

Rotation Periods for Stars in Open Cluster NGC 6819 From *Kepler* IRIS Light Curves

SABINA SAGYNBAYEVA,¹ ISABEL L. COLMAN,^{2,3} AND WILL M. FARR^{1,4}

¹*Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794, USA*

²*Liberal Studies, New York University, 726 Broadway, New York, NY 10003, USA*

³*Department of Astrophysics, American Museum of Natural History, 200 Central Park West, New York, NY 10024, USA*

⁴*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA*

ABSTRACT

We present an updated catalog of stellar rotation periods for the 2.5 Gyr open cluster NGC 6819 using the *Kepler* IRIS light curves from superstamp data. Our analysis uses Gaussian Process modeling to extract robust rotation signals from image subtraction light curves, allowing us unprecedented data access and measurement precision in the crowded cluster field. After applying stringent quality and contamination cuts, we identify 271 reliable rotation periods, representing by far the largest sample of rotators measured in a single intermediate-age cluster. Compared to previous work, which relied on only ~ 30 stars, our catalog extends the gyrochronological sequence of NGC 6819 with an order of magnitude more measurements and improved precision. The expanded dataset reveals both the expected temperature-dependent spin-down trend and substantial scatter at fixed effective temperature, including a bimodal distribution of fast and slow rotators. We also identify a distinct “pile-up” sequence consistent with predictions of weakened magnetic braking at critical Rossby numbers. These results strengthen this cluster’s role as a benchmark for stellar spin evolution, while also highlighting the limitations of traditional gyrochronology at older ages. The final catalog and the model implementations are all available on Zenodo^{a)}.

Keywords: Stellar rotation (1629) — Period determination (1211) — Astronomy data analysis (1858)

1. INTRODUCTION

Accurate stellar age determination remains one of the most significant challenges in stellar astrophysics. Gyrochronology, the connection between stellar surface rotation and age (Barnes 2003, 2007; Mamajek & Hillenbrand 2008), provides a viable method for age estimation along the main sequence. The fundamental assumption underlying gyrochronology — that stellar rotation periods can serve as reliable age indicators — is increasingly called into question by observational evidence that stellar spin evolution is far more complex than previously thought. Traditional gyrochronology models assume that after an initial convergence phase around 500–700 Myr, stars of similar mass and metallicity will follow predictable spin-down trajectories governed by magnetic braking, making the rotation period a monotonic func-

tion of age for single main-sequence stars (Barnes 2003; Mamajek & Hillenbrand 2008).

However, multiple lines of observational evidence now challenge this paradigm. Angus et al. (2020) used velocity dispersions as kinematic age proxies to demonstrate that this framework breaks down for stars older than ~ 1 Gyr, showing that stars with identical rotation periods and colors can have vastly different ages. This suggests that some physical processes, like variable core-envelope angular momentum transport, or changes in magnetic field topology, can fundamentally alter rotational evolution in ways that current gyrochronology calibrations cannot capture. Supporting this conclusion, van Saders et al. (2016) found evidence for weakened magnetic braking in stars with Rossby numbers greater than ~ 2 , approximately corresponding to the Sun, which would cause a dramatic slowdown in rotational evolution for older, slowly rotating stars.

Understanding the pre-main-sequence context is also crucial for interpreting the main-sequence challenges. Matt et al. (2010) demonstrated that magnetic star-

Corresponding author: Sabina Sagynbayeva
sabina.sagynbayeva@stonybrook.edu

^{a)} see <https://doi.org/10.5281/zenodo.17247991>

disk coupling plays a critical role in setting initial main-sequence rotation rates, but the effectiveness of this “disk locking” mechanism depends sensitively on uncertain parameters, including magnetic field strength and the degree of magnetic coupling to the disk. When [Matt et al. \(2010\)](#) include field-line opening due to differential rotation between star and disk (rather than assuming a fully closed field), they find that magnetic torques alone cannot explain young stars observed with surface rotation periods of 1–10 days. This suggests that additional angular momentum loss mechanisms — particularly stellar winds — may be necessary even during the disk accretion phase, which would further complicate the initial conditions for main-sequence gyrochronology. Further, [Matt et al. \(2015\)](#) demonstrated that the mass-dependence of angular momentum evolution in Sun-like stars, with lower-mass stars spinning down faster than solar-mass stars once in the magnetically unsaturated regime, creating the observed gyrochronology sequence, can only be reproduced by incorporating empirically-derived torque scalings that depend on both stellar mass and the Rossby number. This highlights how sensitively rotational evolution depends on magnetic field properties and wind mass-loss rates. Open clusters, where all members are assumed to be co-eval, provide a laboratory to search for signatures of these underlying physics, and build on the past two decades of work aimed at characterizing the angular momentum evolution of Sun-like stars.

Thanks to high-quality data from NASA’s *Kepler* mission ([Borucki et al. 2010](#)), NGC 6819 serves as one of the key gyrochronology calibrator clusters (e.g., [Bouma et al. 2023](#); [Curtis et al. 2020](#)). With an age of 2.5 Gyr ([Hole et al. 2009](#); [Anthony-Twarog et al. 2014](#)) and solar metallicity ($Z = 0.09 \pm 0.03$; [Lee-Brown et al. 2015](#)), this cluster has so far provided crucial constraints for gyrochronological relations in the intermediate age regime. Currently, the stellar rotation periods derived by [Meibom et al. \(2015\)](#) serve as the standard reference for this cluster. While the establishment of the NGC 6819 gyrochronology sequence represents an important milestone, the sample comprises only approximately 30 stars, and now, with the availability of more light curves for this cluster ([Colman et al. 2022](#)), we have the ability to conduct a more complete investigation and strengthen gyrochronological calibrations relying on results from this cluster. We seek to independently verify or refine these detections of surface rotation, as well as to expand the sample size through additional detections.

During the *Kepler* mission, the extreme stellar crowding characteristic of cluster environments limited standard pipeline processing to postage stamp cutouts

for only a small number of bright, isolated stars in NGC 6819, many of which were not confirmed cluster members. To address this limitation, NGC 6819, along with NGC 6791, was targeted with “superstamps” — 200 square pixel images centered on the cluster core that were downloaded at *Kepler*’s long cadence ([Jenkins et al. 2010](#)).

Extracting stellar rotation periods from *Kepler* photometry presents unique data analysis challenges. Standard *Kepler* pipeline light curves cannot be used directly for rotation studies, as the data processing was optimized for exoplanet detection and likely removes or suppresses the photometric variability associated with stellar rotation ([Smith et al. 2012](#); [Reinhold et al. 2013](#)). This limitation does not affect superstamp observations, since custom photometry has to be performed regardless. This work uses the *Kepler* IRIS light curves ([Colman et al. 2022](#)), which represent a comprehensive photometric analysis of the NGC 6819 superstamp data. Several previous studies have successfully leveraged *Kepler* superstamp observations for cluster studies (e.g., [Covelo-Paz et al. 2023](#); [Kuehn et al. 2015](#); [Sanjayan et al. 2022](#)). The IRIS light curve extraction method is described in the data release paper, and our data reduction methodology is detailed in Section 2.

Kepler’s 90-day quarter structure provides a sufficient temporal baseline for rotation period detection. Experience with TESS observations demonstrates that data downlink gaps present a more significant obstacle than sector duration limitations ([Colman et al. 2022](#); [Holcomb et al. 2022](#)), indicating that even a single *Kepler* quarter can theoretically accommodate rotation periods up to 45 days without requiring complete phase coverage. While longer periods could be accessed through careful stitching of multiple quarters, this approach is unnecessary for the present study. Stellar rotation theory predicts that main-sequence stars in a 2.5 Gyr cluster should exhibit rotation periods no longer than approximately 30 days ([Meibom et al. 2011](#)), well within *Kepler*’s detection capabilities. However, the 4-year *Kepler* baseline is sensitive to magnetic activity cycles ([Reinhold et al. 2017](#)), which can manifest as periodic signatures in highly active stars as starspot distributions migrate in latitude (e.g., [Reinhold & Gizon 2015](#); [Morris et al. 2017](#); [Sagynbayeva et al. 2025](#)). Our approach in this work is designed specifically to detect rotational modulation (see Section 3), and our sample is unlikely to be polluted by detections of activity cycles, as these tend to be on the scale of months to years long for FGK dwarfs (e.g., [Böhm-Vitense 2007](#)).

With this expanded photometric dataset, we revisit NGC 6819 to provide an updated catalog of stellar ro-

tation periods, increasing the sample size from the previous 30 detections (Meibom et al. 2015) to rotation periods. A comprehensive analysis of the rotational properties and their implications for gyrochronology is presented in Section 4. We present these measurements to the community as a refined 2.5 Gyr gyrochronological sequence, providing improved statistical constraints for age-rotation relationships. Our methodology demonstrates robust performance for rotation period detection in *Kepler* photometry and shows promise for application to other high-precision photometric surveys, including TESS (Ricker et al. 2015) and future missions such as PLATO (Rauer et al. 2014). This work also highlights the scientific potential of *Kepler*’s superstamp observations for stellar cluster studies.

Our paper is organized as follows: we present the data in Section 2; we show the details of our models for detrending, period determination, and gyrochrone fitting in Section 3; we provide a detailed overview of the results and a discussion in Section 4; and finally present our conclusions in Section 5.

2. DATA

We used light curves from the *Kepler* IRIS catalog (Colman et al. 2022), which provides data for 9,150 stars in and around the open clusters NGC 6791 and NGC 6819, publicly available as a high-level data product on the Mikulski Archive for Space Telescopes (MAST)¹. These light curves were produced using a novel implementation of image subtraction photometry on *Kepler* “superstamp” data. The catalog includes 8,427 stars that were never targeted during the original *Kepler* mission, plus additional quarters of data for 382 previously targeted stars. The IRIS (Increased Resolution Image Subtraction) method follows the traditional image subtraction approach of treating all non-variable signals as background (Alard & Lupton 1998), which optimizes light curves for variable signal, with the additional step of upsampling *Kepler* pixel data to enable more accurate image alignment before background subtraction. This enables the extraction of high-quality photometry from crowded stellar fields, and is also particularly sensitive to faint targets.

Kepler’s coverage was not continuous throughout its four-year mission due to the failure of Module 3 during Quarter 4, which reduced the number of operational CCD channels from 84 to 80 (Van Cleve et al. 2016). This region included NGC 6819, meaning that all of our targets are missing Quarters 6, 10, and 14. Additional

targets received only partial coverage due to detector positional variation between rotations, which affected a small number of targets towards the edge of the superstamps.

For cluster membership determination, we utilized the astrometric membership probabilities from the WIYN Open Cluster Study catalog of NGC 6819 (Platais et al. 2013), which provides proper motion-based membership analysis for 15,750 stars in the cluster field with accuracies ranging from $\sim 0.2 \text{ mas yr}^{-1}$ in the cluster center to 1.1 mas yr^{-1} in the outer regions. This membership list, combined with the IRIS catalog, gave us 2,200 stars in the cluster, to which we applied our analysis.




Figure 1. An example of some IRIS light curves before detrending.

3. METHODS

3.1. Detrending of *Kepler* Light Curves

The *Kepler* IRIS light curves are provided as raw data for each quarter, or one stitched and detrended time series; however, the initial detrending was optimized for asteroseismology at the expense of longer simple-periodic signals. As part of the *Kepler* mission, official data releases also provided co-trending basis vectors (CBVs) used to perform data reduction (Van Cleve et al. 2016). The use of CBVs is dependent on the location of each target on the detector; since the majority of our targets were not included in official *Kepler* data releases, and in the interests of uniformity, we followed a similar procedure in order to detrend the raw IRIS data. We sought to identify a common set of basis vectors that provides an optimal linear fit to our ensemble of light curves. This approach allows us to construct a design matrix M that captures the dominant systematic trends present across the dataset (see some examples of CBVs in Figure 2).

¹ <https://archive.stsci.edu/hlsp/iris>

For a given design matrix M , we fit each light curve y_i (where $i = 1, \dots, N$) using coefficients x_i by maximizing the likelihood:

$$\log \mathcal{L}_i = -\frac{1}{2} (y_i - Mx_i)^T \Sigma_i^{-1} (y_i - Mx_i) + \text{const}, \quad (1)$$

where Σ_i is the covariance matrix of the noise in lightcurve y_i .

The optimal likelihood is achieved when we project out the component of y_i that lies within the column space of M . This optimal likelihood is given by:

$$\log \hat{\mathcal{L}}_i = -\frac{1}{2} (y_i^\perp)^T \Sigma_i^{-1} (y_i^\perp) + \text{const} \quad (2)$$

where $y_i^\perp$ represents the component of y_i orthogonal to the columns of M . Through optimization of the likelihood with respect to x_i , we obtain:

$$y_i^\perp = \left(I - M (M^T \Sigma_i^{-1} M)^{-1} M^T \Sigma_i^{-1} \right) y_i \quad (3)$$

The general problem of finding an optimal M would require sampling within this model; but here we make the reasonable assumption that each lightcurve is subject to white (i.e. independent, identically-distributed) noise, whose amplitude we estimate empirically in the following way. First, we difference the lightcurve to eliminate any long-timescale systematic or stellar trends; then we compute the variance of the differenced lightcurve, and estimate the variance of the lightcurve white noise by dividing by two. Thus $\Sigma_i = \sigma_i^2 I$, where σ_i^2 is the estimated white noise variance of lightcurve i , and I is the identity matrix.

In this simplified scenario, the optimal likelihood for a given design matrix M is:

$$\begin{aligned} \log \hat{\mathcal{L}}_i &= -\frac{1}{2\sigma_i^2} (y_i^\perp)^T y_i^\perp + \text{const} \\ &= -\frac{1}{2} (\tilde{y}_i^\perp)^T (\tilde{y}_i^\perp) + \text{const} \end{aligned} \quad (4)$$

where $\tilde{y}_i = y_i/\sigma_i$ represents the standardized light curve with unit noise variance. Figure 1 shows some standardized light curves before detrending. The total likelihood across all light curves is obtained by summation:

$$\log \hat{\mathcal{L}} = \sum_i \log \hat{\mathcal{L}}_i = -\frac{1}{2} \sum_i (\tilde{y}_i^\perp)^T (\tilde{y}_i^\perp) + \text{const} \quad (5)$$

To determine the optimal design matrix M that maximizes the likelihood, we employ singular value decomposition (SVD). The objective is to identify the subspace of whitened light curves that contains the maximum squared length for a given number of components

in M . The complementary subspace will minimize the sum-of-squares over any $(N - k)$ -dimensional subspace.

We are thus led to compute the SVD of the whitened light curves and select the first k singular vectors to form the detrending basis that becomes the columns of M . This is similar to the procedure used in the *Kepler* pipeline to compute CBVs (Van Cleve et al. 2016), but that procedure did not whiten before computing singular vectors, and so likely over-weights the noisiest lightcurves in the estimation of CBVs; our procedure is optimal under the assumption that (1) the noise is iid and well-characterized by the estimated variance and (2) instrument systematics enter each lightcurve as a linear combination of common “systematic modes.” Figure 2 shows some examples of the CBVs that result from our procedure for Quarter 2 of observations.

To ensure proper treatment of the mean flux, we zero-mean all whitened light curves before applying SVD. This procedure guarantees that: (1) the constant vector has zero singular value, and (2) the first k singular vectors are orthogonal to the constant vector. We then augment the k singular vectors with the constant vector to construct a $(k + 1)$ -dimensional design matrix that simultaneously fits the mean flux and removes the estimated linear trends.




Figure 2. A subset of 9 different CBVs.

It remains to choose the number of CBVs to incorporate in the detrending; the likelihood above provides no direct information on the optimal number, since the optimal likelihood is achieved by including every basis vector spanning the entire collection of lightcurves. We chose to include additional CBVs until the coefficient of the last-added CBV in the median star reached 5-sigma significance above zero. Thus the last-added CBV meaningfully improves the likelihood of the median star.

This procedure leads to $\mathcal{O}(8)$ CBVs in each quarter. Detrended lightcurves lack the dramatic systematic trends and effects seen in Figure 1, but some systematic effects remain; these will be handled by our low-frequency DRW term in the rotation kernel discussed in the next subsection. We bin the final detrended light curves by the bin-width of 0.4 d for computational efficiency. An example of an analysis-ready light curve appears in Figure 3.

3.2. The Gaussian Process Model for Rotation Periods

We implement a Gaussian Process (GP) model to characterize stellar rotation signatures in multi-quarter photometric time series, first introduced in Foreman-Mackey et al. (2017) as a “rotation kernel.” The model represents the stellar variability as a sum of harmonic oscillations corresponding to the fundamental rotation frequency and its first harmonic. We add to it terms approximating a damped-random-walk (“real term”) at low frequency to account for remaining long-timescale systematics (Foreman-Mackey et al. 2017).

For each observation i at time t_i in quarter q_i , we model the observed flux y_i as:

$$y_i = \mu_{q_i} \cdot (1 + \text{GP}(t_i)) + \epsilon_i \quad (6)$$

where μ_{q_i} is the mean flux level for quarter q_i , $\text{GP}(t_i)$ represents the fractional variability described by a Gaussian Process, and $\epsilon_i \sim \mathcal{N}(0, \sigma_{\text{err},i}^2)$ is the measurement noise.

To account for instrumental variations and systematic effects between observing quarters, we introduce quarter-dependent scaling factors. The mean flux for each quarter is modeled as:

$$\mu_q = \bar{\mu}_q \cdot s_q \quad (7)$$

where $\bar{\mu}_q$ is the empirical mean flux in quarter q and s_q is a scaling factor. The log-scaling factors are assigned hierarchical priors:

$$\log s_q \sim \mathcal{N}\left(0, \left(\frac{\hat{\sigma}_q}{\sqrt{n_q}}\right)^2\right) \quad (8)$$

where $\hat{\sigma}_q$ is the relative standard deviation of flux measurements in quarter q and n_q is the number of observations in that quarter.

The GP kernel consists of two components designed to capture different aspects of stellar variability. The primary component is a rotation term that captures quasi-periodic stellar variability, while a secondary component models longer-timescale stochastic variations that could otherwise interfere with period detection.

The rotation kernel employs a **RotationTerm** from **celerite** that combines two damped harmonic oscillators to represent the fundamental rotation frequency and its first harmonic:

$$k_{\text{rot}}(\tau) = \sigma^2 \left[\frac{1}{2}(1+f) \exp\left(-\frac{|\tau|}{Q_0 P}\right) \cos\left(\frac{2\pi|\tau|}{P}\right) + \frac{1}{2}(1-f) \exp\left(-\frac{|\tau|}{Q_1 P}\right) \cos\left(\frac{4\pi|\tau|}{P}\right) \right] \quad (9)$$

where:

- σ^2 is the variance amplitude
- P is the rotation period
- Q_0 and Q_1 are quality factors controlling the damping of the fundamental and first harmonic
- $f \in [0, 1]$ controls the relative amplitude of the two harmonics
- $\tau = |t_i - t_j|$ is the time lag

To prevent the rotation kernel from fitting spurious long-period variations that could mask the true stellar rotation signal, we include an additional stochastic component modeled as a Simple Harmonic Oscillator (SHO) term operating in the heavily damped regime ($Q = 0.1$). This component effectively acts as a **RealTerm** in the power spectral density, providing a mechanism to absorb red noise and long-timescale trends:

$$k_{\text{red}}(\tau) = \text{SHOTerm}(S_0, \omega_0, Q = 0.1) \quad (10)$$

The parameters of this red noise component are constrained based on the observation characteristics:

$$\tau_{\text{longest}} = \frac{1}{f_0/\sqrt{2}} \quad (11)$$

$$\log \tau_c \sim \text{Uniform}(\log \tau_{\text{longest}}, \log T_{\text{obs}}) \quad (12)$$

$$\omega_0 = \frac{2\pi/\tau_c}{2Q} \quad (13)$$

$$S_0 = \frac{\sigma_{\text{red}}^2}{\omega_0 Q} \quad (14)$$

where τ_c is the characteristic timescale, T_{obs} is the total observation duration, and σ_{red} controls the amplitude of the red noise component. The timescale is constrained to be longer than the expected rotation period but shorter than the total observation baseline, ensuring that this component captures long-term variations without interfering with the periodic signal of interest.

The total kernel is the sum of these components:

$$k(\tau) = k_{\text{rot}}(\tau) + k_{\text{red}}(\tau) \quad (15)$$




Figure 3. An example of a light curve after our detrending. We also binned the light curves for computational efficiency. The left panel shows the full light curve, and the right panel shows the zoomed-in light curve between 200 and 400 days.

The kernel’s power spectral density is computed at specified frequencies ω :

$$S(\omega) = \int_{-\infty}^{\infty} k(\tau) e^{-i\omega\tau} d\tau \quad (16)$$

3.2.1. Priors

The model parameters are assigned the following priors:

$$\log P \sim \mathcal{N}(-\log f_0, \sigma_f^2) \quad (17)$$

where f_0 is the prior guess for the fundamental frequency and σ_f is the fractional uncertainty.

$$\sigma \sim \text{HalfNormal}(0.1 \cdot \text{median}(\hat{\sigma}_q)) \quad (18)$$

$$\Delta Q_0, \Delta Q_1 \sim \text{LogNormal}(\log 5, 1) \quad (19)$$

with $Q_0 = 0.5 + \Delta Q_0 + \Delta Q_1$ and $Q_1 = 0.5 + \Delta Q_1$.

$$f \sim \text{Uniform}(0, 1) \quad (20)$$

For the red noise component:

$$\sigma_{\text{red}} \sim \text{HalfNormal}(0.1 \cdot \text{median}(\hat{\sigma}_q)) \quad (21)$$

The centered and scaled observations are defined as:

$$\tilde{y}_i = \frac{y_i}{\mu_{q_i}} - 1 \quad (22)$$

The GP likelihood is computed for these transformed observations:

$$\log p(\tilde{y}|\theta) = -\frac{1}{2} \tilde{y}^T K^{-1} \tilde{y} - \frac{1}{2} \log |K| - \frac{n}{2} \log(2\pi) \quad (23)$$

where K is the covariance matrix derived from the combined kernel.

A Jacobian correction accounts for the transformation from y to $\tilde{y}$:

$$\log p(y|\theta) = \log p(\tilde{y}|\theta) - \sum_{i=1}^n \log \mu_{q_i} \quad (24)$$

The model has difficulty “transitioning” between multiple modes in period space if the range of allowed rotation periods is too large; it tends to get stuck, for example, with the first harmonic assigned to the fundamental frequency, or with the fundamental assigned to a first harmonic. Transitions away from these local maxima require large changes in the parameters, and therefore have low probability. To address this, we run the model multiple times with different priors on the rotation period, always constraining the range of the fundamental to a factor of two in period space, ensuring that, within that prior range, there is no secondary mode in the likelihood that could trap the sampler.

For each star, we ran the model with five priors on frequency (inverse of period) corresponding to five ranges in rotation period: from 1 to 2 days (model **f0.1.2**), from 2 to 4 days (model **f0.2.4**), from 4 to 8 days (model **f0.4.8**), from 8 to 16 days (model **f0.8.16**), and from 16 to 32 days (model **f0.16.32**). In the end, we chose the model with the highest log likelihood and used the posterior PDFs from that model. This is indicated as a “best model” column in the final catalog and in Table 1.

3.3. The gyrochrone fitting model

We developed a Bayesian mixture model to describe the relationship between stellar rotation periods and effective temperatures.

We modeled the observed rotation periods as arising from a mixture of two components: a primary gyrochronological relationship and a background population of stars that do not follow the standard age-rotation-temperature relation.

Our gyrochronology model employs Legendre polynomials as the basis functions for regression analysis. While these polynomials are orthogonal on the domain $-1 \leq x \leq 1$, they provide superior numerical stabil-

ity compared to standard polynomial bases even on unbounded domains. The Legendre polynomials are constructed using the recursion relation:

$$(\ell + 1)P_{\ell+1}(x) = (2\ell + 1)xP_{\ell}(x) - \ell P_{\ell-1}(x) \quad (25)$$

We model the gyrochronological relationship as a function that provides the expected rotation period at any effective temperature:

$$\mu_P(T_{\text{eff}}) \quad (26)$$

The true rotation periods follow a normal distribution around this expected relationship with intrinsic scatter σ :

$$P \sim \mathcal{N}(\mu_P(T_{\text{eff}}), \sigma) \quad (27)$$

To ensure numerical stability, we standardize both observed periods P_{obs} and temperatures T_{obs} to zero mean and unit standard deviation before fitting, then transform results back to physical scales. In this standardized coordinate system, the expected period-temperature relationship is expressed as:

$$\mu_y(x) = \sum_{\ell} C_{\ell} P_{\ell}(x) \quad (28)$$

where C_{ℓ} are the polynomial coefficients to be inferred.

To account for period measurement uncertainties, we introduce the measurement model relating the true periods P to the observed periods,

$$P_{\text{obs}} \sim \mathcal{N}(P, \sigma_{\text{obs}}). \quad (29)$$

Since this model is linear in P , we can analytically marginalize over the latent period variables (see Hogg et al. (2020)):

$$P_{\text{obs}} \sim \mathcal{N}(\mu_P(T_{\text{eff}}), \sqrt{\sigma_{\text{obs}}^2 + \sigma^2}) \quad (30)$$

To accommodate stars that deviate from the standard gyrochronological sequence, we implement a mixture model where each star's true period probabilistically follows either the gyrochronological relation or a background distribution. With probability f_{bg} , a star follows the background model:

$$P \sim \mathcal{N}(\mu_{\text{bg}}, \sigma_{\text{bg}}) \quad (31)$$

With probability $(1 - f_{\text{bg}})$, it follows the gyrochronological relation:

$$P \sim \mathcal{N}(\mu_P(T_{\text{eff}}), \sigma) \quad (32)$$

Then the mixture model probability density is:

$$p(P | \dots) \propto (1 - f_{\text{bg}})\mathcal{N}[\mu_P(T_{\text{eff}}), \sigma](P) + f_{\text{bg}}\mathcal{N}[\mu_{\text{bg}}, \sigma_{\text{bg}}](P) \quad (33)$$

We analytically marginalize over P in each mixture component and constrain $0 \leq f_{\text{bg}} \leq 1/2$ to ensure the gyrochronological component remains dominant.

We also account for uncertainty in temperature measurements that were taken from the Kepler Input Catalog, Cat. V/133 (Kepler Mission Team 2009); the gyrochronological relation above is expressed in terms of the true effective temperature, to which we apply a Normal prior,

$$T \sim \mathcal{N}(\mu_T, \sigma_T). \quad (34)$$

We assume that the observed temperatures are normally distributed around true temperatures:

$$T_{\text{obs}} \sim \mathcal{N}(T, \sigma_{T,\text{obs}}) \quad (35)$$

with $\sigma_T = 200$ K (Kepler Mission Team 2009). The gyrochronological relation depends on the true temperature T , creating nonlinear dependencies that require direct sampling over the latent temperature variables rather than analytical marginalization.

Using the ratio of terms in the mixture model, we can compute for each star the posterior probability that it belongs to the gyrochronological sequence as opposed to the background population. Estimates of stars' true P and T_{eff} are improved by the model compared to the raw measurement error, particularly for stars with high probability of belonging to the gyrochronological sequence in regions where the sequence is a strong function of temperature (where our accurate inferences of the rotation period allow to infer the temperature from the sequence); see Figure 6.

3.4. Vetting

From our membership list and complete sample of detected periods, we applied a series of quality cuts to ensure robust rotation period measurements. First, we implemented a “significance” cut based on the parameter σ , which represents the RMS fluctuation in the light curve attributable to the rotation terms in the GP kernel and is therefore directly related to the variability amplitude (see Section 3.2). We define the “significance” as the ratio $\text{mean}_{\sigma}/\text{std}_{\sigma}$, and retained only objects with $\text{significance} \geq 15$ to ensure sufficient signal-to-noise in the rotational modulation. This value was chosen based on visual inspection. Second, we imposed a precision requirement by excluding all targets with period uncertainties exceeding 1 day, maintaining a sample with well-constrained rotation periods suitable for our analysis. Figure 4 illustrates the effectiveness of our qual-

ity cut procedure in refining the rotation period sample, comparing the complete dataset of detected periods (left panel) with the filtered sample after applying significance, precision, and giants cuts (right panel).

To remove evolved stars from our sample, we employed a color-magnitude cut based on the Gaia color-magnitude diagram. We first calculated absolute G magnitudes using the Gaia G magnitudes and parallax measurements (Gaia Collaboration et al. 2018). For all stars with absolute G magnitude > 0 , we fit a third-order polynomial of the form:

$$M_G = 3.37 + (-2.13) \cdot (G_{BP} - G_{RP}) + 2.02 \cdot (G_{BP} - G_{RP})^2 + (-2.66) \cdot (G_{BP} - G_{RP})^3 \quad (36)$$

To define the main sequence upper envelope, we subtracted 1.5 magnitudes from this polynomial fit based on visual identification of a boundary that excludes evolved stars while retaining main-sequence objects. Stars falling above this adjusted boundary were removed from our sample. This cut was applied only to stars with available Gaia G magnitude, parallax, and effective temperature (T_{eff}) measurements. Figure 5 demonstrates our evolutionary status selection using the Gaia color-magnitude diagram to distinguish between main-sequence stars (retained in the sample) and evolved giants (excluded from the sample).

Finally, we sought to ensure that each measured rotation period is correctly associated with its target by ruling out signal contamination. Contamination is a common issue in *Kepler* data owing to the $4''$ square pixel size; it is estimated that, throughout the *Kepler* field, approximately 11% of sources will fall in the same pixel as a secondary target to which the instrument is also sensitive (Colman et al. 2017), though it should be noted that of course many of these secondary targets are likely to be non-variable. Nevertheless, due to the increased visual density of an open cluster, one would expect the rate of subpixel contamination to be fairly higher, not to mention that, depending on the chosen aperture, a target may also exhibit contamination from bright targets more than a pixel away. We assume that there is minimal contamination from field stars, and, given that many of the targets in this cluster and its field are so faint, we can also assume there is no contamination from stars not in the *Kepler* Input Catalog (KIC). We based our contamination detection on visual inspection of detrended IRIS light curves for 2,217 cluster members. We performed a cross-inspection of each target’s light curve with all light curves found within a radius of $40''$, or approximately ten pixels on either side. This is sufficient to capture contamination even from a bright, saturated

source at a decently large visual separation from the target. Where multiple targets seemed contaminated by the same signal, we also inspected Lomb-Scargle periodograms (Astropy Collaboration et al. 2013). Image subtraction photometry removes contaminating photon noise as background signal, so we are able to say that typically the highest-amplitude manifestation of any one signal is the real one.

After quality and giant cuts, we had 370 targets for visual inspection. There were only three targets where the detected period was clearly not due to stellar rotation: an eclipsing binary signal in the light curve of KIC 5023951, and the known γ -Doradus variable KIC 5024084, which also contaminated the light curve of KIC 5024091. This implies that the quality cuts were largely effective in removing detections of binarity or other stellar variability from the sample. Of the remaining 367 targets, we identified 45 clearly contaminated light curves and the relevant source of contamination; 38 “low visibility” targets where the rotational variability could not be clearly identified by eye; four targets where low data quality (due to systematic detector edge effects that manifest in dramatically different pixel sensitivities between quarters) suggested an inaccurate detection; and six targets affected by both low visibility and data quality issues. We note that the low visibility detections could be due to contamination by field stars, but regardless of their cause, they are omitted from the remainder of the study. Of the 319 clear signals with no data quality issues, this works out to a contamination fraction of 14%, which is in line with expectations for being higher than the field contamination fraction. Our final sample contains 271 reliable detections.

4. RESULTS & DISCUSSION

Figure 6 presents our final sample of rotation periods after applying all quality cuts, plotted against effective temperature in the gyrochronal color-period diagram. Using the model described in Section 3.3, we show the fits using observed temperatures (left panel) and also the inferred true temperatures (right panel). Both panels show rotation period as a function of effective temperature, with the red curve representing the median gyrochronological relationship and the yellow shaded regions indicating the intrinsic scatter (the σ parameter). Data points are color-coded according to their posterior probability of belonging to the gyrochronological sequence (P_{gyro}), with purple points indicating high probability of following the gyrochrone and green points representing likely background objects. Our mixture model successfully identifies stars that deviate from the canonical gyrochronological relation, probabilistically assign-




Figure 4. Two plots showing before and after applying the quality cuts described in Section 3.4. The left panel shows the complete dataset of detected rotation periods, while the right panel displays the same data after applying quality cuts. Quality cuts include: (1) a significance threshold requiring $\text{mean}_\sigma/\text{std}_\sigma$ to ensure sufficient signal-to-noise in the light curve amplitude, (2) a precision cut excluding targets with period uncertainties exceeding 1 day, and (3) a color-magnitude cut to remove evolved giant stars. Data points are colored by the statistical “significance” as we define in Section 3.4.




Figure 5. Color-magnitude diagram of 2,622 NGC 6819 members with available values for T_{eff} , Gaia parallax, and Gaia magnitude. The red dashed line represents a polynomial fitted to all points brighter than 0th absolute magnitude, with 1.5 magnitudes subtracted to align it with the visually identified upper envelope of the main sequence. The black points are main sequence members identified as acceptable candidates for stellar surface rotation detection, and the 238 white points were removed from the sample.

ing them to either the gyrochrone or background population. Stars with high P_{gyro} values (yellow) cluster tightly along the inferred relationship, while those with low

probabilities (blue) are scattered away from the main trend, demonstrating the model’s ability to distinguish between genuine gyrochronological sequence members.

Table 1 provides a representative sample of our final rotation period catalog, showcasing the precision of our measurements.

Our final sample represents the most comprehensive collection of rotation periods for members of a single open cluster to date. Douglas et al. (2019) assembled what was previously the largest catalog of Hyades rotators with 232 stars having measured rotation periods; our study expands this significantly to 271 stars with high-quality rotation period measurements. The combination of our expanded NGC 6819 catalog with the established Hyads sequence from Douglas et al. (2019) builds on the foundations of gyrochronology with strong results for a younger and an older open cluster, at the representative ages of $\sim 600\text{--}800$ Myr (Hyades) and ~ 2.5 Gyr (NGC 6819).

Figure 7 shows the distribution of rotation periods for stars in our NGC 6819 sample. Traditional stellar spin-down theory (e.g. Skumanich 1972; Barnes 2003) predicts a mass-dependent evolution where more massive, hotter stars (F and early G dwarfs) spin down more gradually due to their thinner convective envelopes and weaker magnetic dynamos, while lower-mass K dwarfs should experience rapid spin-down due to their deeper convective zones and more efficient magnetic braking. Our observed distribution exhibits a pronounced bimodality with a primary peak around 10–30 days and a significant population of faster rotators with periods




Figure 6. Gyrochronological relationship between stellar rotation period and effective temperature. The left panel shows the model fit using observed temperatures, while the right panel illustrates inferences of the true temperatures as described in Section 3.3. In all cases, errorbars give the $1-\sigma$ (68%) credible interval. The red curves represent the average of 100 samples with 1σ (intrinsic scatter) uncertainty bands (yellow shading). Data points are colored by their posterior probability of belonging to the gyrochrone (P_{gyro}), from blue (low probability outliers) to yellow (high probability). The mixture model successfully identifies and probabilistically weights stars that follow the canonical gyrochronological relation versus those in the background population.

Table 1. Representative entries from our rotation period catalog showing KIC identifiers, best-fitting GP model variants, derived rotation periods with uncertainties, and stellar parameters. The “Best Model” column indicates the optimal GP kernel configuration selected for each star based on model comparison criteria. Period uncertainties are reported as median values with -1σ and $+1\sigma$ bounds from the posterior distributions. Stellar parameters include effective temperature, surface gravity ($\log g$), and stellar radius in solar units taken from the KIC catalog.

KIC ID	Best Model	Period (days)	T_{eff} (K)	$\log g$	$R/R_{\odot}$
5112874	f0_16_32	$24.916^{+0.256}_{-0.253}$	4830	4.579	0.78
5024947	f0_16_32	$21.019^{+0.161}_{-0.161}$	4867	4.551	0.82
5112628	f0_8_16	$12.862^{+0.266}_{-0.267}$	4876	4.537	0.84
5025235	f0_16_32	$21.806^{+0.152}_{-0.151}$	4885	4.512	0.87
5024079	f0_16_32	$17.554^{+0.034}_{-0.035}$	4900	4.556	0.82

between 1-10 days. We will further discuss the substructure within the primary peak in Section 4.2.

Figure 8 shows the rotation period-temperature relation derived from our Bayesian analysis. The orange solid line represents the mean gyrochronological relation obtained from the Legendre polynomial fit, with the orange shaded region indicating the $\pm 1\sigma$ of the curve. Our sample (gray circles) exhibits the expected trend of increasing rotation periods toward cooler temperatures.

We compare our results with previous literature on NGC 6819: the measurements from Meibom et al.

(2015) (orange stars with error bars), and cluster members that were represented in the Santos et al. (2021) survey of all Kepler data (red diamonds). We also include the Curtis et al. (2020) results for Ruprecht 147 (purple triangles), which provides a comparison at a similar age of ~ 2.7 Gyr. As expected, the Ruprecht 147 data points lie towards the upper envelope of our sample.

The literature periods generally follow our derived mean relation, particularly in the temperature range

5500-6000 K, where the gyrochronological sequence is well-established.




Figure 7. Distribution of rotation periods in our final quality sample. The histogram shows the distribution of measured rotation periods, with the overlaid curve representing a kernel density estimate. The distribution exhibits a bimodality with peaks around 5 days and 20 days.

4.1. Gyrochronology reliability




Figure 8. Comparison of our derived gyrochronological relationship with literature estimates. Our *Kepler* field sample (gray circles, $N=271$) is shown alongside rotation periods from open clusters: NGC 6819 (red diamonds, Santos et al. (2021), $N=51$), Ruprecht 147 (purple triangles, Curtis et al. (2020), $N=35$), and Meibom et al. (2015) measurements (orange stars, $N=30$). The red curve represents our inferred mean gyrochronological relation of 100 samples, with intrinsic scatter indicated by the yellow shaded region. The agreement between our field star calibration and cluster benchmarks validates our approach.

Our in-depth study of NGC 6819 presents an opportunity to probe the observational challenges to traditional gyrochronology. This builds on the work by Curtis et al. (2020) with the closely co-eval clusters NGC 6819 and Ruprecht 147. The breakdown of gyrochronology is particularly evident in the apparent “stalling” of rotational evolution observed in intermediate-age clusters, observed by Curtis et al. (2020) in their combined ~ 2.5 Gyr dataset. Douglas et al. (2019) found that low-mass stars ($\lesssim 0.6 M_{\odot}$) in the 700 Myr Praesepe cluster rotate at essentially the same rate as their counterparts in the ~ 670 Myr Praesepe cluster, despite the age difference. Separately, Curtis et al. (2019) found similar convergence between the 1 Gyr NGC 6811 cluster and younger Hyades/Praesepe.

To further contextualize our observations, we turn to the work by Spada & Lanzafame (2020), who demonstrated that the competing effects of surface magnetic braking and interior angular momentum transport can naturally explain the observed stalling phenomenon. The efficiency of the internal angular momentum transport process exhibits a steep mass dependence, causing different mass stars to undergo re-coupling between their radiative cores and convective envelopes at different ages. During the re-coupling phase, the apparent surface spin-down is dramatically reduced or “stalled” because the angular momentum transported from the interior partially compensates for the angular momentum loss at the surface due to the magnetized wind. The Spada & Lanzafame (2020) model predicts large ranges in rotation period across stellar mass for young coeval groups, but it is precisely around 2–3 Gyr that the model starts to converge for masses in the FGK regime. To return to the example from Curtis et al. (2019), this effect naturally explains why stars in the young clusters NGC 6811 and Praesepe show overlapping rotation periods at low masses while maintaining the expected age ordering at higher masses (where re-coupling occurred earlier), and nicely illustrates why the ~ 2.5 Gyr stars in NGC 6819 display the continued effects of period bimodality in the post-stalling epoch.

The challenges are not limited to individual stellar evolution but extend to population-level analyses. Lu et al. (2021) examined rotation period distributions in different Galactic populations and found that the relationship between rotation and kinematics – a proxy for age – varies systematically with stellar mass and metallicity in ways that simple gyrochronology models cannot reproduce. This suggests that environmental factors, stellar composition, and possibly even Galactic chemical evolution history may influence rotational evolution in ways not currently accounted for in age-dating pre-

scriptions. Additionally, it is important to note that all theoretical gyrochronology calibrations – including those discussed here – assume single-star evolution, but a significant fraction of observed field stars may have experienced binary interactions that would alter their rotation rates in ways not captured by single-star models. Wide binaries may not be spectroscopically detected as binary systems but could still influence rotation evolution through tidal interactions and overall scatter in the data.

The consensus emerging from recent works is that while rotation periods remain valuable age indicators for young (< 1 Gyr) stellar populations, their reliability diminishes significantly for older stars, particularly those with rotation periods exceeding $\sim 15 - 20$ days or in mass ranges experiencing rotational coupling transitions. Accurate gyrochronological age determination requires detailed modeling of both magnetic braking and internal angular momentum transport, with careful consideration of the mass-dependent timescales governing core-envelope coupling and the potentially variable initial conditions set during the pre-main-sequence phase.

Our results for NGC 6819 (~ 2.5 Gyr) provide direct observational evidence for these theoretical predictions at older ages. The substantial scatter observed in our rotation period measurements, with many stars rotating significantly faster than expected from standard gyrochronological relations, demonstrates that the rotational coupling effects identified by Spada & Lanzafame (2020) continue to influence stellar evolution well beyond 1 Gyr. However, while Spada & Lanzafame (2020) demonstrated that rotational coupling effects can explain the stalled spin-down observed between Praesepe and NGC 6811, their model predicts that most stars should achieve quasi-solid-body rotation and enter a stable wind-braking-dominated regime by ages of 2-3 Gyr. The substantial rotational scatter we observe with our expanded sample of NGC 6819 periods therefore suggests that additional physical mechanisms beyond the two-zone coupling model may be operating at these older ages.

4.2. The pile up

Looking at the period-effective temperature diagram (Figure 9), we observe a distinct linear sequence of stars within the longer-period regime that runs roughly parallel to the main stellar relation but offset to shorter periods. Select visual inspection reveals no evidence of incorrectly measured periods. This tight, well-defined sequence spans effective temperatures from approximately 5500K to 6000K and appears to represent a separate population from the broader scatter of stars around the

primary relation. The coherent relationship suggests they share some common physical property or evolutionary characteristic.




Figure 9. Rotation period-temperature relationship showing an anomalous horizontal feature in the data with red points. Black points represent our rotation period measurements with error bars, while red points highlight a distinct population of stars that exhibit unusually uniform rotation periods (~ 16 -17 days) across a wide range of effective temperatures (approximately 5200-5800 K).

This could represent a manifestation of the “long-period pileup” described by David et al. (2022) that provides observational evidence for weakened magnetic braking in Sun-like stars. This feature, which spans effective temperatures from approximately 6250K to 5500K and follows a curve of nearly constant Rossby number ($Ro \approx 1.4$), demonstrates that stars across a range of masses and ages have converged onto the same rotational state, contradicting traditional spin-down models that predict continuous angular momentum loss throughout stellar lifetimes. This pileup also validates the theoretical predictions of van Saders et al. (2019). The scatter in rotation periods at NGC 6819’s age may reflect the “Rossby edge” in the *Kepler* field population described in van Saders et al. (2019). The existence of this linear ridge, populated by stars with ages ranging from $\sim 2 - 6$ Gyr, fundamentally challenges gyrochronology applications for solar-type stars and suggests that stellar rotation evolution involves distinct phases where spin-down effectively stalls, creating these observable “lines” in the period-temperature plane. Curtis et al. (2020) noted that their observed ~ 2.5 Gyr gyrochrone intersects with the “gap” observed in *Kepler* field star periods (Santos et al. 2021; McQuillan et al. 2014). The presence of a pile-up in the expanded NGC 6819 dataset which also crosses the *Kepler*

“gap” supports the idea that the pile-up is caused by a variety of complex factors but is likely not related to age or Galactic formation history.

5. CONCLUSIONS

In this work, we used *Kepler* IRIS light curves to re-measure stellar rotation periods in the open cluster NGC 6819, producing the most comprehensive rotation period catalog for any open cluster. By applying a combination of Gaussian Process modeling, rigorous vetting, and membership constraints, we isolated 271 high-confidence rotation period detections from an initial sample of over two thousand stars. This effort expands the available rotator sample in NGC 6819 by nearly an order of magnitude compared to previous studies.

Our results confirm the overall trend of increasing period toward cooler stars, but also reveal significant intrinsic scatter and a bimodal period distribution that cannot be reconciled with simple spin-down models. The presence of both fast and slow rotators at the same effective temperature suggests that internal angular momentum transport and mass-dependent coupling between stellar interiors and envelopes play a key role in shaping rotational evolution beyond ~ 1 Gyr. Furthermore, we find evidence for a “pile-up” of stars at nearly constant period.

Together, our results reinforce NGC 6819 as a cornerstone for testing gyrochronology at intermediate ages,

but also demonstrate that standard gyrochronological relations break down at the 2-3 Gyr timescale. Rotation periods remain powerful tracers of stellar physics, yet their use as precise chronometers for older field stars requires models that incorporate angular momentum transport and magnetic braking models. Our expanded catalog provides a critical benchmark for calibrating such next-generation gyrochronology models, and will serve as a valuable dataset for comparisons with upcoming missions such as *PLATO*.

ACKNOWLEDGMENTS

S.S. thanks Andy Casey for valuable discussions. S.S. acknowledges support from award 644616 from the Simons Foundation. Software citation information aggregated (in part) using [The Software Citation Station](#) (Wagg & Broekgaarden 2024; Wagg et al. 2025).

Software: [Astropy](#) (Astropy Collaboration et al. 2013, 2018), [Matplotlib](#) (Hunter 2007), [NumPy](#) (Harris et al. 2020), [Pandas](#) (Reback et al. 2022), [SciPy](#) (Virtanen et al. 2021), [PyMC](#) (Abril-Pla et al. 2023), [emcee](#) (Foreman-Mackey et al. 2013), [julia](#) (Bezanson et al. 2017), [Turing.jl](#) (Ge et al. 2018), [Makie.jl](#) (Danisch & Krumbiegel 2021), [ArviZ](#) (Kumar et al. 2019), [xarray](#) (Hoyer & Hamman 2017).

REFERENCES

- Abril-Pla, O., Andreani, V., Carroll, C., et al. 2023, *PeerJ Computer Science*, 9, e1516, doi: [10.7717/peerj-cs.1516](#)
- Alard, C., & Lupton, R. H. 1998, *ApJ*, 503, 325, doi: [10.1086/305984](#)
- Angus, R., Beane, A., Price-Whelan, A. M., et al. 2020, *AJ*, 160, 90, doi: [10.3847/1538-3881/ab91b2](#)
- Anthony-Twarog, B. J., Deliyannis, C. P., & Twarog, B. A. 2014, *AJ*, 148, 51, doi: [10.1088/0004-6256/148/3/51](#)
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: [10.1051/0004-6361/201322068](#)
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: [10.3847/1538-3881/aabc4f](#)
- Barnes, S. A. 2003, *ApJ*, 586, 464, doi: [10.1086/367639](#)
- . 2007, *ApJ*, 669, 1167, doi: [10.1086/519295](#)
- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2017, *SIAM Review*, 59, 65, doi: [10.1137/141000671](#)
- Böhm-Vitense, E. 2007, *ApJ*, 657, 486, doi: [10.1086/510482](#)
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, 327, 977, doi: [10.1126/science.1185402](#)
- Bouma, L. G., Palumbo, E. K., & Hillenbrand, L. A. 2023, *ApJL*, 947, L3, doi: [10.3847/2041-8213/acc589](#)
- Colman, I. L., Bedding, T. R., Huber, D., & Kjeldsen, H. 2022, *ApJS*, 258, 39, doi: [10.3847/1538-4365/ac3a11](#)
- Colman, I. L., Huber, D., Bedding, T. R., et al. 2017, *MNRAS*, 469, 3802, doi: [10.1093/mnras/stx1056](#)
- Covelo-Paz, A., Themeßl, N., Espinoza-Rojas, F., & Hekker, S. 2023, *A&A*, 679, A134, doi: [10.1051/0004-6361/202347211](#)
- Curtis, J. L., Agüeros, M. A., Douglas, S. T., & Meibom, S. 2019, *ApJ*, 879, 49, doi: [10.3847/1538-4357/ab2393](#)
- Curtis, J. L., Agüeros, M. A., Matt, S. P., et al. 2020, *ApJ*, 904, 140, doi: [10.3847/1538-4357/abbf58](#)
- Danisch, S., & Krumbiegel, J. 2021, *Journal of Open Source Software*, 6, 3349, doi: [10.21105/joss.03349](#)
- David, T. J., Angus, R., Curtis, J. L., et al. 2022, *ApJ*, 933, 114, doi: [10.3847/1538-4357/ac6dd3](#)
- Douglas, S. T., Curtis, J. L., Agüeros, M. A., et al. 2019, *ApJ*, 879, 100, doi: [10.3847/1538-4357/ab2468](#)

- Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, *AJ*, 154, 220, doi: [10.3847/1538-3881/aa9332](https://doi.org/10.3847/1538-3881/aa9332)
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306, doi: [10.1086/670067](https://doi.org/10.1086/670067)
- Gaia Collaboration, Babusiaux, C., van Leeuwen, F., et al. 2018, *A&A*, 616, A10, doi: [10.1051/0004-6361/201832843](https://doi.org/10.1051/0004-6361/201832843)
- Ge, H., Xu, K., & Ghahramani, Z. 2018, in *International Conference on Artificial Intelligence and Statistics, AISTATS 2018*, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain, 1682–1690. <http://proceedings.mlr.press/v84/ge18b.html>
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Hogg, D. W., Price-Whelan, A. M., & Leistedt, B. 2020, *arXiv e-prints*, arXiv:2005.14199, doi: [10.48550/arXiv.2005.14199](https://doi.org/10.48550/arXiv.2005.14199)
- Holcomb, R. J., Robertson, P., Hartigan, P., Oelkers, R. J., & Robinson, C. 2022, *ApJ*, 936, 138, doi: [10.3847/1538-4357/ac8990](https://doi.org/10.3847/1538-4357/ac8990)
- Hole, K. T., Geller, A. M., Mathieu, R. D., et al. 2009, *AJ*, 138, 159, doi: [10.1088/0004-6256/138/1/159](https://doi.org/10.1088/0004-6256/138/1/159)
- Hoyer, S., & Hamman, J. 2017, *Journal of Open Research Software*, 5, doi: [10.5334/jors.148](https://doi.org/10.5334/jors.148)
- Hunter, J. D. 2007, *Computing in Science and Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010, *ApJL*, 713, L120, doi: [10.1088/2041-8205/713/2/L120](https://doi.org/10.1088/2041-8205/713/2/L120)
- Kepler Mission Team. 2009, *VizieR Online Data Catalog: Kepler Input Catalog (Kepler Mission Team, 2009)*, *VizieR On-line Data Catalog: V/133*. Originally published in: *Kepler Mission Team, NASA (2009)*
- Kuehn, C. A., Drury, J. A., Bellamy, B. R., et al. 2015, in *European Physical Journal Web of Conferences*, Vol. 101, *European Physical Journal Web of Conferences (EDP)*, 06040, doi: [10.1051/epjconf/201510106040](https://doi.org/10.1051/epjconf/201510106040)
- Kumar, R., Carroll, C., Hartikainen, A., & Martin, O. 2019, *Journal of Open Source Software*, 4, 1143, doi: [10.21105/joss.01143](https://doi.org/10.21105/joss.01143)
- Lee-Brown, D. B., Anthony-Twarog, B. J., Deliyannis, C. P., Rich, E., & Twarog, B. A. 2015, *AJ*, 149, 121, doi: [10.1088/0004-6256/149/4/121](https://doi.org/10.1088/0004-6256/149/4/121)
- Lu, Y. L., Angus, R., Curtis, J. L., David, T. J., & Kiman, R. 2021, *AJ*, 161, 189, doi: [10.3847/1538-3881/abe4d6](https://doi.org/10.3847/1538-3881/abe4d6)
- Mamajek, E. E., & Hillenbrand, L. A. 2008, *ApJ*, 687, 1264, doi: [10.1086/591785](https://doi.org/10.1086/591785)
- Matt, S. P., Brun, A. S., Baraffe, I., Bouvier, J., & Chabrier, G. 2015, *ApJL*, 799, L23, doi: [10.1088/2041-8205/799/2/L23](https://doi.org/10.1088/2041-8205/799/2/L23)
- Matt, S. P., Pinzón, G., de la Reza, R., & Greene, T. P. 2010, *ApJ*, 714, 989, doi: [10.1088/0004-637X/714/2/989](https://doi.org/10.1088/0004-637X/714/2/989)
- McQuillan, A., Mazeh, T., & Aigrain, S. 2014, *ApJS*, 211, 24, doi: [10.1088/0067-0049/211/2/24](https://doi.org/10.1088/0067-0049/211/2/24)
- Meibom, S., Barnes, S. A., Platais, I., et al. 2015, *Nature*, 517, 589, doi: [10.1038/nature14118](https://doi.org/10.1038/nature14118)
- Meibom, S., Barnes, S. A., Latham, D. W., et al. 2011, *ApJL*, 733, L9, doi: [10.1088/2041-8205/733/1/L9](https://doi.org/10.1088/2041-8205/733/1/L9)
- Morris, B. M., Hebb, L., Davenport, J. R. A., Rohn, G., & Hawley, S. L. 2017, *ApJ*, 846, 99, doi: [10.3847/1538-4357/aa8555](https://doi.org/10.3847/1538-4357/aa8555)
- Platais, I., Gosnell, N. M., Meibom, S., et al. 2013, *AJ*, 146, 43, doi: [10.1088/0004-6256/146/2/43](https://doi.org/10.1088/0004-6256/146/2/43)
- Rauer, H., Catala, C., Aerts, C., et al. 2014, *Experimental Astronomy*, 38, 249, doi: [10.1007/s10686-014-9383-4](https://doi.org/10.1007/s10686-014-9383-4)
- Reback, J., Jbrockmendel, McKinney, W., et al. 2022, *pandas-dev/pandas: Pandas 1.4.4, v1.4.4*, Zenodo, doi: [10.5281/zenodo.7037953](https://doi.org/10.5281/zenodo.7037953)
- Reinhold, T., Cameron, R. H., & Gizon, L. 2017, *A&A*, 603, A52, doi: [10.1051/0004-6361/201730599](https://doi.org/10.1051/0004-6361/201730599)
- Reinhold, T., & Gizon, L. 2015, *A&A*, 583, A65, doi: [10.1051/0004-6361/201526216](https://doi.org/10.1051/0004-6361/201526216)
- Reinhold, T., Reiners, A., & Basri, G. 2013, *A&A*, 560, A4, doi: [10.1051/0004-6361/201321970](https://doi.org/10.1051/0004-6361/201321970)
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003, doi: [10.1117/1.JATIS.1.1.014003](https://doi.org/10.1117/1.JATIS.1.1.014003)
- Sagynbayeva, S., Farr, W. M., Morris, B. M., & Luger, R. 2025, *ApJ*, 990, 32, doi: [10.3847/1538-4357/adf6be](https://doi.org/10.3847/1538-4357/adf6be)
- Sanjayan, S., Baran, A. S., Németh, P., et al. 2022, *AcA*, 72, 77, doi: [10.32023/0001-5237/72.2.1](https://doi.org/10.32023/0001-5237/72.2.1)
- Santos, A. R. G., Breton, S. N., Mathur, S., & García, R. A. 2021, *ApJS*, 255, 17, doi: [10.3847/1538-4365/ac033f](https://doi.org/10.3847/1538-4365/ac033f)
- Skumanich, A. 1972, *ApJ*, 171, 565, doi: [10.1086/151310](https://doi.org/10.1086/151310)
- Smith, J. C., Stumpe, M. C., Van Cleve, J. E., et al. 2012, *PASP*, 124, 1000, doi: [10.1086/667697](https://doi.org/10.1086/667697)
- Spada, F., & Lanzafame, A. C. 2020, *A&A*, 636, A76, doi: [10.1051/0004-6361/201936384](https://doi.org/10.1051/0004-6361/201936384)
- Van Cleve, J. E., Christiansen, J. L., Jenkins, J. M., et al. 2016, *Kepler Data Characteristics Handbook*, *Kepler Science Document KSCI-19040-005*, id. 2. Edited by Doug Caldwell, Jon M. Jenkins, Michael R. Haas and Natalie Batalha
- van Saders, J. L., Ceillier, T., Metcalfe, T. S., et al. 2016, *Nature*, 529, 181, doi: [10.1038/nature16168](https://doi.org/10.1038/nature16168)
- van Saders, J. L., Pinsonneault, M. H., & Barbieri, M. 2019, *ApJ*, 872, 128, doi: [10.3847/1538-4357/aafafe](https://doi.org/10.3847/1538-4357/aafafe)
- Virtanen, P., Gommers, R., Burovski, E., et al. 2021, *scipy/scipy: SciPy 1.6.3, v1.6.3*, Zenodo, doi: [10.5281/zenodo.4718897](https://doi.org/10.5281/zenodo.4718897)

- Wagg, T., Broekgaarden, F., & Gültekin, K. 2025,
TomWagg/software-citation-station: v1.3, v1.3, Zenodo,
doi: [10.5281/zenodo.17145205](https://doi.org/10.5281/zenodo.17145205)
- Wagg, T., & Broekgaarden, F. S. 2024, arXiv e-prints,
arXiv:2406.04405. <https://arxiv.org/abs/2406.04405>