

A statistical search for a uniform trigger threshold in solar flares from individual active regions

JULIAN B. CARLIN,^{1,2} ANDREW MELATOS,^{1,2} AND MICHAEL S. WHEATLAND³

¹*School of Physics, University of Melbourne, Parkville, VIC 3010, Australia*

²*OzGrav, University of Melbourne, Parkville, VIC 3010, Australia*

³*Sydney Institute for Astronomy, School of Physics, The University of Sydney, NSW 2006, Australia*

(Received 2022 December 19; Revised 2023 February 20; Accepted 2023 March 10)

Submitted to The Astrophysical Journal

ABSTRACT

Solar flares result from the sudden release of energy deposited by sub-photospheric motions into the magnetic field of the corona. The deposited energy accumulates secularly between events. One may interpret the observed event statistics as resulting from a state-dependent Poisson process, in which the instantaneous flare rate is a function of the stress in the system, and a flare becomes certain as the stress approaches a threshold set by the micro-physics of the flare trigger. If the system is driven fast, and if the threshold is static and uniform globally, a cross-correlation is predicted between the size of a flare and the forward waiting time to the next flare. This cross-correlation is broadly absent from the *Geostationary Operational Environmental Satellite (GOES)* soft X-ray flare database. One also predicts higher cross-correlations in active regions where the shapes of the waiting time and size distributions match. Again there is no evidence for such an association in the *GOES* data. The data imply at least one of the following: i) the threshold at which a flare is triggered varies in time; ii) the rate at which energy is driven into active regions varies in time; iii) historical flare catalogs are incomplete; or iv) the description of solar flares as resulting from a build-up and release of energy, once a threshold is reached, is incomplete.

Keywords: Astrostatistics, solar flares

1. INTRODUCTION

Broad consensus exists that the micro-physical process triggering individual solar flares is magnetic reconnection in the corona. A reconnection event becomes more likely to occur the more magnetic energy accumulates in the time between flares (Priest & Forbes 2002; Benz 2016). This phenomenological class of stress-relax model was popularized by Rosner & Vaiana (1978), and expounded by Wheatland & Glukhov (1998); Wheatland (2000a, 2008); Kanazir & Wheatland (2010); Hudson (2019), among others. A complementary phenomenological description is the avalanche or self-organized criticality model (Lu & Hamilton 1991; Lu et al. 1993; Lu 1995a), inspired by the canonical sandpiles of Bak et al. (1987, 1988). These two classes of model are broadly

compatible, but make different predictions about some long-term statistical observables (Lu 1995b; Boffetta et al. 1999; Wheatland 2000a; Lippiello et al. 2010; Farhang et al. 2019; Aschwanden & Johnson 2021).

Upon aggregating historical datasets, the probability density function (PDF) of the energy released in solar flares is found to be a power law or log-normal over multiple decades (Rosner & Vaiana 1978; Lu & Hamilton 1991; Verbeeck et al. 2019). The PDF of time intervals between successive flares in the same active region, henceforth termed waiting times, is less clear-cut. Wheatland (2000b) found evidence that the power-law-like shape of the tail of the waiting time PDF found by Boffetta et al. (1999) is explained by a sum of exponentials, with individual rates themselves drawn from an exponential PDF. This interpretation is further developed by Aschwanden et al. (2021). Flaring rates that vary in time are also noted by Lepreti et al. (2001) and Gorbets & Messerotti (2012). Cross-correlations between the size of a flare and the subsequent (preceding) wait-

ing time, henceforth termed forward (backward) cross-correlations, are a key differentiator between the stress-relax and avalanche descriptions — the latter predicts no size–waiting-time cross-correlations (Jensen 1998). Forward and backward correlations are broadly absent from solar flares datasets (Biesecker 1994; Hudson et al. 1998; Crosby et al. 1998; Wheatland 2000a; Hudson 2020), with the exception of strong forward cross-correlations found in two active regions (Hudson 2019).

Rotational glitches in rotation-powered pulsars (Lyne & Graham-Smith 2012; Haskell & Melatos 2015) are an analogous astrophysical phenomenon to solar flares, in the limited sense that they are consistent with a stress-relax process even though they do not involve magnetic reconnection as far as one knows (Aschwanden et al. 2018). While the exact process that triggers a glitch is unknown, most models are encompassed by the fundamental idea that “stress” (possibly differential rotation or elastic deformation) builds up secularly between glitches, and is released sporadically and partially at a glitch. The instantaneous glitch rate is assumed to grow with the stress in the system. This idea is formalized phenomenologically in the state-dependent Poisson (SDP) process popularized by Fulgenzi et al. (2017). Precise, falsifiable predictions about size and waiting time PDFs, auto-, and cross-correlations, as well as comparisons to current datasets, show the power and flexibility of the SDP model in the neutron star context (Melatos et al. 2018; Carlin & Melatos 2019a,b; Melatos & Drummond 2019). The same falsifiable predictions also deliver new physical insights when applied to solar flare data, as we show in this paper.

Our goal in this paper is to search the *Geostationary Operational Environmental Satellite (GOES)* soft X-ray flare database for signatures of a threshold-driven stress-relax process. We do this by de-aggregating the data from different active regions, and studying summary statistics of flare waiting times and sizes. In Section 2 we outline the SDP framework, and how it maps to solar flares. In Section 3 we explore various regimes of the SDP process. We also specify precise, falsifiable tests for the question of whether solar flares are triggered once the energy reaches a static threshold, if it obeys a stress-relax process. The *GOES* soft X-ray dataset to which we apply these tests is described in Section 4. In Section 5 we look for associations between matching waiting time and size PDFs and high size–waiting-time cross-correlations, as predicted for the SDP process. We conclude with a discussion of the microphysical implications of the data analysis in Section 6.

2. STATE-DEPENDENT POISSON PROCESS

2.1. Equation of motion

The SDP process is a doubly stochastic renewal process which models the “stress” in the system as a function of time, $X(t)$, as

$$X(t) = X(0) + t - \sum_{i=0}^{N(t)} \Delta X_i, \quad (1)$$

where X and t are expressed respectively in dimensionless units of X_{cr} (the critical stress in the system at which a stress-release event becomes certain), and τ (the time taken for the system to accumulate the critical stress X_{cr} , in the absence of any stress-release events). We attach a physical interpretation to $X(t)$ in the solar flare context in Section 2.3 and Appendix A. The amount of stress released at the i -th event, ΔX_i , is a random variable, drawn from a user-specified PDF, $\eta[\Delta X_i | X(t_i^-)]$, where $X(t_i^-)$ is the stress in the system immediately prior to the i -th event. In the standard configuration, $\eta[\Delta X_i | X(t_i^-)]$ is fixed as a power law, but other options exist (Carlin & Melatos 2019a, 2021). Making η conditional on $X(t_i^-)$ is necessary to ensure that the stress remains positive-definite and is plausible physically. The second random variable in Equation (1) is $N(t)$, a stochastic function which counts the number of events up to time t . It is determined iteratively via the waiting time between each event.

We assume the instantaneous event rate is a monotonically increasing function of the stress in the system, viz.

$$\lambda[X(t)] = \frac{\alpha}{1 - X(t)}, \quad (2)$$

where $\alpha = \lambda_0 \tau$ is a dimensionless control parameter, and $\lambda_0 = \lambda(X = 1/2)/2$ is a reference rate. The long-term statistical output of the SDP process does not depend strongly on the functional form of $\lambda[X(t)]$, so long as it diverges in the limit $X \rightarrow 1$ as in Equation (2) (Fulgenzi et al. 2017; Carlin & Melatos 2019a). Equation (2) implies that the probability of an event occurring approaches one, as the stress approaches the critical stress. As the stress increases deterministically between events, the PDF of waiting times Δt following the i -th event is that of a variable-rate Poisson process (Cox 1955; Fulgenzi et al. 2017)

$$p[\Delta t | X(t_i^+)] = \lambda[X(t_i^+) + \Delta t] \times \exp \left\{ - \int_{t_i^+}^{t_i^+ + \Delta t} dt' \lambda[X(t')] \right\}, \quad (3)$$

where $X(t_i^+)$ is the stress immediately following the event at time t_i .

2.2. Monte Carlo automaton

Analytically solving the coupled equations (1)–(3) to calculate the PDFs of event waiting times or sizes is usually not feasible, except for particular choices of $\eta[\Delta X_i | X(t_i^-)]$; see section 6 of Fulgenzi et al. (2017) for an example. Instead, it is simple to run the following automaton to generate numerical solutions:

1. Pick Δt from Equation (3), given the current stress X .
2. Update the stress to $X + \Delta t$ to account for the deterministic evolution.
3. Pick ΔX from $\eta[\Delta X | X + \Delta t]$, and subtract it from the stress.
4. Repeat from step 1.

Given α and $\eta[\Delta X_i | X(t_i^-)]$ the automaton generates a time-ordered sequence of waiting times and sizes. From the sequence we can calculate the long-term PDFs for waiting times and sizes (Fulgenzi et al. 2017; Carlin & Melatos 2019a), as well as the cross-correlation (Melatos et al. 2018; Carlin & Melatos 2019a), autocorrelations (Carlin & Melatos 2019b), and other observables.

2.3. Mapping to solar flares

We identify the stress that accumulates between flares and relaxes at a flare with the spatially-averaged magnetic energy density in a given active region. One could equally choose a different physical quantity, such as the magnetic shear, depending on the particular microphysics of the flare trigger. We assume that active regions have independent stress reservoirs, i.e. the coronal magnetic fields of different active regions do not interact strongly, and all flares from the same active region extract energy from one reservoir. An idealized toy model relating the foregoing definition of stress to magnetic energy input from the photosphere is outlined in Appendix A, as a simplified but concrete illustration of the physical picture under consideration.

Two major simplifying assumptions are that the rate at which energy is fed into the reservoir, τ^{-1} , is constant in time, and so is the critical, spatially-averaged magnetic energy density, X_{cr} , for a given active region. These assumptions are motivated in Appendix A, but a key goal of the paper is to test their veracity. We also assume that flares reduce $X(t)$ instantaneously. This assumption is defensible as the typical time between flares (typically hours) is much larger than the typical duration of flaring events (typically minutes) (Fletcher et al. 2011). In what follows we do not prescribe a particular

functional form for $\eta[\Delta X_i | X(t_i^-)]$, to keep the SDP process as flexible as possible.

The SDP process generates sequences of dimensionless waiting times, $\{\Delta t_i^{\text{SDP}}\}$ and sizes, $\{\Delta X_i^{\text{SDP}}\}$. It is natural to ask how they correspond to directly observable sequences of waiting times, $\{\Delta t_i^{\text{obs}}\}$ and sizes, $\{\Delta s_i\}$. The waiting times satisfy

$$\Delta t_i^{\text{obs}} = \Delta t_i^{\text{SDP}} \tau, \quad (4)$$

where τ is unknown a priori but can be estimated, as discussed in Appendix B. The sizes are related in a more complicated way, because some energy is released through channels other than the soft X-ray flux. We make the simplifying assumption that the peak soft X-ray flux multiplied by the duration of the flare (henceforth Δs_i for the i -th flare in a given active region) is proportional to the stress released from the reservoir. This implies

$$\Delta s_i \propto \Delta X_i^{\text{SDP}} X_{\text{cr}}. \quad (5)$$

The unknown constant of proportionality in Equation (5) hampers direct parameter estimation of X_{cr} from an observed size PDF.

Full parameter estimation for the SDP process given a set of observed waiting times and sizes is discussed in Melatos & Drummond (2019). It lies outside the scope of this paper. In Appendix B we also describe for completeness an alternative parameter estimation procedure using a hierarchical Bayesian scheme. The advantage of the hierarchical scheme is that it clarifies as a matter of principle the information content and flow in the estimation problem, i.e. the parameter combinations that can be estimated uniquely, and the data components that inform each parameter estimate. The disadvantage is that it can be implemented in practice only if $\eta[\Delta X_i | X(t_i^-)]$ is drawn from a certain class of mathematical functions. The favored class produces waiting time and size PDFs which do not match solar flare observations (see Section 5), so the hierarchical scheme is not applied to real data in this paper. Nonetheless, it is included for the benefit of the reader, who may wish to develop it further for solar flare analysis in the future.

3. OBSERVABLE SIGNATURES OF A RAPIDLY-DRIVEN PROCESS

Broadly speaking, an SDP process operates in one of two regimes: slowly-driven, with $\alpha \gg 1$, and rapidly-driven, with $\alpha \ll 1$. In this section we identify the key dynamics in both regimes, in Sections 3.1 and 3.2 respectively. We then infer, in Section 3.3, a set of observable signatures, which if absent rule out the operation of a rapidly-driven process of the form described in Section 3.2.

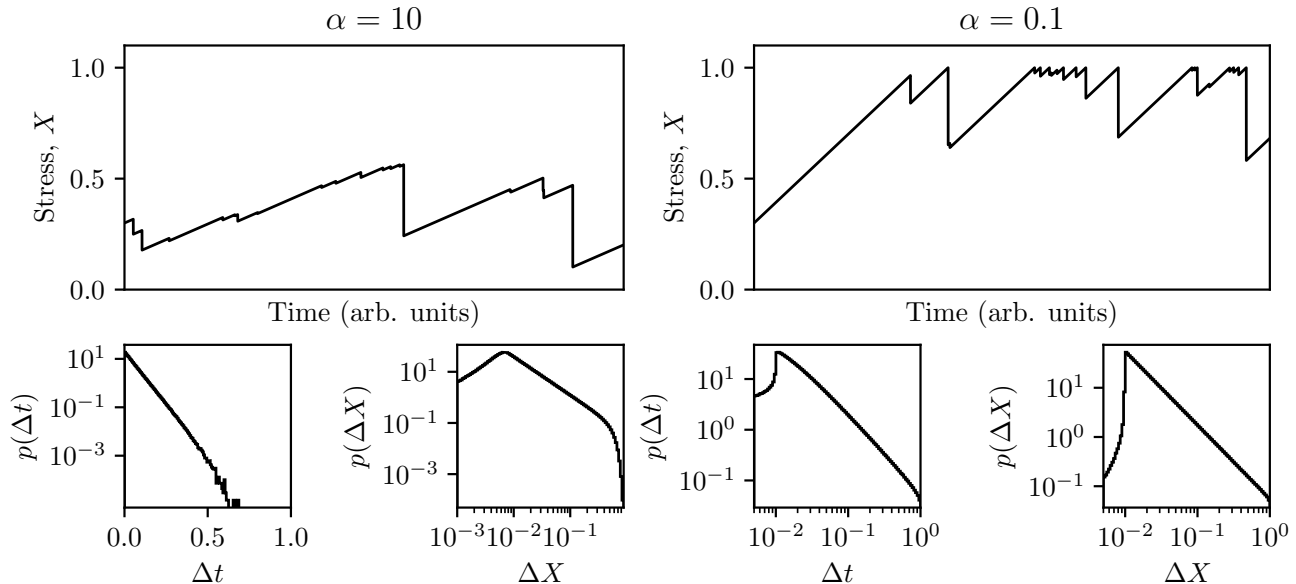


Figure 1. Top panels: qualitative behavior of the evolution of the stress, $X(t)$, in the SDP process in the two main regimes, slowly driven (left panel, $\alpha = 10$), and rapidly driven (right panel, $\alpha = 0.1$). Bottom panels: predicted waiting time and size PDFs $p(\Delta t)$ and $p(\Delta X)$ respectively in the above regimes. For all panels we fix $\eta[\Delta X_i | X(t_i^-)] \propto (\Delta X_i)^{-3/2} H[\Delta X_i - 10^{-2} X(t_i^-)]$, where the Heaviside function $H(\dots)$ enforces a minimum stress-release size of 1% of the stress in the system, ensuring integrability. The histograms in the bottom panels each include $N = 10^7$ events.




Figure 2. Forward, ρ_+ (solid curve), and backward, ρ_- (dashed curve), cross-correlations between event sizes and waiting times for 10^2 values of $10^{-2} \leq \alpha \leq 10^2$, with $N = 10^5$ events generated per value of α . We fix $\eta[\Delta X_i | X(t_i^-)]$ to the same functional form as in Figure 1. A version of this figure first appeared as figure 13 in Fulgenzi et al. (2017).

3.1. Slow driver

When the system is slowly driven, we have $X(t) \ll 1$, as events are usually triggered before the stress accumulates to near the threshold. When $\eta[\Delta X_i | X(t_i^-)]$ is a power law, the predicted waiting time PDF $p(\Delta t)$ is an exponential, and the predicted size PDF $p(\Delta X)$ is a power law over multiple decades. We show the qualitative behavior of the stress in the slowly driven regime in the top-left panel of Figure 1, where for the sequence of 20 stress-release events shown we have $0.1 < X(t) < 0.6$. In the bottom-left panels of the

same figure we show the predicted waiting time and size PDFs. The size PDF $p(\Delta X)$ has a turn-over in logarithmic slope at $\Delta X \approx 10^{-2}$ due to the particular choice of $\eta[\Delta X_i | X(t_i^-)]$ (see caption for the functional form).

Depending on the choice of $\eta[\Delta X_i | X(t_i^-)]$, there may be detectable backward cross-correlations (i.e. a correlation between the size of an event and the waiting time since the preceding event). This is because the size of an event cannot exceed the amount of stress in the system, so longer periods of stress-accumulation allow for the possibility of larger events. This backward cross-correlation is especially pronounced with the choice $\eta[\Delta X_i | X(t_i^-)] \propto \delta[\Delta X - X(t_i^-)]$, where $\delta(\dots)$ is the Dirac-delta function. This choice forces the stress reservoir to empty at each event, and collapses the SDP process to other stochastic processes in the literature, e.g. forest fire models (Daly & Porporato 2006). In the solar flare context, a backward cross-correlation is a long-standing prediction of the Rosner & Vaiana (1978) stress build-up model; and the “reset” model of Hudson et al. (1998); Hudson (2019). If instead $\eta[\Delta X_i | X(t_i^-)]$ prefers small stress-release events, e.g. if it is a power law (Fulgenzi et al. 2017), the backward cross-correlation is small.

3.2. Fast driver

When the system is rapidly driven, we have $X(t_i^-) \lesssim 1$, i.e. the stress is driven close to the threshold before each event. As the dynamics of the system are strongly

influenced by the presence of the threshold, we sometimes call this regime “threshold-driven”. A prediction of the SDP process in this regime is that the size and waiting time PDFs should “match”, i.e. they should be the same distribution, up to a linear scaling (Fulgenzi et al. 2017). For example, if $\eta [\Delta X_i | X(t_i^-)]$ is a power law, both the waiting time and size PDFs are power laws. We show this in the bottom-right panels of Figure 1. We also show the qualitative behavior of the stress versus time in this regime in the top-right panel of the same figure, where we see $X(t) \rightarrow 1$ (i.e. unit critical stress) before every stress-release event.

A strong forward cross-correlation (i.e. a correlation between the size of an event and the waiting time to the next event) is observed in this regime. This is because a large stress-release event results in a longer delay before the system accumulates enough stress to approach the threshold. In the solar flares context, this cross-correlation is predicted by the “saturation” model of Hudson et al. (1998) and Hudson (2019). The smooth transition between the fast-driven regime ($\alpha \ll 1$), with high forward cross-correlations, and the slow-driven regime ($\alpha \gg 1$), is shown in Figure 2.

3.3. Heuristic test for a threshold-driven process

Ideally, we fit the parameters of the SDP process (τ , X_{cr} , λ_0 , and $\eta [\Delta X_i | X(t_i^-)]$) to a paired sequence of solar flare sizes and waiting times, and infer what regime applies. In practice, however, at least some of these parameters vary between active regions, which limits us to sequences of length $N \lesssim 50$ (typical active regions have at most dozens of events; see Section 4), which are insufficient to infer four or more parameters. Instead, we combine the results in Sections 3.1 and 3.2 to deduce qualitative features, which must be present in multiple observables simultaneously (e.g. PDFs, cross-correlations), if the SDP process operates in a particular α regime or indeed operates at all. Specifically, we summarize the behavior in Sections 3.1 and 3.2 into the following observationally testable prediction. If solar flares are well-modeled with a rapidly-driven process (e.g. the SDP process with $\alpha \ll 1$), then we should see large forward cross-correlations, accompanied by waiting time and size PDFs with the same shape in individual active regions. This coincident signature should become more prominent as α (or a proxy thereof) decreases. We quantify this prediction in Section 5.2.

With the test above we also ameliorate the impact of a potentially mis-specified model. Although it is agnostic about the microphysics, the SDP process is not the only possible prescription of stress accumulation and release. If some of the assumptions outlined in Section 2.3

do not hold, a different underlying model may underpin solar flares. For example, one may build a model in which stress accumulates according to a Brownian random walk with some underlying drift, until a threshold is breached, at which point a stress-release event is triggered (Carlin & Melatos 2020). In the Brownian stress accumulation model, $X(t)$ is always driven to the threshold before each event. If drift occurs faster than diffusion, we would predict the same coincident signature as for the SDP process: large forward cross-correlations are accompanied by matching waiting time and size PDFs.

4. GOES SOFT X-RAY OBSERVATIONS

The X-ray Sensor (XRS) on the *Geostationary Operational Environmental Satellites (GOES)* has continuously monitored the Sun in soft X-rays (0.5 to 8 Å) since 1975. In Section 4.1 we introduce the flare summary data analyzed in this paper. In Section 4.2 we remind the reader of the obscuration effect described in Wheatland (2001), explain how it affects our analysis, and briefly touch on other biases the *GOES* flare detection algorithm may have on the completeness of the catalog as a whole. In Section 4.3 we analyze aggregated waiting time and size PDFs across all active regions.

4.1. Flare data

We use the publicly available flare summary data hosted by the National Geophysical Data Center (NGDC)¹ for flares before June 29 2015, and by the Space Weather Prediction Center (SWPC)² for flares after June 28 2015. We combine these two data sources into a homogeneous, cleaned database (henceforth “catalog”), with some anomalies corrected as described in Appendix C. These data are collated with the flare start epochs, t^s , peak epochs, t^p , and end epochs, t^e . The peak epoch is defined as the epoch which contains the highest peak flux, after the start epoch. The end epoch is defined as the epoch at which the flux returns to half of the difference between the peak flux and the background flux. The peak flux (irradiance) of each flare, f^p , is calculated with reference to the recorded flare class, in units of W m^{-2} . The longitude and latitude of the flare are also often available, if the flare is associated with an active region. For clarity, we henceforth denote as $x_{i,k}$ an arbitrary measurement of the variable x in the i -th flare from the k -th active region.

We define the waiting time between two flare epochs as $\Delta t_{i,k} = t_{i+1,k}^p - t_{i,k}^p$. One could equally use the

¹ <ftp://ftp.ngdc.noaa.gov/STP/space-weather/solar-data/solar-features/solar-flares/x-rays/GOES/xrs/>

² <ftp://ftp.swpc.noaa.gov/pub/indices/events/>




Figure 3. Top panel: Probability mass function of the number of flares per active region, N_k , for all active regions with $N_k \geq 5$. Regions with $0 \leq N_k \leq 4$ are excluded, as the number of samples per region is too small for the statistical tests in this paper. Bottom panel: PDF of the flare rate, λ_k , for all active regions with $N_k \geq 5$, binned into 50 uniformly spaced bins between 0 hours^{-1} and 0.5 hours^{-1} .

flare start or end epochs to define the waiting time. We use the flare peak epoch as it is less influenced by the background level, as discussed more in Section 4.2. We define the flare size as $\Delta s_{i,k} = f_{i,k}^p (t_{i,k}^e - t_{i,k}^s)$; that is, we multiply the irradiance by the duration of the flare to obtain a flare size with dimensions of energy per unit area. The duration of an active region is $\Delta T_k = t_{N_k,k}^e - t_{1,k}^s$, where there are N_k flares in the region. The average flare rate in an active region is $\lambda_k = N_k / \Delta T_k$. The average λ_k overestimates systematically the true flare rate, as the duration ΔT_k is between two flare epochs, both of which are counted in N_k . An unbiased estimate would take ΔT_k as the difference between the epochs when the active region appears and disappears, but neither epoch is recorded in the *GOES* flare summary data³. A concern we touch on further in Section 4.3 is that $\lambda_k = N_k / \Delta T_k$ assumes that the flare catalog is complete, i.e. all flares from a given region are detected and correctly attributed to that region.

As of September 1 2022, there are 83283 flares in the catalog, 49328 of which are associated with an active region. There are 2429 active regions with $N_k \geq 5$, accounting for 42982 flares. The median number of flares in an active region (when considering only regions with

$N_k \geq 5$) is 12, but the average is 17.7, as the distribution is peaked at $N_k = 5$ and monotonically decreases with N_k . We show the probability mass function of the number of flares in an active region, $\text{Pr}(N_k)$, in the top panel of Figure 3 for regions with $N_k \geq 5$. The PDF of observed flare rates, $p(\lambda_k)$, across all active regions with $N_k \geq 5$, is displayed in the bottom panel of the same figure. The distribution of λ_k is well-described by a log-normal, with mean 0.08 hours^{-1} , and standard deviation 0.7 hours^{-1} . Henceforth, we typically only include regions with $N_k \geq 5$ in our analysis unless stated otherwise, as many of the subsequent statistical tests, such as the cross-correlation(s), have low statistical power with smaller sample sizes.

4.2. Obscuration and flare size bias of subsequent flares after a large flare

As described in section 2.2 of Wheatland (2001), the *GOES* flare detection algorithm involves a selection effect that obscures the detection of flares following a large flare, due to the enhanced background soft X-ray flux. To detect a flare, the flux must monotonically increase for four consecutive minutes, with the last value 1.4 times the value three minutes earlier. Hence a flare must produce a 40% increase above the background flux. However, while flares typically rise rapidly to peak flux, the soft X-ray emission is observed to decay on the timescales of hours (Benz 2016). Thus, even large flares may be obscured by the enhanced background flux following, say, an X1 ($f^p = 10^{-4} \text{ W m}^{-2}$) class flare.

The above detection algorithm introduces a secondary bias in the catalog. A flare, say with recorded peak flux f_i^p , that occurs during the decay of another flare, say with recorded peak flux f_{i-1}^p , has an enhanced peak flux compared to if the same flare occurs during a period of low background flux, viz.

$$f_i^p \approx f_{i,\text{true}}^p + f_{i-1}^p \exp\left(\frac{-\Delta t_{i-1}}{\tau}\right), \quad (6)$$

where τ is the timescale on which the peak flux from the $(i-1)$ -th flare decays, and $f_{i,\text{true}}^p$ is the “true” peak flux of the i -th flare. A secondary effect is that the duration of the i -th flare will have a reduced duration due to the decaying contribution of the previous flare to the background flux. That is, during the i -th flare, the background flux decays by a factor of $\exp[-(t_i^e - t_i^s)/\tau]$, which reduces the time taken for the flux to decay back to half the difference between the peak and pre-flare fluxes. These two effects run counter to one another, as Δs_i is defined as the product of peak flux and duration. They are likely present in a small subset of the catalog, as only 1.2% of all flares from regions with $N_k \geq 5$ have start times before the end time of the previous flare.

³ The epochs when the active region appears and disappears may be extracted from other records in some instances. Such an analysis lies outside the scope of this paper.

A comprehensive re-analysis of the soft X-ray flux time-series, with appropriate background subtraction, may reveal flares that were not detected with the original flare detection algorithm and correct the biases in flare size outlined above. Such a re-analysis lies outside the scope of this paper. Acknowledging that the *GOES* catalog is incomplete, we mitigate the impact on our analysis by creating a secondary masked catalog that only includes flares with $\Delta s \geq 10^{-3} \text{ J m}^{-2}$ (i.e. class C1 multiplied by the median duration of 10^3 s , and higher), where we know the fraction of missed flares is smaller. This masked catalog is used in Sections 4.3 and 5.1 when we perform comparative studies and parametric fits for the waiting time and size PDFs.

4.3. Aggregate size and waiting time statistics

When we aggregate flare waiting times and sizes across all active regions with $N_k \geq 5$, we obtain the PDFs shown in the top and bottom panels of Figure 4 respectively. The data are shown as the black histograms. The grey region in the bottom panel shows the flare sizes that are not included in the masked catalog, as described in Section 4.2. The difference between the masked and full catalog is not visible in the top panel as the effect on $p(\Delta t)$ is minimal.

Three empirical trial distributions with common analytic forms are overlaid on each PDF for the masked catalog. The overlaid distributions have parameters fixed to their maximum likelihood values, given the data. By eye, it is clear that the log-normal distribution best describes the aggregated waiting time PDF⁴, $p(\Delta t)$, while a power law distribution best describes the aggregated size PDF, $p(\Delta s)$. We formalize this comparison using the corrected Akaike Information Criterion (AICc) (Akaike 1974; Hurvich & Tsai 1989). The AICc calculates the model, among a set of possible models, which minimizes the information loss, while accounting for potential bias due to the number of model parameters and the sample size. When calculated for $p(\Delta t)$, the relative probability for a log-normal describing the data over an exponential or a power law is $e^{3 \times 10^3}$ and $e^{3 \times 10^4}$ respectively. When calculated for $p(\Delta s)$, the relative probability for a power law describing the data over an exponential or a log-normal is $e^{3 \times 10^4}$ and $e^{4 \times 10^3}$ respectively. If we instead only mask flares with peak flux less than 10^{-6} W m^{-2} (i.e. class C1), we find that a log-normal best describes the data. This preference is also noted in Verbeeck et al. (2019). It arises because of the




Figure 4. Top panel: PDF $p(\Delta t)$ of waiting times (thick black stepped curve), Δt , aggregated over all active regions with $N_k \geq 5$, binned into 50 linearly spaced bins between the minimum and maximum Δt in the full, unmasked catalog. Overlaid are the best fit estimates of an exponential (blue dashed curve), power law (orange dotted curve), and log-normal (green dot-dashed curve) PDF. Bottom panel: As for top panel, but for sizes Δs , and the data is binned into 50 logarithmically spaced bins. The grey region in the bottom panel shows the difference between the size PDFs from the masked and full catalogs (see Section 4.2).

small number of flares in this alternatively-masked catalog with $10^{-4} < \Delta s / \text{J m}^{-2} \lesssim 10^{-3}$. If we fit the full, unmasked catalog (i.e. include the flares shaded in grey in the bottom panel of Figure 4) we find that $p(\Delta s)$ is again fitted best with a log-normal. We remind the reader that $p(\Delta s)$ is an observed distribution; it is a function of both the underlying generative physics (i.e. how much energy is released in each flare) and the systematic observational biases (i.e. how many and what flares are detected or not). The masking described in Section 4.2 is a first pass at accounting for some of these biases.

If we assume that the same α and $\eta [\Delta X_i | X(t_i^-)]$ apply to all regions, Figure 4 is inconsistent with the SDP framework for any choice of α and $\eta [\Delta X_i | X(t_i^-)]$. However, as we discuss in Section 2.3, there is good reason to believe that α (and perhaps even $\eta [\Delta X_i | X(t_i^-)]$) may vary region-to-region.

⁴ Waiting times that are best described with a log-normal are seen in many contexts. We provide a brief survey in Appendix D for the interested reader.

5. DISAGGREGATED DATA IN INDIVIDUAL ACTIVE REGIONS

The goal of this section is to search for signatures of a threshold-driven SDP process in individual active regions, rather than considering all flares in aggregate. In Section 5.1 we consider only waiting time and size PDFs, without regard to their potential cross-correlation. In Section 5.2 we calculate these cross-correlations, and in Section 5.3 we apply the test outlined in Section 3.3 by searching for an association between matching waiting time and size PDFs and the cross-correlation, in individual active regions. In Section 5.4 we perform a preliminary investigation of the longer-term memory in the system by calculating the autocorrelation between subsequent waiting times and subsequent sizes. The analysis in Section 5.1 involves parametric fitting of the size PDFs, so we use the masked catalog described in Section 4.2. However in Sections 5.2–5.4 we use the full, unmasked catalog, as the tests performed in these sections are non-parametric.

5.1. Waiting time and size PDFs versus flare rate

For regions with $N_k \geq 5$ we disaggregate the data, and ask what PDF shape best characterizes each region’s waiting time and size PDFs, rather than considering only the aggregated dataset as in Section 4.3. Using the AICc, we find that for waiting time PDFs, 63% of regions are fitted best with an exponential distribution, 20% with a log-normal, and the remainder with a power law. For size PDFs, 63% of regions are fitted best with a power law distribution, 24% with an exponential, and the remainder with a log-normal. These proportions are broadly consistent with previously published results, i.e. that while the aggregated $p(\Delta t)$ is fitted best with a log-normal, individual regions are often fitted best with an exponential (with varying rates) (Wheatland 2000b). Individual regions have $p(\Delta s)$ that is usually fitted best with a power law, but can occasionally be better represented with log-normal or exponential distributions, especially in regions with lower N_k . As an example of the different shapes that distributions of waiting times and sizes can have in different regions, we display four arbitrary but representative per-active-region fits in Appendix E.

One may reasonably ask whether there are clear trends, or predictors, for the waiting time and size PDFs in any given region. An exhaustive search for predictors, e.g. with a multiple regression analysis, lies outside the scope of this paper, but one sensible first step is to see if the shape that fits best evolves with λ_k . This test is performed for the waiting time and size PDFs in the top and bottom panels of Figure 5 respectively.

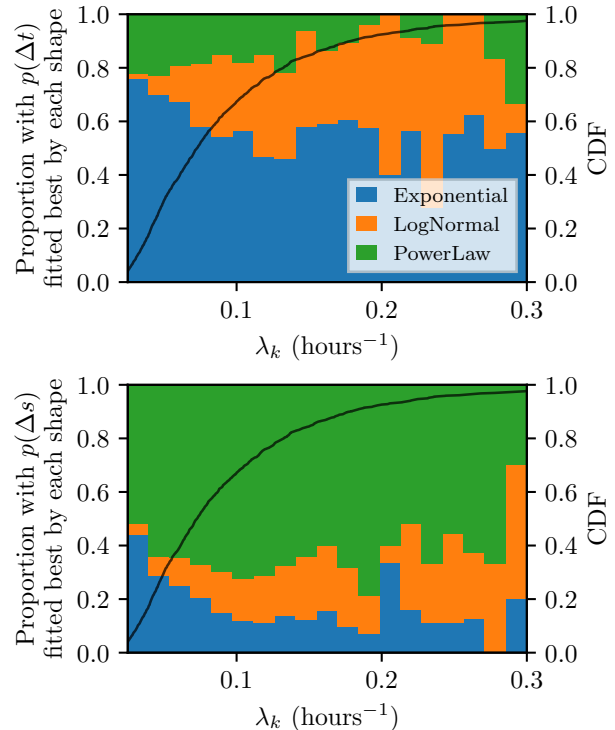


Figure 5. Top panel: Proportion of active regions with waiting time PDF fitted best by three possible shapes (a power law, an exponential, and a log-normal) as determined by the AICc, as a function of the region’s flare rate, λ_k . The CDF of λ_k is overlaid in black. Bottom panel: As in the top panel, but for size PDFs.

This figure is constructed by binning all active regions with $N_k \geq 5$ into 20 uniformly spaced bins between $\lambda_k = 0.025$ hours⁻¹ and $\lambda_k = 0.3$ hours⁻¹, then calculating which distribution fits best the waiting times and sizes using the AICc. For the waiting time PDFs, the proportion of regions fitted best with a power law stays roughly constant as λ_k increases, at around 15%, while the proportion of regions fitted best with a log-normal grows with λ_k , at the expense of the exponential. For the size PDFs we see a similar trend, with the proportion of regions fitted best with a power law staying roughly constant as λ_k increases, this time at around 65%, while the proportion of regions fitted best with a log-normal grows with λ_k , at the expense of the exponential. In both panels we plot the cumulative distribution function (CDF) of λ_k to remind the reader that each λ_k bin does not contain the same number of regions; 90% of regions with $N_k \geq 5$ have $0.028 < \lambda_k / (1 \text{ hour}^{-1}) < 0.27$.

Can the evolution of the proportion of each shape versus λ_k be explained with the SDP model? Under the assumption that λ_k is a tracer of the driving rate, i.e. $\lambda_k \propto \alpha^{-1}$, we expect to see a greater proportion of exponentially distributed waiting times at low λ_k (high

α). This is the regime in which the stress does not approach the threshold at $X = X_{\text{cr}}$ before each event, so waiting times are un-correlated with sizes and are (broadly) Poissonian, i.e. the waiting times are exponentially distributed. This expectation broadly conforms with what we see in the top panel of Figure 5.

The evolution versus λ_k in the bottom panel of Figure 5 is harder to explain with the SDP model, if each region has the same $\eta[\Delta X_i | X(t_i^-)]$. We expect $p(\Delta s) \propto \eta$, when α is low, but we see all of the three possible shapes represented at the highest values of λ_k . This implies at least one of the following: i) η truly varies from one region to the next, which implies different stress-release mechanisms are at play in different regions; or ii) $\lambda_k \gtrsim 0.2 \text{ hours}^{-1}$ does not correspond to $\alpha \ll 1$, and hence $p(\Delta s) \propto \eta$; or iii) the small sample size of events in each region results in the AICc not favoring the “true” size distribution; or iv) the obscuration effects described in Section 4.2 is stronger with higher λ_k , due to the enhanced background flux in regions that have many flares in a short period of time. To test iii), we generate Figure 5 again for regions with $N_k \geq 10$, instead of $N_k \geq 5$. For both the waiting times and sizes, the proportion fitted best by an exponential drops by $\sim 10\%$ in each λ_k bin, while the proportion fitted best by a log-normal increases. Qualitatively, however, the evolution with λ_k remains consistent with what is seen in Figure 5.

5.2. Size–waiting-time cross-correlations

The correlation between flare sizes and subsequent waiting times, i.e. the “forward” cross-correlation, is denoted as $\rho_{+,k}$, while the correlation between flare sizes and the preceding waiting times, i.e. the “backward” cross-correlation, is denoted as $\rho_{-,k}$. We calculate these correlations using the Spearman correlation coefficient (Lehmann & D’Abrera 2006).

The PDFs of forward and backward cross-correlations measured in all active regions with $N_k \geq 5$ are displayed in the top panel of Figure 6. The PDFs are broad, mostly because the median N_k is 12 (i.e. small). However, $p(\rho_{+,k})$ and $p(\rho_{-,k})$ are different; the latter has a median of $\rho_{-,k} = 0.0$, while the former has a median of $\rho_{+,k} = 0.08$. A Kolmogorov-Smirnov (KS) two-sample test provides quantitative evidence of the difference, returning a p -value of 10^{-14} .

The SDP process predicts a high forward cross-correlation for $\alpha \ll 1$, as discussed in Section 3.2. The data do not show evidence in favor of such a correlation in the majority of active regions. In the bottom panel of Figure 6 we do not see a clear visual trend between $\rho_{+,k}$ and λ_k (for regions with $N_k \geq 10$), although a Spearman




Figure 6. Top panel: PDF of forward cross-correlation, $\rho_{+,k}$ (black stepped curve), and backward cross-correlation, $\rho_{-,k}$ (black dashed curve), for all active regions with $N_k \geq 5$. Data binned into 30 uniformly spaced bins between $\rho_{\pm,k} = -1$ and $\rho_{\pm,k} = 1$. Bottom panel: Scatter plot of the forward cross-correlation, $\rho_{+,k}$, and flare rate, λ_k , for all regions with $N_k \geq 10$. The dashed grey line corresponds to $\rho_{+,k} = 0$.

correlation test returns a small but non-zero correlation of 5×10^{-2} (p -value of 0.06). We remind the reader that both $\rho_{+,k}$ and λ_k are empirical estimates for each active region, and the number of events per region is small. We do not plot uncertainties in the bottom panel of Figure 6, but they are typically large, of order the magnitude of the central estimate. These results imply that either i) all active regions have $\alpha \gtrsim 1$, where no forward cross-correlation is expected, i.e. $\lambda_k \gtrsim 0.2 \text{ hours}^{-1}$ does not correspond to the threshold-limited regime; or ii) some active regions have $\alpha \ll 1$, but either the trigger threshold or the driving rate is not constant with time, i.e. the random process triggering flares does not conform to the assumptions made in the SDP framework.

As described in Section 4.2 the *GOES* catalog is not complete, i.e. while a given active region is visible, not all flares that occur are recorded. Wheatland (2000a) noted that the obscuration in Section 4.2 creates an artificial forward cross-correlation, which explains why the median $\rho_{+,k}$ is positive.

5.3. Matching the shapes of $p(\Delta t)$ and $p(\Delta s)$

Another way to probe these data is to ask whether regions with relatively high $\rho_{+,k}$ exhibit “matching” functional forms for $p(\Delta t)$ and $p(\Delta s)$, as the SDP model




Figure 7. Calibrating the cross-correlation test in Section 5.3: two-dimensional KDEs of the relationship between ρ_+ and \mathcal{M} (top panel), and ρ_- and \mathcal{M} (bottom panel), for events simulated from the SDP model with four values of α (see legend for color code). The text in Section 5.3 reports details about the Monte Carlo procedure. Contours correspond to the 10%, 50%, and 90% credible intervals.

predicts for $\alpha \lesssim 1$. We quantify the degree to which the PDFs match via the p -value \mathcal{M}_k of the KS two-sample test applied to the sampled PDFs, $\{\Delta t_{i,k}\}$ and $\{\Delta s_{i,k}\}$. To calibrate, we perform a Monte Carlo simulation using sequences of events drawn from the SDP model. For each value of α we simulate 10^5 “regions”. For each region we generate M events, where M is a random number drawn from $p(N_k)$, restricted to values $N_k \geq 10$, as empirically measured for the *GOES* regions (the distribution for $N_k \geq 5$ shown in Figure 3). From these M events we calculate summary statistics, e.g. \mathcal{M} and ρ_+ . This results in 10^5 pairs of (\mathcal{M}, ρ_+) , from which we construct a two-dimensional kernel density estimate (KDE) (Wand & Jones 1995). The KDE estimates the true joint PDF $p(\mathcal{M}, \rho_+)$, and is qualitatively equivalent to a smoothed, two-dimensional histogram.

The result of this procedure for four values of α is shown in the top panel of Figure 7. For $\alpha = 0.01$ (dark orange lines) the 10% and 50% credible interval contours are invisible in the top-right corner of the panel, as the vast majority of simulated regions have $\rho_+ \approx 1$ and $\mathcal{M} \approx 1$. For $\alpha = 0.1$ the KDE spreads out slightly,




Figure 8. As for Figure 7 but for *GOES* data with black crosses marking all regions in the full, unmasked *GOES* catalog with $N_k \geq 10$. Red contours correspond to the 10%, 50%, and 90% credible intervals inferred from the black crosses.

with the 50% credible interval reaching $\rho_+ \approx 0.8$ and $\mathcal{M} \approx 0.9$. For $\alpha = 1$, the PDF shifts such that the 50% credible interval contour stretches from $\mathcal{M} = 1$ to $\mathcal{M} \approx 0.8$, while we have $0 \lesssim \rho_+ \lesssim 0.8$. For $\alpha = 10$ we find ρ_+ centered around zero, while the 50% credible interval of \mathcal{M} extends to $\mathcal{M} \approx 0.3$. The bottom panel of Figure 7 repeats the exercise for the backward cross-correlation. It confirms that ρ_- and \mathcal{M} are uncorrelated, as predicted elsewhere (Fulgenzi et al. 2017; Carlin & Melatos 2019a,b). Integrating over ρ_{\pm} the reader would find that the marginal distribution of \mathcal{M} is broad in both panels, for $\alpha \gtrsim 1$, as it is difficult to confidently reject the null hypothesis that two sets of samples are from the same distribution, when dealing with small sample sizes.

Monte Carlo calibration in hand, we present the equivalent data for all regions in the *GOES* catalog with $N_k \geq 10$ in Figure 8. To compute \mathcal{M}_k for each region we first normalize both the waiting times and the sizes by their respective means for that region, before applying the KS two-sample test. In the top (bottom) panel we see no clear relationship between $\rho_{+,k}$ ($\rho_{-,k}$) and \mathcal{M}_k . Marginalizing over \mathcal{M}_k we recover the PDFs of $\rho_{+,k}$ and $\rho_{-,k}$, shown in the top panel of Figure 6. If, for the sake of argument, we assume that all regions

have the same value of α , we can compare the KDEs in Figure 8 to those in Figure 7. Under this assumption we are pushed into the regime $1 \lesssim \alpha \lesssim 10$, as the $p(\rho_{+,k}, \mathcal{M}_k)$ KDE is slightly off-set from the horizontal axis (i.e. median $\rho_{+,k} > 0$), but the 50% credible interval for \mathcal{M}_k only extends to $\mathcal{M}_k \lesssim 0.6$. Even if α is not exactly the same in each region, we can interpret the above result as ruling out that a large proportion of regions have $\alpha \lesssim 1$, as if that were the case we would see higher values of \mathcal{M}_k associated with higher $\rho_{+,k}$ more often than in Figure 8.

5.4. Autocorrelations and longer term memory

While there is a wealth of information available in flare waiting time and size PDFs, as in Section 5.1, and cross-correlations, as in Sections 5.2 and 5.3, one may also consider statistics that quantify the longer term memory of the stress in the system. For example, as in the context of neutron star glitches (Carlin & Melatos 2019b), studying the autocorrelation between consecutive waiting times, $\rho_{\Delta t}$, or between consecutive sizes, $\rho_{\Delta s}$, allows one to place constraints on applicable SDP model parameters. In the fast-driven regime, $\alpha \ll 1$, we predict $\rho_{\Delta t} = 0$ and $\rho_{\Delta s} = 0$, as we have $X(t) \rightarrow X_{\text{cr}}$ before every stress-release event resulting in the system resetting at every event. If we have $\eta[\Delta X_i | X(t_i^-)] \propto \delta[X(t_i^-)]$, i.e. all stress in the system is released at each event, the same prediction of no autocorrelations holds — again the system is reset at every event. Observations of non-zero autocorrelations therefore rule out certain regimes in the SDP framework.

When we calculate $\rho_{\Delta t, k}$ and $\rho_{\Delta s, k}$ for all regions in the *GOES* catalog with $N_k \geq 10$, we find that the PDFs $p(\rho_{\Delta t, k})$ and $p(\rho_{\Delta s, k})$ are broad, akin to $p(\rho_{+,k})$ shown in Figure 6. The medians are 0.11 and 0.10 respectively. This result is incongruent with the assumption $\alpha \ll 1$ in all regions, as $\alpha \ll 1$ implies both PDFs should have a median of zero. These results likely stem from the obscuration effects described in Section 4.2. One tentative physical interpretation of a positive $\rho_{\Delta t}$ and $\rho_{\Delta s}$ is via analogy to terrestrial earthquakes which exhibit aftershocks, i.e. large events are often followed by larger-than-average events, with smaller-than-average waiting times (Utsu et al. 1995).

6. CONCLUSION

In this paper, we revisit the long-standing question of whether solar flares in different active regions are triggered by a stress-relax process with a trigger threshold which is constant in time. We do so by mapping the flaring process to a SDP process, which operates independently in each active region. The SDP framework

has been applied in related contexts to forest fires (Daly & Porporato 2006) and solar flares (Wheatland 2008), before being extended in the context of neutron star glitches (Fulgenzi et al. 2017; Melatos et al. 2018; Carlin & Melatos 2019a,b). If one assumes a constant driving rate, and a constant “stress” threshold at which relaxation events are guaranteed to occur, the model makes precise, falsifiable predictions regarding the statistics of sequences of waiting times and sizes, such as their PDFs, and cross-correlations. It is agnostic about the underlying microphysical mechanism.

Analyzing the historical *GOES* soft X-ray flare catalog, using data from 1975 until 2022, we systematically search all active regions for signatures that flares are consistent with the SDP model with $\alpha \ll 1$ (rapidly driven). We find no evidence that this is the case. Specifically, when considering just the waiting time and size PDFs of each active region we find that either i) the conditional PDF of stress-release sizes, $\eta[\Delta X_i | X(t_i^-)]$, varies from one region to another, or ii) a flare rate of $\lambda_k \gtrsim 0.2 \text{ hours}^{-1}$ does not correspond to $\alpha \ll 1$ (assuming that λ_k traces α). The analysis takes into account the selection effect that obscures the detection of flares following a large flare due to enhancement of the soft X-ray background.

On the other hand, if many active regions house a SDP process driven rapidly towards a static-in-time threshold before each event, i.e. have $\alpha \ll 1$, then those regions that have large cross-correlations between event sizes and subsequent waiting times should also have waiting time and size PDFs of the same shape. This prediction is not supported by the data, as proved clearly by Figure 8. The match between $p(\Delta t)$ and $p(\Delta s)$, quantified by \mathcal{M}_k (the p -value from a KS two-sample test) for the k -th active region, does not correlate with the forward cross-correlation, $\rho_{+,k}$. If we assume each region has the same α , this test implies $1 \lesssim \alpha \lesssim 10$.

We emphasize that the prediction that \mathcal{M}_k should correlate with $\rho_{+,k}$ for a process driven to a threshold does not rely on the specific details of the SDP framework, such as the relationship between flare rate and stress in Equation (2), nor the functional form of $\eta[\Delta X_i | X(t_i^-)]$. Any stochastic process that drives the stress towards a static threshold before each event would predict an equivalent observable, for example the Brownian stress-accumulation model (Carlin & Melatos 2020).

There are many ways to interpret the results in Sections 4.3 and 5: i) solar flares may not be triggered by stress (e.g. magnetic energy density) breaching a threshold; a completely different process may trigger a flare. ii) The *GOES* flare catalog is incomplete; flares that

occur in the aftermath of large flares are not recorded, creating an artificial cross-correlation. iii) The threshold at which a glitch is triggered and/or iv) the driving rate at which stress accumulates may vary with time; either of options iii) and iv) could wash out the observable signature. Untangling these explanations entails exploiting the entire wealth of data available in the flare catalog(s), rather than focussing on individual special flares or regions.

In closing, we touch on the following question: if option iii) above is true, is a threshold that varies with time compatible with any plausible microphysical flare triggers? The question falls outside the scope of this paper, which focuses on the microphysics-agnostic analysis in Sections 4 and 5, so we limit ourselves to the following brief remark. Magnetohydrodynamic instabilities relevant to solar flare activity are triggered above a threshold, which typically depends on the detailed geometry of a flaring magnetic loop, not just its bulk properties (e.g. magnetic energy density). The geometry of a flaring loop does vary with time in general. For example, the ideal kink instability occurs, when the total twist $\Phi = lB_\phi(r)/[rB_z(r)]$ satisfies $\Phi > \Phi_{cr}(a) \sim 10$, where l and r are the length and minor radius of the current-carrying loop, $B_\phi(r)$ and $B_z(r)$ are the toroidal

and axial magnetic field components, and a is the loop aspect ratio (Török et al. 2004). As sub-photospheric turbulence perturbs randomly the footpoints of a magnetic flux tube, the variables a and hence $\Phi_{cr}(a)$ fluctuate stochastically. Quantifying the amplitude (drift and diffusion) of the fluctuations is a task for detailed magnetohydrodynamic simulations and lies well outside the scope of this paper, but it is conceivable that the amplitude is sufficiently large to render option iii) above viable.

ACKNOWLEDGEMENTS

We acknowledge helpful discussions with Kai Yang in the early stages of this paper. JBC, AM, and MSW are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project number CE170100004) and ARC Discovery Project DP220102201. JBC is supported by an Australian Postgraduate Award.

Software: Numpy (Harris et al. 2020), Scipy (Virtanen et al. 2020), Matplotlib (Hunter 2007), and Stan (Stan Development Team 2022) through the cmdstanpy interface.

REFERENCES

- Akaike, H. 1974, IEEE Transactions on Automatic Control, 19, 716. <https://ui.adsabs.harvard.edu/abs/1974ITAC...19..716A/abstract>
- Aschwanden, M. J., & Johnson, J. R. 2021, The Astrophysical Journal, 921, 82, doi: 10.3847/1538-4357/ac2a29
- Aschwanden, M. J., Johnson, J. R., & Nurhan, Y. I. 2021, The Astrophysical Journal, 921, 166, doi: 10.3847/1538-4357/ac19a9
- Aschwanden, M. J., Scholkmann, F., Béthune, W., et al. 2018, Space Science Reviews, 214, 55, doi: 10.1007/s11214-018-0489-2
- Bak, P., Tang, C., & Wiesenfeld, K. 1987, Physical Review Letters, 59, 381, doi: 10.1103/PhysRevLett.59.381
- . 1988, Physical Review A, 38, 364, doi: 10.1103/PhysRevA.38.364
- Benz, A. O. 2016, Living Reviews in Solar Physics, 14, 2, doi: 10.1007/s41116-016-0004-3
- Betancourt, M. 2018, arXiv e-prints. <http://arxiv.org/abs/1701.02434>
- Biesecker, D. A. 1994, PhD thesis, University of New Hampshire. <https://ui.adsabs.harvard.edu/abs/1994PhDT.....51B>
- Boffetta, G., Carbone, V., Giuliani, P., Veltri, P., & Vulpiani, A. 1999, Physical Review Letters, 83, 4662, doi: 10.1103/PhysRevLett.83.4662
- Carlin, J. B., & Melatos, A. 2019a, Monthly Notices of the Royal Astronomical Society, 483, 4742, doi: 10.1093/mnras/sty3433
- . 2019b, Monthly Notices of the Royal Astronomical Society, 488, 4890, doi: 10.1093/mnras/stz2014
- . 2020, Monthly Notices of the Royal Astronomical Society, 494, 3383, doi: 10.1093/mnras/staa935
- . 2021, The Astrophysical Journal, 917, 1, doi: 10.3847/1538-4357/ac06a2
- Cox, D. R. 1955, Journal of the Royal Statistical Society, 17, 129, doi: 10.2307/2983950
- Crosby, N., Vilmer, N., Lund, N., & Sunyaev, R. 1998, Astronomy and Astrophysics, 334, 299. <https://ui.adsabs.harvard.edu/abs/1998A&A...334..299C/abstract>
- Daly, E., & Porporato, A. 2006, Physical Review E, 74, 1, doi: 10.1103/PhysRevE.74.041112
- Farhang, N., Wheatland, M. S., & Safari, H. 2019, The Astrophysical Journal, 883, L20, doi: 10.3847/2041-8213/ab40c3

- Fisher, G. H., Welsch, B. T., & Abbett, W. P. 2012, *Solar Physics*, 277, 153, doi: [10.1007/s11207-011-9816-4](https://doi.org/10.1007/s11207-011-9816-4)
- Fletcher, L., Dennis, B. R., Hudson, H. S., et al. 2011, *Space Science Reviews*, 159, 19, doi: [10.1007/s11214-010-9701-8](https://doi.org/10.1007/s11214-010-9701-8)
- Fulgenzi, W., Melatos, A., & Hughes, B. D. 2017, *Monthly Notices of the Royal Astronomical Society*, 470, 4307, doi: [10.1093/mnras/stx1353](https://doi.org/10.1093/mnras/stx1353)
- Gardiner, C. 2009, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th edn. (Berlin Heidelberg: Springer-Verlag).
<https://www.springer.com/gp/book/9783540707127>
- Gavriil, F. P., Kaspi, V. M., & Woods, P. M. 2004, *The Astrophysical Journal*, 607, 959, doi: [10.1086/383564](https://doi.org/10.1086/383564)
- Gelman, A., Carlin, J. B., Stern, H. S., et al. 2013, *Bayesian data analysis*, 3rd edn. (New York: Chapman and Hall/CRC)
- Gorobets, A., & Messerotti, M. 2012, *Solar Physics*, 281, 651, doi: [10.1007/s11207-012-0121-7](https://doi.org/10.1007/s11207-012-0121-7)
- Gourdji, K., Michilli, D., Spitler, L. G., et al. 2019, *The Astrophysical Journal Letters*, 877, L19, doi: [10.3847/2041-8213/ab1f8a](https://doi.org/10.3847/2041-8213/ab1f8a)
- Gögüş, E., Woods, P. M., Kouveliotou, C., et al. 1999, *The Astrophysical Journal*, 526, L93, doi: [10.1086/312380](https://doi.org/10.1086/312380)
- . 2000, *The Astrophysical Journal*, 532, L121, doi: [10.1086/312583](https://doi.org/10.1086/312583)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Haskell, B., & Melatos, A. 2015, *International Journal of Modern Physics D*, 24, 1530008, doi: [10.1142/S0218271815300086](https://doi.org/10.1142/S0218271815300086)
- Hudson, H. S. 2019, *Monthly Notices of the Royal Astronomical Society*, stz3121, doi: [10.1093/mnras/stz3121](https://doi.org/10.1093/mnras/stz3121)
- . 2020, *Solar Physics*, 295, 132, doi: [10.1007/s11207-020-01698-w](https://doi.org/10.1007/s11207-020-01698-w)
- Hudson, H. S., Labonte, B. J., Sterling, A. C., & Watanabe, T. 1998, in *Space Science Library*, Vol. 229, *Observational Plasma Astrophysics: Five Years of Yohkoh and Beyond* (Springer Dordrecht), 237.
<https://doi.org/10.1007/978-94-011-5220-4>
- Hunter, J. D. 2007, *Computing in Science and Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Hurvich, C. M., & Tsai, C.-L. 1989, *Biometrika*, 76, 297, doi: [10.1093/biomet/76.2.297](https://doi.org/10.1093/biomet/76.2.297)
- Jensen, H. J. 1998, *Self-Organized Criticality. Emergent Complex Behavior in Physical and Biological Systems*, *Cambridge Lecture Notes in Physics* (Cambridge: Cambridge University Press)
- Ji, H., & Daughton, W. 2011, *Physics of Plasmas*, 18, 111207, doi: [10.1063/1.3647505](https://doi.org/10.1063/1.3647505)
- Kanazir, M., & Wheatland, M. S. 2010, *Solar Physics*, 266, 301, doi: [10.1007/s11207-010-9623-3](https://doi.org/10.1007/s11207-010-9623-3)
- Keys, P. H., Mathioudakis, M., Jess, D. B., et al. 2011, *The Astrophysical Journal Letters*, 740, L40, doi: [10.1088/2041-8205/740/2/L40](https://doi.org/10.1088/2041-8205/740/2/L40)
- Kingman, J. F. C. 1993, *Poisson processes* (Oxford: Oxford University Press)
- La Roche-Carrier, N., Dituba Ngoma, G., Kocafee, Y., & Erchiqui, F. 2019, *International Journal of Quality & Reliability Management*, 37, 223, doi: [10.1108/IJQRM-01-2019-0035](https://doi.org/10.1108/IJQRM-01-2019-0035)
- Last, G., & Penrose, M. 2017, *Lectures on the Poisson Process* (Cambridge University Press), doi: [10.1017/9781316104477](https://doi.org/10.1017/9781316104477)
- Lehmann, E. L., & D'Abrera, H. J. M. 2006, *Nonparametrics: Statistical Methods Based on Ranks* (New York: Springer-Verlag)
- Lepreti, F., Carbone, V., & Veltri, P. 2001, *The Astrophysical Journal*, 555, L133, doi: [10.1086/323178](https://doi.org/10.1086/323178)
- Li, H., & Fenimore, E. E. 1996, *The Astrophysical Journal*, 469, L115, doi: [10.1086/310275](https://doi.org/10.1086/310275)
- Lippello, E., Arcangelis, L. d., & Godano, C. 2010, *Astronomy and Astrophysics*, 511, L2, doi: [10.1051/0004-6361/200913784](https://doi.org/10.1051/0004-6361/200913784)
- Lu, E. T. 1995a, *The Astrophysical Journal*, 446, L109, doi: [10.1086/187942](https://doi.org/10.1086/187942)
- . 1995b, *The Astrophysical Journal*, 447, 416, doi: [10.1086/175885](https://doi.org/10.1086/175885)
- Lu, E. T., & Hamilton, R. J. 1991, *The Astrophysical Journal Letters*, 380, L89, doi: [10.1086/186180](https://doi.org/10.1086/186180)
- Lu, E. T., Hamilton, R. J., McTiernan, J. M., & Bromund, K. R. 1993, *The Astrophysical Journal*, 412, 841, doi: [10.1086/172966](https://doi.org/10.1086/172966)
- Lyne, A., & Graham-Smith, F. 2012, *Pulsar Astronomy*, 4th edn., *Cambridge Astrophysics* (Cambridge: Cambridge University Press), doi: [10.1017/CBO9780511844584](https://doi.org/10.1017/CBO9780511844584)
- Melatos, A., & Drummond, L. V. 2019, *The Astrophysical Journal*, 885, 37, doi: [10.3847/1538-4357/ab44c3](https://doi.org/10.3847/1538-4357/ab44c3)
- Melatos, A., Howitt, G., & Fulgenzi, W. 2018, *The Astrophysical Journal*, 863, 196, doi: [10.3847/1538-4357/aad228](https://doi.org/10.3847/1538-4357/aad228)
- Millhouse, M., Melatos, A., Howitt, G., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 511, 3304, doi: [10.1093/mnras/stac194](https://doi.org/10.1093/mnras/stac194)

- Mitzenmacher, M. 2003, *Internet Mathematics*, 1, 226.
<https://projecteuclid.org/journals/internet-mathematics/volume-1/issue-2/A-Brief-History-of-Generative-Models-for-Power-Law-and-im/1089229510.full>
- Paxson, V., & Floyd, S. 1994, *ACM SIGCOMM Computer Communication Review*, 24, 257,
 doi: [10.1145/190809.190338](https://doi.org/10.1145/190809.190338)
- Peng, Z., & Zhao, P. 2009, *Nature Geoscience*, 2, 877,
 doi: [10.1038/ngeo697](https://doi.org/10.1038/ngeo697)
- Priest, E., & Forbes, T. 2002, *The Astronomy and Astrophysics Review*, 10, 313,
 doi: [10.1007/s001590100013](https://doi.org/10.1007/s001590100013)
- Rosner, R., & Vaiana, G. S. 1978, *The Astrophysical Journal*, 222, 1104, doi: [10.1086/156227](https://doi.org/10.1086/156227)
- Sahu, S., Joshi, B., Prasad, A., & Cho, K.-S. 2022, Evolution of magnetic fields and energy release processes during homologous eruptive flares, arXiv,
 doi: [10.48550/arXiv.2212.04150](https://doi.org/10.48550/arXiv.2212.04150)
- Shelyag, S., Fedun, V., Keenan, F. P., Erdélyi, R., & Mathioudakis, M. 2011a, *Annales Geophysicae*, 29, 883,
 doi: [10.5194/angeo-29-883-2011](https://doi.org/10.5194/angeo-29-883-2011)
- Shelyag, S., Keys, P., Mathioudakis, M., & Keenan, F. P. 2011b, *Astronomy and Astrophysics*, 526, A5,
 doi: [10.1051/0004-6361/201015645](https://doi.org/10.1051/0004-6361/201015645)
- Shibata, K., & Tanuma, S. 2001, *Earth, Planets and Space*, 53, 473, doi: [10.1186/BF03353258](https://doi.org/10.1186/BF03353258)
- Singhai, R., Joshi, S. D., & Bhatt, R. K. P. 2007, in 2007 15th International Conference on Software, Telecommunications and Computer Networks, 1–5,
 doi: [10.1109/SOFTCOM.2007.4446103](https://doi.org/10.1109/SOFTCOM.2007.4446103)
- Stan Development Team. 2022, *Stan Modeling Language Users Guide and Reference Manual*, 2.30.
<https://mc-stan.org>
- Sun, X., Hoeksema, J. T., Liu, Y., Kazachenko, M., & Chen, R. 2017, *The Astrophysical Journal*, 839, 67,
 doi: [10.3847/1538-4357/aa69c1](https://doi.org/10.3847/1538-4357/aa69c1)
- Sun, X., Hoeksema, J. T., Liu, Y., et al. 2012, *The Astrophysical Journal*, 748, 77,
 doi: [10.1088/0004-637X/748/2/77](https://doi.org/10.1088/0004-637X/748/2/77)
- Török, T., Kliem, B., & Titov, V. S. 2004, *Astronomy and Astrophysics*, 413, L27, doi: [10.1051/0004-6361:20031691](https://doi.org/10.1051/0004-6361:20031691)
- Utsu, T., Ogata, Y., S, R., & Matsu'ura. 1995, *Journal of Physics of the Earth*, 43, 1, doi: [10.4294/jpe1952.43.1](https://doi.org/10.4294/jpe1952.43.1)
- van der Linden, W. J. 2006, *Journal of Educational and Behavioral Statistics*, 31, 181.
<https://www.jstor.org/stable/3701364>
- Verbeeck, C., Kraaikamp, E., Ryan, D. F., & Podladchikova, O. 2019, *The Astrophysical Journal*, 884, 50, doi: [10.3847/1538-4357/ab3425](https://doi.org/10.3847/1538-4357/ab3425)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Wand, M. P., & Jones, M. C. 1995, *Kernel Smoothing* (London; New York: Chapman & Hall)
- Wheatland, M. S. 2000a, *Solar Physics*, 191, 381,
 doi: [10.1023/A:1005240712931](https://doi.org/10.1023/A:1005240712931)
- . 2000b, *The Astrophysical Journal*, 536, L109,
 doi: [10.1086/312739](https://doi.org/10.1086/312739)
- . 2001, *Solar Physics*, 203, 87,
 doi: [10.1023/A:1012749706764](https://doi.org/10.1023/A:1012749706764)
- . 2008, *The Astrophysical Journal*, 679, 1621,
 doi: [10.1086/587871](https://doi.org/10.1086/587871)
- Wheatland, M. S., & Glukhov, S. 1998, *The Astrophysical Journal*, 494, 858, doi: [10.1086/305245](https://doi.org/10.1086/305245)

APPENDIX

A. MAGNETIC ENERGY DENSITY AS A STRESS VARIABLE IN AN ACTIVE REGION

The SDP framework described in Sections 2 and 3 is agnostic about the microphysics of solar flares, beyond the general assumption that the flaring rate increases with the stress in the system and diverges at a stress threshold. In this appendix, we sketch briefly one possible mapping between the SDP model and a specific microphysical flare model, in which the stress variable is the spatially-averaged magnetic energy density in an active region. We emphasize that the model is phenomenological and highly idealized. We do not favor it over the many alternatives; it is merely one possible illustration of how such a mapping may work, as a guide to the interested reader.

Let $X(t)$ correspond to the spatially-averaged magnetic energy density in an active region. The spatial average is taken in order to package the stress into a single variable, noting that the magnetic energy density is nonuniform in reality. Let the active region have characteristic linear dimension L . Let S be the magnetohydrodynamic Poynting flux through the photosphere. We assume that the Poynting flux deposits magnetic energy into the active region without losses and at a constant rate, so that one has $X(t) = X(t_i^+) + St/L$ between flares ($t_i^+ \leq t \leq t_{i+1}^-$). Recent vector magnetogram measurements allow direct observation of the Poynting flux, and the magnetic energy density of an active region, which are at-odds with our assumption that S is steady in time (Fisher et al. 2012; Sun et al. 2012, 2017; Sahu et al. 2022). Putting aside the latter consideration for the moment, we can write the control parameter α as

$$\alpha = \frac{\lambda_0 S}{X_{\text{cr}} L}, \quad (\text{A1})$$

where X_{cr} is the magnetic energy density threshold for a magnetohydrodynamic instability, for example, and λ_0 is the instability's trigger reference rate.

Several plausible magnetohydrodynamic instabilities have been suggested as solar flare triggers in the literature (Ji & Daughton 2011). Some of them do not involve a magnetic energy density threshold at all. For example, the kink instability discussed at the end of Section 6 is triggered when the field-line twist Φ exceeds a threshold $\Phi_{\text{cr}}(a)$, whose value depends on the aspect ratio a of the flaring loop (Török et al. 2004). However, instabilities triggered by a magnetic energy density threshold do exist. One example is plasmoid-induced magnetic reconnection via tearing modes in a fractal current sheet (Shibata & Tanuma 2001), whose threshold depends on fractional powers of the Alfvén speed (or equivalently the Lundquist number) and hence on the magnetic energy density; see Ji & Daughton (2011) or section 5 in Shibata & Tanuma (2001) for example. In the latter reference, the threshold condition also depends on the aspect ratio a (width divided by length) of the current sheet, which can vary with time, as a flaring loop responds to sub-photospheric turbulence.

What are the time-scales on which S and X_{cr} vary? In the SDP picture, both variables are steady, but the *GOES* analysis in Section 5 implies that one or both may vary in reality (although other scenarios are possible too, as discussed in Section 6). As far as S is concerned, one expects to find statistical fluctuations on the eddy turnover time-scale of sub-photospheric turbulence and flux emergence, τ_{flux} , as observed with vector magnetograms and Doppler measurements (Fisher et al. 2012). Magnetohydrodynamic simulations and G-band radiative signatures suggest $\tau_{\text{flux}} \sim$ minutes for magnetic features associated with solar granulation, with peak-to-peak fluctuation amplitude $\lesssim 30\%$ (Keys et al. 2011; Shelyag et al. 2011b,a). As far as X_{cr} is concerned, in the plasmoid-induced reconnection picture as one illustrative example, the magnetic energy density threshold for a Sweet-Parker current sheet to undergo secondary tearing depends not only on the magnetic diffusivity, which fluctuates on the same time-scale as the local temperature, but also on a and L , which fluctuate on the turnover time-scale τ_{flux} , e.g. equation (15) in (Shibata & Tanuma 2001). Flare waiting times are typically comparable to or longer than τ_{flux} , so it is conceivable that the constant- X_{cr} approximation in the SDP theory does not apply to every active region.

We emphasize again that the mapping in this appendix is idealized and illustrative. None of the statistical analysis in Sections 4.3 and 5 is predicated on the microphysics in this appendix.

B. HIERARCHICAL BAYESIAN FRAMEWORK

This appendix lays out a complementary approach to the heuristic tests with which we analyze individual active regions in Section 5. We first write down the likelihood for a set of observed waiting times and sizes in an individual region, given a set of model parameters, in Appendix B.1. We introduce the hierarchical Bayesian framework which combines inference in different regions to estimate population-level parameters, in Appendix B.2. In Appendix B.3 we test the efficacy of this approach with synthetic data, and explain why it is not appropriate in this paper to apply this framework to the real *GOES* catalog despite its efficacy. Finally, in Appendix B.4 we propose an alternative approach using only the average waiting time in each region, as a motivation for future studies. The recipes in this appendix are included for completeness and as a starting point for readers, who wish to develop hierarchical methods of solar flare analysis further.

B.1. Likelihood and Bayes' theorem

Solving Equations (1)–(3) for the long-term observable PDFs, $p(\Delta t)$ and $p(\Delta X)$, is intractable analytically for most choices of $\eta [\Delta X_i | X(t_i^-)]$. However, for the special case

$$\eta [\Delta X_i | X(t_i^-)] \propto [X(t_i^-) - \Delta X_i]^\delta, \quad (\text{B2})$$

we obtain the analytic result

$$p(z) = (\alpha + \delta + 1)(1 - z)^{\alpha + \delta}, \quad (\text{B3})$$

where the variable z is either Δt or ΔX , i.e. the waiting time and size PDFs are identical (Fulgenzi et al. 2017). In Equation (B2) we have $\delta > 0$ and the constant of proportionality is set by the condition that η must integrate to unity between $\Delta X_i = 0$ and $\Delta X_i = X(t_i^-)$; see section 6 and appendix D of Fulgenzi et al. (2017) for a full derivation. Equation (B2) is a monotonically decreasing function of ΔX_i , i.e. small stress-release events are preferred over large events. Its specific functional form is reasonable but arbitrary; it is not inferred from solar flare data. We adopt it here as a pedagogical device to illustrate the advantages and disadvantages of a hierarchical Bayesian approach to analyzing flare data.

When we restore the dimensions to Equation (B3), and consider the set of observations in the k -th active region, $D_k = \{\Delta t_{i,k}, \Delta s_{i,k}\}$, with $1 \leq i \leq N_k$, we can write the likelihood

$$\mathcal{L}(D_k | \boldsymbol{\theta}_k) = \prod_{i=1}^{N_k-1} p(\Delta t_{i,k} | \boldsymbol{\theta}_k) \prod_{i=1}^{N_k} p(\Delta s_{i,k} | \boldsymbol{\theta}_k) \quad (\text{B4})$$

$$= \tau_k^{1-N_k} \xi_k^{-N_k} (\beta_k + 1)^{2N_k-1} \prod_{i=1}^{N_k-1} \left(1 - \frac{\Delta t_{i,k}}{\tau_k}\right)^{\beta_k} \prod_{i=1}^{N_k} \left(1 - \frac{\Delta s_{i,k}}{\xi_k}\right)^{\beta_k} \quad (\text{B5})$$

where ξ_k is the constant of proportionality in Equation (5) which translates the drop in magnetic energy density ΔX to the observed flare size, Δs , and one has $\beta_k = \lambda_{0,k} \tau_k + \delta_k$. While there are N_k flare sizes in the region, there are only $N_k - 1$ waiting times. The three model parameters $\boldsymbol{\theta}_k = \{\tau_k, \xi_k, \beta_k\}$ are assumed constant in time for the lifetime of the active region. Bayes' theorem calculates the posterior probability distribution (henceforth “posterior”), $p(\boldsymbol{\theta}_k | D_k)$, i.e. the probability density of parameter vector $\boldsymbol{\theta}_k$ given the data D_k , using the prior probability distribution (henceforth “prior”), $\pi(\boldsymbol{\theta}_k)$, viz.

$$p(\boldsymbol{\theta}_k | D_k) \propto \mathcal{L}(D_k | \boldsymbol{\theta}_k) \pi(\boldsymbol{\theta}_k). \quad (\text{B6})$$

The normalizing constant of proportionality is often called the evidence and, while essential for model comparison, is not relevant for the parameter estimation exercise below (Gelman et al. 2013).

B.2. Population-level parameter estimation

Suppose, for the sake of illustration, that X_{cr} , λ_0 , and every component of $\boldsymbol{\theta}_k$ are approximately the same in all active regions. In the language of hierarchical Bayesian inference, this corresponds to assuming that the value of $\boldsymbol{\theta}_k$ for each region is a random number drawn from a population-level distribution known as a “hyper-prior”, with narrow extent. For concreteness, we assume the hyper-prior for each model parameter is a Gaussian with mean and standard deviation μ_a and σ_a respectively, with $a \in \{\tau, \xi, \beta\}$. While β_k does depend on τ_k , the inference remains accurate so




Figure 9. Posterior predictive check demonstrating the recovery of injected population-level parameters. The weakly informative priors set on the hyper-parameters τ , ξ , and β are shown in green, histograms of samples from the injected population parameters are shown in red, while histograms of samples from the posterior distributions are shown in blue. See text in Appendix B.3 for details on the procedure.

long as the covariance between τ_k , $\lambda_{0,k}$, and δ_k is minimal. The marginal posterior distribution for the parameters describing the hyper-priors is calculated as

$$p(\mathbf{\Lambda} | \mathcal{D}) \propto \pi(\mathbf{\Lambda}) \prod_k^M \int d\boldsymbol{\theta}_k \mathcal{L}(D_k | \boldsymbol{\theta}_k) \pi(\boldsymbol{\theta}_k | \mathbf{\Lambda}), \quad (\text{B7})$$

with $\mathbf{\Lambda} = \{\mu_a, \sigma_a\}$. In Equation (B7), $\pi(\mathbf{\Lambda})$ is the prior on $\mathbf{\Lambda}$, and $\mathcal{D} = \{D_1, \dots, D_k, \dots, D_M\}$ is the set of all data, D_k , from M different active regions.

In practice, we can use a Monte Carlo sampler to estimate Equation (B7), and thus infer the posterior predictive distributions, $p(\boldsymbol{\theta})$, where we drop the subscript k to signify that these distributions bring together information from all active regions to inform the inference. We opt to sample using a Hamiltonian Monte Carlo No U-Turn Sampler (Betancourt 2018), as implemented in the Stan programming language (Stan Development Team 2022).

B.3. Validation with synthetic data

To estimate the efficacy of the scheme in Appendix B.2 to infer population-level parameters, we first apply it to a set of synthetic data generated directly from the model. That is, we generate $M = 50$ fake “active regions”, each with $N_k = 100$ events, with flare sizes and waiting times generated from the SDP framework with $\eta [\Delta X_i | X(t_i^-)]$ as in Equation (B2). In what follows we use the compact notation $x \sim \mathcal{N}(\mu, \sigma)$ to denote that the random variable x is drawn from a Gaussian distribution with mean μ and standard deviation σ . To generate the data for each region we select a value $\tau_k \sim \mathcal{N}(\mu_\tau = 3 \text{ day}, \sigma_\tau = 0.5 \text{ day})$, a value $\xi_k \sim \mathcal{N}(\mu_\xi = 2 \text{ arb. units}, \sigma_\xi = 0.3 \text{ arb. units})$, and a value $\beta_k \sim \mathcal{N}(\mu_\beta = 5, \sigma_\beta = 0.3)$. Each Gaussian is truncated such that all variates are positive.

After running the sampler with the synthetic data we perform a posterior predictive check, i.e. compare the samples from our posterior distributions $p(\tau)$, $p(\xi)$, and $p(\beta)$ with the injected distributions and the priors on the hyper-parameters. We show the results in Figure 9. We see that the posterior samples (blue) appropriately overlap with the injected ground truth (red) for τ , ξ , and β .

Despite the encouraging results in Figure 9, we do not apply this hierarchical Bayesian scheme to the *GOES* catalog in this paper. This is because in most regions neither the waiting time nor the size PDFs follow the functional form in Equation (B3). When we include Equation (B3) in the set of options available for the AICc to select between, as in Section 5.1, fewer than 0.1% of active regions are fitted best with Equation (B3). Therefore Equation (B5) is not an appropriate likelihood for the data. Other choices of the jump distribution in Equation (B2) may alleviate this problem but they lie outside the scope of this paper.

B.4. Likelihood based on the average waiting time

An alternative approach is to build the likelihood out of a summary statistic, such as the average waiting time in a given region, $\langle \Delta t \rangle_k$, rather than each observed waiting time and/or size. In the SDP framework, we find (Fulgenzi et al. 2017; Millhouse et al. 2022)

$$\langle \Delta t \rangle^{\text{SDP}} = \frac{1}{\alpha + \alpha_0}, \quad (\text{B8})$$

where $\alpha_0 \sim 1$ is a dimensionless constant, whose exact value depends on the particular form of $\eta[\Delta X_i | X(t_i^-)]$. Restoring the dimensions to Equation (B8) we can write the likelihood as

$$\mathcal{L}(\langle \Delta t \rangle_k | \boldsymbol{\theta}_k) = \frac{1}{\sqrt{2\pi}\sigma_k} \exp \left[-\frac{\left(\langle \Delta t \rangle_k - \frac{\tau_k}{\lambda_{0,k}\tau_k + \alpha_{0,k}} \right)^2}{2\sigma_k^2} \right], \quad (\text{B9})$$

with $\boldsymbol{\theta}_k = \{\tau_k, \lambda_{0,k}, \alpha_{0,k}, \sigma_k\}$, where we assume the residual of $\langle \Delta t \rangle_k$ with the model subtracted is normally distributed, with zero mean and standard deviation σ_k . With the likelihood in Equation (B9) one could perform population-level parameter estimation, as described in Section B.2. We leave this, and an exploration of how $\boldsymbol{\theta}_k$ depends on parameters intrinsic to each active region, to future work.

C. CLEANING *GOES* FLARE SUMMARY DATA

The SWPC and NGDC File Transfer Protocol servers that host the *GOES* flare summary data, as well as other data portals such as the Heliophysics Event Knowledgebase⁵, contain numerous typographical anomalies found when collating the data into a homogeneous catalog. These anomalies are found in the recorded active region number by checking whether one has $\Delta t_{i,k} < 14$ days for $1 \leq i \leq N_k - 1$ for each active region. The limit of 14 days applies because active regions are typically only visible for two weeks due to the rotation period of the Sun. The one exception in the *GOES* flare summary data is eight flares that occurred on 2002-11-02, and are assigned to active region number 10198. Two weeks later the active region appears on the eastern limb of the Sun, and 33 additional flares are assigned to active region number 10198, from 2002-11-17 until 2002-11-28. We opt not to consider the former eight flares as part of active region number 10198, and do not include them in the catalog.

The anomalies, and hand-corrected values, are tabulated in Table 1. The corrected values are determined manually by referring to the context of surrounding flares in the database. For example, the flare starting at 1981-07-18 11:46 is recorded with active region number 3121. Yet the other 11 flares associated with the latter region occurred between 1981-05-24 and 1981-06-01, while active region number 3221 has 32 flares recorded between 1981-07-17 and 1981-07-29, indicating that the anomalous flare should be associated with the latter active region. The latitude of the anomalous flare is also within $\pm 5^\circ$ of other flares from active region number 3221, while flares in active region number 3121 are $\gtrsim 10^\circ$ higher in latitude.

D. LOG-NORMAL WAITING TIME DISTRIBUTIONS

A sequence of instantaneous events ordered in time is known as a point process. If the instantaneous event rate does not change with time, and if the process is memoryless, it is called a Poisson point process, and exhibits exponentially distributed waiting times (Kingman 1993). If the instantaneous rate does change with time it is a non-homogeneous Poisson process, which can exhibit waiting times that are distributed as a log-normal for certain rate functions (Gardiner 2009; Last & Penrose 2017). The SDP is one example of a non-homogeneous Poisson process that can generate waiting times distributed as a log-normal, for certain choices of $\eta[\Delta X_i | X(t_i^-)]$ and α . One can alternatively generate log-normally distributed values via a multiplicative process, e.g. an organism that grows in proportion to its current size multiplied by a random variable generates (at least approximately) log-normally distributed sizes, due to the Central Limit Theorem (Mitzenmacher 2003).

Log-normal waiting time distributions are observed in many other contexts besides solar flares, including but not limited to gamma-ray bursts (Li & Fenimore 1996), X-ray bursts (Gögüş et al. 1999, 2000; Gavriil et al. 2004), fast radio bursts (Gourdji et al. 2019), network traffic (Paxson & Floyd 1994; Singhai et al. 2007), mining equipment failure (La Roche-Carrier et al. 2019), earthquake aftershocks (Peng & Zhao 2009), and test response times (van der Linden 2006).

⁵ Accessible through http://hec.helio-vo.eu/hec/hec_gui.php.

Table 1. Anomalous active region numbers found in the *GOES* flare summary data by checking whether one has $\Delta t_{i,k} < 14$ days for $1 \leq i \leq N_k - 1$ for each active region. Corrected values are determined manually, by considering the context of what active regions are present at the time of the anomalous flare at a similar latitude. Corrected values of “—” indicate that we cannot identify a reasonable active region to associate with the anomalous flare.

Flare start time	Anomalous active region number	Corrected active region number
1978-05-30 06:19	1000	1134
1981-07-18 11:46	3121	3221
1981-08-22 06:58	366	3266
1983-03-01 18:24	2102	4102
1983-03-01 18:54	2102	4102
1983-07-04 06:09	4135	4235
1983-07-29 03:53	4236	4263
1993-09-27 01:35	7500	7590
2000-11-09 21:13	9125	—
2002-06-14 20:18	1	10001
2003-07-31 07:59	422	10422
2003-12-21 04:10	10000	—
2011-12-22 13:04	11281	11381
2017-07-11 01:09	12655	12665
2017-07-16 10:25	12655	12665
2021-05-10 23:46	12282	12822
2021-09-01 03:03	12680	12860
2021-09-01 04:27	12680	12860
2021-11-06 22:01	12984	12894
2021-12-18 11:17	12807	12907
2021-12-18 17:27	12807	12907

E. EXAMPLE FITS FOR INDIVIDUAL ACTIVE REGIONS

The reader may be curious to examine some examples of the variety of distributions fitted to individual active regions. To this end, in Figure 10 we show the complementary cumulative distribution function (CCDF) for the flare waiting times and sizes from four arbitrary but representative active regions with $N_k \geq 20$. We display the distributions as CCDFs instead of PDFs to avoid binning the data. The black stepped curves indicate the empirical CCDF constructed from the *GOES* catalog, while the colored curves are the various fits to the data, with parameters fixed to their maximum likelihood values. The distribution with bold type-face in the legend is the shape that best fits the data according to the AICc.

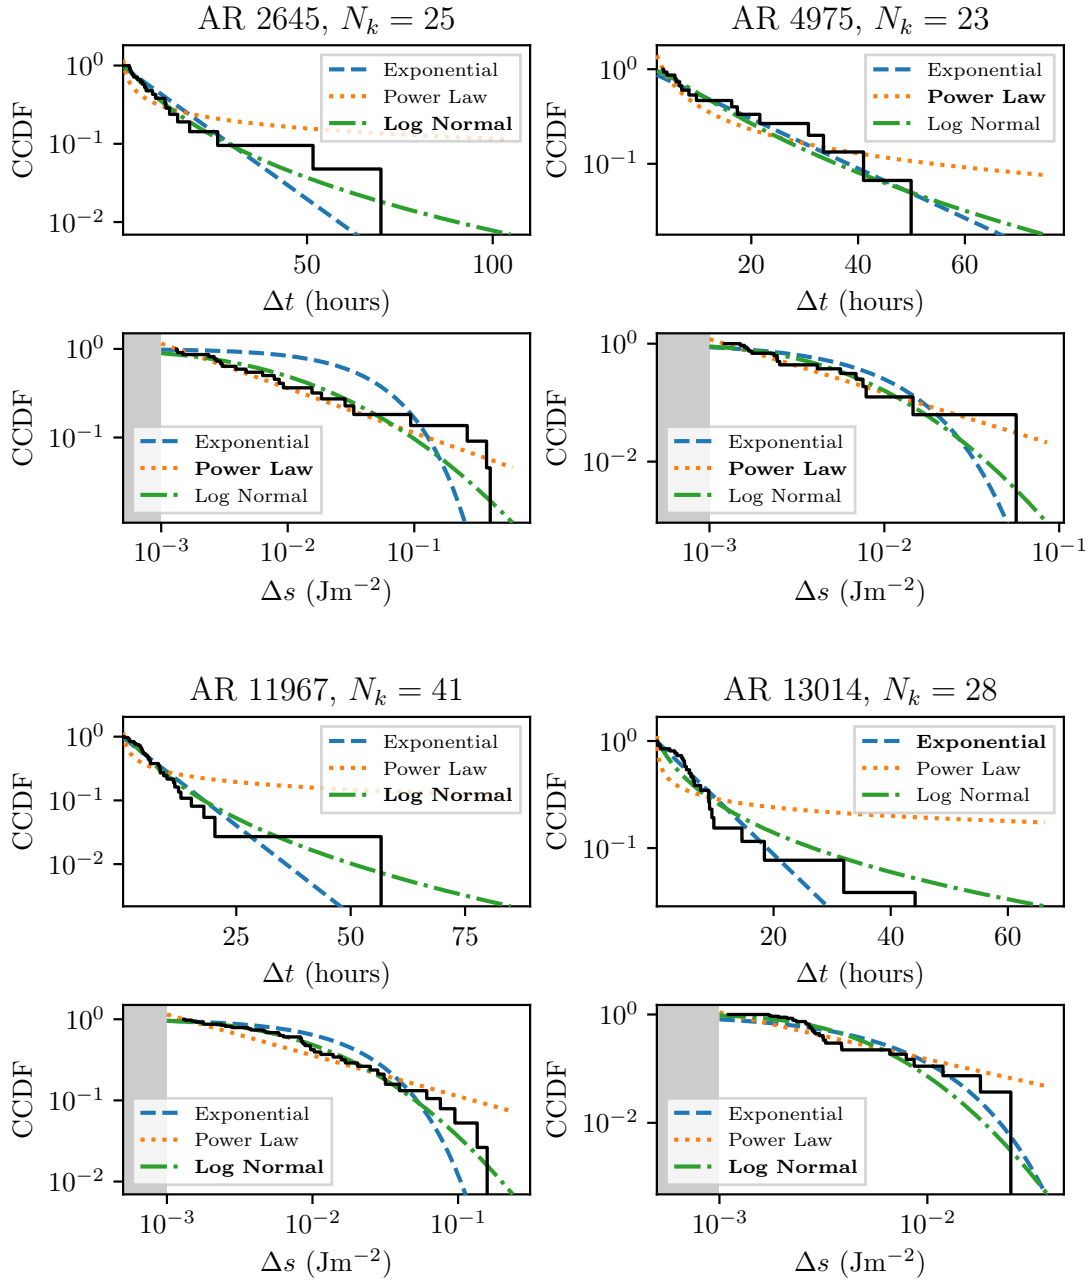


Figure 10. Complementary cumulative distribution functions (CCDFs) for four arbitrary but representative active regions with $N_k \geq 20$. The distribution that best describes the data in each panel (black stepped curve), according to the AICc, is in bold type-face in the legend for each panel. The grey regions in the bottom panels indicate $\Delta s \leq 10^{-3} \text{Jm}^{-2}$, i.e. where the masking described in Section 4.2 is applied.