A Study of Flares in the Ultra-Cool Regime from SPECULOOS-South

C. A. Murray[1]⋆, D. Queloz[1], M. Gillon[2], B. O. Demory[3], A. H. M. J. Triaud[4], J. de Wit[5],
A. Burdanov[5], P. Chinchilla[2,6], L. Delrez[2,7], G. Dransfield[4], E. Ducrot[2], L. J. Garcia[2],
Y. Gómez Maqueo Chew[8], M. N. Günther[9]†, E. Jehin[7], J. McCormac[10], P. Niraula[5],
P. P. Pedersen[1], F. J. Pozuelos[2,10], B. V. Rackham[5,9]‡, N. Schanche[3], D. Sebastian[4],
S. J. Thompson[1], M. Timmermans[2], R. Wells[3]

[1]*Cavendish Laboratory, JJ Thomson Avenue, Cambridge CB3 0HE, UK*
[2]*Astrobiology Research Unit, University of Liège, Allée du 6 août, 19, 4000 Liège (Sart-Timan), Belgium*
[3]*University of Bern, Center for Space and Habitability, Gesellschaftsstrasse 6, 3012 Bern, Switzerland*
[4]*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom*
[5]*Department of Earth, Atmospheric and Planetary Sciences, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*
[6]*Instituto de Astrofísica de Canarias (IAC), Calle Vía Láctea s/n, 38200, La Laguna, Tenerife, Spain*
[7]*Space sciences, Technologies and Astrophysics Research (STAR) Institute, University of Liège, Belgium*
[8]*Instituto de Astronomía, Universidad Nacional Autónoma de México, Ciudad Universitaria, Ciudad de México, 04510, México*
[9]*Department of Physics, and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[10]*Department of Physics, University of Warwick, Coventry, CV4 7AL, UK*

ABSTRACT

We present a study of photometric flares on 154 low-mass ($\leq$ 0.2M$_\odot$) objects observed by the SPECULOOS-South Observatory from 1st June 2018 to 23rd March 2020. In this sample we identify 85 flaring objects, ranging in spectral type from M4 to L0. We detect 234 flares in this sample, with energies between $10^{29.2}$ and $10^{32.7}$ erg, using both automated and manual methods. With this work, we present the largest photometric sample of flares on late-M and ultra-cool dwarfs to date. By extending previous M dwarf flare studies into the ultra-cool regime, we find M5–M7 stars are more likely to flare than both earlier, and later, M dwarfs. By performing artificial flare injection-recovery tests we demonstrate that we can detect a significant proportion of flares down to an amplitude of 1 per cent, and we are most sensitive to flares on the coolest stars. Our results reveal an absence of high-energy flares on the reddest dwarfs. To probe the relations between rotation and activity for fully convective stars, we extract rotation periods for fast rotators and lower-bound period estimates of slow rotators. These rotation periods span from 2.2 hours to 65 days, and we find that the proportion of flaring stars increases for the very fastest rotators. Finally, we discuss the impact of our flare sample on planets orbiting ultra-cool stars. As stars become cooler, they flare less frequently; therefore, it is unlikely that planets around the very reddest dwarfs would enter the 'abiogenesis' zone or drive visible-light photosynthesis through flares alone.

Key words: stars: flare – stars: rotation – planet-star interactions – planets and satellites: terrestrial planets

1 INTRODUCTION

In the search for planets capable of supporting life, ultra-cool dwarfs (UCDs) make compelling hosts. In our local stellar neighbourhood

UCDs are plentiful (Chabrier 2003; Henry 2004) and predicted to host large numbers of planets (Cantrell et al. 2013; Dressing & Charbonneau 2015; Ballard & Johnson 2016). Moreover, their small sizes and low temperatures make it easier to detect Earth-sized, habitable-zone planets around these objects, and to probe those planets' atmospheres for biosignatures (Kaltenegger & Traub 2009; Seager et al. 2009; de Wit & Seager 2013), than for any other type of star.

⋆ E-mail: cam217@cam.ac.uk
† Juan Carlos Torres Fellow
‡ 51 Pegasi b Fellow

Despite their promise, serious questions remain about the habitability of planets around red dwarfs. Several authors find it unlikely that these extremely cool, red stars would provide their planets with enough UV photons to initiate specific prebiotic chemistry pathways (Rimmer et al. 2018) or enough visible light for photosynthesis to occur (Mullan & Bais 2018; Lehmer et al. 2018; Covone et al. 2021). UCDs are also especially active objects (West et al. 2015; Williams et al. 2015; Gizis et al. 2017; Paudel et al. 2018; Günther et al. 2020), producing energetic stellar flares and large-scale photometric variability, which results in treacherously variable conditions for their planets. The dynamic relationship between a planet and its host is a major factor in evaluating how conducive a planet's environment is for life. For UCDs, however, this relationship is poorly understood.

Stellar flares may be a major determinant in whether a planet can initiate and sustain life. Flares are explosive events caused by magnetic recombination in the upper atmosphere of a star (Benz & Güdel 2010). This sudden eruption of magnetic energy usually occurs around active regions, such as stellar spots, and causes bursts of particles and electromagnetic radiation. The spectra of this radiation resembles a black body with an effective temperature of 9000 K (Shibayama et al. 2013). Often, though not always, powerful flares will come accompanied with a coronal mass ejection (CME) event, where clouds of charged particles are directionally ejected from the star. While flares can be destructive, through atmospheric erosion (Lammer et al. 2007), ozone depletion (Segura et al. 2010; Tilley et al. 2019), and even extinction events, they can also be an essential power source for life. It is possible that flares could provide the missing energy at the bluer end of the spectrum needed for cool, red dwarfs to initiate prebiotic chemistry (Buccino et al. 2007; Ranjan et al. 2017; Rimmer et al. 2018) and for photosynthesis (Mullan & Bais 2018; Lingam & Loeb 2019a). The additional UV energy may also affect the evolution of a planet's atmospheric chemistry (Segura et al. 2010; Vida et al. 2017). In summary, stellar flares have extreme and far-reaching consequences on the planets hosted by cool stars. Therefore, it is essential to study these stars' flaring activity, and assess how applicable our current understanding of stellar activity is to the ultra-cool regime.

Stellar activity and rotation are also closely connected. Magnetised stellar winds dictate the loss of angular momentum, which slows a star's rotation over time. These stellar winds are dependent on the structure and properties of the magnetic field. As the rotation slows, this decreases the magnetic activity in a process known as spin down (Skumanich 1972; Noyes et al. 1984). Due to this effect, rapidly rotating stars flare much more frequently than slow rotators (Skumanich 1986; Davenport et al. 2019; Mondrik et al. 2018; Medina et al. 2020). Spots and faculae on the surface of stars and brown dwarfs cause periodic photometric variations as they come in and out of view, on the same timescale as the object's rotation. This allows the rotation period to be deduced directly from the photometry. Therefore, the rotation of a star can provide valuable insights into the magnetic dynamo, the mechanism which generates a star's magnetic field.

This magnetic dynamo is poorly constrained for fully convective low-mass objects (with masses $\leq 0.35 M_\odot$, Chabrier & Baraffe (1997)), where it is believed to differ significantly from the solar model. Despite this predicted difference, recent work has shown that relationships between activity and rotation remain consistent from partially to fully convective stars (Wright & Drake 2016; Newton et al. 2017; Wright et al. 2018). Spin down is, however, believed to occur on slower timescales for fully convective stars, which accounts for the enhanced activity of mid-to-late M dwarfs (West et al. 2008; Newton et al. 2017; Jackman et al. 2021). Therefore, rotation,

activity, and the relationship between the two, are extremely useful probes of the underlying magnetic activity of ultra-cool dwarfs.

Due to the promising nature of M dwarfs as planetary hosts, there have been several detailed studies of their flaring activity within the past decade. Space telescopes, such as the Kepler/K2 missions (Borucki et al. 2010; Howell et al. 2014), allowed the first insights into the flares of bright M dwarfs (Davenport et al. 2014; Hawley et al. 2014; Lurie et al. 2015; Silverberg et al. 2016) and of (small numbers of) ultra-cool dwarfs (Paudel et al. 2018; Gizis et al. 2013, 2017). Additionally, the recently launched Transiting Exoplanet Survey Satellite (TESS, Ricker et al. (2015)) has facilitated studies of the flaring and rotating activity of cool stars (Günther et al. 2020; Medina et al. 2020; Seli et al. 2021). Several ground-based photometric surveys, such as MEarth (Nutzman & Charbonneau 2008), the Next Generation Transit Survey (NGTS, Wheatley et al. (2018)), the All-Sky Automated Survey for Supernovae (ASAS-SN, Shappee et al. (2014)) and Evryscope (Law et al. 2015) have also carried out detailed flare studies that include M dwarfs (West et al. 2015; Mondrik et al. 2018; Jackman et al. 2021; Schmidt et al. 2019; Martínez et al. 2019; Howard et al. 2019, 2020a,b).

The most frequent stellar flares are small, fast and difficult to detect above photometric scatter (Lacy et al. 1976). Due to the intrinsic faintness of UCDs in the visible, it is difficult to achieve the high photometric precisions necessary to constrain their flaring activity. However, it is possible to perform small, dedicated studies of the much rarer, high-energy flares on UCDs (Gizis et al. 2013; Paudel et al. 2018; Jackman et al. 2019). Therefore, to obtain a sufficient sample size, previous large flare studies have focused on hotter stars, up to mid-M dwarfs, where photometric precisions are much higher. This has resulted in limited flare statistics for the coolest stars, for which we would require a large, high-cadence photometric survey optimised for UCDs.

The SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars) project (Gillon 2018; Burdanov et al. 2018; Delrez et al. 2018; Jehin et al. 2018; Sebastian et al. 2021) aims to search for transiting planets around the nearest (within 40 pc) ultra-cool dwarf stars. SPECULOOS's motivation is to provide temperate, terrestrial planets for detailed atmospheric characterisation with the James Webb Space Telescope (JWST, Gardner et al. (2006)) and future extremely large telescopes. The SPECULOOS target catalogue of ultra-cool objects is defined in Sebastian et al. (2021). SPECULOOS comprises a network of 1-m class telescopes spread across the Northern and Southern Hemispheres. The largest facility, the SPECULOOS-Southern Observatory (SSO) in Cerro Paranal, Chile, consists of four identical, robotic telescopes. These telescopes are named Io, Europa, Callisto and Ganymede. Each telescope operates independently and in robotic mode following plans written by SPECULOOS's automatic scheduler, SPeculoos Observatory sChedule maKer (sᴘᴏᴄᴋ, Sebastian et al. 2021). Additionally, the SPECULOOS-Northern Observatory (Niraula et al. 2020) in Tenerife and the SAINT-EX (Search And characterIsatioN of Transiting EXoplanets) facility (Demory et al. 2020) in San Pedro Mártir each have one 1-m telescope. The SSO began official scientific operations in January 2019, followed by SAINT-EX in March 2019 and the SNO in June 2019.

In this paper we present a study of flares and rotation of SSO targets, observed over a span of almost two years. Section 2 defines the SSO data sample. Section 3 outlines how we generate and clean our global lightcurves and how we combine an automated flare detection algorithm with manual vetting to obtain our final flare sample. We describe modelling the flares, calculating the flare

energies, and estimating flare rates. This section also includes an assessment of the flare sample's completeness. In Section 4 we measure the rotation periods in our sample. The results of our flare and rotation analyses are presented in Section 5 and discussed in Section 6, which includes a contextualisation of the impact of flares on the potential for life on planets around ultra-cool dwarfs.

2 SPECULOOS-SOUTH DATA SAMPLE

For this study we defined a dataset spanning from 1st June 2018 to 23rd March 2020, using observations from all four telescopes in the SSO. While this start date is before official scientific operations began, it marks a point of stability in the commissioning phase. After this date we performed no major maintenance to the observatory; DONUTS (McCormac et al. 2013), our auto-guiding software for precise pointing, was in use; and our operating strategy had been finalised (see Sebastian et al. (2021) for details). The end of the data package was defined as the date that ESO Paranal Observatory shutdown due to the coronavirus pandemic.

In the time period between 1st June 2018 and 23rd March 2020, there are 661 potential nights of observation. Ganymede, the final telescope installed in the SSO, started commissioning on 30th September 2018, therefore, between the four telescopes there is a cumulative total of 2523 nights. However, primarily due to weather loss, the SSO has lost 23–24 per cent of observing time, resulting in 1931 nights of observation. Over the 1931 (combined) nights in this sample, we have observed 176 unique photometric targets for at least one night. These observations have typical exposure times of 20–60 seconds. The majority, though not all, of this data sample are in the SPECULOOS target list (86 per cent), as defined in Sebastian et al. (2021).

This target list is divided into three main scientific programs, which are described in more detail in Sebastian et al. (2021). The objects in the SSO sample which are not in the target list are exclusively objects with spectral types around or earlier than M6. It is likely these objects were part of the commissioning phase of the telescopes or were removed from the target list (due to reclassifying their spectral type) before official scientific operations began. For the objects in the target list, we extract radii, masses, effective temperatures, and spectral type classifications from Sebastian et al. (2021). The stellar parameters for non-target-list objects are calculated in the same way.

We remove any objects that are not part of SPECULOOS's usual "survey mode", such as targets observed for follow-up and monitoring of TRAPPIST-1's transits. We note that in SPECULOOS's "survey mode", there are no simultaneous observations with multiple telescopes during the course of a night. During this time period the operational temperature of the CCD was increased from −70 to −60°C in October 2018. The choice to raise the temperature of the CCD was due to the effect on the quantum efficiency of the detector, improving our sensitivity at the red limit, while the increase in dark current was found to be negligible. This temperature change introduces an offset in the differential flux between the nights before and after. Whilst this offset has no consequence on our flare program, it affects long-term photometric trends. Therefore, when recovering rotation periods (in Section 4), we split the lightcurves that straddle this temperature change and analyse the before and after sections independently. As this temperature change happens during the first few months of our dataset, the majority of targets were either observed entirely before or after it. Therefore, we only need to split a handful of our lightcurves. If a target has been ob-
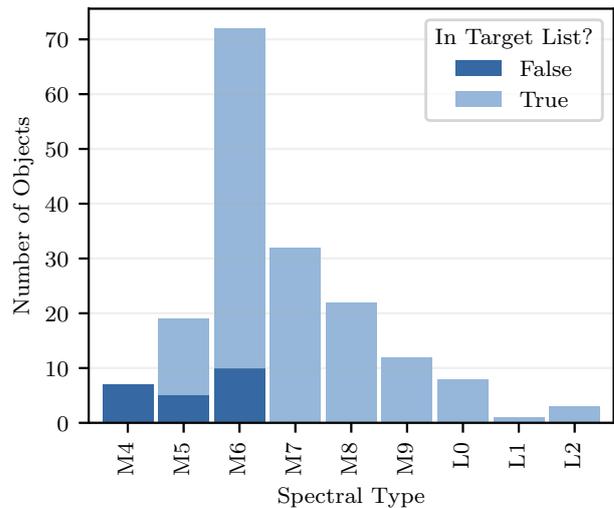
Figure 1. Spectral type distribution of the objects in the SPECULOOS-South data sample from 1st June 2018 to 23rd March 2020. We indicate which objects are in the SPECULOOS target list (Sebastian et al. 2021).

served by multiple telescopes we do not combine their lightcurves, as each will experience slightly different instrumental systematics. However, we do check that the rotation periods estimated in Section 4 are present in observations from all telescopes.

If we choose only the objects which have been observed more than 20 hours with one telescope, this provides us instead with 154 targets (of which 134, or ~87 per cent, are in the SPECULOOS target list). We define the observation time as the sum of the span of all observation nights (start of night to end of night), where we exclude any gap longer than 15 minutes. These 154 objects define the SPECULOOS-South data sample. This sample covers a range of M dwarfs in spectral type, extending from M4 into the early L dwarf regime (with masses of 0.07–$0.2\,M_\odot$), as shown in Figure 1. All objects in this sample will therefore be fully convective, as the convection limit occurs around $0.35\,M_\odot$ (Chabrier & Baraffe 1997). While M4 and M5 objects are not considered ultra-cool dwarfs, we include them in this sample to explore any differences between mid-M, late-M, and L dwarfs.

In this sample, 80 per cent of the objects have been observed for less than 77 hours (Figure 2), and 50 per cent have been observed for less than 51 hours. The short observation times in this sample will inevitably bias us towards detecting flares on targets with high flare rates in Section 3.3 and detecting the clearest rotation periods for fast rotators in Section 4.

From Sebastian et al. (2021) we confirm that none of the target list objects in the SSO data sample are known binaries. We could not find any record in the literature of binarity for any of the 20 non-target list objects, except from the M4 star, Gaia DR2 4450376396936878336 (Lépine & Bongiorno 2007). We flag this object, however, we find no evidence that this object behaves anomalously. This star is flaring, however, it only flares twice, therefore, it is not included in our analysis in Section 5.3. We also do not find a rotation period for this object, therefore it does not impact our results in Section 4.
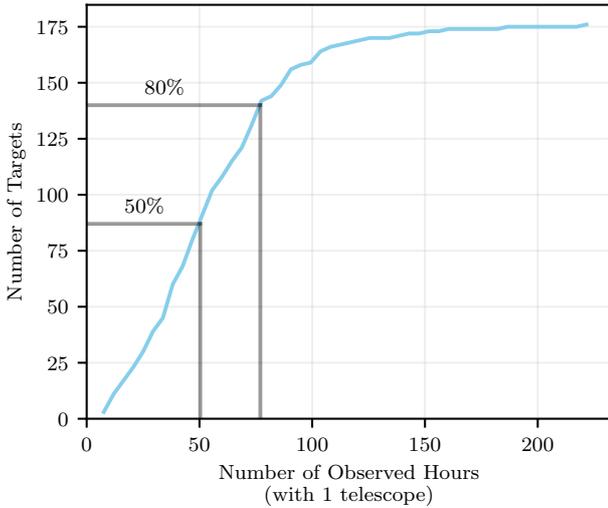
Figure 2. A cumulative histogram of the number of days observed for the objects in the SSO data sample. Of the 176 objects, 14 per cent (24) have been observed for less than 20 hours, 50 per cent have been observed for less than 51 hours, and 80 per cent have been observed for less than 77 hours.

3 GENERATING THE SPECULOOS-SOUTH FLARE SAMPLE

3.1 Global lightcurves

To search for flares, we extract global lightcurves for all targets in the data sample, using the SSO Pipeline described in Murray et al. (2020). The SSO Pipeline performs automated differential photometry by generating an 'artificial lightcurve' from a normalised, weighted average of reference stars in the field. Only objects above a brightness threshold (optimised for each field) are chosen as reference stars. Reference stars whose differential lightcurves exhibit photometric variability, or that are far in spatial distance on-sky from the target, are then weighted down in an iterative algorithm, based on Broeg et al. (2005). It is worth noting that the target star is removed from this process. This means that, except for the distance weighting, this 'artificial lightcurve' is entirely independent of the target star's lightcurve. This pipeline then corrects for the second-order effects of atmospheric water absorption on our differential photometry, using measurements of precipitable water vapour from the ground-based LHATPRO (Low Humidity and Temperature PROfiling radiometer) instrument (Kerber et al. 2012).

To generate global lightcurves, we perform our differential photometry process on the entire time series of a target at once. We obtain differential lightcurves for 13 different aperture sizes. For each aperture, the aperture size, comparison stars, and weightings do not change over the time span of observations. Using global lightcurves allows us to study their long-term photometric variability, but restricts our ability to optimise the photometry night-by-night.

To mitigate any systematics caused by changing atmospheric conditions, we carefully select the 'best' aperture for each target. As we only use one aperture size across all nights, it is important to take time at this step. Choosing the best aperture for a target involves a fine balance between an aperture that's large enough to avoid losing flux for nights with sub-optimal observing conditions (e.g., a larger seeing), and an aperture that's small enough to prevent blending from nearby stars and extra white-noise contamination from the

background. We select the aperture by eye which minimises the correlation between the lightcurve flux and the seeing for the whole lightcurve, and has no clear contamination from neighbouring stars. In addition, we implement a bad weather flag and thresholds for the background sky level, airmass, and seeing (above which the data are severely impacted), described in more detail below.

3.2 Cleaning the lightcurves

3.2.1 Removing 'bad' observations

In the context of this paper, 'bad' observations are defined as those which are significantly affected by the observing conditions of the night, where distinguishing real stellar variability from ground-based systematics would be extremely difficult.

To reduce the impact of these observations on our flare study, we implement a bad weather flag, as defined in Murray et al. (2020). As previously mentioned, the artificial lightcurve has no relation to the target's lightcurve, therefore it is not impacted by flares on the target star. This distinction means that we can use the artificial lightcurve as an independent reflection of the photometric conditions at that time to identify where the quality of our lightcurves deteriorates. Data points in our lightcurves are flagged based on the RMS of the section of the artificial lightcurve surrounding that data point (local RMS). We define the length of the local section to be ±0.01d. We make the assumption here that there is a threshold for photometric scatter in the artificial lightcurve (which therefore must be present in the highly weighted reference stars' lightcurves) above which we cannot obtain good photometric precision. We define this threshold as a local RMS of 8 per cent, as defined in Murray et al. (2020).

We also include strict cuts on specific observation parameters. We exclude data points in our lightcurves where the sky background level is greater than 4000 counts per pixel, the airmass is greater than 2.5, or the full width at half maximum (FWHM) of the point spread function (PSF) is greater than 2.6 arcsec (the SSO cameras have a pixel scale of 0.35 arcsec pixel^{-1}).

From 1st May to 19th June 2019 we experienced an issue with moving dust on the CCD window of Ganymede. While stationary dust can be easily corrected by flat images taken at the start of each night, if dust moves across the CCD window during the night, this leads to residuals in the flat correction and structures in the final differential lightcurves. These structures can mimic planetary transits, though the frames affected by dust are easy to identify from the raw images. However, as we do not currently have a robust correction technique for moving dust, we removed all observations taken by Ganymede during this period. This amounts to less than 2 per cent of total observations, therefore the impact is marginal.

3.3 Flare detection

As the SSO data sample is constrained to only 154 targets, it is possible to identify flares manually. Therefore we did not find it necessary to implement a fully automated, complex flare-detection algorithm. Instead we decided to perform flare detection in two parts: a simple, automated flare-detection algorithm to extract all the flare candidates, followed by a manual vetting process to confirm them. Both stages of this flare detection process are demonstrated in Figure 3 for target Gaia DR2 3200303384927512960. We note that this target is one of the most frequently flaring objects studied by Seli et al. (2021). This object is a particularly challenging case due

to its rapid variability, where flare decays are difficult to separate from photometric modulations.

3.3.1 Automatically searching for flares using lightcurve gradients

For the first part of our flare detection process, we chose to implement the automatic flare detection method set out in Lienhard et al. (2020), due to its simplicity, speed, and robustness. This method was developed on lightcurves from the TRAPPIST (TRAnsiting Planets and PlanetesImals Small Telescope) telescopes (Gillon et al. 2011; Jehin et al. 2011). Lienhard et al. (2020) capitalise on the changing gradients in the asymmetric structure of a flare, which involves a sharp increase to a peak followed by a slower exponential decay (Moffett 1974). They evaluate the following two criteria:

$$\frac{2f_j - f_{j-2} - f_{j+3}}{\sigma_j} \times \frac{|f_j - f_{j-2}| - |f_{j+1} - f_j|}{\sigma_j} > A_{\text{thresh}} \quad (1)$$

and

$$2f_j - f_{j-2} - f_{j+3} > 0, \quad (2)$$

where f_j is the flux of the j^{th} data point in the target's lightcurve, and σ_j is the RMS of the section of the lightcurve comprised of the nearest 60 data points. By only considering small neighbouring sections of the lightcurve in our RMS, and the flux for a few points either side of the flare in our criteria, we are able to detect small flares on rapidly rotating or frequently flaring objects. Equation 1 assesses the quality of the flare's shape. If we assume that the peak of the flare is at j then the first half of Equation 1 confirms that we are at a peak by ensuring the flux at j is greater than the flux a few exposures before and after. The steeper the peak, or larger the flux difference between the peak and surrounding points, the higher its value. The second half of Equation 1 hinges on the asymmetry of the flare by requiring that the gradient before the peak, the fast rise, is larger than the gradient after the peak, the slow decay. For this second half of the equation, the greater the asymmetry, the higher its value. By dividing by the local RMS, the aim is to remove flare structures caused by photometric scatter, which is especially problematic for ground-based observing with rapidly changing weather conditions. Equation 2 simply prevents the case where both halves of Equation 1 are negative.

It is possible to have significant time gaps within the lightcurve sections of 60 points due to bad weather, however, this scenario is rare. Including data points that are far apart in time would lower the reliability of the RMS in this section of the lightcurve, however, when weather conditions are adverse enough to close the telescope, or seriously affect photometric precision, the RMS directly before or after is likely to be enhanced (due to poor observing conditions that do not quite meet the threshold to be removed). We note that the "nearest 60 points" are only used for this flare detection method, following Lienhard et al. (2020), whereas time boxes are used for calculating the running median and RMS in Sections 3.3.2, 3.3.3 and the rest of this work.

In their paper, Lienhard et al. (2020) determine $A_{\text{thresh}} = 12$ by inspection of the smallest flares that they intended to remove; however, this value is specific for the photometric precision and typical exposure times of TRAPPIST lightcurves. We would expect the precisions obtained by TRAPPIST's 60-cm telescopes to differ significantly from those achieved by SPECULOOS's 1-m telescopes. Instead, for the SSO lightcurves, we derived a value of $A_{\text{thresh}} = 5$. We find a lower value for A_{thresh} partly due to the increased photometric precision of the SSO and partly because we are choosing to

manually verify the flare candidates that this algorithm detects, so we can afford to over-detect at this stage.

By choosing a flare detection method based on the shape of a flare, and not outlier detection, this may limit the diversity of flare morphologies that we will detect. Since there have been few large-scale studies done to date on ultra-cool dwarf flares, such as those in our SSO sample, any difference in flare structure between earlier- and later-spectral-type objects is still unknown. Therefore, we decided to focus on the flares most resembling the standard flare shape (Moffett 1974), while flagging more unusual lightcurve behaviour in the manual vetting stage (see Section 3.3.4).

3.3.2 Small vs. large flares

To both maximise the detection of small flares, and optimise the modelling of high-amplitude flares with slow recovery times, we split the flare candidate sample into two categories: *small* and *large* flares. We classify *large* flare candidates as those with at least 2 data points in the flare region more than 7 times the running local standard deviation (standard deviation of the surrounding 80 data points) above the running local median (defined similarly as the median of the surrounding 80 data points); otherwise, we classified it a *small* flare candidate. The flare region starts at least two points before the peak. We approximate the end of the flare region by using least-squares to fit the flare decline with a sum of a fast and a slow-decaying exponential, as in Davenport et al. (2014). The decay time of the slower exponential decay is used to estimate the time of the end of the flare region.

3.3.3 Validating and vetting flares

Once a flare candidate is classified as *large*, then we use the surrounding 320 points (defined as global), instead of 80 points (local), to calculate the running median and standard deviation for the next step of validation.

If the criteria in Section 3.3.2 is met, we then run several additional automatic quality checks to separate flare candidates from flare-like signals replicated by noise. By ensuring that there is more than one data point in the flare region after the peak, we remove cosmic rays from the flare candidates. Additionally at least 2 points in the flare region must be more than 2 times the running local standard deviation above the running local median of the lightcurve. The local standard deviation and median are as defined in Section 3.3.2. When calculating the running local median and standard deviation, we mask all potential flare candidates in the lightcurve, and then interpolate over the flare region.

This automated algorithm provides us with a collection of flare candidates to follow up with manual vetting to then obtain our final flare sample.

3.3.4 Manual vetting

Once we have a collection of flare candidates from the method detailed in Sections 3.3.1 and 3.3.3, then the second part of our flare detection process is to manually inspect these candidates and obtain our final flare sample. We consider each flare candidate and only confirm those which can be clearly identified as matching the standard flare shape of a sharp flux increase followed by a slower, exponential flux decay (Moffett 1974; Davenport et al. 2014). We also ensure that the flares cannot be attributed to rapid changes in atmospheric conditions.

As a final validation step, we check through the global lightcurves in the SSO sample. The only flare-like structures remaining are due to the following events: a flare occurs before the start of the night (where we catch the tail end of the flare) or at the very end of a night, the structures are low-amplitude and so cannot be distinguished from noise, the structure diverges significantly from the standard flare model, or they are correlated with systematics. Examples of structures identified in the SSO data set that deviate from the standard flare model, and are not clearly overlapping flares, are symmetric flares (with similar rise and decay times) and flares with a sudden increase and linear decline. During this manual vetting stage we do not add any of these potential missed flares into the flare sample, as they do not meet the criteria to be clearly identified as a flare with the standard flare shape.

3.4   Modelling flares

In order to model each flare, we coarsely flatten the lightcurve surrounding that flare by dividing the lightcurve by a median filter, using the local or global criteria as in Section 3.3.2.

We chose to model our flares by fitting the empirical flare template described in Davenport et al. (2014). In this paper the authors generate a median flare template from Kepler observations of 885 'classical' flares on the M4 star GJ 1243. They found a sharp flux increase, modelled with a fourth-order polynomial, followed by an initial fast exponential decay and subsequent slower exponential decay. This model requires a time for the peak of the flare, a relative amplitude on the normalised lightcurve, and a full width half maximum ($FWHM_{\mathrm{flare}}$), which corresponds to the flare's decay time. We demonstrate fitting this flare model to flares of differing amplitudes and decay times in Figure 3.

While this model works well for the more classically shaped flares, it struggles to represent complex flares that do not fit the standard flare morphology (Davenport et al. 2014). This template has yet to be tested on a statistically large number of low-mass objects, such as UCDs, whose flare profiles may differ significantly. Complex flares include those with multiple peaks (Davenport et al. 2014), oscillations (Anfinogentov et al. 2013), and those which are closely entangled with variability. The overlapping of multiple flares is likely the culprit of more unusual lightcurve structure. During this manual vetting, we divide flare regions containing clear multiple peaks into separate flares and each of the new separated flares are added to the flare list. We also flag (but do not remove) 27 flares that either do not fit the standard flare model, or are not easily separable into multiple flares. This corresponds to 11 per cent of the total flare sample. An example of a possible flare overlap is shown in the bottom left plot in Figure 3.

3.5   Extracting flare energies

To measure the flare energies, we follow the technique described in Shibayama et al. (2013). We model the flare as blackbody with an effective temperature, T_{flare}, of $9000 \pm 500\,\mathrm{K}$ (Kowalski et al. 2013), and we assume this temperature remains constant. We calculate the luminosity of each star and flare, as seen through the SSO's I+z' filter, accounting for the overall system efficiency, as follows:

$$L_*' = \pi R_*^2 \int_{I+z'} R_\lambda B_\lambda(T_{\mathrm{eff}})\,\mathrm{d}\lambda \tag{3}$$

$$L_{\mathrm{flare}}' = A_{\mathrm{flare}}(t) \int_{I+z'} R_\lambda B_\lambda(T_{\mathrm{flare}})\,\mathrm{d}\lambda. \tag{4}$$

Here, R_λ is the total SSO response function, which is the product of the transmission in the $I + z'$ filter, quantum efficiency of the CCD, CCD window, and reflectivity of the mirror coatings. $B_\lambda(T_{\mathrm{eff}})$ and $B_\lambda(T_{\mathrm{flare}})$ are the Planck functions evaluated for the star's effective temperature, T_{eff}, and the flare temperature, respectively. Finally, $A_{\mathrm{flare}}(t)$ is the area of the flare. We do not include atmospheric transmission in our response function due to its rapid variability. Since the relative flare amplitude can be found directly from the normalised lightcurve, $F(t) = (\Delta F/F_{\mathrm{mean}})(t) = L_{\mathrm{flare}}'(t)/L_*'$, we can solve for A_{flare} as follows:

$$A_{\mathrm{flare}}(t) = F(t)\,\pi R_*^2 \, \frac{\int_{I+z'} R_\lambda B_\lambda(T_{\mathrm{eff}})\,\mathrm{d}\lambda}{\int_{I+z'} R_\lambda B_\lambda(T_{\mathrm{flare}})\,\mathrm{d}\lambda}. \tag{5}$$

To best estimate the flare amplitude, we use the local or global running median for F_{mean}, as defined at the end of Section 3.3.2. The bolometric flare luminosity, L_{flare}, can then be calculated with the assumption that the star radiates as a blackbody:

$$L_{\mathrm{flare}} = \sigma_{\mathrm{SB}} T_{\mathrm{flare}}^4 A_{\mathrm{flare}}, \tag{6}$$

where σ_{SB} is the Stefan-Boltzmann constant. The total bolometric energy of a flare (E) is then the integral of L_{flare} over the flare duration:

$$E = \int_{\mathrm{flare}} L_{\mathrm{flare}}\,\mathrm{d}t. \tag{7}$$

Using Equation 6, and assuming the flare temperature is constant, we can separate the time-dependent and independent components:

$$E = C \int_{\mathrm{flare}} A_{\mathrm{flare}}\,\mathrm{d}t, \tag{8}$$

where $C = \sigma_{\mathrm{SB}} T_{\mathrm{flare}}^4$. Using Equation 5 we can further separate out the time-independent components that depend on the properties of the star and the flare, and the time-dependent integral:

$$E = C_{\mathrm{star}}(T_{\mathrm{eff}}, R_*) C_{\mathrm{flare}}(T_{\mathrm{flare}}) \int_{\mathrm{flare}} F(t)\,\mathrm{d}t. \tag{9}$$

The motivation for calculating the energy in this way is that it allows us to calculate these three components separately. When generating a large pool of synthetic flares for injection-recovery tests, as in the following section, this drastically reduces the processing time. Every object in our SSO data sample will have a constant value for C_{star}, which only needs to be calculated once per star. The value for C_{flare} also becomes constant with our assumptions. It is then only the integral of the flare in the normalised lightcurve that remains to be calculated.

For this integral we decided to use the normalised lightcurve (corrected with the same local or global median filter described in Section 3.4) directly for the energy calculation, rather than using the flare model. In doing so we note that the energies we calculate provide a lower bound on the flare energy, as we may miss the flare's peak. It is also possible that if there was a large flare on a rapidly-rotating star, the underlying F_{mean} may 'smooth out' photometric structure, leading to the flare energy being over-estimated. However, to be classed a 'large' flare, a flare needs to be more than 7 times the running global standard deviation (as defined in Section 3.3.2). This standard deviation would be highly inflated from the periodic photometric modulations, and therefore this flare would have to be very high energy for this to become an issue. As high energy flares are rare (Gershberg 1972; Lacy et al. 1976), this scenario is not a major concern. We chose not to use the flare model because the flares in our lightcurves are well-sampled, and when there is increased photometric noise in the lightcurve, or more complex
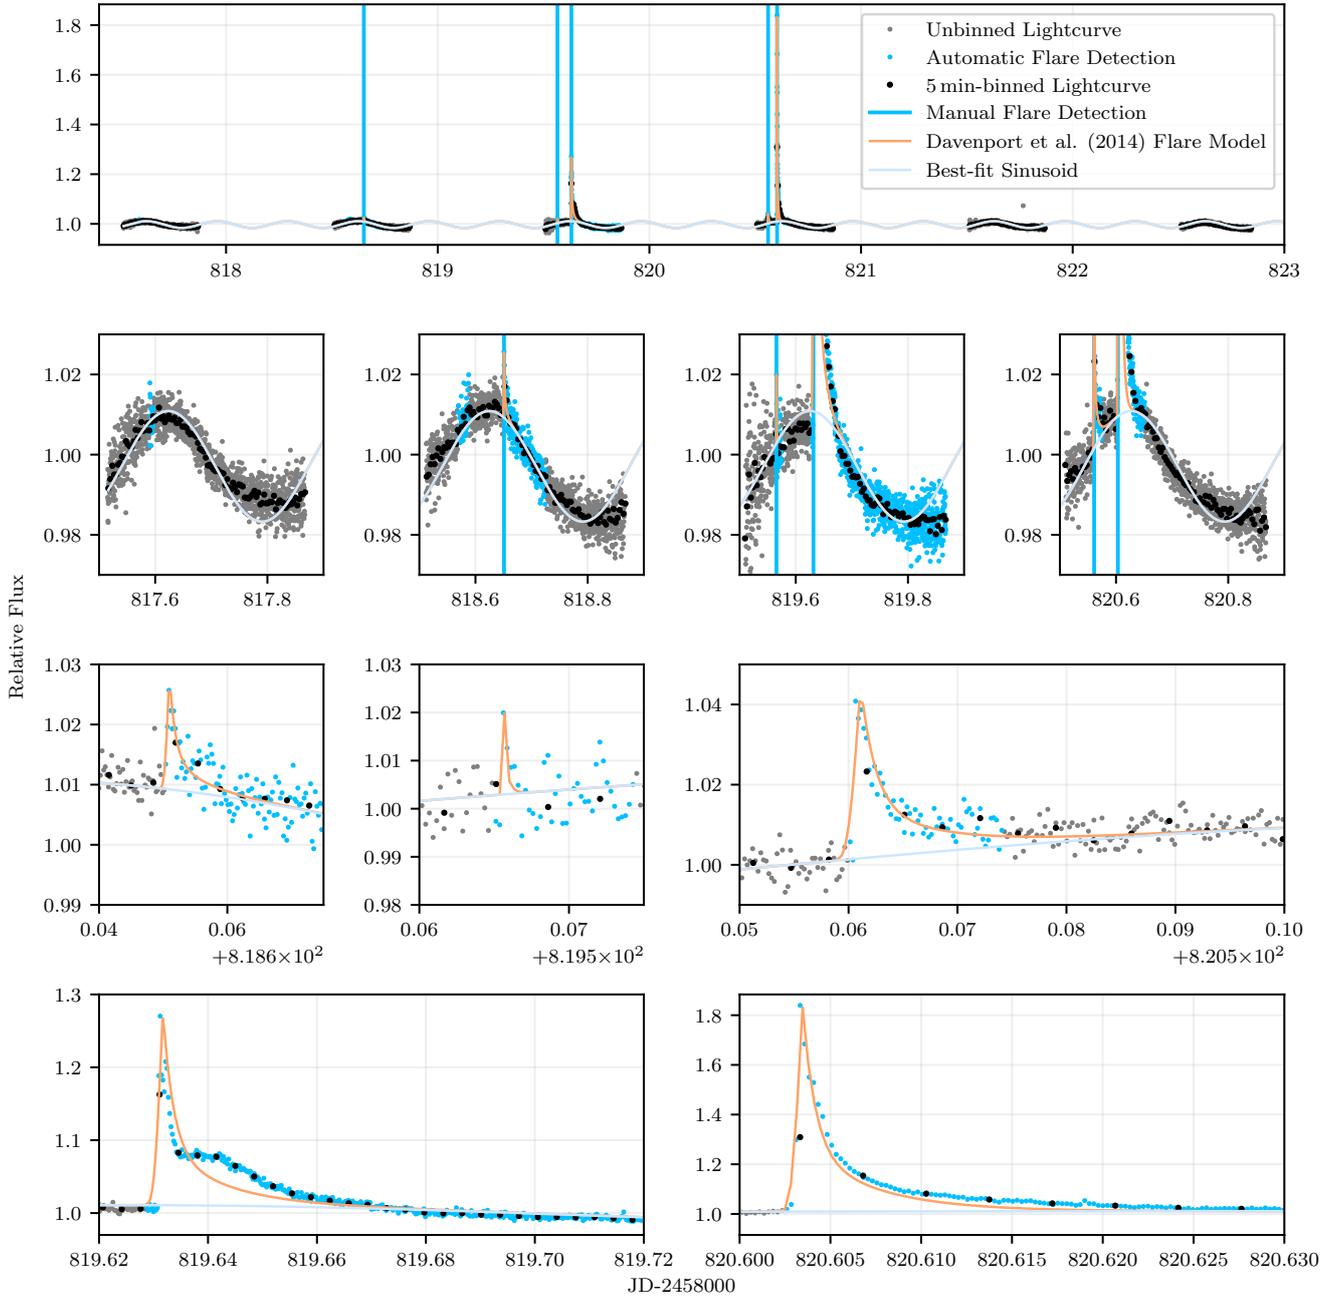
Figure 3. Demonstration of the flare detection process on six nights of observations of an M7 object (J = 10.7 mag), which exhibits both frequent flaring and short-period rotation. The four rows show different views of the same global lightcurve, zoomed in to individual nights on the second row, and only the five manually-verified flares on the bottom two rows. The grey points are the unbinned data, while the black points are binned every 5 minutes. We do a simple least-squares fit of a sine wave with a period of 0.33d, shown in light blue. It is clear from this fit that the rotation pattern is not perfectly sinusoidal; however, it provides a way to visualise the periodicity. The medium-blue data points are those initial flare regions flagged by the automatic flare detection algorithm described in Section 3.3.1. The vertical, medium-blue lines are the flare candidates that have also been confirmed manually. During visual inspection of this lightcurve, we remove two flares (one in the first plot on the second row and the first flare in the second plot on the second row), due to their small amplitude, which were found too difficult to detach from the photometric scatter. The best-fit Davenport et al. (2014) models for each flare are shown in orange. The case where we have a poor fit to the template is shown in the bottom left plot.

flare shapes from overlapping flares (see Section 3.4), the fit to the flare model can be unreliable. Integrating the lightcurve directly means that bursts of flares occurring in short succession are counted as one flare of increased energy. While this simplification may affect the calculation of our flare rates and energies, it should not

significantly affect our results as we only detect 27 unusually-shaped flares that may be a blend of multiple flares. We note that while flares are typically approximated as blackbodies with temperatures from 9000–10000 K, the flare temperature has been observed to vary

outside this range, both between and within flare events (Howard et al. 2020b).

3.6 Calculating Flare Rates

As flares are stochastic events, the gaps in observations due to the day-night cycle or bad weather loss do not hinder our ability to calculate flaring rates, which are only dependent on the total time on-sky. It is likely that we underestimate flaring rates for the several reasons flares are not found by our automatic detection algorithm, outlined in Section 3.3.4.

We calculate our flare rates as the number of flares detected per target divided by the observation time (summed across all telescopes). Almost all of our targets have been observed for less than 200 hours, therefore we have a lower limit on individual stars' flaring rates of ~0.12 d^{-1}.

3.7 Completeness of the flare sample

To evaluate the completeness of our flare sample (the minimum energy flares recoverable by our flare detection process for different spectral types), we perform artificial flare injection-recovery tests. Using the global lightcurves from the 154 different targets with at least 20 hours of observation, we mask all of the flares detected in Section 3.3. We use the Davenport et al. (2014) flare model to generate 100,000 artificial flares with amplitudes drawn from a log-normal distribution between relative fluxes of 0.001 and 5, and FWHM$_{\mathrm{flare}}$ drawn from a uniform distribution between 30 seconds (typical exposure time for SPECULOOS) and 1 hour. We calculate the energy of the resulting flares for each star, and divide into 6 energy bins in $\log_{10}$ space from 10^{28} to 10^{34} erg.

We note that here we have followed the work of Davenport et al. (2014); Davenport (2016); Günther et al. (2020), among others, and decided not to consider any relationship between amplitude and duration. While we may be injecting "unphysical" flares, this approach allows us to explore a larger parameter space of different flare morphologies. The amplitude-duration relationship for flare has also not been well studied for the case of UCDs, where it may differ from that seen in earlier M dwarfs. We do not see a clear relationship between fitted amplitude and duration in our sample (and so cannot confirm any other observed amplitude-duration relationships, e.g. Hawley et al. 2014), though our estimates of amplitude and duration rely on a good fit to the Davenport template, which is significantly more challenging with messy ground-based data. Our purpose with these tests is to examine the decay time, amplitude, energy and spectral type limits of our flare detection method, not to reproduce a realistic flare frequency distribution. To model a realistic flare distribution we would need to consider that low energy flares are more common than high energy flares, as discussed in Section 5.3. However, then we would have significantly more low energy flares injected and recovered, making our sensitivity results more robust for small amplitudes/decay times, and less robust for high energy flares.

For each energy bin, treated independently, we then randomly select 5 flares and inject them into our differential lightcurves at a random observed time, taking care to ensure the ±5 data points surrounding the time of the flare's peak are within 0.01 d (14.4 mins). This prevents those 5 flares from occurring too close to the start or end of night or during gaps in our observations, which we would always miss. We allow flares to overlap to reflect the scenario we see often in real lightcurves, however as there are only ever a maximum of 5 flares in a lightcurve at once, the impact on our recovery results is limited.

We then run the automatic part of our flare detection method, described in Sections 3.3.1 and 3.3.3, on all injected lightcurves and record the recovered flares. We consider a flare to be recovered if the recovered time for its peak flux is within 1 FWHM$_{\mathrm{flare}}$ (of that flare) of the injection time. We allow for a more flexible recovery time because of the decision to allow flares to overlap, which can lead to more complex structures with multiple peaks. If two or more flares overlap within 1 FWHM$_{\mathrm{flare}}$, they may be detected as only one flare. We assume if they occur further apart in time then they should be easily separable. This process is then repeated for each energy bin 5 times, making sure to use different flares from our artificial sample in each iteration.

The manual vetting step of our flare detection technique was not feasible for checking the approximately 23,000 flares we injected. This may introduce a bias into our results. It is more likely during manual vetting that we would conflate low-amplitude flares or flares on high photometric scatter targets with noise, therefore by removing this stage we have potentially inflated our recovery rates for the lowest energy flares, or for the faintest target stars.

Figure 4 shows the results of the injection-recovery tests, carried out for approximately 23,000 artificial flares. We are able to detect a significant proportion of flares with amplitudes above 1 per cent. Due to the short exposure times of SPECULOOS, we have the advantage of being very sensitive to flares with a short duration (with FWHM$_{\mathrm{flare}}$ < 5 minutes). However, we are more limited in detecting longer duration flares (with FWHM$_{\mathrm{flare}} \geq 60$ minutes) due to the day-night cycle, with typical uninterrupted observation windows of 4–8 hours. There will also be some stars which have nights of observation less than 4 hours, due to weather, or limited visibility, and these will be the most difficult targets on which to detect slowly decaying flares. However, we note that during the manual vetting (Section 3.3.4) we inspect the lightcurves with flares removed and we do not find any undetected high energy flares remaining in our sample. We also assess our limitations with the photometric RMS of our lightcurves. When the RMS of a global lightcurve exceeds ~ 0.7 per cent (for 5-minute binning) we begin to see a drop in detection efficiency, resulting in an increase in the minimum detectable amplitude. Only three targets in our dataset exceed this RMS.

For a given flare amplitude and FWHM$_{\mathrm{flare}}$, the flare energy will vary depending on the effective temperature and radius of the star (Equation 9). The recovery fraction for different flare energies across all stars is shown in Figure 5. However, the minimum detectable energy of recovered flares varies with spectral type, as shown in Figure 6, which facilitates the detection of lower energy flares on our coolest dwarfs. For low-energy flares, we have two competing biases. The lightcurves for the later, and fainter, M and L dwarfs have higher photometric scatter, which makes it difficult to detect small-amplitude (and therefore lower energy) flares. However, as flares have a strong white light component (Namekata et al. 2017), there is an increased contrast between flares and the stellar spectra of red dwarfs, which should make lower energy flares easier to identify in our coolest stars (Allred et al. 2015; Schmidt et al. 2019). SPECULOOS is also optimised to observe ultra-cool objects, and due to this specificity, we only see a very minimal increase in RMS with later spectral types.

The results from our artificial flare injection-recovery tests allow us to calculate the recovery fraction of our flare detection method, $R(E)$, that reflects the low recovery rates of low energy flares. The average $R(E)$ for all stars is shown in Figure 5, as well as
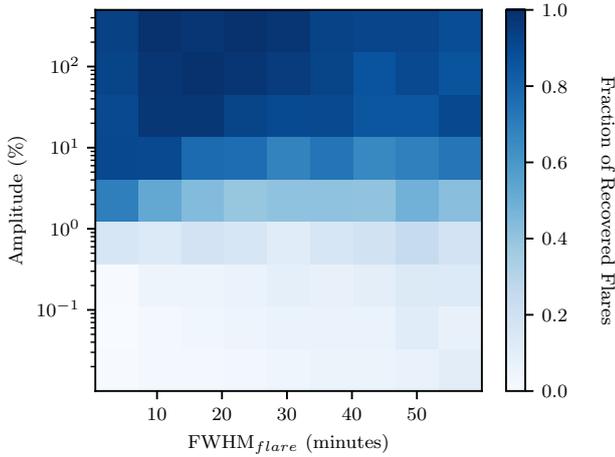
Figure 4. Fraction of injected flares recovered as a function of flare amplitude and FWHM$_{flare}$ using the automatic flare detection algorithm described in Section 3.3.

Figure 6. Fraction of injected flares recovered as a function of flare energy and the object's spectral type. We find that as we move to cooler stars we are able to detect lower energy flares. The upper and lower flare energy limits, determined by the amplitude and FWHM$_{flare}$ limits of the artificial flare sample and the radius and T$_{eff}$ of each star, are shown in orange. We are only impacted by the sample limits for spectral types earlier than M5 and later than L0.

Figure 5. Fraction of injected flares recovered as a function of flare energy.

its spectral dependence in Figure 6. As our sample is already small we are unable to reduce detection-sensitivity effects by simply limiting our sample to flares with energies above a minimum recoverable energy (Paudel et al. 2018), or energies where a sufficient fraction of flares are recovered (Davenport 2016). Instead, we decided to calculate flare occurrence rates in Section 5.4 using the method in Jackman et al. (2021), which models the decline in sensitivity for low energy flares.

4 IDENTIFYING ROTATION PERIODS

With the previously identified flares and bad observations masked, we search for rotation periods in the 154 low-mass objects with more than 20 hours of observation. We apply the Lomb-Scargle periodogram analysis (Lomb 1976; Scargle 1982) to our global lightcurves, binned every 20 minutes. The SSO lightcurves are not uniformly sampled, as we have gaps in our data, not only from the day/night cycle, but also from bad weather, masked flares, and changes to our observation strategy. Therefore, careful treatment of the resulting periodograms is essential to remove aliases.

We searched for periods between the Nyquist limit (twice the bin size) and the entire observation window. While we cannot be

fully confident in periods greater than half their observation window, by expanding the period range we search, we can estimate a lower bound on long-period rotation. We then remove all peaks in the periodogram with a false alarm probability below 3σ (0.0027). Any target which has at least one significant Lomb-Scargle peak undergoes visual inspection of their periodogram and the corresponding phase-folded lightcurves for all possible rotation periods (all significant peaks in the periodogram). In addition, by comparing with periodograms of the time stamps, airmass, and FWHM, we can eliminate signals which arise from the non-uniform sampling or ground-based systematics.

We decided to apply a similar classification system as Newton et al. (2016) to our rotating objects. Classifying our rotators helps account for the difficulties in observing from the ground, both through atmospheric systematics, and through irregular, non-uniform sampling. When examining each lightcurve, we ask ourselves several questions:

(i) Is the period clearly visible in the phase-folded lightcurve?

(ii) Was the object observed for long enough to span multiple periods?

(iii) Is the frequency an alias of the "day signal", seen as integer multiples in frequency space (periods of 1, 0.5, 0.33 d etc.)?

(iv) Is there a correlation with systematics (as seen in the airmass or FWHM periodograms)?

(v) Can the period be seen by eye in the un-phased lightcurve?

(vi) If we observed this object with more than one telescope, does this period fit them all?

(vii) Can we easily disentangle the 'real' period from its one day aliases?

(viii) Is the amplitude of the periodic signal above the level of noise in the lightcurve?

If a rotator passes all the above criteria, then we class it as 'A'. If it fails any of the above, but the rotation still seems likely then we class it as a 'B' grade rotator. Most commonly, the 'B' rotators are convincing, but we do not observe multiple cycles, or

we cannot easily choose between the period and its one-day aliases. Any lightcurves for which we see some periodic structure, but we cannot easily determine a period, we class as 'U'. This can result from broad peaks in the Lomb-Scargle periodogram, or multiple possible periods due to lack of observations, or a very low amplitude periodic signal that is comparable to the noise level. Because of these ambiguities in the period measurements, we do not attempt to place errors on our period estimates. If we do not detect any periodic signal or we cannot remove correlations with systematics (such as for very crowded fields), then we class the object as 'N'. Finally, in addition to the classes defined in Newton et al. (2016), we add an extra 'L' grade. This is for the lightcurves where we see clear long-period rotation, but the period is similar or longer than the time window observed. For these objects the best we can do is to estimate the lowest possible period we measure, and acknowledge that there are large uncertainties on these period values.

5 RESULTS

The SSO data sample, along with the results of the flare and rotation analyses, are presented as supplementary material.

5.1 SPECULOOS-South Flare Sample

From the SSO dataset described in Section 2, we identify 234 flares. We find that of our 154 unique targets, 78 are flaring (50 per cent). These flaring stars span the spectral type range from M4 (T_{eff} = 3160K) to L0 (T_{eff} = 2313K). Figure 7 shows the spectral type distribution and proportions of the flaring objects identified in the SSO sample. From this figure we see the proportion of flaring stars stays consistently above 60 per cent for objects of spectral type M5–M7. We see that the rate of flaring stars begins to decline around M8 (~ 30 per cent), with no detected flares for any object beyond L0. The coolest flaring star we detect is a 2313 K, M9.6V object (which is rounded to L0 in Figure 7). However, for the L dwarfs we are limited by the small sample size, and therefore cannot make any conclusions about whether the fraction of flaring objects continues to decrease beyond late M dwarfs. Likewise, we do not have any objects earlier than M4; however, these earlier M dwarfs have been well studied by TESS, Kepler, and MEarth. Despite our small sample of M4 dwarfs, we see an apparent reduction in activity for stars earlier than M5 that agrees with the previous results (~30 per cent for M4 with TESS, Günther et al. (2020), 25% for M5 with ASAS-SN, Martínez et al. (2019)). Several authors find a steep rise in flaring fractions around spectral type M4 (West et al. 2004; Yang et al. 2017; Günther et al. 2020; Martínez et al. 2019). Specifically, West et al. (2004) analysed the Hα emission of cool stars (which roughly correlates with flaring activity, Yang et al. 2017; Martínez et al. 2019) and determined that the fraction of active stars rose monotonically from spectral type M0 to M8, peaked at M8 (~80 per cent active) and then declined to L4. Our results show an increase from M4 to a broad peak from M5–M7, followed by a decrease from M8 to L2.

We probe the parameter space of high-flare-rate stars with low-to-mid energy flares. Due to the relatively short baselines of SPECULOOS observations, if a target does not flare frequently, then it is unlikely we would identify it as flaring. This also means we are less likely to detect the rarer, high-energy flares (Gershberg 1972; Lacy et al. 1976). In conjunction with the difficulties arising from the day-night cycle (typical night observations are 4–8 hours), which complicates our detection of very slowly decaying flares (see Figure 4), it is unlikely that superflares with energies
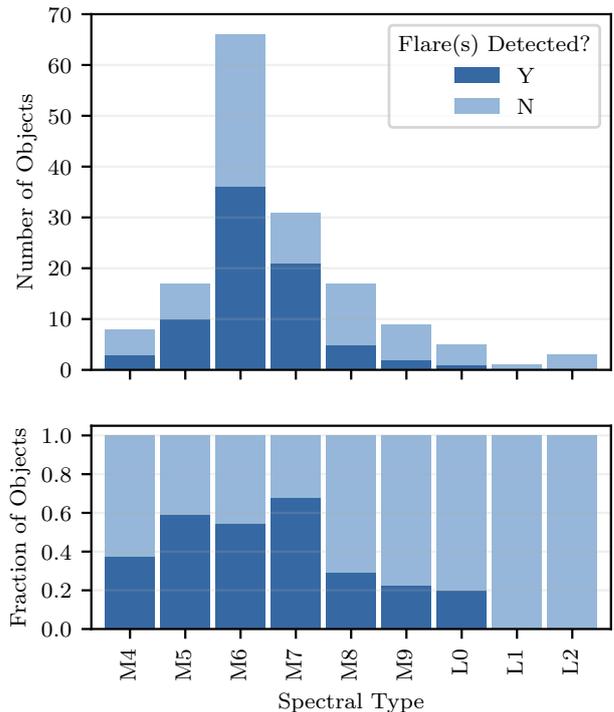
Figure 7. Top: Histogram of objects in the SSO data sample as a function of spectral type. The flaring objects are marked in light blue. **Bottom:** Fraction of flaring objects as a function of spectral type.

> 10^{33} erg (Shibayama et al. 2013) will be detected by our survey. From studying the completeness of our flare sample, we also conclude it is unlikely that we will detect any flares below $10^{29.5}$ erg for any spectral type.

5.2 Flare energies

In this flare sample, we recover flares ranging in energy from $10^{29.2}$ to $10^{32.7}$ erg, with a median of $10^{30.6}$ erg. We do not detect any superflares, with energies between 10^{33}–10^{38} erg (Shibayama et al. 2013).

From our flare injection-recovery analysis we see that we have a spectral dependence in the flares we can retrieve. This dependence means that we will struggle to recover flares on the earliest M dwarfs below an energy of ~10^{31} erg. As we move towards later spectral type objects, we are able to recover significantly higher fractions of lower energy flares, as shown in Figure 6.

5.3 Flare frequency distribution

Flare Frequency Distributions (FFDs) explore how often a star will flare with at least a certain energy. The FFD assumes that the following power law applies (Gershberg 1972; Lacy et al. 1976; Hawley et al. 2014)):

$$dN(E) = kE^{-\alpha}dE, \tag{10}$$

where N is the flare occurrence rate, E is the flare energy, and k and α are constants. This can also be represented as:

$$\log(\nu) = C + \beta \log(E), \tag{11}$$

where ν is the cumulative frequency of flares of energies $\geq E$, $C = \log(\frac{k}{1-\alpha})$ and $\beta = 1 - \alpha$. We linearly fit for the two coefficients, C and β, using a least-squares optimisation. Physically, where $\beta > 1$ low-energy flares have the largest contribution to the total energy emitted by flares, whereas when $\beta < 1$, it is the highest energy flares that make up the majority of this total energy.

To generate FFDs we extract only the objects for which we have at least three flares to obtain a good linear fit whilst including as many stars as possible. There are 31 objects in the SSO flaring sample (of 78 objects) which have at least three flares (the 31 objects have between 3–19 flares, with an average of 5 flares). We bin the energies per 0.1 erg and calculate the cumulative frequency of flares greater than that energy. In Figure 8 we present the best-fit power laws from Equation 11. We also present the individual FFDs and line fits for each of the 31 objects in Appendix A, to demonstrate the diversity of power law fits between stars of the same spectral type. We also performed the same FFD fitting using only the actual flare energies (unbinned) which did not change our results. We note that these results are not corrected for the detection efficiency as in Section 5.4 due to the complexity of applying it to the cumulative flare rates for low numbers of flares. Instead we correct for the recovery rate in the next section, where we combine flares from multiple stars to yield a much larger dataset.

By extrapolating the relationship between the flare energy and cumulative frequency for each star, we can predict the amount of energy that flares would deliver to the planets orbiting those stars. However, extrapolating the power-law relationship from our parameter space into the high-energy-flare regime is dangerously unreliable, as it is common to see breaks in these power-law relationships (Silverberg et al. 2016; Paudel et al. 2018). If we instead treat these linear fits as upper limits, then we can give estimates for the maximum frequency of high-energy flares. We note that the power laws also vary with the energy binning and ranges chosen for fitting.

We see a tentative trend that as the stars get cooler, the flares we detect become less energetic and less frequent. However, we see a large diversity of FFD profiles even within the same spectral type. This could be due to variation in stellar age or metallicity. Since we are biased towards detecting lower energy flares for cooler stars, it is possible we more significantly underestimate the rates of lower energy flares for our earlier M stars. This effect may result in steeper power-law slopes, further separating the mid-M dwarfs from the late-M and L dwarfs. However, we only have one M4, M9 and L0 object with at least 3 flares, and within the larger sample of M5–M7 stars, we see little distinction. The decline in flare rates for cooler spectral types could be partially explained by our decreased sensitivity for early-M stars. If only the infrequent, high energy flares are within reach, we only pick up the objects that flare more often. Conversely, since we are more sensitive to detecting lower energy flares on late-M and L dwarfs, we are able to identify stars with lower flaring frequencies. This effect explains the lack of low frequency flaring stars of early spectral types, but not the lack of high frequency flaring stars of later spectral types. As we have a small sample of flaring ultra-cool stars (see Figure 7), we cannot confirm whether they, as a whole, flare less frequently. Paudel et al. (2018) similarly find that the flare rates for low-energy flares decrease as they move towards later spectral types, with L0 and L1 dwarfs having the lowest flaring rates. Critically, however, they find shallower slopes for cool stars (also in agreement with Gizis et al. (2013) and Mullan & Bais (2018)), implying that they have higher occurrence rates for the high-energy flares that are inaccessible to us. High frequency of these flares could be sufficient for a star to enter the abiogenesis zone. However, we are unable to confirm this trend, due to our small
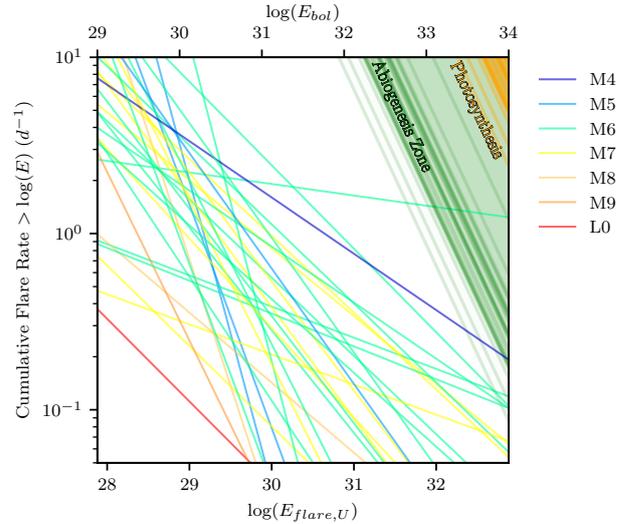
Figure 8. Flare frequency distributions for every star with at least 3 flares. We find the best power-law fit to the cumulative flare rates against flare energy for every star, shown by straight lines. The spectral type of each star is indicated by its colour. The abiogenesis zones (Rimmer et al. 2018) are also shown for each star as the regions shaded in green (similarly to Günther et al. 2020), and the photosynthesis thresholds (Lingam & Loeb 2019a) are shown in orange.

parameter space and the large uncertainties on our individual power law slopes.

If we take a more conservative threshold of at least 5 flares to fit the power law, then we extract FFDs for only M5-M7 stars (for which we have a much larger flare sample). For this sub-sample we find values for α in the range 1.2–2, roughly in agreement with Paudel et al. (2018), who measure 1.3–2 with their sample of 10 UCDs (for a smaller sample, but a similar spectral type range). Paudel et al. (2018) discuss various reasons for the variations in FFD slopes that are not dependent on age or spectral type, such as rotation, stellar spot coverage, and magnetic field topology.

5.3.1 Converting Bolometric Flare Energies to U-band

We calculate the energy in the U band, E_U, by integrating the flux density in the U-band spectral response function, as in Günther et al. (2020). Similarly to Section 3.5, we calculate:

$$E_U = \int_{\text{flare}} A_{\text{flare}} \, dt \int_U R_\lambda B_\lambda(T_{\text{flare}}) \, d\lambda. \tag{12}$$

Therefore,

$$E_U = E \frac{1}{\sigma_{\text{SB}} T_{\text{flare}}^4} \int_U R_\lambda B_\lambda(T_{\text{flare}}) \, d\lambda, \tag{13}$$

where R_λ is now the Johnson U-band response function. The normalisation of the response function is of greater importance here than when calculating the bolometric energies, as it does not cancel out. From this we estimate that 7.6 per cent of the flare's bolometric energy falls in the U band.

5.3.2 Prebiotic chemistry

Here we consider the laboratory work of Rimmer et al. (2018) in defining "abiogenesis zones" around each of our potential planet

hosts, outside of which it is unlikely a planet would receive enough energy for the following prebiotic photochemistry scenario to occur. Specifically, the pathway considered in this work is the synthesis of ribonucteotides, as a precursor to ribonucleic acid (RNA) (Patel et al. 2015; Sutherland 2017; Xu et al. 2018). By considering stellar flares as a mechanism for providing this UV energy, we can calculate the abiogenesis zones around the stars in our sample from their FFDs. Whether or not this chemistry is possible on a planet does not tell us if life has originated there, but instead, whether this mechanism can allow that planet to generate this first building block for RNA. Conversely, if a planet does not receive the necessary energy for this reaction it does not rule out alternate prebiotic pathways for the origins of life.

Günther et al. (2020) determine the necessary flare frequency to power prebiotic chemistry for a planet receiving the same amount of flux from its host as the Earth, by adapting the abiogenesis zone equations from Rimmer et al. (2018) as follows:

$$\nu \geq 25.5 \text{day}^{-1} \left(\frac{10^{34} \text{ erg}}{E_{\text{U}}} \right) \left(\frac{R_*}{R_\odot} \right)^2 \left(\frac{T_*}{T_\odot} \right)^4. \tag{14}$$

From Section 5.3 we calculate the flare energies in the U-band from each flare's bolometric energy, using $E_U = 0.076E$.

If we extrapolate the FFD power laws to $E = 10^{34}$ ergs, we have only one star that provides the necessary UV flux to reach the abiogensesis zone (see Figure 8). This star is an M6 object. If we extrapolate the FFD power laws instead to $E = 10^{36}$ ergs, then we have 13 objects that reach the abiogenesis zone (with spectral types consisting of 1 M4, 1M8, 8 M6 and 3 M7). However, as our sample is confined to probing the low energy flare regime, we have to be careful to not over-extrapolate to high energies. Therefore, we do not extend the power law predictions further than $E = 10^{34}$ ergs. Almost all the FFDs will eventually reach the abiogenesis zone, but at very high energy the uncertainties in our power law fit would be too large to draw any conclusions.

5.3.3 *Photosynthesis*

Lingam & Loeb (2019a) define a threshold for sustaining a biosphere on an Earth-like planet using flare-driven photosynthesis. By considering "photosynthetically active radiation" in the region of 400–750 nm, they find a similar functional form to Equation 14 for the minimum flare rates necessary to receive the same photon flux on a temperate planet as on Earth:

$$\nu \geq 9.4 \times 10^3 \text{day}^{-1} \left(\frac{10^{34} \text{ erg}}{E} \right) \left(\frac{R_*}{R_\odot} \right)^2 \left(\frac{T_*}{T_\odot} \right)^4. \tag{15}$$

This threshold is a significantly greater inhibitor than that for prebiotic chemistry. We plot this threshold for each star with our FFDs in Figure 8. None of our stars exceed this threshold when extrapolating to $E = 10^{34}$ erg. Because the condition for oxygenic photosynthesis is more strict than Equation 14, if this an object satisfies Equation 15, it must also meet the requirements for the abiogenesis zone.

5.4 **Applying our Sensitivity to the Flare Sample**

To compare the average flare occurrence rates for different spectral types, we must account for the incompleteness of our sample. Due to the difficulty of detecting small amplitude flares above photometric scatter, we likely underestimate the frequency of low energy flares. This is often seen as a non-linear 'tail-off' from the expected power law in log-log space (Equation 11).

Jackman et al. (2021) derive the following equation (equivalent to their Equation 3) for the number of flares, N, with energies, E, greater than E_{flare}:

$$N(E \geq E_{\text{flare}}) = \frac{k}{\alpha - 1} \left(R(E_{\text{flare}}) E_{\text{flare}}^{-\alpha+1} + \int_{E_{\text{flare}}}^{E_{\text{max}}} R'(E) E^{-\alpha+1} \, dE \right), \tag{16}$$

where k and α are the same constants in Equation 10, $R(E)$ is the flare recovery fraction, $R'(E)$ is the differentiated flare recovery fraction, and E_{max} is the energy at which the recovery fraction saturates. For high energy flares, where $R = 1$, Equation 16 reduces to the power law in Equation 10, whereas for low energy flares, where $R = 0$, it will reduce to the constant value of the integral (equivalent to the 'tail-off' effect).

Due to the small population of flares on the earliest and latest stars in our sample, we build up flare numbers by combining spectral types into the following five bins: M4–5, M6, M7, M8 and M9–L0. We calculate every star's unique recovery fraction, $R(E)$, based on the results of the injection-recovery tests. We bin flare energies into 12 logarithmically spaced bins from 10^{28} to 10^{34} ergs. For every star we find a value for the recovery fraction in each energy bin by calculating the fraction of flares with energies in that bin that were recovered. We smooth the recovery fraction using a Wiener filter of three bins. We average the recovery fractions within each spectral type bin to get the recovery fraction of an "average" star.

We then extract the observed flare occurrence rates from our flare sample. First, we isolate only the stars in each spectral type bin. We include all flaring and non-flaring stars in the SSO sample, to avoid overestimating our flaring rates. For every energy bin, E, we sum the total number of flares with energies $\geq E$. Then we divide the total number of flares by the sum of the stars' total observation times to produce flaring rates. In doing so we make the assumption that 10 stars observed for 10 hours is equivalent to 1 star observed for 100 hours.

We fit Equation 16 to our observed flare occurrence rates and energies using a Markov Chain Monte Carlo (MCMC) parameter optimisation procedure. To implement MCMC we use the EMCEE Python package (Foreman-Mackey et al. 2013). We used 32 walkers for 10000 steps and the last 2000 steps to sample the posterior distribution. As in Ilin et al. (2019) and Jackman et al. (2021), we multiplied our errors by $R(E)^{-0.5}$ to account for larger uncertainties in recovering the smallest energy flares. The observed flare occurrence rates for each spectral type bin are presented in Table 1 and the results of the best-fit power laws for each spectral type bin are shown in Figure 9. We obtain similar values of $\alpha = 1.88 \pm 0.05$, 1.72 ± 0.02, 1.82 ± 0.02, 1.89 ± 0.07, and 1.81 ± 0.08 for M4–5, M6, M7, M8, and M9–L0 respectively. This work demonstrates that on average the power-law relationship does not change with spectral type in the mid- to late-M regime, despite the large variations within each spectral type. FFDs with similar gradients, but offset (different y-intercepts), implies that these spectral types produce similar relative proportions of high and low energy flares, but the coolest stars have lower rates of flares of all energies.

Our results are consistent with three similar studies of the relation between flaring rate and energy for cool stars. As previously mentioned, Paudel et al. (2018) measured a range of α from 1.3–2.0 for 10 UCDs from K2 lightcurves. Raetz et al. (2020) compiled K2 lightcurves and derived a value for α of 1.83 ± 0.05 for FFDs of the coolest stars in their sample (with spectral types M5–M6). Gizis et al. (2017) found α=1.8 for three UCDs. However, our results have noticeably lower values of α than several other works on mid-to-late M-dwarfs. From Kepler lightcurves, Yang et al. (2017)

Table 1. The observed flare occurrence rates per day for each spectral type and flare energy bin. The parameter ν_{28} represents the rate of flares per day with $E \geq 10^{28}$ ergs. There are no flares detected for any spectral type with $E \geq 10^{33}$ ergs.

Spectral Type	ν_{28}	$\nu_{28.5}$	ν_{29}	$\nu_{29.5}$	ν_{30}
M4–M5	0.7 ± 0.9	0.7 ± 0.6	0.7 ± 0.4	0.7 ± 0.2	0.68 ± 0.17
M6	0.7 ± 0.4	0.7 ± 0.2	0.66 ± 0.17	0.64 ± 0.11	0.60 ± 0.07
M7	0.53 ± 0.08	0.6 ± 0.5	0.6 ± 0.3	0.58 ± 0.16	0.56 ± 0.12
M8	0.15 ± 0.03	0.2 ± 0.2	0.17 ± 0.17	0.17 ± 0.07	0.17 ± 0.04
M9–L0	0.052 ± 0.009	0.14 ± 0.09	0.14 ± 0.05	0.14 ± 0.04	0.10 ± 0.02

Spectral Type	$\nu_{30.5}$	ν_{31}	$\nu_{31.5}$	ν_{32}	$\nu_{32.5}$
M4–M5	0.44 ± 0.08	0.24 ± 0.05	0.055 ± 0.014	0.037 ± 0.009	0.018 ± 0.007
M6	0.43 ± 0.05	0.24 ± 0.02	0.112 ± 0.009	0.047 ± 0.004	0.024 ± 0.003
M7	0.30 ± 0.04	0.140 ± 0.016	0.060 ± 0.008	0.020 ± 0.002	–
M8	0.068 ± 0.014	0.017 ± 0.005	–	–	–
M9–L0	0.017 ± 0.006	–	–	–	–

Table 2. The best-fit power law α, for each spectral type bin, as determined using MCMC.

Spectral Type	α
M4–M5	1.88 ± 0.05
M6	1.72 ± 0.02
M7	1.82 ± 0.02
M8	1.89 ± 0.07
M9–L0	1.81 ± 0.08
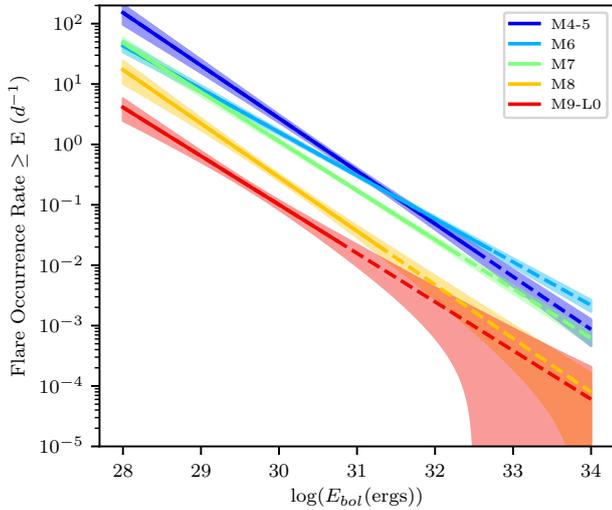
Figure 9. The computed flare occurrence rates (Equation 10), and their 1σ uncertainties, against bolometric flare energy for an average star of each spectral type bin. These flare rates have been corrected for our incompleteness at low flare energies and represent the 'intrinsic' flaring rate (equivalent to Equation 16 at very high energies). The solid lines mark up to the highest energy flare for each spectral type bin, after that the "extrapolated" region is marked by the dotted line.

and Yang & Liu (2019) generated flare catalogues and found on average M dwarfs had a power law indices of $\alpha = 2.09 \pm 0.10$ (for fully convective stars) and $\alpha = 2.07 \pm 0.35$ respectively. For a sample of cool dwarfs observed by TESS, Medina et al. (2020) found $\alpha = 1.98 \pm 0.02$. Seli et al. (2021) specifically isolated TESS's observations of stars similar in spectral type to TRAPPIST–1, for

which they found α=2.11. This work also modified the FFD for TRAPPIST-1 generated by Vida et al. (2017) to include recovery rate, updating their value of α from 1.59 to 2.03 ± 0.02. Lin et al. (2021) used EDEN and K2 lightcurves to analyse the flaring activity of the nearby active M dwarf Wolf 359, for which they found $\alpha = 2.13 \pm 0.14$. We could find a lower value of α if the incompleteness of the sample is not fully described by Equation 16. We see a slightly earlier tail-off at low energies compared to the Jackman et al. (2021) model for all spectral types, which may be evidence of this. Alternatively, previous works may find higher values for α than this work as they have studied higher energy flares on UCD targets using space-based observations. We are limited in this study to only frequent, low energy flares due to our observing strategy. Therefore, it is possible that power laws may steepen at higher flaring energies.

This analysis agrees with the results of Section 5.3, which show that there is a decline in flaring rates as we move to the coolest stars. The injection-recovery tests demonstrate that if flares were present on the lowest-mass stars we would be able to detect them; therefore, these flares must occur too infrequently to be detected with our survey strategy.

5.5 Flares and rotation

From the SSO data sample, we recover 69 (24 A, 22 B, and 23 L) rotators, with periods ranging from 2.2 hours to 65 days. We also find 29 U and 60 N class objects. From here on we refer to 'rotators' as only grade A, B and L.

We find 41 targets for which we can detect both flares and rotation. Therefore, of the 78 flaring objects, described above, 53 per cent have clear rotation. Alternatively, we detect flares on 59 per cent of our rotators. It appears that while the fraction of rotators across all spectral types stays consistently between 20–50 per cent,
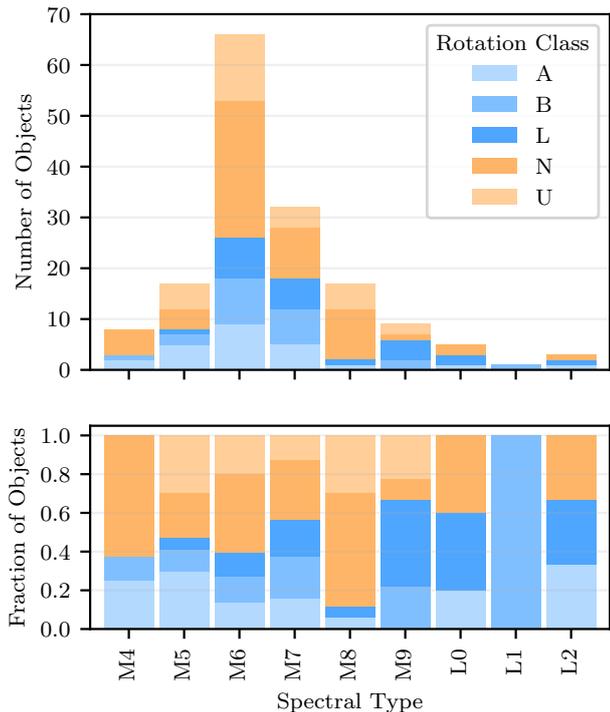
Figure 10. Top: Histogram of objects in the SSO data sample as a function of spectral type, divided by rotation class. The 'rotators' are shown in varying shades of blue, and the 'non-rotators' are shown in shades of orange. **Bottom:** Fraction of different rotator classes for each spectral type.
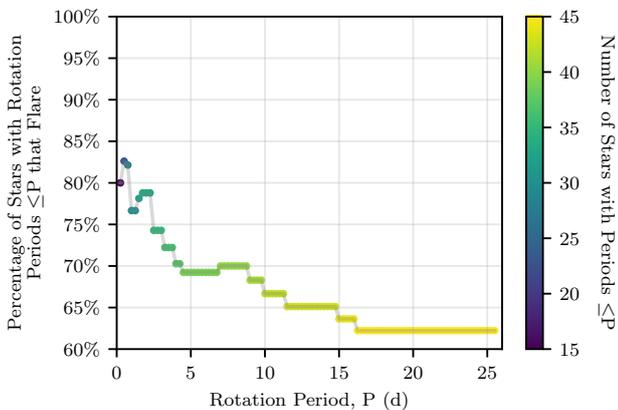
Figure 11. The proportion of stars with rotation periods $\leq P$ that have detected flares, against rotation period P, in steps of 0.25 d. We only include our 46 'A' and 'B' rotators, which range in period from 0.09 d to 25.7 d. The colour scale shows the number of stars with rotation period $\leq P$.

there are increasing proportions of slow rotators for the later M9 and L0 stars.

We see clearly that the very fast rotators are much more likely to flare than slow rotators. By comparing the rotation period in steps of 0.25 d, P, with the proportion of stars with rotation periods $\leq P$ that flare (see Figure 11), we demonstrate that the likelihood that we detect a star as flaring decreases as the rotation slows. In Figure 11 we only include our 46 'A' and 'B' grade rotators, as the periods of the 'L' (long period) grade rotators have very large uncertainties. In

this sample at least 74 per cent of fast-rotating stars, with $P \leq 2$ d, flare. Comparatively, 59 per cent of all rotators flare, 63 per cent of 'A' and 'B' rotators flare, 42 per cent of stars with no detected rotation flare and 50 per cent of all stars flare. We note that the majority (33/46) of 'A' and 'B' rotators have periods ≤ 2 d.

5.5.1 Comparison with MEarth

Within the SSO Sample, 20 targets also appear in the Newton et al. (2018) MEarth-South rotation sample. Of these 20 objects, our results agree for eight, and they suggest periods not found by MEarth for another seven. For the eight objects in agreement: two have long period rotation measured by MEarth, with periods too long to be measured by SPECULOOS (for which we find no rotation); four have no detected rotation with either survey; and two have a similar measured period by both surveys (though one has only a tentative detection with the SSO). For the remaining 12 objects, five have clear, short periods detected by SPECULOOS but not by MEarth, and two have long-period, low-amplitude estimates given by SPECULOOS with no detection in MEarth. We discount the remaining five objects classified as U/N in both surveys, with missing or disagreeing periods. All of the rotation periods measured by the two surveys are shown in Figure 12, and their overlapping observations are reported in Table 3.

The difference in rotation periods measured by the two surveys is likely a result of their different observing strategies. While SPECULOOS performs continuous monitoring of every target for 4–8 hours each night over several weeks, MEarth cycles through multiple targets during a night, returning to each at 20–30 minute cadences. This advantages SPECULOOS to observe very short, <5 hour period rotators, and MEarth to measure very long, >50 day rotation periods. We also derive a long-period rotation estimate with SPECULOOS for two objects with small amplitudes that may not be evident in the MEarth data, due to limitations on the lowest mass stars. MEarth experiences a drop-off in recovery rate for stars with $M < 0.2M_{\odot}$ (which includes all of the objects in our sample) likely due to systematics and the precipitable water effect (Newton et al. 2018).

The Newton et al. (2018) sample shows a clear dichotomy of fast rotators, with periods less than 10 days, and slow rotators, with periods greater than 70 days. We are unable to confirm this, as the majority of our sample have been observed for less than a 10-day span. For the clear slow rotators, which we observe for less than a full phase, we can only assign a long-period estimate (L). The apparent gap in rotation periods from our sample between one and two days is likely a bias from the difficulty in disentangling real rotation from 1-day aliases.

5.5.2 Comparison with TESS

Seli et al. (2021) study the relationships between age, activity and rotation for a sample of 248 "TRAPPIST-1 analogues" from 30-minute TESS Full Frame Image (FFI) observations. They detect a total of 94 flare events on 21 stars. We compare their target catalogue to the SSO data sample and find 35 objects that appear in both. Of these 35 objects, we detect 19 as flaring including five of the seven targets that Seli et al. (2021) identify as flaring. The five objects that both surveys identify as flaring stars are Gaia DR2 2331849006126794880, 2349207644734247808, 3200303384927512960 (as shown in Figure 3), 4825880783419986432 and 5055805741577757824. The two flaring stars detected from TESS lightcurves, on which SPECULOOS does not detect flares are Gaia DR2 4967628688601251200

Table 3. Comparison of SPECULOOS and MEarth rotations for overlapping objects in rotation analyses of both surveys (Newton et al. 2018).

Gaia ID	SSO Class	SSO Period (d)	MEarth Class	MEarth Period (d)
4127182375667696128	U	0.57	A	0.571
5072067381112863104	L	12.04	B	11.067
6056881391901174528	N	N/A	B	156.741
6421389047155380352	N	N/A	U	51.202
6340981796172195584	A	0.22	N	N/A
6405457982659103872	A	0.14	N	N/A
2631857350835259392	A	0.13	N	N/A
5565156633450986752	B	1.65	N	N/A
6914281796143286784	B	1.28	N	N/A
6504700451938373760	L	13.56	N	N/A
2393563872239260928	L	9.53	N	N/A
3005440443830195968	U	3.25	U	0.913
4349305645979265920	U	3.31	N	N/A
4654435618927743872	U	1.23	N	N/A
4854708878788267264	U	3.09	N	N/A
6494861747014476288	U	2.02	U	166.165
6385548541499112448	N	N/A	N	N/A
3175523485214138624	N	N/A	N	N/A
3474993275382942208	N	N/A	N	N/A
3562427951852172288	N	N/A	N	N/A

and 5637175400984142336. For 14 of the targets in this overlapping sample of 35 objects neither survey detects any flares. The 21 overlapping objects (excluding the 14 that have no flares detected by either study) are presented in Table 4.

For 42 stars in their sample Seli et al. (2021) detect rotation periods. Out of the stars that appear in both UCD samples, Seli et al. (2021) find rotation periods for four objects that appear in our catalogue. We find rotation periods for 16 of the 35 objects (with rotation classes A, B or L). Therefore, we can also compare both our flaring and rotation results for this handful of targets. For three of these objects we detect an alias of the period extracted from TESS 30-minute FFIs. For each of these three objects we inspected the SSO lightcurves and periodograms. In each case, the rotation period measured by Seli et al. (2021) is the second most promising period, however, the peaks in the periodogram are very close in power, and the chosen period usually gives a better phase coverage. This is the reason for the 'B' grade we assign to rotators Gaia DR2 2349207644734247808 and 4967628688601251200. Gaia DR2 3200303384927512960 is assigned a grade of 'A', however, it is a unique case. For this target we detect a rotation period of 0.334 d, while Seli et al. (2021) measure a period of 0.50062 d. Rotation periods of 0.33, 0.5 or 1 d are special scenarios for ground-based observations where, over short baselines, we only see the same part of the phase on every observation night. Therefore, due to the limit imposed by the day-night cycle, we are unable to distinguish between these periods. The target for which we do not detect any rotation period, Gaia DR2 4971892010576979840, has an amplitude for periodic photometric variation of 2.1 mmag (measured by Seli et al. (2021)). Upon inspection of this object's SSO lightcurve this periodicity is too small to detect above the level of photometric scatter, and residual systematics.

Since TESS is a space-based mission which does not contend with the day-night cycle or ground-based systematics, and they perform 27 day continuous monitoring, this provides a large number of phases for fast-rotating stars, which can counteract lower photometric precisions for red dwarfs. Therefore, we would expect rotation periods derived from TESS lightcurves to be more accurate than those from ground-based observations.

For their catalogue of stars, Seli et al. (2021) find that ~8 per cent are flaring. This flaring fraction contrasts this work, where we detect flares on 28 of the 57 M7–M9 objects within the SSO sample (49 per cent). There may be several reasons for this discrepancy. Firstly, Seli et al. (2021) detect flares from lightcurves with a low time resolution of 30 minutes. TESS is also not optimised for very red objects such as UCDs, therefore the photometric precision is lower than for SPECULOOS. Both of these factors inhibit a flare survey on TESS data from detecting low energy flares (which have small amplitudes and short durations) on UCDs. Instead, continuous 27-day monitoring of each TESS field allows Seli et al. (2021) to probe the rarer, high energy (even superflare) flare regime. This is demonstrated by their recovery fraction, which is less than 20 per cent for flares with energies below $10^{31.5}$ ergs, compared to SSO recovery fractions of 50–70 per cent (depending on spectral type) for the same energy. Additionally, in this paper we demonstrate that cooler stars flare less frequently at all energies, therefore it is likely that a survey focused on high energy flares on cool stars would detect significantly fewer flares. Finally, space-based data is less impacted than ground-based data by time-varying systematics due to the Earth's rapidly-changing atmosphere, therefore the risk of false positives is presumably lower.

6 DISCUSSION

6.1 Extension into the ultra-cool regime

The breadth of the SSO flare sample allows us to compare the flaring and rotating behaviour of mid-M stars with ultra-cool dwarfs.

Table 4. Comparison of this work and Seli et al. (2021) for 21 overlapping objects with flares detected in either flaring analysis.

Gaia ID	SSO Flaring?	SSO Period (d)	SSO Amp. (mmag)	SSO Rotation Class	TESS Flaring?	TESS Period (d)	TESS Amp. (mmag)
2331849006126794880	Y	N/A	N/A	N	Y	N/A	N/A
2349207644734247808	Y	0.687	17.0	B	Y	2.13867	20.9
3200303384927512960	Y	0.334	24.1	A	Y	0.50062	3.8
4825880783419986432	Y	65.468	26.1	L	Y	N/A	N/A
5055805741577757824	Y	25.662	6.8	L	Y	N/A	N/A
4967628688601251200	N	0.682	6.5	B	Y	0.40391	4.2
5637175400984142336	N	0.573	5.2	U	Y	N/A	N/A
3421840993510952192	Y	4.297	22.3	B	N	N/A	N/A
3493736924979792768	Y	0.142	8.8	A	N	N/A	N/A
3504014060164255104	Y	16.326	5.4	L	N	N/A	N/A
4404521333221783680	Y	N/A	N/A	N	N	N/A	N/A
4928644747924606848	Y	11.819	7.8	L	N	N/A	N/A
4971892010576979840	Y	N/A	N/A	N	N	0.70254	2.1
5047423236725995136	Y	N/A	N/A	N	N	N/A	N/A
5156623295621846016	Y	6.752	3.9	B	N	N/A	N/A
5392287051645815168	Y	0.182	8.6	A	N	N/A	N/A
5469802724480366848	Y	N/A	N/A	N	N	N/A	N/A
56252256123908096	Y	0.614	29.4	A	N	N/A	N/A
5657734928392398976	Y	N/A	N/A	N	N	N/A	N/A
6357834388848708224	Y	2.034	15.7	U	N	N/A	N/A
6733860940302404864	Y	0.27	4.4	A	N	N/A	N/A

We predominantly detect flares on M5–M7 stars (~60-70 per cent flare), and decreasing proportions of flaring stars for both earlier and later M dwarfs. Since there is a spectral dependency to our flare detection (Figure 6), it is likely that we are missing the lower energy flares on mid-M dwarfs, and therefore underestimating the proportion of flaring stars. On the other hand, these higher mass objects do produce rarer, high-energy flares more frequently than late-M dwarfs (Lacy et al. 1976). These high-energy flares should be easier to detect. Our results also agree with other surveys, with around ~ 30 per cent of M4 stars flaring. It is possible that these surveys are limited in different ways for mid-M stars, which exist at the redder end of their samples, rather than at the bluer end of ours.

Lacy et al. (1976) showed that the flares on mid-M dwarfs were more frequent and lower energy than on early-M dwarfs. Hilton et al. (2010) found that this trend continued into the late-M dwarf regime (M6–M8). From our findings we see little difference between spectral types up to M7. However, our results extend beyond M8, where we find that flares of all energies appear to become less frequent. From our injection-recovery tests, we conclude that if higher energy flares were present on the reddest stars in our sample, then we would detect them; therefore, the lack of high-energy flares for late-M and L dwarfs on the sampled timescales of 0.8 to 8.7 days is likely real.

We see a relationship between flaring activity and rotation, especially for our very fastest rotators. The stars with rapid rotation periods of <5 days, are much more likely to flare than the rest of the SSO sample. At least 74 per cent of the stars with rotation periods faster than 2 days flare, compared to around 59 per cent of all rotating stars and 42 per cent of non-rotating stars. Therefore, we are more likely to detect flares on stars where we can determine a rotation period than stars we cannot. Possible explanations are that fast rotators flare much more frequently, more energetically, or are more likely to flare than stars with no detectable rotation

period. Being unable to detect a rotation period may be a symptom that a star has a period too long to be measured by our survey, or the photometric variations on the rotation timescale are too small to be seen, possibly due to small spot covering fractions, a low spot-photosphere contrast, or a particular orientation of spots (e.g. distributed at the poles). This contrast between rotators and non-rotators is further enhanced for the very fastest periods ($P \leq 2$ d). Assessing the spectral distribution of our fastest rotators shows that almost all of the M4–M7 stars with periods less than 2 days flare. Comparatively, very few objects with spectral types later than M8 flare, fast-rotating or otherwise. Stars with fast rotation will have more magnetic energy available; therefore, these results are consistent with faster rotators exhibiting more active behaviours (Davenport 2016; Allred et al. 2015; Newton et al. 2017; Yang & Liu 2019; Günther et al. 2020; Martínez et al. 2019). Specific activity-rotation studies on cool stars, such as those carried out by West et al. (2015), Stelzer et al. (2016), Medina et al. (2020), and Raetz et al. (2020), have also found that fast rotating stars flare more frequently than slow rotators. Seli et al. (2021) found that ~50 per cent of their fast-rotating ($P < 5$ days) late-M stars flared, compared to ~70 per cent of our sample with the same periods.

In this work, we have compared results for SPECULOOS targets that also appear in the rotation study performed by Newton et al. (2018) on MEarth data, and in the activity-rotation study from Seli et al. (2021) on TESS observations. SPECULOOS and MEarth are two ground-based surveys focused on cool stars, but with very different observing strategies. SPECULOOS is a fast-cadence survey that targets ultra-cool dwarfs for relatively short baselines of 4–8 days. MEarth, however, revisits its M dwarf targets only every 20–30 minutes, over a period of several months. With SPECU-LOOS's lightcurves, therefore, we are able to study the short-term photometric variability of an object, such as its flaring and fast rotation (<10 d), whereas with MEarth's lightcurves we can monitor
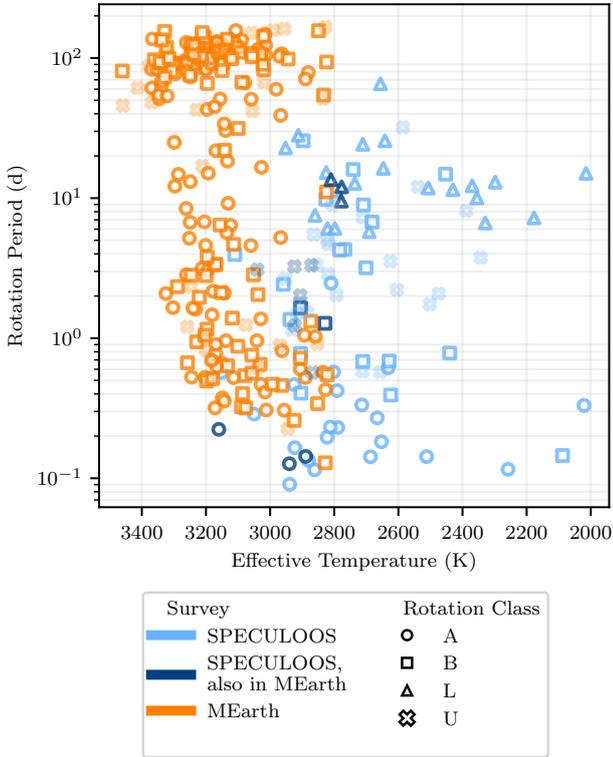
Figure 12. A comparison of the rotation periods measured in this SSO data sample with the rotation periods from MEarth (Newton et al. 2018). The objects observed only by SPECULOOS are in light blue and the objects observed only by MEarth are in light orange. The dark blue points are the measured period in SPECULOOS for objects that have also been observed by MEarth. The different rotation classes are indicated by the marker shapes, with the uncertain (U) rotation periods made transparent. Note that the L rotators, a class unique to SPECULOOS, provide only lower limit estimates of the periods, due to our shorter observing baselines.

an object's long-term variability, such as any slow rotation (>50 d) and stellar surface evolution. SPECULOOS and MEarth, however, are both ground-based surveys which have to contest with the challenges of observing through the Earth's atmosphere and the loss of data from the day-night cycle. TESS, a space telescope, is able to perform continuous 27 d monitoring of each field, however ultra-cool dwarfs are not its ideal targets. Therefore, TESS lightcurves of UCDs are mostly only available from the 30-minute FFIs, and have generally lower precisions. SPECULOOS has the short exposure times and high precisions necessary to identify frequent low energy flares, however for rotation, the irregular data sampling and large gaps (limited by weather conditions and the observing cycle) make it difficult to identify the correct period above its aliases. For this reason we recommend caution when interpreting the rotation results from this survey, and provide 'rotation classes' to help identify targets with rotation period uncertainties. TESS is only able to detect the rarer high energy flares (with large amplitudes and slow decay times), however, its ability to provide continuous phase coverage of fast-rotating objects, and to observe up to hundreds of cycles, make it well suited for rotation studies. Comparisons such as these, between ground and space-based surveys, and between fast and slow-cadence observing strategies highlight that they are highly complementary. Leveraging the differences and the overlaps between different surveys allows us to fully explore the photometric

activity parameter space of UCDs, and unearth more information about the subsequent impact on any potential planets they might host.

6.2 Planetary Impact

Often used to put stellar flares into context, the largest solar flare ever directly observed on Earth was the 'Carrington Event' in 1859 (Carrington 1859; Hodgson 1859). This flare released an energy of 10^{32} erg and the associated coronal mass ejection hit the Earth's magnetosphere, resulting in the largest geomagnetic storm on record. This storm disrupted telegraph systems and produced auroras reaching almost as far as the Equator. While our flare sample is relatively low-energy compared to the wealth of recorded flares on other stellar types, in our sample we detect 14 flares with energies greater than 10^{32} erg on 10 different stars. Characterising these interactions between planets and their host stars, especially those which have the capability to initiate or destroy life, is vital to our understanding of planetary habitability.

6.2.1 Habitability of Earth-sized planets around UCDs

From calculating FFDs for the stars in our sample, we can place energetic bounds on whether flares can drive photosynthesis (Section 5.3.3) or provide their planets with enough UV flux for the prebiotic chemistry described in Section 5.3 to occur. If we extrapolate the FFD power laws to $E = 10^{34}$ erg, we have only one M6 star that provides this necessary UV flux. However, there are several caveats to these results. As we see a non-linear "tail-off" effect for low energy flares due to our detection limits we likely underestimate low energy flares rates, resulting in shallower FFD power laws in Section 5.3. Therefore, without accounting for the completeness of our flare sample this work only places upper bounds on the stars whose planets can reach the abiogensesis zone, or which can sustain photosynthesis. As we move towards cooler stars, we see that the frequency of flaring decreases (see Figure 8). However, there are only one of each of the M4, M9 and L0 stars, with at least 3 flares detected. Due to our low detection efficiency for mid-M dwarfs, we are biased to detect only the most frequently flaring stars; therefore, it is likely that we would see an artificial increase in flare rate for M4 objects. In contrast, we should be more sensitive to flares of all energies on our late-M and L dwarf stars, so this bias does not explain the decrease in flare frequency for these stars.

By including our detection sensitivity in Section 5.4 and considering an "average" star for each spectral type, we can conclude more general relationships between flare frequency and spectral type. While there is a wide diversity of flare frequencies for individual stars' within a given spectral type, on average cooler stars have lower flaring rates. If this trend holds for a larger sample of UCDs, it may have severe consequences on the search for life around the very lowest mass objects. Similar conclusions have been reached by work on the flaring activity of TRAPPIST-1 (Glazier et al. 2020; Ducrot et al. 2020; Seli et al. 2021) and "TRAPPIST-1 analogues" (Seli et al. 2021).

If the majority of late-M and L dwarfs have flaring rates too low to initiate the synthesis of ribonucleotides, then for life to originate it must happen through a different mechanism. Possible alternatives are surface hydrothermal vents (Rimmer & Shorttle 2019), aerial biospheres (Lingam & Loeb 2019b), impact-shock synthesis (Furukawa et al. 2015) or extraterrestrial delivery (Rimmer & Shorttle 2019; Krijt et al. 2017). We also find that none of our sample meets

the requirement set out by Lingam & Loeb (2019a) to sustain an Earth-like biosphere, due to the significantly lower photosynthetically active radiation (with wavelengths of 400–750 nm) incident within UCD habitable zones. These results are consistent with the work of Mullan & Bais (2018) and Lingam & Loeb (2019a). However, we are considering an Earth-centric viewpoint, and it has been suggested that photosynthesis could proceed using deep-sea hydrothermal vents (Beatty et al. 2005) or at infrared wavelengths, where UCDs are brightest (Scalo et al. 2007; Takizawa et al. 2017; Claudi et al. 2021).

The activity of M dwarfs is predicted to decrease as they age (West et al. 2008, 2015; Paudel et al. 2019). Therefore, it is possible that earlier in these stars' life cycles they would produce more energetic and frequent flares, with enough UV flux to trigger prebiotic chemistry on their planets, or even for photosynthesis. The results in this paper only represent a snapshot of each object at its current point in time. We do not consider age in this work due to the huge uncertainties on stellar age estimates for these objects.

Modelling of the impact of stellar flares on modern Earth-like analogues by Tilley et al. (2019) shows that stars with at least 0.1 flare per day with energies above 10^{34} erg will strip the ozone layers from any terrestrial planets they host. We do not extend our FFDs to energies above 10^{34} erg as we are probing the low energy flare regime; therefore, we do not consider ozone depletion in our analysis.

Here we have only discussed two aspects of how flares could impact the habitability of planets around UCDs. In reality, we would need to apply a more holistic approach. There are a multitude of other factors (such as orbit dynamics, stellar variability, tidal locking, atmospheric composition, stellar age, etc.) that influence the environments on planets hosted by UCDs, and other possible pathways for life that haven't been considered here.

7 CONCLUSIONS

By analysing the high-precision photometry produced by SPECU-LOOS, we can monitor the activity of late-M and ultra-cool dwarfs. This allows us to probe further into the ultra-cool regime, expanding our understanding of the relationships between activity and rotation into a new parameter space. This work can be used in conjunction with previous and ongoing flare studies on earlier M dwarfs, such as those carried out on TESS, K2, and MEarth observations, to assess the diversity of M dwarfs and how well we can extrapolate the results of these studies to lower mass objects. We demonstrate the benefit of a large-scale flare study of UCDs, such as this, to extending existing flare catalogues.

Using our flare sample to make more general predictions about flaring activity as we move to very low mass stars, we address two different scenarios in which flares may assist the origin and sustenance of life on planets they host. Firstly, the quiescent UV flux of cool stars is too low for the synthesis of ribonucleotides, a major step in prebiotic chemistry. Secondly, the visible light incident on planets around cool stars is also found to be too low for photosynthesis to occur. In both cases, while we find it is possible that flares on active M dwarfs could provide the necessary extra radiation, the frequency of flares of all energies appears to decrease with spectral type. These two considerations alone suggest UCD systems may not be favourable sites for abiogenesis, though other factors—known and unknown—certainly impact the likelihood of finding life in these systems as well.

DATA AVAILABILITY

The data underlying this article is not currently public. The SPECULOOS-South Consortium will make all SPECULOOS-South Facility reduced data products available to the ESO Science Archive Facility following the regular Phase 3 process as described at http://www.eso.org/sci/observing/phase3.html.

REFERENCES

Allred J. C., Kowalski A. F., Carlsson M., 2015, Astrophysical Journal, 809
Anfinogentov S., Nakariakov V. M., Mathioudakis M., Van Doorsselaere T., Kowalski A. F., 2013, The Astrophysical Journal, 773, 156
Ballard S., Johnson J. A., 2016, The Astrophysical Journal, 816, 66
Beatty J. T., et al., 2005, Proceedings of the National Academy of Science, 102, 9306
Benz A. O., Güdel M., 2010, Annual Review of Astronomy and Astrophysics, 48, 241
Borucki W. J., et al., 2010, Science, 327, 977
Broeg C., Fernández M., Neuhåuser R., 2005, Astronomische Nachrichten, 326, 134
Buccino A. P., Lemarchand G. A., Mauas P. J., 2007, Icarus, 192, 582
Burdanov A., Delrez L., Gillon M., Jehin E., 2018, in Deeg H. J., Belmonte J. A., eds, , Handbook of Exoplanets. p. 130, doi:10.1007/978-3-319-55333-7_130
Cantrell J. R., Henry T. J., White R. J., 2013, The Astronomical Journal, 146, 99
Carrington R. C., 1859, Monthly Notices of the Royal Astronomical Society, 20, 13
Chabrier G., 2003, Publications of the Astronomical Society of the Pacific, 115, 763
Chabrier G., Baraffe I., 1997, Astronomy and Astrophysics, 327, 1039
Claudi R., et al., 2021, Life, 11, 1
Covone G., Ienco R. M., Cacciapuoti L., Inno L., 2021, Monthly Notices of the Royal Astronomical Society, 505, 3329
Davenport J. R. A., 2016, The Astrophysical Journal, 829
Davenport J. R., et al., 2014, Astrophysical Journal, 797, 122
Davenport J. R. A., Covey K. R., Clarke R. W., Boeck A. C., Cornet J., Hawley S. L., 2019, The Astrophysical Journal, 871, 241
Delrez L., et al., 2018, Ground-based and Airborne Telescopes VII, 10700, 49
Demory B. O., et al., 2020, Astronomy and Astrophysics, 642, A49
Dressing C. D., Charbonneau D., 2015, The Astrophysical Journal, 807, 45
Ducrot E., et al., 2020, Astronomy and Astrophysics, 640
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, Publications of the Astronomical Society of the Pacific, 125, 306
Furukawa Y., Nakazawa H., Sekine T., Kobayashi T., Kakegawa T., 2015, Earth and Planetary Science Letters, 429, 216
Gardner J. P., et al., 2006, Space Science Reviews, 123, 485
Gershberg R. E., 1972, Astrophysics and Space Science, 19, 75
Gillon M., 2018, Nature Astronomy, 2, 344
Gillon M., Jehin E., Magain P., Chantry V., Hutsemékers D., Manfroid J., Queloz D., Udry S., 2011, EPJ Web of Conferences, 11, 06002
Gizis J. E., Burgasser A. J., Berger E., Williams P. K. G., Vrba F. J., Cruz K. L., Metchev S., 2013, Astrophysical Journal, 779
Gizis J. E., Paudel R. R., Schmidt S. J., Williams P. K. G., Burgasser A. J., 2017, The Astrophysical Journal, 838, 22
Glazier A. L., Howard W. S., Corbett H., Law N. M., Ratzloff J. K., Fors O., Ser D. d., 2020, The Astrophysical Journal, 900, 27
Günther M. N., et al., 2020, The Astronomical Journal, 159, 60
Hawley S. L., Davenport J. R., Kowalski A. F., Wisniewski J. P., Hebb L., Deitrick R., Hilton E. J., 2014, Astrophysical Journal, 797, 121
Henry T. J., 2004, in Spectroscopically and Spatially Resolving the Components of the Close Binary Stars. pp 159–165
Hilton E. J., West A. A., Hawley S. L., Kowalski A. F., 2010, The Astronomical Journal, 140, 1402
Hodgson R., 1859, Monthly Notices of the Royal Astronomical Society, 20, 15
Howard W. S., Corbett H., Law N. M., Ratzloff J. K., Glazier A., Fors O., Ser D. d., Haislip J., 2019, The Astrophysical Journal, 881, 9
Howard W. S., et al., 2020a, The Astrophysical Journal, 895, 140
Howard W. S., et al., 2020b, The Astrophysical Journal, 902, 115
Howell S. B., et al., 2014, Publications of the Astronomical Society of the Pacific, 126, 398
Ilin E., Schmidt S. J., Davenport J. R. A., Strassmeier K. G., 2019, Astronomy & Astrophysics, 622, 133

Jackman J. A., et al., 2019, Monthly Notices of the Royal Astronomical Society: Letters, 485, L136
Jackman J. A. G., et al., 2021, Monthly Notices of the Royal Astronomical Society, 000, 1
Jehin E., et al., 2011, The Messenger, 145, 2
Jehin E., et al., 2018, The Messenger, 174, 2
Kaltenegger L., Traub W. A., 2009, Astrophysical Journal, 698, 519
Kerber F., et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. p. 84463N, doi:10.1117/12.924340
Kowalski A. F., Hawley S. L., Wisniewski J. P., Osten R. A., Hilton E. J., Holtzman J. A., Schmidt S. J., Davenport J. R., 2013, Astrophysical Journal, Supplement Series, 207, 15
Krijt S., Bowling T. J., Lyons R. J., Ciesla F. J., 2017, The Astrophysical Journal, 839, L21
Lacy C. H., Moffett T. J., Evans D. S., 1976, Astrophysical Journal Supplement Series, 30, 85
Lammer H., et al., 2007, Astrobiology, 7, 185
Law N. M., et al., 2015, Publications of the Astronomical Society of the Pacific, 127, 234
Lehmer O. R., Catling D. C., Parenteau M. N., Hoehler T. M., 2018, The Astrophysical Journal, 859, 171
Lépine S., Bongiorno B., 2007, The Astronomical Journal, 133, 889
Lienhard F., et al., 2020, Monthly Notices of the Royal Astronomical Society, 497, 3790
Lin C.-L., et al., 2021, The Astronomical Journal, 162, 11
Lingam M., Loeb A., 2019a, Monthly Notices of the Royal Astronomical Society, 485, 5924
Lingam M., Loeb A., 2019b, The Astrophysical Journal, 883, 143
Lomb N. R., 1976, Astrophysics and Space Science, 39, 447
Lurie J. C., Davenport J. R., Hawley S. L., Wilkinson T. D., Wisniewski J. P., Kowalski A. F., Hebb L., 2015, Astrophysical Journal, 800, 95
Martínez R. R., Lopez L. A., Shappee B. J., Schmidt S. J., Jayasinghe T., Kochanek C. S., Auchettl K., Holoien T. W. S., 2019, The Astrophysical Journal, 892, 144
McCormac J., Pollacco D., Skillen I., Faedi F., Todd I., Watson C. A., 2013, Publications of the Astronomical Society of the Pacific, 125, 548
Medina A. A., Winters J. G., Irwin J. M., Charbonneau D., 2020, The Astrophysical Journal, 905, 107
Moffett T. J., 1974, Technical Report 273, UV CETI FLARE STARS: OBSERVATIONAL DATA
Mondrik N., Newton E., Charbonneau D., Irwin J., 2018, The Astrophysical Journal, 870, 10
Mullan D. J., Bais H. P., 2018, The Astrophysical Journal, 865, 101
Murray C. A., et al., 2020, Monthly Notices of the Royal Astronomical Society, 495, 2446
Namekata K., et al., 2017, The Astrophysical Journal, 851, 91
Newton E. R., Irwin J., Charbonneau D., Berta-Thomspon Z. K., West A. A., 2016, in Proceedings of the International Astronomical Union. Cambridge University Press, pp 124–125, doi:10.1017/S174392131500602X
Newton E. R., Irwin J., Charbonneau D., Berlind P., Calkins M. L., Mink J., 2017, The Astrophysical Journal, 834, 85
Newton E. R., Mondrik N., Irwin J., Winters J. G., Charbonneau D., 2018, The Astronomical Journal, 156, 217
Niraula P., et al., 2020, The Astronomical Journal, 160, 172
Noyes R. W., Hartmann L. W., Baliunas S. L., Duncan D. K., Vaughan A. H., 1984, Astrophysical Journal, 279, 763
Nutzman P., Charbonneau D., 2008, Publications of the Astronomical Society of the Pacific, 120, 317
Patel B. H., Percivalle C., Ritson D. J., Duffy C. D., Sutherland J. D., 2015, Nature Chemistry, 7, 301
Paudel R. R., Gizis J. E., Mullan D. J., Schmidt S. J., Burgasser A. J., Williams P. K. G., Berger E., 2018, The Astrophysical Journal, 858, 55
Paudel R. R., Gizis J. E., Mullan D. J., Schmidt S. J., Burgasser A. J., Williams P. K., Youngblood A., Stassun K. G., 2019, Monthly Notices of the Royal Astronomical Society, 486, 1438

Raetz S., Stelzer B., Damasso M., Scholz A., 2020, Astronomy & Astrophysics, 637, A22

Ranjan S., Wordsworth R., Sasselov D. D., 2017, The Astrophysical Journal, 843, 110

Ricker G. R., et al., 2015, Journal of Astronomical Telescopes, Instruments, and Systems, 1, 14003

Rimmer P. B., Shorttle O., 2019, Life, 9, 12

Rimmer P. B., Xu J., Thompson S. J., Gillen E., Sutherland J. D., Queloz D., 2018, Science Advances, 4, eaar3302

Scalo J., et al., 2007, Astrobiology, 7, 85

Scargle J. D., 1982, The Astrophysical Journal, 263, 835

Schmidt S. J., et al., 2019, The Astrophysical Journal, 876, 115

Seager S., Deming D., Valenti J. A., 2009, Astrophysics and Space Science Proceedings, 10, 123

Sebastian D., et al., 2021, Astronomy & Astrophysics, 645, A100

Segura A., Walkowicz L. M., Meadows V., Kasting J., Hawley S., 2010, Astrobiology, 10, 751

Seli B., Vida K., Moór A., Pál A., Oláh K., 2021, Astronomy & Astrophysics, 650, A138

Shappee B. J., et al., 2014, Astrophysical Journal, 788, 48

Shibayama T., et al., 2013, Astrophysical Journal, Supplement Series, 209, 5

Silverberg S. M., Kowalski A. F., Davenport J. R. A., Wisniewski J. P., Hawley S. L., Hilton E. J., 2016, The Astrophysical Journal, 829, 129

Skumanich A., 1972, The Astrophysical Journal, 171, 565

Skumanich A., 1986, The Astrophysical Journal, 309, 858

Stelzer B., Damasso M., Scholz A., Matt S. P., 2016, Monthly Notices of the Royal Astronomical Society, 463, 1844

Sutherland J. D., 2017, Nature Reviews Chemistry, 1, 0012

Takizawa K., Minagawa J., Tamura M., Kusakabe N., Narita N., 2017, Scientific Reports, 7, 1

Tilley M. A., Segura A., Meadows V., Hawley S., Davenport J., 2019, Astrobiology, 19, 64

Vida K., Kővári Z., Pál A., Oláh K., Kriskovics L., 2017, The Astrophysical Journal, 841, 124

West A. A., et al., 2004, The Astronomical Journal, 128, 426

West A. A., Hawley S. L., Bochanski J. J., Covey K. R., Reid I. N., Dhital S., Hilton E. J., Masuda M., 2008, The Astronomical Journal, 135, 785

West A. A., Weisenburger K. L., Irwin J., Berta-Thompson Z. K., Charbonneau D., Dittmann J., Pineda J. S., 2015, Astrophysical Journal, 812

Wheatley P. J., et al., 2018, Monthly Notices of the Royal Astronomical Society, 475, 4476

Williams P. K., Berger E., Irwin J., Berta-Thompson Z. K., Charbonneau D., 2015, Astrophysical Journal, 799, 192

Wright N. J., Drake J. J., 2016, Nature, 535, 526

Wright N. J., Newton E. R., Williams P. K., Drake J. J., Yadav R. K., 2018, Monthly Notices of the Royal Astronomical Society, 479, 2351

Xu J., Ritson D. J., Ranjan S., Todd Z. R., Sasselov D. D., Sutherland J. D., 2018, Chemical Communications, 54, 5566

Yang H., Liu J., 2019, The Astrophysical Journal Supplement Series, 241, 29

Yang H., et al., 2017, The Astrophysical Journal, 849, 36

de Wit J., Seager S., 2013, Science, 342, 1473

APPENDIX A: FLARE FREQUENCY DISTRIBUTIONS

In this appendix, we present the FFDs for each object with at least three flares. This includes the binned cumulative frequency against bolometric flare energy data and the best fit lines.

This paper has been typeset from a T_EX/L^AT_EX file prepared by the author.
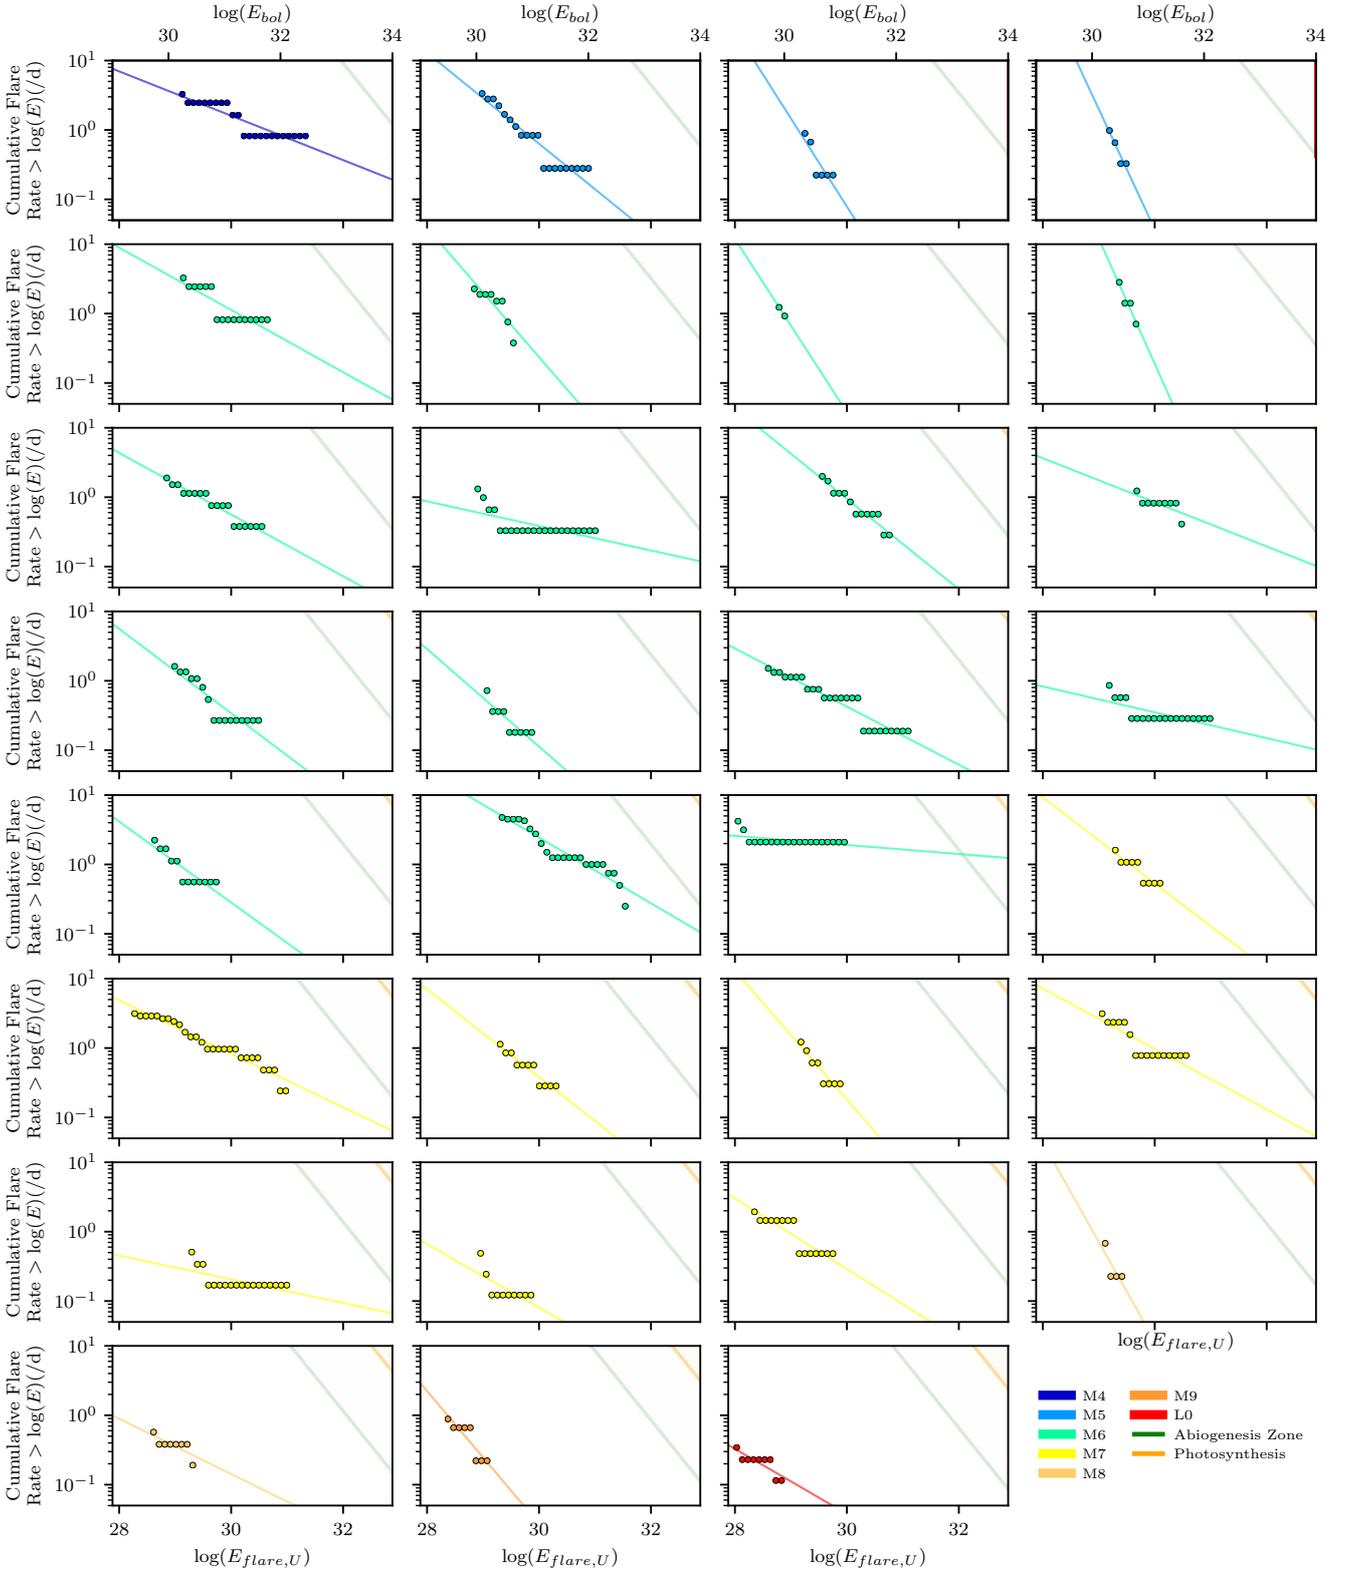
Figure A1. The FFDs for each of the 31 objects with at least three flares. The cumulative flare rates against bolometric flare energies (binned every $10^{0.1}$ ergs in log space) are shown by the points as well as the best-fit power law for each star. The plots are ordered in increasing spectral type and the colour of the points and lines correspond to that object's spectral type. The thresholds to reach the abiogenesis zone and perform oxygenic photosynthesis are shown for each object by the green and orange lines respectively.