

A 4 Gyr M-dwarf Gyrochronone from CFHT/MegaPrime Monitoring of the Open Cluster M67

Ryan Dungee¹ , Jennifer van Saders² , Eric Gaidos³ , Mark Chun¹ , Rafael A. García⁴ , Eugene A. Magnier² , Savita Mathur^{5,6} , and Ângela R. G. Santos⁷ 

¹ Institute for Astronomy, University of Hawai'i, 40 North A'ohōkū Place, Hilo, HI, 96720, USA; rdungee@hawaii.edu

² Institute for Astronomy, University of Hawai'i, 2680 Woodlawn Drive, Honolulu, HI, 96822, USA

³ Department of Earth Sciences, University of Hawai'i at Mānoa, 1680 East-West Road, Honolulu, HI, 96822, USA

⁴ AIM, CEA, CNRS, Université Paris-Saclay, Université de Paris, Sorbonne Paris Cité, F-91191 Gif-sur-Yvette, France

⁵ Instituto de Astrofísica de Canarias (IAC), E-38205 La Laguna, Tenerife, Spain

⁶ Universidad de La Laguna (ULL), Departamento de Astrofísica, E-38206 La Laguna, Tenerife, Spain

⁷ Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Rua das Estrelas, PT4150-762 Porto, Portugal

Received 2022 February 22; revised 2022 August 10; accepted 2022 August 23; published 2022 October 19

Abstract

We present stellar rotation periods for late K- and early M-dwarf members of the 4 Gyr old open cluster M67 as calibrators for gyrochronology and tests of stellar spin-down models. Using Gaia EDR3 astrometry for cluster membership and Pan-STARRS (PS1) photometry for binary identification, we build this set of rotation periods from a campaign of monitoring M67 with the Canada–France–Hawaii Telescope's MegaPrime wide-field imager. We identify 1807 members of M67, of which 294 are candidate single members with significant rotation period detections. Moreover, we fit a polynomial to the period versus color-derived effective temperature sequence observed in our data. We find that the rotation of very cool dwarfs can be explained by simple solid-body spin-down between 2.7 and 4 Gyr. We compare this rotational sequence to the predictions of gyrochronological models and find that the best match is Skumanich-like spin-down, $P_{\text{rot}} \propto t^{0.62}$, applied to the sequence of Ruprecht 147. This suggests that, for spectral types K7–M0 with near-solar metallicity, once a star resumes spinning down, a simple Skumanich-like relation is sufficient to describe their rotation evolution, at least through the age of M67. Additionally, for stars in the range M1–M3, our data show that spin-down must have resumed prior to the age of M67, in conflict with the predictions of the latest spin-down models.

Unified Astronomy Thesaurus concepts: Open star clusters (1160); Stellar evolution (1599); Stellar ages (1581); Stellar rotation (1629)

Supporting material: figure set, machine-readable tables

1. Introduction

A critical piece of understanding the evolution of any system—be it stars, planets, or the Milky Way itself, is understanding both the order in which events occur and their timescales. To do this properly, one requires precise, reliable ages for the stars involved. M dwarfs are the most numerous stars in the galaxy (Gould et al. 1996; Bochanski et al. 2010) and have higher occurrence rates of small planets compared to higher-mass stars (Dressing & Charbonneau 2015; Hardegree-Ullman et al. 2019). They also do not fuse heavy elements, and many tens of Gyr must pass before they show perceptible signs of evolution on a Hertzsprung–Russell diagram (Laughlin et al. 1997). As a result, M dwarfs can serve as particularly excellent tracers of galactic chemical evolution.

However, M dwarfs are also resistant to most methods commonly used for measuring a star's age. Their evolution on the main sequence is undetectable (Laughlin et al. 1997), there are no observable asteroseismic oscillations (Chaplin et al. 2011; Berdinis et al. 2017; Mathur et al. 2019), and their deep convective envelopes burn Li within the first \sim 50 Myr (Bildsten et al. 1997). The age–peculiar velocity relation is only statistical, making it unreliable for individual stars, and it

breaks down for stars that are too young or too old (Aumer & Binney 2009; Lu et al. 2021). The age–metallicity relationship for the Milky Way has flattened out over the past few Gyrs (Holmberg et al. 2007).

Fortunately, rotation period–age relations, or *gyrochronology*, show promise for M-dwarf age-dating (Barnes 2003). Gyrochronology relies on the fact that a star spins down over time owing to the interaction of its magnetic field with stellar winds, causing a loss of angular momentum (Weber & Davis 1967; Skumanich 1972; Barnes 2007). Observations of Sun-like stars have shown that this angular momentum loss rate, $\frac{dJ}{dt}$, scales strongly with the angular rotation velocity (ω) of the star, $\frac{dJ}{dt} \propto \omega^3$ (Skumanich 1972; Kawaler 1988; Mamajek & Hillenbrand 2008; Meibom et al. 2009; Angus et al. 2015; Gallet & Bouvier 2015), and as a result a star's initial rotation period (P_{rot}) becomes less important with age (Epstein & Pinsonneault 2014; Gallet & Bouvier 2015). The availability of independent age-dating techniques for Sun-like stars such as asteroseismology and isochrone analysis has facilitated the calibration of this age–rotation relationship. This can be extended to M dwarfs but requires coeval populations of stars (i.e., clusters) of established ages in which the rotation periods of M dwarfs can be obtained.

The rotation period of a star as it hits the main sequence largely depends on two factors: the rotation period the star was born with (P_{init}), and the protostellar disk lifetime. Using



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

observations of the youngest pre-main-sequence clusters to fix the values of P_{init} indicates that slower rotators are likely to originate from longer disk lifetimes (Somers et al. 2017; Roquette et al. 2021). The diversity of P_{init} values and disk lifetimes leads to a spread of rotation periods (~ 0.2 –8 days) at the zero-age main sequence. Regardless of the rotation period a star has once it reaches the main sequence, the evolution is dominated by spin-down. As these stars forget their initial conditions, they converge onto the slow rotator sequence, a well-defined sequence in temperature–period space. Prior to this, the presence of any stars on a fast rotator sequence or those still heavily influenced by their P_{init} value will make any rotation–age relation ambiguous. Since accurate gyrochronology relies on a star’s initial rotation period becoming less important over time, it relies on convergence onto the slow rotator sequence.

One of the mechanisms by which these fast rotators delay their convergence is saturated spin-down. Saturated spin-down occurs for stars with rotation rates greater than a critical value (ω_{crit} , a function of mass; Epstein & Pinsonneault 2014), where spin-down scales as $\frac{dI}{dt} \propto \omega \cdot \omega_{\text{crit}}^2$ (Krishnamurthi et al. 1997). This is driven by a saturation of magnetic activity, which can be quantified using the Rossby number (Ro) defined as $(\omega \tau_{\text{cz}})^{-1}$, where τ_{cz} is the convective overturn timescale. Generally, smaller Rossby numbers indicate that a star is more magnetically active. However, below a certain value ($\text{Ro} \lesssim 0.1$) stars appear to reach a maximal amount of activity, where decreasing values of Ro no longer correspond to increases in magnetic activity indicators (Wright et al. 2011; Matt et al. 2015, and references therein). This indicates that particularly fast rotators undergo a decoupling of their rotation rate and their magnetic field strength, resulting in a weaker scaling of torque with rotation rate. Because τ_{cz} increases with mass, M dwarfs can remain in the saturated regime longer than their higher-mass counterparts, which can be seen in the high number of M-dwarf fast rotators in clusters such as Praesepe (670 ± 67 Myr; Douglas et al. 2017; Rebull et al. 2017) and the Hyades (728 ± 71 Myr; Douglas et al. 2019). Knowing when these fast rotators finally converge is critical for M-dwarf gyrochronology.

Observations of solar-mass stars younger than the Hyades ($\lesssim 600$ Myr) have shown that models that assume that the entire star rotates with a uniform angular velocity (i.e., solid-body rotation) fail to match the observed convergence onto a slow rotator sequence and subsequent evolution (Keppens et al. 1995; Krishnamurthi et al. 1997; Allain 1998). Models that incorporate the internal transport of angular momentum (i.e., differential rotation) relax the assumption of solid-body rotation. In particular, core–envelope decoupling models take a simplified approach of treating the core and envelope as two separate, rotationally solid bodies with a mechanism that transports angular momentum between the two (MacGregor & Brenner 1991; Denissenkov et al. 2010; Lanza & Spada 2015). The critical parameter, then, is the timescale over which torques act to equilibrate the rotation rates ($\tau_{\text{c-e}}$). Fits to cluster data have shown that $\tau_{\text{c-e}}$ is ~ 20 Myr for solar-mass stars and a strong function of mass ($\tau_{\text{c-e}} \propto M^{-7.28}$ or $M^{-9.1 \pm 1.8}$ by Lanza & Spada 2015; Somers & Pinsonneault 2016, respectively).

Recent observations of a collection of open clusters, namely, Praesepe (Douglas et al. 2017; Rebull et al. 2017), the Hyades (Douglas et al. 2019), NGC 6811 (Curtis et al. 2019), and NGC

752 (Agüeros et al. 2018), have shown that K and early M dwarfs appear to halt their spin-down for a period of time—a striking departure from a standard spin-down model with solid-body rotation, but a phenomenon that can be explained by core–envelope decoupling models (Spada & Lanza 2020). In Spada & Lanza (2020) the apparent stalling is caused by the angular momentum loss of the envelope being balanced by transport from the core. The result is a net loss of angular momentum from the star, while the envelope continues to rotate at a roughly constant rate. However, their model does not predict the same degree of stalling as observed in open clusters, as it predicts that stars later than K5 should be rotating ~ 5 days slower than they are in Ruprecht 147 (Curtis et al. 2020). We know that K and M dwarfs must resume spinning down, as field samples show K and M dwarfs with rotation periods that are many tens of days (McQuillan et al. 2013; Newton et al. 2016, 2017; Santos et al. 2019). Such rotation periods are roughly consistent with a Skumanich-type spin-down over the age of the galactic disk (van Saders et al. 2019). Knowing when these stars resume spinning down and the timescales over which internal angular momentum exchange occurs both directly affect the mapping of a rotation period to an age.

Calibrating gyrochronology and testing spin-down models for M dwarfs requires a larger sample of older, well-dated M dwarfs. Only a handful of such stars are currently available, the majority of which are in young clusters (700 Myr at the oldest; Douglas et al. 2017; Rebull et al. 2017) or are limited by the use of kinematic ages (Newton et al. 2016; Popinchalk et al. 2021). Previously, Barnes et al. (2016) used K2 to obtain calibrators for gyrochronology of solar-type stars in M67, but observing faint M dwarfs in crowded fields has proved impossible for missions such as K2 or TESS. The Canada–France–Hawaii Telescope’s (CFHT) MegaPrime (Boulade et al. 2003) instrument allows us to overcome the limitations of K2 and TESS in sensitivity, without significant loss in field of view. In this paper we present the rotation periods of late K- and early M-dwarf members of the 4 Gyr old cluster M67 (3.5–5.0 Gyr; Nissen et al. 1987; Demarque et al. 1992; Montgomery et al. 1993; Carraro & Chiosi 1994; Fan et al. 1996; VandenBerg & Stetson 2004; Balaguer-Núñez et al. 2007; Stello et al. 2016). We present the oldest K- and M-dwarf gyrochronone to date and compare it to literature gyrochronology relations.

2. Observations and Data Reduction

Our campaign used CFHT’s MegaPrime to monitor M67 (center coordinates: $\alpha = 8^{\text{h}}51^{\text{m}}18^{\text{s}}$, $\delta = +11^{\circ}48'00''$) from 2018 October 15 to 2021 March 5 (UT). MegaPrime is the MegaCam imager placed at the prime focus of CFHT; it has a 1 deg^2 field of view sampled by 40 CCDs arranged in four rows of 9, 11, 11, and 9 detectors each (Figure 1). In total we obtained 694 exposures of the cluster in discrete runs of 1–2 weeks, producing 131 epochs of data for our light curves. All data were collected with 121 s integration times using the Sloan i filter, red enough that M dwarfs are not too faint but blue enough to observe spot variability. Observations were taken five at a time in a cross-like dither pattern with $10.^{\prime\prime}4$ offsets. Bias subtracting, flat-fielding, fringe correction, and bad-pixel masking were all performed by version 3.0 of the CFHT *Elixir* pipeline (Magnier & Cuillandre 2004). For our data reduction we treated the five exposures in a dither pattern independently, only combining them when we averaged the

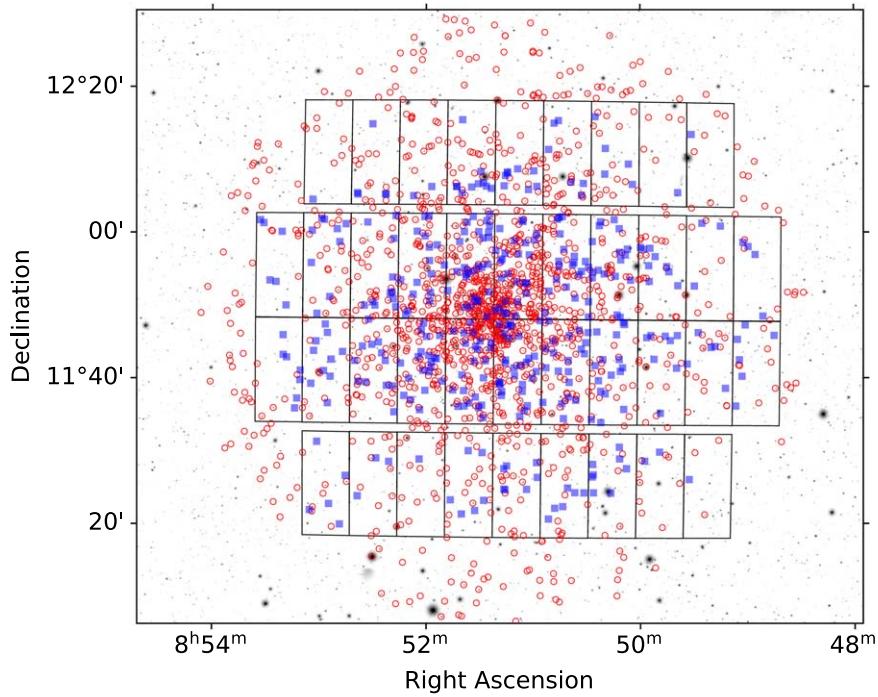


Figure 1. A Sloan Digital Sky Survey DR9 *i*-band image of the M67 field, with all of our candidate M67 cluster members identified by circles. Blue circles identify the stars with reported rotation periods in Table 2. The black boxes are the MegaPrime footprint for one of the pointings in our dither pattern.

five photometric measurements together in the later steps of the pipeline. An image of the field taken from *i*-band images of the Sloan Digital Sky Survey DR9 is shown in Figure 1. Included are the candidate cluster members of M67 (see Section 3 for further details).

2.1. Sky Background

We started with determining and removing the sky background from each detector’s image. For this we used `MMBackground`, a python implementation of the DAO-PHOT MMM algorithm contained in the `photutils` package (Stetson 1987; Bradley et al. 2020). We divided each image into an 8×10 grid of equally sized subregions, and for each subregion we estimated the background level through an estimation of the mode by the equation $\text{Mode} \approx 3 \times \text{Median} - 2 \times \text{Mean}$. This 8×10 grid was then interpolated to the size of the original image using a bi-cubic spline, and the resulting sky background was subtracted from the image. To estimate the uncertainty on this sky background, we also computed the sigma-clipped standard deviation of each subregion in the grid, which was similarly interpolated to produce an estimated uncertainty for each pixel in the sky background.

2.2. Source Finding

With the sky background subtracted, we then used `DAOStarFinder`, a python implementation of the DAO-FIND algorithm in the package `photutils`, to find the location of every source in the field for each individual image (Stetson 1987; Bradley et al. 2020). The threshold was set relatively low, at three times the sigma-clipped standard deviation of all pixel values in the image, and the FWHM was set at the seeing value reported in the image header. Using the World Coordinate System values in the image headers, we

converted the pixel coordinates reported by `DAOStarFinder` to R.A. and decl. (J2000.0). This enabled us to cross-match our detected sources with external catalogs.

We downloaded a catalog of every Gaia EDR3 (Gaia Collaboration et al. 2021) source in the field of view and converted the R.A. and decl. coordinates to the epoch J2000.0. For each Gaia source we then found the nearest neighbor match reported by `DAOStarFinder`. A handful of Gaia sources did not have a detection within the cutoff of $0.^{\circ}75$ and were considered nondetections. Any remaining sources found by `DAOStarFinder` that were not paired up with a Gaia source were considered false positives and discarded from our catalog. A nearest neighbor search, with the same distance cutoff, was also used to match every Gaia source to a Pan-STARRS1 (PS1) DR2 source. After this cross-matching, our catalog contained 8287 sources, all of which were matched to a source in both Gaia EDR3 and PS1 DR2. The limiting magnitudes of our observations ($i \sim 21$) compared to that of Gaia EDR3 ($G \sim 21$) and PS1 DR2 ($i \sim 22$) limit the number of real sources that were discarded by this method.

2.3. Photometry

Next, we performed aperture photometry on the background-subtracted images for every source in our catalog. The aperture diameter was set at four times the seeing value for the image, a value computed from the average empirical FWHM of bright sources scattered throughout the field. This diameter was chosen after analyzing the effect of aperture size on the noise properties of the light curves (see related discussion in Section 2.4). Any sources with overlapping apertures were flagged and excluded from the calculation of the zero-point corrections described in this section owing to the source confusion introduced by their overlap. Instrumental magnitudes were computed by a sum over the aperture divided by exposure time, and an uncertainty was estimated from the quadrature

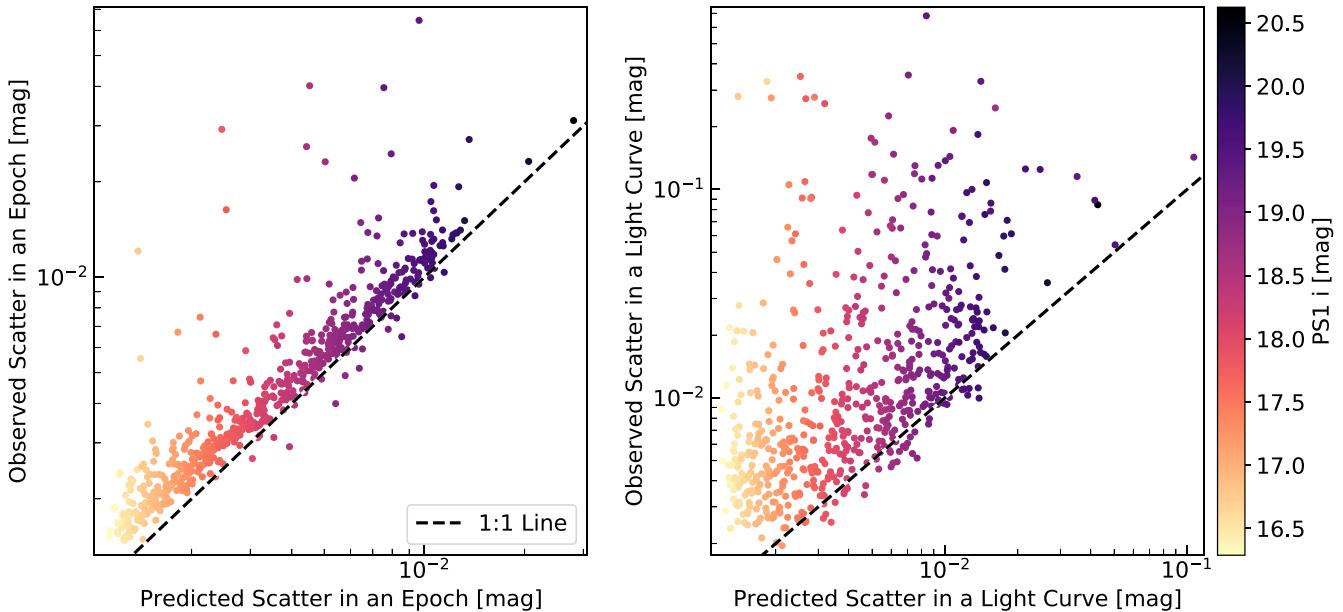


Figure 2. Left panel: a scatter plot of the predicted vs. observed scatter of the zero-point-corrected magnitudes that are averaged together to form one epoch in our light curves. The expectation is that these points lie on the one-to-one line. Right panel: a scatter plot of the predicted vs. observed scatter of the zero-point-corrected magnitudes in our light curves; the one-to-one line is the expected lower limit. In both panels the predicted values are derived from our estimated uncertainties.

sum of the photon noise on the flux in the aperture, the read noise of the MegaPrime detectors, and the previously estimated sky uncertainty for each pixel in the aperture. With this process repeated for each of our observations, we then began constructing the light curves for each target.

In order to construct the light curves, we first averaged together the measurements within an epoch (i.e., the set of five exposures that make up one dither pattern). To do this, we corrected for small changes in the photometric zero-point that may have occurred between exposures. The correction, which we call Δ_i , was taken to be the median difference between the stars with low scatter in their instrumental magnitudes. In equation form, the correction applied to the i th observation relative to the first ($i = 0$) is given by

$$\Delta_i = \text{Median}(M_{ls,0} - M_{ls,i}), \quad (1)$$

where $M_{ls,i}$ represents the set of stars with less than median scatter in their n measurements. Since any measurement where the aperture included bad pixels is discarded, $n \leq 5$. The epoch magnitude was then computed from the average of the zero-point-corrected measurements. If m_i is the instrumental magnitude from the i th observation within the j th epoch, then \bar{m}_j , which will become a point in a light curve, is given by

$$\bar{m}_j = \frac{1}{n} \sum_{i=0}^{n-1} (m_i + \Delta_i). \quad (2)$$

We repeated this calculation for every star and every epoch.

Finally, we corrected for the changes in the photometric zero-point between epochs. We took the same approach as before, computing this correction from the low scatter stars. The zero-point correction, z_{p_j} , of the j th epoch relative to the first ($j = 0$) is given by

$$z_{p_j} = \text{Median}(\bar{M}_{ls,0} - \bar{M}_{ls,j}), \quad (3)$$

where $\bar{M}_{ls,j}$ represents the set of stars with less than median scatter in their computed \bar{m} values. Each \bar{m} was then corrected

by this zero-point correction:

$$m_{\text{epoch},j} = \bar{m}_j - z_{p_j}. \quad (4)$$

In total, we obtained 4396 light curves that met our completeness criterion of at least 99 epochs of available data (see Section 4 for further details).

2.4. Validation

In order to validate our model of the photometric noise, we performed a comparison of the observed scatter in the zero-point-corrected magnitudes to the scatter expected from the estimated uncertainties alone. We made two important assumptions for these tests: (1) the uncertainties were the standard deviation of independent Gaussian distributions, and (2) all of these distributions had the same mean (i.e., the uncertain measurement was the only source of variability). Thus, any sources with additional variability in their magnitudes would fall above the one-to-one line on a plot of the theoretical scatter versus the observed scatter.

First, we performed this test on the $n \leq 5$ measurements that form an epoch, comparing the standard deviation of these points to the scatter expected from the uncertainty on their average. Exposures in an epoch were collected over a period of roughly 20 minutes, short enough that we expected each star not to vary. As a result, a scatter plot of how the noise was modeled versus the observed scatter should follow a one-to-one line, as we see in our data (left panel of Figure 2). For sources with magnitudes of $i \lesssim 18$, the data show a departure from the one-to-one line that we attribute to either nonlinearity in the detector as it approaches saturation or a fractional measurement error, such as flat-fielding errors.

Second, we performed this test on full light curves, comparing the standard deviation of the epoch magnitudes to the scatter expected from their uncertainties. The one-to-one line is expected to be the lower limit of photometric scatter; therefore, many sources will show variability beyond the case of random variations due to uncertain measurements of the

magnitude. We stress that a source falling above the one-to-one line in the right panel of Figure 2, meaning that it has more variability than expected from uncertainty in the photometry alone, is not proof of the source having an astrophysical process driving that variability. That the one-to-one line is indeed the lower limit in the right panel of Figure 2 demonstrates that our model of the noise is correct.

3. Cluster Membership and Stellar Properties

An important aspect of the results presented in this paper is that the rotation periods reported can be used as benchmarks for stellar spin-down models. Critical to this is knowing the age and T_{eff} . The age determination comes from their membership in the open cluster M67, whose age has been previously determined to be 4 Gyr (3.5–5 Gyr; Nissen et al. 1987; Demarque et al. 1992; Montgomery et al. 1993; Carraro & Chiosi 1994; Fan et al. 1996; Vandenberg & Stetson 2004; Balaguer-Núñez et al. 2007; Stello et al. 2016). The effective temperatures are derived using a color– T_{eff} relation. In this section, we provide the details on both of these critical aspects.

3.1. Cluster Membership

Cluster membership was determined by using a clustering algorithm, HDBSCAN (McInnes et al. 2017), on Gaia EDR3 proper motions and parallaxes for every star in the MegaPrime field of view, regardless of whether or not it appeared in our catalog of light curves. HDBSCAN works by using the density of points to estimate a probability density function (pdf) that describes the full distribution of values in the data. Clusters are then defined by the peaks in this pdf. The primary advantage of HDBSCAN is that it relies on fewer assumptions about the data than more traditional clustering algorithms such as K-means, which assumes Gaussian distributions. It also does not require that every point in the data set be assigned to a cluster, reducing the risk that outlier field stars might incorrectly be assigned M67 membership. The `python` package of the same name provides many different parameters to tune the performance of the clustering.⁸ We found that the defaults for the version we used, v0.8.18, were acceptable with one exception: `min_samples`. This parameter can be thought of as determining the level of detail in HDBSCAN’s estimation of the underlying pdf. Too small a value and each data point produces its own peak in the estimated pdf; too large, and the finer details of the estimated pdf are washed out. Given that the distribution of parallaxes and proper motions is effectively a two-peaked distribution (Figure 3), we found that a value of 200 gave suitable clustering results compared to the default of 5.

A disadvantage of the HDBSCAN algorithm is that it does not make use of the uncertainties on any input data. To incorporate these into our cluster membership determination, we performed Monte Carlo sampling. We ran the clustering algorithm on our list of Gaia sources, recorded the results, and resampled every star assuming Gaussian uncertainties, repeating this process 1000 times. HDBSCAN cannot be instructed to find a cluster with specific properties, so with each realization we computed the median parallax and proper motions for each grouping it found in the data and used the one with the closest match to literature values for M67 ($\pi = 1.1327 \pm 0.0018$ mas, $\mu_{\alpha} \cos \delta = -10.9738 \pm 0.0078$, $\mu_{\delta} = -2.9465 \pm 0.0074$ mas yr⁻¹; Gao 2018).

The difference was never more than a few percent. We then calculated a “kinematic membership probability” from the fraction of realizations in which a star was assigned membership to M67. A final membership criterion of $>50\%$ was selected based on an inspection of the Gaia color–magnitude diagrams (CMDs; Figure 4) that were produced for various thresholds.

Another aspect we considered in our use of HDBSCAN was the Bayesian nature of this approach. Too small a field of view and the algorithm may not have had the leverage needed to separate M67 members from the field; too large, and the diversity of field stars may have encouraged labeling true members as field stars. To address this, we repeated our membership determination on a Gaia EDR3 catalog including stars out to twice the radius of the MegaPrime field of view and compared the two membership lists. Of the 1807 members within the MegaPrime field of view, 76 were considered field stars when using the larger catalog, and none of the field stars gained membership in M67. None of these 76 stars are outliers on our CMDs (Figure 4), so we have kept them in our final list of members. However, we flagged them as potentially suspect, so that the interested reader may remove them from the sample if they wish.

We compared our list of M67 members to that of Gao (2018), who applied a Gaussian-mixture-model-based approach to Gaia DR2 astrometry. They found a list of 1502 likely members, whereas we have found 1807. In common between the two catalogs are 1241 members, leaving 261 stars that are unique to the Gao (2018) catalog and 566 that are unique to our catalog. There are several factors that contributed to these differences. First, our search for members was limited to the field of view of CFHT MegaPrime, and this truncated our search at a radius of $\sim 30'$; all 261 stars that are only in the Gao (2018) catalog were outside our field of view and thus were not included in our clustering. Second, 431 of the 566 stars that appear only in our catalog have parallaxes and/or proper motions that are closer to the literature values for M67 in EDR3 than in DR2. Third, 99 of the 566 stars that appear in our catalog are new in EDR3 and thus could not have been included in the Gao (2018) catalog. Finally, there are 36 stars that are unique to our catalog for otherwise unknown reasons; we attribute these to the differences between the two methods used. The members that our two catalogs have in common are denoted by the “Gao Member” column in Table 2.

3.1.1. Single versus Binary Members

Unresolved binaries bias our inferred stellar parameters, and close binaries have spin-down influenced by tidal forces (Simonian et al. 2020); as a result, we also need to identify whether or not the M67 members have a companion. With a parallax of $\pi = 1.1327 \pm 0.0018$ mas, Gaia is able to resolve binaries that are separated by $\gtrsim 600$ au. However, due to the size of our apertures, stars with physical separations $\lesssim 7000$ au have overlapping apertures. As such, their photometry was potentially limited by confusion with their nearest neighbor. For completeness, we included these stars in our catalog of reported rotation periods (Section 5), but we excluded them from our subsequent analysis. Work done by Deacon & Kraus (2020) indicates there are no wide binaries separated by $\gtrsim 3000$ au in clusters; thus, our analysis was focused only on single members of M67. Binary systems that are not resolved require a different method of detection. Common approaches include (1) spectroscopy that resolves double-lined absorption features, (2) excess astrometric

⁸ Details on the parameters can be found under Parameter Selection for HDBSCAN* in the docs.

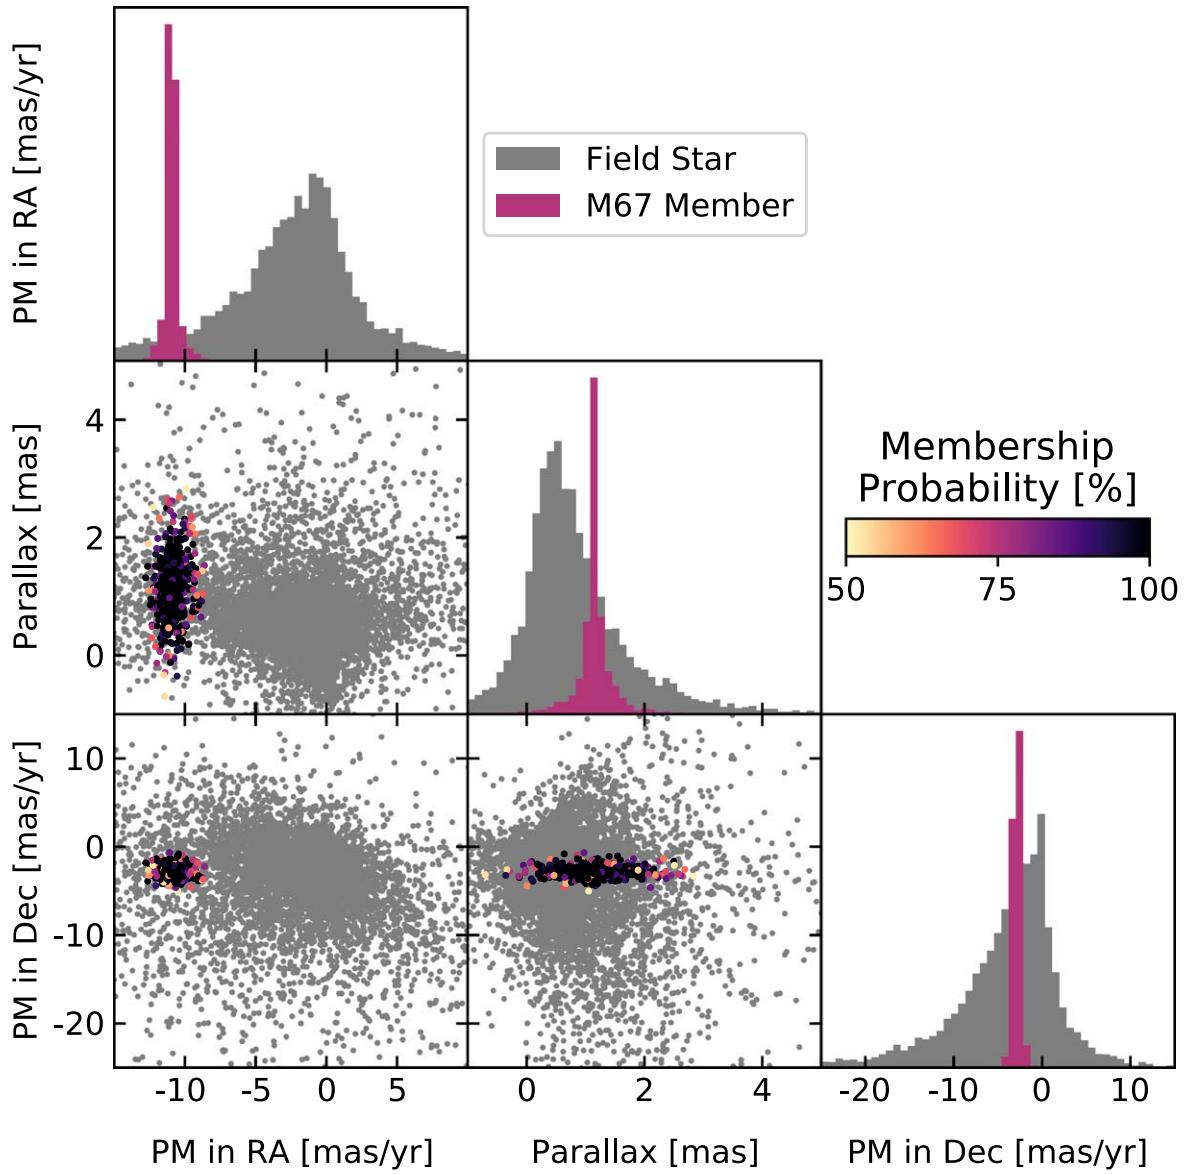


Figure 3. Scatter plots and histograms of the three parameters used for cluster membership determination (proper motions, or PM, in R.A. and decl., and parallaxes). In gray, the distribution of these parameters for the noncluster stars is shown as a point of comparison. For cluster members color indicates the probability of membership, showing that the members on the “outskirts” of the distribution are less likely to be considered members of M67 (see Section 3.1 for further details).

noise (quantified by the renormalized unit weight error, or RUWE, for Gaia astrometric solutions; Belokurov et al. 2020), and (3) photometric excess, stars that appear brighter than the main sequence on a CMD. For our data we used the photometric excess approach, calculated from PS1 photometry (Magnier et al. 2020; Flewelling et al. 2020). The effectiveness of this approach was confirmed by the finding that all the sources that exhibited excess astrometric noise (i.e., $\text{RUWE} \gtrsim 1.4$) were also found to show photometric excess.

In PS1 DR2 the saturation limit is 12–14 mag, depending on seeing and filter; for M67 we found that a cut at $i = 13.75$ removed these problematic sources. From the photometry in Gaia EDR3 and PS1 DR2 we found that PS1 r and i were the two filters with the highest signal-to-noise ratios for the faintest members in our catalog. Therefore, we used 3σ iterative outlier rejection to fit an eighth-order polynomial to the main sequence of the cluster on a PS1 $r - i$ versus i CMD. We then categorized each source by its vertical distance from the main sequence on the

CMD based on the distribution of residuals after subtracting out our fit to the main sequence (Figure 5). Stars within ± 0.3 mag of the main-sequence fit were classified as single members, whereas sources outside these bounds were categorized as photometric binaries. The small secondary peak of binary members is broad, which suggests that there may be a small number of high-contrast binaries contaminating our sample of candidate single M67 members. Sources fainter than the main sequence are thought to be binaries with a white dwarf component. Further observations are required to confirm this, which will be included in future work. The fraction of sources labeled as binaries by this method is 26%, in line with the expectation for M-dwarf multiplicity rates (Duchêne & Kraus 2013; Winters et al. 2019). Both the PS1 $r - i$ versus i CMD that was used for this binary classification and an additional Gaia $G - RP$ versus G CMD can be seen in Figure 4. The stars identified as binaries by this method were set aside for future analysis. They do not have rotation periods reported in this paper.

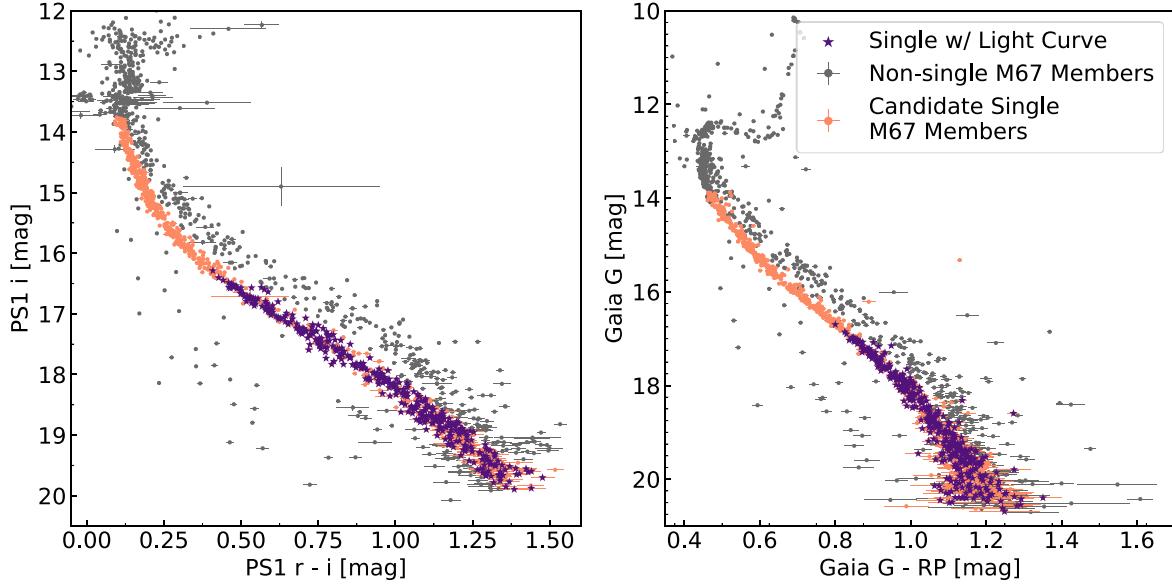


Figure 4. The CMD of kinematically selected M67 members. The left panel is using the Pan-STARRS DR2 photometry, and the right panel is using Gaia EDR3 photometry. In both panels orange points represent main-sequence single members (subject to a brightness cutoff at lower magnitudes), gray points represent photometric binaries, and purple points indicate stars with rotation periods reported in our results.

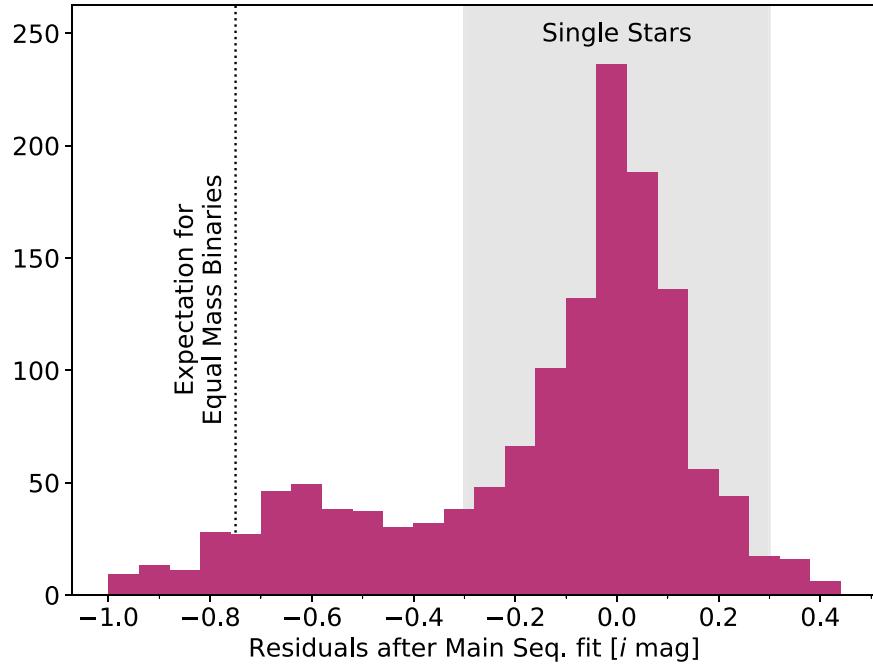


Figure 5. The distribution of residuals after subtracting out a polynomial fit to the main sequence from each star in the catalog. The shaded region denotes the stars chosen as single stars on the main sequence. The -0.3 offset represents the approximate location of the valley between the two peaks in this distribution, while accounting for the wide spread about the binary peak centered at ~ -0.6 . The vertical dotted line denotes the expected excess brightness for equal-mass binaries.

3.2. Effective Temperatures

In order to calculate the effective temperature (T_{eff}) for the stars in our catalog, we used an $(r - i)$ versus T_{eff} relation derived from the sample of late K and M dwarfs analyzed by Mann et al. (2015). We converted the synthetic Sloan r and i photometry provided into the PS1 r and i passbands using the Tonry et al. (2012) relations and applied corrections for reddening ($E(B - V) = 0.041 \pm 0.004$ mag; Taylor 2007), as well as a conversion to apparent magnitudes for the distance to M67 ($\pi = 1.1327 \pm 0.0018$ mas; Gao 2018). We trimmed the sample to stars with metallicities of $-0.07 \leq [\text{Fe}/\text{H}] \leq 0.07$, a range chosen to cover various values reported for the

metallicity of M67 in the literature (Pace et al. 2008; Santos et al. 2009; Önehag et al. 2011; Liu et al. 2016; Sandquist et al. 2018). The parameters of the trimmed sample are included in Table 1. Finally, we fit a second-order polynomial to the $(r - i) - T_{\text{eff}}$ pairs to obtain our relation:

$$T_{\text{eff}} = 139.8(r - i)^2 - 1062.1(r - i) + 4502.1. \quad (5)$$

The residual dispersion of 46 K is small compared to the errors on the temperatures (Figure 6). Adding this in quadrature with the spectroscopic errors provided by Mann et al. (2015) yields a T_{eff} uncertainty of ~ 75 K.

Table 1
M Dwarfs Used to Derive $T_{\text{eff}}(r - i)$

| Name | Table Header | Units | Description |
|---------------------------------|--------------|-------------|--|
| Object name | name | ... | Source name used in Mann et al. (2015) |
| R.A. | raDeg | deg | ... |
| decl. | deDeg | deg | ... |
| Gaia G magnitude | gaiaGmag | mag | The synthetic Gaia G magnitude |
| Gaia BP magnitude | gaiaBPMag | mag | The synthetic PS1 Gaia BP magnitude |
| Gaia RP magnitude | gaiaRPMag | mag | The synthetic PS1 Gaia RP magnitude |
| Mass | solMass | M_{\odot} | Mass of the star |
| Error on the mass | e_solMass | M_{\odot} | ... |
| Metallicity | [Fe/H] | ... | Metallicity of the star |
| Error on the metallicity | e_[Fe/H] | ... | ... |
| T_{eff} | teff | K | Effective temperature |
| PS1 <i>g</i> magnitude | ps1gmag | mag | The synthetic PS1 <i>g</i> magnitude |
| Error on PS1 <i>g</i> magnitude | e_ps1gmag | mag | ... |
| PS1 <i>r</i> magnitude | ps1rmag | mag | The synthetic PS1 <i>r</i> magnitude |
| Error on PS1 <i>r</i> magnitude | e_ps1rmag | mag | ... |
| PS1 <i>i</i> magnitude | ps1imag | mag | The synthetic PS1 <i>i</i> magnitude |
| Error on PS1 <i>i</i> magnitude | e_ps1imag | mag | ... |
| PS1 <i>z</i> magnitude | ps1zmag | mag | The synthetic PS1 <i>z</i> magnitude |
| Error on PS1 <i>z</i> magnitude | e_ps1zmag | mag | ... |
| K_s magnitude | Ksmag | mag | The synthetic K_s magnitude |
| Error on K_s magnitude | e_Ksmag | mag | ... |

Note. A description of each column in the table of stars used to derive our color– T_{eff} relationship.

(This table is available in its entirety in machine-readable form.)

4. Measuring Rotation Periods

In our data set there are 7222 sources that contain at least one epoch of data. To reduce complications with recovering periodic signals, we applied a conservative cut to our data, requiring that a light curve have a minimum completion of 99 out of the possible 131 epochs of data. After applying this cut, we were left with a sample of 4674 stars, with a mean completeness of 129 epochs. Of these 4674 stars, 3607 have light curves with 131 epochs. For the 636 candidate cluster members that made these cuts the mean completeness is 128 epochs, with 444 having a light curve that has 131 epochs. Due to the irregular sampling of our light curves, we used Lomb–Scargle (LS) periodograms (Lomb 1976; Scargle 1982; Press & Rybicki 1989; Zechmeister & Kürster 2009) for the detection of periodic signals in our light curves. In each case the rotation period we report was the period of maximum power in the periodogram.

A common method for quantifying the uncertainty of LS periodograms is the false-alarm probability (FAP). The FAP is a measure of probability that data with no signal would produce a peak in the periodogram of equivalent height (for further details see Section 7.4.2 of VanderPlas 2018). We required an FAP value of less than 1% for a periodic signal to be considered significant. To maintain the computational feasibility of our injection and recovery tests (see Section 4.1), we report FAP values estimated using the Baluev (2008) method. As a test of the validity of using the Baluev estimates, we performed a comparison of the FAP values estimated by the Baluev method to those computed using a bootstrapping ($N = 10^4$) of all of our light curves. Using the bootstrapping method, it is roughly expected to find $\text{FAP} \times N \pm \sqrt{\text{FAP} \times N}$ false positives (VanderPlas 2018). Accounting for this uncertainty on the FAP value, every one of our Baluev-estimated FAP values is consistent with its bootstrapped equivalent. We also required that a rotation period have at least

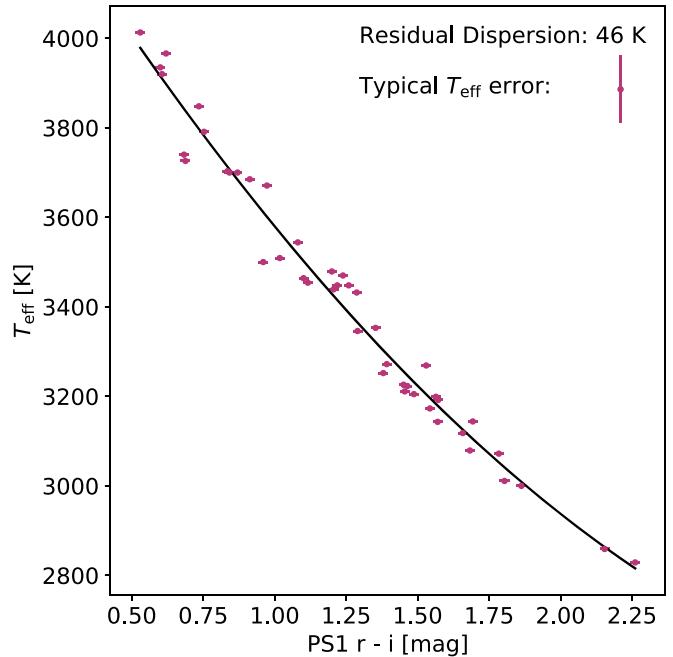


Figure 6. The sample of late K and M dwarfs of Mann et al. (2015) used to derive our color– T_{eff} relationship. We chose stars to have a metallicity within the range of literature values for M67. The T_{eff} errors are of order 75 K and do not affect the dispersion.

five complete periods within the light-curve duration in order to be considered a detection. This placed an upper limit of 175 days on any rotation periods used in our analysis.

4.1. Injection and Recovery Tests

We performed injection and recovery tests to determine the detection efficiency and false-positive rates for our recovered periods. The Kepler long-cadence data provided a database of

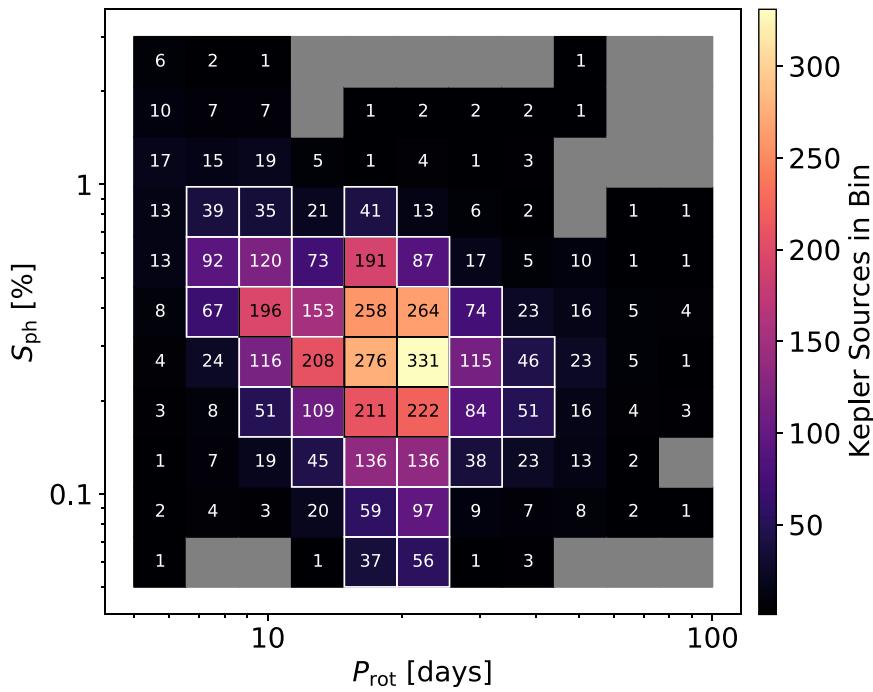


Figure 7. The number of Kepler light curves used in the injection and recovery tests to establish completeness as a function of rotation period (P_{rot}) and photometric variability (S_{ph}). Bins with a black or white (color chosen for optimal contrast) box drawn around them are the ones where over half (i.e., ≥ 25) of the injected signals are Kepler light curves.

real astrophysical signals of rotation for us to test our rotation recovery. For the injected signals, we used the KEPSEISMIC light curves of K- and M-dwarf main-sequence stars obtained with the Kepler Asteroseismic Data Analysis and Calibration Software (KADACS; García et al. 2011, 2014; Pires et al. 2015). The rotation periods for these stars have been derived by Santos et al. (2019). We applied a few additional cuts of our own: First, we checked for the completeness of the Kepler light curve, rejecting any star with fewer than 11 continuous quarters of data. Second, we applied a cut on the height of the autocorrelation function peak (H_{ACF} in Santos et al. 2019), the average difference between the peak height and the two adjacent local minima) requiring that $H_{\text{ACF}} \geq 1.0$, a value typical of stable signals. Finally, we applied a cut on the effective temperature of $T_{\text{eff}} < 5270$ K, so that the observed spot pattern evolution in the Kepler sample would more closely match the expectation for our targets in M67. Given the precision of the Kepler photometry relative to our data, we made the assumption that these light curves contained noiseless rotation signals.

This gave us a sample of 4599 signals with known rotation periods for injection. Each signal was characterized by two values: the rotation period (P_{rot}) and the photometric activity index (S_{ph}), a measure of the amplitude of variability. The value of S_{ph} was calculated by dividing a light curve into subseries, each five times the length of the star's rotation period, and then taking the mean of the standard deviations of each of the subseries (Mathur et al. 2014). One of the advantages of S_{ph} over other measures of photometric variability is its correlation with proxies of magnetic activity (Salabert et al. 2016, 2017). We created logarithmically spaced bins for the injections: $5 \leq P_{\text{rot}} \leq 100$ [days] and $0.05 \leq S_{\text{ph}} \leq 3$ [%Flux] using 11 bins along each axis. However, this left some of the outlier bins (see Figure 7) with very few, if any, injections. To compensate for this deficiency, we

also generated a set of synthetic light curves. These synthetic light curves were simple sinusoids:

$$\text{Flux}(t) = \sqrt{2} S_{\text{ph}} \sin\left(\frac{2\pi}{P_{\text{rot}}} t + \phi\right), \quad (6)$$

where ϕ is a uniformly distributed phase, and the factor of $\sqrt{2}$ comes from the fact that S_{ph} is calculated from a standard deviation. The synthetic light curves were sampled with the same cadence as the Kepler data. For every bin with less than 50 Kepler light curves we generated a sample of up to 50 synthetic ones with P_{rot} and S_{ph} values uniformly distributed (in linear space) within the bounds of that bin. In total we used 3934 synthetic light curves.

Each injection and recovery test involved taking a Kepler (or synthetic) light curve and sampling it to match the cadence and length of our CFHT observations. We then added the signal into one of our CFHT light curves and computed an LS periodogram for the combined data. If the period of maximum power in the resulting LS periodogram was within 10% of the injected period (see Figure 8) and had an estimated FAP of less than 1%, then we considered this a successful recovery. If the period of maximum power was more than 10% different from the injected period and the estimated FAP was less than 1%, we considered this a false positive. All other cases were considered nondetections. We did not want to assume that the period of maximum power in our periodograms was due to rotation; thus, we did not remove any preexisting signal from the light curves before injection. To prevent confusion with the signal already present in the CFHT light curves, we removed any case where the injected period was within 10% of the signal already detected in the light curve. This filtered out no more than 11% of the tests in any given bin, with every bin having at least 45,000 tests. This approach enabled us to incorporate the actual

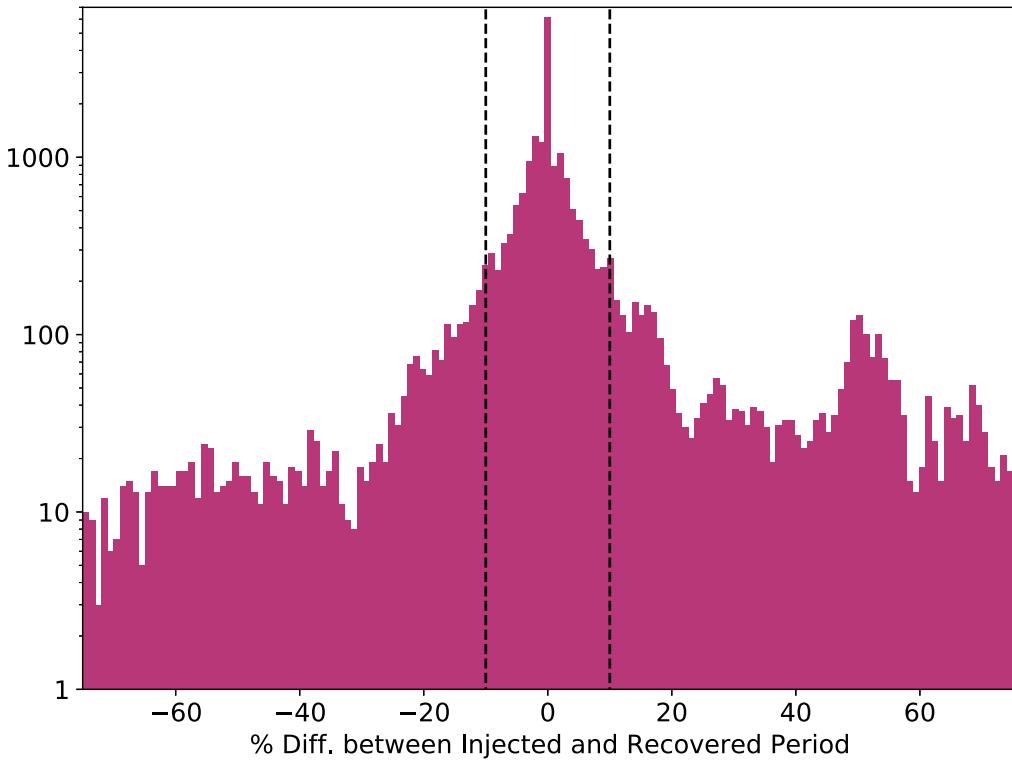


Figure 8. The distribution of the percent differences between the injected and recovered period (irrespective of the FAP of the recovery). The precision on our rotation periods is set by the standard deviation of this distribution: 10%.

systematics present in our CFHT photometry that may have limited the recovery of periodic signals.

Since we are only interested in the rotation periods of the members of M67, we limited the sample of CFHT light curves to a subset of the candidate cluster members and a matching number of randomly selected field stars. We selected the light curves for the injection and recovery testing by applying three criteria. First, cluster members were required to have very high (i.e., =1.0) membership probability, while field stars must have had very low (i.e., =0.0) membership probability. Second, they must have had at least 99 epochs of data available. Finally, they must not have had an overlapping aperture. We also required that the selected field stars have similar $r - i$ colors and apparent i magnitudes to our selected cluster members, to mitigate the impact of any systematics that depended on color. This yielded 740 total light curves, 370 cluster members, and 370 field stars, into which we injected each of our 4599 Kepler and 3934 synthetic light curves. Each injection and CFHT light-curve pairing was repeated with three or four different phases, depending on how many CFHT light curves fit within the injection light curve. This was done to capture the shift in phase due to spot pattern evolution over the years of observations.

We compiled the results of these tests into our completeness diagram (Figure 9), as well as our false-positive diagram (Figure 10). We have plotted the results from the cluster members and field stars separately. Injections into the light curves of cluster members served as a direct test of our ability to recover rotation signals in the cluster member data, while injections into the field stars served as a control sample. The underlying distribution of P_{rot} is different for each of these populations, and thus each is expected to impact the completeness diagram in different ways. Trends that are

common to both figures are thus reflective of the pipeline's recovery capabilities in general. Any differences between the two panels that cannot be attributed to different P_{rot} distributions would reflect issues in the pipeline, but we do not see any such differences.

The completeness diagram (Figure 9) shows the major trends we would expect: (1) as the amplitude of the rotation signal decreases, our ability to recover the correct period also decreases; and (2) the evolving spot patterns in the Kepler light curves reduced our ability to recover the correct period. Our false-positive diagram (Figure 10) also shows the major trends that we expected. In particular, we highlight the 10%–20% difference in false-positive rates between cluster members and field stars for low-amplitude injections ($S_{\text{ph}} \lesssim 0.25\%$). Many of the light curves we injected signals into already had an existing periodic signal, and when injecting low-amplitude signals, we expected to instead recover the already-present signal. This explains both the high percentage of false positives for low-amplitude injections and the difference in false-positive rates. The cluster members were generally expected to show periodic variability due to their spot patterns. On the other hand, a smaller fraction of field stars were expected to show rotational variability, and those that do span a much wider range of timescales (e.g., background evolved stars).

5. Results and Analysis

We present the full rotation catalog in Table 2. The information includes Gaia EDR3 and PS1 DR2 source identifiers, the Gaia EDR3 astrometry and photometry, the PS1 DR2 $griz$ photometry, the recovered rotation periods (if available), their estimated FAP values, the derived T_{eff} values, the percentage probability we calculated for M67 membership, and whether or not the star was flagged as a candidate binary.

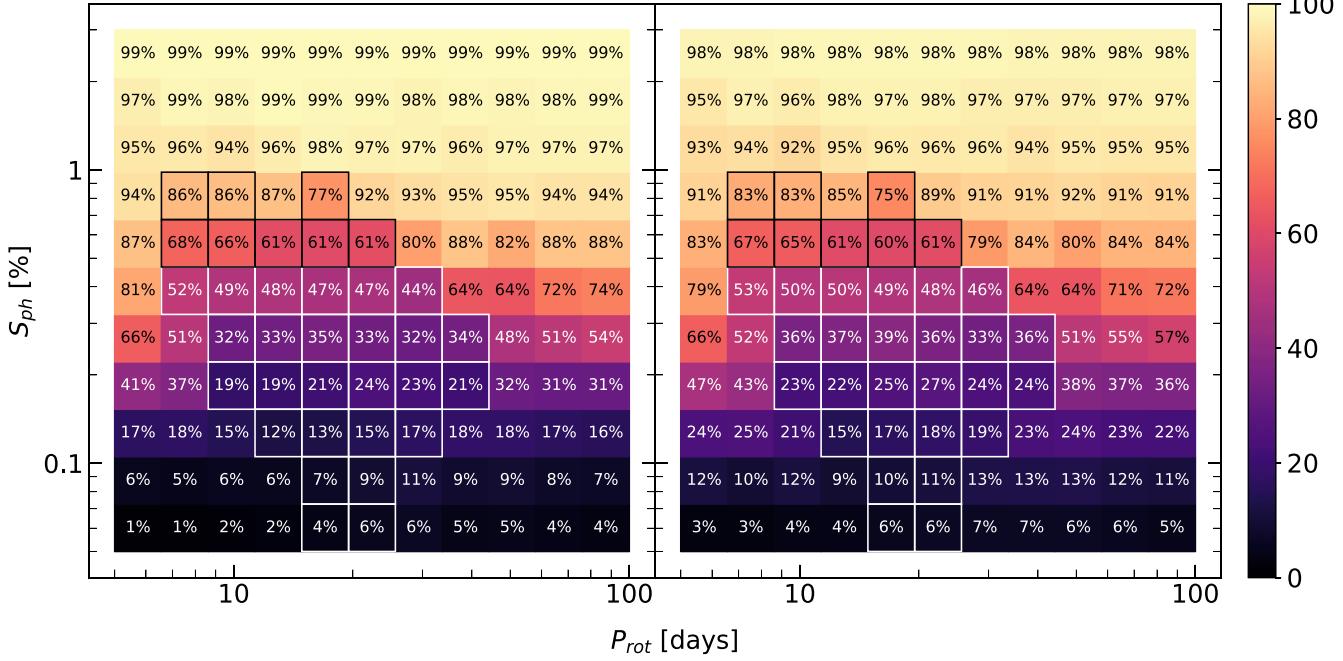


Figure 9. Left panel: the recovery rate (%) of the period of the signal injected into a cluster member’s light curve. Right panel: same as the left panel, but for field stars rather than cluster members. The gradient from top to bottom demonstrates that the lower a signal’s amplitude is, the harder it is to recover. In both panels bins with a black or white (color chosen for optimal contrast) box drawn around them are the ones where over half (i.e., ≥ 25) of the injected signals are Kepler light curves.

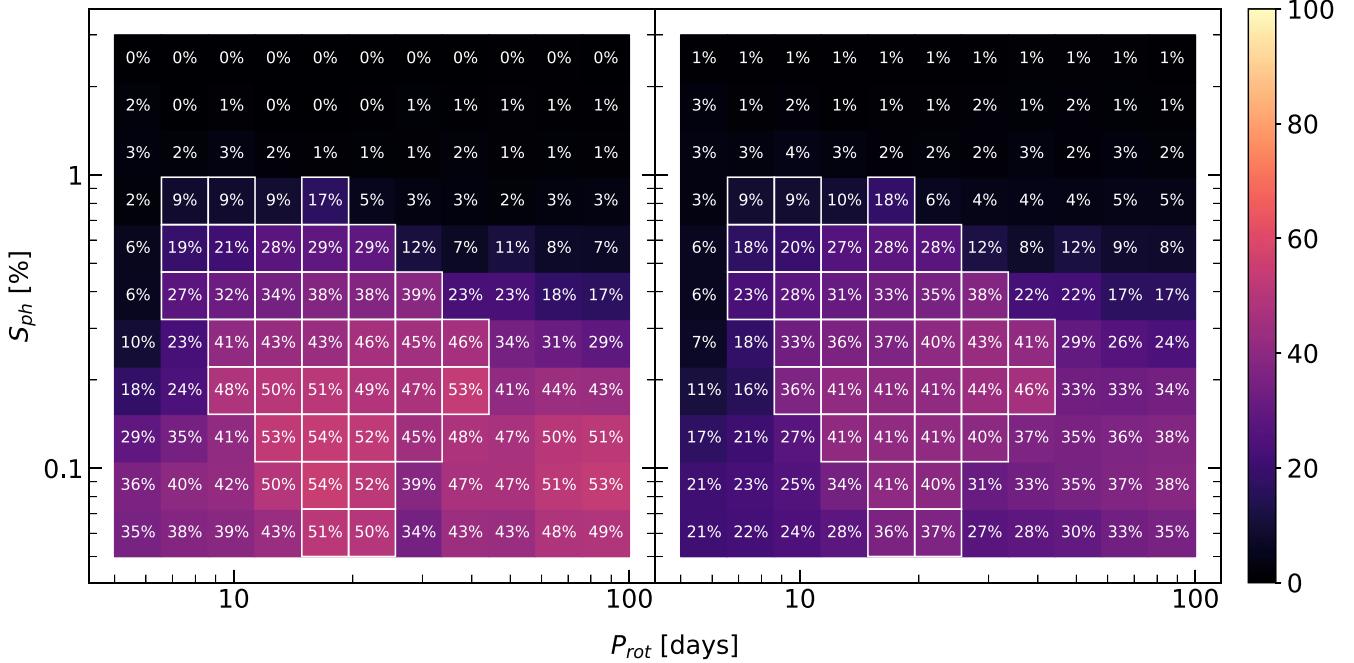


Figure 10. Left panel: the percentage of the time we find a significant (FAP < 0.01) rotation signal but the period does not match what we injected into a cluster member’s light curve (to within 10%). Right panel: same as the left panel, but for injections into a field star instead of a cluster member. Many of these light curves, especially those of the cluster members, already have periodic signals in them. Thus, in the case of a low-amplitude injection we often recover what already existed in the data. Since this does not match what was injected, this is marked as a “false positive.” In both panels bins with a white box drawn around them are the ones where over half (i.e., ≥ 25) of the injected signals are Kepler light curves.

The Appendix contains plotted light curves and periodograms for each star in Table 2 with a reported rotation period.

We have plotted the measured rotation periods versus effective temperature for the 294 candidate single members of M67 with significant rotation detections in Figure 11. For the analysis, we applied two extra cuts on the rotation periods, requiring that the stars have not been flagged as having an

overlapping aperture (Section 2.3) or as a field-of-view-dependent member (Section 3.1). Despite some scatter in the periods, they are concentrated about a locus in $T_{\text{eff}} - P_{\text{rot}}$ space. In an effort to describe this sequence, we performed a polynomial fit to the data using iterative outlier rejection, where at each step outliers were defined as the data greater than three median absolute deviations away from the median of the

Table 2
Catalog of M67 Members

| Name | Table Header | Units | Description |
|----------------------------------|--------------|----------------------|--|
| Gaia source ID | gaiaid | ... | Gaia EDR3 source_id |
| R.A. (R.A.) | RAdeg | deg | ... |
| Error on R.A. | e_RAdeg | mas | ... |
| Decl. (decl.) | DEdeg | deg | ... |
| Error on decl. | e_DEdeg | mas | ... |
| Parallax | plx | mas | ... |
| Error on parallax | e_plx | mas | ... |
| Proper motion in R.A. | pmRA | mas yr ⁻¹ | ... |
| Error on proper motion in R.A. | e_pmRA | mas yr ⁻¹ | ... |
| Proper motion in decl. | pmDE | mas yr ⁻¹ | ... |
| Proper motion in decl. | e_pmDE | mas yr ⁻¹ | ... |
| Renormalized unit weighted error | ruwe | ... | ... |
| Gaia <i>G</i> magnitude | gaiaGmag | mag | ... |
| Error on Gaia <i>G</i> magnitude | e_gaiaGmag | mag | ... |
| Gaia BP magnitude | gaiaBPmag | mag | ... |
| Error on Gaia BP magnitude | e_gaiaBPmag | mag | ... |
| Gaia RP magnitude | gaiaRPmag | mag | ... |
| Error on Gaia RP magnitude | e_gaiaRPmag | mag | ... |
| PS1 source ID | pslid | ... | Pan-STARRS1 DR2 Obj ID |
| PS1 <i>g</i> magnitude | ps1gmag | mag | ... |
| Error on PS1 <i>g</i> magnitude | e_ps1gmag | mag | ... |
| PS1 <i>r</i> magnitude | ps1rmag | mag | ... |
| Error on PS1 <i>r</i> magnitude | e_ps1rmag | mag | ... |
| PS1 <i>i</i> magnitude | ps1imag | mag | ... |
| Error on PS1 <i>i</i> magnitude | e_ps1imag | mag | ... |
| PS1 <i>z</i> magnitude | ps1zmag | mag | ... |
| Error on PS1 <i>z</i> magnitude | e_ps1zmag | mag | ... |
| Probability of membership | memberprob | ... | Probability of membership based on kinematics (Section 3.1) |
| Photometric single star | single | ... | Star was determined to be single (Section 3.1.1) |
| Photometric binary | binary | ... | Star was determined to be a multiple system (Section 3.1.1) |
| Member in Gao's M67 catalog | gaomember | ... | Star is also listed as a member by Gao (2018) |
| Potentially suspect member | suspect | ... | Star's membership depended on catalog size (Section 3.1) |
| Used in fit | converged | ... | Is used in the polynomial fit after outlier rejection. |
| Effective temperature | teff | K | Effective temperature derived from (<i>r</i> − <i>i</i>) color (Section 3.2) |
| Rotation period | prot | d | Rotation period derived from Lomb–Scargle periodograms (Section 4) |
| False-alarm probability | fap | ... | The estimated false-alarm probability of the rotation period |

Note. A description of the columns in the table of results available in machine-readable format. Astrometric and Gaia photometries are taken from Gaia EDR3. PS1 photometry is taken from PS1 DR2.

(This table is available in its entirety in machine-readable form.)

residuals. We did this for both T_{eff} versus P_{rot} and PS1 (*r* − *i*) versus P_{rot} , finding that both approaches converged to the same solution: a subset of 64 stars, for which the least-squares best fits are

$$P_{\text{rot}}(T_{\text{eff},4\text{K}}) = 9.66 \times 10^{-10} \cdot T_{\text{eff},4\text{K}}^4 + 8.25 \times 10^{-7} \cdot T_{\text{eff},4\text{K}}^3 + 2.69 \times 10^{-4} \cdot T_{\text{eff},4\text{K}}^2 + 0.016 \cdot T_{\text{eff},4\text{K}} + 25.9 \quad (7)$$

or

$$P_{\text{rot}}(r - i) = 292 \cdot (r - i)^4 - 895 \cdot (r - i)^3 + 1054 \cdot (r - i)^2 - 543 \cdot (r - i) + 127.9, \quad (8)$$

where $T_{\text{eff},4\text{K}} = T_{\text{eff}} - 4000$ K. We used bootstrapping ($N = 10,000$) to calculate confidence intervals about our fit, fitting a polynomial to 64 stars sampled with replacement from the 253 stars that passed all quality cuts. As a point of comparison, we have also taken the approach of binning the 253 stars in T_{eff} ,

computing a median P_{rot} for each bin, and fitting a polynomial to these medians. The medians are plotted as red squares in Figure 11, and their fit is plotted as a red line, which we have found is in agreement with the iterative outlier approach. We favor the results of the iterative outlier rejection owing to its exclusion of points we believe are aliases from the fit to the T_{eff} versus P_{rot} sequence (see related discussion in Section 5.2).

5.1. Lomb–Scargle Failure Modes

In addition to the 10% uncertainties determined from our injection and recovery tests (Section 4.1), there are systematic uncertainties contributing to the scatter in our results (Figure 11). These are the failure modes of the LS periodogram, originating from the irregular sampling in time. For the purposes of this discussion we will be using the term “window function” in a slightly different manner than in more classical time-series analysis discussions. Instead of describing a traditional window function, such as the Hann window, we take an approach similar to that of VanderPlas (2018), where the window function

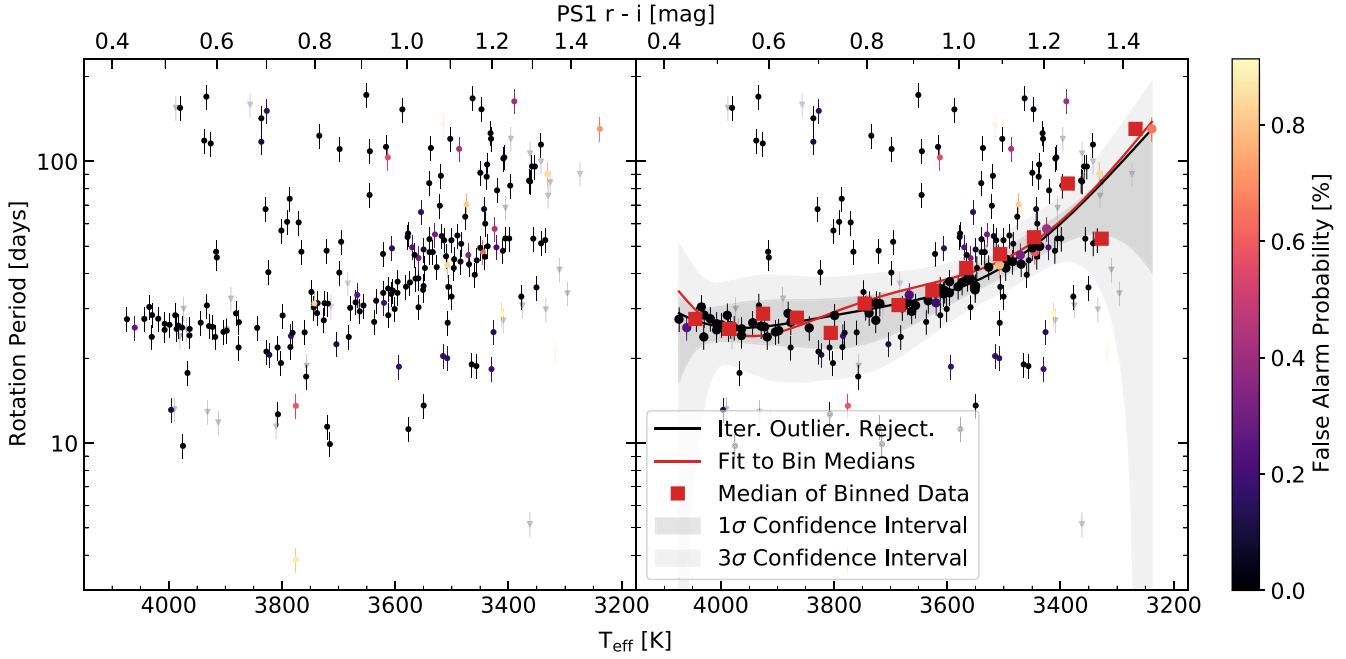


Figure 11. Left panel: recovered rotation periods for members of the open cluster M67 plotted vs. their color-derived effective temperatures. All points have a low FAP (<1%) for the detected rotation period. Right panel: we fit a polynomial to the sequence by iterative outlier rejection. The larger points are used in the polynomial fit after outlier rejection. The shaded region is the confidence interval of this fit. Red squares are the median values of the data binned in T_{eff} , and the red line is a polynomial fit to these values.

describes how the light curve was sampled in time. This window function has predictable effects on the LS periodograms computed from the data, which can all be combined into one equation (Equation (47) in VanderPlas 2018):

$$P_{\text{obs}} = \left| \frac{m}{P_{\text{true}}} + \frac{n}{\delta P} \right|^{-1}, \quad (9)$$

where P_{obs} is the observed peak in the periodogram, P_{true} is the true period of the underlying signal, and m and n are integers. Parameters $m=1$ and $n=0$ yield the true period, and $m=2$ and $n=0$ represent the classic case of half-period aliasing; however, they can both take any integer value, positive or negative. The final term, δP , is the period of a peak in the window function's periodogram; there are typically more than one, and for our CFHT observations there were two dominant ones. They were the “month window peak” ($\delta P \approx 29.5$ days), arising from only observing during the bright lunar phases, and the “year window peak” ($\delta P \approx 380.8$ days), arising from only observing when the cluster is up. The exact values of each depend on the precise sampling in time (i.e., on the completeness of the light curve). The effects of these window peaks can be easily seen in a scatter plot of P_{true} versus P_{obs} , which we have plotted using the results of a complete set of injections into one of our CFHT light curves (Figure 12). There is no way to determine whether the period of maximum power in a periodogram corresponds to P_{true} or one of its failure modes, P_{obs} , with absolute certainty. Moreover, because m and n are integers, there is no continuum of window effects, meaning that a standard deviation computed across the peaks in a periodogram is a poor description of the uncertainty. VanderPlas (2018) provides a prescription for how one can use detected failure modes to improve the accuracy of

interpreting periodograms. We did not use his prescription; instead, we found that many of the stars identified as likely to be failure modes as opposed to true rotation periods by his method are rejected in our iterative outlier rejection and thus already excluded from our analysis. We have included a table of all the detected potential failure modes associated with each of our reported P_{rot} values in Table 3. Readers who are interested in further trimming to create their own subset of the M67 rotation periods reported here may use these values in a prescription like that of Section 7.2 of VanderPlas (2018).

5.2. Deviations from the Sequence

There are a number of mechanisms that can result in an incorrectly measured value of P_{rot} , both observational and astrophysical. Astrophysically, spot pattern evolution can spread the power from a rotation signal into multiple peaks in the periodogram, as well as shift the central location of these peaks. By using Kepler light curves as a part of our injection and recovery tests (Section 4.1), we are able to quantify the effect this has on our recovery. Comparing the bins with majority Kepler light curves to their neighbors with majority synthetic light curves in Figure 9 indicates that spot pattern evolution among the Kepler light curves led to a ∼15%–20% drop in recovery. The same comparison using Figure 10 shows an equivalent uptick in false positives, highlighting the impact of spreading the power across multiple peaks in the periodogram. Additionally, Basri & Nguyen (2018) have shown that stars with lower T_{eff} and longer P_{rot} tend to favor a “double dip” spot pattern that lends itself to half-period aliasing. However, the results of our injection and recovery tests suggest that this is a relatively minor effect for our data set (see the half-period alias line in Figure 12). Finally, close binary systems will have rotation periods that appear as outliers in the data. Such systems are affected by both the confusion of brightness

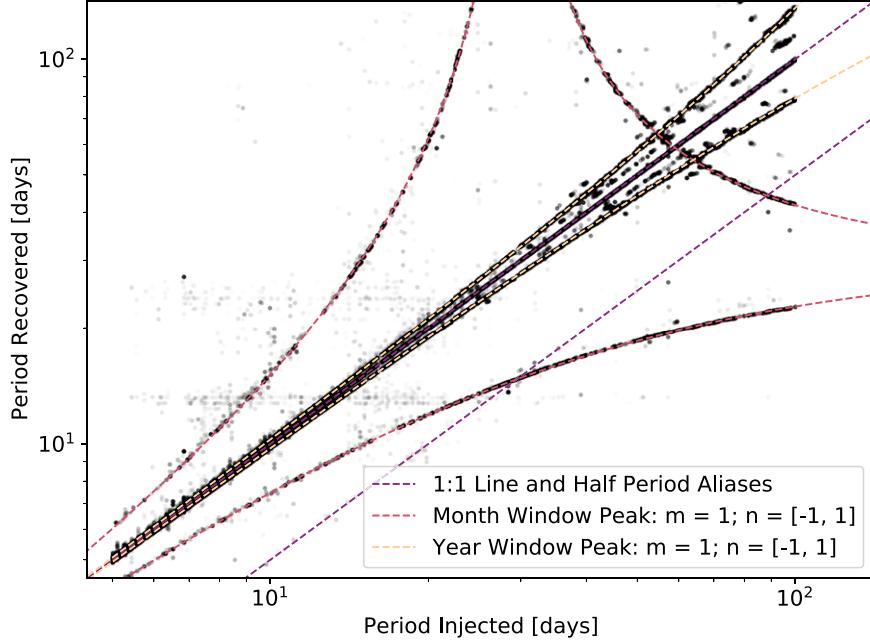


Figure 12. A scatter plot of the injected period vs. the recovered period for one CFHT light curve. Each point represents the result of one Kepler (or synthetic) light curve being added to this cluster member’s light curve. Low-amplitude injections are dominated by the existing signal in the data, resulting in the horizontal features on this diagram. Dashed lines represent the 1:1 line for successful recoveries and the most prominent window function effects: half-period aliasing and the Month and Year window peaks (see Section 5.1 for detailed discussion). Additional peaks in the window function align with other trends in this figure but are not plotted to reduce figure crowding (e.g., $m = 1$, $n = \pm 1$, and $\delta P \sim 180$ would fall between the plotted year-based effect and the 1:1 line).

Table 3
Detected Lomb–Scargle Failure Modes

| Name | Table Header | Units | Description |
|--------------------------------------|--------------|-------|--|
| Gaia source ID | gaiaid | ... | Gaia EDR3 source_id |
| False-alarm probability | fap | ... | The estimated false-alarm probability of the rotation period |
| Rotation period (P_{rot}) | prot | days | The measured rotation period of the star |
| Half-period alias | $m = 2$ | days | $m = 2$, and $n = 0$ |
| Third period alias | $m = 3$ | days | $m = 3$, and $n = 0$ |
| First month failure mode | monthn=-2 | days | $m = 1$, $n = -2$, and $\delta P \approx 29.5$ days |
| Second month failure mode | monthn=-1 | days | $m = 1$, $n = -1$, and $\delta P \approx 29.5$ days |
| Third month failure mode | monthn=+1 | days | $m = 1$, $n = 1$, and $\delta P \approx 29.5$ days |
| Fourth month failure mode | monthn=+2 | days | $m = 1$, $n = 2$, and $\delta P \approx 29.5$ days |
| First year failure mode | yearn=-2 | days | $m = 1$, $n = -2$, and $\delta P \approx 380.8$ days |
| Second year failure mode | yearn=-1 | days | $m = 1$, $n = -1$, and $\delta P \approx 380.8$ days |
| Third year failure mode | yearn=+1 | days | $m = 1$, $n = 1$, and $\delta P \approx 380.8$ days |
| Fourth year failure mode | yearn=+2 | days | $m = 1$, $n = 2$, and $\delta P \approx 380.8$ days |

Note. A description of the columns that are in the machine-readable table. All alias and failure mode values are the output of Equation (9) with $P_{\text{true}} = P_{\text{rot}}$ and are only included in the table if a peak was detected at that period in the periodogram.

(This table is available in its entirety in machine-readable form.)

modulations in both stars and tidal forces changing their rotational evolution relative to single stars. We have mitigated the contamination from binary stars through our CMD cuts (Section 3.1.1). Observationally, we were also limited by our irregular sampling in time; such effects are described by the LS failure modes (see Section 5.1). This means that even a perfectly stable sinusoid can be recovered incorrectly as the signal-to-noise ratio on the data decreases. With these effects in mind, we believe that the use of iterative outlier rejection for our reported fit to the T_{eff} versus P_{rot} sequence was justified.

To demonstrate this, we substituted Equation (7) for P_{true} in Equation (9) and plotted the resulting sequences alongside our original results in Figure 13. Given the difference in the effect of each window peak, we plotted the month and year effects

separately. For the collection of stars rotating faster than our fit to the T_{eff} versus P_{rot} sequence, they align well with two possible cases: (1) they fall along the half-period alias of our fit, or (2) they fall along a sequence that is associated with the month window peak (left panel of Figure 13). The sequences associated with the year window peak show that these failure modes contribute to the scatter about our fit to the sequence. This, combined with the precision on our P_{rot} values derived from the injection and recovery tests (Figure 8), is what prevents us from measuring a T_{eff} versus P_{rot} sequence that is as sharply defined as the slow rotator sequences observed in younger clusters (e.g., Praesepe and NGC 6811; Douglas et al. 2017, 2019; Curtis et al. 2019).

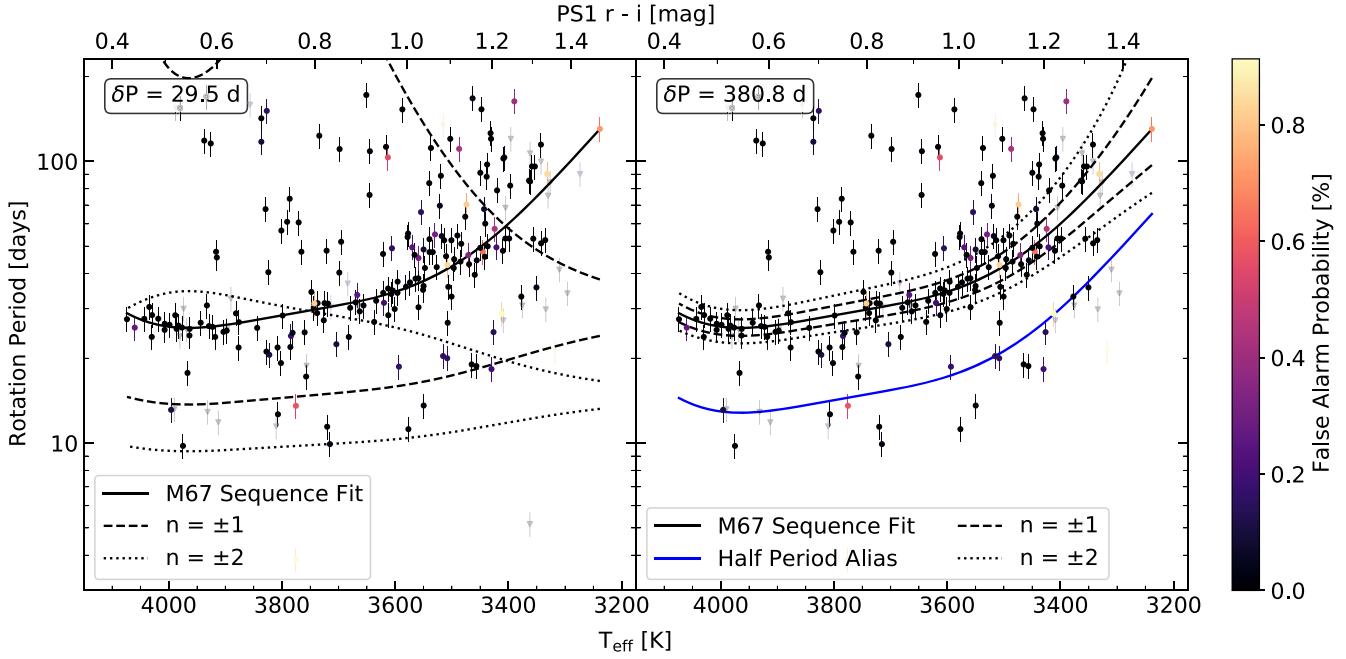


Figure 13. The rotation period vs. effective temperature of M67 members with our fit to the M67 sequence (solid black). Dashed lines are the result of assuming that our fit to the sequence is the true period in Equation (9), along with $m = 1$ and $n = \pm 1$. Dotted lines are the same, but with $n = \pm 2$. Left panel: using 29.5 days for δP Equation (9). Right panel: using 380.8 days for δP Equation (9). Additionally, the half-period alias of Equation (7) is in solid blue.

Finally, there remain a number of stars with relatively long rotation periods ($P_{\text{rot}} \gtrsim 40$ days) and high temperatures ($T_{\text{eff}} \gtrsim 3600$ K) that are inconsistent with our fitted sequence and its expected failure modes. To gain some insight into the origin of these inconsistent stars, we binned our P_{rot} values and computed the fraction of stars in a period bin that were inconsistent with the fitted sequence. Then, we estimated the S_{ph} values for these stars, allowing us to compare the computed fractions to the false-positive rates of Figure 10. The stars inconsistent with the sequence have a mean S_{ph} of 0.3% with a standard deviation of 0.2%, compared to a mean S_{ph} of 0.4% with a standard deviation of 0.1% for the stars of equivalent temperature consistent with the sequence. Given these S_{ph} values, the inconsistency fractions of 20%–50% align well with the false-positive rates. Moreover, the inconsistent stars all have detected failure modes that are consistent with the fitted sequence itself, whereas those on the sequence primarily have detected failure modes consistent with a half-period alias. A multiterm model would aid in clarifying the true rotation period of these stars; however, the sparse sampling in our data made those fits poorly conditioned. Therefore, we are satisfied with their exclusion from our fit to the sequence. We postulate that these stars were affected by either spot pattern evolution or the signal-to-noise ratio of the data, both of which contribute to how likely a failure mode is to be recovered instead of the true rotation period. Confirmation would require follow-up observations at much higher cadence and with more regular sampling in time.

6. Discussion

We first compared our observations (the 64 stars our iterative outlier rejection converged to; Section 5) against the predictions of two classes of gyrochronological models: empirical and theoretical (Figure 14). Empirical models are agnostic to the physics of magnetic braking and angular momentum transport, fitting a relation to a set of periods and ages often as a

function of color. Generally they follow a Skumanich-like relation of $P_{\text{rot}} \propto t^n$, using cluster data and the Sun as anchor points to tune the value of the exponent. On the other hand, theoretical models make assumptions regarding the underlying physics and spin-down that manifests from their description. Whether they are empirical or theoretical in nature, all models are calibrated against objects of known age and rotation period, relatively few of which are young M dwarfs, meaning that their predictions for the age of M67 are an extrapolation.

There are four empirical relations included in our comparison. First is the Barnes (2010) model, where dP_{rot}/dt is parameterized in terms of the Rossby number (Ro) and two dimensionless constants calibrated on the Sun and young open cluster observations. Second is the Angus et al. (2019) empirical relation, which is a broken power law with mass fit to the sequence of Praesepe and a spin-down law tuned to replicate the Sun. Finally, we evolved the sequences of Praesepe (670 Myr; Douglas et al. 2017, 2019) and Ruprecht 147 (2.7 Gyr; Curtis et al. 2020) forward in time through the use of a simple Skumanich-like spin-down: $P_{\text{rot}}(4 \text{ Gyr}) = P_{\text{rot},0}(4 \text{ Gyr}/t_0)^{0.62}$. For the hotter stars in our sample all of these empirical relations, with one exception, predict that the stars of M67 should be rotating ~ 10 –20 days slower than observed. The exception is the Skumanich-like spin-down relation launched from the stars of Ruprecht 147, which provides an excellent match to our data for the earlier M dwarfs ($T_{\text{eff}} \gtrsim 3700$ K); for the later M dwarfs there were no data available in Ruprecht 147.

The first of the two theoretical models we considered is from van Saders & Pinsonneault (2013) and described therein. Similar to our Skumanich-like empirical relation, we launched this model from two starting places: Praesepe and Ruprecht 147. This model is calibrated to match the spin-down of solar-mass stars and the Sun and then applied to our low-mass M67 members. Braking laws of this form perform well on the Sun but fail to capture the rotational evolution during the first few

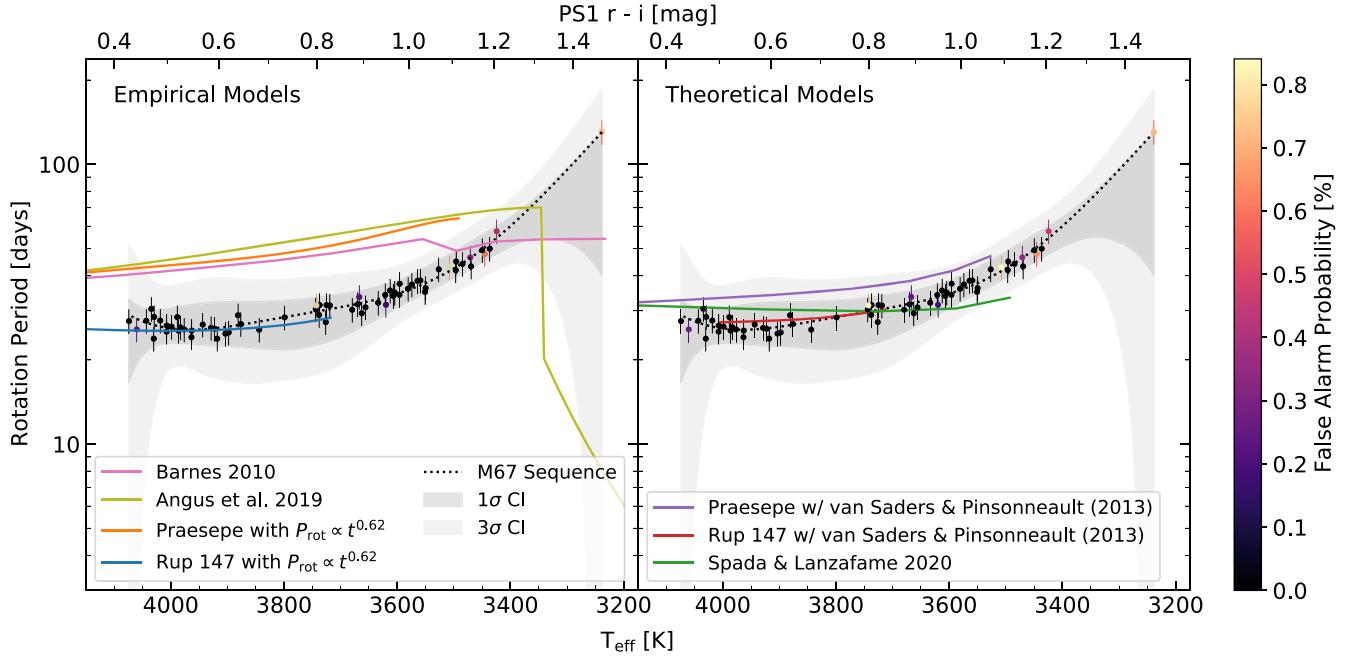


Figure 14. Left panel: the selected subset (from iterative outlier rejection; Section 5) of M67 rotation periods vs. their effective temperatures. The shaded region in gray is the confidence interval of the fit, plotted as a dotted black line. Predictions for the sequence from the empirical models are plotted as colored lines. Right panel: same as the left panel, but with predictions from the theoretical models.

hundreds of Myr in young clusters (see Douglas et al. 2017; Breimann et al. 2021; Roquette et al. 2021); it is unclear to what degree the issue is with the braking law itself (saturation, core–envelope decoupling, etc.) or with the assumed distribution of initial rotation periods (Roquette et al. 2021). We manage these early-time uncertainties by starting our models as solid-body rotators at Praesepe age and rotation rate and evolving them forward to M67. The model evolved forward from Praesepe is a better match than the empirical relations, as it is within the 3σ confidence interval, but still predicts that the stars of M67 should be rotating slower than we observed. The model evolved forward from Ruprecht 147 provides an excellent match, being very closely aligned with the empirical relation launched from the same starting point, although this is again limited by a lack of later M dwarfs in the Ruprecht 147 data.

The second theoretical model we included is the model of Spada & Lanzafame (2020). This iteration on the model includes some minor adjustments from that of Lanzafame & Spada (2015), a core–envelope decoupling model. Their model incorporates a mass-dependent wind-braking law that follows the classical rotation rate dependence of Kawaler (1988), $\frac{dJ}{dt} \propto \omega^3$. It also uses a two-zone approach to the interior, treating the core and envelope as two separate rotationally solid bodies that are allowed to exchange angular momentum. As a result, this model explains the apparent stalling of spin-down as an epoch during which significant angular momentum transport occurs from the core to the envelope, balancing out the angular momentum the envelope loses to wind braking. The prediction of this model agrees with the observations quite well down to a T_{eff} of around 3600 K. At this point, the model’s prediction is too fast compared to our observations, though still consistent at the 3σ level.

In total, we have found three models that provide an excellent (i.e., within 1σ) match to our observations of M67. Of these models, two are solid-body spin-down (one theoretical,

one empirical) applied to the stars of Ruprecht 147, which we interpret as a sign that the late stages of low-mass stellar spin-down are dominated by solid-body rotation. The other model is that of Spada & Lanzafame (2020), which is both the only core–envelope decoupling model tested and the only model launched from the birth of the star. The excellent agreement between it and our observations makes a compelling case for the core–envelope decoupling theory.

6.1. The Case for Core–Envelope Decoupling

The evidence for core–envelope decoupling goes deeper than the agreement between the model of Spada & Lanzafame (2020) and our observations. In the core–envelope decoupling framework, after the epoch of significant angular momentum transport occurs, the expectation is that the core and envelope of the star have equalized in angular velocity. At this point, the star spins down as a solid body. If the stars of Ruprecht 147 have resumed their spin-down (Curtis et al. 2020), core–envelope decoupling would predict that they are now spinning down as solid bodies. The precise agreement between our solid-body and empirical models launched from Ruprecht 147 and our observations implies that this is the case, at least through the age of M67.

Another important test of core–envelope decoupling lies in the behavior of spin-down for stars that are nearly or fully convective. The diminishing size of the core limits the amount of angular momentum it can store relative to the envelope, reducing the length of time the star’s spin-down would stall. Furthermore, stars with no radiative core should not stall their spin-down at all. Curtis et al. (2020) provide an empirical relation for the age at which stars resume spinning down. They find

$$t_R = 0.202 \text{ Gyr} \times \left(\frac{T_{\text{eff}}}{5770 \text{ K}} \right)^{-5.11}, \quad (10)$$

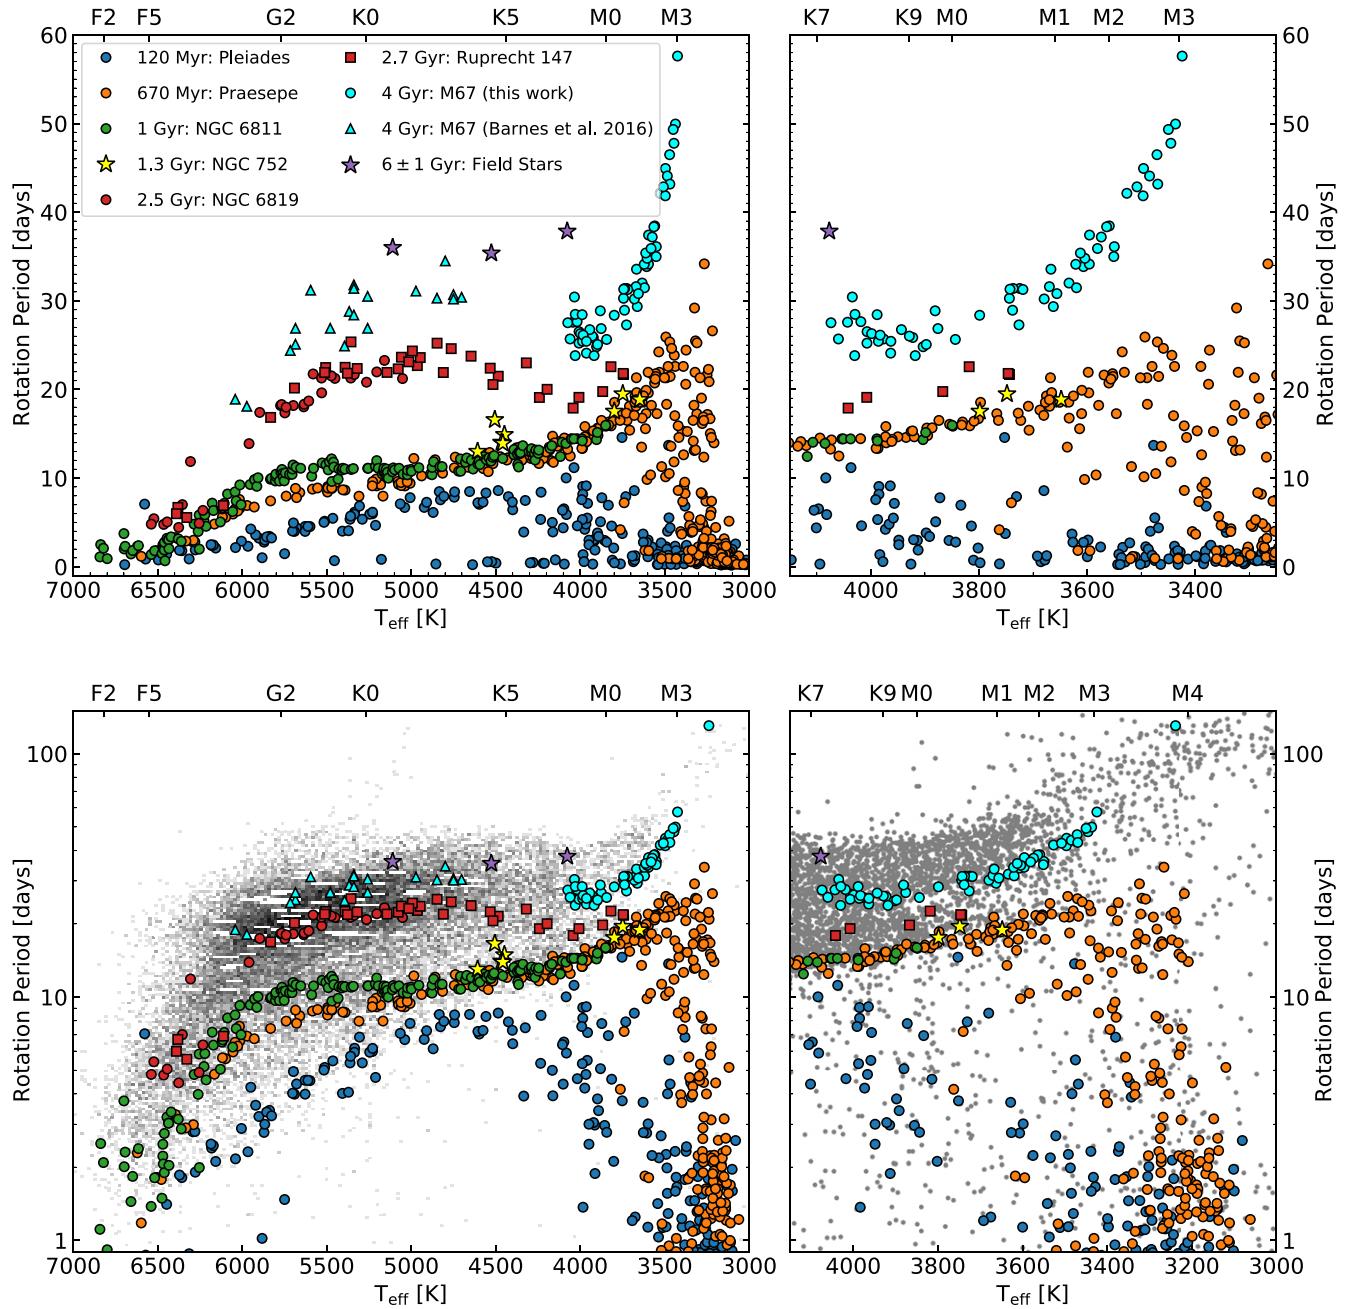


Figure 15. Top left panel: a replication of the first panel of Figure 7 of Curtis et al. (2020), but with our results and the results of Barnes et al. (2016) added. Top right panel: a subset of the top left panel, zoomed in on the T_{eff} range of the results presented in this paper. Bottom left panel: Same as the top left panel, but now plotted over a 2D histogram of field star rotation periods. Bottom right panel: same as the top right panel, but now plotted over field star rotation periods. In all cases, the colored points are a collection of open clusters that have been used as gyrochronological benchmarks. Benchmarks include the Pleiades (120 Myr; Rebull et al. 2016), Praesepe (670 Myr; Douglas et al. 2017, 2019), NGC 6811 (1 Gyr; Curtis et al. 2019), NGC 752 (1.4 Gyr; Agüeros et al. 2018), NGC 6819/Ruprecht 147 (2.5 Gyr projected forward by Curtis et al./2.7 Gyr; Meibom et al. 2015; Curtis et al. 2020), M67 (4 Gyr; Barnes et al. 2016; this work), and three field stars: α Cen B and 61 Cyg A and B (Table 3 of Curtis et al. 2020, and references therein). Field stars are taken from a collection of literature sources: Kepler (Santos et al. 2019), the PS1 Medium Deep Survey (Kado-Fong et al. 2016), MEarth (Newton et al. 2016, 2018), CARMENES (Díez Alonso et al. 2019), Evryscope (Howard et al. 2020), and the K2SDSS sample (Popinchalk et al. 2021).

based on a simplified model where spin-down comes to a full stop and then suddenly resumes after a mass-dependent length of time. If we extrapolate this relation to later spectral types (i.e., beyond M0), we find that by M1.5 the age at which spin-down resumes is approaching that of M67 (3.6 Gyr vs. 4 Gyr), and by M2 this age has potentially exceeded that of M67 (4.6 Gyr). However, the stars in this range of temperatures (M1–M3) are rotating \sim 10–30 days slower than their younger

counterparts (top two panels of Figure 15), implying that they have been spinning down for at least part of the intervening \sim 3 Gyr. This is suggestive of a need for a t_{R} relation that has a turnover as it approaches the fully convective boundary, as expected in the core–envelope decoupling framework. Observations of younger M dwarfs of these spectral types (e.g., those in Ruprecht 147 or NGC 752) will be a critical test for determining when these stars resumed their spin-down.

6.2. M67 and the Field

We have also compared our observations to an ensemble of field star rotation periods collected from a variety of sources. The largest contributor to this collection is the Kepler sample, with temperatures and rotation periods from Santos et al. (2019). The rest are predominantly M dwarfs with rotation periods from the PS1 Medium Deep Survey (Kado-Fong et al. 2016), MEarth (Newton et al. 2016, 2018), CARMENES (Díez Alonso et al. 2019), Evryscope (Howard et al. 2020), and the K2SDSS sample (Popinchalk et al. 2021). Popinchalk et al. (2021) provided the Gaia DR2 identifiers for the targets from all of these surveys, which we used to obtain temperatures from v8 of the TESS Input Catalog (TIC; Stassun et al. 2019). We then plotted these field stars with the open cluster data (bottom two panels of Figure 15), ignoring any stars that did not have a temperature in TIC.

Field M dwarfs follow a bimodality in their rotation periods (Kado-Fong et al. 2016; Newton et al. 2016; Howard et al. 2020). Using kinematic ages, Newton et al. (2016) speculated that the transition between the fast and slow populations must be quick and must occur between the ages of 2 and 5 Gyr. The stars of M67 fall along the lower envelope of the slow rotator population, suggesting that they represent convergence onto a slow rotator sequence for M dwarfs. The age of M67 (4 Gyr) is consistent with the bounds for the transition. Higher-cadence observations are needed to confirm whether or not there are still rapid rotators in M67, a lack of which would make M67 a fully converged slow rotator sequence. Since accurate gyrochronology depends on stars converging to a slow rotator sequence, the age of M67 serves as a lower bound for accurate gyrochronological ages of M dwarfs.

Another interesting feature seen in the distribution of field star rotation periods is the intermediate period gap. This is a bimodal distribution of stars with T_{eff} values less than 5000 K and intermediate rotation periods (15–25 days; McQuillan et al. 2013). A number of explanations have been put forth to explain this gap, including a lull in star formation (Davenport 2017), a transition to faculae-dominated photospheres (Reinhold et al. 2019), or an epoch of accelerated spin-down during the recoupling of the core and envelope (McQuillan et al. 2013; Gordon et al. 2021). Open cluster data show that any explanation relying on the gap stars having a common age is incorrect (Curtis et al. 2020). Instead, the mechanism that causes this gap must occur at different times for stars of different masses. The stars of M67 appear along the upper envelope of the intermediate period gap, suggesting an upper bound of 4 Gyr for the age by which this mechanism has occurred. Furthermore, if the gap is indeed caused by accelerated spin-down during core–envelope recoupling, then the stars along the upper envelope of the intermediate period gap should be composed of stars that are spinning down as solid bodies, in line with our observations.

6.3. Evidence of a Unique Spin-down History

While this description is compelling, some caution is important. Somers & Pinsonneault (2016) identified M67 as an outlier among open clusters in terms of its lithium abundance. Having demonstrated that Li depletion is a strong test of core–envelope recoupling, they concluded that the most likely scenario explaining M67’s Li abundances is an “intrinsically different mixing history” driven by a surplus of

rapid rotators in the cluster’s early years. Observations of young clusters and associations show that massive stars in large clusters can drive photoevaporation of the disks of nearby lower-mass stars, shortening disk lifetimes and resulting in a larger population of rapid rotators (Roquette et al. 2021). Such a surplus of rapid rotators would shift the mean sequence of M67 to faster rotation periods compared to stars of equivalent ages until the initial conditions are forgotten. However, this will not affect the braking laws describing their spin-down. We can control for M67’s unique initial rotation periods by modeling a variety of cases for the initial conditions, as well as observing other clusters of similar ages.

7. Conclusions

In this paper we have done the following:

1. Generated a new catalog of 1807 M67 members based on Gaia EDR3 parallaxes and proper motions, identified potential unresolved binaries by their location on the cluster’s CMD, and calculated the color-based effective temperatures for the late K- and early M-dwarf single members of M67.
2. Reported the rotation periods for 294 of these M67 members, providing a sample of 4 Gyr old late K and early M dwarfs for calibrating gyrochronological models and a polynomial fit to the sequence they form in T_{eff} versus P_{rot} for use as a gyrochronone:

$$P_{\text{rot}}(T_{\text{eff},4\text{K}}) = 9.66 \times 10^{-10} \cdot T_{\text{eff},4\text{K}}^4 + 8.25 \times 10^{-7} \cdot T_{\text{eff},4\text{K}}^3 + 2.69 \times 10^{-4} \cdot T_{\text{eff},4\text{K}}^2 + 0.016 \cdot T_{\text{eff},4\text{K}} + 25.9 \quad (11)$$

or

$$P_{\text{rot}}(r - i) = 292 \cdot (r - i)^4 - 895 \cdot (r - i)^3 + 1054 \cdot (r - i)^2 - 543 \cdot (r - i) + 127.9. \quad (12)$$

3. Having compared the gyrochronological models to our gyrochronone, we found that late K and early M dwarfs spin down as solid bodies between 2.7 and 4 Gyr of age. This behavior is broadly consistent with core–envelope decoupling models of stellar spin-down.

R.D. would like to acknowledge the organizers of the Fifty Years of the Skumanich Relation conference for facilitating some enlightening discussions around these data, especially those with Jason Curtis. R.D., J.v.S., and E.G. acknowledge support from NSF Astronomy & Astrophysics grant AST-1817215. A.R.G.S. acknowledges the support from the FCT and FEDER/COMPETE2020 through work contract No. 2020.02480.CEECIND/CP1631/CT0001 and grants UIDP/04434/2020, PTDC/FIS-AST/30389/2017, and POCI-01-0145-FEDER-030389. R.A.G. acknowledges the support from a PLATO CNES grant. S.M. acknowledges support by the Spanish Ministry of Science and Innovation with the Ramon y Cajal fellowship No. RYC-2015-17697 and grant No. PID2019-107187GB-I00. Based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA/DAPNIA, at the Canada–France–Hawaii Telescope (CFHT), which is operated by the National Research Council (NRC) of Canada, the Institut National des Sciences de l’Univers of the Centre National de la Recherche Scientifique (CNRS) of

France, and the University of Hawaii. The observations at the Canada–France–Hawaii Telescope were performed with care and respect from the summit of Maunakea, which is a significant cultural and historic site. The Pan-STARRS1 Surveys (PS1) and the PS1 public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, Johns Hopkins University, Durham University, the University of Edinburgh, the Queen’s University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission

Directorate, the National Science Foundation grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation.

Facilities: CFHT (MegaPrime), Gaia, PS1.

Software: astropy (Astropy Collaboration et al. 2013, 2018), photutils (Bradley et al. 2020), numpy (Harris et al. 2020), scipy (Virtanen et al. 2020), matplotlib (Hunter 2007), hdbscan (McInnes et al. 2017), jupyter (Kluyver et al. 2016), astroquery (Ginsburg et al. 2019).

Appendix

Example Light Curves and Periodograms

A light curve, phase-folded light curve, and periodogram are available for every star in Table 2; included here are two examples (Figure 16). The full set is available in the online journal.

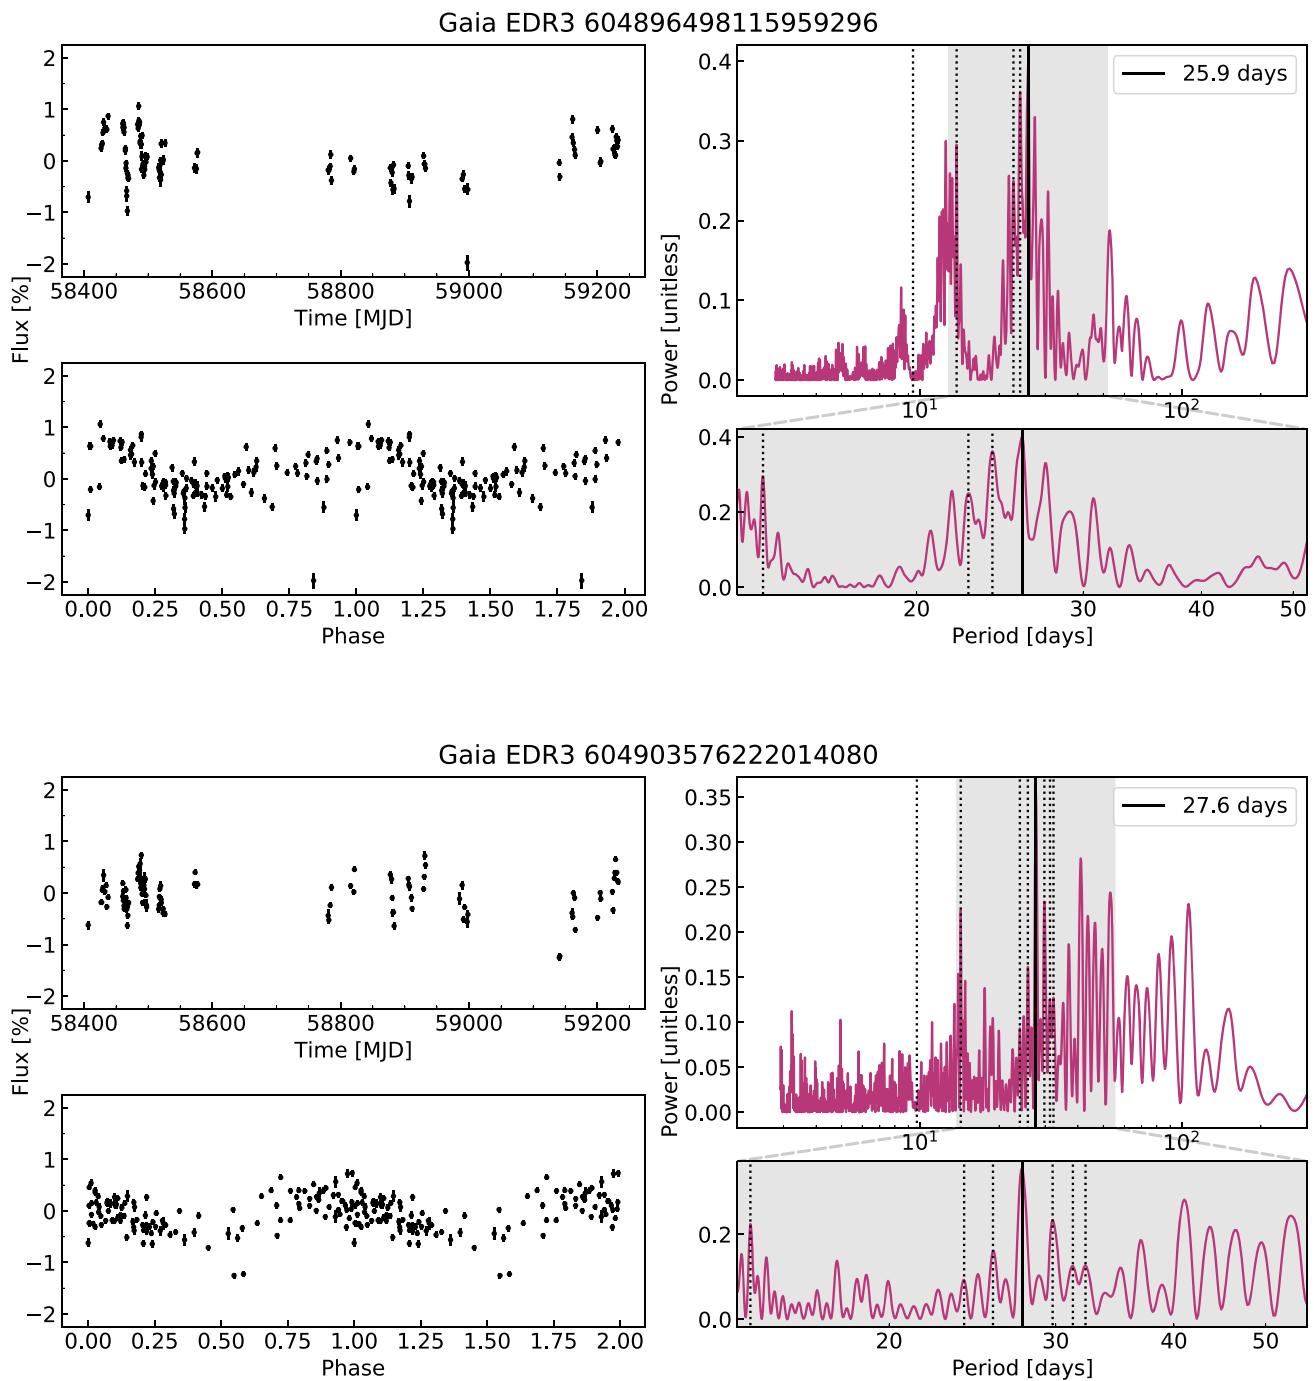


Figure 16. Example light curves and periodograms for select cluster members. In the top left of each target’s set is the light curve of the target. Bottom left is the phase-folded version, folded on the period of max power in the periodogram. On the right, the periodogram and a subset centered on the period of max power. Window peaks are denoted with vertical dotted lines. The complete figure set (383 images) is available in the online journal.

(The complete figure set (383 images) is available.)

ORCID iDs

Ryan Dungee  <https://orcid.org/0000-0001-6669-0217>
 Jennifer van Saders  <https://orcid.org/0000-0002-4284-8638>
 Eric Gaidos  <https://orcid.org/0000-0002-5258-6846>
 Mark Chun  <https://orcid.org/0000-0002-8462-0703>
 Rafael A. García  <https://orcid.org/0000-0002-8854-3776>
 Eugene A. Magnier  <https://orcid.org/0000-0002-7965-2815>

Savita Mathur  <https://orcid.org/0000-0002-0129-0316>
 Ângela R. G. Santos  <https://orcid.org/0000-0001-7195-6542>

References

Agüeros, M. A., Bowsher, E. C., Bochanski, J. J., et al. 2018, *ApJ*, **862**, 33
 Allain, S. 1998, *A&A*, **333**, 629
 Angus, R., Aigrain, S., Foreman-Mackey, D., & McQuillan, A. 2015, *MNRAS*, **450**, 1787
 Angus, R., Morton, T. D., Foreman-Mackey, D., et al. 2019, *AJ*, **158**, 173

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33

Aumer, M., & Binney, J. J. 2009, *MNRAS*, **397**, 1286

Balaguer-Núñez, L., Galadí-Enríquez, D., & Jordi, C. 2007, *A&A*, **470**, 585

Baluev, R. V. 2008, *MNRAS*, **385**, 1279

Barnes, S. A. 2007, *ApJ*, **669**, 1167

Barnes, S. A. 2003, *ApJL*, **586**, L145

Barnes, S. A. 2010, *ApJ*, **722**, 222

Barnes, S. A., Weingrill, J., Fritzewski, D., Strassmeier, K. G., & Platais, I. 2016, *ApJ*, **823**, 16

Basri, G., & Nguyen, H. T. 2018, *ApJ*, **863**, 190

Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, *MNRAS*, **496**, 1922

Berdinas, Z. M., Rodriguez-Lopez, C., Amado, P. J., et al. 2017, *MNRAS*, **469**, 4268

Bildsten, L., Brown, E. F., Matzner, C. D., & Ushomirsky, G. 1997, *ApJ*, **482**, 442

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

Boulade, O., Charlot, X., Abbon, P., et al. 2003, *Proc. SPIE*, **4841**, 72

Bradley, L., Sipőcz, B., Robitaille, T., et al. 2020, astropy/photutils: v1.0.0, Zenodo

Breimann, A. A., Matt, S. P., & Naylor, T. 2021, *ApJ*, **913**, 75

Carraro, G., & Chiosi, C. 1994, *A&A*, **288**, 751

Chaplin, W. J., Bedding, T. R., Bonanno, A., et al. 2011, *ApJL*, **732**, L5

Curtis, J. L., Agüeros, M. A., Douglas, S. T., & Meibom, S. 2019, *ApJ*, **879**, 49

Curtis, J. L., Agüeros, M. A., Matt, S. P., et al. 2020, *ApJ*, **904**, 140

Davenport, J. R. A. 2017, *ApJ*, **835**, 16

Deacon, N. R., & Kraus, A. L. 2020, *MNRAS*, **496**, 5176

Demarque, P., Green, E. M., & Guenther, D. B. 1992, *AJ*, **103**, 151

Denissenkov, P. A., Pinsonneault, M., Terndrup, D. M., & Newsham, G. 2010, *ApJ*, **716**, 1269

Díez Alonso, E., Caballero, J. A., Montes, D., et al. 2019, *A&A*, **621**, A126

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

Douglas, S. T., Curtis, J. L., Agüeros, M. A., et al. 2019, *ApJ*, **879**, 100

Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, **807**, 45

Duchêne, G., & Kraus, A. 2013, *ARA&A*, **51**, 269

Epstein, C. R., & Pinsonneault, M. H. 2014, *ApJ*, **780**, 159

Fan, X., Burstein, D., Chen, J. S., et al. 1996, *AJ*, **112**, 628

Flewelling, H. A., Magnier, E. A., Chambers, K. C., et al. 2020, *ApJS*, **251**, 7

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, *A&A*, **649**, A1

Gallet, F., & Bouvier, J. 2015, *A&A*, **577**, A98

Gao, X. 2018, *ApJ*, **869**, 9

García, R. A., Hekker, S., Stello, D., et al. 2011, *MNRAS*, **414**, L6

García, R. A., Mathur, S., Pires, S., et al. 2014, *A&A*, **568**, A10

Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, **157**, 98

Gordon, T. A., Davenport, J. R. A., Angus, R., et al. 2021, *ApJ*, **913**, 70

Gould, A., Bahcall, J. N., & Flynn, C. 1996, *ApJ*, **465**, 759

Hardegree-Ullman, K. K., Cushing, M. C., Muirhead, P. S., & Christiansen, J. L. 2019, *AJ*, **158**, 75

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357362

Holmberg, J., Nordström, B., & Andersen, J. 2007, *A&A*, **475**, 519

Howard, W. S., Corbett, H., Law, N. M., et al. 2020, *ApJ*, **895**, 140

Hunter, J. D. 2007, *CSE*, **9**, 90

Kado-Fong, E., Williams, P. K. G., Mann, A. W., et al. 2016, *ApJ*, **833**, 281

Kawaler, S. D. 1988, *ApJ*, **333**, 236

Keppens, R., MacGregor, K. B., & Charbonneau, P. 1995, *A&A*, **294**, 469

Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87

Krishnamurthi, A., Pinsonneault, M. H., Barnes, S., & Sofia, S. 1997, *ApJ*, **480**, 303

LanzaFame, A. C., & Spada, F. 2015, *A&A*, **584**, A30

Laughlin, G., Bodenheimer, P., & Adams, F. C. 1997, *ApJ*, **482**, 420

Liu, F., Asplund, M., Yong, D., et al. 2016, *MNRAS*, **463**, 696

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

Lu, Y. L., Angus, R., Curtis, J. L., David, T. J., & Kiman, R. 2021, *AJ*, **161**, 189

MacGregor, K. B., & Brenner, M. 1991, *ApJ*, **376**, 204

Magnier, E. A., & Cuillandre, J. C. 2004, *PASP*, **116**, 449

Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., et al. 2020, *ApJS*, **251**, 6

Mamajek, E. E., & Hillenbrand, L. A. 2008, *ApJ*, **687**, 1264

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

Mathur, S., García, R. A., Bugnet, L., et al. 2019, *FrASS*, **6**, 46

Mathur, S., Salabert, D., García, R. A., & Ceillier, T. 2014, *JSWSC*, **4**, A15

Matt, S. P., Brun, A. S., Baraffe, I., Bouvier, J., & Chabrier, G. 2015, *ApJL*, **799**, L23

McInnes, L., Healy, J., & Astels, S. 2017, *JOSS*, **2**, 205

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

Meibom, S., Barnes, S. A., Platais, I., et al. 2015, *Natur*, **517**, 589

Meibom, S., Mathieu, R. D., & Stassun, K. G. 2009, *ApJ*, **695**, 679

Montgomery, K. A., Marschall, L. A., & Janes, K. A. 1993, *AJ*, **106**, 181

Newton, E. R., Irwin, J., Charbonneau, D., et al. 2016, *ApJ*, **821**, 93

Newton, E. R., Irwin, J., Charbonneau, D., et al. 2017, *ApJ*, **834**, 85

Newton, E. R., Mondrik, N., Irwin, J., Winters, J. G., & Charbonneau, D. 2018, *AJ*, **156**, 217

Nissen, P. E., Twarog, B. A., & Crawford, D. L. 1987, *AJ*, **93**, 634

Önehag, A., Korn, A., Gustafsson, B., Stempels, E., & Vandenberg, D. A. 2011, *A&A*, **528**, A85

Pace, G., Pasquini, L., & François, P. 2008, *A&A*, **489**, 403

Pires, S., Mathur, S., García, R. A., et al. 2015, *A&A*, **574**, A18

Popinchalk, M., Faherty, J. K., Kiman, R., et al. 2021, *ApJ*, **916**, 77

Press, W. H., & Rybicki, G. B. 1989, *ApJ*, **338**, 277

Rebull, L. M., Stauffer, J. R., Bouvier, J., et al. 2016, *AJ*, **152**, 113

Rebull, L. M., Stauffer, J. R., Hillenbrand, L. A., et al. 2017, *ApJ*, **839**, 92

Reinhold, T., Bell, K. J., Kuszlewicz, J., Hekker, S., & Shapiro, A. I. 2019, *A&A*, **621**, A21

Roquette, J., Matt, S. P., Winter, A. J., Amard, L., & Stasevic, S. 2021, *MNRAS*, **508**, 3710

Salabert, D., García, R. A., Beck, P. G., et al. 2016, *A&A*, **596**, A31

Salabert, D., García, R. A., Jiménez, A., et al. 2017, *A&A*, **608**, A87

Sandquist, E. L., Mathieu, R. D., Quinn, S. N., et al. 2018, *AJ*, **155**, 152

Santos, A. R. G., García, R. A., Mathur, S., et al. 2019, *ApJS*, **244**, 21

Santos, N. C., Lovis, C., Pace, G., Meléndez, J., & Naef, D. 2009, *A&A*, **493**, 309

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

Simonian, G. V. A., Pinsonneault, M. H., Terndrup, D. M., & van Saders, J. L. 2020, *ApJ*, **898**, 76

Skumanich, A. 1972, *ApJ*, **171**, 565

Somers, G., & Pinsonneault, M. H. 2016, *ApJ*, **829**, 32

Somers, G., Stauffer, J., Rebull, L., Cody, A. M., & Pinsonneault, M. 2017, *ApJ*, **850**, 134

Spada, F., & LanzaFame, A. C. 2020, *A&A*, **636**, A76

Stassun, K. G., Oelkers, R. J., Paegert, M., et al. 2019, *AJ*, **158**, 138

Stello, D., Vanderburg, A., Casagrande, L., et al. 2016, *ApJ*, **832**, 133

Stetson, P. B. 1987, *PASP*, **99**, 191

Taylor, B. J. 2007, *AJ*, **133**, 370

Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, *ApJ*, **750**, 99

van Saders, J. L., & Pinsonneault, M. H. 2013, *ApJ*, **776**, 67

van Saders, J. L., Pinsonneault, M. H., & Barbieri, M. 2019, *ApJ*, **872**, 128

VandenBerg, D. A., & Stetson, P. B. 2004, *PASP*, **116**, 997

VanderPlas, J. T. 2018, *ApJS*, **236**, 16

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261

Weber, E. J., & Davis, L. J. 1967, *ApJ*, **148**, 217

Winters, J. G., Henry, T. J., Jao, W.-C., et al. 2019, *AJ*, **157**, 216

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

Zechmeister, M., & Kürster, M. 2009, *A&A*, **496**, 577