The Astrophysical Journal, 913:70 (14pp), 2021 May 20 https://doi.org/10.3847/1538-4357/abf63e

Stellar Rotation in the K2 Sample: Evidence for Modified Spin-down

Tyler A. Gordon[1], James R. A. Davenport[1], Ruth Angus[2], Daniel Foreman-Mackey[3], Eric Agol[1], Kevin R. Covey[4], Marcel A. Agüeros[5], and David Kipping[5]

[1] Department of Astronomy, University of Washington, Box 351580, U.W., Seattle, WA 98195-1580, USA; tagordon@uw.edu
[2] Department of Astrophysics, American Museum of Natural History, 200 Central Park West, Manhattan, NY, USA
[3] Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010 USA
[4] Department of Physics & Astronomy, Western Washington University, MS-9164, 516 High Street, Bellingham, WA, 98225 USA
[5] Department of Astronomy, Columbia University, Manhattan, NY, USA

Received 2021 January 19; revised 2021 April 1; accepted 2021 April 6; published 2021 May 26

Abstract

We analyze light curves of 284,834 unique K2 targets using a Gaussian process model with a quasi-periodic kernel function. By cross-matching K2 stars to observations from Gaia Data Release 2, we have identified 69,627 likely main-sequence stars. From these we select a subsample of 8977 stars on the main sequence with highly precise rotation period measurements. With this sample we recover the gap in the rotation period−color diagram first reported by McQuillan et al. While the gap was tentatively detected in Reinhold & Hekker, this work represents the first robust detection of the gap in K2 data for field stars. This is significant because K2 observed along many lines of sight at wide angular separation, in contrast to Kepler's single line of sight. Together with recent results for rotation in open clusters, we interpret this gap as evidence for a departure from the $t^{1/2}$ Skumanich spin-down law, rather than an indication of a bimodal star formation history. We provide maximum likelihood estimates and uncertainties for all parameters of the quasi-periodic light-curve model for each of the 284,834 stars in our sample.

Unified Astronomy Thesaurus concepts: Stellar properties (1624); Stellar rotation (1629); K dwarf stars (876); G dwarf stars (556); Gaussian Processes regression (1930); Star clusters (1567); Stellar astronomy (1583); Stellar ages (1581)

1. Introduction

Stellar rotation is a key physical property for understanding individual stars as well as stellar populations. Rotation drives the stellar dynamo that produces surface magnetic fields. These magnetic fields in turn give rise to stellar activity (e.g., starspots and flares). The rotation period (P_{rot}) of a star is tied to its age through magnetic braking, which slows the star's rotation over time (Durney 1972; Skumanich 1972). Age is a fundamental stellar parameter but is difficult to determine from the position of a star on a color–magnitude diagram, especially for stars on the main sequence. Any information we can extract about the age of a star from its rotation period is therefore very valuable. This is the subject of gyrochronology, which seeks to measure stellar ages by observing the star's rate of rotation (Barnes 2003).

The Kepler mission (Borucki et al. 2010) revolutionized the study of stellar rotation by producing high-precision light curves for hundreds of thousands of stars, from which over 34,000 rotation periods have been inferred (Nielsen et al. 2013; McQuillan et al. 2014). The distribution of rotation periods measured by McQuillan et al. (2013) showed an unexpected bimodality in the field M dwarfs, which was found to extend to K dwarfs by McQuillan et al. (2014). This bimodality was recovered in G dwarfs by Davenport & Covey (2018), who used Gaia astrometry to limit their analysis to main-sequence stars with well-determined Gaia photometric solutions, removing contamination by subgiants.

Several explanations have been put forward to explain this bimodal period distribution. Davenport & Covey (2018) and McQuillan et al. (2013, 2014) suggest that the bimodality may be the result of a bimodal star formation history, with a recent burst of star formation accounting for the fast-rotating branch of the bimodality and an older population of stars forming the slow-rotating branch. Reinhold et al. (2019) propose that the gap between modalities may represent a minimum in detectability of rotation periods due to the transition from spot-dominated to faculae-dominated stellar activity.

A third possibility, to be discussed in more detail in Section 4, is that the gap results directly from the spin evolution of G, K, and M dwarfs. An epoch of stalled spin-down followed by a period of rapid angular momentum loss before the resumption of Skumanich spin-down may be able to explain such a feature. To date, rotation periods from open clusters have provided the most compelling evidence that modified spin evolution is indeed the cause of the gap. These fixed-age populations have shown that rotation periods for low-mass stars break from the expected Skumanich spin-down model, such as the cluster of rotation periods at $P_{rot} \approx 10$ days found in the 1 Gyr old cluster NGC 6811 by Meibom et al. (2011). Similar deviations from the traditional Skumanich spin-down profile have been seen in, e.g., Praesepe at 650 Myr (Douglas et al. 2017) and NGC 752 at 1.3 Gyr (Agüeros et al. 2018). Curtis et al. (2019) note in their analysis of NGC 6811 that the stall in spin-down appears to be mass and age dependent. Further, Curtis et al. (2020) show compelling evidence that this deviation from Skumanich spin-down indeed corresponds to the gap in field rotation periods, with individual cluster sequences "crossing" the rotation period gap.

This scenario may be explained in terms of time-variable and mass-dependent rotational coupling between the core and envelope of the star (Spada & Lanzafame 2020). Magnetic braking slows the rotation of the convective envelopes of stars. However, if the core and envelope are only weakly coupled, the stellar core may continue to spin rapidly even as the envelope slows down. A decoupled core and envelope with reduced angular momentum exchange is expected for young stars

(Endal & Sofia 1981; MacGregor & Brenner 1991; Bouvier 2008; Denissenkov et al. 2010; Gallet & Bouvier 2013; Lanzafame & Spada 2015; Somers & Pinsonneault 2016). After a period of time, which appears to depend on stellar mass, the core and the envelope begin exchanging angular momentum. When angular momentum transport is happening efficiently, the core's angular momentum transferred to the envelope would offset magnetic braking, allowing the envelope to maintain a constant rotation period. After the core and envelope have coupled and the star rotates as a solid body, the star would resume spinning down. This process would result in a departure from the Skumanich spin-down law (Skumanich 1972), which prescribes a smooth spin-down over time following the relation

$$P_{\mathrm{rot}} \propto t^{-1/2}, \tag{1}$$

where P_{rot} is the stellar rotation period as a function of the age of a star, t. This scenario may explain the convergence of cluster sequences below the gap. Explaining the underdensity within the gap requires us to posit an additional stage in stellar spin evolution consisting of a period of accelerated spin-down immediately after the epoch of stalled spin-down and before the resumption of Skumanich spin-down (Curtis et al. 2020). This accelerated spin-down is not predicted by the coupling scenario presented here, and its explanation will likely require further theoretical work.

The Kepler data alone gives us a limited ability to explore these various hypotheses due to its single pointing, which admits the possibility that the bimodality is unique to the Kepler field. In contrast, K2 observed the sky in 18 separate campaigns, each having different lines of sight (save for a few overlapping campaigns). As van Saders et al. (2019) note, if the Kepler line of sight happened to point directly through a late burst of star formation, thereby accounting for the fast-rotating branch of the bimodal period distribution, we would not expect this feature to be visible in all 18 K2 campaigns.

In this work, we measure and report probabilistic constraints on periodic signals for 284,834 K2 stars from all 18 campaigns and analyze a subset of 8943 highly accurate rotation periods. For those stars appearing in multiple campaigns, we run our analysis separately for each light curve. We use a modification of the Gaussian Process (GP) regression method described in Angus et al. (2018) to measure periodic signals. We find that the bimodality is visible in all K2 campaigns, lending support to the idea that the feature is related to stellar physics rather than being a product of the star formation history within the Kepler field.

2. Measuring Rotation Periods

We begin by describing the model that we use to infer probabilistic rotation periods from the EVEREST light curves (Luger et al. 2018). Stellar magnetic activity induces starspots and faculae on the star's surface. As the star rotates, these features are carried into and out of view, introducing periodicity into the light curve. If the starspots and faculae were static over time, we would observe a perfect periodicity, with the star returning to the same luminosity once every period. However, starspots and faculae are not static but rather evolve over time, emerging, changing shape, and disappearing as the star's rotation brings them into and out of view. As a result, the light curves do not display perfectly periodic variations but rather a quasi-periodic variability with the shape and amplitude of the variability changing from period to period. This means that inferring rotation periods using straightforward sinusoidal variability models does not give good results. Instead, non-inference-based methods such as autocorrelation functions (ACFs; McQuillan et al. 2013), or Lomb-Scargle periodograms (Reinhold et al. 2013) can be employed. An alternative to these non-inference methods is to use a stochastic variability model such as a GP (Angus et al. 2018). In this work, we use the ACF to derive a multimodal "prior" over the period. We then use a Markov chain Monte Carlo (MCMC) method to estimate posteriors for the parameters of the quasi-periodic GP variability model defined in Section 2.2. The use of quotation marks around the word "prior" references the fact that this is not technically a Bayesian prior because it does not strictly depend on our prior beliefs about the period distribution. We explain this further and describe the effect that this has on our analysis in the next section.

2.1. Autocorrelation Function Analysis

We use long-cadence K2 EVEREST light curves from Luger et al. (2018) as the starting point for our analysis. We use the cotrending basis-vector-corrected flux (keyword FCOR in the EVEREST FITS files), which removes systematic trends from the raw light curves. To further remove long-term trends, we subtract a third-order polynomial from each light curve before computing the ACF. This has the effect of flattening the overall decay of the ACF at short time lags, improving our ability to detect the rotation period from the ACF peak. We note that the third-order polynomial may overfit and remove some rotation signals for the slowest rotators in our sample. This primarily affects rotation periods longer than about 25 days, which does not interfere with our detection or analysis of the period gap in Section 3. In our initial experiments, we found that a second-order polynomial did not sufficiently flatten the decay of the ACF, and higher-order polynomials were too likely to overfit the rotation signal. We remove outliers from our light curves by masking all flux observations greater than 3σ from a running median with a kernel width of five K2 long cadences.

After preprocessing the light curve as described above, we compute the ACF for each light curve using the implementation provided in exoplanet, which wraps the astropy ACF function (Astropy Collaboration et al. 2018; Foreman-Mackey et al. 2019). We smooth the ACF with a Gaussian filter with a kernel width of 0.5 days. We then use the smoothed ACF to construct a unique multimodal period "prior," which we find aids in convergence during the MCMC step. We place "prior" in quotes because, strictly speaking, a prior should only reflect our prior beliefs about the period distribution rather than depending on the data itself. The distribution we derive here is not technically a prior although we use it as such. Because we are building our "prior" from the same data that we use to fit the GP, we risk underestimating the uncertainty on the period. Because our analysis focuses on point estimates of the period rather than the full posterior, we elected to accept this risk in exchange for the benefit of recovering more rotation periods.

Our prior is a Gaussian mixture with 3N components, where N is given by

$$N = \begin{cases} N_{\mathrm{peaks}}(p > 0.01) & N_{\mathrm{peaks}}(p > 0.01) < 10 \\ 10 & N_{\mathrm{peaks}}(p > 0.01) \geq 10, \end{cases} \tag{2}$$

where $N_{\mathrm{peaks}}(p)$ is an integer corresponding to the number of peaks in the ACF with topographical prominence greater than

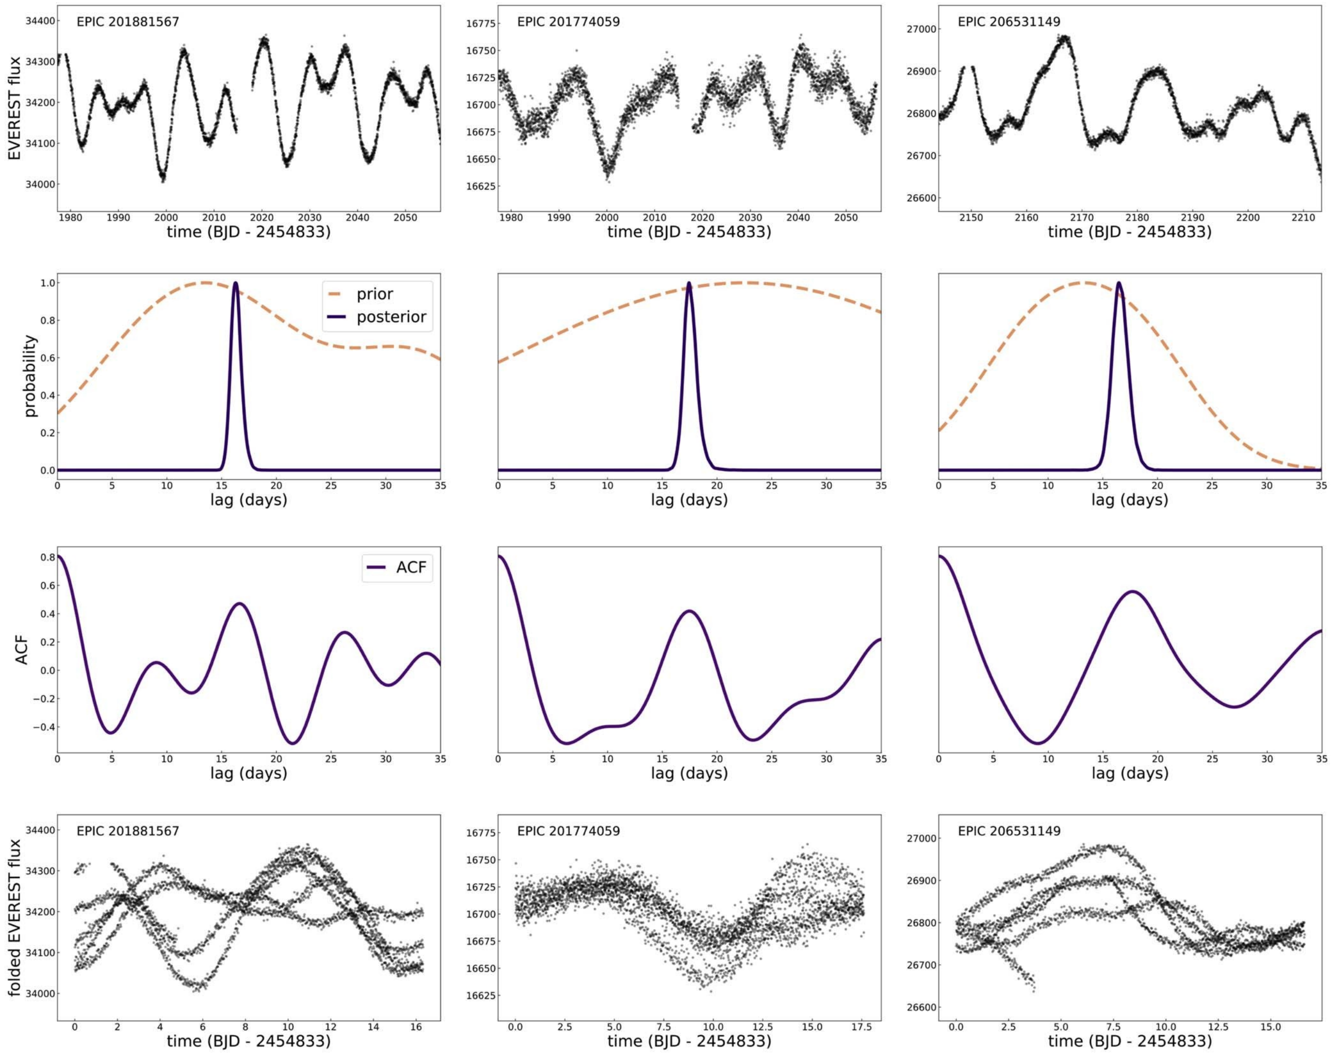

Figure 1. Sample output from our period detection procedure for three K2 stars with well-determined rotation periods. The top row of panels shows the cotrending basis-vector-detrended EVEREST flux. The second row shows the period prior and MCMC-estimated posterior. The third row shows the autocorrelation function, a the bottom row shows the light curve folded on the mean of the posterior for the period. For visibility of the prior, both the prior and posterior are normalized such th their maximum probability is 1.

p. The topographical prominence is computed with respect to the adjacent ACF minima, by scipy's signal library (Virtanen et al. 2020). We take each of these peaks to represent a candidate period, recognizing that for a well-defined periodic signal there will be multiple peaks corresponding to the same period.

The factor of 3 in 3N arises from our inclusion of candidate periods at $\tau_i/2$ and $2\tau_i$, where τ_i is the lag of the ith peak so that for each peak we have three candidate periods. The weight of the component of the Gaussian mixture prior corresponding to each peak is given by

$$w_i = h_i \sqrt{p_i}, \quad (3)$$

where h_i is the height of the peak and is the same for the candidate periods at $\tau_i/2$ and $2\tau_i$ as for the candidate period at τ_i itself. The standard deviation of each Gaussian component, σ_i, is given by the width of the peak at half of the peak height. This means that the width of the Gaussian component of the prior is wider by a factor of approximately 2.35 than the standard deviation of the Gaussian equivalent to the ACF peak. The standard deviation associated with the candidate period at τ_i is also used for the candidate periods at $2\tau_i$ and $\tau_i/2$. In the case where no peaks are detected in the ACF, we adopt a uniform prior over the range $P = (0, \Delta T/2)$, where ΔT is the total duration of the light curve.

The choices we made in computing w and σ are motivated by the logic that a higher peak should be given more weight in the mixture than a lower peak and that the component in the mixture should have a smaller standard deviation if the peak in the ACF is sharper, reflecting the smaller uncertainty on the corresponding period. We also aim to construct a prior that is informative but that does not prohibit the MCMC from exploring periods not identified as period candidates by our algorithm. As can be seen in Figure 1, the period priors we construct tend to have wide regions of high probability and therefore limited influence over our point estimates of the rotation period.

2.2. Gaussian Processes and MCMC Analysis

A GP can be thought of as a distribution from which we may draw random functions with a given covariance structure. They are commonly used in astrophysics to model stochastic variability in light curves (see Dawson et al. 2014, Barclay et al. 2015, and Chakrabarty & Sengupta 2019 for examples from studies of transiting exoplanets and MacLeod et al. 2010 for an example in which a GP is used to model AGN variability). A GP can be split into two components: a kernel function $k(\tau)$ that describes the covariance of the functions in the distribution and a mean function $\mu(t)$. The kernel function defines the covariance matrix of a multidimensional Gaussian distribution by specifying the covariance between every pair of flux measurements. The covariance matrix is given by

$$K_{ij} = k(\tau_{ij}), \tag{4}$$

where $\tau_{i,j} = |t_i - t_j|$ is the absolute value of the separation between times t_i and t_j and $K_{i,j}$ is the ith, jth entry of the covariance matrix K.

The log of the likelihood function of the GP is given by

$$\ln \mathcal{L} = -\frac{1}{2}(\boldsymbol{y} - \boldsymbol{\mu})^{\mathrm{T}} K^{-1}(\boldsymbol{y} - \boldsymbol{\mu}) - \frac{1}{2}\ln\det(K) - \frac{N}{2}\ln(2\pi), \tag{5}$$

where y is a vector of observations and μ is a mean vector with the same length as y. Both the kernel and mean are parameterized by a set of hyperparameters. GP regression is the process of finding the hyperparameters that maximize the GP's likelihood with respect to a set of observations. For a more detailed primer on GP in astronomy, see Foreman-Mackey et al. (2017), or, for a more complete resource on GP across fields, we refer the reader to Rasmussen & Williams (2006).

To construct a GP stellar rotation model, we take the mean function of the GP to be constant and allow the kernel function to model the correlated variability introduced into the star's light curve by spots and faculae as they rotate in and out of view and evolve over time. Our GP model has three terms: two quasi-periodic terms to capture the rotationally induced variability and one that is aperiodic to capture any leftover variability originating from other astrophysical sources or instrumental effects. The power spectrum of each term is given by

$$S_i(\omega) = \sqrt{\frac{2}{\pi}} \frac{S_i \omega_i^4}{(\omega^2 - \omega_i^2)^2 + 2\omega^2 \omega_i^2}. \tag{6}$$

For the periodic terms, we follow Foreman-Mackey et al. (2017) in setting

$$\begin{aligned} Q_1 &= 1/2 + Q + \Delta Q \\ Q_2 &= 1/2 + Q \\ \omega_1 &= \frac{4\pi Q_1}{P\sqrt{4Q_1^2 - 1}} \\ \omega_2 &= \frac{8\pi Q_1}{P\sqrt{4Q_1^2 - 1}} \\ S_1 &= \frac{\sigma^2}{(1 + f)\,\omega_1 Q_1} \\ S_2 &= \frac{f\sigma^2}{(1 + f)\,\omega_2 Q_2}, \end{aligned} \tag{7}$$

where Q is the quality factor, ΔQ specifies the offset in the quality factor between the two oscillators, P is the period of the oscillator, σ^2 is the variance of the oscillation, and f specifies the fractional contribution of the oscillator at the half period P/2 compared to the oscillator at the full period. For the aperiodic term we set

$$Q_3 = 1/\sqrt{2}, \tag{8}$$

while ω_3 and S_3 are free parameters. Setting $Q_3 = 1/\sqrt{2}$ means that the third oscillator is critically damped and will not display periodic oscillations. For this term the power spectrum simplifies to

$$S_3(\omega) = \sqrt{\frac{2}{\pi}} \frac{S_3}{(\omega/\omega_3)^4 + 1}. \tag{9}$$

The full variability model including both the quasi-periodic and aperiodic terms has the power spectrum

$$S(\omega) = \sum_i^3 S_i(\omega). \tag{10}$$

To compute the GP model, we use the celerite GP method (Foreman-Mackey et al. 2017) as implemented in exoplanet (Foreman-Mackey et al. 2019). We maximize the GP likelihood with respect to the EVEREST cotrending basis-vector-detrended flux for the parameters {P, Q, ΔQ, A, f, S_3, ω_3}. We then use the maximum likelihood solution as a starting point for our MCMC analysis. We use uninformative priors for all GP hyperparameters except the period, for which we use the multimodal Gaussian mixture prior described previously. We use the NUTS sampler provided by PyMC3 (Salvatier et al. 2016) to run 1000 tuning samples followed by 500 production samples on each of 28 cores for a total of 28,000 tuning and 14,000 production samples. We have found that a relatively large number of tuning samples is helpful for achieving convergence when using a multimodal period prior in order to allow the sampler to fully explore the multimodal likelihood space.

In Figure 2 we show the variation in the binned mean of the period P, maximum quality factor $Q_{\max} = \max(Q_1, Q_2)$, the logarithm of the ratio between the periodic and aperiodic components of the model, and the logarithm of the fractional uncertainty in rotation period.

2.3. Selecting Main-sequence Stars

We begin by making selections based on the quality of the Gaia DR2 photometric solutions (Gaia Collaboration et al. 2018). We require that the following conditions be met:

1. $\sigma(G)/G < 0.01$
2. $\sigma(G_{RP})/G_{RP} < 0.01$,

where G and G_{RP} refer to the passbands used in Gaia DR1 and DR2.

In order to reduce contamination from giants, subgiants, and unresolved binary stars, we require that stars in our sample be on or near the main sequence, as defined by a MIST isochrone (Paxton et al. 2011, 2013, 2015; Choi et al. 2016; Dotter 2016) with an age of 200 Myr and a metallicity of [Fe/H] = +0.25, to identify the nominal main sequence, and we select stars within 0.3 mag below and 0.9 mag above the isochrone, as shown in Figure 3. This wide slice of magnitude space allows us to encompass different ages and metallicities while reducing contamination from the giant and subgiant branches. The cost of selecting such a wide slice is that we likely incorporate a significant number of unresolved binaries into our final sample,

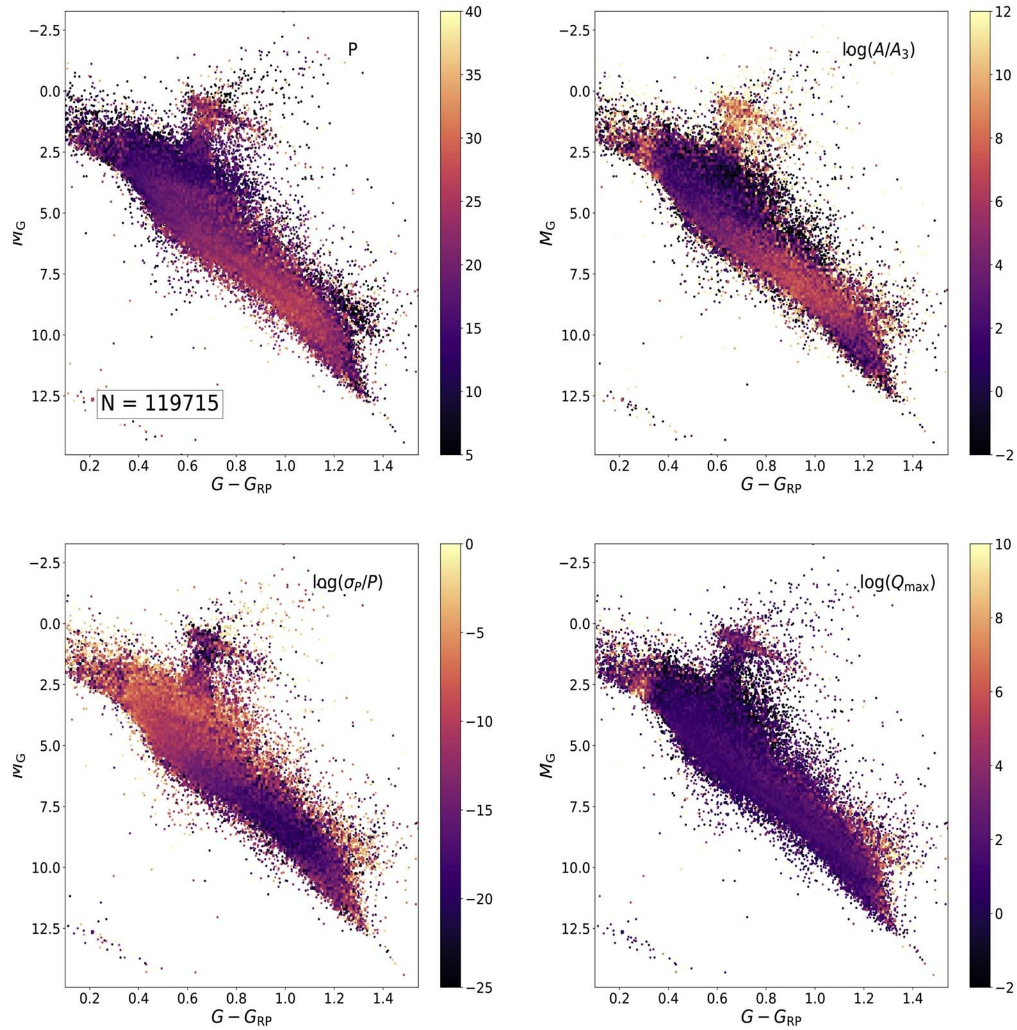

Figure 2. Selected hyperparameters plotted over the Gaia color–magnitude diagram. For each plot the color in a bin indicates the mean of the quantity given in the upper-right-hand corner of the plot and the scattered points are colored by that quantity in regions where the density of stars is low. Upper left: rotation period. Upp right: log of the ratio between the periodic variance, A, and the variance of the aperiodic component, A_3. Lower left: log of the fractional uncertainty for the inferred period. Lower right: log of the mean quality factor, $Q_{\max}$, with larger $Q_{\max}$ indicating stronger periodicity.

which will add some amount of contamination. We find this acceptable because we do not expect this contamination to have a systematic influence on the overall shape of the period–color diagram.

It should be emphasized that we made no attempt to choose an isochrone that represents the actual main sequence for stars in our sample, which would be infeasible because the K2 sample contains stars with a wide range of ages and metallicities. The choice of [Fe/H] = +0.25 was made on the basis that this isochrone does an adequate job of matching the trend of the main sequence. We have found that the exact choice of age and metallicity does not have a significant impact on our results, so long as the isochrone and the width of the box in M_G selects a sufficient number of stars for our edge finding algorithm to perform well. Our final sample consists of 8943 stars near the main sequence, passing our Gaia photometry cuts and possessing well-determined periodicity.

2.4. Vetting Rotation Periods

We select a final sample of well-measured rotators from the main-sequence sample based on MCMC convergence and period measurement precision. For inclusion of a star in our final sample, we require the following conditions be met:

1. $P/\sigma_P > 15$
2. $\log_{10}(A/A_3) > -3$
3. $0.9 < \hat{R}_P < 1.1$,

where P is the measured period, σ_P is the error on the period derived from MCMC, A is the variance of the periodic GP component (the amplitude is $= \sqrt{A}$, and $A_3 = S_3\omega_3Q_3$ is the amplitude of the aperiodic GP component), $\hat{R}_P$ is the Gelman-Rubin statistic (Gelman & Rubin 1992), which compares the variance of samples for an individual parameter (in this case the period, P) within a chain to the variance between chains. For chains that have converged to the same solution, these values will be approximately the same, and their ratio, $\hat{R}$, will be close to 1. The cutoff on $\log_{10}(A/A_3)$ is meant to exclude stars for which the periodic component is very small compared to the nonperiodic variability on the basis that these stars are more likely to be showing periodicity due to contamination or systematics, rather than rotation. We find that when we do not include this cutoff, a pileup of stars at a period of around two days is observed. This pileup, as shown in Figure 4, spans a

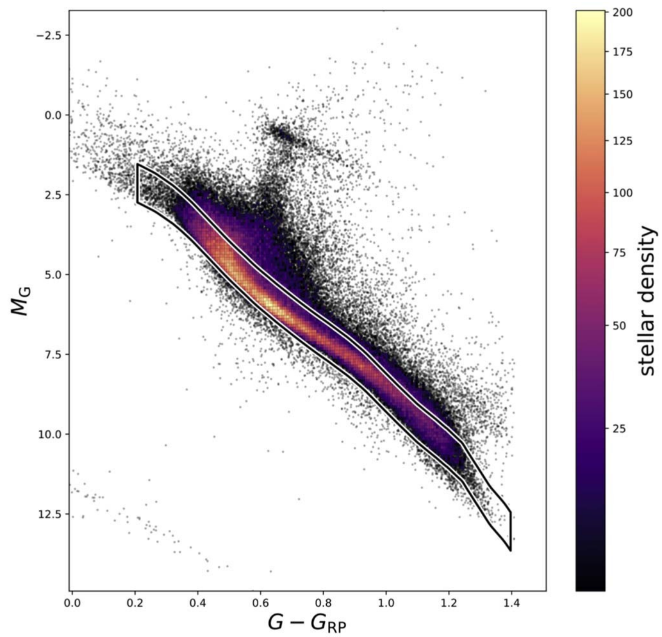

Figure 3. K2 stars on the Gaia color–magnitude diagram. The boxed region shows the area selected as the main sequence in order to exclude, e.g., evolved stars and unresolved binaries from our final sample. The main sequence is defined by a MIST isochrone with an age of 10^9 yr and [Fe/H] = +0.5. We identify 123,079 stars belonging to the main sequence, of which 8943 meet the requirements for our final sample.

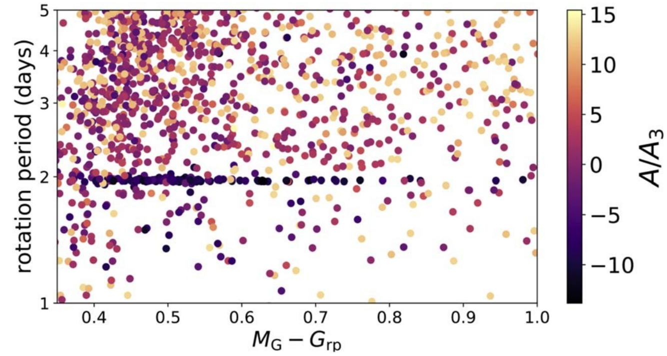

Figure 4. A segment of the period–color diagram showing the presumed artificial pileup at a period of two days. Points are colored by A/A_3, the amplitude of the periodic component of the GP relative to the aperiodic component. Stars in the pileup are notable for having a very small value for this ratio relative to the rest of the stars in the sample, allowing us to effectively remove this feature by imposing a cutoff in A/A_3 for our final sample.

range of stellar masses but is extremely localized in period space and therefore appears to be artificial, though its origin is not known. Figure 5 shows our final sample including the main-sequence cuts described in Section 2.3 in blue compared to the full main-sequence sample in the P/σ_P versus $\log_{10}(A/A_3)$ plane.

We have chosen these specific conditions with the goal of being conservative about selecting only the highest-quality period measurements in order to highlight structure in the color–period diagram. As a result, many periodic signals are excluded from our final sample because they do not meet the condition $P/\sigma_P > 15$. A less selective criteria could be used to obtain a much larger sample size for applications that do not require such high precision, such as an analysis of the rotation period–metallicity relation reported by Amard et al. (2020). The exclusion of signals on the basis of P/σ_P affects the completeness of our sample most at long periods. In Figure 6 this means that the upper envelope of the rotation period–color distribution is not as well defined as it might appear in the plot, especially toward the fainter K and M dwarfs toward the right-hand side of the plot.

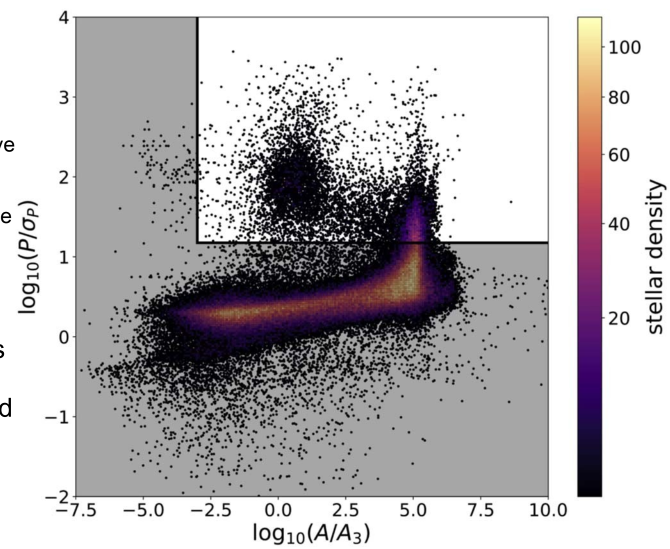

Figure 5. Main-sequence stars plotted in the P/σ_P vs. $\log(A/A_3)$ plane. The full main-sequence sample is shown with the grayed-out area representing the region excluded by the cuts detailed in Section 2.4. Our final sample consists of stars in the highlighted region.

3. Features in Period–Color Space

Figure 6 shows the distribution of rotation periods in period–color space for our final sample. There are several prominent features in this space, among them the aforementioned gap, the upper edge of the envelope of rotation periods, the lower edge of the same envelope, and the overdensity of M dwarfs with short rotation periods in the lower-right-hand corner. In this section we focus on the gap and the overdensity of fast-rotating M dwarfs. The upper edge of the envelope is not well measured in our sample as it is largely determined by the exclusion of rotation periods longer than 32 days, as well as by our cutoff in the precision of the rotation period measurement, which preferentially excludes slow rotators.

3.1. The Period Gap

The gap in rotation periods extends from ~15 days at $G - G_{RP} = 0.75$ to ~25 days at $G - G_{RP} = 1.1$. This gap has been extensively studied in the Kepler data, first by McQuillan et al. (2013) for M dwarfs in the Kepler sample. In McQuillan et al. (2014) the same feature was found in Kepler K dwarfs. Davenport (2017) identified the gap for the G dwarfs, and Davenport & Covey (2018) showed that the gap is present in stellar populations out to 525 pc, beyond which the feature can no longer be recovered due to the difficulty in recovering rotation periods at large distances. Besides the tentative detection by Reinhold & Hekker (2020), our work represents the first robust measurement of this feature outside of the Kepler data. In contrast to the Kepler data, the gap detected here appears wider and has more sharply defined boundaries than is seen in the Kepler sample.

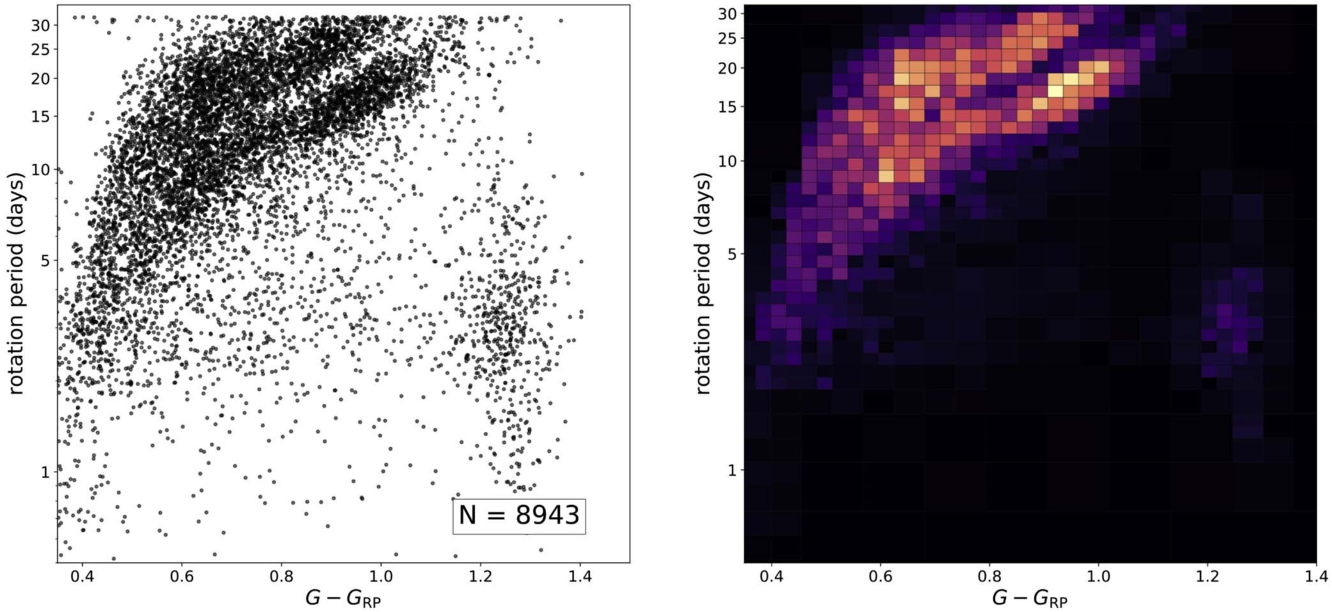

Figure 6. Inferred rotation periods for the 8943 main-sequence K2 stars, plotted against Gaia $G - G_{RP}$ color. Left: scatter plot showing the measured periods. Right: the same rotation periods presented as a two-dimensional histogram in order to highlight the variations in the density of stars across period–color space.

The robustness of our detection of this gap allows us to constrain new details of the feature. We use an edge-detection procedure based on the Canny edge-detection algorithm (Canny 1986) to find the edges of the period gap. The position of the gap edges at select values of $G-G_{RP}$ are given in Table 2. We then fit a parametric model to the gap edges as a function of color. We measure the locations of the gap edge in each campaign individually in order to verify the presence of the gap in each campaign. We confirm that the gap properties do not appear to depend on the direction of the K2 pointing.

We begin by applying our edge-detection algorithm to the rotation period–color distribution for all campaigns combined. Our edge-detection algorithm is based on the Canny edge-detection algorithm, which operates on a two-dimensional array. The Canny algorithm first applies a Sobel operator to the image, which produces an approximation to the gradient. It then identifies local maxima and minima of the approximated gradient, which correspond to edges (points where the intensity of the image is changing most quickly) in the original image. Our modification replaces the first step of applying the Sobel operator with the computation of a kernel gradient estimator, allowing us to apply the algorithm directly to the distribution of stars in rotation period–color space, which is not possible to do directly with the Canny algorithm as it requires a two-dimensional image grid rather than a set of points in the plane. The kernel gradient estimator is defined to be the gradient of the kernel density estimator. The kernel gradient estimator is defined

$$\nabla \hat{f}(x, y) = \frac{1}{nh} \sum_{i}^{n} \nabla K_i(xh^{-1}, yh^{-1}), \tag{11}$$

which is the gradient of the more widely known kernel density estimator. K_i is the kernel function taken here to be a Gaussian centered at the coordinates of the ith data point, n is the number of data points used to make the estimate, h is the width of the kernel, which we set to 0.04, and (x, y) are the coordinate at which the kernel estimate is computed. We then apply the second part of the Canny algorithm as implemented in scikit-image (van der Walt et al. 2014) to identify local maxima and minima of the kernel gradient estimate. We take these local extrema to be the edges of the distribution. Figure 7 shows the output of this algorithm applied to our sample.

We parameterize the gap edges using a function of the form

$$P_{\mathrm{upper}} = A(G - G_{\mathrm{RP}} - x_0) + B(G - G_{\mathrm{RP}} - x_0)^{1/2} \tag{12}$$

using the edges identified in the slice of color space given by

$$0.8 < G - G_{\mathrm{RP}} < 1.05. \tag{13}$$

which corresponds to the stellar mass range

$$0.57 M_{\odot} < M < 0.76 M_{\odot}. \tag{14}$$

Equation (12) is taken from the gyrochronology model of Barnes (2003). Our decision to fit the gap edges with this equation is motivated by the observation that the gap edges appear to have a similar trend to the gyrochrones from that work, but this choice is not meant to imply that the gap edges occur at constant age. The best-fit parameters are given in Table 1, and Figure 7 shows the best-fit models in period–color space.

We have estimated the locations of the gap edges in each campaign individually in terms of their offset from the best-fit model for the full sample. Because the sample sizes are small for some campaigns, we collapse the problem to one dimension rather than considering the full two-dimensional period–color diagram. For each edge of the gap (upper and lower) in each campaign, we first subtract off the gap trend and then sum over color in the range of colors for which the best-fit edge model is valid:

$$0.8 < G - G_{\mathrm{RP}} < 1.05. \tag{15}$$

This gives us the one-dimensional period distribution in terms of the offset from the edge locations defined in Table 1. Treating each gap edge separately, we apply the Gaussian kernel derivative estimator, which is the one-dimensional analog of Equation (11), to the period distribution. We then

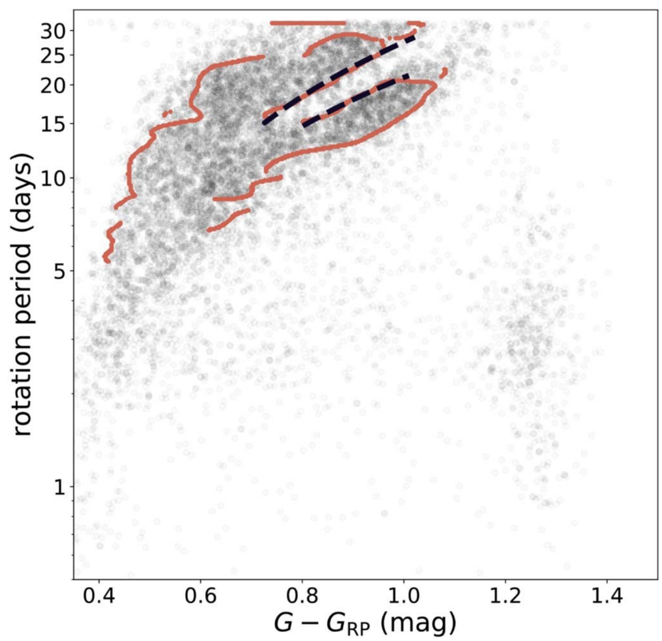

Figure 7. Detected edges of the rotation period diagram using a modified version of the Canny edge-detection algorithm. Best-fit models to the gap edges are shown with blue dashed lines. The model used to fit the edges is given in Equation (12) and the best-fit parameters are in Table 1.

Table 1
Best-fit Parameters from Equation (12)

	A (days)	B (days)	x_0
Upper edge	68.2277	−43.7301	−0.0653
Lower edge	34.0405	−2.6183	0.3150

Table 2
Measured Gap Edges and Widths

$G - G_{RP}$ (mag)	P_{lower} (days)	P_{upper} (days)	Gap Width (days)
0.80	15.20	17.97	2.771
0.85	15.92	19.90	3.98
0.90	17.97	22.26	4.29
0.95	19.54	24.89	5.35
1.00	20.66	28.61	7.95

identify the local maximum (in the case of the upper edge) or minimum (in the case of the lower edge) nearest to the zero-offset point and take this to be an estimate of the location of the edge relative to the edges in Figure 7. Figure 8 illustrates this procedure for Campaign 8.

Figure 9 shows the superimposed kernel density estimates for each campaign, and Figure 10 shows the lower edge locations plotted against the upper edge locations for 16 of the 18 campaigns. We have excluded Campaigns 0, 2, and 11, which have too few stars to make an accurate determination of the edge locations. The error bars are determined by bootstrap resampling from the full sample. We do not observe an obvious correlation between the upper and lower gap edges, which would indicate a shifting of the gap toward longer or shorter periods for some campaigns. There are a few outlier campaigns, with Campaign 18 being the most significant. Campaigns 4, 8, 10, 13, and 15 also deviate noticeably from at least one of the measured gap edges for the all-campaign sample. In Figure 11 we show the period–color diagrams for these outliers. We note that Campaigns 5 and 18 observed the Praesepe cluster, which imprints a visible sequence of stars corresponding to a 600–700 Myr gyrochrone onto the diagram. This does not appear to influence the gap measurement for Campaign 5, but for Campaign 18 the Praesepe cluster is likely responsible for the displacement of the lower gap edge from the expected position. For the rest of the outliers we do not see any obvious evidence of a systematic displacement in the gap edges from the locations derived from the all-campaign data. In general we do not find the outliers to be significant, and we conclude that the rotation period gap shows no dependence on the direction of the K2 pointing.

To further illustrate this point, Figure 12 shows the locations of the K2 footprints for these outlying campaigns on the sky relative to the rest of the campaigns. We note that the outlier campaigns are in general widely separated on the sky, appear both above and below the galactic plane, and show no evidence of clumping. We therefore state our conclusion that the rotation period gap appears to be an isotropic feature of the stellar populations in the nearby Milky Way.

3.2. Other Features in the Rotation Period–Color Diagram

In addition to the prominent period gap, a major feature of the period–color distribution seen here is the overdensity in the bottom-left corner, which represents a population of fast-rotating M dwarfs. This population has been studied in young open clusters (Rebull et al. 2016, 2018) as well as in the MEarth sample (Newton et al. 2016). The break between the slow-rotating M dwarfs for which the rotation period increases with decreasing mass, and the fast-rotating sequence for which the rotation period decreases with decreasing mass occurs at approximately the mass at which M dwarfs become fully convective and corresponds to a change in the morphology of the surface magnetic field from a more complex toward a more simple configuration (Morin et al. 2010; Garraffo et al. 2018).

We find that the light curves of these fast-rotating M dwarfs show rotational modulation that is more periodic than those of the other stars in our sample. For our GP rotation model, the degree to which a light curve shows periodic variations is measured by the parameters Q_1 and Q_2 in Equation (7). Larger values of Q_1 and Q_2 mean that the power spectrum of the variability is more sharply peaked about ω_1 and ω_2. In this analysis we consider the maximum of (Q_1, Q_2), which we call Q_{max}. Figure 13 demonstrates how this parameter affects the appearance of the light curve. We interpret a larger Q_{max} value to indicate that features on the star's surface are stable over a longer period of time, giving rise to variations that are coherent across many periods.

The right panel of Figure 14 shows how Q_{max} varies across the period–color diagram. From this figure we see that stars with large Q_{max}, indicating stronger periodicity, cluster in the fast-rotating M dwarfs. These stars are, however, not limited to this cluster and occur in lower densities across the full range of $G - G_{RP}$ at short rotation periods. The left panel of Figure 14 shows our sample in period–Q_{max} space with stars colored by their $G - G_{RP}$ color. In this space we observe a distinct population of strongly periodic rotators with a negative correlation between rotation period and Q_{max}.

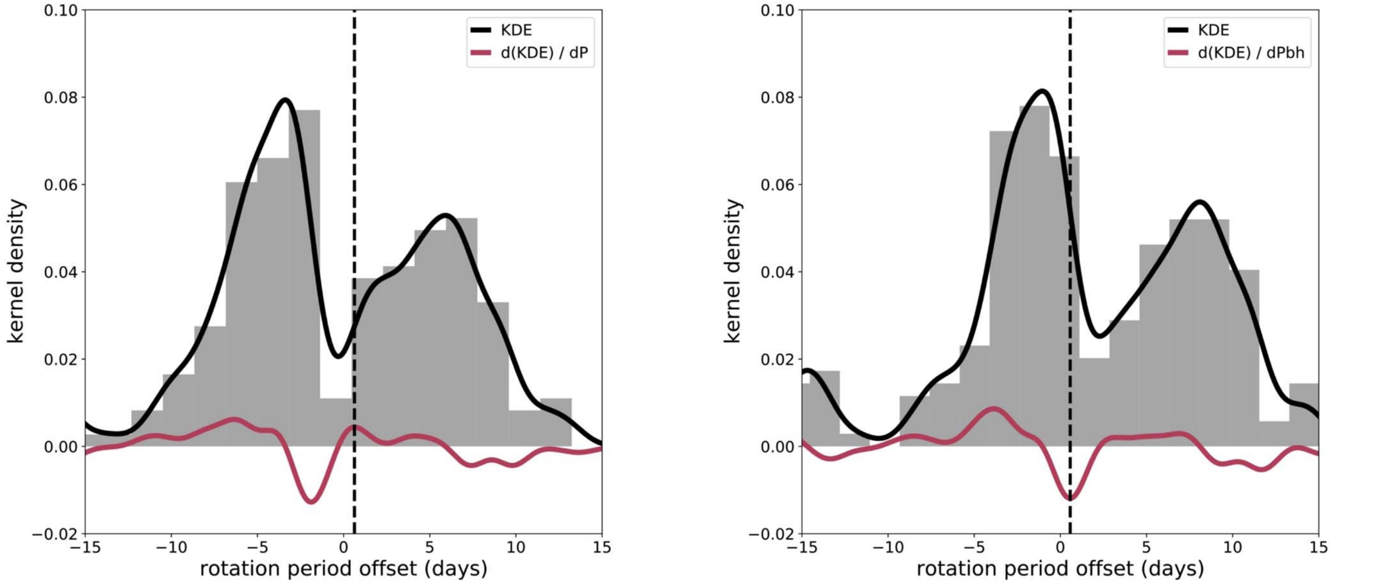

Figure 8. Histograms showing the distribution of rotation periods for stars in Campaign 8, in the color grange $0.8 < G - G_{RP} < 1.05$. The kernel density estimate is shown in black, and its derivative is shown in red. The locations determined for the gap edges are shown by the dashed vertical line. The left panel shows the uppe gap edge and the periods are given as the difference between the observed rotation period and the trend of the upper gap edge. The right panel shows the same f lower gap edge.

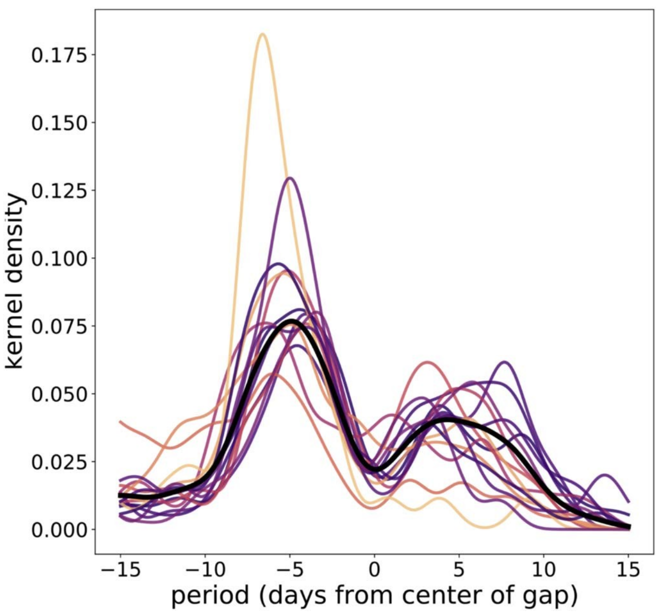

Figure 9. Kernel density estimates for the 16 campaigns with N > 200 for stars with $G - G_{RP}$ in the range defined in Equation (15). The thick black curve is the kernel density estimate for all 16 campaigns combined.

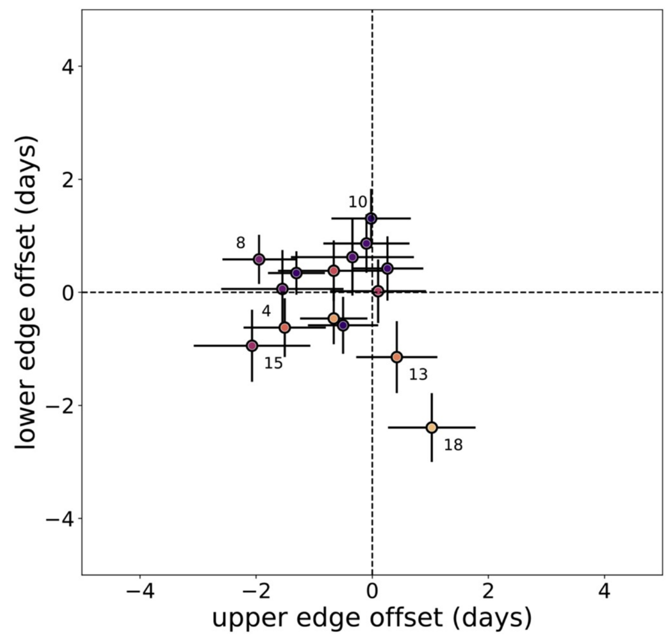

Figure 10. Location of the lower edge of the gap plotted against the location of the upper edge. Both gap edge locations are given as the displacement in days from the best-fit gap edges in Figure 7 and are determined from the slope of the one-dimensional kernel density estimates in Figure 9. The four most notable outliers are labeled and their period–color diagrams are shown in Figure 11.

4. Discussion

There has been much interest in, and discussion of, the origin of the rotation period gap, with several promising possible explanations having been put forward since its discovery (Reinhold et al. 2013; Davenport & Covey 2018; Reinhold et al. 2019; Angus et al. 2020). We now consider these potential explanations in light of our new measurements, as well as taking into account recent work on the 2.7 Gyr cluster Ruprecht 147 by Curtis et al. (2020) and Gruner & Barnes (2020), which crosses the rotation period gap.

McQuillan et al. (2014) and Davenport & Covey (2018) propose that the gap may be an artifact of a recent (<500 Myr) burst of star formation in either the solar neighborhood or in the direction of the Kepler field, which would have produced a population of young, fast-rotating stars that make up the lower branch of the observed bimodality. The single pointing of the Kepler mission admitted the possibility that this feature is confined to that field. Our sample has the benefit of K2's multiple pointings, which has allowed us to demonstrate that the bimodality is present in all directions and is therefore not unique to the stellar population observed by Kepler. The possibility remains that the bimodal star formation history

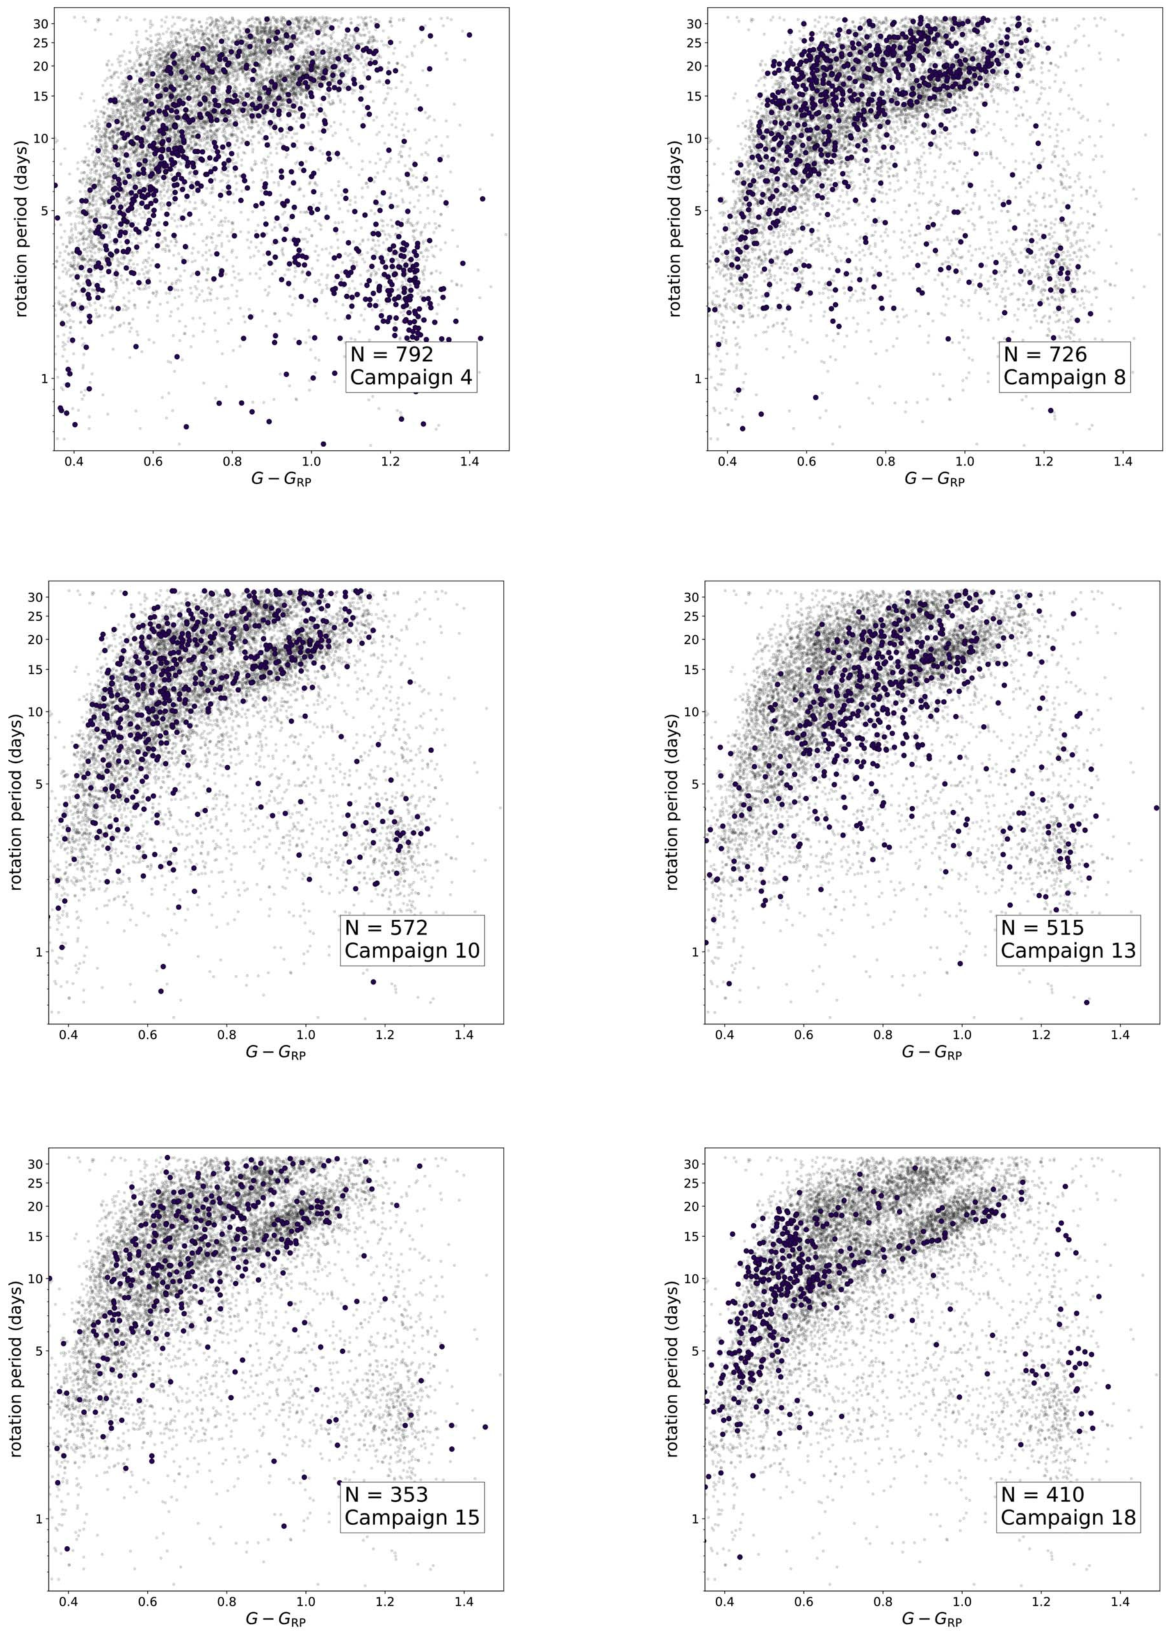

Figure 11. Period–color plots for the outliers identified in Figure 10. The dark purple points show the stars from the individual campaign, while the gray points are th full sample of 8943 stars. The locations of the K2 footprints on the sky for these campaigns is shown in Figure 12. The Praesepe sequence can be seen in Campai but it only appears to impact the gap edge detection for Campaign 18, as seen in Figure 10. With the possible exception of Campaign 18, the gap appears to be respected by the subsamples for each campaign, which indicates to us that the outliers in Figure 10 are the result of stochastic variations within the sample and ar significant.

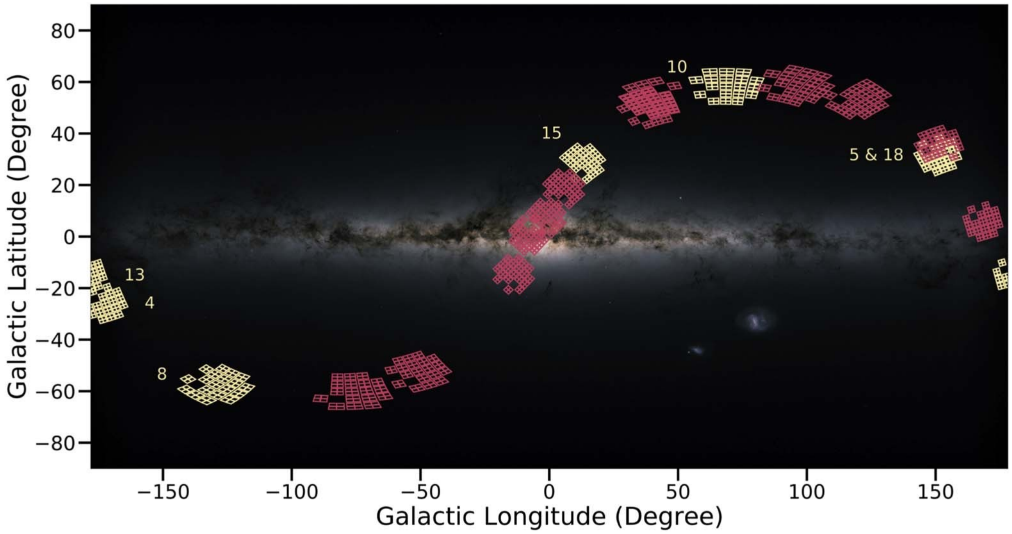

Figure 12. Positions of the campaigns shown in Figure 11 with Milky Way as seen by Gaia DR2 for reference. There is no obvious correlation between the direction of the K2 pointing and the change in shape or position of the gap for the outlier campaigns. Background image credit: ESA/Gaia/DPAC.

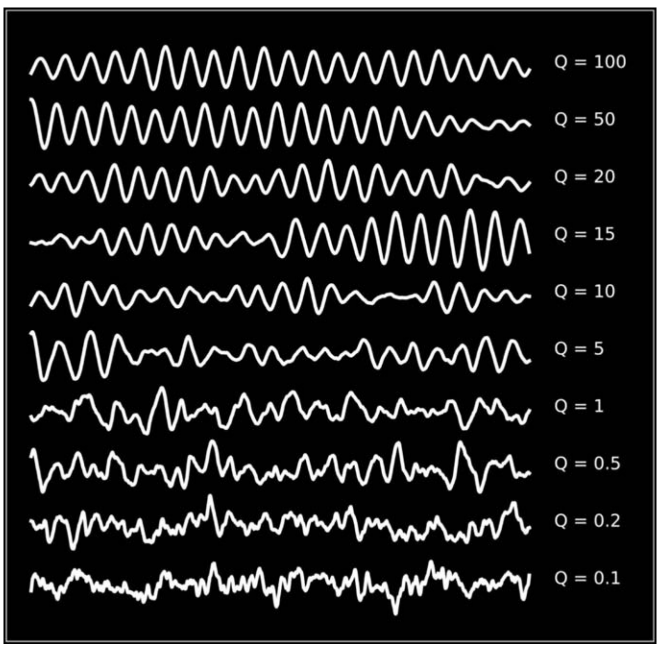

Figure 13. Samples drawn from our GP model showing the effect of increasing the quality factor, Q, on a light curve. All light curves have the same period and amplitude. We have set $\Delta Q = 0$ for these simulations so that $Q = Q_{max}$. A higher Q value means that the light curve shows stronger periodicity. In terms of stellar rotation, this likely indicates that surface features are stable for a longer period of time when Q is large.

suggested by McQuillan et al. (2014) and Davenport & Covey (2018) might be isotropic. However, the position and shape of the gap revealed by our sample make this explanation untenable, as the trend of the gap shows a sharper slope than the sequences associated with constant age populations from Praesepe and NGC 6811 (e.g., Curtis et al. 2019). It is interesting to note that the sequence of stars associated with the 2.7 Gyr cluster Ruprecht 147 appears to cross the gap around $G - G_{RP} \sim 0.7$. While caution should be exercised due to the fact that this sequence has a relatively smaller number of stars than for the younger clusters and shows a large intrinsic scatter, this apparent crossing of the gap lends further evidence against the hypothesis that the gap represents a feature at constant age. Indeed, as Curtis et al. (2020) note in their analysis of periods for Ruprecht 147 and other clusters, this gap crossing seems to occur at a roughly fixed Rossby number, rather than at a single fixed age, which agrees with our assessment.

Reinhold et al. (2013, 2019) suggest that the gap is an artifact of the transition from spot-dominated to faculae-dominated photospheres as stars age. In this explanation, the gap would result from a minimum in the detectability of rotation periods for stars at the point in this transition where neither spot- nor facula-induced variability is able to dominate the light curves of these stars. Our measurements of the gap do not rule out this explanation. To do this would require more work on the evolution of stellar activity over a range of ages and spectral types.

Our preferred explanation is that the gap emerges from a period of accelerated spin-down immediately after the stalled spin-down noted by Curtis et al. (2020). This explanation was first put forth by McQuillan et al. (2013), but the hypothesis was dismissed in favor of the "two populations" hypothesis preferred by McQuillan et al. (2014) and Davenport & Covey (2018).

In this scenario, a young star with its envelope initially decoupled from its core would experience magnetic braking, reducing the spin of the envelope while the decoupled core would be allowed to continue its faster rotation. At a later time the core and envelope would begin to exchange angular momentum. At this point the transfer of angular momentum from the core to the envelope would slow or even halt spin-down by offsetting magnetic braking at the surface, resulting in an overdensity of stars just below the period gap. The underdensity making up the gap itself could then be explained by a period of increased spin-down once this coupling is complete and before the star resumes ordinary Skumanich spin-down. This could be due to a temporary increase in magnetic activity. Lanzafame & Spada (2015) and Spada & Lanzafame (2020) have developed a spin-down model featuring a mass-dependent core–envelope coupling timescale that reproduces the stalling behavior and has been applied to observations of open clusters. Curtis et al. (2020) found the stall in spin-down corresponds to a track of roughly constant Rossby number.

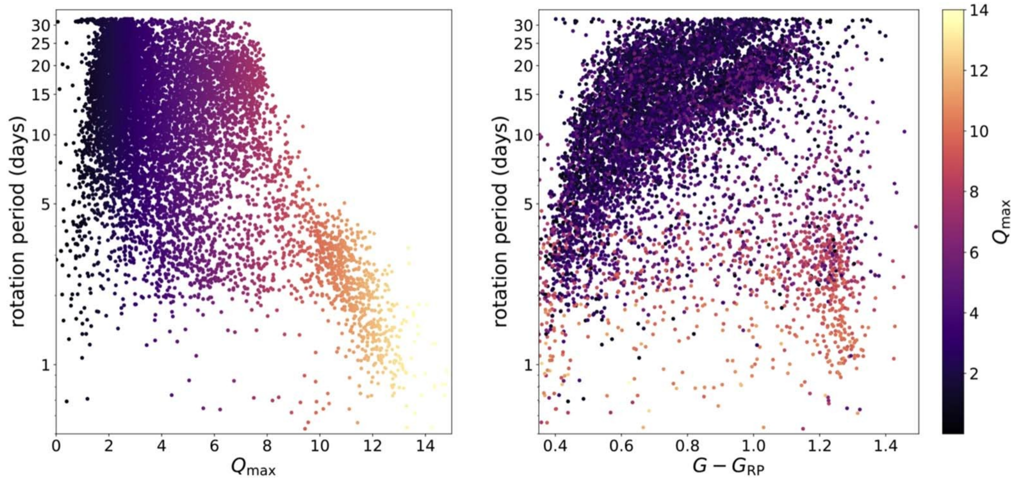

Figure 14. Left: sample in period–Q_{max} space, showing that sinusoidal rotators (high-Q stars) cluster separately from the main population and preferentially occur at short rotation periods. Right: sample in period–color space with stars colored by the maximum quality factor Q_{max}. While stars with higher Q_{max} values cluster in the fast-rotating M dwarfs, they also occur across all colors, and hence across all stellar masses in our sample.

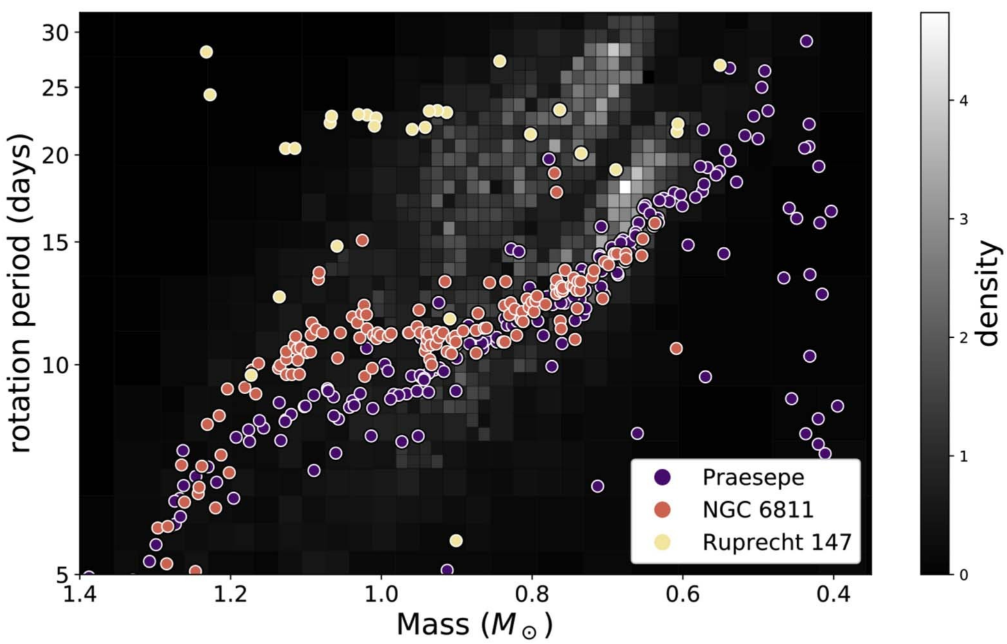

Figure 15. Clusters Praesepe, NGC 6811, and Ruprecht 147 superimposed on the distribution of field stars. For Praesepe, we use our own rotation measurements membership in the cluster taken from Douglas et al. (2019). Rotation periods for NGC 6811 are from Curtis et al. (2019) and for Ruprecht 147 from Curtis et al. (2020).

Angus et al. (2020) suggest that this mechanism may explain the period gap as a break between a "young" regime in which rotation periods increase with decreasing mass from an "old" regime in which rotation periods are nearly constant or even decreasing with decreasing mass, with the gap representing a period of relatively fast spin evolution during the transition between these regimes. In Figure 15 we plot rotation measurements for several important clusters over the distribution of K2 field stars to show the correspondence between the location of the gap in the field stars and the apparent stalling of spin-down in clusters. We use the open clusters Praesepe, which has an age of 600–700 Myrs, NGC 6811 which has an age of approximately 1 Gyr (Agüeros et al. 2018; Curtis et al. 2019), and Ruprecht 147, which is older than both at 2.7 Gyrs. The sequences for Praesepe and NGC 6811 sit on top of each other for low-mass stars but have diverged for stars more massive than about 0.9 $M_\odot$, suggesting that lower-mass stars have stalled in their angular momentum loss while higher-mass stars have continued to spin down. By the time we reach the age of Ruprecht 147, spin-down has resumed for stars down to about 0.7 $M_\odot$. The location of the gap in the rotation measurements for the K2 field stars coincides with the point at which the clusters transition from stalled spin-down at low masses to resumed Skumanich spin-down at higher masses. This supports the notion that the gap represents a discontinuity between these two regimes of spin-down.

There is still much work to be done to determine whether core–envelope coupling and decoupling fully explain these observations. On the observational side it will be important to continue to benchmark clusters of ages between 1 Gyr and solar age, as clusters in this age range may cross the gap (similar to Ruprecht 147). On the theoretical side, models of rotational evolution that can explain the period of rapid spin-down after the epoch of stalled spin-down and in doing so reproduce the shape and trend of the gap will be important for testing this explanation. Another promising avenue of investigation may be the kinematic dating of field stars (Angus et al. 2020). Proper motion measurements from Gaia may provide us with the ability to estimate ages of stellar populations by the vertical component of their motion with respect to the galactic disk because stars become excited in this direction by dynamical interactions over time. This would allow for independent calibration of gyrochronological relations, which may shed light on how stars evolve across the gap.

Finally, while the population of stars within the gap is small, it appears to be nonzero, which opens up the possibility of targeted studies of stars that are currently crossing the gap. Detailed observations of individual stars in the gap or near the lower boundary of the gap may reveal interesting aspects of their activity and the processes that shepherd them across this span of the color–period diagram.

5. Conclusions

We have measured precise rotation periods for 8943 main-sequence K2 stars by GP regression. We perform MCMC simulations on each light curve to obtain estimates of the GP hyperparameters and their uncertainties. We detect and measure the gap in the rotation period distribution and show that this feature appears in all K2 campaigns and is thus unlikely to result from a peculiarity of the stellar populations observed by Kepler. We review several explanations for the gap and argue that the most likely is that the gap results from stalled spin-down on the fast-rotating sequence for low-mass stars, followed by rapid evolution across the gap to the slow-rotating sequence. This evolution may be governed by time-variable core–envelope coupling, which controls the rate of transfer of angular momentum from the core to the surface of the star.

In the future, TESS observations will provide a large sample of light curves for field stars. We expect that a similar distribution of rotation periods will be observed for this sample. One key observation that TESS may enable is whether or not the gap extends to stars more massive than $\sim 0.8\,M_\odot$. If the gap represents the space between two separate stellar populations at different ages, then it should extend to higher mass stars, but if the gap emerges from the physics of core–envelope coupling, then we may expect to observe a mass dependence for the phenomenon.

Finally, one dimension that has been left out of this work is that of metallicity. Amard et al. (2020) report a metallicity dependence on stellar rotation in Kepler, which may be detectable in our K2 sample as well. As the inner structure and the evolution of a star are known to be dependent on its chemical composition, this dependence of rotation period on metallicity may help to illuminate the relationship between interior structure and spin-down. We leave the task of exploring this relationship to future work.

By making the full results of our MCMC simulations available to the community, we hope to make it possible for other researchers to make different choices about which periods to include and exclude. Machine-learning techniques such as convolutional neural networks or random forests may also be useful for identifying rotation signals in EVEREST light curves. Combining these techniques with our period measurements may result in a larger final sample without sacrificing quality. Our sample may also be of use as a training set for machine-learning algorithms seeking to identify stellar rotation signals.

This research was supported by the National Aeronautics and Space Administration (NASA) under grant number 80NSSC18K1660 issued through the NNH17ZDA001N Astrophysics Data Analysis Program (ADAP).

T.A.G. was supported by NSF grant AST-1907342

This work was facilitated through the use of the advanced computational, storage, and networking infrastructure provided by the Hyak supercomputer system at the University of Washington.

J.R.A.D. acknowledges support from the DIRAC Institute in the Department of Astronomy at the University of Washington. The DIRAC Institute is supported through generous gifts from the Charles and Lisa Simonyi Fund for Arts and Sciences, and the Washington Research Foundation.

This paper includes data collected by the Kepler mission and obtained from the MAST data archive at the Space Telescope Science Institute (STScI). Funding for the Kepler mission is provided by the NASA Science Mission Directorate. STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 526555.

This work has made use of data from the European Space Agency (ESA) mission Gaia (https://www.cosmos.esa.int/gaia), processed by the Gaia Data Processing and Analysis Consortium (DPAC, https://www.cosmos.esa.int/web/gaia/dpac/consortium). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement.

Software: Python, IPython (Pérez & Granger 2007), NumPy (Oliphant 2007), Matplotlib (Hunter 2007), SciPy (Virtanen et al. 2020), Pandas (Wes McKinney 2010), Astropy (Astropy Collaboration et al. 2013), Exoplanet (Foreman-Mackey et al. 2019), Scikit-image (van der Walt et al. 2014).

ORCID iDs

Tyler A. Gordon https://orcid.org/0000-0001-5253-1987
James RA. Davenport https://orcid.org/0000-0002-0637-835X
Ruth Angus https://orcid.org/0000-0003-4540-5661
Daniel Foreman-Mackey https://orcid.org/0000-0002-9328-5652
Eric Agol https://orcid.org/0000-0002-0802-9145
Kevin R. Covey https://orcid.org/0000-0001-6914-7797
Marcel A. Agüeros https://orcid.org/0000-0001-7077-3664
David Kipping https://orcid.org/0000-0002-4365-7366

References

Agüeros, M. A., Bowsher, E. C., Bochanski, J. J., et al. 2018, ApJ, 862, 33
Amard, L., Roquette, J., & Matt, S. P. 2020, MNRAS, 499, 3481
Angus, R., Beane, A., Price-Whelan, A. M., et al. 2020, AJ, 160, 90
Angus, R., Morton, T., Aigrain, S., Foreman-Mackey, D., & Rajpaul, V. 2018, MNRAS, 474, 2094
Astropy Collaboration, Price-Whelan, A. M., SipH ocz, B. M., et al. 2018, AJ, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Barclay, T., Endl, M., Huber, D., et al. 2015, ApJ, 800, 46
Barnes, S. A. 2003, ApJ, 586, 464

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Sci, 327, 977
Bouvier, J. 2008, A&A, 489, L53
Canny, J. 1986, ITPAM, PAMI-8, 679
Chakrabarty, A., & Sengupta, S. 2019, AJ, 158, 39
Choi, J., Dotter, A., Conroy, C., et al. 2016, ApJ, 823, 102
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
Davenport, J. R. A., & Covey, K. R. 2018, ApJ, 868, 151
Dawson, R. I., Johnson, J. A., Fabrycky, D. C., et al. 2014, ApJ, 791, 89
Denissenkov, P. A., Pinsonneault, M., Terndrup, D. M., & Newsham, G. 2010, ApJ, 716, 1269
Dotter, A. 2016, ApJS, 222, 8
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
Durney, B. 1972, in Evidence for Changes in the Angular Velocity of the Surface Regions of the Sun and Stars - Comments, ed. C. P. Sonett, P. J. Coleman, & J. M. Wilcox, Vol. 308 (Washington, DC: NASA), 282
Endal, A. S., & Sofia, S. 1981, ApJ, 243, 625
Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, celerite: Scalable 1D Gaussian Processes in C++, Python, and Julia, Astrophysics Source Code Library, ascl:1709.008
Foreman-Mackey, D., Czekala, I., Luger, R., et al. 2019, dfm/exoplanet v0.2.3, doi:10.5281/zenodo.1998447
Gaia Collaboration, Katz, D., Antoja, T., et al. 2018, A&A, 616, A11
Gallet, F., & Bouvier, J. 2013, A&A, 556, A36
Garraffo, C., Drake, J. J., Dotter, A., et al. 2018, ApJ, 862, 90
Gelman, A., & Rubin, D. B. 1992, StaSc, 7, 457
Gruner, D., & Barnes, S. A. 2020, A&A, 644, A16
Hunter, J. D. 2007, Computing In Science Engineering, 9, 90
Lanzafame, A. C., & Spada, F. 2015, A&A, 584, A30
Luger, R., Kruse, E., Foreman-Mackey, D., Agol, E., & Saunders, N. 2018, AJ, 156, 99
MacGregor, K. B., & Brenner, M. 1991, ApJ, 376, 204
MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, ApJ, 721, 1014
McQuillan, A., Aigrain, S., & Mazeh, T. 2013, MNRAS, 432, 1203
McQuillan, A., Mazeh, T., & Aigrain, S. 2014, ApJS, 211, 24
Meibom, S., Barnes, S. A., Latham, D. W., et al. 2011, ApJL, 733, L9
Morin, J., Donati, J. F., Petit, P., et al. 2010, MNRAS, 407, 2269
Newton, E. R., Irwin, J., Charbonneau, D., et al. 2016, ApJ, 821, 93
Nielsen, M. B., Gizon, L., Schunker, H., & Karoff, C. 2013, A&A, 557, L10
Oliphant, T. E. 2007, CSE, 9, 10
Paxton, B., Bildsten, L., Dotter, A., et al. 2011, ApJS, 192, 3
Paxton, B., Cantiello, M., Arras, P., et al. 2013, ApJS, 208, 4
Paxton, B., Marchant, P., Schwab, J., et al. 2015, ApJS, 220, 15
Pérez, F., & Granger, B. E. 2007, CSE, 9, 21
Rasmussen, C. E., & Williams, C. K. I. 2006, Gaussian Processes for Machine Learning (Cambridge, MA: MIT Press)
Rebull, L. M., Stauffer, J. R., Bouvier, J., et al. 2016, AJ, 152, 113
Rebull, L. M., Stauffer, J. R., Cody, A. M., et al. 2018, AJ, 155, 196
Reinhold, T., Bell, K. J., Kuszlewicz, J., Hekker, S., & Shapiro, A. I. 2019, A&A, 621, A21
Reinhold, T., & Hekker, S. 2020, A&A, 635, A43
Reinhold, T., Reiners, A., & Basri, G. 2013, A&A, 560, A4
Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, PeerJ Computer Science, 2, e55
Skumanich, A. 1972, ApJ, 171, 565
Somers, G., & Pinsonneault, M. H. 2016, ApJ, 829, 32
Spada, F., & Lanzafame, A. C. 2020, A&A, 636, A76
van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., et al. 2014, PeerJ, 2, e453
van Saders, J. L., Pinsonneault, M. H., & Barbieri, M. 2019, ApJ, 872, 128
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nat. Method., 17, 261
Wes McKinney 2010, in Proc. 9th Python in Science Conf (Austin, TX) ed. S. van der Walt & J. Millman, 56