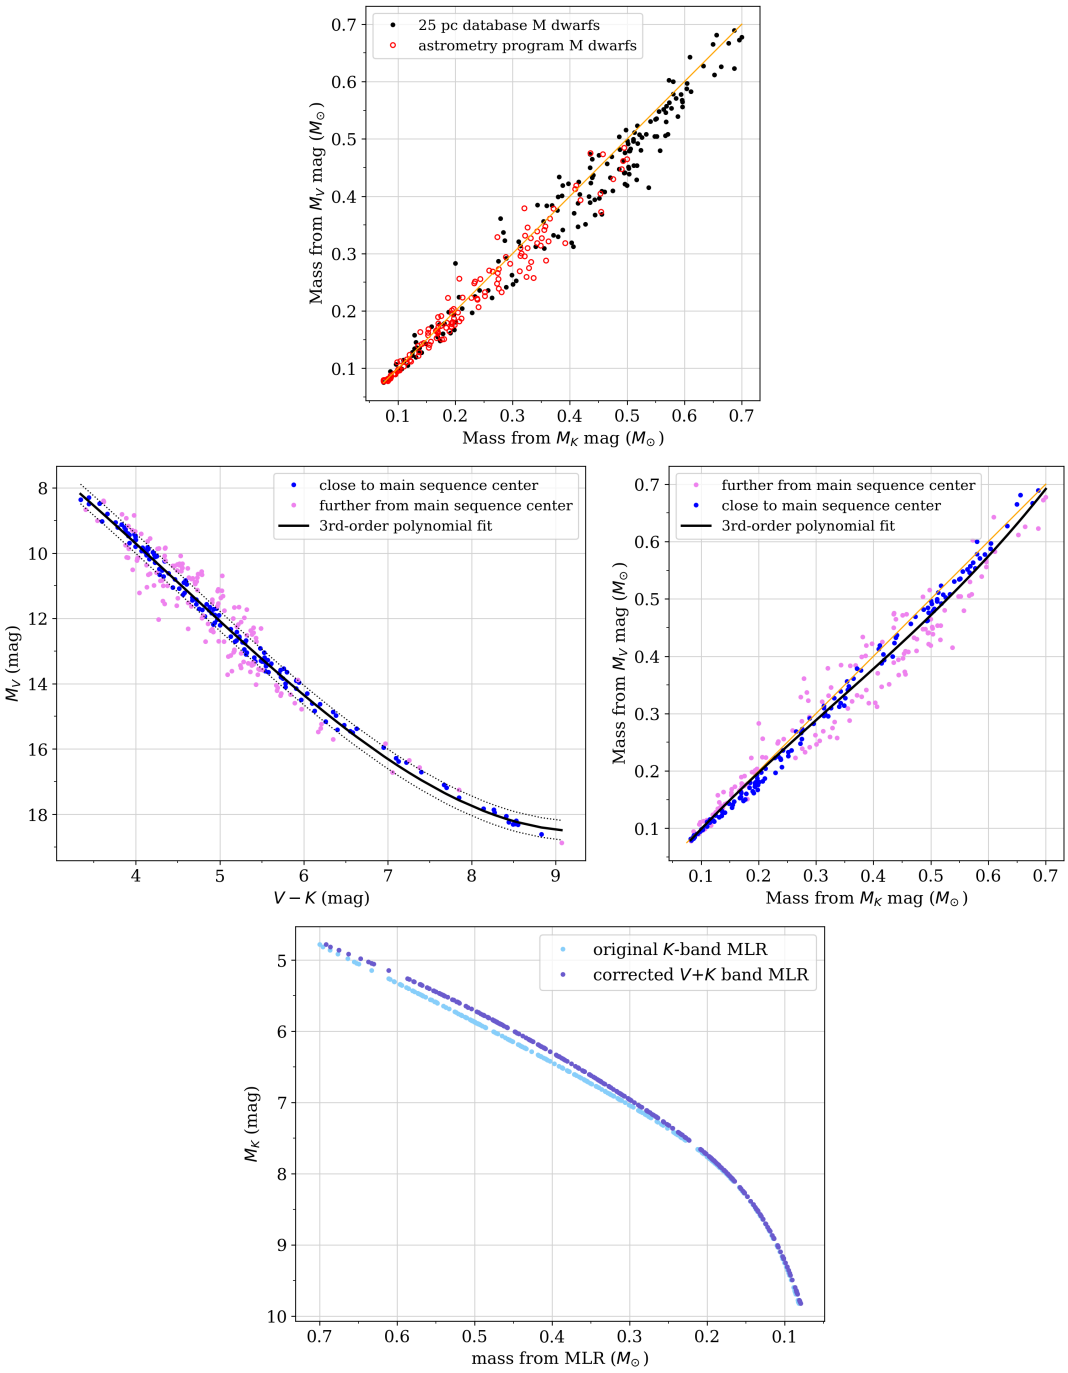

Figure 2.2 *Top panel:* Masses of 289 *bona fide* single M dwarfs within 25 pc, as determined from their M_V vs. as determined from M_K, both using the Benedict et al. (2016) MLRs. The symbol color indicates the source from which the targets were drawn. The orange line indicates the 1:1 relation. *Middle left panel:* The HRD of the same targets shown at left, with a 3$^{\rm rd}$-order polynomial fit marked in black. The symbol colors indicate deviation from that fit: stars within 0.3 mag of the fit are blue, and further stars are in pink. *Middle right panel:* The same mass vs. mass plot as shown at top, but with colors marked as in the middle left panel. A 3$^{\rm rd}$-order polynomial fit to the full set of points is marked in black, and the orange line again indicates the 1:1 relation. *Bottom panel:* The original Benedict et al. (2016) MLR (light blue) in K-band and the corrected V+K-band MLR (dark blue).

The magnitude of the scatter in Figure 2.2 indicates that it would be hazardous to use either the V-band or K-band MLR alone; the mass estimate for any single system could vary by up to 0.1 $M_{\odot}$ depending on the bandpass used, and the direction of this variation is not predictable. To compromise between these relations, we fit a 3$^{\rm rd}$-order polynomial to the full set of M_K mass vs. M_V mass values, including the scattered points as well as the tighter main sequence points. This fit is shown in Figure 2.2, middle right panel, as the thick black curve. The polynomial order was set at 3$^{\rm rd}$ to capture the broad trends without inducing any additional features. In testing, 4$^{\rm th}$-order (and greater) polynomials added curves that were not visible by eye in these data; thus, higher orders did not provide significantly better fits. Going forward, then, a system's mass can be estimated by first determining its mass from M_K using the MLR in Benedict et al. (2016), then choosing the corrected value of mass at that M_K from the polynomial fit of M_K mass vs. M_V mass.

2.2.2 Extending the MLR to New Bandpasses

Although the Benedict et al. (2016) MLRs provide mass estimates for M dwarfs using their M_V and M_K values, later in the sample construction we need to draw a deeper pool of targets from larger catalogs with measurements in other, different bandpasses. In this section we determine absolute magnitudes at the 0.62 $M_{\odot}$ highest mass and 0.075 $M_{\odot}$ lowest mass points for M dwarfs in the *Gaia* G, B_G, and R_G bands, as well as in the Johnson-Kron-Cousins I band. As discussed later in this section, a different process will be needed for the 0.62 $M_{\odot}$ limit vs. the 0.075 $M_{\odot}$; we thus begin with the 0.62 $M_{\odot}$ limit.

For stars in the *bona fide* single M dwarf sample bright enough to apply the Benedict

et al. (2016) MLRs ($M_V < 19$ mag and $M_K < 10$ mag), we estimated their masses using the V+K corrected MLR discussed in §2.2.1. We then collected their *Gaia* and I band photometry measurements, and Figure 2.3 shows these stars' absolute magnitudes in G, B_G, R_G, and I with respect to those mass estimates. For each of these bandpasses, we identified the nine stars with mass estimates within 10% of 0.62 $M_\odot$, and took the median of their absolute magnitudes. We then adopted the 0.62 $M_\odot$ absolute magnitude limits as the median magnitudes of "similar" stars in each bandpass. Table 2.1 lists the ranges of those similar stars' magnitudes in each band and the subsequent adopted absolute magnitudes for 0.62 $M_\odot$.

We applied the same procedure to the V and K bands as well, rather than taking the values offered by the Benedict et al. (2016) MLRs or the single V+K corrected MLR derived in §2.2.1. This approach ensured we were applying a consistent technique to derive each mass limit. The Benedict et al. (2016) MLR offers V- and K-band absolute magnitude limits marking the most and least massive M dwarfs — here we find that the resulting V-band absolute magnitude for 0.62 $M_\odot$ is nearly identical to the Benedict et al. (2016) MLR (8.99 mag instead of 9.03 mag) and the K-band absolute magnitude is 0.06 mag brighter (5.26 mag instead of 5.32 mag). The difference between the original and corrected relations rises to $\sim$0.25 mag between between 0.35 $M_\odot$ and 0.60 $M_\odot$ (Figure 2.2, bottom panel).

For the absolute magnitudes marking the end of the main sequence, we used the above process as a starting point, although not as the final result. Dieterich et al. (2014) established that the hydrogen-burning limit occurs around spectral type L2.5V. Comparing the M_V and

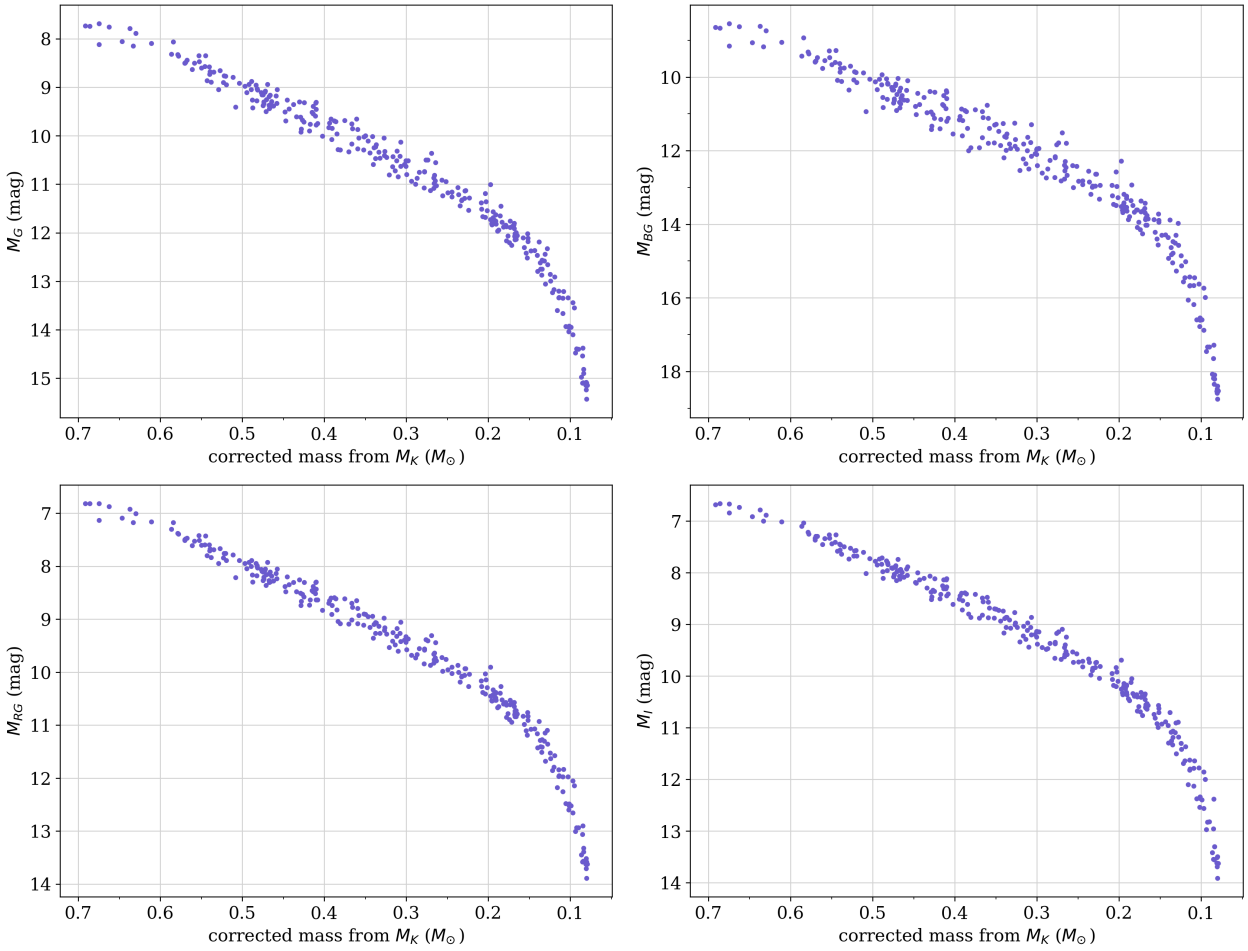

Figure 2.3 Estimated masses from the corrected V+K-band MLR (§2.2.1) for 289 single M dwarfs plotted against their absolute magnitudes in (*left to right panels*) G, B_G, R_G, and I bands.

M_K magnitudes of spectroscopically confirmed L dwarfs with the predictions for 0.075 $M_\odot$ by the Benedict et al. (2016) MLRs shows that the MLRs' predicted magnitudes of $M_V \sim 19$ and $M_K \sim 10$ are too bright for L2.5V. This is evident in the HRD of Figure 2.1, as the L2.0V–L2.5V dwarfs span $M_V = 20.0$–21.5, and the sources around $M_V = 19$ all have M spectral types. The discrepancy between the MLR value and the spectroscopy value adopted here is unsurprising given that there are few stars with dynamical masses near the end of

Bandpass	Abs. mags within 10% of 0.62 $M_\odot$	**Adopted** abs. mag at 0.62 $M_\odot$	Abs. mag fit at $M_V = 21.825$	Abs. mag fit at $M_K = 11.318$	**Adopted** abs. mag at 0.075 $M_\odot$
G	7.686–8.632	**8.23**	17.841	17.733	**17.79**
V	8.299–9.551	**8.99**	—	21.575	**21.82**
B_G	8.547–9.758	**9.24**	20.627	20.580	**20.60**
R_G	6.818–7.611	**7.24**	16.256	16.153	**16.20**
I	6.669–7.451	**7.07**	16.411	16.296	**16.35**
K	4.858–5.398	**5.26**	11.334	—	**11.32**

Table 2.1 Absolute magnitude ranges and key values used to determine the specific absolute magnitudes corresponding to 0.62 $M_\odot$ and 0.075 $M_\odot$, the highest and lowest masses adopted for M dwarfs. The final adopted magnitudes for those mass limits are highlighted in yellow.

the main sequence used in the MLR, coupled with the significant drop in flux for stars as masses decrease toward the end of the main sequence.

The L2.0V–L2.5V dwarfs significantly overlap with the L0.0V–L1.5V dwarfs, making it difficult to pinpoint precise absolute magnitudes corresponding to any one of those spectral types. In M_V as well as M_K there is, however, a significant gap between those sources and three of the four plotted L3.0–L4.5 dwarfs.[4] This main-sequence gap is the expected observational result for the hydrogen-burning limit because the brown dwarfs below that point should cool and lose luminosity quickly over time, moving down and to the right on the HRD. Further establishing the stellar/substellar line is beyond the scope of this dissertation, thus we set the dividing line as the midpoint between the last object above the gap and the first object below the gap: at $M_V = 21.82$ mag and $M_K = 11.32$ mag.

We also need to establish this stellar/substellar line in the G, B_G, R_G, and I bandpasses. Figure 2.4 shows that most of the HRDs in these bands show similar gaps as seen in M_V and

[4]Beyond type L2.5V, the V denoting main sequence stellar dwarf is not used.

M_K, but the width and clarity of the gaps vary. In particular, the B_G HRD shows a turn-back in color near the end of the main sequence instead of a gap, wherein progressively less massive objects become more blue starting around ∼L2.0V. This turn-back occurs primarily because there are features in the spectra of the lowest mass red dwarfs that cause $B_G - K$ colors to reach a maximum value, before lower mass objects appear bluer in that bandpass combination. In addition, *Gaia* B_G measurements become less reliable for sources fainter than B_G ∼19 mag, as illustrated in Figure 2.5. The variation in *Gaia* reliability with brightness means we need a more careful approach than the simple gap-midpoint applied above for V and K.

The stellar/substellar lines in the bandpasses must be consistent with each other — e.g., an object identified as stellar according to M_V should not be marked as substellar according to M_{BG}. With this in mind, we used the stellar/substellar lines already set for M_V and M_K to determine the lines for M_G, M_{BG}, M_{RG}, and M_I. Figures 2.6 and 2.7 shows each of these absolute magnitudes plotted against M_V (left column) and M_K (right column) for the *bona fide* single M and L dwarf sample, creating absolute magnitude-absolute magnitude relations for each bandpass. Dotted lines in the panels of Figures 2.6 and 2.7 indicate the stellar/substellar lines already established in M_V and M_K; the goal here is to establish the alternate absolute magnitude corresponding to that M_V or M_K limit.

In each bandpass X, the stars' M_X values become more scattered around the M_V and M_K values at the end of the main sequence, making it hazardous to choose the M_X limit using a simple median or mean. Instead, we selected the intrinsically faintest 33% of the *bona fide*

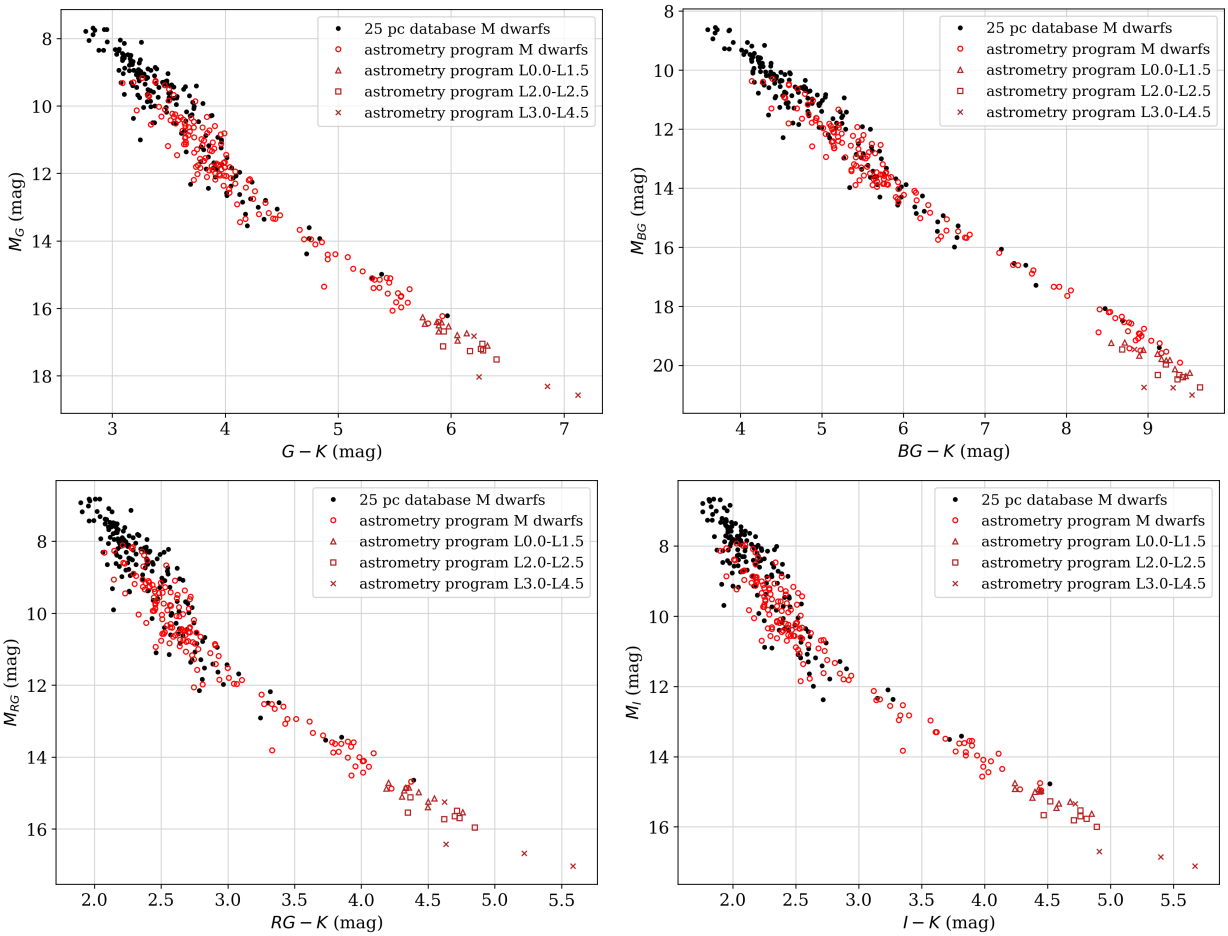

Figure 2.4 Hertzsprung-Russell Diagrams of the same set of nearby single M dwarf stars, shown in different filters with respect to K. From *top left* to *bottom right* are M_G vs. $G-K$, M_{BG} vs. B_G-K, M_{RG} vs. R_G-K, and M_I vs. $I-K$. The colors and symbols are the same as in Figure 2.1.

sample and fit a 2$^{\text{nd}}$-order polynomial to those points' magnitude-magnitude relation (M_X-M_V or M_X-M_K). The fits are shown in each panel of Figures 2.6 and 2.7 as orange lines. We did not fit the full sample because the goal was to replicate the behavior for only these faint stars, regardless of how the brighter 66% were related. The relations in those areas are thus given by the orange polynomial fits, and the intersections with the established M_V or M_K lines mark the M_X corresponding to those M_V or M_K. We adopted the stellar/substellar

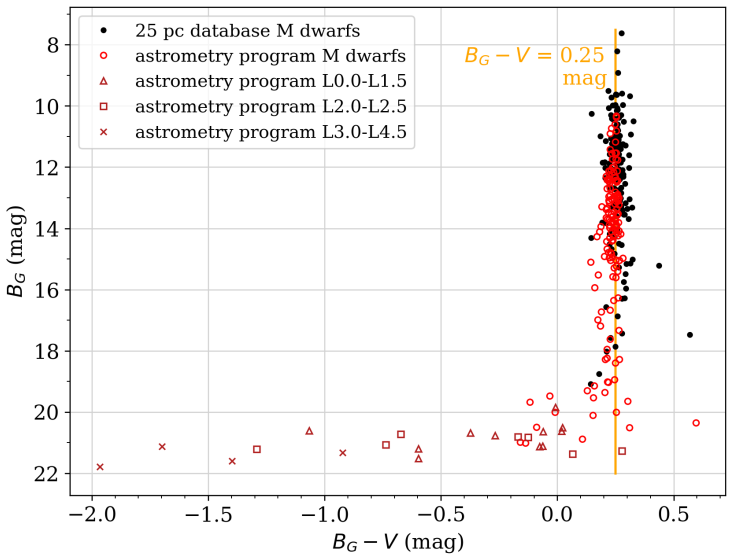

Figure 2.5 Differences between apparent B_G and V magnitudes as a function of B_G. The colors and symbols are the same as in Figure 2.1.

line for each bandpass X as the average of the two M_X-M_V and M_X-M_K intersections. The final values are listed in Table 2.1.

2.3 Verifying the Photometric M Dwarf Definitions

Applying the photometric limits in Table 2.1 to the *Gaia* DR2 catalog (requiring that each star meet the criteria in at least one bandpass) yields 5866 sources across the entire sky with parallaxes indicating they are within 25 pc. Because the goal of defining these limits was to establish samples for observing from facilities in the Southern Hemisphere, we drew only the 4565 sources in these limits below Decl. +25°. Finally, we omitted white dwarfs from the sample by excluding sources with $B_G - R_G < 1.80$, leaving 3415 potential M dwarfs.

Not every source in this extraction of the *Gaia* catalog corresponds to a true star, however, as *Gaia* DR2 contains many spurious entries due to source confusion and sources near the

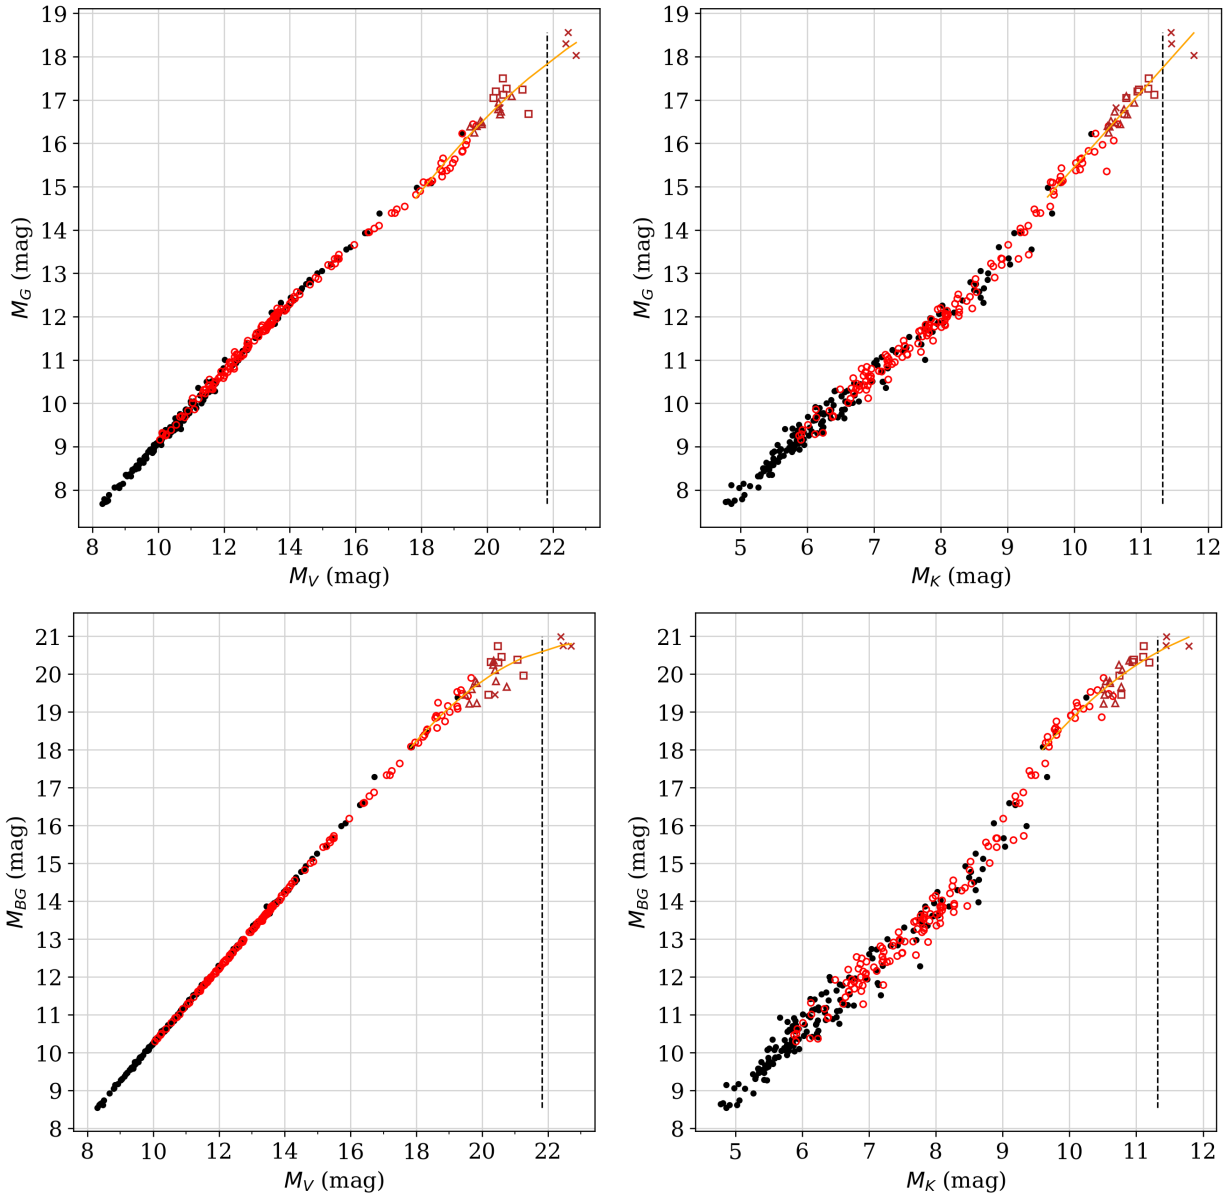

Figure 2.6 Absolute magnitudes in G and B_G against those in V and K for the *bona fide* single M dwarfs. The colors and symbols are the same as in Figure 2.1. Dashed black lines mark the adopted M_V and M_K values for the stellar/substellar lines. Orange curves are fits for each relation derived using the 33% faintest stars in the samples. The analogous plots for R_G and I are shown in Figure 2.7.

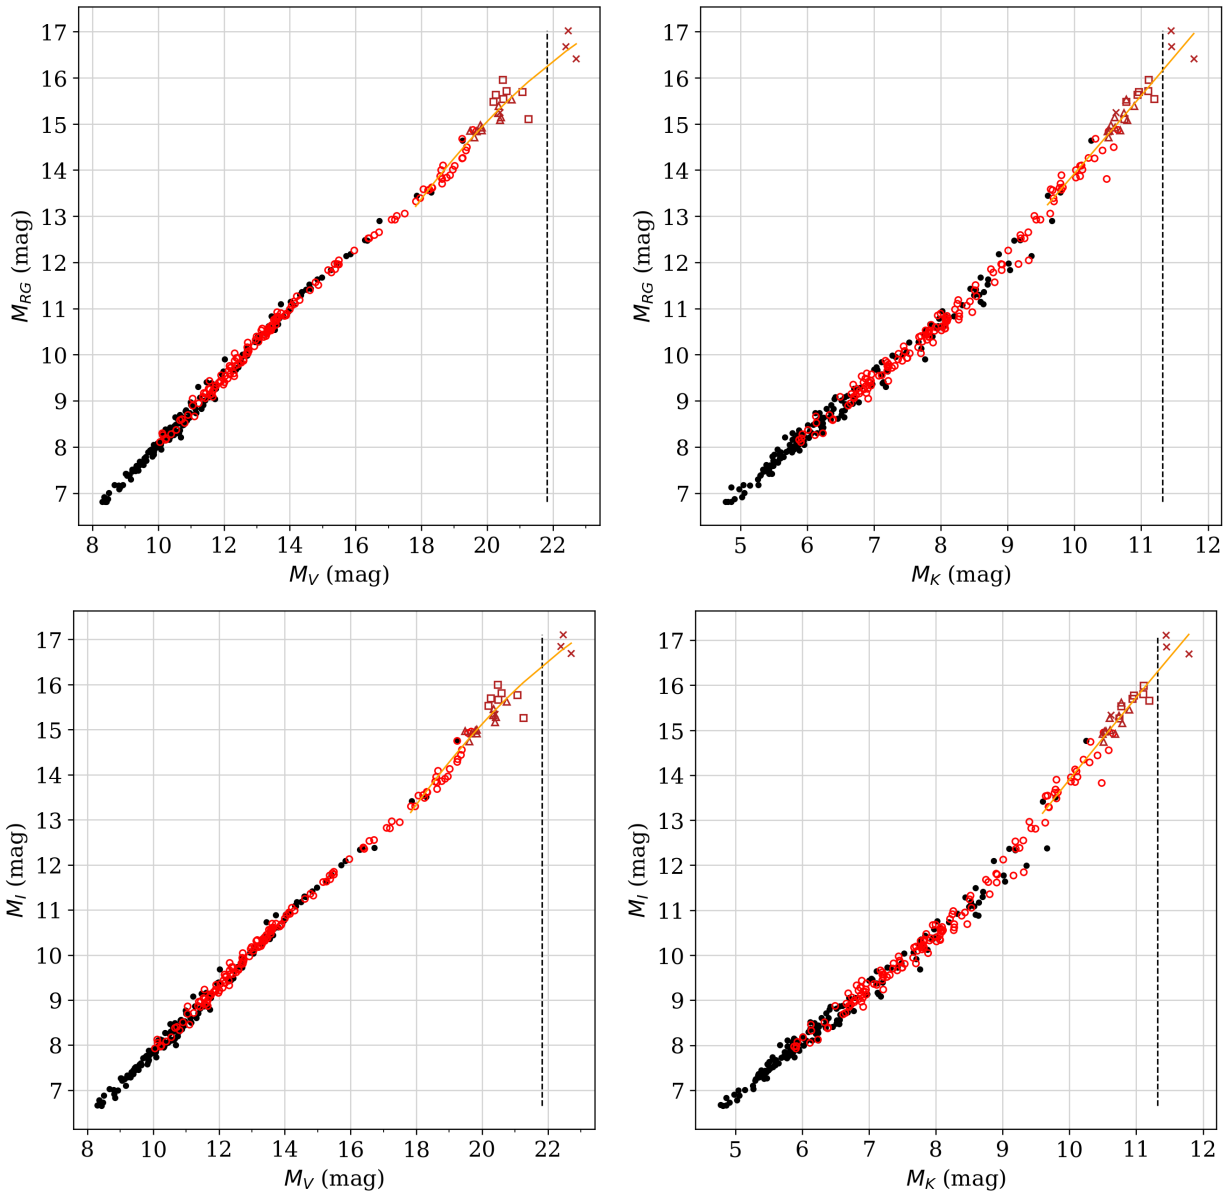

Figure 2.7 Absolute magnitudes in R_G and I against those in V and K for the *bona fide* single M dwarfs. The colors and symbols are the same as in Figure 2.1. Dashed black lines mark the adopted M_V and M_K values for the stellar/substellar lines. Orange curves are fits for each relation derived using the 33% faintest stars in the samples. The analogous plots for G B_G are shown in Figure 2.6.

detection limits of the instruments. Although there are several methods to "clean" the *Gaia* catalog in an automated way by removing entries with poor-quality data flags, such as the "ABC" cuts (Lindegren et al. 2018, Appendix C), those methods are dangerous to apply to our particular target sample. This is because we are gathering very nearby stars for which photocentric motions from orbiting companions could be large enough to look like poor quality astrometric fits to the data. In fact, many of the targets we desire for orbit mapping are binaries that would be culled using the ABC cuts. Additionally, some of our lowest-mass M dwarfs are too faint to have reliable *Gaia* magnitudes, which are generally limited to $G \lesssim 20$ and $B_G \lesssim 19$.

To clean the sample, we identified the sources in the 2MASS catalog (Skrutskie et al. 2006) that correspond to each *Gaia* source, as most M dwarfs are bright in the 2MASS K_s bandpass. This process required using each source's proper motion (as measured by *Gaia*) to convert its *Gaia* DR2 2015.5 coordinates to epoch 2000.0, then matching it to the nearest 2MASS source on the sky. For most of the *Gaia* sources, a corresponding 2MASS source was located less than $0\farcs05$ away. To verify the matches that were not so close, we checked each visually in Aladin using the *Gaia* catalog overlaid on the 2MASS images and catalog positions, supplemented by the CDS-composited Digitized Sky Survey II images (Lasker et al. 1996) for bluer or very high proper motion sources. This process yielded 2592 M dwarfs in 2461 systems within 25 pc.

Finally, we added 45 systems from the RECONS astrometry program that were not already represented in the above *Gaia* sample. Most of these systems were astrometric

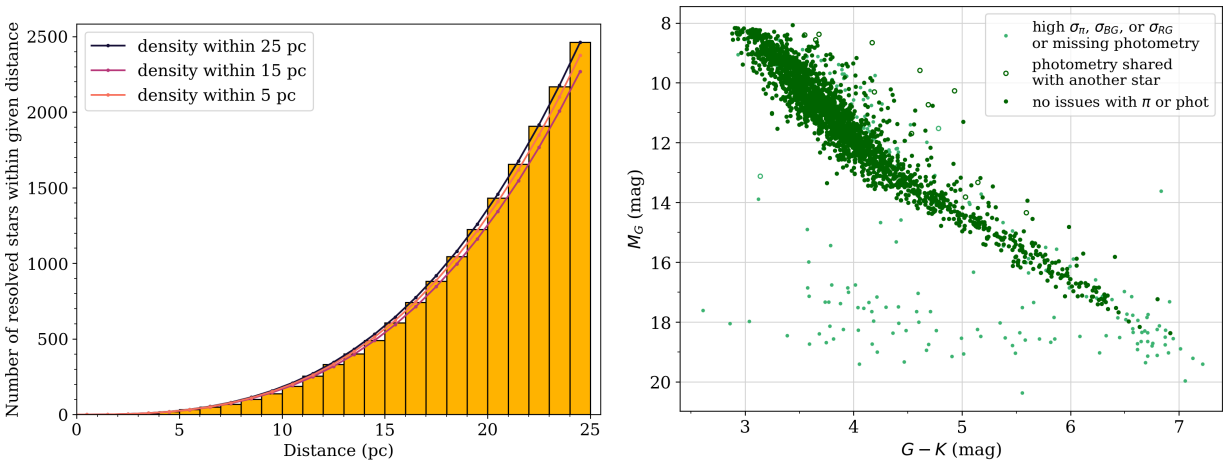

Figure 2.8 *Left panel:* Number of systems within a given distance in the 25 pc sample. Curves of constant density in black, dark red, and light red illustrate the number of systems expected given the average density within 25 pc, 15 pc, and 10 pc, respectively. *Right panel:* HRD of the 25 pc sample. Each point color roughly indicates the quality of that system's *Gaia* data (and thus the reliability of its HRD position): light points have parallax error $\geq$ 0.40 mas or are missing B_G or R_G, dark points have no such issues with their data, and open points (of either color) have more than one star included in their 2MASS K_s measurement (joint photometry).

multiples or had $\pi < 40$ mas in *Gaia* DR2 but $\pi \geq 40$ mas from RECONS data. The final set of 2637 sources in 2506 systems comprises an effectively volume-complete sample of M dwarf stars within 25 pc south of Decl. +25°, as demonstrated in Figure 2.8. The only systems that may be missing would be nearby red dwarfs with no entry in *Gaia* that remain unknown to the RECONS team, even after more than two decades targeting exactly these systems in the southern sky. From this sample of 2506 systems, we drew the target lists to add to the RECONS astrometry program (§3) and to begin the new SOAR speckle interferometry program (§4). The full list is given in Table 2.2.

Table 2.2: List of M dwarf systems within 25 pc drawn from *Gaia* DR2, as described in §2.3. These systems have not been vetted to distinguish single stars from multiples, but in cases where a system was noted as resolved in *Gaia* but not 2MASS, a "J" flag is listed in column 12. The full table of 2506 systems is available at `http://www.astro.gsu.edu/~vrijmoet/PhDthesis/sample_25pc.txt`

R.A. J2000.0 (1)	Decl. J2000.0 (2)	*Gaia* DR2 ID (3)	2MASS ID (4)	π (mas) (5)	μ_α (mas yr^{-1}) (6)	μ_δ (mas yr^{-1}) (7)	G (mag) (8)	B_G (mag) (9)	R_G (mag) (10)	K_s (mag) (11)	Joint phot? (12)
00 01 15.82	+06 59 35.50	2745860763717896448	00011579+0659355	42.75 ± 0.08	−436.74	−83.49	14.693	16.801	13.339	10.418	–
00 02 06.22	+01 15 36.14	2738415592529108096	00020623+0115360	48.05 ± 0.14	463.38	61.41	16.337	19.449	14.804	11.129	–
00 03 16.71	−55 16 29.10	4923590430409730048	00031670-5516292	42.07 ± 0.05	448.53	−235.67	13.702	15.403	12.450	9.849	–
00 04 36.44	−40 44 02.89	4996141155411983744	00043643-4044020	81.23 ± 0.11	677.73	−1505.12	11.499	13.103	10.307	7.737	–
00 04 41.47	−20 58 29.88	2340736324254735488	00044144-2058298	66.33 ± 0.16	758.23	85.20	16.869	20.484	15.249	11.396	–
00 04 57.55	−17 09 36.96	2414623952318068224	00045753-1709369	58.68 ± 0.12	145.14	−8.50	14.605	16.971	13.201	10.084	–
00 05 24.99	−50 02 52.98	4976609705736840960	00052498-5002529	41.35 ± 0.04	−155.91	−11.36	11.008	12.230	9.958	7.627	–
00 05 34.87	−06 07 06.87	2442074771933927424	00053484-0607070	40.21 ± 0.06	177.31	−55.30	11.979	13.395	10.822	8.411	–
00 06 19.19	−65 50 25.92	4899957905439144320	00061920-6550262	57.34 ± 0.07	197.07	−549.22	11.103	12.442	9.988	7.631	–
00 06 39.24	−07 05 35.93	2441755840548967424	00063925-0705354	50.87 ± 0.62	−100.95	118.25	13.102	14.971	11.768	8.958	–
00 06 43.20	−07 32 17.02	2441630500517079808	00064325-0732147	206.21 ± 0.13	−811.48	−1892.90	11.780	14.096	10.405	7.439	–
00 07 05.45	−56 05 04.62	4922569671300905088	00070543-5605045	40.08 ± 0.04	339.79	−7.32	12.315	13.754	11.150	8.710	–
00 08 17.38	−57 05 52.91	4919497979411495296	00081737-5705528	78.12 ± 0.06	−355.45	−44.45	10.935	12.390	9.793	7.395	–

CHAPTER 3

The RECONS Astrometry Contribution

In order to form a big-picture view of orbit sizes and shapes for the smallest stars, we need orbits on multi-decade scales to complement the days- and years-long orbits available in the literature. Accurate long-period orbits are difficult to accumulate, however, because to map a significant portion of these orbits requires more years than typically allocated to observing programs. To alleviate this problem we have turned to the astrometry program of the REsearch Consortium On Nearby Stars (RECONS), which has been taking data on the nearest M dwarfs at the CTIO/SMARTS 0.9m telescope since 1999. Having more than 20 years of data from a single observing program with consistent protocols gives us access to multi-decade orbits while simplifying the later task of analyzing observational biases affecting the results. For this thesis, we are focused on M dwarf primaries with companions orbiting in periods of 0–30 years.

3.1 Targets Observed by the RECONS Astrometry Program

The RECONS astrometry program is a broad survey that primarily targets a volume-complete sample of M dwarfs with parallaxes $\geq$ 60 mas (i.e., within $\sim$16.67 pc) and Decl. $\leq 0^\circ$ (the southern sky). The program continues to observe $\sim$300 low-mass targets outside these limits, including targets as far north as $+30^\circ$ Decl., M dwarfs of interest out to 25 pc, a few young M dwarfs at larger distances, and a few dozen white dwarfs and brown dwarfs. The "interesting" M dwarf additions are generally systems with orbital motion or with evidence of stellar activity and cycles, hence the multiplicity or photometric variability statistics

must be interpreted with care for RECONS systems beyond 16.67 pc. But, out to that horizon, 97% of the final sample systems described in Chapter 2 are currently targeted at the CTIO/SMARTS 0.9m. Only 32 stars are not being observed for the volume-complete survey — these are too bright (generally having $V < 10$ mag), have very poor reference star fields, and/or are corrupted by nearby sources.

Prior to 2019, the volume-complete RECONS survey was limited to M dwarfs with $5.0 < V - K < 12.0$, corresponding to stars with masses $\lesssim$0.3 $M_{\odot}$. As this thesis began in August 2018, the volume-complete survey was redefined to include more massive M dwarfs up to 0.62 $M_{\odot}$, set by M_V and M_K values that correspond roughly to $3.9 \leq V - K \leq 10.5$, to enable stronger comparisons between partially convective and fully convective stars. This broadening of the survey required defining an M dwarf in terms of absolute magnitudes, which is discussed in detail in Chapter 2.

Expanding the RECONS astrometry survey required adding ∼126 targets to the program's observing list for ongoing monitoring, bringing the total survey to 461 systems as of May 2023. The final member of this sample was set up for monitoring on August 31, 2021, as the CTIO/SMARTS 0.9m was closed for most of 2020 during the global pandemic. Considerations of the astrometric technique does prevent this sample from being truly 100% volume-complete. As mentioned above, systems brighter than ∼10 mag in V, R, and I usually cannot be monitored because the short exposure times for those targets leave the potential reference stars in those fields too faint to be used. Further details of this requirement are discussed in §3.2. This issue disproportionally affects the more massive M dwarfs

due to their intrinsic brightness.

The addition of 126 systems to the observing program for this thesis effort was not the only time the RECONS observing list was modified. Over the 20+ years of this program, targets have been added to or removed from the observing list as new nearby M dwarfs were discovered (or refuted) and as the active goals of the RECONS group evolved and shifted. The end result is that targets on the astrometry program have been observed over a range of timescales, as illustrated in Figure 3.1, which represents the 782 M dwarf systems that the program has ever monitored. Of these systems, 539 have enough data to be considered for this thesis, as an astrometric fit requires at least 2 years of data and 60 frames total.

For the purposes of this dissertation, our goal in observing these 539 M dwarfs is to characterize the orbits of the multiples within that sample. The varied distribution of time coverage (Figure 3.1) does affect the type of orbits we are sensitive to for each system. Despite this complication, our observations allow us to fit orbits to most systems with $P_{\rm orb}$ up to 20 years, and in some cases even longer (discussed further in §5.1.2).

3.2 Observing at the CTIO/SMARTS 0.9m for RECONS astrometry

Observations for the RECONS astrometry program are carried out at Cerro Tololo Inter-American Observatory (CTIO) in Chile at the CTIO/SMARTS 0.9m telescope. The program is discussed in detail in Jao et al. (2005) and Henry et al. (2018); here we give a brief summary of the observing procedures.

Observing runs are carried out 4–6 times per year, with runs lasting 10–20 nights each.

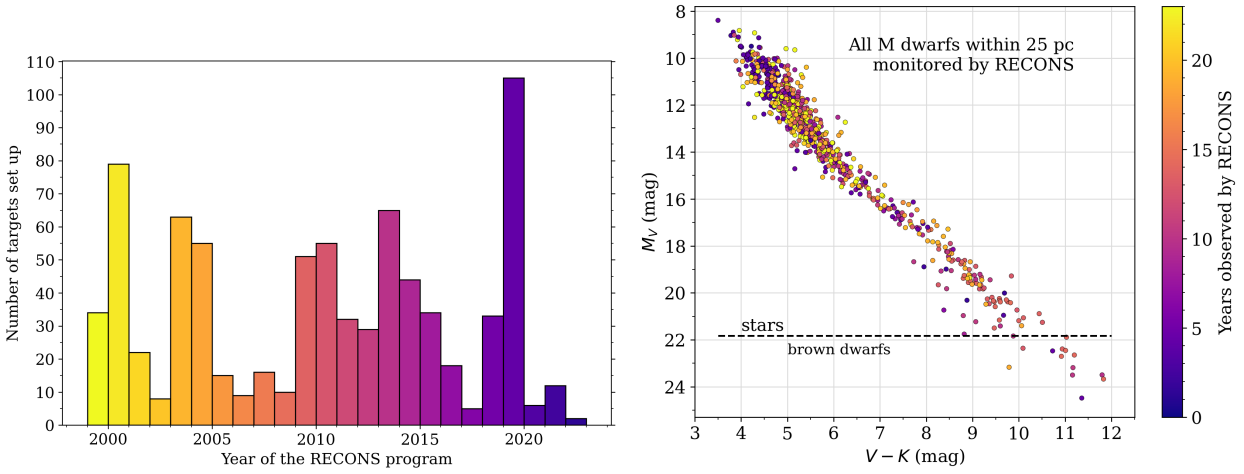

Figure 3.1 *Left panel:* Distribution of targets added to the RECONS astrometry program each year since it began in 1999, considering only the M dwarfs within 25 pc. Targets are added to (or subtracted from) the program to keep the core survey volume-complete as new nearby M dwarfs are detected, with additional modifications as the RECONS research priorities shifted over time. Nearly all systems added earlier than 2021 (539 total) have enough data to be monitored for orbits for this thesis. *Right panel:* HRD of the same sample of 25 pc M dwarfs, with points colored according to their observing baselines on the RECONS program. Brown dwarfs followed by RECONS are also shown here to emphasize our coverage, but we do not consider them for this thesis analysis; these systems are below the horizontal dashed "stars-brown dwarfs" line.

During a night, typically ∼20 targets are observed at airmasses less than 2.0. Throughout a run, 200–300 targets are typically observed at least once, with a visit consisting of 3–5 images spanning roughly 30 minutes. These images are routinely combined into nightly mean astrometric points because the astrometric positions of stars do not change markedly during a night. The end result is that most nearby M dwarfs receive 2–4 total observations per year.

The raw data are images taken on the Tek2k CCD, which has a full extent of 2048×2046 pixels and $0\farcs401$ per pixel plate scale. Only the central quarter of the chip is used for this program, so images are 1024×1024 pixels and the field size is $6.8' \times 6.8'$. The filters used

are V, R, or I on the Johnson-Kron-Cousins system, depending on the brightness of the target and its surrounding field stars. A single filter is used for all astrometric frames on a given target, and is selected to provide an exposure time giving ∼50,000 peak counts to the brightest star used in the field (nearly always the target star) in typical seeing of $1\farcs2$. A set of 5–29 stars surrounding the target star is used as reference stars during data reduction (§3.3), and ideally, at least a few have $\gtrsim$10,000 peak counts. Exposures are generally 30–300 seconds long, with some very faint targets requiring 600 or 900 seconds per image. Exposure times are adjusted on-the-fly during observations to ensure maximum counts during changing sky conditions, e.g., seeing fluctuations and clouds.

For each target, the filter used remains consistent through the entire span of its observations — for example, if we began observing a star in 1999 in V filter, every astrometric observation on that target since then has been carried out in V. This consistency is essential to obtaining reliable astrometry from the data, as it ensures the specific shifts in the field induced by the different colors of the stars as seen through a specific filter remain consistent from frame to frame and epoch to epoch.

Calibration images are taken either nightly or every two nights during the run, and consist of 17 bias frames and 11 dome flats in each of the three filters, for a total of 50 frames. After each observing run, the data are bias-subtracted and flat-fielded using a set of IRAF scripts, then organized into target-specific subfolders so a given target can have all its entire time series of data reduced as a single continuous data set.

3.3 Data Reduction for RECONS Astrometry

Every time a target's data are reduced, all suitable images — going back to the first image taken of that target for this program — are reduced together as one continuous data set, rather than processing the new data separately and adding it to the previous solution. This is essential to ensure that (1) the parallactic ellipse is as well-sampled by the data as possible, which improves the parallax fit; and (2) the time baseline of the data is as long as possible, which ensures the most accurate proper motion fit.

In general, the data reduction procedure consists of:

1. measuring the positions of the target and several (5–29) reference stars in each image,
2. using those positions to calibrate the target's position in R.A./Decl. space in those images, then
3. fitting those positions in time series to an astrometric model that includes proper motion and parallactic motion.
4. If any orbital motion is evident in the residuals of that first astrometric fit, we re-fit the data with a more comprehensive model that includes orbital motion in addition to proper motion and parallax (§3.4).

What follows in this section is a more detailed description of each of the above steps.

The first time a target's data are reduced, reference stars are chosen for that field. These stars are distant (beyond 300 pc) and collectively have small proper motions so they move minimally over time, which allows them to be used to calibrate the grid of R.A./Decl. in

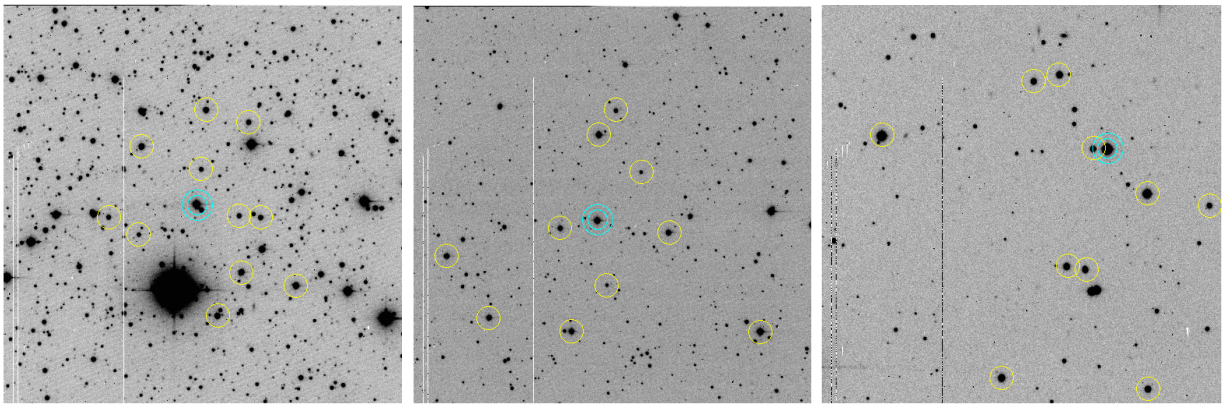

Figure 3.2 Examples of target star fields from the CTIO/SMARTS 0.9m for three targets: WT 2180 (*left*), LP 754-8 (*middle*), and LP 717-36 AB (*right*). Colors have been inverted such that darker pixels indicate more counts (i.e., brighter). In each panel, the target star is marked with a cyan double circle and the reference stars are marked with yellow single circles. The field in the left panel is excellent, with stars clustered around the target and near to it. The middle field is not as good because the reference stars are more dispersed, but the distribution is still reasonable — this field is typical of most targets. The right panel's field is poor, as the references are spread far across the image and are generally all to one side of the target star.

the image. The ideal reference stars are therefore located near to the target on the image (minimizing the area of the CCD to be calibrated) and distributed evenly in R.A. and Decl. (ensuring the CCD is well-sampled). Figure 3.2 gives examples of excellent, good, and poor reference star distributions.

The same set of reference stars is used for every image, with minor adjustments made for individual images if a telescope pointing error placed a reference star off the field or on a bad pixel on the CCD. If the reference set is revised, that new set is then used for all of the target's images, following the requirement that the all the old and new data for a given field be processed uniformly together.

First, each image is checked for poor seeing (worse than $2\farcs4$ FWM) or errors in telescope

tracking, both of which affect all stars of interest. In addition, checks are made for cosmic rays, bad pixels, or other anomalies that alter the PSFs of either the target star or any of the designated reference stars. If such an issue affects a reference star in a given image, then that star can later be excluded from that image's data, but if the issue affects the target star then the whole frame must be discarded. Every frame must be checked using this procedure, although if this target's data have been reduced before, then previous frames were already verified and only the new/additional frames need to be checked.

Next, the rough positions of the target star and reference stars are noted ("tagged") in each image. This is generally done via a script that asks the user to mark the stars in one representative image, then propagates these positions to the rest of the images, taking into account the proper motion of the target star supplied by the user. These tagged positions can be manually adjusted later by the user if necessary.

The stars' tagged positions are then used to determine the precise positions of these stars. SExtractor (Bertin & Arnouts 1996) fits a 2D Gaussian window function to each tagged star; this is an efficient technique that approximates fitting a 3D Gaussian to each star's PSF. This procedure outputs the precise positions of those stars' centroids in each image, as well as their instrumental magnitudes. At this stage, reference stars that had issues in some frames (often noted in the first step during the frame-checking) can be excluded from those frames to ensure their spurious data are not used.

Those stellar positions are used next to calibrate the field to the R.A./Decl. grid. The user selects a high-quality frame taken at minimal hour angle, in good seeing, and with

maximized signal-to-noise, to use as the "trail plate." The trail plate's stellar positions are then matched to the 2MASS catalog (Cutri et al. 2003; Skrutskie et al. 2006) to determine the rotation and scaling of that frame with respect to R.A./Decl. space. That solution is applied to the rest of the frames in the data set, so the measured positions now reflect those stars' motions due to proper motion and parallax only.

Next, the target and reference star positions in each frame are corrected for the effects of differential color refraction (DCR). This effect is a result of each star's light being refracted — by different amounts, based on the star's color — as it passes through Earth's atmosphere. The DCR correction was empirically determined for the RECONS program in each of the V, R, and I filters as described in Jao et al. (2005), and ultimately amounts to a small ($\sim$few mas) shift in the measured positions of the target and reference stars in each frame. Observations are typically made within one hour of the target's transit in order to minimize DCR effects. At this stage, we also use the precise observation times of each image to generate the parallax factors in R.A. and Decl., which indicate the Earth's position in its orbit during the observation. These factors must be taken into account when fitting the astrometric model to the data.

Finally, the GaussFit program (Jefferys et al. 1988) uses a least-squares algorithm to determine the plate constants that describe the scale, rotation, and stretch for each frame, as well as the relative proper motions of the reference stars and science target star. This solution assumes that the reference stars' proper motions and parallaxes sum to zero[1]. The

[1]Alternatively, the positions could be matched to the *Gaia* catalog, and *Gaia* parallaxes and proper motions could be used to account for the stars' motions during the observations. This substantial improvement is planned for a future upgrade to the RECONS data pipeline.

relative parallax determined here for the target star is corrected to absolute parallax using the photometric distances of the reference stars, which are derived using VRI photometry of the fields also acquired at the CTIO/SMARTS 0.9m.

We evaluate the results by computing the residuals between the parallax + proper motion model and the target star's position data over the duration of the data set. Some representative examples are shown in Figure 3.3. If any data point has an unusually high residual relative to the rest, we check that frame for common issues (cosmic rays or bad pixels on stars, bad background, misshapen PSFs). If any are found, we exclude those frames and re-compute the fit.

3.4 Fitting Orbits to RECONS Astrometry

When a target star has a bound stellar or brown dwarf companion, its orbital motion is often visible in the residuals to the astrometric fit described above (§3.3). Figure 3.4 shows two examples of such signals. These perturbations with respect to the astrometric model, or "PBs," are usually tracing orbital motion, but occasionally a false or ambiguous PB is induced by instrumental effects or when the target star passes very near a background star on the CCD over time. §4.1.2 revisits this topic of false/ambiguous PBs.

3.4.1 Principles of Photocentric Orbits

The motion evident in these residuals is not that of the primary star nor secondary star; rather, it is motion of the photocenter (center of light) of the two stars. The photocenter position with respect to the system's center of mass, p, is the same as that of the secondary

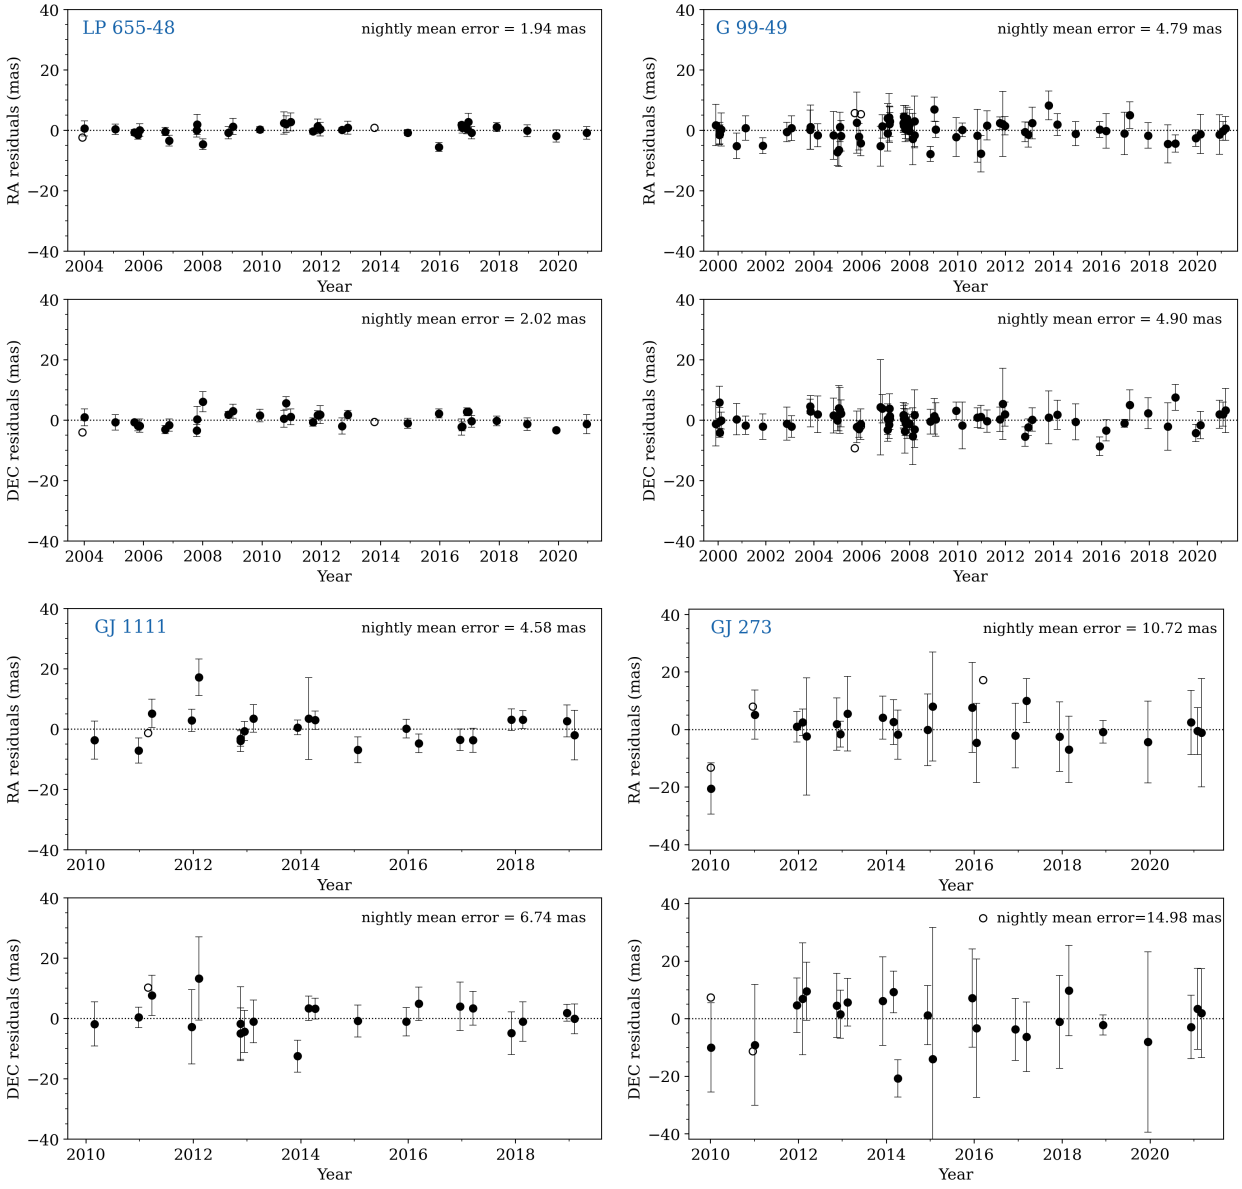

Figure 3.3 Examples of residuals to the proper motion and parallax fit for four target stars: (*upper left* to *lower right*): LP 655-48, G 99-49, GJ 1111, and GJ 273. Each system's residuals are represented by two panels: the upper showing residuals with respect to R.A. over time, and the lower showing residuals with respect to Decl. over time. The nightly (per-epoch) uncertainties of LP 655-48 (the first system) are unusually low (∼2 mas), and those of GJ 273 (the last system) are unusually high (10–15 mas), while G 99-49 ad GJ 1111 are typical.

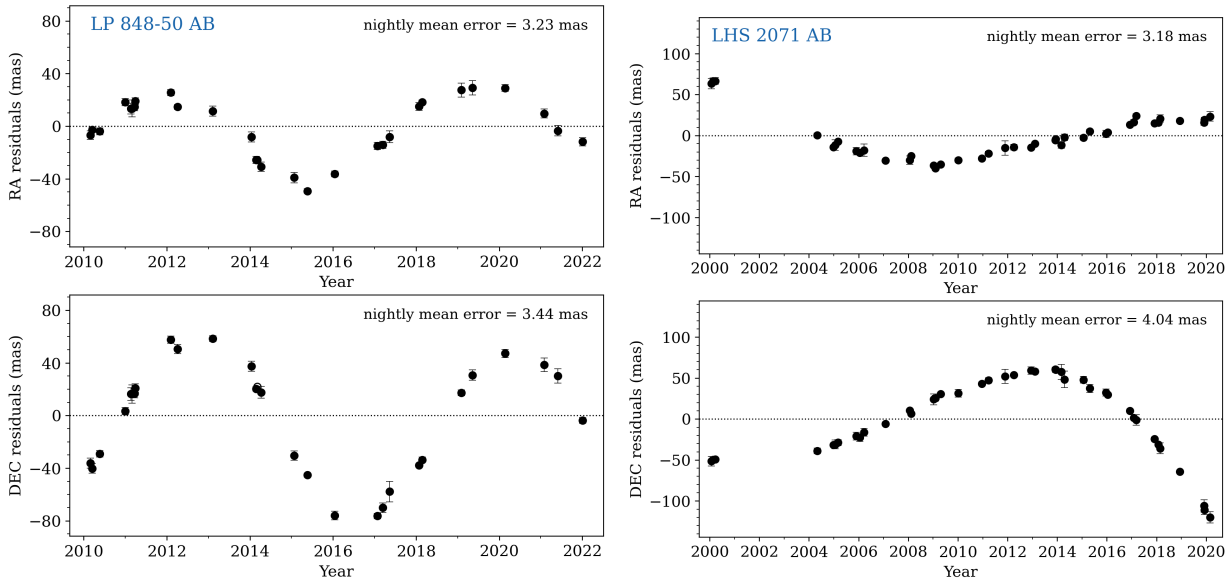

Figure 3.4 Two examples of RECONS astrometry perturbations (PBs) characteristic of orbital motion; the systems shown here are the binaries LP 848-50 AB (*left*) and LHS 2071 AB (*right*). These PBs are with respect to the "single-star" astrometric model, i.e., they are the residuals left after fitting for the target star's proper motion and parallax. In each column, the top panel shows residuals with respect to R.A. over time, and the bottom panel shows residuals with respect to Decl.

companion (ρ) but scaled down by the difference between the companion's fractional mass (B) and fractional luminosity (β) with respect to the primary star (Van De Kamp 1967):

$$p = \rho \cdot \left(\frac{M_B}{M_A + M_B} - \frac{F_B}{F_A + F_B} \right) = \rho \cdot (B - \beta) \tag{3.1}$$

Following this principle, the semi-major axis of the photocentric orbit (α) is proportional to the semi-major axis of the relative orbit (a) by the same factor:

$$\alpha = a \cdot (B - \beta). \tag{3.2}$$

The shape of the photocentric orbit is thus identical to that of the relative orbit, with only the size reduced by the $(B - \beta)$ scale factor, whereas the orientation is similar but

flipped by 180°. This proportionality makes some binary orbits more difficult to detect than others — for example, two equal-mass stars with equal fluxes would exhibit no apparent photocentric orbit ($\alpha = 0$). The consequences of this selection function on the sensitivity of the RECONS program are discussed in §**??**.

Although orbital motion is visible in the residuals to the proper motion and parallax fit, an accurate orbital fit requires fitting the orbit to the data simultaneously with fitting the proper and parallactic motions, rather than fitting the sets of residuals. Fitting all three motions simultaneously rather than in sequence allows the algorithm to balance them equally, rather than optimizing the first model fit (parallax+proper motion) at the expense of the second (orbital motion). In testing these two approaches, we found that fitting proper motion and parallax first, separately from orbital motion, resulted in orbits with systematically lower eccentricities than fits of the same systems using relative astrometry data (which reflects only orbital motion, no proper motion and parallax).

3.4.2 The MCMC Orbit Fitting Algorithm for RECONS Astrometry

To fit the proper, parallactic, and orbital motions simultaneously, we used the Markov chain Monte Carlo (MCMC) IDL code of Dieterich et al. (2018) designed specifically for RECONS astrometry data taken at the CTIO/SMARTS 0.9m. This code is open-source and available on GitHub (Dieterich 2018). The inputs for the Dieterich et al. (2018) orbit fitter are the target star's R.A. and Decl. displacements from its initial position after DCR and the trail plate calibration have been applied (§3.3). The user specifies hard limits on each of the 10 parameters to fit — proper motions in R.A. and Decl., parallax, and the 7 orbital parameters

— as well as ranges from which each parameter will be drawn. The MCMC code draws values uniformly from within these ranges, thus they constitute uniform priors.

Formally, the astrometric model computed by the code gives the displacement at time t of the photocenter in R.A. and Decl. ($\Delta\alpha$ and $\Delta\delta$) from the center of mass as:

$$\Delta\alpha = \mu_\alpha\,(t - t_0) + \pi\,p_\alpha + (B\,X + G\,Y), \tag{3.3}$$

$$\Delta\delta = \mu_\delta\,(t - t_0) + \pi\,p_\delta + (A\,X + F\,Y), \tag{3.4}$$

where t_0 is the time of the first observation, μ_α and μ_δ are the proper motions in R.A. and Decl., π is the parallax, and p_α and p_δ are the parallax factors in each direction. Several of the orbital parameters are encoded in the Thiele-Innes constants A, B, F, and G above via:

$$A = a\,(\cos\Omega\cos\omega - \sin\Omega\sin\omega\cos i)$$

$$B = a\,(\sin\Omega\cos\omega + \cos\Omega\sin\omega\cos i)$$

$$F = a\,(-\cos\Omega\sin\omega - \sin\Omega\cos\omega\cos i)$$

$$G = a\,(-\sin\Omega\sin\omega + \cos\Omega\cos\omega\cos i)$$

where a is the semi-major axis, Ω is the longitude of the ascending node, ω is the argument of periastron, and i is the inclination of the orbit such that $i = 90°$ is along the line of sight. Figure 1.1 illustrates these parameters of the orbit. Finally, the rectilinear coordinates X

and Y above are

$$X = \cos E - e,$$

$$Y = \sqrt{1 - e^2} \sin E,$$

with e indicating the orbital eccentricity and E the eccentric anomaly, as given by Kepler's equation:

$$E - e \sin E = \frac{2\,\pi}{P}(t - T). \tag{3.5}$$

The orbital period is P and the time of periastron passage is T.

The goal of the algorithm is thus to find the combination of orbital parameters a, e, i, Ω, ω, T, P, and astrometric parameters μ_α, μ_δ, π, that best matches the model given the values $\Delta\alpha$, $\Delta\delta$, p_α, p_δ, and t for each data point. The MCMC algorithm effectively accomplishes this goal by exploring the parameter spaces and building distributions of where it explored, i.e., the "posterior" distributions, that indicate the most likely values for the best-fit model. For this process to work, the algorithm needs to draw models from the entire possible range of each parameter, but also draw more values around models that are relatively better fits to the data than the others.

The user specifies two ranges of values for each parameter: an "initial values" range, from which the MCMC algorithm uniformly draws random starting values for each parameter, and a "hard limit" range, from which the parameters are not allowed to escape. After drawing parameter values on each iteration, the MCMC algorithm chooses a small step for each

parameter and evaluates whether or not that new model is an improved fit to the data over the previous model. To make this decision, the algorithm uses the χ^2 to compute the probability that each model matches the data, then takes the ratio of these probabilities to determine if the new model is a sufficient improvement. A modified Metropolis-Hastings procedure is used to determine this threshold (see Dieterich et al. 2018, Appendix A, for further details). Improved models are accepted, while worse models are rejected and a new random step is chosen. These incremental steps form a "chain" of models, and after a specified number of them — for this thesis, 200 steps — the chain is declared finished and a new chain is begun with new random parameter values drawn again from the "initial values" ranges. The MCMC process terminates when 51 chains have completed.

The algorithm varies the step size for each parameter based on the partial derivatives at those points, which avoids the potential of one parameter dominating the effect of each step. This issue is particularly a danger for the orientation angles Ω, ω, and i, as certain values of these angles would cause small changes in motion to be particularly large or small on the sky through their projection effects.

To avoid over-exploring one small part of the parameter space, every 200 steps the Dieterich et al. (2018) algorithm also takes a "spider step" for one of the parameters (selected at random) wherein it chooses a completely new random value within the hard-limits range. After varying the chain normally for 100 more steps, the algorithm evaluates whether the spider step was an improvement; if not, it reverts that parameter to its pre-spider value and continues the chain. At any point, if a step generates a value that it outside the "hard

limit" range set by the user, then the algorithm rejects that step and computes a new one at random. The spider step mechanism allows the user to set very broad ranges for the parameters, as it ensures the algorithm will not get sidetracked by a local minimum. This avoids the potential of the posteriors reflecting the user's parameter starting ranges (i.e., the user's initial guesses).

The code terminates when it reaches a specified number of chains — for this thesis, 51 chains were usually enough to reach convergence. Then a separate code is used to extract the last 100,000 models and plot the distributions of parameter values of these models, i.e., the posteriors (also referred to as the probability density functions or PDFs). If they show Gaussian peaks around specific parameter values, the chains have converged and the best-fit model can be inferred from the median of those PDFs. Examples of these PDFs are shown in Figure 3.5.

All the orbits fit from the RECONS astrometry program are shown and discussed in §5.1.1.

3.4.3 Fitting Orbits Longer than the Astrometric Time Baseline

Our orbit fitting procedure works reliably as long as the entire orbit is visible in the astrometric observations — i.e., for systems with $P_{\rm orb}$ shorter than the observed time baseline. Orbits that are not completely observed in these data are difficult or even impossible to constrain because the orbital motion cannot be distinguished from proper motion. For example, a short curved arc of motion could be fit by a decade-long orbit or a century-long orbit, depending on the proper motion (both angle and total amount) assumed to be contributing

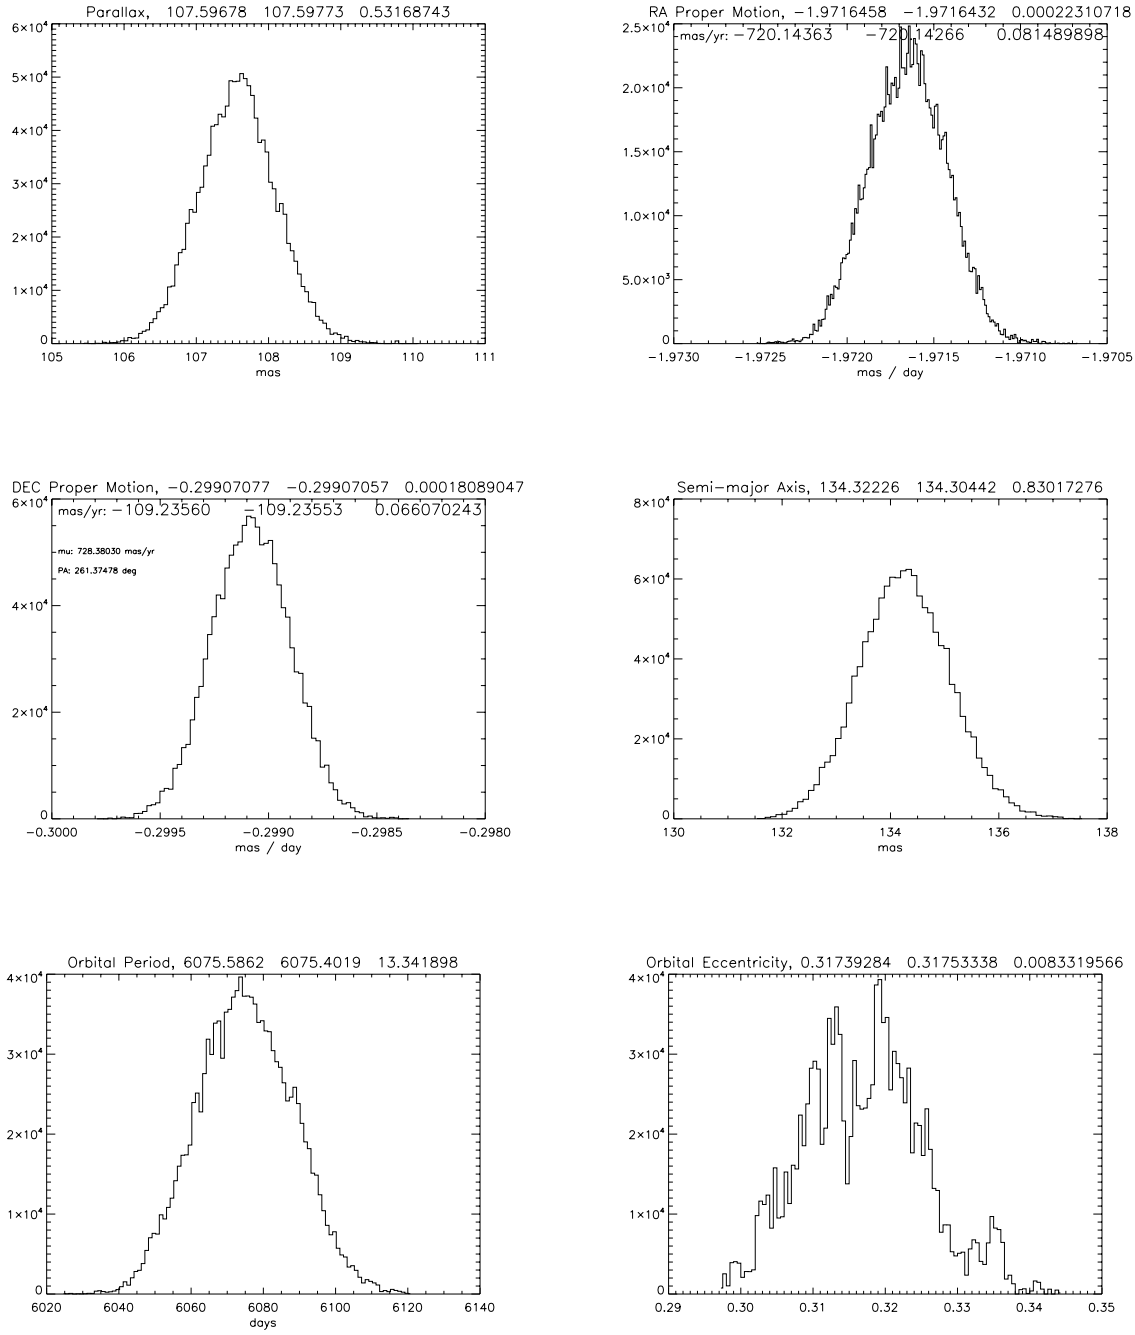

Figure 3.5 Examples of well-convered probability distribution functions (PDFs) for one astrometric binary (WT 460 AB) fit by the Dieterich et al. (2018) MCMC code. For space reasons, only a subset of the ten fit parameters are shown here. *Top left* to *bottom right*: parallax (in mas), proper motion in R.A. (mas yr^{-1}), proper motion in Decl. (mas yr^{-1}), semi-major axis of the photocentric orbit (mas), orbital period (days), and eccentricity. At the top of each plot, the mean of the distribution is printed on the left, the median in the middle, and the standard deviation on the right.

to that arc. This degeneracy is thus unique to absolute astrometry such as the RECONS data, and is not a factor when fitting relative astrometry (such as from the SOAR program; Chapter 4).

In practice we have found that this proper vs. orbital motion degeneracy means we can only fit orbits that are less than ∼110% the length of a given system's observational baseline. If sufficient motion is visible (at least two "turns" in R.A. and one in Decl., or vice versa), we can fit longer orbits by holding the proper motion fixed while fitting the rest of the orbital parameters and parallax. For this dissertation, this method was only applied when an initial 10-parameter fit showed:

1. orbit size parameters (period and semi-major axis) were constrained even if not converged,

2. PDFs of μ_α and μ_δ were well-constrained to within a few tenths of mas per day (the medians of these distributions were chosen as the fixed values),

3. eccentricity well converged despite the orbit size parameters (period and semi-major axis) being poorly constrained.

Given these caveats, this technique of holding proper motion fixed was only applied when $\gtrsim$80% of the orbit was observed over the given time baseline. Shorter orbital arcs did not allow these parameters to be well enough constrained. This technique also induces some bias toward less-eccentric orbits, as those show more orbital motion over shorter time period; partially observed eccentric orbits look more linear unless the observations happen to catch

them at periastron. Biases related to orbit shape for RECONS astrometry are discussed in more detail in §5.1.2.

CHAPTER 4

The SOAR Speckle Contribution

The speckle program was initiated in 2019 at the Southern Astrophysical Research (SOAR) 4.1m telescope to map orbits that complement the RECONS astrometry results. Because the SOAR effort has been underway for nearly years, it is particularly good for determining orbits with periods less than a few years, while the RECONS astrometry effort is best in providing systems with periods of 3 years and longer. Although the cadence of the SOAR observations is similar to that of the RECONS program, these data map relative orbits directly — proper motion and parallactic motion thus need not be considered because the motion seen is only the orbital motion of the secondary star about the primary. This feature also makes it straightforward to combine SOAR speckle data with other high-resolution imaging data in the literature to form a longer or more complete observing time baseline for a given system.

In this chapter we begin with the target selection for the SOAR program, which includes multiples and suspected multiples from the literature (§4.1.1), from the RECONS astrometry program (§4.1.2), and from *Gaia* DR2 (§4.1.3). We then explain the procedures for observing and reduction of these data (§4.2), and finally describe how orbits are fit to those results (§4.4).

Many of the details shown here regarding the SOAR speckle target list and observations were first presented in Vrijmoet et al. (2022), which described results from the first 1.5 years of our SOAR speckle observing program.

4.1 Choosing Targets for SOAR Speckle

To increase the likelihood of companion detections and to maximize efficiency during the allocated observing time, the SOAR program targeted only known or very likely multiples. These included systems identified in the literature (§4.1.1), from the RECONS astrometry program (§4.1.2), and from *Gaia* DR2 (§4.1.3). These three subsets had substantial overlap; for example, many systems with perturbations in RECONS data were also known multiples in the literature (sometimes from previous RECONS work). The overlap between these subsamples is illustrated in Figure 4.1. In total there were 337 systems targeted with M dwarf primaries, with the majority — but not all — being confirmed multiples at the beginning of the program. Table 4.1 gives the full target list, with the subsets from Figure 4.1 (and discussed below) marked in columns 10–12.

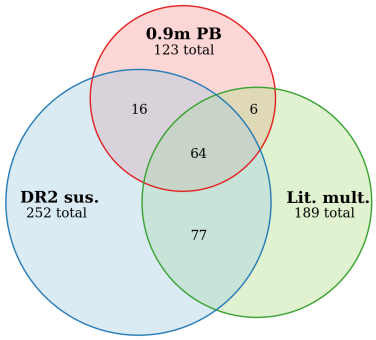

Figure 4.1 Venn diagram illustrating the three subsets of systems with M dwarfs targeted in the SOAR speckle survey, comprising 337 systems total. Each circle is labeled according to the subset's source: "0.9 m PB" for targets showing perturbations (PBs) in the RECONS astrometry program at the CTIO/SMARTS 0.9m (§3), "Lit. mult." for known multiples from the literature, and "DR2 sus." for systems suspected to be multiples based on their *Gaia* DR2 results. The number of targets is given under each subset name, and the numbers in the overlapping sections indicate the number of targets common to multiple subsets. The area of each circle is proportional to the number of targets in that subset, although the overlapping regions are not to scale.

Table 4.1: Systems targeted for the M dwarf speckle interferometry program at SOAR. The WDS code (column 3) is the name each system *would* have in the Washington Double Star Catalog (Mason et al. 2001) if it were included there (some are there already, some not). The subsets justifying each target's inclusion are marked in columns 10–12, and are discussed in depth in §4.1.1, §4.1.2, and §4.1.3. This table is given in full at `http://www.astro.gsu.edu/~vrijmoet/PhDthesis/sample_soar.txt`

R.A. J2000.0 (1)	Decl. J2000.0 (2)	WDS (3)	Discov. code (4)	Name (5)	π (mas) (6)	π ref. (7)	V (mag) (8)	V ref. (9)	$V-K$ (mag) (10)	0.9m PB (11)	Lit. mult. (12)	DR2 sus. (13)	SOAR res. (14)	Orbit ref. (15)
00 06 39.24	−07 05 35.9	00067−0706	JNN 11	2MA0006-0705AB	46.960 ± 0.403	EDR3	14.72	APdr9	5.76		✓	✓	N	
00 08 53.92	+20 50 25.6	00089+2050	BEU 1	G131-026AB	55.256 ± 0.761	DR2	13.52	Rie14	5.51	✓	✓	✓	Y	Vri22
00 09 45.04	−42 01 39.3	00098−4202		LEHPM1-0255AB	60.889 ± 0.350	EDR3	13.62	Win15	5.40			✓	Y	
00 13 46.60	−04 57 37.2	00138−0458		LHS1042	42.627 ± 0.219	EDR3	17.98	estim	7.50			✓	N	
00 15 27.99	−16 08 01.8	00155−1608	HEI 299	GJ1005AB	169.522 ± 0.969	Vri20	11.48	Win15	5.09	✓	✓	✓	Y	Ben16
00 15 58.07	−16 36 57.6	00160−1637	BWL 2	2MA0015-1636AB	56.096 ± 0.093	EDR3	13.20	Win19	5.29		✓	✓	Y	Vri22
00 16 01.97	−48 15 39.1	00160−4816	TOK 808	L290-072AB	40.672 ± 0.525	EDR3	11.55	Koe10	4.44		✓	✓	Y	
00 16 14.63	+19 51 37.5	00162+1952		GJ1006AC	65.108 ± 0.041	EDR3	12.26	Wei96	5.17		✓	✓	Y	
00 21 37.26	−46 05 33.4	00216−4606		L290-028	51.569 ± 0.045	EDR3	12.24	Koe10	4.79			✓	N	
00 24 44.19	−27 08 24.2	00247−2653	LEI 1AB	GJ2005AB	129.317 ± 0.126	EDR3	15.28	Win15	7.04	✓	✓	✓	Y	Koe12
00 24 44.10	−27 08 24.0	00247−2653	LEI 1BC	GJ2005BC	129.317 ± 0.126	EDR3	15.28	Win15	7.04	✓	✓	✓	Y	Man19
00 25 04.31	−36 46 17.9	00251−3646	BRG 2	LTT00220AB	49.871 ± 0.110	EDR3	12.48	Win15	4.65		✓	✓	Y	
00 27 55.99	+22 19 32.8	00279+2220	FRV 1	LP349-025AB	70.781 ± 0.427	EDR3	18.04	Win17	8.47	✓	✓	✓	Y	Dup10b
00 32 53.14	−04 34 07.0	00329−0434	JNN 12AB	GIC0050AB	52.853 ± 0.100	EDR3	13.97	Win15	5.62	✓	✓	✓	Y	
00 32 53.14	−04 34 07.0	00329−0434	JNN 12AC	GIC0050AC	52.853 ± 0.100	EDR3	13.97	Win15	5.62	✓	✓	✓	Y	
00 43 26.00	−41 17 34.0	00434−4118		LHS1134	97.661 ± 0.032	EDR3	13.00	Win15	5.29	✓			N	
00 48 13.33	−05 08 07.7	00482−0508		LTT00453	40.227 ± 0.030	EDR3	12.03	Win15	4.10			✓	N	
00 58 27.94	−27 51 25.4	00585−2751		GJ0046	84.979 ± 0.462	EDR3	11.77	Bes90	4.88			✓	N	

Regardless of their subset, the instrument configuration on SOAR required that all targets be brighter than $I = 14.0$ mag in the Johnson-Cousins I band, which was the filter used for these observations. For targets with no previous I measurements, brightnesses were estimated by assuming the *Gaia* R_G magnitude was roughly equivalent. Collaborator and SOAR/HRCam+SAM Instrument Scientist, Andrei Tokovinin, carried out the observations and found that none of these targets were too faint to observe, thus the $R_G \approx I$ estimate was sufficiently valid.

4.1.1 189 SOAR Targets from the Literature

Roughly 50% of the SOAR speckle targets were already identified as known multiples in the literature, some via high-resolution imaging and others through spectroscopy or (unresolved) astrometry, e.g., the RECONS astrometry program. In observing these known multiples with HRCam+SAM, our goals were:

1. Obtain separations and position angles for companions that were previously unresolved from their primary stars.

2. Map each companion's motion through relative astrometry to enable fitting its relative orbit, incorporating any previous data available to extend its time baseline.

3. Improve upon any existing orbit fits, in particular by improving the accuracy of the orbital elements.

To select known multiples that would have pre-existing data, we matched the sample of 25 pc M dwarfs (Chapter 2) against the Sixth Catalog of Orbits of Visual Binary Stars (ORB6;

Hartkopf et al. 2001) and the Washington Double Star Catalog (WDS; Mason et al. 2001). Some additional targets were identified through private communication with collaborators.

Cross-matching with known orbits in ORB6 allowed us to exclude systems with preliminary orbits much longer than the 0–30 year regime targeted by this thesis. For the systems without previous orbits (e.g., from WDS), we excluded potentially long orbits by requiring each pair to be separated by $\lesssim 2\farcs0$ at its widest separation recorded to date. Matching against ORB6 also ensured that many of these M dwarf multiples already had characterized orbits, supporting our goal of improving on previous characterizations, and allowing us to supplant those with new measurements of magnitude differences to better characterize each component. Roughly one-third of the systems in this literature sample also had never been resolved in high-resolution imaging before, as their orbits were spectroscopic or astrometric. Resolving these systems gives valuable information about the components' physical separation and flux ratio (in this case, in the I band), both of which can be used to help constrain their masses.

Through this process, the SOAR target list included 189 known systems with M dwarf primaries, 81 of which already had orbits in the literature.

4.1.2 123 SOAR Targets from RECONS Astrometry

About one-third of the targets on the SOAR observing list were drawn from the RECONS astrometry program (Chapter 3), which reveals multiples based on perturbations detected in the positions of target stars' photocenters. These SOAR targets all had evidence of potential multiplicity based on those data, but many did not yet have a companion confirmed — and

a substantial number had ambiguous signals that could not be clearly classified as orbital, as discussed below. Adding speckle data for these targets accomplishes several aims to aid this thesis:

1. Every result — both resolutions and non-resolutions — constrains orbital parameters for orbits that are incompletely mapped in the RECONS astrometry.

2. When no companion is resolved around a given star, this non-resolution places constraints on the potential companion's luminosity (and thus its mass) and its separation from the primary star.

3. For systems with ambiguous signals in RECONS astrometry (discussed below), the speckle data either alleviates that ambiguity with a detection or mitigates it by revealing the absence of a companion within specific luminosity (thus mass) and separation limits.

4. When an orbit is completely mapped by RECONS astrometry and the companion is also resolved in speckle data, the combination of those data allows the determination of the components' individual dynamical masses.

A multiple is identified in the RECONS data when the changing orbital position of the two components causes a perturbation (or "PB") of the source's position after solving for proper motion and parallax. The position shifts observed are that of the photocenter (center of light) of the unresolved pair of stars; as discussed in §3.4, this photocentric orbit is identical to the relative orbit of the secondary star about the primary, but is rotated 180°

(in argument of periastron) and is smaller in size (semi-major axis), typically by a factor of 2–3.

Observing photocentric motion creates some limitations on the types of companions that can be detected in the RECONS data. On the plus side, a multiple can generally be *identified* with this technique when as little as 10% of its orbit is observed, as long as those observations have captured some of the orbit's curved shape. An orbit thus need not be fully mapped in RECONS data to be a good candidate for our SOAR speckle follow-up program. On the other hand, no perturbation is detected in cases of planetary-mass companions and for systems where the components have nearly-identical fluxes (in which case the center of light is halfway between the two sources at all times and hence, does not shift; §3.4.1). Orbital motion is also more difficult to detect when the astrometric precision is poor, such as for bright stars ($\lesssim$10 mag for RECONS data), or when the orbit is poorly sampled, e.g., when the orbital period is short so that many orbits are completed between observations.

Although many systems show a PB with clear orbital motion that can be fit by an astrometric model (§3.4), a similar number have a PB that is clearly evident but more ambiguous in nature. Two examples are shown in Figure 4.2. These ambiguous PBs could represent true orbits that are poorly sampled or unfortunately oriented on the sky, or they could be non-astrophysical artefacts, such as star positions that have been mis-measured because the target has moved in front of a background star or if its PSF overlaps bad CCD pixels. Another notable class of non-astrophysical PBs were created by changing the V filter to a supposedly identical copy. The "new V" filter was photometrically identical, but

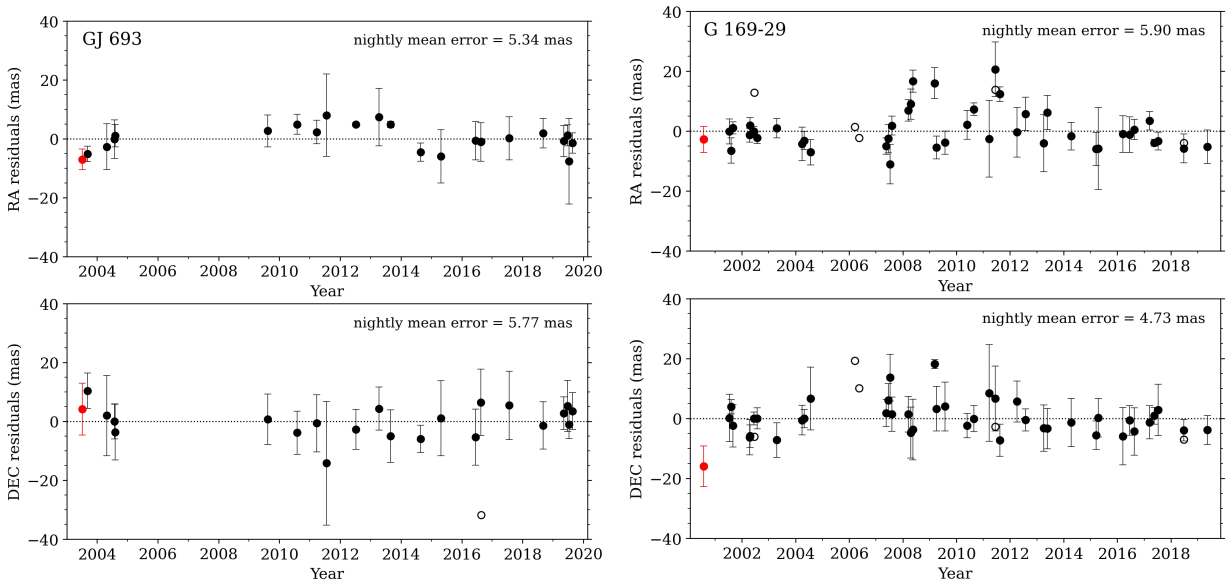

Figure 4.2 Two examples of astrometric perturbations (PBs) in RECONS data that are ambiguous, or not clearly astrophysical: GJ 693 (*left column*) and G 169-29 (*right column*). To contrast, compare these to the PBs showing clear orbital motion in Figure 3.4. For each system, each panel shows residuals in R.A. or Decl. with respect to time, computed after fitting proper motion and parallax to the system's astrometry.

its bandpass curve — and thus the way stars shifted in the astrometric field due to their different colors — was slightly different than the previous one, so the positions of the stars as measured in the "old V" were several milliarcseconds different than those measured in the "new V."

Ultimately, of the 123 systems contributed to the SOAR target list by RECONS astrometry, 62 were resolved with SOAR (§5.2).

4.1.3 252 SOAR targets from Gaia *DR2*

To supplement the SOAR target list beyond known M dwarf multiples, we added 95 systems that showed evidence of *potential* multiplicity based only on the quality of their astrometry in

Gaia DR2. These are the targets among the 252 total systems from DR2 represented in the blue region of Figure 4.1 that are not in either the literature and RECONS lists. Applying the criteria outlined below to the entire SOAR target list showed that 75% had this evidence of multiplicity in *Gaia*, indicating these criteria were reliable for selecting potential multiples. The targets added from *Gaia* that were not yet known multiples were included in the SOAR campaign to:

1. Identify and confirm new multiples and begin building their long-term data sets for the astronomical community.

2. Measure magnitude differences at I to characterize components of any new multiples.

3. Map orbits of multiples with $P_{\rm orb} \lesssim 6$ years, to which these *Gaia* DR2 data should be particularly sensitive because DR2 only observed for 22 months.

The specific criteria used to construct this list of "Gaia suspects" (potential multiples) are outlined in the analysis presented in Vrijmoet et al. (2020). In that work, we compared RECONS astrometry results to *Gaia* DR2 to identify the most useful parameters for potential multiplicity in the *Gaia* catalog. The steps of that analysis are summarized as follows:

1. We matched all the systems — singles or multiples — ever published by RECONS to their counterparts in *Gaia* DR2. This allowed us to consider the *Gaia* results for ∼500 systems for which we already had access to their time-series astrometry data.

2. *Gaia* provides a substantial suite of parameters describing different aspects of the data and the quality of their fit to the astrometric and photometric models. From this list,

we considered eight parameters that might be affected by the presence of a companion.

3. We distinguished the known multiples from the systems with perturbations in RECONS astrometry (as described in §3.4 and §4.1.2) and plotted the values of those eight parameters. These plots are shown in Figures 4.3 and 4.4.

4. The four parameters shown in Figure 4.3 revealed clear differences in the distributions for multiples vs. single stars and could be used to identify unresolved multiples in *Gaia* DR2 data. The remaining four parameters, shown in Figure 4.4, were not so helpful.

5. Based on the distributions of "helpful" parameters in Figure 4.3, we identified the value of each parameter for which 75% of the stars above that value were unresolved multiples. Stars below that value were not necessarily single, but stars above that value were overwhelmingly likely to be multiple. These "cutoff" values form the criteria for selecting likely unresolved multiples in *Gaia* DR2.

The above analysis indicates that three out of four stars that fulfilled the following criteria in *Gaia* DR2 were unresolved multiples:

1. missing *Gaia* DR2 parallax or missing DR2 catalog entry,

2. `parallax_err` ≥ 0.32 mas for $G \lesssim 18$ mag (≥ 0.40 mas otherwise),

3. `astrometric_gof_al` ≥ 56.0,

4. `astrometric_excess_noise_sig` ≥ 108.0, and

5. `ruwe` ≥ 2.0.

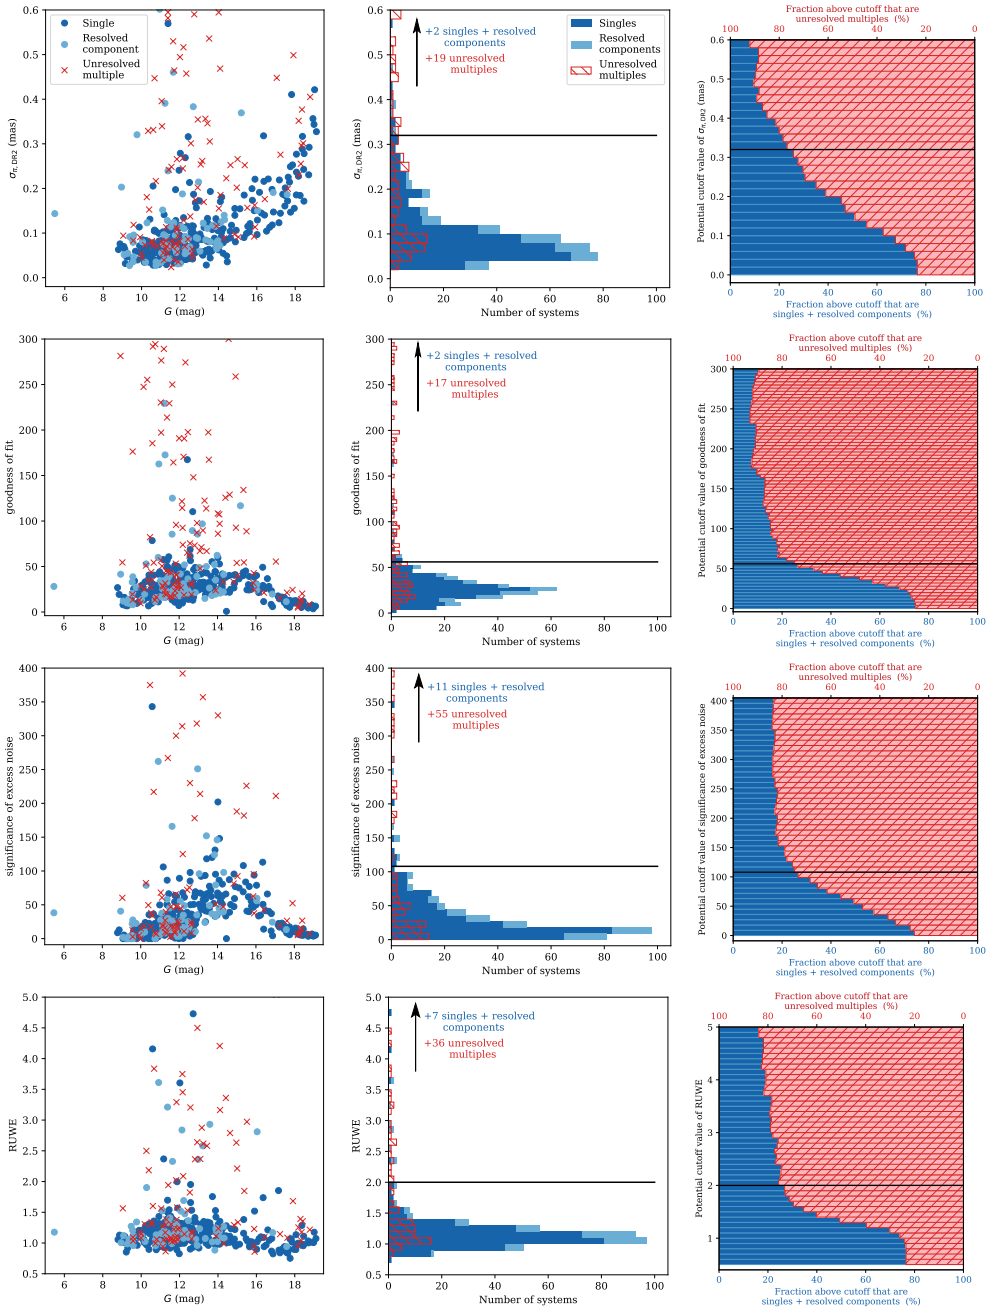

Figure 4.3 Four astrometric fit parameters in DR2 that are useful for selecting potential unresolved multiples. Unresolved multiples are indicated with red symbols and hatched bars, resolved components are light blue dots/bars, and presumed singles are dark blue dots/bars. *Left column:* parameters for singles, resolved components, and unresolved multiples in *Gaia* DR2, plotted against their G mags. Unresolved multiples tend to have higher values here, indicating their poor astrometric fits in DR2. *Middle column:* distributions of these parameters for systems from the left, separated again by multiplicity. Systems not shown are noted in each panel with an arrow and text. Although both distributions peak at low values of each parameter, unresolved multiples extend to higher values. *Right column:* potential cutoff values for each parameter, showing the fraction of systems above each cutoff belonging to singles and resolved components (blue) and unresolved multiples (red). The values above which 75% of our systems are unresolved multiples are indicated with a black line.

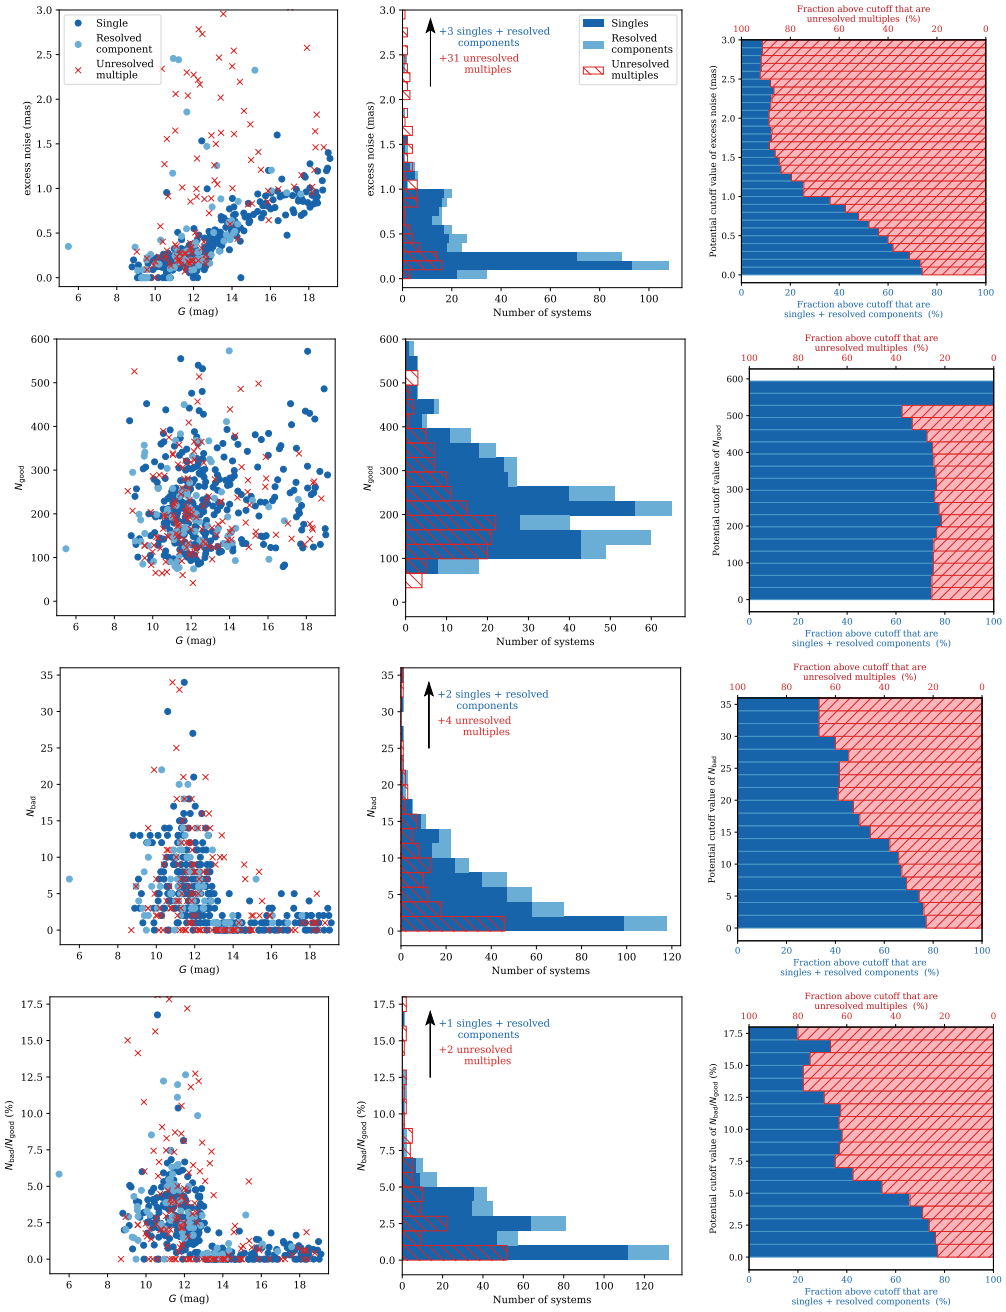

Figure 4.4 Four astrometric fit parameters in *Gaia* DR2 that are less useful for choosing potential unresolved multiples. The color schemes and columns are the same as for Figure 4.3, and any systems with values exceeding these ranges are noted with an arrow and text in each panel. For these quantities, the distributions of single and unresolved systems do not differ as significantly as those in Figure 4.3, making them less useful for identifying potential unresolved systems. The parameter `astrometric_excess_noise` does show a distinction between singles and multiples, but it is less useful than the very similar `astrometric_excess_noise_sig` because it strongly depends on G magnitude (faint singles have values similar to brighter unresolved multiples).

Note that the above statement does not *guarantee* that a system meeting these criteria has a companion — 25% of stars in that group had no evidence of multiplicity in RECONS data or the literature. Other factors such as high proper motion or very crowded fields could have affected the quality of those systems' *Gaia* DR2 astrometric fits instead. Note also that systems *not* meeting these criteria could still be unresolved multiples. These could include systems with very small astrometric signals (§3.4), with very slow orbital motions (long $P_{\rm orb}$), or with orbital motions that mimicked a single star's linear motion during the *Gaia* observations (such as a minimally curved portion of an eccentric orbit).

While selecting targets for the "*Gaia* suspects" subset, we also chose several dozen targets that had parameters slightly lower than the formal criteria from Vrijmoet et al. (2020). This supplementary list of 32 systems was included to leave the door open for possibly revising the criteria later once the SOAR speckle campaign revealed or confirmed new multiples.

4.2 SOAR Speckle Observing and Data Reduction

Data for this program were collected at the 4.1m Southern Astrophysical Research (SOAR) telescope by the high-resolution camera (HRCam) mounted on the SOAR adaptive optics module (SAM; Tokovinin et al. 2016). Since 2007 this instrument+telescope configuration has been imaging stellar multiples using the speckle interferometry technique (Tokovinin 2018), with regular results published in a yearly paper series (e.g., Tokovinin et al. 2020, 2021, 2022). Although M dwarfs have been occasionally included in those results, no dedicated M dwarf speckle observing program existed on SOAR before this thesis effort.

The observations for this thesis program were conducted from July 2019 to March 2023, with all data taken and reduced by HRCam Instrument Scientist Andrei Tokovinin. The telescope time for the program was allocated by NSF's NOIRLab through several standard and long-term proposals. The allocated time was always combined with that of other observing programs on HRCam+SAM to increase the opportunities for timely observations of fast-orbiting systems. Typically, targets received 2–5 observations per year. In preparation for each observing run, we considered previous SOAR observations and RECONS astrometry to prioritize systems that exhibited rapid orbital motion. This procedure improved the likelihood that defining features of the orbit shapes, e.g., epochs of periastron, would not be missed.

These HRCam+SAM observations were conducted in the instrument's seeing-limited mode, meaning that no laser guide star was used during image acquisition. Frames were taken almost exclusively in the Kron–Cousins I filter, usually in 2–3 sets (data cubes) of 400 frames per target, with integrations typically 24 ms per frame. These sets were each processed independently to verify results. Most observations used the HRCam narrow $3''$ field of 200×200 pixels, whereas pairs known to have separations of $1''.4$ or more were observed with the $6''$ field of 400×400 pixels. The formal resolution limit in I on the SOAR 4.1m is 45 mas, but the ultimate separation limit reached depends on target brightness and sky conditions. In some cases, good conditions and the data reduction procedure (discussed below) allow a system to be "super-resolved" down to 35 mas, beyond the formal definition of resolution (as shown in Figure 1 of Tokovinin et al. 2020). Targets that are unresolved in

the first two attempts are usually observed a third time, then retired from the program if still unresolved.

The data are processed and reduced for this program using the standard procedures described in Tokovinin et al. (2010) and Tokovinin (2018), and representative images of the reduced data products are shown in Tokovinin (2018). In brief, for each target the power spectrum and autocorrelation function are calculated, and companions are noted via power spectrum fringes or secondary peaks in the autocorrelation function. Fitting an empirical model to the power spectrum yields the parameters of each detected pair: the separation between components (ρ), the position angle (θ) of the secondary with respect to primary star (north = 0° through east = 90°), and the difference in magnitude between components (Δm). In addition to the those three measurements from the data reduction, important details about these results are:

1. The position angle determined through this procedure is only ascertainable modulo 180°, leaving some ambiguity in the secondary's true position on the sky. This ambiguity has been eliminated whenever possible by applying a shift-and-add procedure to each target's data (Tokovinin 2018); this process reveals the true quadrant for companions that are not too faint but still have some magnitude difference with their primary star ($\Delta m \gtrsim 0$ mag).

   These results are noted with the "q" flag in Table 4.2, indicating that the quadrant has been determined.

2. For some observations of wider pairs, a separate procedure is used to determine the

magnitude difference using the average image for a target (described in detail in Tokovinin et al. 2010). This method produces more reliable photometry for cases where the stars' separations are greater than the size of the isoplanatic patch, which is the angular size of coherent turbulent cells in the atmosphere.

Observations with Δm determined with this method are marked by a "p" in Table 4.2, indicating that this photometric method has been used.

3. For observations in which no companion was detected, a contrast curve is computed to report the detection (magnitude) limits as a function of the distance from the primary star on the sky (for example, see Figure 5 of Tokovinin 2018). The parameters of this curve are reported in the results (§4.3 and Table 4.2) as the minimum separation resolvable for pairs with $\Delta m < 1$ mag, as determined from the maximum spatial frequency of the power spectrum, and the maximum detectable magnitude difference at separations of $0\farcs15$ and $1\farcs0$ (the dynamic range).

4.3 SOAR Speckle Data Results

Here we briefly present the data obtained from the SOAR speckle campaign, as the results of interest for the remaining dissertation chapters will be the orbits fit to these data rather than the data points themselves.

Overall, of the 337 M dwarf systems in the SOAR program, 1289 total speckle observations were made on all 337 systems. Of these, 216 (64%) were resolved. Notably, this includes new companions for 76% of the subset selected from Gaia DR2. Among the supplemental sample of 32 stars, 13 (41%) were resolved, indicating that stretching the criteria

slightly does reveal additional multiples. In all, we report the first direct detections of 97 new stellar companions to the observed M dwarfs. The system-by-system results are given in Table 4.2.

To augment the SOAR speckle results, we also searched the literature for additional high-resolution imaging data on these targets to add to their data sets. These literature results were essential for (1) extending the time baseline of the observations to enable fitting longer orbits, and (2) informing or confirming the correct quadrant for the SOAR speckle data's ambiguous position angles. The full list of references contributing data to each orbit fit is given alongside those orbits in Chapter 5.

Table 4.2: Results of the SOAR speckle interferometry campaign for nearby M dwarfs. The meaning of columns 8–10 differ depending on whether or not the companion was resolved in that observation, hence the two lines in those columns' headers. The ":" flag in the last column indicates highly uncertain results; the rest of the flags are described in §4.2. This table is given in full at `http://www.astro.gsu.edu/~vrijmoet/PhDthesis/results_soar.txt`

WDS	Name	Date of obs. (yr)	N	Fil.	Res? (Y/N)	ρ (if Y) $\rho_{\rm min}$ (if N) (mas)	θ (if Y) $\Delta m_{0\farcs15}$ (if N) (° or mag)	Δm (if Y) $\Delta m_{1\farcs0}$ (if N) (mag)	flag
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
00067-0706	2MA0006-0705AB	2019.8568	2	*I*	N	0.0768	2.3	2.9	
00067-0706	2MA0006-0705AB	2020.8342	2	*I*	N	0.0594	2.3	2.8	
00067-0706	2MA0006-0705AB	2022.4419	4	*I*	Y	0.1019	204.8	0.7	q
00067-0706	2MA0006-0705AB	2022.6823	2	*I*	Y	0.1135	201.9	0.9	:
00067-0706	2MA0006-0705AB	2023.0063	2	*I*	Y	0.1349	198.1	1.0	q
00089+2050	G131-026AB	2019.5397	2	*I*	Y	0.1110	154.4	0.0	
00089+2050	G131-026AB	2019.8564	2	*I*	Y	0.1233	130.0	0.3	
00089+2050	G131-026AB	2020.8340	2	*I*	Y	0.1522	77.5	0.0	
00089+2050	G131-026AB	2020.9241	2	*I*	Y	0.1514	74.2	0.2	
00089+2050	G131-026AB	2021.5684	2	*I*	Y	0.1436	42.8	0.3	
00089+2050	G131-026AB	2022.4420	2	*I*	Y	0.1243	349.3	0.2	
00098-4202	LEHPM1-0255AB	2019.6133	2	*I*	Y	0.0522	339.0	0.8	
00098-4202	LEHPM1-0255AB	2019.8567	2	*I*	N	0.0525	2.5	3.9	
00098-4202	LEHPM1-0255AB	2020.8341	2	*I*	Y	0.0959	115.7	1.0	
00098-4202	LEHPM1-0255AB	2020.9270	2	*I*	Y	0.1089	117.4	0.9	q
00098-4202	LEHPM1-0255AB	2021.5684	2	*I*	Y	0.1498	128.7	0.9	q

4.4 Fitting Orbits to SOAR Speckle

An orbit is ready for its first fit when its shape becomes apparent in its imaging data. This state typically occurs when ∼40% of the orbit is mapped by the data. Before the orbit's true period and orientation are known, this coverage is estimated by computing the position angle change between the first and most recent observations and comparing the coverage to the 360° of a full orbit. Some consideration is also made for the companion's speed — if the companion has considerably slowed or sped up over the course of the observations, the available observations (regardless of position angle percentage) may constrain the orbital eccentricity and thus period. A revised orbital fit is attempted when we acquire more data for a system, as the richer data set invariably improves the certainty of the result.

Each orbit is fit with the `ORBIT` code of Tokovinin (2016). This IDL routine uses a Leavenberg-Marquardt least-squares algorithm to find the orbit model that minimizes the differences between the model and data in R.A. and Decl. Each data point is weighted by the inverse square of its uncertainty, with the uncertainty set to nominal values depending on the instrument used for that observation. If an orbit is well-constrained by the data, varying these uncertainties by ±1 mas does not significantly change the resulting best-fit orbit model.

The initial positions of the orbit parameters for `ORBIT` are set by the user. For this thesis, we began with either a previous fit from the literature, a photocentric orbit from RECONS astrometry, or a rough estimate of the orbital period and parameters set to nominal values ($e = 0.5$, $i = 45°$, etc.). The fit was checked for robustness by adjusting these parameters

(to $e = 0.3$, $i = 80°$, etc.) and computing a new fit. Well-constrained orbits would quickly converge to the same result regardless of these adjustments.

If there are any points for which the position angle is still ambiguous modulo 180° — i.e., that flipping the position angle 180° produces an equally plausible path of motion — then we make that adjustment, re-compute the orbit fit, and compare that result with the prior fit. One of these two models must clearly fit the data better in order to be considered robust; if both are equally plausible, then the data are not constraining the models well enough for any model to be considered reliable.

For every model computed (including the data-flipping variations discussed above), we check the result by computing the total mass of the system and comparing that to the masses expected from the absolute magnitudes of those individual stars. Spectral types are also considered if they are available — in particular, the combination of I magnitudes and/or spectral types for a pair of stars can be used to determine whether or not either component is a white dwarf. This procedure provides a very rough check on the feasibility and robustness of the result. Occasionally, the procedure at this stage reveals that the code is converging to a $P_{\rm orb}$ that is significantly too large or too small, and several speckle data points need to be flipped 180° to find the orbit that makes physical sense.

All the orbits we have fit from the SOAR speckle program are illustrated and discussed in §5.2.1.

CHAPTER 5

Results from Three Sources of Orbits

With our sample defined (§2) and observations completed and characterized from the RECONS astrometry program at the CTIO/SMARTS 0.9m (Chapter 3) and the speckle imaging program at the SOAR 4.1m (Chapter 4), we are now in a position to synthesize the results. In this chapter we describe the RECONS astrometry orbits (§5.1) and SOAR speckle orbits (§5.2), including our procedure for assessing the sensitivity of both surveys in $P_{\rm orb}$ vs. e space. Then we augment the sample with results selected from the literature (§5.3).

5.1 The RECONS Astrometry Results

Of the 696 systems on the RECONS astrometry program on 1 April 2023, 132 showed some sort of perturbation (PB) in their single-star fit residuals (§3.4) potentially due to a companion, and 78 of those PBs clearly traced out orbital motion. Here we discuss 19 of those systems to which we fit orbits. Of the remaining systems with orbital motion but no RECONS orbit presented here, many already had orbits published in the literature; if a literature orbit for a system was more precise, we used that in this dissertation in lieu of our RECONS characterization. Many clearly orbiting systems also have no fit here because the orbit was not yet fully mapped by the data, a situation which causes considerable difficulty with obtaining a reliable fit (§3.4.3). §5.1.1 presents the 19 unique RECONS orbits and their general statistics, and §5.1.2 analyzes how well those orbits represent the $P_{\rm orb}$ vs. e parameter space for this thesis, driven primarily by our desire to map orbits on timescales of 0–30 years.

5.1.1 Orbits from RECONS Astrometry

The RECONS astrometry program yielded 19 systems for which we determined a reliable orbit fit. These orbits span 5–29 years in orbital period and 0.06–0.87 in e, as shown in Table 5.1. Roughly half of the systems have primary stars with masses $\lesssim$0.35 $M_{\odot}$, as suggested by the HRD in Figure 5.1, in which $M_V = 11.37$ mag corresponds to 0.35 $M_{\odot}$. The astrometric technique of measuring the target's position with respect to reference stars generally favors the lower mass M dwarfs (12–16 mag, for the 0.9m program) because fainter stars allow longer exposures, which in turn result in smoother point spread functions (PSFs) and better centroids (§3.2). Hence, somewhat counterintuitively, the sensitivity to companions for faint stars is better than for bright stars.

Orbits that could be reliable determined include every orbit with:

- $P_{\rm orb} \gtrsim 3$ years and signal-to-noise (S/N) $\geq 3\sigma$ in R.A. and/or Decl., where σ represents the average uncertainty of the per-epoch position measurement — resulting in clearly traced orbital motion considering our cadence of 2–4 data points per year per target, and

- rbital period shorter than the time baseline of observations — allowing proper motion to be clearly disentangled from orbital motion (§3.4).

A few systems that violated the time baseline rule were added by holding the proper motion fixed to the most likely value determined from a preliminary fit of proper motion, parallax, and orbital motion.

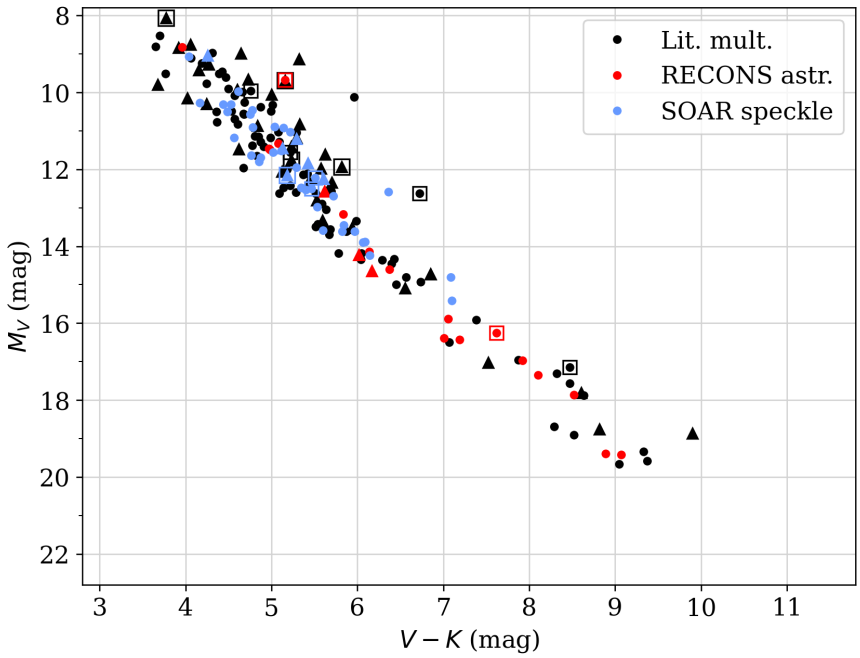

Figure 5.1 Hertzsprung-Russell diagram of every system with an orbit presented in this work. One point is plotted per system, representing combined magnitudes for unresolved pairs. Circular points are binaries and triangular points are subsystems of higher-order multiples, and boxes indicate systems that are likely young. Colors refer to the source of the orbit: RECONS astrometry (red; Chapter 3), SOAR speckle interferometry (blue; Chapter 4), or the literature (black; §5.3).

Systems with projected separations exceeding ∼0.″8 at maximum were usually excluded, as this is the limit is where the two stars' PSFs look like two peaks rather than one unresolved source, depending on the seeing. This consideration is important because the overlap in the PSFs shifts the peaks' positions toward each other, and as the seeing changes the PSFs' shapes, so too does it change those peak positions. The typical seeing at the CTIO/SMARTS varies over 0.″8–2.″0 from night to night, thus point sources with separation in that range will be unresolved in bad seeing and resolved (but with PSFs superimposed) in good seeing. This issue may be alleviated by the labor-intensive process of fitting the PSFs of both stars

simultaneously in every image, but this is only done in particularly compelling cases, such as the nearby young binary AT Mic AB (Holden et al. 2023).

Figures 5.2–5.5 show all 19 orbits determined from RECONS astrometry, and their orbit parameters are given in Table 5.1. We fit each orbit with the Dieterich et al. (2018) MCMC code, as discussed in §3.4. Three of these orbits also have fits published in the literature from other groups using other techniques. Results are consistent between previous orbits and the new results shown here, and for the cases shown here our RECONS orbits are improvements over those previously published results, usually because we have significantly longer data timelines. The remaining 16 orbits are new, with no orbit characterization published before this work.

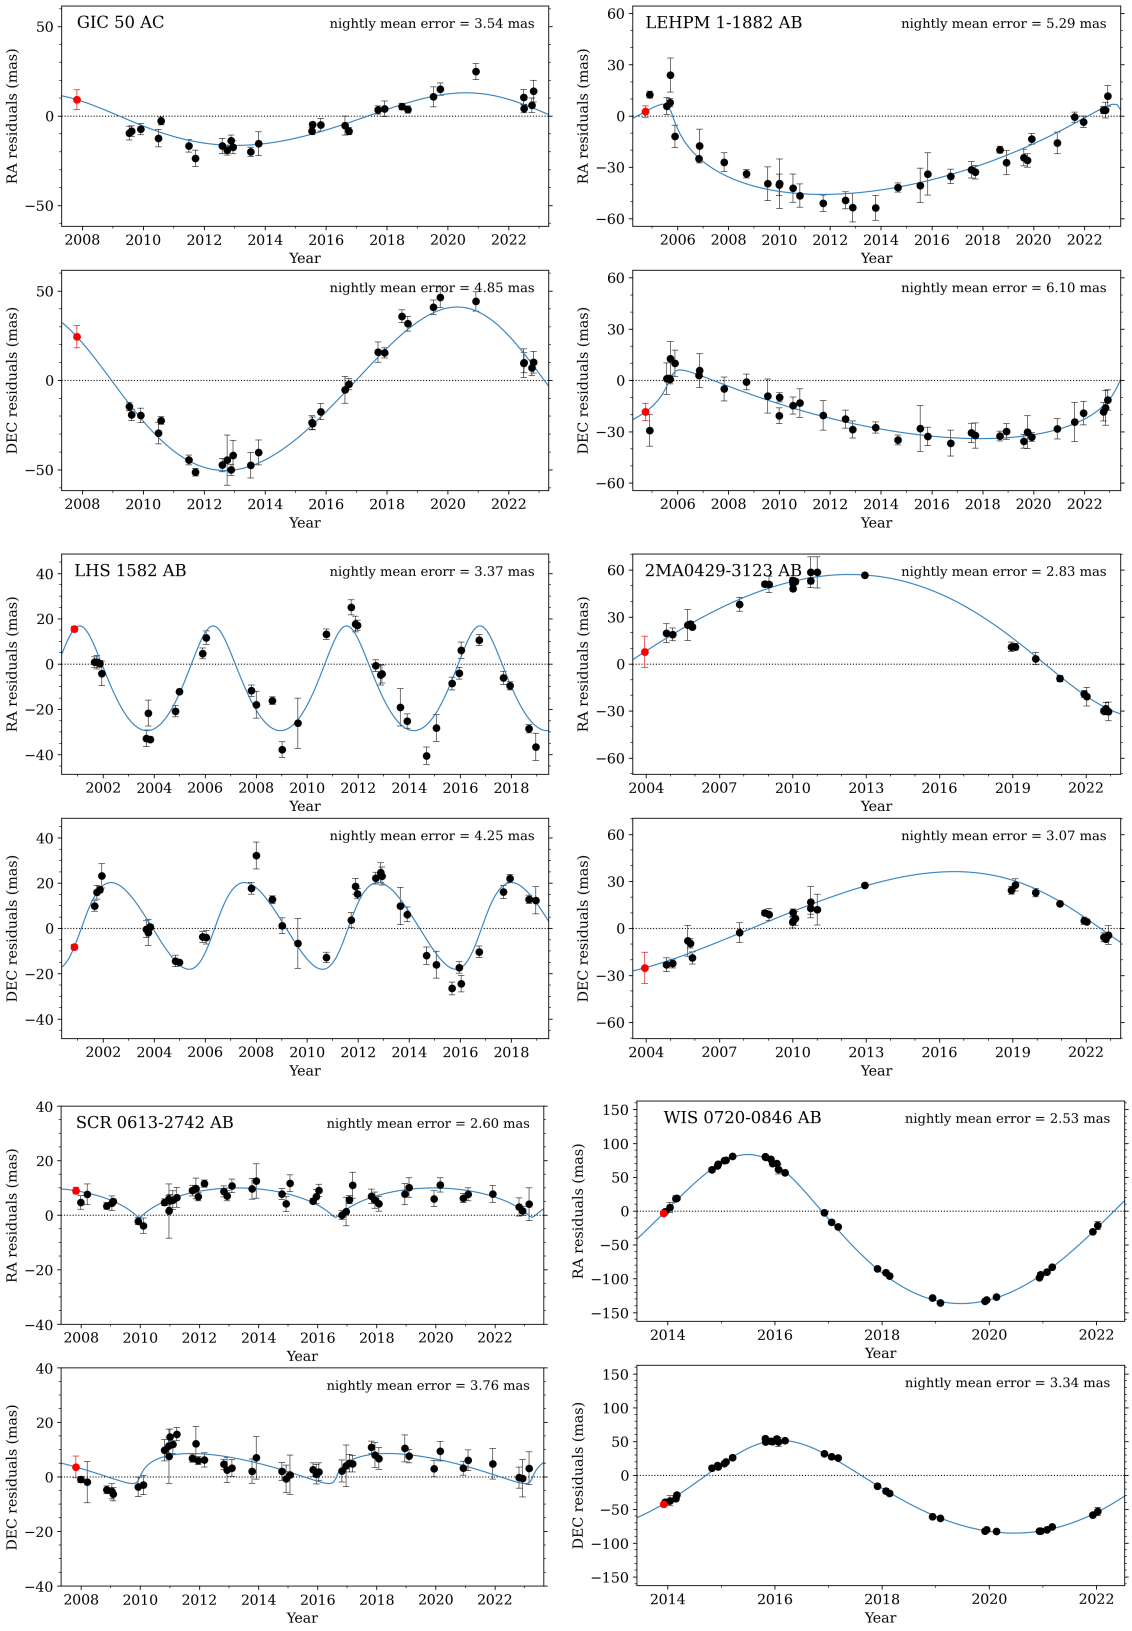

Figure 5.2 Orbits fit from RECONS astrometry, shown here in R.A. order. For each system, the upper panel shows residuals with respect to R.A. over time, the bottom panel shows residuals with respect to Decl. over time, and the blue curve shows the orbit model fit. For every system, proper motion and parallax were fit simultaneously with orbital motion (§3.4).

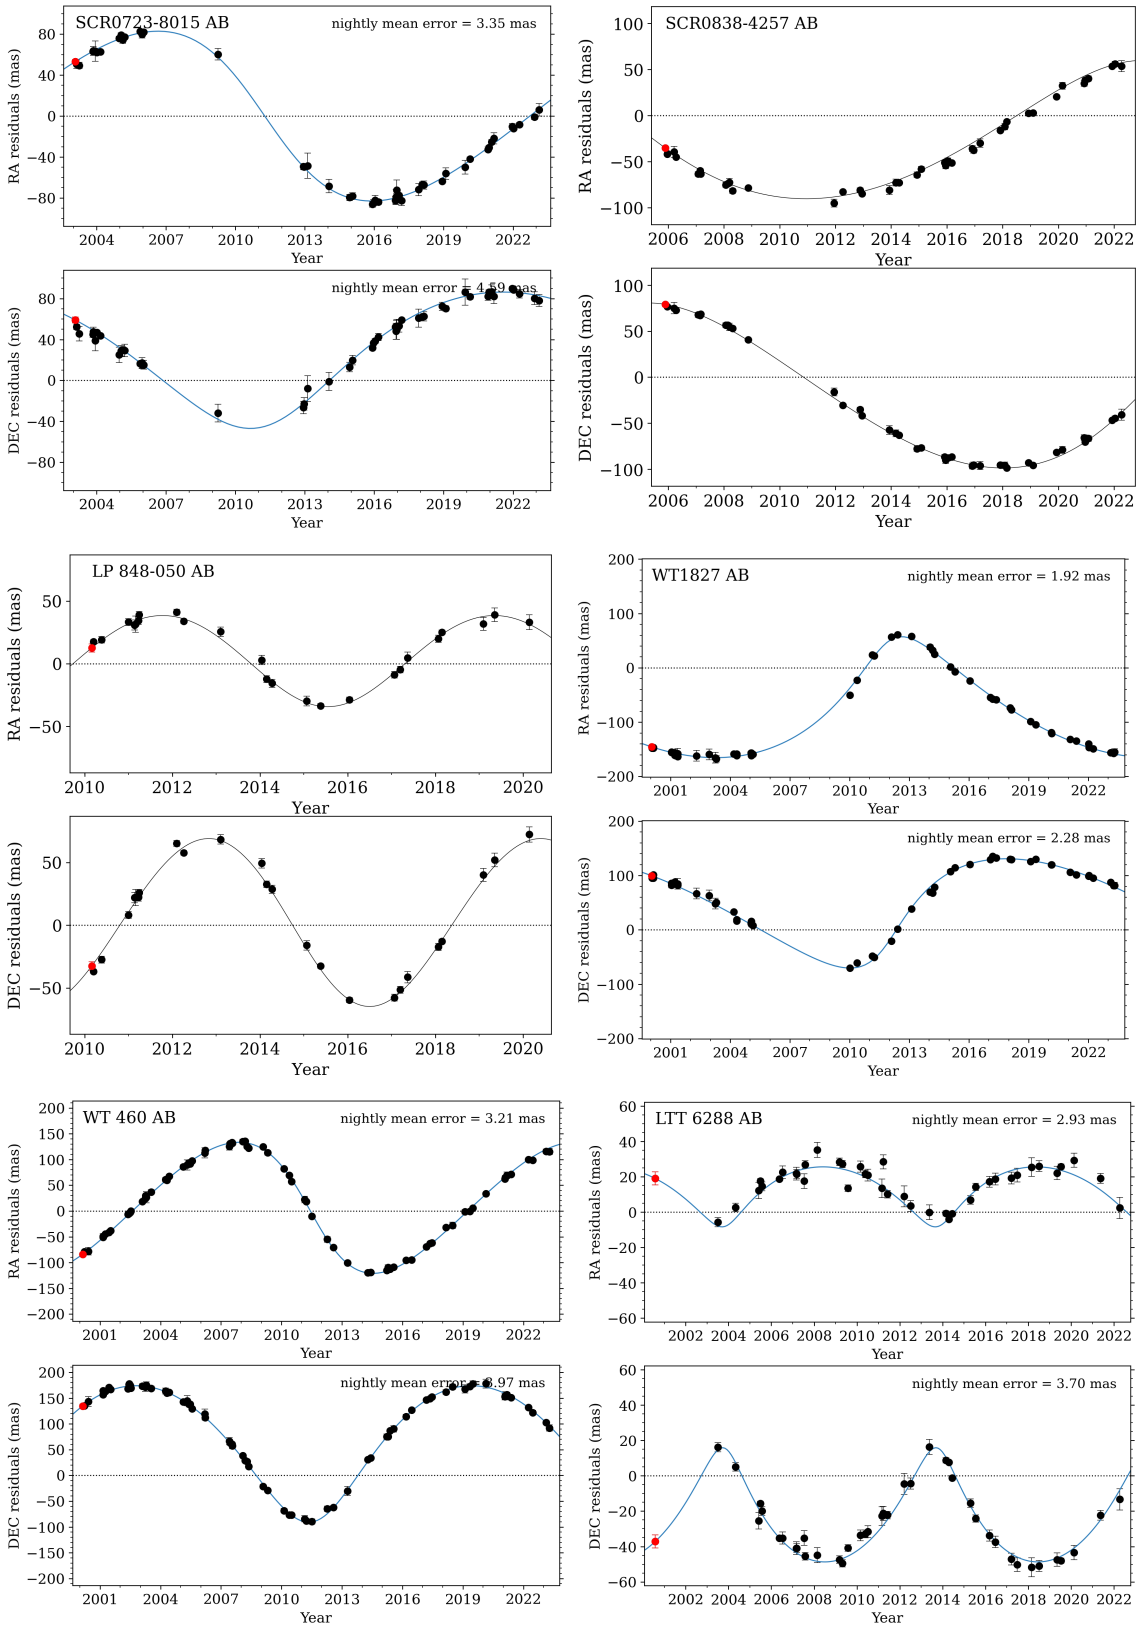

Figure 5.3 Orbits fit from RECONS astrometry, shown here in R.A. order. For each system, the upper panel shows residuals with respect to R.A. over time, the bottom panel shows residuals with respect to Decl. over time, and the blue curve shows the orbit model fit. For every system, proper motion and parallax were fit simultaneously with orbital motion (§3.4).

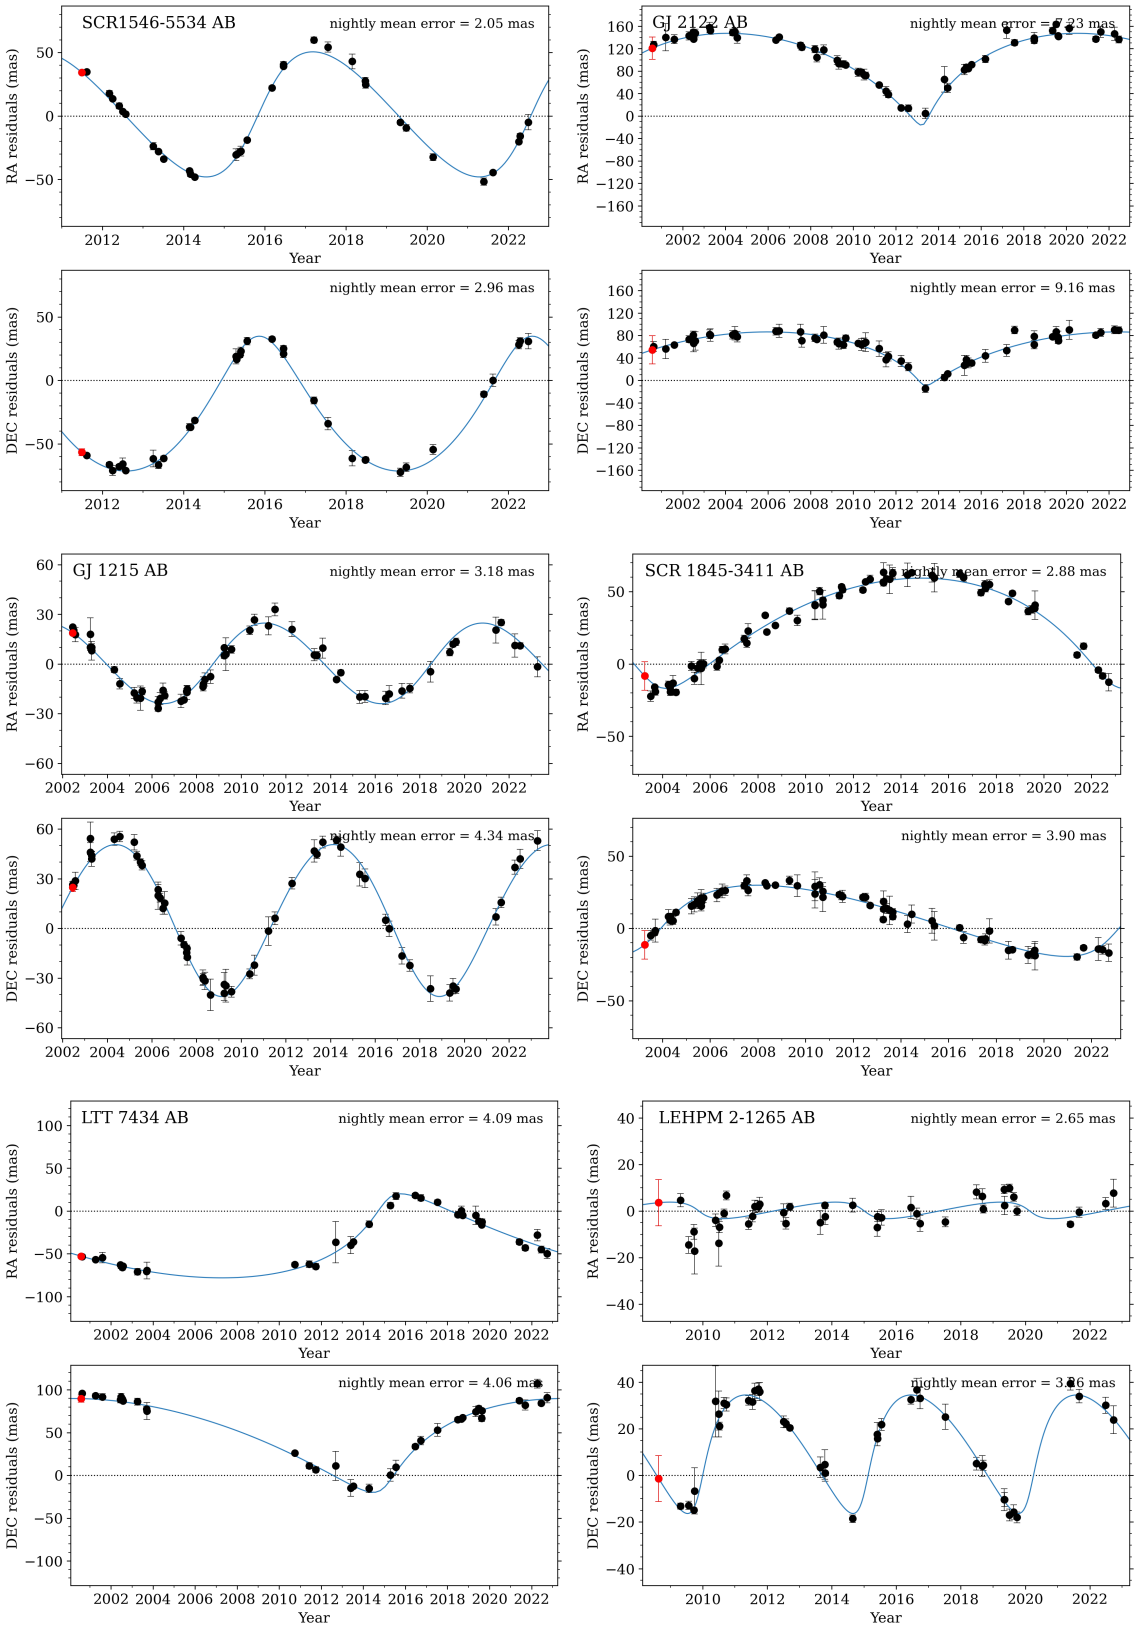

Figure 5.4 Orbits fit from RECONS astrometry, shown here in R.A. order. For each system, the upper panel shows residuals with respect to R.A. over time, the bottom panel shows residuals with respect to Decl. over time, and the blue curve shows the orbit model fit. For every system, proper motion and parallax were fit simultaneously with orbital motion (§3.4).

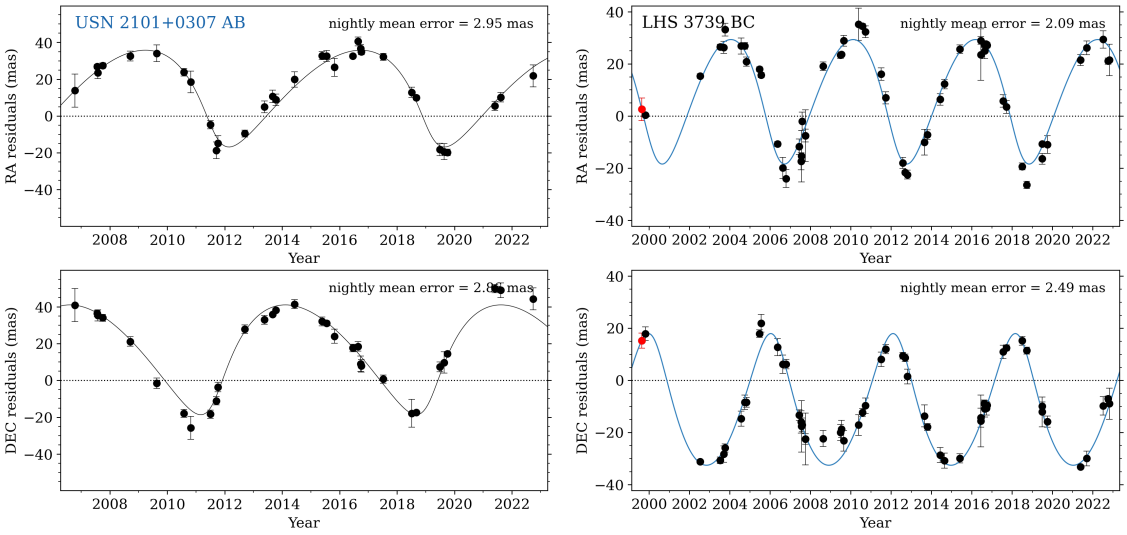

Figure 5.5 Orbits fit from RECONS astrometry, shown here in R.A. order. For each system, the upper panel shows residuals with respect to R.A. over time, the bottom panel shows residuals with respect to Decl. over time, and the blue curve shows the orbit model fit. For every system, proper motion and parallax were fit simultaneously with orbital motion (§3.4).

Table 5.1: Parameters for every orbit in this work, with the source of the orbit listed in column (1): RECONS refers to RECONS astrometry (Chapter 3), SOAR is SOAR speckle interferometry (Chapter 4), and the 5-letter reference codes are defined in Table 1. Note that for all the RECONS orbits and many of the literature orbits, the semi-major axis given is photocentric rather than relative. This table is given in full at `http://astro.gsu.edu/~vrijmoet/PhDthesis/orbit_params.txt`.

Orbit source (1)	System name (2)	Pair (3)	Num. comp. (4)	a (5)	a unit (6)	$P_{\rm orb}$ (7)	$P_{\rm orb}$ unit (8)	e (9)	i (deg) (10)	Ω (deg) (11)	ω (deg) (12)	T_0 (13)	T_0 unit (14)
SOAR	G131-026	AB	2	0.1475	as	5.91174	yr	0.0767	142.32	83.05	247.73	2019.0110	yr
				± 0.0017		± 0.01612		± 0.0104	± 1.37	± 2.79	± 7.39	± 0.0940	
SOAR	LEHPM1-0255	AB	2	0.0956	as	3.35393	yr	0.5936	70.38	130.91	177.86	2020.0631	yr
				± 0.0018		± 0.01862		± 0.0144	± 1.17	± 1.46	± 5.22	± 0.0394	
Ben16	GJ1005	AB	2	0.3037	as	1666.1	d	0.364	146.1	62.8	-13.4	49850.4	JD
				± 0.0005		± 2.5		± 0.001	± 0.2	± 0.4	± 0.5	± 0.8	
SOAR	2MA0015-1636	AB	2	0.1109	as	4.24139	yr	0.0196	67.83	90.98	66.90	2020.8500	yr
				± 0.0038		± 0.04438		± 0.0246	± 3.21	± 1.42	± 44.10	± 0.5172	
Bar18	GJ1006	AC	3	——	–	3.956523	d	0.00220	——	——	28	709.24	JD
				——	–	± 0.000092		± 0.00096	——	——	± 23	± 0.25	
Xia19	GJ2005	BC	3	3.80	au	18.2	yr	0.079	63.8	13.6	184.0	2005.5	yr
				± 0.05		± 0.5		± 0.04	± 0.05	± 0.05	± 0.5	± 0.5	
Dup17	LP349-025	AB	2	145.99	mas	7.698	yr	0.0468	117.36	36.64	262.2	2457758	JD
				± 0.18		± 0.014		± 0.0019	± 0.11	± 0.10	± 1.8	± 15	
Ben16	GJ0022	AC	3	0.5106	as	5694.2	d	0.163	43.7	178.3	104.5	49850.4	JD
				± 0.0007		± 14.9		± 0.002	± 0.2	± 0.2	± 0.5	± 0.8	
RECONS	GIC0050	AC	3	48.301859	mas	5191.2065	d	0.15382901	87.631763	17.662439	47.942086	2449184.4	JD
				± 2.5628136		± 223.84258		± 0.074894763	± 1.3273227	± 1.2328539	± 19.733370	± 489.94430	

5.1.2 Sensitivity of the RECONS Astrometry Data

To understand how well the RECONS contribution of orbits represents the full distribution of orbits in the solar neighborhood, we have undertaken an analysis of the sensitivity of the RECONS astrometry as a function of $P_{\rm orb}$ and e. This sensitivity analysis required simulating 2500 orbits for each system observed and counting how many of those orbits would create astrometric signals we could detect — effectively an injection-recovery procedure.

The simulations for the sensitivity analysis were generated as follows:

1. or every target in the RECONS volume-complete sample ($\pi < 60$ mas and Decl. $< 0°$), we calculated its absolute K magnitude. We then identified every combination of primary and secondary mass (M_1 and M_2) that could have a combined M_K matching that measured value according to our corrected K-band MLR (§2.2.1).

   To ensure reasonable computation time, before beginning this analysis we generated a grid of combined M_K values corresponding to each possible M_1 and M_2 combination. M_1 was varied between 0.60 $M_\odot$ and 0.08 $M_\odot$, and M_2 between 0.60 $M_\odot$ and 0.06 $M_\odot$, both in steps of 0.01 $M_\odot$. No mass combinations were duplicated in this grid — i.e., for the row corresponding to a given M_1, only the columns for which $M_2 < M_1$ were populated (the grid was thus triangular in shape).

   An extended grid was also calculated to use when a system could not find a match on the standard grid. This extended grid included masses every 0.001 $M_\odot$ below 0.1 $M_\odot$, to ameliorate the fact that absolute magnitude changes so rapidly with respect to mass below 0.1 $M_\odot$.

2. For every compatible M_1 and M_2 combination, we calculated the threshold at which a PB would be detectable as three times the target's mean epoch uncertainty (i.e., 3σ). This threshold was chosen based on experience evaluating these signals by-eye throughout the program.

3. Each compatible M_1 and M_2 combination also corresponding to a specific mass and flux ratio, given as B and β in Eq. 3.1 (§3.4). These values were later required to calculate the photocentric PB of each simulated unresolved pair.

4. For every compatible (M_1, M_2) pair, we then stepped through each potential ($P_{\rm orb}$, e) combination in our $P_{\rm orb}$ vs. e parameter space of interest (0–30 years in $P_{\rm orb}$ and 0.0–1.0 in e). To reduce computation time, the step sizes were 2 years for $P_{\rm orb}$ and 0.1 for e. For each of those potential orbits we calculated the expected PB of 2500 random snippets of randomly-oriented orbits using the following procedure:

   (a) The semi-major axis of the simulated orbit was computed from Kepler's law given the M_1, M_2, and $P_{\rm orb}$ being considered. The distance to the system was then used to scale this semi-major axis from AU to arcseconds.

   (b) An orbit orientation was chosen by selecting a random value of i and Ω between 0° and 180° and a random ω between 0° and 360°. To reduce computation time, these angles were restricted to multiples of 10°. These random angles were then converted to the corresponding Thiele-Innes constants (§3.4.2).

   (c) A random point within the orbit was chosen as the starting point of simulated

observations. Because a companion moves quickly through periastron and thus is more likely to be observed at apastron, the potential starting point angles were evenly spaced but were each weighted by the time the companion takes to reach the following angle. The random starting point was thus drawn from this weighted distribution of angles. The procedure thus takes into account the role of e in the orbit's observability.

(d) With all orbit parameters now determined, the companion's positions on the sky were computed from the chosen starting point to the end of the simulated observations, with that end point determined by the number of RECONS observations actually obtained for that target to date. An example for one simulated circular system is shown in the top panel of Figure 5.6 and for one simulated eccentric system in the top panel of Figure 5.7.

(e) To mirror the process we use with real observations to detect PBs, we computed residuals of a single-star astrometric fit to this simulated orbit. We omitted parallactic motion from these simulations because experience has shown that its contribution to the single-star residuals is minimal even when there is dramatic orbital motion[1] The proper motion fit was then assigned by a fitting a line to the simulated motion in R.A. vs. time and another line to Decl. vs. time. These fits are illustrated in the middle panels of Figures 5.6 and 5.7 with red dashed lines.

[1]This is primarily because the parallactic residuals are much smaller than the semi-major axis of the photocentric orbit ($\sim$1.5 mas vs. $\gtrsim$20 mas). These residuals are small because the parallactic ellipse wraps exactly once per year and we typically have many years of observations, so the ellipse is exceedingly well-determined.

(f) We then multiplied the residuals to those fits, shown in the bottom panels of Figures 5.6 and 5.7, by the factor $B - \beta$ determined earlier for this (M_1, M_2) combination. That scaling shrunk the residuals to represent the photocentric orbit and create the final simulated PB.

(g) Finally, we computed the amplitude of the simulated PB residuals and compared it to the threshold of detectability that was determined earlier for this system based on its mean nightly uncertainties (step 2). This random snippet of randomly-oriented orbit was thus counted as either a detection or non-detection for this ($P_{\rm orb}$, e) combination.

(h) After the above process was completed for one random starting point, steps 4c–4g were repeated until 50 random starting points (and associated orbit snippets) had been simulated. Then a new random orientation was selected, restarting the process at step 4b, until 50 random orientations had been explored. The end result was a count of the total orbits detected, out of 2500 simulated data sets, for that single ($P_{\rm orb}$, e) combination.

5. The result of step 4 was a full accounting of what percentage of orbits would be detected for every ($P_{\rm orb}$, e) for that (M_1, M_2). The process was repeated for the next (M_1, M_2) combination compatible with the measured combined M_K. There were 5–15 of these compatible mass combinations for most systems.

6. After all compatible mass combinations had been simulated, we took the average of

their $P_{\rm orb}$ vs. e sensitivity results, generating a single $P_{\rm orb}$ vs. e sensitivity plot for each target. Two examples are given in Figure 5.8.

This process produced an assessment of sensitivity in $P_{\rm orb}$ vs. e space for each individual target. To assess the observing program as a whole, we then combined these single-target results by averaging over all the plots from the RECONS volume-complete sample that had enough data to be considered, which numbered 258 M dwarfs within 16.67 pc. The result, shown in Figure 5.9, indicates that for the volume-complete sample the RECONS program detects $\gtrsim$70% of companions as long as they have $P_{\rm orb} < 30$ years. The overall coverage is quite good, with only a few "softer spots" at $P_{\rm orb} < 3$ years and at the longest-period + highest-eccentricity orbits. This validates that the volume-complete sample can be used to determine multiplicity statistics.

Not every detectable orbit in Figure 5.9 can be characterized, however, as the length of the observing baseline places a hard limit on our ability to distinguish proper motion from orbital motion (discussed further in §3.4). Additionally, the RECONS orbits presented in this dissertation are drawn not only from the volume-complete 16.67 pc sample, but also from targets we monitor out to 25 pc. The targets between 16.67 pc and 25 pc cannot be used to infer multiplicity statistics because they are only included if they are "interesting" (e.g., if they are multiple or exhibit significant photometric variability); thus, the more distant sample is affected by numerous selection biases. We can sidestep that potential overrepresentation of multiples, however, if we redefine our question as: *Given that a system is multiple, in which regions of $P_{\rm orb}$ vs. e space are we able to fit orbits?*

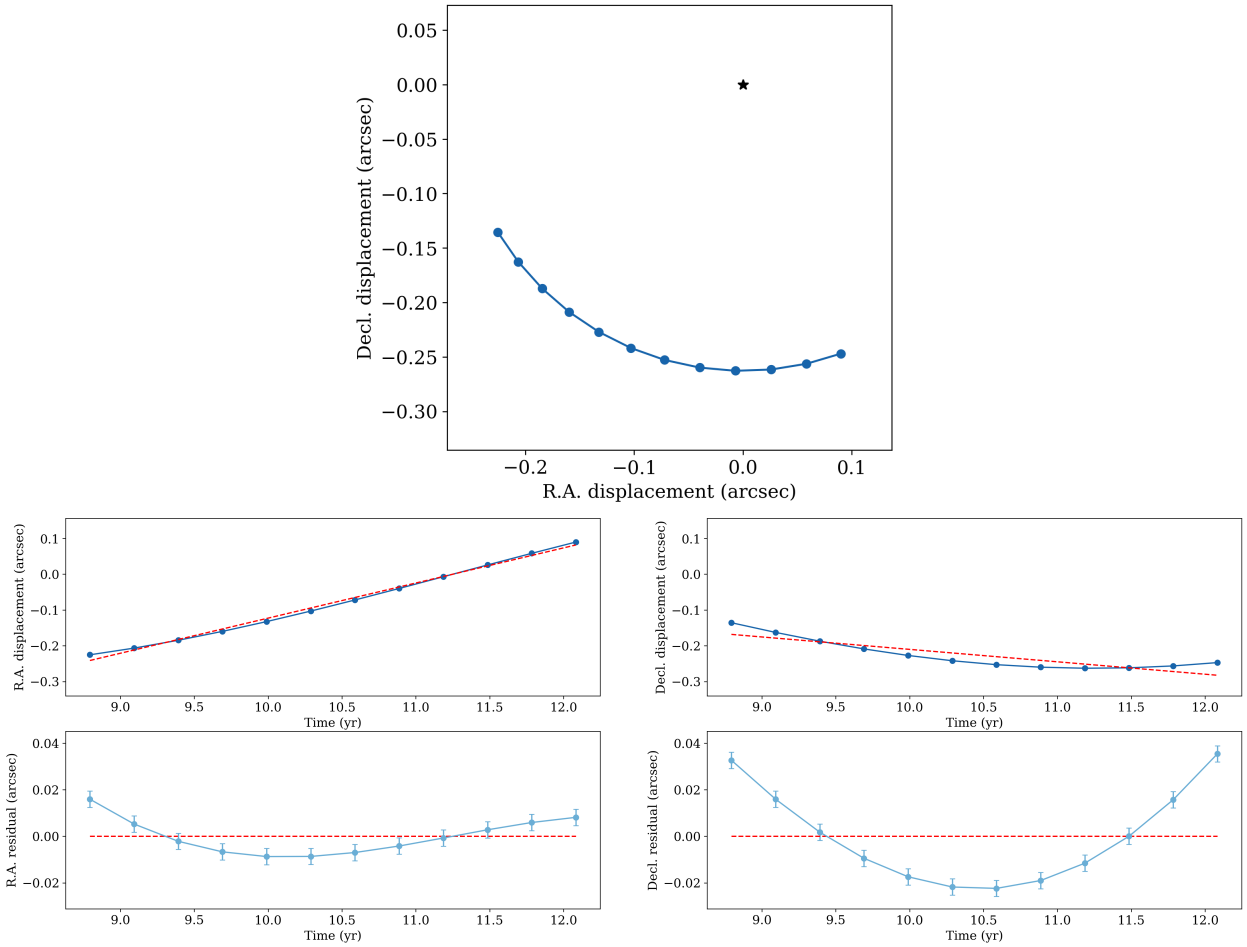

Figure 5.6 Example of a circular orbit used in the simulations that underpin the RECONS sensitivity plots. *Top panel:* simulated positions of a companion in a randomly oriented orbit on the sky, with each point representing its position at the time of RECONS observation. For the purpose of this demonstration, the random inclination angle were restricted to ±30° of face-on. *Middle panels:* simulated R.A. and Decl. displacements vs. time for the orbit snippet shown in the top panel. The red dashed line is the linear fit to the astrometry, representing the RECONS single-star model fit to the photocenter's proper motion. Note that parallactic motion is not simulated nor fit here because those residuals are usually much smaller than the orbital motion (see step 4e). *Bottom panels:* simulated R.A. and Decl. residuals to the single-star fit, with the red line the same as in the middle panels. Error bars representing typical uncertainties are also illustrated now (and omitted from the top and middle panels for clarity). This residual is then multiplied by the $B - \beta$ factor for this (M_1, M_2) combination to produce the final simulated astrometric PB (not shown).

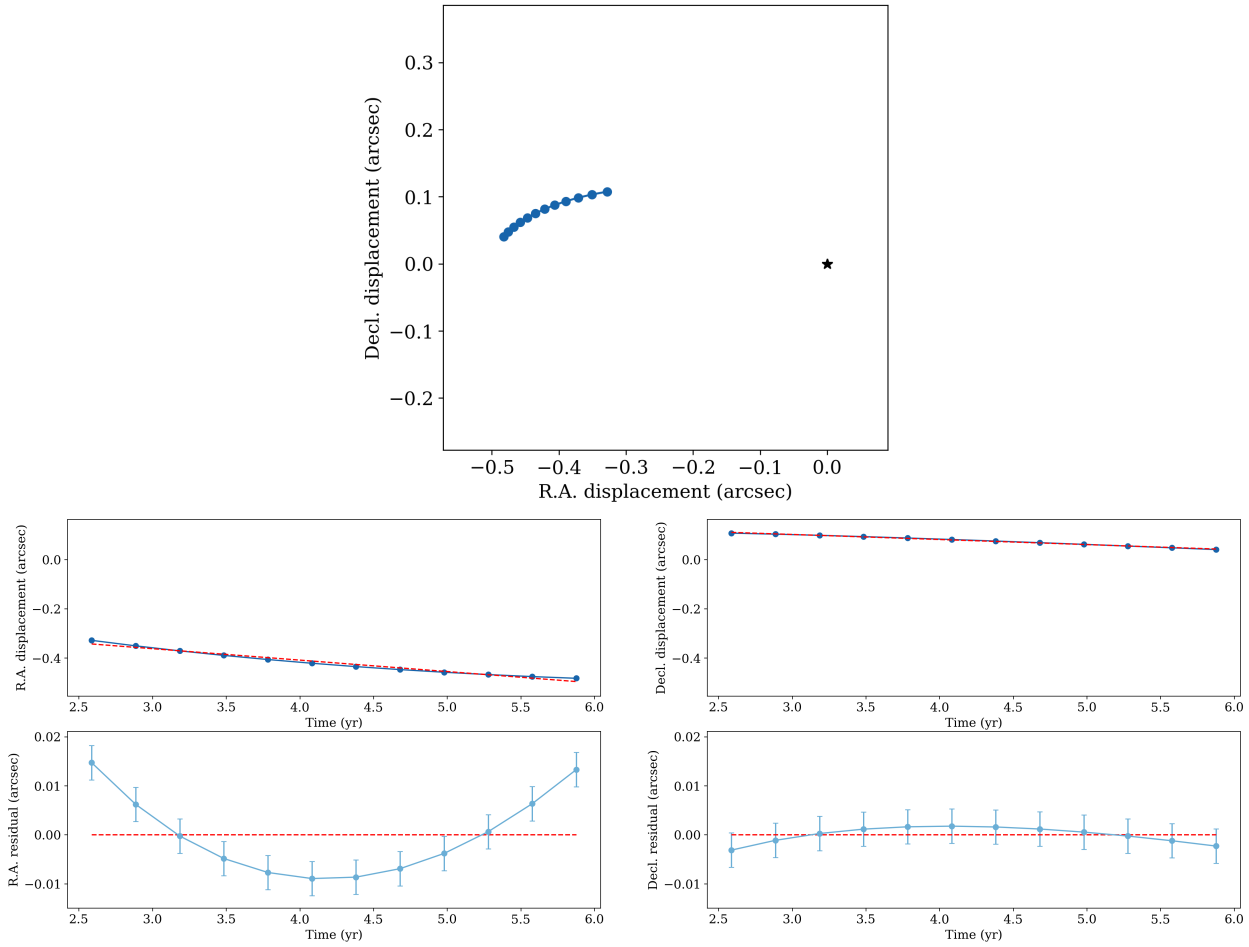

Figure 5.7 Example of an $e = 0.9$ orbit used in the simulations that underpin the RECONS sensitivity plots. Each panel is the same as in Figure 5.6, with the only differences being this orbit's higher eccentricity and different randomly selected orientation angles.

To address this revised question of fitting capability, we completed the above sensitivity analysis using only confirmed and suspected multiples monitored by RECONS. This analysis included 127 multiples, after excluding 5 systems separated by $\gtrsim 0\farcs8$ because our current procedure for these systems does not produce astrometry suitable for orbit characterization (see §5.1.1)[2]. We then zeroed out each multiple's sensitivity plot for $P_{\rm orb}$ greater than

[2]This expanded sample of 127 systems is used in lieu of the 78 with clearly traced orbital motion because many of these systems are retained on the program only because of their astrometric PBs, regardless of

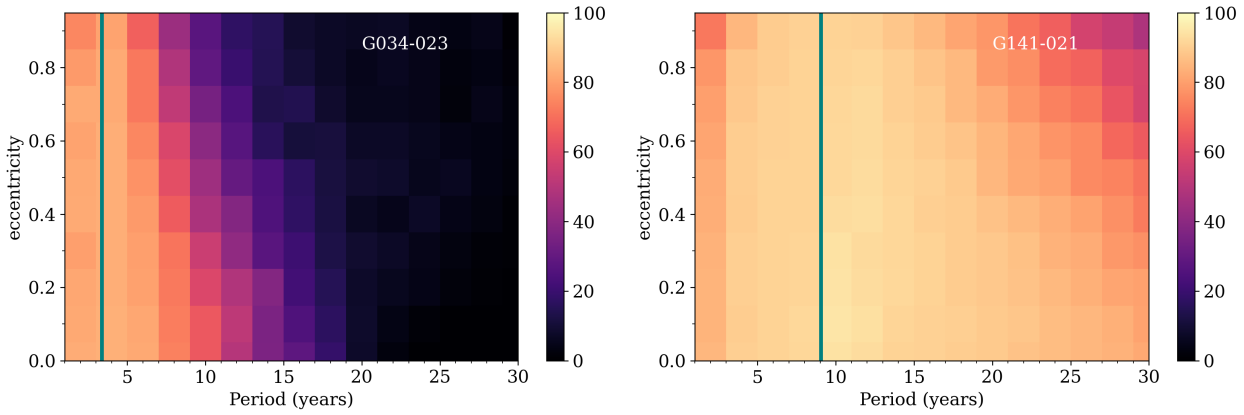

Figure 5.8 Examples of average sensitivity plots for two single stars: (*left panel*) G034-023, a star with limited coverage over only ∼3.5 years, and (*right panel*) G141-021, a star with modest coverage over ∼9 years. The color indicates the percentage of simulated orbits that were detectable given that system's typical nightly astrometric uncertainty. The teal vertical line in each plot indicates the length of RECONS observations for that system to date; orbits longer than those baselines may be detectable, but cannot usually be fit reliably.

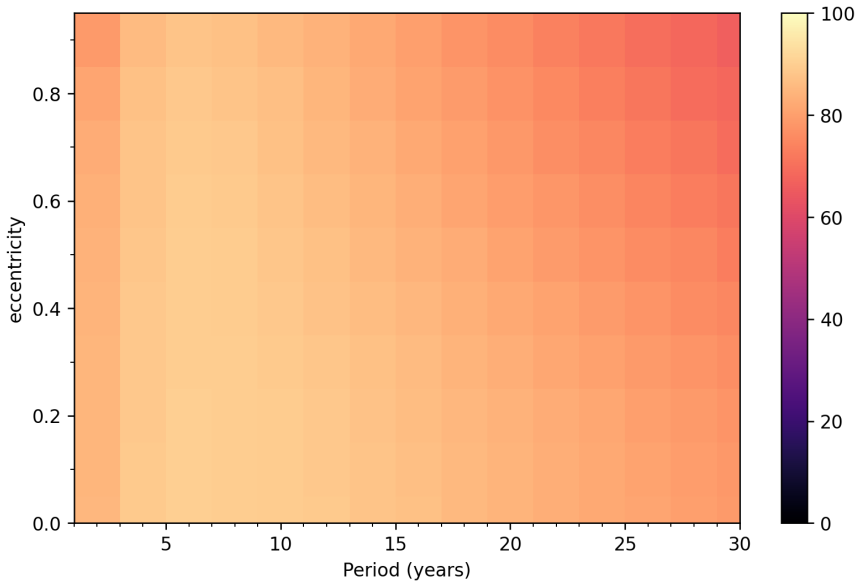

Figure 5.9 Average detection sensitivity in $P_{\rm orb}$ vs. e space for the subset of the RECONS 16.67 pc volume-complete sample which had enough data to be considered (258 red dwarfs). The colors represent the percentage of orbits that were detectable, calculated via the injection-recovery process described in §5.1.2 averaged over all systems in the sample. It is evident that at least 70% of orbits would have been detected for orbital periods of $\leq$30 years at effectively all eccentricities.

110% of that target's observed baseline. This effectively transformed the plot from *detection sensitivity* to *capability of an orbit fit*. We then re-computed the average over all the multiples monitored.

The left panel of Figure 5.10 shows the fraction of orbits we are able to fit, given that a system is multiple (or potentially multiple), across $P_{\rm orb}$ vs. e space. The plot varies very little in e because the detection sensitivity for this sample of multiples is nearly identical to Figure 5.9 due to our generally excellent time coverage of these systems (right panel of Figure 5.10); cutting off systems at 110% their time baseline to translate that plot to fit capability then creates the striped pattern of the final plot. This result shows that as expected, we are best at fitting binaries with orbital periods of $\sim$5–10 years, and worst at fitting very short and very long $P_{\rm orb}$. This efficacy distribution is driven by the decades-long RECONS data sets for most of these systems; overall, the median observing baseline for these multiples is 18.26 years. The right panel of Figure 5.10 shows the multiples considered in this analysis on the HRD, to give a sense of how well they represent the M dwarf main sequence.

Overplotted on Figure 5.10 are the $P_{\rm orb}$ vs. e values for the 19 orbits characterized for this dissertation. Although the distribution is sparse, the number of stars in each 5 year bin of $P_{\rm orb}$ does increase somewhat from 30 $\rightarrow$ 5 years, following the prediction of the fit capability map. Orbits with $P_{\rm orb} \lesssim 3$ years are missing because the MCMC fitting algorithm currently does not usually converge on e in those cases. This issue is related to the low

whether those PBs are clearly orbital (78 systems) or questionable in nature but not consistent with single-star motion (49 systems). Internally within RECONS these sets are known as "PB!" and "PB?", respectively.

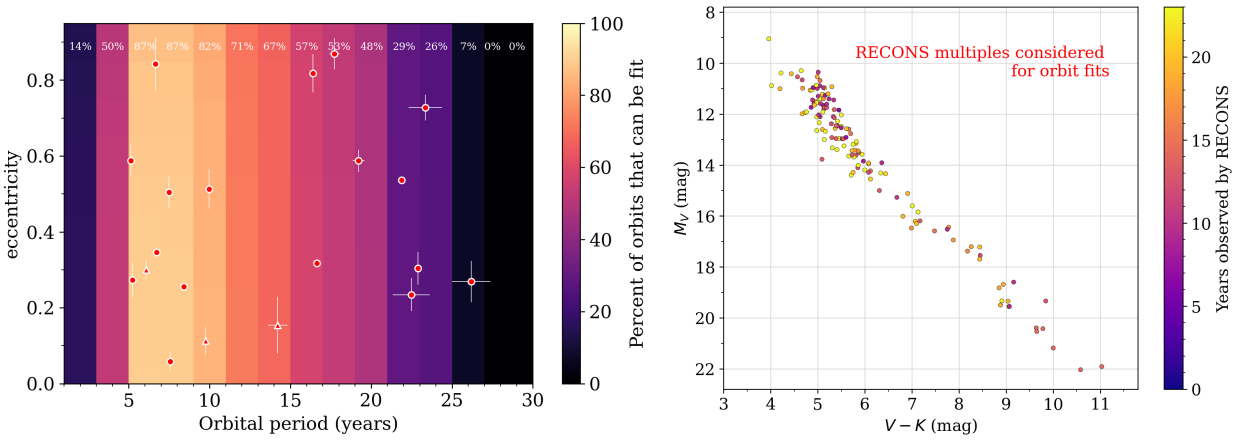

Figure 5.10 *Left panel:* average percentage of orbits that can be fit, as a function of $P_{\rm orb}$ vs. e, for 127 multiples and suspected multiples followed by the RECONS astrometry program. For clarity, the exact values for the patches in the topmost row are printed in white. This fit capability map was created by averaging the fit capability maps of the individual multiples (following §5.1.2). The variation is quite low in e because coverage is generally excellent for these systems and their detectability plot was nearly uniform, as noted in the text. Overplotted are $P_{\rm orb}$ vs. e values for the orbits fit from RECONS astrometry, with circles for binaries and triangles for subsystems of triples and quadruples. *Right panel:* Hertzsprung-Russell diagram of the multiples analyzed in this sensitivity analysis, with each point's color indicating the time baseline over which it has been observed.

S/N of those systems (their orbits have relatively small photocentric semi-major axes) and our sparse observing cadence of only 2–4 epochs per year. Because we could not model this unsolved issue in our sensitivity analysis, here we have assumed that only 20% of orbits with $P_{\rm orb} \leq 3$ years can be fit, as experience has shown that that is the fraction that are unaffected by this non-convergence issue.

5.2 The SOAR Speckle Data

As outlined in Chapter 4, the SOAR speckle program imaged 337 potential multiples between July 2019 and March 2023, and continued observing 225 of these systems to map their

companions' motions. Of those resolved pairs, 42 yielded enough data (in combination with other imaging in the literature) to enable reliable orbit fits. In §5.2.1 we present these SOAR speckle orbits and their basic demographics, and in §5.2.2 we analyze how well those orbits represent the population we were sensitive to in $P_{\rm orb}$ vs. e space.

5.2.1 Orbits from SOAR Speckle Interferometry

The SOAR orbits span 0.7–28 years in $P_{\rm orb}$ and 0.02–0.90 in e, providing an excellent complement to the RECONS orbits by extending toward shorter orbital periods. Many of the systems have orbits mapped using SOAR data previous to our $\sim$4 year program, supplemented by data from the literature. Figure 5.1 shows that these systems are concentrated in the $\sim$0.15–0.60 $M_{\odot}$ area of the main sequence, balancing the RECONS sample in the mass regime as well. Many of the lowest-mass targets on the SOAR program were supplied by the RECONS astrometry program (Chapter 3), which can reach much fainter stars. The SOAR sample was limited mainly by the $I = 14$ mag faintness limit of HRCam+SAM (§**??**).

For this sample from SOAR, we fit every orbit with more than four observations and $\gtrsim$144° of motion in position angle, corresponding to at least 40% of a complete orbit. We omitted orbit fits for a handful of systems that were well-observed but had short periods ($\lesssim$1 year) and already had high-quality orbits available in the literature, as the 180° ambiguity in position angle that affects these speckle data proved to be considerably challenging in those cases. The SOAR program also yielded several orbits with $P_{\rm orb}$ too long to be included in this study (>30 years) that will be incorporated into future efforts.

The SOAR orbits are shown in R.A. vs. Decl. space in Figures 5.11–5.14. As discussed

in §4.4, each system was fit with the IDL code `ORBIT` (Tokovinin 2016). This code fits relative astrometry and (optionally) RV data simultaneously using a Leavenberg-Marquardt algorithm to find the best-fit orbit model.

5.2.2 Sensitivity of the SOAR Speckle Data

To assess the representation of the SOAR $P_{\rm orb}$ vs. e results, we completed a sensitivity analysis that closely mirrored the procedure used for the RECONS astrometry program sensitivity (§5.1.2). For this analysis we omitted the program's detectability assessment (i.e., Figure 5.9), as every system on the SOAR target list was either a suspected multiple or known multiple established from a diverse array of previous data. We thus proceeded to the fitting step, computing the average fit capability plot for SOAR using only the targets with successfully detected companions. The goal of this analysis was to address the question: *For the given set of multiples monitored by the SOAR program, in which areas of $P_{\rm orb}$ vs. e space are we able to fit orbits?*

The simulations were conducted exactly as described in §5.1.2, with a few key differences for this relative astrometry data:

- Because the SOAR data provided ΔI magnitudes for every resolved system, for each pair we calculated the M_I mag of each component and estimated their masses using the I-band pseudo-MLR derived in §2.2.2 (Figure 2.3). This step replaced the grid of combined M_K values and related steps used for the RECONS astrometry simulations (i.e., step 1 and the repetition for different compatible mass combinations).

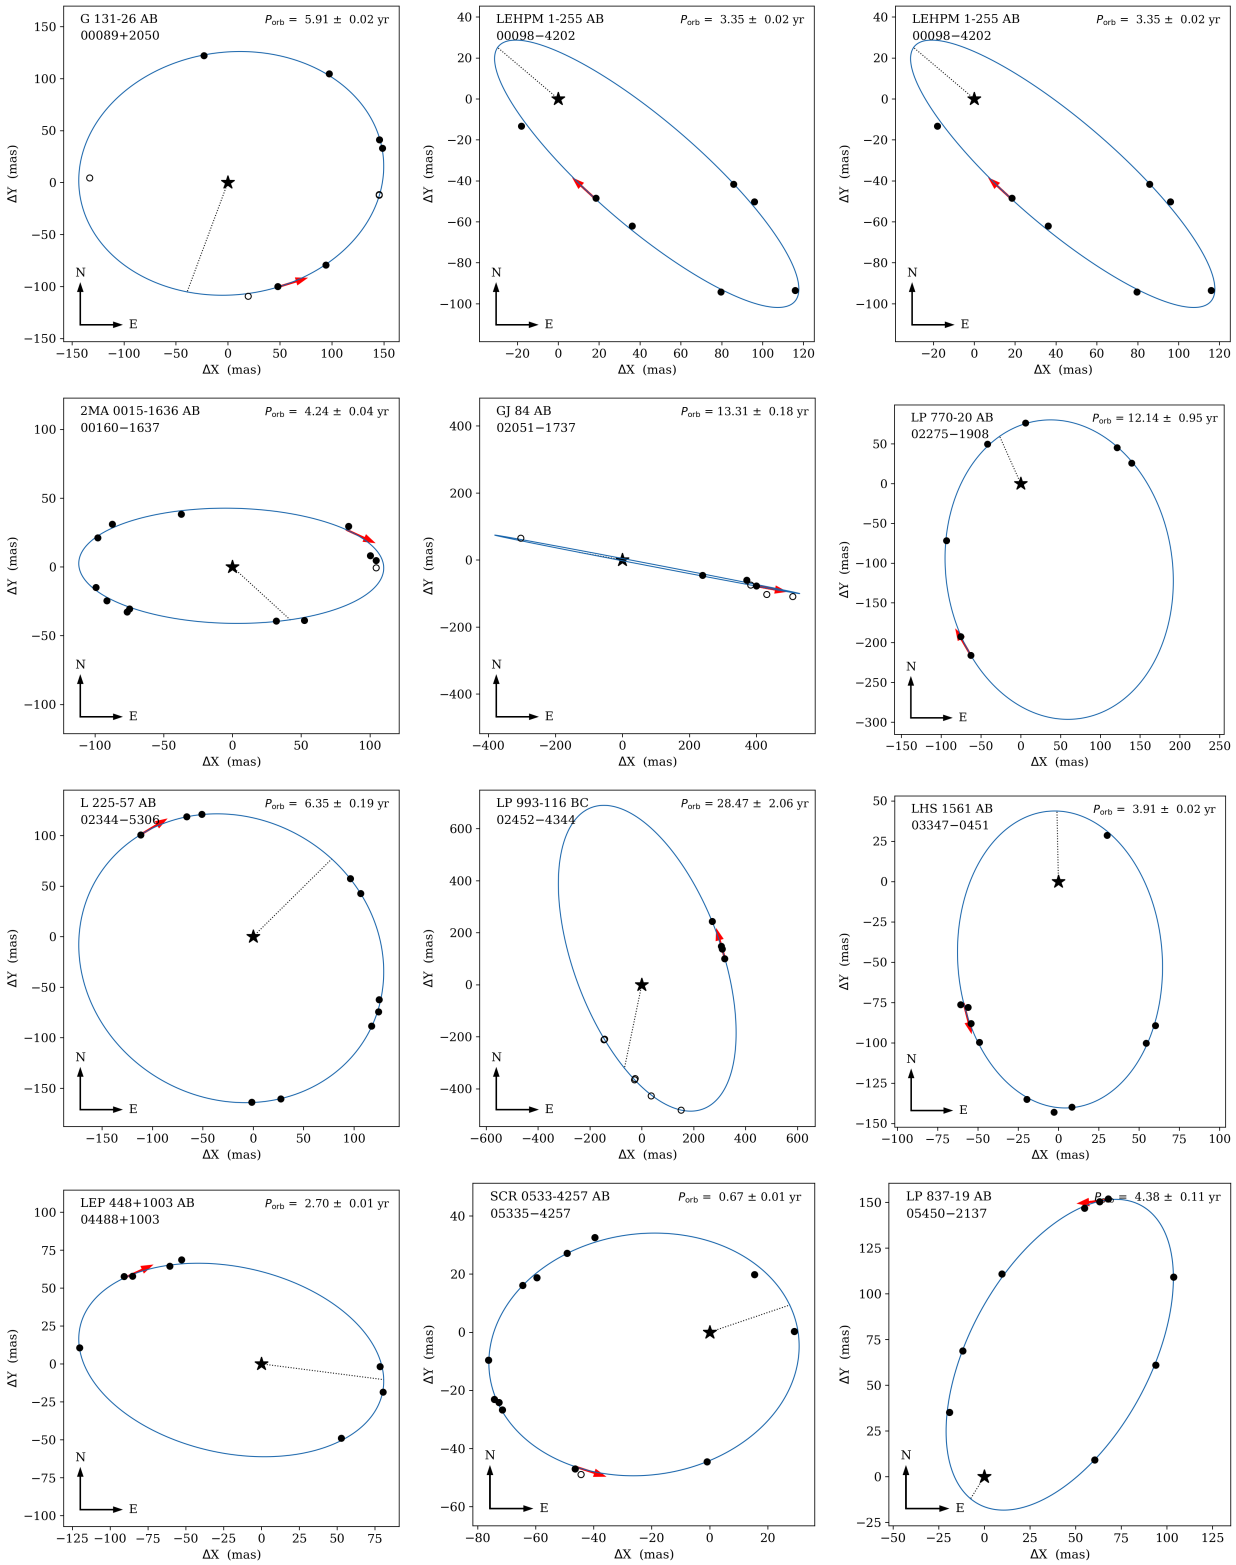

Figure 5.11 Orbits fit to SOAR speckle interferometry combined with existing high-resolution imaging data from the literature (if available). The systems are shown here in order of R.A., with each panel showing R.A. vs. Decl. for the system's SOAR data (black filled points) and literature data (open points) as well as its orbit model fit (blue curve).

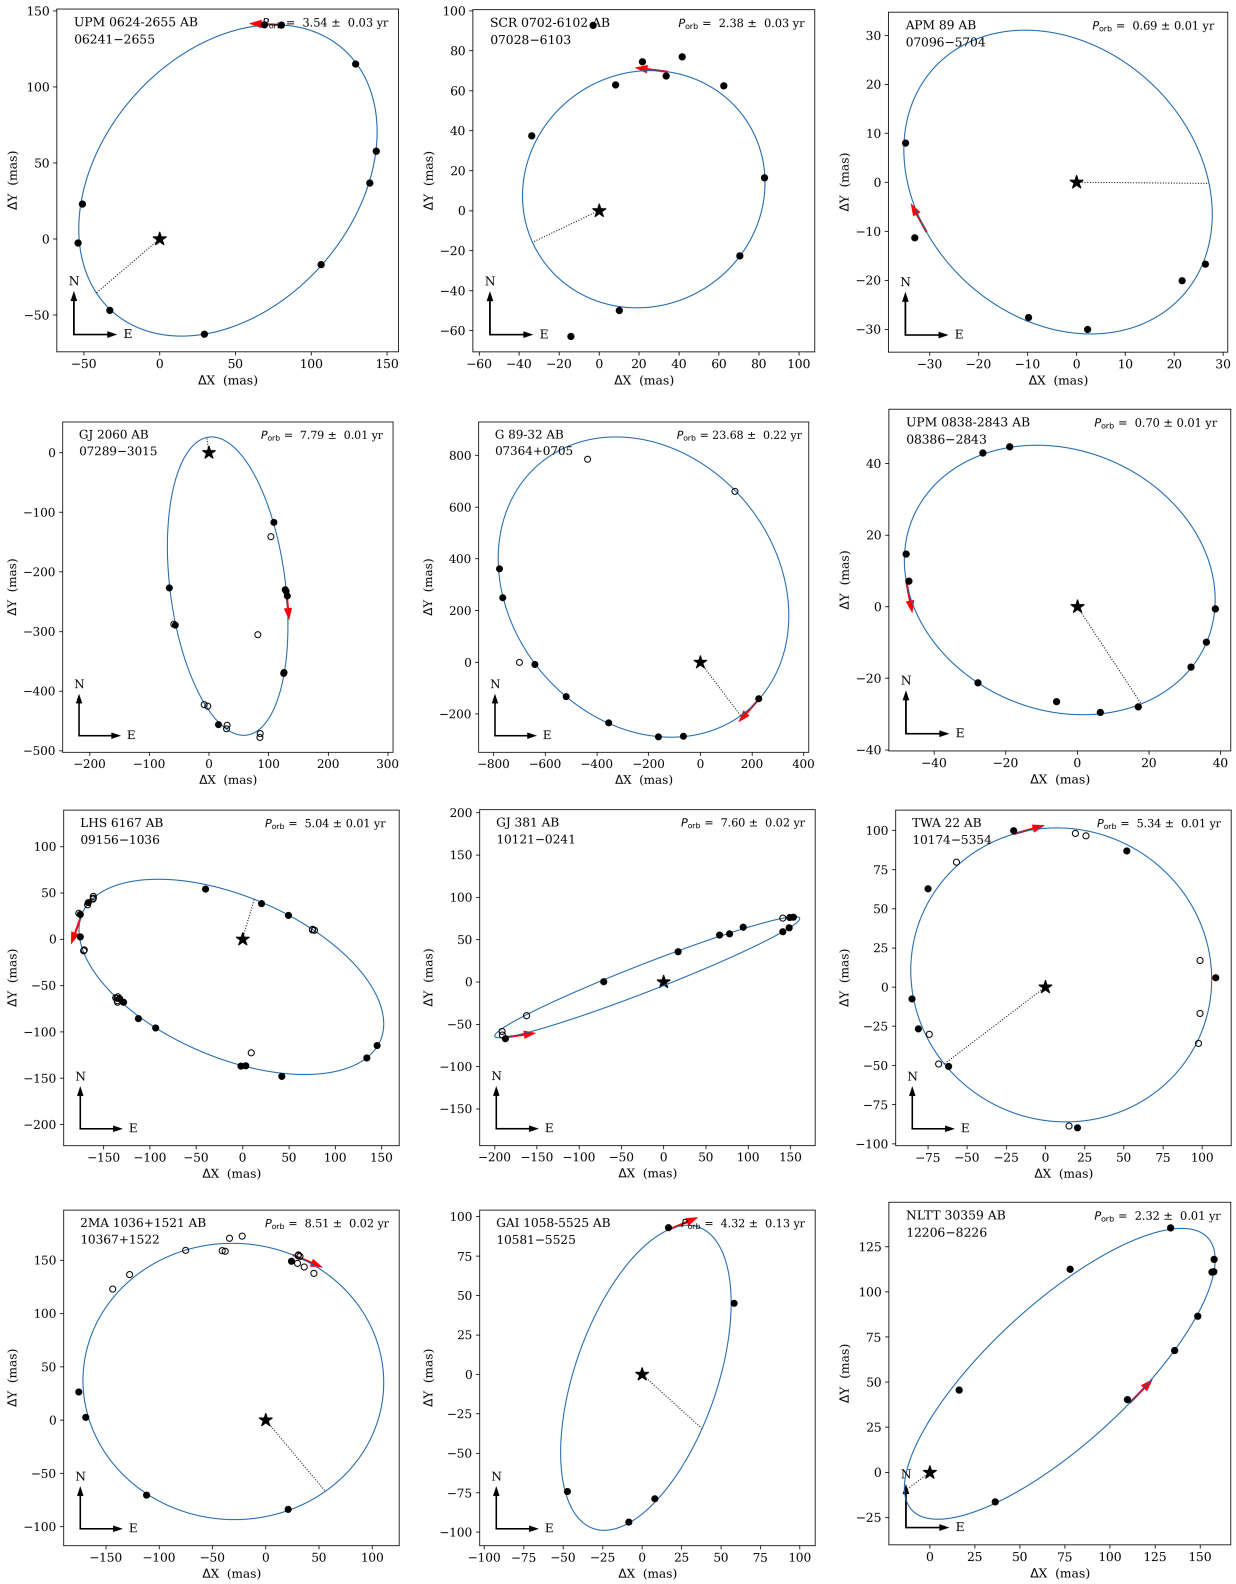

Figure 5.12 Orbits fit to SOAR speckle interferometry combined with existing high-resolution imaging data from the literature (if available). The systems are shown here in order of R.A., with each panel showing R.A. vs. Decl. for the system's SOAR data (black filled points) and literature data (open points) as well as its orbit model fit (blue curve).

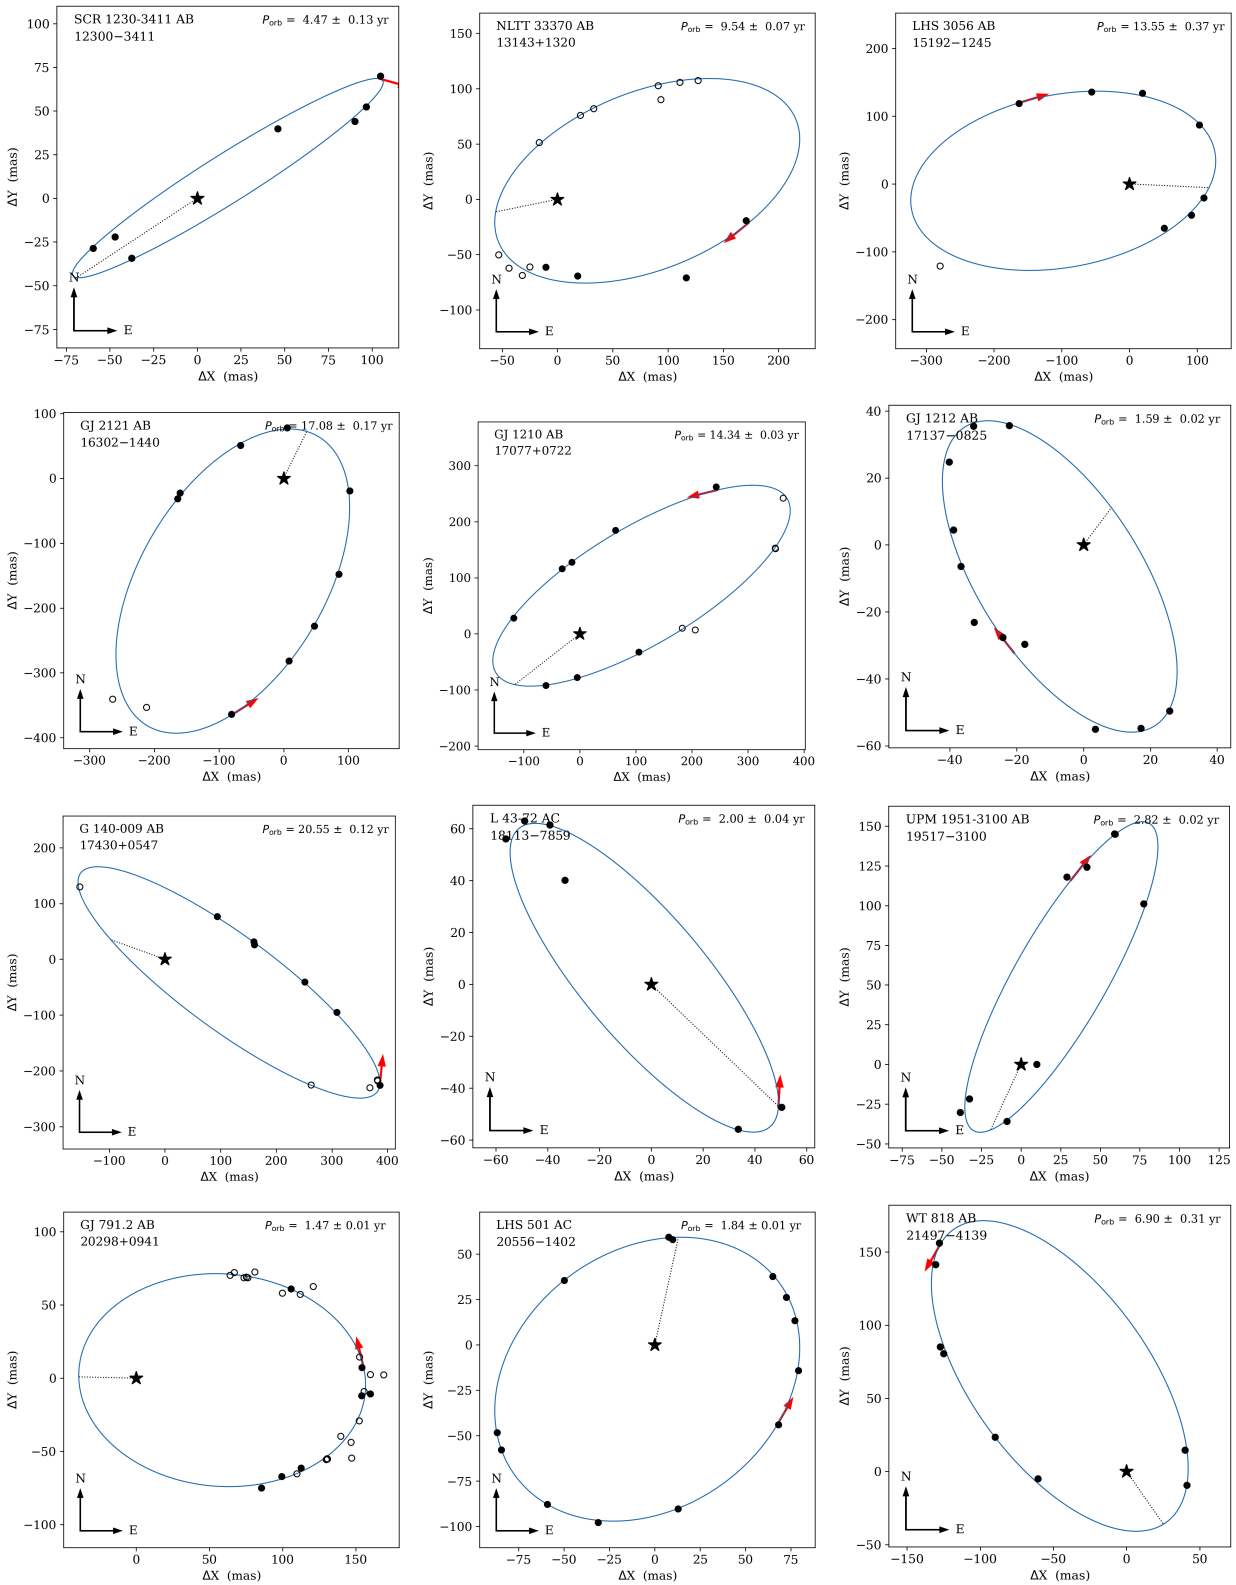

Figure 5.13 Orbits fit to SOAR speckle interferometry combined with existing high-resolution imaging data from the literature (if available). The systems are shown here in order of R.A., with each panel showing R.A. vs. Decl. for the system's SOAR data (black filled points) and literature data (open points) as well as its orbit model fit (blue curve).

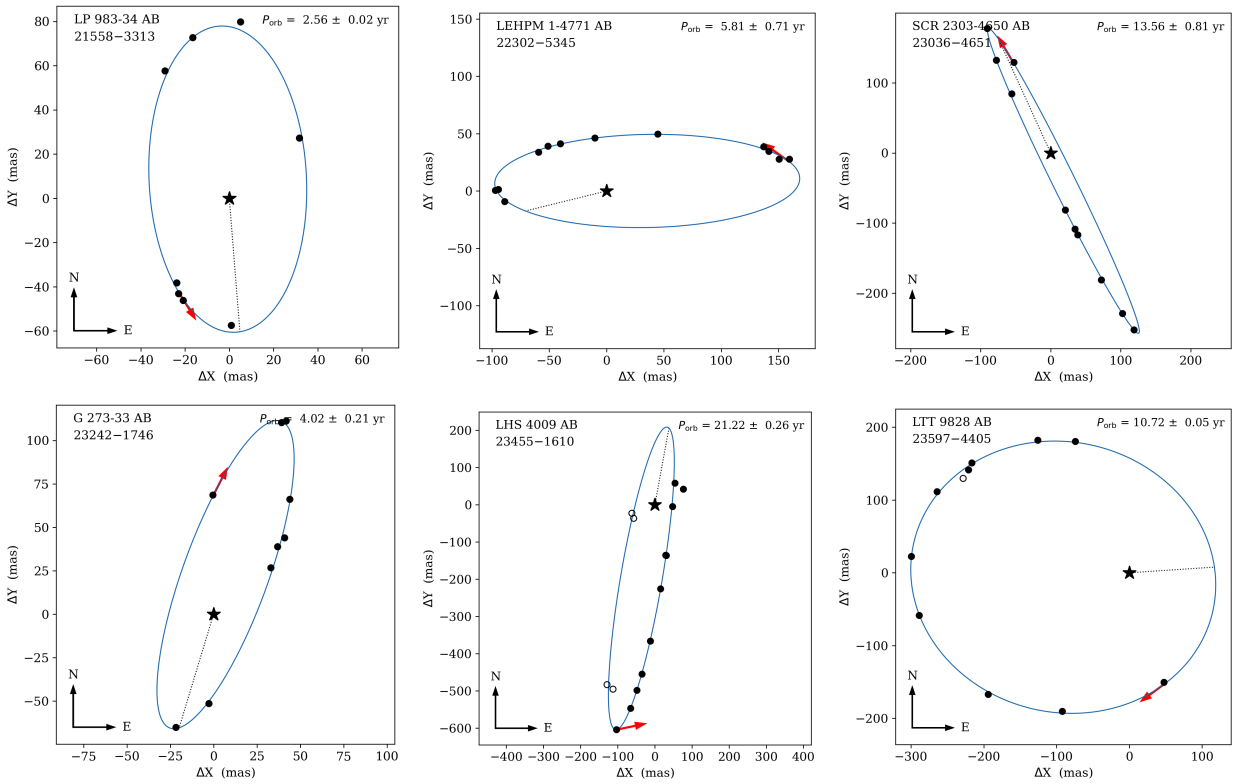

Figure 5.14 Orbits fit to SOAR speckle interferometry combined with existing high-resolution imaging data from the literature (if available). The systems are shown here in order of R.A., with each panel showing R.A. vs. Decl. for the system's SOAR data (black filled points) and literature data (open points) as well as its orbit model fit (blue curve).

Because we did not fit a mathematical function to those pseudo-MLRs, for each star in the pair we identified every point in the pseudo-MLR with M_I within 0.02 mag of that star, then averaged those points' mass estimates to arrive at a single mass for the star of interest.

- A system's simulated relative astrometry was deemed able to be fit if the simulated companion moved at least 144° in position angle (i.e., 40% of 360°) over its entire data set. These data included the 3.67 years of SOAR as well as imaging results from the literature if those were available. The fit capability criterion of 40% position angle

coverage was chosen based on prior experience fitting orbits to these data (§4.4).

This simulated position angle calculation entirely replaced the simulated PB calculation (i.e., the linear fit and amplitude of residuals) in steps 4e–4g of §5.1.2.

As for the RECONS sensitivity analysis, at each ($P_{\rm orb}$, e) point 50 randomly oriented orbits were simulated for each of the binaries, then each orbit was repeated for 50 simulated observing start times. This meant for each ($P_{\rm orb}$, e) a total of 2500 orbit snippets were simulated and evaluated for fit capability.

Figure 5.15 presents the results of that sensitivity analysis for the SOAR data. It shows that our program should fit nearly all orbits with $P_{\rm orb} \lesssim 8$ years, with 50% of orbits able to be fit out to 15 years. A dropoff in sensitivity is evident for $e \gtrsim 0.75$ regardless of $P_{\rm orb}$, which matches the bias first described by Harrington & Miranian (1977) based on the difficulty of fitting very eccentric orbits. Tokovinin & Kiyaeva (2016) also confirmed this trend in the context of binaries of solar-type stars.

Overplotted on Figure 5.15 is $P_{\rm orb}$ vs. e for the 42 orbits we have fit from the SOAR program. That distribution mainly follows the expectation from our fit capability plot, although for $P_{\rm orb} \gtrsim 10$ years we lack orbits with $e \lesssim 0.3$ despite no strong decrease in sensitivity in that area. This suggests a lack of low-e orbits in the underlying population, the implications of which are discussed further in Chapter 6.

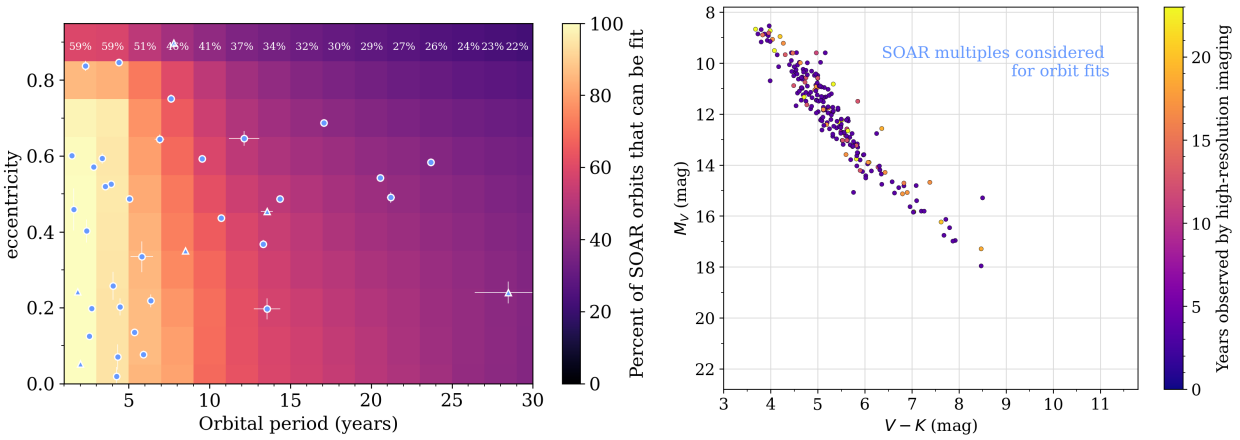

Figure 5.15 *Left panel:* average percentage of orbits that can be fit, as a function of $P_{\rm orb}$ vs. e, for multiples followed by the SOAR speckle program. This plot is analogous to that of Figure 5.10, but with fit capability defined as a companion moving through 40% of 360° over the observations (described in detail in §5.2.2). Overplotted is $P_{\rm orb}$ vs. e for the orbits fit from the SOAR speckle program, with circles for binaries and triangles for subsystems of triples and quadruples. *Right panel:* Hertzsprung-Russell diagram of the multiples analyzed in this sensitivity analysis, with each point's color indicating the time baseline over which it has been observed, including literature data. The lack of low-mass systems (compared to, e.g., the RECONS program in the right panel of Figure 5.10) is mainly due to the $I \leq 14$ mag limit of the SOAR speckle program.

5.3 Adding Orbits from the Literature

Although our two observing campaigns have provided an avalanche of data, the rich history of binary star astronomy means the literature holds many well-characterized orbits as well. Because these orbits are from campaigns with different techniques or faster cadences than the observing programs carried out for this dissertation, these literature orbits fill critical gaps in the orbital period, eccentricity, and mass ratio parameter spaces.

5.3.1 Selecting Literature Orbits of Nearby M Dwarfs

To select literature orbits, we relied on the database of the Sixth Catalog of Orbits of Visual Binaries ("ORB6"; Hartkopf et al. 2001). We extracted systems with primary stars having masses $\leq$0.6 $M_{\odot}$ from ORB6 by matching the coordinates of the 25 pc sample (from Chapter 2) to the ORB6 catalog coordinates. Each match was allowed a wide radius (1.″5) to account for potential coordinate errors in individual publications, then any coordinates that were different by $\gtrsim$ 0.″5 between the catalogs were displayed over 2MASS images in Aladin to check (by eye) if they could belong to two separate sources instead of one.

Two sources of published orbits not included in ORB6 were also added to the selection. We identified M dwarfs in Mann et al. (2019) by computing their absolute K magnitudes and verifying they matched the M_K criteria in Table 2.1 (§2.2.2). Next, we identified astrometric M and L dwarf binaries from the *Gaia* DR3 Non-Single Star catalog (NSS; Gaia Collaboration et al. 2022) via a cross-matching process identical to that described above for ORB6. Spectroscopic orbits given by *Gaia* NSS were also checked, but none had absolute magnitudes consistent with red dwarf primaries.

After this selection process, we located the publication cited for each literature orbit and verified the quality of their model fits. This process allowed us also to note the techniques used for their observations, any flags for system youth, and any masses or mass estimates for the system's components. Orbits were omitted at this stage if their authors noted them as preliminary or if the orbital eccentricity or period had to be fixed (rather than treated as a free parameter) to reach a solution.

Through this process we selected 132 orbits from published works. In the HRD of Figure 5.1, these literature systems (black points) span the entire main sequence, including near the hydrogen-burning limit of $M_V = 21.82$ mag (although they are elevated above that line in Figure 5.1 because that plot uses combined magnitudes).

5.3.2 Sensitivity of the Literature Results

The orbits published in the literature were obtained with a diverse array of techniques using a variety of instruments, adding considerable complexity to the task of assessing their observational biases. Some were the results of systematic surveys that reported all they detected, while many others were more sporadic efforts to characterize specific systems. We therefore did not conduct an extensive sensitivity analysis of this subset.

Instead, we can gain a sense of the parameter space explored by the literature by considering the methods used to map those orbits. The 132 orbits reported by the literature were obtained using the following techniques:

- relative astrometry from high-resolution imaging (e.g., speckle interferometry, adaptive optics, *HST* Fine Guidance Sensors)
- radial velocities (RVs) obtained via spectroscopy (both single-lined/SB1 and double-lined/SB2 systems)
- absolute astrometry of photocentric (unresolved) orbits (mainly the *Gaia* DR3 Non-Single Stars catalog and the CFHT/WIRCam survey of Dupuy & Liu (2017))
- light curves of eclipsing binaries (often in conjunction with RV data)

Table 5.2 Techniques used to observe and characterize the literature orbits used in this work. Relative astrometry refers to observations of a companion orbiting a primary star, while absolute astrometry refers to observing the pair's unresolved photocenter orbiting its center of mass.

Technique	Number of orbits
absolute astrometry	24
relative astrometry	16
relative astrometry + absolute astrometry	15
radial velocities	42
radial velocities + eclipsing light curves	15
radial velocities + rel. or abs. astrometry	18

A significant fraction of the characterizations used a combination of these methods. The complete breakdown is given in Table 5.2.

These efforts ranged from three to 10 years in length, and 80% of the orbits they produced had $P_{\rm orb} \leq 6$ years. These orbits also had to be at least 40% mapped by observations to pass our informal quality check (§5.3.1). Together, these factors suggest that these literature sources were unable to fit orbits longer than 7.5–25 years (for campaigns 3–10 years long) — and indeed it is the case that 78% of the literature orbits used here have $P_{\rm orb}$ within that minimum 7.5 year limit (and 98% have $P_{\rm orb}$ within 25 years).

The most significant *caveat* of the literature sample is that the short $P_{\rm orb}$ regimes are dominated by orbits determined from RVs. Obtaining RV data requires a bright primary star — and, in SB2 cases, a similarly bright secondary star — which means a $P_{\rm orb}$ regime dominated by RV orbits may overrepresent massive M dwarfs and equal-mass pairs. In this dissertation, literature authors used RVs to characterize all of the orbits with $P_{\rm orb} \leq 10$ days, and 55% of the orbits with 10 days $\leq P_{\rm orb} \leq 1$ year. Neither the SOAR nor RECONS observing program are very sensitive for $P_{\rm orb} \leq 1$ year, as those orbits are too tight for

our speckle imaging and too fast for our sparse RECONS astrometry cadence. We must therefore apply caution in our later discussion (§6.4) about the mass ratio distributions and parameters that depend on mass ratios.

5.3.3 Distinguishing Binaries from Higher-Order Multiples

Finally, in this work we only briefly consider the dynamical evolution of systems with three or more components (triples, quadruples, etc.), which may proceed differently than that of binaries. With this potential issue in mind, we used *Gaia* DR3 to search each orbit reported here for common proper motion companions separated by 1–5 arcmin from the primary star. *Gaia* reports resolved companions with separations as low as $0\farcs7$, and reliably detects ∼100% of companions separated by at least $1\farcs0$ within *Gaia*'s magnitude limit and for applicable brightness ratios. A potential companion meeting these criteria was deemed to be a gravitationally bound companion if its parallax was within 8.0 mas of the primary star and its proper motion was within 50 mas/year of the primary's proper motion in both R.A. and Decl. directions. These specific limits were set by the values found in *Gaia* for widely-separated companions already evaluated during the RECONS program for binaries within 50 pc. This search for wide companions identified several dozen systems as triples or higher-order multiples, all of which were previously noted in the literature.

Overall, these thesis results include 138 binaries and 51 triples or higher-order multiples. These results are not indicative of the true fraction of higher-order multiples vs. simple binaries, as systems with more than two components are often noted but not characterized in orbit fitting work. Throughout this dissertation, the triples and higher-order multiples

are always distinguished with triangular symbols or, for histograms, hatched bars.

CHAPTER 6

The $P_{\rm orb}$ vs. e Distribution for M Dwarfs and its Implications

In this chapter we finally present the full results of our accumulation of M dwarf multiples' orbits. We discuss the overall assessments of orbital periods and eccentricities in §6.2, as well as the short- and long-period regimes in detail. Next, in §6.3, we estimate masses of the systems, and in §6.4 we discuss the $P_{\rm orb}$ vs. e results in light of those mass estimates. In §6.4.4 we compare these M dwarf multiple results to those of solar-like and very low-mass multiples. Finally, in §6.5 we summarize the implications for M dwarf formation and dynamical evolution.

6.1 193 M Dwarf Orbits on One Plot

With orbits from the RECONS astrometry program, the SOAR speckle program, and the literature, Figure 6.1 shows the full $P_{\rm orb}$ vs. e for 189 M dwarf systems. This includes 19 orbits from RECONS, 42 from the SOAR program, and 132 from the literature. This figure shows the plot in terms of log $P_{\rm orb}$ (left panel) and linear $P_{\rm orb}$ (right panel) to highlight the behaviors in the short- and long-period regimes, respectively. At very short $P_{\rm orb}$ nearly all the orbits are circular, and at longer periods their eccentricities vary widely between an upper envelope and loose lower envelope of e.

As shown in the left panels of Figure 6.2, there are many systems with $P_{\rm orb} \leq 1$ year. There is then a gentle dropoff in numbers from 1–8 years and a smattering of orbits out to 30 years. This pattern most likely reflects the longevity of observing programs focused on orbit determinations, because short-period orbits are easier to map completely and char-

acterize well than long period orbits. Nonetheless, it is encouraging that nearly 60 orbits are now available for M dwarf pairs with $P_{\rm orb}$ = 10–30 years, allowing us to evaluate their distribution for comparisons to formation scenarios.

In concert, the eccentricity histograms shown in the right panels of Figure 6.2 reveal a relatively flat overall distribution for e = 0.1–0.6, with a rather precipitous dropoff beyond e = 0.6. It appears that high eccentricity orbits are present, but not common, in populations of of M dwarf multiple systems. These overall characteristics for $e > 0.1$ should be robust, given the detection capability evaluations described in Chapter 5 showing that we are sensitive to orbits of all eccentricities.

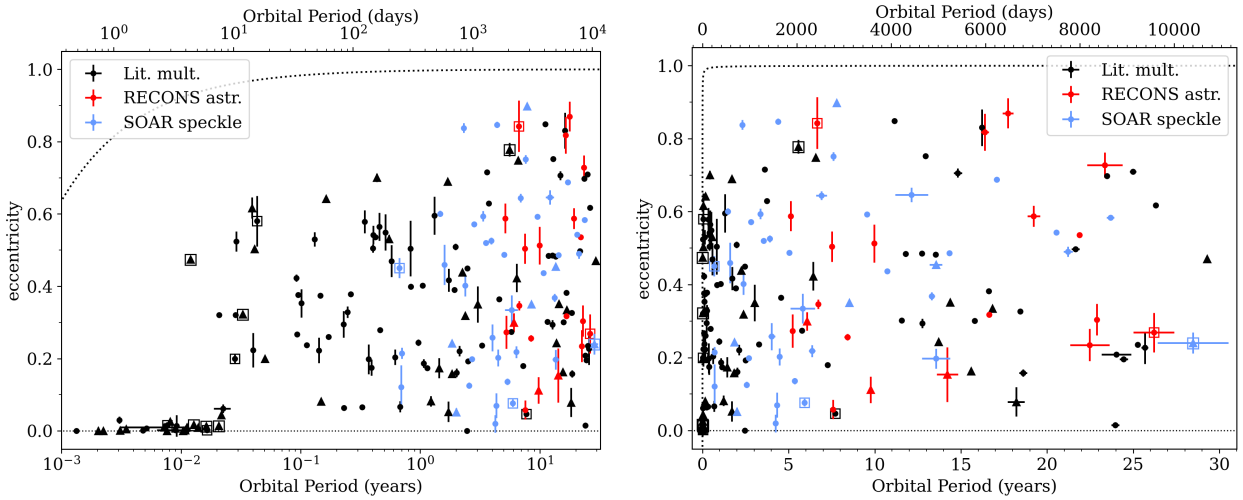

Figure 6.1 The distribution of $P_{\rm orb}$ vs. e for 193 systems with red dwarf stellar primaries, shown in terms of log $P_{\rm orb}$ (*left panel*) and linear $P_{\rm orb}$ (*right panel*). The color of each point indicates the source of its data: fit from RECONS astrometry (red, 19 points), fit from SOAR speckle interferometry (blue, 42 points), and published in the literature (black, 132 points). The shape of the point indicates the multiplicity, with binaries as filled circles and the small subset of triples and quadruples as filled triangles. Points noted in the literature as young systems are marked with open squares. The dotted curve marks the limit at which two mid-M dwarfs would collide at periastron (§6.2.2).

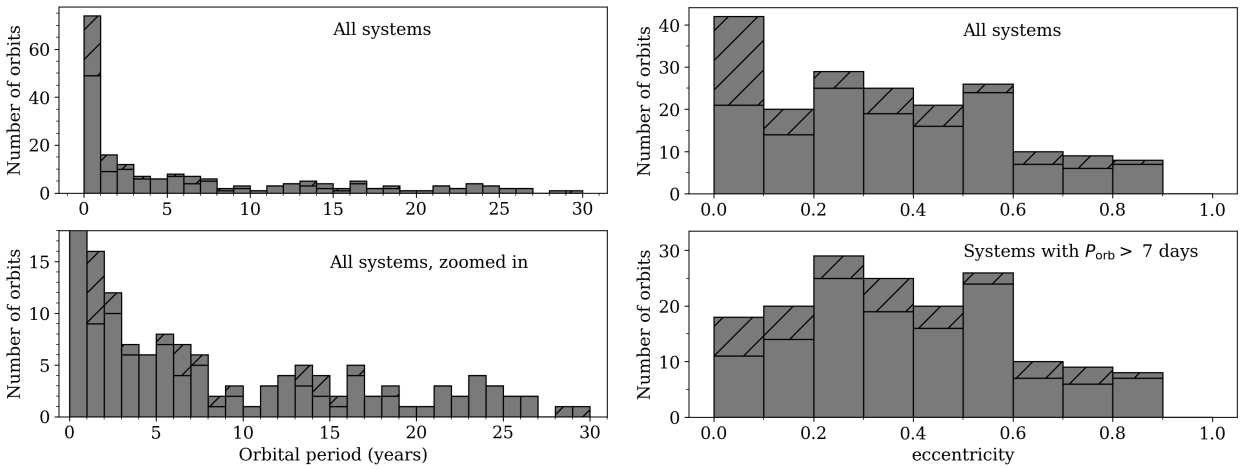

Figure 6.2 Distributions of orbital period (*left panels*) and eccentricity (*right panels*) for the orbits presented in this work (e.g., as shown in Figure 6.1). In every panel, the non-hatched bars indicate binaries and hatched bars represent subsystems of higher-order multiples. For $P_{\rm orb}$ in the left column, both panels are the full set of orbits, with the bottom panel as a zoom-in to better show the distribution for systems with $P_{\rm orb} > 1$ year. For e in the right column, the top panel is the full set of orbits and the bottom panel is the set with $P_{\rm orb} < 7$ days excluded, as those shortest-period orbits all fall in the lowest e bin.

6.2 Zoom-in on $P_{\rm orb}$ vs. e Regimes

Specific regions of $P_{\rm orb}$ vs. e space can reveal important information about how these systems evolved, and we will later note differences in these regions with stellar mass. Here we introduce the "zoom-in" regions of most significance.

6.2.1 Empirical $P_{\rm circ}$ for M Dwarf Systems

Given sufficient time, every two-body system in an eccentric orbit will eventually migrate toward a circular configuration with tidally-locked components because every periastron passage raises tides that sap orbital energy. Thus, for the sample of M dwarfs with orbits presented here, there should be an orbital period below which all aged M dwarf systems are

found in circular orbits. The log $P_{\rm orb}$ view of $P_{\rm orb}$ vs. e (left panel of Figure 6.1) shows that orbits with $P_{\rm orb} \lesssim 7$ days are almost exclusively circular. The one exception to that rule is the young system HIP47133 AB, which is a β Pictoris moving group member and is noted by Sperauskas et al. (2019) as synchronized but not circularized. That result is consistent with the prediction that tidal torques synchronize a binary pair's rotation more quickly than it circularizes their orbit (§1.3.3 and Zahn & Bouchet (1989)).

Although a line at $P_{\rm orb} \approx 7$ days cleanly separates eccentric systems from those with nearly circular orbits, considering their characterizations more precisely reveals more nuance. Table 6.1 shows 21 systems with 3 days $\lesssim P_{\rm orb} \lesssim$ 15 days, with binaries highlighted in yellow to easily distinguish them from higher-order multiples. Column 5 of that table gives each system's lower bound of e given its published uncertainty, allowing us to determine which systems have e consistent with zero. Note that subtracting more than 1σ (i.e., if the published uncertainties are underestimated) does not change the following discussion in any meaningful way.

Including only binaries in the sample, the longest period for a circular orbit is 4.08 days. If we exclude the five young systems because they may not have had time to circularize, the shortest-period eccentric system is NGT2143-3801 AB at $P_{\rm orb} = 7.62$ days. Thus, $P_{\rm circ}$ for these binaries is in the 4–8 days range. Despite this wide range of uncertainty, the M dwarfs' $P_{\rm circ}$ is firmly below the 12-day $P_{\rm circ}$ that Raghavan et al. (2010) demonstrated for binaries of solar-type stars.

Table 6.1: Systems with $P_{\rm orb}$ between 3 and 10 days, with binaries highlighted in yellow. The precision shown is as given in the original publications for these systems, given in the last column. The full references for the reference codes shown here are given in Table 1.

Name	$N_{\rm comp}$	$P_{\rm orb}$ (days)	e	$e - \sigma_e$	Flag	Ref. code
(1)	(2)	(3)	(4)	(5)	(6)	(7)
GJ0815 AC	3	3.276188 ± 0.00005	0.007 ± 0.007	0.000		Duq88
G203-060 AB	2	3.29 ± 0.99	0.002 ± 0.002	0.000		Ski18
HAT-TR-318-007 AB	2	3.3439539 ± 0.0000002	0.014 ± 0.03	-0.016		Har18
GJ0866 AC	3	3.786516 ± 0.000005	0.0 ± 0.0	0.0		Seg00
GJ1006 AC	3	3.956523 ± 0.000092	0.00220 ± 0.00096	0.00124		Bar18
LP827-020 AB	2	4.077017 ± 0.000001	0.002 ± 0.002	0.000		Zho15
GJ0867 AC	4	4.08322 ± 0.00004	0.010 ± 0.010	0.000		Her65
HIP47133 AB	3	4.38804 ± 0.00001	0.474 ± 0.005	0.469	young	Spe19
PARANAGO-1802 AB	3	4.67390 ± 0.00006	0.017 ± 0.003	0.014	young?	Gom12
GJ1230 AC	3	5.06880 ± 0.00005	0.009 ± 0.001	0.008		Del99
LP790-002 AC	3	5.922845 ± 0.000061	0.01350 ± 0.0012	0.0123	young	Bar18
PTF0850+1948 AB	2	6.015742 ± 0.000002	0.0017 ± 0.0006	0.0011	young?	Kra17
G041-014 AC	3	7.5555 ± 0.0002	0.014 ± 0.002	0.012	young	Del99
NGT2143-3801 AB	2	7.61793 ± 0.00000544	0.32034 ± 0.00120	0.31914		Act20
G229-018 AD	4	7.9461 ± 0.0002	0.044 ± 0.007	0.037		Spe19
G212-034 AB	2	8.17 ± 1.31	0.062 ± 0.012	0.050		Ski18
GJ0735 AB	2	10.3191 ± 0.0008	0.200 ± 0.012	0.188	young	Duq88
GJ0268 AB	2	10.42672 ± 0.00006	0.3203 ± 0.0009	0.3194		Bar12
LHS1610 AB	2	10.588495 ± 0.001312	0.5239050 ± 0.027422	0.49683		Gaia3
LP476-207 AC	3	11.9623 ± 0.0005	0.323 ± 0.006	0.317	young	Del99
LHS1955 AC	3	14.325696 ± 0.001731	0.6166959 ± 0.029570	0.587126		Gaia3
G093-033 AB	2	14.646202 ± 0.003492	0.2232439 ± 0.047969	0.175275		Gaia3

Including higher-order multiples allows GJ1230 AC to set the upper $P_{\rm circ}$ limit to 5.07 days, although that system's third component is close enough to have possibly influenced the inner pair — both Delfosse et al. (1999) and *Gaia* DR3 place that third companion at 5″ (∼50 AU), giving it a potential $P_{\rm orb}$ of only several centuries. We also note that all higher-order multiples below $P_{\rm orb} = 4.08$ days are circular with the exception of GJ1006 AC, which has an orbit that is only barely eccentric at $e = 0.00220$. Triples with their third component at ∼a few hundred AU will evolve under the additional effects of Kozai-Lidov cycles with

tidal friction (Fabrycky & Tremaine 2007), which ultimately drives the inner pair toward circular. This makes them not comparable to simple binaries in our evaluation of $P_{\rm circ}$. If the third component is further out (e.g., beyond 1000 AU) then the inner pair may effectively function as a simple binary, but not many M dwarf systems with those configurations are known and have accurate orbits for their inner pairs.

Another unusual feature is that in the $P_{\rm orb}$ = 4–8 days range an unusually high number of systems are noted as young (marked with boxes in Figure 6.1). This feature should be interpreted with caution, as the entire sample has not been uniformly searched for signs of youth. The results in this short-$P_{\rm orb}$ regime are dominated by orbits determined from radial velocities, thus these systems all have spectroscopic data available to assess their potential ages — which is not the case for every system on the $P_{\rm orb}$ vs. e plot.

On the other hand, if there is a formation path that sets some systems quickly migrating inward toward circular orbits with tidally-locked components, we would expect the region of $P_{\rm circ}$ orbits to be dominated by those young, quickly-migrating systems. Models of disk fragmentation show migration inward can happen on astronomically short timescales of a few Myr, but this process also results in accretion of mass onto the companion. In the bottom panel of Figure 6.4 (discussed in §6.4) we see that high mass ratios are not overrepresented in this $P_{\rm orb}$ regime, so this is not a compelling formation path for the shortest period systems.

6.2.2 The Observed Upper Envelope of e

A clear envelope of maximum eccentricity extends from the $P_{\rm circ}$ line to our survey limit of 30 years. As $P_{\rm orb}$ increases from ∼10 days to 30 years, this upper envelope extends from

$e \approx 0.65$ to $e \approx 0.95$.

This envelope is well below the limit at which two M mid-M dwarfs would pass close enough to collide, represented in the left panel of Figure 6.1 with the dotted black curve. This limit was drawn assuming two stars of 0.30 $M_\odot$ and 0.33 $R_\odot$ (Mann et al. 2015) each collide when their periastron distance is $\leq$0.66 $R_\odot$. We tested additional combinations of masses and radii from the M dwarf sequence and found no significant difference in the result — i.e., in all cases, the observed points fall well below the collision curve in $P_{\rm orb}$ vs. e space.

The line of stars along the upper edge of the e envelope are almost exclusively subsystems of triples or higher-order multiples. These systems also populate lower eccentricities, thus they do not necessarily have higher eccentricities than binaries; rather, they have a somewhat higher-e upper envelope. This is presumably a sign of an eccentricity-pumping mechanism only accessible to higher-order multiples, such as the Kozai-Lidov mechanism (§1.3.2).

Interpreting the specific shape of the upper e envelope — such as its gradual slope upward with increasing $P_{\rm orb}$ — is treacherous for this data set, as most of the shape is set by our diverse set of literature orbits. Roughly 50% of the orbits between 10 days and 1 year were obtained with radial velocities (RVs), while the rest were from absolute astrometry and eclipsing binaries. Without a careful analysis of those surveys' sensitivities, we cannot be confident that they can characterize high-e orbits just as easily as low-e orbits, so the true slope of the M dwarf multiples' upper envelope remains uncertain. Its somewhat abrupt rise in the SOAR-sensitive region ($P_{\rm orb} \gtrsim 1$ year) validates these suspicions.

6.2.3 Eccentricities at Long $P_{\rm orb}$

Beyond $P_{\rm orb}\sim$10 years, the density of orbits in $P_{\rm orb}$ vs. e space decreases, although the bottom left panel of Figure 6.2 shows that the distribution is only slightly decreasing beyond that threshold. This lack of long-period orbits is due to the declining sensitivity of the RECONS and SOAR observational programs in this regime (see Figure 5.10 and Figure 5.15), as well as the decline of the literature efforts' sensitivities (§5.3.2) because most observing programs are short in duration.

That said, orbits with $P_{\rm orb} \geq 7$ years are generally contained between $e \approx 0.1$ and $e \approx 0.9$, with only a few exceptions to that lower bound. This distribution is demonstrated in the bottom panel of Figure 6.3. To interpret this feature, we must consider the sensitivity of our observing programs. The SOAR speckle program shows a notable decrease in fitting capability with increasing e; for example, at $P_{\rm orb} = 24$ years, the percentage of orbits that can be fit decreases from 45% to 26% from $e = 0$ to $e = 0.9$. The RECONS astrometry program does not have such a strong dependence on e because its sensitivity is mostly flat for multiples (nearly identical to Figure 5.9) and its fit capability depends only on $P_{\rm orb}$, not e, resulting in the vertical striped pattern in Figure 5.10. Considering those analyses together suggests that the lack of circular orbits at long $P_{\rm orb}$ is reflective of the underlying population, while the lack of very eccentric orbits may or may not be real.

6.3 Estimating Masses for Every M Dwarf

Because in Chapter 1 we noted so much variation in $P_{\rm orb}$ vs. e between high-mass, solar-type, and very low-mass systems (§1.4), our M dwarf results will be richer if we interpret their distributions in light of their masses as well. For 52 of our 193 systems we have dynamical masses of the individual stars, but for the remaining systems we need to estimate their masses using alternative means.

When dynamical masses are not available, our most reliable technique to estimate component masses in a system is to use the absolute magnitudes of its individual components. Of our 193 pairs, 69 have magnitude difference (Δm) measurements of their components, either from our SOAR speckle program or from various efforts published in the literature. For these systems, we estimated component masses using either the corrected M_K MLR described in §2.2.1 or the I-band pseudo-MLR presented in §2.2.2. A few systems had only Δm measurements from alternative bandpasses at optical or near-infrared wavelengths, often from narrow-band filters used in speckle surveys, and in these cases we considered the shapes of those filters' throughput functions and chose either K or I, depending on which was most similar.

For six systems there were no Δm measurements or dynamical masses, but we did have a photocentric orbit determined by RECONS. In those cases we used that photocentric orbit to constrain the individual masses through the following procedure:

1. We began by finding every potential mass combination that could have produced the measured M_K for the combined pair. This process was identical to that used for step 1

of §5.1.2: we created a grid of mass combinations between 0.08 $M_\odot$ and 0.6 $M_\odot$ in steps of 0.01 $M_\odot$, computed the M_K expected from each component mass on that grid using the corrected MLR (§2.2.1), then computed the M_K expected for each combination of those absolute magnitudes. These simulated combined magnitudes needed to be within 0.02 mag of the measured M_K to be considered a match. The grid contained no duplicated mass combinations.

As we did in §5.1.2, if the measured M_K did not match any grid values, we matched it against an extended grid that included masses every 0.001 $M_\odot$ below 0.1 $M_\odot$. This was necessary because the change in absolute magnitude with respect to mass becomes quite steep in that region.

2. Next, we computed B and β (the fractional mass and luminosity; §3.4.1) of each potential mass combination from the previous step.

3. We also computed the semi-major axis a of each mass combination given their combined masses and the orbit's measured $P_{\rm orb}$.

4. Finally, we used the B, β, and a computed above for each mass combination to determine $\alpha = a\,(B - \beta)$ (Eq. 3.1), where α is the semi-major axis of the photocentric orbit. The mass combination for which α was closest to the orbit's measured value was considered most likely, and its component masses were adopted for the system.

After checking each of the above possibilities for estimating masses, there were 62 systems that did not have appropriate data to use any of the above mass estimating procedures. For

these systems, we assumed $\Delta K \approx 0.8$ mag and followed the procedure described above for when magnitude differences were available. This specific value was the average for the entire sample with measured Δm in any bandpass, after converting Δm in other bandpasses to Δm at K.

The result is that dynamical masses make up slightly less than 1/3 of the sample, masses estimated for systems with Δm measurements make up slightly more than 1/3, and systems with combined magnitudes assumed to have $\Delta K = 0.8$ mag make up slightly more than 1/3. Mass estimates from photocentric orbits only account for a few percent. These masses and mass estimates are given in Table 6.2.

Table 6.2 Masses and mass estimates used for the analysis described in §6.4 and thereafter. Column 4 indicates the technique used to produce the values in the preceding columns; masses that are not dynamical are estimates based on either magnitude differences, photocentric orbits from RECONS, or combined magnitudes assuming $\Delta K = 0.8$ mag (described in §6.3). The precision shown here for non-dynamical masses was used for the calculations in this analysis (e.g., mass ratios) but is not representative of the true precision on these estimates. This table is given in full at `http://astro.gsu.edu/~vrijmoet/PhDthesis/table_masses.txt`.

Name	M_1 ($M_\odot$)	M_2 ($M_\odot$)	mass type
(1)	(2)	(3)	(4)
G131-026 AB	0.23248	0.21992	dmag I
LEHPM1-0255 AB	0.22574	0.14749	dmag I
GJ1005 AB	0.17900	0.11200	dynamical
2MA0015-1636 AB	0.23983	0.23799	dmag I
GJ1006 AC	0.43000	0.15100	dynamical
GJ2005 BC	0.08440	0.08040	dynamical
LP349-025 AB	0.15850	0.00930	dynamical
GJ0022 AC	0.40500	0.15700	dynamical
GIC0050 AC	0.24688	0.10040	dmag I
GJ0046 AB	0.30722	0.20212	combined phot

6.4 $P_{\rm orb}$ vs. e by Mass and Mass Ratios

With masses (or mass estimates) available for every system, we now have the values needed to look for trends with respect to this fundamental parameter. Figures 6.4 and 6.5 show the $P_{\rm orb}$ vs. e plot of Figure 6.1, but with the points colored by either the primary's mass (top plot in each Figure), secondary's mass (middle plot), or mass ratio (bottom plot). In the mass ratio plots, systems with masses estimated by assuming $\Delta K = 0.8$ mag are shown as open symbols, as those masses are the least reliable.

Distributions of eccentricity with respect to primary mass are shown in Figure 6.6, with $P_{\rm orb} < 7$ days excluded to avoid tidally circularized systems. These distributions show that the highest-mass M dwarf multiples have a roughly flat distribution of eccentricity, but this distribution skews toward lower e for the lower-mass M_1 cohorts. The distribution shapes are qualitatively different between the 0.15–0.30 $M_\odot$ vs. 0.075–0.15 $M_\odot$ groups, but we caution against interpreting this feature because the peak numbers of those histograms are somewhat low. Of the 1/3 of systems with masses estimated from combined photometry (§6.3), ∼40% fall into these two lower bins, but in general those estimated masses' eccentricity distributions match that of the full set (see the bottom panels of Figures 6.4 and 6.5). The right panels of Figure 6.6 also show statistically too few systems to interpret, but they do hint that the systems with $P_{\rm orb} > 7$ years match the trend of the full set in the left panels.

Figure 6.7 shows the eccentricity distributions with respect to secondary mass, again with $P_{\rm orb} < 7$ days excluded. The trends are consistent with the primary mass plots of Figure 6.6, with the 0.075–0.15 $M_\odot$ cohort now strongly peaked around $e \approx 0.35$ and the more massive

0.15–0.30 $M_{\odot}$ group peaked at $e \approx 0.2$. Estimated masses are evenly distributed between the top three mass groups, so they should have less of an effect here. The shift in peak from lower to higher e runs counter to the overall trend of e decreasing with stellar mass (discussed in §1.4 and later in §6.4.4 and §6.4.5). It is likely, however, that the decrease in low-mass orbits at $P_{\rm orb} \lesssim 3$ years is playing a role in this trend, as e generally decreases in that region as well (Figure 6.4), and the trend is not significant when the orbits are restricted to $P_{\rm orb} \gtrsim 7$ years (right panel of Figure 6.7). We have not conducted an extensive analysis of the biases with respect to mass and mass ratio for the literature orbits that dominate this region, thus we cannot draw conclusions about whether the decrease in low-mass systems below $P_{\rm orb} \lesssim 3$ years represents the underlying population — and, similarly, whether the eccentricity trend that follows from it is representative.

Adding the dimension of masses to the $P_{\rm orb}$ vs. e plots reveals a few important features to specific regions of $P_{\rm orb}$ vs. e space, discussed next.

6.4.1 Masses Among the Shortest-$P_{\rm orb}$ Orbits

Figure 6.4 shows that among the circular or nearly-circular orbits with $P_{\rm orb} \lesssim 7$ days, many systems have primaries with masses of 0.3 $M_{\odot}$ or higher. This feature is likely a selection effect, as every one of those systems were characterized using RVs from spectroscopy, which requires brighter targets than the other techniques contributing to this plot.

On the other hand, it is important to note that the migration that would bring a companion closer to its primary star — before, or while tidal forces circularize the orbit — would cause that companion to accrete material and drive its mass closer to that of the primary

(§1.3.1). If this scenario is the only one by which companions could migrate inward, then we would expect this short-$P_{\rm orb}$ regime to be exclusively populated by systems with high mass ratios. The bottom panel of Figure 6.4 shows that this is indeed the case because all but one of the 25 systems with $P_{\rm orb} \lesssim 7$ days have q values of at least 0.5, and the middle panel of that Figure confirms that almost none of those $q \gtrsim 0.5$ systems are higher-mass M dwarfs with very low-mass companions. With this result we must also tread carefully, however, as more massive companions are more detectable via RV than less massive companions.

We therefore conclude that for short-$P_{\rm orb}$ orbits of M dwarfs, the evidence is consistent with migration+tides being responsible for populating that regime, but the evidence is also not sufficient to disprove that scenario. To form stronger conclusions requires a uniform survey of companions that targets these short-$P_{\rm orb}$ orbits, can easily characterize orbits of the lowest-mass companions, and has its fitting capability space well mapped (analogous to §5.1.2 and §5.2.2).

6.4.2 Masses Among the Longest-$P_{\rm orb}$ Orbits

In the $P_{\rm orb} \gtrsim 7$ years region, the primary masses represented span the entire M dwarf sequence, as shown in the top panel of Figure 6.5 and the right column of Figure 6.6. The lowest-mass cohort is slightly more prevalent, which is unsurprising given that many of these orbits are from RECONS astrometry, which is generally better suited to fainter stars (§3.2) and has shorter time baselines for M dwarfs with masses $\gtrsim$0.3 $M_{\odot}$ because they were added to the observing program relatively recently (Figure 5.10, right panel).

The most notable feature among these long orbits ($\gtrsim$ 7 years) is that seven of the 11

systems with $e < 0.2$ have primaries with masses of 0.075–0.15 $M_{\odot}$. In the same plot showing secondary masses (Figure 6.5, middle panel), nine of the 11 systems have secondaries ≤ 0.15 $M_{\odot}$. Only two systems in this long-$P_{\rm orb}$, low-e region have both stars ≥ 0.15 $M_{\odot}$. In summary, for $P_{\rm orb}$ longer than a few years, nearly all systems with $e \lesssim 0.2$ are those with at least one very low-mass star or brown dwarf.

This result is consistent with the findings of Dupuy & Liu (2017), who showed that the very low-mass systems have lower eccentricities than solar-type systems. Our comparison with their results is discussed in more depth in §6.4.5. Those authors suggested that this difference could be due to the lower-mass stars and brown dwarfs having longer-lived disks that impart additional eccentricity damping. The bottom panel of Figure 6.5, however, shows that four of the eleven systems with $P_{\rm orb} \geq 7$ years and $e < 0.2$ have $q < 0.5$. It would be surprising for so many companions to move through disks enough to dampen their eccentricities without accreting significant material.

On the other hand, three of those four high-$P_{\rm orb}$, low-e, low-q systems have brown dwarf rather than stellar companions. Furthermore, in the bottom panels of Figure 6.7 we also note that brown dwarf companions exclusively have $e < 0.5$. Brown dwarfs' lower luminosities and potentially different formation processes are most likely playing a role here.

6.4.3 Mass Ratios

The mass ratios of the M dwarf multiples are potentially their most telling property. Figure 6.8 presents the mass ratio distribution for all systems in our sample excluding those that required assuming $\Delta K = 0.8$ mag to estimate mass. This distribution reveals that twin

multiples ($q \geq 0.9$) vastly outnumber all other pairs in our sample. Figure 6.9 demonstrates that this twin prevalence does not depend on mass of the primary or secondary — this mass ratio is common in the sample regardless of the type of system. The only panel in Figure 6.9 that does not show these excess twins is that for brown dwarf secondaries, which will never have many similar-mass primaries in this sample because we excluded systems with primaries of mass $\leq 0.075\ M_{\odot}$.

This twins excess could have strong implications for M dwarf multiples' dynamical evolution, but there is also high potential for observational bias to contribute to this result. The bottom panel of Figure 6.4 shows that the $q > 0.9$ systems become prevalent for $P_{\rm orb} \gtrsim 1$ year, which is precisely the region where high-resolution imaging methods become sensitive (e.g., see Figures 5.15 and 6.1 for SOAR speckle). These imaging techniques excel when companions are similar in brightness. Additionally, as we consider longer-period orbits that rely on data from older instruments, we might expect brighter companions to be better represented due to those instruments' reduced sensitivities. Figure 6.10 shows the distributions of q for different regimes of $P_{\rm orb}$, and indeed the twin pairs become less overrepresented as $P_{\rm orb}$ decreases. To determine the full extent of the role observational bias is playing here requires a more complete sensitivity analysis of each observing program — including literature efforts — with respect to mass ratio.

Figure 6.11 shows the eccentricity distribution for systems in different cohorts of mass ratio. Considering either the full sample (left column) or the long-period orbits alone (right column), the eccentricity distribution is roughly flat regardless of mass ratio. In most panels

there are not enough orbits to discern details of the distributions' shapes. Figure 6.12 confirms this lack of structure and shows that there is no correlation between eccentricity and our estimated mass ratios.

6.4.4 Comparing M Dwarfs to Solar-like Systems

The $P_{\rm orb}$ vs. e distribution for systems of solar-type stars was presented by Raghavan et al. (2010), who collected multiples from a sample of 454 F, G, and K stars to determine the multiplicity of that population, as discussed in §1.4.2. In doing so, they also showed that the solar-type multiples circularize at $P_{\rm orb} \approx 12$ days and rarely have circular orbits for $P_{\rm orb} \gtrsim 100$ days. Figure 6.13 shows their distribution (yellow points) over our M dwarf distribution (black points), with $P_{\rm circ}$ for the solar-type stars marked at 12 days with the dashed yellow line.

There are no fewer than eight eccentric M dwarf systems to the left of that 12-day line, along with a handful of nearly (but not quite) circular systems. This firmly demonstrates that $P_{\rm circ}$ for M dwarfs is shorter than $P_{\rm circ}$ for solar-type stars, validating our analysis in §6.2.1.

On an intuitive level it seems reasonable that the M dwarf multiples would have a shorter $P_{\rm circ}$ than solar-type multiples, as their lower mass means the stars must be closer together to feel the same amount of gravitational force and thus comparable tides. In order for a 0.6+0.3 $M_{\odot}$ pair to feel the same gravity as a 1.0+0.5 $M_{\odot}$ pair with $P_{\rm orb} = 12$ days and semi-major axis 0.117 AU, the M dwarfs need to orbit at 0.071 AU, corresponding to $P_{\rm orb} = 7.2$ days. For less massive M dwarfs, a 0.3+0.15 $M_{\odot}$ pair would need to orbit at

3.6 days. The observed $P_{\rm circ}$ of 4–8 days is thus what we would expect over the wide range of M dwarf masses, although important factors such as their stellar radii and internal structures should also be considered.

Just beyond the $P_{\rm circ}$ lines, the solar-type multiples exhibit a somewhat tighter correlation of e with respect to $P_{\rm orb}$ than the M dwarf multiples. For $P_{\rm orb} \gtrsim 100$ days, the solar-type e distribution is roughly the same width as the M dwarf distribution, but is shifted toward higher e; this is also evident in the top panel of Figure 6.14 as compared to the bottom right panel of Figure 6.2. Low-e orbits are entirely lacking among the solar-type multiples, while the M dwarfs only see that behavior past $P_{\rm orb}$ of $\sim$7 years (§6.2.1), and systems that include very low-mass components can violate this rule. These differing features suggest differences in the migration paths that led to these orbits and/or differences in stable orbits available after the migration process. We elaborate upon these scenarios in §6.5.

Finally, the upper envelope of the solar-type multiples' eccentricity distribution is somewhat higher than the M dwarfs (see Figure 14 of Raghavan et al. 2010). Considering that for the solar-type systems this envelope is dominated by higher-order multiples just as it is for our M dwarfs, it is again likely due to the Kozai-Lidov mechanism, or a similar evolutionary process in systems with more than two components.

6.4.5 Comparing M Dwarfs to Very Low-Mass Systems

For very low-mass (VLM) systems with $M_1 \lesssim 0.1\ M_\odot$, Dupuy & Liu (2017) showed that the eccentricity distribution firmly skews toward lower e at high $P_{\rm orb}$. This distribution is reproduced in the bottom panel of Figure 6.14. Their sample is somewhat small (25

multiples) because VLM systems are intrinsically faint and thus difficult to observe. We compare those 25 multiples to our M dwarfs in Figure 6.15, in which blue points indicate the VLM systems and black points are once again our M dwarf multiples. The VLM sample overlaps the M dwarfs (and hence some points are aligned), but the VLMs also include the systems with brown dwarf primaries that we excluded from our sample.

The logarithmic view of $P_{\rm orb}$ (Figure 6.15, left panel) is not particularly revealing, as the VLM systems lack observations of short-period orbits. The two short-period orbits both have $P_{\rm orb} < 12$ days and are still moderately eccentric ($e > 0.3$), consistent with our conclusion that lower-mass systems have a shorter $P_{\rm circ}$ than solar-type systems. Although one of these VLM systems has an eccentric orbit shorter than our 7-day upper limit for the M dwarf $P_{\rm circ}$, that orbit is a young system (Pleiades member; Basri & Martín 1999), thus may not have had time to circularize yet. We thus cannot draw more conclusions about $P_{\rm circ}$ for VLM systems.

The linear $P_{\rm orb}$ view (Figure 6.15, right panel) demonstrates that these systems have representatives out to similar $P_{\rm orb}$ as our M dwarf multiples, and yet their upper envelope of eccentricity remains firmly below that of the M dwarfs. Only two VLM systems have $e > 0.5$, whereas 27% of our M dwarfs exceed that limit. Although this upper envelope is composed of higher-order multiples among the M dwarfs, the highest-e VLM systems are binaries, not triples or quadruples (shown in Figure 18 of Dupuy & Liu 2017). The maximum eccentricity being higher for the higher-mass population matches what we observed in our comparison of solar-type vs. M dwarf multiples. On the other hand, it is surprising that the upper e

envelope of VLM systems is *not* dominated by higher-order multiples. This suggests that eccentricity-pumping mechanisms for subsystems of higher-order multiples are less efficient for these VLM systems. The orbital evolution of Kozai-Lidov cycles, for example, has several dependencies on the components' masses (Fabrycky & Tremaine 2007). The necessary *caveat* to this conjecture is, of course, that there not many VLM systems with well-determined orbits.

The most notable feature of the VLM $P_{\rm orb}$ vs. e distribution is that 88% of the orbits have $e < 0.5$. In contrast, 70% of the M dwarf orbits and and 58% of the solar-type orbits are below the $e = 0.50$ line. There is thus a clear trend for orbits to be less eccentric as the stellar masses decrease. We explore the implications of this result in §6.5.

6.5 The Picture of M Dwarf Formation

With evidence in hand for the M dwarfs as well as the two mass regimes that bookend them, we are well positioned to draw comparative conclusions regarding how these systems form and dynamically evolve.

As outlined in Chapter 1, current models suggest systems of multiple stars form with components separated by tens to tens of thousands of AU. Considering that all of the M dwarf systems presented here have $P_{\rm orb} \leq 30$ years and thus $a \lesssim 10$ AU, every companion in our sample has undergone significant migration to arrive at its orbit that we observe today.

At $P_{\rm orb} \gtrsim 7$ years, orbits are generally eccentric among solar-type multiples, somewhat less eccentric among M dwarf multiples, and appear to be even less eccentric for very low-

mass (VLM) multiples. This trend creates the characteristic lack of nearly-circular orbits at long $P_{\rm orb}$ for solar-type and higher-mass M dwarf systems. Previous work (Dupuy & Liu 2017) suggested that this trend is a consequence of disk lifetimes increasing as stellar mass decreases; lower-mass companions and their primary stars spend more time swimming through circumstellar material, losing more angular momentum to those disks than higher mass stars do before the disks dissipate.

An essential piece of that orbital evolution process is accretion of disk material onto the secondary star. If the overabundance of twins observed for M dwarfs (§6.4, esp. Figures 6.9 and 6.10) is not entirely due to observational bias, we can be assured that most companions experience a significant amount of accretion to arrive at their current positions. On the other hand, the lack of correlation between eccentricity and mass ratio (Figure 6.12) throws a wrench into this interpretation — why would some very low-mass systems evolve to low eccentricity without significant accretion?

We obtain more clues as we look toward shorter $P_{\rm orb}$. The eccentricity distribution of solar-type multiples narrows as $P_{\rm orb}$ decreases (Figure 6.13), whereas the M dwarfs' distribution narrows by an insignificant amount. This is another counterintuitive set of features: if M dwarfs spend *more* time exchanging angular momentum with disk material than solar-type stars, we might expect their final orbits to be *more* strictly regulated. The simplest explanation may be that M dwarf companions arrive at those orbits with less disk interaction than solar-type companions. This may mean the disk fragmentation radius is smaller or the disks themselves are smaller or less dense. This scenario could also explain why we find a few

low-mass stars with low eccentricities and low mass ratios, assuming the odds of a multiple forming with $e = 0.1$ are not very different than one forming with $e = 0.8$.

The final clue to the M dwarf binary formation process comes from their tidal circularization period ($P_{\rm circ}$). In §6.2.1 we noted that the observed range of potential $P_{\rm circ}$ is remarkably close to what we would expect given the solar-type multiples' $P_{\rm circ}$ and the difference in gravitational force between those two populations. The timescale of tidal orbital evolution depends on other important factors, however, such as the stars' radii, internal structures, and initial eccentricity at the moment the tidal force becomes significant. The agreement of reality with our back-of-the-envelope calculation thus suggests that the post-migration distributions of M dwarf systems and solar-type systems are different by roughly the amount needed to cancel out their differences in the other important tidal factors. The finer details of the tidal evolution timescales could thus be constrained by a very complete, well-characterized set of observed multiples and their orbits. The dissertation work presented here may qualify if we update our orbits with more reliable mass ratios for systems with $P_{\rm orb} \lesssim 1$ year.

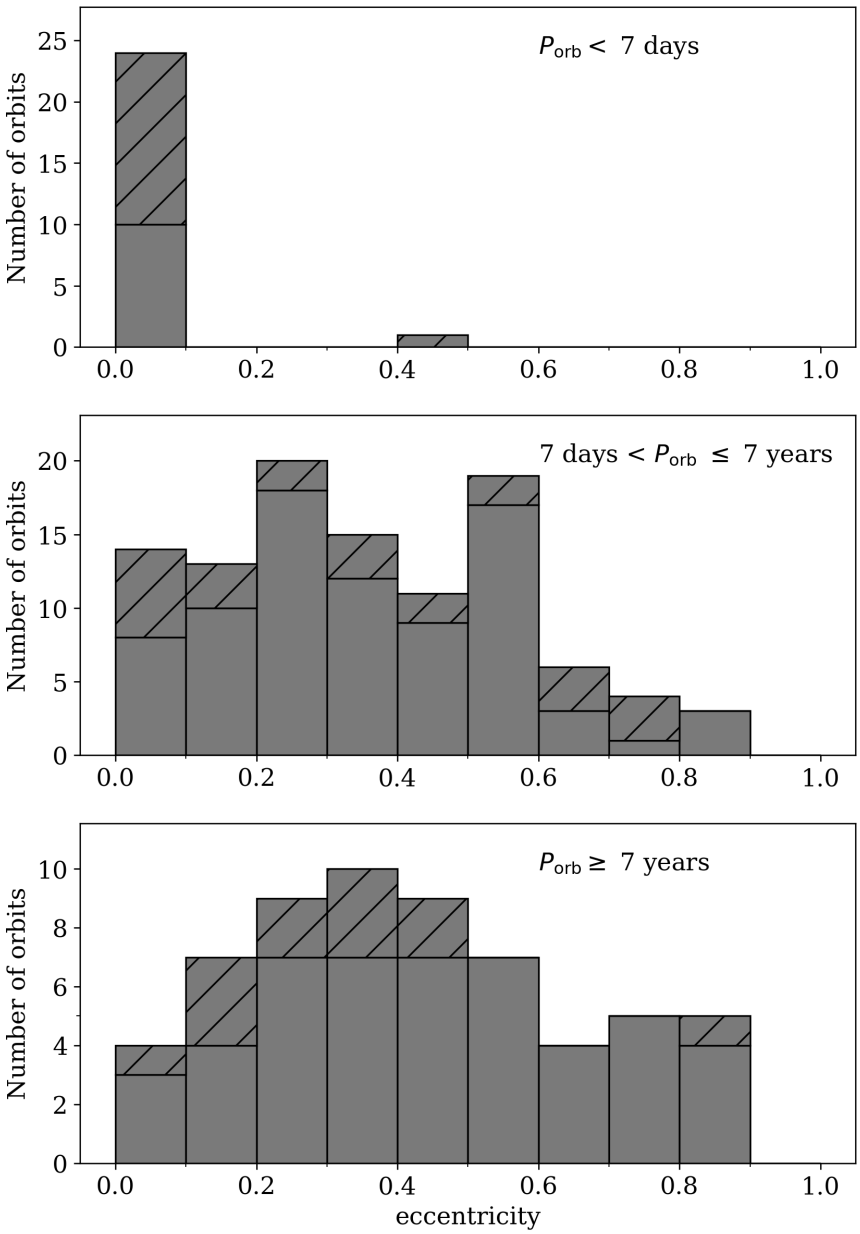

Figure 6.3 Distribution of eccentricity in three regimes of $P_{\rm orb}$: $P_{\rm orb} < 7$ days (*top panel*), 7 days $\leq P_{\rm orb} < 7$ years (*middle panel*), and $P_{\rm orb} \geq 7$ years (*bottom panel*). In every panel, the non-hatched bars indicate binaries and the hatched bars indicate subsystems of higher-order multiples.

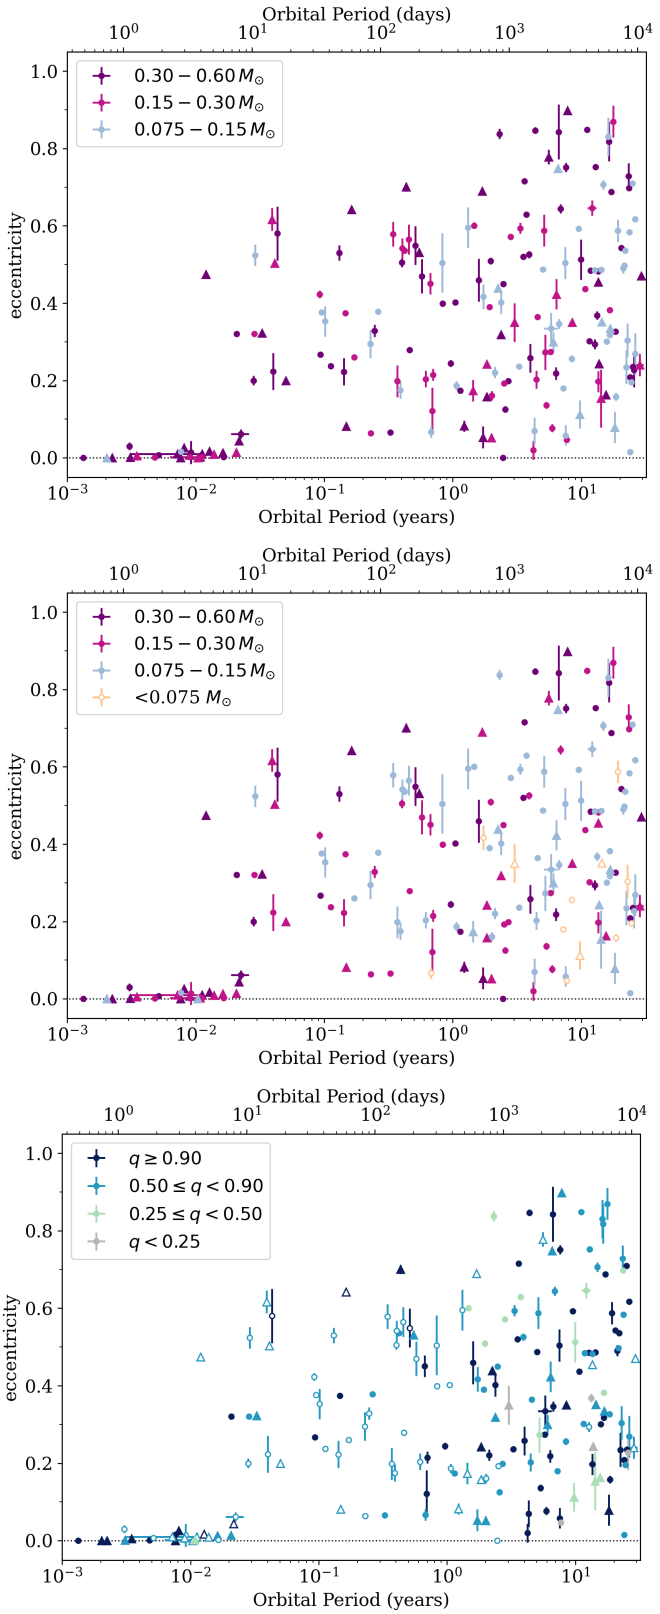

Figure 6.4 The $P_{\rm orb}$ vs. e plot showing the same systems as in Figure 6.1, but with points colored by primary mass (*top*), secondary mass (*middle*), or mass ratio (*bottom*). Open symbols indicate masses estimated assuming $\Delta K = 0.8$ mag, as those estimates are the least reliable.

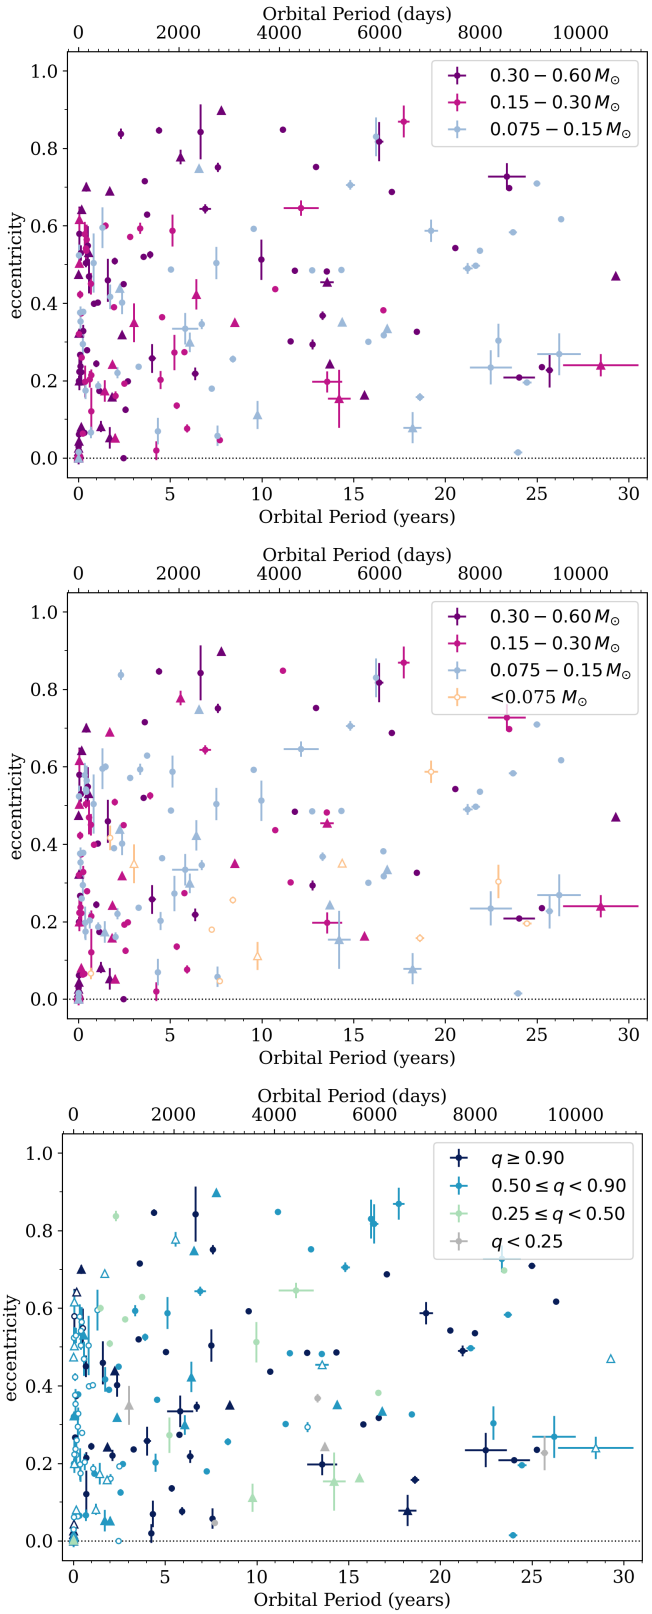

Figure 6.5 These $P_{\rm orb}$ vs. e plots are identical to those in Figure 6.4, but here the plots are in terms of linear $P_{\rm orb}$ instead of log $P_{\rm orb}$. Again, the points are colored by primary mass (*top*), secondary mass (*middle*), or mass ratio (*bottom*). Open symbols indicate masses estimated by assuming $\Delta K = 0.8$ mag, as those estimates are the least reliable.

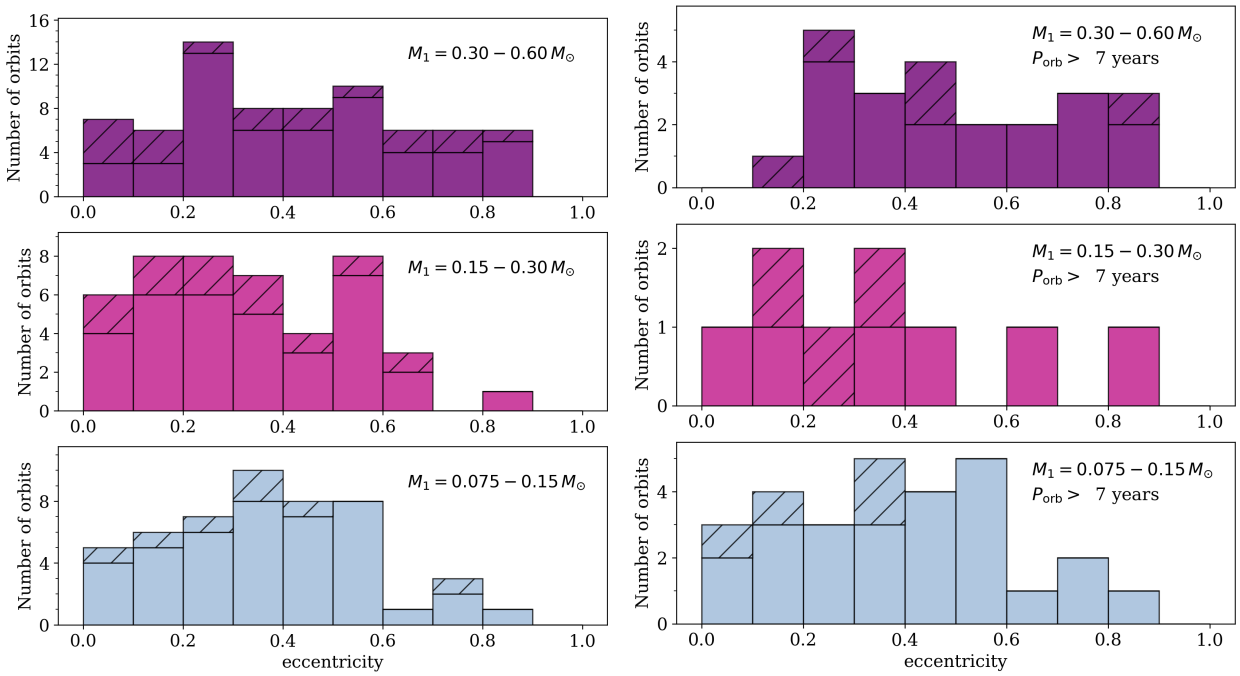

Figure 6.6 Distributions of eccentricity for multiples with different M dwarf primaries, shown for all systems (*left column*) and those with $P_{\rm orb} > 7$ years (*right column*). Systems with $P_{\rm orb} < 7$ days were excluded. Binaries are shown with non-hatched bars and higher-order multiples are indicated with hatched bars. Masses were either dynamical or estimated as described in §6.3. The analogous plot organized by M dwarf secondary mass is given in Figure 6.7.

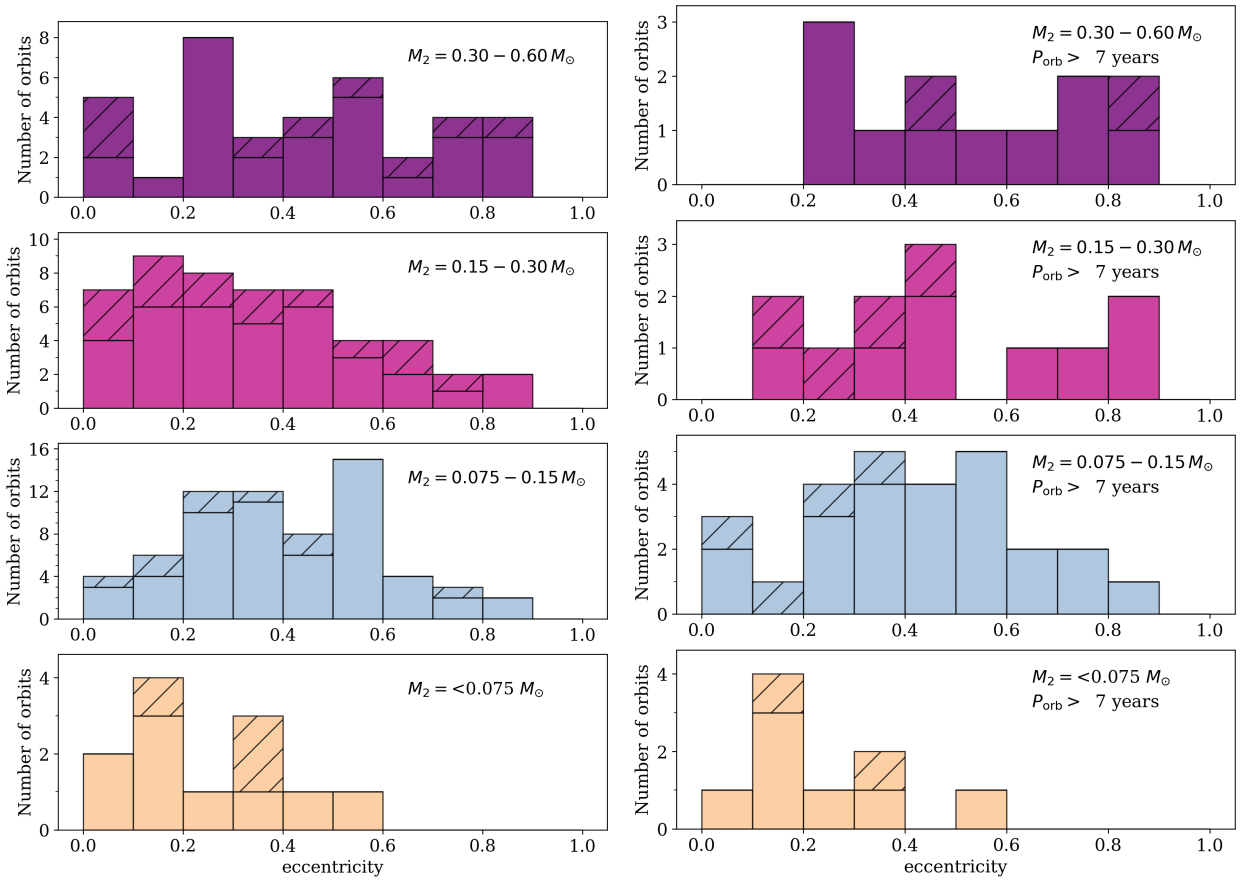

Figure 6.7 Distributions of eccentricity for multiples with different M dwarf secondaries, shown for all systems (*left column*) and those with $P_{\rm orb} > 7$ years (*right column*). Systems with $P_{\rm orb} < 7$ days were excluded. Binaries are shown with non-hatched bars and higher-order multiples are indicated with hatched bars. Masses were either dynamical or estimated as described in §6.3. The analogous plot organized by M dwarf primary mass is given in Figure 6.6.

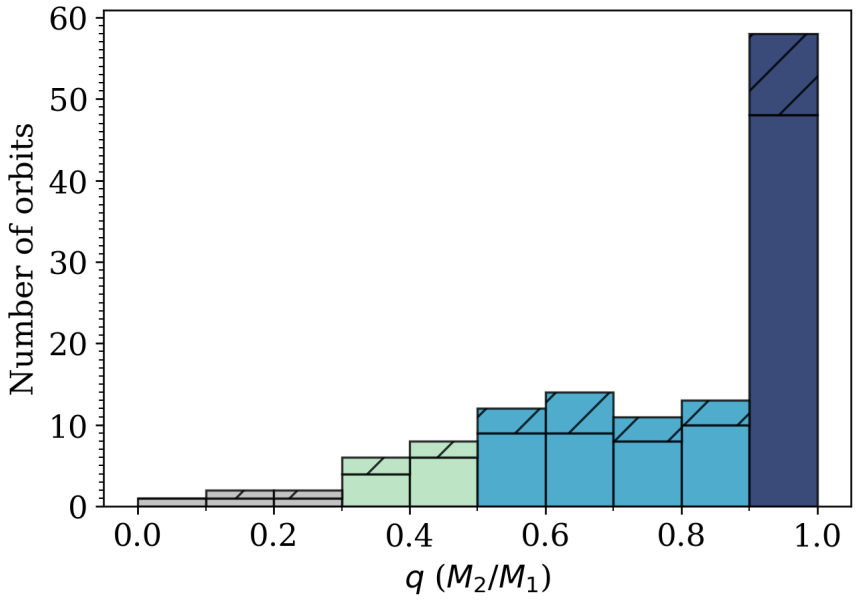

Figure 6.8 Distribution of mass ratios for M dwarf systems presented in this work. For systems without dynamical mass, masses were estimated following §6.3. Systems were excluded if their mass estimates required assuming $\Delta K = 0.8$ mag, as those masses are the least reliable. Binaries are represented here with non-hatched bars and higher-order multiples are indicated with hatched bars. The bars are colored to match the color scheme for q used in other plots in this chapter.

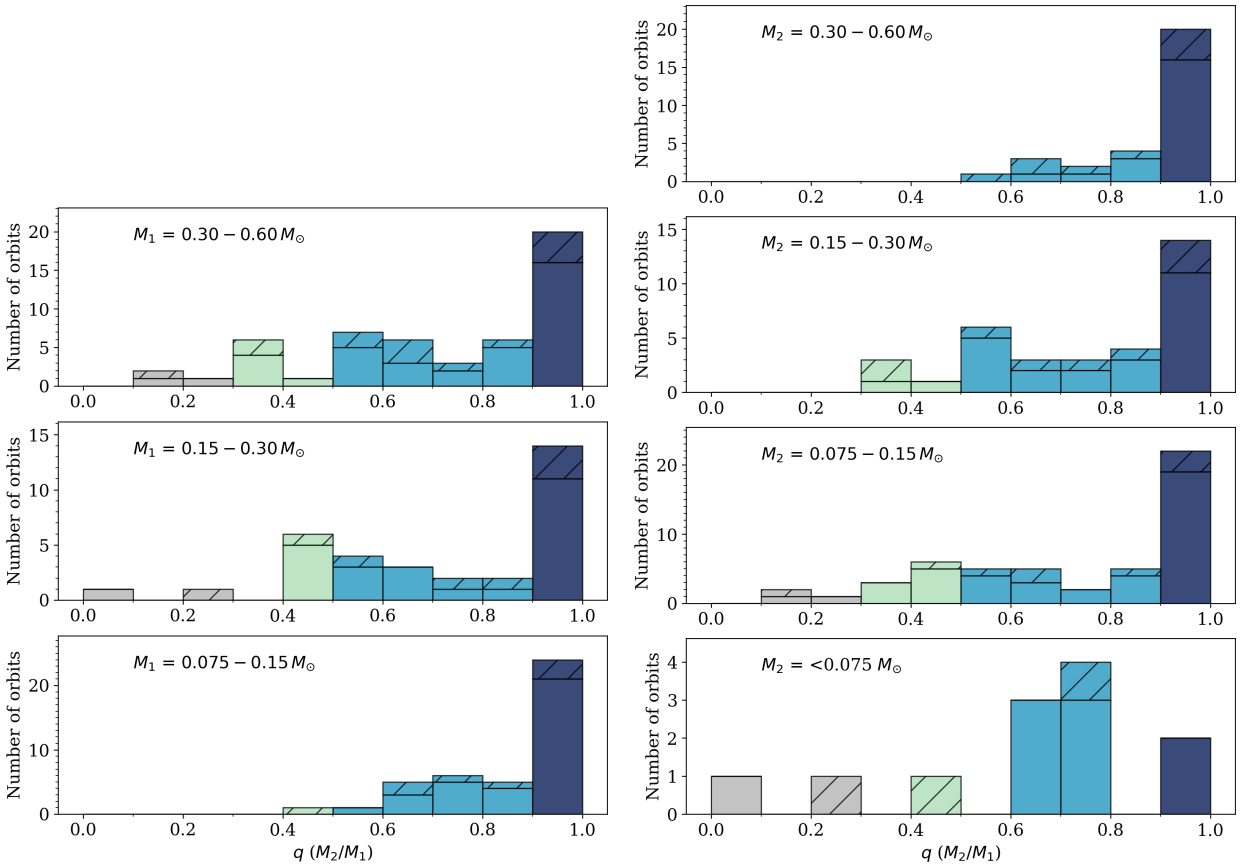

Figure 6.9 Distributions of mass ratios for multiples with different primary masses (*left column*) and secondary masses (*right column*). In every panel, the non-hatched bars represent binaries and the hatched bars indicate subsystems of higher-order multiples.

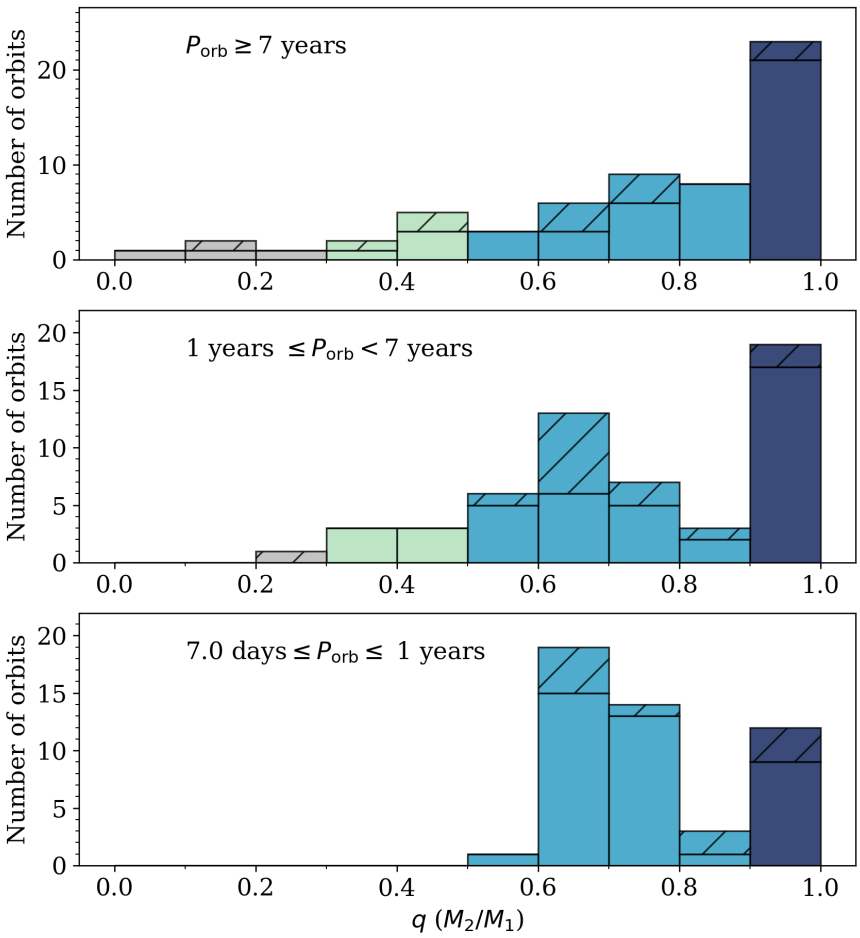

Figure 6.10 Distributions of mass ratios for multiples of different $P_{\rm orb}$. In every panel, the non-hatched bars represent binaries and the hatched bars indicate subsystems of higher-order multiples.

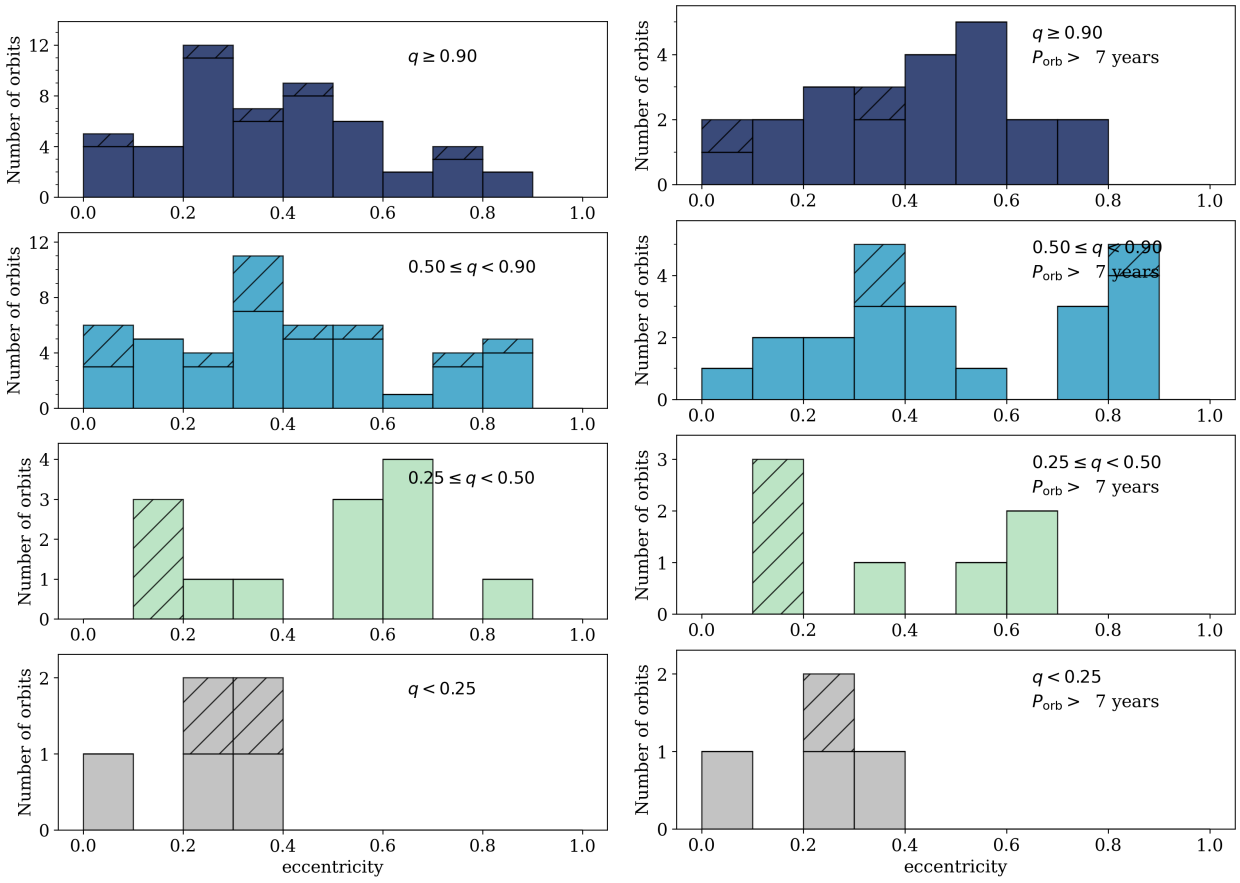

Figure 6.11 Distributions of eccentricity for systems of different mass ratios, given for all systems (*left column*) and those with $P_{\rm orb} > 7$ years (*right column*). All masses (and thus mass ratios) that were not dynamical masses were estimated using the procedure in §6.3.

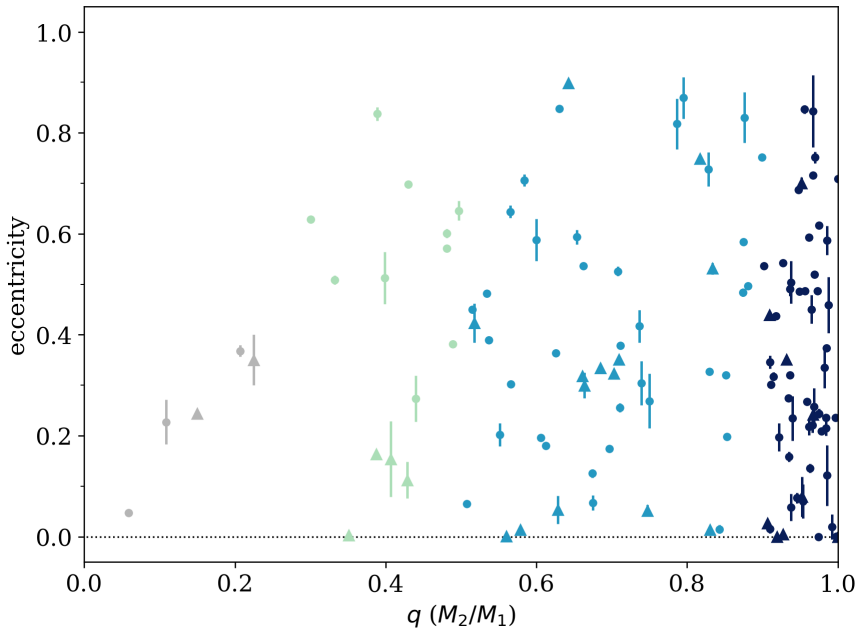

Figure 6.12 Distribution of eccentricity with respect to mass ratio as q vs. e, excluding systems for which masses were estimated by assuming $\Delta K = 0.8$ mag. All masses (and thus mass ratios) that were not dynamical masses were estimated using the procedure in §6.3. The symbol shapes indicate the types of system: circles are binaries and triangles are subsystems of higher-order multiples.

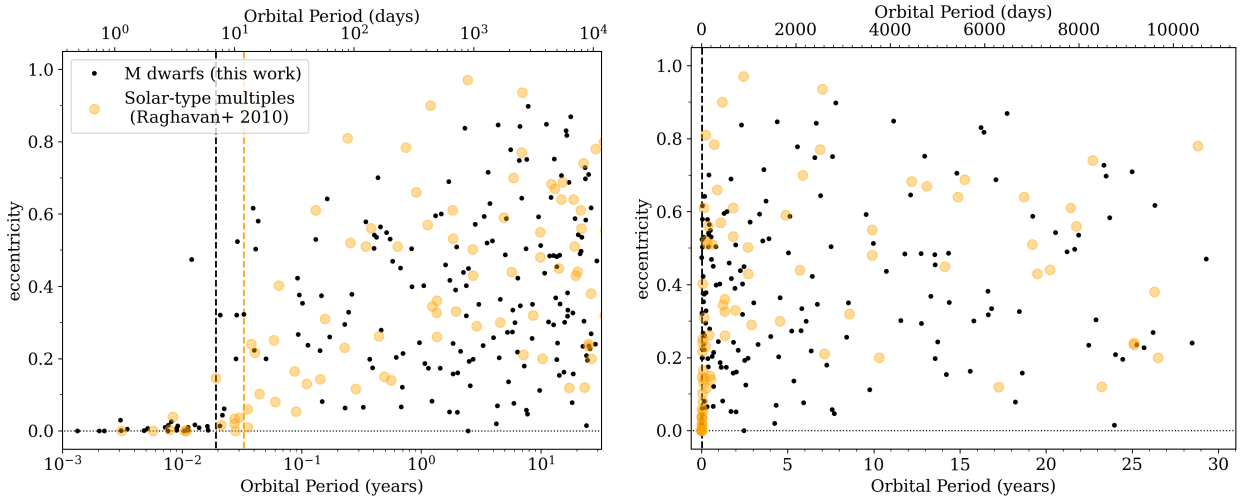

Figure 6.13 $P_{\rm orb}$ vs. e for multiples of solar-type stars (yellow points; Raghavan et al. 2010) and M dwarfs (black points; this work). The same distribution is shown in terms of log $P_{\rm orb}$ (*left panel*) and linear $P_{\rm orb}$ (*right panel*). The 12-day tidal circularization period for solar-type stars is shown with the yellow dashed line, and the black dashed line indicates the 7-day period that is the upper limit of $P_{\rm circ}$ for M dwarfs.

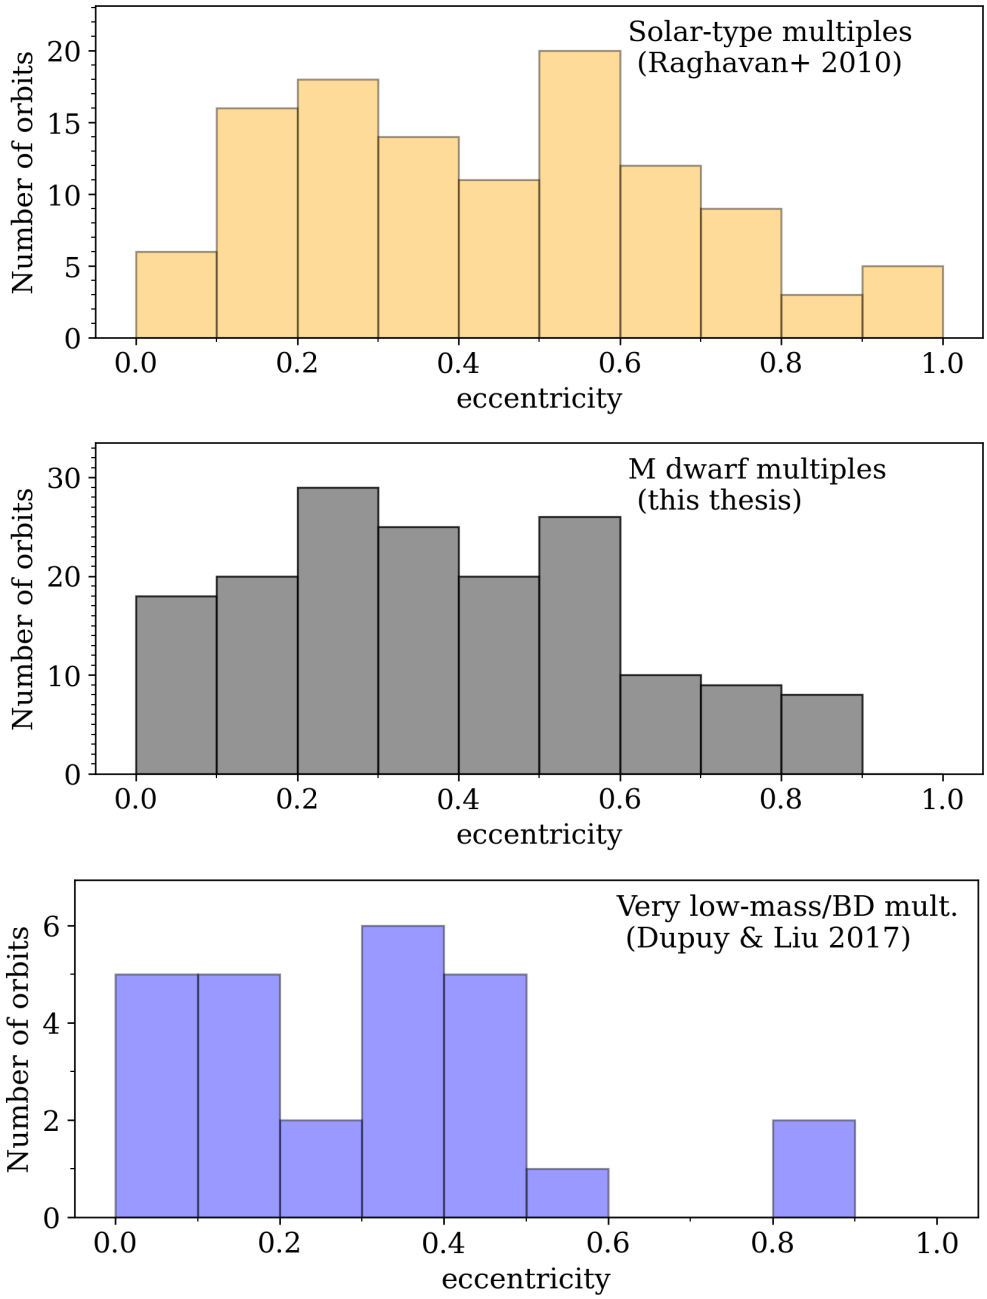

Figure 6.14 Comparison of the eccentricity distributions for three populations of multiples. *Top panel:* multiples of F, G, and K dwarfs ("solar-type") from Raghavan et al. (2010); *middle panel:* multiples of M dwarfs from this work (repeated from Figure 6.2); *bottom panel:* multiples of very low-mass (VLM) stars and brown dwarfs from Dupuy & Liu (2017). Note that the scale for the VLM distribution is much smaller, as they have fewer orbits than the other sets. In every panel, the bars represent the sum of all multiples (binary, triple, etc.), and systems with $P_{\rm orb} < P_{\rm circ}$ for that population are excluded (except for the VLM set, which contains no tidally circularized systems).

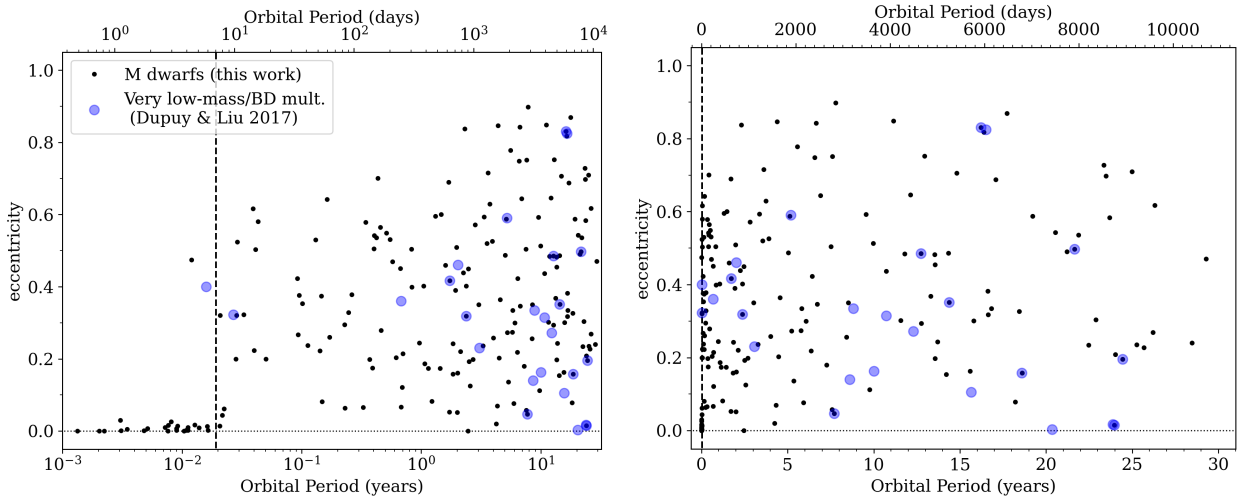

Figure 6.15 $P_{\rm orb}$ vs. e for multiples of very low-mass (VLM) stars and brown dwarfs (blue points; Dupuy & Liu 2017) and M dwarfs (black points; this work). The distribution is given in terms of log $P_{\rm orb}$ (*left panel*) as well as for linear $P_{\rm orb}$ (*right panel*). The black dashed line indicates the 7-day limit below which all M dwarf systems are nearly circular; $P_{\rm circ}$ is not indicated for the VLM sample because there are not enough of those systems with short $P_{\rm orb}$ observed to constrain it.

CHAPTER 7

Conclusions and the Future

We have presented here a comprehensive study of the orbits of M dwarf multiples, focusing primarily on the population within 25 pc (augmented by a few systems beyond that radius). By combining data from two observing programs as well as the literature, we have mapped the distribution of $P_{\rm orb}$ vs. e between 0.3 days and 30 years. These results show that the low-mass systems demonstrate tidal circularization at very short orbits and exhibit a diverse range of shapes for longer orbits, with the most eccentric systems being higher-order multiples. The M dwarf multiples' orbital eccentricity distributions lie intermediate between those of the solar-type multiples and brown dwarf multiples, forming a continuum across stellar mass.

The multiples considered for this study most likely formed at wide separations (at least tens of AU) and migrated to their current configurations ($\lesssim$10 AU). The only efficient method for this process is if the stars travel through each others' circumstellar disks and lose angular momentum to that material before it disperses. This process shrinks the stars' orbits as they accrete material from the disks, which is also expected to drive the orbits toward circular architectures. Among our M dwarfs we observed a high incidence of very high mass ratios ($q \gtrsim 0.9$), a result that is consistent with the majority of the systems migrating through accretion, although we caution that observational biases are undoubtedly present because low mass/low flux companions are more difficult to detect than higher mass/higher flux companions. These selection effects drive up the number of roughly equal mass systems in our sample.

Future improvements to this study should address those observational biases first and foremost. The continuation of the RECONS volume-complete astrometric survey over the next 10 years should ameliorate some of the biases described here, as many of the higher-mass M dwarfs in that survey currently have extremely short timelines of observation that limit their data's usefulness. More impactful, however, would be if the sensitivity space of the diverse set of literature orbits were better characterized by re-fitting each publication's data in a uniform matter, although some variations will be required for the different combinations of types of data employed. Each data set could also be characterized in terms of its sensitivity to different mass ratios, which would allow us to calibrate the observed mass ratio distributions to determine the underlying populations. Our current interpretations are somewhat conservative because of the poorly constrained selection effects for mass ratio.

Another significant dimension to this work could be added if the entire sample were uniformly assessed for indicators of system youth or age. Although most orbital evolution occurs quickly, we do occasionally find a system that is in the midst of these rapid processes (such as HIP47133 AB, discussed in §6.2.1). Separating the systems that are decidedly post-evolution from those that are potentially still mid-evolution could add an unusually empirical view of how these processes change the distributions of orbits. Although it is notoriously difficult to determine an age for an M dwarf system, a first step in estimating ages for the overall M dwarf sample would be to derive UVW space motions for all of the systems considered here.

Of course, as is the case with nearly every astrophysical study, more data points on

virtually every plot presented here would further clarify the trends we have already outlined. Nonetheless, this study presents 193 M dwarf orbits that have been well-characterized and used to reach the conclusions given. This is four times more than any previous study for these stars and extends out to three decades of orbital coverage, matching the orbit of Saturn around our Sun.

Significant improvements could also be made to the RECONS data if we incorporated the new results from the *Gaia* mission into our data reduction process. If we matched our reference stars to the *Gaia* catalog instead 2MASS, we could use the *Gaia* parallaxes and proper motions to determine their true positions at each observation, and from those determine a more accurate and precise measurement of the non-astrophysical shape of the field (i.e., the shape imparted by the optics between the CCD and sky). This upgrade would improve the RECONS program's sensitivity to equal-mass and very low-mass companions down to Jovian planet masses, and reduce the number of signals of ambiguous origin (such as those shown in Figure 4.2 in §4.1.2). Characterizing the non-astrophysical field shape so precisely would also allow RECONS data to be combined with the *Gaia* time-series data when that is released in the coming years (Data Release 4, planned for late 2025). This upgrade is already planned for the coming years at RECONS by members of the group.

Finally, the most telling extension to this study would be to include even lower-mass systems using various techniques, including both brown dwarfs and gas giant planets. Both of these populations are expected to be rare companions to M dwarfs, and a comprehensive survey of their occurrence rates would add another dimension to the results of their dynamics.

This study is already planned for the coming years. Observing programs are also underway to search for giant planets on short orbits around M dwarfs (Cañas et al. 2022, 2023; Kanodia et al. 2023), and the RECONS data will complement these by opening the door to long-period orbits once it is improved with the *Gaia* results described above. Having the $P_{\rm orb}$ vs. e parameter space mapped for planetary companions as well would allow us to compare the stellar and brown dwarf results directly with a population that is well known to form and evolve in circumstellar disks. We would then have one big picture of formation and dynamical evolution across the entire mass spectrum of the main sequence — and beyond.

Appendices

This is the appendix!

Ref. code	Reference
Act20	Acton et al. (2020)
Bal10	Balega et al. (2010)
Bar12	Barry et al. (2012)
Bar18	Baroch et al. (2018)
Ben16	Benedict et al. (2016)
Bla10	Blake et al. (2010)
Cas18	Casewell et al. (2018)
Dav16	David et al. (2016)
Del99	Delfosse et al. (1999)
Dia07	Díaz et al. (2007)
Die18	Dieterich et al. (2018)
Dup17	Dupuy & Liu (2017)
Duq88	Duquennoy & Mayor (1988)
Gaia3	Gaia Collaboration et al. (2022)
Gom12	Gómez Maqueo Chew et al. (2012)
Har18	Hartman et al. (2018)
Har96	Harlow (1996)
Her65	Herbig & Moorhead (1965)
Ire08	Ireland et al. (2008)
Irw11	Irwin et al. (2011)
Kor16	Koren et al. (2016)
Kra17	Kraus et al. (2017)
Kur08	Kürster et al. (2008)
Laz18	Lazorenko & Sahlmann (2018)
Lop05	López-Morales & Ribas (2005)
Lub17	Lubin et al. (2017)
Man19	Mann et al. (2019)
Mar09	Martinache et al. (2009)
Mat97	Mathieu et al. (1997)
Maz01	Mazeh et al. (2001)
Mor09	Morales et al. (2009)
Nid02	Nidever et al. (2002)
Sah15	Sahlmann et al. (2015)
Seg00	Ségransan et al. (2000)
Ski18	Skinner et al. (2018)
Spe19	Sperauskas et al. (2019)
Tok15	Tokovinin et al. (2015)
Tok97	Tokovinin (1997)
Win20	Winters et al. (2020)
Xia19	Xia et al. (2019)
Zho15	Zhou et al. (2015)

Table 1 Five-letter reference codes and their corresponding formal references, as used in tables throughout this dissertation.

REFERENCES

Abt, H. A., Gomez, A. E., & Levy, S. G. 1990, ApJS, 74, 551

Acton, J. S. et al. 2020, MNRAS, 494, 3950

Balega, I. I., Balega, Y. Y., & Malogolovets, E. V. 2010, Astrophysical Bulletin, 65, 250

Baroch, D. et al. 2018, A&A, 619, A32

Barry, R. K. et al. 2012, ApJ, 760, 55

Basri, G., & Martín, E. L. 1999, AJ, 118, 2460

Bate, M. R. 2012, MNRAS, 419, 3115

Bate, M. R. 2015, in Astronomical Society of the Pacific Conference Series, Vol. 496, Living Together: Planets, Host Stars and Binaries, ed. S. M. Rucinski, G. Torres, & M. Zejda, 37

——. 2018, MNRAS, 475, 5618

Bate, M. R., & Bonnell, I. A. 1997, MNRAS, 285, 33

Benedict, G. F. et al. 2016, AJ, 152, 141

Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393

Blake, C. H., Charbonneau, D., & White, R. J. 2010, ApJ, 723, 684

Brandner, W., Calissendorff, P., & Kopytova, T. 2023, MNRAS, 518, 662

Burn, R., Schlecker, M., Mordasini, C., Emsenhuber, A., Alibert, Y., Henning, T., Klahr, H., & Benz, W. 2021, A&A, 656, A72

Cañas, C. I. et al. 2022, AJ, 164, 50

——. 2023, arXiv e-prints, arXiv:2302.07714

Casewell, S. L. et al. 2018, MNRAS, 481, 1897

Chabrier, G., & Baraffe, I. 1997, A&A, 327, 1039

Clarke, C. J., Gendrin, A., & Sotomayor, M. 2001, MNRAS, 328, 485

Clarke, C. J., & Pringle, J. E. 1991a, MNRAS, 249, 584

——. 1991b, MNRAS, 249, 588

Cuello, N., Ménard, F., & Price, D. J. 2023, European Physical Journal Plus, 138, 11

Cutri, R. M. et al. 2003, VizieR Online Data Catalog, II/246

David, T. J., Hillenbrand, L. A., Cody, A. M., Carpenter, J. M., & Howard, A. W. 2016, ApJ, 816, 21

Delfosse, X., Forveille, T., Beuzit, J. L., Udry, S., Mayor, M., & Perrier, C. 1999, A&A, 344, 897

Delfosse, X., Forveille, T., Ségransan, D., Beuzit, J. L., Udry, S., Perrier, C., & Mayor, M. 2000, A&A, 364, 217

Díaz, R. F., González, J. F., Cincunegui, C., & Mauas, P. J. D. 2007, A&A, 474, 345

Dieterich, S. 2018, Sergedieterich/Mcmc_Sd: Mcmc_Sd V1.0, Zenodo

Dieterich, S. B., Henry, T. J., Jao, W.-C., Winters, J. G., Hosey, A. D., Riedel, A. R., & Subasavage, J. P. 2014, AJ, 147, 94

Dieterich, S. B., Simler, A., Henry, T. J., & Jao, W.-C. 2021, AJ, 161, 172

Dieterich, S. B. et al. 2018, ApJ, 865, 28

Douglas, S. T., Agüeros, M. A., Covey, K. R., & Kraus, A. 2017, ApJ, 842, 83

Dupuy, T. J., & Liu, M. C. 2017, ApJS, 231, 15

Duquennoy, A., & Mayor, M. 1988, A&A, 200, 135

——. 1991, A&A, 248, 485

El-Badry, K., Rix, H.-W., Tian, H., Duchêne, G., & Moe, M. 2019, MNRAS, 489, 5822

Fabian, A. C., Pringle, J. E., & Rees, M. J. 1975, MNRAS, 172, 15

Fabrycky, D., & Tremaine, S. 2007, ApJ, 669, 1298

Fisher, R. T. 2004, ApJ, 600, 769

Gaia Collaboration et al. 2022, arXiv e-prints, arXiv:2206.05595

——. 2016, A&A, 595, A1

Gammie, C. F. 2001, ApJ, 553, 174

Gómez Maqueo Chew, Y., Stassun, K. G., Prša, A., Stempels, E., Hebb, L., Barnes, R., Heller, R., & Mathieu, R. D. 2012, ApJ, 745, 58

Gray, R. O., Corbally, C. J., Garrison, R. F., McFadden, M. T., Bubar, E. J., McGahee, C. E., O'Donoghue, A. A., & Knox, E. R. 2006, AJ, 132, 161

Gray, R. O., Corbally, C. J., Garrison, R. F., McFadden, M. T., & Robinson, P. E. 2003, AJ, 126, 2048

Guszejnov, D., Grudić, M. Y., Hopkins, P. F., Offner, S. S. R., & Faucher-Giguère, C.-A. 2021, MNRAS, 502, 3646

Guszejnov, D., Hopkins, P. F., & Krumholz, M. R. 2017, MNRAS, 468, 4093

Harlow, J. J. B. 1996, AJ, 112, 2222

Harrington, R. S., & Miranian, M. 1977, PASP, 89, 400

Hartkopf, W. I., Mason, B. D., & Worley, C. E. 2001, AJ, 122, 3472

Hartman, J. D. et al. 2018, AJ, 155, 114

Hawley, S. L., Gizis, J. E., & Reid, I. N. 1996, AJ, 112, 2799

Henry, T. J., Franz, O. G., Wasserman, L. H., Benedict, G. F., Shelus, P. J., Ianna, P. A., Kirkpatrick, J. D., & McCarthy, Donald W., J. 1999, ApJ, 512, 864

Henry, T. J., Jao, W.-C., Subasavage, J. P., Beaulieu, T. D., Ianna, P. A., Costa, E., & Méndez, R. A. 2006, AJ, 132, 2360

Henry, T. J. et al. 2018, AJ, 155, 265

Henry, T. J., & McCarthy, D. W., J. 1990, ApJ, 350, 334

Henry, T. J., & McCarthy, Donald W., J. 1993, AJ, 106, 773

Herbig, G. H., & Moorhead, J. M. 1965, ApJ, 141, 649

Hilditch, R. W. 2001, An Introduction to Close Binary Stars (Cambridge, UK: Cambridge University Press)

Holden, N., Henry, T., Vrijmoet, E., Couperus, A., Jao, W.-C., Dieterich, S., & Recons Team. 2023, in American Astronomical Society Meeting Abstracts, Vol. 55, American Astronomical Society Meeting Abstracts, 302.17

Hoogerwerf, R., de Bruijne, J. H. J., & de Zeeuw, P. T. 2001, A&A, 365, 49

Ireland, M. J., Kraus, A., Martinache, F., Lloyd, J. P., & Tuthill, P. G. 2008, ApJ, 678, 463

Irwin, J., Charbonneau, D., Nutzman, P., & Falco, E. 2009, in Transiting Planets, ed. F. Pont, D. Sasselov, & M. J. Holman, Vol. 253, 37–43

Irwin, J. M. et al. 2011, ApJ, 742, 123

Jao, W.-C., Henry, T. J., Gies, D. R., & Hambly, N. C. 2018, ApJ, 861, L11

Jao, W.-C., Henry, T. J., Subasavage, J. P., Brown, M. A., Ianna, P. A., Bartlett, J. L., Costa, E., & Méndez, R. A. 2005, AJ, 129, 1954

Jefferys, W. H., Fitzpatrick, M. J., & McArthur, B. E. 1988, Celestial Mechanics, 41, 39

Kanodia, S. et al. 2023, AJ, 165, 120

Kirk, H. et al. 2017, ApJ, 838, 114

Koren, S. C., Blake, C. H., Dahn, C. C., & Harris, H. C. 2016, AJ, 151, 57

Kozai, Y. 1962, AJ, 67, 591

Kratter, K. M. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 447, Evolution of Compact Binaries, ed. L. Schmidtobreick, M. R. Schreiber, & C. Tappert, 47

Kratter, K. M., & Murray-Clay, R. A. 2011, ApJ, 740, 1

Kraus, A. L. et al. 2017, ApJ, 845, 72

Kuffmeier, M., Calcutt, H., & Kristensen, L. E. 2019, A&A, 628, A112

Kürster, M., Endl, M., & Reffert, S. 2008, A&A, 483, 869

Lasker, B. M., Doggett, J., McLean, B., Sturch, C., Djorgovski, S., de Carvalho, R. R., & Reid, I. N. 1996, in Astronomical Society of the Pacific Conference Series, Vol. 101, Astronomical Data Analysis Software and Systems V, ed. G. H. Jacoby & J. Barnes, 88

Law, N. M., Kraus, A. L., Street, R. R., Lister, T., Shporer, A., Hillenbrand, L. A., & Palomar Transient Factory Collaboration. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 448, 16th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun, ed. C. Johns-Krull, M. K. Browning, & A. A. West, 1367

Lazorenko, P. F., & Sahlmann, J. 2018, A&A, 618, A111

Lee, A. T., Offner, S. S. R., Kratter, K. M., Smullen, R. A., & Li, P. S. 2019, ApJ, 887, 232

Lee, Y.-N., & Hennebelle, P. 2018, A&A, 611, A89

Lidov, M. L. 1962, Planet. Space Sci., 9, 719

Lindegren, L. et al. 2018, A&A, 616, A2

López-Morales, M., & Ribas, I. 2005, ApJ, 631, 1120

Lubin, J. B. et al. 2017, ApJ, 844, 134

Lubow, S. H., & Artymowicz, P. 1996, in NATO Advanced Study Institute (ASI) Series C, Vol. 477, Evolutionary Processes in Binary Stars, ed. R. A. M. J. Wijers, M. B. Davies, & C. A. Tout, 53

Lubow, S. H., & Artymowicz, P. 2000, in Protostars and Planets IV, ed. V. Mannings, A. P. Boss, & S. S. Russell, 731

Ma, B., & Ge, J. 2014, MNRAS, 439, 2781

Mamajek, E. E. 2009, in American Institute of Physics Conference Series, Vol. 1158, Exoplanets and Disks: Their Formation and Diversity, ed. T. Usuda, M. Tamura, & M. Ishii, 3–10

Mann, A. W. et al. 2019, ApJ, 871, 63

Mann, A. W., Feiden, G. A., Gaidos, E., Boyajian, T., & von Braun, K. 2015, ApJ, 804, 64

Martinache, F., Rojas-Ayala, B., Ireland, M. J., Lloyd, J. P., & Tuthill, P. G. 2009, ApJ, 695, 1183

Mason, B. D., Gies, D. R., Hartkopf, W. I., Bagnuolo, William G., J., ten Brummelaar, T.,

& McAlister, H. A. 1998, AJ, 115, 821

Mason, B. D., Hartkopf, W. I., Gies, D. R., Henry, T. J., & Helsel, J. W. 2009, AJ, 137, 3358

Mason, B. D., Wycoff, G. L., Hartkopf, W. I., Douglass, G. G., & Worley, C. E. 2001, AJ, 122, 3466

Mathieu, R. D., & Mazeh, T. 1988, ApJ, 326, 256

Mathieu, R. D., Stassun, K., Basri, G., Jensen, E. L. N., Johns-Krull, C. M., Valenti, J. A., & Hartmann, L. W. 1997, AJ, 113, 1841

Mazeh, T. 2008, in EAS Publications Series, Vol. 29, EAS Publications Series, ed. M. J. Goupil & J. P. Zahn, 1–65

Mazeh, T. et al. 2001, MNRAS, 325, 343

McDonald, J. M., & Clarke, C. J. 1995, MNRAS, 275, 671

Meibom, S., & Mathieu, R. D. 2005, ApJ, 620, 970

Moe, M., & Di Stefano, R. 2017, ApJS, 230, 15

Moe, M., & Kratter, K. M. 2018, ApJ, 854, 44

Morales, J. C. et al. 2009, ApJ, 691, 1400

Muñoz, D. J., Kratter, K., Vogelsberger, M., Hernquist, L., & Springel, V. 2015, MNRAS, 446, 2010

Muirhead, P. S., Dressing, C. D., Mann, A. W., Rojas-Ayala, B., Lépine, S., Paegert, M., De Lee, N., & Oelkers, R. 2018, AJ, 155, 180

Newton, E. R., Irwin, J., Charbonneau, D., Berta-Thompson, Z. K., Dittmann, J. A., &

West, A. A. 2016, ApJ, 821, 93

Nidever, D. L., Marcy, G. W., Butler, R. P., Fischer, D. A., & Vogt, S. S. 2002, ApJS, 141, 503

Offner, S. S. R., Dunham, M. M., Lee, K. I., Arce, H. G., & Fielding, D. B. 2016, ApJ, 827, L11

Offner, S. S. R., Klein, R. I., McKee, C. F., & Krumholz, M. R. 2009, ApJ, 703, 131

Offner, S. S. R., Moe, M., Kratter, K. M., Sadavoy, S. I., Jensen, E. L. N., & Tobin, J. J. 2022, arXiv e-prints, arXiv:2203.10066

Olczak, C., Pfalzner, S., & Spurzem, R. 2006, ApJ, 642, 1140

Öpik, E. 1924, Publications of the Tartu Astrofizica Observatory, 25, 1

Pascucci, I. et al. 2016, ApJ, 831, 125

Pass, E. K., Charbonneau, D., Irwin, J. M., & Winters, J. G. 2022, ApJ, 936, 109

Perryman, M. 2012, European Physical Journal H, 37, 745

Pineda, J. E. et al. 2015, Nature, 518, 213

Popper, D. M. 1980, ARA&A, 18, 115

Raghavan, D. et al. 2010, ApJS, 190, 1

Reid, I. N. et al. 2003, AJ, 126, 3007

Reid, I. N., Hawley, S. L., & Gizis, J. E. 1995, AJ, 110, 1838

Reynolds, N. K. et al. 2021, ApJ, 907, L10

Ribas, Á., Bouy, H., & Merín, B. 2015, A&A, 576, A52

Ribas, I. et al. 2023, A&A, 670, A139

Rizzuto, A. C. et al. 2013, MNRAS, 436, 1694

Rodriguez, J. E. et al. 2018, ApJ, 859, 150

Sahlmann, J. et al. 2015, A&A, 579, A61

Ségransan, D., Delfosse, X., Forveille, T., Beuzit, J. L., Udry, S., Perrier, C., & Mayor, M. 2000, A&A, 364, 665

Shatsky, N., & Tokovinin, A. 2002, A&A, 382, 92

Skinner, J. et al. 2018, AJ, 156, 45

Skrutskie, M. F. et al. 2006, AJ, 131, 1163

Sperauskas, J., Deveikis, V., & Tokovinin, A. 2019, A&A, 626, A31

Tokovinin, A. 2014, AJ, 147, 86

——. 2016, Orbit: Idl Software For Visual, Spectroscopic, And Combined Orbits, Zenodo

——. 2018, PASP, 130, 035002

Tokovinin, A., Cantarutti, R., Tighe, R., Schurter, P., Martinez, M., Thomas, S., & van der Bliek, N. 2016, PASP, 128, 125003

Tokovinin, A., & Kiyaeva, O. 2016, MNRAS, 456, 2070

Tokovinin, A., Mason, B. D., & Hartkopf, W. I. 2010, AJ, 139, 743

Tokovinin, A., Mason, B. D., Hartkopf, W. I., Mendez, R. A., & Horch, E. P. 2015, AJ, 150, 50

Tokovinin, A., Mason, B. D., Mendez, R. A., & Costa, E. 2022, AJ, 164, 58

Tokovinin, A., Mason, B. D., Mendez, R. A., Costa, E., & Horch, E. P. 2020, AJ, 160, 7

Tokovinin, A., Mason, B. D., Mendez, R. A., Costa, E., Mann, A. W., & Henry, T. J. 2021,

AJ, 162, 41

Tokovinin, A., & Moe, M. 2020, MNRAS, 491, 5158

Tokovinin, A. A. 1997, A&AS, 121, 71

Udry, S., Mayor, M., Delfosse, X., Forveille, T., & Perrier-Bellet, C. 2000, IAU Symposium, 200, 158

Valtonen, M., & Mikkola, S. 1991, ARA&A, 29, 9

Van De Kamp, P. 1967, Principles of astrometry (San Francisco, CA: Freeman)

Vrijmoet, E. H., Henry, T. J., Jao, W.-C., & Dieterich, S. B. 2020, AJ, 160, 215

Vrijmoet, E. H., Tokovinin, A., Henry, T. J., Winters, J. G., Horch, E., & Jao, W.-C. 2022, AJ, 163, 178

Weis, E. W. 1991, AJ, 101, 1882

——. 1996, AJ, 112, 2300

Whitworth, A., Bate, M. R., Nordlund, Å., Reipurth, B., & Zinnecker, H. 2007, in Protostars and Planets V, ed. B. Reipurth, D. Jewitt, & K. Keil, 459

Winters, J. G. et al. 2019, AJ, 157, 216

——. 2020, AJ, 159, 290

Xia, F., Fu, Y., & Wang, X. 2019, ApJ, 882, 147

Zahn, J. P. 1977, A&A, 57, 383

Zahn, J. P., & Bouchet, L. 1989, A&A, 223, 112

Zhao, B., & Li, Z.-Y. 2013, ApJ, 763, 7

Zhou, G. et al. 2015, MNRAS, 451, 2263