Applications of Photospheric Spot Temperature Models to the Study of Stellar Angular Momentum Evolution in the Orion Nebula Cluster

Swarthmore College Senior Thesis in Astronomy

Swarthmore College Department of Physics and Astronomy,

500 College Avenue, Swarthmore PA 19081

Matthew Miller

Advisors: Eric L. N. Jensen and Keivan G. Stassun

Contents

Abstract

We apply a simple photospheric spot temperature model to photometric variability measurements of T Tauri stars in the Trapezium region of the Orion Nebula Cluster. Our aim is to search for the relationship, if any, between spot temperatures and stellar rotation periods to better understand the relationship between accretion and angular momentum regulation in T Tauri stars. Current magnetic disk-locking models of young stars ascribe spot temperatures hotter than the photosphere to active accretion from a circumstellar disk. If accretion acts to brake stellar rotation, spot temperatures hotter than the photosphere should be more prevalent among slow rotators. From the variability amplitudes at four wavelengths (B, V, R, I), we determine spot temperatures and the areal coverage of the spot on the stellar surface. The results of our model show that we can unambiguously distinguish spots hotter than the photosphere from spots cooler than the photosphere for most stars. Comparisons between spot temperatures and previously determined rotation periods reveal no significant correlation between spot temperatures and stellar rotation. Comparisons between spot temperatures and measurements of near-IR excess, $\triangle$(I-K), reveal possible problems with the accuracy of $\triangle$(I-K) as an indicator of circumstellar disk presence. We confirm this problem by further comparing spot temperatures to Ca II equivalent widths and comparing $\triangle$(I-K) measurements with Ca II equivalent widths corresponding to evidence for accretion from a circumstellar disk. We propose that future studies of the angular momentum evolution of T Tauri stars take considerable care in distinguishing stars with circumstellar disks and possible magnetic disk locking from stars without circumstellar disks or magnetic disk locking. Finally, we argue that our spot temperature and rotation period comparisons support new theories of angular momentum evolution in T Tauri stars that relate the angular momentum evolution of pre-main-sequence stars to the core-envelope decoupling that occurs during the convective-radiative transition.

1 Introduction

The apparent brightness, or magnitude, of a star is perhaps its most easily measured property. Simple measurements of the amount of light emitted from a star at different wavelengths can provide valuable insights into the physical properties of the star. In young, pre-main-sequence (PMS) stars, the magnitude of a star often varies, to a point where most PMS stars show some evidence of variability (Herbst et al. 1994). Understanding the physical mechanisms behind the variability of these young stars is very important to our understanding of stellar evolution theory as well as planet formation theory. Current research investigating the physical properties of PMS stars have constrained our understanding of stellar evolution theory. In particular, this research has shown that stellar evolution theory cannot explain the angular momentum evolution of stars as they move from the PMS phase to the main-sequence phase in their lifetimes (Stauffer & Hartmann 1987). More specifically, astronomers are puzzled by the angular momentum distribution of T Tauri stars (TTSs), stars within the PMS phase that have already formed from their natal clouds of gas and dust and are the precursors of main-sequence dwarfs like the the sun. Understanding the physical mechanisms governing the evolution of T Tauri stars is particularly valuable to planet formation theory since TTS are quite often accompanied by circumstellar disks of material from which planets form.

The purpose of this thesis is to investigate the physical properties of T Tauri stars by measuring the temperature of stellar photospheric spots on individual stars in the Trapezium region of the Orion Nebula Cluster (ONC). The temperature of these spots is intrinsically linked to various physical properties of the star believed to be governing the star's rotation. By determining the temperature of these spots, the angular momentum distribution of the ONC can be directly related to the physical activity associated with the spots, and, from this relationship, it may be possible to determine what mechanisms, if any, affect the rotational evolution of T Tauri stars. In order to better explain the why angular momentum distribution of TTSs is so puzzling, and how photospheric spot temperatures can aid us in our understanding of the angular momentum evolution of PMS stars, it is necessary to provide more background for the issues at hand. We will first discuss theoretical models describing the rotation of PMS stars, then explain what spot properties and rotation properties have been observed in young star forming regions, and finally discuss how we can describe observations of stellar variability by the rotational modulation of spots hotter and cooler than the stellar photosphere on the stellar surface. Following this digression, we will move on to discuss the methods used to quantify the variability of stars in the ONC in Sections 2 and 3. The interpretation of this variability will be given in Section 4, and a discussion of the implications of our results will be given in Section 5.

1.1 The Magnetic Disk-Locking Model and the Angular Momentum Conundrum

It is fairly well known that stars form from collapsing clouds of rotating gas and dust that are massive enough to spontaneously collapse, satisfying the so-called the Jeans mass criterion. As these clouds collapse, conservation of angular momentum tells us that highly condensed clouds should be rotating faster than clouds that have just begun their gravitational collapse. Furthermore, we would also expect that not all the material from the natal cloud will immediately end up at the center of mass of the cloud where the star will most likely form. The effects of the centrifugal force from rotation, coupled with the force of gravity pulling material inward, will bring some of the material into a circumstellar disk. At this stage of stellar evolution, the star is gaining most of its energy from gravitational contraction, rather than hydrogen fusion. Low-mass stars in this evolutionary stage are known as T Tauri stars. Low-mass stars pass through the T Tauri phase with continuing gravitational contraction to the main sequence, at which point hydrogen fusion takes over as the dominant method energy transport and the hydrostatic equilibrium stops the collapse of the star.

One would expect stellar rotation rates between the PMS phase and the main-sequence phase to increase given the continued gravitational contraction. Observational studies comparing PMS stars and stars that have just reached the main sequence, so-called zero-age main-sequence (ZAMS) stars, show that the rotation rate distribution of T Tauri stars is not consistent with the rotation rate distribution of ZAMS stars. Given angular momentum conservation, the rotation rates of T Tauri stars projected to the main sequence are much higher than the observed rotation rates of ZAMS stars (Queloz et al. 1998). Therein lies the angular momentum conundrum. To have the observed distribution of rotation rates for ZAMS stars, we must expect that some physical mechanism is removing angular momentum from PMS stars. The physical mechanism that has garnered the most attention for its ability to explain the angular momentum loss is the magnetic disk-locking model.

Current magnetic disk-locking models for PMS stars are derived from the pioneering work of Ghosh and Lamb and their models involving the interaction between accretion disks and neutron stars (Ghosh, Lamb, & Pethick 1977; Ghosh & Lamb 1979a, 1979b). Their models conclude that the interaction of the stellar magnetic field and the circumstellar disk allow the accretion of material from the disk along magnetic field lines to the stellar surface. This result is attractive in two ways: first, the accretion of material onto the stellar surface can explain many observed characteristics of TTSs, namely ultraviolet excesses and veiling (Basri & Bertout 1989), hot photospheric spots (Herbst, et al. 1994), and truncated circumstellar disks (Meyer, Calvet, & Hillenbrand 1997). Second, the accretion of material onto the star can occur in the model only if angular momentum is transferred out from the star to the disk. This latter result is particularly attractive since it could possibly explain the apparent angular momentum loss that occurs between the PMS and ZAMS phases of stellar evolution. The quantitative details of the magnetic disk-locking model are beyond the scope of this thesis, but a qualitative explanation is useful for understanding the accretion processes that may create the photospheric hot spots observed and analyzed in this thesis.

There are a variety of processes that can cause the magnetic field to thread the disk plasma, creating the star-disk interaction that leads to accretion. Overall, there are many models for star-disk interaction (for further reference see Ghosh & Lamb 1977, 1979a, and 1979b; Königl 1991; Ostriker & Shu 1995, Balbus & Hawley 1998). In general, most models find similar results: the magnetic field penetrates the disk at its inner

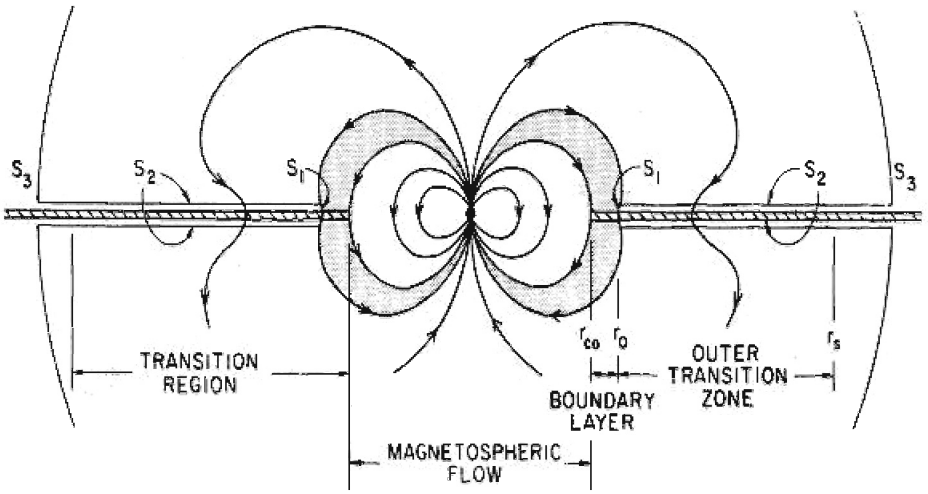

Figure 1: Ghosh and Lamb's view of magnetic disk-locking. Here, the stellar magnetic field threads the circumstellar disk out to the edge of the transition zone. Within the boundary layer, the disk material is forced to co-rotate with the star. Energy losses within boundary layer allow disk material to diffuse onto magnetic field lines and fall onto the star. Beyond the edge of the boundary layer, the disk rotates at Keplerian velocities.

edge to some outer radius where the disk energy density is sufficiently large enough to allow for turbulent diffusion between the disk and the magnetic field. Located near the inner radius of the disk is a narrow boundary layer, where the strength of the magnetic field stress is large enough to force the disk material to co-rotate with the star and its magnetic field. From the outer edge of this narrow boundary layer to the outer radius, where star-disk interaction ends, lies what is known as the outer transition zone (see Figure 1).

In Figure 1 we see the basic structure of the star-disk interaction as described by the Ghosh and Lamb model. Here, the magnetic field threads the disk out to some radius. The boundary layer consists of the

primary region where we see accretion along magnetic field lines. It is this region where the stress of the magnetic field forces the disk material to co-rotate with the star. Beyond the boundary layer (beyond r_o) the disk is free to rotate at its Keplerian value.

The Ghosh and Lamb model provides a simple interpretation of the angular momentum transfer between the star and disk: the torque between the magnetic field and the disk acts to spin up the star inside the radius at which the stellar rotational velocity matches the orbital velocity, the so-called corotation radius of the disk (not to be confused with r_{co} in Figure 1), while the torque between the magnetic field and the disk acts to spin down the star outside the corotation radius. The pioneering model of Ghosh and Lamb provides a simple way to examine how angular momentum is transferred between the disk and the star, although subsequent work (Königl 1991; Shu et al. 1994a; Shu et al. 1994b; Najita & Shu 1995; Ostriker & Shu 1995) has expanded the Ghosh and Lamb model to more specifically identify the the role of angular momentum transfer between the star and disk. These recently developed models question many of the assumptions of the Ghosh and Lamb model. These models differ significantly from the Ghosh and Lamb model in their conclusion that the stellar magnetic field can only connect to the disk in a small region near the corotation radius (see R_x in Figure 2). The basic idea is that the magnetic field near the inner radius of the disk will be frozen into the disk plasma and will become twisted as the disk material rotates a rate different than the stellar rotation rate. Reconnection events cannot necessarily maintain the connection between the magnetic field and the disk, and the magnetic field will preferentially connect to the disk near the corotation radius (Shu et al. 1994). At large radii, the disk material will not be perfectly conductive (since distance from the star decreases the disk temperature and the amount of free ions). Instead, the disk material at large radii will be diamagnetic and will repel the stellar magnetic field, causing the field lines to bow outward from the disk (see Figure 2).

The more recent disk model assumptions consequently show that angular momentum is removed from the star via accretion. Diffusion of disk material onto the stellar field lines requires a loss of energy from the disk material. The energy loss in the disk decreases the angular momentum of the disk material. This creates a problem because the magnetic field lines frozen into the disk want to continue to rotate at the same velocity. The result is a torque that transfers angular momentum from the star to the disk to maintain corotation

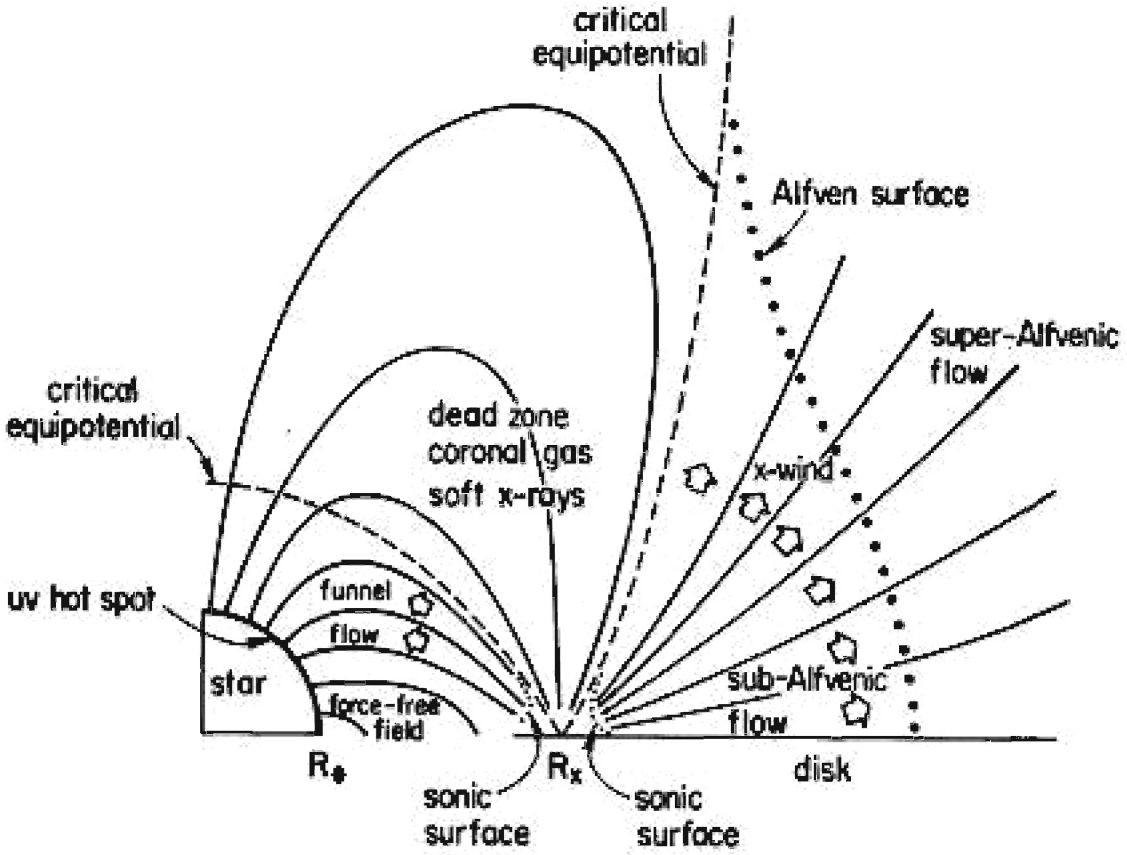

Figure 2: A more recent schematic of magnetic disk locking. Unlike the Ghosh and Lamb model, the stellar magnetic field can only remain connected to the disk if it does so near the corotation radius (R_x). At the corotation radius, disk material falls along stellar magnetic field lines to the star. In order for material to make the jump from the disk to the magnetic field lines, the disk material must lose angular momentum. The angular momentum loss drives disk material away from the disk in the X-wind.

(Shu et al. 1994). Angular momentum loss from the star and disk also occur through an intensification of the stellar wind (an X-wind–see Figure 2) beyond the corotation radius, thereby removing material from the star and the disk (Shu et al. 1994).

Even without of the details of the magnetic disk-locking models, we know that the possibility for a causal link between magnetic fields and angular momentum exists. On the microscopic level, the details of magnetic disk-locking are complex will never fully be explained by models. Nevertheless, we know that a star-disk interaction should affect the angular momentum of the star; it is with this knowledge that research has sought to establish empirical relationships between accretion activity and stellar rotation.

1.2 Observed Rotation Period Distributions of CTTS and WTTS

1.2.1 Bimodal vs. Unimodal

If we expect magnetic disk-locking to be retarding the rotation rates of stars actively accreting from circumstellar disks, we should also expect stars with accretion signatures to have a longer average rotation period (i.e. slower rotation rate) than stars without evidence of accretion. Numerous studies of the Orion Nebula Cluster have searched for correlations between rotation period and accretion signatures, often with varying results. The first major studies of rotation period in young stellar clusters were performed by Attridge and Herbst (1992), Bouvier et al. (1993), and Herbst et al. (1994). These studies were the first to interpret rotation periods using the magnetic disk-locking model. These earlier studies were also the first to find a correlation between disk presence and rotation rate, which, at the time, was thought to confirm the active role of magnetic disk locking in regulating the angular momentum of T Tauri stars. Later studies proved to be mixed. While some studies (Choi & Herbst 1996) confirmed the results of these earlier studies, others found no correlation between disk presence and rotation rate (Stassun et al. 1999, hereafter SMMV; Rebull 2001), and yet others found that correlations between disk presence and rotation rate only exist for stars less massive than $0.25 M_{\odot}$(Herbst et al. 2002). These varying studies conclude one of two results: the rotation period distribution of T Tauri stars is bimodal, or the rotation period distribution of T Tauri stars is uniform or unimodal.

Those studies detecting a bimodal rotation period distribution find that stars with disks tend to distribute around one rotation period, while stars without disks tend to distribute around another rotation period. For those studies reporting a bimodal distribution, the favored period modes are two days for stars without disks (weak-line T Tauri stars–WTTSs) and eight days for those with disks (classical T Tauri stars–CTTSs) (Attridge & Herbst 1992; Herbst et al. 1994; Choi & Herbst 1996; Herbst et al. 2000). Still others find a unimodal distribution of rotation periods (SMMV). Upon examining the unimodal distribution of SMMV, Herbst et al. (2000 & 2002) note that their results also exhibit a unimodal distribution when limited to stars less massive than 0.25 solar masses. Rebull (2001) complicates this issue further by finding a distribution statistically unimodal at a 50% confidence level for $M < 0.25M_{\odot}$ and at an 80% confidence level for $M > 0.25M_{\odot}$. Furthermore, supposing that the distribution is not unimodal, Rebull finds that the modes of the bimodal distribution do not correspond with the two and eight day modes of the Herbst distributions. Overall, no researchers find evidence for bimodality for $M < 0.25M_{\odot}$ (Hartmann 2002). Outside the ONC, Bouvier et al. (1993) find that WTTSs in the Taurus-Auriga dark cloud (another young star population) preferentially rotate faster than CTTSs in that same cloud. The Bouvier study, however, only concerns a total 25 WTTSs and CTTSs, a statistically small distribution, and does not control for the masses of stars in the distribution.

At the very least, it is easy to say that the magnetic disk-locking model is plagued by observational studies of rotation periods, as there is no dominant correlation between rotation rate and disk accretion. Researchers have not abandoned the magnetic disk-locking theory, however, for it is still attractive as a theory that can explain many observed characteristics of TTSs, particularly processes that can result from accretion along magnetic field lines like UV excess and hot photospheric spots.

1.2.2 Determining Disk Presence and Active Accretion

Given the ambiguity of reported rotation period and accretion correlations, the underlying question of all recent studies of rotation period distributions has been: is the distribution unimodal or bimodal? Answering this question is not as straightforward as it may sound, for one first has to identify those stars with signatures of active accretion from a circumstellar disk. A first step toward isolating those stars with active accretion

from a disk is to detect the presence of a circumstellar disk. The presence of a disk may be inferred by the presence of an infrared (IR) excess in a star's spectrum. The physics behind this observation relies on simple Planck blackbody radiation. Each star will emit its radiation in the form of a continuous blackbody function peaked at a wavelength determined by the photospheric temperature of the star. For stars with circumstellar disks, the gas and dust in the disk will absorb some of the star's radiation and re-radiate it as a continuous blackbody function centered at a longer wavelength (generally located in the infrared, corresponding to a temperature cooler than the star's photosphere). The resultant spectrum will contain radiation from both the star and the disk, and there will be a noticeable excess of radiation above that expected from the star at infrared wavelengths.

While IR excesses offer direct evidence for disk presence, angular momentum studies will benefit from rotation period comparisons of CTTSs and WTTSs only if there is some indication that CTTSs are accreting matter from their circumstellar disks. Observational studies of TTS rotation periods have focused on two accretion diagnostics: Hα emission and ultraviolet (UV) excess. In a few cases, evidence for accretion is drawn from Ca II emission (Hillenbrand et al. 1998). All of these observable characteristics are the direct result of the same physical mechanism.

Observed Hα emission, UV excess, and Ca II emission result from the energy released as circumstellar disk material accretes onto the stellar surface. As disk material traveling in orbit around the star reaches the stellar surface (it does this through some accretion process: gravity, flow along magnetic field lines, etc.), it must promptly slow to a speed regulated by the rotation rate of the star. The energy released by the friction to slow the speed of disk material is manifested observationally by Hα emission, UV excess, and Ca II emission. In the case of Hα emission, the collision of disk material with the stellar gas ionizes the neutral hydrogen in the disk material. Upon recombination, we preferentially see transitions from the n=3 to n=2 states in hydrogen, the Balmer α transition, also denoted as Hα emission. We must also keep in mind that the disk material is relatively dense (in comparison to the interstellar medium), and will therefore radiate much of the energy of friction as a blackbody. Thus, in the same way we see an IR excess from the disk material in orbit, the disk material falling onto the star emits a blackbody peaked a wavelength shorter than the peak wavelength of the stellar continuum. This blackbody contributes more continuum radiation at ultraviolet

wavelengths than other wavelengths, thereby producing ultraviolet radiation in excess of what the star itself produces. Indications of excess continuum radiation at ultraviolet wavelengths are visible through observed UV continuum excess and the decreasing strength of absorption features. This latter effect, often referred to as veiling, is a direct consequence of the increased UV continuum, for the increased amount of continuum photons decreases the size of an absorption or emission feature relative to the continuum level. Thus, the "strength" of an absorption or emission feature as measured through this relative size (the formal calculation of this relative size is called the equivalent width of the spectral line) decreases even though the number of photons removed from or added to the line of sight through the absorption/emission process remains constant.

Ca II emission results from the ionization of calcium atoms in the photosphere of the star upon impact of accretion material. The collision of the disk material with the stellar surface ionizes neutral calcium in the photosphere to the first ionized state. The recombination of ionized calcium back into the ground state is seen as an emission feature in the stellar spectrum. It is therefore believed that emission lines of Ca II are excellent indicators of accretion activity. Comparisons of Ca II equivalent widths and IR excess (indicative of disk presence) show that active accretion and disk presence is not a one-to-one correlation. Whereas IR excess always appears where there is Ca II emission, Ca II emission does not always appear when there is IR excess (Hillenbrand et al. 1998). Moreover, correlations between stellar X-ray luminosity and Ca II emission show that the two are often found in tandem, whereas no such correlation appears to exist between X-ray luminosity and IR excess, implying that there might be causal links between accretion activity and X-ray activity but no causal link between disk presence and X-ray activity (Flaccomio et al. 2003). These results imply that stars do not ubiquitously accrete from their circumstellar disks. In fact, correlating disk presence and rotation periods may not be enough to draw conclusions about the magnetic disk-locking model, although the first step in recognizing the impact of magnetic disk-locking is to first determine which stars have disks.

We have four excellent indicators of disk accretion: Ca II emission, Hα emission, UV excess, and veiling. (Among the latter two, observation of UV excess is somewhat more reliable because one may know the expected level of UV continuum from the photosphere better than one knows the expected amount of absorption or emission for a given spectral feature.) Analysis of angular momentum regulation in TTSs has

proceeded by comparing rotation rates to disk presence and to signatures of active accretion. Typically, distinctions between WTTSs and CTTSs are made with near-IR excesses and Hα emission. Disk presence is commonly associated with a near-IR excess of Δ(I-K) $>$ 0.3 mag (Carpenter, Hillenbrand, & Skrutskie 2001; Rebull 2001; Herbst et al. 2002) while active accretion is commonly associated with an Hα equivalent width of 10Å or more (Bouvier et al. 1993, Herbst et al. 2002) or 20Å or more where more conservative demarcations are needed to avoid nebular Hα contamination (SMMV). Interestingly, those who find correlations between disk presence and rotation rates find them regardless of whether they define disk candidates by IR excess or Hα emission (Bouvier et al. 1993, Herbst et al. 1994, Herbst et al. 2002). Moreover, those who find no correlation between disk presence and rotation rates repeatedly find no correlation regardless of whether they define disk candidates by IR excess, Hα emission, or UV excess (SMMV, Rebull 2001). In fact, SMMV find that 10% of rapid rotators (rotation period $\leq$ 4 days) show evidence for accretion while 13% of slow rotators (rotation period $>$ 4 days) show evidence for accretion.

1.2.3 Variability Amplitudes and Spot Lifetimes

While much research has focused on disk-rotation correlations, a considerable amount of research has also attempted to correlate the degree of observed variability and photospheric spot temperatures with disk and accretion indicators. The general conclusion of many of these studies is that WTTSs exhibit spots cooler than the photosphere while CTTSs exhibit spots both cooler and hotter than the photosphere. (At this point it might be worth noting that these studies, and our study as well, use the photometric variability of a star at different wavelengths to infer spot temperatures and spot sizes. For instance, how the variability for a star changes as one moves from the optical to the infrared regions of the electromagnetic spectrum can be described by a spot, with a given size and temperature, on the stellar surface coming in and out of view as the star rotates on its axis. In this sense, we never really see the spot, but instead ascribe a spot of a certain size and temperature to the measured variability. For further discussion of the derivation of spot temperatures and sizes, see Section 1.3.) The result that WTTSs exhibit cool spots while CTTSs exhibit both hot and cool spots is what we would expect to observe from magnetic accretion theory. We should expect cool photosphere spots for both types of stars. These cool spots arise from the same interaction between the star and the stellar magnetic field that leads to the cool photospheric spots we see on

the sun. The hot spots commonly seen on CTTSs result from the accretion process. Given signatures of accretion (Ca II emission, Hα emission, UV excess, and veiling) we should expect to see spots hotter than the photosphere for those stars with disks (CTTS), in addition to the stellar magnetic field cool spots that we see on those stars without disks (WTTSs). Observationally, hot spots are found almost exclusively on CTTSs (Bouvier et al. 1993). Observational evidence of these spots has shown that typical spot sizes are much larger than those seen on the sun. This is not, however, unusual. Large photospheric spots are seen on RS CVn stars, where the tidal force between stars in a binary system generates huge magnetic fields that are visually manifested in rather large cool photospheric spots (Vogt 1981). The fact that sunspots and the observed RS CVn spots are cooler than the photosphere is a good indication that the observed hot spots in TTSs result from physical processes, namely accretion from a circumstellar disk, that are absent from our sun and RS CVn stars (Herbst et al. 1994).

Initial studies of hot and cool spots on TTSs found that the amplitude of variability observed and the duration of variability strongly depend on the spot temperature. Observations show that variability due to rotational modulation of cool photospheric spots tends to more stable than variability due to rotational modulation of hot photospheric spots. It is for stars with hot spots that we see period wandering, where the period of variability for a star changes over the course of a few years, and changes in the lightcurve variability amplitude (Vrba et al. 1989, 1993). Herbst et al. (1994) attribute the period wandering and aperiodic variability of hot spots to coronal flares, unsteady accretion, and variations in circumstellar extinction. An interesting correlation to note is one between amplitude of variability and spot temperature. Observationally, some researchers find that evidence for a circumstellar disk (through near-IR excess) correlates well with variability amplitudes greater than 0.5 magnitudes (Choi & Herbst 1996; Herbst et al. 2002). Many interpretations of this correlation assume that the large variability amplitude results from hot spots formed from disk accretion. An important question to ask, then, is whether or not stars with variability amplitudes associated with hot photospheric spots have any preferred rotation distribution. In other words, does a variability amplitude greater than 0.5 magnitudes imply accretion from a circumstellar disk and a retarded rotation rate from magnetic disk locking? A few researchers have confronted this question, but with mixed results. While some researchers find that no rapid rotators (interpreted as WTTSs if under the auspices of the magnetic disk locking theory) have large variability amplitudes (Choi & Herbst 1996), others find

that large variability amplitude does not imply accretion signatures visible through UV excess even though there is an overall weak correlation between UV excess and large variability amplitudes (Rebull 2001). The lack of consistency between observations once again exhibits the need for more observations as well as more comprehensive measures of accretion among a large group of TTSs.

1.2.4 Comparison to the Pleiades V sin(i) Distribution

When a direct measurement of a stellar rotation period is not available, the rotational velocity of a star may be inferred from a measurement of the projected rotational velocity. Like the rotation period analysis in Section 1.2.1, the distribution of projected rotational velocities can be used to glean information about the angular momentum of a stellar cluster. We can find the projected rotational velocity of a star by relating the width of a spectral feature to the rotational speed of a star, for, as the star rotates on its axis, various portions of the stellar photosphere are Doppler shifted from our view. At the edges of the star, we will see the extremes of this Doppler shift: at one edge the photosphere will be redshifted from our view, and at another edge the photosphere will be blueshifted from our view. The Doppler shift over entire surface of the star creates a spectral feature, broadened by stellar rotation. What we can measure is the stellar rotation along our line of sight, also called the projected rotational velocity, or v sin(i). The v sin(i) distribution of low-mass stars from the Pleiades cluster is similar to, if not indistinguishable from, the v sin(i) distribution of stars in the ONC (SMMV). The similar v sin(i) distributions of the Pleiades, a cluster of ZAMS stars approximately 10^8 years old, and the ONC, a cluster of PMS stars approximately 10^4 - $10^{6.5}$ years (Hillenbrand 1997), reveal that the distribution of specific angular momentum for stars upon contraction to the main-sequence remains the same, even though the specific angular momentum of these stars should be increasing with contraction. This poses a problem for the magnetic disk-lock theory, for the time of contraction to the main-sequence from the age of the ONC to the age of the Pleiades is much longer than the lifetime of circumstellar disks, implying that angular momentum loss occurs beyond the lifetime of the disk (SMMV). Solar-analog winds can remove angular momentum from the stars as they contract toward the main-sequence, but the winds themselves cannot remove as much angular momentum as is observed to be lost by PMS stars over the relatively short time of contraction to the main-sequence (SMMV). Consequently, the similarity between the Pleiades and ONC v sin(i) distributions provides evidence that angular momentum loss cannot occur solely through

the magnetic disk lock model. To understand the causal connection between rotation periods of TTSs and accretion activity, it is imperative to first establish more empirical connections between accretion activity, disk presence, and rotation periods. The more empirical understanding of interconnected physical mechanisms in TTSs we have, the more likely we are to understand disk-regulated angular momentum evolution. Currently, the magnetic disk lock model nicely explains many of the observational characteristics of TTSs, but there is no authoritative observational understanding of correlations between accretion properties and rotation periods. We hope to further the understanding of empirical relationships between accretion and rotation by studying spot temperatures, enabling us to determine if spot temperature correlates with disk presence, rotation period, or indicators of accretion activity.

1.3 Characterizing Stellar Variability with Hot and Cool Spots on the Stellar Surface

The premise of much of the discussion in the introduction to this thesis is that photospheric hot spots on young stars are related to accretion processes from a circumstellar disk while photospheric cool spots are related to physical processes not associated with a circumstellar disk. It is our hope to determine what correlations exist, if any, between spot temperatures and the angular momentum evolution of young stars, but, before we can begin to answer this question, we must first address the essential question that has been lingering in the mind of the observational astronomer reading this thesis: how do we know if there is a hot or cool spot on the stellar surface? The short answer to this question is that we can ascribe the photometric variability to a hot or cool spot on the surface of a rotating star.

In our investigation of young stars in the Orion Nebula Cluster, we hope to use the photometric variability of stars to determine the physical properties associated with the observed variability. Namely, we hope to use the photometric measurements of a star's brightness over time to create a lightcurve (the measured brightness of the star on the astronomical magnitude scale versus time) from which we can quantify how the star varies with time. For our purposes, we wish to quantify the variability of a star by the amplitude of variability seen in the lightcurve. It is this amplitude, as well as some other properties of the lightcurve, that can be described by a hot or cool spot on a rotating star. To see why, let us consider two hypothetical cases. In the first case we have a star without a spot on its surface. In the second case we have a star with a hot

spot on the stellar surface. In the first case, we should see no variability in our lightcurve (i.e., the magnitude of our photometric measurements will remain constant with time). This case is simple to understand, for, as the star rotates on its axis, the side of the star that we see is always of one single temperature. In the second case, things become more complicated by the hot spot. Supposing we have a spot hotter than the photosphere on the stellar surface that is some small fraction of the stellar surface (say, 1%), we should expect to see a lightcurve where the magnitude of our photometric measurements will vary with time. The reason for this is simple: the star will appear brighter when the spot facing us and dimmer, or at its unaltered brightness, when the spot is not facing us. The difference between the maximum and minimum brightness defines the variability amplitude for our star. How our variability amplitudes change at different regions in the electromagnetic spectrum (i.e., how our variability amplitudes appear in the optical versus the infrared) depends on the spot temperature and spot size (see Section 3.1). The entire next chapter will be devoted to the measurement of variability amplitudes and the detailed process of making the lightcurves from which we draw these variability amplitudes. The chapter that follows (Chapter 3) is devoted to how these variability amplitudes produce the spot temperatures we need for the analysis discussed in Chapters 4 and 5.

2 Observations and Methods

In this section we describe the choices made regarding the observations and methods we used to extract data for our rotation period and spot temperature analysis. We begin by describing the nature of our observations, and how these observations are used to make the lightcurves used in our spot model. A detailed description of the photometry used to extract stellar variability data is then described, followed by a discussion of the types of variability observed in our lightcurves and the method through which we can extract the variability amplitudes used in our spot model from these lightcurves.

2.1 Observing in the Orion Nebula Cluster

One of the most carefully studied young star regions is the Orion Nebula Cluster (ONC). The attention focused on the ONC arises more from practicality than anything else. Located at a distance of 470 pc, the

proximity of the ONC accommodates the study of faint young stellar objects. Large stellar cluster studies, such as ours, are further aided by the location and surroundings of the ONC. At a galactic latitude of -20°, the ONC is located 160 pc out of the plane of the Galaxy. With a location outside the stellar disk, observations of the Orion cluster have limited contamination by foreground stars not considered to be members of the ONC and, more importantly, not considered to be at the same evolutionary state as stars in the ONC. Contamination of background stars is also reduced because the ONC is located directly in front of a molecular cloud with over 100 magnitudes of visual extinction (Hillenbrand 1997). We choose to observe in the ONC because the abundant number of observational studies conducted there have generated a large database of detailed information for thousands of stars in the cluster. This stellar database includes information about spectral type, infrared excess, stellar age, interstellar reddening, and rotation period, among other things.

Our observations come from the Trapezium region of the Orion Nebula Cluster. To perform our multi-bandpass photometry campaign, Keivan Stassun observed four fields in this region during a 13-night observing run at the WIYN 0.9m telescope at Kitt Peak from January 3rd to January 15th in 1998. The four fields of our study are 23' square fields centered on 5:35:24.0 -5:30:19.0, 5:35:04.0 -5:11:05.0, 5:35:15.0 -4:49:55.0, and 5:34:47.0 -5:52:15.0, epoch and equinox 1998. The data for our multi-bandpass photometry come from observations of these four fields in the Harris B, V, R, and I filters. Observations of each field were made two to three times per night, and, with 10 nights of data (nights where we have no photometry data are January 3rd, 4th, and 10th), we typically have 20-30 images of each field taken through each filter. Standard reduction procedures (flat-field correction, bias subtraction, etc.) were completed using IRAF before photometric flux measurements were made.

2.2 Method of Ensemble Differential Photometry

In our analysis of spot temperatures and sizes we are only concerned with the amplitude of variability; as the amplitude of variability is found through the difference between two flux measurements of the same star, we have no need to evaluate the absolute flux from each star (since, as it were, any absolute flux offset will be subtracted out when calculating the amplitude–see Section 2.4.2). Therefore, our lightcurves can be derived from measurements of differential photometry rather than absolute photometry. To produce lightcurves for

the thousands of stars in our data set, our method of photometry must be somewhat automated. An ideal method of photometry for our data is that prescribed by Honeycutt (1992).

Honeycutt's method for differential photometry uses a set of comparison stars of constant magnitude from image to image to define the ensemble that sets the flux offset for each image. Honeycutt's method differs from strict ensemble photometry in that the exact number and identity of comparison stars varies for each image. In our data, choosing a varying set of comparison stars is imperative, for changes in observing conditions (such as differences in sky background brightness) and changes in limiting magnitudes for images taken in different filters can affect the set of comparison stars chosen (here we switch to a discussion of magnitudes rather than flux, since our flux measurements are converted to a magnitude scale). An added bonus of Honeycutt's method is that we can find a large number of the comparison stars scattered throughout an image to produce the magnitude offset. The large number of stars will increase the statistical accuracy of our final comparison magnitude, and the presence of comparison stars throughout the image will average out residual flat-fielding errors (Honeycutt 1992).

The first step in producing differential magnitudes is to identify the set of comparison stars for an image. Candidates for this set are those stars that exhibit no intrinsic variability across all images. At first glance, it may seem difficult to separate stars of constant magnitude from those that vary in magnitude from image to image, especially since image to image variations in overall brightness may lead the naïve observer to believe that no stars in the image are of constant magnitude. We can, however, automate our search for comparison stars by using a least-squares analysis to identify the constant stars needed to assign the magnitude offset for an image. For instance, suppose we have x images. On all x images we find y stars. For now, we can assume that every star in the set y appears in every image in the set x, although this is not always the case. Each star has a measured instrumental magnitude in every image. This instrumental magnitude is the combination of some instrumental magnitude that would have been measured in the absence of exposure-to-exposure variations (such as a brighter sky background in a region of nebulosity or the presence of thin cirrus clouds on a given night of observations) plus the image magnitude for that particular

exposure:

$$m(e,s) = m0(s) + em(e)$$

where *m(e,s)* is the instrumental magnitude for star s on image e, *m0(s)* is the instrumental magnitude that would have been measured for star s in the absence of exposure-to-exposure variations and *em(e)* is the image, or exposure, magnitude representing the variability for the entire image, e. Now, the only information we have is that of *m(e,s)*. If we want to find those stars of constant magnitude, we can instead define *m0(s)* as the mean instrumental magnitude of star s across all images. Those stars with constant magnitude should have *m(e,s)* - *m0(s)* - *em(e)* = 0. We thus want to minimize the least-squares solution that follows:

$$\beta = \sum_{e=1}^{x} \sum_{s=1}^{y} [m(e,s) - m0(s) - em(e)]^2 w(e,s)$$

where *w(e,s)* is a statistical weight set equal to $w_1(e)w_2(s)w_3(e,s)w_4(e,s)$. The weights w_1, w_2, w_3 are assigned values of 0 or 1 (w_1 equals 0 if the entire image e is to be excluded from the solution, w_2 equals 0 if all exposures of the star s are to be excluded from the solution, w_3 equals 0 if the star s on image e is to be excluded from the solution) and $w_4 = 1/\sigma^2[m(e,s)]$, the error on the measured instrumental magnitude. The solution to this linear least-squares equation can be found by solving the normal equations $\frac{\partial}{\partial em(e)}$ and $\frac{\partial}{\partial m0(s)}$. After solving these normal equations to find *em(e)*, the exposure magnitude for each image, we iteratively remove stars considered to be variables from the least-squares solution. This can be done through an analysis of the deviation of exposure-to-exposure magnitude measurements, *m(e,s)*, from the mean magnitude measurement, *m0(s)*. As seen in Figure 3, once a first solution to the exposure magnitudes is obtained, we can check the standard deviation of magnitude measurements for a star over all images against the mean magnitude for that star. Those stars with high standard deviations can generally be considered to be variables. We can exclude these stars from our data set and reiterate our least-squares solution and acquire a better constrained estimate of exposure magnitudes.

Once we have definitive values for the exposure magnitude for each image, we can define the magnitude of

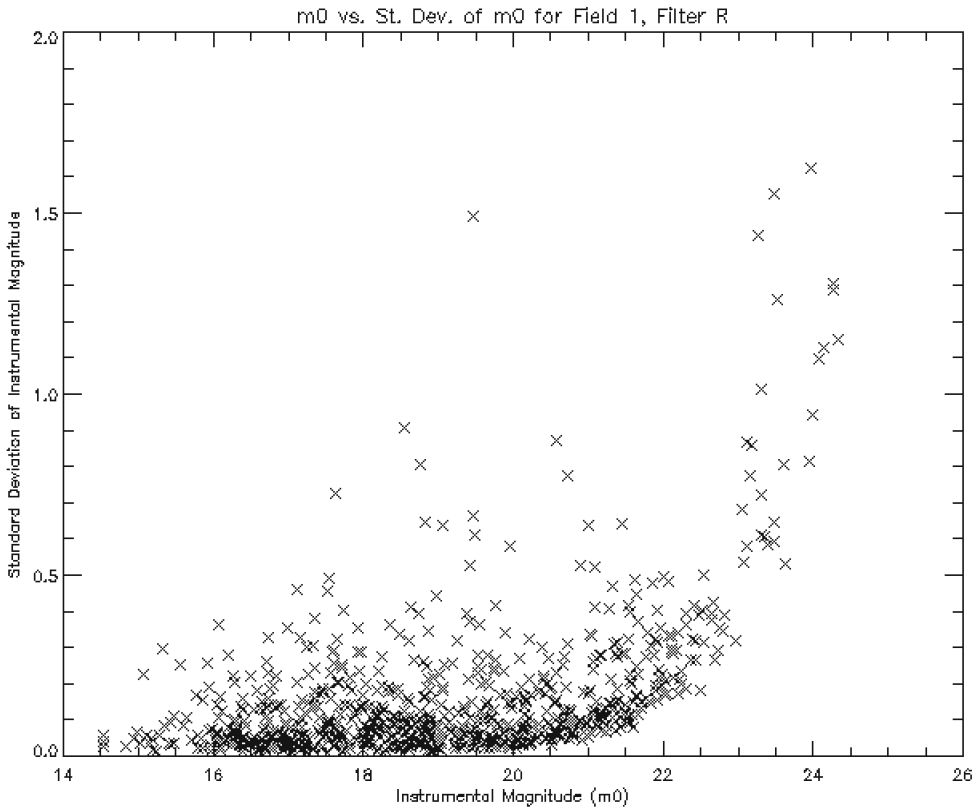

Figure 3: The standard deviation of magnitude measurements across all images as a function of mean magnitude, where each star is represented by a hatch mark. We consider those stars with the highest standard deviation to be variables. The variable stars are omitted from the least-squares solution such that only stars with constant magnitude are used to recalculate the exposure magnitude for all of the images. (Notice that the standard deviation increases with increasing magnitude. This effect is due to the increased photometric noise that comes with faint star observations.)

any star s as

$$M(e,s) = m(e,s) - em(e).$$

We then compute differential magnitudes by comparing this magnitude to the mean magnitude measurement for star s across all images. The error on our magnitude measurement *M(e,s)* is defined by

$$\sigma^2_{M(e,s)} = \sigma^2_{m(e,s)} + \sigma^2_{mean(em(e))}$$

where $\sigma^2_{mean(em(e))}$ is just the square of the total error produced by the exposure magnitude divided by the number of stars used to define the exposure magnitude. The comparison of M(e,s) and m0(s) defines the differential magnitude scale for our lightcurves. The error on each point in the lightcurve is found by $\sigma^2 = \sigma^2_{M(e,s)} + \sigma^2_{m0(s)}$ (Honeycutt 1992).

2.3 Making Lightcurve Solutions

Numerous steps must be completed before our field images are ready for the Honeycutt photometry analysis. These steps begin with identifying stars on our CCD images. The objects we label as stars are those objects with a Gaussian-shaped stellar point-spread-function (PSF) (i.e., a Gaussian profile with the sharpness ratio typical of a star). This latter constraint works to eliminate cosmic ray detections that trigger a Gaussian distribution of counts on the CCD. Once we have a list of stars available, we can then transform x and y pixel positions to right ascension and declination values given a list of bright stars with known values for right ascension and declination. From this point, we can begin photometry on all stars in all images. The photometry used in our lightcurve analysis is aperture photometry, the principle of which is to measure the flux from a star by finding the flux within a circular aperture centered on the star. An annulus surrounding the star is then used to compute the contribution of the background sky flux to be subtracted out of the circular aperture flux measurement. The size of the aperture used in our photometry is determined by the FWHM of the stellar PSF. Once we have performed photometry on all stars and in all images, we then match the stars across all images by searching for stars with the same right ascension and declination values. At this

point we can begin to find exposure magnitudes using the Honeycutt method for differential photometry. We begin our analysis with 4693 stars, but approximately 1500 of these stars are removed from our final analysis because they do not appear frequently enough in all images. (For instance, in our data analysis we remove stars that appear fewer than 10 times in the over 20 images we may have for that field taken in a particular filter. This is necessary for several reasons. First, it is difficult to compute the least-squares solution if the photometry consists of a majority of null values. Second, it is likely that our star, if indeed the object that appears only 10 times is a star, is at our magnitude detection limit, and any photometric calculation based on this faint star will have a low signal-to-noise ratio. Third, we cannot derive characteristics from a lightcurve with only 10 points with much certainty.)

Further examination of the Honeycutt analysis can tell us which images should be removed from the analysis. For instance, one image may have an exposure magnitude standard deviation that is considerably different from the rest of the images. Examination of observing logs for this image may yield information that leads us to remove the image from the least-squares solution (information such as a plane passing overhead or considerable cloud coverage). At the end of the Honeycutt photometry analysis, we are prepared to make lightcurves (Julian date vs. differential magnitude) of approximately 3,000 stars with anywhere from 10 to 26 data points (10 to 26 images from which we compute photometry).

2.4 Lightcurve Analysis

As mentioned before, the primary goal of differential photometry is to make lightcurves from which we can measure variability amplitudes in the B, V, R, and I bandpasses for each star. These amplitudes are then used in our photospheric spot model to determine the spot temperature and spot size that best describes the observed variability. Before discussing the details of measuring variability amplitudes, it is necessary to first describe the general characteristics of our lightcurves. Understanding the types of variability seen in our data set will allow the reader to visualize the types of variability typically associated with the rotational modulation of hot and cool spots as well as some forms of variability not associated with hot and cool spots rotating with the stellar surface. These latter types of variability are not the dominant form of variability seen in our data set, but are useful for conveying the complexity involved in stellar variability.

2.4.1 Characteristics of Derived Lightcurves

For our investigation, only lightcurves for stars with measured rotation periods are of interest. Nevertheless, our photometric monitoring campaign does not preclude us from analyzing the lightcurves of stars with no known rotation period. Our lightcurves can therefore be lumped into two broad categories: periodic and aperiodic. Many stars within the periodic category have sinusoidal light curves, which is exactly what we would expect if the variability is caused by a cool or hot spot on a rotating star. Within these periodic cases, lightcurves vary from showing what appear to be approximately perfect sinusoids (see Figure 4a) to others showing distorted sinusoids (see Figure 4b). The explanation for the latter of these cases is a combination of our viewing position and the possibility of multiple spots of roughly the same temperature on the stellar surface. For instance, we can imagine viewing a star whose rotation axis is inclined toward or away from our viewing position (the two extreme cases of this are viewing the star pole-on or at the equator) with a spot, or multiple spots, at varying latitudes on the stellar surface. Such a composition would place us, the observer, in the position of seeing the spots coming into view at varying times (actually, what we really see here is the change in flux, the indication of a spot coming into and out of view). The resulting lightcurve in the multiple spot case would be a complex combination of sinusoids all with the same period as seen in Figure 4b.

Among the non-periodic cases, we predominantly see irregular variability. In such cases, the typical lightcurve has no discernible sinusoidal pattern, much like that seen in Figure 5a. The causes of aperiodic variability such as those seen in Figure 5 are highly speculative, but mechanisms such as coronal flares, irregular accretion of material onto the star, and temporal variations in circumstellar extinction are very likely contributors to the observed irregular variability (Herbst et al. 1994; Ménard & Bertout 1999). Occasionally we see cases within the aperiodic category in which the lightcurve shows a star with a steadily increasing or decreasing brightness (see Figure 5b). Again, the causes for this type of variability are highly speculative, but variability studies conducted over the period of several months and/or years have observed cases in which the brightness of a star increases by several magnitudes in less than a year. Stars with observed increases in brightness over the course of a few months or years are labeled FU Orionis-like stars, named after the first T Tauri star observed to show such characteristics (Hartmann & Kenyon 1996). The presumed cause of this

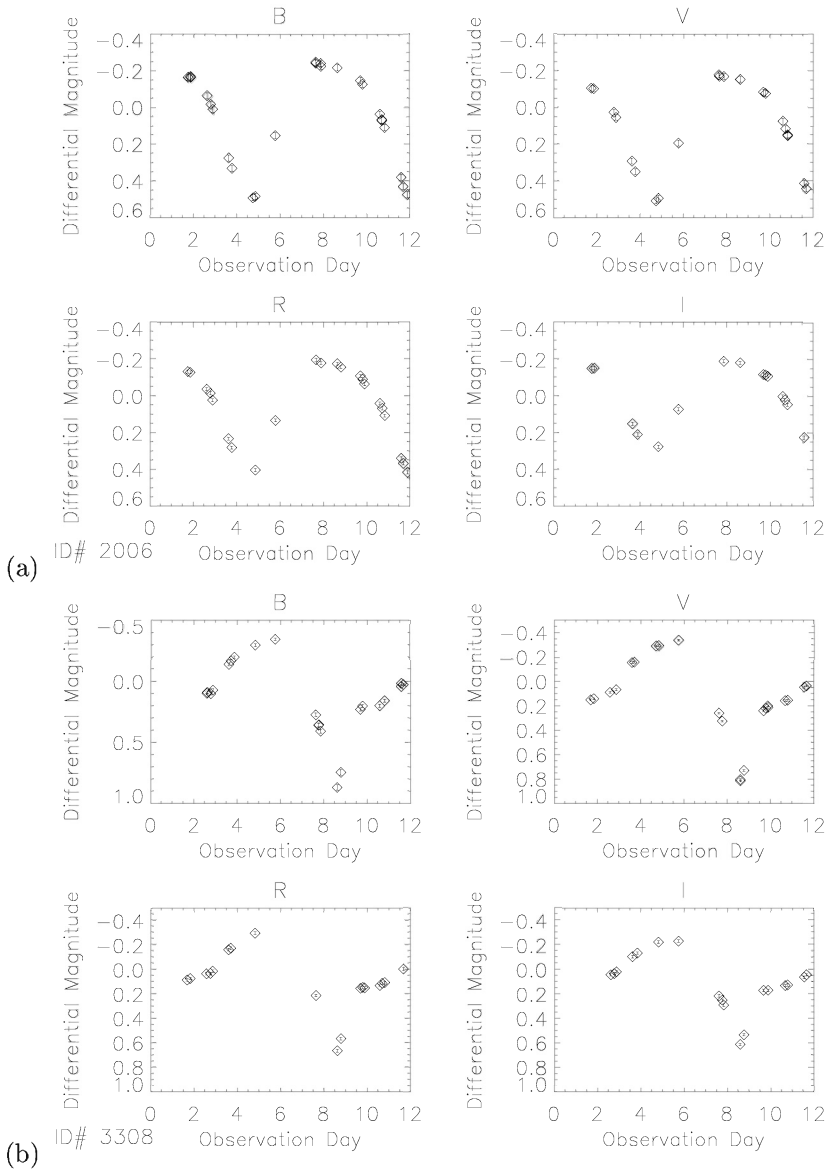

Figure 4: Examples of (a) a near-perfect sinusoidal lightcurve and (b) a lightcurve that appears to be a linear combination of sinusoidal lightcurves of the same period but different phase. (Note: in this plot and the lightcurve plots that follow, error bars are plotted for each point, but in many cases the error bars are smaller than the symbol size.)

steady increase in brightness is an increase in the accretion rate, possibly driven by large-scale instabilities in the disk structure (Hartmann, Kenyon, & Hartigan 1993). On the whole, our lightcurves show variability amplitudes that range from 0.01 magnitudes (the lower limit to our ability to detect variability with any statistical significance) to an astonishing 4 magnitudes (see Figure 6). In our study, stars with aperiodic variability will be analyzed by our spot temperature model, but they will not be included in our rotation period analysis, for obvious reasons. We discuss aperiodic variability here because our study is, in effect, a study of stellar variability. We do note, however, that previous studies attempting to describe aperiodic variability with variability models invoking hot spots, extinction, and unsteady accretion fail to encompass all types of observed variability for the aperiodic cases (Carpenter, Hillenbrand, & Skrutskie 2001).

2.4.2 Measuring Variability Amplitudes

To reiterate a point made in Section 2.2, differential photometry, not absolute photometry, is all that is necessary for our spot model. Differential photometry is sufficient because our spot model matches the variability amplitudes in four bandpasses (ΔB, ΔV, ΔR, ΔI) to a spot temperature and spot size. These variability amplitudes are simply the difference between the point of maximum brightness and minimum brightness in a given lightcurve. With differential photometry, our measured magnitudes are instrumental rather than absolute, but the constant offset applied to a given lightcurve is subtracted out during the calculation of the variability amplitude. Thus, the final measurement is a reflection of the difference between the true maximum brightness and the true minimum brightness even though we never have any knowledge of the true absolute brightness of the star. (To be more precise, the final measurement is a reflection of the difference between the true maximum brightness and the true minimum brightness on the magnitude scale; as a difference in magnitudes corresponds to a ratio of fluxes, the final measurement reflects the ratio of the maximum flux to the minimum flux.) For instance, our calculation of ΔB can be described as follows:

$$\Delta B = B_{instrumentalmax} - B_{instrumentalmin}$$

$$B_{instrumental} = B_{true} - B_{offset}$$

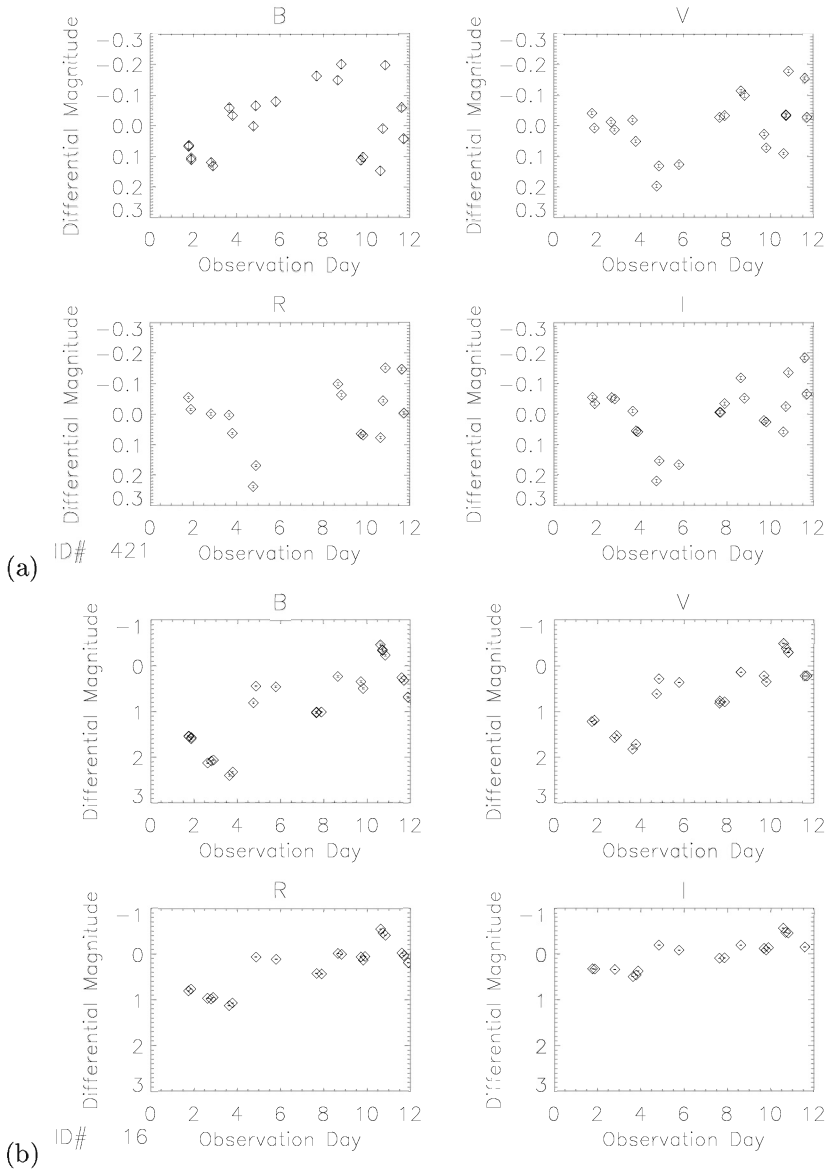

Figure 5: Two examples of lightcurves exhibiting irregular variability. In (a) we see an aperiodic lightcurve. In (b) we see an example irregular variability in a lightcuve where the stellar brightness is steadily rising (the magnitude is becoming more negative)

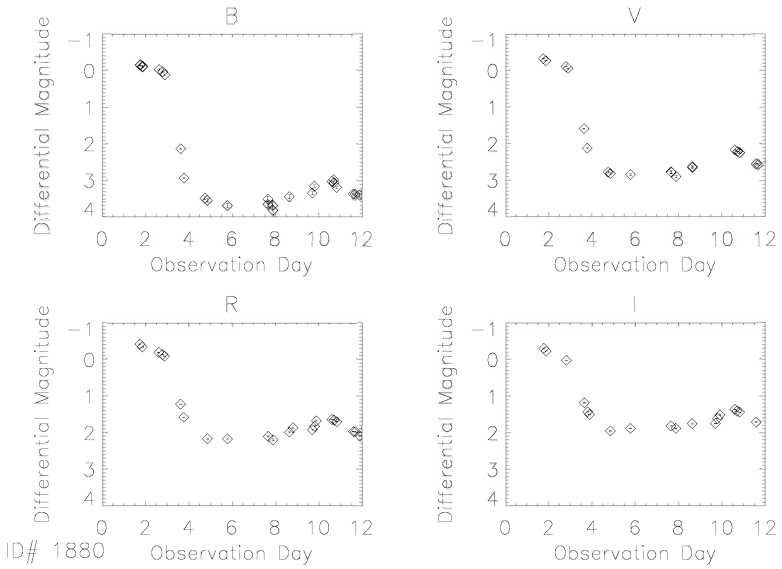

Figure 6: A lightcurve with the largest change in magnitudes measured in our data set. On the magnitude scale, a difference in four magnitudes (seen in the B filter lightcurve) corresponds to a relative change in flux of 40 (i.e., the point of maximum brightness if 40 times brighter than the point of minimum brightness).

$$\Delta B = (B_{truemax} + B_{offset}) - (B_{truemin} + B_{offset}) = B_{truemax} - B_{truemin}.$$

We can thus utilize our lightcurve data to produce variability amplitudes in all four bandpasses by subtracting the point of maximum brightness from the point of minimum brightness (giving us an intrinsically positive quantity since faint magnitudes are more positive than relatively brighter magnitudes) for each star in each bandpass.

While our method for measuring variability amplitudes is simple and requires little calculation (all that is involved is the subtraction between two data points), it is not without its drawbacks. Our first concern when considering this method was that our measurements would not reflect the true difference between the maximum and minimum brightness (i.e., it is possible that the actual maximum or minimum brightness for a given star is missed by our observing frequency). To quantify this concern, we generated perfectly sinusoidal lightcurves with phases in the range of our stellar periods; we then calculated the probability for

one observation in our observing window to fall at the maximum or minimum brightness given the errors on our measurements. For each simulated lightcurve, we can assign error bars to the points on the lightcurve by assuming the amplitude of variation to be at least 3 times greater than root-mean-square error on the two measurements used to calculate the variability amplitude. (For instance, on our real lightcurves one may say that we have a statistically significant measurement of the variability amplitude if, for example, $\triangle \mathrm{B} \geq 3\sigma_{\triangle B}$, where $\sigma^2_{\triangle B} = \sigma^2_{Bmax} + \sigma^2_{Bmin}$, and so we take the 3σ approximation as the upper limit to our error bars on the simulated lightcurve.) From this assumption, we can calculate the theoretical error bars on our simulated lightcurves, and, from that, calculate the amount of time a star with the simulated lightcurve would spend at its maximum or minimum brightness within the range of the error bars. Given this time, we can compute a probability that at least one of our observations falls within the range of error bars around the maximum or minimum brightness. As an example, the average star will spend approximately 25% of its time within the error range surrounding its minimum brightness. At this point, we could naïvely assume that our sampling is random and just compute a Poisson statistic based on the average amount of time a star would spend in within the range around the maximum and minimum brightness with a total of 20 observations taken over the course of the observing run. In reality, our sampling is not random. Our observations are random during the course of the night (the observations were taken in an order to assure such a random sampling), but the fact that we can only observe at night limits the amount of random sampling we can achieve. For instance, it is possible that a star we observe always reaches one of its extreme values of flux during the night and the other during the day. In that case, we are surely missing both the maximum and minimum brightness in the lightcurve. For long period stars, however, our sampling is random in the sense that we have no fear that our sampling rate and period align with each other such that we could miss one of the extreme values of flux. Thus, for a short period star we may have as few as three random samples (for three observations during the night), while for a long period star we may have as many as 20 random samples (for the 20 observations made over the course of the observing run). Poisson statistics tell us that the probability of having at least one observation fall within the error range surrounding the minimum or maximum is 52.8% for the short period stars and 99.3% for the long period stars. From this analysis we see that our method becomes more accurate for stars with longer periods. We note, however, that there is little we could do to stop the possibility of missing one of the extreme values of flux for the short periods stars; even other methods of finding the maximum and minimum brightness, such as fitting a sinusoid to the

lightcurve and then using the sine function to find the maximum and minimum brightness, are plagued by the sampling problem of the short period variables. In short, this analysis reduces some of the concern that our measured maximum and minimum brightness do not reflect the actual maximum and minimum brightness of a lightcurve. To gain more confidence in our ability to accurately measure the variability amplitude of a lightcurve we must address some other issues regarding our method.

A second major issue concerning the effectiveness of the ability of our method to measure the true variability amplitude is the effect of stray data points on our measurements of maximum and minimum brightness. As our data set is rather large, we sought an automated method to determine the points of maximum or minimum brightness in our lightcurves; nevertheless, this automated method cannot rule out the possibility that a point deemed to be the maximum or minimum brightness is an erroneous data point not taken out during our Honeycutt derivation of lightcurves. Often, these erroneous data points appear to be distinct from the rest of the points in the lightcurve by visual inspection, and, as such, they often appear several magnitudes (or tenths of magnitudes for less variable stars) brighter or fainter than the next brightest or faintest point. The problem of erroneous data points is complicated by the fact that there are certain types of variability, caused mostly by coronal flares, that show a one-point jump in brightness by several magnitudes or tenths of magnitudes. We can distinguish these true data points from erroneous data points through a comparison of the points of maximum or minimum brightness for the lightcurves of a star in each filter. For instance, a coronal flare causing a dramatic one-point change in variability will also cause a dramatic one-point change in variability in all four lightcurves (one lightcurve for each filter), whereas an erroneous data point is more likely to appear in only one lightcurve among the four lightcurves we have for each star. We automated our search for these erroneous data points through the use of a variability statistic modified from the work of Welch and Stetson (1993). Our statistic takes the measured maximum (or minimum) magnitude in one filter, computes the difference between the maximum magnitude and the weighted mean magnitude of measurements for the lightcurve in that filter, and then compares that difference to the sum of the differences between the magnitude of every data point in the lightcurve and the weighted mean magnitude. The comparison we are left with is a ratio, denoted below as R, that relates the variability of the point of maximum brightness to the total variability of the lightcurve, computed, by example in the B band,

as follows:

$$R = \frac{\delta B_{max}}{\sum_{i=1}^{n} \frac{|\delta B_i|}{n}}$$

where n equals the total number of data points in the lightcurve and

$$\delta B_i = \frac{B_i - B_{mean}}{\sigma_i}, \delta B_{max} = \frac{B_{max} - B_{mean}}{\sigma_{max}}$$

with

$$B_{mean} = \frac{\sum_{i=1}^{n} \frac{B_i}{\sigma_i^2}}{\sum_{i=1}^{n} \frac{1}{\sigma_i^2}}$$

We can then compute this ratio for a star in every filter. If a dramatic one-point change in variability occurs in all filters (as it would for a real variability event such as a coronal flare), the ratio computed by our variability statistic should be of roughly the same value in all filters. If a dramatic one-point change in variability occurs in just one filter (which is what we expect from erroneous data points), the ratio computed by our variability statistic will differ significantly. For a real data point, a dramatic one-point change in variability will appear in each lightcurve such that $R_B/R_V \approx 1$ (any two filters, not just B and V, should also have a ratio of R's approximately equal to one). With this statistic, we can isolate variability measurements drawn from lightcurves with one or two spurious data points by throwing out those measurements yielding a ratio greater than two or less than one-half.

Finally, we believe our method for measuring variability amplitudes is justified by comparison to the effectiveness of other methods for measuring variability amplitudes. In particular, we considered whether or not our variability amplitude would be more accurate if we first fit a sinusoidal function to the data and then used that sinusoid to compute the variability amplitude. The function fitting method is problematic for two major reasons. First, by assuming all lightcurves to be sinusoidal, we would be ignoring lightcurves exhibiting irregular variability. Granted, the lightcurves we are most interested in are those for stars with known periods, which are unlikely to exhibit irregular variability, but we would still lose valuable information about the accretion activity of stars with irregular variability. (Note: At some level the model for stars with

irregular variability cannot be described by a single spot rotating on the stellar surface, but a change in amplitudes from bandpass to bandpass will still tell us if the observed variability is being caused by some process increasing or decreasing the temperature of the photosphere.) Second, sinusoidal fits to stars with short periods would be worse than sinusoidal fits to stars with long periods. Like the method we use to calculate variability amplitudes, the accuracy of the method of fitting a sinusoidal function to the data is biased toward longer periods.

Having discussed in detail the benefits and drawbacks to our method of calculating variability amplitudes from our lightcurves, we can use our measured variability amplitudes, Δ, and the error for each amplitude, $\sigma_{\Delta}^2 = \sigma_{max}^2 + \sigma_{min}^2$, as comparison data for our spot model.

3 The Spot Model and Its Applications

3.1 Spot Model Theory

The spot model we employ on our photometric variability measurements is simple in concept. Supposing we have a perfect lightcurve (i.e. a lightcurve wherein we have a measurement of a star's magnitude at every point in time), we can measure the amplitude of variability by subtracting the point of maximum brightness from the point of minimum brightness (a positive value since magnitudes corresponding the brighter points (more flux) are smaller than magnitudes corresponding to the fainter points (less flux)). This amplitude is a ratio of fluxes just like any other magnitude difference. If we can represent this flux ratio as some function of the spot temperature, stellar photospheric temperature (effective temperature), and spot size, then we can use the measured variability amplitude of a single star from lightcurves at several bandpasses to determine the spot temperature and spot size.

In the simple case where one extreme of the lightcurve represents our view of the star with a spot while the other extreme represents our view of the star without a spot, the flux ratio between our view of the star

with the spot and the star without the spot is given by:

$$\frac{B_\lambda(T_{star}) - f * B_\lambda(T_{star}) + f * B_\lambda(T_{spot})}{B_\lambda(T_{star})} = 1 - f\left[1 - \frac{B_\lambda(T_{spot})}{B_\lambda(T_{star})}\right]$$

where $B_\lambda(T_{star})$ is the Planck blackbody function for the stellar photospheric temperature at a given wavelength , $B_\lambda(T_{spot})$ is the Planck blackbody function for the spot temperature at the given wavelength and the temperature of the spot, and f is the areal coverage of the spot as a fraction of the total stellar area. If we expand this simple case to one where we assume we know nothing about the location of the spot, including whether or not it disappears completely from our view, we can express our variability amplitudes as:

$$\Delta m(\lambda) = -2.5log\left\{\frac{1 - f_2[1 - Q(\lambda)]}{1 - f_1[1 - Q(\lambda)]}\right\}$$

where $Q(\lambda) = B_\lambda(T_{spot}) / B_\lambda(T_{star})$ and f_1 and f_2 represent the minimum and maximum spot areal coverage we see, respectively (Vrba, Herbst, & Booth 1988; Bouvier et al. 1993; Stassun & Wood 1999). We can reduce this equation from one of three variables to two variables by combining our minimum and maximum spot areal coverage into some effective spot areal coverage. Note that, if we make a substitution for an equivalent spot size as

$$F_{eq} = \frac{f_2 - f_1}{1 - (1 - Q(\lambda)) * f_1}$$

then our variability amplitudes can be re-expressed as:

$$\Delta m(\lambda) = -2.5log[1 - (1 - Q(\lambda)) * F_{eq}(\lambda)].$$

We can expand the detail of the model given above (Bouvier et al. 1993) by incorporating stellar limb-darkening through $F(\lambda)$. If we note that $F(\lambda) = \frac{F'}{\pi * (1 - \mu(\lambda)/3)}$, where F'/π is the the projected spot area as a fraction of the stellar disk surface and $\mu(\lambda)$ is the limb-darkening coefficients, we have an expression for the spot size that changes appropriately with wavelength (Torres & Ferraz Mello 1973). For the B, V, R, and I bands, the stellar limb-darkening coefficients are 1.0, 0.89, 0.68, and 0.53, respectively (Bouvier, Bertout, &

Bouchet 1986). The final form of the variability amplitude equation we use in our model is:

$$\Delta m(\lambda) = -2.5log\left[1 - (1 - Q(\lambda)) * \frac{F'}{\pi * (1 - \mu(\lambda)/3)}\right].$$

Our model ignores the effects of opacity differences between the stellar photosphere and the spot, and inclination effects (Stassun & Wood 1999; Carpenter, Hillenbrand, & Skrutskie 2001). Our model also assumes that both the star and spot radiate as perfect blackbodies. We assume nothing about the size, location, or number of spots; rather, we assume that we see some minimum and maximum spot area coverage for one spot or a number of spots with the same temperature. Given variability amplitudes at multiple wavelengths, the spot temperature and equivalent spot areal coverage can be determined by matching theoretical variability amplitudes, calculated from a range of spot temperatures and spot sizes, to observed variability amplitudes. One can see that the sign of the magnitude measurement depends on the temperature of the spot relative to the photospheric temperature. For instance, spots cooler than the photosphere will yield Q(λ) less than one and thus give a flux ratio less than one, producing a positive amplitude measurement; spots hotter than the photosphere will have the opposite result, producing a negative amplitude measurement. Without any a priori knowledge of the true spot temperature, this sign ambiguity can be overcome by comparing the absolute magnitude of variability amplitudes calculated in the model to those values measured in our lightcurves. Earlier applications of this model to lightcurves from a small sample of stars or to lightcurves from theoretical stellar photospheres with artificially added spots find that variability amplitudes always decrease with longer wavelengths, but the rate of decrease is unique to the spot temperature relative to the photosphere, i.e. photometric amplitudes due to hot spots decrease with wavelength at a rate that differs from the decreasing rate of cool spots (Bouvier et al. 1993). Stassun and Wood (1999) find that this result allows them to consistently distinguish hot spots from cool spots with little uncertainty for model stellar photospheres with artificial spots added to the surface. Unlike previous studies, this thesis will apply this model to a large set of stars (approximately 2,000), of which approximately 500 have known periods. (Previous studies examining the spot temperatures of T Tauri stars (Vrba et al. 1989; Bouvier et al. 1993; Vrba et al. 1993) have applied this model to fewer than 40 stars.) The purpose is to provide a data set from which details of stellar accretion may be more accurately determined. The details of accretion determined from the model spot temperatures (hot spots can be interpreted as a sign of active accretion from the circumstellar

disk and active accretion interpreted as a sign of a magnetic field threading a disk) can then be used in comparison with rotation period data to investigate what links, if any, exist between stellar rotation and star-disk interaction. The ultimate purpose is to interpret the accretion and rotation period data in regard to the questions raised by the angular momentum conundrum. The hope is that we may gain further insight about the role of circumstellar disks in the rotational evolution of stars and the effectiveness of the magnetic disk-locking model.

3.2 Applying the Spot Model to Measured Variability Amplitudes

3.2.1 Non-linear Least-Squares Fit to Data

There are several methods through which one could find the ideal combination of spot parameters (temperature and equivalent spot size) for which the observed variability amplitudes matches the variability amplitudes returned by the spot model. Our approach to finding the spot parameters that best describe our observed variability is to calculate chi-squared fits to our data for the variability amplitudes generated by the spot model at a variety of spot parameters, with the lowest value of chi-squared corresponding to the spot parameters that best describe the observed variability. A brute-force way to find the best-fit spot parameters is to run the model on a fine grid of spot temperatures and spot sizes. The brute-force method is perhaps the simplest way to characterize the observed variability with spot parameters, but it is not without its drawbacks. The primary drawback to the brute-force method is the computation time required to compute the chi-squared values for a fine grid of spot temperatures and spot sizes. On the other hand, the brute-force method allows us to know the value of chi-squared at every point in parameter space, depending on how fine we choose to make our grid of spot temperatures and sizes. With every value of chi-squared known for every point in parameter space, we can make accurate judgments of the range of best fit spot parameters. These judgments are also time consuming, and so we look to other methods that can mathematically map the chi-squared space and determine the uncertainties on our best-fit spot parameters. The CURVEFIT program in IDL provides us with this mathematical method, and so we choose to use it to compute the spot parameters that will best match the spot model to our observations.

In short, IDL's CURVEFIT program computes a non-linear least-squares fit to our data. The program starts with an initial guess for the spot temperature and spot size and then uses the partial derivatives of the variability amplitude function given in section 3.1 with respect to the spot temperature and the spot size to compute the next step in parameter space. At each step, the program computes a chi-squared fit of the data to the spot model for the parameters at that step. As the program iterates through steps in parameter space, it records the chi-squared information about its previous steps and arrives at a parameter solution that yields the lowest value of chi-squared. The spot parameters corresponding to the lowest value of chi-squared are then returned to the user. The errors on these spot parameters are a reflection of the chi-squared values for the numerous steps made in the program, and, as such, are a mapping of the chi-squared space surrounding the minimum found in the parameter space (Bevington 1969).

For each lightcurve and set of variability amplitudes, we compute the best-fit spot parameters by supplying IDL's CURVEFIT program with a coarse grid of initial guesses for spot temperatures and spot sizes. As our rotation period and accretion analysis relies on our ability to characterize spots as either hot or cold (both designations are with respect to the temperature of the stellar photosphere), we want to compare the best hot spot model and best cool spot model for each star by forcing our non-linear least-squares fit to converge on only hot spot temperatures if the initial guess supplied is a hot spot and only on cool spot temperatures if the initial guess supplied is a cool spot. For instance, if the photospheric temperature is 3500K and our initial guess is at 2000K, we restrict CURVEFIT to converge on the best-fit spot parameters with a spot temperature less than 3500K. Thus, for each star we receive two best-fit spot parameters, one for the best-fit hot spot model and one for the best-fit cool spot model. In each spot temperature regime, we find that a large number of guesses (usually more than half) converge on the same spot parameters. Other guesses tend to converge on spot parameters with chi-squared values much too high to be considered reasonable models of our observed variability amplitudes. We view these results positively, for they favor the interpretation that there is one global minimum in our chi-squared parameter space corresponding to the best-fit hot spot model and the best-fit cool spot model.

3.2.2 Classification of Model Results

Once we have the spot temperature, spot size, and chi-squared for the hot spot model and the cool spot model of a given lightcurve, the only task we have left to complete before beginning our rotation period and accretion analysis is the classification of spot temperatures. For each lightcurve, we use the values of chi-squared for the hot spot model and the cool spot model to assign the spot temperature that best describes the observed variability amplitudes. If the reduced chi-squared returned for either model is less than our cutoff value of one, we consider that model to be a candidate for describing the observed variability. In some cases, one type of spot is clearly more likely to be the cause of the variability than another. In Figure 7 we see a cool spot model that best describes the observed variability amplitudes. At the top of Figure 7 we see the B, V, R, and I lightcurves for a star. In the lower left-hand corner, we see the variability amplitudes for each lightcurve as well as the matching data for the best-fit cool spot model (dotted lines) and the best-fit hot spot model (dashed lines). In Figure 8 we see the opposite effect, a hot spot model that clearly describes the observed variability. Note that, in each case, the rate of change of the hot spot data is visually distinct from the rate of change of the cool spot data, which is consistent with previous work that finds the rate of change of variability among the two spot regimes to be distinctly different (Bouvier et al. 1993; Stassun & Wood 1999). In some cases, the spot model appears to work well for both the hot spot model and the cool spot model, but, upon closer inspection, the best-fit parameters for one of the models are questionable (see Figure 9). In these cases, we often see that one of the models returns a value for the spot temperature very close to the photospheric temperature and a value for the spot size of unreasonably large size. At spot sizes nearing 100% of the surface of the star one may conclude that the spot model would rather have no spot or a different photospheric temperature than describe the variability in that temperature regime with any type of spot. In the cases where we see suspect model parameters but reduced chi-squared values less than 1.25, we choose the model that returns a spot size less than 50% of the stellar surface. In addition to the cases mentioned above, we see two more types of model fits: those where the spot model fit is truly ambiguous (i.e., both models return reasonable chi-squared values and have reasonable model parameters) and those where neither model yields a good fit to the observed variability. In the former case (see Figure 10), the observed variability is often small and both models return spot temperatures that are relatively close to the photospheric temperature of the star (within 500K).

At this point it is worthwhile to briefly discuss the hot spot temperatures that we find with our model. Upon examining the examples of our model fits and the analysis presented in Chapter 4, the astute reader may wonder why some of our hot spot temperatures are so low. Typical spot temperatures derived from accretion are several thousand Kelvin above the photosphere temperature, yet, while we derive spot temperatures that are several thousand Kelvin hotter than the photosphere, we also derive spot temperatures that are only a few hundred Kelvin hotter than the photosphere. The astronomer familiar with accretion activity may claim that the spot temperatures that are only several hundred Kelvin hotter than the photosphere are not spots related to accretion activity. To this we remind the reader that our model describes only one spot on the photosphere that describes all possible spots that may be on the photosphere. (For instance, a star may have hot and cool spots on its surface, but our model assumes only one spot of one temperature. This single spot temperature may reflect some combination of temperatures for the multiple spots on the surface.) Also, in our analysis, we choose to run correlations of spot temperatures versus rotation period, disk presence, and accretion activity, and we make no a priori assumption that we are correlating accretion activity versus rotation period and disk presence. The analysis presented in Chapter 4 will better explain our methodology, and the reader will see that many of our hot-spotted stars show evidence for accretion activity, regardless of the spot temperature relative to the photosphere. Finally, we would like to remark that previous studies involving our model (namely, Bouvier et al. 1993, 1995) have assumed all hot spot temperatures, including those that are only a few hundred Kelvin, to be the result of some accretion activity. At the very least, our model tells us that any hot spot, regardless of absolute temperature, is the result of some physical process increasing the temperature of the photosphere.

4 Examining the Distribution of Spot Temperatures with Rotation Period, Disk Signatures, and Accretion Signatures

The purpose of this chapter is to address the issues raised in the introduction to this thesis. In particular, we hope to further investigate the angular momentum evolution of young stellar clusters with our spot temperature model. It is our intention to use our derived spot temperatures to perform analysis similar

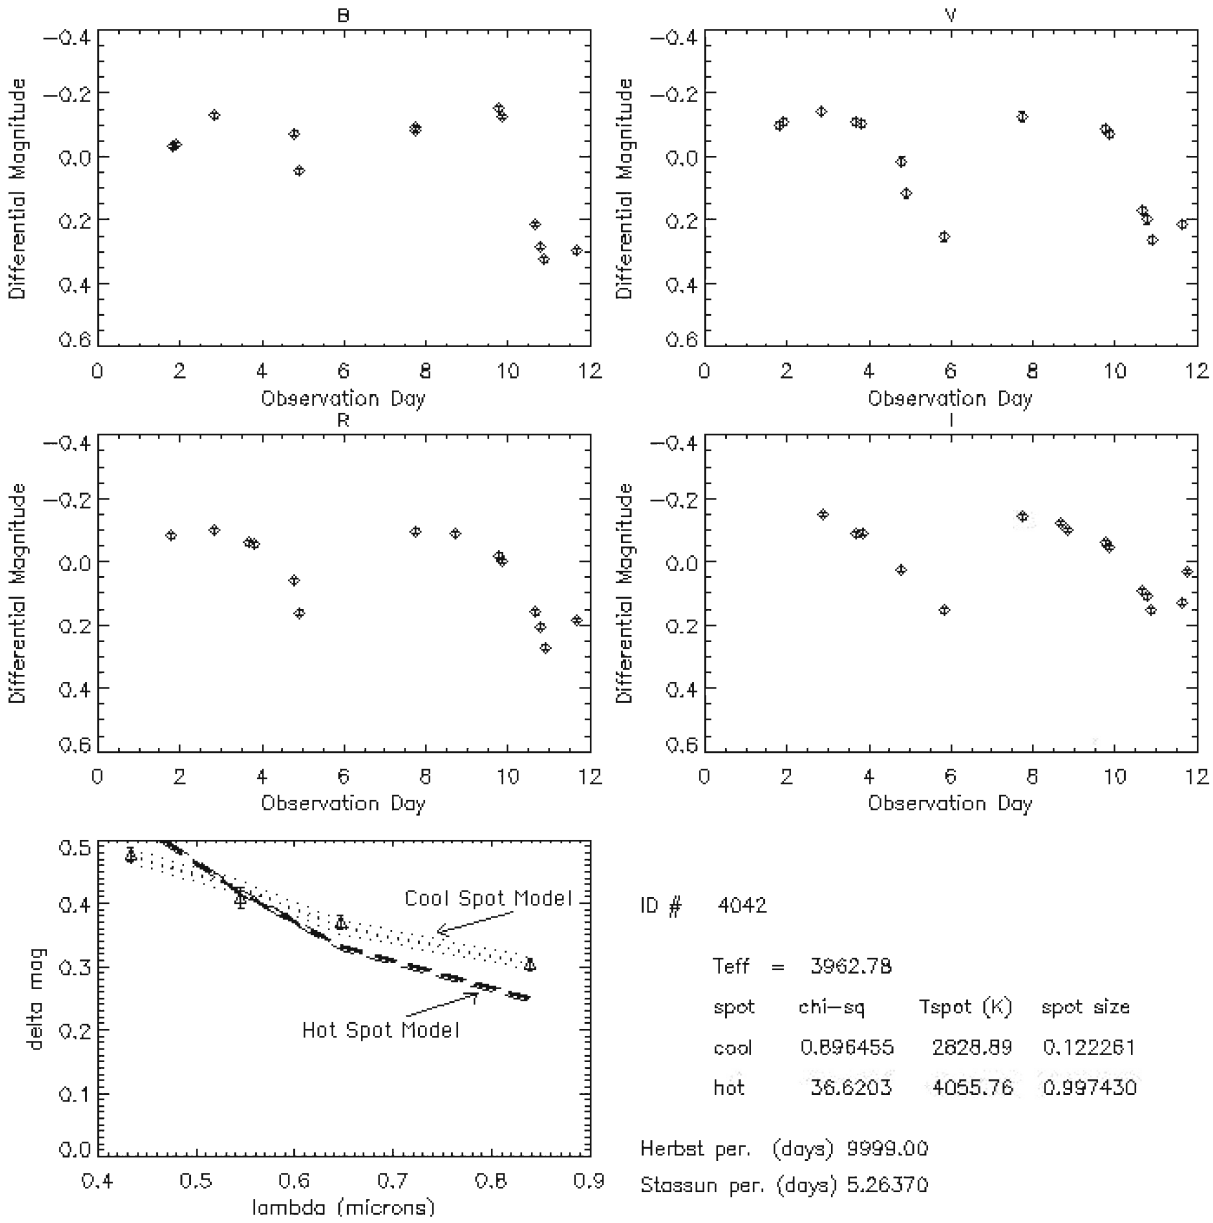

Figure 7: An example of a model fit to data where the best-fit cool spot model clearly describes the variability better than the best-fit hot spot model. In the top four panels we see the B, V, R, and I lightcurves for a given star. In the lower left hand panel we see the measured variability amplitudes for the lightcurves at each bandpass (each bandpass corresponds to a wavelength in the bottom graph, with B having the smallest wavelength, V the next smallest wavelength, and I the smallest wavelength). Overlaid on the variability amplitudes in the lower left hand graph are the data for the best fit cool spot model (dotted lines) and the best fit hot spot model (dashed lines). The two outer lines for each model correspond to our errors on the model calculations. In the lower right hand portion of the graph is information regarding the star's rotation period, effective temperature, as well as the parameters returned for the best-fit hot and cool spot models, including the reduced chi-squared for each model fit.

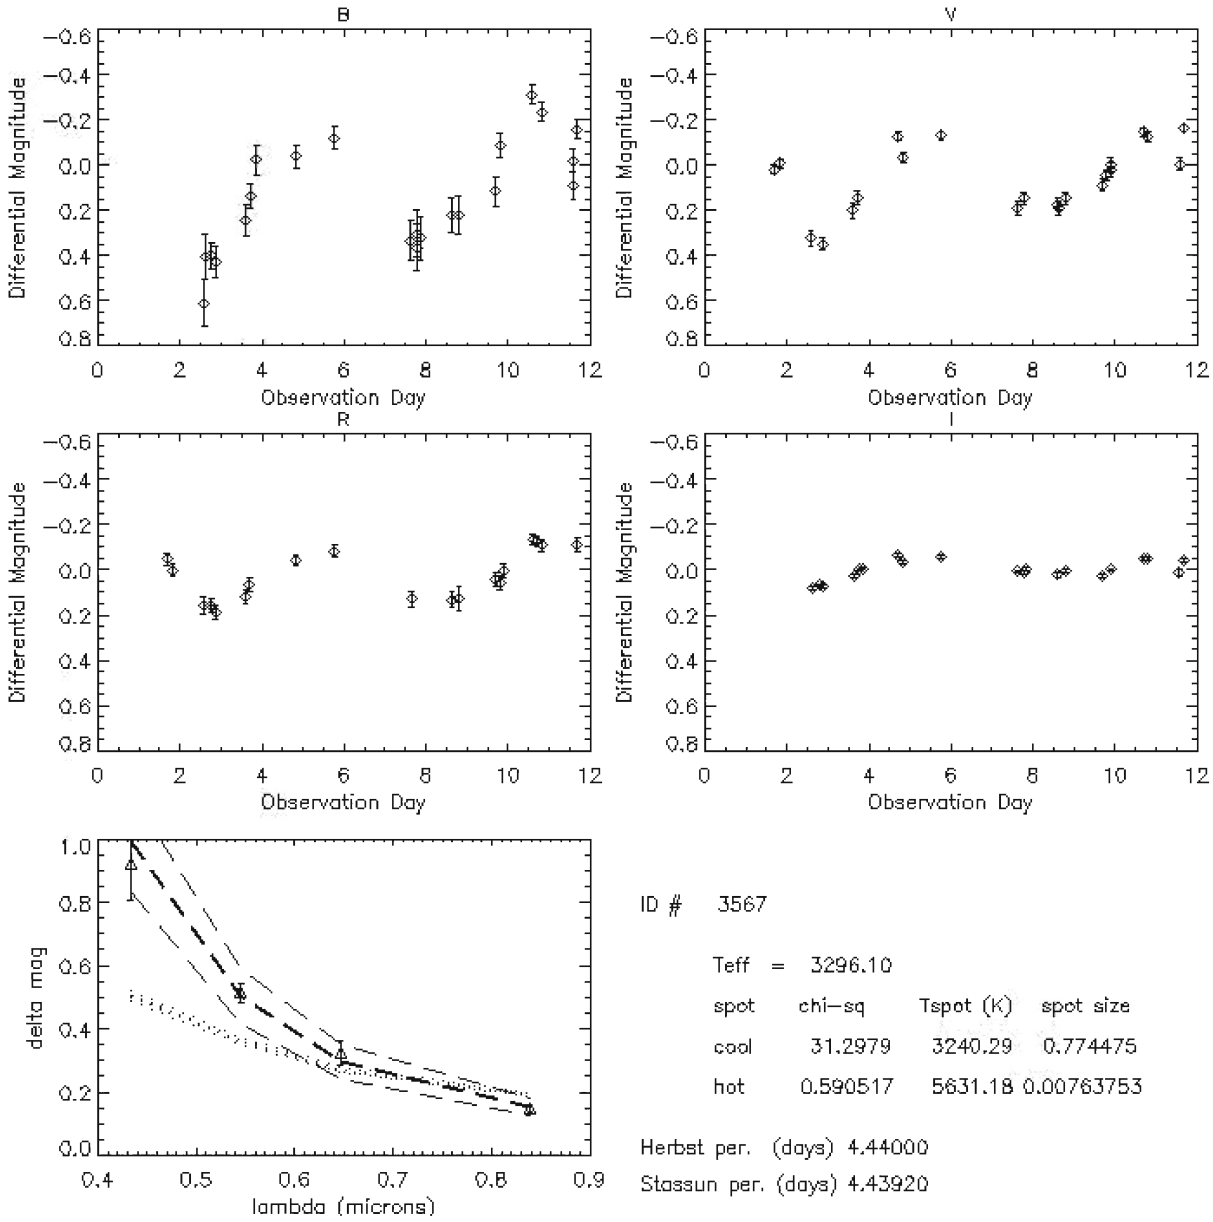

Figure 8: An example of a model fit to data where the best-fit hot spot model clearly describes the variability better than the best-fit cool spot model. See the caption of Figure 7 for a description of the graphs in this figure.

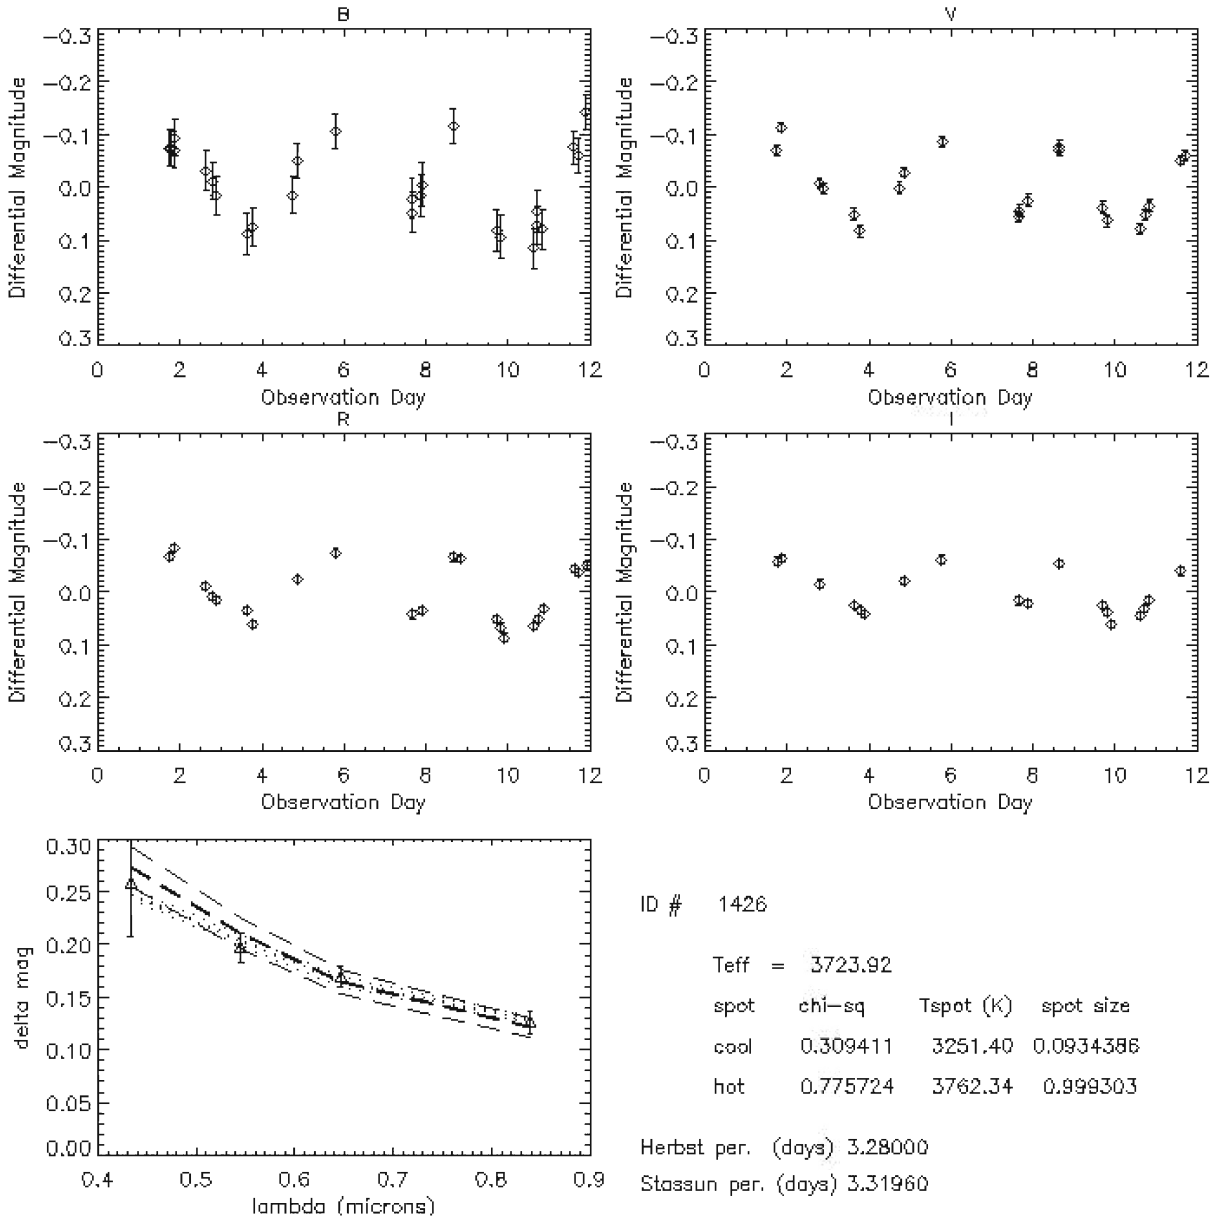

Figure 9: An example of a fit where both the cool spot model and hot spot model appear to describe the variability well, yet the unusual spot parameters returned for the hot spot model (a spot temperature nearly the same as the photospheric effective temperature and a spot size nearly the size of the entire star) lead us to believe that the cool spot model better describes the observed variability. See the caption of Figure 7 for a description of the graphs in this figure.

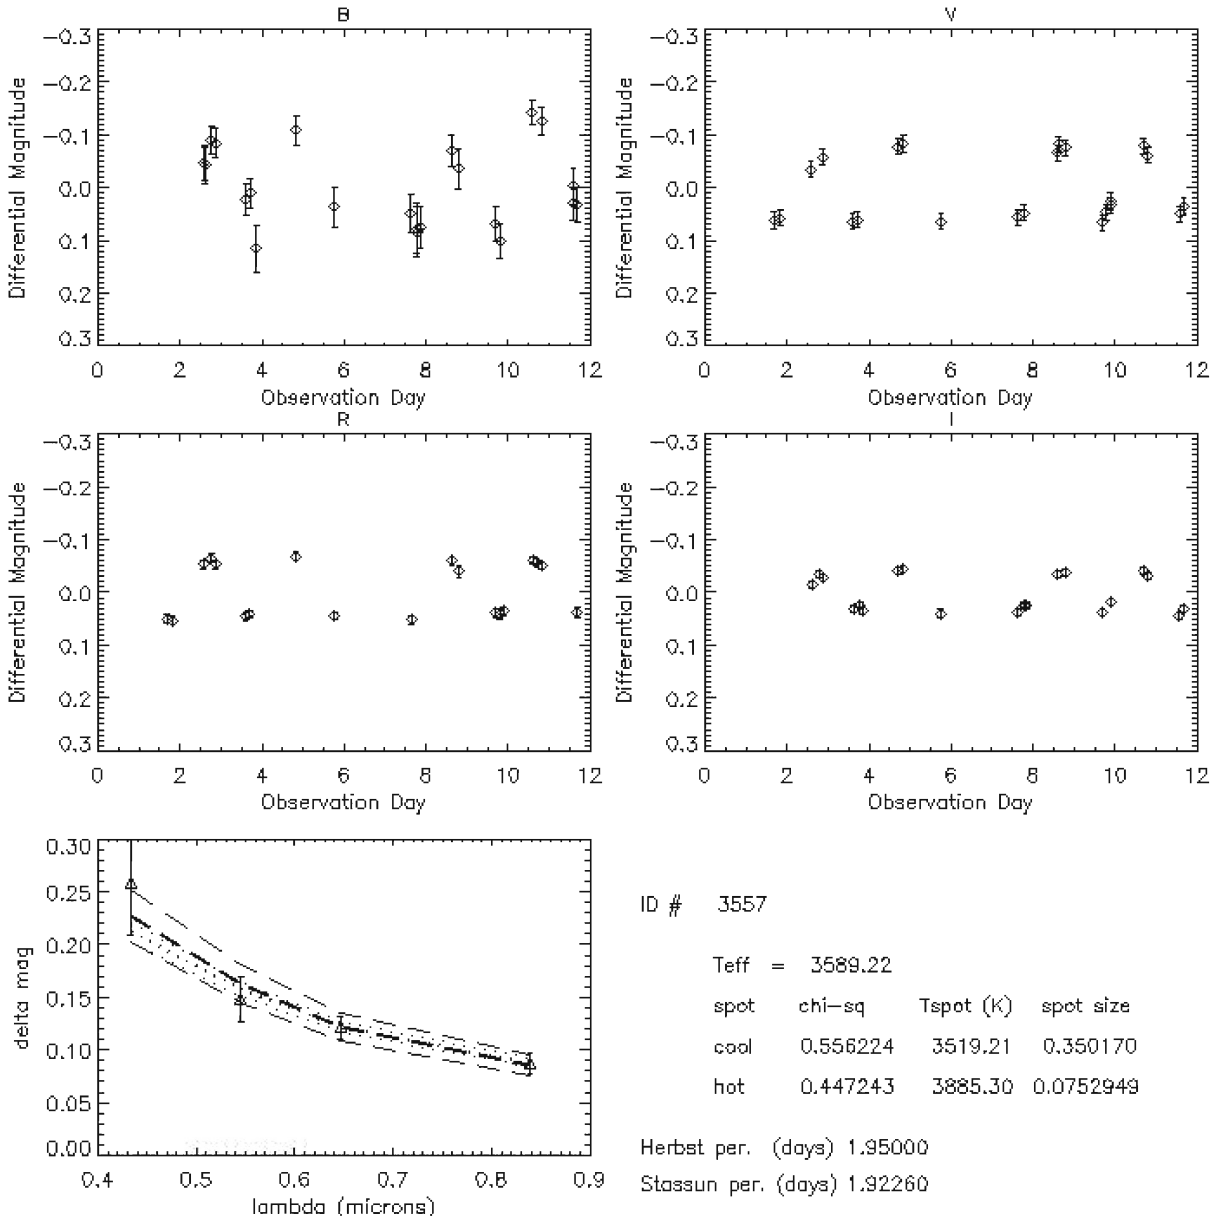

Figure 10: An example of a fit where both the hot spot model and the cool spot model describe the variability well. In this case, a final spot temperature is left unassigned. See the caption of Figure 7 for a description of the graphs in this figure.

to the analysis presented in earlier studies of the angular momentum evolution of the ONC. To properly address the role of magnetic disk locking in the angular momentum evolution of young stars, our primary concern will be to check for correlations between spot temperatures and rotation periods. Like many studies investigating the angular momentum evolution of young stars, the question we seek to answer is whether or not the rotation period distribution is bimodal or unimodal. In other words, do stars with hot spots show a rotation period distribution that is different from stars with cool spots, or is the rotation period distribution of stars with hot spots similar to the rotation period distribution of stars with cool spots? We will then ask whether or not evidence for a circumstellar disk through near-infrared excess (as measured by Δ(I-K)) correlates with spot temperature. Finally, we will check to see if evidence for accretion through Ca II emission correlates with spot temperature.

To look for relationships between spot temperature and rotation period, near-IR excess, and Ca II emission requires a database of information from which we can draw information about stellar rotation rates, disk presence, and accretion activity. Our study has focused on the Orion Nebula Cluster precisely because such a database exists. In our analysis, we use the rotation periods taken from the work of Stassun and collaborators (SMMV) and Herbst and collaborators (Choi & Herbst 1996), stellar masses derived by Hillenbrand (1997), and Δ(I-K) and Ca II equivalent widths measured by Hillenbrand and collaborators (Hillenbrand et al. 1998).

The analysis presented in the next three sections relies heavily upon the comparison between two distributions (a distribution for the stars with hot spots and a distribution for the stars with cool spots). We choose to compare these two distributions by performing a Kolmogorov-Smirnov test. The general purpose of the K-S test is to determine whether or not two distributions come from the same parent distribution. For instance, we would certainly be interested to know whether or not the rotation period distributions of stars with different spot temperatures are different, or if the two distributions are likely to be the result of the statistical sampling of a single parent distribution. Because the K-S test makes no assumptions about the distributions under comparison (i.e., the test will not assume that the distribution of rotation periods for stars with hot spots, for example, is Gaussian, uniform, etc.), we consider the K-S test to be a robust statistic for our purposes (Press et al. 1986). As the K-S test utilizes the cumulative distribution function (CDF) of each of the two distributions under comparison to determine the probability that the two distributions come

from the same parent distribution, we choose to show the cumulative distribution function for the hot and cool spot distributions under comparison. The CDF plots displayed in the rest of this section therefore serve as a visual complement to the K-S test probabilities. When looking at the CDF plots, the reader should focus on the point on the x axis where there is the greatest difference between the two distributions, for the difference between the two CDFs at this point is the basis for determining the probability returned by the K-S test.

4.1 Relationships between Spot Temperatures and Rotation Period

Figure 11 shows the distribution of rotation periods for stars with hot spots and stars with cool spots. This figure, like many of the figures presented in this section, shows the hot- and cool-spotted star distributions as they are (top two plots) as well as normalized to the total number of stars with a model fit that yields either a hot or cool spot (the normalized plots are the middle two plots and the total number of stars with a hot or cool spot model fit is the bottommost plot). In Figure 12 we see the cumulative distribution function for the cool-spotted star rotation period distribution (solid line) and the hot-spotted star rotation period distribution (dash-dot line). The K-S test on the hot-spotted star period distribution and the cool-spotted star period distribution tells us that the probability of the two distributions coming from the same parent distribution is 9%. While this probability is low, it is only suggestive of the fact that there might be a difference between the rotation rates of stars with hot spots versus the rotation rates of stars with cool spots. The canonical interpretation would say that the two distributions are distinctly different only if the probability of them coming from the same parent distribution is less than 0.03%. This percentage probability corresponds to the 3σ-level of detection (i.e., if the two distributions did come from the same parent distribution and we were to draw a random sample of rotation periods of the hot- and cool-spotted stars, the likelihood that we would see the two distributions from our random sampling is less than 0.03%). Our 9% probability corresponds to a 1.7σ-level of detection, which does not logically rule out the possibility that our hot spot and cool spot rotation period distributions come from the same parent distribution.

We also find the mean hot spot rotation period (4.5 ± 3.2 days) to be similar to the mean cool spot rotation period (5.4 ± 3.3 days). If we are to interpret the stars with hot spots as stars with active accretion

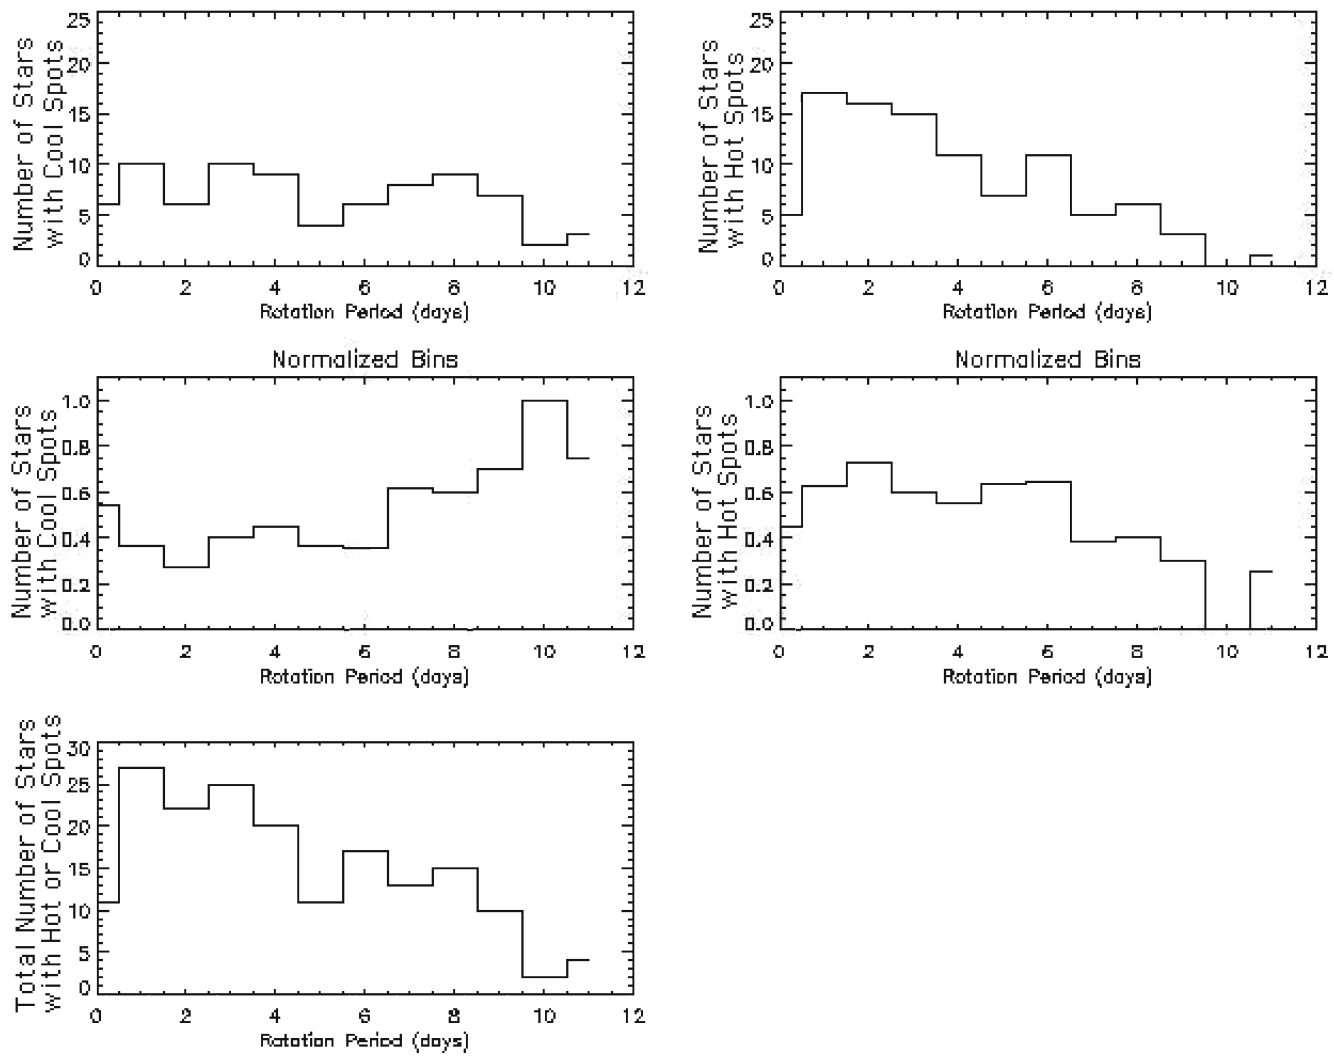

Figure 11: The rotation period distributions of stars with hot spots and stars with cool spots. The top two plots show the total number of stars in a given rotation bin. The middle two plots show the number of stars in a given rotation bin as a fraction of the total number of hot- and cool-spotted stars in that rotation bin. (One can think of the two middle plots as the top two plots divided by the plot at the lower left, where the plot at the lower left is the total number of stars with a hot spot or a cool spot as a function of rotation period.)

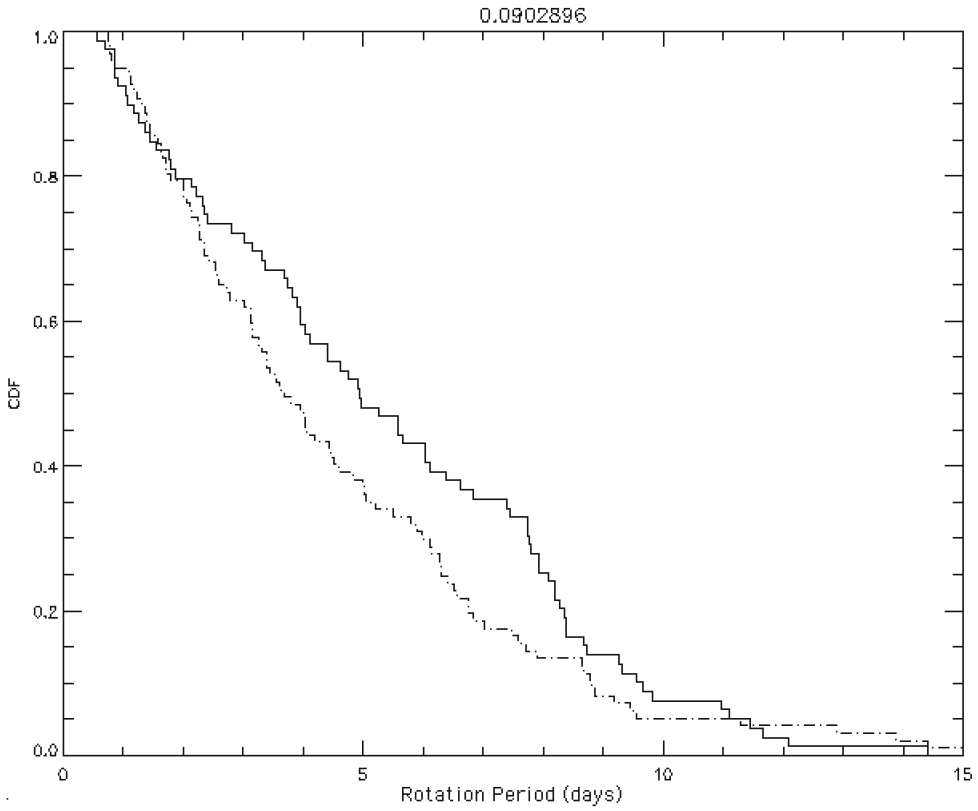

Figure 12: The cumulative distribution function (CDF) for the top two plots in Figure 11 as a function of rotation period. The solid line represents the CDF for the cool-spotted stars while the dash-dot line represents the CDF for the hot-spotted stars. The probability returned by the K-S test run on the distributions is 9%. While the K-S test probability of 9% suggests that the hot- and cool-spotted star rotation period distributions are different, it is far from the 0.3% probability that would allow us to report with confidence that the two are indeed different.

and possible magnetic disk locking, the mean rotation periods derived from our distributions are rather surprising, for our results imply that the stars with probable magnetic disk locking are rotating at a rate similar to, or perhaps even faster than, the stars without disk locking. (Of course, one must take this statement with a grain of salt, for the standard deviation on our mean rotation periods are rather high, suggesting that the two means are indistinguishable.) With our K-S test analysis, we must conclude that our derivation of spot temperatures suggest no correlation of spot temperatures with rotation periods in the ONC. If we interpret our hot spot temperatures as evidence for accretion activity and disk presence then we can say that we see no correlation between disk presence and rotation period. To stop at this point in our rotation period analysis, however, would ignore the arguments made in the literature in favor of a unimodal rotation period distribution for stars with masses below $0.25M_{\odot}$ (Herbst et al. 2000 & 2002; SMMV; Rebull 2001) and in favor of a bimodal rotation period distribution for stars with masses above $0.25M_{\odot}$ (Herbst et al. 2000 & 2002).

4.1.1 Rotation Period Correlations for Mass Constrained Above and Below the $0.25M_{\odot}$ Threshold

Figure 13 shows the rotation period distributions for the hot- and cool-spotted stars with $M < 0.25M_{\odot}$, while Figure 14 shows the corresponding CDF for these distributions. The K-S test on these distributions reveals that there is an 86% (approximately 0.18σ) probability that the two distributions come from the same parent distribution. The mean rotation period for these two distributions also reveal that hot and cool spot distributions are similar: our measured mean of the hot spot distribution is 4.0 ± 2.9 days while the mean of the cool spot distribution is 3.8 ± 3.1 days. Our result is therefore consistent with the findings of Herbst and collaborators (2002), Stassun and collaborators (1999), and Rebull (2001), all of whom report a uniform period distribution for $M < 0.25M_{\odot}$.

In Figure 15 we see the rotation period distributions for the hot- and cool-spotted stars with $M > 0.25M_{\odot}$. Figure 16 shows the corresponding CDF for these distributions. The K-S test for these distributions reveals that there is a 58% (approximately 0.55σ) probability that the two distributions come from the same parent distribution. Like the comparison of the hot- and cool-spotted stars for $M < 0.25M_{\odot}$, the K-S does not support the observation of a bimodal period distribution, and is therefore inconsistent with the findings of

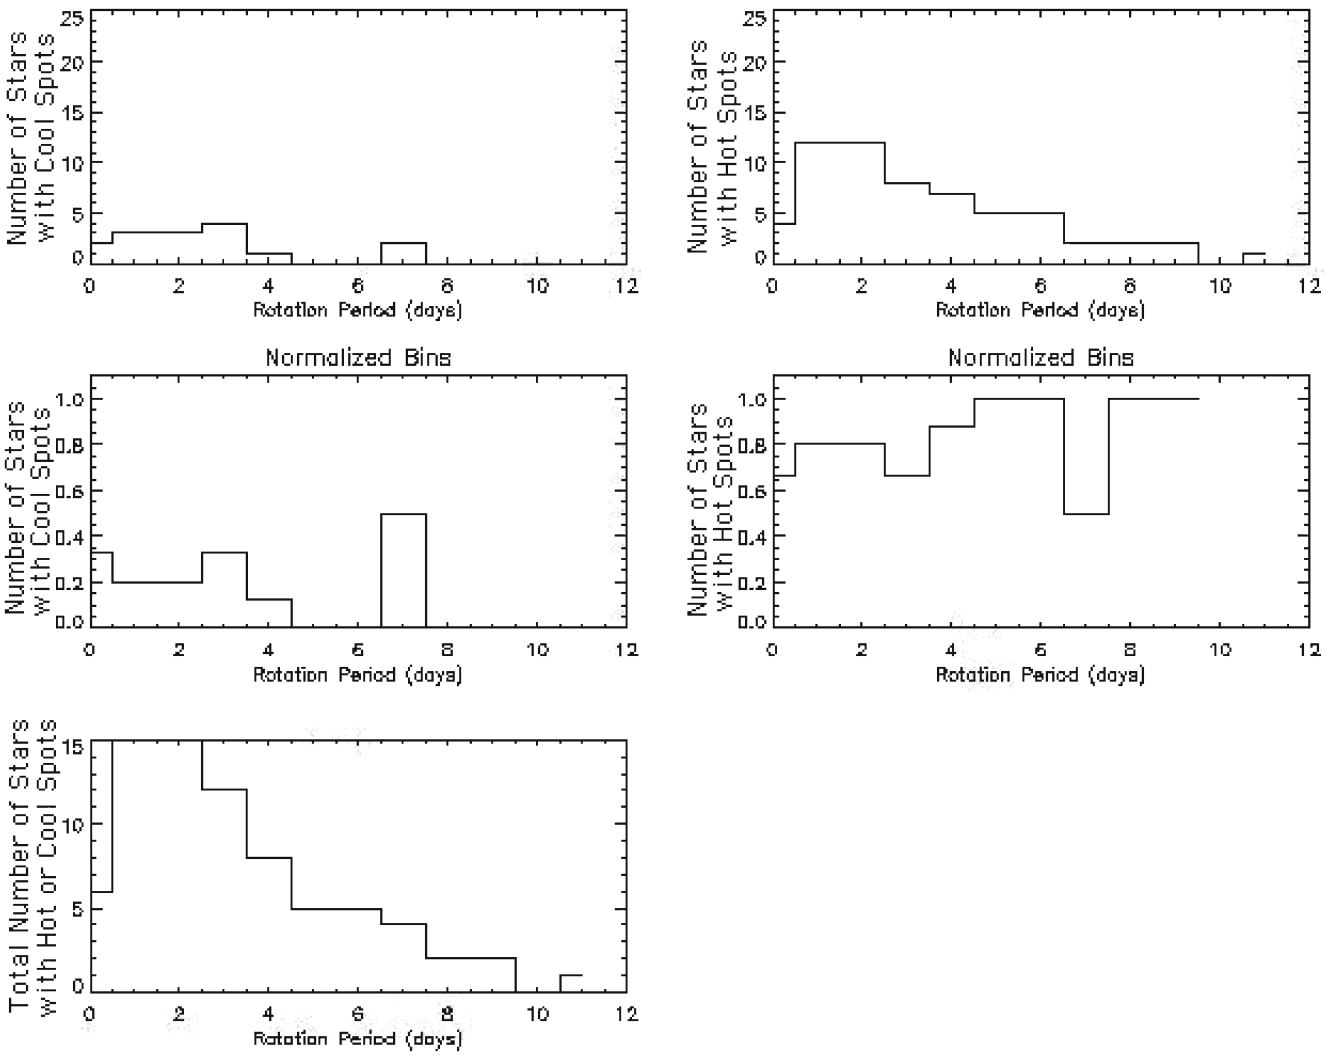

Figure 13: Same as Figure 11, but for stars with masses less than $0.25M_{\odot}$.

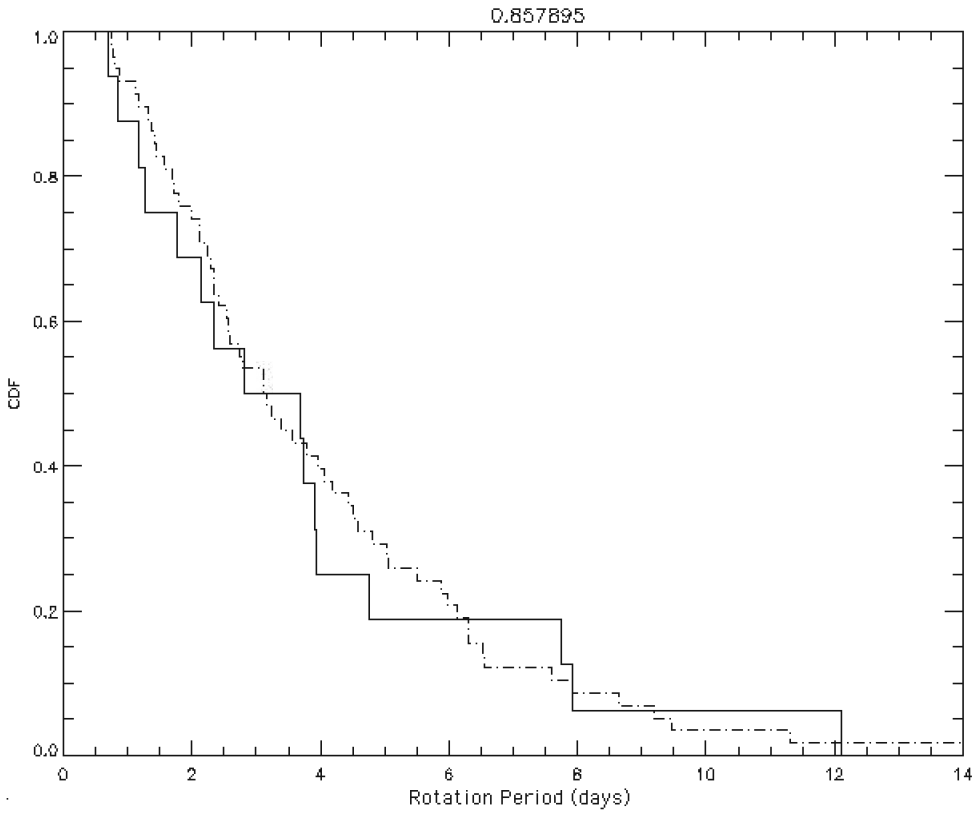

Figure 14: The CDF describing the distributions in Figure 13. As before, the solid line represents the cool-spotted star CDF, and the dash-dot line represents the hot-spotted star CDF. The K-S test for these two distributions suggest that there is no difference between the rotation period distributions for hot- and cool-spotted stars with $M < 0.25 M_{\odot}$.

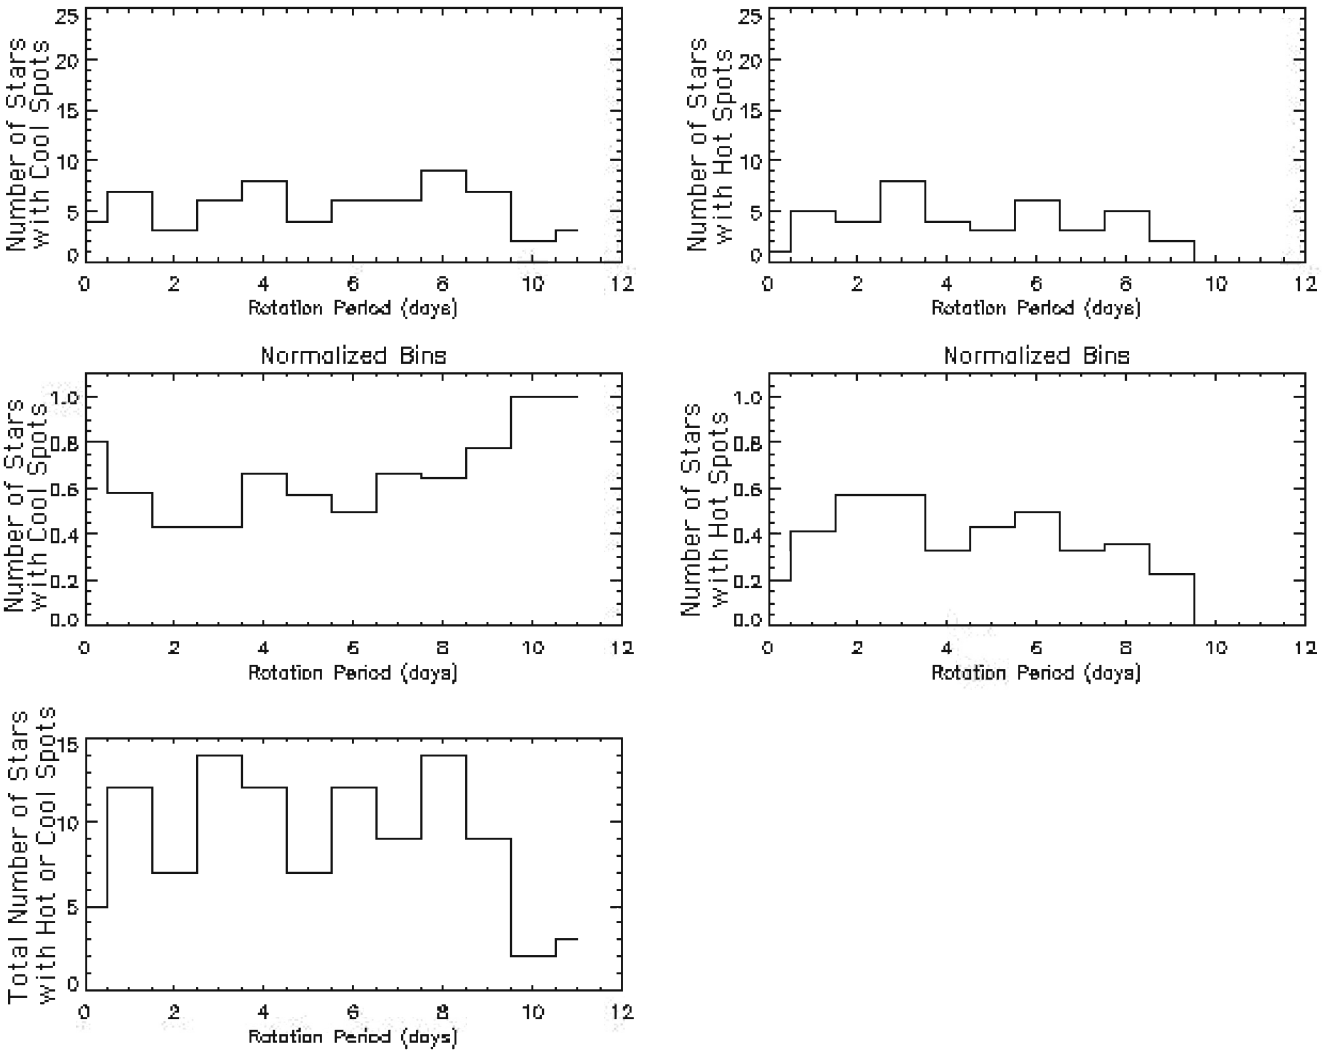

Figure 15: Same as Figure 11, but for M $> 0.25 M_{\odot}$.

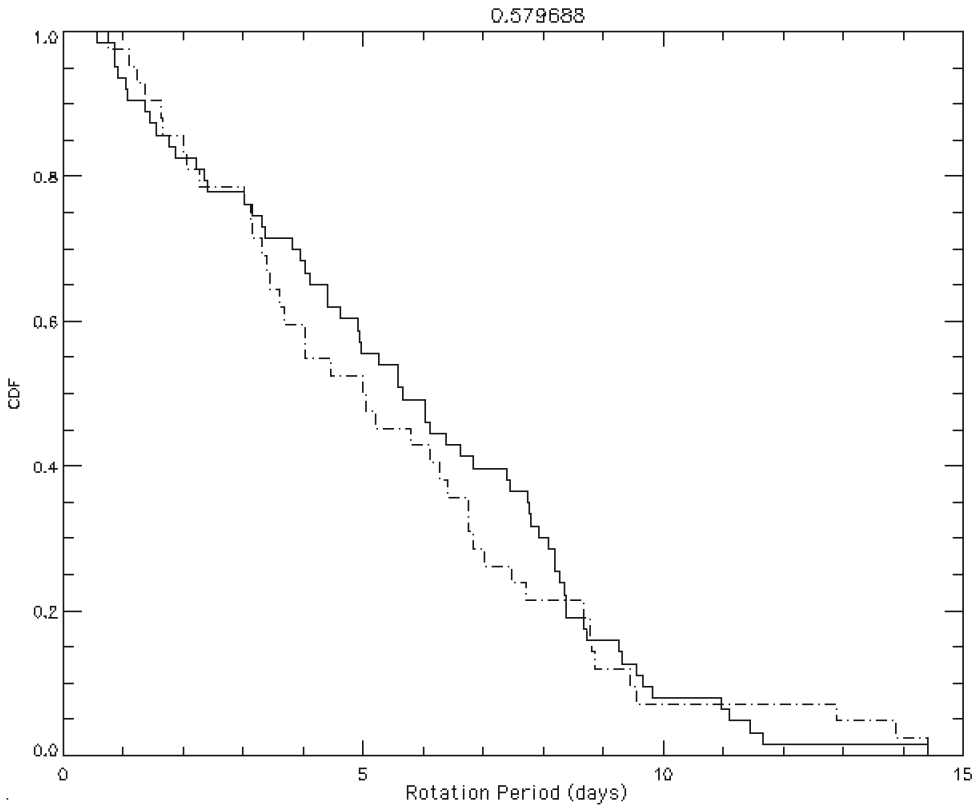

Figure 16: The CDF describing the distributions in Figure 15. As before, the solid line represents the cool-spotted star CDF, and the dash-dot line represents the hot-spotted star CDF. The K-S test for these two distributions suggest that there is no difference between the rotation period distributions for hot- and cool-spotted stars with $M > 0.25M_{\odot}$.

Herbst and collaborators (1996; 2002). Again, the means are indistinguishable from each other. (For the hot spots we find a mean rotation period of 5.5 ± 3.4 days; for the cool spots we find a mean rotation period of 5.8 ± 3.3 days.)

From our relationships between spot temperatures and rotation period, we see no evidence for a correlation between spot temperature and rotation period in the ONC under any of the constraints applied in previous studies. Before we can interpret our results with regard to the magnetic disk locking model, however, we must first address the exact meaning of our derived spot temperatures. The meaning of our derived spot temperatures with regard to the magnetic disk locking model is a delicate issue. If we interpret spot temperatures as in Bouvier et al. (1993; 1995) and Stassun & Wood (1999), thereby ascribing hot spots to evidence for accretion from a circumstellar disk, we cannot ignore the possibility that the accretion we see is the result of some mechanism other than magnetic disk locking. Overall, we believe that accretion resulting from disk material flowing along magnetic field lines to the stellar surface has a strong theoretical justification. Given observations of T Tauri stars with magnetic fields several kilogauss in strength (Guenther et al. 1999), we expect many of these stars to have strong magnetic fields. Also, models describing the magnetic accretion of disk material in classical T Tauri star systems predict the types of periodic variability observed for stars in the ONC and the Taurus-Auriga cloud (Mahdavi & Kenyon 1998). Even considering this evidence, we do not want to rule out the possibility that the accretion we see, under the interpretation of previous studies, results from the gravitational infall of disk material, rather than accretion along magnetic field lines. Thus, before we make any conclusions about the rotation period distribution of stars in the ONC, we wish to consider some other properties of the hot- and cool-spotted stars. Namely, we look to the mass distribution of these stars, the distribution of disk presence as given by the near-IR excess indicator $\triangle$(I-K), and the distribution of accretion activity given by the equivalent width of the Ca II triplet. These properties will help us better understand the meaning of the hot and cool spots derived in our model, particularly in relationship to the meaning ascribed to hot and cool spots derived in previous studies.

4.1.2 Mass Distribution with Spot Temperature

Figure 17 shows the mass distribution for the hot- and cool-spotted stars; the corresponding CDFs for these distributions are shown in Figure 18. The K-S test on these distributions yields a probability of 1.6×10^{-9}%

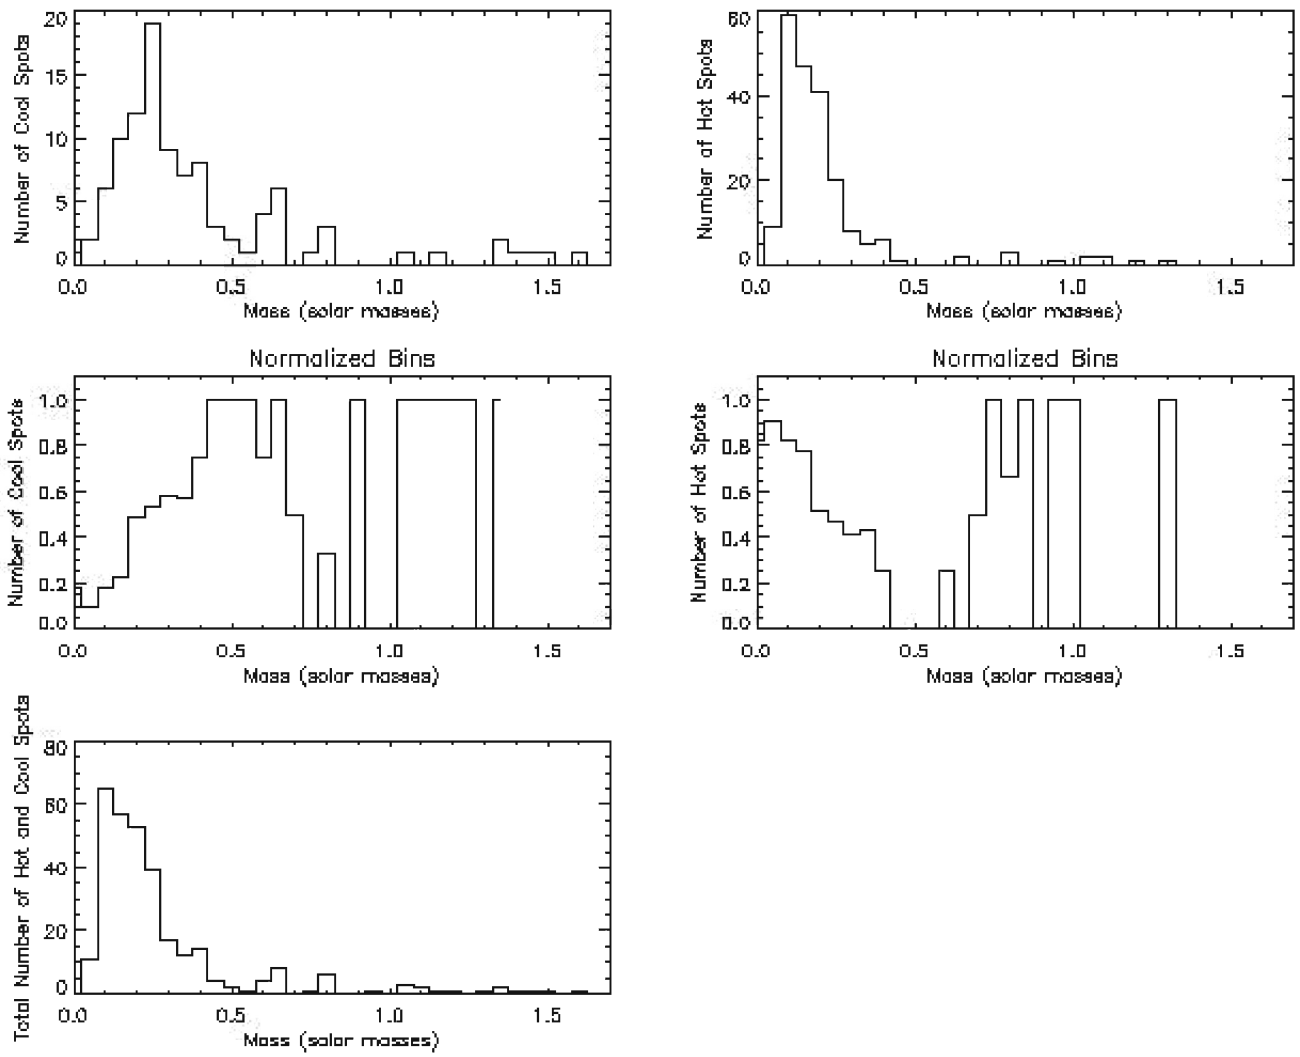

Figure 17: The distribution of hot- and cool-spotted stars with stellar mass. Like Figure 11, the top two plots show the total number of hot- and cool-spotted stars within a stellar mass bin, while the middle two plots show the distributions normalized to the total number of stars with hot and cool spots in the mass bin (shown in the lower left plot). The low-mass distribution of hot-spotted stars may arise from the relatively fast accretion rates of high-mass stars. The fast accretion rates of high-mass stars dissipate circumstellar disks on a faster timescale than low mass stars, and, as such, we should expect there to be less observable accretion activity in high mass stars.

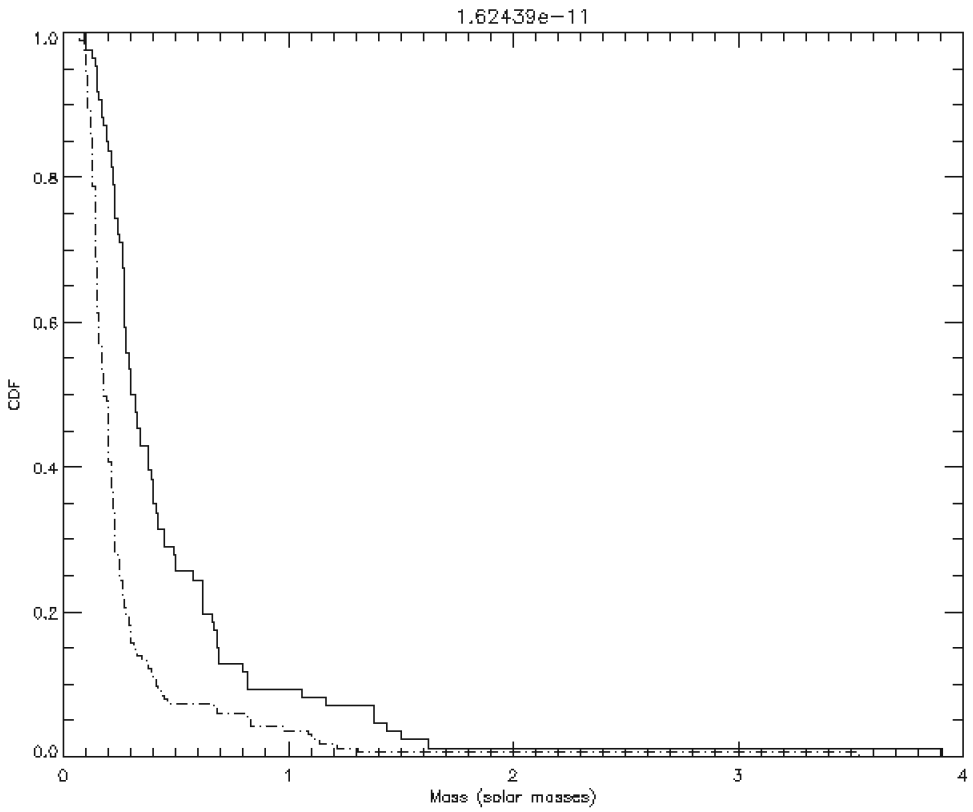

Figure 18: The CDF for the distributions in Figure 17. The low probability returned by the K-S test for the hot- and cool-spotted star mass distributions tells us that the mass distribution of hot-spotted stars is distinctly different from the mass distribution of cool-spotted stars.

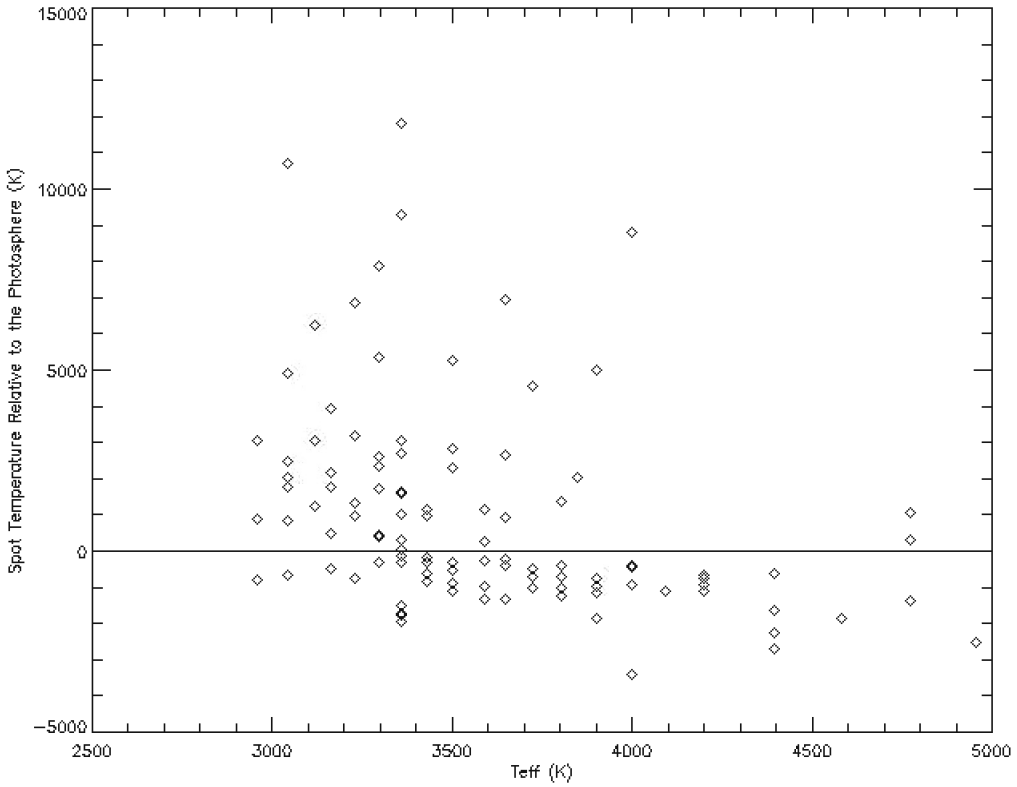

Figure 19: The distribution of spot temperatures as a function of stellar effective temperature. Looking at hot- and cool-spotted stars separately, we see an even distribution of spot temperatures for all effective temperatures (i.e., for all effective temperatures with a considerable number of one type of spot; the apparent lack of hot spots for effective temperatures greater than 4000K may be due to a correspondence between mass and disk presence).

that these two distributions come from the same parent distribution. In this case, the K-S test proves with considerable certainty that the mass distribution of stars with hot spots is different from the mass distribution of stars with cool spots. At first glance, the astute reader may question the reliability of masses derived by the Hillenbrand (1997) study and the possible effects uncertain mass derivations might have on our observed distributions. After all, one might wonder how the presence of spots of various temperatures on the stellar surface affect the measured luminosities used to derive the stellar masses from a PMS H-R diagram. Hillenbrand is aware of the variability problem and notes that the uncertainties on the measured luminosities have little to no effect on the derived stellar masses for $M < 1.5M_{\odot}$. Given this awareness, we have great confidence in the accuracy of the stellar masses used in our distributions. The accuracy of the photospheric temperatures used to derive spot temperatures and sizes in our model is another issue. It is possible that the conversion from spectral type to stellar effective temperature may show some systematic effect for low-mass stars. In this scenario, we might be systematically underestimating the photosphere temperature, thereby affecting our derived spot temperatures. If we use a photosphere temperature that is less than what it should be, we might see a small hot spot in our model where there should either be a cool spot or no spot at all. In Figure 19 we see the distribution of spot temperatures with the stellar effective temperature. If a spectral type to effective temperature conversion problem did exist, we might expect to see an uneven distribution of spot temperatures for the hot-spotted stars and the cool-spotted stars. (Note that we do not want to compare hot-spotted stars to cool-spotted stars. Our only interest lay in the distribution of spot temperatures for a hot-spotted star or the distribution of spot temperatures for a cool-spotted star.) Because we see a similar range of spot temperatures for a given type of spot at every value of stellar effective temperature, we see no overwhelming influence of a spectral type to effective temperature conversion problem. We therefore interpret the different mass distributions to be the result of some physical process rather than some systematic error.

The fact that stars with signatures of active accretion (if we choose to interpret our hot spots as evidence for active accretion) are generally lower in mass than stars without signatures of active accretion is consistent with the theory that more massive stars are less likely to have disks. Because more massive stars reach the main sequence sooner than lower-mass stars forming at the same time, we would expect the more massive stars, if they form with a circumstellar disk, to dissipate their circumstellar disks faster than lower-mass

stars also forming with circumstellar disks. We expect the more massive stars to dissipate their disks faster because of the higher accretion rates typically seen for more massive stars. Our interpretation, therefore, favors the notion that we should expect the distribution we see in Figure 17 given that fewer higher mass stars are expected to have circumstellar disks from which we would see accretion and, by association, hot spots in our model. The true test of this interpretation will be to look at the distributions of disk presence against mass and accretion activity against mass. Should higher-mass stars be less likely to have circumstellar disks, we should see fewer higher-mass stars with evidence for a circumstellar disk through the $\triangle$(I-K) diagnostic as well as fewer higher-mass stars with evidence for active accretion through the measurement of Ca II equivalent widths.

4.2 Relationships between Spot Temperatures and $\triangle$(I-K)

Previous studies of the ONC have found a correlation between slow rotating stars and the presence of a circumstellar disk through measurements of near-IR excess (Herbst et al. 2002); Bouvier and collaborators (1993; 1995) have also found a correlation between slow rotating stars and those stars labeled CTTSs (stars believed to be accreting from their circumstellar disks) in the Taurus-Auriga molecular cloud. Still others have found no evidence for correlations between disk presence and rotation period in the ONC (SMMV; Rebull 2001). If we wish to provide a complete interpretation of our rotation period analysis in reference to other angular momentum studies, it is necessary to compare the spot temperatures derived from our model to evidence of circumstellar disks through the near-IR excess indicator, $\triangle$(I-K).

In Figure 20 we see the distribution of spot temperatures relative to the photosphere temperature as a function of $\triangle$(I-K) (measured by Hillenbrand et al. 1998). The logic of our argument up until now has associated spots hotter than the photosphere (positive spot temperatures in Figure 20) to accretion from a circumstellar disk. Thus, we should expect spot temperatures hotter than the photosphere to have values of $\triangle$(I-K) corresponding to the presence of a circumstellar disk. The canonical interpretation of the $\triangle$(I-K) measurement is that values of $\triangle$(I-K) greater than 0.3 correspond to a near-IR excess from a circumstellar disk (Hillenbrand et al. 1998; Herbst et al. 2002; SMMV). Upon inspection of Figure 20, the presence of stars with $\triangle$(I-K) < 0 is problematic. Our primary understanding of the results presented in Figure 20 is

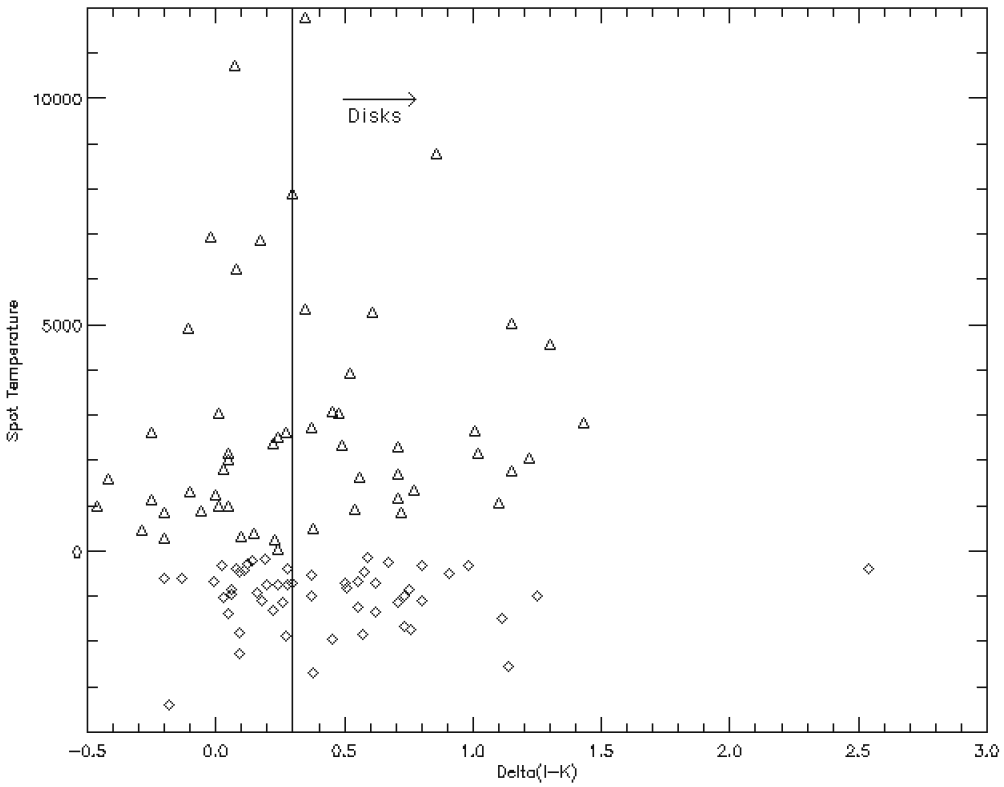

Figure 20: The distribution of spot temperatures values as a function of $\triangle$(I-K). Stars with hot spots have positive spot temperatures relative to the photosphere and are represented graphically by triangles while stars with cool spots have negative spot temperatures relative to the photosphere and are represented graphically by diamonds. Disk presence is indicated by values of $\triangle$(I-K) > 0.3. The large number of hot-spotted stars with $\triangle$(I-K) < 0 leads us to question the accuracy of the $\triangle$(I-K) measurements.

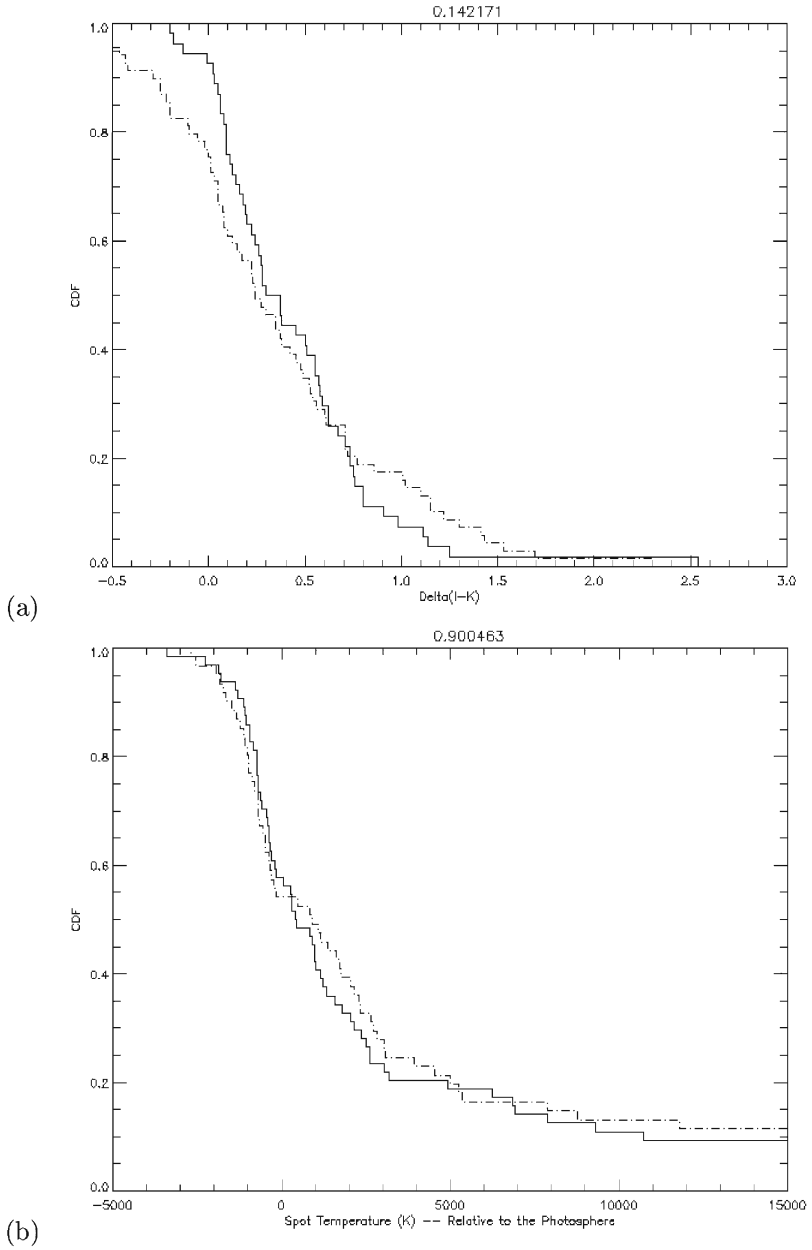

Figure 21: The CDF for the spot temperature and $\triangle$(I-K) distribution sliced two ways. (a) The CDF for cool-spotted stars as a function of $\triangle$(I-K) (solid line), and the CDF for hot-spotted stars as a function of$\triangle$(I-K) (dash-dot line). (b) The CDF for stars with $\triangle$(I-K) < 0.3, indicating no evidence of a circumstellar disk, as a function of spot temperature (solid line), and the CDF for stars with $\triangle$(I-K) > 0.3, indicating evidence of a circumstellar disk, as a function of spot temperature (dash-dot line). In both cases, (a) and (b), the K-S test finds no evidence for a difference between the distribution of disk-candidate stars and spot temperatures.

that the measurement of $\triangle$(I-K) is plagued by the stellar variability in the cluster. In order to calculate an infrared excess ($\triangle$(I-K)), one must measure the observed stellar radiation in the infrared bands I and K and then subtract from that measurement the expected contribution of stellar radiation from the star based on the star's spectral type; the result of this calculation is the excess radiation coming from the warm circumstellar disk. The values of $\triangle$(I-K) calculated by Hillenbrand et al. 1998 are given by

$$\triangle(I-K) = (I-K)_{observed} - 0.5A_v - (I-K)_{photosphere}$$

where $0.5A_v$ is the contribution of interstellar reddening (color excess toward longer wavelengths) calculated from the extinction values determined from the color ratio, V-I (i.e., the $0.5A_v$ term corrects for reddening due to interstellar extinction, and so we must subtract out this reddening, as well as the true color of the photosphere, if we are to calculate the amount of reddening due to the circumstellar disk). Within the equation used to calculate $\triangle$(I-K) we find two places in which stellar variability can affect the outcome of the $\triangle$(I-K) reported in Figure 20. The first measurement that might skew $\triangle$(I-K) is the measurement of the interstellar reddening. Given that the the interstellar reddening is calculated from the visual extinction, A_v, which is calculated from the color ratio V-I, one might question how reliable a calculation of interstellar reddening is for a star with hot spots on its surface, for a star with a hot spot on its surface will show an excess of flux in the V band. An excess flux in the V band will correspond to less extinction, which will produce larger values of $\triangle$(I-K). Another factor that can skew the calculated values of $\triangle$(I-K) is the determination of the photospheric contribution of (I-K). The photospheric contribution of (I-K) for each star is determined by the spectral energy distribution for a blackbody with temperature set by the photospheric temperature of the star. In our variability study, it is likely that photospheric temperatures used in our analysis are affected by the presence of hot and cool spots on the stellar surface. In this case, the expected contribution of I-K for the photosphere may be affected by misleading photosphere temperatures, affecting the calculated values of $\triangle$(I-K). (At this point in might be worthwhile to reassure the reader that inaccurate photosphere temperatures should have little effect on the spot temperatures relative to the photosphere. We would like to remind the reader that the rate of change of variability determines the type of spot, and that the rate of change is different for hot spots than it is for cool spots (see Section 3.1). Where inaccurate photosphere temperatures will affect the spot temperatures is in their absolute values (i.e., a hot spot will

still be a hot spot, but a misleading photosphere temperature may change a 3000K hot spot to a 2900K hot spot). Overall, misleading photosphere temperatures will have more effect on the measured values of $(\text{I-K})_{\text{photosphere}}$ but little influence on our model results.)

Given the uncertainty on $\triangle$(I-K), we can draw few conclusions from Figure 20. As stated before, we should expect different distributions of $\triangle$(I-K) for hot- and cool-spotted stars given our interpretation of the causes of hot spots and cool spots. Regardless of the probable errors on $\triangle$(I-K) , we compared the distribution of hot and cool spots as a function of $\triangle$(I-K) and the distribution of $\triangle$(I-K) as a function of spot temperature. The CDFs from the K-S test corresponding to these two comparisons are shown in Figures 21a and 21b, respectively. For the distribution of hot- and cool-spotted stars as a function of $\triangle$(I-K), the K-S test reveals a 14% probability that the two distributions come from the same parent distribution. While this probability is suggestive of the possibility that the hot-spotted stars and the cool-spotted stars have different probabilities for disk presence, we cannot say with certainty that the hot- and cool-spotted stars have distinctly different distributions for $\triangle$(I-K). The K-S test for the distribution of $\triangle$(I-K) as a function of spot temperature shows that spot temperature has little to do with the observed value of $\triangle$(I-K). The K-S test for $\triangle$(I-K) as a function of spot temperature was run on two distributions, one being the distribution of spot temperatures for $\triangle$(I-K) > 0.3 and one being the distribution of spot temperatures for $\triangle$(I-K) < 0.3. The K-S test comparing these two distributions reveals a 90% probability that the spot temperatures for stars with evidence for circumstellar disks come from the same parent distribution as the spot temperatures for stars without evidence for circumstellar disks.

Earlier, in Section 4.1.2, we looked at the distribution of spot temperatures as a function of stellar mass and found that hot-spotted stars appear to be lower in mass than cool-spotted stars. A possible cause of this result may be the tendency for more-massive stars to dissipate their circumstellar disks faster than less-massive stars. The test of this conjecture is presented in Figure 22, where we have plotted the observed values of $\triangle$(I-K) versus stellar mass. At first glance, one notices that stars with disks ($\triangle$(I-K) > 0.3) appear at all masses. The most interesting result of Figure 22 is the location of stars with $\triangle$(I-K) < 0. Almost all stars with physically impossible values of $\triangle$(I-K) appear to be low-mass stars. As noted in Figure 20, most values of $\triangle$(I-K) < 0 occur for hot-spotted stars, and so the presence of low-mass stars in the $\triangle$(I-K) < 0

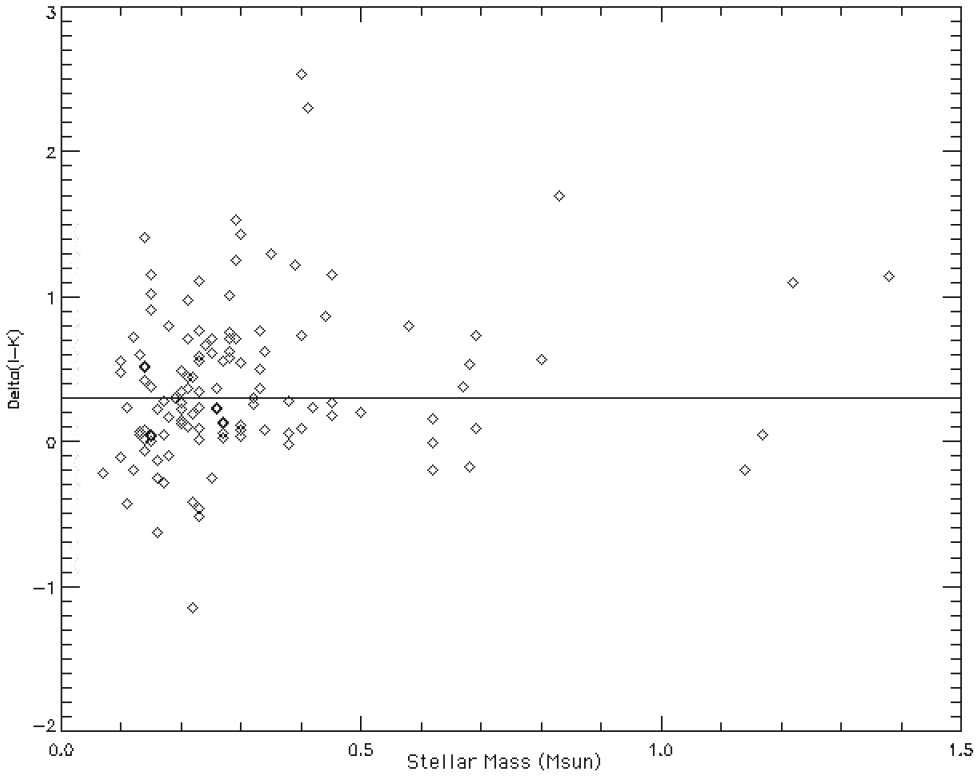

Figure 22: The distribution of $\triangle$(I-K) with stellar mass, with $\triangle$(I-K) > 0.3 interpreted as evidence for disk presence. From this figure we see that disks are seen for all masses. Almost all physically impossible values for $\triangle$(I-K), $\triangle$(I-K) < 0, occur at low masses, indicating that the $\triangle$(I-K) disk indicator may be inaccurate for low-mass stars.

zone is not unexpected. The results of our analysis of $\triangle$(I-K) versus spot temperature and stellar mass call into question the reliability of $\triangle$(I-K) as true measure of disk presence.

Given the uncertainty on the calculated values of $\triangle$(I-K) and the results of our K-S tests on the spot temperature and $\triangle$(I-K) distributions, we propose that the results of rotation period analysis comparing disk presence to rotation period be examined further. In particular, we propose that future studies re-examine the criteria for labeling a star as a disk candidate, perhaps using more than just $\triangle$(I-K) as a disk indicator.

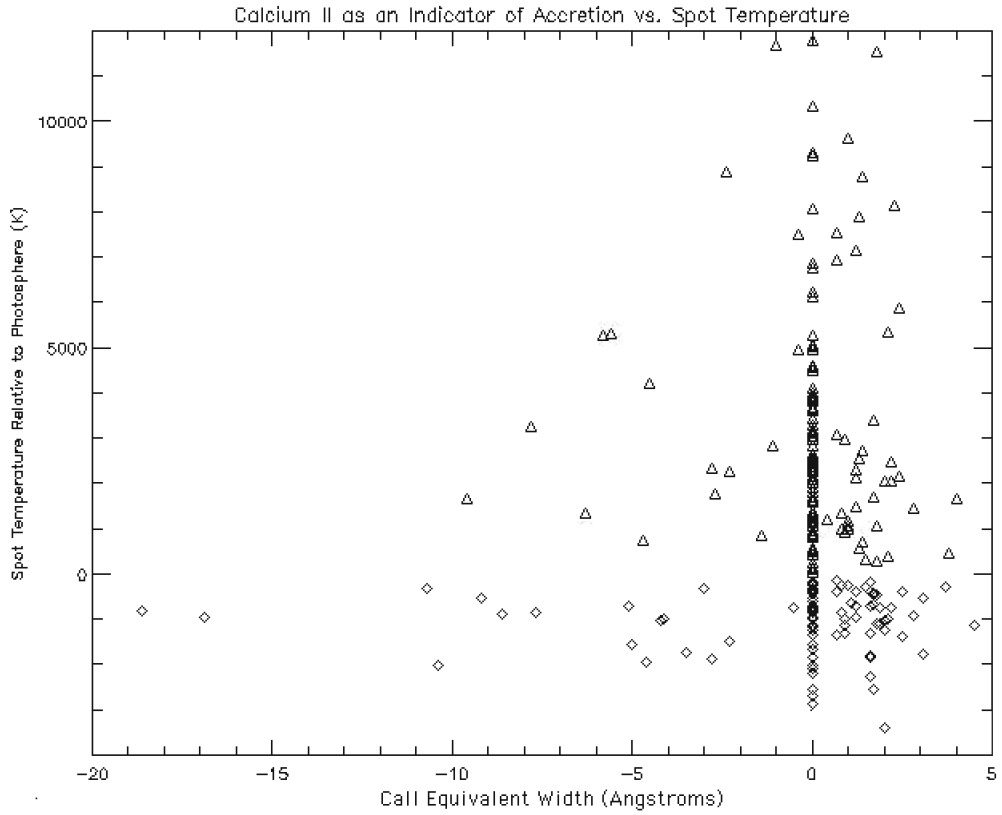

Figure 23: Spot temperatures relative to the photosphere (hot spots are positive and represented by triangles; cool spots are negative and represented by diamonds) as a function of Ca II equivalent width. Negative equivalent widths indicate emission of Ca II and are evidence for accretion while positive equivalent widths indicate absorption of Ca II and do not support heavy accretion from a circumstellar disk. An equivalent width of zero represents those cases where no emission or absorption feature is observed and may result from accretion-related veiling.

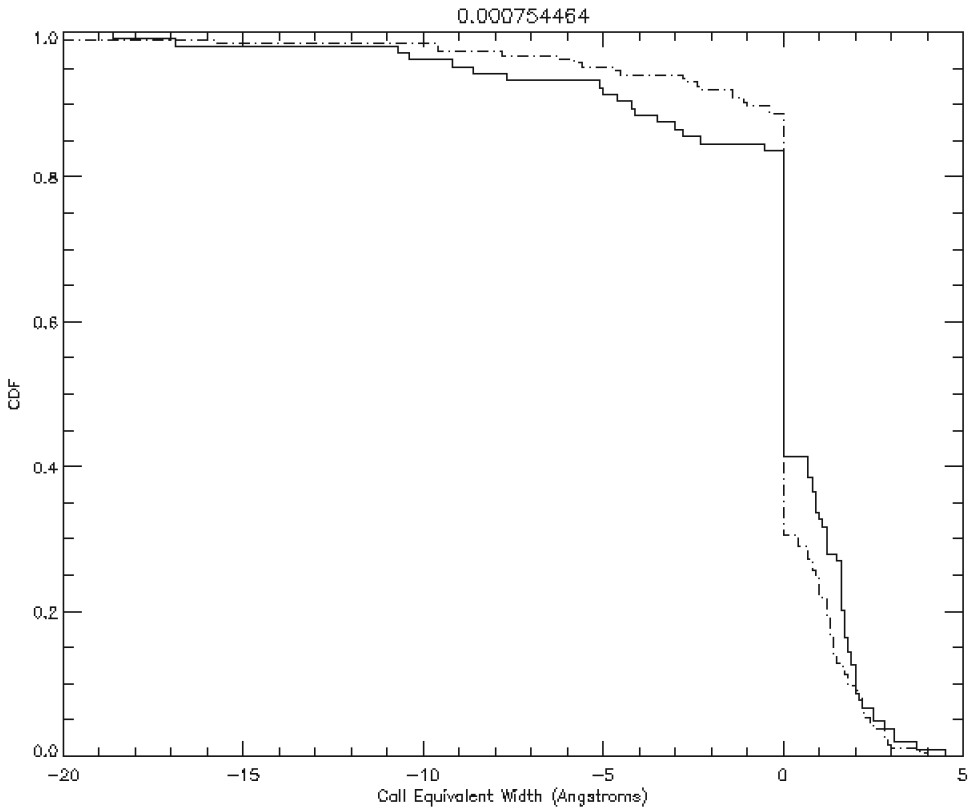

Figure 24: The CDF for the distribution in Figure 23 (cool-spotted stars are the solid line while hot-spotted stars are the dash-dot line). The K-S test on these distributions tells us that the hot- and cool-spotted stars most likely do not come from the same parent distribution. We should expect this result in the case that our hot photospheric spots derive from accretion activity, yet we caution that the strong dependence of the distributions on the zero Ca II equivalent width may affect our results. (We do note, however, that the zero Ca II equivalent widths are valid measurements and should not be removed from the distributions.)

4.3 Relationships between Spot Temperatures and Ca II Emission

Figure 23 shows the distributions of spot temperatures as a function of Ca II equivalent widths. Stars with evidence for accretion from a circumstellar disk have negative Ca II equivalent widths, corresponding to Ca II emission, whereas stars without evidence for accretion from a circumstellar disk have positive Ca II equivalent widths, corresponding to Ca II absorption.

As discussed in Section 1.2.2, Ca II emission arises from the impact of disk material onto the stellar surface. From the energy of impact, neutral calcium in the photosphere of the star is ionized; upon recombination we see photons contributing to Ca II emission in the stellar spectrum. The spectral feature resulting from accretion will be a broad emission feature. It is possible that we never visibly see an emission feature in the stellar spectrum even with accretion from a circumstellar disk. In some cases, we are likely to see "filled-in" spectral lines, i.e., spectral lines that show no Ca II emission but also show no Ca II absorption, even though at the very least we would expect absorption from the stellar spectrum. The "filled-in" spectral lines can be interpreted in one of three ways. The first interpretation holds that there is absorption from the stellar photosphere, but also emission from accretion that reduces the Ca II spectral line to the continuum level. A more likely interpretation, one that does not rely so heavily on a unique balance, is that veiling raises the continuum level such that a broad spectral feature, such as a Ca II emission feature, will likely be lost in the new continuum level. The reader may recall that veiling is intrinsically linked to ultraviolet continuum excess emission, and that ultraviolet continuum excess is directly related to accretion activity (Section 1.2.2). Thus, the absence of a Ca II spectral line (equivalent width equal to zero) may arise from accretion from a circumstellar disk. Another possible cause of zero Ca II equivalent widths is spectral line broadening caused by stellar rotation. As a star rotates on its axis, one half of the star will be moving away from us while another half of the star will be moving toward us; the result of this motion is a Doppler shift of the photons from the star toward shorter and longer wavelengths. In the case of Ca II, we may have photons from a rapidly rotating star that are redshifted and blueshifted to such extreme values that the spectral feature is smeared to the level of the continuum. Thus, in the case of rapid rotation, we may see a "filled-in" Ca II line. To investigate the influence of rotation on Ca II zero equivalent widths, we plotted Ca II equivalent widths as a function of rotation period in Figure 25. Although there are more "filled-in" Ca II lines at shorter rotation periods, there are a considerable number of "filled-in" Ca II lines at longer rotation periods. We conclude

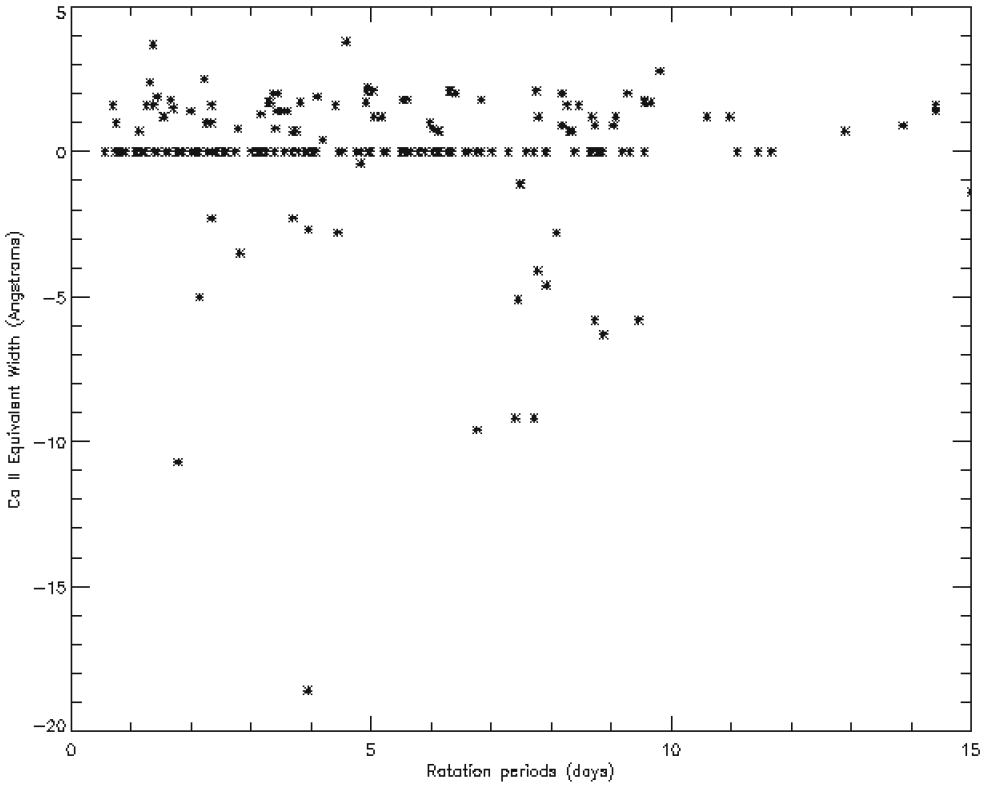

Figure 25: Calcium II equivalent widths as a function of rotation period. A considerable number of Ca II equivalent widths equal to zero fall at short rotation periods (fast rotators), but a large number also occur at long rotation periods (slow rotators), suggesting that rotationally-influenced line broadening may have a slight, although not overwhelming, effect on the number of Ca II equivalent widths measured to be zero.

from Figure 25 that rotation rate does not have a serious effect on the number of reported Ca II equivalent widths equal to zero. One may also notice in Figure 25 that stars with definitive accretion activity (Ca II equivalent width less than zero) appear at short and long rotation periods. (Note: In Figure 23 we look at all hot- and cool-spotted stars regardless of whether or not they have measured rotation periods. In Figure 25 we look only at hot- and cool-spotted stars that have rotation periods. A close examination between the two figures reveals that there are a few hot- and cool-spotted stars without rotation periods that appear in Figure 23 that do not appear in Figure 25.)

In Figure 23 we see a number of stars, both with hot spots and cool spots, with Ca II equivalent widths less than or equal to zero, indicating probable accretion processes. The K-S test on our spot temperature

distributions as a function of Ca II equivalent width reveals 0.08% probability that the hot-spotted star distribution and the cool-spotted star distribution come from the same parent distribution. It is clear that the CDFs for the hot- and cool-spotted star distributions are strongly dependent on the number of stars with Ca II equivalent widths equal to zero (see Figure 24). As it is likely that the zero Ca II equivalent widths strongly direct the result of the K-S test, we are hesitant to say that the hot-spotted star distribution is distinctly different from the cool-spotted star distribution. To forge a better understanding of the distribution of spot temperatures with Ca II equivalent widths, one could also involve a correlation with disk presence. Unfortunately, as noted in Section 4.2, disk presence as measured by $\triangle$(I-K) can be skewed by hot spots on the stellar surface, which may result in values of $\triangle$(I-K) less than 0.3 and even less than zero for cases where there is a circumstellar disk present. Because the physical process that causes a decrease in the measured value of $\triangle$(I-K) is the same process that raises the continuum level of the star and causes measured Ca II equivalent widths of zero, we checked to see what fraction of zero equivalent width Ca II measurements correspond to $\triangle$(I-K) measurements less than zero. For stars with hot spots, 30 of 54 stars (55%) with negative $\triangle$(I-K) measurements have zero equivalent width Ca II measurements. For stars with cool spots, the ratio is lower as only 7 of 45 stars (16%) with negative $\triangle$(I-K) measurements have zero equivalent width Ca II measurements. We believe this calculation strengthens our cautionary assessment of the reliability of $\triangle$(I-K) for determining disk presence, for the calculation suggests that the distribution of spot temperatures as a function of $\triangle$(I-K) is skewed by inaccurate measurements of $\triangle$(I-K), especially by stars showing evidence for accretion through hot spots on the stellar surface.

Looking back to the mass issue raised in Section 4.1.2, we want to know how mass is related to accretion activity. In Figure 17, we saw a tendency for low-mass stars to have hot spots on their surface. We then wondered if the mass distribution of hot- and cool-spotted stars was in any way related to the faster evolution of high mass stars and the likely chance that high-mass stars dissipate their circumstellar disks faster than low-mass stars. To check for any relationship between accretion and mass we plotted Ca II equivalent widths as a function of mass in Figure 26. When we look at the accretion activity through Ca II equivalent widths for low-mass and high-mass stars we find results suggestive of a relationship between mass and disk presence. For stars with $M < 0.25M_{\odot}$, we find that 96 of 138 stars (70%) with measured Ca II equivalent widths have an equivalent width less than or equal to zero (we have included the zero equivalent

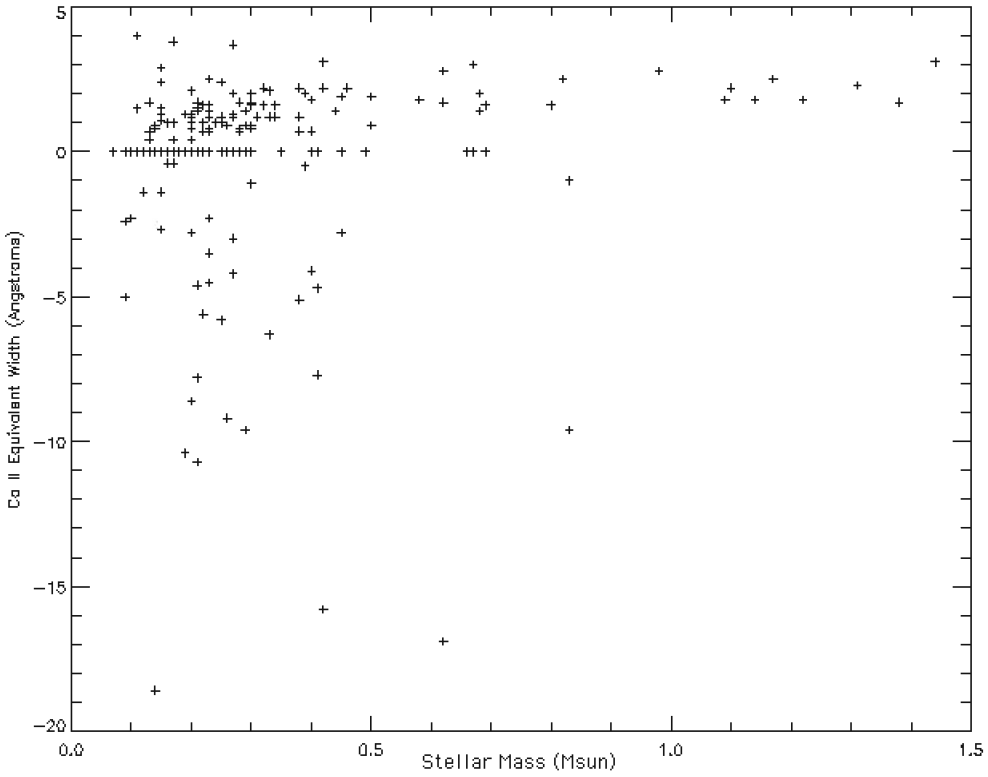

Figure 26: Calcium II equivalent widths as a function of stellar mass. Most stars showing some evidence for accretion (EW $\leq$0) appear at lower masses, reinforcing the belief that high-mass stars dissipate their circumstellar disks faster than low-mass stars.

width because of its connection to accretion activity), indicating that as many as 70% of low-mass stars are actively accreting. For stars with M $>$ 0.25M$_{\odot}$, we find that 35 of 92 stars (38%) with measured Ca II equivalent widths have an equivalent width less than or equal to zero, indicating that as many as 38% of high-mass stars are actively accreting.

Given our discussion of spot temperature distributions with $\triangle$(I-K) and Ca II equivalent widths, we have strong doubts about the results of rotation period analyses that invoke correlations between rotation period and disk presence. In our data, we find no evidence for a bimodal rotation period distribution. Instead, we find what appears to be a uniform distribution of rotation periods that does not depend on accretion activity or stellar mass. We also find independent correlations between spot temperature and mass, spot temperature

and accretion activity, and accretion activity and mass, suggesting that the hot spots we observe are the result of accretion activity that is prevalent among low-mass stars.

5 Conclusions

Our photospheric spot temperature analysis of young stars in the Trapezium region of the Orion Nebula Cluster finds little evidence for a correlation between rotation rate and accretion activity. The distribution of rotation periods for stars with hot spots on the stellar surface is statistically indistinguishable from the distribution of rotation periods for stars with cool spots on the stellar surface, regardless of whether or not the distributions are constrained to be above or below the $0.25M_{\odot}$ threshold quoted in numerous angular momentum studies (SMMV; Herbst et al. 2002;Rebull 2001). Because we interpret stars with hot spots in our model to be the result of some accretion process, our results do not support a model in which rotation rates are governed by disk locking. Instead, we find that the $\triangle$(I-K) disk indicator, used to distinguish stars with circumstellar disks and possible magnetic disk locking from stars without circumstellar disks and magnetic disk locking, is not a perfect marker for circumstellar disks. This might explain why studies using one disk criterion to study rotation rates of young stars find evidence for a bimodal distribution of rotation rates (Bouvier et al. 1993, 1995; Herbst et al. 2002) while studies using multiple criteria for disk detection find no bimodal distribution of rotation rates (Rebull 2001).

Our results may seem to complicate what appears to be an already confused debate about the angular momentum evolution of young stars. In truth, our results seem to fall more toward the recent opinion that magnetic disk locking, while present, is not solely responsible for governing the angular momentum evolution of young stars. This new idea has grown from observations contradicting the initial Attridge and Herbst bimodal period distribution (SMMV; Rebull 2001; Herbst et al. 2002) as well as new research investigating the effects of core-envelope decoupling on the angular momentum evolution of young stars (Wolff et al. 2004; Barnes 2003). The latter research argues that young stars lose angular momentum as they evolve along their convective tracks. When the stars transition from using convection as their primary method of energy transport to radiation, the stellar core becomes radiative while the outer regions of the star (the envelope)

remain convective. It is then possible that the envelope may be rotating at a different rate than the stellar core, in which case observations of stellar rotation rates via periodicity in stellar variability may not reflect the rotation and angular momentum of the envelope and not of the whole star (Wolff et al. 2004). Our research, which finds no preferred rotation rate for stars with evidence for accretion may reflect the fact that the rotation rate of the star is not wholly dependent on the star-disk interaction.

In closing, I would like to acknowledge my wonderfully supportive and patient advisers, Eric Jensen and Keivan Stassun. I also would like to thank David Cohen, Chris Burns, and Rachel Akeson for excellent conversations regarding this thesis. Funding for this thesis has been provided by grants from the Howard Hughes Medical Institute, NASA Delaware Valley Space Grant Consortium, Vanderbilt University, and Swarthmore College.

References

Attridge, J. M., & Herbst, W. 1992, ApJ, 398, L61

Balbus, S. A., & Hawley, J. F. 1998, RvMP, 70, 1

Barnes, S. A. 2003, ApJ, 586, 464

Basri, G., & Bertout, C. 1989, ApJ, 341, 340

Bevington, P. R. 1969, Data Reduction and Error Analysis for the Physical Sciences (New York: McGraw-Hill Book Company), 237

Bouvier, J., Bertout, C. & Bouchet, P. 1986, A&A, 158, 149

Bouvier, J., Cabrit, S., Fernandez, M., Martín, E. L., & Matthews, J. M. 1993, A&A, 272, 176

Bouvier, J., Covino, E., Kovo, O., Martín, E. L., Matthews, J. M., Terranegra, L., Beck, S. C. 1995, A&A, 299, 89

Carpenter, J. M., Hillenbrand, L. A., & Skrutskie, M. F. 2001, AJ, 121, 3160

Choi, P. I., Herbst, W. 1996, AJ, 111, 283

Flaccomio, E., Damiani, F., Micela, G., Sciortino, S., Harnden, F. R., Jr., Murray, S. S., Wolk, S.J. 2003, ApJ, 582, 398

Ghosh, P., & Lamb, F. K. 1979, ApJ, 232, 259 (1979a)

Ghosh, P., & Lamb, F. K. 1979, ApJ, 234, 296 (1979b)

Ghosh, P., Lamb, F. K., & Pethick, C. J. 1977, ApJ, 217, 578

Guenther, E. W., Lehmann, H., Emerson, J. P., Staude, J. 1999, A&A, 341, 768

Hartmann, L. 2002, AJ, 566, L29

Hartmann, L., & Kenyon, S. J. 1996, ARA&A, 34, 207

Hartmann, L, Kenyon, S., & Hartigan, P. 1993, in Protostars and Planets III, eds. E.H. Levy & J.I. Lunine (Tuscon: Univ. Arizona Press), 497

Herbst, W., Bailer-Jones, C. A. L., Mundt, R., Meisenheimer, K., & Wackermann, R. 2002, A&A, 396, 513

Herbst, W., Herbst, D. K., Grossman, E. J., & Weinstein, D. 1994, AJ, 108, 1906

Herbst, W., Rhode, K. L., Hillenbrand, L. A., & Curran, G. 2000, AJ, 119, 261

Hillenbrand, L. A. 1997, AJ, 113, 1733

Hillenbrand, L. A., Strom, S. E., Calvet, N., Merrill, K. M., Gatley, I., Makidon, R. B., Meyer, M. R., Skrutskie, M. F. 1998, AJ, 116, 1816

Honeycutt, R. K. 1992, PASP, 104, 435

Königl, A. 1991, ApJ, 370, L39

Ménard, F. & Bertout, C. 1999, in The Origin of Stars and Planetary Systems, eds. C.J. Lada & N.D. Kylafis (Dordrecht: Kluwer), 341

Meyer, M. R., Calvet, N., & Hillenbrand, L. A. 1997, AJ, 114, 288

Najita, J. R., & Shu, F. H. 1994, ApJ, 429, 808

Ostriker, E. C., & Shu, F. H. 1995, ApJ, 447, 813

Press, W. H., Flannery, B. P., Teukolsky, S. A., Vetterling, W. T. 1986, Numerical Recipes (Cambridge: Cambridge University Press), 472

Queloz, D., Allain, S., Mermilliod, J. C., Bouvier, J., Mayor, M. 1998, A&A, 335, 183

Rebull, L. M. 2001, AJ, 121, 1676

Shu, F., Najita, J., Ostriker, E., Wilken, F., Ruden, S., & Lizano, S. 1994, ApJ, 429, 781 (1994a)

Shu, F. H., Najita, J., Ruden, S. P., & Lizano, S. 1994, ApJ, 1994, 797 (1994b)

Stassun, K. G., Mathieu, R. D., Mazeh, T., & Vrba, F. J. 1999, AJ, 117, 2941 (SMMV)

Stassun, K., & Wood, K. 1999, ApJ, 510, 892

Stauffer, J. R., & Hartmann, L. W. 1987, ApJ, 318, 337

Torres, C. A. C., & Ferraz Mello, S. 1973, A&A, 27, 231

Vogt, S., 1981, ApJ, 250, 327

Vrba, F. J., Chugainov, P. F., Weaver, W. B., & Stauffer, J. S. 1993, AJ, 106, 1608

Vrba, F. J., Herbst, W., & Booth, J. F. 1988, AJ, 96, 1032

Vrba, F. J., Rydgren, A. E., Chugainov, P. F., Shakovskaya, N. I., & Weaver, W. B. 1989, AJ, 97, 483

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

Wolff, S. C., Strom, S. E., & Hillenbrand, L. A. 2004, ApJ, 601, 979