Spots and Flares: Stellar Activity in the Time Domain Era

James R. A. Davenport

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

Doctor of Philosophy

University of Washington

2015

Reading Committee:

Suzanne L. Hawley, Chair

Leslie Hebb

Željko Ivezić

Program Authorized to Offer Degree:
Astronomy

University of Washington

Abstract

Spots and Flares: Stellar Activity in the Time Domain Era

James R. A. Davenport

Chair of the Supervisory Committee:
Professor Suzanne L. Hawley
Department of Astronomy

Time domain photometric surveys for large numbers of stars have ushered in a new era of statistical studies of astrophysics. This new parameter space allows us to observe how stars behave and change on a human timescale, and facilitates ensemble studies to understand how stars change over cosmic timescales. With current and planned time domain stellar surveys, we will be able to put the Sun in a Galactic context, and discover how typical or unique our parent star truly is.

The goal of this thesis is to develop techniques for detecting and analyzing the most prominent forms of magnetic activity from low-mass stars in modern time domain surveys: starspots and flares. Magnetic field strength is a fundamental property that decays over a star's life. As a result, flux modulations from both flares and starspots become smaller amplitude and more infrequent in light curves. Methods for detecting these forms of magnetic activity will be extensible to future time domain surveys, and helpful in characterizing the properties of stars as they age.

Flares can be detected in sparsely sampled wide field surveys by searching for bright single-point outliers in light curves. Using both red optical and near infrared data from ground-based surveys over many years, I have constrained the rate of flares in multiple wavelengths for an ensemble of M dwarfs. Studying flares in these existing ground-based datasets will enable predictions for future survey yields.

Space-based photometry enables continuous and precise monitoring of stars for many years, which is crucial for obtaining a complete census of flares from a single star. Using 11 months of 1-minute photometry for the M dwarf GJ 1243, I have amassed over 6100 flare events, the largest sample of white light flares for any low-mass star. I have also created the first high fidelity empirical white light flare template, which shows three distinct phases in typical flare light curves. With this template, I demonstrate that complex multi-peaked flares can be decomposed into their constituent flare events. This is the first modern study of the detailed white light morphology of stellar flares.

Space-based survey data is also ideal for studying starspots, whose photometric modulation amplitude is typically much smaller than for flares. Using 4 years of 30 minute photometry for GJ 1243, I have traced the sizes and longitudes for multiple large starspots. A primary starspot is found that is stable in position and size over the 4 years of data, as well as secondary starspot features that decay on 100 to 500 day timescales and evolve in longitude. The secular longitude evolution of the secondary starspots indicates a very low rate of differential rotation on this rapidly rotating low-mass star.

The presence of a transiting exoplanet can provide a great deal of information about the sizes and locations of starspots on the host star. When the planet crosses in front of a starspot, a small deviation in the predicted transit light curve is observed. By tracing these transit anomalies, I have detected more than 100 distinct starspots in 4 years of data on the young G2 dwarf, Kepler 17. These starspots are up to an order of magnitude larger than those on the Sun, and this star shows an almost four times larger amplitude of differential rotation than on the Sun.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

This document represents only a small piece of many years of work. However, none of this work would have been possible without the patience, love, and guidance from an enormous group of people. I am indebted to everyone who helped me both personally and professionally along this journey, and would like to highlight a few specific individuals:

Thanks to the members of the "Red Army" research group who first inspired me to pursue research and ask better questions, including Andrew West, Kevin Covey, John Bochanski, Lucianne Walkowicz, and Chris Laws. Thanks to my fellow graduate students in this group who helped keep me sane, and stayed up for many late nights of observing and jocularity, including Eric Hilton, Adam Kowalski, and Sarah Schmidt.

I have been fortunate to travel to telescopes with nearly every one of my fellow graduate students during my time at UW. Thank you all for the fascinating conversations on the mountain and fun-filled long drives, which made each trip a joy. I have especially loved the many deep and wandering conversations about science and life with Cliff Johnson, and about stars and side projects with John Lurie.

For over 5 years I had the absolute joy of sharing the office PAB341, playing countless foosball games, and going on long lunches with John Ruan. You have been my dearest friend in graduate school, and I will miss our constant scheming and brainstorming. I hope we make good on even 1% of the hair-brained ideas we came up with.

Thanks to Andy Becker who not only generously funded my research for several years and multiple conference trips, but expanded my scientific horizons and encouraged wild ideas and methods. I owe so many of my skills to your patient guidance. Thanks to Željko Ivezić who helped me learn the value of being interested in everything.

Thank you to my advisor, Leslie Hebb, for flying to far-flung places in the world just to meet with me. Your passion and keen guidance helped focus my thesis into a coherent body of work. More importantly, you reignited a flame of interest in science when I was most uncertain in my progress.

For the last 10 years I have had the greatest mentor, role model, and dear friend in Suzanne Hawley. I have loved our many trips to conferences, the beach, observing, and even lunch. Each week you pushed me to do more and do better. Thank you for taking a chance on me, giving me a career, spotting every typo, and teaching me so many valuable lessons in science and life.

Thank you to my family, Emery, Joan, and Phillip Davenport, and Steve, Kathy, David, and Erin Miner, who loved me every day of this journey.

Most of all, thank you to my wife and best friend, Sarah Davenport. Thank you for picking me up every time I fell, and cheering every time I succeeded. You give meaning to all that I do, and inspire me every day with your strength and love.

DEDICATION

To my best friend, Sarah.

Chapter 1

INTRODUCTION

"If the sun didn't have a magnetic field, then it would be as boring a star as most astronomers think it is." [attributed to R. B. Leighton]

The goal of this thesis is to advance statistical studies of magnetic activity in cool stars towards the next generation of large scale, multi-band photometric survey databases. Activity is observed here as photometric brightness variations, which originate from rapid flares (brightening) and cool spots (dimming) on the stellar surface. Through time domain studies of both specific objects and large ensemble datasets, I have worked to develop the tools and methodology that will be useful in mining future datasets for these signatures of magnetic activity. This includes analyzing ground-based wide-field optical and infrared datasets that contain sparsely sampled multi-band photometry over many years, as well as cutting edge space-based photometry of select objects with densely sampled single-band light curves. Importantly, as I have worked to develop new tools and techniques, I have also determined fundamental properties of stellar activity for many low-mass stars. These properties provide constraints for stellar dynamo and flare models, and establish benchmarks for future observational studies.

In this first chapter I provide a broad introduction to the concepts discussed throughout the thesis. I begin with a general overview of low-mass stars, followed by a discussion of stellar dynamos and the role of rotation and differential rotation. I then describe the observable forms of magnetic activity used in my research, and a brief history of their study in stars beyond the Sun. Finally, I describe the various datasets I have used, and provide a general outline for the rest of the thesis.

1.1 Overview of Magnetic Fields in Low-Mass Stars

Magnetic fields are ubiquitous among low-mass stars, and are a fundamental component of their main sequence lives. In this thesis I study some of the observable surface manifestations of stellar magnetic fields, known generally as "magnetic activity". The surface activity is driven by the underlying magnetic dynamo, which on the Sun we observe on both global and very small scales. Here I provide an overview of the properties of low-mass stars, and of two basic magnetic dynamo models at work in low-mass stars.

The most familiar low-mass star is the Sun, an apparently typical main sequence star with spectral type G2 (Morgan et al., 1943; Morgan & Keenan, 1973). Besides anchoring the main sequence spectral classification system (from hottest to coolest: OBAFGKMLTY), the Sun is the cornerstone on which our intuition for the physics of stellar atmospheres and magnetic activity in other stars rests. It is the only star whose magnetic field structure can be spatially resolved, and local features tracked throughout their entire evolution. The Sun is also unique in that its internal structure has been studied in detail, and is the only star where the boundary or interface layer between the radiative core and convective outer envelope (known as the "tachocline") has been directly measured via helioseismology (Spiegel & Zahn, 1992; Mazumdar & Antia, 2001).

In this thesis I adopt the broad definition of "low-mass stars" to be those having convective outer envelopes, which are critical for the generation of solar type stellar magnetic fields. Outer convection zones are present in main sequence stars with masses below $\sim$1.2 $M_\odot$ (late F dwarfs). Approximate mass ranges for low-mass stars are 0.8–1.04 $M_\odot$ for G dwarfs, 0.45–0.8$M_\odot$ for K dwarfs, and 0.08–0.45$M_\odot$ for M dwarfs. In all low-mass stars, core hydrogen fusion proceeds primarily via the proton-proton chain process. This fusion process is less sensitive to the core temperature and less efficient at burning hydrogen than the CNO cycle found in higher mass stars, and as a result low-mass stars have long lifetimes on the main sequence, from $\sim 10^{10}$ years in G dwarfs to $>10^{12}$ years for M dwarfs. The lower mass limit for stars represents the boundary between the stellar and brown dwarf regime (spectral

types LTY), where core temperatures and pressures can no longer support hydrogen fusion.

Spectral types G, K, and early M are thought to have a solar-like internal structure, with an outer convective envelope and inner radiative zone around the core. For the Sun the convection zone reaches a depth of 0.287 $R_\odot$ below the surface (Kosovichev & Fedorova, 1991). The fractional radial extent of the convective envelope increases with later spectral type (decreasing stellar mass), until the convection zone extends fully from the surface to the stellar core at a mass around 0.35 $M_\odot$ (spectral type M3–M4) (e.g. Reiners & Basri, 2009). The "fully convective" limit, and the differences between stars on either side of this limit, are not well defined observationally. The very central core in these fully convective stars may still be radiative, but directly studying their detailed internal structure with observations has not been possible.

1.1.1 Measuring Fundamental Properties of Low-Mass Stars using Photometry

The origin and strength of the magnetic dynamo is fundamentally tied to the internal structure of a star. However, because stellar atmospheres become optically thick very quickly (indeed, this is the definition of the photospheric surface), and we lack a scalpel long enough to reach them, we cannot directly observe the interior of a star. At present, the only means to probe the interior of stars is by using seismic information, characterizing the propagation of waves throughout the stellar interior to infer the radial density profile and rotation structure. This is analogous to seismology used in geological studies of the Earth's interior, and is called "helioseismology" for the Sun, and "asteroseismology" for other stars. This blossoming field of asteroseismology will be one of the great legacies of the planet-hunting era, and will produce constraints on models of stellar interiors. Since the surface magnetic field morphology is strongly influenced by the internal rotation and density structure of the star, it is important to understand what asteroseismology is teaching us. Therefore, as a short aside I provide a brief overview of how the fundamental properties of low-mass stars may be studied with this fascinating technique.

The Sun is the only star whose internal radial structure has been accurately determined

using spatially resolved seismic data. For others stars whose surfaces are unresolved, asteroseismology currently provides precise *model dependent* estimates of stellar mass, radius, and age (Chaplin & Miglio, 2013). These general stellar properties are determined from scaling relations calibrated on observations of the Sun, which can be written:

$$\frac{M}{M_{\odot}} \approx \left(\frac{\nu_{max}}{\nu_{max,\odot}}\right)^{3} \left(\frac{\Delta\nu}{\Delta\nu_{\odot}}\right)^{-4} \left(\frac{T_{eff}}{T_{eff,\odot}}\right)^{3/2}, \tag{1.1}$$

$$\frac{R}{R_{\odot}} \approx \left(\frac{\nu_{max}}{\nu_{max,\odot}}\right) \left(\frac{\Delta\nu}{\Delta\nu_{\odot}}\right)^{-2} \left(\frac{T_{eff}}{T_{eff,\odot}}\right)^{1/2}, \tag{1.2}$$

where ν_{max} is the frequency of the asteroseismic waves in the time series (a.k.a. light curve) power spectrum with the maximum power, and $\Delta\nu$ the characteristic separation between peaks in the power spectrum (Silva Aguirre et al., 2011). These relations are calibrated using the well-known 5-min primary pulsation mode of Sun, with hotter or more massive stars exhibiting larger amplitude and higher frequency pulsations (Christensen-Dalsgaard & Frandsen, 1983). Since *Kepler* has produced long duration light curves with high precision (Borucki et al., 2010), these pulsations have been successfully detected for giant and subgiant stars using the long cadence (30-min) observations. With the shortest cadence *Kepler* observations (1-min), pulsations in solar-mass and slightly higher mass main sequence stars have been detected. However, this model-dependent asteroseismology technique has not been successfully applied to stars with masses smaller than the Sun, as their pulsation amplitudes are too weak to detect, even with several years of space-based photometric monitoring.

To determine the internal rotational structure of stars, as has been done in the Sun, one must measure fine structure in the power spectrum in the form of the splitting of single peaks (pulsation modes). In the Sun, helioseismology has shown the radiative zone to rotate as a solid body, while the convective envelope shows latitude- and radius-dependent differential rotation throughout (see for example the review by Thompson et al., 2003). Observational asteroseismology studies (model *independent*) have started to directly detect the internal structure and core rotation for slowly rotating higher mass stars (Kurtz et al., 2014). This measurement has only been possible for a handful of stars, which have revealed that the stellar surface appears to rotate more rapidly than the core. However, these are especially

slow rotators, and may not be typical among low-mass stars.

Asteroseismology has begun to detect power spectrum frequency splitting for solar type stars (e.g. Nielsen et al., 2014), but has not yet been able to determine the internal differential rotation (core versus envelope) or detailed internal structure. Measuring the basic properties (masses, radii) and detailed internal structure for lower mass K and M stars using asteroseismology is still not possible even with the precision data from *Kepler*. Note also the depth of the convective zone for low-mass stars, which is so critical in the generation of magnetic activity, as well as regulating the stellar angular momentum evolution (discussed below), has not been constrained observationally for any star besides the Sun. Still, asteroseismology from future generations of time domain studies may be our best hope for detecting the internal structure of other low-mass stars.

1.1.2 Mean-Field Dynamo Theory

The basic mechanics of the mean-field dynamo theory in the Sun were laid down first by Parker (1955). Here "mean-field" refers to a dynamo model that only predicts large scale magnetic features, which are necessarily much larger than the surface turbulence size scales or individual sunpots. This is one of the oldest, simplest, and most popular dynamo models used to explain the general magnetic field properties of the Sun. While mean-field dynamo theory is certainly not the *only* model used in current studies of the Sun (Charbonneau, 2005), it provides an excellent framework for understanding the global properties of the solar dynamo. For solar-type stars, mean-field dynamo theory postulates that differential rotation will convert an existing dipolar field into an almost fully toroidal one. This is known as the Ω effect because the Greek letter Ω is commonly used in physics to represent rotation. This differential rotation occurs both at the tachocline interface between the convective envelope and the (solid-body rotating) radiative interior, and radially throughout the convective zone to the surface where it is most readily observed. In Figure 1.1 I show schematic examples of the solar surface and radial differential rotation.

At or near the tachocline, toroidal flux tubes eventually become strong enough that they

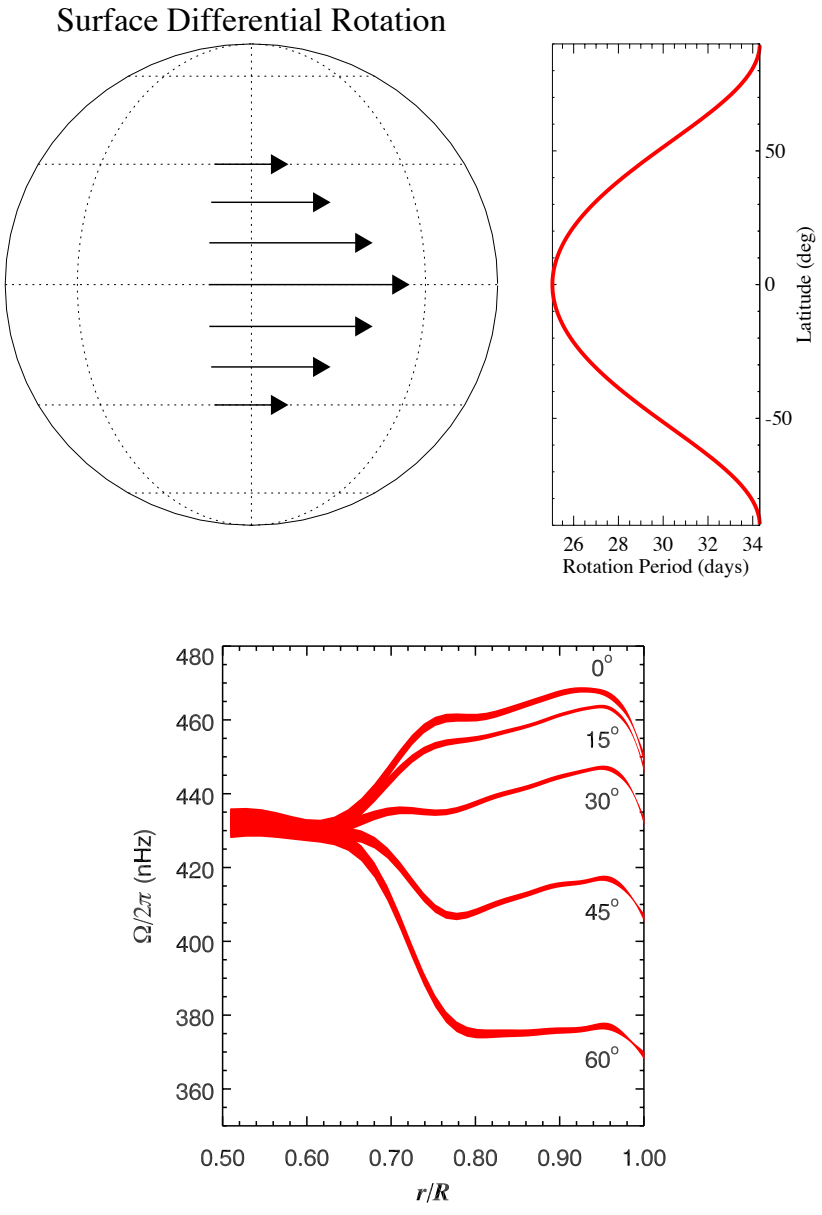

Figure 1.1: Top: Schematic diagrams of the surface differential rotation on the Sun, where the equator rotates at ∼25.05 days, and the poles at ∼34.1 days. Bottom: Radial differential rotation determined in the Sun as a function of latitude (Figure courtesy of the National Solar Observatory and the NSF).

exert a buoyant force on the surrounding gas and begin to rise towards the surface. Convective or cyclonic motions throughout the envelope work to twist the rising flux tubes (also called "helicity"). This is known as the α effect because of the resulting twisted shape of the magnetic field lines at the surface, which resemble the Greek letter α. These flux tubes pierce the photosphere, creating starspot pairs with opposite polarity at the footpoints of their emergence. This twisting in the flux tube also produces a preferential, latitudinally dependent *tilt* of the starspot pair, where the leading spot is closer to the equator (slightly lower latitude), and where the inclination of the tilt becomes more pronounced with increasing latitude of the spot pair (known as Joy's law), and is a critical component of the famous Babcock–Leighton mean-field solar dynamo model (Babcock, 1961; Leighton, 1964).

While the majority of these small magnetic features or spots neutralize each other, a small fraction migrate to the poles via surface meridional flow. This builds up magnetic flux at the stellar poles that is preferentially of the opposite polarity from the initial dipolar field. The result is the gradual conversion of the toroidal magnetic field back to a poloidal field, which allows the process to begin again.

The combination of these effects at work in the Sun, i.e. the differential rotation and convectively driven helicity, forms the $\alpha\Omega$ dynamo model. The $\alpha\Omega$ dynamo requires both rotation and differential rotation to operate, as well as the presence of a tachocline interface region to create additional shear and store the toroidal flux until the flux tubes are strong enough to buoyantly rise. Many observations of the Sun are successfully explained within this theoretical framework. Joy's law of increasing spot inclinations as a function of latitude occurs due to differential rotation winding the surface magnetic field. As sunspots migrate towards the poles, magnetic flux is deposited from the spot pairs with polarity preferentially opposite that of the poles. Over time this opposite polarity builds up, and gives rise to the 11-years periodic magnetic polarity reversal (solar cycle). However, the specific nature of the α effect, whether it is driven by convective cell motions throughout the envelope, or through the Coriolis force acting on the flux tubes due to the stellar rotation, is still debated (Charbonneau, 2005).

In the case of fully convective stars, those with masses smaller than 0.35 $M_{\odot}$, the $\alpha\Omega$ dynamo model likely can not function. These stars lack a tachocline interface region to build up toroidal magnetic fields, and are predicted to have less differential rotation both radially and on the surface. The dynamo may instead be driven by the convection that dominates the stellar interior (Durney et al., 1993). In the case of the $\alpha\Omega$ dynamo, the Ω mechanism is most responsible for converting poloidal magnetic flux to toroidal flux via shear and differential rotation, while the α mechanism generates large scale poloidal magnetic flux via convection and rotation. In fully convective low-mass stars, an α^2 dynamo is thought to operate (Christensen et al., 2009), which relies on the convective turbulence to convert between both toroidal and poloidal magnetic flux (Brandenburg & Subramanian, 2005). The Ω effect may also contribute to the global or mean magnetic fields in these stars, however. Thus, an $\alpha^2\Omega$ dynamo where convection and rotation both play dominant roles, while weak differential rotation helps to organize the large scale field, cannot be ruled out.

The surface morphology of magnetic fields produced by an α^2 dynamo are still debated. Many simulations have shown that an α^2 dynamo will produce non-axisymmetric fields on the surface for even modest rotation rates (e.g. Chabrier & Küker, 2006). Other models, however, find that weak differential rotation is present in fully convective stars, and that these dynamos still possess significant axisymmetric magnetic fields (Browning, 2008).

1.1.3 The Role of Rotation and Differential Rotation in Low-Mass Stellar Dynamos

As mentioned above, rotation plays a key role in both the $\alpha\Omega$ and α^2 dynamo models as a key part of the α effect. For both dynamo models, faster rotation is also linked to larger total magnetic field strengths. This result is found both observationally and theoretically (e.g. Browning et al., 2010). However, for fully convective stars with very rapid rotation rates, the dynamo appears to reach a maximum strength (also known as "saturating"), and is not affected by the specific rotation rate until the star slows below some critical value (Chabrier & Küker, 2006; Reid & Hawley, 2005; Wright et al., 2011). An example of magnetic activity saturation, as exhibited by coronal X-ray emission for rapidly rotating stars, is shown in

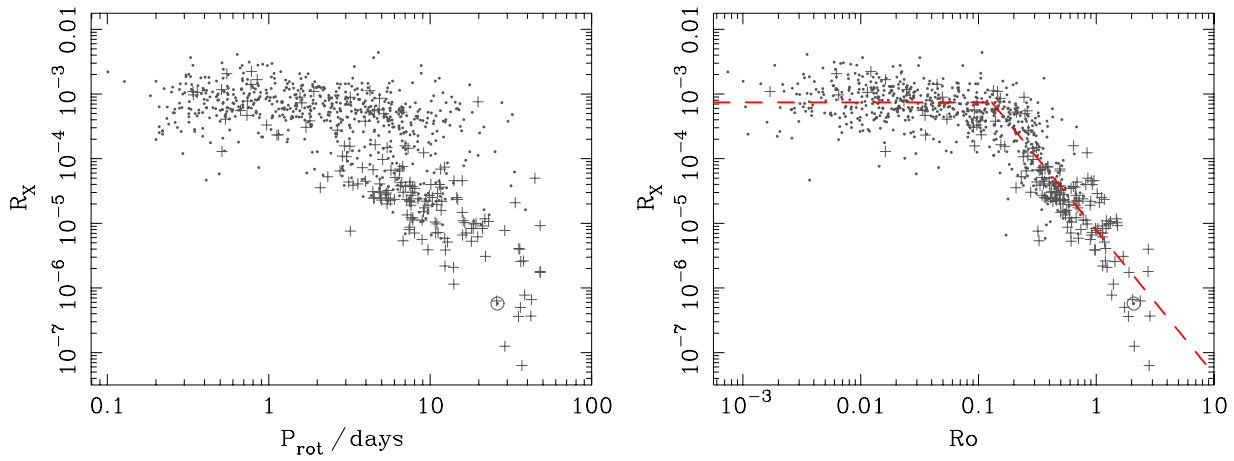

Figure 1.2: Left: ratio of X-ray to bolometric luminosity (R_X) versus rotation period. Right: R_X versus the stellar Rossby number, Ro = P_{rot}/τ, where τ is the convective turnover timescale. Note the saturation in the X-ray flux ratio at low Rossby numbers or rapid rotation. In both panels, known binary stars are shown as crosses. Figure from Wright et al. (2011), © AAS. Reproduced with permission.

Figure 1.2 from Wright et al. (2011).

For all low-mass stars, differential rotation in the convective envelope should occur as a direct consequence of combining stellar rotation and convection. However, the convective turnover timescale (τ_c) increases with decreasing stellar mass (Noyes et al., 1984; Küker & Rüdiger, 2008). In fully convective stars the Rossby number ($Ro = P_{rot}/\tau_c$) is very small, so the star is in a "fast" rotation state even for slowly rotating stars. In this case horizontal (as opposed to radial) angular momentum transport dominates throughout the convective envelope, and a much lower amplitude of differential rotation is predicted (Küker & Rüdiger, 2008; Kitchatinov & Olemskoy, 2011). Additionally, strong dynamos in the case of rapidly rotating stars may also suppress differential rotation by providing magnetic resistance against shearing layers (Browning, 2008), and so rapidly rotating M dwarfs are predicted to exhibit almost no surface differential rotation. For solar-type stars with rotation periods greater

than a few days, the amount of surface differential rotation shear is expected to decrease with slower rotation (Küker & Rüdiger, 2011).

1.2 Observable Manifestations of Stellar Magnetic Activity

For the Sun we are able to measure the spatial structure, strength, and polarity of the surface magnetic field directly, with instruments such as the Helioseismic and Magnetic Imager (Scherrer et al., 2012), and even trace the magnetic evolution of specific active regions. From this data we may also trace the global magnetic field strength and total solar irradiance throughout the Solar cycle (e.g. Norton et al., 2013). Still, these sophisticated measurements only reveal the surface magnetic field strength and topology, and to understand the internal magnetic field structure (both for the Sun and other stars) we must carefully interpret the story that these surface features tell.

The study of magnetic activity on the Sun has historically included observations of the number and position of sunspots dating back to ∼1610 (Hoyt et al., 1994). This has led to the discovery of both periodic activity such as the 11-year solar cycle, and (thus far) aperiodic events such as the Maunder Minimum (Eddy, 1976). Tracing the location, motion, and evolution of these cool surface features has been a cornerstone in our understanding of the stellar dynamo. The Sun also produces frequent flares, which have been studied for over a century and a half (Carrington, 1859). These rapid outbursts are typically small in energy (relative to the quiescent solar flux), and occur physically near sunspots.

In the case of stars beyond our Sun where we cannot resolve the stellar surface, the topology of the magnetic field is very difficult to constrain. For nearby active low-mass stars the Zeeman-Doppler imaging (ZDI) technique can be used to construct the best available maps of large scale magnetic field strengths and orientations for individual stars (Semel, 1989). This method relies on extensive spectro-polarimetry observations to measure the modulations of polarized absorption lines as the star rotates, and can only be carried out on nearby stars with fast rotation rates (Donati & Brown, 1997). Total magnetic field strengths for low-mass stars can be inferred via the strength of selected absorption bands (e.g. FeH)

or chromospheric emission lines (e.g. Hα) in their spectrum (Reiners & Basri, 2007).

While direct measurement of the magnetic field strength for single stars requires time consuming spectroscopy or spectro-polarimetry observations, the presence of stellar magnetic activity can be easily seen using photometry. For rapidly rotating stars, approximate positions and sizes of starspots are now frequently inferred due to small decreases in the integrated light of the star (flux modulations up to a few %). Similarly, flares are observed as transient increases in both the total flux (white light) and specific spectroscopic features, and can reach peak photometric amplitudes of several magnitudes in blue optical bands (see for example Hawley & Pettersen, 1991; Kowalski et al., 2010).

In this section I will briefly describe the generation of stellar spots and flares, which I have studied in detail for this thesis using time domain databases for low-mass stars. I will also review other observable characteristics of magnetic activity in surveys of low-mass stars.

1.2.1 Flares

The first observation of a solar flare in the modern astronomical literature was recorded by Carrington (1859), who referred to it as a "Singular Appearance". This was one of the largest energy flare events within the past 200 years, and the event generated what is known as a Solar storm in space weather terminology. The associated coronal mass ejection hit Earth, and induced aurorae so powerful they were reportedly visible during the day and at low latitudes, and disrupted telegraph lines worldwide (Green & Boardsen, 2006). Thankfully solar flares and solar storms of this magnitude occur infrequently, perhaps only once per $\sim 10^2$ years.

Flares, as I discuss them in this thesis, are called "white-light flares", due to their detection in broadband visible light photometry. The standard flare model that is presumed to be at work on all low-mass stars describes a flare as the result of a sudden and explosive release of energy due to magnetic reconnection events above the stellar surface (Martens & Kuin, 1989; Benz, 2008). This event rapidly rearranges magnetic loops in the upper chromosphere. These loops originate from starspot pairs on the stellar surface. Particles trapped within

the pre-flare chromosphereic loops are then accelerated along the newly "reconnected" loops toward the photosphere, forming a beam of energetic electrons that impacts and explosively heats the stellar surface. This series of events causes emission at nearly every wavelength, including radio, X-ray, optical, and UV. The specific origin of the white light continuum emission commonly observed in stellar flares is poorly constrained on the Sun, due in part to a lack of comparable white-light "Sun-as-a-star" observations. However, this standard flare model usually acknowledges the impact foot-points of the electron beam as the location of the white light emission.

Flares on low-mass stars besides the Sun have been studied for almost 75 years. The canonical variable in the class of flare stars is UV Ceti, also known as Luyten 726-8 (Joy & Humason, 1949), a nearby M5 dwarf and a member of a binary star system. Active dMe dwarfs, those with the chromospheric Hα line and Ca II H & K lines in emission (e.g. see Giampapa et al., 1989) were recognized early on for exhibiting dramatic brightening in their blue continuum flux, as well as a strong increase in the chromospheric emission lines (Joy & Humason, 1949; Solomon, 1966). Because of this chromospheric emission and the hot effective temperatures implied during the flare events by the increased blue continuum flux, a connection to the Sun was made almost immediately, with early studies speculating that these stochastic events were probably related to Solar flares (Kron, 1950a; Schatzman, 1959). By the 1970's, a general foundation for the energies and timescales at work in stellar flares, and their relation to processes observed on the Sun had been proposed (Kunkel, 1970; Gershberg, 1975).

Spectroscopy is important for understanding the underlying physics of the white light flare emission, as it gives a constraint on the flare color (effective temperature of the continuum flux), as well as the emission line energies. These features must also be measured as a function of time throughout the flare event to constrain the total energy and affected coverage area on the stellar surface. However, obtaining high quality spectroscopy throughout flares requires laborious monitoring of active stars for many nights (e.g. Kowalski et al., 2013). The largest of these flares for active M dwarfs can reach flux increases in the blue optical of 4–5 magnitudes,

having total event energies of $>10^{34}$ erg, and showing emission in the entire hydrogen Balmer series and dozens of other lines (Hawley & Pettersen, 1991; Kowalski et al., 2010). Observing these large stellar flares is rare, and capturing such events with spectroscopy is equal parts dedication and luck.[1]

Observing stellar flares with photometry, and especially multi-band photometry, is also very fruitful for studies of their color and energy evolution. Flare stars can be identified in very large area photometric catalogs by searching for single- or few-epoch brightness increases, which constrains the frequency of such stars throughout the Galaxy (Kowalski et al., 2009). Extensive flare studies from photometric monitoring can also constrain the distribution of flare event energies, which reveal an inverse power-law relationship between occurrence frequency and event energy (Lacy et al., 1976). This power-law distribution is found to hold over 9 orders of magnitude in flare energy (Schrijver, 2011; Schrijver et al., 2012).

The classic white light flare light curve profile is known colloquially as a "FRED": Fast Rise, Exponential Decay, shown in Figure 1.3. Indeed, some authors simply assume a light curve model that follow an instantaneous rise and single exponential decay profile when searching for flares in time domain datasets (e.g. Loyd & France, 2014). However, it has long been known that "normal", single-peaked flares *generally* follow two phases of exponential decay, a rapid or impulsive decay followed by a gradual decay in the luminosity (Houdebine et al., 1991). The initial rapid decay is typically of the same timescale as the flare rise phase, and this rapid rise and decay together are known as the "impulsive phase" of a white light flare. Hawley & Pettersen (1991) found that the impulsive phase can encompass approximately 2/3 of total emitted energy throughout the event, while only lasting for less than half of the event duration in large M dwarf flares. The remainder of the flare decay is known as the gradual phase, where the light curve follows a shallower exponential cooling curve.

---

[1] *"In the fields of observation, chance favors only the prepared mind."* [Louis Pasteur, 1854]

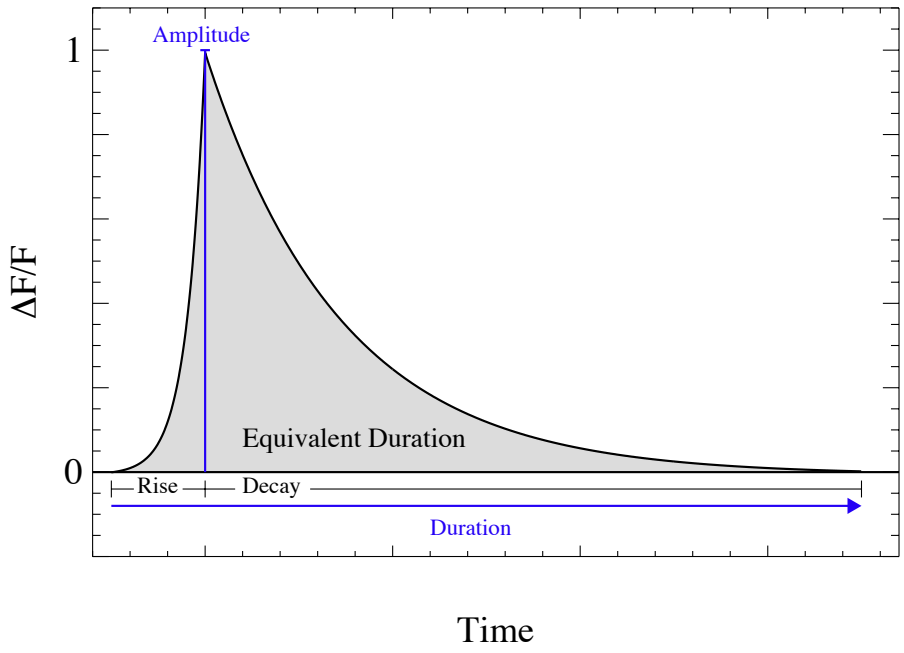

Figure 1.3: Idealized classical flare light curve "FRED" profile, with measurable quantities labeled – amplitude, rise and decay phases, total duration, and equivalent duration (integral of flare flux). Figure from Hawley et al. (2014), © AAS. Reproduced with permission.

1.2.2 Starspots

Starspots are cool regions on the stellar surface, generated as a byproduct of the stellar dynamo. A popular model for the formation of starspots is the local suppression of convection by magnetic flux tubes, which in turn cools a small portion of the stellar surface. These flux tubes appear as magnetic loops on the Sun, and the footprints of the loop form a pair of spots with a North/South magnetic dipole on the surface. A high resolution optical image of the Sun with a very large spot group is shown in Figure 1.4.

The archetype of this phenomena are sunspots, which have been cataloged with regularity for over 400 years (e.g. Hoyt et al., 1994). As with dramatic and colorful atmospheric aurorae from geomagnetic storms, spots on the Sun have been observed and written about throughout human history. Naked-eye sunspot observations have been found in historical chronicles dating back more than 2100 years (Wittmann & Xu, 1987; Vaquero, 2007). The earliest known drawing of a sunspot, however, dates to an observation from 1128 AD by John of Worcester (Stephenson & Willis, 1999). This active region must have been much larger than those typically observed on the Sun, and likely produced the corresponding geomagnetic aurorae recorded in Korean and Chinese chronicles that same year (Willis & Stephenson, 2001). The largest sunspot ever photographed to date was observed in 1947 (Hoge, 1947), and at maximum had a surface coverage area 3 times greater than the very large spot group seen in late 2014, shown in Figure 1.4.

Sunspots are characterized by two concentric regions, both of which are visible in the large central spot group in Figure 1.4. The central region of a starspot is known as the umbra, and is the darkest portion of the starspot. This is the area where the magnetic flux tubes are aligned nearly vertically, and create the greatest inhibition of convection. For spots on the Sun the umbra typically has a temperature contrast with the photosphere of $\sim$1700 K (Berdyugina, 2005). The outer portion of the spot is known as the penumbra, which surrounds the entire umbra. The penumbra appears lighter than the umbra, though still darker than the photosphere with a temperature contrast of $\sim$750 K. Here the magnetic flux

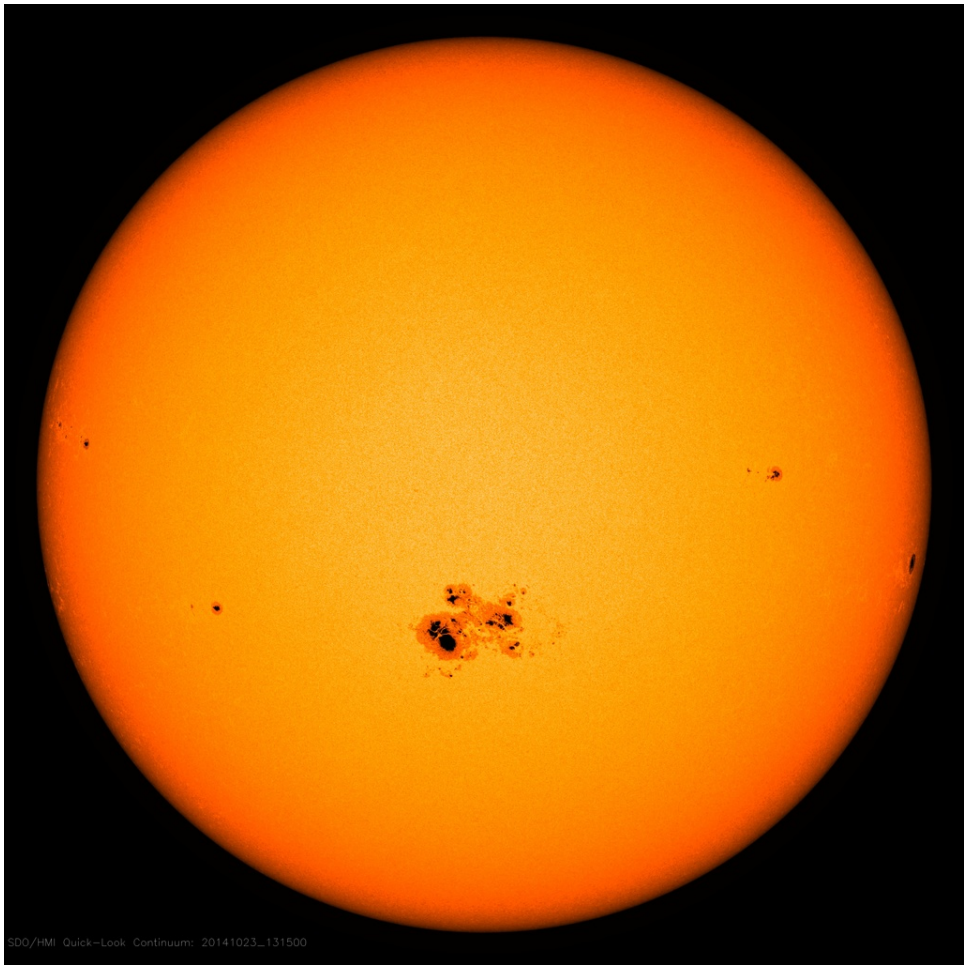

Figure 1.4: Large sunspot group from October 2014, the largest over the past 2 Solar cycles, observed by the space-based Solar Dynamics Observatory. Image from NASA/SDO.

tubes are more inclined towards the photospheric surface. The main spot group in Figure 1.4 shows several small dark umbra from individual starspot pairs, all encompassed in a large region of nearly continuous penumbra. The smaller, isolated spots in Figure 1.4 show the more traditional configuration of a single circular umbra surrounded by an annulus of penumbra, with a roughly equal radius for both.

Spots on the surface of other, more distant stars cannot be resolved. Instead we infer their presence by variations in the brightness of the star over time, due to the cool spots moving in and out of view as the star rotates. This results in periodic modulations in the stellar light curve, which are larger in amplitude if the starspots are larger or darker (Strassmeier, 2009). Starspots are also known to affect the light curves of eclipsing binary systems, causing the eclipse depths to be asymmetric or to change over time as the spots evolve. Photometric data indicating brightness modulations due to starspots dates back almost 70 years; they were first discovered on RS CVn type eclipsing binary systems that had very high amplitudes of magnetic activity (Kron, 1947, 1950b).

Sunspots are the sites referred to as "active regions" on the Sun, due to the strong magnetic fields present. The sunspots mark where the magnetic loops pierce the surface, and are often the footpoints of large coronal loops. Flare occur in or near these active regions, when the magnetic fields violently rearrange themselves. Indeed, the massive Carrington flare of 1859 was observed to occur coincident with a large spot complex present on the Sun at the time, and the apparent position and size of the spot shifted in response to the flare event.

1.2.3 Faculae and Plage

Besides explosive flare events, there are two other localized activity phenomena of note that can increase the stellar brightness. "Faculae" are often located near sunspots or active regions on the solar photosphere, and are small areas of strong magnetic field between convection flows (granules). These can be seen as either hot spots or networks of hot regions between convective cells. "Plage" are hot chromospheric regions that occur above and surround

sunspots, and are also due to concentrated magnetic fields. I describe them here briefly for completeness of thought, and with hope that they might be studied more in the future for other stars.

Both of these magnetic activity features create local bright spots that move in and out of view as the star rotates, and could result in sinusoidal "peaks" in a light curve instead of "dips" from starspots. We can not know a priori if sinusoidal modulations in a light curve are due to localized bright or dark regions, creating the "Zebra Effect" paradox coined by Pettersen et al. (1992). However, in the case of stars with transiting exoplanets (discussed in Chapter 5) we can definitively rule out hot spots as the cause of the light curve modulations.

Interestingly, during the phase of solar maximum when the Sun produces the highest occurrence rates of spots (and flares), faculae and plage actually cause the total solar irradiance to be slightly higher. The resulting increase in the optical brightness of the Sun during maximum is $\sim$0.04% (Foukal & Lean, 1986; Willson & Hudson, 1988). Networks of faculae can rival individual sunspot groups in terms of impact on the solar brightness. However, very high levels of sunspot activity has been correlated with a *decrease* in the relative facular coverage on the Sun (Foukal, 1993). This may indicate that for highly active stars with large starspot coverage fractions, as observed on active M dwarfs for example, bright facular regions may be smaller and have a weaker impact on light curve modulations.

Detecting faculae on other stars has been fraught with difficultly. Gondoin (2008) searched for signatures of faculae on two active stars with light curves from the MOST space telescope (Walker et al., 2003). However, no robust signal from faculae was detected using light curve modeling, and the observed modulations appeared degenerate with cool starspots, possibly indicating faculae had a weaker impact on the light curves for these active stars. An entirely different approach to inferring the presence of faculae from high precision light curves was attempted by Karoff et al. (2013). This method was based on the Solar observations of faculae networks having longer decay timescales than their corresponding sunspots, resulting in lingering bright regions after the sunspot had decayed away. Using the power spectrum of *Kepler* light curves, Karoff et al. (2013) searched for characteristic

timescales that would only be the result of this slower decay. They find encouraging evidence for both faculae and surface granulation, in addition to the expected pulsation and starspot modulations in the light curves of three solar-like stars.

1.3 Sources of Data

In this section I provide information on the various time domain surveys I have used extensively throughout this thesis. In Chapter 2 I utilize near infrared and optical ground-based surveys to characterize flare rates and properties for M dwarf stars. In Chapters 3, 4, and 5 I study the flare and starspot properties from the space-based planet hunting mission *Kepler*. Here I will give a more general background for these three publicly available databases, while in each chapter I will only give a brief review of the data sources as needed.

1.3.1 2MASS

The Two Micron All Sky Survey, hereafter referred to as 2MASS (Skrutskie et al., 2006) was an all sky photometric survey conducted between 1997 and 2001. This remarkable survey collected data using two telescopes (one in the north, and one in the south), each simultaneously imaging in three near-infrared bandpasses: J (1.25 μm), H (1.65 μm), and K_s (2.16 μm). The all sky data release contained 471 million point sources, and has been one of the most widely used and cited surveys in modern astronomy, with over 4158 citations recorded in the SAO/NASA Astrophysics Data System for Skrutskie et al. (2006) at the time of this thesis. 2MASS has been recently used as the astrometric calibration for newer surveys such as the space-based Wide-field Infrared Survey Explorer (WISE; Wright et al., 2010).

Along with the point source and extended source catalogs, 2MASS contained several supplementary or engineering data products that have become useful databases in their own right. Here I highlight one such dataset that has been instrumental to my thesis work. Each night throughout the four years of survey operations, both the northern and southern 2MASS telescopes would image one of the 35 standard reference fields to provide photometric

calibration. These 35 fields were spaced throughout the sky, each with an area of 8.5' in Right Ascension (RA), and 1° in Declination (Dec). A reference field visit consisted of a standard imaging sequence used throughout the survey, with a burst of 6 images each with an exposure time of 1.3 seconds. As in the full all sky survey, this 6 image burst was combined into a single point source measurement. This standard sequence was repeated 6 times per visit, providing 6 independent point source measurements for the reference field, spaced out over approximately 7 minutes.

Throughout the 4 years of survey operations, these fields received between 562 and 3692 epochs of data. Additionally, 5 extra calibration tiles were placed on or around the Large and Small Magellanic Clouds during the final year of survey operations. Over 191 million point source measurements were made in these 40 calibration fields over the 4 years of the survey, 40% as many as in the full all-sky survey. This is a remarkable achievement, considering the 40 calibration regions totaled only ∼5 sq deg. This dataset is known as the 2MASS Calibration Point Source Working Database, or Cal-PSWDB, and contains some of the most robust and densely sampled (in time) near-IR light curves ever captured. These light curves have produced benchmark observations of many variable sources, from RR Lyr stars (Szabó et al., 2014) to blazars and quasars (Davenport et al., 2015b).

1.3.2 SDSS

The Sloan Digital Sky Survey, hereafter SDSS (York et al., 2000) is an ambitious survey consisting of both photometric and spectroscopic data. The 2.5-m Sloan Telescope at Apache Point Observatory in New Mexico, USA, has a large 3° field of view. Photometric data were collected in five bands, from shortest to longest wavelength $ugriz$, which were imaged nearly simultaneously (in the order $riuzg$). Each filter had an exposure time of 53.9 seconds. The maximum temporal separation between bands in a given field was therefore ∼5 minutes, between the r and g filters. This was achieved using a drift scanning technique following great circle strips, which imaged adjoining stripes in all five filters over a quarter of the sky covering the northern and southern galactic caps. The CCDs were constantly read out at a

speed that matched the projected drift scan motion along the field of view. The photometric data were then carefully calibrated (Tucker et al., 2006; Davenport et al., 2007), achieving 1% photometric errors down to $r \sim 20.5$ mag (Ivezić et al., 2007; Padmanabhan et al., 2008).

While the primary goal of the photometric catalog for SDSS was to gather homogenous and contiguous multi-band imaging for a quarter of the sky (π steradians), one region of sky was surveyed differently. An equatorial stripe, known as "Stripe 82" for the SDSS internal stripe ID number assigned to it, was repeatedly scanned from 1998–2007. This stripe spanned $-1.27^\circ < \mathrm{Dec} < 1.27^\circ$ (the characteristic $\sim$2.5° width of an SDSS stripe), and ranged from about $22^\mathrm{h} - 4^\mathrm{h}$ in RA, or about 290 deg^2 (Ivezić et al., 2007). The driver for this repeat imaging was both as a calibration region for the entire SDSS survey, and to discover supernova and other high-value time domain science. In total between $\sim$40 and $\sim$60 epochs of multi-band imaging are available for this entire stripe, creating one of the deepest and most precise optical time domain catalogs publicly available. This unique dataset has been used to characterize the variability of many classes of objects, including distant quasars (MacLeod et al., 2010), pulsating variable stars such as RR Lyr (Sesar et al., 2007), and eclipsing binary stars (Becker et al., 2011).

While the photometric imaging campaign for SDSS was finished in 2010 (Aihara et al., 2011), the SDSS telescope continues to utilize its large field of view using state-of-the-art optical and infrared multi-object spectrographs, and multiple fiber integral field unit spectroscopy bundles. At the time of this thesis, the Sloan 2.5-m telescope is by some estimates the most highly cited facility in modern astronomy (Madrid & Macchetto, 2009). Further, the influence of SDSS is generally comparable to facilities such as the Hubble Space Telescope that have much higher (exponentially in some cases) construction and operations costs.

1.3.3 Kepler

The *Kepler* satellite was launched in 2009 into an Earth-trailing orbit (Koch et al., 2010), and was designed to detect transiting Earth-like exoplanets orbiting Sun-like stars (Borucki et al., 2010). When the orbital plane of a planet around its host star is aligned with respect

to an observer (*Kepler* in this case), transits are observed as a periodic decrease in the brightness of the system when the planet passes in front of the host star. Observing a star–planet system like the Sun–Earth, transits would occur once per year for a duration of ∼13 hours, and with a minuscule change in optical brightness of 88 parts per million.

This incredible level of photometric precision can only be achieved outside of the Earth's atmosphere. Indeed, every aspect of the *Kepler* spacecraft was designed to increase the chance of capturing such small changes in brightness. The satellite features a small aperture telescope (1.4-m diameter), with very "fast" optics (low focal ratio) and a wide field of view of 115 deg^2 (see Figure 1.5). The focal plane is covered by a mosaic of 42 CCD chips, and the *Kepler* filter is over 4000Å wide to allow as much signal from the stars as possible. Since the likelihood of a planet's orbital plane being oriented exactly along the line of sight is very low, *Kepler* observed more than 150,000 stars within its field of view. Furthermore, these stars were observed continuously for several years, using a typical cadence of 30-minutes for most targets, and a rapid 1-minute cadence for a few (∼1000) select targets. To unambiguously detect a periodic signal in such long light curves requires three transits by the planet, and so to detect an Earth–Sun analogue required a mission duration of at least 3.5 years. *Kepler* observed this initial field for approximately 4 years before mechanical failures forced an end to the data collection. To date, over 1000 confirmed transiting exoplanets have been discovered by *Kepler*, making this the most effective exoplanet-hunting mission to date.

The continuous photometric monitoring of more than 150,000 stars for 4 years makes *Kepler* ideal for studies of stellar variability (Basri et al., 2010; Walkowicz et al., 2011). Besides transiting exoplanets, *Kepler* has discovered over 2000 eclipsing binary stars (Slawson et al., 2011), including some very rare systems such as eclipsing white dwarf – main sequence pairs (e.g. Kruse & Agol, 2014). Active flaring stars in the *Kepler* dataset have produced the first catalogs of "superflares" on Solar-type stars (e.g. Shibayama et al., 2013), and the most complete studies of flare rates for individual stars ever amassed (Hawley et al., 2014; Davenport et al., 2014). Additionally, brightness fluctuations due to starspots have been discovered for over 30,000 stars (McQuillan et al., 2014), and the fields of stellar rotation

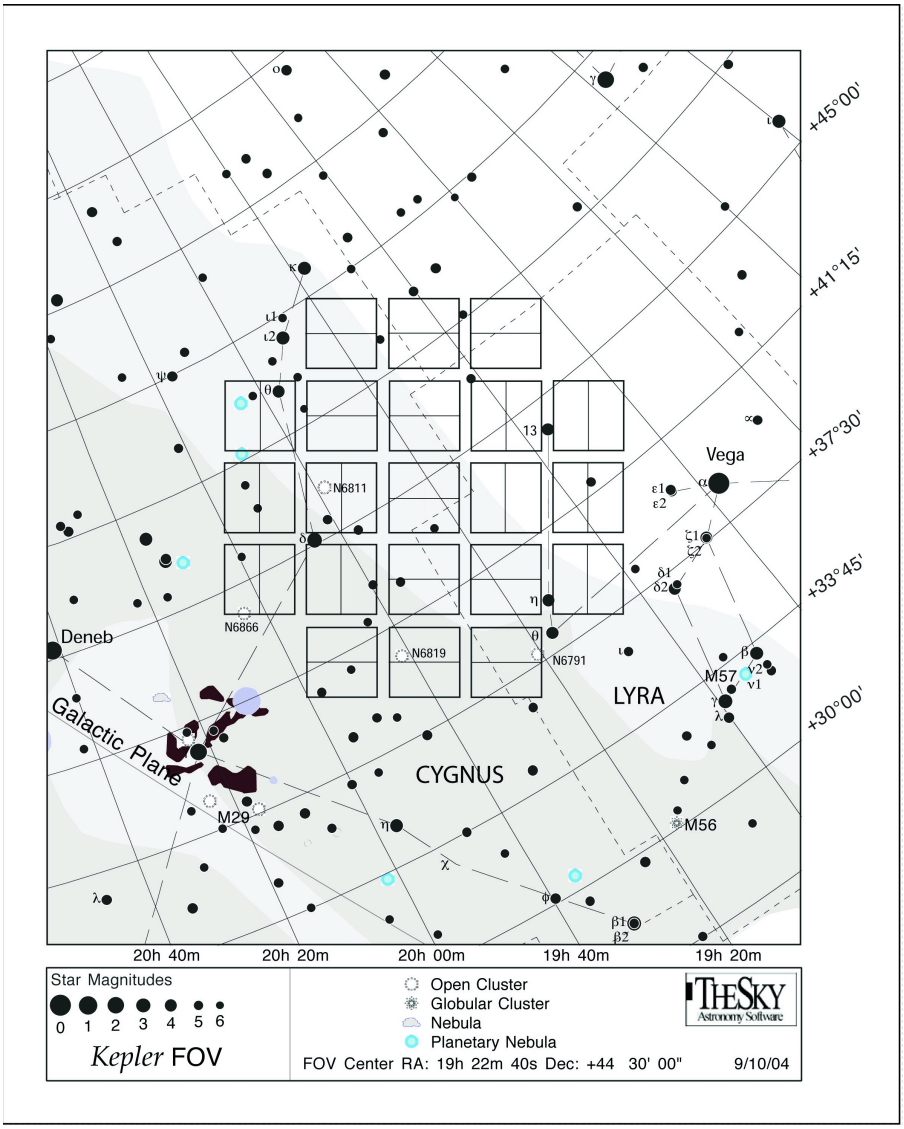

Figure 1.5: Projection on the sky of the original *Kepler* field of view, relative to major constellations and nearby bright stars. Each square shows the projected coverage of the 21 pairs of CCD modules (two CCD chips each). Image from NASA: http://kepler.nasa.gov/Science/about/targetFieldOfView/

and asteroseismology have been opened to statistical studies (Chaplin et al., 2010).

1.4 Outline of this Thesis

The outline for the remainder of my thesis is as follows. In Chapter 2, I will discuss results from mining existing multi-band ground-based surveys for flares on low-mass stars. Flares are discovered as single epoch outliers in sparse light curves. This work studies for the first time both red optical and near-IR photometry for flares on M dwarfs, and puts important constraints on the expected yields for future flare studies in time domain surveys at these wavelengths.

Next, in Chapter 3, the study of flares is narrowed to the single active M4 dwarf GJ 1243, whose flare properties are determined in exquisite detail. Using the high fidelity, 1-minute cadence light curves from *Kepler*, I have gathered the largest sample of flares for any single star besides the Sun. This unique sample has enabled development of a detailed empirical flare light curve template. Determining this "typical" temporal morphology of flares in white light will enable more rigorous flare detection in future surveys.

Then in Chapter 4 the focus of the thesis pivots from flares to starspots, with a detailed study of the starspot properties and evolution for the same active, rapidly rotating M dwarf, GJ 1243. Using the 30-minute cadence *Kepler* data, I have tracked the sizes and longitudes for two separate starspot groups over 4 years. This work helps constrain the surface magnetic field topology and starspot lifetimes, as well as directly estimates the surface differential rotation for this active star.

I then further narrow the focus on starspots in Chapter 5 to the very special case of studying starspots properties in a transiting exoplanet system. In these systems, we use the transiting exoplanet to break degeneracies between starspot sizes, locations, and temperature contrasts. In this chapter, I outline the light curve modeling approach used in my research, and demonstrate its use on the young G dwarf exoplanet host system Kepler 17. Star+planet systems like Kepler 17 are the most valuable for constraining starspot properties using photometry alone, but in turn are the most difficult to analyze.

Finally, in Chapter 6 I provide a short discussion on the current and future context of my thesis research, and important directions this work will take in the next few years.

Chapter 2

MULTI-WAVELENGTH CHARACTERIZATION OF STELLAR FLARES ON LOW-MASS STARS USING SDSS AND 2MASS TIME DOMAIN SURVEYS

In this chapter I present the first rates of flares from M dwarf stars in both red optical and near infrared (NIR) filters. We have studied ∼50,000 M dwarfs from the SDSS Stripe 82 area, and 1,321 M dwarfs from the 2MASS Calibration Scan Point Source Working Database that overlap SDSS imaging fields. We assign photometric spectral types from M0 to M6 using $(r-i)$ and $(i-z)$ colors for every star in our sample. Stripe 82 stars each have 50-100 epochs of data, while 2MASS Calibration stars have ∼1900 epochs. From these data we estimate the observed rates and theoretical detection thresholds for flares in eight photometric bands as a function of spectral type. Optical flare rates are found to be in agreement with previous studies, while the frequency per hour of NIR flare detections is found to be more than two orders of magnitude lower. An excess of small amplitude flux increases in all bands exhibits a power-law distribution, which we interpret as the result of flares below our detection thresholds. In order to investigate the recovery efficiency for flares in each filter, we extend a two-component flare model into the NIR. Quiescent M0–M6 spectral templates were used with the model to predict the photometric response of flares from u to K_s. We determine that red optical filters are sensitive to flares with u-band amplitudes $\gtrsim$2 mag, and NIR filters to flares with $\Delta u \gtrsim$4.5 mag. Our model predicts that M0 stars have the best color-contrast for J-band detections, but M4–M6 stars should show the highest rate of NIR flares with amplitudes of $\Delta J \sim$0.01 mag. Characterizing flare rates and photometric variations at longer wavelengths is important for predicting the signatures of M dwarf variability in next-generation surveys, and we discuss their impact on surveys such as LSST.

Material from this chapter was previously published in collaboration with Andrew C. Becker, Adam F. Kowalski, Suzanne L. Hawley, Sarah J. Schmidt, Eric J. Hilton, Branimir Sesar, and Roc Cutri, in the March 2012 edition of the Astrophysical Journal (Davenport et al., 2012), and has been reproduced here with permission of the American Astronomical Society.

2.1 Introduction

The study of flares on M dwarfs has a rich history, spanning more than seven decades (e.g. van Maanen, 1940; Joy & Humason, 1949). The similarity of these events to resolved flares on the Sun was recognized many years ago (e.g. Lovell, 1969). The seminal paper by Kunkel (1970) laid the groundwork for our present understanding of the broad multi-wavelength properties of M dwarf flares. The basic two-component flare spectral model (H recombination + warm blackbody continuum) presented in Kunkel (1970) is useful for estimating the total energy and stellar surface extent of flares in the optical regime (Hawley & Pettersen, 1991; Kowalski et al., 2010). However, this model does not specifically address the detailed spectral features observed during flares, such as line emission from H, He I, He II, and Ca II H&K, nor the evolution of the spectral energy distribution (SED) as the flare cools and decays (Fuhrmeister et al., 2008). A partnership of multi-wavelength photometric and spectroscopic time-domain observing campaigns (e.g. Byrne, 1989; Hawley et al., 2003), and detailed radiative hydrodynamical modeling (Allred et al., 2006) are required to extend our understanding of stellar flare emission physics.

West et al. (2008) found that the fraction of M dwarfs that show evidence of magnetic activity as defined by Hα emission decreases with height above the Galactic plane. This is thought to be an age effect, with the older, less active stars lying further from the plane due to dynamical heating and radial migration within the disk (West et al., 2008; Loebman et al., 2011). Flares may play an important role in the main sequence stellar evolution for low-mass stars, and correlating M dwarf flare rates with mass and metallicity promises to yield an effective constraint on the ages of stellar populations.

Models and observations of the initial mass function indicate that M dwarfs are the primary stellar component in the Galaxy by number (Bochanski et al., 2010). Recent large scale surveys, such as the Sloan Digital Sky Survey (York et al., 2000, hereafter SDSS), have photometrically identified millions of M dwarfs (Bochanski et al., 2010). These stars make excellent tracers of nearby galactic structure due to their high number densities (Bochanski et al., 2007a), and future survey missions, like the Large Synoptic Survey Telescope (LSST; Ivezić et al., 2008) will detect M dwarfs to nearly the full extent of the Galaxy, using longer wavelength bandpasses and deep repeat imaging. LSST will map the sky approximately 1000 times over 10 years using *ugrizy* bands, and is expected to detect flares both as variability in previously quiescent stars, and as transient emission from previously undetected sources. This foreground "fog" of characteristically blue transient events will prove a significant contaminant to any statistical measurements of variability from cosmological sources (Rau et al., 2008).

Kowalski et al. (2009, hereafter K09) used the SDSS Stripe 82 time domain photometric data (Sesar et al., 2007) to look for flares from ∼50,000 M dwarf light curves in the u and g-bands. From these sparse light curves they were able to recover a flare rate consistent with those found from dedicated photometric monitoring campaigns (e.g. Lacy et al., 1976). The rate of flares was found to decrease with peak flare luminosity, and increase with stellar spectral type (redder colors). Hilton et al. (2010) found serendipitous flares in the time-resolved SDSS spectroscopic data from ∼38,000 M dwarfs across the entire SDSS Data Release 6 footprint (Adelman-McCarthy et al., 2008). They confirmed the increase of the flare duty cycle with later spectral type, and further investigated the Galactic height dependence of flaring activity, finding that the stars that flared were even more closely confined to the Galactic plane than the stars with Hα activity.

The few studies of M dwarf flares in the near infrared (NIR) that have been previously attempted have not found a consistent relationship to the optical emission. Rodonò & Cutispoto (1988), for example, measured a marginal anti-correlation between optical U-band and infrared K-band photometry of two flares. Similarly, Panagi & Andrews (1995)

found no K-band enhancement during 26 M dwarf flares that were observed simultaneously in the U-band. Observations of flares on the Sun in the NIR offer a more promising avenue. Xu et al. (2006) studied the luminosity of solar flares at 1.6μm, the H^- opacity minimum, and found a correlation with hard X-ray and white-light optical emission for two X-class flares. Their data indicated that the NIR band reached peak luminosity after the X-ray, and was responding to secondary heating, possibly from chromospheric back-warming.

Modern M dwarf flare studies are expanding both in wavelength and temporal coverage to gain insight into stellar flare physics, as well as the contaminating impact these events have on detecting other sources of variability. Hawley et al. (2003) found a correlation between flare energies in the optical, near-UV, and X-ray. Schmidt et al. (2012) have observed flaring M dwarfs with NIR spectroscopy to characterize the correlation between NIR emission lines, such as the H Paschen series, and blue optical broadband flux during flares. Dedicated high time-cadence white-light flare studies from the Kepler mission (Koch et al., 2010) will allow an order-of-magnitude improvement in the study of flare frequency distributions (Walkowicz et al., 2011) and their possible evolution. Tofflemire et al. (2012) have undertaken the most comprehensive simultaneous optical and NIR photometric observing campaign for flares to date to determine the effects of flares on exoplanet detection around M dwarfs. These high cadence and high precision data place upper limits on the broadband NIR emission from several medium u-band flares from M3 and M4 stars with $\sim$3.9 milli-mag photometric precision.

In this chapter I seek to extend the analysis of K09 using the time domain data from the 2MASS Calibration Point Source Working Database (hereafter Cal-PSWDB) as well as the SDSS Stripe 82 data. The combination of these data allows us to statistically characterize for the first time the rate of broadband M dwarf flares in red optical and NIR wavelengths. Our data are outlined in §2.2. Variability statistics are discussed in §2.3. Flare rates as a function of both spectral type and observed passband are presented in §2.4. A spectroscopic model of a M dwarf flare is given in §2.5, and the predicted sensitivity to finding flares in each band in §2.6. Finally a concluding discussion is given in §2.7.

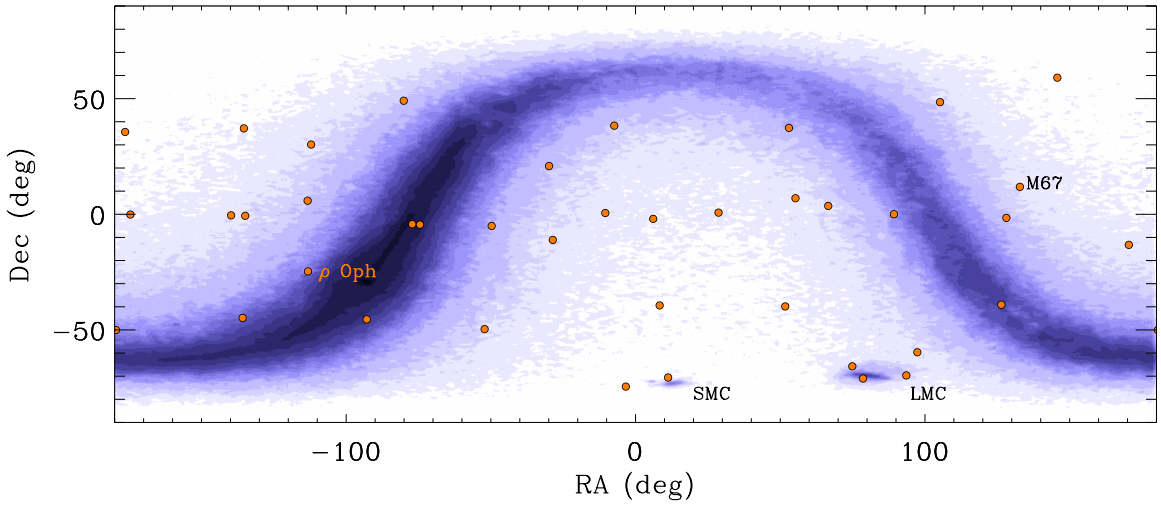

Figure 2.1: The distribution of the Cal-PSWDB tiles across the sky (filled red circles), with a few notable fields labeled. One million randomly selected point sources from the 2MASS database are shown for galactic reference (blue contours).

2.2 Time Domain Databases

Our data are drawn from the SDSS and 2MASS photometric time domain databases. Every M dwarf in our study has single-epoch SDSS Data Release 7 (DR7 Abazajian et al., 2009) photometry, used to determine photometric spectral types and distances. Here we describe the selection criteria and spatial matching of our time domain samples.

2.2.1 2MASS Cal-PSWDB

Our near infrared data come from the 2MASS Cal-PSWDB. Full details are given in the online Explanatory Supplement and Cutri et al. (2006). In summary, the repeat imaging was obtained to provide consistent photometric calibration for the 2MASS survey over the four years of its operation. There were 35 regular calibration tiles spread across the sky, with each tile spanning approximately 8.5' in RA, and 1° in Dec. One of these 35 tiles was visited

hourly during survey operations by the 2MASS telescopes, with simultaneous imaging in the J, H, K_s bands. This sequence produced between 562 and 3692 epochs for these calibration fields. Five additional tiles were imaged in the LMC and SMC with fewer epochs. A total of 191 million point source measurements are contained in the database. A pilot study with the Cal-PSWDB by Becker et al. (2008) revealed a 2.6 day low-mass eclipsing binary system. Plavchan et al. (2008b) used the Cal-PSWDB data to search for periodic variability of low mass stars, finding several eclipsing binary candidates. Additionally, Plavchan et al. (2008a) found a periodic signature in Cal-PSWDB from a possible triple young stellar object in the ρ Ophiuchus star forming region.

We entered the Cal-PSWDB point source measurements into a MySQL database to allow rapid matching of these data to other datasets. The 191 million measurements were clustered into individual objects using the spatial clustering algorithm, `OPTICS` (Ankerst et al., 1999), creating light curves for 113,030 unique objects. Of the 35 primary Cal-PSWDB tiles 16 were coincident with the SDSS Data Release 7 photometric catalog (Abazajian et al., 2009, hereafter DR7), and we found 11,445 matched point source objects using a 2.5" radius. Figure 2.1 shows the positions of the 40 Cal-PSWDB tiles as compared to the SDSS photometric footprint. An optical–NIR $(g - K_s, i - J)$ color-color diagram for the matched point source objects is shown in Figure 2.2, with 1,963 galaxies and 9,369 stars. The SDSS OBJ_TYPE flag was used to distinguish between stars and galaxies. We removed the 113 objects with unknown object type (OBJ_TYPE=0). The distribution of stars in our sample is well fit by the Covey et al. (2007) fiducial stellar locus.

We imposed magnitude limits of $r < 23$ and $z < 23$, and standard SDSS photometric flag cuts to ensure good SDSS DR7 photometry. M dwarfs were selected and spectral types assigned using $(r - i)$ and $(i - z)$ colors and the covariance matrix technique as outlined in K09. We found 4,860 stars with photometric spectral types later than M0, with very few stars in the M7-M9 range due to their faint apparent magnitudes. The resulting 4,731 M0–M6 stars had a total of 5.75×10^6 point source measurements in Cal-PSWDB. Only 12 of our stars were in the SDSS spectroscopic M dwarf catalog (West et al., 2011) so we rely

Table 2.1. Outline of sample selection cuts for the Cal-PSWDB data.

Step	Number	Cut(s)	Description
1.	11,445 (objects)		Cal-PSWDB – DR7 Match
2.	5.75×10^6 (epochs) 4,731 (stars)	0 $\leq$SpType $\leq$ 6, OBJ_TYPE = 6	M dwarfs only
3.	2.51×10^6 (epochs) 1,321 (stars)	PH_QUAL = A,	High S/N Epochs
		$\sigma_{RA,Dec} < 0.25''$	Deblending
4.	1.14×10^6 (JH epochs)	$J \leq 17$, $H \leq 16.5$, $\sigma_J < 0.2$, $\sigma_H < 0.25$	Good Photometry
5.	9.08×10^5 (HK epochs)	$H \leq 16.5$, $K \leq 16$, $\sigma_H < 0.25$, $\sigma_K < 0.3$	Good Photometry

solely on the SDSS photometry for classification.

We restricted our analysis to epochs with the Cal-PSWDB flag PH_QUAL = A, which required every JHK_s measurement to have a signal-to-noise ratio of at least 10. To ensure that each DR7 source matched to only a single Cal-PSWDB object, we conservatively removed all objects with a spatial standard deviation in their Cal-PSWDB data of $\sigma_s = \sqrt{\sigma_{RA}^2 + \sigma_{Dec}^2} \geq 0.25''$. This cut limited the effects of source confusion and de-blending, but also removed any high proper motion objects from our sample. We additionally excluded epochs where photometric errors exceeded $\sigma_J \geq 0.2$, $\sigma_H \geq 0.25$, or $\sigma_K \geq 0.3$. Our final Cal-PSWDB sample contained 1,321 M dwarfs with $\sim2\times10^6$ epochs, for an average of $\sim$1,900 good epochs per star. The selection of our Cal-PSWDB sample is detailed in Table 2.1.

2.2.2 SDSS Stripe 82

In addition to the matched time domain JHK_s data from Cal-PSWDB, we analyzed M dwarfs from the SDSS Stripe 82 repeat scan database (Sesar et al., 2007) to obtain time domain data in the *ugriz* bands (Fukugita et al., 1996). These data were previously used by K09 to search for flares in the u- and g-bands, and we adopt their sample as described in their Table 2. Briefly, these data contain 50,130 photometrically selected M dwarfs, with

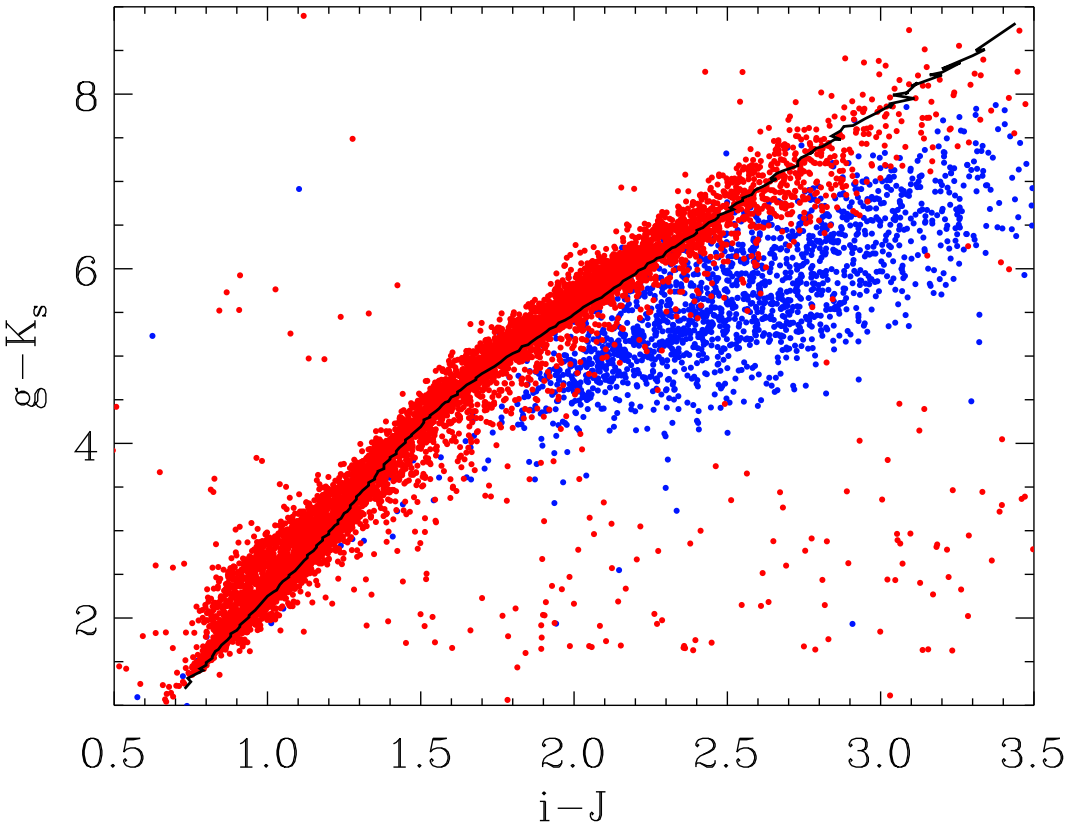

Figure 2.2: A color-color diagram for the 10,000 point source objects sptypech2with matches in both DR7 and Cal-PSWDB. Stars, as determined by the SDSS OBJ_TYPE flag, are shown as red circles, galaxies as blue. These two populations nicely separate. The Covey et al. (2007) fiducial stellar locus is shown for reference, and follows closely the stellar objects in our sample.

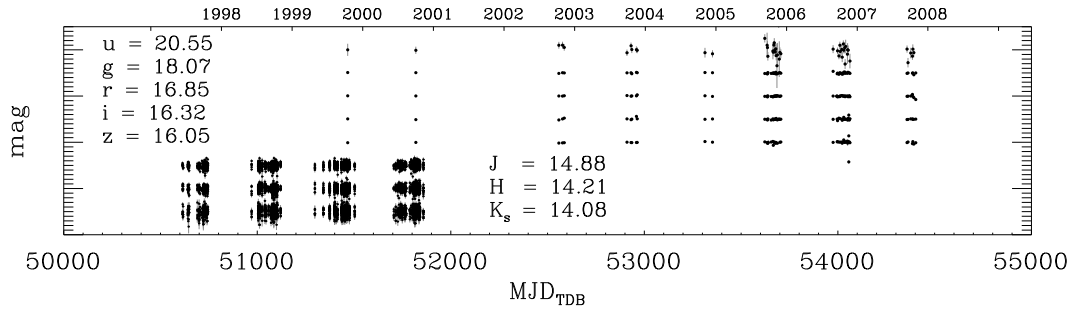

Figure 2.3: An eight-band light curve for SDSS J231820.76+010141.3, which had both Stripe 82 and Cal-PSWDB time-domain photometry. This object had a photometric spectral type of M0 assigned. Time has been converted for both surveys into BJD TDB. The median good magnitude for each filter is given beside each light curve. For reference the calendar year is given on the top-axis.

1.9 million total photometric measurements. The stars were required to have good u-band magnitudes ($u < 22$), which limited the sample to relatively bright M dwarfs. Only 429 of these objects have matches in the Cal-PSWDB database, due to the limited spatial overlap. We measure the flare rate independently in each filter to build statistics using the entire population from each time-domain survey.

The spatial location of the SDSS equatorial Stripe 82 data is shown in Figure 2.1 in blue. Figure 2.3 gives an example of the eight-band light curves we have produced for a representative M0 star with both Stripe 82 and Cal-PSWDB photometry, illustrating the differences in time coverage and cadence between the two surveys. This star had 2,922 epochs of data in each of the JHK_s bands, and 70 epochs in each of $ugriz$. The time units for all of our SDSS and 2MASS data have been converted to barycentric dynamical time ($\mathrm{BJD_{TDB}}$), and are presented as a modified barycentric Julian date (BJD-2400000.5). Figure 2.4 shows the numbers of stars in each spectral type bin for our Stripe 82 and Cal-PSWDB samples.

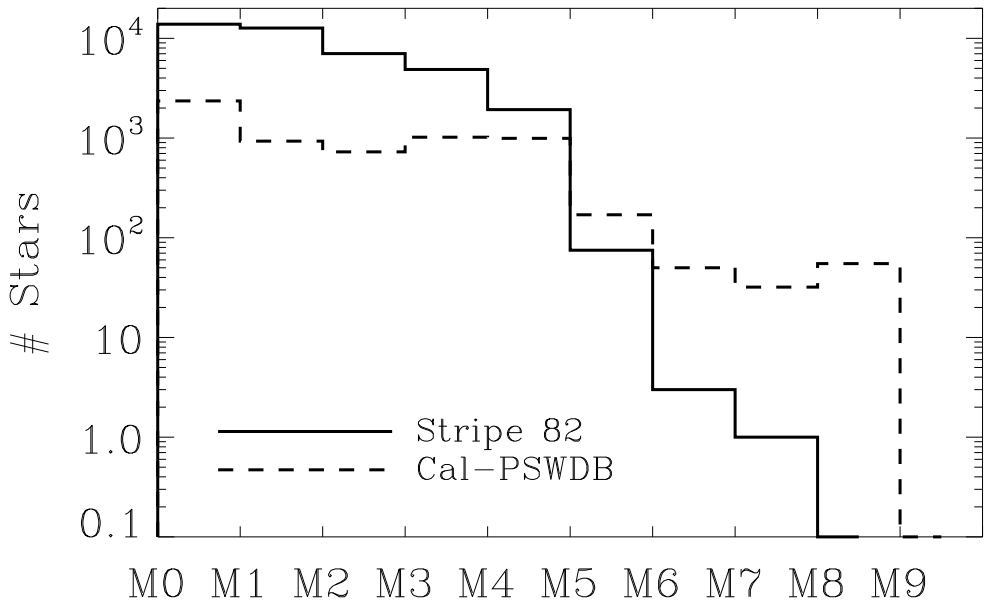

Figure 2.4: The distribution of spectral types for our Stripe 82 and Cal-PSWDB samples. Due to the lack of stars in the M7–M9 bins we have only performed our analysis on M0 through M6 sub types.

2.3 The Φ Statistic

Discerning true variability from background noise requires sophisticated statistical techniques. Following K09, we utilize a modified Welch–Stetson variability index, Φ, to search for flares in our sample (Welch & Stetson, 1993; Stetson, 1996). This statistic uses measurements in two bands to separate real stellar variability from the intrinsic scatter in a light curve. For the n-th measurement of a light curve in J- and H-bands, for example, Φ is defined as:

$$\Phi_{JH}(n) = \left[\frac{m_J(n) - \langle m_J \rangle}{\sigma_J(n)}\right]\left[\frac{m_H(n) - \langle m_H \rangle}{\sigma_H(n)}\right], \tag{2.1}$$

where $m(n)$ is the apparent magnitude at epoch n, $\langle m \rangle$ the median magnitude over all epochs, and $\sigma(n)$ the photometric error at epoch n. When the magnitude increases or decreases for *both* bands simultaneously, Φ yields a positive value. Conversely, when the magnitude from one band is higher than normal while the other is diminished, Φ is negative. For purely random noise between the two filters, this index would yield a symmetrical distribution about $\Phi = 0$. Excesses of positive Φ values represent correlated variability, from events such as flares or eclipses.

We computed the Φ statistic for all pairs of bandpasses "adjacent" in wavelength space (e.g. Φ_{ug}, Φ_{JH}) for every epoch in our samples. Since the SDSS and 2MASS data were not obtained contemporaneously (see Figure 2.3) we do not compute Φ_{zJ}. Figure 2.5 shows the Φ histograms for all objects in our Stripe 82 and Cal-PSWDB samples. The black line indicates the $\Phi > 0$ data, hereafter denoted Φ^+, while the blue line shows the $|\Phi < 0|$ data, hereafter Φ^-.

As the sampling cadence of Stripe 82 and Cal-PSWDB surveys was much longer than typical flare timescales of minutes to hours (Moffett, 1974), our flare search relies on single-epoch outliers in the light curves. In order to act as an effective variability discriminant the two bandpasses making up the Φ statistic must be imaged nearly simultaneously to capture the same physical event. The 2MASS JHK_s bands were imaged simultaneously (Skrutskie et al., 2006), allowing the photometry to be directly compared for each 7.8 second exposure.

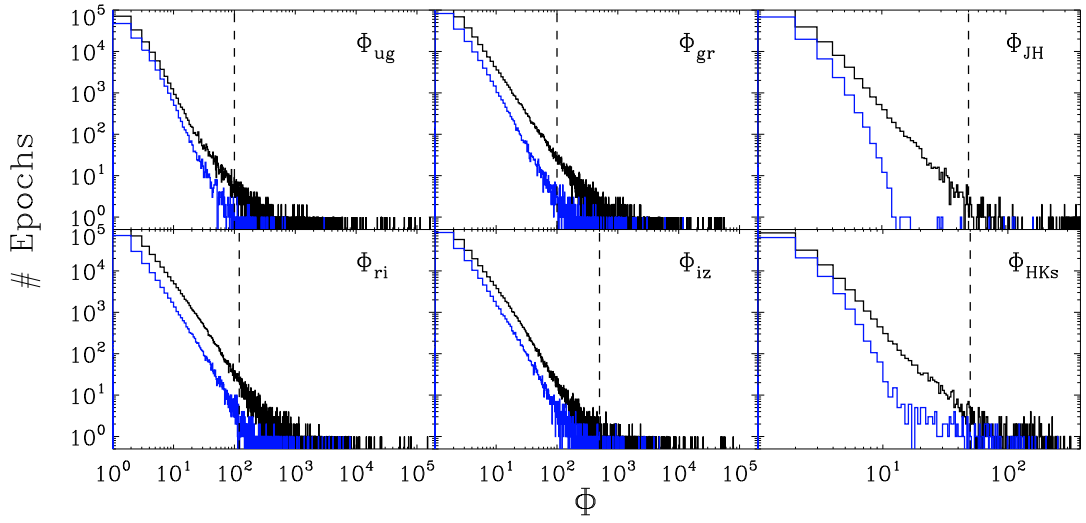

Figure 2.5: Histograms of the two-band variability index, Φ, for each filter combination. The Φ^+, where the flux in both filters is greater than their respective median values, are shown in black. The Φ^- null distribution has been flipped about $\Phi = 0$ for comparison. We subtracted the null distribution from the positive distribution to measure the residual numbers of epochs with correlated increases in flux between bands. The Φ thresholds used in each distribution to define flare candidates are shown as vertical dashed lines. These are described in the text.

The SDSS used a drift-scanning technique that produced separate 52.4 second exposures in each filter taken in the order r, i, u, z, g (York et al., 2000). As a result the Φ_{ri} values, for example, come from measurements which were separated by only $\sim$1 minute, while the Φ_{gr} values are derived from data which were separated by almost 5 minutes. We searched for flares in each of the six Φ distributions separately, but recognize the limitation that flares with timescales of only a few minutes may not have been imaged by all SDSS bands. As such, these rapid events would not yield large Φ values, and may fall below our detection thresholds.

To remove large Φ values from the sample that resulted from a *decrease* in flux in both bands, for example due to photon loss from clouds or eclipses, we required all epochs in our investigation to have $m < \langle m \rangle$ for the bluer (shorter wavelength) filter in each of the Φ^+ and Φ^- distributions. This removed approximately half the epochs in each distribution, as expected. As discussed in K09, the Φ^- distribution, also referred to as the "null" distribution, is representative of uncorrelated noise in the light curve. The Φ^- distribution was still well modeled by gaussian noise in all filters after the removal of half the epochs. Comparison of the Φ^+ distribution with the Φ^- yields excesses in the Φ^+ due to candidate flaring epochs.

2.4 Optical and NIR Flare Rates

In this section we explore the selection of flares from our light curves, and the luminosities of the flare candidate epochs.

2.4.1 Observed Rates

As discussed above, given our sparse sampling we assume that each flare is observed in only one epoch of our data. A threshold value for Φ^+ must be chosen from the null distributions to separate flare epochs from the quiescent epochs with statistical fluctuations. If too high a threshold is chosen, we may lose valuable weak signal in the redder bandpasses by disregarding real flare events. However, too low a threshold will add significant false positives to the resulting flare rates. Using the False Discovery Rate method (Miller et al., 2001) on the

Φ_{ug}^{-} distribution, K09 determined that a conservative cutoff for flares was $\Phi_{ug}^{+} \geq 100$. This ensured less than 10% contamination from non-flaring epochs would be included in their final analysis. We performed the False Discovery Rate analysis to limit the contamination to 10% in all six of our Φ^{-} distributions, and determined Φ thresholds of (100, 100, 120, 500, 50, 50) for $(\Phi_{ug}, \Phi_{gr}, \Phi_{ri}, \Phi_{iz}, \Phi_{JH}, \Phi_{HK_s})$ respectively. The Φ^{-} distributions using 2MASS filters in Figure 2.5 decline rapidly at $\Phi \sim 20$, nearly an order of magnitude lower values than for the optical filters. This is indicative of the lower intrinsic stellar variability seen at these wavelengths, and the larger photometric errors in the 2MASS survey. The Φ thresholds are shown as dashed vertical lines in Figure 2.5.

The number of flaring epochs for each filter combination was very small compared to the total number of observations. Previous studies such as K09 have shown that flare rates for M dwarfs increase as a function of spectral type. To examine this trend we grouped our sample into three spectral type bins (M0-M1, M2-M3, M4-M6). The ratio of flare candidate epochs to the total number of epochs for each bin is shown in Figure 2.6. As expected, all the optical SDSS filter combinations show an increasing numbers of flares with later spectral type (K09). There were considerably fewer flare epochs detected in the NIR passbands, although the Φ_{JH} and Φ_{HK} flare rates fall within the range of the SDSS data in Figure 2.6. This is likely due to the weak signatures of flares in these NIR bands, as discussed in §2.5. Table 2.2 presents the total numbers of flares in each band for the three spectral type bins.

2.4.2 Flare Luminosities

To compare our results in many photometric bands to those of previous studies, we must transform the flares from apparent magnitude enhancements to luminosities. The apparent magnitudes for every epoch and filter in our time-domain Stripe 82 and Cal-PSWDB samples were first converted to observed flux, in units of erg s^{-1} $cm^{-2}\mu m^{-1}$, using the calibrations from Ivezić et al. (2007) and Cohen et al. (2003) respectively. To estimate the quiescent luminosity for each epoch, we obtained distance estimates for our stars from the $(\mu_r, r-z)$ photometric parallax relation of Bochanski et al. (2010), using the $(r-z)$ color from DR7.

Table 2.2. Summary of flare candidates in each spectral type bin per wavelength pair.

Φ	Spectral Type	Candidate Flares	Number Epochs	Flare Fraction
ug	M0-M1	81	1.5×10^{6}	5.4×10^{-5}
ug	M2-M3	116	6.8×10^{5}	1.7×10^{-4}
ug	M4-M6	262	1.1×10^{5}	2.3×10^{-3}
gr	M0-M1	722	1.6×10^{6}	4.6×10^{-4}
gr	M2-M3	322	7.2×10^{5}	4.5×10^{-4}
gr	M4-M6	251	1.2×10^{5}	2.1×10^{-3}
ri	M0-M1	567	1.6×10^{6}	3.6×10^{-4}
ri	M2-M3	283	7.3×10^{5}	3.9×10^{-4}
ri	M4-M6	149	1.2×10^{5}	1.2×10^{-3}
iz	M0-M1	125	1.6×10^{6}	8.0×10^{-5}
iz	M2-M3	100	7.2×10^{5}	1.4×10^{-4}
iz	M4-M6	53	1.2×10^{5}	4.3×10^{-4}
JH	M0-M1	41	6.3×10^{5}	6.5×10^{-5}
JH	M2-M3	10	3.0×10^{5}	3.2×10^{-5}
JH	M4-M6	48	2.0×10^{5}	2.4×10^{-4}
HK	M0-M1	99	5.1×10^{5}	1.9×10^{-4}
HK	M2-M3	13	2.4×10^{5}	5.5×10^{-5}
HK	M4-M6	69	1.7×10^{5}	4.1×10^{-4}

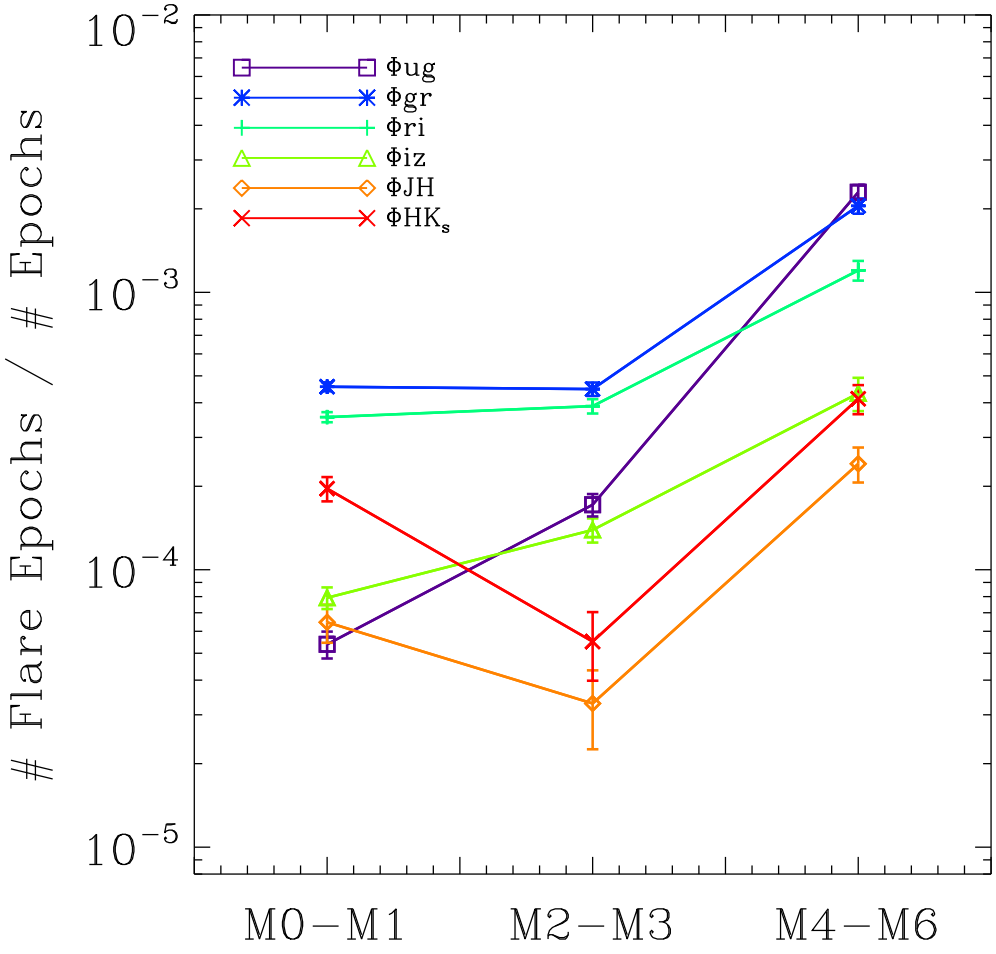

Figure 2.6: The fraction of flaring epochs, defined as the number of good epochs above the Φ threshold divided by the total number of good epochs, for each filter combination and spectral type bin.

Errors in the computed luminosities are dominated by the distance uncertainties.

Many previous studies have shown that the frequency of flares decreases as a function of increasing flare energy (e.g. Lacy et al., 1976) or flare luminosity (K09). For each bandpass we selected the epochs identified as flare candidates in Table 2.1, and found a cumulative distribution in flare luminosity. The frequency per hour that these flares are observed was calculated by dividing the cumulative number distribution by the total observing time in hours for each filter, giving the flare frequency distributions plotted in Figure 2.7. These show for each filter, f, the number of flares per hour of observation, ν_f, as a function of observed luminosity, L_f. The YZ CMi flare frequency distribution from Lacy et al. (1976) is shown, and corresponds well with our optical distributions.

The optical ($ugriz$) distributions are very similar, indicating that flares produce smooth continuous variations in these bandpasses. As with Figure 2.6, the JHK_s-band flare frequency distributions in Figure 2.7 are highly under-sampled. NIR flares were recovered more than two orders of magnitude less frequently than for SDSS optical colors for the M4–M6 bin.

The turnover at low luminosities seen in the SDSS filters is due to incompleteness in the low-luminosity flare census. To explore whether the observed power-law distributions (from log L_f=29 to 31 erg s^{-1}) extend to even lower luminosity flares than our Φ thresholds recover, we moved the Φ^+ detection thresholds to higher and lower values in the u-band. In Figure 2.10 the power-law slope appears to extend to much lower luminosities as a result of the lower Φ thresholds. We recovered many additional flare candidate epochs, but with a much higher rate of false positives. The high-luminosity end of Figure 2.10 (above log L_u $\sim$29.1 erg s^{-1}) is insensitive to the choice of Φ threshold. We present our flare frequency distributions as observed, but note that it is possible to correct the frequencies for the efficiency in recovering flares as a function of luminosity (E. J. Hilton 2011 PhD Thesis).

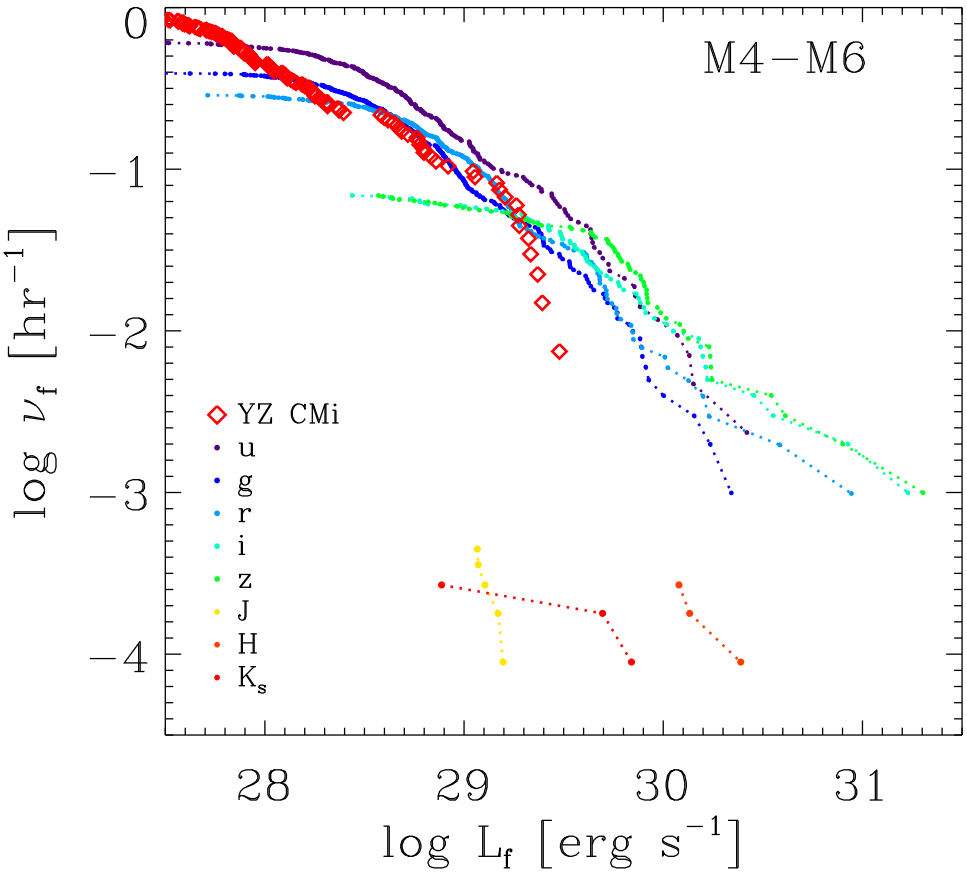

Figure 2.7: The cumulative flare frequency distributions as a function of observed luminosity in each passband for a sample spectral type bin. Errors in luminosity are dominated by photometric errors and assumptions in the distance. Incompleteness in recovering flares at the low-luminosity end results in the flattening seen clearly in each of the SDSS bands.

2.4.3 Residual Φ Distributions

The Φ distributions in Figure 2.5 indicate that there are many more epochs in the Φ^+ distribution than in the Φ^-, showing that the two bandpasses which make up each Φ distribution preferentially increased in flux together. The excess between the Φ^+ and Φ^- is also evident at values of Φ well below the flare candidate thresholds. By subtracting the Φ^+ from the Φ^- distributions for each of the six filter combinations in Figure 2.5 we are left with residual Φ^+ distributions, as shown in Figure 2.8 for Φ_{ri}. These residual distributions follow a roughly power-law profile below the Φ^+ thresholds, with a turn-over at low Φ^+ amplitudes. We fit the residual distributions using the equation

$$\log_{10} Y = \alpha + (\beta \log_{10} X)(10^{-\gamma/(X+\delta)}) \, , \tag{2.2}$$

where Y is the $\Phi^+ - \Phi^-$ residual, α is the peak amplitude for the residual, β a power-law slope at larger luminosities, δ is the turnover Φ between the power-law and the flattening at low Φ, γ the exponential plateau slope at low Φ, and X is the Φ^+ value. Figure 2.8 also illustrates where in the residual distribution each of the fit coefficients constrains. This functional form was chosen to simultaneously fit the residual power-law tail at large Φ, and the turnover at low Φ values due to random noise dominating. The power-law coefficient β is plotted for each of the distributions in Fig 2.9. The β values ranged between -1.9 and -3.1. The higher slope in β in the Φ_{gr} fit may be due to the better photometric precision in the g and r bands, compared to the u band, allowing those filters to be more sensitive to flares below the Φ threshold. The Cal-PSWDB Φ_{HKs} conversely shows the smallest observed residual, again suggesting that these longer wavelength filters are less affected by intrinsic stellar variability. However, all the fits are consistent with a value of $\beta \approx -2$ within the errors.

The physical cause of this power-law slope, β, at Φ values less than our flare cutoff threshold, may be the extension of the increasing flare frequency at lower luminosities. Lacy et al. (1976) showed that the flare frequency distribution has a power-law shape, with a turnover at lower energies due to confusion with the quiescent variability, and a fall-off at

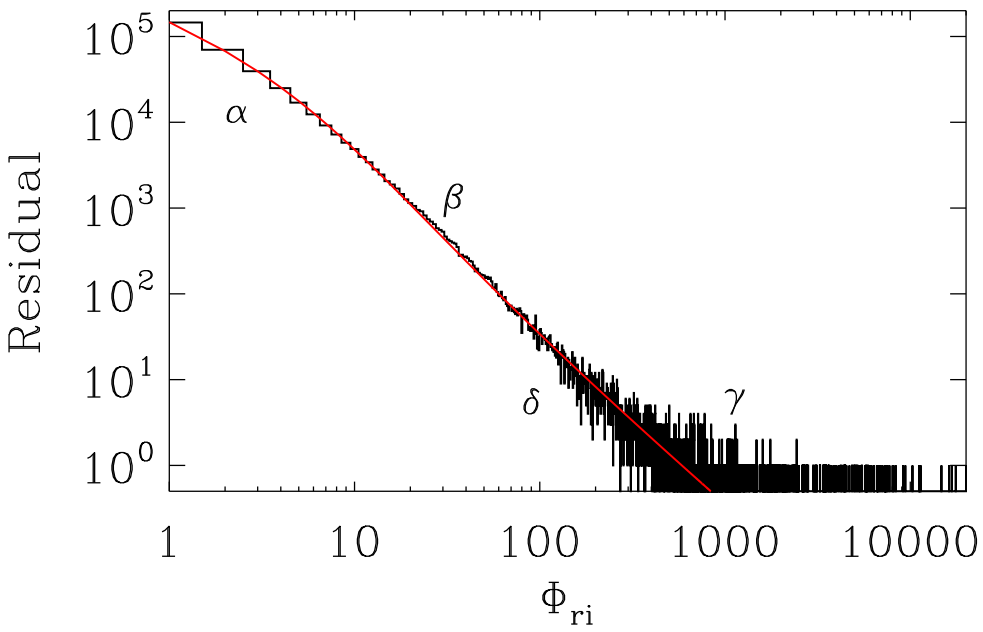

Figure 2.8: The Φ_{ri} residual (black), and the corresponding fit (red) defined in Eqn 2.2.

higher energies due to incompleteness. Their flare detection extended to lower energy flares than our analysis was sensitive to. This trend has been verified for M dwarf flares more recently by K09 and Hilton et al. (2010).

There is not a direct transformation to compare between our observed Φ residual distributions and the flare frequency distributions calculated by Lacy et al. (1976) and others. Our residual Φ distributions contain all the M sub-types combined, for stars across a large range of distances. Hilton et al. (2011) have shown that the flare frequency distribution slopes are similar, though not constant, for different spectral type bins, and between active and inactive stars

Robinson et al. (1995) and Robinson et al. (1999) found that "microflares," very frequent low-energy rapid flares, were detectable using HST for two very active M dwarfs. By comparing the distribution of fluxes about the median quiescent value in their light curves, the microflares, which are of too low amplitude to be detected individually, could be detected statistically. They found that these microflare events followed a power-law distribution which was steeper than that of the larger resolved flares previously detected. Thus there are even greater numbers of microflares than would be predicted by simply extending the power-law flare frequency to lower energies.

Our method of finding the residual Φ distribution yields a similar result, by determining an asymmetry in the light curves via the Φ statistic due to excess flux enhancements. However we caution that we are not detecting the extremely low-amplitude microflares seen in the HST observations from Robinson et al. (1995). Instead, the residual power in the Φ distributions are likely a result of the underlying cacophony of low amplitude flares that are below our single-epoch detection limits in each bandpass, but which would be visible in high-precision continuous monitoring programs. It is especially notable that the significant asymmetric distributions were seen in all wavelengths studied (u–K_s), although we are sampling very different contrasts with the underlying quiescent stellar spectrum, which supports our interpretation of these events as flares. While the flare frequency distributions in Figure 2.7 are under sampled for the NIR filters, the residual epochs with excess flux seen in the

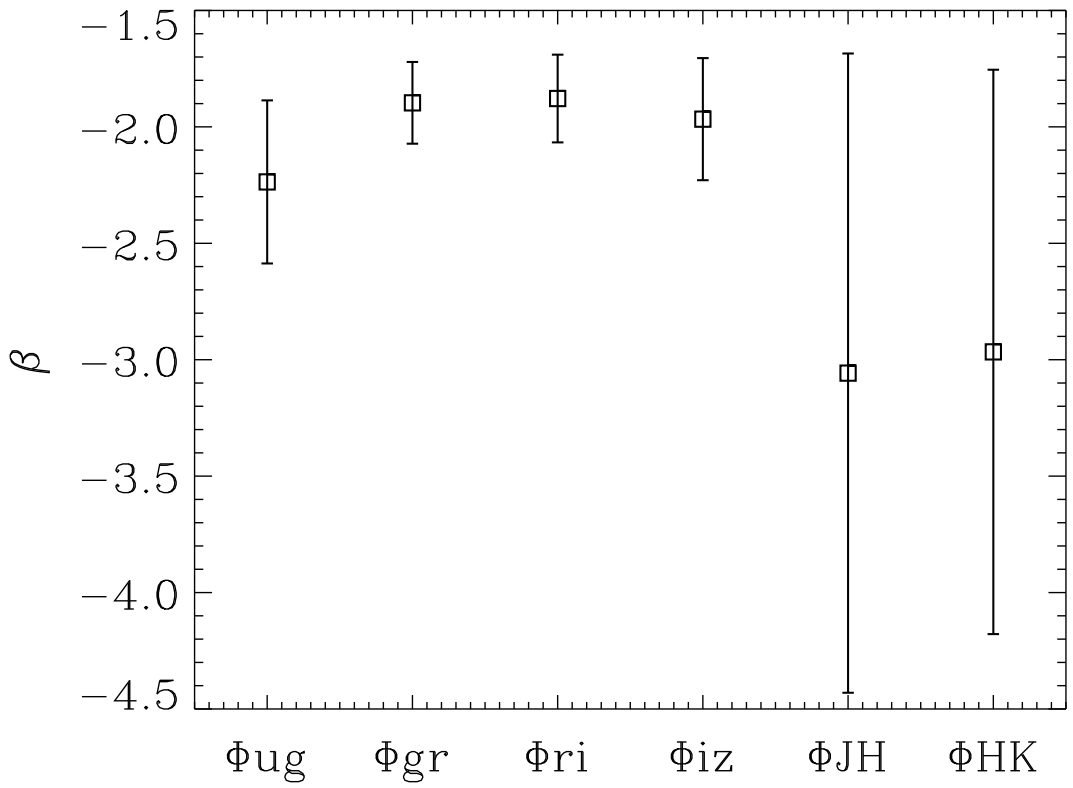

Figure 2.9: The slope of the residual Φ distributions, as defined in Eqn 2.2, for each filter combination. The steeper slope at increasing wavelength is due to the lower intrinsic rate of variability within the photometric errors.

Φ_{JH} and Φ_{HKs} in Figure 2.5, and their power-law slopes in Figure 2.9, provide a statistically significant detection of small-amplitude flux enhancements.

2.5 Optical and NIR Flare Model

We have employed a semi-analytical two-component flare model, as first advocated by Kunkel (1970), to explore the completeness of our search for flares in every filter and to compare the observed flare luminosities between filters. We extend the Kowalski et al. (2010) model

which was used in blue optical wavelengths to red optical and NIR bandpasses.

In order to simulate flux enhancements due to flares in each filter using the two-component model, we built quiescent template M dwarf spectra in the optical and NIR. We utilized flux calibrated spectra from the IRTF stellar catalog (Cushing et al., 2005; Rayner et al., 2009) for stars with spectral type M0 through M6. The M0 through M6 SDSS templates from Bochanski et al. (2007b) were used for the most of the optical regime. The IRTF and Bochanski spectra were normalized and joined at 8500Å. Since the Bochanski optical templates only go as short as 3825Å, to fully cover the M dwarf spectra in the u-band we used templates from Pickles (1998). As these are lower resolution than the Bochanski templates, we used them only from 2900Å to 5000Å. The Pickles and Bochanski templates are both published in normalized flux units, while the flux-calibrated IRTF data we used in the NIR are in observed flux. We found distances to each of the IRTF M dwarf standard stars from Jenkins (1952), van Altena et al. (1995), van Leeuwen (2007), and Jenkins et al. (2009) using the SIMBAD database, and standard radii from N. Reid & S. L. Hawley (2000) to convert the observed flux into surface luminosity for each star in the NIR. The normalized optical templates were then scaled to match at 8500Å, and or resultant optical/NIR spectral energy distributions for the M0 through M6 sub types are shown in Figure 2.11.

The lower panel of Figure 2.12 presents our M3 spectrum in arbitrary flux units in black, with the $ugrizJHK_s$ transmission curves from Fukugita et al. (1996) and Cohen et al. (2003) shown for reference. The overall flare SED shape was defined by a 10,000K blackbody continuum. Previous studies have found temperatures as low as ∼8500K (Hawley & Fisher, 1992) for flares, while Kowalski et al. (2010) found evidence for temperatures as high as 13,000K. Fuhrmeister et al. (2008) fit a blackbody temperature of 20,000 K during the brief impulsive phase of an enormous flare on CN Leo. These temperatures rapidly declined, however, with values around 7,000–10,000 K for the remainder of the flare event.

Shortward of 3646Å, we added a second component of H Balmer continuum line emission, which was computed from the Allred et al. (2006) radiative hydrodynamic M dwarf atmosphere model. The Balmer continuum was scaled to have 10 times larger surface area

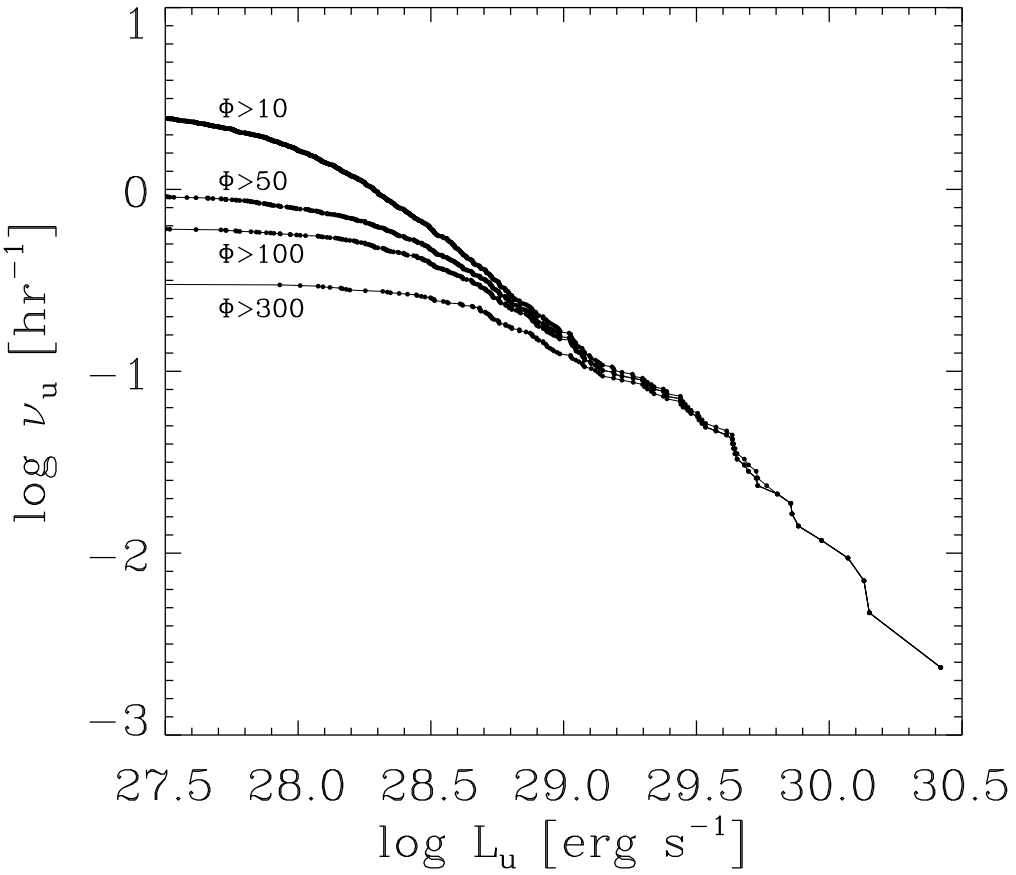

Figure 2.10: The cumulative flare frequency distribution, as in Figure 2.7, for the u-band using four different Φ thresholds to select flare candidate epochs. Lower thresholds allow more contaminating epochs into our flare distribution, but the power-law slope seen at luminosities above $\sim 10^{29}$ erg s^{-1} remains the same.

coverage than the blackbody component. The 10:1 scaling was approximately an average of those used by previous studies, with Kowalski et al. (2010) using values from 3:1 to 16:1. Our idealized model does not include many detailed emission mechanisms that would alter these fluxes, such as line emission from other elements, or other levels of hydrogen. Using measured spectral line emission from M dwarf flares in the NIR from Schmidt et al. (2012), we estimate that line emission would contribute an order of magnitude less than our predicted blackbody continuum emission in JHK_s bands, and is therefore negligible for our analysis. Continuum emission from higher levels of hydrogen in the flare atmosphere, and effects from chromospheric back-warming such as discussed in Xu et al. (2006), likely provide the greatest changes to the red optical and NIR flare SED.

To investigate the effect of the flare model, we scale it by a surface coverage fraction, ranging from $5{\times}10^{-7}$ to 0.5, and add it to the M0–M6 template spectra. The photometric response to flares in each filter was then calculated by convolving the SDSS and 2MASS filter transmission curves (Fukugita et al., 1996; Cohen et al., 2003) to the M0–M6 model spectra with and without the flare. In Figure 2.13 we show the transformations between the Δu-band and the $\Delta(grizJHK_s)$ band amplitudes for all the flare models we ran.

The top panel of Figure 2.12 shows the changes in each photometric bandpass as computed by our model for a flare on an M3 star with 1% surface coverage for the Balmer continuum and 0.1% coverage for the blackbody component. The resulting increases in flux for the $ugrizJHK_s$ bands were approximately (2.7, 0.62, 0.19, 0.06, 0.03, 0.01, 0.01, 0.01) magnitudes respectively.

The extremely blue response of the flare is due to a contrast effect between the intrinsically blue flare model and the underlying red stellar spectrum. Since the M3 spectrum is much redder than the flare model, even large amplitude flares do not give a strong singal in the reddest filters. However, K09 recovered a very bright flare in the Stripe 82 sample, visible in all five passband light curves, on an active M6 star. This flare had magnitude excursions of (5.50, 3.07, 2.34, 0.77, 0.29) for the $ugriz$ bands respectively. Our M6 model, scaled to match the u-band excursion, predicted $ugriz$ magnitude excursions of (5.5, 2.5, 1.2, 0.3,

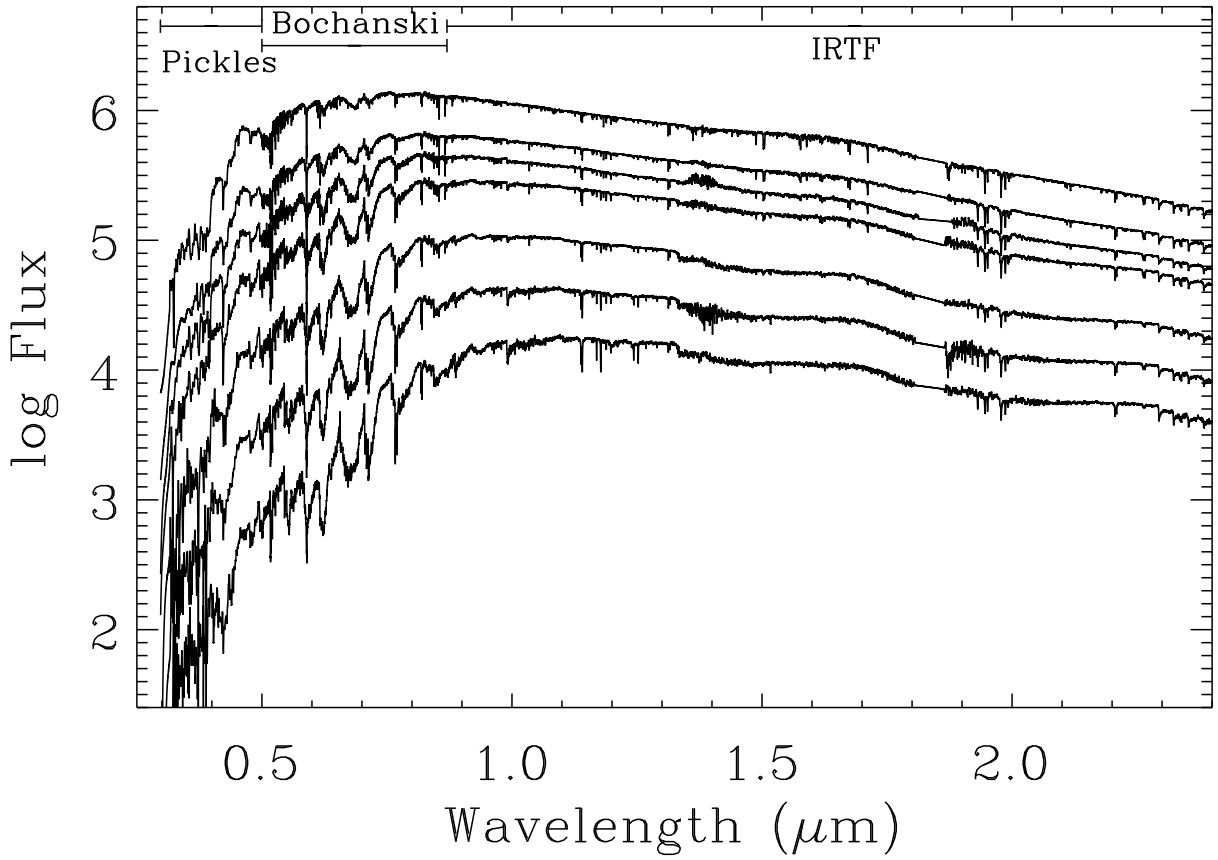

Figure 2.11: The combined 2900Å to 2.5μm spectral energy distributions for our M0-M6 templates. The ranges in wavelength over which each of the template sources contributes are shown at top.

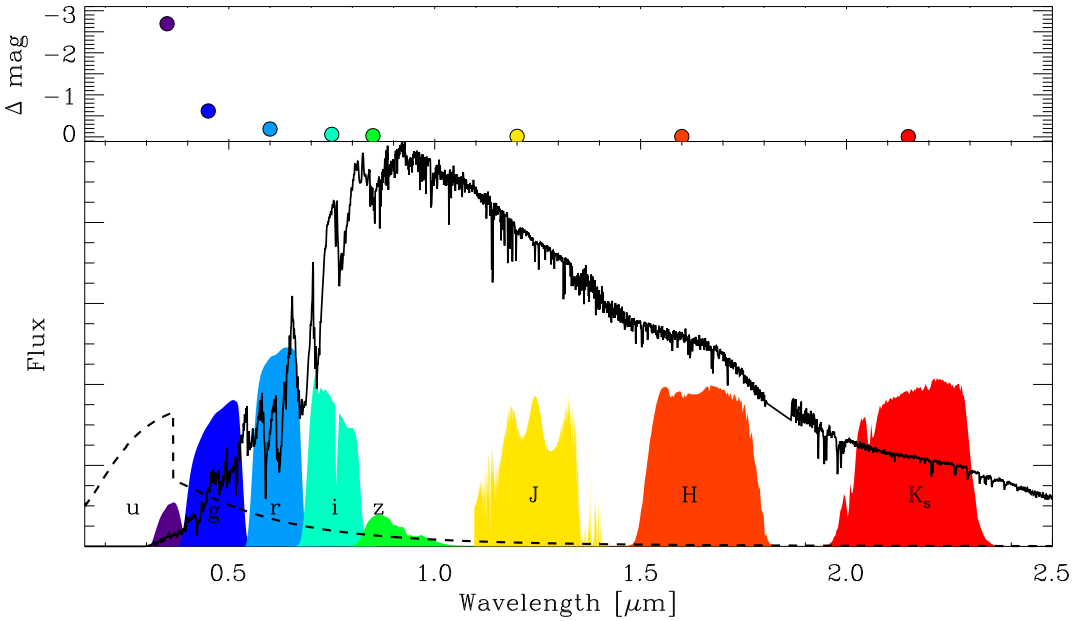

Figure 2.12: Bottom: Our two-component flare model (dashed line) shown against the quiescent M3 optical-NIR stellar SED (black line). The $ugrizJHK_s$ filter passbands are shown for reference. Top: The resulting change in magnitude for each filter from a sample flare with 1% surface coverage on an M3.

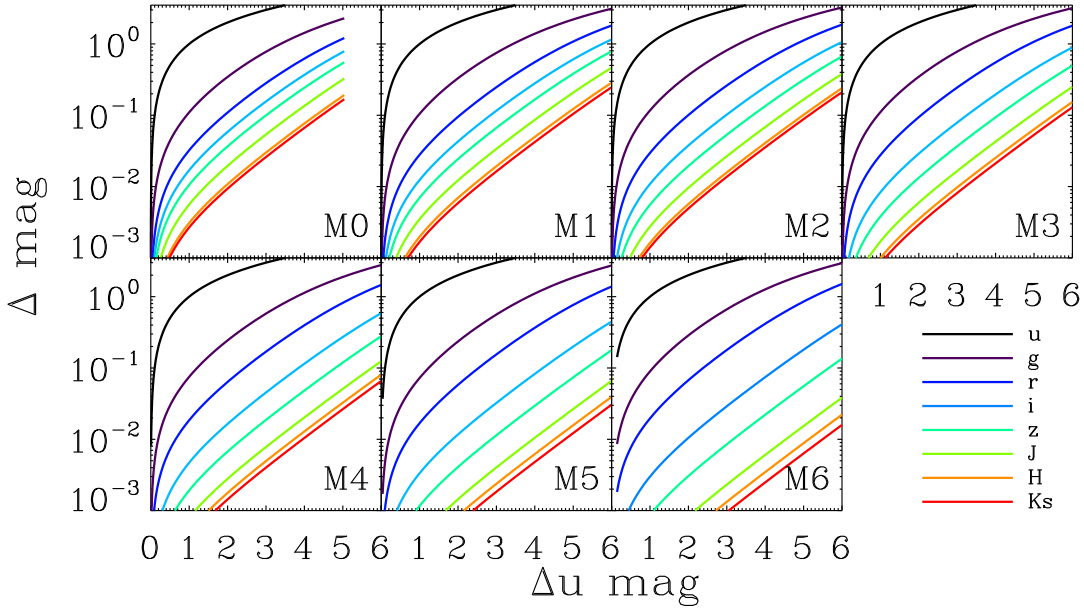

Figure 2.13: The transformation between the predicted u-band response and the $grizJHK_s$-band response for each spectral type template.

0.1). In the 2MASS JHK_s filters we predicted excursions of approximately (0.03, 0.01, 0.01). The coverage for this simulated flare was 0.06% of the stellar surface. As the two-component model does not include detailed effects, as discussed above, the model excursions agree with the observations to within a factor of two in flux. The K09 flare may also have evolved in temperature over the course of the SDSS observations (Fuhrmeister et al., 2008). We allowed the blackbody temperature within our model to vary from 8,000K to 13,000K. Lower temperatures yielded only a marginally improved fit for the r-band, and required a larger effective surface area for the flare.

2.6 Model Flare Recovery

Using our flare model with the M3 spectral template, we investigated the efficiency of recovering flares from our sample with the Φ statistic. We first computed the sigma-clipped mean

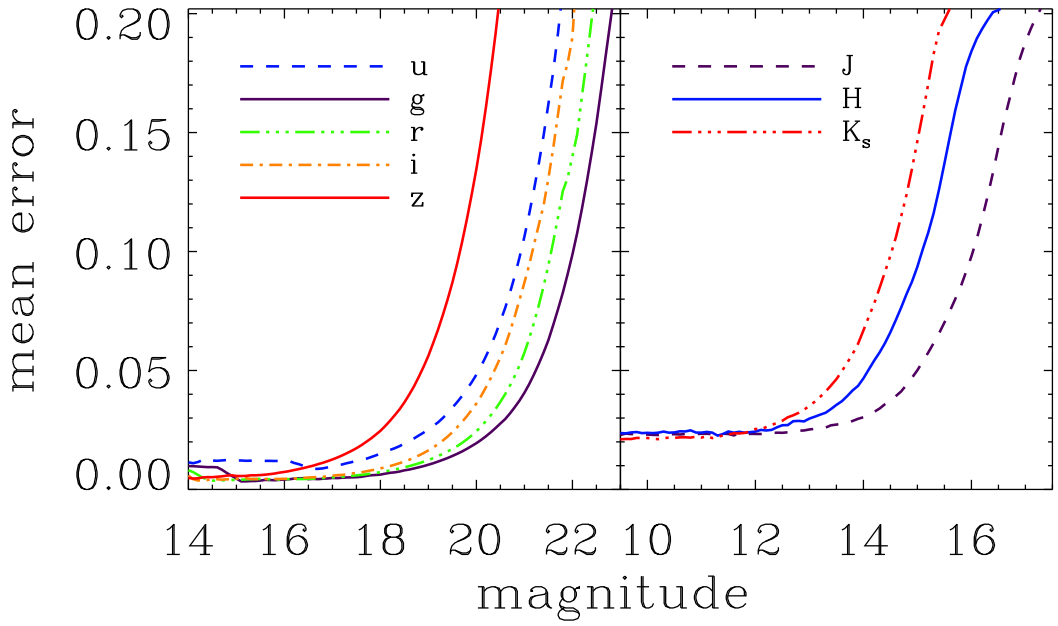

Figure 2.14: Mean photometric error as a function of apparent magnitude for all the Stripe 82 and Cal-PSWDB measurements in our database.

photometric error profiles as a function of observed magnitude for each of the $ugrizJHK_s$ bands in increments of 0.1 mag, shown in Figure 2.14. For each of the M3 stars in our NIR time domain sample, we simulated 200 artificial flares using our model from §2.5, with a logarithmically increasing surface area coverage fraction from 10^{-7} to 10^{-1}. The mean error profiles were then used to select an appropriate photometric error for the resultant apparent magnitudes of each simulated flare in all eight passbands. Finally, we calculated the Φ values for each simulated flare on every star.

The result was a "Φ profile", which was a function of both the quiescent apparent magnitude of the star, and the flare amplitude. These represent the two competing physical quantities, the distance to the star and the size of the flare, which conspire to limit our detection of small flares or flares on distant stars. This envelope is shown in Figure 2.15 as a function of quiescent magnitude on the x-axis, with Φ contour values increasing from

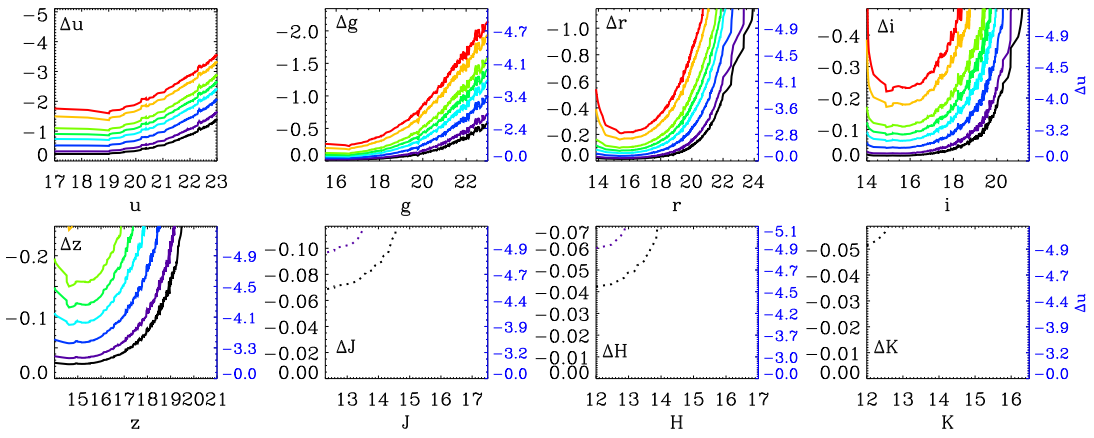

Figure 2.15: Lines of constant Φ as a function of both mean apparent magnitude (distance) and change in magnitude (flare amplitude) in each filter, using the M3 spectral flare model above, and the mean photometric error profiles from Figure 2.14. The predicted change in u-band which corresponds to the change in each other passband is shown on the right axis in blue. Contours from black to red are Φ=(50, 100, 300, 700, 1200, 2000, 5000, 8000). These represent our detection thresholds for a given flare amplitude as a function of apparent magnitude in these surveys. The dotted-line threshold contours for 2MASS filters are Φ=5 and 10.

black to red. The Φ profile was smoothed with a kernel of 0.1 magnitudes. Because flares are typically characterized by their enhancement of blue flux, the corresponding predicted u-band excursion from our model is shown on the right axis in blue for the $grizJHK_s$ bands.

Using our described Φ thresholds from §2.3, we found that small flares in the u-band ($\Delta u \sim 0.5$mag) would be recovered for every star in the sample with a quiescent magnitude of $u \leq 20.5$. The completeness is a direct function of the photometric error profile, which parallels the Φ threshold curves shown in Figure 2.15. Interestingly, note that in the $griz$ bands in Figure 2.15, the Φ profiles "turn up" for bright stars, since the error profiles begin to increase as stars approach the saturation limit for the SDSS. Consider a star with a moderate

quiescent luminosity in the u-band (e.g. $u \sim 17$). A very large flare such those observed in Hawley & Pettersen (1991) and Kowalski et al. (2010) would result in an observed u-band magnitude too bright for the SDSS photometry, and would therefore be excluded from our analysis. This would be possible in the g, and even r-band, for bright stars with very large flares. Incompleteness is therefore also expected at the brightest flare luminosities, although these events are rare with duty cycles estimated to be $\sim$1 per month on most active stars (Hilton et al., 2010).

The error profiles in the g, r, and i filters plotted in Figure 2.14 are best suited to monitor small amplitude variations such as eclipses and spot modulation. Figures 2.15 and 2.13 clearly show that we are only completely sampling flares larger than $\Delta u \sim 2$ in these bands for most of our sample, and thus the u-band is the most suitable to search for the frequent low-amplitude flares in these stars.

As expected from the model discussion above, only very large flares would be detectable in the NIR passbands for the brightest stars in our sample using the 2MASS photometric error profile. Our model predicts that a flare with a nearly ~ 5 magnitude increase in the u-band would be required to detect any significant increase in flux in the Cal-PSWDB for a nearby M3 type star. No simulated flares on M3 stars showed enough increase in flux to register as a $\Phi = 50$ response in the JHK_s bands, however. The flare emission in the NIR must therefore be greatly influenced by physics not included in our two-component model, and definitive simultaneous detection of flares in the NIR and optical is needed to calibrate this model.

2.6.1 Optimal NIR Flare Detection

Our analysis has extrapolated a well-known two-component flare model into a new wavelength regime, and we therefore provide a prediction for the "best-case" to detect M dwarf flares in the NIR. An enhancement of $\Delta J = 0.01$ magnitudes is easily observable for many nearby M dwarfs with present day ground based differential photometry. Using the Δu to ΔJ flare amplitude transformations shown in Figure 2.13, we estimated the Δu enhancement

for each spectral type required to achieve $\Delta J = 0.01$. For spectral types M–M6 this yielded a Δu =(1.4, 1.8, 2.0, 2.5, 3.3, 4.0, 4.5) mag. The changing shape in the underlying stellar SED in the NIR indicates that J-band flares are detectable from smaller Δu flares at earlier spectral types.

The frequency of flaring has been previously found to increase with spectral type, such as in our Figure 2.6. The characteristic lifetime of magnetic activity is also highly correlated with spectral type (West et al., 2008), and later spectral types M dwarfs are expected to show not only more but larger amplitude flares (Hilton 2011 PhD Thesis). This highlights the contrasting effects at work in detecting broadband NIR flares: Persistently active M0 stars that show large amplitude flares are exceedingly rare, as compared to those at later spectral types. The higher bolometric luminosity of earlier spectral type stars additionally means much larger energy flares are needed to achieve these large amplitudes.

Figure 2.16 shows the observed u-band flare frequency distribution for each spectral type bin as a function of Δu magnitude. The total rate of flaring increased with spectral type, and the slopes were very similar for all the distributions, in agreement with results from E. J. Hilton (2011 PhD Thesis). A linear fit was found for each $\log_{10}\nu(\Delta u)$ distribution between $0.4 \leq \Delta u \leq 3.0$ to avoid false-positives at the low-amplitude end, and incompleteness at high Δu amplitudes. The predicted frequency at which we expect to observe a $\Delta J = 0.01$ mag flare, from the corresponding Δu amplitudes, at each spectral type was then calculated from these fits, and is shown in Figure 2.17. The optimal spectral type to search for these NIR flare detections is in our latest type stars, M6. The slight up-turn at M0 is due to the more favorable color-contrast between the flare and the star. The overall trend from M1–M6, however, shows that the strongest effect for all spectral types is the total flaring duty cycle. This provides one possible reason that the recovered flare rates are higher than expected for M0–M1 stars in Figure 2.6. A new method for calculating the duty cycle of a given amplitude excursion, known as the Magnitude Likelihood Distribution, has been recently put forward by E. J. Hilton (2011 PhD Thesis). Using these Magnitude Likelihood Distributions we predict a peak duty cycle of $\sim 3 \times 10^{-4}$ for the magnetically active M4 and M5 stars. For

the 54 sec exposures of SDSS, this is equal to a frequency of $\nu \approx 3 \times 10^{-2}$ hr^{-1}, which is consistent with our estimation. We would encourage future NIR flare observers to monitor bright/nearby targets in the latest spectral types possible, at least M4 or later. Promising stars might include GJ 1061, CN Leo (Wolf 359), and especially the highly active YZ CMi.

2.7 Summary and Conclusions

We have shown that the signatures of flares from M dwarfs are not confined to the bluest bandpasses in large scale photometric surveys. There is significant response from medium to large flares in *gri* bands, and the largest flares are observable in z and in some cases into the NIR passbands. This is the first statistical characterization of M dwarf flares in the NIR regime.

A self-consistent result between all eight filters shows that while the frequency of flares decreases with increasing luminosity in all filters, the total number of flares recovered decreases with wavelength. This suggests that like our simplified flare model the underlying SED for flares of all sizes is intrinsically blue. Furthermore, the total rate of flares in all optical bands appears to increase with increasing spectral type, but low-number statistics prevent us from robustly exploring this distribution at NIR wavelengths. The slope of the flare frequency distribution is found to be insensitive to the flare detection threshold below $\sim 10^{29}$ erg s^{-1}.

We have also presented an implementation of the two-component spectral flare model from Kowalski et al. (2010) to make predictions of the red-optical and NIR response of flares in M dwarfs. These conversions will be useful in future broadband photometric studies of M dwarfs to predict the contamination from flares across a wider wavelength space. Flares in the J-band are predicted to be detected most frequently in spectral types from M4 to M6, with rates of $\sim 2 \times 10^{-2}$ hr^{-1}.

Medium sized flares larger than 0.5 magnitudes in u were found to occur at a rate of one per few thousand SDSS exposures for M dwarfs. Low-amplitude flare rates are limited by the photometric precision of the survey in question, and LSST will detect even smaller

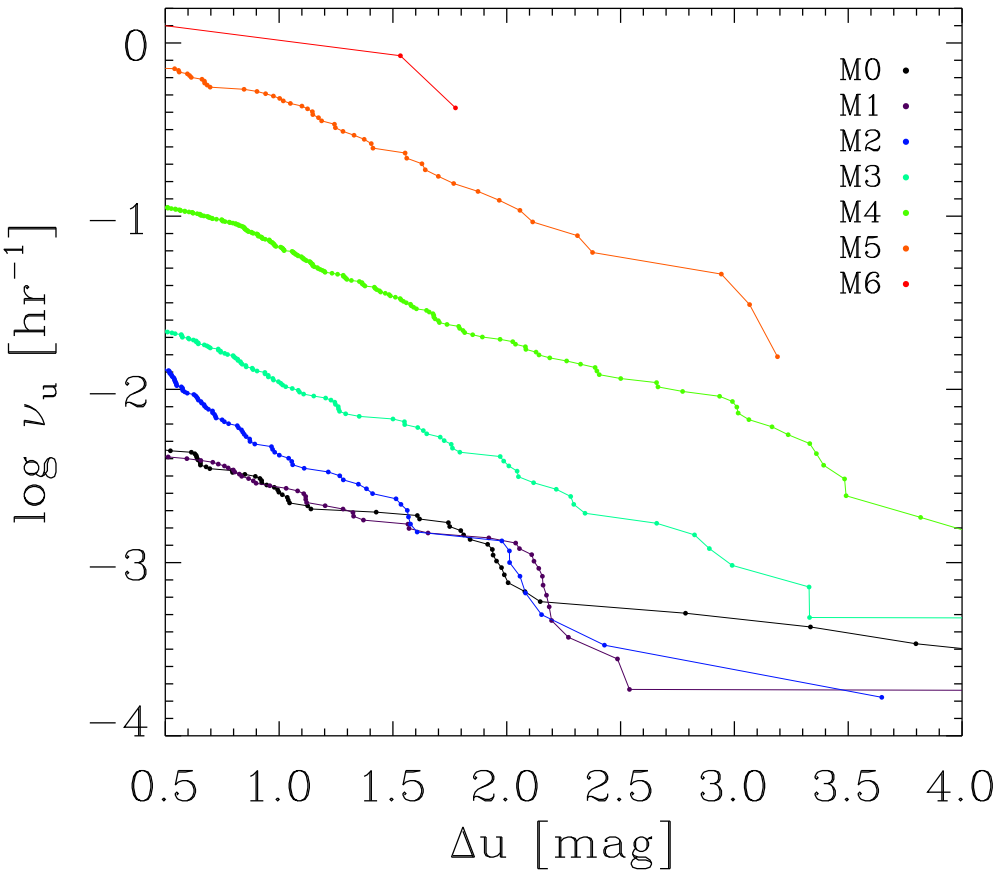

Figure 2.16: The observed flare frequency distribution for u-band flares at each spectral type. Each distribution was fit with a linear polynomial in the range $0.4 \leq \Delta u \leq 3.0$.

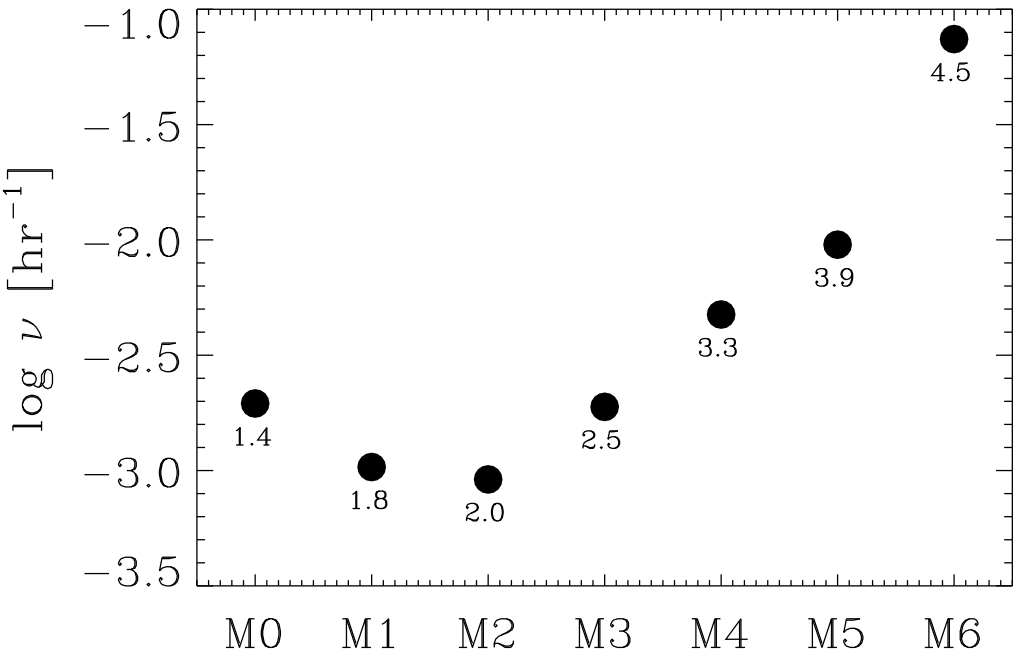

Figure 2.17: The rate per hour at which a $\Delta J = 0.01$ mag amplitude flare is expected to be seen on each spectral type, based on our model and measured flare rates. The numbers below each point are the Δu amplitudes that our flare model predicts from a $\Delta J = 0.01$ mag flare at each spectral type.

amplitude flares from these stars, and will increase the rate at which variability from these sources is detected. LSST will also employ a y-filter, with a central wavelength between that of z and J. Using the "$y4$" filter curve for LSST, our two-component flare model from Figure 2.12 predicts a flare amplitude of $\Delta y4 \sim 0.02$ mag. The increased photometric sensitivity will therefore allow LSST to routinely detect large flares in this NIR bandpass for bright stars, providing useful constraints on the SED structure of flares at long wavelengths

Finally, we have shown that the intrinsic SED of a flare, convolved with the photometric characteristics of a given survey, can greatly affect the sample of flares recovered. The cadence and integration times will also affect the recovery of flares. Short exposures such as in 2MASS capture the near instantaneous luminosity of an event, while longer exposure times may encapsulate entire flare events. We anticipate a day when this problem can be inverted, using orders of magnitude larger samples of flare candidate epochs in many bands, as well as spectra, from next-generation surveys to probe the underlying flare SED (or its deviations) across the population of M dwarfs in the Galaxy.

2.8 Acknowledgements

The authors would like to thank C. MacLeod and J. Ruan for their insight in characterizing astrophysical variability, B. Tofflemire and J. Wisniewski for fruitful discussions regarding NIR flares, and P. Kundurthy for his assistance in converting times to BJD. ACB and JRAD acknowledge support from NASA ADP grant NNX09AC77G. EJH, SLH, and AFK acknowledge support from NSF grant AST 08-07205.

This publication makes use of data products from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center, California Institute of Technology, funded by the National Aeronautics and Space Administration and the National Science Foundation.

This research has made use of the SIMBAD database, operated at CDS, Strasbourg, France.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation,

the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is `http://www.sdss.org/`.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

Chapter 3

THE TEMPORAL MORPHOLOGY OF WHITE-LIGHT FLARES ON GJ 1243 USING *Kepler*

In this chapter I present the largest sample of flares ever compiled for a single M dwarf, the active M4 star GJ 1243. Over 6100 individual flare events, with energies ranging from 10^{29} to 10^{33} erg, are found in 11 months of 1-minute cadence data from Kepler. This sample is unique for its completeness and dynamic range. We have developed automated tools for finding flares in short-cadence Kepler light curves, and performed extensive validation and classification of the sample by eye. From this pristine sample of flares we generate a median flare template. This template shows that two exponential cooling phases are present during the white-light flare decay, providing fundamental constraints for models of flare physics. The template is also used as a basis function to decompose complex multi-peaked flares, allowing us to study the energy distribution of these events. Only a small number of flare events are not well fit by our template. We find that complex, multi-peaked flares occur in over 80% of flares with a duration of 50 minutes or greater. The underlying distribution of flare durations for events 10 minutes and longer appears to follow a broken power law. Our results support the idea that sympathetic flaring may be responsible for some complex flare events.

Material from this chapter was previously published in collaboration with Suzanne L. Hawley, Leslie Hebb, John P. Wisniewski, Adam F. Kowalski, Emily C. Johnson, Michael Malatesta, Jesus Peraza, Marcus Keil, Steven M. Silverberg, Tiffany C. Jansen, Matthew S. Scheffler, Jodi R. Berdis, Daniel M. Larsen, and Eric J. Hilton in the December 2014 edition of the Astrophysical Journal (Davenport et al., 2014), and has been reproduced here with permission of the American Astronomical Society.

3.1 Introduction

M dwarfs have long been known for their magnetic activity, most famously in the form of powerful and frequent UV Ceti-type flares. These explosive events are thought to be analogs to the flares we observe on the Sun, but with much larger energies and higher occurrence rates. On the Sun, flares form as the result of violent magnetic reconnection events, and are ultimately a byproduct of the solar magnetic dynamo that forms from the shearing interface between the radiative core and convective envelope. The greater power and frequency with which M dwarfs flare may be due to their turbulent magnetic dynamos. Whether the creation or cooling mechanisms of M dwarf flares are truly the same as on the Sun is unknown.

A primary limitation of observational flare studies has traditionally been the challenge in gathering statistically complete samples of flares with detailed light curves. Since flares occur stochastically, acquiring a large and detailed sample of M dwarf flares from the ground necessitates studying many active stars of similar spectral types over multiple nights (e.g. Moffett, 1974; Hilton, 2011). These laborious monitoring campaigns may produce light curves for a few hundred flares, with varying degrees of completeness between targets.

Automated surveys that repeatedly image large portions of the sky provide an attractive alternative, as they can efficiently yield millions of individual photometric (or even spectroscopic) measurements of M dwarfs in which to search for flares (Kowalski et al., 2009; Hilton et al., 2010; Davenport et al., 2012; Berger et al., 2013). However, such aggregate studies usually do not provide temporal information for individual flares, nor complete flare rates for individual stars. Instead, these studies rely on the assumption that the characteristics of flares are common amongst large groups of stars.

With the introduction of dedicated space-based monitoring we can overcome many of these observational challenges. The Kepler satellite (Borucki et al., 2010) provides a nearly ideal platform to build statistically complete samples of stellar flares. With round-the-clock white light monitoring for over 150,000 stars spanning nearly 4 years, and remarkable photometric precision, Kepler is the portent to a new era of statistical completeness in stellar

activity studies across the main sequence (Basri et al., 2010; Walkowicz et al., 2011).

The occurrence rates and energy distributions for flares, and so-called "superflares", for many late type stars in Kepler has already been studied to some extent (Notsu et al., 2013). However, the detailed information content from the events themselves has yet to be realized. Understanding the temporal evolution ("light curve morphology") of white light flares is a great utility for planet hunting, where precision flare templates helps improve planet detection efficiency. Stellar flares are also interesting astrophysical phenomena in their own respect. The physics of stellar flare heating and radiation has been studied for many decades (e.g. Gershberg & Shakhovskaya, 1973; Houdebine et al., 1991; Hawley et al., 1995). The source of the white light continuum emission in flares, and its relationship to the high energy and emission line behavior during these events, is an active area of research (Butler et al., 1988; Hawley & Pettersen, 1991; Hawley et al., 2003; Kowalski et al., 2013).

In this chapter, we expand on the sample of M dwarf flares described in Hawley et al. (2014, hereafter Paper 1) and use Kepler 1-minute data to develop a statistical understanding of the morphological characteristics of flare light curves. The detailed generation and cooling (decay) of flares is likely to be strongly dependent on the properties of the host star, such as temperature, age, mass, or surface gravity (Pettersen et al., 1984). To control for these physical differences we concentrate on a single target, the active M4 star GJ 1243 (Kepler ID # 09726699). The dramatic level of stellar activity from this star in the Kepler data, manifested in both long-lived starspots and frequent flares, has previously been noted (Savanov & Dmitrienko, 2011; Ramsay et al., 2013). Our boutique analysis of this single M dwarf allows us to compile a very large sample of flares, unprecedented in its completeness, and provides the foundation for studying flares from many targets in the Kepler database.

The outline of this chapter is as follows. We build our sample with automatic selection and manual validation in §3.2. Using this clean sample of flares, we create a high fidelity flare template in §3.3. The flare template provides us a robust way to decompose complex multi-peaked flare events in §3.3. In §3.5 we perform tests to determine the completeness and limitations of our sample. We discuss the small number of unusual flares that are not fit

by our template in §3.7. Finally, in §3.7 we conclude with a brief discussion of the context and implications of this work.

3.2 Flare Sample

As one of the few highly active mid-M dwarfs known in the Kepler field, GJ 1243 was the focus of several Kepler "Guest Observer" campaigns (GO programs 20016, 20028, 20031, 30002, 30021). Due to both the uniqueness of the Kepler data, and the high level of magnetic activity present, this is a benchmark object for future studies of starspots, rotation, and flares. A rapid rotation period of $P = 0.592$ days was established from light curve modulations by two starspots in the first release of Kepler data (Savanov & Dmitrienko, 2011), and independently discovered from ground-based data by Irwin et al. (2011). We are conducting a parallel study of the long-term evolution of these starspots, as well as recovering the signature of weak differential rotation in this rapidly rotating star (Davenport et al. 2014, in prep). Preliminary flare rate analysis for GJ 1243 was conducted by Ramsay et al. (2013), and a detailed investigation comparing the flare rates between several Kepler active and inactive M dwarfs is presented in Paper 1.

For this study, we have utilized all 11 months of short cadence data for GJ 1243 available from the primary Kepler mission. We used the latest reduction of the Kepler light curves, which includes the PDC-MAP Bayesian de-trending analysis from Smith et al. (2012). The short cadence data are provided online[1] in individual months, as opposed to long cadence data, which are provided in whole quarter increments.

The raw PDC-MAP light curve is shown in the top panel of Figure 3.1. The frequent flares are apparent, visible as positive flux excursions throughout the data. Month-to-month discontinuities in the light curve are present. The smooth trends in the light curve over many quarters, after Time$\approx$ 1100 days, are suggestive of potentially real changes in the total luminosity of the star on $>$ 100 day timescales. However, the trends may also be due to

[1]http://archive.stsci.edu/kepler/

errors in the de-trending and calibration, and we speculate that the real starspot and flare variability may be the primary cause for these errors. We subtracted a linear fit from the light curve of each individual data-month to remove any long term calibration errors. The resulting "flat" light curve is shown in the bottom panel of Figure 3.1.

We next describe the process to select a reliable sample of flares from the light curve. An iterative approach was used, first employing an automated flare-finding method, followed by a manual ("by-eye") validation and classification for every flare. The final selection of flare events represents agreement between many such by-eye validations.

3.2.1 Automatic Selection

The first step in our flare selection from the short cadence light curve was an unsupervised detection of candidate flare events. The starspots produce $\sim$3% flux variations, which are smooth sinusoidal features that evolve slowly (timescales greater than $\sim$100 days). A great many flares exist in the light curve with amplitudes smaller than the starspot, and as such a simple flux threshold selection would only be useful in detecting the largest flares.

Instead, we subtracted the starspot features using a custom smoothing function. The smoothing function used a variable span smoothing method, inspired by the Supersmoother algorithm (Friedman, 1984), and with a three-pass iterative approach. The light curve was first smoothed with a large boxcar kernel of 75 data points. All data falling more than 1σ away from the boxcar smoothed light curve were removed, and the entire light curve was then fit using a cubic spline. This process was repeated twice more, using progressively smaller boxcar kernels. The resulting smoothed model light curve was only minimally affected by the presence of large flares, and effectively traced the starspot. The starspot curve was then subtracted from the original light curve. We emphasize that while this iterative "sigma-clipping" method is not always the favored approach, it provided a rapid and effective means to *identify* flares, whose properties were then measured on the original data.

From the flattened (starspot-free) data, we selected all single epochs with positive flux excursions greater than 2.5 times the standard deviation (σ) of the light curve. Examples of

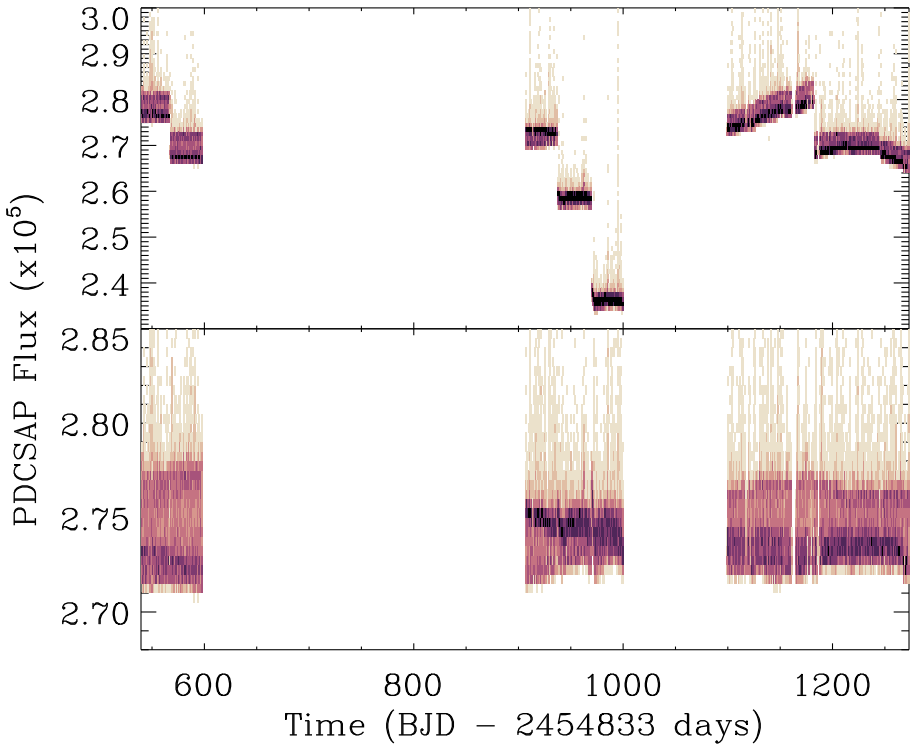

Figure 3.1: Top: Raw PDC-MAP light curve for 11 months of short cadence data for GJ 1243. These data span Quarter 6 to Quarter 13. Bottom: Resulting light curve after our additional linear corrections. Density of points is represented by pixel shade, increasing from light to dark. Note the positive flux excursions due to flares. The median error on the photometry is 78 counts $\sec^{-1}$, or $\sigma_F/\bar{F} = 2.9 \times 10^{-4}$.

these epochs are highlighted in Figure 3.2 as green diamonds. Flare candidates were required to have at least 2 consecutive epochs that passed this threshold cut. To avoid inadvertently splitting flares with complex morphologies (e.g. multi-peaked events) in to separate events, neighboring candidate events were merged in to a single flare event if they were only separated by one epoch below the 2.5σ threshold. We initially found that the recovered flare candidates systematically had a shortened duration compared to those defined by visual inspection of the first month of data examined. This was due to over-smoothing during the starspot removal. A final empirical correction to the end-time of the candidate flares accounted for this truncation, which extended the candidate flare duration by $\sim$57%. The effect of this correction is evident in the flares shown in Figure 3.2, with the event duration extending beyond the epochs highlighted with green diamonds. The completeness of recovering flares with this automatic algorithm is presented in detail in §3.5.

A crude flare type classification was also automatically assigned. By default, all flare candidates were considered "classical" (one peak) unless they passed some conservative criteria. Candidate events with durations shorter than four minutes were labeled as "maybe" flares, indicating uncertainty in their identification. Events with a duration of at least 20 minutes, and having a secondary peak with amplitude greater than 30% the maximum peak (either before or after the maximum peak) were labeled as "complex". This simple type classification was not robust enough for scientific analysis, frequently missing true complex events, but was helpful in the human validation stage.

3.2.2 By-Eye Validation with FBEYE

The sample of automatically detected flare candidates was then validated by manual inspection. To facilitate this validation, we developed a tool called Flares By EYE (hereafter FBEYE), an IDL suite of programs to view light curves and identify and classify flares. FBEYE also contains the auto-finding routines described above. Users interacted with FBEYE using an interactive graphical window. Examples of some of FBEYE's graphical interface functionality are seen in Figures 3.2, 3.3, 3.4, and 3.5.

Briefly, the workflow for processing data with FBEYE is as follows: Users first load the light curve of interest, typically a single month of short cadence data. FBEYE runs the smoothing and auto-finding algorithms described in §3.2.1, and identifies candidate flares. Then, starting at the beginning of the light curve, the user will scan through the entire month of data, inspecting each flare. The user can choose to look at the smoothed, starspot-subtracted light curve, or the original data (shown in Figure 3.2. For each time window, the user will validate all the flares present. Flares that have been incorrectly selected are deleted. Events with incorrect start or end times are modified as needed. Flares that were missed entirely by the auto-finder are then defined. Flares with the wrong type classification, e.g. classical versus complex, have the correct classification assigned. The user then steps forward by half-window increments and repeats the validation processes. Every operation with FBEYE is automatically saved, updating the master flare list for the month.

Typical users would validate a single month of GJ 1243 data in 1–2 hours, while the auto-finder takes only a few seconds to run on a typical workstation. This highlights the great expense in human labor needed to validate these events in a dataset as rich as Kepler, and for a star as active as GJ 1243. Each of the 11 months of short cadence GJ 1243 data was inspected by at least five separate users, with a maximum of eight users, producing over 37,000 individual (though not necessarily unique) flare identifications.

Once the data were manually validated, we selected the final flares from a composite of all the user identifications. Figure 3.6 shows a representative day of data from GJ 1243, with epochs manually selected as containing a flare highlighted, and agreement between users indicated with colors. Many flares show a color gradient, particularly in the decay phase, indicating the regions of strongest disagreement between users. We selected our final flare start and end times to include epochs with at least two users indicating the presence of a flare. As with the auto-finding procedure, flare events were defined as continuous sequences of epochs. Events that were separated by a single non-flaring epoch were merged. The final flare classifications were simplified in to “classical” and “complex”. Events that had been tagged as “maybe”, either by hand or automatically, where called “classical”, while the few

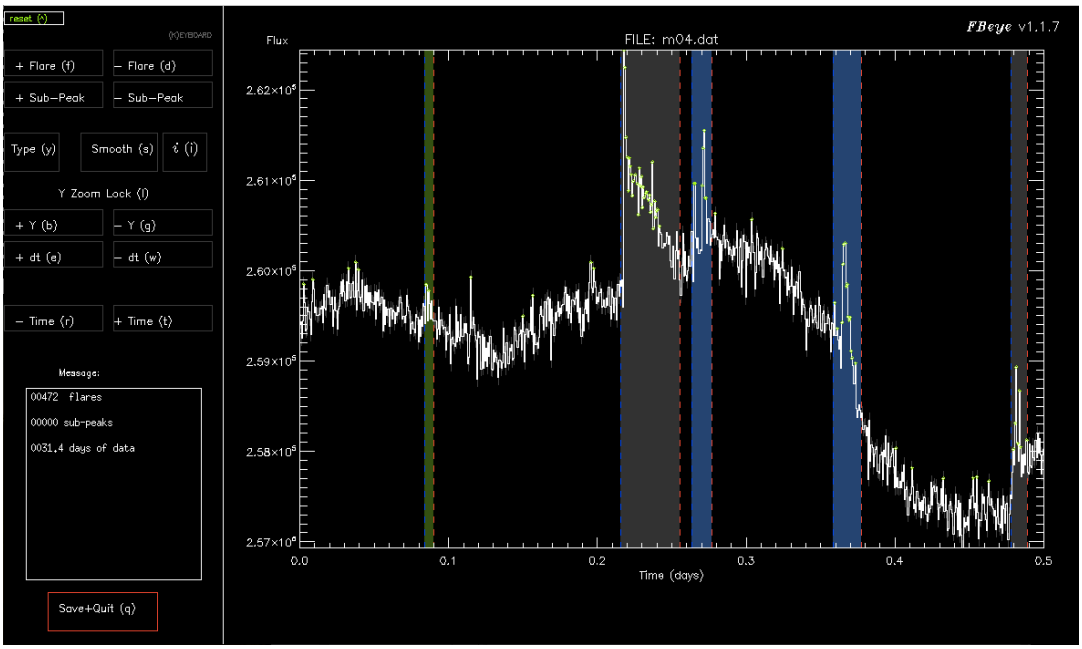

Figure 3.2: Screenshot of our interactive flare-finding suite, FBEYE. User controls to manipulate the light curve display and tag and classify flares are presented in the panel on the left. One half day window of the light curve for GJ 1243 is shown on the right. Displayed are the flares for this time window identified by the auto-finding method detailed in §3.2.1. The flare start times (blue dashed lines) and stop times (red dashed lines) define the candidate flare events. Epochs that the starspot removal algorithm identified as positive 2.5σ flux outliers are marked with green diamonds. The background shading behind each flare indicates one of the four possible flare type classifications: "classical" (grey), "complex" (blue), "unusual" (pink, not shown), and "maybe" (green).

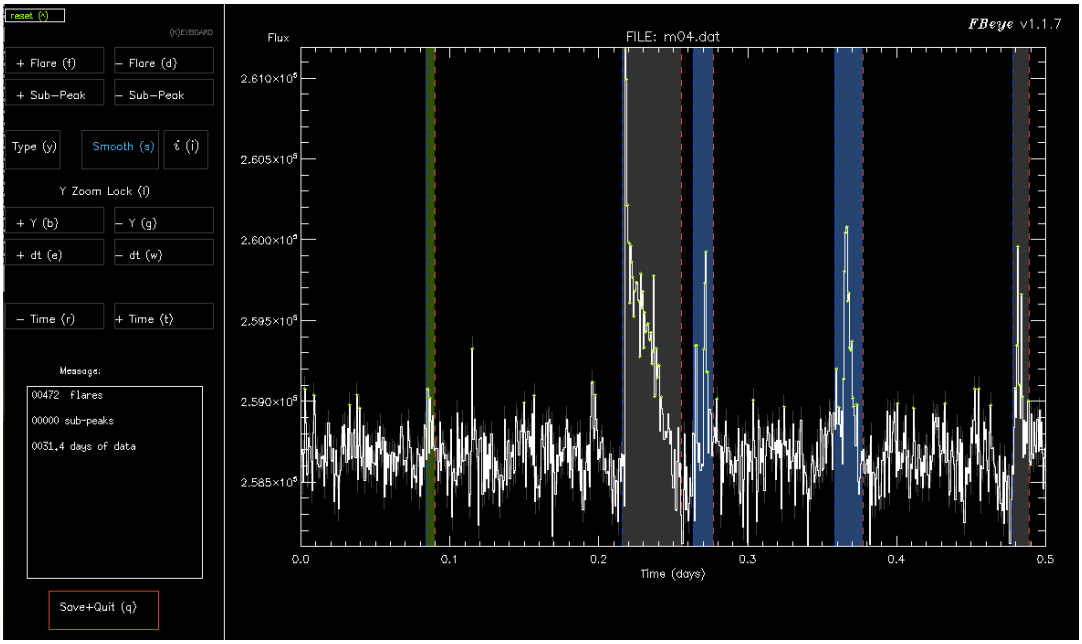

Figure 3.3: Same as Figure 3.2, but demonstrating the smoothing function available to users. This local detrending has been tuned to the GJ 1243 light curve, but has been used for other stars with short cadence data without modification. Smoothing is done automatically when the user loads FBEYE, and is stored in memory for the entire session to allow rapid switching between the raw and smoothed light curve versions.

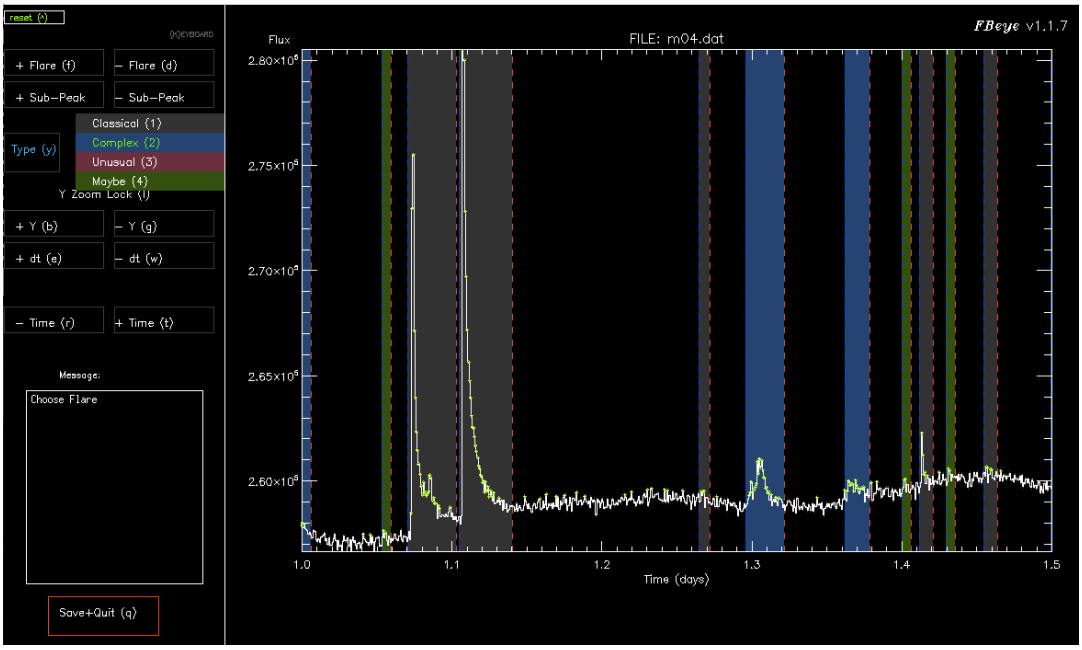

Figure 3.4: Example showing how users can change between flare "types". In this example the flare ending at Time = 1.1 days was mis-classified as a classical type by FBEYE's automatic algorithm. User would typically select "Complex" from the menu (blue) and click within the event to reclassify it. Flare events with durations shorter than 3 epochs are automatically classified as "Maybe" (green).

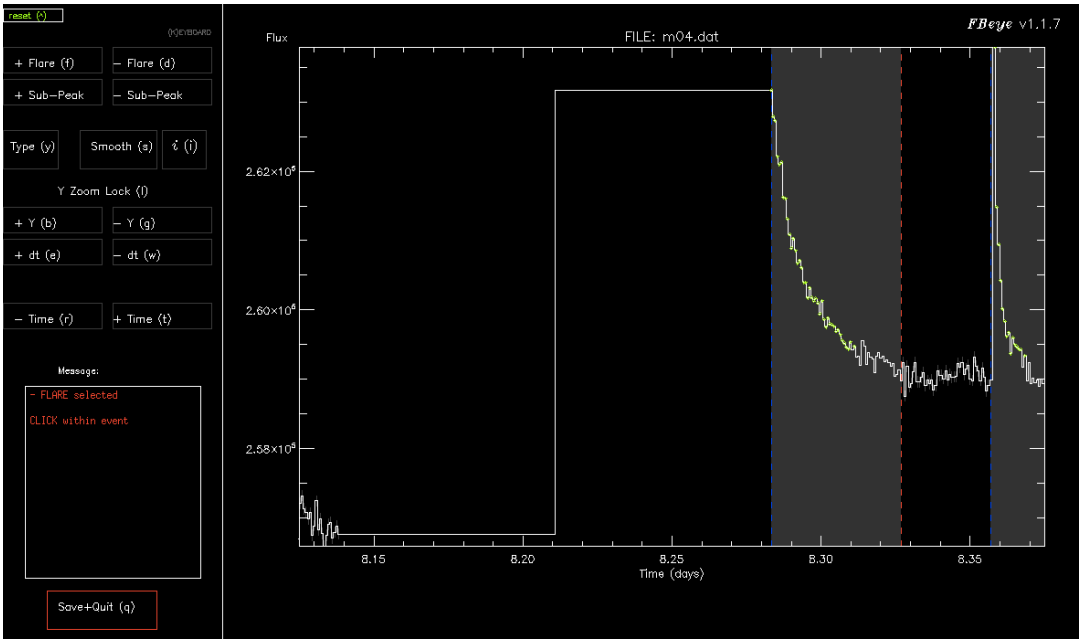

Figure 3.5: Example of a *Kepler* data systematic that was misidentified as a flare. In this case about 0.14 days worth of data was missing, and the flux appeared temporarily brighter after data collection resumes (Time $\sim$ 8.28 days). FBEYE erroneously identified this feature as a flare. The user in this case should select the "- Flare" button, and click within the event centered at Time = 8.30 days to remove it from the output database.

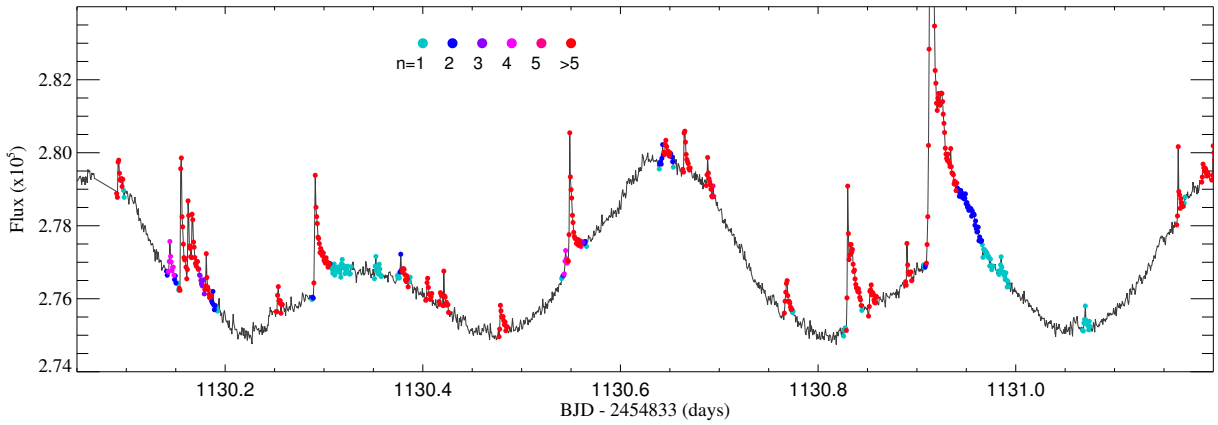

Figure 3.6: A portion of the GJ 1243 light curve from Quarter 12 (February 2012), with epochs that users of FBEYE identified as having flares highlighted. The number of users who identified each epoch is denoted by colors, from blue to red, as indicated in the legend. The end of the gradual decay phase for large flares is the region with the most user disagreement.

"unusual" events were called "complex". The mode of the users' flare classifications was used for each event. The maximum number of users who agreed on each event, as well as the number of users who inspected the corresponding month's data, were also saved in the flare database.

3.2.3 Sample Properties

The final sample of flares from our FBEYE analysis of GJ 1243 contained 6107 unique events. Of these, 5162 were classified as "classical", and 945 as "complex" (15.5%). The general properties of these flares, such as the correlation between flare duration and energy from ∼1000 flares (2 months of data) are presented in Paper 1. The sample of >6000 flares presented here is the largest catalog of stellar flares for a single object (excluding the Sun) that we are aware of, and will be a benchmark for future studies of stellar activity across the main sequence. In Table 3.1, we present the total number of flares and number of users who

Table 3.1. Monthly breakdown statistics for our sample of flares on GJ 1243. Duration is given in units of days.

Month #	Quarter	Duration	# Users	# Flares
1	6a	27.1	8	560
2	6b	30.9	7	680
3	10a	30.1	6	618
4	10b	31.4	7	594
5	10c	30.2	7	534
6	12a	26.5	6	499
7	12b	27.4	7	568
8	12c	27.1	6	455
9	13a	32.3	6	505
10	13b	29.2	5	595
11	13c	27.2	5	499

examined each month. While the month-to-month number of flares varies considerably, no overall trend in the flare rate was observed.

In Paper 1 we found the quiescent luminosity for GJ 1243 to be $\log \mathrm{L} = 30.66$ erg s^{-1} in the Kepler bandpass using ground-based spectrophotometry. By integrating the fractional flux under each flare light curve, we computed the equivalent duration (hereafter ED; Gershberg, 1972; Hunt-Walker et al., 2012), which has units of seconds. Multiplying the ED by the quiescent luminosity for GJ 1243, we found that the largest flares in our sample had energies of $\sim 10^{33}$ erg, while the smallest were $\sim 10^{27}$ erg.

3.3 Empirical Flare Template

Stellar flares are believed to share a common underlying formation mechanism. We wish to investigate if the observed flare morphology can be described by a small number of free parameters and a sufficiently accurate model. Such a template may differ between individual stars, due to a possible dependence on properties such as stellar effective temperature or magnetic field strength. We limit our analysis and discussion, however, to this single star where such variations are assumed to be insignificant. Our sample of flares on GJ 1243 provides a unique dataset for studying the fundamental morphology of flares. Most previous studies of flare statistics have focused on specific measurable properties of individual flares, such as amplitude and duration (see Paper 1). We took a different approach, combining our large sample of individual events to produce a single, high-fidelity template. In this section we describe the creation of this template, whose temporal morphology is described using only two free parameters.

We limited the data to the best observed flares, with an estimated total duration of at least 20 minutes, and with a by-eye classification of "classical". We also omitted flares with durations greater than 75 minutes, as these had a higher likelihood of being complex events. This yielded 885 flares for use in our empirical template. We experimented in adjusting this minimum duration limit, using values ranging from 5 minutes to 50 minutes. The final resulting shape of the empirical flare template was insensitive to the choice of this limit.

For each flare we first subtracted the local quiescent flux level using a linear fit between small time windows before and after the flare start and stop times, respectively. This was done using the non-smoothed initial version of the light curve, detailed in §3.2, and effectively subtracted the local effect of the starspot modulations. The starspot-subtracted light curve used in flare detection was not used for building our empirical flare model, as the flares themselves (especially large amplitude flares) could skew the local smoothing prescription used to remove the starspot, and thus affect the resulting template flare shape. After subtracting the local continuum, we then divided each flare by the maximum flux within the

event, normalizing the flare to have a relative flux range from 0 before and after the event, to 1 at the flare peak. This scaled amplitude is the first free parameter in our template.

Since the decay phase of a flare (defined as the time between the flare's peak and its return to quiescence) dominates the observed timescale, as well as the total flare energy emitted in white light, care must be taken in normalizing the flare timescales to a common range (Kowalski et al., 2013). The choice of a timescale factor is not as straight forward as for the flux amplitude, however. The physics involved in the rise and decay phases are undoubtedly different, and only weakly correlated (Paper 1). This might indeed suggest the need for separate light curve timescales (or "stretches") for the impulsive rise, impulsive decay, and slow decay phases, rather than a single timescale for the entire flare event. This is reminiscent of the 1- versus 2-stretch discussion in normalizing supernovae light curves (Perlmutter et al., 1999; Hayden et al., 2010). However, the 1-minute cadence of Kepler was too coarse to effectively study the shape of the very rapid rise-phase of flares (Moffett, 1974; Houdebine et al., 1991), with most flares having rise times (start to peak duration) of only a few minutes. As such we used a single timescale, and encourage future flare studies with higher cadence to specifically revisit the detailed morphology of the rise phase.

We measured the light curve full time width at half the maximum flux, denoted $t_{1/2}$, which included both the impulsive rise and decay components. This metric was used in Kowalski et al. (2013), and is the second free parameter in our template. Since the rise and impulsive decay phases are so rapid, we linearly interpolated each flare to a 10X higher time resolution of 0.1 minutes to find a more accurate value of $t_{1/2}$. Each flare was set to a relative timescale, centered at the time of peak flux, and then normalized by the characteristic timescale $t_{1/2}$. We note that $t_{1/2}$ is dominated by the impulsive decay phase of the flare, while the total energy emitted throughout the decay phase (in the Kepler bandpass) is split between the impulsive and gradual decay phases. We emphasize that the choice of $t_{1/2}$ as the time normalization factor assumes that the entire flare is governed by a single characteristic timescale, and therefore the impulsive and gradual decay phases are connected and do not vary independently. We discuss this further in §3.7.

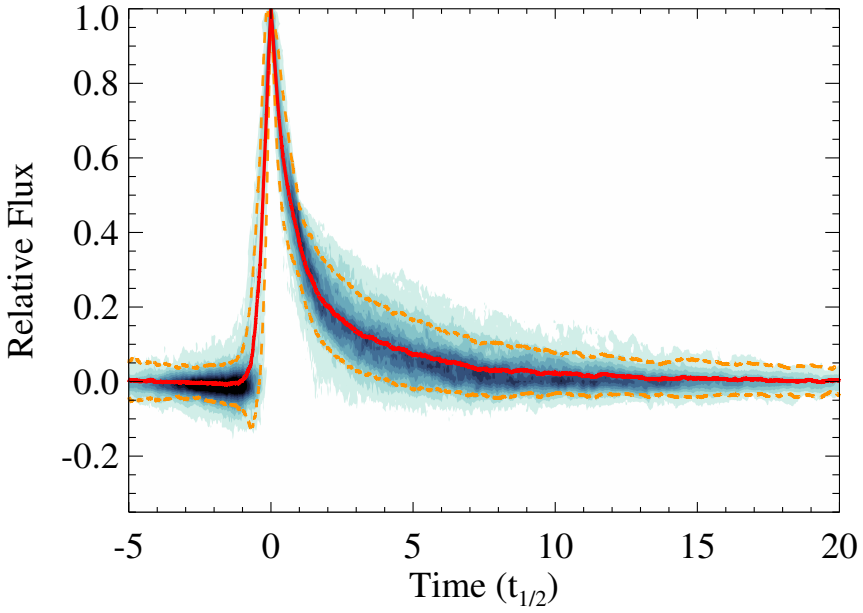

Figure 3.7: Overlay of all 885 classical flares used for the template construction, scaled to relative time and amplitude, and resampled to $\delta t = 0.001 t_{1/2}$ time resolution (blue contours). Contour levels increase from light to dark in units of 50. The median of all 885 flares in each time step (red solid line), as well as the robust standard deviation (orange dotted lines), are overlaid. Outlier points are primarily due to errors in local de-trending of the starspot by other nearby flares.

The entire sample of 885 flares is shown in scaled flux versus scaled time units in Figure 3.7. Since the sample of flares spans several orders of magnitude in duration, each flare has been resampled to a common relative time resolution. A linear interpolation was used to resample the flares to a time resolution of $\delta t = 0.001$, ranging from $t = -5$ to $t = 20$ in $t_{1/2}$ scaled time units. We then computed the median flux value for all 885 flares at each δt bin, shown as the red line in Figure 3.7, which defined our fiducial flare template, described by only two free parameters: the amplitude and the scale time ($t_{1/2}$). We chose to place the flare peak at $t_{1/2} = 0$, so the rise phase occurs in negative time units, and the decay phase in positive time units. We note that a very slight dimming may be visible in the median template before the initial impulsive rise. Pre-flare dimming has been reported by several authors (e.g. Hawley et al., 1995), and has been attributed to a temporary elevation of the Balmer continuum absorption in the upper chromosphere, due to an increase in the local electron density caused by ionization from the electron beam (Abbett & Hawley, 1999; Allred et al., 2006). This pre-flare dimming was observed over much shorter timescales than the small effect in our template, and so we can not rule out other phenomena such as filament eruption. However, the possible slight dimming in our flare template could also be due to artifacts in the starspot subtraction.

3.3.1 Analytic Rise Phase Fit

The median flare template can be described to a very high precision with low-order analytic functions, which in turn makes fitting real flares straight forward. We have divided the flare template into two regimes: the rise and decay phases, or equivalently in Figure 3.7, $t \leq 0$ and $t > 0$ respectively. The resulting continuous functions can also be compared with theoretical predictions for flare heating and cooling curves.

The observed rise time for most flares in our sample was between 1 and 5 minutes (see Paper 1), yielding very little information about the morphology of the rise phase with the Kepler 1-minute cadence for any particular flare. However, with our high-fidelity template, details about the rise phase do emerge as shown in Figure 3.8a. We fit the rise phase with

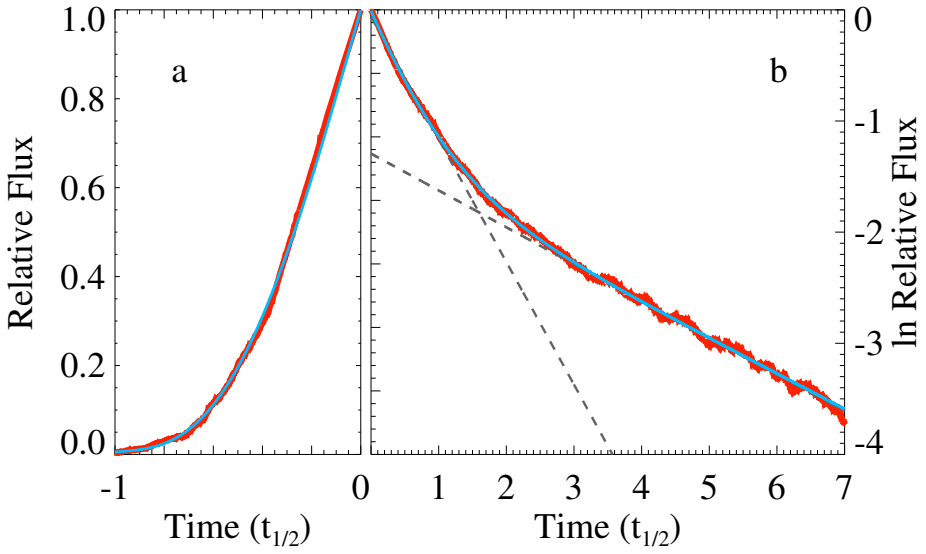

Figure 3.8: **a)** Rise phase of the flare template (red line) in relative flux units, fit with a fourth order polynomial (blue line). The fit was forced to go through relative flux of 1 at time 0. **b)** Decay phase of the flare template (red), in natural logarithm flux units. Straight lines in this space correspond to exponential functions. Two decay regimes are present, and are fit with single exponential functions (grey dashed lines) and a double exponential curve (blue line).

a 4th order polynomial[2] spanning the time region $-1 < t_{1/2} \leq 0$, chosen to reproduce the initial slow rise and nearly linear behavior near the peak. The fit was forced to go through the point (0,1), in other words to equal a relative flux of 1 at maximum light. The best-fit solution took the form:

$$\begin{aligned} F_{\rm rise} = 1 \quad & +1.941(\pm 0.008) t_{1/2} - 0.175(\pm 0.032) t_{1/2}^2 \\ & -2.246(\pm 0.039) t_{1/2}^3 - 1.125(\pm 0.016) t_{1/2}^4 \,, \end{aligned} \tag{3.1}$$

where values in parentheses indicate uncertainties on each fit coefficient. This analytic model is shown in Figure 3.8a, and we emphasize it is valid only for the time region $-1 < t_{1/2} \leq 0$. At times before $t_{1/2} = -1$ the analytic model was forced to equal zero relative flux.

The initially gradual rise, followed by a very rapid climb towards peak flux is similar to the morphology seen in ground-based white light photometry for many of the "Impulsive Flares" observed in Kowalski et al. (2013). However, we were not able to resolve any significant "roll over" or flattening of the median flare shape as it approached maximum light, as has been observed with much higher time cadence observations (e.g. Kowalski et al., 2011). This is due to both the coarse sampling of our light curve compared to the duration of this peak phase, as well as our method of stacking flares based on their observed peak time (the mid-point of the exposure of maximum flux). We believe the latter reason is the primary cause for the "sharpness" of our empirical template at maximum light. Our flare model could be used in future studies to provide a robust peak time estimate for flares, and thus investigate the detailed morphology at the light curve maximum with higher time resolution data.

3.3.2 Analytic Decay Phase Fit

Previous efforts to fit the decay phases (often termed "cooling curves") for flares have used a variety of parameterizations. A generic exponential decay with time is frequently assumed in searches for flares in large catalogs (Walkowicz et al., 2011; Loyd & France, 2014). A

[2]using the IDL non-linear least squares package MPFIT (Markwardt, 2009)

departure from a single exponential (or "teapot") cooling curve towards a more detailed two-phase model for large M dwarf flares was explored as early as Andrews (1965), who parameterized the two phases as a steep linear decline, followed by a gradual inverse square shape. The linear regime was then postulated to be due to Bremsstrahlung radiation of fully (or near fully) ionized hydrogen in a small region, while the gradual phase indicated radiative recombination. Hilton (2011) parameterized the decay phase observationally with an initial linear decline followed by an exponential profile. Recently, Kowalski et al. (2013) determined that the impulsive cooling phase in the white light was dominated by the decay of a 10^4 K blackbody, which cools rapidly to ∼8000 K. A red continuum component, dubbed the "Conundruum", is often present during this impulsive cooling phase, and begins to dominate during the gradual cooling phase. Balmer continuum radiation is also observed throughout the flare cooling, but at a lower level.

The decay portion of our median flare template exhibited two clear exponential regimes, seen in Figure 3.8b: an initially rapid decay, followed by a longer timescale gradual phase. The flare template is displayed using the natural log of the flux in Figure 3.8b, so that straight lines correspond to exponential functions of the form $F(t) = a\, e^{(b\, t_{1/2})}$. A single exponential curve is clearly inappropriate in representing the flare decay. The shape was also not well fit using a single power law.

We fit the decay profile using two separate parameterizations. Our initial model used least-squares minimization to fit the median flare template with two exponential curves. The fit was computed in two well-separated time regions: $0 < t_{1/2} < 0.5$ and $3 < t_{1/2} < 6$. The resulting exponential fits were:

$$F_1 = \quad 0.948\, e^{-0.965\, t_{1/2}} \tag{3.2}$$

$$F_2 = \quad 0.322\, e^{-0.290\, t_{1/2}} \tag{3.3}$$

respectively. The two curves intersected at $t_{1/2} = 1.60$, and relative flux = 0.20 (ln flux $= -1.595$), where the flare is presumed to instantaneously switch between decay profiles. This corresponds to the time of transition between the impulsive and gradual decay phases,

and occurs at nearly the same relative flux value found by Hilton (2011). Such an abrupt threshold between the impulsive and gradual phases would represent a state change in the flare cooling, or rapid change between the dominant emission components.

The energy budget of the template flare shape can then be divided into three regimes: the impulsive rise, impulsive decay, and gradual decay, with the transition between decay phases set using the intersection of the two exponential curves defined above. Integrating the flare template within each regime, we found that the rise phase contains 19.9%, the impulsive decay phase 41.1%, and the gradual decay phase 38.9% of the total energy. By combining the impulsive rise and decay phase values, the fraction of total energy emitted during the impulsive phase observed in the Kepler band is 61%. This is close to the V- and R-band continuum energy fractions found in a very large flare on AD Leonis (Table 6 of Hawley & Pettersen, 1991).

However, the template flare decay profile in Figure 3.8b indicates a smoother transition between decay phases. We thus fit the entire decay profile of the flare template with a continuous function, using the sum of two exponential curves:

$$\begin{aligned} F_{\rm decay} = \;& 0.6890(\pm 0.0008)\, e^{-1.600(\pm 0.003)\, t_{1/2}} + \\ & 0.3030(\pm 0.0009)\, e^{-0.2783(\pm 0.0007)\, t_{1/2}} \,. \end{aligned} \tag{3.4}$$

This parameterization would represent two physically distinct regions, each with its own exponential cooling profile, radiating throughout the entire flare decay. The initial decay would then be dominated by a brighter (presumably hotter) region that cools more quickly, and the gradual decay to a cooler region that cools more slowly, similar to the physical description of the emission components from the spectroscopic analysis of Kowalski et al. (2013).The transition time would correspond to the time when the two regions were equal in luminosity.

Likely neither of these simple parameterizations is entirely correct for describing the evolution of white light emission region(s) in a flare. Thus, we have provided both sets of equations for use in comparing to future flare atmosphere models. For the remainder of this

chapter we adopt the latter parameterization, given in Equation 3.4.

3.4 Complex Flares

In our final, by-eye validated sample of over 6000 flare events, more than 15% were classified by users as "complex". Observationally this meant a significant secondary peak was present in the light curve before the flux returned to the quiescent level. While our automatic flare-finding algorithm attempted to quantitatively define which events were "complex", as described in §3.2.1, the ultimate choice was determined by the average classification selected by the users. This subjective classification was biased towards recovering complex events with large secondary peaks that are well separated from the primary peak. Conversely, when the secondary structure has a small amplitude, or is close to the primary peak, human classification tended to be less accurate. In this section, we seek to produce a quantitative, objective classification for complex flare events using the flare template developed above.

3.4.1 Fitting Complex Flares

The analytic (classical) flare template we have developed can be used as a model or basis function to decompose complex events. In this ansatz, complex events are described as the superposition of several classical flares. Linearly adding a series of our models then reproduces the observed complex morphology. For a given complex event, the task is to determine both the specific properties of each constituent flare, namely (t_{peak}, amplitude, $t_{1/2}$), as well as the *number* of classical flares needed to describe the event.

We fit model flares to every event in the flare sample with a (by-eye determined) duration of 10 minutes or greater. This sub-sample included 3737 individual flare events. For each event, we fit 1 to 10 individual flare models. This was done iteratively, and the final "best-fit" number of flares was determined using the Bayesian Information Criterion. Two example flare events fit by this procedure, which we describe in detail below, are shown in Figure 3.9.

Every flare event was fit in fractional flux units relative to the global median flux (Figure 3.1), and using the de-trended starspot-removed light curve. The flare region to fit for each

event included 3 minutes before the user-defined t_{start} and 15 minutes after t_{stop}, to ensure that the flare fully returned to the quiescent flux level within the fitting window. As was done in constructing the median flare model, we subtracted any residual starspot signal using a linear fit to small windows of time before and after each flare region.

For the first pass ($n = 1$ model) we fit each event with a single flare template. We again utilized the IDL non-linear least squares fitting suite MPFIT (Markwardt, 2009) to minimize the empirical model to the flare event. We seeded this minimization using the peak flux amplitude and time, and 15% of the full duration for $t_{1/2}$. The resulting $n = 1$ best-fit was subtracted from the observed event, and the maximum peak in the residual flux was assumed to be an additional flare. To seed the $n = 2$ model minimization, we used the best-fit $n = 1$ model values, as well as the amplitude and time of the largest positive flux residual, and again 15% of the full event duration for $t_{1/2}$. The $n = 2$ model was then minimized, allowing both the primary and secondary flare to be fully solved to the observed data. We repeated the process of fitting and subtracting for $n = 1$ through $n = 10$ models on each flare event, solving for every flare component in each iteration.

Boundary conditions on the fit parameters were also imposed. Each component flare amplitude was required to be larger than twice the average photometric uncertainty within the flare event. We also required t_{peak} for each component to occur within the boundaries of the flare time window, and $t_{1/2}$ for each component to be larger than 1 minute and smaller than 50% of the observed flare duration. No priors on the relationships between flare amplitude, $t_{1/2}$, and t_{peak} were included in our fitting.

We then choose the "best" model to represent each event using the Bayesian Information Criteria (BIC). This statistic attempts to determine the improvement of the fit (decreased χ^2) while penalizing the increasing number of free parameters used in subsequent models. We computed the BIC for all $n = 1$ through $n = 10$ models, which took the form

$$\mathrm{BIC}_n = \chi^2 + k_n \ln(M)\,,$$

where M was the number of observed data points that fell within the flare event window,

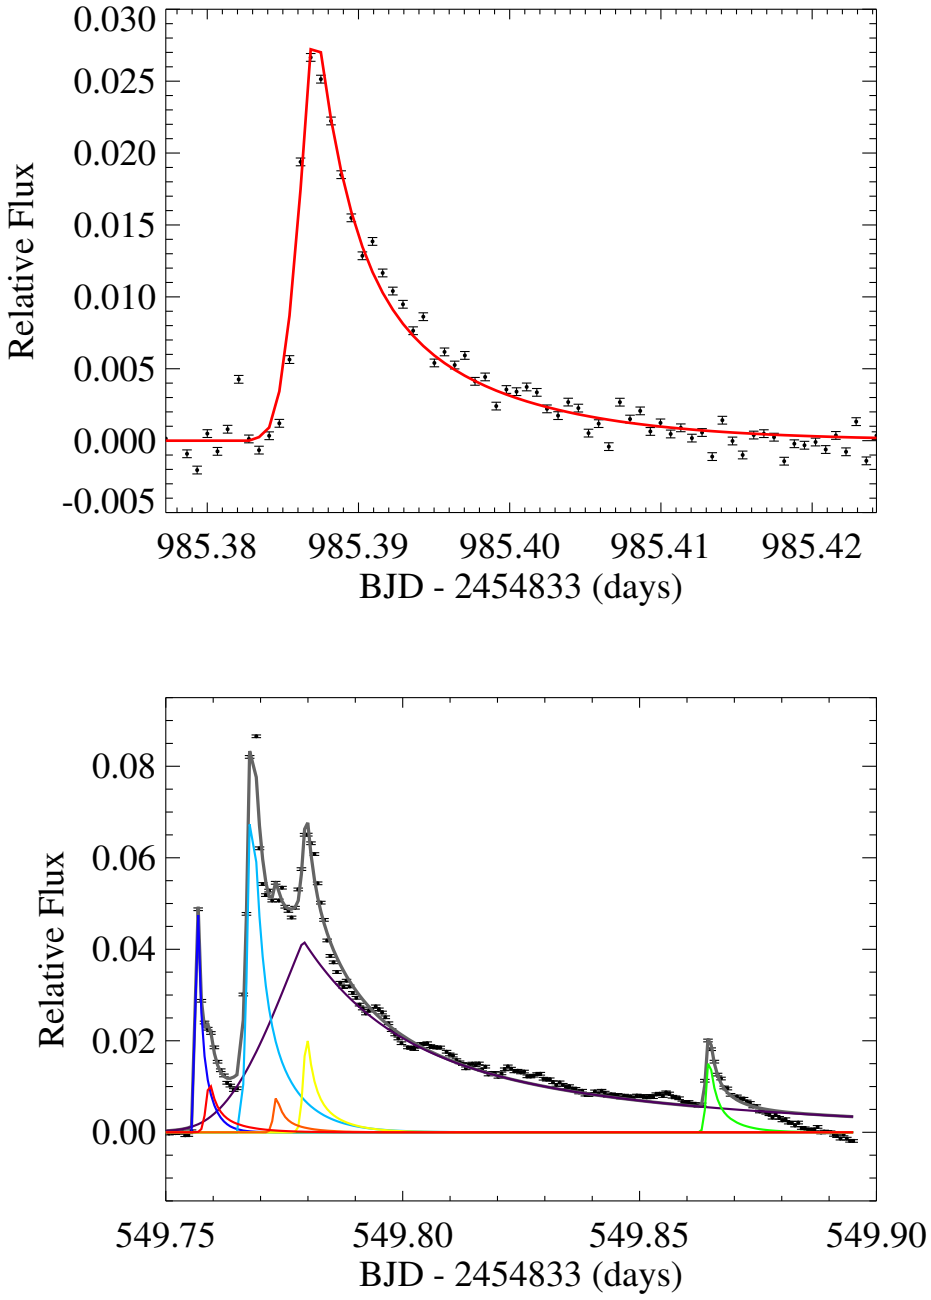

Figure 3.9: Two examples of model fits to flare data. Top: a classical flare event that is well-fit by the template. Bottom: a complex flare event that required seven template flares to produce a good fit to the light curve.

and k_n the number of degrees of freedom in the n'th model. The "best fit" solution was then selected as the n'th model with the smallest BIC parameter, where we additionally required the BIC to have decreased by at least 10% from the previous $(n-1)$ model. Finally, complex flares were defined to be any event best fit with a $n > 1$ model. The choice of a 10% BIC improvement threshold was determined by manual inspection of repeated fits to complex flares, and ensures we are not over-fitting these events.

In Figure 3.9, we show two examples of this fitting procedure. The top panel demonstrates a classical flare, very well fit by the $n = 1$ solution shown in red. The bottom panel contains an example of a very complex event, with four clearly separated peaks in the light curve. Our best model decomposed this event in to 7 distinct components, shown individually as colored lines. While the total morphology of this event is well fit by our procedure, a few discrepancies are apparent. For example, the broadest component flare (purple) does not account for the 4 or so low-amplitude secondary flares in the gradual decay phase. Additionally, the t_{peak} for this component aligns with a much smaller amplitude component (yellow), which may not be physically realistic. This broad component flare over-estimates the flux near the end of the flare event, further indicating that this component should have a shorter timescale, and additional component flares should be fit in the decay phase. This example illustrates both the utility of our template in decomposing these events and classifying them as complex, as well as the intrinsic degeneracies in such a procedure.

3.4.2 Rate of Complex Flares

We have thus far described the light curve structure of complex flares using superpositions of our empirical flare template. This methodology implies the varied substructure seen in complex events is the result of additional flares, which have the same morphological properties as classical events. From this approach two physical interpretations for complex flare events are possible: either 1) the multiple component-flares seen in complex events are physically associated within a single or nearby active regions on the star, or 2) complex events occur due to random superpositions of unassociated flares from separate active regions on the star.

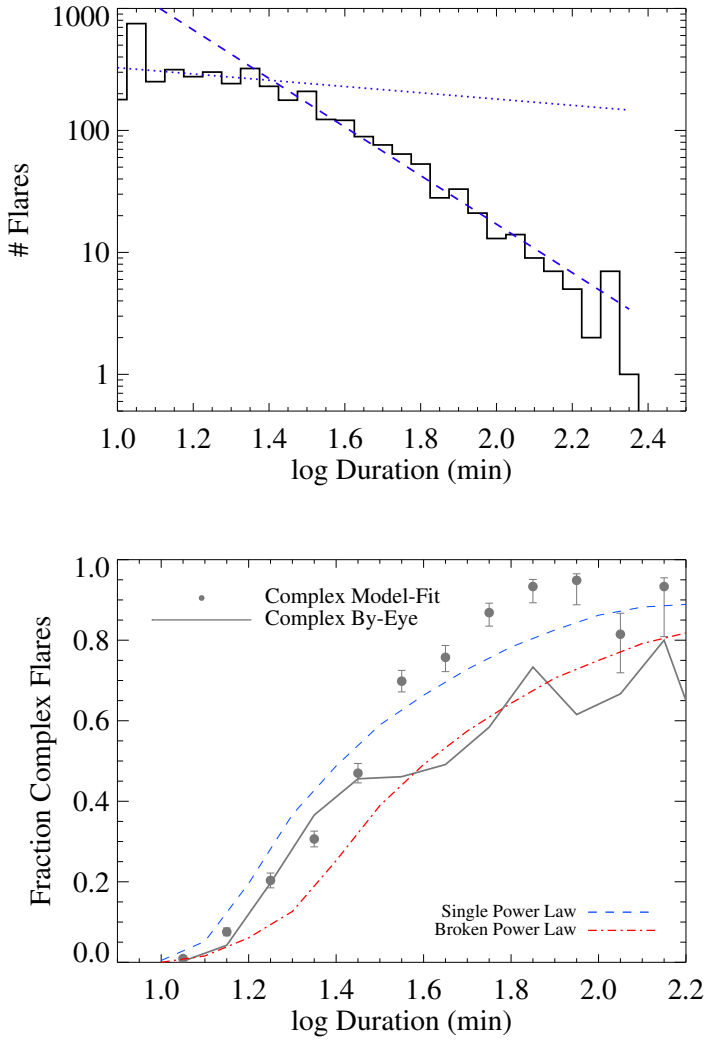

Figure 3.10: Top: Distribution of flare durations for all events longer than 10 minutes in our sample. Single power law (blue dashed line) and broken power law (blue dashed and dotted lines) fits are shown. The total event duration as selected in our by-eye analysis in §3.2 was used for each event. Bottom: Fraction of flares identified as complex using our iterative flare-fitting technique (black circles with error bars, see text for description), compared to those selected in our by-eye analysis (grey solid line). Two Monte Carlo models are overlaid (red and blue dashed lines), corresponding to the single and broken power-law durations distributions, see text.

The former interpretation may include phenomena such as homologous flares, induced or sympathetic flares, and large tangled or multi-loop structures, all of which are seen on the Sun. Large complex flare events on M dwarfs, such as the one in Kowalski et al. (2010), have previously been interpreted in this manner (Anfinogentov et al., 2013). Given the high level of magnetic activity and flaring on GJ 1243, the latter interpretation that some complex flares are chance superpositions likely accounts for at least *some* of the observed complex events. Larger energy, longer duration flares clearly have a higher likelihood of overlapping other non-associated flare events.

Our user-validated flare sample included a flare classification of complex versus classical, as described in §3.2.2. However, deciding if an event is complex requires that the two (or more) underlying flares be sufficiently separated in time, at least by a few minutes, such that the peaks are distinct. Additionally, the amplitudes of secondary flares within complex events must be high enough to clearly be distinguished above the morphology of the primary flare. These two considerations mean that the complex flare rate determined by eye is a lower limit. In Figure 3.10 (top) we show the distribution of flare durations for all events in our by-eye validated sample. The fraction of events classified by eye as complex as a function of their duration is presented in Figure 3.10 (bottom, solid grey line). As naively predicted above, the rate of complex flares does increases with the duration.

As described in §3.3.1, we fit all 3737 flare events with durations of 10 minutes or longer with $n = 1$ through $n = 10$ template fits. This produced a sample of 1141 flares (30.5%) that were best fit by an $n > 1$ model, as in Figure 3.9 (bottom), which we then classified as complex events. The fraction of model-fit complex events as a function of duration is also shown in Figure 3.10 (bottom, black points). The uncertainties were computed from the binomial error within each duration bin. The rate of complex flares from the model was higher than in our by-eye sample, as expected. Flares with durations greater than $\sim$50 minutes (log duration $\sim$1.7) show more than 80% chance of being complex events from the model, but only 60% from the users.

To compare with the observed fraction of complex flares as a function of duration, we

generated two Monte Carlo models for creating complex flares in our data. In the first model, we generated random flares with durations 10 minutes or greater, drawn from a single power law with slope -1.99 ± 0.02, which was determined by fitting the observed flare duration distribution for events longer than 20 minutes, shown in Figure 3.10 (top, dashed blue line). A single power law fit for flare energy distributions is commonly used (see Paper 1). In the second model, we used a broken power law, with a slope of -0.25 ± 0.03 for flares with duration between 10 and 20 minutes, and -1.99 ± 0.02 for flares greater than 20 minutes to more closely fit the data, as shown in Figure 3.10 (top, dashed and dotted blue lines). This two component, broken power law model is similar to the flare energy distribution model of Kashyap et al. (2002), though we emphasize we do not extend our simulation to unobservable "microflares".

The procedure for both the single and broken power law Monte Carlo models was the same. We began each model with a series of 50 trials, simulating a fixed number of flare events with durations ranging between 10 and 200 minutes. For these 50 trials, we adjusted the number of simulated flares between 3000 to 7500 events, with the first trial having 3000 flares and each subsequent trial increasing the number of flare events by 90. In each trial the specified number of flares was drawn from the respective duration distribution, and then placed at random start times throughout a blank light curve with the same time sampling as our dataset. Any flare events that overlapped in time were combined, and the resulting event was classified as complex. We saved the resulting total number of events, both complex and classical, for each trial.

We used these 50 trials to determine the number of simulated flares required to reproduce the observed total number of flare events. A second order polynomial was fit to the number of resulting versus number of simulated flares. For these models we required the resulting number of flares to match observations for the number of events (both classical and complex) with durations of 20 minutes or greater, which included 1750 flares from our by-eye sample. For the single power law model, generating this number of $\geq$20 minute events required an input of 6587 simulated flares, while the broken power law model required 4147 flares.

Using the respective number of input flares needed to generate the observed number of $\geq$ 20 minute events, we repeated this procedure 1000 times for both models. The fraction of complex events as a function of the resulting flare durations were recorded for each of the 1000 trials. The averaged complex flare fractions for both the single and broken power law models are shown in Figure 3.10 (bottom, blue and red dashed lines, respectively).

Our data appear to rule out a single power law for the underlying duration distribution. While the single power law model more closely reproduced the template-fit complex flare fraction curve, and was tuned to match the number of flares with durations 20 minutes or greater, it required far more flares than were seen in our data. The total number of short duration events (less than 20 minutes) resulting from this model exceeded our observation by over 1000. According to Paper 1, flares with durations between 10 and 20 minutes would be expected to have energies of log E $\approx$ 30.5–31 erg, and are easily detectable in the Kepler data for GJ 1243.

Instead, a broken power law is favored by our observations. Using the slopes fit from the observed duration distribution in Figure 3.10, the broken power law model was able to reproduce the user-selected complex flare fractions. This model, however, somewhat underproduced the 10–20 minute duration flares as compared to our sample. However, we have made very simple assumptions about the duration distribution. The true, underlying power law slopes for the distribution of individual flare durations are almost certainly different than what we observe, as many of the flares overlap and thus skew the observed duration distribution from its true shape.

The preferred (broken power law) model did not fully reproduce the the model-fit complex fraction. We believe this indicates that some of the complex structure our model-fitting scheme recovered must come from actual sympathetic flaring, for example from the same or other nearby active regions. We note that for a given complex event, from the Kepler observations alone we cannot differentiate whether the complex structure was due to random superpositions or sympathetic flaring. The overall rate of sympathetic flaring is likely represented by the excess complex flare fraction above the broken power law model. Obser-

vations with higher time cadence, and/or bluer wavelength coverage, may be able to detect additional lower amplitude complex flare structure, particularly for shorter duration flares. Further study on the complex flare fraction for other active stars would also be useful to constrain these model results. Studies of less active stars, where random superpositions are less frequent, will also help constrain the rate of sympathetic flaring.

3.5 Sample Completeness

As the sample of flares defined in this work is the largest such dataset for a single active M dwarf, our flare catalog will be a benchmark for many future empirical and theoretical investigations. Our flare auto-finding methodology is also generally useful for single-event detection in light curves. As such, it is critical that we accurately characterize the completeness and limitations of our sample.

We first quantified the completeness of the auto-finder algorithm by injecting artificial flares in to our light curves and testing for their recovery. A total of 500,000 artificial flares were tested in our 11-month dataset. Each artificial flare was generated using the flare template described by Equations 3.1 and 3.4. The $t_{1/2}$ and amplitude parameters for each artificial flare were drawn randomly from ranges of 1–60 minutes and 4×10^{-5}– 4×10^{-2} relative flux, respectively. The flare peak-time was then placed at a random time within the 11-month light curve.

Flares were considered recovered if the peak time of the input artificial flare was contained within any flare event found by the auto-finder. Care was taken to prevent the input artificial peak times from overlapping known flare events in the light curve. However, artificial flares that were placed near real large flares in the light curve could be combined into a complex event by the auto-finder. These were also considered recovered events. We did not test recovery of the input flare duration or amplitude. In Figure 3.11 (solid line) we show the fraction of recovered flares as a function of their event energies (in units of equivalent duration). The median fraction of recovered flares was computed in bins of log ED = 0.2 sec. The auto-finder alone recovers events with log ED $>$ 0 sec, corresponding to flares with

durations of 10 minutes (Paper 1), at 70% completeness.

Our final flare sample was selected by a two-step procedure, first by the automated detection and then refined by human validation. As such, measuring completeness of the sample is a difficult task, since it would require measuring both the performance of the auto-finder, and studying the agreement between humans. This could theoretically be accomplished by injecting artificial flares into the light curves and tracing their recovery in both the auto-finder and the human validator steps, but would require a severe increase in human labor that was not practical. We therefore look at the user results separately.

In Figure 3.11 (dashed line) we show the fraction of users who selected the flares in our final vetted sample as a function of their energies. Every flare in our sample was checked by at least five users. We used the peak time for each flare to compute the number of users who selected the event, and normalized by the number of users who validated the respective month of data. The median fraction of users who identified flares was again computed in bins of log ED = 0.2 sec. Users identified flares with 90% agreement for events with log ED > -0.5 sec, corresponding to flares with durations of $\sim$5 minutes (Paper 1). This illustrates that our final sample is 90% complete, and that the auto-finding methods, while useful as a first pass, still need additional work to approach the confidence level of identifying by eye.

3.6 Unusual Flares

We have shown that our flare template is able to reproduce the morphology for many types of flare events, both classical and complex. However, a small fraction of flare events in our sample were not well fit by our template. To identify unusual flares in our sample, we selected events whose best-fit model had a reduced χ^2 of at least 15. This yielded 49 events (1.3%) that were poorly fit by our iterative model approach. These were preferentially higher energy flares: median log E = 32.4 erg for unusual events compared to median Log E = 30.8 erg for the entire sample of flares of duration $\geq$10 minutes. We note that the rate of unusual flares is highly dependent on the details of our model fitting scheme, and changes to the BIC criteria or the iterative approach may result in fewer flares being classified as unusual.

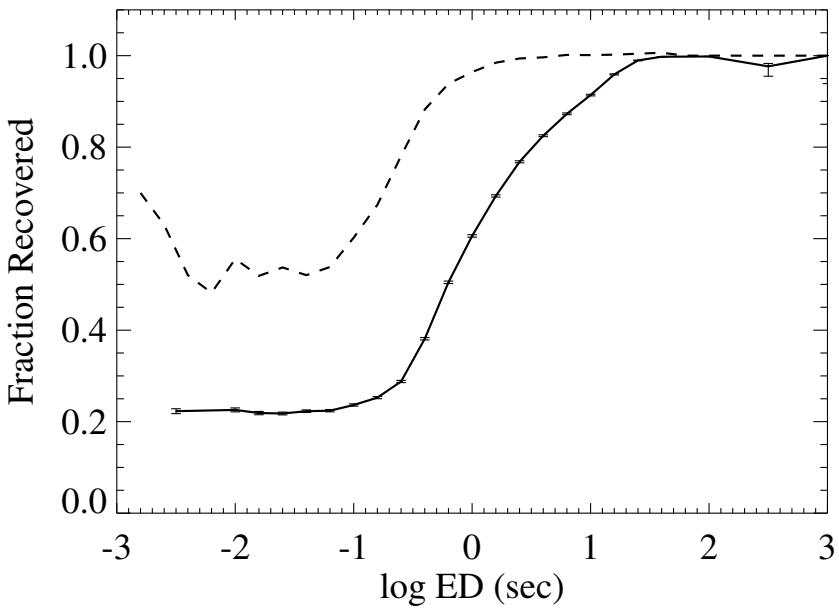

Figure 3.11: Median fraction of recovered flares by the automatic flare-finding algorithm from tests using 500,000 artificial flares, calculated as a function of the event energies (solid black line). Error bars on each bin represent the binomial errors on flare recovery. Our method is 90% complete for flares with ED greater than 10 seconds, and 70% complete for flares with ED greater than 1 second (approximately 10 minute duration events). For comparison, the fraction of users who confirmed the flares in our final sample is shown for the same bins of event energy (dashed line).

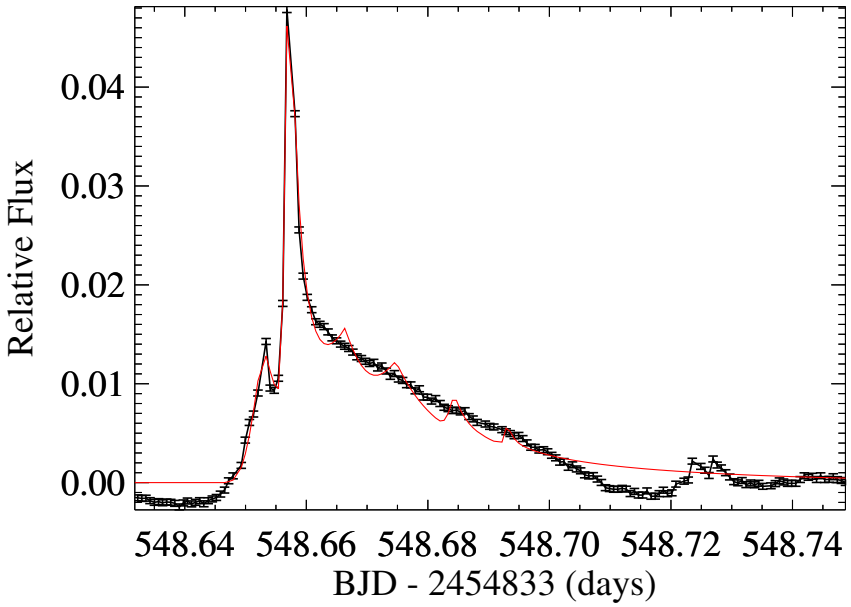

Figure 3.12: An example of a flare with an unusual decay phase profile that is not well-fit by a combination of classical flare templates (black points). The best-fit model for this event (red line) included six component flares, but does not reproduce the structure observed in the decay phase. Negative flux before and after the event was an artifact of over-subtraction of the starspot signal, and did not affect our flare model fitting.

Our model is able to fit events classified as "impulsive" and "hybrids" by Kowalski et al. (2013). Fitting their "gradual" flares (specifically GF1) would require either a large number of small component flares or a different template entirely. Without multi-wavelength or spectral data for such events, we are unable to address the underlying physical differences between these unusual or gradual flares, and the majority of the impulsive flares that are similar to our template.

In Figure 3.12, we show an example of a high-energy complex flare that was poorly fit by our model. The best-fit complex flare model, which had six component flares and a reduced $\chi^2 = 15$, is shown. Negative relative flux near the beginning and end of the flare is

an artifact from subtracting the local starspot modulation using a linear fit, and does not affect our analysis. The two main peaks in the light curve are well fit by two impulsive flare components. The slow decay phase, however, does not follow the model template (see Figure 3.7). Instead, after the largest amplitude peak, this flare exhibits a very slow, almost linear decay in flux for over an hour. Our model attempted to reproduce this decay with a series of decreasing energy impulsive flare events, seen as the subsequent ripples in the model light curve. While such ripples have been observed in the decay of complex events (e.g. Kowalski et al., 2010), they do not follow the data for this flare.

Manually fitting the long decay using a single low-amplitude and large $t_{1/2}$ flare component also could not reproduce the linear flux decay observed. The gradual decay profile of this flare could be reproduced using a much larger number of short timescale flares with decreasing energies. This is reminiscent of flare models used to reproduce the gradual soft X-ray decay in Solar flares (Warren, 2006), or of large flares from reconnection along arcades of loop structures (Grigis & Benz, 2005). Such an unusual decay profile might also be produced as a result of a flare occurring near, and possibly rotating over, the stellar limb (Tovmassian et al., 2003). Given the rapid rotation rate for GJ 1243 this would appear a reasonable possibility for long duration flares, though it should not affect the majority of flares in our sample. We may speculatively ascribe the evolution of the flare in Figure 3.12 to an event initially occurring relatively close to the limb of the star, which then rotated partially or fully out of view.

3.7 Discussion

We have presented the largest sample of flares ever compiled for a single star besides our Sun, totaling 6107 unique flaring events from over 11 months of photometric monitoring. Of these, 15% were classified as complex flares. In Paper 1 we described the general properties of flares on GJ 1243, including correlations between the decay time, duration, and flare energy for both complex and classical flares, and the waiting time distributions of the flares. Paper 1 found no correlation between the rotation phase of the star and the flare rate or average

flare energy. We note here that the ratio of complex to classical flares also is not correlated with the stellar rotation phase.

Using a subset of 885 well-measured classical (single peak) events, we generated an empirical flare template, which we fit with a fourth-order polynomial rise and double exponential decay. This template has only two morphological parameters: the amplitude and the characteristic timescale, $t_{1/2}$. We demonstrated the utility of this template in decomposing complex, multi-peaked flare events into their constituent classical flares. This empirical template will be a powerful tool for flare investigations on other stars. Future studies will show whether the morphological properties observed in GJ 1243 flares (e.g. the impulsive versus gradual decay profiles) as well as the choice of a single characteristic timescale are generic for all stellar flares.

We have restricted ourselves to studying the flare morphology for a single star in the Kepler dataset, GJ 1243. A detailed spectroscopic and photometric characterization of this fascinating star is also underway (Wisniewski et al. 2015, in prep.) Modeling the starspot modulations in the light curve, which we have thus far treated as a noise source, has also produced constraints on the differential rotation rate, spot geometry, and spot lifetimes (Davenport et al., 2015a)

The degree of self-similarity between the classical flares in our template sample is remarkable. The "impulsive phase" (rise and decay) of the flare dominates the characteristic timescale ($t_{1/2}$) used in generating the template. The slow decay phase, which contains roughly half the flare energy, is also well traced. This indicates that our single timescale parameterization is appropriate for flare data at this cadence and in this wavelength regime, and that the two cooling phases do not appear to be independent, at least for flares on this star. Higher cadence investigations will be useful to uncover the universality of this template shape for the rise and both decay phases.

However, as Andrews (1965) first noted, the shape of the slow decay phase may depend on the energy or temperature of the flare event. Gershberg & Shakhovskaya (1973) described star-independent properties for the slopes of the decay phases of flares, using power-law decay

shapes rather than the exponential functions we found in our template. Star-independent flare morphology was also described by Shakhovskaya (1989). These results have been disputed by Kunkel (1974), however. Building on the work in Paper 1, and using the methods developed in this chapter, we intend to look at flares on stars across all spectral types in the Kepler database. This will address the possible dependence of flare light curve morphology on the underlying star, as well as characterize the predicted yield and properties of flares in next generation time domain surveys.

3.8 Acknowledgments

We gratefully acknowledge support for this work from NASA Kepler Cycle 2 GO grant NNX11AB71G, NASA Kepler Cycle 3 GO grant NNX12AC79G. SLH, JRAD and LH acknowledge support from NSF grant AST13-11678. EJH, AFK, and SLH acknowledge support from NSF grant AST08-07205. JRAD wishes to thank Andrew C. Becker for valuable insights in model fitting, and John J. Ruan for discussions of time series analysis.

This chapter includes data collected by the Kepler mission. Funding for the Kepler mission is provided by the NASA Science Mission directorate. Some of the data presented in this chapter were obtained from the Mikulski Archive for Space Telescopes (MAST). STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. Support for MAST for non-HST data is provided by the NASA Office of Space Science via grant NNX13AC07G and by other grants and contracts.

Chapter 4

DETECTING DIFFERENTIAL ROTATION AND STARSPOT EVOLUTION ON THE M DWARF GJ 1243 WITH *Kepler*

We present an analysis of the starspots on the active M4 dwarf GJ 1243, using four years of time series photometry from *Kepler*. A rapid $P = 0.592596 \pm 0.00021$ day rotation period is measured due to the $\sim$2.2% starspot-induced flux modulations in the light curve. We first use a light curve modeling approach, using a Monte Carlo Markov Chain sampler to solve for the longitudes and radii of the two spots within 5-day windows of data. Within each window of time the starspots are assumed to be unchanging. Only a weak constraint on the starspot latitudes can be implied from our modeling. The primary spot is found to be very stable over many years. A secondary spot feature is present in three portions of the light curve, decays on 100-500 day timescales, and moves in longitude over time. We interpret this longitude shearing as the signature of differential rotation. Using our models we measure an average shear between the starspots of 0.0047 rad day^{-1}, which corresponds to a differential rotation rate of $\Delta\Omega = 0.012 \pm 0.002$ rad day^{-1}. We also fit this starspot phase evolution using a series of bivariate Gaussian functions, which provides a consistent shear measurement. This is among the slowest differential rotation shear measurements yet measured for a star in this temperature regime, and provides an important constraint for dynamo models of low mass stars.

Material from this chapter was previously published in collaboration with Leslie Hebb and Suzanne L. Hawley June 2015 edition of the Astrophysical Journal (Davenport et al., 2015a), and has been reproduced here with permission of the American Astronomical Society.

4.1 Introduction

For low-mass, fully convective stars, the nature of the magnetic dynamo and the role of differential rotation is not so clear. Some radial and surface differential rotation is expected to exist, due to the combination of rotation and convection. However, despite the deep convective zones of M dwarfs, their long convective turnover timescales result in a lower amplitude of surface differential rotation or shear (Küker & Rüdiger, 2008; Kitchatinov & Olemskoy, 2011), (the Ω effect). Since these stars are nearly or fully convective from surface to core, and therefore lack a "tachocline" interface region in which to store toroidal magnetic field, the dynamo mechanism must be fundamentally different than the popular $\alpha\Omega$ dynamo model for the Sun (e.g. Parker, 1955; Schrijver & Zwaan, 2000). Instead, this convectively driven process is known as an α^2 dynamo. For rapidly rotating M dwarfs, the magnetic field strength is expected to be increased, which suppresses differential rotation and forces nearly solid-body rotation (Browning, 2008). Without strong radial or surface differential rotation to organize the global magnetic field, activity cycles may not be present, and the surface magnetic topology is predicted by some models to be highly non-axisymmetric and multipolar (e.g. Chabrier & Küker, 2006).

However, observations of many low mass stars reveal highly organized, strongly poloidal magnetic fields (e.g. Morin et al., 2008a) and prominent long-lived starspot features (e.g. Barnes et al., 2005). Some rapidly rotating low-mass stars show indications of polar "starspot caps" possibly due to this large scale dipolar field (Donati & Collier Cameron, 1997; Morin et al., 2008b), while others do not (Barnes et al., 2004; Morin et al., 2010). Though differential rotation is expected to play a lesser role for these rapidly rotating low-mass stars, even small amounts of differential rotation may help to organize the chaotic, α^2 driven magnetic fields in to a coherent, axisymmetric field (Kitchatinov & Olemskoy, 2011), capable of producing very long-lived polar spot features. Therefore, given the wide variety of observed surface magnetic topologies, and the complex inter-dependence of rotation, differential rotation, and the magnetic field, measuring differential rotation rates for low mass stars is a high priority

for constraining dynamo theory.

Rotation can now be measured with relative ease for many stars, for example using spectral line broadening that produces v sin i measurements, or periodic flux modulations due to starspots in precision space-based time series photometry. Data from the *Kepler* mission (Borucki et al., 2010) has revolutionized the study of stellar rotation using starspot modulations, with tens of thousands of stars having measuring rotation periods (McQuillan et al., 2014; Reinhold et al., 2013), and has revealed starspot properties for stars ranging from solar mass (Bonomo & Lanza, 2012) to brown dwarfs (Gizis et al., 2013).

Differential rotation, however, is notoriously difficult to detect for stars. Spectral techniques can trace active regions at different latitudes for stars with lower activity levels such as the Sun (Bertello et al., 2012). Detecting differential rotation via Zeeman Doppler Imaging (ZDI; Semel, 1989; Donati & Brown, 1997) requires comparing complex surface magnetic reconstructions or maps between subsequent visits. Photometric surveys may be able to produce differential rotation rates for an ensemble of active stars (Reinhold & Reiners, 2013). A recent blind survey of competing techniques for detecting rotation and differential rotation from model photometry showed excellent agreement in recovering rotation periods from active stars. However, a complex degeneracy was found between differential rotation rate, starspot lifetimes, and the number of starspots present, and little agreement between competing methods and the model light curves (Aigrain et al., 2015). In general, methods for detecting differential rotation in photometry follow one of two approaches: 1) Fourier methods, which measure the broadening or splitting of peaks in the power spectrum, autocorrelation function, or periodogram, or equivalently by decomposition of the light curve using sine functions (Reinhold et al., 2013). These methods utilize the entire light curve at once, and are efficient for analyzing large volumes of data from many stars, but may suffer more from the degeneracies mentioned above. 2) Tracking specific starspot features either via light curve inversion (Roettenbacher et al., 2013), or light curve modeling for individual starspots (Frasca et al., 2011). These methods are more computationally expensive, but their results seem robust for rapidly rotating stars with long-lived spots.

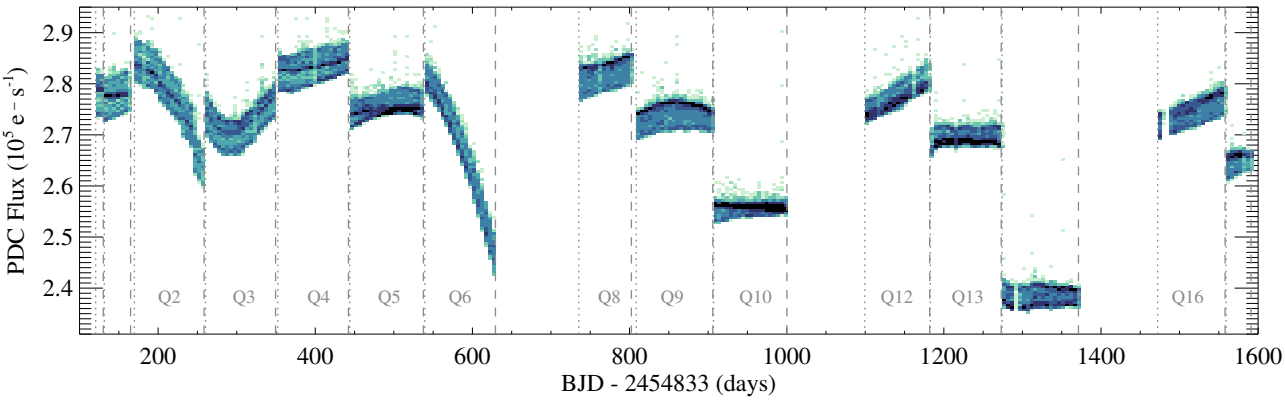

Figure 4.1: The long cadence PDC MAP light curve for GJ 1243. Pixel shade (light to dark) indicates the density of epochs. Breaks in the light curve due to quarterly spacecraft rolls are indicated (grey dashed lines).

In this chapter we venture into a relatively new region of starspot evolution parameter space, detecting very gradual differential rotation and spot decay for a rapidly rotating M dwarf. The fast time cadence and continuous monitoring provided by *Kepler*, along with a short stellar rotation period, allow us to trace small changes in starspot phase and amplitude over long periods of time. In Section 4.2.2 we describe our target, the active M dwarf GJ 1243, and the previous investigations of this low-mass star with *Kepler*. Our detailed light curve modeling is presented in Section 4.3. We trace small changes in starspot phase over four years, and interpret this as a signature of differential rotation in Section 4.4. A simpler approach to detect this slow differential rotation by modeling the phase evolution with Gaussians is given in Section 4.5. We place the differential rotation signal from GJ 1243 in the context of other cool stars, and compare the *Kepler* photometric results with older ground-based data in Section 4.6. Finally, in Section 4.7 we provide a summary of our results, and discuss the great potential for understanding starspots and the stellar dynamo still to be realized from the unique photometric *Kepler* database and future missions.

4.2 GJ 1243

The target of our study is the nearby mid M dwarf, GJ 1243 (Kepler ID # 09726699). This star has a short rotation period of 0.5926 days that has been noted in previous studies of *Kepler* light curves (Savanov & Dmitrienko, 2011; McQuillan et al., 2013). The spectral type has been measured as M4 (Hawley et al., 2014), placing GJ 1243 near the fully-convective boundary where stars are expected to remain magnetically active for many Gyr (N. Reid & S. L. Hawley, 2000; West et al., 2008). Using the parallax distance of 11.9 pc from Lépine & Shara (2005), the apparent K-band magnitude from Zacharias et al. (2013), and the M_K-mass relation from Delfosse et al. (2000), we estimate a mass for GJ 1243 of 0.24 $M_{\odot}$. The convective turnover timescale for stars in this mass range (assuming M=0.235$M_{\odot}$) from Kiraga & Stepien (2007) is quite slow at $\tau_{conv} \approx 70$ days. Comparing this timescale to the rotation period, we find GJ 1243 has a very low Rossby number of $R_o = P_{rot}/\tau_{conv} \approx 0.008$. Lucky imaging of GJ 1243, as well as ground based spectroscopy, have shown no indications of a binary companion (Wisniewski 2015 in preparation). In addition, GJ 1243 has been the subject of detailed flare activity studies with *Kepler* data, producing the largest catalog of M dwarf flares ever observed for a single star (Hawley et al., 2014; Davenport et al., 2014). In this section we describe our treatment of the *Kepler* data for this active M dwarf, removing systematic trends from the light curve, and detecting a periodic signal.

4.2.1 Kepler Long Cadence Data

The *Kepler* light curve for GJ 1243 contains dramatic stellar variability in the form of flares and starspots. Ramsay et al. (2013) have examined the flare energy distribution using one quarter (Q6) of data from *Kepler*. Davenport et al. (2014) and Hawley et al. (2014) used 11 months of *Kepler* short cadence (1-minute) data for GJ 1243, over 300 days worth in total, to robustly measure the flare rate and develop a statistical understanding of the flare morphology from this very active dwarf. For these flare studies the starspot signature had been treated as a noise source to be smoothed out.

In the present investigation, we utilized all available long cadence (30-minute) *Kepler* data for GJ 1243 to study the evolution of the starspots while minimizing the impact of small amplitude flares. GJ 1243 was observed in 14 separate quarters of *Kepler* data (Q0–Q6, Q8–Q10, Q12–Q14, and Q16–Q17), spanning over four years of observation (MJD 54953.04 through 56423.50). We used the most recent reduction of the *Kepler* data available, including the "PDC-MAP" Bayesian de-trending analysis from (Smith et al., 2012). The entire 4-year catalog PDC-MAP light curve for GJ 1243 is shown in Figure 4.1. Data from Q7, Q11, and Q15 was not available due to the failure of CCD Module 3 in 2010, which GJ 1243 resided on for one quarter of the year.

In Figure 4.1, large discontinuities in the flux are apparent between quarters, as well as systematic trends in the mean flux within quarters. These long timescale variations are systemic to *Kepler* data, due to spacecraft drift and calibration limitations, and are not astrophysical. For every quarter, we fit and subtracted low order (linear or quadratic) polynomials from the data to remove these systematic errors and discontinuities. Because the stellar rotation period is so short, and each quarter contains on average $\sim$150 rotations, these polynomial fits do not affect the starspot signal on the timescales we are interested in.

Large amplitude flares were also present in our data, visible as positive flux excursions throughout the light curve in Figure 4.1. While the short cadence *Kepler* data for GJ 1243 is a treasure trove for flare studies (e.g. Davenport et al., 2014), only the largest energy flares are visible in the 30-minute data (see Walkowicz et al., 2011). To remove the flares from our analysis, we smoothed the light curve with a 12-hour "boxcar" filter, and then discarded epochs with fluxes that deviated by more than 0.3% from the smooth flux. This boxcar smoothing was only used to remove outlying epochs, and was not used in our starspot analysis. These smoothing values were arrived at by eye to remove the most dramatic flares and outliers in the data. As this was not a comprehensive outlier removal scheme, some small amplitude flares and data systematics remained in the light curve. These small amplitude excursions occurred stochastically throughout the light curve, had no dependence on rotational phase, and therefore did not affect our spot modeling results. As discussed in

Lurie et al. (2015), saturation can affect the *Kepler* light curves for flare stars during the brightest flare events. However, the starspot modulations for GJ 1243 were low amplitude, and the quiescent flux level was not near the saturation limit. While the brightest flux excursions due to flares may be affected by saturation, our starspot analysis is not. Our final, inter- and intra-quarter polynomial detrended, flare-cleaned light curve for GJ 1243 contained 47,478 epochs of data over the four years of *Kepler* long cadence observations.

4.2.2 Periodic Signal

The rapid rotation of GJ 1243 was first detected from periodic flux modulations due to starspots by Irwin et al. (2011) using ground-based photometry from the MEarth project (Nutzman & Charbonneau, 2008; Irwin et al., 2009). Following the initial Q0 release of *Kepler* data, Savanov & Dmitrienko (2011) published the first analysis of the starspots on GJ 1243, using 44 days of continuous long cadence data. They reported a rotation period of 0.593 days for GJ 1243, and found that GJ 1243 exhibited two starspot features, separated in longitude by 203°, and both stable in position over the 44 days of observation (equal to ∼74 rotation periods). The starspots covered 3.2% of the visible stellar surface, with a modest amount of growth reported over Q0.

A study searching for rotation periods using the autocorrelation function for ∼2500 *Kepler* M dwarf stars was carried out by McQuillan et al. (2013) using 10 months of *Kepler* photometry. They reported a rotation period of 0.593 days for GJ 1243 as well. However, a larger scale analysis of over 40,000 active *Kepler* stars by Reinhold et al. (2013), using the Lomb-Scargle periodogram method (Lomb, 1976; Scargle, 1982), did not report a rotation period for GJ 1243, as the star's rapid rotation was below their period cutoff.

These previous studies of GJ 1243 only reported the stellar rotation period to an accuracy of 0.001 day (∼86 seconds). With such a short rotation period for GJ 1243, an error of 0.001 days would result in phase-folded data being out of phase by an entire rotation within one year. Thus, to measure any real phase evolution of the starspot features over four years we must determine the most accurate mean rotation period possible. We computed the

normalized Lomb-Scargle periodogram using the entire 4-year detrended long cadence light curve, using no frequency oversampling or smoothing. The strongest peak in the resulting periodogram was very narrow, and had a period of 0.592596 days. We then computed the Lomb-Scargle periodogram over each of the 14 quarters of data individually. The mean period from all quarters we recovered was 0.592673 days, which was only ∼6.5 seconds longer than the period found from all quarters simultaneously. These 14 period estimates had a standard deviation of 0.00021 days, or about 18 seconds, which we adopt as the period uncertainty. Since the rotation period is very stable over the course of the *Kepler* observations, we assume the period determined from the entire light curve, $P = 0.592596 \pm 0.00021$ days, for our analysis.

We then empirically defined the ephemeris of the flux minimum by phase-folding the entire *Kepler* light curve at this rotation period. The phase of flux minimum was fit using a least squares regression with a Gaussian function, which determined an ephemeris of $t_0 = 2454833.11567807 \pm 0.00015$. In Figure 4.2 we show median-smoothed 10-day windows of the entire GJ 1243 light curve, phased using this rotation period and ephemeris. The primary dip in brightness stays fixed near Phase=0 over the 4 years of observation, which is due to the primary starspot. Slow evolution in both phase and amplitude of the secondary starspot feature is clearly seen. The secondary starspot is almost entirely absent at Time $\sim$ 700 days (using units of time as BJD - 2454833.11567 days), while the primary and secondary starspots appear to have nearly equal amplitudes at Time $\sim$ 1100.

Given the long starspot evolution timescale and short rotation period for GJ 1243, along with the nearly continuous *Kepler* light curve for most of the 4-year timespan, we are able to study the change in starspot properties in much higher temporal detail than illustrated in Figure 4.2. Using 10-day windows of time, we show the 4 year continuous phase evolution of flux from GJ 1243 in Figure 4.3. For visual clarity the data is folded twice in phase. White vertical gaps correspond to quarters with no *Kepler* data, as seen in Figure 4.1. Each column of pixels in this phase versus time flux map contains data spanning 10 days. This binning resulted in ∼16 rotation periods per column, with an average of over 400 data points. Each

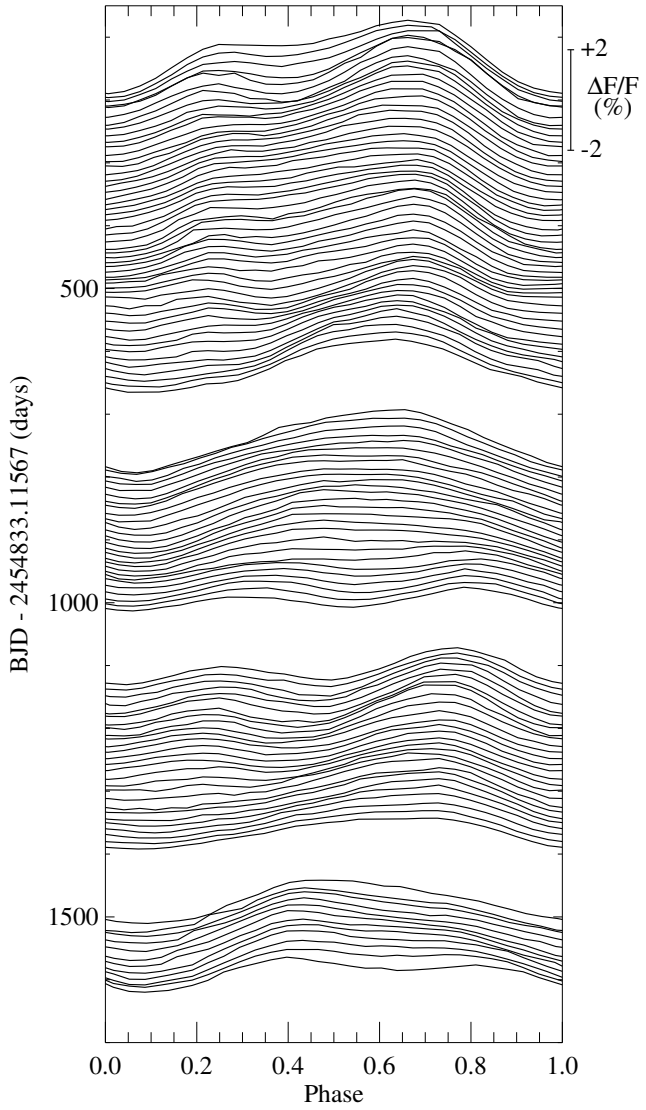

Figure 4.2: Phase-folded, median smoothed light curves for GJ 1243 from 10-day windows of time, showing the slow evolution of the starspot modulations over time. The vertical position for each curve corresponds to the start time of the 10-day window on the left axis. Each time window is scaled to the same relative flux, shown on the right axis. The primary dip, centered at Phase=0 corresponds to the long-lived starspot.

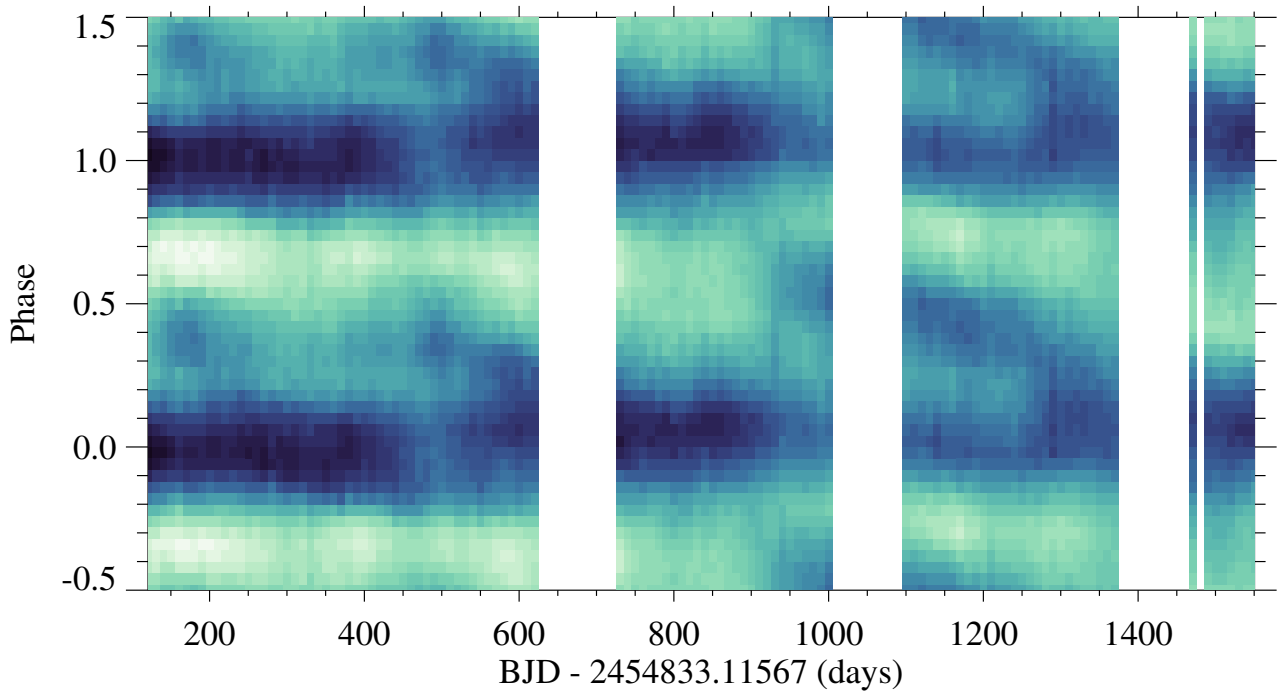

Figure 4.3: Continuous phased light curve map for the entire *Kepler* long cadence dataset. Pixel shade, from dark to light, indicates the median flux in each (time, phase) bin. Vertical white gaps correspond to times with no *Kepler* data, as in Figure 4.1. Pixels span 10 days in time and 0.04 in phase. The starspots are seen as dark regions in this diagram, which evolve in time from left to right.

row spans 0.04 in Phase, or equivalently a 14.4° slice in longitude. The median value for the flux within each (time, phase) pixel corresponds to the shading, with the darkest regions corresponding to a flux 1.5% below the median value, and the lightest pixels 1.5% above the median flux.

The dark band centered at Phase=0 in Figure 4.3, which extends throughout the timespan of the data, is due to the primary starspot. This feature does not significantly change in phase over the course of our data. The flux amplitude for the primary spot is also nearly constant. There is an apparent change in the starspot flux amplitude around day ∼500 and day ∼900, due to the the presence of the secondary starspot combined with the systematic errors in the flux calibration. The starspot features seen in Figure 4.3 are very large compared to spots seen on the Sun, appearing to span 50–90° in longitude. The detailed geometries of these features cannot be determined from this phase versus time flux map, and each observed "starspot" may in fact be a large spot group. Additionally, we cannot constrain the total starspot coverage, which may include many smaller spots and active regions across the entire stellar surface. Instead we are observing the total flux asymmetry due to these spots or spot groups.

The secondary starspot feature continuously changes in both phase and flux amplitude (equivalently pixel shade) in this diagram. This secondary feature seems to emerge and decay at least three times over the 4-year dataset, each time appearing nearly on the opposite hemisphere of the star and evolving towards the primary starspot. Note, a decrease in phase corresponds to a starspot *advancing* in longitude in the direction of rotation over time. We interpret the slow, linear phase evolution of the secondary starspot to be the signature of differential rotation.

4.3 Modeling the Light Curve

To quantitatively trace the differential rotation on GJ 1243, we must determine the precise sizes and positions of the starspots over time. To accomplish this we performed a detailed fit to the *Kepler* light curve using the starspot modeling software from L. Hebb (2015 in

preparation). Here we give a brief overview of this light curve modeling program, as well as our specific use with the GJ 1243 system.

The starspot modeling code simulates the star as a sphere with uniform surface brightness and limb darkening onto which circular, gray starspots are fixed. Limb darkening is implemented by treating the star as a series of overlapping, concentric circles with brightness values defined by the 4-coefficient limb darkening model of Claret & Bloemen (2011). Starspots are modeled as non-moving circular regions with a fixed flux contrast relative to the photosphere, and may be placed anywhere on the stellar surface. At each time step in the input light curve, as the model star rotates, the code calculates the flux blocked by the spots rotating in and out of view, and thus generates a synthetic light curve.

The program can generate a synthetic light curve for a single star, with or without a transiting exoplanet, and with the spin-axis of the star and orbital axis of the planet in any orientation (aligned or misaligned). To derive the properties (latitude, longitude and radius) for a number of spots that best reproduces the observed flux modulations, a χ^2 comparison is made between the observed data and a synthetic light curve. The model light curve generating engine is wrapped with several types of Markov Chain Monte Carlo (MCMC) samplers, including an affine invariant MCMC based on Foreman-Mackey et al. (2013), which explores the parameter space to find the lowest χ^2, and thus the optimum spot properties.

The program requires that we choose the number of spots on the star a priori, and that the spot distribution remains static. We only analyze a subset of the *Kepler* data at any one time, using a "window" to model a timescale over which we do not expect the spots to evolve. By sliding this window over the full length of the light curve and running the code many times, we fit the entire light curve and determine the evolution of the spots. We emphasize that the MCMC runs are done independently, generating a unique best-fit spot solution within each window. This approach of multiple discrete models over time avoids parameterizing the starspot evolution with analytic functions as has been done previously (e.g. Kipping, 2012), which in turn allows us to track non-linear behavior in the size and

position evolution of the spots. We refer the reader to L. Hebb (2015 in preparation) for a description of the full details and capabilities of this program, and briefly describe our specific use below.

We split the GJ 1243 light curve into windows with 5 day durations, or approximately 8.4 rotation periods at the *Kepler* 30-minute cadence. The short rotation period, combined with the slow evolution of spot features seen in Figure 4.3, resulted in many stellar rotations for each window, minimizing the effect of spurious light curve features such as flares or small data gaps. Each time window was required to contain at least 100 data points, or equivalently $\sim$3.5 rotation periods. Each subsequent time window was advanced by 2.5 days, providing two independent MCMC solutions for each datum. A total of 447 such time windows were used spanning the 14 quarters of data.

We assumed a fixed flux contrast value of 0.7 for the starspots, which is consistent with contrast values seen for spots on active giants, as well as the average contrast of the solar umbra (e.g. Berdyugina, 2005). Note that while resulting spot sizes are directly dependent on the contrast value used in our model, the longitude and therefore the differential rotation is not affected. The default value of 100 annuli was used to compute the limb darkening. Based on v sin i measurements from echelle spectroscopy of GJ 1243 (Wisniweski 2015 in preparation) and our measured rotation period, we used a fixed inclination of 32 degrees.

For each of the 447 windows of time, we modeled the GJ 1243 light curve using two starspots where each starspot is defined by fitting three parameters: its latitude, longitude and radius. This was the simplest model that was able to reproduce the observed flux modulations for all time windows to high accuracy. We note that some time windows were well fit using a single spot solution, particularly at Time $\sim$ 800 in Figures 4.2 and 4.3, where the flux modulation was dominated by a single sine-curve like feature. Models with higher numbers of spots (three or more) were tested and could easily reproduce the observed flux modulations, but were not preferred when properly compared to the two-spot models with fewer free parameters.

Constraining the latitudes for large starspots is often difficult when deriving 2-dimensional

starspot configurations from 1-dimensional light curves. There exists a well known degeneracy between spot latitude and radius, resulting in families of solutions for spots at a given *longitude* but a range of latitudes and spot sizes that provide equally good fits to the observed light curve. Therefore, we chose to fix the latitudes for the two starspots to break this degeneracy in our model runs. This does not affect our final conclusions because the differential rotation measurements depend only on the derived longitudes of the spots. To select latitudes at which to fix the two spots in our model, we ran our entire light curve modeling analysis for 1/10th of the time windows, and using five configurations of starspot latitudes. For each model configuration one spot was fixed at the stellar equator (0°), and one at a higher latitude towards the inclined pole. The five higher latitudes spot positions tested were (72.8°, 55.6°, 38.4°, 26.9°, 9.8°). Note, given the inclination of 32°, spots above ∼58° would be partially or fully visible during the entire stellar rotation, and therefore would produce less flux modulation. The starspot longitudes and sizes were allowed to vary in configuration.

Each resulting set of MCMC solutions produced comparably good fits to the light curve, and had the same number of free parameters. The individual sine-like modulations seen in the light curve were not required to correspond to the higher or lower latitude spot in any given model configuration. As a result, some models would exhibit a "flip" between spot latitudes for a given feature at nearly the same longitude between subsequent time windows. This flipping was observed for the two configurations with higher latitude spots (72.8° and 55.6°). For our final analysis we chose the solution set with the highest latitude configuration that did not exhibit this flipping in spot latitudes between subsequent time windows. The two starspots in our analysis were therefore fixed at latitudes of 38.4° and 0°. We note our resulting longitudinal shear results were insensitive to the latitudes chosen.

The affine invariant sampler based on Foreman-Mackey et al. (2013) was employed for each time window, with random starting values for the spot radius and longitude, but fixed latitudes as described above. For each window of data, the MCMC was run for 300 steps using 100 walkers, and the "a scale" parameter was set to 2.0. To carry out these independent MCMC realizations efficiently in parallel, we used `CONDOR` (Litzkow et al., 1988; Thain et al.,

2005) to distribute the 447 MCMC explorations across 180 Linux workstation computer cores. Each window's MCMC chain was converged after 300 steps, and the starspot configuration that produced the best-fit (lowest χ^2) solution for each time window was adopted. We note that the phase-folded data within each window had a scatter about 10 times greater than the typical photometric uncertainty given in the *Kepler* data. This was due to errors in the underlying light curve and limitations of our detrending algorithm, as well as small amplitude evolution of the starspot features within each window. The starspot flux modulation signal was more than 20 times greater than this scatter. Average reduced χ^2 values were $\sim$2 per time window, assuming a 10 times increase in the photometric uncertainty.

In Figure 4.4 we show phase-folded light curves for two representative time windows of data with their best fit solutions overlaid, along with orthographic map projections of the model stellar surface showing the best-fitting spot configurations. This map projection demonstrates the inclination of the star as well as the relatively large size of the starspots. The higher latitude spot appears at nearly the same phase (longitude) and size, while the lower latitude spot shrinks in radius and advances in longitude (lower phase) in the direction of the stellar rotation.

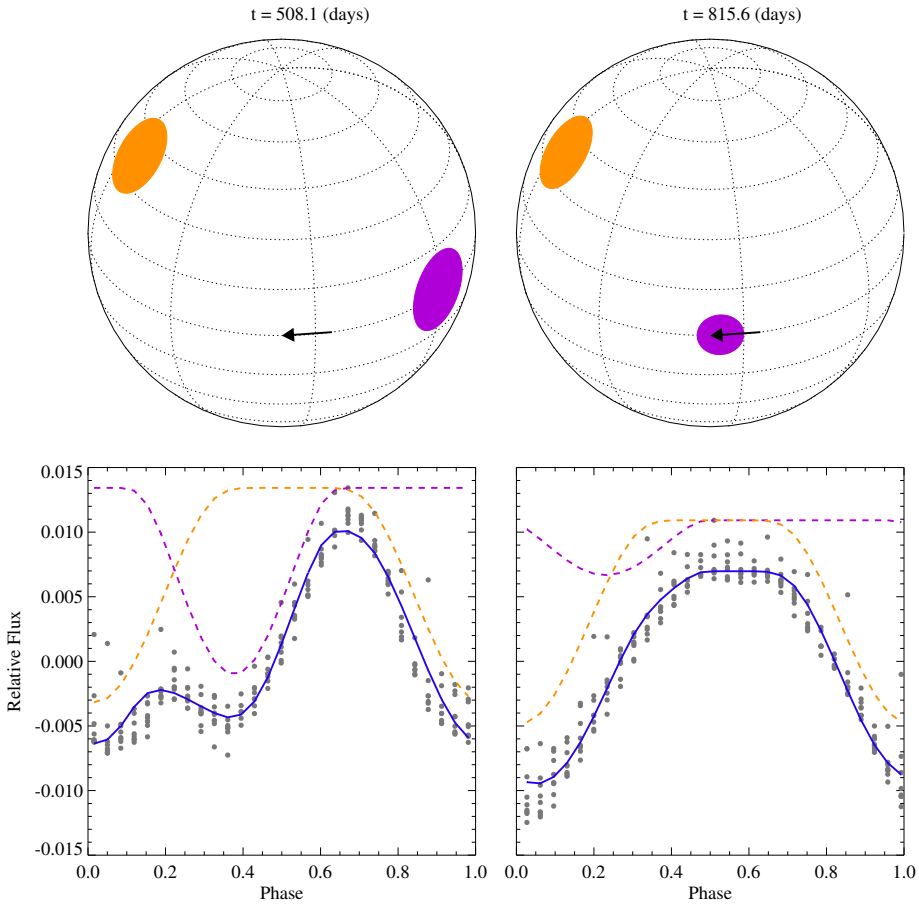

Figure 4.4: Top: orthographic projections of the model star, with an inclination of 32°, and the best-fit positions for two circular spots for the 5 day time window starting at BJD - 2454833.11567 = 508.1 days (left) and 815.6 (right). The direction of stellar rotation is indicated by the black arrow. Bottom: phase-folded light curve for the data in the same 5 day time windows, with the best-fit two-spot models overlaid (blue solid line), and the contributions from both the higher latitude (orange dashed line) and equatorial (purple dashed line) starspots offset for clarity.

4.4 Quantifying the Starspot Evolution

Using the best-fit parameters from each stationary, independent MCMC model, as in Figure 4.4, we were able to trace the sizes and longitudes of two starspots over the entire span of our *Kepler* data. In Figure 4.5 we show the rotational phase (equivalently the longitude facing the observer) for both starspots as a function of time. The higher latitude starspot indicated in Figure 4.4 (orange) is very stable in Figure 4.5 in phase (longitude) with a standard deviation in longitude of only 4.5% (16 degrees) over the four years of data, and traces the dark band seen in Figure 4.3 centered at Phase = 0. The amplitude of this higher latitude spot on the light curve changes slowly over the data, with a standard deviation of 34% in fractional flux in Figure 4.5. We refer to this feature as the "primary starspot".

The "secondary starspot" (purple), however, evolves significantly in phase across the stellar surface over time in Figure 4.5. This feature corresponds to the lower latitude, equatorial starspot in Figure 4.4, and traces the transient secondary features seen in Figure 4.3. Between Time ∼750 and 900, the two best-fit starspot locations were very close in phase, and the variance between solutions in subsequent time steps increased for both the primary and secondary spots. These correspond to time windows where a one-starspot model would be preferred.

We manually identified two regions in Figure 4.5 that displayed nearly constant linear evolution in the secondary starspot longitude: Time = 510–630, and Time = 945–1400. We interpret these to be the signatures of differential rotation, with secular spot motions in time. Within these time windows we used a non-linear least squares first order polynomial fit to measure the linear slopes. Lines of best fit for these two regions are shown in Figure 4.5 as dashed and solid black lines, and had slopes of -0.000927 and -0.000569, respectively. The occurrence of these secondary starspots at multiple times within our data may in fact be due to a single lower latitude feature lapping the primary starspot, but we note the slopes and separations in these features in Figure 4.5 are not consistent with a single spot at a fixed rate of differential rotation.

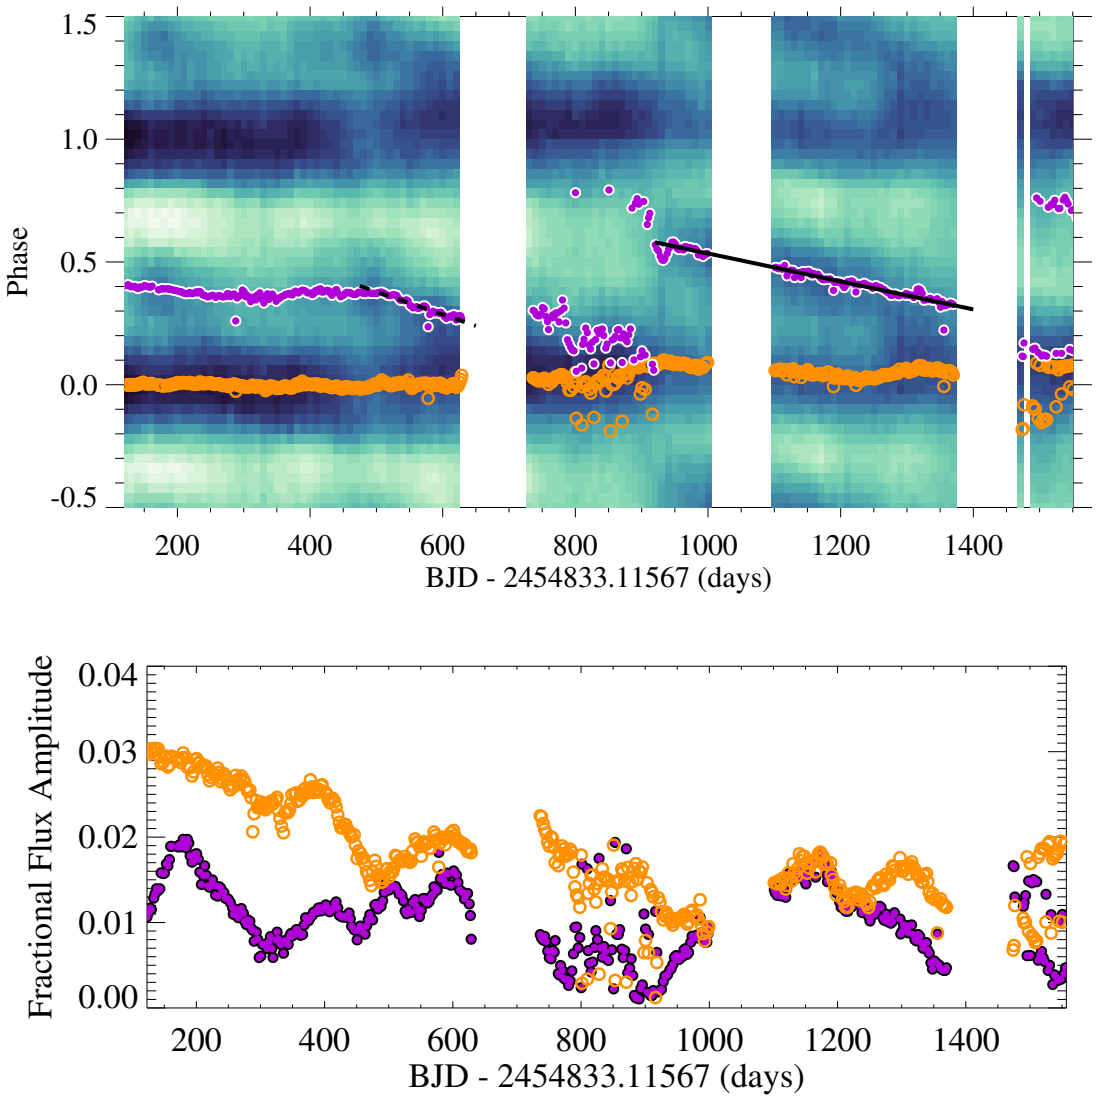

Figure 4.5: Top: Continuous phased light curve map, as in Figure 4.3, with the best-fit solutions from our two spot model overlaid. The higher latitude spot shown in Figure 4.4 (orange open circles) remains nearly constant in phase, while the secondary lower latitude spot (purple filled circles) evolves significantly. Linear fits to the phase evolution for the secondary spot are overlaid (black solid and dashed lines), which we interpret as differential rotation. Bottom: Fractional flux amplitude of each starspot as a function of time for the best-fit solutions from our two spot model. Colors are the same as above.

The measured slopes were in units of phase day^{-1}, and corresponded to a rotation shear of $\Delta\Omega = 2\pi/t_{lap} = 0.0058$ and 0.0036 rad day^{-1}, using the definition from Küker & Rüdiger (2008). Note however this does not include any consideration of the starspot latitudes. Given the slope of a linear feature in the longitude versus time diagram, we can compute the rotation period of a secondary starspot feature. The longitude versus time diagram requires phase-folding the data at a constant primary period (P_0). A flat feature, with slope m= 0, would be constant in phase over time, meaning it rotated at the same rate as the mean period. Features that evolve linearly in phase (or equivalently longitude) with time represent stable secondary periods (P_i in this case), while features that change slope or show non-linear behavior in phase over time would be due to changing their rotation period with time. The slope in this diagram has units of phase per day. The inverse of the slope yields the number of days it takes the a spot rotating at the second period to lap a spot rotating at the mean period. The slope is therefore defined as:

$$m_i = \frac{P_i - P_0}{P_0}\frac{1}{P_i}, \tag{4.1}$$

where P_0 is the mean rotation period that the diagram is phase-folded at, and P_i the rotation period of a secondary spot. The first term corresponds to the fractional difference between the secondary period and the mean rotation period, and is unit-less. Including the second term gives the needed units of time^{-1}. Thus, given the mean rotation period P_0 used to phase-fold the data to make this figure, and the linear slope of an evolving starspot feature m_i, the secondary rotation period can be calculated by

$$\begin{aligned} m_i &= \frac{P_i - P_0}{P_0}\frac{1}{P_i} \\ m_i P_0 &= \frac{P_i - P_0}{P_i} \\ &= 1 - \frac{P_0}{P_i} \\ \frac{P_0}{P_i} &= 1 - m_i P_0 \\ P_i &= \frac{P_0}{1 - m_i P_0} \end{aligned} \tag{4.2}$$

Note by this definition a negative slope yields a shorter rotation period than the phase-folding period.

Differential rotation is generally parameterized (e.g. Henry et al., 1995) as:

$$P_\phi = P_{eq}/(1 - k \sin^2 \phi)\,, \tag{4.3}$$

where P_ϕ is the rotation period at a given latitude (ϕ), P_{eq} is the rotation period at the equator, and $k \equiv \Delta\Omega/\Omega_{eq}$ governs the rate of differential rotation as a function of latitude. Our model results indicate that the period used to phase fold the data in Figure 4.5 corresponds to the higher latitude (38.4°) starspot, and assumes the secondary starspot features are on the stellar equator. Using an average slope from Figure 4.5 of $m = -0.000748$, and the phase-folding period from §4.2.2, we estimated an equatorial rotation period of $P_{eq} = 0.5923336$ days via Equation 4.2. We then solved for the unitless differential rotation parameter using Equation 4.3, finding $k = 0.00114$, which corresponds to $\Delta\Omega = 0.012 \pm 0.002$ rad day^{-1}. The uncertainty we quote here is propagated from the errors in the linear least squares fits in Figure 4.5.

Assuming the primary spot for GJ 1243 is indeed at a higher latitude than the faster rotating secondary spots, this behavior is consistent with Solar-like surface differential rotation where the equator rotates faster than the poles. If the primary and secondary starspots are well separated in latitude as our model indicates, such a low value of shear indicates very weak differential rotation, with the star rotating nearly as a solid body. For comparison, the Sun's surface differential rotation is much stronger, with $\Delta\Omega = 0.055$ rad day^{-1} (Berdyugina, 2005).

4.5 Fitting with Gaussians

In the previous two sections we have focused on measuring starspot evolution using a series of sophisticated stationary models, and finding the differential rotation rate by comparing the position of spots in subsequent model realizations. In this section we explore an alternative method of explicitly determining the starspot time evolution, and thus the differential

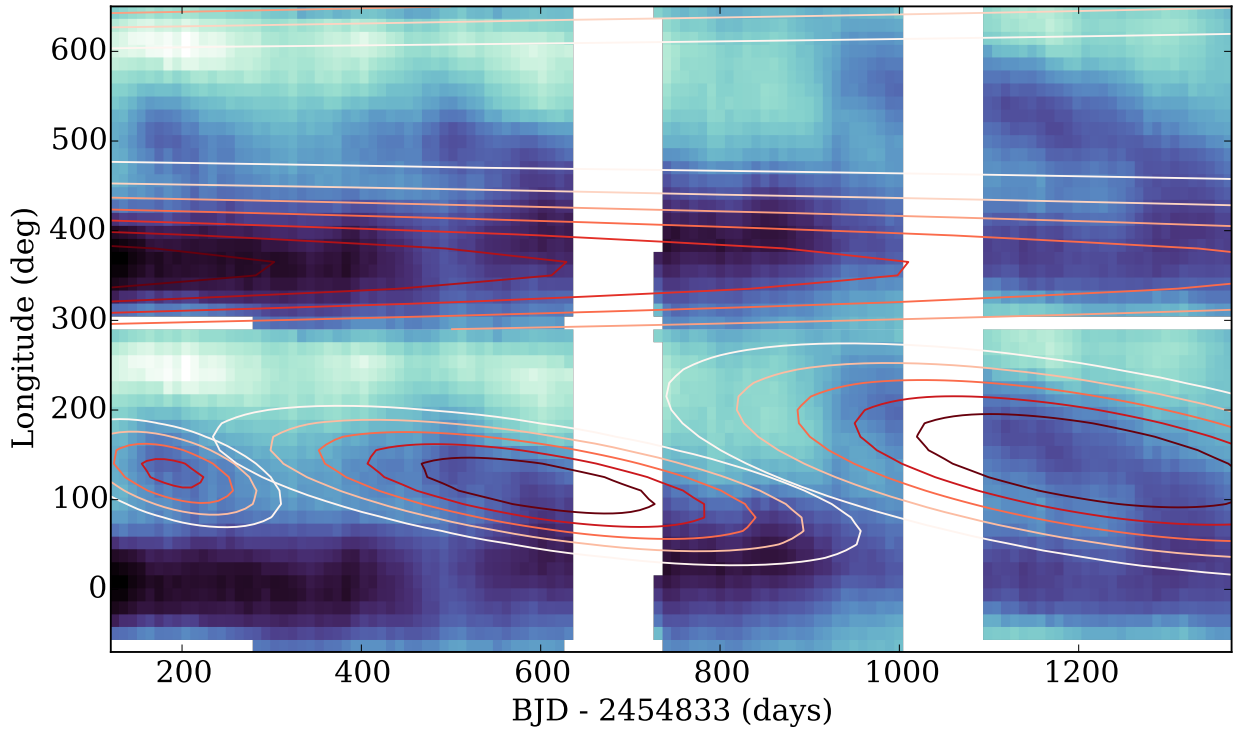

Figure 4.6: Bivariate Gaussian models of starspot evolution in phase and time (open contours), overlaid on the time–phase flux map from Figure 4.3. A total of four Gaussians defined by Eqn 4.4 were fit, representing one primary and three secondary starspot features. For visual clarity we have offset the Gaussian that corresponds to the primary starspot by 1 phase.

rotation rate, using Gaussian functions.

Rather than modeling the entire light curve directly to infer the starspot sizes and positions, as in Figure 4.4, we analyzed the three-dimensional "surface" shown in Figures 4.3 and 4.5, which traces the flux as a function of both time and rotational phase. The data were binned in both time and phase (longitude), using bin sizes of 10 days and 14.4 degrees, respectively. To model this flux map we used 2-dimensional bivariate Gaussian distributions of the form:

$$\begin{aligned} t' &= (t - t_0)\cos\theta - (l - l_0)\sin\theta \\ l' &= (t - t_0)\sin\theta - (l - l_0)\cos\theta \\ F(t, l) &= A \exp\left(-\frac{(t'/\tau)^2 + (l'/b)^2}{2}\right) \end{aligned} \tag{4.4}$$

where $F(t, l)$ is the flux as a function of time t and longitude l, A is the flux amplitude of the starspot, τ is the lifetime of the starspot, b is the scale width of the large starspot in longitude, t_0 is the center time of the starspot, l_0 is the center longitude of the starspot, and θ is the slope of the spot evolution in units of degrees day^{-1}. Here the longitude is a circular coordinate, with range between 0 and 360 degrees, and defined to continuously wrap from 360 back to 0 degrees. This definition enables long-lived starspots with large rates of differentially rotation to "lap" the stellar surface multiple times. Each starspot's evolution is defined by evaluating Equation 4.4 over the entire time span of our data, and the full range in longitude. The entire flux map in Figure 4.4 is reproduced by summing the 2-D Gaussian functions.

For our Gaussian analysis we used the data spanning from the beginning of the available *Kepler* data (Q0) through Quarter 14 (Time $\sim$ 1370). This time range was chosen to focus our analysis on the secondary spot evolution and differential rotation signal measured above. We discarded data from Q16 and Q17 data which showed no sign of the secondary starspot.

We solved for the positions and evolution of four starspots over the entire duration of the data, using the Python MCMC sampler `emcee` from Foreman-Mackey et al. (2013) to

explore parameter space for all four simultaneously. In this four Gaussian model, we consider the largest spot (also with the smallest slope in θ) as the primary, and the three secondary spots as independent spot features, or repeat occurrences of the secondary spot discussed before. A third occurrence of the secondary spot feature was needed to account for the small feature seen around Phase$\sim$0.4 at Time$\sim$100 days in Figure 4.3, which was not chosen in our conservative by-eye selection above. We used rough values for the parameters to seed `emcee` with. The initial seeds for the primary spot were $\theta = 0$, $t_0 = 800$, $\tau = 1300$ days, $l_0 = 75$ degrees, and $b = 15$. The three secondary spots were all seeded with θ = -0.2 degree day^{-1}, $\tau = 100$ days, $l_0 = 200$ degrees. The secondary starspot center times seeds were set to $t_0 = 200$, 530, and 1200.

We then ran `emcee` with 50 walkers for 2000 steps. The best fit model from this parameter space search for the three secondary starspots is shown in Figure 4.6. The best fit slope for the primary spot was θ = -0.0011 deg day^{-1}, and for the secondary spots (in time order) was θ = -0.20, -0.14, and -0.15 deg day^{-1}. These secondary spot shear rates corresponded to $2\pi/t_{lap} = \Delta\Omega = 0.0036$, 0.0025, and 0.0027 rad day^{-1}, or an average of 0.0029 rad day^{-1}, somewhat lower amplitude than measured from the linear features in Figure 4.5. Note again this method does not constrain the latitudes of the starspots, and so the measurement of shear is only a lower limit on the true differential rotation rate.

This approach assumes a priori that the starspot evolution in both time and longitude can be represented by a Gaussian function, meaning the spots may only evolve linearly in longitude over time. We note the resulting estimate for the differential rotation shear rates for the secondary starspots are very similar to the values determined when fitting many time-stationary MCMC instances. The Gaussian modeling approach used flux data that was binned in time and phase, greatly reducing the number of data points to be fit. This entire MCMC analysis took only a few minutes to compute using a standard Linux workstation. We thus propose this to be an efficient means of estimating the starspot lifetimes and differential rotation rates in certain cases, which we discuss further in Section 4.6 below.

4.6 Discussion

We have produced two significant results in this work. The first is the identification of a long-lived starspot, which we attribute to a higher latitude starspot (possibly due to a spot cap or group) on the rapidly rotating M dwarf, GJ 1243. The second is a robust measurement of weak differential rotation for this star due to a spot closer to the stellar equator. In this section we provide additional context and discussion of these results, and their implications for the magnetic field's surface topology.

4.6.1 A Long Lived Starspot

To further illustrate the remarkable stability of the higher latitude spot on GJ 1243, we retrieved ground-based light curves from SuperWASP which predated the *Kepler* mission by $\sim$2 years (Butters et al., 2010). This SuperWASP public archive photometry was phased using the period and ephemeris we determined from our *Kepler* light curve. The phase of flux minimum matches between these two datasets to within 1%, indicating this large starspot has been stable in longitude for more than 6 years. The amplitudes of the flux modulations between *Kepler* and SuperWASP are only slightly different, with the median *Kepler* variation of 2.19% (averaged over all 4 years of data), and for SuperWASP of 2.86%. The SuperWASP data was taken in the V-band, which is more narrow than the very wide *Kepler* filter. The V-band also is centered at a shorter wavelength than the *Kepler* filter, which is weighted more towards the R-band. As a result, we would expect to find larger flux contrast between cool spots and the stellar photosphere in the V-band than compared to the *Kepler* filter. However, since these observations were not concurrent we cannot rule out small differences in the starspot's physical size over time.

4.6.2 Differential Rotation in Cool Stars

The average starspot shear for GJ 1243 observed in this chapter of 0.0047 rad day^{-1} corresponds to a differential rotation rate of $\Delta\Omega = 0.012 \pm 0.002$ rad day^{-1} (assuming the spot

configurations used in our models), and is one of only a few such measurements yet obtained for low mass, rapidly rotating, fully convective stars. In Figure 4.7, we place this measurement in the context of other existing observations of stellar differential rotation (Barnes et al., 2005; Morin et al., 2008b) along with the empirical extrapolation to cool stars from Reiners (2006); Collier Cameron (2007), and models from Küker & Rüdiger (2011). One of the few other objects with a robust differential rotation measurement in this regime is the cool, rapidly rotating star, V374 Peg (Morin et al., 2008b). These authors employ Doppler Imaging, a completely different technique to our own, to derive a value for the surface differential rotation of V374 Peg ($\Delta\Omega = 0.0063 \pm 0.0004$ rad day^{-1}) that is similar to that of GJ 1243 we have measured. This Doppler Imaging method assumes a Solar-like differential rotation profile as in Equation 4.3, and simultaneously fits for the starspot positions, sizes, and shear rates. Lurie et al. (2015) have also estimated the starspot shear for both rapidly rotating components of the M5+M5 binary GJ 1245AB, using *Kepler* data and the 2D Gaussian modeling approach detailed in our Section 4.5, finding shear rates that are comparable to GJ 1243 and V374 Peg.

The various methods for constraining differential rotation each have unique limitations in their sensitivity and degeneracies. Photometric phase-tracking methods such as ours, frequency splitting approaches (Reinhold et al., 2013), as well as spectroscopic line broadening techniques (i.e. the Fourier Transform Method; Reiners & Schmitt, 2003) have been considered as only providing lower limits on the true differential rotation rate, due to a lack of constraint on the starspot latitudes (Barnes et al., 2005). Statistical corrections for this latitudinal uncertainty have been developed (Hall & Henry, 1994), and our models of GJ 1243 provide a weak constraint on the spot latitudes. Together, these varied observations indicate that rapidly rotating, mid M-dwarf stars exhibit differential rotation that is significantly weaker than on the Sun (0.055 rad day^{-1}) by up to an order of magnitude.

These observational results of low $\Delta\Omega$ values support recent theoretical work in this area. Mean field theory models predict that for stars at a fixed temperature surface differential rotation rates decrease with faster rotation (shorter periods), down to periods of a few

days (Küker & Rüdiger, 2005). In addition, the total amplitude of the differential rotation decreases with decreasing temperature, as models indicate lower mass main sequence stars ($T_{eff} < 6000$ K) should exhibit lower amounts of surface differential rotation than hotter stars at a fixed rotation period (Küker & Rüdiger, 2011).

Dipolar dynamos with strong magnetic field and quenched differential rotation have been reported in global dynamo models of rapidly rotating low-mass stars (e.g. Gastine et al., 2013). Furthermore, recent global dynamo modeling efforts have also been successful in producing polar starspots self consistently (Yadav et al., 2015) Furthermore, in hydrodynamic simulations of rotating, non-magnetic, solar-type stars, Browning (2008) found that convection in the outer zone of the star redistributes angular momentum quickly giving rise to solar-like differential rotation. However, when strong magnetic fields are introduced, the field lines act to reduce the differential rotation by linking together individual regions of the stellar interior. As the strength of the magnetic field is increased, as is typically seen for more rapidly rotating stars, the differential rotation is suppressed to almost negligible values.

The observations from Lurie et al. (2015) also support this connection between rotation, magnetic field strength, and suppressed differential rotation. For the M5+M5 binary system GJ 1245AB, the more rapidly rotating component (GJ 1245A, P=0.26 days) has a slightly higher total chromospheric Hα emission flux, and significantly less phase evolution of its starspot modulations compared with GJ 1245B, (P=0.71 days).

4.6.3 Magnetic Field Topology

Dynamo models and observations of stars like GJ 1243 are in agreement that for rapidly rotating, low-mass stars, the magnetic field strength is increased and surface differential rotation is suppressed. The detailed surface topology of the magnetic field is less certain. The convective dynamo models from Browning (2008), for example, indicate that rapid rotation produces strong axisymmetric magnetic fields. Earlier models of convective envelopes (Dobler et al., 2006) also show net axisymmetric magnetic fields. Other models of fully convective stars, however, have produced fully non-axisymmetric fields at the surface (Küker & Rüdiger,

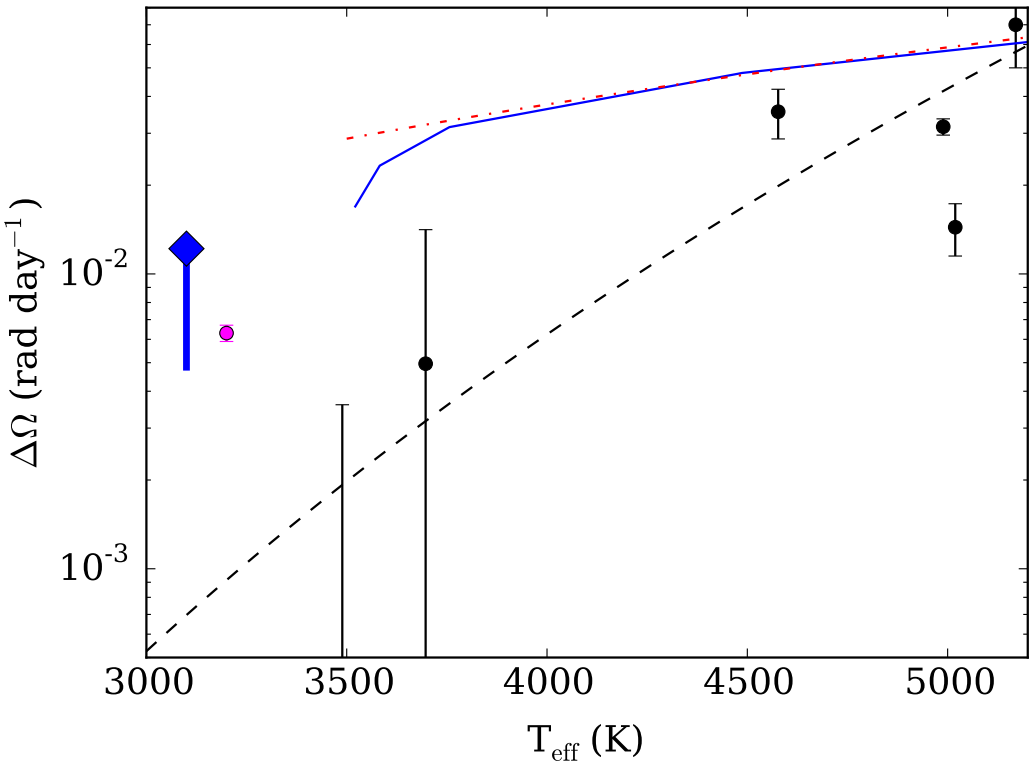

Figure 4.7: Average starspot shear from the two linear fits to the MCMC light curve models in Figure 4.5, assuming the primary spot was at a latitude of 38.4° and secondary spot at the equator (blue diamond). The blue bar extends to the minimum possible amplitude of differential rotation for GJ 1243, assuming the primary and secondary spots are at the pole and equator, respectively. For comparison, the Collier Cameron (2007) observed fit for cool stars (black dashed line), theoretical prediction from Küker & Rüdiger (2011) (red dot-dash and blue solid lines), individual stars from Barnes et al. (2005) (black circles), and the estimated shear rate for V374 Peg determined using doppler imaging from Morin et al. (2008b) (purple filled circle) are shown.

2005; Chabrier & Küker, 2006).

The smooth rotational flux modulation observed is due to a local feature on the stellar surface. The magnetic activity from this star as traced by flares in *Kepler* data has also been well studied (Hawley et al., 2014; Davenport et al., 2014). No correlation between rotational phase (or equivalently longitude) and the occurrence rate or energy emitted from flares has been found. The star may be uniformly covered by many smaller active regions and spots that would not create observed modulations in the light curve. The flares would be a result of these small scale multipolar magnetic field structures. Similarly, the stellar pole may be entirely covered by a polar spot "cap", resulting from strong poloidal magnetic flux geometry due to the rapid stellar rotation. This poloidal component of the magnetic field could also be slightly misaligned from the stellar rotation axis, as seen in other convective stellar and planetary dynamos (e.g. Christensen et al., 2009; Hull et al., 2013), resulting in the observed light curve asymmetries.

As we noted in Section 2, the Rossby number for GJ 1243 is $R_o \approx 0.008$. According to the ZDI observations of M dwarfs aggregated in Gastine et al. (2013), for stars with $R_o < 0.1$ both dipole and multipole fields are possible. In this low Rossby number regime, they also find surface differential rotation should be stronger when the magnetic field is multipolar. However, we find for GJ 1243 a very low rate of differential rotation, and large, long-lived starspot modulations. As a result, in this context we predict a highly organized and stable dipolar magnetic field geometry. This is in agreement with the ZDI observations of the similar star V374 Peg, which has a large scale dipolar field and long-lived starspots (Morin et al., 2008b).

4.7 Summary

In this chapter, we have presented a classic approach of phase tracking starspots in light curves, made new by the exquisite photometric monitoring from *Kepler*. By tracing the phase evolution for two starspot regions on GJ 1243 we have found the smallest amplitude of differential rotation rate ever robustly measured for a cool star. This phase-tracking

technique is similar to Henry et al. (1995), and is sensitive to a comparable amplitude differential rotation signal to Morin et al. (2008b). The large starspots, or starspot groups, on GJ 1243 are very long-lived, with the primary high-latitude spot found to be constantly aligned in phase for over 6 years. The secondary starspot features evolve on timescales of hundreds of days in both phase and amplitude.

There remain many challenges in modeling the starspots using broadband light curves alone. For example, we have almost no constraint on the actual latitudes of the spots. Modeling starspots on stars with transiting exoplanets (e.g. Sanchis-Ojeda et al., 2013) may help break many of these degeneracies. As most transiting systems in the *Kepler* data are around G dwarf stars, more M dwarf systems like *Kepler* 186 (Quintana et al., 2014) are needed to better understand the detailed starspot characteristics of stars across the main sequence.

Both the light curve modeling MCMC and 2D Gaussian phase-tracking techniques used to measure surface differential rotation in this work are best suited for tracking long-lived spots on rapidly rotating stars, as in the GJ 1243 system where we are able to average over hundreds of stellar rotation periods during a starspot's lifetime. We believe the light curve modeling approach provides the most robust estimates for starspot properties, but note the 2D Gaussian approach is orders of magnitude faster to execute. This methodology could be applied to hundreds of rapidly rotating active stars in the *Kepler* dataset for which rotation periods are already known (e.g. see Reinhold et al., 2013; McQuillan et al., 2014). There are $\sim$20 other stars in the *Kepler* data with M dwarf colors and estimated rotation periods shorter than 1 day (McQuillan et al., 2013). A cursory look at these light curves reveals many with dramatic flare activity and sinusoidal starspot modulations, as found on GJ 1243 and GJ 1245AB. The phase versus time diagrams (as in our Figure 4.5 for GJ 1243) for these other rapidly rotating stars show a diverse set of morphologies, ranging from even more stable spots than on GJ 1243, to stars with faster shear rates and shorter spot lifetimes.

Finally, we have introduced an efficient technique for empirically tracking starspot evolution in the phase versus time flux map, by fitting bivariate Gaussians to model the spot

motion and evolution. Studying the phase-evolution of the starspots with *Kepler* for single field stars appears to be feasible for stars with fast rotation periods and long spot lifetimes. The Gaussian-fitting method presented here has already been applied to the *Kepler* data for the active M5+M5 binary system, GJ 1245 AB (Lurie et al., 2015). We have pointed out many other rapidly rotating low-mass stars in the *Kepler* archive that may be studied with this technique, and hope this work will be the beginning of a larger observational understanding of surface differential rotation in cool stars.

4.8 Acknowledgments

The authors wish to thank the anonymous referee, whose helpful comments greatly improved the quality of this manuscript. This work was supported by *Kepler* Cycle 2 GO grant NNX11AB71G and Cycle 3 GO grant NNX12AC79G and NSF grant AST13-11678. We are pleased to thank John Gizis for advance use of his GO program's long-cadence data on GJ 1243. We thank John P. Wisniewski and Adam F. Kowalski for continuing discussions on stellar activity, and Rakesh Yadav for very useful comments on an early draft of this manuscript. Davenport thanks Morgan Fouesneau for help with Python code.

Chapter 5

MODELING STARSPOT EVOLUTION ON THE TRANSITING EXOPLANET HOST STAR, KEPLER 17

In this chapter I extend the analysis of starspots to *Kepler* systems with transiting exoplanets, focusing on one simulated light curve and a real *Kepler* system. The addition of a planet in the analysis greatly increases the complexity of the light curve modeling, but simultaneously provides much greater constraints on the starspot properties. When the planet passes in front of the star, deviations from the uniform brightness stellar surface due to starspots are detected as "bumps" in the transit light curve. The planet then tomographically maps out the contrast of the stellar surface along the projected transit path.

I use this spot-crossing technique, along with an improved version of our starspot modeling code, to study the exoplanet host star Kepler 17. I test our code's ability to recover starspot properties and evolution using an independent model of an active star with a transiting exoplanet provided by Dr. Joe Llama, based on the models in Llama et al. (2012). This model light curve is processed in the same manner as real *Kepler* data with our code, splitting the light curve into windows of time and assuming static starspot solutions within each window. The synthetic data are used to develop our procedure for interpreting the outputs of our code, and to validate our ability to recover the stellar differential rotation rate and track the sizes and position of starspots crossed by the planet. I then extend this procedure to study Kepler 17, estimating it's differential rotation rate and spot diffusion timescale.

5.1 Using Transiting Exoplanets to Study Starspots

In the previous chapter I introduced our method for studying starspots and their evolution in *Kepler* data using a light curve modeling program and MCMC sampling to explore the starspot parameter space. I also presented a new Gaussian fitting approach to estimate starspot lifetimes and longitudinal evolution that is applicable to data from large numbers of stars. With both of these methods we have (effectively) limited our analysis to studying the radii and longitudes of spots, due to the strong degeneracies in determining latitudes. Moreover, both of these methods have been limited to studying at most two spots within a given portion of the light curve.

There has been great debate on the uniqueness and accuracy of constraining starspot models using broadband light curves, as well as the many "light curve inversion" methods. Some authors believed the degeneracies were too fundamental to break, and that most attempts using light curves alone were destined to fail or worse: over interpretation (Kopal, 1982; Hall, 1994). For a single spot, the position (both latitude and longitude) and spot temperature can theoretically be recovered using light curves alone given high signal to noise data (Eker, 1996). The problem becomes much more difficult and degenerate with each additional spot added to the model, though Eker (1996) claim it is still feasible to recover the starspot properties given data with sufficiently high fidelity.

Starspot temperatures (or equivalently flux contrast values) are difficult to measure, and there is a natural degeneracy with the filling factor (spot size relative to the observable stellar surface). Typically the best method to estimate starspot temperatures on low-mass stars has been through spectroscopic indicators, specifically molecular bands such as OH and TiO (O'Neal et al., 2001, 2004). These starspots are often constrained for eclipsing short-period binary stars, where photometric color can also be used to estimate spot temperatures (O'Neal et al., 2004; Frasca et al., 2008). Molecular band studies find a large range of spot contrasts for stars at a given temperature, and model fits indicate very large spot filling factors up to ∼30–50% (Berdyugina, 2005). Doppler Imaging (DI) methods seem to produce

smaller spot filling factors, which may be due to their dependence on rotational modulation of spectroscopic features versus the molecular band strength measurements. However, DI shows a remarkable ability to recover surface structure. For example, the seminal work of Vogt et al. (1987) demonstrated the reconstruction power of DI using a starspot pattern that traces the author's name, "VOGT", across the surface. Using DI methods on the simulated data (and assuming a scenario with extremely high signal to noise data), Vogt et al. (1987) were able to recover the spot distribution, clearly discerning each "letter" in the spot map.

Starspot lifetimes are dependent on their size, as well as the amplitude of the stellar differential rotation (Hall & Henry, 1994). On the Sun, lifetimes depend on the spot latitude and radius. Sunspots appear to roughly obey a single diffusive decay profile, where the decay rate is linearly correlated with the spot area (Hathaway & Choudhary, 2008). Sunspots at higher latitudes also appear to decay faster (Hathaway & Choudhary, 2008; Javaraiah, 2011), which may be due to a latitude dependence on differential rotation and meridional flow at the stellar surface (Javaraiah, 2015). Studies that track the light curve phase minimum over time, as in Chapter 4 or Henry et al. (1995), are limited in their ability to discern the decay and lifetimes of individual starspots versus evolution and merging of multiple surface spots. ZDI (and DI) can map the positions of magnetic features or spots very precisely, but the difficult spectroscopic observations mean these data do not have high time fidelity, often only mapping the stellar surface once every year.

As we have seen in Chapter 4, and has been discussed by many authors (e.g. Vogt, 1981), the most difficult aspect in modeling starspots for single stars with broadband photometry comes in constraining the starspot latitudes. In the era of ground-based, modest signal to noise stellar data, spectroscopic techniques like DI and ZDI provided the best methods for constraining the starspot distributions, but they lack the temporal resolution to measure starspot evolution. Even for light curves with high signal to noise, there remains a degeneracy in starspot latitudes with both fundamental spot properties of temperature and filling factor. The ability to constrain the latitudes, however, is key to accurately measuring the differential rotation profile of the star.

Space-based photometric monitoring can now beat many of the limitations on estimating starspot properties. For spotted stars with transiting exoplanets, we may use the exoplanet itself to measure smaller scale starspot structures on the star's surface. This method was first detailed by Silva (2003), who used images of the Sun with a cool starspot + planet model to interpret transit data from HD 209458. Significant transit anomalies were detected in space-based photometry of TrES-1, which Rabus et al. (2009) have also attributed to cool starspots. Recently, spot-crossings by Kepler 63b were used to validate the high degree of obliquity, and detect a bona fide polar starspot structure on the surface of Kepler 63 (Sanchis-Ojeda et al., 2013).

In the simplest case, as the planet passes in front of its parent star, the amount of light blocked (or the transit depth) depends solely on the ratio of the planet to star radius (r_p/r_s), assuming a totally opaque planet that crosses the projected center of the star. The detailed transit ingress and egress shape also depends on the stellar limb darkening and surface effective temperature, while the transit duration (relative to the orbital period) depends on the ratio of the stellar radius to the orbital semi-major axis. Given physical parameters of the star–planet system, the transit light curve can be approximated to a high level of precision using the equations of Mandel & Agol (2002).

If the star is not uniform in brightness, due to cool starspots on the surface, the transit light curve shape will deviate from this ideal model. During the transit, as the planet passes in front of the cool starspot, it will block *less* light than if it were passing over just the ambient stellar surface. As a result the transit depth is temporarily shallower, and we observe a characteristic "bump" in the transit light curve during the spot-crossing event. The height of this bump is directly related to the temperature contrast between the starspot and the surrounding photosphere. This is demonstrated in Figure 5.1, using light curves models of a single temperature (grey) circular starspot. As the planet passes over the starspot region, the temperature contrast is mapped out. The saturated, flat-topped bump seen in this model is due to this large spot having a single temperature contrast, unlike sunspots that show umbral and penumbral structure as described in Chapter 1.2.2. The model in Figure 5.1 also

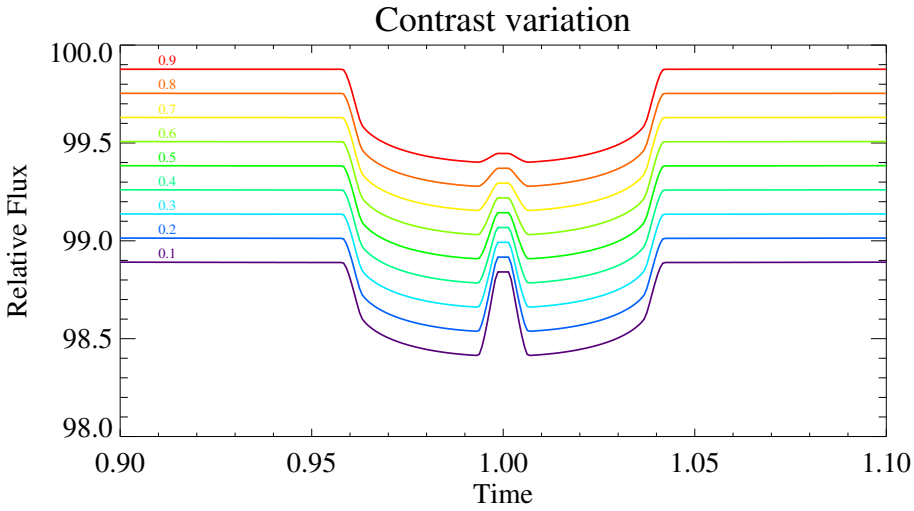

Figure 5.1: Results from modeling a single starspot at the stellar equator (latitude $\phi = 0°$) with a range of flux contrast values (indicated at left). Darker starspots correspond to smaller values of flux contrast. The height of the resulting bump is directly correlated to this flux contrast. The out-of-transit flux level also varies significantly. Note the flat-topped bump is due to our starspot model having a single contrast.

clearly shows the spot crossing ingress and egress structures, at the beginning and end of the occultation. These are due to the starspot being larger in projected size than the transiting planet, and are analogous to (though inverted) the star–planet transit ingress and egress. Importantly, by measuring the maximum amplitude of spot-crossing bumps, we provide a direct constraint on the temperature contrast of starspots that is *model independent*.

In Figures 5.2 and 5.3, I also demonstrate the effect of varying the spot radius and latitude on the spot-crossing bumps. As the spot radius increases in Figure 5.2 the total bump duration also increases, while the bump ingress and egress durations are unchanged. The out-of-transit flux near the transit is naturally affected by the increased starspot filling factor as well. If the starspots are comparably sized to, or smaller than the transiting planet the bumps will not "saturate", as the ingress and egress dominate the light curve. This

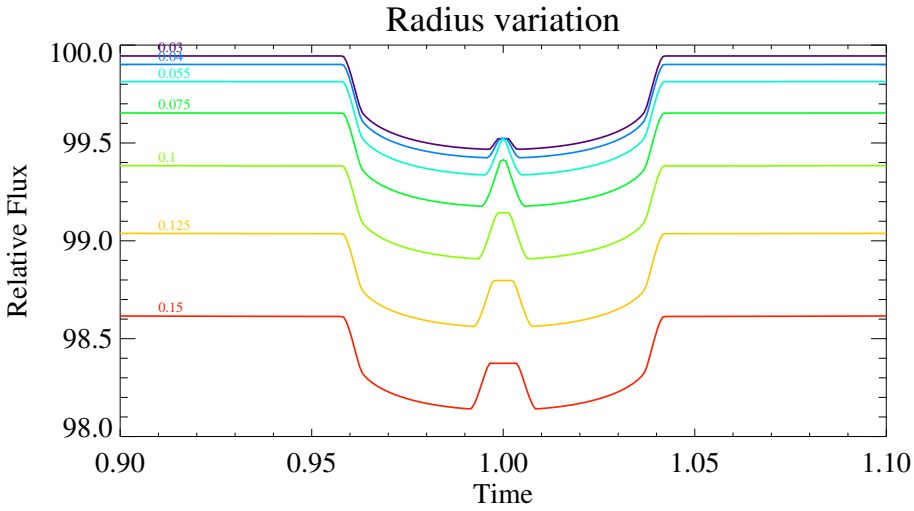

Figure 5.2: Results from modeling a single starspot at the stellar equator with fixed contrast, but varying the spot radius from 3% (dark purple curve) to 15% (red curve) of the stellar radius.

creates a more "V"-shaped bump, and is analogous to stellar binary eclipses in morphology. The relative size of the exoplanet (r_p/r_s) places a lower limit on the filling factors that can be reliably measured for starspots.

The effect in Figure 5.3 of only changing the spot latitude is more complex on the bump morphology. When the spot is moved to high enough latitude that it does not fully intersect the planet, in other words a grazing spot occultation, the bump becomes rounded. This shape is more similar to the bumps observed in real data, as described in later sections. The amount of light blocked by the spot out-of-transit along the line of sight to the observer *decreases* with increasing spot latitude due to the projection of the circular spot on a sphere.

While light curve studies using photometry alone have previously been limited to estimating properties of one or two spots, the spot-crossing transit method can definitively show the presence of many spots along the transit path. The maximum number of spots that the transit crossing method can recover is also dependent on the ratio of the observing

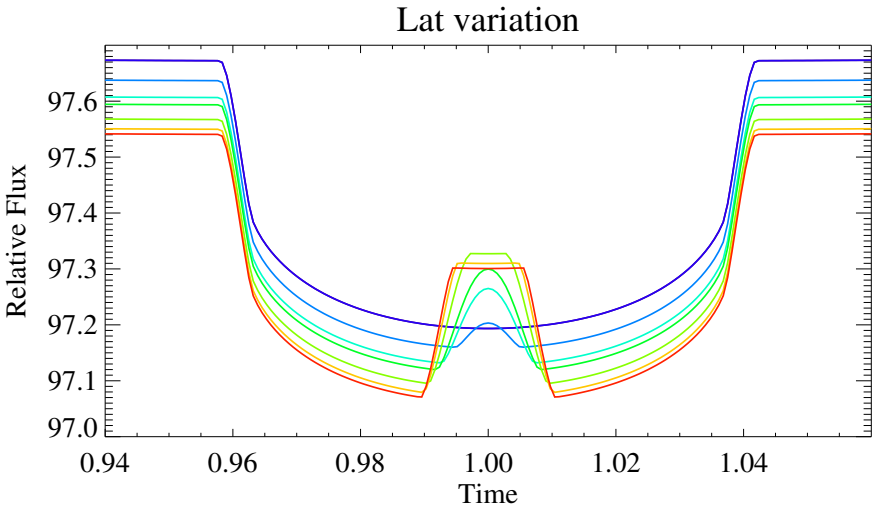

Figure 5.3: Results from modeling a single, fixed contrast and radius starspot, but varying the latitude from the stellar equator $\phi = 0°$ (red curve) to $\phi = 75°$ (dark blue curve). The star here is modeled with an inclination angle of 0°, with the equator face-on. The out-of-transit flux level changes monotonically due to spherical projection.

cadence to the transit duration, which along with the planet radius governs the filling factor that is resolvable. The number of transits per stellar rotation determines the fraction of the surface that is mapped by the planet. For stars with high cadence observations (1-minute from *Kepler*) and many transits per stellar rotation, our tests are able to correctly recover the presence of at least 10 spots along the path of the planet. By modeling both the in-transit and out-of-transit observations, one or two additional starspot groups can be inferred out of the path of the transit to account for the observed flux modulation. These additional starspots have approximately the same degeneracies in their recovered properties as the spots on stars without transits.

5.2 Modeling the Light Curves

We use the same approach as in Chapter 4 for GJ 1243 to estimate the starspot properties from *Kepler* light curves with transits. Our light curve synthesis code has been enhanced to include a planet orbiting the star with tunable size and orbital properties. The orbital and geometric elements of the star–planet system are passed as fixed arguments to the code via an input parameter file. These include the stellar rotation period, stellar inclination, planet orbital period, transit ephemeris, planet radius, impact parameter, and the obliquity (angle between the planetary orbit and stellar rotation axis). These values are well constrained by the exoplanet community for *Kepler* target planets, and with the exception of planet radius (r_p/r_s) we typically take them as given in the literature. We also used the published values for the stellar temperature, mean density, and (4th-order polynomial) limb darkening coefficients, which significantly affect the transit shape.

Our light curve modeling code assumes that the starspots are static on the stellar surface, and models both the in-transit and out-of-transit simultaneously. To capture spot evolution and motion we chop the light curve into windows of time, fitting each window individually with an MCMC sampler. For GJ 1243, where a small number of starspots or spot groups appeared to change very slowly, we used relatively large windows of time relative to the fast stellar rotation period. In this chapter we focus on stars with transiting exoplanets, where the

stellar rotation period of several days is many times longer than the planet's orbital period. This allows the planet to potentially cross the same spot over multiple transits, giving much stronger constraints on the spot positions and sizes. For these transiting systems we binned the light curve into time windows that were 95% the stellar rotation period. Removing the trailing 5% of the star's rotation did not negatively affect the model results, but instead allowed a tolerance for the edges of each window's light curve to not match perfectly. Such mismatching at the edges is noticeable if the starspot evolution timescale is not much larger than the rotation period.

As with the study of GJ 1243, the number of starspots used to fit the light curve in each time window much be selected. We experimentally selected this number for each star studied, and for this study we did not vary it between time windows. The preferred number of spots within each time window can be determined programmatically by doing many independent model fits to each time window using different numbers of starspots, for example N=[4,5,6,7,8,9] spots. Within each time window a model-comparison technique like the Bayesian Information Criteria is then employed to select between the independent results while accounting for the extra degrees of freedom with increasing numbers of spots. Such a procedure greatly increases the computation time required for our analysis. However, for this work we carry out model runs with a fixed number of spots within each time window, and instead focus on our ability to trace their evolution between subsequent time windows.

5.3 Testing our Methodology with Simulated Data

To ensure the final starspot parameters we determine for Kepler 17 were reliable, we first developed our analysis technique and intuition using simulated data. We performed this exercise in a "blind" fashion, only knowing the minimal amount of information that would normally be determined by the exoplanet community for *Kepler* targets. This exercise proved to be very instructive. To approximately match the configuration of Kepler 17 we specified this simulated target should represent a typical G star rotating faster than the Sun, the star should be more active and have larger starspots than the Sun, possess Solar-like differential

rotation (but no meridional starspot flow), host an opaque hot Jupiter that transited many times per stellar rotation, that the planet should cross along the equator, and that the transits did show spot-crossing features. The data was also simulated at a high cadence similar to the *Kepler* data, and spaned many years of observation.

The simulated light curve was kindly provided to us by Dr. Joe Llama, using the light curve forward modeling code detailed in (Llama et al., 2012). These models are capable of simulating light curves from Solar-like activity cycles, which includes the evolution of starspot emergence latitudes, surface differential rotation, and spot evolution and decay. This simulated light curve generation code was also recently used in a multi-team differential rotation blind-recovery experiment (Aigrain et al., 2015). Dr. Llama provided us with 4 years of simulated data, sampled at a 5 minute cadence. The simulation was stated to have a stellar equatorial rotation period of 10 days, a planet orbital period of 2 days, and a relatively large planet size of $r_p/r_s = 0.1$. Photometric uncertainties were assigned by drawing from a Gaussian distribution with a width of 0.002 fractional flux units.

While the Llama model was quoted as having an equatorial rotation period of 10 days, it is worth noting that a Lomb Scargle periodogram of this simulation light curve yielded a best-guess rotation period of about 9.8 days. A shorter rotation period would be physically impossible given the known equatorial rotation period and the constraints of Solar-like differential rotation imposed. This spurious period recovery is due to starspot evolution (primarily rapid spot emergence), as discussed in the starspot and differential modeling exercise carried out by Aigrain et al. (2015). For this exercise we take the equatorial rotation period of 10 days as a given, though the final results were found to be insensitive to this choice.

The spot-crossing bumps occur for only a portion of the transit, and the transit duration in turn is only a small fraction of the total rotation period. Bumps are often only a few times larger than the typical photometric uncertainty, even for very precise *Kepler* data. As such, the bumps carry very little weight in a traditional χ^2 metric, since they represent very few data points. To minimize this effect we down-sampled the Llama model out-of-transit data by a factor of 10 (or 50 minute cadence). The full simulation cadence of 5 minutes was used

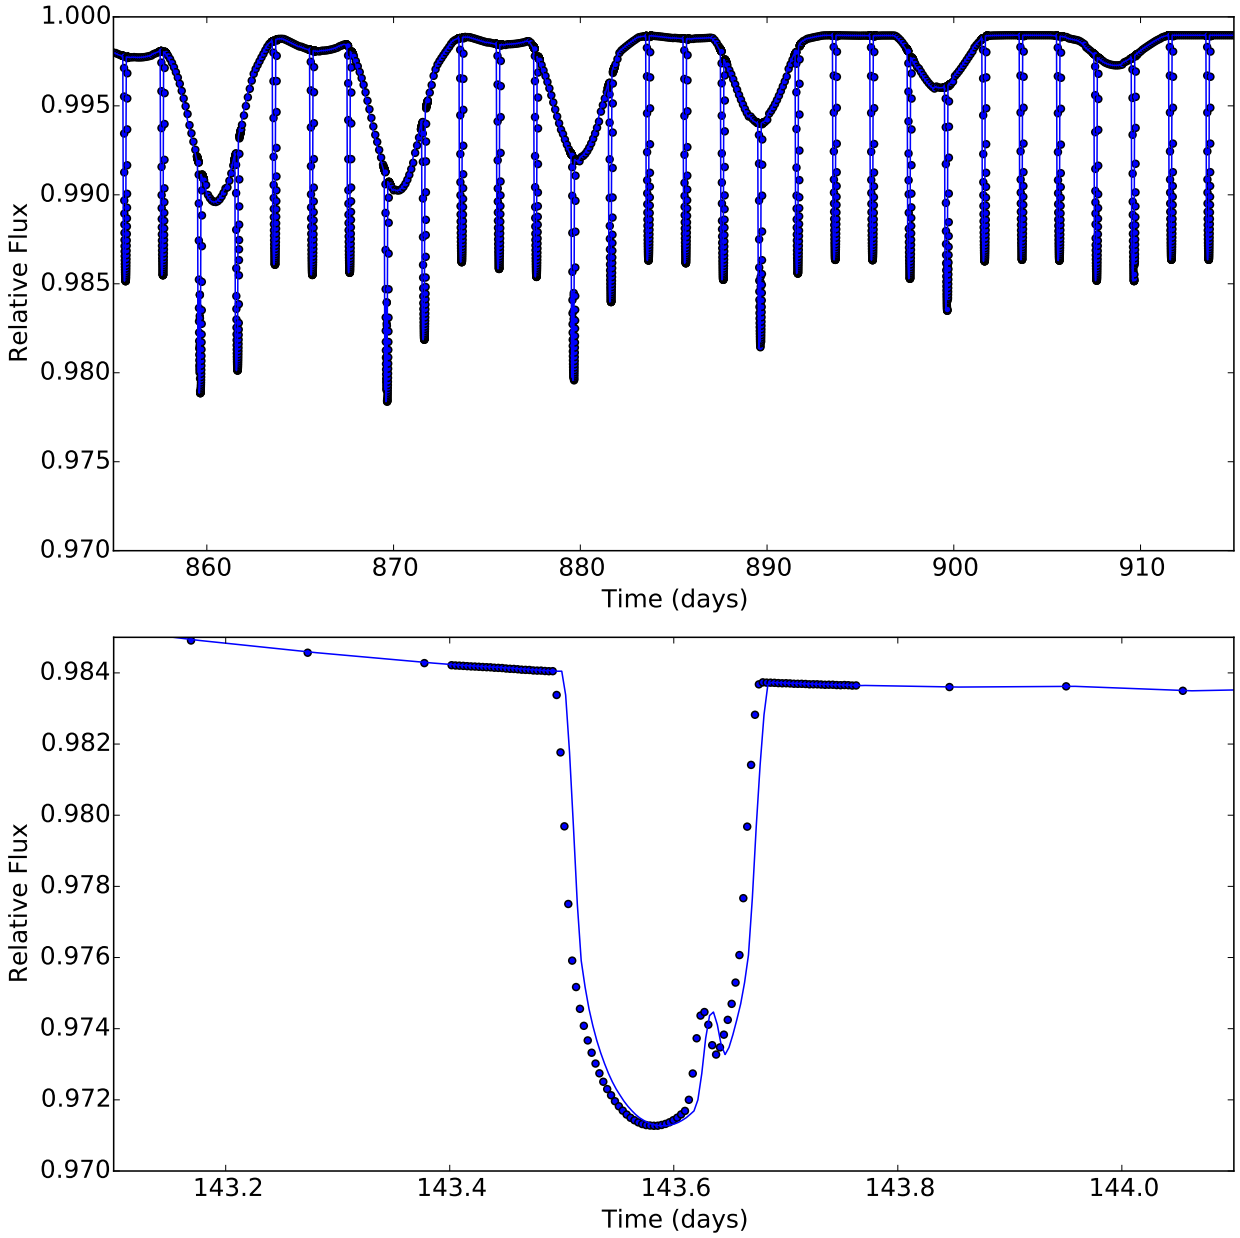

Figure 5.4: Top: 60 days from the 4-year simulated active G dwarf with a transiting hot Jupiter light curve, generated by Dr. Joe Llama for testing our method. This time window highlights the out-of-transit light curve evolution due to the decay of one large starspot group. Bottom: Zoomed in light curve of a single transit with a prominent starspot bump. A 5-minute cadence for in-transit regions ($\pm 1 t_{dur}$ from mid-transit) was used, while the out-of-transit data was down-sampled by a factor of 10.

for windows of time spanning twice the transit duration centered at the mid-point of each transit. In Figure 5.4 we show a 60 day portion of this simulated light curve, and a transit featuring a prominent spot-crossing bump.

The simulated, down-sampled light curve was split into 9.5 day long windows of time, with each subsequent window moved forward by 5.0 days (half a rotation period). This produced a total of 219 time windows to run through our MCMC fitting. We note that most transits in the simulated light curve had no spot-crossing bumps. Transits typically had at most 2 spot-crossing bumps, indicating the presence of 2 starspots in the path of the planet on a given hemisphere. We experimentally found that only 4 starspots were needed to adequately fit the Llama model data, typically with only 1 or 2 in the path of the transit.

Our MCMC modeling was run on each time window using 1000 steps and 300 independent parameter space walkers. The starspot contrast value was set to 0.7, as with GJ 1243. An a-scale length of 2.5 was used to govern the random Monte Carlo jump length. Our affine invariant sampler was run using a random seed, and all starspot parameters (radius, latitude, longitude) were allowed to vary. A maximum flux value was assigned by boxcar smoothing the light curve within each time window using a kernel of 11 data points and choosing the peak flux. This smoothed light curve was not used for the model fitting or analysis.

Once the MCMC fitting was complete, the best-fit (lowest χ^2) solution was chosen for each time window. For time windows with no recovered spot-crossing events, all 4 starspots were placed out of the path of the transit. The latitudes and radii of these non-occulted spots were weakly constrained. As such we only focused our analysis on the starspots that were flagged as having bumps in the transit. Figure 5.5 shows the longitude over time for starspots in every window that were determined to be in the path of the transit. The mid-time of each window was used as the abscissa, and all spots with detected bumps within each window are shown along the ordinate. Uncertainties from our MCMC modeling on the spot longitudes, as well as the radii and latitudes, were very small on average. Isolated spots in Figure 5.5 are considered a source of noise for our analysis.

Many linear features are clearly seen in this diagram, indicating starspots that were

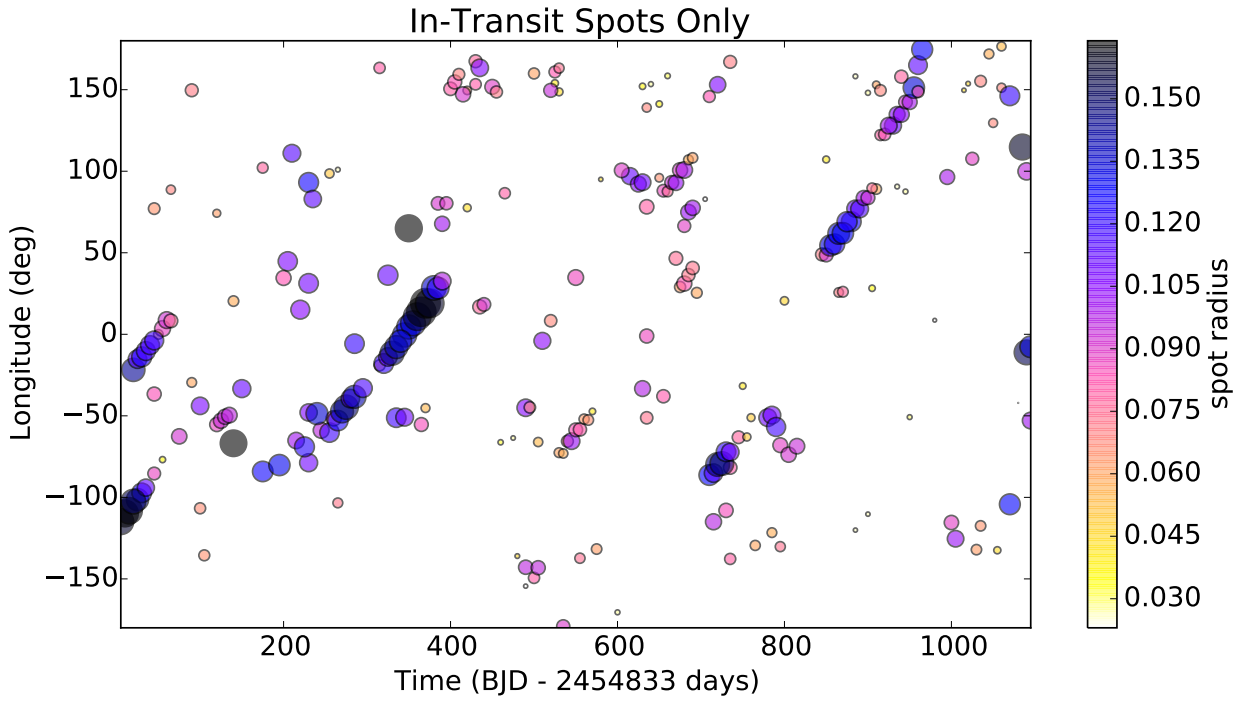

Figure 5.5: Longitude versus time over the 4 years of simulated data for starspots that were determined to be in the path of the planet from our MCMC models. Both the symbol color (from light to dark) and symbol size (from small to large) indicate the best-fit radius relative to the stellar radius of the starspots. Time here was arbitrarily set to match the *Kepler* mission.

recovered in subsequent windows, but that were rotating at a different rate than the rotation period used to phase-fold the data - 10 days in this case. This is the hallmark of differential rotation. Linear features with positive slope are the result of spots with *longer* rotation periods, where the feature appears at later times (larger latitudes) with each rotation period. Assuming Solar-like differential rotation, where the poles rotate slower than the equator, these linear features would point to spots at higher latitudes. Starspots located at the stellar equator, the centerline for the transit path, would produce linear features with a flat slope. However, no such features are seen in these model fits to the simulated data, which indicate the starspots must all be at latitudes above (or by symmetry, below) the equator. Note also these slopes are equivalent to the linear spot tracking features from Chapter 4, and can be converted into new rotation periods using Equation 4.2.

Starspot decay can be clearly seen for several of the linear features in Figure 5.5. The stellar activity model of Llama et al. (2012) produces very fast starspot emergence, with a more gradual, exponential decay. Linear features, such as beginning at Time$\sim$850 days, Longitude$\sim$50 deg, show clear emergence within one stellar rotation (time window), and a gradual decay tracked over $\sim$60 days.

5.3.1 Determining Starspot Evolution

To quantitatively study the starspot evolution for these transiting systems, we must develop a programatic way of choosing which of the recovered in-transit spots in Figure 5.5 are associated with each-other. In Chapter 4 we tracked the evolution of only two starspots over time. The spans of time showing clear differential rotation were chosen manually, and linear trends were fit to the starspot model longitudes. This more simple procedure was appropriate for exploring slow starspot shear on a rapidly rotating M dwarf (GJ 1243). In this Chapter we first study starspots from a simulated G dwarf, where the spots shear and evolve on timescales only a few times longer than the rotation period. We also wish to track the evolution of a larger number of starspots simultaneously: four spots in the case of the Llama simulation. Further, the procedure must be robust against possible small gaps in the

linear features due to missing time windows, data errors, or spurious MCMC results.

Our first attempt at grouping the starspots into unique associates used the simple Gaussian Mixture Model (GMM) classifier available in the Python machine learning package, *scikit-learn* (Pedregosa et al., 2011). This family of probabilistic models are appealing for their simplicity and popular use. The GMM classifier will associate all data points to a (user defined) number of Gaussian clusters, providing a probability for every data point of membership to each cluster. A major drawback of using GMM for this problem is the need for pre-defining the number of clusters to represent our data. The GMM classifier performs best on data with large variations in the density of points, and for clusters that have roughly Gaussian distributions in density. Unfortunately, GMM did not perform well on the data in Figure 5.5, as the linear features do not have Gaussian-like densities. Adding additional dimensions to the classifier (spot radius and latitude) did not improve the performance of the GMM.

We then turned to the wider range of "clustering" algorithms available in *scikit-learn*, and settled on the Density-Based Spatial Clustering of Applications with Noise algorithm, or DBSCAN, for our uses (Ester et al., 1996). DBSCAN is recommended in *scikit-learn* for applications where the clusters have similar densities, and works by finding the high density cores and expending the clusters outward using a kernel density tree method. DBSCAN can cluster data in an arbitrary number of dimensions, and does not prefer a certain morphology for the clusters. The predicted number of clusters and their properties are returned when DBSCAN is trained on a dataset. We tried a variety of the other available *scikit-learn* clustering algorithms, including K-means, Agglomerative, and Mean Shift; the performance of DBSCAN for our needs was consistently superior.

In Figure 5.6 we show the result of running the DBSCAN clustering algorithm on the Llama model starspot results from Figure 5.5. The DBSCAN machine learning algorithm was trained on the in-transit starspot times, longitudes, and radii resulting from our MCMC fitting. We repeated the procedure including also the spot latitudes, but they provided no additional constraint on the clusters. A minimum threshold of 3 data points per cluster,

and a maximum point separation of 15°, was used to define the DBSCAN clusters. The DBSCAN algorithm produced 21 possible clusters of starspots over time, with each cluster assigned an arbitrary color in Figure 5.6. Grey points show spots or groups of spots that did not pass these thresholds, and thus did not belong to a cluster feature. Training the DBSCAN clustering algorithm on these data took under 2 seconds to compute on a standard Linux workstation.

The longitude axis here was defined from -180° to 180° to accommodate the large cluster of features seen in the first 400 days. If a starspot was long-lived and differentially rotating from the equatorial period used to define phase (equivalently longitude) in Figures 5.5 and 5.6, it could wrap through the longitude axis over time. The DBSCAN algorithm, however, cannot fit features over the longitude axis as a circular coordinate system. We manually adjusted the range of the longitude axis in this case.

To determine the rotation period of each evolving starspot feature detected by the DBSCAN algorithm, we fit the clusters with linear trends. The longitude evolution of spot clusters with at least 5 independent measurements was determined using a linear least squares fit, shown in Figure 5.6 as black lines. The slopes of these linear fits have units of degrees day^{-1}, and with the phase-folding period of 10 days the slopes could be converted to new rotation periods using Equation 4.2. The distribution of starspot rotation periods from the DBSCAN clusters is shown in Figure 5.7. Nearly all the spot cluster features had slower rotation rates (larger periods, or positive slopes in longitude evolution), likely from starspots at latitudes away from the equator. Only two starspot clusters were recovered with rotation rates that appeared to be faster than the phase-folding rotation rate (10 days). Since we know from the input parameters of the Llama simulation that the equatorial rotation period was defined defined to be the fastest rotation speed on the star (Solar-type differential rotation), these clusters are considered spurious detections. Our estimation of the underlying differential rotation law must therefore be robust against such outliers.

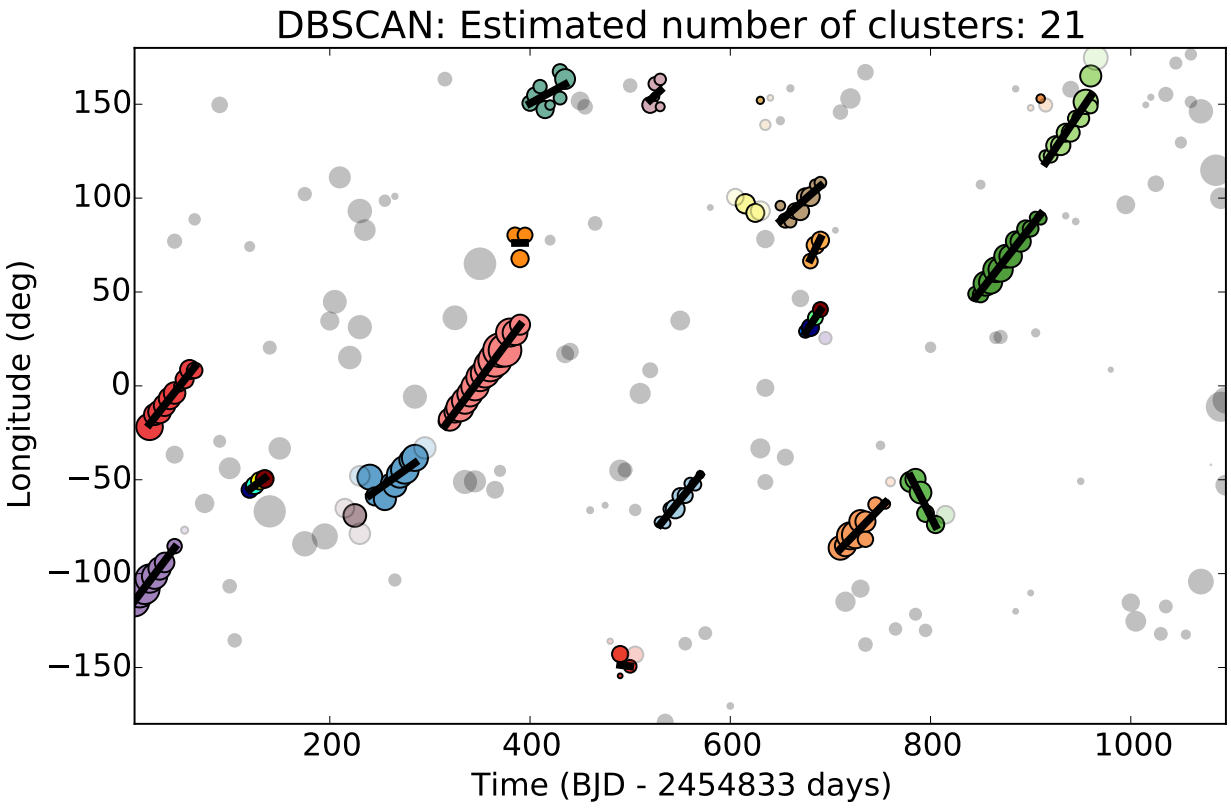

Figure 5.6: The longitude versus time map for starspots in the path of the transit, with the symbol size indicating spot size, as in Figure 5.5. Colored symbols indicate starspot measurements that belong to a cluster from the machine learning DBSCAN clustering algorithm. Grey symbols indicate spot measurements that were not associated with any other spot features from DBSCAN. Spot clusters with at least 5 epochs were fit with a first-order polynomial to quantify their evolution in longitude over time (black lines). Positive slopes correspond to larger rotation periods, which are due to higher latitude starspots.

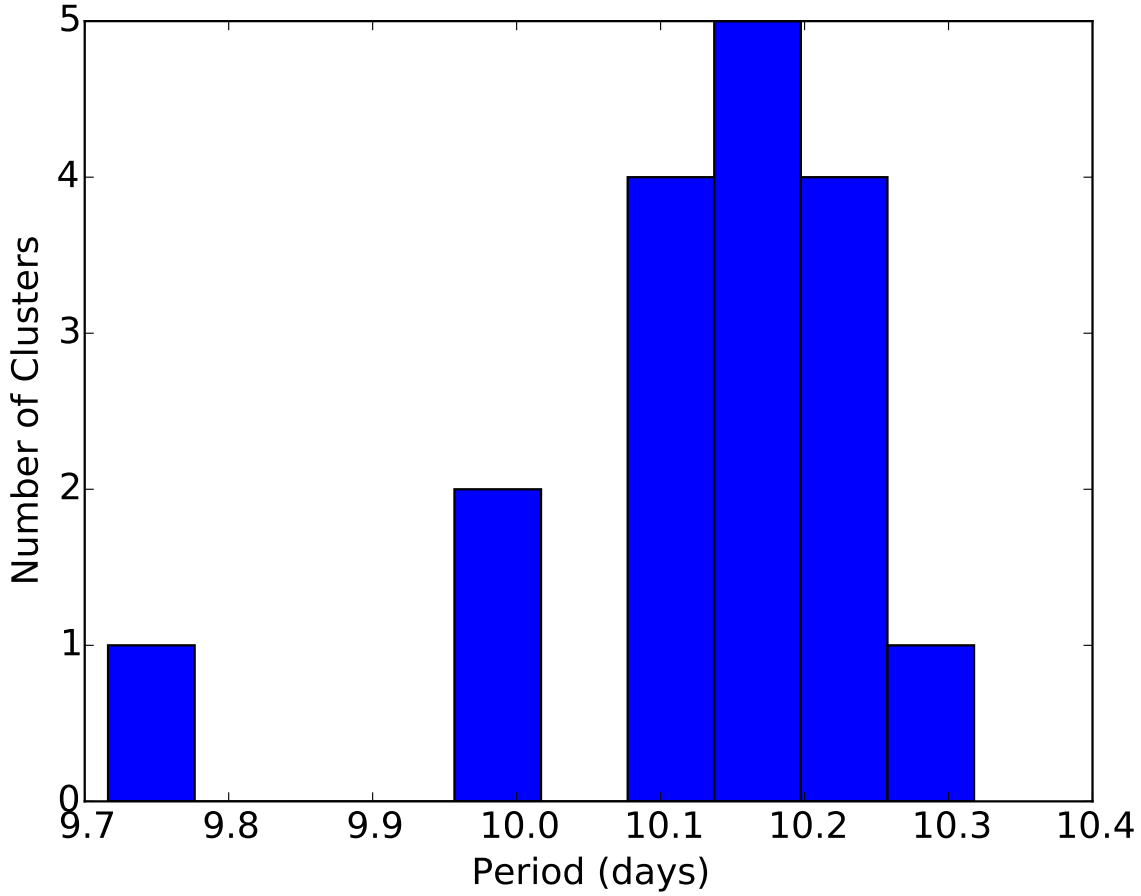

Figure 5.7: Distribution of periods determined from the linear fits to spot clusters in Figure 5.6. For reference, the phase folding period was 10 days. Lines with positive slope in longitude evolution from Figure 5.6 yielded larger rotation periods, indicative of spots at higher latitudes, assuming a Solar-like differential rotation law.

5.4 Estimating Physical Parameters

5.4.1 Differential Rotation Law

The scientific goal for our modeling these data is not to simply reproduce the light curves, but instead to estimate physical parameters of the star and starspots. Specifically we aim to measure the differential rotation law characterized by the coefficient $k \equiv \Delta\Omega/\Omega$, and the starspot diffusion timescale traced by the starspot decay profile. Both of these measurements require tracking starspots sizes and positions (or rotation periods) over time in detail.

To measure differential rotation *directly* from the results of our light curve modeling and DBSCAN clustering, we would look for a correlation between the rotation period and latitude of a starspot. In Figure 5.8 we show the median latitude versus the rotation period (derived from the feature slopes in Figure 5.6) for the 17 in-transit DBSCAN clusters detected in the previous section. For comparison we also show a solar-type differential rotation curve overlaid, using values of $k = 1$ and $P_{eq} = 10$ days. Since the Llama simulation was generated explicitly with solar-type differential rotation, we expect the DBSCAN cluster properties to follow this type of curve. While our light curve modeling procedure determines these starspots to be in the path of the transit (approximately $\pm 10^\circ$ in latitude), and the starspot periods are very well determined for most of the DBSCAN features, no coherent relationship is seen between starspot latitude and rotation period. This indicates our light curve fitting procedure does not yield latitudes more accurately than being contained within the transit path.

Without accurate latitudes for each starspot, we instead rely on a statistical determination for the differential rotation law. If we assume starspots form over the entire range of latitudes covered by the transit path, the differential rotation law could be computed by assuming the maximum and minimum rotation periods measured correspond to starspots located at the maximum and minimum observable latitudes. In the case of the Llama simulation we know the planet transit path is directly along the stellar equator, that a planet radius of $r_p/r_s = 0.1$ can detect spot features with latitudes up to $\pm 9^\circ$, and that the equato-

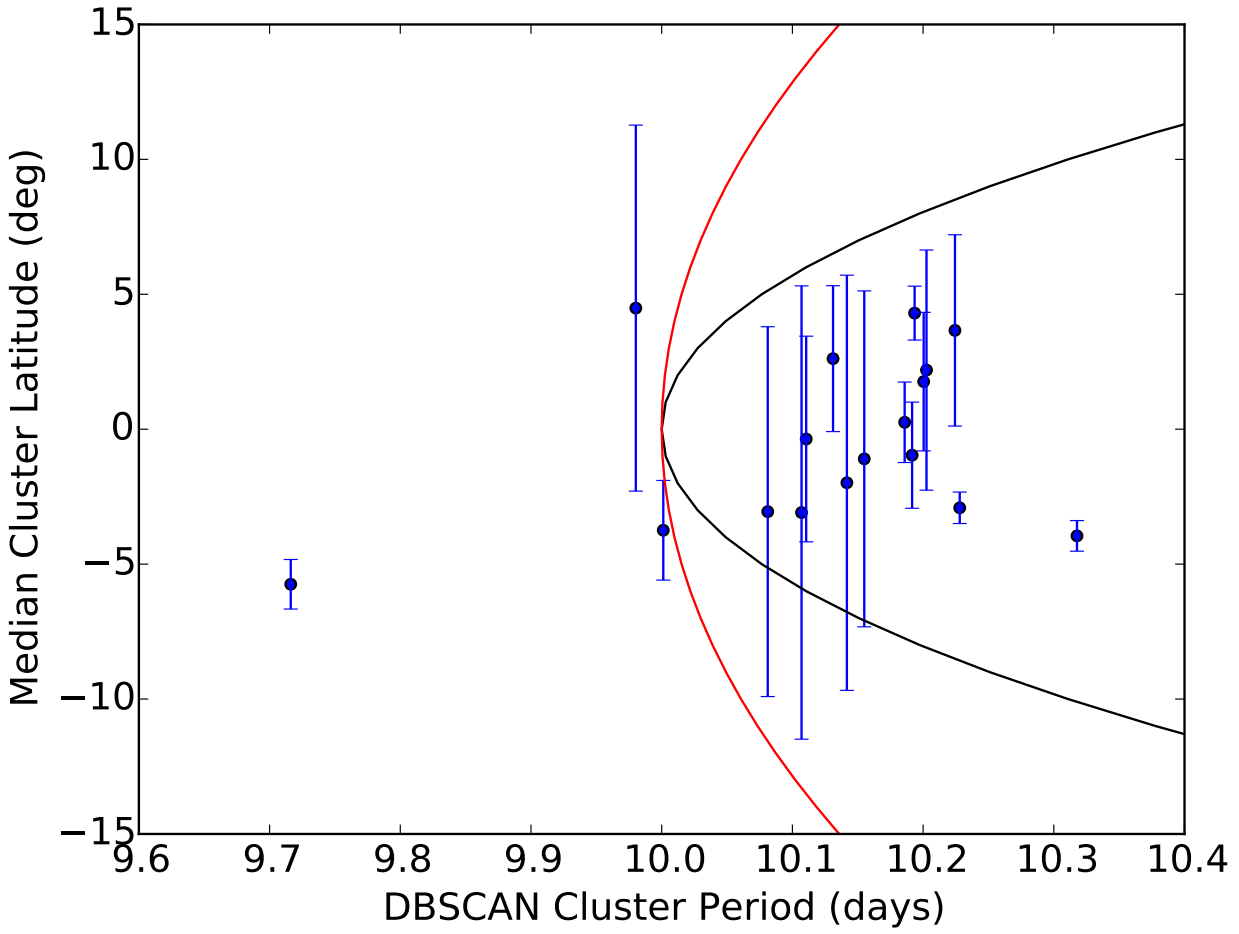

Figure 5.8: Median best-fit latitudes on the stellar surface for spots within each DBSCAN cluster versus the rotation period determined from the linear fits to the DBSCAN cluster longitudes over time. Error bars denote the standard deviation in the latitude for spots within each DBSCAN cluster. No correlation between the cluster period and latitude is seen, indicating the latitudes cannot be constrained beyond the planet radius. Solar-type differential rotation with $k = 1$ (black curve) and the Solar value of $k = 0.2$ (red curve) are shown for reference.

rial rotation period is 10 days. The maximum rotation period determined in Figure 5.7 was 10.317 days. Starting with Equation 4.3, we can rearrange terms to solve for the differential rotation coefficient:

$$\begin{aligned} P_\phi = &\quad P_{eq}/(1 - k\sin^2\phi) \\ k \equiv &\quad \frac{\Delta\Omega}{\Omega} = \frac{1 - P_{eq}/P_\phi}{\sin^2\phi} \end{aligned} \tag{5.1}$$

Assuming $P_{eq} = 10$ days, the maximum latitude $\phi = 10°$, and $P_\phi = 10.317$ days, we determined the differential rotation coefficient to be $k = 1.256$. This was reasonably close to true value used in the Llama simulation of $k = 1.0$, shown in Figure 5.8. Our estimate of k was larger than the true value because the starspot with highest rotation period had a central latitude of 10.106°, *slightly higher* than the limit of the transit path.

For a real star whose equatorial rotation period is not known a priori, to estimate the differential rotation law we must rely on the range of recovered rotation periods, and assume that starspots sample the entire latitude range covered by the transit. Note that the minimum period recovered in our analysis (9.716 days) was significantly shorter than the stated equatorial period of 10 days, which is physically impossible given the solar-type differential rotation imposed in the simulation. If we remove the two spurious DBSCAN clusters recovered with rotation periods shorter than 10 days, the two next shortest rotation periods (10.0012 and 10.081 days) did not quite reach the equatorial period. This implies none of the starspots that we recovered from the Llama simulation formed exactly at the equator, and so there was no way to measure the rotation period at this lowest latitude with the transit. If either of these spurious minimum periods were kept in the analysis, and now assuming a maximum starspot latitude of ±10°, the estimated differential rotation coefficient would be $k = 1.93$ for $P_{eq} = 9.716$ days, and $k = 1.085$ for $P_{eq} = 9.98$ days. The resulting differential rotation coefficient is very sensitive to the choice of the minimum and maximum periods.

If very few starspots are recovered, the likelihood is small that the entire range of rotation periods has been sampled. Clearly this method for determining the differential rotation law favors heavily spotted stars (many bumps in the transit light curves), as well as long

duration datasets where many unique starspots are observed. Henry et al. (1995) describe an equivalent problem for determining the differential rotation law from starspots tracked using light curve modulations (as in our analysis of GJ 1243 in Chapter 4), but where the latitude of each starspot is entirely unknown. Henry et al. (1995) quantify the differential rotation constant from the range of observed rotation periods as

$$kf = (P_{max} - P_{min})/P_{avg}\,, \tag{5.2}$$

where f is a "distribution function" (really a correction factor) computed by Hall & Henry (1994) for different numbers of starspots. This distribution function corrects for the limited sampling of latitudes by small numbers of starspots, and assumes that spots emerge randomly over the entire stellar surface. By generating n starspots at uniformly spaced latitudes, the distribution function is computed as

$$f(n) = \sin^2 \phi_{max}(n) - \sin^2 \phi_{min}(n)\,, \tag{5.3}$$

where $\phi_{max}(n)$ and $\phi_{min}(n)$ are the largest and smallest latitudes of the n spots. The correction factor then ranges from $f = 0.5$ for $n = 2$ to $f = 0.986$ for $n = 18$ in Hall & Henry (1994).

In estimating of the differential rotation law for a given number of in-transit spots with poorly determined latitudes we might invoke a similar correction factor to that of Hall & Henry (1994), and thus account for the limited span of latitudes sampled by a finite number of starspots. The latitude range sampled by the transit path is well determined for most *Kepler* planetary systems, and even for large hot Jupiter systems the planet typically only spans 10–20° in latitude on the stellar surface. Thus, assuming starspots are randomly located within this latitude band seems more feasible than across the entire stellar surface, and the correction factor would be appropriate to use for a limited number of spots. For in-transit spots, the differential rotation correction factor as a function of the number of individual observed spot features can be computed as

$$f(n) = \frac{\sin^2 \phi_{max}(n) - \sin^2 \phi_{min}(n)}{\sin^2 \phi_{tr,max}}\,, \tag{5.4}$$

where $\phi_{max}(n)$ and $\phi_{min}(n)$ are the largest and smallest latitudes sampled by the n uniformly spaced spots within the transit path, and $\phi_{tr,max}$ is the maximum latitude probed by the transit. Note that unlike the Hall & Henry (1994) distribution function from Equation 5.3 that was applicable to all stars with spots presumably spanning the full stellar disk, our correction factor $f(n)$ must be calculated for each star–planet system, as the latitude ranges are dependent on the system parameters. For the Llama simulation we recovered 17 starspot features with the DBSCAN algorithm, (including the spurious detections with periods < 10 days), and had a maximum probed latitude of $\phi_{tr,max} = 10^{\circ}$ based on the planet radius. This results in a correction factor of $f(17) = 0.8898$, which would further increase the implied differential rotation coefficient. Example correction factors for various numbers of starspots and different limiting latitudes are given in Table 5.1. For very active stars having more than 100 starspots with measured rotation periods (such as Kepler 17), this correction factor becomes effectively negligible.

5.4.2 Starspot Lifetimes

Besides constraining the differential rotation law for a star, the DBSCAN clusters trace the size evolution of in-transit starspot groups over time. For the Sun, the rate of spot decay has been measured for individual active regions over the past century (Javaraiah, 2011). The decay rate of sunspots scales with spot size so that larger area spots live longer (Hathaway & Choudhary, 2008). A linear relationship between spot area and duration was determined more than 50 years ago (Bumba, 1963). This connection between the maximum area and lifetime of a sunspot is known as the "Gnevyshev-Waldmeier rule" (Petrovay & van Driel-Gesztelyi, 1997).

In Figure 5.9, I show the duration of the DBSCAN starspot clusters versus the maximum starspot area recovered within the cluster for the Llama simulation results. The starspot area here is given in the traditional solar units of 10^{-6} times the area of the visible hemisphere (μHem), or approximately $3{\times}10^{16}$ cm^2. A stellar radius of 1 R$_{\odot}$ was assumed for the Llama simulation. The peak sizes of these simulated starspots are enormous compared to spot

Table 5.1. Differential rotation coefficient correction factor values ($f(n)$) as a function of number of spots (n) and various maximum spot latitudes ($\phi_{tr,max}$).

n	$\phi_{tr,max} = 5°$	$\phi_{tr,max} = 10°$	$\phi_{tr,max} = 20°$	$\phi_{tr,max} = 30°$
2	0.3337	0.3348	0.3394	0.3473
7	0.7504	0.7517	0.7567	0.7654
12	0.8465	0.8474	0.8511	0.8574
17	0.8891	0.8898	0.8927	0.8976
22	0.9132	0.9138	0.9162	0.9201
27	0.9287	0.9292	0.9312	0.9345
32	0.9395	0.9400	0.9417	0.9445
37	0.9475	0.9479	0.9494	0.9519
42	0.9536	0.9539	0.9553	0.9575
47	0.9584	0.9587	0.9599	0.9620
52	0.9624	0.9626	0.9637	0.9656
57	0.9656	0.9659	0.9669	0.9686
62	0.9683	0.9686	0.9695	0.9711
67	0.9707	0.9709	0.9717	0.9732
72	0.9727	0.9729	0.9737	0.9751
77	0.9744	0.9746	0.9754	0.9767
82	0.9760	0.9761	0.9769	0.9781
87	0.9773	0.9775	0.9782	0.9793
92	0.9785	0.9787	0.9793	0.9804
97	0.9796	0.9798	0.9804	0.9814
102	0.9806	0.9808	0.9814	0.9823

groups found on the Sun, which typically reach maximum areas less than 1000 μHem (e.g. Hathaway & Choudhary, 2008).

A significant correlation between the maximum starspot area and the spot duration is clear in Figure 5.9, which seems to satisfy the "Gnevyshev-Waldmeier rule". However, given the uncertainty in the starspot latitudes, both the maximum starspot area and starspot duration may be lower limits. For example, if a larger area starspot was centered above or below the planetary transit path, at it's peak size the spot might extend far enough to be grazed by the transit. The resulting bump in the transit light curve would too small to accurately indicate the total spot size. As the spot decayed it might then be too small to be in the path of the transit, and thus the starspot's apparent duration would be truncated. Indeed, the Llama simulation was generated with every starspot having the same maximum area, and their decay rates only dependent on the amplitude of differential rotation shear at their (randomly chosen) latitudes. Still, the ability to directly compare measurements of spot area and lifetime between the Sun and another star for starspots of many sizes is unique. It is tantalizing to study such ensemble results for possible information about the underlying stellar diffusion timescale, which may be primarily responsible for governing the starspot lifetimes (Krause & Ruediger, 1975).

In Figure 5.10 I show the detailed area evolution for the 10 largest starspots recovered from the DBSCAN clustering. The starspot radius from our MCMC code has been again converted to area in the conventional Solar units of μHem. The 5-day spacing of data points in this figure is due to each area value coming from an independent time window of the light curve. Multiple spot values at single time window can occur due to spurious results from the DBSCAN clustering algorithm. As every spot in the Llama simulation was generated with the same initial size, the variation in areas is due to starspots being located at a range of latitudes as mentioned above. The largest spots show a rapid emergence (Time < 0) and gradual decay (Time > 0), as observed for sunspots. The diffusion timescale, which governs the starspot decay rate as a function of area, is slower in this simulation than for the Sun, as evidenced by the more gradual evolution in the simulated starspots than for the scaled-up

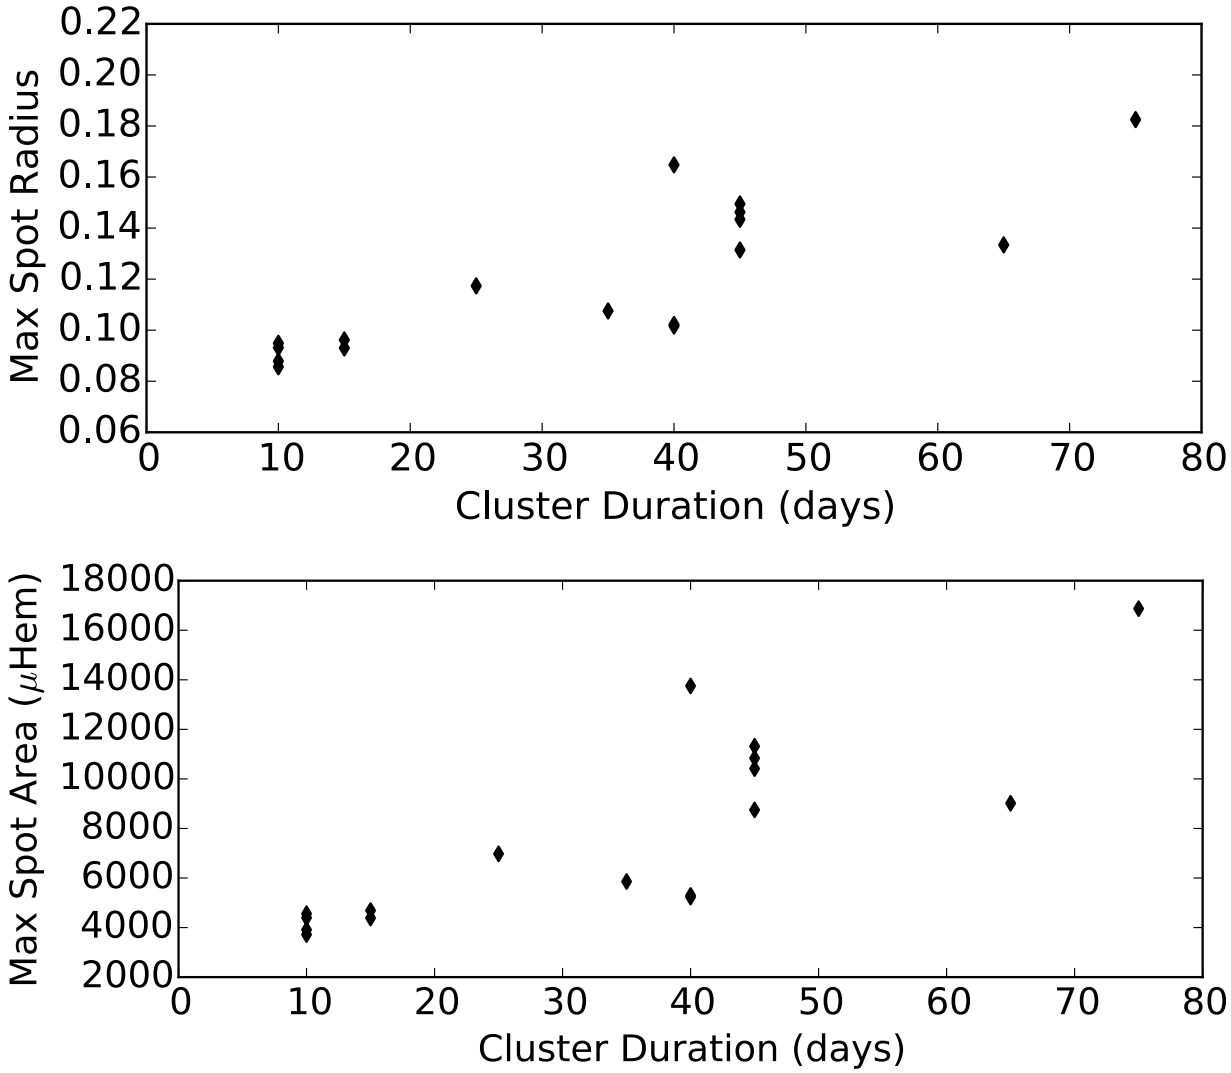

Figure 5.9: Top: Duration of transit-crossing starspot features recovered from our DBSCAN clustering versus the maximum starspot radius determined with each cluster. Bottom: The same data as above, converting the maximum starspot radii to areas, and presented in the traditional solar units of millionths of a solar hemisphere. This trend matches the expectation from the solar "Gnevyshev-Waldmeier rule" of spot lifetimes increasing with surface area. However, the durations may be inaccurate due to larger spots at higher latitudes grazing the transit for less than the full spot lifetime.

sunspot decay shown in Figure 5.10. Determining the diffusion timescale from tracing this area decay directly may prove more robust for our in-transit analysis than from measuring the starspot lifetimes, as starspot durations can be shortened at the beginning and end of the spot lives by low signal to noise in the light curve.

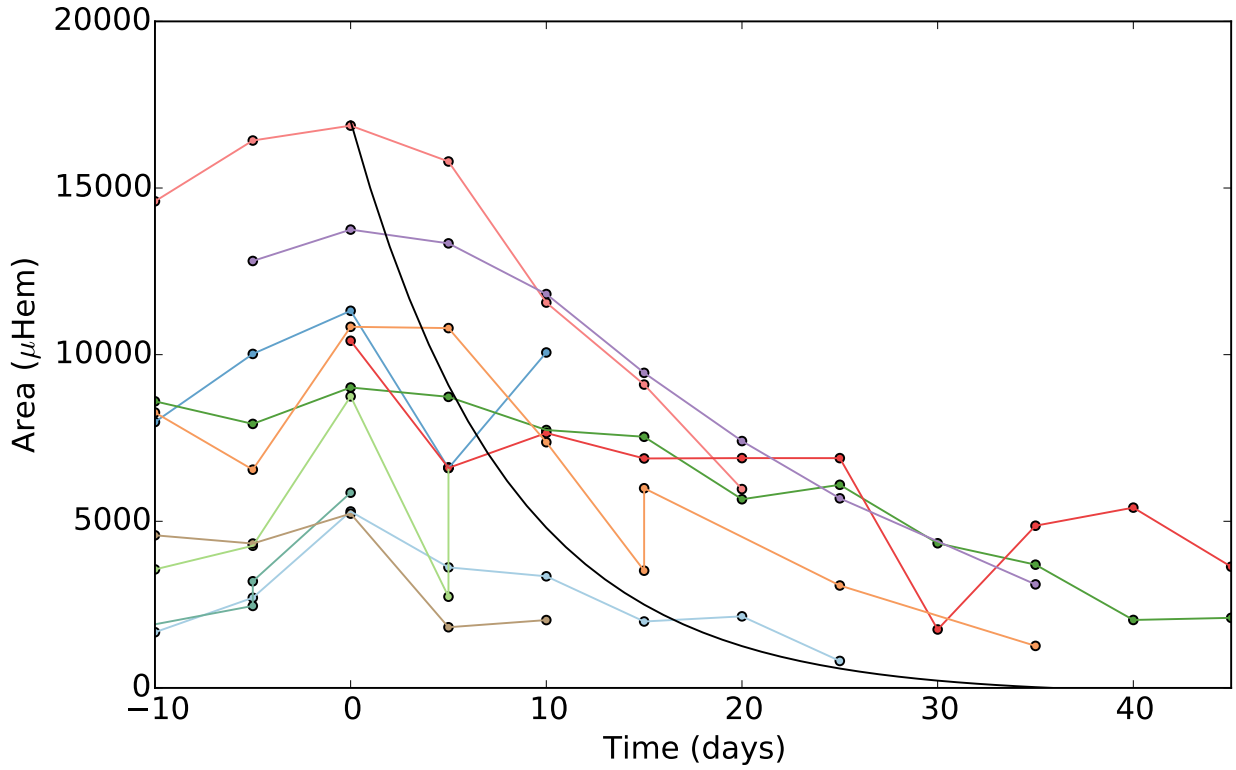

Figure 5.10: Decay profiles for the 10 largest, in-transit starspot DBSCAN clusters, with the time axis centered to the time of maximum recovered spot area. Spot area values here are given in the conventional solar units of millionths of a solar hemisphere. The sunspot decay profile from Equation 3 of (Hathaway & Choudhary, 2008) is shown for comparison, scaled to the largest area recovered from any starspots in the Llama simulation (black solid line).

5.5 Kepler 17

Thus far we have built up our analysis of the starspot evolution and stellar differential rotation based on modeling and analyzing a simulated light curve. Starspots were robustly traced over 10's to 100's of days in this model. By studying only the in-transit spots, we are able to statistically recover the stellar differential rotation coefficient used to generate the simulated light curve. The starspot evolution and decay profiles recovered may shed light on the diffusion timescales for the star as well.

With these tools at our disposal, we are now able to model and interpret the light curve from a real star + planet system: Kepler 17. This light curve is significantly more complex than the Llama simulation, with a wide range of systematic trends in the flux, discontinuities and gaps in the data between quarters, and entire missing quarters of data due to a failed CCD module on board *Kepler* after the first year. Additionally, the transits show many more starspot features, with nearly all transits exhibiting one or two (or more) distinct spot bumps.

Kepler 17 is an ideal system for analyzing with our spot modeling code. The planet has been well characterized from the *Kepler* light curve, as well as follow-up ground and space-based data by Désert et al. (2011). Kepler 17 is a Solar analog, with a mass of 1.06 $M_{\odot}$. The transiting planet, Kepler 17b, is a "hot Jupiter", with an orbital period of 1.4857108 days (Désert et al., 2011). The nearly 4% (at maximum) out-of-transit quasi-periodic light curve modulations are due to strong starspot activity, and Désert et al. (2011) note the presence of at least 5 individual in-transit spot bumps in the first 6 Quarters of *Kepler* data. Bonomo & Lanza (2012) further characterized the starspots for Kepler 17 using the same Q0–Q6 data as Désert et al. (2011). This system is unique in that the planet orbits almost exactly 8 times faster than the stellar rotation period, meaning each starspot is probed by multiple transits each time it appears on the facing side of star.

In Figure 5.11 we show the nearly 4-year light cure for Kepler 17. For each quarter of data we removed a linear trend to correct for systematic trends in the light curve. As in

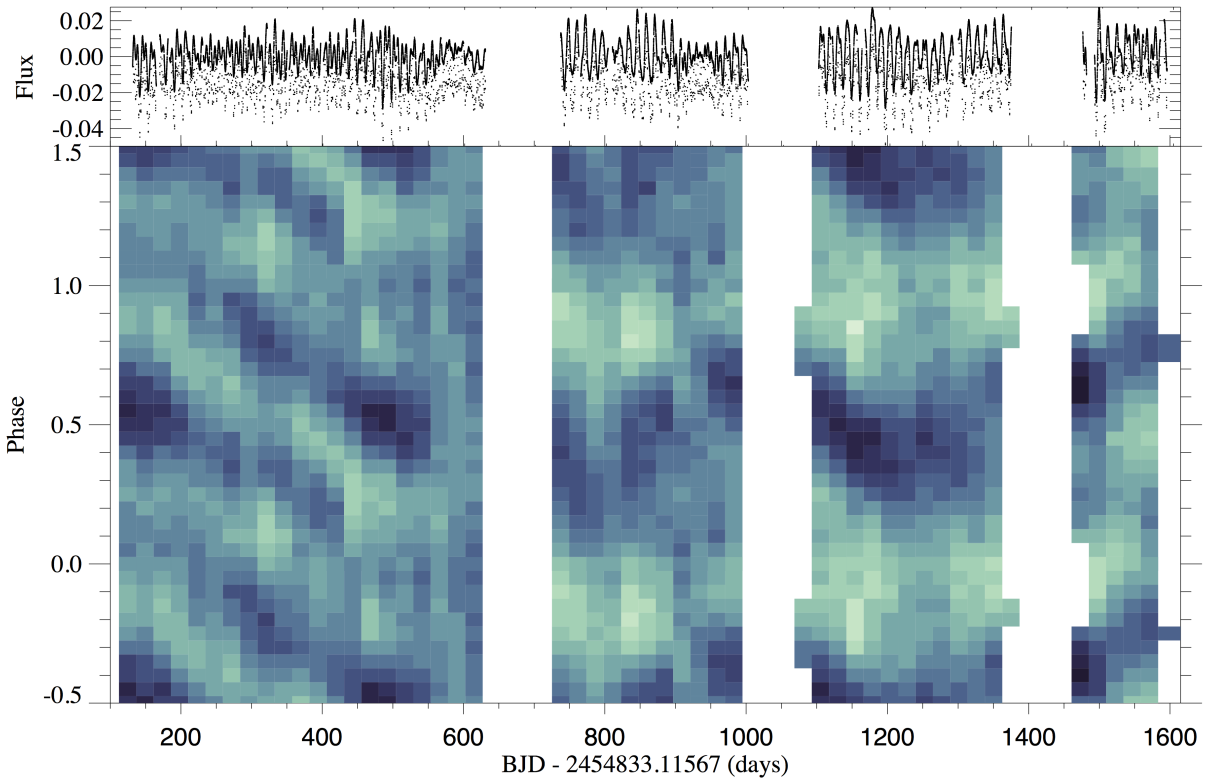

Figure 5.11: Top: Fractional flux light curve for Kepler 17, including the long cadence out-of-transit data and short cadence in-transit. Bottom: Continuous phased light curve map for the out-of-transit Kepler 17 data. This is the same as Figure 4.3. Pixels span two rotation periods in time (24.516 days), and 0.05 in phase. Pixel shade (dark to light) indicates increasing median flux within each (time, phase) bin. Vertical white gaps are due to quarters with no *Kepler* data. Starspot evolution is seen as dark bands migrating in time and phase.

the Llama simulation, we use a higher sampling of data around the transits to gain the most information about the in-transit spot bumps. The *Kepler* long (30 minute) cadence PDCMAP data was used to trace the out-of-transit light curve. Within ±1 transit duration of the transit mid-points we used the *Kepler* short (1 minute) cadence data.

Désert et al. (2011) used a Lomb-Scargle periodogram to calculate a rotation period for Kepler 17 of $P_{rot} = 11.89$ days, almost exactly 8 times slower than the very well constrained planetary orbital period. Bonomo & Lanza (2012) analyze the same *Kepler* data (Q0–Q6), but find $P_{rot} = 12.01$ days by minimizing the migration of spot features in the out-of-transit photometry. We used nearly three times more data for Kepler 17 to estimate it's mean rotation period. Using a Lomb-Scargle periodogram of the entire detrended long cadence light curve (omitting the transits) we find the strongest period to be at 12.258 days, significantly slower than found by these previous studies.[1] This may be due to higher latitude starspots being present at later times in the *Kepler* light curve. However, as found in the rotation and differential rotation recovery exercise we participated in (Aigrain et al., 2015), the periodogram from out-of-transit modulations can be skewed simply due to the emergence and evolution of spots at multiple phases, even if no latitude dependent differential rotation is present.

In Figure 5.11 we also show the phase versus time map of Kepler 17's relative flux for the 4 years of *Kepler* long cadence data. The rotation has been phased here at the rotation period of 12.258 days. Many coherent starspot features are seen as dark bands, which evolve in phase over time. The most significant features in the first 2 years of the light curve appear to last for over 200 data and have a negative slope in phase over time, indicating they rotate at a faster period than the phase-folding period (akin to the secondary spots seen on GJ 143 in Chapter 4). This reinforces the notion that the average rotation period of 12.258 days represents the rotation at a higher latitude, and that spot features at latitudes closer to the equator are present in the first 2 years of the light curve. Manually tracing the largest three

[1]The exact value used in our analysis was 12.25817669188 days.

starspot tracks in Figure 5.11 before Time = 650, produces approximate slopes of -0.00195, -0.00162, and -0.00224 phase per day. Converting these to rotation periods using Equation 4.2, we find an average rotation period for the out-of-transit modulations in the first 2 years of the Kepler 17 light curve of ∼11.97 days, very close to the value recovered by Bonomo & Lanza (2012) of 12.01 days.

5.5.1 Refining the Planet Radius

While characterizing the transiting exoplanet itself was not initially a goal of this study, we have found that the relative planet radius (r_p/r_s) reported by Désert et al. (2011) is too small. We discovered that if we fit the Kepler 17 light curve with our starspot modeling code using the published planet radius, the resulting starspot bumps were far too small, and the MCMC parameter space search for spot properties did not yield comparable fits to the in-transit and out-of-transit data. This was due to the presence of starspot bumps in nearly every transit, which caused the average transit depth to be too shallow, and thus the implied planet radius to be too small in Désert et al. (2011). In Figure 5.12 we show the combined light curve for all 611 transits with short cadence data in the entire 4 year *Kepler* light curve. The scatter in the out-of-transit data is nearly 2.5 times smaller than at transit center, due to the presence of evolving and decaying starspots. This effect has been noted before (Czesla et al., 2009), and was even partially accounted for in the analysis of Q0–Q6 data for Kepler 17 by Désert et al. (2011). However, our much longer light curve provided over three times more transits to use in the re-determination of the transit depth.

A typical method for addressing this is to chose the deepest available transit, or similarly the deepest portions of several transits. These are assumed to be "spot free", and the Mandel & Agol (2002) transiting planet model is then fit. However, as we'll see in the next section, Kepler 17 does not show a period of time that is absent of in-transit starspots. The true transit depth in this data can only be approximated by finding the "lower boundary" of the combined transits in Figure 5.12 from the deepest points of every transit.

We fit the Mandel & Agol (2002) transit model to the lowest points of all the transits for

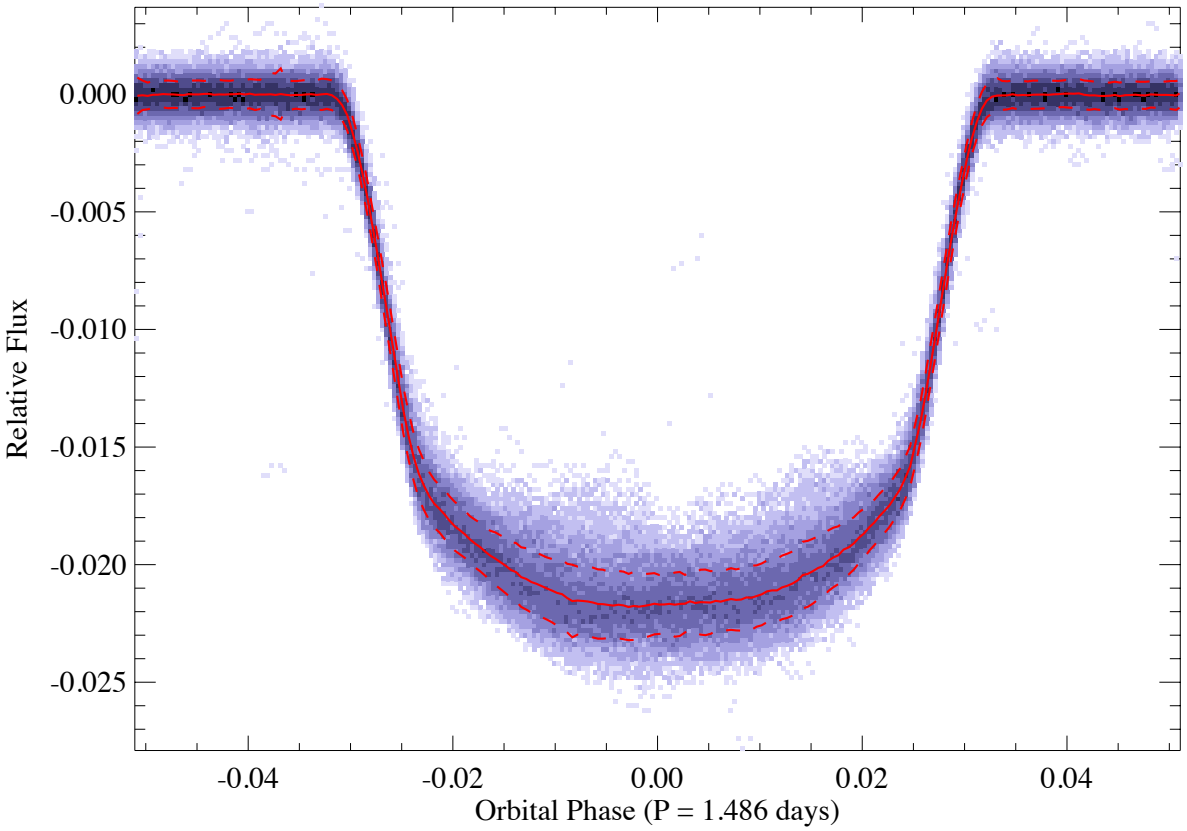

Figure 5.12: Stacked light curve for all 611 transits of Kepler 17b with short cadence data in the 4 years of *Kepler* data. Pixel shade (light to dark) indicates increasing numbers of observations. The median transit shape (red line) and standard deviation (red dashed lines) is shown. The out-of-transit scatter is $\sim$2.5 times smaller than the in-transit data.

Kepler 17. This resulted in a planet radius of $r_p/r_s = 0.13414$, slightly larger than the radius provided in Désert et al. (2011) of $r_p/r_s = 0.13031$. This resulted in sufficiently deep transits to model the light curve with. Our re-fitting for the planetary parameters also produced a small but significant update to the planet's orbital inclination as well, increasing it from the Désert et al. (2011) value of $i = 87.2°$ to $i = 88.95°$. The stellar inclination was fixed at 0° (equator facing the observer).

We emphasize the importance of this effect in constraining the sizes of transiting exoplanets for future research. For stars that show out-of-transit, starspot driven modulations, the implied planet radii are certainly affected by surface starspot activity at optical wavelengths. In the infrared, where the starspot temperature contrast is less important, the planet radius should be closer to the true value. The infrared is also the wavelength regime that is most interesting for transit transmission spectroscopy, where determining the exoplanet's radius from variations in the depths of transits as a function of wavelength is used to estimate the planetary atmosphere's composition. Determining the "true" value of r_p/r_s remains an open challenge.

5.6 Fitting the Kepler 17 Light Curve

Using the mixed short and long cadence data described above, we again binned the light curve in to individual windows of time for our spot modeling analysis. Time windows with a duration of 95% the rotation period were used, and each subsequent window was advanced by half a rotation period. Each individual time window contained 7 transits to fit.

In each time window we fit for the longitudes, latitudes, and radii of all 8 starspots. The stellar and planet properties remained fixed. As in our previous analyses, we used a fixed flux contrast of 70% for every starspot. The spots were given random starting positions, and the affine invariant MCMC sampler was run with 1000 walkers for 300 steps. The MCMC was then re-seeded using the best-fit parameters from these 300 steps, and the sampler was then run for 1000 steps. The initial 300 steps were effectively a burn-in stage, and the following 1000 steps used for finding converged solutions.

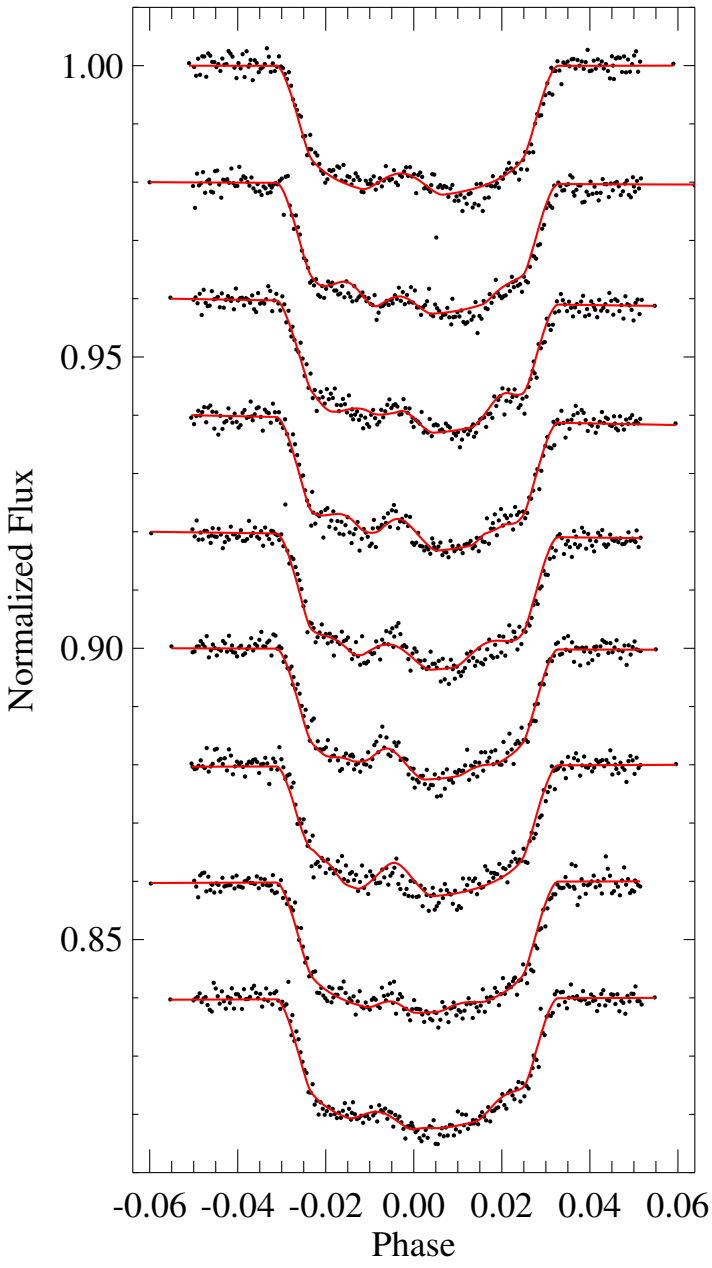

Figure 5.13: An example of the high quality fits of our MCMC modeling (red line) to the high cadence in-transit data (black points), and demonstrating the recovered spot bumps evolving over time. This shows every 8th transit, such that nearly the same longitude is probed within each transit. Multiple starspots can be seen to emerge and decay, while slowly shifting in their longitudes.

These models produced very good fits to the in-transit and out-of-transit data. In Figure 5.13 we show the evolution of in-transit spot bumps over a period of ∼100 days, plotting every eight transit. Note that each of the transits shown comes from a totally independent MCMC modeling time window. Since the mean stellar rotation period is approximately eight times greater than the planet's orbital period, every eighth transit passes over nearly the same stellar longitudes, and thus the same starspots can be visually tracked. The prominent spot bump that emerges in the first transit in Figure 5.13 grows and then decays, while slowly shifting to earlier phases within the transit (or larger values of stellar longitude).

The results of our MCMC modeling for each transit are also shown in Figure 5.13. Since each transit may be included in up to three separate neighboring time windows (typically only two, however), we have chosen the best-fit (lowest χ^2) model on a per-transit basis. An average of three starspot bumps were detected within each of these transits.

As with the Llama simulation, we tracked the evolution only for starspots determined to be in the path of the transit. We further only studied spots within each time window that produced at least one bump of duration 3 minutes (3 data points) in any of the 7 transits. As the starspot sizes and positions may evolve significantly within a single rotation period, we utilized the overlapping time windows as independent estimates of the starspot properties within each transit. By stepping through each transit and picking the model that best fit the data, it was possible to have time windows where the positions of more than 8 spots would be recovered. We show the longitudes and radii for the per-transit selected starspots in Figure 5.14. Note we have excluded the first ∼120 days of data as they only contained long cadence observations and did not have enough sampling in-transit to constrain the starspot properties.

The morphology of the starspot evolution in Figure 5.14 is much more complicated than was seen in the Llama simulation. At all times throughout the data several large starspots are found in the path of the transit. Many of the spots appear to move linearly in longitude with time, while some appear to evolve and possibly merge with other spots. This is the first observation of such high fidelity starspot evolution using optical broadband imaging alone.

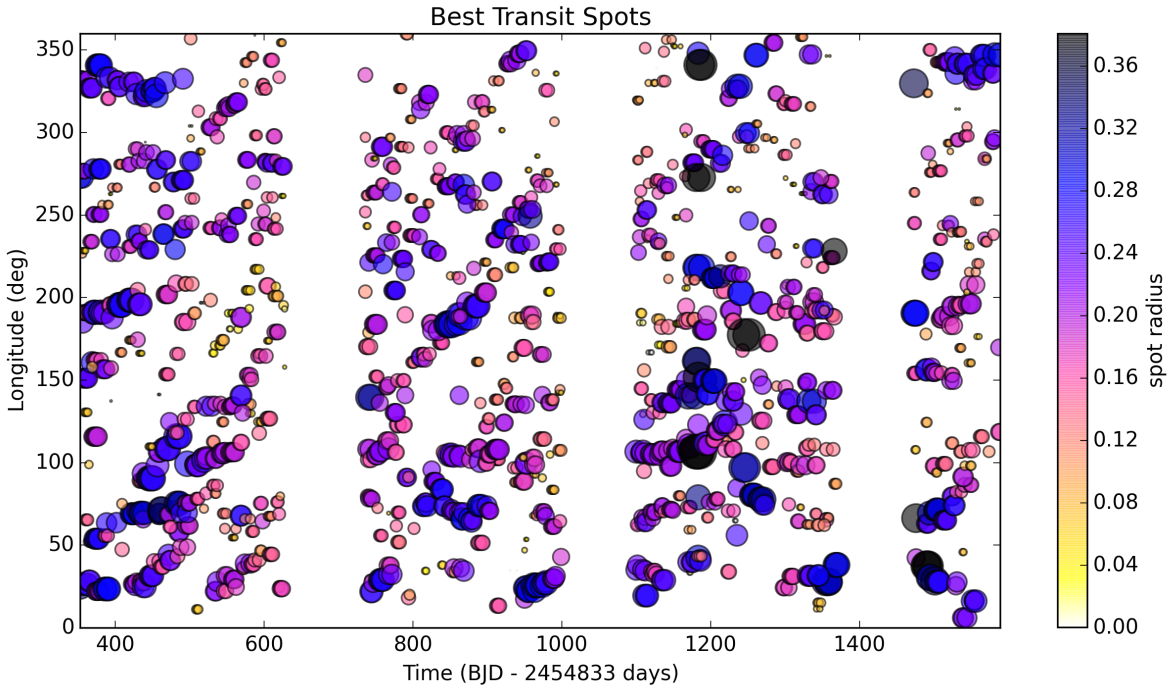

Figure 5.14: Longitude versus time map for starspots that were determined to be in the path of the planet from our 8 spot MCMC models. Both the symbol color (from light to dark) and symbol size (from small to large) indicate the best-fit radius relative to the stellar radius of the starspots.

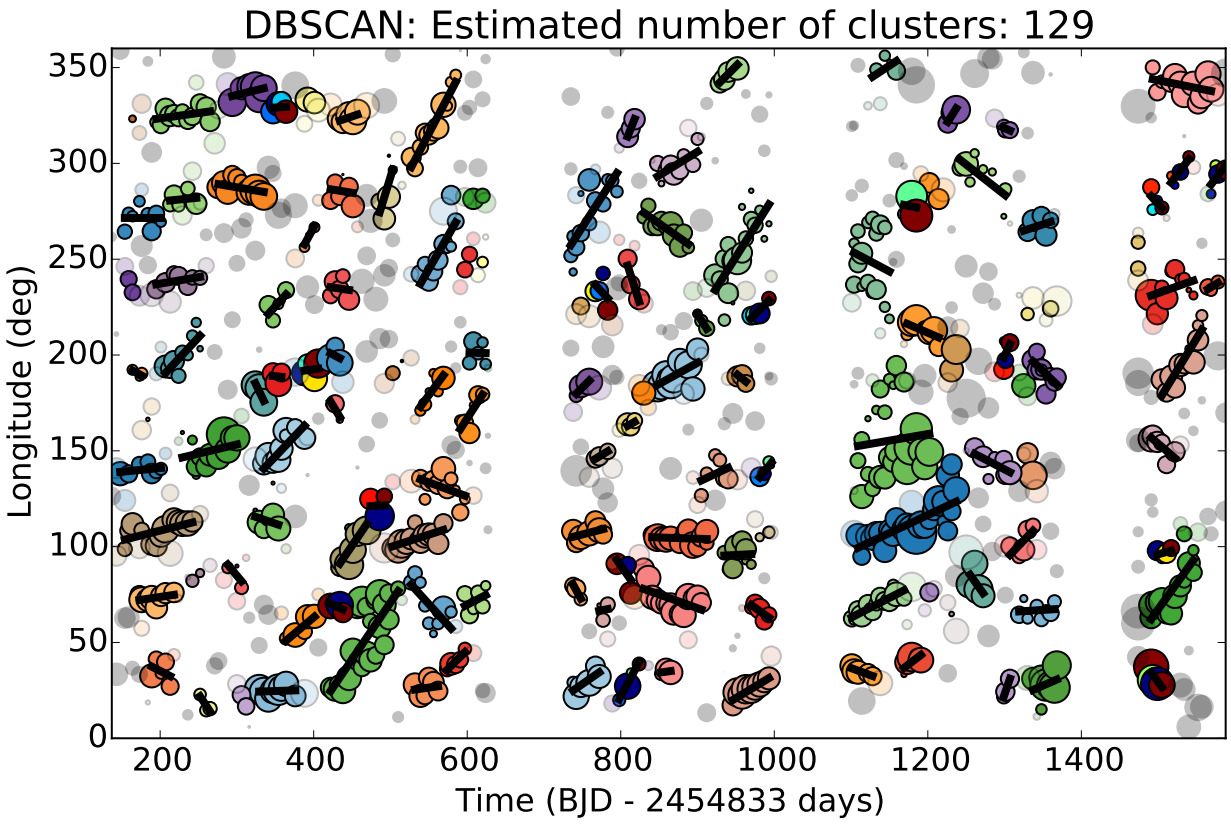

Figure 5.15: The longitude versus time map for starspots in the path of the transit, with the symbol size indicating spot size, as in Figure 5.14. Colored symbols indicate starspot measurements that belong to a cluster from the machine learning DBSCAN clustering algorithm. Grey symbols indicate spot measurements that were not associated with any other spot features from DBSCAN. Spot clusters with at least 5 epochs were fit with a first-order polynomial to quantify their evolution in longitude over time (black lines).

We again used the DBSCAN clustering algorithm from the Python *scikit-learn* machine learning package to generate unique clusters of starspot measurements over time. The results from the DBSCAN clustering for Kepler 17 is shown in Figure 5.15. This analysis recovered 129 individual starspot features, almost an order of magnitude more starspot clusters than in the Llama simulation. For starspot clusters containing at least 5 data points (measurements from at least 5 individual transits) we computed lines of best fit in longitude over time, also shown in Figure 5.15. In total, 107 DBSCAN clusters were fit, which we use to characterize the differential rotation and starspot evolution of Kepler 17.

Unlike the starspot features seen from the out-of-transit light curve modulations with durations of hundreds of days, the in-transit starspots appeared to have much shorter lifetimes. The longer lived starspots observed in first 600 days of the out-of-transit data (Figure 5.13) are evidently not in the path of the transit. The typical (mean) lifetime of the 107 DBSCAN clusters recovered with at least 5 data points was 35 ± 23 days, where the uncertainty is the standard deviation of all cluster lifetimes. By requiring the clusters to have at least 5 data points, the minimum duration would typically be $\sim$7.4 days. However, since multiple starspots within the same time window could be associated within the DBSCAN algorithm, the shortest observed durations were close to 6 days.

For each of the 107 starspot clusters detected with the DBSCAN algorithm and fit with a linear trend, a single rotation period can be determined using the slope of the fit and the phase-folding period of 12.258 days. We determined the rotation period for each of these starspot clusters using Equation 4.2. As with the results from the Llama simulation (Figure 5.8), no correlation between the rotation period and the median latitude for each starspot cluster was found in Figure 5.16 for Kepler 17. We therefore cannot measure the differential rotation law directly from the outputs of our light curve modeling procedure.

Instead we employ a statistical method to constrain the differential rotation coefficient for Kepler 17. In Figure 5.17 we show the distribution of 107 rotation periods determined for the DBSCAN clusters. To minimize the impact of spurious outlier periods from the DBSCAN clusters, we assumed a period range that spans the 5th to 95th percentile of this

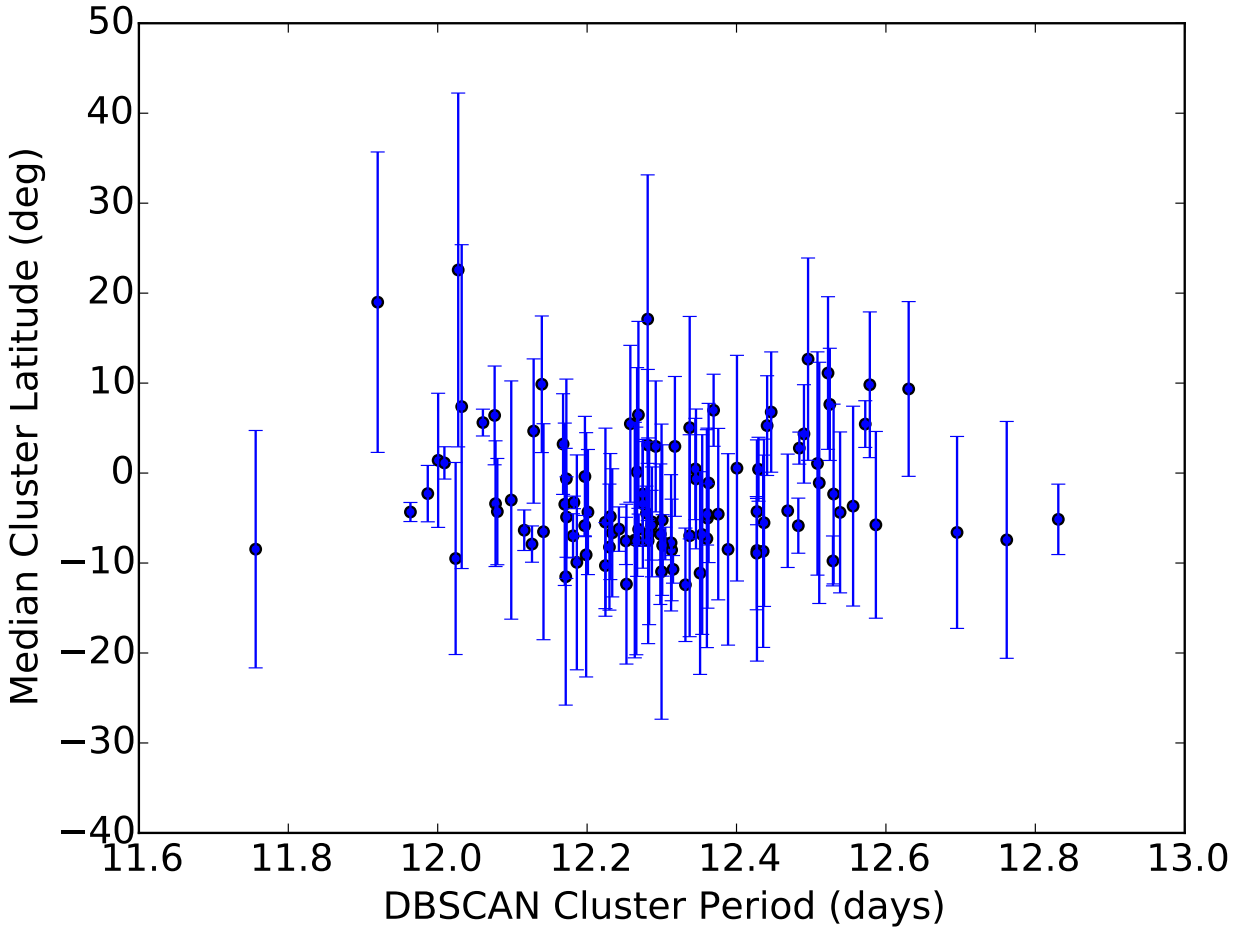

Figure 5.16: Median latitude versus rotation period for the 107 DBSCAN starspot clusters from Figure 5.15. Error bars shown are the standard deviation of the latitudes within each starspot cluster. As in the Llama simulation, no significant correlation is found between the latitude and rotation rate, meaning we cannot directly fit a solar-type differential rotation law to the spot features. This is due to the latitude of the in-transit starspots being poorly constrained.

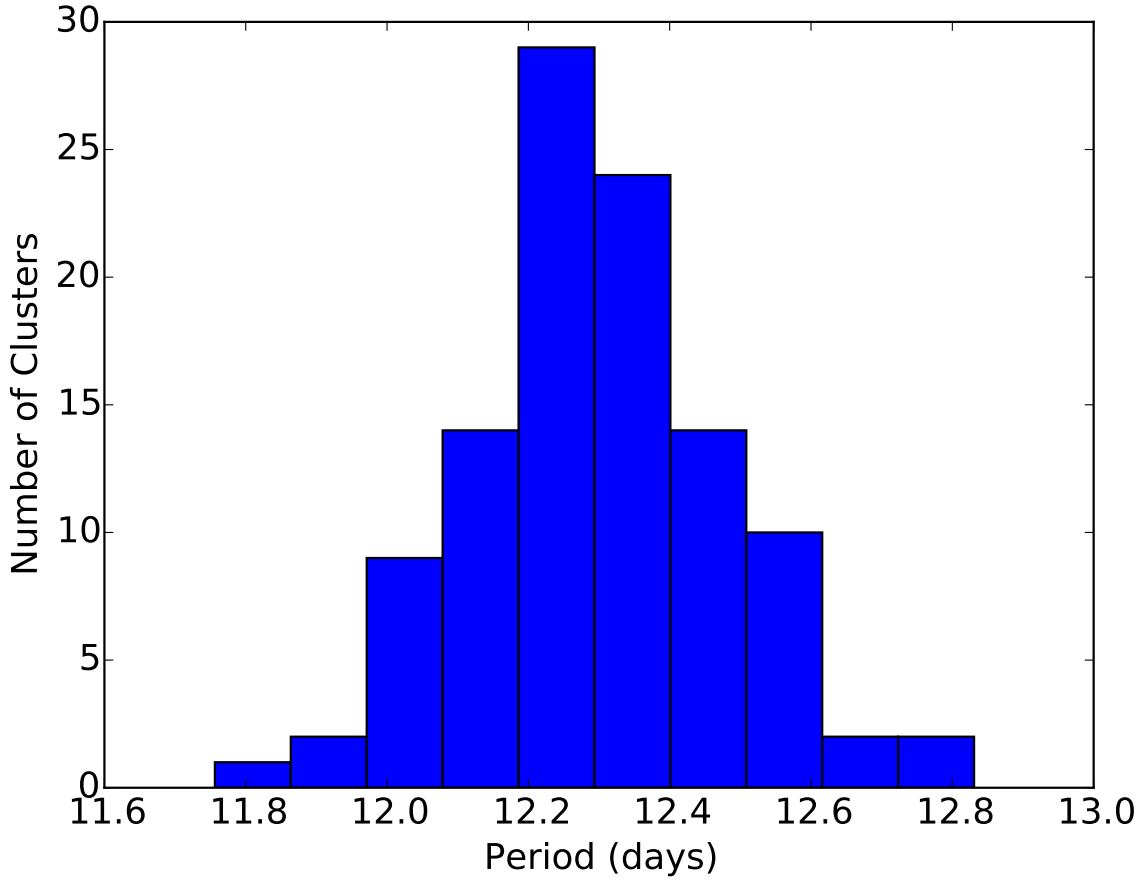

Figure 5.17: Distribution of rotation periods determined from the linear fits to the 107 starspot clusters found in Figure 5.15. For reference, the average out-of-transit rotation period used in our analysis was 12.258 days.

distribution, or between 12.01 and 12.58 days respectively.

Since the planet's orbital inclination is so large ($i = 88.95°$), the lowest latitude that the transit samples is 0° (the stellar equator). Given the large size of the planet ($r_p/r_s = 0.13414$), the and assuming we are again sensitive to large starspots up to $\sim$1° away in latitude grazing the transit path, we can estimate the maximum latitude sampled by the transit path to be $\sim$13.1°. With this range of rotation periods measured and latitudes sampled, we estimate the differential rotation coefficient for Kepler 17 using Equation 5.1 to be $k = 0.88$. This indicates Kepler 17 has less differential rotation than in the Llama simulation, but much stronger shear as a function of latitude than is observed on the Sun (with $k \approx 0.2$). Since the estimated period range came from sampling over 100 starspots, the distribution function correction factor listed in Table 5.1 is negligible.

Interestingly, the 5 in-transit starspots identified by Désert et al. (2011) and shown in Bonomo & Lanza (2012) for Kepler 17 do not seem to agree with our starspot period distribution in Figure 5.17. We traced the linear features in Fig 4. of Bonomo & Lanza (2012) to find an approximate slope of 3.195 deg day^{-1} = 0.00887 phase per day. Using a phase-folding period of 12.01 days (close to the out-of-transit starspot rotation period we found in the first 2 years of data) and our Equation 4.2, we can translate their longitude evolution rate in to a new rotation period for these starspots of 13.44 days. This period is much slower than any of our DBSCAN cluster periods, and is at odds with the spot evolution recovered in our analysis. A slower rotation period along the equator for Kepler 17 is also inconsistent with solar-type differential rotation. Bonomo & Lanza (2012) however state that the equatorial spots rotate *faster* than the mean rotation period, consistent with solar-type differential rotation. If we then assume a sign error occurred in the presentation of Fig 4. of Bonomo & Lanza (2012), the slope of 3.195 deg day^{-1} for the in-transit starspots then becomes a rotation period of 10.85 days. This faster period is consistent with solar-type differential rotation, but is much faster than any feature recovered in our DBSCAN clusters.

In Figure 5.18 we show the maximum radius versus the lifetime for the 107 clusters from our DBSCAN analysis of Kepler 17. The typical radius of these starspots are more than

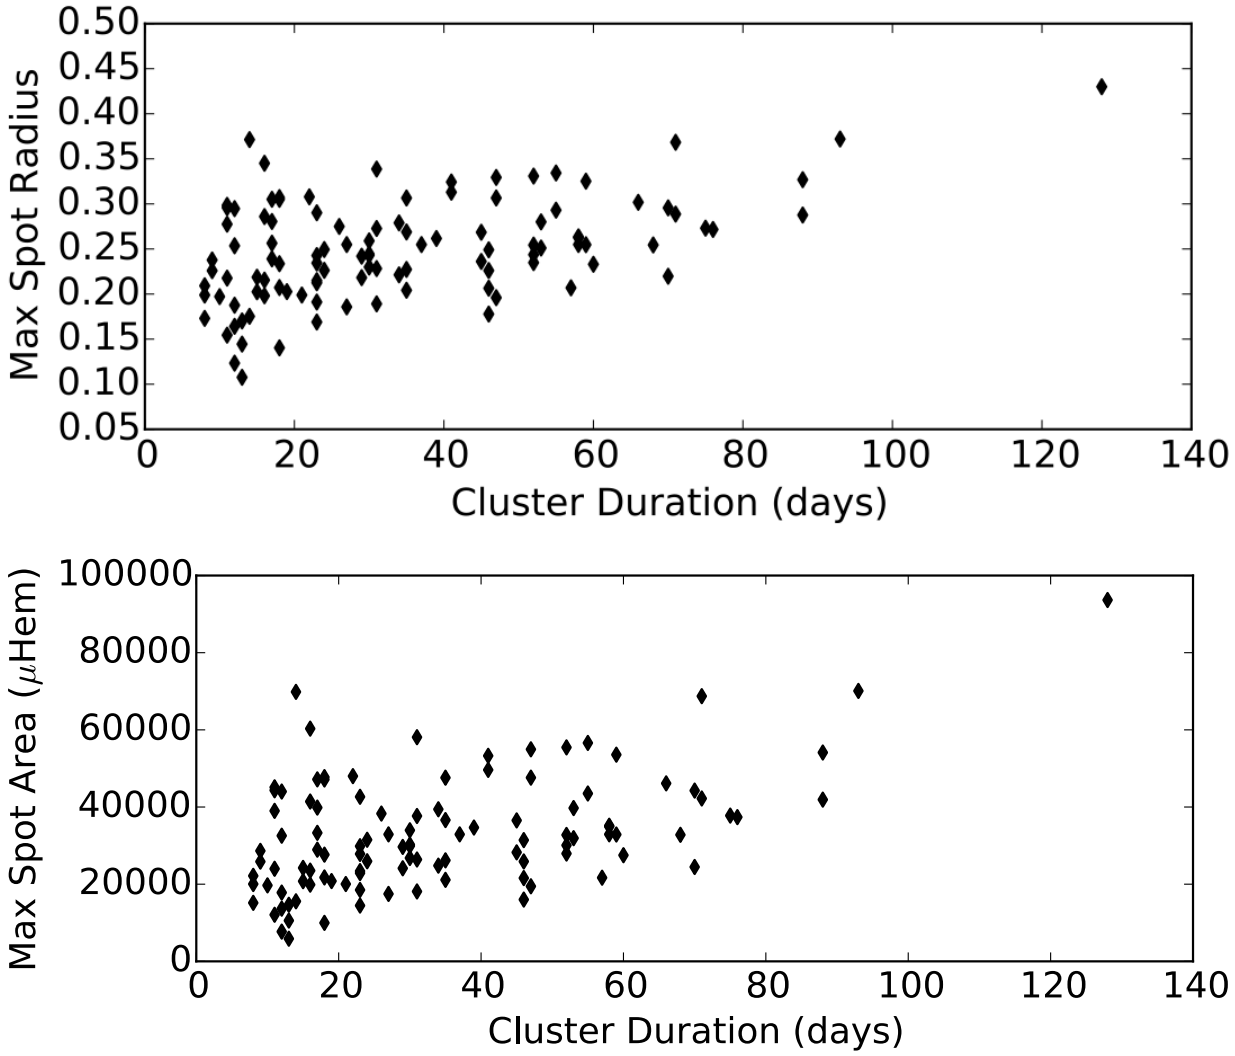

Figure 5.18: Top: Duration of transit-crossing starspot features recovered from our DBSCAN clustering versus the maximum starspot radius determined with each cluster. Bottom: The same data as above, converting the maximum starspot radii to areas, and presented in the traditional solar units of millionths of a solar hemisphere. This trend matches the expectation from the solar "Gnevyshev-Waldmeier rule" of spot lifetimes increasing with surface area. However, the durations may be inaccurate due to larger spots at higher latitudes grazing the transit for less than the full spot lifetime.

twice as large as those recovered in the Llama simulation, and dwarf any spots seen on the Sun. The almost linear correlation between the maximum starspot area and the feature lifetime in Figure 5.18 qualitatively matches the "Gnevyshev-Waldmeier rule" on the Sun.

Most of the starspots traced with the DBSCAN clustering do not show simple emergence and decay. Instead, some clusters appear to monotonically grow in size, while many others appear to not change size significantly over their recovered lifetimes. A few of the starspot clusters do show a more classical evolution, with a rapid emergence and gradual decay as seen in sunspots. The best candidate for this sunspot-like profile is the tan colored cluster located in Figure 5.15 with a center time of 970 days and a median longitude of 25.7°. The area of this starspot, shown in Figure 5.19, appears to decay almost linearly with time. However, this spot's decay was truncated by the end of data collection for *Kepler* Quarter 10. This starspot also exhibits remarkably linear evolution in longitude indicative of differential rotation.

5.7 Summary

In this chapter we have demonstrated the ability to recover starspot evolution and stellar differential rotation by modeling light curves where a transiting exoplanet occults starspots. Modeling the in-transit bumps over time enables a more robust determination of the starspot's rotation period than from tracing the out-of-transit flux modulations. We have tested our light curve fitting code and analysis procedures using a simulated light curve, and then extended this methodology to a *Kepler* planet host. Despite the enhanced ability to trace the longitudes and radii for many more starspots as a function of time, constraining the latitudes beyond simply "in" versus "out" of the transit path remains elusive.

Kepler 17 is very active G dwarf, and is consistent with being a young Solar analogue. Despite 4 years of starspot activity with relative flux modulation amplitudes comparable or greater to that of the active M dwarf GJ 1243, Kepler 17 produced no obvious stellar flares. There were no windows of time within the 4 years of *Kepler* data where Kepler 17 did not appear to have significant starspots present in the path of the transit.

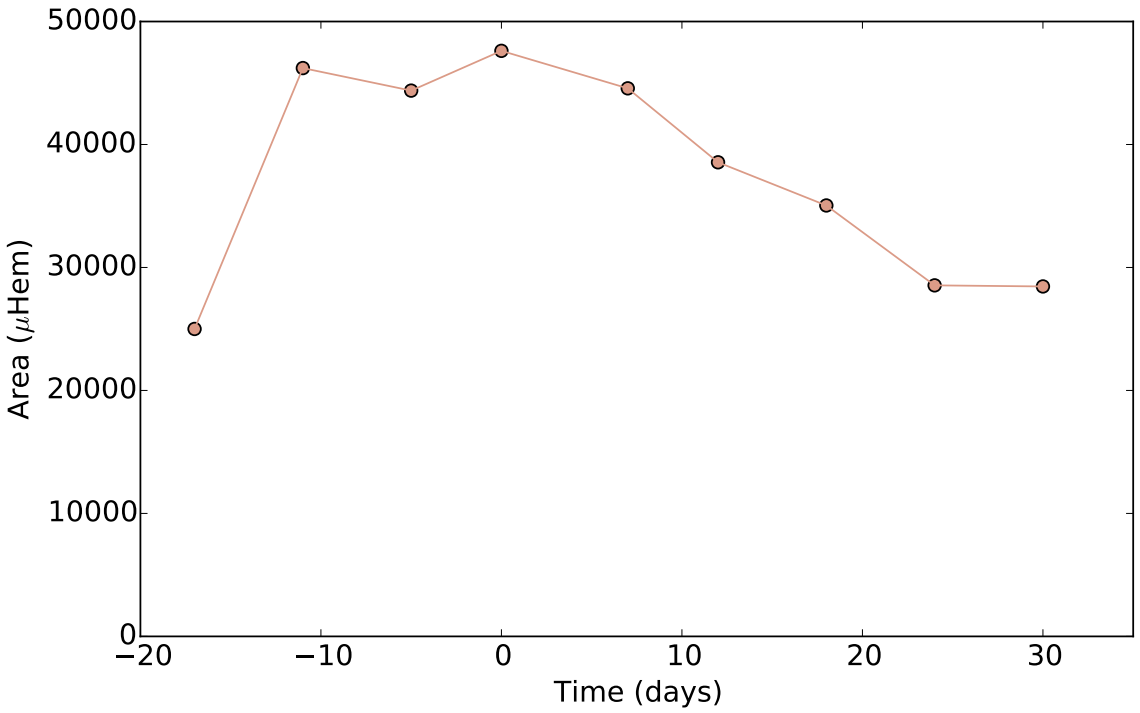

Figure 5.19: Emergence and decay profile for the starspot feature that most closely resembles sunspot-like evolution. This starspot cluster is located in Figure 5.15 with a median time of 970 days and a median longitude of 25.7°. The color used for this figure matches that of the same starspot cluster in Figure 5.15.

Kepler 17 rotates with a mean period of 12.258 days, more than twice as fast as the Sun. The differential rotation coefficient for Kepler 17, statistically estimated from robust maximum and minimum rotation periods measured form the in-transit starspots, was $k = 0.88$. This is approximately four times more differential rotation than is observed on the Sun. This result is contrary to solar-type dynamo models that predict more rapidly rotating stars should have *less* differential rotation.

Tracing of starspots using transiting exoplanets with our code could be extended to other stars in the *Kepler* database, such as Kepler 63 or Hat-P-11. These other targets are known to have strong starspot signatures in both their in-transit and out-of-transit light curves (Sanchis-Ojeda & Winn, 2011; Sanchis-Ojeda et al., 2013). A misalignment between the stellar rotation and planetary orbital axis (as found for Kepler 63) would improve our ability to calculate differential rotation, as the transit would sample different latitude bands over time.

For Kepler 17 a complex family of starspot evolution profiles was recovered. In general, starspots that reach the largest areas live the longest in our analysis, consistent with the evolution profiles of sunspots. Characterizing the detailed decay rates of individual starspots, and correlating the starspot areas versus lifetimes for many spots, may help constrain the stellar diffusion timescale.

Chapter 6

CONCLUSIONS

In this final chapter I briefly summarize and discuss the results from the various investigations contained in each chapter of this thesis. I will then describe some possible avenues for future work in the exciting and timely field of stellar activity.

6.1 Summary

The main theme throughout this thesis has been the study of stellar activity from low-mass stars in time domain photometric databases. Activity here has been characterized by the two most notable phenomena in optical light curves: stellar flares and starspots. The thesis contained four individual studies of active stars in the time domain, two of which focused on flares, and two on starspots. Within each study I have worked to develop new techniques for analyzing temporal data, with an emphasis on making instructive visualizations. Plots like Figure 4.3 have led to "ah ha!" moments, and were incredibly important for my own comprehension of the processes at work.

In Chapter 2, we studied the sparse light curves of M dwarfs in large ground-based optical and near infrared time domain datasets for stellar flares. We produced the first predictions for broadband flare response from M0–M6 stars. The lack of strong flare emission at infrared wavelengths reinforces the utility for exoplanet searches at longer wavelengths around M dwarfs, even for active stars with dramatic optical variability (see also Tofflemire et al., 2012). We also studied the occurrence frequency of flares as a function of spectral type. This showed for the first time, using such a large dataset, that flare rates increase with later spectral type across the "fully convective boundary" (at type ∼M4) for stars at low galactic latitudes. This work helped pave the way for ensemble studies of flares from next

generation multi-wavelength ground-based datasets like LSST, which will study the flare rate throughout the Galaxy. Understanding stellar activity as a function of galactocentric position will in turn help establish a detailed age map of the Milky Way.

In Chapter 3, we changed the scope of our flare studies from the broad to the specific, focusing on a single, highly active M dwarf. I developed a graphical interface (FBEYE) to manually search through *Kepler* light curves and identify flares. By searching 11 months of short cadence data with the help of many colleagues, we assembled a sample of over 6,100 flares for the M dwarf GJ 1243. This is the largest dataset of homogeneously observed flares for any star besides the Sun, and will be an ideal training set for developing new algorithms for automatic flare finding from high cadence light curves. In this chapter I also developed the first empirical white light flare template from such a large sample of precisely measured flares. The temporal morphology showed a clear rapid rise phase, followed by a two-phase cooling curve, and can be used both for characterizing single-peaked flares, and deconstructing complex, multi-peaked flare events.

We then moved from flares to studying stellar activity using starspots. For both Chapters 4 and 5 I focused on making robust constraints of the stellar differential rotation rate for active stars. In Chapter 4, I again studied the rapidly rotating, active M4 dwarf, GJ 1243. The large starspots on its surface were seen to evolve in phase and size very slowly. The primary spot was determined to be almost perfectly stable in phase and size over the entire *Kepler* dataset, while secondary starspot features emerged and decayed on timescales of hundreds of days. GJ 1243 has very little apparent differential rotation, and its starspots live more than 10 times longer than sunspots. We traced the surface differential rotation for GJ 1243 using two unique phase-tracking methods. The first involved computationally expensive light curve modeling, but produced reliable starspot longitudes and sizes over the entire four years of data. The second used an easy to implement 2D Gaussian fitting approach to studying the phase evolution map to produce coarser constraints on the starspot lifetimes and evolution. The latter method for tracing of the starspot phase evolution on GJ 1243 has recently been applied to studies of other rapidly rotating stars in the *Kepler*

dataset (e.g. Lurie et al., 2015).

Lastly, in Chapter 5 we studied starspot evolution on a unique system with a transiting exoplanet. We have demonstrated the ability to trace the simultaneous evolution of at least 8 starspots. The analysis methods were validated using a simulated light curve, and the techniques were then applied to an active, planet-hosting G dwarf, Kepler 17. Understanding the active region lifetimes and locations on the surface, as well as the stellar rotation and differential rotation rates, are key for constraining convective and interface dynamo models. Kepler 17 twice as fast as the Sun, but shows a weaker degree of differential rotation, which is consistent with current interface dynamo models. We demonstrated that larger starspots live longer, as seen on the Sun, and speculated that the stellar diffusion timescale may be recoverable from this type of analysis. We traced the detailed size evolution for many starspot groups, and used machine learning techniques to characterize the longitude evolution. Our study using 4 years of *Kepler* data also found that the exoplanet Kepler 17b has a larger radius than previously published. Our results emphasize that accurate characterization of exoplanets requires detailed treatment of the starspot activity on the host star, which can significantly change the depth and shape of the transit.

6.2 Future Work

In the process of working on a PhD thesis, the generation of new ideas naturally outpaces the opportunity for exploring them. The *Kepler* dataset, much like the SDSS, has been fruitful for generating new ideas, testing new techniques, and discovering unique stellar phenomena (e.g. Kruse & Agol, 2014). In this brief section I present an incomplete list of projects or ideas that have come to mind while working on this thesis, but which I have not yet had time to pursue.

6.2.1 Starspot Activity from Exoplanet Transits Without "Bumps"

There are tight observational requirements for an active, planet hosting star to exhibit in-transit starspot bumps. The bump amplitude may only be $\sim$10% the transit depth, and

therefore close to the photometric uncertainty limit of the *Kepler* observations. Our method requires a large planet around a relatively bright star, creating deep transits and probing a relatively wide range of latitudes. To detect distinct bump features we also require high time cadence observations. The size scale of starspot features that can be mapped using the transit method depends on the ratio of transit duration to the observing cadence (t_{dur}/t_{exp}). *Kepler*'s 30 minute (long cadence) observations are generally too slow to measure these in-transit features, but still give accurate transit depths.

For active stars with transiting planets observed by *Kepler* that are too faint or do not have the required short cadence transit depth to trace starspot bumps, we can use the ensemble information from all transits to constrain the starspot contrast and variability. By aligning and stacking every observed transit for a system, starspot variability can be estimated by the scatter in the transit depth over many years. In Figure 5.12 we found for Kepler 17 that the scatter in-transit was more than 2.5 times greater than the scatter immediately out-of-transit. This method could provide general, first-order information about starspots for hundreds of additional planet candidates, and even for detached stellar binaries. A systematic study of the in-transit versus out-of-transit variability for all *Kepler* planet hosts is a high priority.

Such a study would also constrain the change in transit depth with time, due to starspot evolution. The 4 year light curves from *Kepler* are too short to resolve an entire 11 year Solar-like activity cycle. Modulations in the out-of-transit light curve are typically dominated by the 1-2 largest starspot groups, and may not reflect the total spot coverage on the stellar surface. The depth of transits, however, should unambiguously reflect changes in the stellar surface brightness. By detecting coherent trends in the transit depth over the 4 year *Kepler* data set, we may be able to make progress on identifying and characterizing activity cycles on low-mass stars.

6.2.2 Robust Differential Rotation Constraints for Many Stars

The number of active stars that are appropriate for carrying out our spot modeling, which requires high signal to noise and a large transiting exoplanet, is small even in the *Kepler* database. Also, our out-of-transit MCMC-driven light curve modeling approach to studying starspots is prohibitively slow for analyzing the thousands of stars in the *Kepler* data with starspot activity. Mean rotation periods can be reliably and quickly estimated using ensemble light curve methods like the Lomb-Scargle Periodogram or the autocorrelation function. However, the "Hares and Hounds" experiment we participated in (Aigrain et al., 2015) showed these methods do note reliably recover accurate differential rotation rates for most stars.

Our method of modeling the phase versus time flux map with bivariate Gaussians (see Chapter 4) to determine the approximate starspot evolution is relatively fast, and could be extended to analyze a large sample of rapidly rotating active stars from *Kepler*. Already we have used this approach to characterize the starspot evolution from the active M5+M5 binary system, GJ 1245AB (Lurie et al., 2015). This technique is ideal for stars with rapid rotation periods and long starspot lifetimes, where the phase and size evolution is traced slowly over time.

As a simple demonstration, I searched through the catalog of rotating *Kepler* M dwarfs from McQuillan et al. (2013). I selected stars with colors consistent with M0 or later spectral types, and published rotation periods shorter than 1 day, yielding 24 targets for analysis including GJ 1243. I used only the 30-minute, long cadence light curves for this demonstration. For each light curve I fit and removed low-order polynomials to detrend the quarter-to-quarter variations. The rotation period was then independently determined using the IDL Lomb Normalized Periodogram package (`LNP_TEST`). Two of the systems appeared to be pulsating variables (likely M sub-giants or giant stars that were mis-classified as dwarfs in the Kepler Input Catalog) with no notable phase evolution. At least one pulsating variable appeared to show a slight, sinusoidal like phase evolution across the 4 years of data, possibly

indicating an unseen orbiting companion. Many of the systems, however, looked similar to GJ 1243, with clear starspot evolution and frequent flares. In Figure 6.1 I present four example diagnostic plots that showcase the fascinating range of phase evolution morphologies present in this test. These plots show examples of strong differential rotation, long-lived spots, and many short-lived features that move in phase relative to the mean rotation rate. Systematically characterizing and tracking the phase evolution of the starspot features for every rotating low-mass star in the *Kepler* and K2 datasets will reveal a wealth of basic information on starspot lifetimes and differential rotation across the lower main sequence.

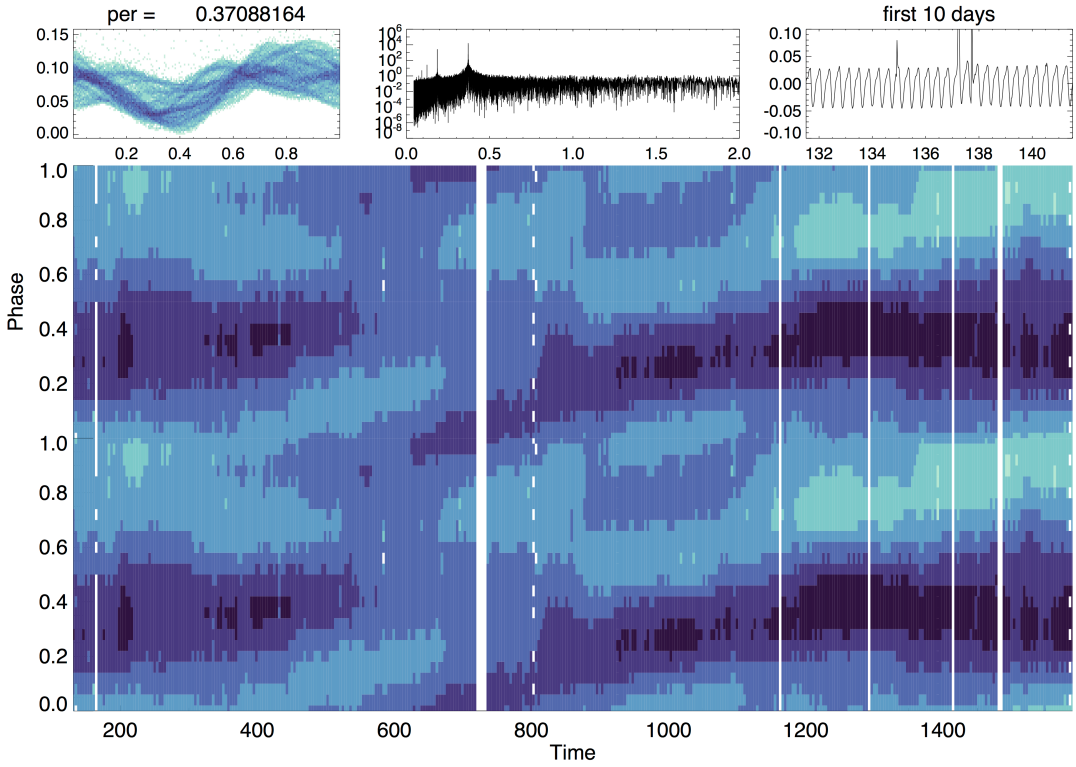

Figure 6.1: A series of diagnostic panels for the low-mass star, KIC 1572802l. Top Left: the phase-folded light curve in relative flux units, folded at the period indicated above (listed as "per", in days). Top Center: the Lomb Normalized Periodogram for the entire 4 year light curve over the range of 0 to 2 day periods. Top Right: the first 10 days of the long cadence light curve, showing clear rotation signal and several flares. Bottom: Continuous phased light curve map for the entire *Kepler* long cadence dataset, where pixel shade indicates the median flux in each (time, phase) bin (dark to light). Pixels span 10 times the rotation period in time, and 4% of the phase. Vertical white bands are due to gaps in the data from *Kepler*.

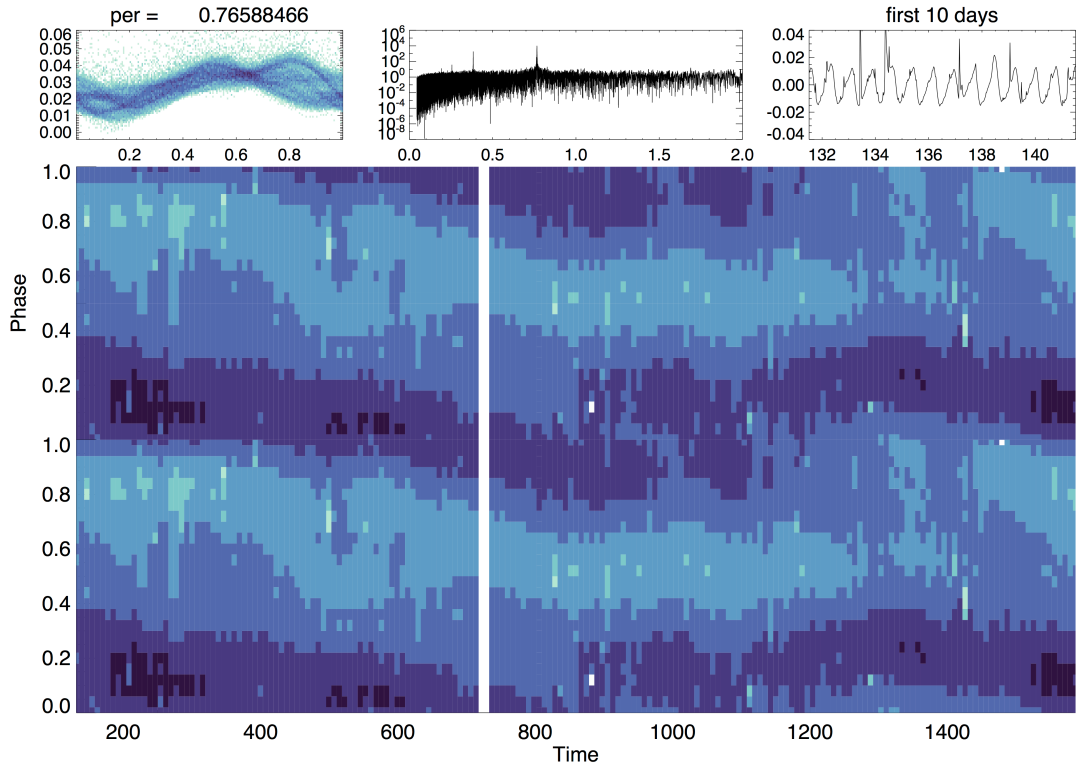

Figure 6.1: Continued – The same series of diagnostic panels for the low-mass star, KIC 5791720l.

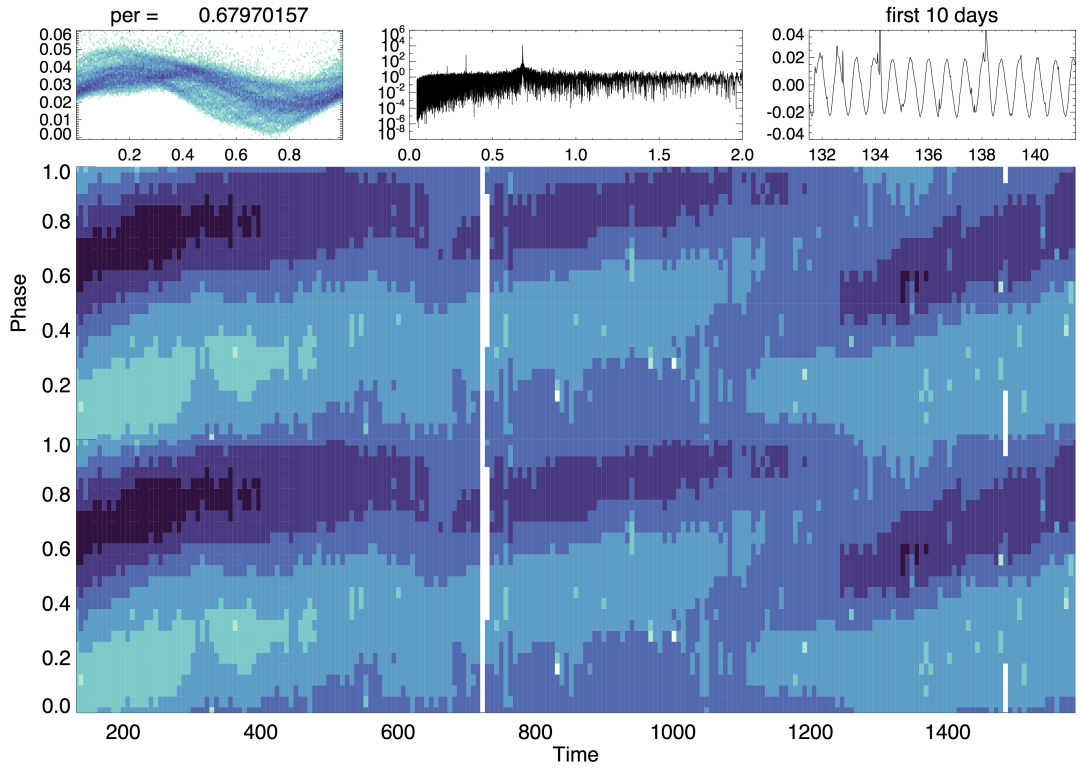

Figure 6.1: Continued – The same series of diagnostic panels for the low-mass star, KIC 9761113l.

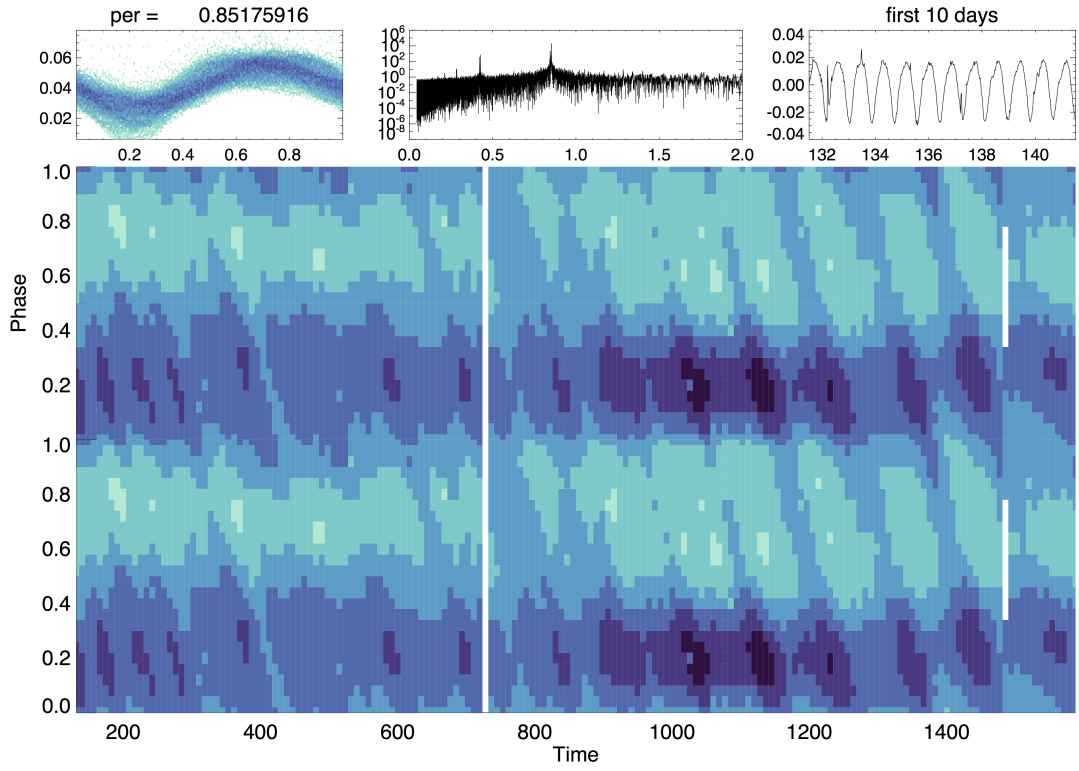

Figure 6.1: Continued – The same series of diagnostic panels for the low-mass star, KIC 12365719l.

6.2.3 Detecting Bright Spots

In Chapter 5 we have exploited the spot-occulting exoplanet transits to trace the sizes and locations of dark features on the surface of Kepler 17. However, as discussed in Chapter 1, bright facular networks are often seen surrounding spots on the Sun. Despite their importance to the total solar irradiance, these bright features have largely evaded detection for stars beside the Sun. Transiting exoplanets that occulted facular regions would show temporary "dips" in the transit shape, the opposite of the "bumps" seen for dark spot groups. Since faculae decay at slower rates than starspots, "dips" may linger in the transit path long after the spot "bumps" have faded.

For faculae (and plage) surrounding large starspots, the detected transit shape may reveal these dips before and after the bumps. Using many transits from systems like Kepler 17, it might be possible to detect this structure. First we must subtract the smooth transit profile light curve from every transit, for example using the Mandel & Agol (2002) model. The large in-transit bumps (e.g. those identified by our spot modeling code) would then be aligned and co-added. If a faint "W" shape emerges around each transit, this would signify the presence of bright features leading and trailing the dark starspots along the path of the transit.

Further, if the total stellar irradiance increase during times of high starspot activity due to faculae, as seen for the Sun, we might also use the mean transit depth to infer the presence of bright regions. For a Sun-like star, as the out-of-transit starspot activity grows over time, the mean transit depth would *increase*, indicating a net increase in the total surface brightness. If the transit depths *decreased* with increasing out-of-transit modulations, then the magnetic activity for the star must be dominated by the cool spot regions alone, and less contribution from faculae and plage as seen on the Sun.

6.3 Closing Thoughts

During my time in graduate school, the *Kepler* mission has opened a new era of ultra-precise time domain photometric studies of stars, making this is a thrilling and unique time for studying their magnetic activity. Through this and other planned surveys, we may finally answer some of the most basic questions about the history of our own star: Is the Sun a typical star? What was it like when it was younger? What will it be like when it is older? When did the Milky Way's stellar disk form? The scale of these planned surveys will also enable us to place the Sun and its activity in a truly Galactic context. New observations of stellar clusters with K2 and TESS, and sparsely sampled photometry from LSST combined with kinematic data from Gaia, will build a connection between stellar age, activity, and rotation for entire populations of stars. *Kepler* and SDSS have repeatedly reminded us to ponder more deeply and expect the unexpected, and so I close with a reminder to myself:

"You can't depend on your eyes when your imagination is out of focus" – Mark Twain

BIBLIOGRAPHY

Abazajian, K. N., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2009, ApJS, 182, 543

Abbett, W. P., & Hawley, S. L. 1999, ApJ, 521, 906

Adelman-McCarthy, J. K., Agüeros, M. A., Allam, S. S., et al. 2008, ApJS, 175, 297

Aigrain, S., Llama, J., Ceillier, T., et al. 2015, MNRAS, 450, 3211

Aihara, H., Allende Prieto, C., An, D., et al. 2011, ApJS, 193, 29

Allred, J. C., Hawley, S. L., Abbett, W. P., & Carlsson, M. 2006, ApJ, 644, 484

Andrews, A. D. 1965, Irish Astronomical Journal, 7, 20

Anfinogentov, S., Nakariakov, V. M., Mathioudakis, M., Van Doorsselaere, T., & Kowalski, A. F. 2013, ApJ, 773, 156

Ankerst, M., Breunig, M. M., Kriegel, H.-P., & Sander, J. 1999, in SIGMOD 1999, Proceedings ACM SIGMOD International Conference on Management of Data, June 1-3, 1999, Philadephia, Pennsylvania, USA, ed. A. Delis, C. Faloutsos, & S. Ghandeharizadeh (ACM Press), 49–60

Babcock, H. W. 1961, ApJ, 133, 572

Barnes, J. R., Collier Cameron, A., Donati, J.-F., et al. 2005, MNRAS, 357, L1

Barnes, J. R., James, D. J., & Collier Cameron, A. 2004, MNRAS, 352, 589

Basri, G., Walkowicz, L. M., Batalha, N., et al. 2010, ApJ, 713, L155

Becker, A. C., Bochanski, J. J., Hawley, S. L., et al. 2011, ApJ, 731, 17

Becker, A. C., Agol, E., Silvestri, N. M., et al. 2008, MNRAS, 386, 416

Benz, A. O. 2008, Living Reviews in Solar Physics, 5, 1

Berdyugina, S. V. 2005, Living Reviews in Solar Physics, 2, 8

Berger, E., Leibler, C. N., Chornock, R., et al. 2013, ApJ, 779, 18

Bertello, L., Pevtsov, A. A., & Pietarila, A. 2012, ApJ, 761, 11

Bochanski, J. J., Hawley, S. L., Covey, K. R., et al. 2010, AJ, 139, 2679

Bochanski, J. J., Munn, J. A., Hawley, S. L., et al. 2007a, AJ, 134, 2418

Bochanski, J. J., West, A. A., Hawley, S. L., & Covey, K. R. 2007b, AJ, 133, 531

Bonomo, A. S., & Lanza, A. F. 2012, A&A, 547, A37

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Science, 327, 977

Brandenburg, A., & Subramanian, K. 2005, Phys. Rep., 417, 1

Browning, M. K. 2008, ApJ, 676, 1262

Browning, M. K., Basri, G., Marcy, G. W., West, A. A., & Zhang, J. 2010, AJ, 139, 504

Bumba, V. 1963, Bulletin of the Astronomical Institutes of Czechoslovakia, 14, 91

Butler, C. J., Rodono, M., & Foing, B. H. 1988, A&A, 206, L1

Butters, O. W., West, R. G., Anderson, D. R., et al. 2010, A&A, 520, L10

Byrne, P. B. 1989, Sol. Phys., 121, 61

Carrington, R. C. 1859, MNRAS, 20, 13

Chabrier, G., & Küker, M. 2006, A&A, 446, 1027

Chaplin, W. J., & Miglio, A. 2013, ARA&A, 51, 353

Chaplin, W. J., Appourchaux, T., Elsworth, Y., et al. 2010, ApJ, 713, L169

Charbonneau, P. 2005, Living Reviews in Solar Physics, 2, 2

Christensen, U. R., Holzwarth, V., & Reiners, A. 2009, Nature, 457, 167

Christensen-Dalsgaard, J., & Frandsen, S. 1983, Sol. Phys., 82, 469

Claret, A., & Bloemen, S. 2011, A&A, 529, A75

Cohen, M., Wheaton, W. A., & Megeath, S. T. 2003, AJ, 126, 1090

Collier Cameron, A. 2007, Astronomische Nachrichten, 328, 1030

Covey, K. R., Ivezić, Ž., Schlegel, D., et al. 2007, AJ, 134, 2398

Cushing, M. C., Rayner, J. T., & Vacca, W. D. 2005, ApJ, 623, 1115

Czesla, S., Huber, K. F., Wolter, U., Schröter, S., & Schmitt, J. H. M. M. 2009, A&A, 505, 1277

Davenport, J. R. A., Becker, A. C., Kowalski, A. F., et al. 2012, ApJ, 748, 58

Davenport, J. R. A., Bochanski, J. J., Covey, K. R., et al. 2007, AJ, 134, 2430

Davenport, J. R. A., Hebb, L., & Hawley, S. L. 2015a, ApJ, 806, 212

Davenport, J. R. A., Ruan, J. J., Becker, A. C., Macleod, C. L., & Cutri, R. M. 2015b, ApJ, 803, 2

Davenport, J. R. A., Hawley, S. L., Hebb, L., et al. 2014, ApJ, 797, 122

Delfosse, X., Forveille, T., Ségransan, D., et al. 2000, A&A, 364, 217

Désert, J.-M., Charbonneau, D., Demory, B.-O., et al. 2011, ApJS, 197, 14

Dobler, W., Stix, M., & Brandenburg, A. 2006, ApJ, 638, 336

Donati, J.-F., & Brown, S. F. 1997, A&A, 326, 1135

Donati, J.-F., & Collier Cameron, A. 1997, MNRAS, 291, 1

Durney, B. R., De Young, D. S., & Roxburgh, I. W. 1993, Sol. Phys., 145, 207

Eddy, J. A. 1976, Science, 192, 1189

Eker, Z. 1996, ApJ, 473, 388

Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. 1996, in Proc. of 2nd International Conference on Knowledge Discovery and, 226–231

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306

Foukal, P. 1993, Sol. Phys., 148, 219

Foukal, P., & Lean, J. 1986, ApJ, 302, 826

Frasca, A., Biazzo, K., Taş, G., Evren, S., & Lanzafame, A. C. 2008, A&A, 479, 557

Frasca, A., Fröhlich, H.-E., Bonanno, A., et al. 2011, A&A, 532, A81

Friedman, J. H. 1984, A Variable Span Smoother, Tech. Rep. 5, Department of Statistics, Stanford University

Fuhrmeister, B., Liefke, C., Schmitt, J. H. M. M., & Reiners, A. 2008, A&A, 487, 293

Fukugita, M., Ichikawa, T., Gunn, J. E., et al. 1996, AJ, 111, 1748

Gastine, T., Morin, J., Duarte, L., et al. 2013, A&A, 549, L5

Gershberg, R. E. 1972, Ap&SS, 19, 75

Gershberg, R. E. 1975, in IAU Symposium, Vol. 67, Variable Stars and Stellar Evolution, ed. V. E. Sherwood & L. Plaut, 47–64

Gershberg, R. E., & Shakhovskaya, N. I. 1973, Nature Physical Science, 242, 85

Giampapa, M. S., Cram, L. E., & Wild, W. J. 1989, ApJ, 345, 536

Gizis, J. E., Burgasser, A. J., Berger, E., et al. 2013, ApJ, 779, 172

Gondoin, P. 2008, A&A, 478, 883

Green, J. L., & Boardsen, S. 2006, Advances in Space Research, 38, 130

Grigis, P. C., & Benz, A. O. 2005, ApJ, 625, L143

Hall, D. S. 1994, International Amateur-Professional Photoelectric Photometry Communications, 54, 1

Hall, D. S., & Henry, G. W. 1994, International Amateur-Professional Photoelectric Photometry Communications, 55, 51

Hathaway, D. H., & Choudhary, D. P. 2008, Sol. Phys., 250, 269

Hawley, S. L., Davenport, J. R. A., Kowalski, A. F., et al. 2014, ApJ, 797, 121

Hawley, S. L., & Fisher, G. H. 1992, ApJS, 78, 565

Hawley, S. L., & Pettersen, B. R. 1991, ApJ, 378, 725

Hawley, S. L., Fisher, G. H., Simon, T., et al. 1995, ApJ, 453, 464

Hawley, S. L., Allred, J. C., Johns-Krull, C. M., et al. 2003, ApJ, 597, 535

Hayden, B. T., Garnavich, P. M., Kessler, R., et al. 2010, ApJ, 712, 350

Henry, G. W., Eaton, J. A., Hamer, J., & Hall, D. S. 1995, ApJS, 97, 513

Hilton, E. J. 2011, PhD thesis, University of Washington

Hilton, E. J., Hawley, S. L., Kowalski, A. F., & Holtzman, J. 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, 197

Hilton, E. J., West, A. A., Hawley, S. L., & Kowalski, A. F. 2010, AJ, 140, 1402

Hoge, E. R. 1947, PASP, 59, 109

Houdebine, E. R., Butler, C. J., Panagi, P. M., Rodono, M., & Foing, B. H. 1991, A&AS, 87, 33

Hoyt, D. V., Schatten, K. H., & Nesme-Ribes, E. 1994, Geophys. Res. Lett., 21, 2067

Hull, C. L. H., Plambeck, R. L., Bolatto, A. D., et al. 2013, ApJ, 768, 159

Hunt-Walker, N. M., Hilton, E. J., Kowalski, A. F., Hawley, S. L., & Matthews, J. M. 2012, PASP, 124, 545

Irwin, J., Berta, Z. K., Burke, C. J., et al. 2011, ApJ, 727, 56

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

Ivezić, Ž., Smith, J. A., Miknaitis, G., et al. 2007, AJ, 134, 973

Ivezić, Ž., Tyson, J. A., Acosta, E., et al. 2008, ArXiv e-prints, # 0805.2366, arXiv:0805.2366

Javaraiah, J. 2011, Sol. Phys., 270, 463

—. 2015, New A, 34, 54

Jenkins, J. S., Ramsey, L. W., Jones, H. R. A., et al. 2009, ApJ, 704, 975

Jenkins, L. F. 1952, General catalogue of trigonometric stellar parallaxes., ed. Jenkins, L. F.

Joy, A. H., & Humason, M. L. 1949, PASP, 61, 133

Karoff, C., Campante, T. L., Ballot, J., et al. 2013, ApJ, 767, 34

Kashyap, V. L., Drake, J. J., Güdel, M., & Audard, M. 2002, ApJ, 580, 1118

Kipping, D. M. 2012, MNRAS, 427, 2487

Kiraga, M., & Stepien, K. 2007, Acta Astron., 57, 149

Kitchatinov, L. L., & Olemskoy, S. V. 2011, MNRAS, 411, 1059

Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, ApJ, 713, L79

Kopal, Z. 1982, Ap&SS, 87, 149

Kosovichev, A. G., & Fedorova, A. V. 1991, Soviet Ast., 35, 507

Kowalski, A. F., Hawley, S. L., Hilton, E. J., et al. 2009, AJ, 138, 633

Kowalski, A. F., Hawley, S. L., Holtzman, J. A., Wisniewski, J. P., & Hilton, E. J. 2010, ApJ, 714, L98

Kowalski, A. F., Hawley, S. L., Wisniewski, J. P., et al. 2013, ApJS, 207, 15

Kowalski, A. F., Mathioudakis, M., Hawley, S. L., et al. 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, 1157

Krause, F., & Ruediger, G. 1975, Sol. Phys., 42, 107

Kron, G. E. 1947, PASP, 59, 261

—. 1950a, AJ, 55, 69

—. 1950b, Leaflet of the Astronomical Society of the Pacific, 6, 52

Kruse, E., & Agol, E. 2014, Science, 344, 275

Küker, M., & Rüdiger, G. 2005, Astronomische Nachrichten, 326, 265

—. 2008, Journal of Physics Conference Series, 118, 012029

—. 2011, Astronomische Nachrichten, 332, 933

Kunkel, W. E. 1970, ApJ, 161, 503

—. 1974, Nature, 248, 571

Kurtz, D. W., Saio, H., Takata, M., et al. 2014, MNRAS, 444, 102

Lacy, C. H., Moffett, T. J., & Evans, D. S. 1976, ApJS, 30, 85

Leighton, R. B. 1964, ApJ, 140, 1547

Lépine, S., & Shara, M. M. 2005, AJ, 129, 1483

Litzkow, M., Livny, M., & Mutka, M. 1988, in Proceedings of the 8th International Conference of Distributed Computing Systems

Llama, J., Jardine, M., Mackay, D. H., & Fares, R. 2012, MNRAS, 422, L72

Loebman, S. R., Roškar, R., Debattista, V. P., et al. 2011, ApJ, 737, 8

Lomb, N. R. 1976, Ap&SS, 39, 447

Lovell, B. 1969, Nature, 222, 1126

Loyd, R. O. P., & France, K. 2014, ApJS, 211, 9

Lurie, J. C., Davenport, J. R. A., Hawley, S. L., et al. 2015, ApJ, 800, 95

MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, ApJ, 721, 1014

Madrid, J. P., & Macchetto, D. 2009, in Bulletin of the American Astronomical Society, Vol. 41, Bulletin of the American Astronomical Society, 913–914

Mandel, K., & Agol, E. 2002, ApJ, 580, L171

Markwardt, C. B. 2009, in Astronomical Society of the Pacific Conference Series, Vol. 411, Astronomical Data Analysis Software and Systems XVIII, ed. D. A. Bohlender, D. Durand, & P. Dowler, 251

Martens, P. C. H., & Kuin, N. P. M. 1989, Sol. Phys., 122, 263

Mazumdar, A., & Antia, H. M. 2001, A&A, 368, L8

McQuillan, A., Aigrain, S., & Mazeh, T. 2013, MNRAS, 432, 1203

McQuillan, A., Mazeh, T., & Aigrain, S. 2014, ApJS, 211, 24

Miller, C. J., Genovese, C., Nichol, R. C., et al. 2001, AJ, 122, 3492

Moffett, T. J. 1974, ApJS, 29, 1

Morgan, W. W., & Keenan, P. C. 1973, ARA&A, 11, 29

Morgan, W. W., Keenan, P. C., & Kellman, E. 1943, An atlas of stellar spectra, with an outline of spectral classification

Morin, J., Donati, J.-F., Petit, P., et al. 2010, MNRAS, 407, 2269

—. 2008a, MNRAS, 390, 567

Morin, J., Donati, J.-F., Forveille, T., et al. 2008b, MNRAS, 384, 77

N. Reid & S. L. Hawley, ed. 2000, New light on dark stars : red dwarfs, low mass stars, brown dwarfs

Nielsen, M. B., Gizon, L., Schunker, H., & Schou, J. 2014, A&A, 568, L12

Norton, A. A., Jones, E. H., & Liu, Y. 2013, Journal of Physics Conference Series, 440, 012038

Notsu, Y., Shibayama, T., Maehara, H., et al. 2013, ApJ, 771, 127

Noyes, R. W., Hartmann, L. W., Baliunas, S. L., Duncan, D. K., & Vaughan, A. H. 1984, ApJ, 279, 763

Nutzman, P., & Charbonneau, D. 2008, PASP, 120, 317

O'Neal, D., Neff, J. E., Saar, S. H., & Cuntz, M. 2004, AJ, 128, 1802

O'Neal, D., Neff, J. E., Saar, S. H., & Mines, J. K. 2001, AJ, 122, 1954

Padmanabhan, N., Schlegel, D. J., Finkbeiner, D. P., et al. 2008, ApJ, 674, 1217

Panagi, P. M., & Andrews, A. D. 1995, MNRAS, 277, 423

Parker, E. N. 1955, ApJ, 122, 293

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825

Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, ApJ, 517, 565

Petrovay, K., & van Driel-Gesztelyi, L. 1997, Sol. Phys., 176, 249

Pettersen, B. R., Coleman, L. A., & Evans, D. S. 1984, ApJS, 54, 375

Pettersen, B. R., Hawley, S. L., & Fisher, G. H. 1992, Sol. Phys., 142, 197

Pickles, A. J. 1998, PASP, 110, 863

Plavchan, P., Gee, A. H., Stapelfeldt, K., & Becker, A. 2008a, ApJ, 684, L37

Plavchan, P., Jura, M., Kirkpatrick, J. D., Cutri, R. M., & Gallagher, S. C. 2008b, ApJS, 175, 191

Quintana, E. V., Barclay, T., Raymond, S. N., et al. 2014, Science, 344, 277

Rabus, M., Alonso, R., Belmonte, J. A., et al. 2009, A&A, 494, 391

Ramsay, G., Doyle, J. G., Hakala, P., et al. 2013, MNRAS, 434, 2451

Rau, A., Ofek, E. O., Kulkarni, S. R., et al. 2008, ApJ, 682, 1205

Rayner, J. T., Cushing, M. C., & Vacca, W. D. 2009, ApJS, 185, 289

Reid, I. N., & Hawley, S. L. 2005, New light on dark stars : red dwarfs, low-mass stars, brown dwarfs, ed. Reid, I. N. & Hawley, S. L., doi:10.1007b98494

Reiners, A. 2006, A&A, 446, 267

Reiners, A., & Basri, G. 2007, ApJ, 656, 1121

—. 2009, A&A, 496, 787

Reiners, A., & Schmitt, J. H. M. M. 2003, A&A, 398, 647

Reinhold, T., & Reiners, A. 2013, A&A, 557, A11

Reinhold, T., Reiners, A., & Basri, G. 2013, A&A, 560, A4

Robinson, R. D., Carpenter, K. G., & Percival, J. W. 1999, ApJ, 516, 916

Robinson, R. D., Carpenter, K. G., Percival, J. W., & Bookbinder, J. A. 1995, ApJ, 451, 795

Rodonò, M., & Cutispoto, G. 1988, in Astrophysics and Space Science Library, Vol. 143, Activity in Cool Star Envelopes, ed. O. Havnes, J. E. Solheim, B. R. Pettersen, & J. H. M. M. Schmitt , 163–166

Roettenbacher, R. M., Monnier, J. D., Harmon, R. O., Barclay, T., & Still, M. 2013, ApJ, 767, 60

Sanchis-Ojeda, R., & Winn, J. N. 2011, ApJ, 743, 61

Sanchis-Ojeda, R., Winn, J. N., Marcy, G. W., et al. 2013, ApJ, 775, 54

Savanov, I. S., & Dmitrienko, E. S. 2011, Astronomy Reports, 55, 890

Scargle, J. D. 1982, ApJ, 263, 835

Schatzman, E. 1959, in IAU Symposium, Vol. 9, URSI Symp. 1: Paris Symposium on Radio Astronomy, ed. R. N. Bracewell, 552

Scherrer, P. H., Schou, J., Bush, R. I., et al. 2012, Sol. Phys., 275, 207

Schmidt, S. J., Kowalski, A. F., Hawley, S. L., et al. 2012, ApJ, 745, 14

Schrijver, C. J. 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, 231

Schrijver, C. J., & Zwaan, C. 2000, Solar and Stellar Magnetic Activity (New York : Cambridge University Press, 2000 (Cambridge astrophysics series ; 34))

Schrijver, C. J., Beer, J., Baltensperger, U., et al. 2012, Journal of Geophysical Research (Space Physics), 117, 8103

Semel, M. 1989, A&A, 225, 456

Sesar, B., Ivezić, Ž., Lupton, R. H., et al. 2007, AJ, 134, 2236

Shakhovskaya, N. I. 1989, Sol. Phys., 121, 375

Shibayama, T., Maehara, H., Notsu, S., et al. 2013, ApJS, 209, 5

Silva, A. V. R. 2003, ApJ, 585, L147

Silva Aguirre, V., Chaplin, W. J., Ballot, J., et al. 2011, ApJ, 740, L2

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163

Slawson, R. W., Prša, A., Welsh, W. F., et al. 2011, AJ, 142, 160

Smith, J. C., Stumpe, M. C., Van Cleve, J. E., et al. 2012, PASP, 124, 1000

Solomon, L. H. 1966, SAO Special Report, 210

Spiegel, E. A., & Zahn, J.-P. 1992, A&A, 265, 106

Stephenson, F. R., & Willis, D. M. 1999, Astronomy & Geophysics, 40, 6.21

Stetson, P. B. 1996, PASP, 108, 851

Strassmeier, K. G. 2009, A&A Rev., 17, 251

Szabó, R., Ivezić, Ž., Kiss, L. L., et al. 2014, ApJ, 780, 92

Thain, D., Tannenbaum, T., & Livny, M. 2005, Concurrency - Practice and Experience, 17, 323

Thompson, M. J., Christensen-Dalsgaard, J., Miesch, M. S., & Toomre, J. 2003, ARA&A, 41, 599

Tofflemire, B. M., Wisniewski, J. P., Kowalski, A. F., et al. 2012, AJ, 143, 12

Tovmassian, H. M., Zalinian, V. P., Silant'ev, N. A., Cardona, O., & Chavez, M. 2003, A&A, 399, 647

Tucker, D. L., Kent, S., Richmond, M. W., et al. 2006, Astronomische Nachrichten, 327, 821

van Altena, W. F., Lee, J. T., & Hoffleit, E. D. 1995, The general catalogue of trigonometric [stellar] parallaxes, ed. van Altena, W. F., Lee, J. T., & Hoffleit, E. D.

van Leeuwen, F. 2007, A&A, 474, 653

van Maanen, A. 1940, ApJ, 91, 503

Vaquero, J. M. 2007, Advances in Space Research, 40, 929

Vogt, S. S. 1981, in The Physics of Sunspots, ed. L. E. Cram & J. H. Thomas, 455–479

Vogt, S. S., Penrod, G. D., & Hatzes, A. P. 1987, ApJ, 321, 496

Walker, G., Matthews, J., Kuschnig, R., et al. 2003, PASP, 115, 1023

Walkowicz, L. M., Basri, G., Batalha, N., et al. 2011, AJ, 141, 50

Warren, H. P. 2006, ApJ, 637, 522

Welch, D. L., & Stetson, P. B. 1993, AJ, 105, 1813

West, A. A., Hawley, S. L., Bochanski, J. J., et al. 2008, AJ, 135, 785

West, A. A., Morgan, D. P., Bochanski, J. J., et al. 2011, AJ, 141, 97

Willis, D. M., & Stephenson, F. R. 2001, Annales Geophysicae, 19, 289

Willson, R. C., & Hudson, H. S. 1988, Nature, 332, 810

Wittmann, A. D., & Xu, Z. T. 1987, A&AS, 70, 83

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868

Wright, N. J., Drake, J. J., Mamajek, E. E., & Henry, G. W. 2011, ApJ, 743, 48

Xu, Y., Cao, W., Liu, C., et al. 2006, ApJ, 641, 1210

Yadav, R. K., Gastine, T., Christensen, U. R., & Reiners, A. 2015, A&A, 573, A68

York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, AJ, 120, 1579

Zacharias, N., Finch, C. T., Girard, T. M., et al. 2013, AJ, 145, 44

VITA

James R. A. Davenport was born in Charlottesville, VA in 1983 to loving parents, Joan and Emery. He graduated from Naches Valley High School in 2002. James earned dual BS degrees in Astronomy and Physics from the University of Washington in 2007, and a MS in Astronomy from San Diego State University in 2009, after which he returned to the University of Washington to pursue his PhD. He completed his doctorate in 2015, and is looking forward to next working at Western Washington University as a National Science Foundation Astronomy Postdoctoral Fellow.