



X-Ray, UV, and Optical Observations of Proxima Centauri's Stellar Cycle

Bradford J. Wargelin¹ , Steven H. Saar¹ , Zackery A. Irving², Jonathan D. Slavin¹ , Peter Ratzlaff¹, and José-Dias do Nascimento, Jr.^{1,3}

¹Center for Astrophysics—Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA; bwargelin@cfa.harvard.edu

²School of Physics and Astronomy, University of Southampton, University Road, Southampton SO17 1BJ, UK

³Univ. Federal do Rio G. do Norte, UFRN, Dep. de Física, CP 1641, 59072-970, Natal, RN, Brazil

Received 2024 June 21; revised 2024 November 4; accepted 2024 November 4; published 2024 December 10

Abstract

Proxima Cen (GJ 551; dM5.5e) is one of only about a dozen fully convective stars known to have a stellar cycle, and the only one to have long-term X-ray monitoring. A previous analysis found that X-ray and mid-UV observations, particularly two epochs of data from Swift, were consistent with a well-sampled ~ 7 yr optical cycle seen in All Sky Automated Survey project (ASAS) data, but not convincing by themselves. The present work incorporates several years of new ASAS-SN optical data and an additional 5 yr of Swift XRT and UVOT observations, with Swift observations now spanning 2009–2021 and optical coverage from late 2000. X-ray observations by XMM-Newton and Chandra are also included. Analysis of the combined data, which includes modeling and adjustments for stellar contamination in the optical and UV, now reveals clear cyclic behavior in all three wavebands with a period of 8.0 yr. We also show that UV and X-ray intensities are anticorrelated with optical brightness variations caused by the cycle and by rotational modulation, discuss possible indications of two coronal mass ejections, and provide updated results for the previous finding of a simple correlation between X-ray cycle amplitude and Rossby number over a wide range of stellar types and ages.

Unified Astronomy Thesaurus concepts: [Stellar activity \(1580\)](#); [Late-type dwarf stars \(906\)](#); [M dwarf stars \(982\)](#)

1. Introduction

Despite accounting for $\sim 70\%$ of the stellar population, M stars have, until very recently, been poorly represented in studies of magnetic activity cycles because of their intrinsic faintness. In the pioneering HK Project at Mount Wilson Observatory (S. L. Baliunas et al. 1995), which began in 1966 and monitored chromospheric Ca II H and K lines (3969 and 3934 Å) in roughly 300 stars, only one of the objects of study was an M star (Lalande 21185; dM2). Technological progress, however, has steadily brought an increasing number of M stars under scrutiny in spectroscopic programs such as HARPS (M. Mayor et al. 2003), MDO Planetary Search (W. D. Cochran & A. P. Hatzes 1993), and CASLEO HKalpha (C. Cincunegui & P. J. D. Mauas 2004), and in photometric monitoring programs such as the All Sky Automated Survey project (ASAS; G. Pojmanski 1997, 2002), ASAS for Supernovae (ASAS-SN; B. J. Shappee et al. 2014; C. S. Kochanek et al. 2017), and ATLAS (J. L. Tonry et al. 2018). Although discovery of planets via radial velocity measurements or detection of transient behavior such as supernovae are generally the focus of such projects, their sustained measurements over many years often lend themselves to studies of cyclic behavior as well.

Using roughly a decade of ASAS data, A. Suárez Mascareño et al. (2016) reported on apparent cycles in around 40 stars, half of which were M stars. Of those, around a dozen were fully convective, with stellar type M3.5 or later. This was a surprising result, since most theories of stellar magnetism predict that cyclic behavior can only be supported by solar-type $\alpha\Omega$ dynamos, which are driven by magnetic shear at a

radiative/convective boundary, or tachocline (M. Dikpati & P. Charbonneau 1999). Fully convective stars, of course, do not have tachoclines, and instead their magnetic fields are expected to be driven by α^2 dynamos. Some theoretical work, however, suggests that α^2 dynamos can in fact support activity cycles under certain conditions (G. Rüdiger et al. 2003; G. Chabrier & M. Küker 2006; T. Gastine et al. 2012; P. J. Käpylä et al. 2013; R. K. Yadav et al. 2016), and observations show that fully convective stars follow the same rotation–activity relation as partially convective stars (N. J. Wright & J. J. Drake 2016). The presence of cyclic behavior in the \sim dozen fully convective stars noted by A. Suárez Mascareño et al. (2016) was recently confirmed by Z. A. Irving et al. (2023), who combined the original ASAS data with several years of later ASAS-SN data. Three of those cycles plus one other (GJ876; M3.5) were also reported by L. Mignon et al. (2023) based on chromospheric emission lines (H α , Ca II H&K, or the Na D doublet).

A common target of those studies is Proxima Cen (dM5.5). A. Suárez Mascareño et al. (2016) analyzed 9 yr of Proxima ASAS-3 data and found a cycle with a period of 6.8 ± 0.3 yr, and B. J. Wargelin et al. (2017) found a period of 7.05 ± 0.15 yr after including five additional years of ASAS-4 data. The latter paper also analyzed data from several X-ray missions, particularly two seasons of Swift observations from 2009/2010 and 2012/2013, and reported that measured intensities were consistent with a stellar cycle opposite in phase to the optical cycle, with hints of the same anticorrelation for rotational modulation.

As the star nearest the Sun at a distance of 1.302 pc, Proxima is by far the most easily studied of late-type M's, and is a subject of particular interest for its possession of one confirmed planet (G. Anglada-Escudé et al. 2016) and two additional candidates (M. Damasso et al. 2020; J. P. Faria et al. 2022). Proxima is also notable as the only fully convective star that has been monitored in X-rays over at least one full stellar cycle,

making it unique among an already small set of stars to have their cycles measured at high energies, where cycle amplitudes are much larger than in the optical and more directly tied to magnetic activity.

In Section 2 we describe the assembly and analysis of optical data from ASAS and ASAS-SN (2.1), and X-ray (2.2) and UV data (2.3) from Swift and other missions, followed by a study of brightness correlations among those data sets on rotational timescales (2.4). In Section 3 we measure cycle properties in all three wave bands, compare the derived X-ray cycle amplitude with those from other stars of varying rotational rates and stellar types, and discuss possible signatures of coronal mass ejections, followed by a summary in Section 4.

2. Observational Data and Analysis

2.1. Optical Data

A previous study of Proxima’s cycle (B. J. Wargelin et al. 2017) used *V*-band optical data from ASAS, specifically the publicly available ASAS-3 data covering 2000 December to 2009 September, plus 5 yr of ASAS-4 data (2010 July to 2015 August; G. Pojmański 2024, private communication) that were cross-calibrated with ASAS-3 using 33 nearby stars. For the present work, we include additional ASAS-4 data extending into 2019 that were also provided for M. Damasso et al. (2020), but most of the new optical data come from the ASAS-SN program. ASAS-SN *V*-band observations of Proxima’s field began in 2016 March and ran through 2018 August, overlapping with a switch to the *g* band beginning in 2018 June.

V-band ASAS-3 data were downloaded from <https://www.astrouw.edu.pl/asas/>, and ASAS-4 data were provided by the program leader, G. Pojmański. As recommended, we used only measurements with A or B quality grades, and chose brightness measurements using the 1’ aperture (MAG_2; 4 pixel diameter), which had the lowest scatter among the five aperture choices. Typical ASAS resolution is 23’’ FWHM.

V- and *g*-band ASAS-SN data were downloaded from the ASAS-SN Sky Patrol site (<https://asas-sn.osu.edu/>) in multiple steps using coordinates accurate to within 1’’ over discrete time intervals (generally two intervals per observing season) and then collated into a single file. (ASAS extractions automatically account for proper motion, which is 3’’/86 yr for Proxima.) ASAS-SN brightness extractions use a 32’’-diameter aperture, approximately double the typical telescope point-spread function (PSF) FWHM.

To remove outliers, mostly from flares, we clipped measurements more than 2σ from the seasonal average in ASAS data. For ASAS-SN, we clipped measurements more than 2.5σ from the seasonal average, and then more than 2.5σ from the average in 7 day bins.

The next step was to combine the ASAS and ASAS-SN data, which were collected using different filters and, for ASAS-SN, several telescopes around the world. The *V* and *g* bands are not far apart (central wavelengths of 551 and 520 nm), and for a first approximation, we used the relations of S. M. Kent (1985) and R. A. Windhorst et al. (1991),

$$V = g - 0.03 - 0.42(g - r), \quad (1)$$

where $g - r$ is the $g - r$ color index, given by

$$g - r = 1.02(B - V) - 0.22, \quad (2)$$

where $B - V$ is the color index. Errors for that color correction, however, can be significant for a very red star like Proxima, as

can differences in the responses of filters, optics, and detectors among the several telescopes used here. Fortunately, ASAS observations overlap with ASAS-SN for about 3 yr, ASAS-SN *V* band with *g* band for about 3 months, and the several ASAS-SN telescopes with each other, allowing very accurate cross-calibration.

For internal ASAS-SN cross-calibration, we follow the procedure described by Z. A. Irving et al. (2023), taking the telescope with the most measurements as our reference and then sequentially (in order of their measurements’ temporal overlap with the reference set) applying small offsets to other telescopes so that the average difference among overlapping measurements is zero. As each telescope is added, the combined data set is then used as the reference. The same process was applied to merge ASAS data with the ASAS-SN set, yielding a combined optical light curve spanning 23 yr.

As seen in the first panel of Figure 1, optical data show the clear sinusoidal pattern of the ~ 7 yr cycle reported by A. Suárez Mascareño et al. (2016) and B. J. Wargelin et al. (2017), but in recent years the cycle seems to have become somewhat weaker, and there is an overall brightening trend, which Z. A. Irving et al. (2023) suggested is caused by Proxima’s proper motion into a more densely populated region of the sky.

To correct for the presumed stellar contamination, we convolved an optical image of the sky with the ASAS and ASAS-SN PSFs, simulated extractions along Proxima’s path, and adjusted the measured magnitudes. To construct the reference image, we used *g*-band CCD observations from SkyMapper (DR4 doi:10.25914/5M47-S621; C. A. Onken et al. 2024) collected on 2015 March 9 and 10 (nominally 5 s exposures, but effectively a little less than 4 s) and 2019 April 11 (100 s). The longer exposure provides a good signal-to-noise ratio (S/N) even for weak sources, but Proxima itself was saturated, so its brightness versus other stars was calibrated from the two short exposures, yielding a scaling factor of 26.5. We removed an ellipse around Proxima in the 100 s image, replacing it with a flat background and manually editing some pixels around the ellipse edges to restore the PSF wings of a few adjacent stars. In the earlier short exposures, collected when Proxima was 15’’ away due to 4 yr of proper motion, we measured the intensities of three stars that were obscured by Proxima in the long exposure and restored them in the reference image, scaled up $26.5\times$. The resulting reference image is shown in panel (c) of Figure 1, and after subtracting a flat background, it was convolved with Gaussians of $\text{FWHM} = 23''$ (for ASAS; panel (d)) and $\text{FWHM} = 16''$ (for ASAS-SN, not shown).

Counts within circles matching the ASAS and ASAS-SN extraction regions, with diameters of 60’’ and 32’’, respectively, were then measured as a function of date along Proxima’s proper-motion path and compared to the counts extracted from Proxima itself in the 5 s exposures (scaled up $26.5\times$ to match the 100 s exposure). The date-dependent fractional contamination of Proxima’s measured intensity caused by nearby stars was then subtracted from the uncorrected data, and cross-calibration of ASAS-4 and ASAS-SN in their overlapping interval (2016–2019) was repeated, yielding the corrected light curve in panel (b) of Figure 1. Cyclic behavior is now more apparent, but still relatively muted in recent years. Further discussion of the optical cycle is presented in Section 3, preceded by analysis of rotational modulation in Sections 2.4.1 and 2.4.2.

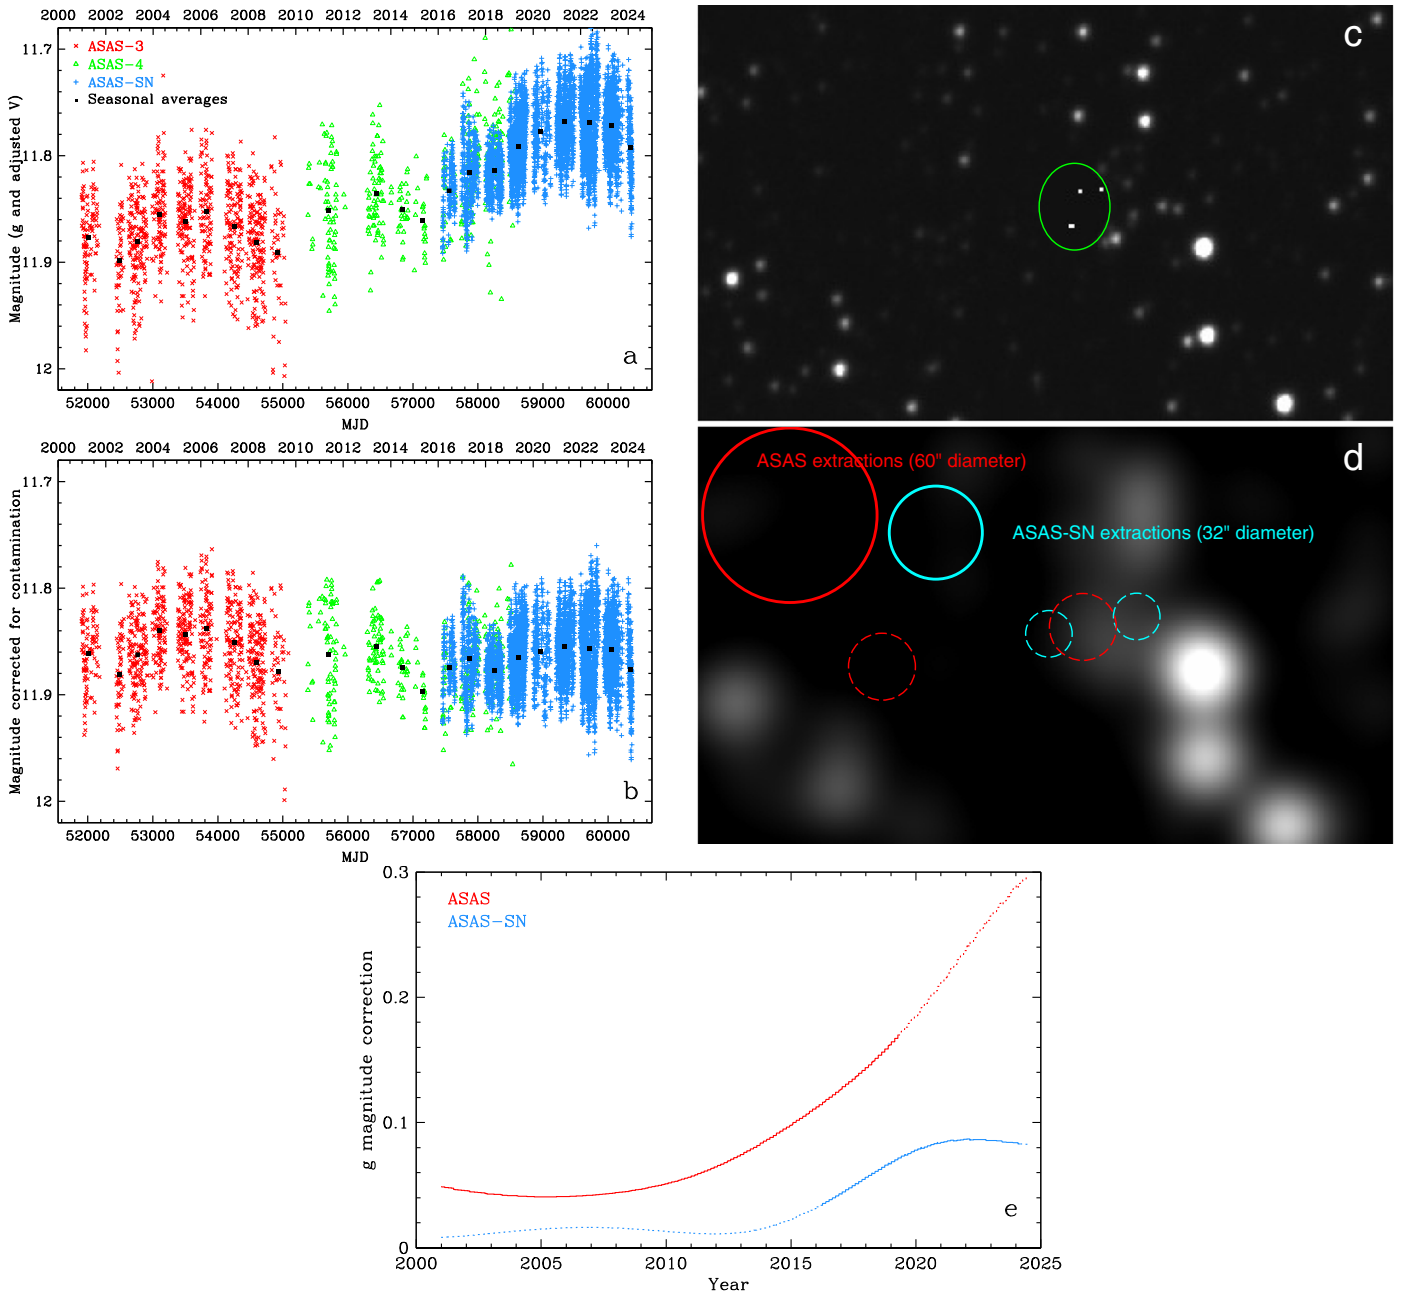


Figure 1. Optical light curves and corrections for stellar contamination. (a): ASAS-3, ASAS-4, and ASAS-SN data, after cross-calibration procedures described in the text. Seasonal average for 2011 includes the few measurements from 2010 and 2012. (b) Light curve after the correction process illustrated in panels (c)–(e). (c): SkyMapper *g*-band image (linear scaling) with Proxima removed (green ellipse) and three obscured sources added back using data from an observation when Proxima was at a different location. (d): Image from panel (c) convolved with the ASAS $23''$ -FWHM PSF. Pairs of dashed circles show the FWHM for ASAS (red) and ASAS-SN (blue) at Proxima’s position for the first (left) and last (right) observations in each data set. Solid circles in the upper left show sizes of ASAS and ASAS-SN brightness extraction regions. (e): Model results for stellar contamination of Proxima, with solid lines denoting observation intervals for both data sets.

2.2. X-Ray Data

B. J. Wargelin et al. (2017) analyzed X-ray and UV data from the Swift Observatory spanning 2009–2013, primarily from AOs 5 and 8 with a few observations from AOs 6 and 7. Additional observations presented here cover AOs 12–17 (excluding AO 14), 2016–2021. X-ray data are from the X-Ray Telescope (XRT; D. N. Burrows et al. 2005) and UV from the UltraViolet/Optical Telescope (UVOT; P. W. A. Roming et al. 2005) in the ~ 1000 Å wide W1 band centered around the chromospheric Mg II h&k lines at 2803.5 and 2796.3 Å.

The h&k transitions, $3P_{1/2,3/2} \rightarrow 3S_{1/2}$, are analogs of the $4P_{1/2,3/2} \rightarrow 4S_{1/2}$ Ca II H&K transitions widely used in studies of stellar cycles and other magnetic activity.

B. J. Wargelin et al. (2017) also included X-ray data from other missions back to 1994, but observations prior to 2009 were of questionable value because of calibration uncertainties and/or source flaring. The most useful measurements were made by XMM in 2009 and by the Chandra High Resolution Camera (HRC) in 2012 and 2015. Those and newer measurements by the same instruments will be discussed in Sections 2.2.2 and 2.2.3.

Table 1
X-Ray and UV Observations

Mission	Epoch	ObsIDs ^a	Dates	Exposure Time (s) ^b	
				X-Ray	UVOT/W1
XMM	5X	0551120301,201,401	2009 Mar 10–2009 Mar 14	75772.0	...
Swift	5	90215002–90215022	2009 Apr 23–2010 Apr 9	38631.0	20938.3 ^c
Swift	6	31676001–31676003	2010 Jul 10–2011 Mar 12	7908.0	8023.5
Swift	7	31676004–31676005	2011 Sep 4–2011 Sep 8	2685.3	2684.5
Chandra	8C	14276	2012 Jun 15	49626.4	...
Swift	8	31676006–31676022,91488001–91488003	2012 Mar 30–2013 Feb 18	38973.5	39386.8
Chandra	12Ca	17377	2015 Dec 9	35900.0	...
Swift	12	31676023–31676038	2016 Jul 11–2016 Dec 8	36687.8	34469.2
Chandra	12Cb	19788–19790,19793	2016 Sep 26–2016 Dec 8	36814.2	...
Swift	13	93156002–93156012	2017 Aug 25–2017 Nov 5	30829.8	27637.1
XMM	13X	0801880201,401,501	2017 Jul 27–2018 Mar 11	58819.8	...
Swift	15	95121001–95121007,31676040–31676043, 95159001–95159015	2019 Apr 10–2019 Jul 1	82409.9	29382.1 ^d
Swift	15D	95121008–95121015	2019 Dec 26–2020 Feb 27	18877.5	18875.0
Swift	16	95673001–95673012	2020 Apr 8–2020 Oct 1	37333.7	37381.8
Swift	17	96044001–96044018	2021 Apr 18–2021 Aug 18	31020.2	28530.9

Notes.

^a Swift ObsID ranges may include gaps from rescheduled or unsuccessful observations.

^b Swift times exclude periods of bad aspect, too-large SSS effects, or too-short exposures. HRC times exclude periods of telemetry saturation. XMM times are live-times for the pn detector operating with large window (deadtime fraction $\sim 7.5\%$); Swift and HRC times are ONTIMES (deadtime $< 0.5\%$).

^c Swift ObsIDs 90215002–90215010 used the UV grism with UVOT.

^d Swift ObsIDs 95159001–95159015 used the M2 filter with UVOT.

2.2.1. Swift

Swift observation information is listed in Table 1. XRT and UVOT observations are essentially simultaneous, but their exposures for AO 5 differ because the first eight observations used the UVOT UV grism before switching to the W1 filter. (The hope was to measure Mg II line intensities but the star field was too crowded, often leading to overlapping dispersed spectra.) Also, in AO 15 there was a separate Proxima observing program focused on flaring (M. A. MacGregor et al. 2021) that provided enough additional exposure time that AO 15 could be split into two epochs for our analysis, but most of those observations (all of the 95159 series) used the UVOT M2 filter instead of W1. Observation cadences were generally multiples of 4 days (4, 8, 12, or 16 depending on the year) to match the Swift “filter of the day” schedule, with enough observations during each epoch (except for AOs 6 and 7) to monitor Proxima fairly evenly over one or two stellar rotations.

All Swift data (XRT and UVOT) were downloaded from the Leicester data archive at https://www.swift.ac.uk/swift_portal/ where they were processed using HEASoft v6.28, particularly the FTOOLS package (<https://heasarc.gsfc.nasa.gov/ftools/>; J. K. Blackburn 1995). We used the Chandra Interactive Analysis of Observations (CIAO) software suite (A. Fruscione et al. 2006) for most analysis tasks, particularly the `dmcopy`, `dmlist`, and `dmstat` commands. For each observation, we computed approximate source positions (with proper-motion corrections) and then performed iterative centroiding using circular regions with 20 pixel ($47''.2$) radii and 300–2500 eV energy filtering (`pi = 30 : 249`), followed by unfiltered source and background (60–110 pixel annular) extractions. The XRT PSF has a half-power radius of $9''$ and so the source region includes some power from nearby blended sources, but inspection of a Chandra image (ObsID 49899) with much better resolution ($< 0''.5$ FWHM) indicates that no more than 1% of the extracted XRT counts come from these

other sources, even when restricted to times of Proxima’s quiescent emission.

The Swift observations last ~ 800 s on average, with the longest being 1850 s. Each consists of one or more “snapshots” usually with intervals of one or more 95 minute orbits. We separated the 128 XRT observations into their 397 component snapshots, further subdividing 10 snapshots into pre-flare and flare components, and then used the light-curve function of the “Swift-XRT data products generator” (P. A. Evans et al. 2009) at https://www.swift.ac.uk/user_objects/ to determine net 300–2500 eV source event rates for each snapshot (or pre/flare subdivision), corrected for dead pixels, event pileup, and other smaller effects. Those rates were then used to determine correction factors for our own rate analysis of each snapshot, so that appropriate rates could be calculated for an arbitrarily binned light curve, which is necessary for our analysis (see below) but not a capability that the products generator can directly provide.

Typical uncorrected quiescent event rates are ~ 0.07 cts s^{-1} , so to obtain a statistically adequate number of counts per bin in light curves, we divided snapshots into ~ 400 s pieces. This also fits well with the typical ~ 800 s snapshot exposure and is significantly less than the typical flaring timescale of several hundred seconds. To make optimal use of as much exposure time as possible, we divided each snapshot into equal pieces ranging from $400/\sqrt{2}$ to $400 \times \sqrt{2}$ (283 to 566) s. Note that 2.6% of the total exposure time occurs in snapshots shorter than 283 s and 1.0% in those shorter than 200 s. We discarded bins shorter than 200 s (3132.3 s in total), leaving 325,356 s of XRT exposure divided among 836 bins, averaging 389 s per bin.

Corrected background-subtracted rates were then computed for each bin; Figure 2 shows concatenated light curves for each epoch. Rates during large flares are less accurate because finer time binning would be required to compute their highly rate-dependent correction factors; our focus is on the quiescent rates that characterize the “baseline” emission of Proxima and its




Figure 2. Concatenated Swift X-ray light curves. Time bins average ~ 400 s but vary so time axes are not exact. Dotted horizontal lines mark each epoch’s 10th and 60th percentile rates (see Figure 3). Vertical orange bars denote periods of overlap with Chandra HRC observations (see Figure 5). Epochs 6 and 7 were not included in further analysis because of their short exposures and high proportion of flaring.

changes over the stellar cycle. Of the few stars with X-ray cycle monitoring, Proxima has by far the most flaring, and including its full light curve in such an analysis would create relatively large and poorly estimated uncertainties because of the variable number of large flares. Our measurements of quiescent rates therefore include only time bins in the 10th–60th percentile of event rates, where rate distributions from one epoch to another are most similar apart from overall normalization (see Figure 3). With this choice, relative rates from one epoch to another are virtually the same (typically $\lesssim 1\%$ differences) over a fairly large percentile sampling range.

Source and background event files from each of those time bins were merged into composite quiescent emission files for each epoch, and their spectra fit with Sherpa (P. Freeman et al. 2001) using the detector Ancillary Response File (`swxs6_20010101v001.arf`) and appropriate Response Matrix File (RMF) for each epoch (`swxpc0to12s6_20090101v014.rmf` for AO 5, `*20110101v014.rmf` for AO 8, and `*20130101v014.rmf` for AO 12 and later). There was too little exposure time in AOs 6 or 7, even when combined, to obtain a reliable estimate of their quiescent emission level, especially since their snapshots’ light curves indicate most of their limited exposure time probably occurred during significant flaring (see Figure 2).

We used the Sherpa `xsphabs` model in combination with either two or three `xsapec` components at different temperatures for our fits. The `xsphabs` H column density was frozen at 10^{18} cm^2 and had no effect in the 0.3–5 keV band used for fitting. Two temperatures generally gave formally acceptable results in terms of the reduced statistic, but 3-T fits were visibly better. Three temperatures also follows the rationale of S. Orlando et al. (2017) and M. Coffaro et al. (2020) in their analyses of X-ray stellar cycles in HD 81809 and ϵ Eri, respectively, that three kT ’s are necessary to model the emission from active regions (AR; lowest kT), AR cores (CO; medium kT), and higher-temperature regions that produce flaring emission that is subtle enough not to be obvious in light curves (FL; highest kT). A fourth component at lower temperatures representing




Figure 3. XRT and UVOT/W1 rate distributions for each epoch, with vertical lines marking the 10th–60th percentile range used for calculating average quiescent rates. UVOT rates are scaled down by a factor of 4 in the plot.

plasma like that in the “quiet Sun” contributes negligibly to detected X-ray emission from these stars and is not included.

Following the 3-T paradigm, we fit each epoch’s spectrum with frozen temperatures and abundances, allowing only the normalization of each temperature component to vary (see examples in Figure 4). The fixed temperatures and abundance were taken from fits to the combined spectra of all epochs that could be fit with a single RMF (that for AOs 12–17): $kT_1 = 0.22$ keV ($T = 2.55 \times 10^6$ K), $kT_2 = 0.56$ keV ($T = 6.50 \times 10^6$ K), $kT_3 = 0.96$ keV ($T = 11.14 \times 10^6$ K), and an abundance of 0.17 times solar. When leaving the abundance thawed in individual fits for each epoch, values ranged from 0.10–0.24 times solar photospheric values with typical 1σ

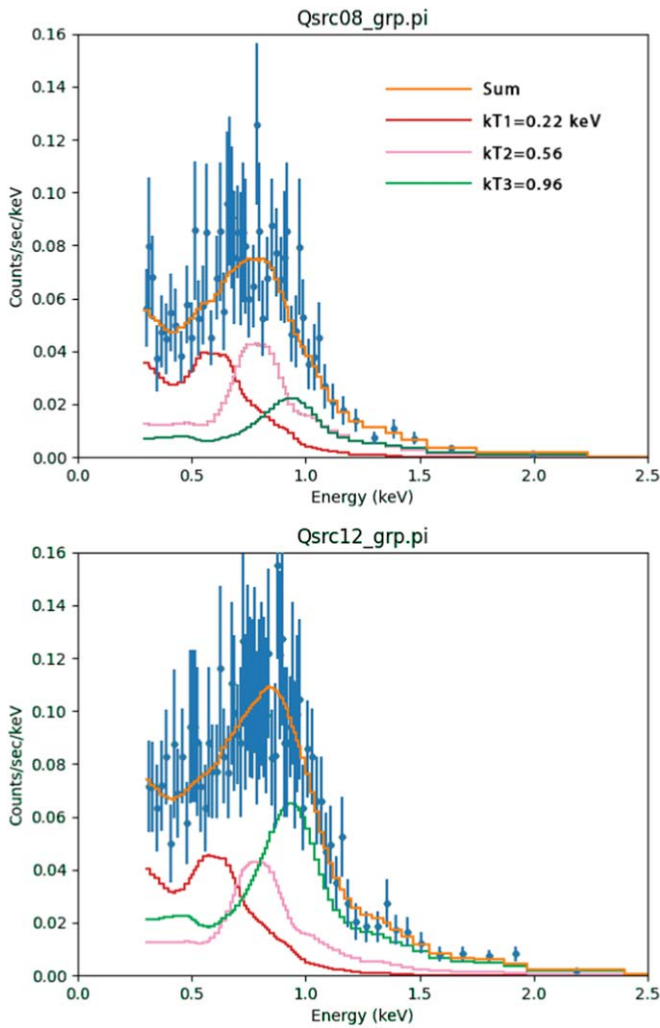


Figure 4. Sherpa fits to quiescent X-ray data from epochs 8 and 12 (near cycle minimum and maximum, respectively) with three components at fixed temperatures corresponding to active regions, active region cores, and subdued flaring regions. Note the large difference in high-temperature emission between the two epochs.

uncertainties of ± 0.05 . There were not enough counts to meaningfully investigate how abundances might vary among elements with different first ionization potentials, but the overall low abundance for Proxima’s corona is consistent with previous findings for quiescent emission in, e.g., S. Lalitha et al. (2020; $0.23^{+0.86}_{-0.12}$ from Chandra LETG observations in 2017, and $0.38^{+0.31}_{-0.16}$ from contemporaneous AstroSat measurements), and in B. J. Wargelin et al. (2017; 0.25 from XMM-Newton observations in 2009).

Using the above temperatures and abundance in fits to individual epochs, we obtained the results in Table 2, which have been adjusted for enclosed energy fractions and missing pixels by applying exposure-weighted averages of the previously calculated XRT rate corrections, ranging from 1.291 (AO 12) to 1.378 (AO 15). Statistical uncertainties from the number of counts in each spectrum are $\sim 3\%$, which we believe are relatively unimportant compared to uncertainties arising from sampling of Proxima’s highly variable emission. Estimating the “true” uncertainties of the measured fluxes is a fraught task, but longer total exposure times will generally provide a more accurate measure than shorter times. We therefore assigned

uncertainties proportional to $1/\sqrt{T_{\text{exp}}}$, with 20 ks of quiescent exposure scaled to yield 5% error. The main importance of those errors is providing the relative weights used in fitting the stellar cycle, as discussed in Section 3.

Converting the fitted fluxes in Table 2 to luminosities, we see that quiescent emission (0.3–2.5 keV) ranges from 3.2 to 4.5×10^{26} erg s^{-1} , agreeing well (keeping in mind that our criteria for quiescence are generally stricter) with levels found by S. Lalitha et al. (2020; 3.5×10^{26} from LETG, and 5.1×10^{26} from AstroSat), B. Fuhrmeister et al. (2011; 5.0×10^{26} from XMM-Newton in 2009), and M. Güdel et al. (2004; $4\text{--}28 \times 10^{26}$ from XMM in 2001). The table also shows that higher total emission is associated with a higher fraction of emission from the hottest (kT_3) component, which is particularly well illustrated in Figure 4.

2.2.2. XMM

XMM first observed Proxima in 2001 but flaring made those data unsuitable for measuring the quiescent emission level. A set of three observations in 2009 March was analyzed in B. J. Wargelin et al. (2017), and their quiescent emission level agreed well with that from the close-in-time Swift AO5 epoch. We have reanalyzed those data from the EPIC pn detector along with a set of four newer observations from 2017 and 2018 (see Table 1). ObsID 0801880301 in the latter set consisted entirely of strong flare emission and was excluded from further analysis.

We downloaded data from the XMM archive and reprocessed using the SAS software (version `xmmsas_20230412_1735-21.0.0`) with standard pn filtering. The source was extracted using a circle with 400 pixel ($20''$) radius, with background ($\sim 0.5\%$ of the source rate) taken from a region of the same size offset by $100''$. Using 200 s bins, we applied the same quiescent emission criteria as for the Swift analysis, created rate filters using the `tabgtigen` tool to generate the quiescent event files, and then made RMF and ARF files using `rmfgen` and `arfgen`.

The resulting spectra were fit using Sherpa with the same three-temperature `xsphabs` model used in the Swift XRT analysis. Abundance was likewise fixed at 0.17; freeing it resulted in values of 0.165 and 0.148 for epochs 5X and 13X, respectively, with negligible flux changes. The fitted fluxes, listed in Table 2, are a little higher than contemporaneous Swift results. XMM has longer exposure times for each epoch than Swift, but because rates in each time bin are not truly independent, particularly during flares, Swift’s snapshot sampling will tend to produce a more representative and generally slightly lower rate distribution than XMM for a given total exposure time. Net uncertainties on the XMM fluxes are not listed in Table 2 because it is not clear how their values relative to Swift should be computed.

2.2.3. Chandra HRC-I

Chandra’s High Resolution Camera for Imaging (HRC-I) microchannel plate detector observed Proxima in 2012, 2015, and 2016. All measurements were made off axis in “Next In Line” mode with restricted telemetry; our analysis removed periods of telemetry saturation caused by background flares, which occurred at the beginning and end of the 2015 observation, and during several intervals in the last two of the four observations in 2016 (see Figure 5). There were also short periods in 2012 and 2016 when Swift and the HRC

Table 2
Fit Results for Quiescent X-Ray Emission

Epoch	Total Flux	kT_1 Fraction	kT_2 Fraction	kT_3 Fraction	Source Counts	Average MJD	Quies. Exp. (s)
5X	2.336	0.480	0.258	0.263	48621	54902	37962
05	1.952 ± 0.096	0.414	0.305	0.282	1369	55167	19213
8C	1.22 ± 0.24		Fixed to Epoch 08 values		6602	56093	24317
08	1.564 ± 0.075	0.420	0.356	0.224	1089	56083	19840
12Ca	4.55 ± 1.07		Fixed to Epoch 12 values		4809	57365	17551
12	2.235 ± 0.113	0.322	0.239	0.438	1497	57630	18048
12Cb	2.24 ± 0.51		Fixed to Epoch 12 values		3332	57682	17889
13	2.133 ± 0.118	0.337	0.286	0.377	1229	58029	15134
13X	2.460	0.482	0.291	0.227	40539	58053	29563
15	2.104 ± 0.071	0.416	0.237	0.347	2967	58620	41158
15D	2.168 ± 0.151	0.345	0.322	0.332	768	58861	9482
16	1.614 ± 0.081	0.455	0.178	0.367	1049	59038	18502
17	1.618 ± 0.086	0.376	0.292	0.332	957	59381	16266

Note. Fluxes are for $E = 0.3\text{--}2.5$ keV in units of 10^{-12} erg cm $^{-2}$ s $^{-1}$ (equivalent to a luminosity of 2.0283×10^{26} erg s $^{-1}$ at Proxima’s distance). Fluxes are adjusted for vignetting and enclosed energy fractions and (for Swift) missing pixels but do not include adjustments for rotational modulation (see Section 2.4 and Table 3). Uncertainties for Chandra HRC fluxes (the “C” epochs) represent normalization statistics only (see the text). XMM flux uncertainties (“X” epochs) from counting statistics are very small; systematic uncertainties are difficult to estimate but larger than for Swift measurements.

observed Proxima simultaneously, allowing rates observed by the HRC, which has essentially no energy resolution, to be normalized against fluxes detected by Swift.

For the 2012 and 2016 observations, made 15'0 off axis, the source was extracted using an ellipse with radii of 207 and 315 pixels (27.''3 and 40.''2), background from a surrounding elliptical annulus of equal area. The 2015 observation, made 25.'62 off axis, used source radii of 460 and 740 pixels and a surrounding background annulus. Periods of quiescent emission were again identified using 10th–60th percentile rates, although without energy filtering, and quiescent event files were constructed for each of the three epochs.

Spectral fitting cannot be used for HRC data, so we used the CIAO `srcflux` tool. We specified the same source models used for Swift and XMM fitting, except that the relative normalizations of the three temperature components were necessarily fixed, leaving only the overall flux as a free parameter. For the HRC 8C epoch, we used the kT fractions from Swift epoch 8 (see Table 2), and Swift epoch 12 fractions for 12Ca and 12Cb. Swapping the distributions changed the derived fluxes by only $\sim 8\%$.

`srcflux` results are for the HRC energy range of 0.1–10 keV, so to estimate the 0.3–2.5 keV flux, we used the Sherpa `calc_energy_flux` tool on Swift fitting results from epochs 8 and 12, deriving scaling factors of 0.675 and 0.692, respectively. The resulting fluxes for epochs 8C, 12Ca, and 12Cb are then 6.10, 11.53, and 4.74×10^{-12} erg s $^{-1}$ cm $^{-2}$, which are all much larger than the fluxes measured by Swift or XMM. It is difficult to understand these results, but they are presumably related to the very different energy dependent responses of the HRC microchannel plate detector versus the Swift and XMM CCDs, a subject discussed by T. R. Ayres et al. (2008) to explain the initially puzzling “fainting” of α Cen A. The CCD detectors of Swift and XMM have effectively no useful response below ~ 300 eV because of electronic noise, but the HRC response extends to about 100 eV and therefore detects emission within the C-K transmission window below 284 eV (see Figure 12 of B. J. Wargelin et al. 2017).

The short periods of simultaneous observation by Swift and the HRC, however, provide the opportunity for direct

cross-calibration; see the orange time bands in Figures 2 and 5. The first of three periods of overlap during Swift epoch 8 and HRC epoch 8C/ObsID 14276 in 2012 occurred when the source was flaring, and we therefore do not use it in our normalizations for quiescent emission; the total overlap during the other two intervals was 1103 s. The total during the 2016 overlap (Swift epoch 12 and HRC epoch 12C/ObsID 19793) was 458 s.

Since the quiescent Swift rates and fluxes have already been measured, we can use HRC versus Swift rates during the times of overlap to convert the HRC quiescent rates to fluxes with

$$F_{Hq} = R_{Hq} \frac{R_{Ho} F_{Sq}}{R_{So} R_{Sq}} \quad (3)$$

where F is flux and R is measured event rate, and the subscripts H , S , o , and q denote HRC, Swift, overlap, and quiescence, respectively. Applying the 2012 normalization to HRC epoch 8C and the 2016 normalization to epochs 12Ca and 12Cb, we obtain fluxes of (1.22 ± 0.24) , (4.55 ± 1.07) , and $(2.24 \pm 0.51) \times 10^{-12}$ erg s $^{-1}$ cm $^{-2}$ where the listed uncertainties are from only the normalization statistics and do not include systematic errors.

The middle measurement (epoch 12Ca) is clearly an outlier, and the light curve in Figure 5 suggests that observation was dominated by flaring with only a few very brief periods of quiescent emission. The measurement from epoch 12Cb should be the most reliable, since its four separate observations come closest to approximating the nearly random sampling of Swift’s many snapshots. Epoch 8 (ObsID 14276) comprises only one observation, but it is relatively long, does not have the strong flares seen in ObsID 17377, and includes substantial periods with rates near the overall minimum, so we believe it is probably a good indicator of the quiescent emission level.

Excluding the highly suspect epoch 12Ca, the cross-calibrated HRC results are indeed consistent with those from Swift (see Table 2). We note, however, that the normalization factors R_{Ho}/R_{So} for epochs 8 and 12 differ by a factor of 2.697 ± 0.807 , a surprisingly large difference. Like the puzzling `srcflux` results, this is probably due to significant




Figure 5. X-ray light curves for XMM pn (epochs 5X and 13X) and Chandra HRC (epochs 8C, 12Ca, and 12Cb) observations. Dotted horizontal lines mark quiescent rate limits. Vertical orange bars denote periods of overlap with Swift observations (see Figure 2), allowing for flux normalizations. The first overlap was not used for normalization, because it was during a flaring period.

differences in spectral distributions during the two epochs, presumably at low energies that the Swift CCD is not sensitive to.

There are two other sets of X-ray observations since the work done in B. J. Wargelin et al. (2017), four with Chandra HRC-S/LETG in 2017, and two with ACIS-S/HETG in 2019. The use of the LETG and HETG gratings, however, decreases rates by nearly an order of magnitude, and ACIS/HETG has very little of its sensitivity below ~ 1 keV. Because of their very low count rates, we do not include these observations in our work here.

2.3. UVOT Data

Analysis of Swift UVOT data is similar to that for the XRT, but with additional steps to correct for contamination by nearby stars, and for temporal and spatial variations in quantum efficiency (QE). After splitting each observation into its component snapshots, we extracted 240×160 pixel ($120.''5 \times 80.''3$) ellipses around Proxima to reduce file sizes by 96% for subsequent steps. Source count rates are roughly 100 times higher than in the XRT, so we used ~ 100 s binning instead of ~ 400 s to divide up the snapshots, which is short enough that there was no need to manually separate pre-flare and flare intervals. A few snapshots were removed or time-filtered because good time intervals were incorrect (usually because the spacecraft had not finished slewing), the target fell outside the field of view, or exposures were too short ($< 100/\sqrt{2}$ s).

The intrinsic UVOT resolution is $2.''5$ FWHM (A. A. Breeveld et al. 2010), but there is usually a little pointing drift during observations, with errors during snapshots sometimes reaching $4''$. We can recapture the intrinsic resolution by centroiding the source in each 100 s time bin, which enables tighter extraction regions to be used, resulting in less contamination by nearby sources. Consistent source sizes also permit more accurate corrections for source contamination when it is unavoidable, which is a larger concern than with X-ray observations because Proxima does not completely dominate its local field in the

UVOT/W1 band as it does in X-rays, and its cycle amplitude in W1 is smaller.

Our choice of UVOT source extraction radius was informed by the Swift website’s “UVOT Aperture Correction” discussion.⁴ S/N is maximized with extraction radii equal to the PSF FWHM, but Proxima is relatively bright, and we are more interested in the relative accuracy of event rates, which primarily means reducing sensitivity to centroiding errors and PSF variations. We therefore chose a larger extraction radius of $5''$ (10 pixels), which allows for shifts in the source center of ± 1 pixel with negligible change in the derived source rate. For background, we used the recommended annulus with radii of $27.''5$ and $35''$ centered on the source. Event pileup effects are negligible for source rates of < 10 ct s $^{-1}$, comfortably above Proxima’s typical quiescent rates of 6–8 ct s $^{-1}$.

As noted in Section 2.1, Proxima’s proper motion is carrying it through a crowded region of the sky, and even with UVOT’s resolution of $2.''5$, some source contamination is unavoidable. In order to correct for this, we stacked images of all W1 observations to improve S/N and quantify the brightness of each nearby source, and then determined the fraction of each star’s emission that fell within Proxima’s (moving) source extraction region.

To align images for stacking, we first computed Proxima’s well-determined date-dependent R.A. and decl. for each snapshot and measured its iteratively centroided source position (in pixel coordinates) in each 100 s time bin. An image file was then created for each time bin, centered on pixel coordinates corresponding to a common reference position in R.A. and decl., near the center of Proxima’s motion over 12 yr; alignment of each image was accurate to ± 0.5 pixels in both dimensions, limited by the necessity to use integer pixel coordinates for image extractions. Finally, the resulting images were summed using `dmimgcalc` into composite files for individual and multiple epochs. The combined image of all observations is shown in Figure 6. Analysis showed that only

⁴ https://swift.gsfc.nasa.gov/analysis/threads/uvot_thread_aperture.html




Figure 6. Summed UVOT/W1 data (highly stretched) from AOs 5 to 17. Proxima’s proper motion ($47.''56$) over that time interval is shown in red. Green circles (radius of 6 pixels, or $3''$) mark sources that can interfere with Proxima or its background-region annuli; UVOT PSF is $\sim 2.''5$ FWHM. Rate corrections for interference were all below 2%.

Proxima has detectable proper motion so the other stars’ positions can be treated as fixed.

From those combined images, we determined the absolute position and average count rate of all stars that could interfere with Proxima’s source or background regions. Their positions relative to Proxima during each observation were then calculated, along with the fraction of each star’s counts that would fall within Proxima’s 10 pixel-radius source region. That fraction was calibrated separately using a summed stellar image and measuring the extracted counts while varying the distance between the star and the extraction circle. Corrections for source contamination were $\sim 2\%$ for AOs 15, 16, and 17, $<1\%$ for AOs 8 and 13, and $<0.4\%$ for AOs 5 and 12. Uncertainties from those corrections, assuming a 1 pixel error in Proxima’s source centroid, were all $<1\%$. Proxima’s annular background regions were re-extracted when necessary to exclude 10 pixel regions around any contaminating sources, and net rates were adjusted for the slow decline in UVOT/W1 QE (about 30% over the last 17 yr; CALDB file `swusenscorr20041120v006.fits`), using QE at the beginning of 2005 as the baseline.

One final UVOT QE consideration is the small scale sensitivity (SSS) issue,⁵ in which a small fraction of the detector area has significantly lower QE than the rest. In the W1 band, the QE can be up to 17% low over several percent of the central $5' \times 5'$ of the field (SWIFT-UVOT-CALDB-17-02).⁶ We wrote a program that reads the `swulssens20041120v003.fits` CALDB file and produces images of where SSS regions fall with respect to our source and background extractions (J. D. Slavin 2024), and found 12 snapshots with affected source regions. We discarded only the five of them that had more than 5% of their counts in the SSS areas, corresponding to no more than 1% rate suppression. SSS affected 22 snapshots’ background regions, but only five of them over more than 5% of their area. For those five, this would introduce rate errors of no more than 0.1%, but we conservatively re-extracted the background regions, excluded the SSS zones, and recomputed the net rates.

In the end, 102 time bins (3.9% of the total) were removed because of SSS within Proxima’s source extraction regions, leaving 2486 bins, and 143 of those bins required adjustments

⁵ https://swift.gsfc.nasa.gov/analysis/uvot_digest/sss_check.html

⁶ https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/swift/docs/uvot/uvotcaldb_sss_02b.pdf

Table 3
Rates for Quiescent UVOT/W1 Emission

Epoch	W1	Quiesc.	Rot. Mod.	X-Ray Rot.
	Rate	Exp. (s)	Adj. (%)	Adj. (%)
05	6.54 ± 0.18	10493	0	0
08	6.19 ± 0.24	19745	0	0
12	6.87 ± 0.24	17297	-0.19	+0.56
13	7.19 ± 0.23	13830	+0.26	-0.42
15	7.34 ± 0.24	14743	+0.44	-2.04
15D	6.82 ± 0.18	9453	+0.85	+4.11
16	7.02 ± 0.26	18626	+0.42	-2.69
17	6.74 ± 0.22	14297	-1.17	-2.90

Note. W1 rates include corrections for contamination by nearby stars and time-dependent QE; rotational modulation adjustments are listed separately, for both W1 and X-ray. (Epochs 5 and 8 did not have enough contemporaneous optical data to analyze modulation effects.) W1 rate uncertainties are scaled in proportion to the inverse square root of the quiescent exposure time such that 20 ks yields 2% error.

to their background regions. As was done for the XRT, we then selected bins with rates between the 10th and 60th percentiles to determine the quiescent rate averages for each epoch (see Figure 3 and Table 3). Similarly to how uncertainties were assigned to X-ray fluxes, we set errors to scale as $1/\sqrt{T_{\text{exp}}}$, with 20 ks of quiescent exposure yielding 2% error.

For completeness, we note that the UVOT UV filters (W1 and W2) have nonnegligible transmission beyond their central wavelength ranges.⁷ To estimate contamination in the W1 band, we rescaled a quiescent UV grism spectrum presented in Figure 4 of B. J. Wargelin et al. (2017; primarily taken from three snapshots in observation 009, with small pieces from other observations to fill in wavelength gaps) using the effective areas for the grism and W1 filter from the CALDB, and calculated that $\sim 16\%$ of the observed W1 signal comes from wavelengths with $\lambda > 3300 \text{ \AA}$. As noted in that Figure, an increasing fraction of the observed grism spectrum arises from higher orders toward longer wavelengths, so the actual contamination is less. Its main effect is to increase the bias level of the observed W1 signal and therefore slightly reduce the apparent relative amplitude of W1 cycle variations.

2.4. Correlations and Corrections: Rotational Modulation

In addition to uncertainties in our determination of quiescent rates arising from the intrinsically variable nature of emission from a flare star, which we addressed by using many short observations and 10th–60th percentile sampling, there is also the issue of systematic error caused by biased sampling of the stellar surface. Emission from the star is not spatially uniform—which is of course the reason for the rotational modulation seen in the optical light curve (Figure 7)—and our Swift observations may be capturing a nonrepresentative sample of that emission despite a fairly even cadence during each epoch and typically monitoring over at least one full rotation period. As described below, we therefore studied the correlation of optical, UV, and X-ray emission variations to see if adjustments are needed for the Swift measurements, as well as determine if spot darkening or faculae brightening is the primary driver of optical modulation.

⁷ https://swift.gsfc.nasa.gov/analysis/uvot_digest/redleak.html




Figure 7. Optical data (with corrections for stellar contamination) and GPR fit showing rotational modulation, with the dates of Swift observations marked by gray shading. The dotted curve is from a sliding boxcar average of the GPR fit over the rotation period; subtraction of the dotted curve yields the rotational “residuals” used to study correlation with X-ray and UV variations (see Figure 8).

2.4.1. Gaussian Process Smoothing of Optical Data

To create a smooth continuous optical light curve that can be used to estimate optical brightness during any of the comparatively few Swift observations, we use Gaussian process regression (GPR; C. E. Rasmussen & C. K. I. Williams 2006) on the data shown in Figure 1, implementing our GP models using the `tinygp`⁸ Python package. GPR requires specification of a mean and covariance function, both of which should capture observed or expected trends within the data. Our combined light curve contains several thousand data points, and GPR scales as N^3 , so a thorough comparison of different models would be prohibitively time consuming. Instead, we attempt to make informed or otherwise justifiable choices to produce a model that interpolates between the data in an insightful way.

For the mean function, about which the stellar emission fluctuates, we use a simple constant:

$$m(t) = c, \quad (4)$$

where t is time, and c is the constant. When choosing the covariance function, we considered the following: (1) Proxima has a rotation period of ~ 84 days, with rotational modulations clearly visible in some parts of the light curve (Figure 7); (2) Proxima exhibits a (quasi)periodic stellar cycle (A. Suárez Mascareño et al. 2016; B. J. Wargelin et al. 2017); (3) residual stellar contamination may still be present after the corrections illustrated in Figure 1. With these points in mind, we use a two-component covariance function. The first component is a slightly modified version of the quasiperiodic (QP) covariance

function described in B. A. Nicholson & S. Aigrain (2022):

$$k_{QP}(t, t') = \sigma^2 \exp \left[-\sin^2 \left(\frac{\pi |t - t'|}{P} \right) - \left(\frac{|t - t'|^2}{2\ell^2} \right) \right], \quad (5)$$

where σ^2 is the variance, P is the period of the oscillation, and ℓ is the “length-scale.” If P is set to the stellar rotation period, then ℓ relates to the starspot evolution time, making this covariance function well-suited to modeling the rotational modulations of active stars. Equation (5) differs from the QP kernel described in B. A. Nicholson & S. Aigrain (2022) by a Γ factor, which we have implicitly assumed to be 1. The Γ factor, sometimes called the “harmonic complexity” parameter, is difficult to interpret physically, but larger values produce oscillations that deviate more from a simple sinusoid. In our experience, this parameter is often highly degenerate, and oscillations can still deviate significantly from a simple sinusoid when $\Gamma = 1$. For these reasons, we omit the Γ factor.

The second component is intended to capture any variability unassociated with Proxima’s rotation. For this we use the squared exponential covariance function, sometimes referred to as the radial basis function or Gaussian kernel,

$$k_{SE}(t, t') = \sigma^2 \exp \left(-\frac{|t - t'|^2}{2\ell^2} \right), \quad (6)$$

where σ and ℓ share the same definitions as in Equation (5). The hyperparameters of the squared exponential are again difficult to interpret physically, but it is well-suited to modeling smoothly varying stochastic processes.

With our GP specified, we then used nested sampling (J. Skilling 2004, 2009) to construct posterior probability distributions for our (hyper)parameters, implemented with the

⁸ <https://github.com/dfm/tinygp>

Table 4
(Hyper)Parameter Prior Probability Distributions

(Hyper)Parameter	Equation	Prior
c	(4)	$\mathcal{U}(\min(f), \max(f))$
σ	(5), (6)	$\log \mathcal{U}(\min(\sigma_f), 10\Delta f)$
P	(5)	$\mathcal{U}(80, 90)$
ℓ	(5), (6)	$\log \mathcal{U}(\min(\delta t), \Delta t)$

Note. $\mathcal{U}(a, b)$ denotes a uniform prior covering the interval a to b , and $\log \mathcal{U}(a, b)$ denotes a prior that is logarithmically uniform. f represents the flux, σ_f represents the errors on the fluxes, and t represents time. We use δ to denote the minimum (nonzero) difference between any two variables, and Δ to denote the maximum difference.

MLFriends algorithm (J. Buchner 2016, 2019) using the UltraNest⁹ (J. Buchner 2021) package. We used UltraNest’s ReactiveNestedSampler with `min_ess=1000` to yield at least 1000 posterior samples, and all other parameters left to their default values. Table 4 lists the prior probability distributions for our (hyper)parameters.

Finally, we used the maximum likelihood sample to obtain the GP posterior in Figure 7 showing Proxima’s rotational modulations, which can deviate considerably from a simple sinusoid. The posterior for the rotation period hyperparameter, P , is 85.4 ± 0.7 days, essentially the same as the 85.1 ± 1.2 day period reported by Z. A. Irving et al. (2023) but with a slightly smaller uncertainty.

2.4.2. Optical-X-Ray/UV Correlations

With a continuous model of the optical light curve, we can now compare optical brightness with the X-ray and UV intensities measured by Swift. To isolate rotational modulation from longer-term, especially cyclical, intensity variations, we subtract epoch-average rates from UV and X-ray measurements, and a sliding 84 day boxcar average from optical magnitudes. “Residual” intensities are plotted in Figure 8. Colors distinguish between cases where there was an optical measurement within 21 days (1/4 rotation period) of a Swift observation so the GPR fit should be reliable, 42 days (less reliable, but only AO 13 had a significant number of such cases), or longer (not reliable; AOs 5 and 8 were dominated by these).

To fit the optical/Swift correlations, we used 10th–60th percentile rates of the Swift data to represent its quiescent emission and varied the slope of a line through the origin, minimizing the least-squares difference between the line and the quiescent points. Only points within a 21 day optical-Swift interval were fit, to minimize the introduction of errors from the optical GPR curve.

Results are shown in Figure 8. The optical/UV correlation is quite strong, with the fit yielding

$$\Delta R_{W1} = (12.40 \pm 0.61)\Delta g. \quad (7)$$

The true significance of the correlation, however, is somewhat less because we had to estimate the optical brightness at the time of each Swift measurement using the GPR fit. Second, the measurements are not entirely independent—flares take some time to decay—and the UV rates do not follow a normal distribution since flares skew it to higher rates and we only

sample the “quiescent” core. As seen in Figure 8 (upper right panel), however, the distribution is close to normal when ignoring the high-rate tail, as we do. If we mirror the low-rate side (0th–10th percentile) to the high side (>60th percentile) and use that pseudo-Gaussian full distribution (as opposed to just the core) to calculate errors, we obtain an uncertainty of 0.97 in the slope, for a formal significance of 12.8σ .

We can then use the anticorrelation relationship between optical intensity and UV rates to adjust the latter for rotational modulation and reveal cleaner measurements of long-term intensity variations arising from the stellar cycle. As hoped for, given our efforts to sample Proxima’s emission evenly from all sides and balance any rotationally induced excess or deficit emission, net effects on quiescent UV rates for each epoch are minor, ranging from -1.17% (AO 17) to $+0.85\%$ (AO 15D). Adjustments were applied to all Swift UV measurements in AO 13, including those with optical-Swift gaps of up to one-half rotation period (42 days); rates for AOs 5 and 8 were not modified because of the lack of sufficiently contemporaneous optical measurements.

The same analysis was applied to optical/X-ray data, finding weaker anticorrelation. The fit slope was 0.346, and error calculations employing an intensity distribution with mirrored 0th–10th percentiles (as was done for the UV data) indicate an uncertainty of 0.095, or a 3.6σ result. Net effects on average quiescent fluxes are somewhat larger than in the W1 band but still small, ranging from -2.9% to $+4.1\%$ (see Table 3). One would expect that W1 and X-ray adjustments should roughly track each other, but even when ignoring Epoch 15 (because of the difference in W1 and X-ray sampling; see Table 1), correlation is poor. We also found that using a slightly different, “peakier,” but similarly justifiable GPR function yielded only a 1.6σ optical/X-ray correlation. (Negligible difference in optical/UV correlation was seen using that function.) Because of the relatively low and uncertain significance of the X-ray/optical correlation, we consider any adjustments of questionable value and use uncorrected X-ray intensities in subsequent analysis.

A possible explanation for the relatively weak correlation of X-ray emission with optical/UV is that X-rays come from the optically thin and spatially extended corona whereas optical (photospheric) and UV emission (chromospheric) are essentially surface phenomena and therefore more subject to rotational modulation. X-ray modulation only occurs for the fraction of emission that originates close to the stellar surface, and would be especially weak in systems where the corona is relatively large compared to the stellar diameter (as in M stars), and also those with smaller inclinations (more “pole on”); a recent paper (B. Klein et al. 2021) reports Proxima’s inclination as $47^\circ \pm 7^\circ$.

3. Results

3.1. Lomb–Scargle Period Analysis

With the final adjusted values for the optical, UV, and X-ray measurements in hand, we can examine Proxima’s long-term behavior and compare it with previous reports of a 7 yr stellar cycle. We employed a floating-mean Lomb–Scargle (L–S) periodogram for the analysis (J. D. Scargle 1982) with false-alarm probabilities (FAPs) computed following J. D. Scargle (1982) and J. H. Horne & S. L. Baliunas (1986). The result (Figure 9) shows a dominant peak at 83.3 days corresponding

⁹ <https://johannesbuchner.github.io/UltraNest/>




Figure 8. Residual rates for UV/W1 (left) and XRT (center) vs. residual optical intensities for the GPR fit (excluding AOs 6 and 7). Upper limits on vertical axes are 95th percentiles. Solid black lines show least-squares fits to data within the 10th–60th percentiles (bounded by dotted lines) and only when an optical measurement was made within 21 days of a Swift measurement (black points). Cyan lines show 1σ uncertainties on the fits. Right panels show binned rates using only the black points, with and without adjustments for rotational modulation derived from the fits. Dotted vertical lines mark the 10th and 60th rate percentiles after adjustments for rotational modulation.

to the rotation period, with multiple nearby peaks and perhaps parts of the cluster around 110 days likely deriving from various manifestations of differential rotation and starspot movement, formation, and dissipation. Some peaks, though, we identify as beat effects between rotation and the year and half year observing windows and the cycle period. Other peaks are from the window function, harmonics, and aliases, except for the peak at 2919 days (7.99 yr), which we take to be the stellar cycle period. The broad bump at several thousand days is likely caused by the ~ 23 yr data interval.

3.2. Stellar Cycle Period and Amplitudes

The X-ray and UV data sets have too few points for useful L–S analysis, but they and the optical data are suitable for sine fits to the multiyear cycle (provided that sinusoidal behavior is sufficiently stable over multiple cycles; see below). ASAS-3 data yield a period of 7.34 ± 0.37 yr with an amplitude of 0.019 mag (3.8% peak-to-peak cycle variation); adding the relatively sparse ASAS-4 data increases the period to 10.7 yr but with a very large increase in uncertainty (± 4.7 yr). ASAS-SN data yield $P_{\text{rot}} = 7.88$ yr and half the amplitude found for ASAS-3, while the full data set, spanning three cycles, yields $P_{\text{rot}} = 7.99 \pm 0.17$ yr (the same as the L–S periodogram result) and an amplitude of 0.0107 ± 0.0006 mag (Figure 10).

As seen in the middle panel of that Figure, the HRC X-ray results are consistent with Swift results, but their uncertainties are quite large, even though the error bars shown are only statistical and do not include systematic errors (primarily from sampling). The two XMM points are also fairly consistent with Swift, particularly when considering that they are expected to be a bit higher than the Swift points (see discussion at end of Section 2.2.2), and their systematic errors are likely larger than those for Swift. Because of the general consistency of results




Figure 9. Floating-mean Lomb–Scargle periodogram of the full optical data set showing periodogram power vs. period. The strongest rotational peak (83.3 days) and the cycle peak (2919 days) are labeled, as is the data time span (8464 days).

from XMM, HRC, and Swift, and the larger, less-well-understood uncertainties on the XMM and HRC points, we only use the Swift data in our cycle fits.

Independent fits to the Swift UV and X-ray data gave periods of 11.84 and 9.74 yr, respectively. A simultaneous fit to both sets of data gave poor results for the UV, with a visibly suppressed amplitude. To apply equal weights to both wave bands, we therefore scaled the uncertainties to yield reduced χ^2 of 1 in separate fits, which required UV errors to be a factor of 1.17 smaller than listed in Table 3 and X-ray errors to be 1.25 times larger than in Table 2. The combined fit then gave

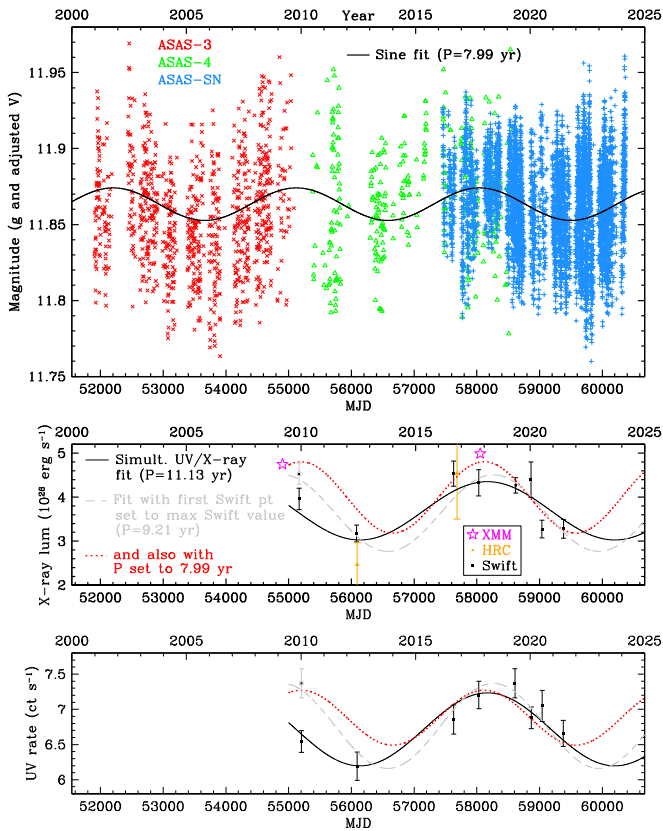


Figure 10. Cycle fits for optical, X-ray, and UV data. Optical brightness scale is inverted to aid comparison with the anticorrelated X-ray and UV intensities. X-ray and UV data were fitted simultaneously, with uncertainties adjusted to provide equal statistical weight. UV/X-ray period agreement with the optical data is much improved if the first (~ 2010) data points are raised to match peak values near the ~ 2018 maximum (see text). XMM and Chandra HRC data points are not used for fitting.

$P_{\text{rot}} = 11.13 \pm 0.53$ yr, with $L_X^{\text{max}}/L_X^{\text{min}} = 1.44$, and a cycle amplitude of 7.7% (max/min = 1.167) for the W1 band (see black curves in the bottom panels of Figure 10).

The fitted UV/X-ray cycle period of 11 yr is longer than the optical 8 yr period, but that from UV/X-ray data is much more uncertain, mostly because of the fit’s sensitivity to the first point in each band, which is because of the limited span of the Swift measurements and data sparseness prior to 2016. Based on the more reliable optical 8 yr period, the UV/X-ray cycle should be near its maximum in 2010. We suggest that the high-energy AO 5 intensities are lower than they “should be” either because of statistical fluctuations or more likely because the cycle amplitude at that time is truly lower than at later times; such cycle variability is common for the Sun and other stars. We also note from the previously mentioned separate fits to ASAS and ASAS-SN data that Proxima’s optical cycle amplitude was twice as large for earlier times (prior to c. 2010) as later (2017 and after), plausibly consistent with (opposite sense) cycle amplitude variations at UV/X-ray energies.

If one therefore assumes that the AO 5 measurements around 2010 were indeed made near a cycle maximum and sets the UV and X-ray intensities to match those around the 2018 cycle maximum to accommodate the implicit assumption of constant amplitude in sinusoidal fitting, the simultaneous UV/X-ray fit then yields $L_X^{\text{max}}/L_X^{\text{min}} = 1.63$, W1 amplitude of 8.9%, and a significantly shorter period of 9.21 ± 0.33 yr (dashed gray

curves in Figure 10). A fit using the altered AO 5 values and freezing the period at the optical fit’s 7.99 yr reduces $L_X^{\text{max}}/L_X^{\text{min}}$ to 1.55 and the W1 amplitude to 5.7% (dotted red curve). Figure 10 also shows that the optical and UV/X-ray cycles are opposite in phase, as one would expect, given the anticorrelation seen in rotational modulation.

In summary, from the 23 yr of optical data, we find a well-constrained cycle period of 8.0 yr, with an average amplitude of 0.011 mag (0.022 mag peak-to-peak) that has declined from 0.019 mag in the 2000s. The cycle is less well constrained in the W1 and X-ray bands but consistent with an 8 yr period. Quiescent X-ray luminosity averages 3.7×10^{26} erg s $^{-1}$ with $L_X^{\text{max}}/L_X^{\text{min}} \sim 1.5$ over the past decade, and perhaps a little less prior to c. 2013. Cycle amplitude in the UV W1 band is $\sim 8\%$ ($\sim 17\%$ peak-to-peak), and like X-ray emission, was likely weaker in the prior cycle.

3.3. X-Ray Cycle Amplitudes and Rossby Number

Although only seven stars had measured X-ray cycles at the time, B. J. Wargelin et al. (2017) noted that cycle amplitudes were approximately proportional to Rossby number ($\text{Ro} = P_{\text{rot}}/\tau_C$, where τ_C is the timescale for convection), and that this correlation held even for fully convective Proxima. This relationship is similar to the finding by N. J. Wright & J. J. Drake (2016) that the $L_X/L_{\text{bol}} \propto \text{Ro}^{-2.7}$ rotation–activity relationship for partially convective stars below the saturation regime (N. J. Wright et al. 2011) also applies to Proxima and three other fully convective stars. Some of the seven stars now have updated amplitudes, and an additional three stars have had their X-ray cycles measured since then (see Table 5): τ Boo (M. Mittag et al. 2017), ϵ Eri (M. Coffaro et al. 2020), and the rapid rotator, AB Dor (G. Singh & J. C. Pandey 2024).

A number of F stars such as τ Boo have recently been discovered from Ca II measurements to have very short cycles of under 1 yr. M. Mittag et al. (2019) found cycles ranging from 180–309 days in three or four stars, and τ Boo has an even shorter cycle of ~ 120 days (M. Mittag et al. 2017). There are also more than 30 XMM observations of τ Boo, with at least one per year from 2000–2011. X-ray data are unfortunately not dense enough to provide reliable independent measurements of cycle periodicity, but the range of intensities does provide a good indication of cycle amplitude. Photometric (V-band) data likewise provide a high-confidence measurement of cycle period for AB Dor (G. Singh & J. C. Pandey 2024), but X-ray data (primarily from XMM) are again too sparse to support confident period analysis, while abundant enough to provide an upper limit on X-ray cycle amplitude.

In addition to adding new stars, we have also updated all stars’ values for τ_C , and include consideration of recent suggestions that dynamos in active, partially convective stars may not be solely driven by the “local” tachocline dynamo but also, or perhaps even instead, by a “global” full convection zone dynamo; see discussion in Z. A. Irving et al. (2023) and references therein. To determine τ_C for all of those stars, we follow the procedure of Z. A. Irving et al. (2023), and estimate T_{eff} from $V - K_S$ (M dwarfs) or $B - V$ (all others) from the tables in M. J. Pecaut & E. E. Mamajek (2013). These T_{eff} were then used to interpolate local τ_C and global τ_C from N. R. Landin et al. (2023). For the lowest-mass stars (nearly fully and fully convective objects; here, only Proxima), because of difficulties in detailed interior models (see Z. A. Irving et al. 2023, for discussion), we prefer the scaling relations of

Table 5
Stars with Published X-Ray Cycles

Star Name	Type	$B - V$ ($V - K_s$)	P_{cyc} (yr)	P_{rot} (day)	τ_C^a (day)	$\text{Ro} =$ P_{rot}/τ_C	$L_X^{\text{max}}/$ L_X^{min}	References
τ Boo	F6V	0.508	0.33, 11.6 ^b	3.05	7.6 ^c	0.40	1.77	M. Mittag et al. (2017)
ι Hor	F8V	0.561	1.6	8.2	11.8, 23.5	0.695, 0.349	1.9	J. Sanz-Forcada et al. (2019)
Sun	G2V	0.653	11	25.4	17.2	1.48	3.91	T. R. Ayres (2020)
HD 81809	G1.5IV-V	0.642	8.2	40.2	22.8 ^d	1.76	4.9	S. Orlando et al. (2017)
α Cen A	G2V	0.697	19	28	19.9	1.41	2.7	T. R. Ayres (2020)
AB Dor	K0V	0.830	19.2	0.51	26.4, 53.5	0.0193, 0.0095	$\lesssim 1.4^e$	G. Singh & J. C. Pandey (2024)
α Cen B	K1V	0.902	8.4	37	29.8	1.24	3.5	T. R. Ayres (2020)
ϵ Eri	K2V	0.881	2.9	11.1	29.1, 57.9	0.38, 0.19	1.5	M. Coffaro et al. (2020)
61 Cyg A	K5V	1.158	7.3	35.4	38.0	0.93	3.0	J. Robrade et al. (2012)
Proxima	M5.5V	(6.75)	8.0	84	326.2	0.26	1.5	This work

Notes.

^a Following the method of Z. A. Irving et al. (2023); see the text for details. The second value (where present) is the full convective zone τ_C (except Proxima, where it is the only value), and the first value is the local (tachocline) τ_C .

^b From S. L. Baliunas et al. (1995).

^c τ Boo has a close, massive, hot Jupiter; P_{rot} is synchronized to P_{orbit} , possibly affecting activity.

^d R. Egeland (2018) argues that the cycling component of this binary is evolved; our scaling to account for the star's evolution makes its τ_C value less certain. See the text.

^e G. Singh & J. C. Pandey (2024) found multiple superposed cycles in AB Dor, which prohibits a firm determination of the dominant cycle's amplitude.

E. Corsaro et al. (2021), which we scaled to match the τ_C of N. R. Landin et al. (2023) at their juncture. Thus, for these stars, $\tau_C = 2.9 \times 10^7 (M/M_\odot)^{1/3} (T_{\text{eff}}/T_{\text{eff},\odot})^{-4/3}$ days, and the mass was similarly interpolated from M. J. Pecaut & E. E. Mamajek (2013) using the estimated T_{eff} .

R. Egeland (2018) found that HD 81809 is a binary whose emission is dominated by a subgiant. We determined a τ_C appropriate for its evolved state by taking the N. R. Landin et al. (2023) value for a $1.6 M_\odot$ main-sequence star and then applying a scaling factor equal to the ratio between the values listed by A. G. Gunn et al. (1998) for a star with HD 81809's current T_{eff} and for its main-sequence counterpart (their Figure 3).

Figure 11 shows the updated results for X-ray cycle amplitude versus Rossby number. The same trend is found as before, but with the additional stars, there is now a clear transition to a limiting amplitude at small Ro (as must occur since amplitudes cannot go below (1)). Proxima lies near the transition between low Ro/active and higher Ro/less active stars. The shape of the curve is again reminiscent of the rotation–activity relationship, which saturates around $L_X/L_{\text{bol}} \sim 10^{-3}$ at small Ro; with the τ_C scale used here, total X-ray emission saturates at $\text{Ro} \approx 0.11$. Thus, it would appear that cycle amplitude saturates at somewhat slower rotation: at $\text{Ro} \approx 0.4$ if active stars are dominated by a tachocline dynamo (and $\tau_{C,L}$ is appropriate, as sketched in Figure 11) or $\text{Ro} \approx 0.2$ if a full convection zone dynamo dominates (and $\tau_{C,G}$ is better).

As suggested by B. J. Wargelin et al. (2017), the correlation of high Ro with high amplitude (now with saturation at low Ro) can be plausibly explained by more active stars (smaller Ro) having a greater covering fraction of X-ray-emitting active regions even at cycle minimum, with less room available for additional emission at cycle maximum. Note that in the most extreme case, maximum amplitude is achieved with half the star active and half quiet, which has exactly half the total flux of the maximum total flux case (full star covered). This scenario is also supported, for Proxima, by a lack of I -band variation coupled with a strong trend in $V - I$ versus V and the star growing redder during the optically fainter part of its cycle,




Figure 11. X-ray cycle amplitude vs. Rossby number, using values listed in Table 5. Three stars have pairs of points, reflecting uncertainty in whether the local (tachocline) or global (full convective zone) dynamo is dominant (see Z. A. Irving et al. 2023). Amplitudes for Proxima and ϵ Eri have been slightly shifted for clarity.

implying a large, fairly even covering fraction of cool starspots (B. J. Wargelin et al. 2017).

3.4. Possible Coronal Mass Ejections

When studying light curves and the rate distributions in Figure 3, we noticed a few instances of extremely low rates. Focusing on the 10 X-ray bins with rates $< 0.02 \text{ ct s}^{-1}$ (out of 856 total for the XRT), three are from snapshots where the source is within a few pixels of dead CCD columns (so rates are suspect), and the other bins within the same snapshot are not particularly low, suggesting that the low rates are instrumental artifacts. A fourth case is from a snapshot with a single bin, with unremarkable rates in nearby snapshots.

The remaining six low-rate bins, however, are particularly interesting because:

1. The source is far away from bad pixels (so rate corrections are reliable).
2. They come in two groups (three from 31676018a and three from 31676033bP, where a, b, c... distinguish snapshots within an observation, and P and F denote pre-flare and flare intervals within a snapshot) in which all bins have low rates, indicating a real low rate and not a statistical fluke.
3. The low-rate bins are followed by large flares ($>1 \text{ ct s}^{-1}$), which occurs for only 1.8% of the 856 bins.

The three low-rate bins from 018a (dropping the 31676 prefix) contain 10 counts, 6.4σ below the 60 counts expected from the AO 8 quiescent average. Their average rate is $0.0121 \pm 0.0038 \text{ ct s}^{-1}$, followed 1.37 hr later by a decaying flare in snapshot 018b, with an initial rate of 1.46 ct s^{-1} . The three bins in 033bP average $0.0137 \pm 0.0037 \text{ ct s}^{-1}$ (10.5σ below the quiescent average), immediately followed by the remainder of the snapshot in a single 69 s bin with a rate of 1.09 ct s^{-1} .

We also examined UVOT data for similar cases, with the only notable occurrences paralleling what was seen in the XRT. For the 018a/b pairing, six of the ten 018a bins (recall the typical four UVOT time bins per XRT bin) had rates within the lowest 0.6% of the 2486 total bins, and all 11 of the 018b bins are among those with the 3% highest rates. The 12 033bP bins had rates in the 2nd–41st percentiles (average 6.364 ct s^{-1}), while the sole 033bF bin (at 70 ct s^{-1}) was the second highest in the entire Swift data set.

A plausible explanation for low rates is that a coronal mass ejection (CME) expands into the stellar corona and leaves an evacuated low-emission volume, as is seen on the Sun where CMEs can cause a reduction in total coronal emission (L. K. Harra et al. 2016). With this mechanism in mind, A. M. Veronig et al. (2021) analyzed extreme UV and X-ray archival data from other stars and found 21 candidate CMEs where flares were followed by significant decreases in coronal emission. The two cases we are examining here, however, feature low rates preceding flares, either immediately or within 1.4 hr, which is difficult to understand in the context of a CME.

Low emission immediately preceding a flare has been seen in the Sun in the UV, particularly the 171 Å band (J. P. Mason et al. 2014), but the degree of dimming was small compared to what occurred after the main flare. If the low rates seen in Proxima observations are indeed connected with CMEs, it is more likely that the low rates are preceded by flares, but that the flares occurred in the gaps (8 days and 1.2 hr, respectively) preceding the two groups of low-rate snapshots and were not seen.

4. Summary and Discussion

This study of Proxima Cen has presented an analysis of optical photometry (23 yr), Swift X-ray and UV measurements (8 epochs over a 12 yr span), and supplementary XMM and Chandra HRC X-ray measurements, finding a clear 8.0 yr stellar cycle with average amplitude of 0.011 mag (0.022 mag peak-to-peak) in the optical data and similar periodicity in the Swift data. Proxima is by far the smallest of the few stars to have \sim regular X-ray monitoring over many years, and the clarity of its cycle and consistency of results across optical, UV, and X-ray energies strongly supports the growing evidence for stellar cycles among fully convective stars.

Proxima’s X-ray cycle is ~ 1.5 times brighter at maximum than at minimum. Among the 10 stars with published X-ray cycles, a strong correlation between cycle amplitude and Rossby number is found, with amplitude decreasing toward smaller Ro before plateauing at a value not far above one.

In the UV W1 band centered around 2800 Å, cycle amplitude is $\sim 8\%$ (max/min = 1.17). Over a cycle, X-ray/UV intensity is anticorrelated with optical brightness, as is also true for rotational modulation, which has very strong optical periodicity around 84 days. We applied corrections for contamination by other stars in the optical and UV bands as Proxima moves across its relatively crowded field, yielding clean light curves for the periodicity and correlation analyses.

Significant flaring is present in the UV and X-ray bands, and quiescent emission levels were computed using the 10th–60th percentiles. Quiescent X-ray emission over a cycle averages $3.7 \times 10^{26} \text{ erg s}^{-1}$, and inspection of the Swift X-ray light curves yielded two instances of statistically significant anomalously low emission that may be associated with CMEs.

Acknowledgments





This work was supported by NASA’s Swift Guest Investigator program under grants 80NSSC17K0332, 80NSSC20K1111, and 80NSSC22K0039, and by NASA’s XMM-Newton Guest Observer program under grant 80NSSC18K0398 and the Chandra Guest Observer program under grant DD6-17086X. B.J.W. and P.R. were also supported by NASA contract NAS8-03060 to the Chandra X-ray Center. S.H.S. gratefully acknowledges additional support from NASA XRP grant 80NSSC21K0607 and NASA EPRV grant 80NSSC21K1037. Z.A.I. acknowledges support from the UK Research and Innovation’s Science and Technology Facilities Council grant ST/X508767/1.

We thank the ASAS and ASAS-SN collaborations for providing optical photometry data, the UK Swift Science Data Centre at the University of Leicester for providing Swift data and processing tools, and the Chandra X-ray Center for the CIAO and Sherpa analysis programs. We also thank E. Shkolnik for helpful discussions about CMEs. The five Chandra HRC observations used in our analysis can be accessed via doi:[10.25574/cdc.306](https://doi.org/10.25574/cdc.306).

Lastly, we acknowledge use of optical data from DR4 of the SkyMapper Southern Survey, which has been funded through ARC LIEF grant LE130100104 from the Australian Research Council, awarded to the University of Sydney, the Australian National University, Swinburne University of Technology, the University of Queensland, the University of Western Australia, the University of Melbourne, Curtin University of Technology, Monash University and the Australian Astronomical Observatory. SkyMapper is owned and operated by the Australian National University’s Research School of Astronomy and Astrophysics. The survey data were processed and provided by the SkyMapper Team at ANU. The SkyMapper node of the All-Sky Virtual Observatory (ASVO) is hosted at the National Computational Infrastructure (NCI). Development and support of the SkyMapper node of the ASVO has been funded in part by Astronomy Australia Limited (AAL) and the Australian Government through the Commonwealth’s Education Investment Fund (EIF) and National Collaborative Research Infrastructure Strategy (NCRIS), particularly the National eResearch Collaboration Tools and Resources (NeCTAR) and the Australian National Data Service Projects (ANDS).

Facilities: Swift (XRT and UVOT), ASAS, ASAS-SN, CXO (HRC-I), Skymapper, XMM (EPIC pn).

ORCID iDs

Bradford J. Wargelin  <https://orcid.org/0000-0002-2096-9586>
 Steven H. Saar  <https://orcid.org/0000-0001-7032-8480>
 Jonathan D. Slavin  <https://orcid.org/0000-0002-7597-6935>
 José-Dias do Nascimento, Jr  <https://orcid.org/0000-0001-7804-2145>

References

- Anglada-Escudé, G., Amado, P. J., Barnes, J., et al. 2016, *Natur*, **536**, 437
 Ayres, T. R. 2020, *ApJS*, **250**, 16
 Ayres, T. R., Judge, P. G., Saar, S. H., & Schmitt, J. H. M. M. 2008, *ApJL*, **678**, L121
 Baliunas, S. L., Donahue, R. A., Soon, W. H., et al. 1995, *ApJ*, **438**, 269
 Blackburn, J. K. 1995, in ASP Conf. Ser. 77, *Astronomical Data Analysis Software and Systems IV*, ed. R. A. Shaw, H. E. Payne, & J. J. E. Hayes (San Francisco, CA: ASP), 367
 Breeveld, A. A., Curran, P. A., Hoversten, E. A., et al. 2010, *MNRAS*, **406**, 1687
 Buchner, J. 2016, *S&C*, **26**, 383
 Buchner, J. 2019, *PASP*, **131**, 108005
 Buchner, J. 2021, *JOSS*, **6**, 3001
 Burrows, D. N., Hill, J. E., Nousek, J. A., et al. 2005, *SSRv*, **120**, 165
 Chabrier, G., & Küker, M. 2006, *A&A*, **446**, 1027
 Cincunegui, C., & Mauas, P. J. D. 2004, *A&A*, **414**, 699
 Cochran, W. D., & Hatzes, A. P. 1993, in ASP Conf. Ser. 36, *Planets Around Pulsars*, ed. J. A. Phillips, S. E. Thorsett, & S. R. Kulkarni (San Francisco, CA: ASP), 267
 Coffaro, M., Stelzer, B., Orlando, S., et al. 2020, *A&A*, **636**, A49
 Corsaro, E., Bonanno, A., Mathur, S., et al. 2021, *A&A*, **652**, L2
 Damasso, M., Del Sordo, F., Anglada-Escudé, G., et al. 2020, *SciA*, **6**, eaax7467
 Dikpati, M., & Charbonneau, P. 1999, *ApJ*, **518**, 508
 Egeland, R. 2018, *ApJ*, **866**, 80
 Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2009, *MNRAS*, **397**, 1177
 Faria, J. P., Suárez Mascareño, A., Figueira, P., et al. 2022, *A&A*, **658**, A115
 Freeman, P., Doe, S., & Siemiginowska, A. 2001, *Proc. SPIE*, **4477**, 76
 Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, *Proc. SPIE*, **6270**, 62701V
 Fuhrmeister, B., Lalitha, S., Poppenhaeger, K., et al. 2011, *A&A*, **534**, A133
 Gastine, T., Duarte, L., & Wicht, J. 2012, *A&A*, **546**, A19
 Güdel, M., Audard, M., Reale, F., Skinner, S. L., & Linsky, J. L. 2004, *A&A*, **416**, 713
 Gunn, A. G., Mitrou, C. K., & Doyle, J. G. 1998, *MNRAS*, **296**, 150
 Harra, L. K., Schrijver, C. J., Janvier, M., et al. 2016, *SoPh*, **291**, 1761
 Horne, J. H., & Baliunas, S. L. 1986, *ApJ*, **302**, 757
 Irving, Z. A., Saar, S. H., Wargelin, B. J., & do Nascimento, J.-D. 2023, *ApJ*, **949**, 51
 Kämpylä, P. J., Mantere, M. J., & Brandenburg, A. 2013, *GApFD*, **107**, 244
 Kent, S. M. 1985, *PASP*, **97**, 165
 Klein, B., Donati, J.-F., Hébrard, É., et al. 2021, *MNRAS*, **500**, 1844
 Kochanek, C. S., Shappee, B. J., Stanek, K. Z., et al. 2017, *PASP*, **129**, 104502
 Lalitha, S., Schmitt, J. H. M. M., Singh, K. P., et al. 2020, *MNRAS*, **498**, 3658
 Landin, N. R., Mendes, L. T. S., Vaz, L. P. R., & Alencar, S. H. P. 2023, *MNRAS*, **519**, 5304
 MacGregor, M. A., Weinberger, A. J., Loyd, R. O. P., et al. 2021, *ApJL*, **911**, L25
 Mason, J. P., Woods, T. N., Caspi, A., Thompson, B. J., & Hock, R. A. 2014, *ApJ*, **789**, 61
 Mayor, M., Pepe, F., Queloz, D., et al. 2003, *Msngr*, **114**, 20
 Mignon, L., Meunier, N., Delfosse, X., et al. 2023, *A&A*, **675**, A168
 Mittag, M., Robrade, J., Schmitt, J. H. M. M., et al. 2017, *A&A*, **600**, A119
 Mittag, M., Schmitt, J. H. M. M., Hempelmann, A., & Schröder, K. P. 2019, *A&A*, **621**, A136
 Nicholson, B. A., & Aigrain, S. 2022, *MNRAS*, **515**, 5251
 Onken, C. A., Wolf, C., Bessell, M. S., et al. 2024, *PASA*, **41**, e061
 Orlando, S., Favata, F., Micela, G., et al. 2017, *A&A*, **605**, A19
 Pecaui, M. J., & Mamajek, E. E. 2013, *ApJS*, **208**, 9
 Pojmanski, G. 1997, *AcA*, **47**, 467
 Pojmanski, G. 2002, *AcA*, **52**, 397
 Rasmussen, C. E., & Williams, C. K. I. 2006, *Gaussian Processes for Machine Learning* (Cambridge, MA: MIT Press)
 Robrade, J., Schmitt, J. H. M. M., & Favata, F. 2012, *A&A*, **543**, A84
 Roming, P. W. A., Kennedy, T. E., Mason, K. O., et al. 2005, *SSRv*, **120**, 95
 Rüdiger, G., Elstner, D., & Ossendrijver, M. 2003, *A&A*, **406**, 15
 Sanz-Forcada, J., Stelzer, B., Coffaro, M., Raetz, S., & Alvarado-Gómez, J. D. 2019, *A&A*, **631**, A45
 Scargle, J. D. 1982, *ApJ*, **263**, 835
 Shappee, B. J., Prieto, J. L., Grupe, D., et al. 2014, *ApJ*, **788**, 48
 Singh, G., & Pandey, J. C. 2024, *ApJ*, **966**, 86
 Skilling, J. 2004, in AIP Conf. Ser. 735, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), 395
 Skilling, J. 2009, in AIP Conf. Ser. 1193, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. P. M. Goggans & C.-Y. Chan (Melville, NY: AIP), 277
 Slavin, J. D. 2024, *Python Routines to Aid in Evaluating the Effects of Small Scale Sensitivity for Swift UVOT Sources*, Zenodo, doi:10.5281/zenodo.13356340
 Suárez Mascareño, A., Rebolo, R., & González Hernández, J. I. 2016, *A&A*, **595**, A12
 Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, *PASP*, **130**, 064505
 Veronig, A. M., Odert, P., Leitzinger, M., et al. 2021, *NatAs*, **5**, 697
 Wargelin, B. J., Saar, S. H., Pojmański, G., Drake, J. J., & Kashyap, V. L. 2017, *MNRAS*, **464**, 3281
 Windhorst, R. A., Burstein, D., Mathis, D. F., et al. 1991, *ApJ*, **380**, 362
 Wright, N. J., & Drake, J. J. 2016, *Natur*, **535**, 526
 Wright, N. J., Drake, J. J., Mamajek, E. E., & Henry, G. W. 2011, *ApJ*, **743**, 48
 Yadav, R. K., Christensen, U. R., Wolk, S. J., & Poppenhaeger, K. 2016, *ApJL*, **833**, L28