On the Statistical Significance of Foreshock Sequences in Southern California

M. P. A. van den Ende[1] and **J.-P. Ampuero[1,2]**

[1]Université Côte d'Azur, IRD, CNRS, Observatoire de la Côte d'Azur, Géoazur, Valbonne, France, [2]Seismological Laboratory, California Institute of Technology, Pasadena, CA, USA

Abstract Earthquake foreshocks may provide information that is critical to short-term earthquake forecasting. However, foreshocks are far from ubiquitously observed, which makes the interpretation of ongoing seismic sequences problematic. Based on a statistical analysis, Trugman and Ross (2019, https://doi.org/10.1029/2019GL083725) suggested that as much as 72% of all mainshocks in Southern California is preceded by foreshock sequences. In this study, we reassess the analysis of Trugman and Ross (2019, https://doi.org/10.1029/2019GL083725), and we evaluate the impact of the assumptions made by these authors. Using an alternative statistical approach, we find that only 15 out of 46 mainshocks (33%) are preceded by significantly elevated seismicity rates. When accounting for temporal fluctuations in the background seismicity, only 18% of the analyzed foreshock sequences remain unexplained by the background seismicity. These results imply that even in a highly complete earthquake catalog, the majority of earthquakes do not exhibit detectable foreshock sequences.

1. Introduction

Prior to large earthquakes, precursory signals, such as accelerating aseismic creep (see Roeloffs, 2006) and elevated tremor and foreshock activity (Dodge et al., 1996; Jones & Molnar, 1979; Marsan & Enescu, 2012), may be recorded by GPS and seismic stations. However, the majority of observations of transient creep and elevated seismicity have only been identified as precursory signals in hindsight. Without prior knowledge of the timing and location of the mainshock, unambiguous identification of earthquake precursors is enigmatic. While it has been argued that some, if not most, large earthquakes are preceded by detectable foreshock sequences (Bouchon et al., 2013), (locally) elevated seismicity rates do not uniquely signify the advent of an earthquake, as many of these excursions simply return to the (average) background seismicity rate. These observations prohibit the utilization of seismicity rates in earthquake forecasting attempts to date.

Recently, Trugmann and Ross (2019, T&R from here on) performed an in-depth statistical analysis of the seismicity rates preceding 46 mainshock events in Southern California employing a highly complete earthquake catalog (the QTM catalog; Ross et al., 2019). By comparing short-term seismicity rates with the background rate over 1 year prior to a selected mainshock, T&R concluded that over 70% of all analyzed mainshocks were preceded by a statistically significant increase in seismicity rates. These authors further alluded to the possibility that in practically all cases foreshock sequences may be detected, provided that the earthquake catalog is sufficiently complete. If this claim holds true, this has the implication that the nucleation process of (large) earthquakes emits detectable information with potential application in short-term earthquake forecasting. However, the reported results do not immediately illuminate the significance of elevated seismicity rates when taking into perspective the ubiquitous fluctuations in the background seismicity rate: if similar elevations in the seismicity rate are observed at random 70% of the time, the presence of a foreshock sequence may simply be due to random chance.

In this work, we revisit the analysis of T&R to assess the significance of their findings, additionally taking into account fluctuations in the background seismicity rate. We first describe the procedure to reproduce the results reported by these authors and comment on some of the assumptions made in this procedure. We subsequently propose alternative approaches that relax these assumptions, and we reassess the statistical significance and interpretation of elevated seismicity rates prior to large earthquakes. We find that the number of mainshocks preceded by statistically significant foreshock activity (taken in a broad sense) is

substantially less than that reported by T&R (20–30% versus 70%, respectively). Only about half of these foreshock sequences can be explained by random fluctuations in the background seismicity rate, which suggests that in some cases elevated seismicity rates are uniquely associated with periods preceding mainshocks.

2. Statistical Methods

2.1. Reproducing the Results of Trugman and Ross (2019)

To set a reference, we begin by reproducing the results reported by Trugman and Ross (2019). Although the approach taken by these authors has been described extensively in their main manuscript and in their supporting information, certain details of the procedure were omitted for brevity. Below we briefly describe the procedure including these details, based on personal communication with D. Trugman.

First, we extract the 46 mainshock events as identified by T&R from the QTM catalog (Ross et al., 2019). For each of these events, we collect all earthquakes that are located within a rectangular box extending ± 10 km around the mainshock with no depth cutoff. We then compute the interevent time (IET) between each pair of subsequent earthquakes that occurred within 380 days and up to 20 days prior to each mainshock. A time window of 20 days prior to the main event is excluded to avoid bias of the estimation of the background rate by potential foreshock activity. The distribution of IET tends to follow a gamma distribution, of which the probability density function ρ is defined as

$$\rho(\tau; \gamma, \mu) = \frac{\tau^{\gamma-1} e^{-\mu\tau}}{\mu^{-\gamma} \Gamma(\gamma)} \tag{1}$$

where τ is the interevent time, γ the shape parameter, μ the background rate, and $\Gamma(\gamma)$ the gamma function. The ratio of independent events over the total number of events (including clustered events) is represented by the shape parameter γ. A shape parameter of $\gamma = 1$ indicates that all events in the catalog are mutually independent, and accordingly equation (1) reduces to the exponential distribution (associated with a Poisson process). When fitting equation (1) to the earthquake catalog, the maximum likelihood estimate of the background rate μ is simply $\gamma / \langle \tau \rangle$, where $\langle \tau \rangle$ is the mean IET. The corresponding maximum likelihood estimate of γ results from maximizing the log-likelihood function:

$$\mathcal{L} = (\gamma - 1) \sum_{i=1}^{N} \ln(\tau_i) - N \left[\gamma + \gamma \ln \left(\frac{\langle \tau \rangle}{\gamma} \right) + \ln(\Gamma(\gamma)) \right] \tag{2}$$

Using these maximum likelihood estimates, T&R fitted the gamma distribution to the IETs of events within $-380d \leq t < -20d$ prior to each individual mainshock, yielding γ and μ that characterize the background seismicity rate associated with each mainshock event.

Next, T&R assumed a Poisson distribution of observing N_{obs} earthquakes in a given 20-day time interval, of which the probability mass function is defined as

$$P(N_{\text{obs}}) = \frac{1}{N_{\text{obs}}!} \lambda^{N_{\text{obs}}} e^{-\lambda} \tag{3}$$

where $\lambda = \mu T$, with T being the observation time window of 20 days. The probability of observing at least N_{obs} events over the 20 days preceding the mainshock is given by the survival function:

$$p(N \geq N_{\text{obs}}) = 1 - \sum_{n=0}^{N_{\text{obs}}-1} \frac{\lambda^n}{n!} e^{-\lambda} \tag{4}$$

Low values of p imply that the number of events in the 20-day window is unlikely to be observed given the background rate μ, and hence are an indication of anomalously high seismicity rates. Following T&R, we adopt a value of $p < 0.01$ as a threshold for the statistical significance of the seismicity rates.

2.2. Alternative Approaches Based on Random Sampling

A key assumption in the above procedure of estimating the expected 20-day event count (i.e., equation (4)), is that the seismicity rate follows a Poisson distribution. This implicitly requires that each event in the analyzed sequence be statistically independent (and accordingly $\gamma = 1$). While this requirement may be met for individual mainshock events, the QTM catalog contains numerous occasions of (short) aftershock sequences associated with events of magnitude $M_w < 4$ (which are therefore not considered as mainshocks). To esti-

mate the background seismicity rate based on the interevent times does not require special treatment of correlated events (Hainzl et al., 2006), and so the estimates of μ obtained by fitting a gamma distribution are representative (for any $\gamma \leq 1$). On the other hand, to subsequently assess the significance of observing N_{obs} events in a given time-window based on a Poisson survival function, declustering of the earthquake catalog is essential (e.g., Reasenberg, 1985). To circumvent this, we propose two alternative methods below that do not assume Poissonian behavior.

First, after fitting the gamma distribution to the IET data (see previous section), we draw N random samples of IET from the resulting probability density function. The total duration Δt of a sequence of N earthquakes is therefore

$$\Delta t = \sum_{}^{N} \hat{\tau} \tag{5}$$

where $\hat{\tau}$ is a random sample drawn from $\rho(\tau; \gamma, \mu)$ (equation (1)). The number of events that are observed within a 20-day time window T is thus defined as the largest value of N for which $\Delta t \leq T$. Note that each sequence begins with an earthquake at $t = 0$, but that this does not affect the statistics of N. By generating 50,000 realizations of Δt, we obtain a distribution of N, that is, the distribution of earthquakes occurring in a time window of 20 days based on the measured background seismicity rate μ and shape factor γ. Since Δt results from the summation of random samples drawn from a gamma distribution, we make use of the following property (for a proof, see Appendix A):

$$\sum_{}^{N} \hat{\tau} \sim \rho(\Delta t; N\gamma, \mu) \tag{6}$$

that is, the summation of samples drawn from a gamma distribution itself follows a gamma distribution of which the shape factor γ is multiplied by N. By introducing the criterion that $\Delta t \approx T = \text{const.}$, it is to be expected that N also follows a gamma distribution (which we will demonstrate numerically in the next section). Under the assumption that N is gamma distributed, we fit a gamma distribution to the random realizations of N and calculate the probability of observing at least N_{obs} events in a 20-day window from the corresponding survival function. A probability lower than a threshold of 0.01 signifies an elevated seismicity rate that is not expected based on the background seismicity rate.

Second, while a gamma distribution may be a significant improvement over a Poisson distribution in capturing the IET statistics, it may not be fully adequate in accounting for temporal clustering. Ideally, the true distribution underlying the observed number of earthquakes is sampled to assess the significance of a given earthquake sequence comprising N_{obs} events. To this end, we create an empirical distribution of N by counting the number of events observed during a random 20-day period within $-380d \leq t < -20d$ prior to each mainshock event. We draw 50,000 random samples uniformly distributed over the year leading up to the mainshock, which unavoidably oversamples cases in which the background seismicity rate is low (mainshock #14601172 with a total of 18 events being an extreme example). For any empirical distribution, the discrete survival function can be obtained based on the notion that, in an ordered set $\mathbf{N}$, the number of elements $\mathbf{N} > N_i$ decreases by 1 for each increment of i, which leads to the following survival function:

$$p(N \geq N_i) = 1 - \frac{i}{\{\mathbf{N}\}} \tag{7}$$

with $\{\mathbf{N}\}$ being the number of elements in $\mathbf{N}$ (i.e., $\{\mathbf{N}\} = 50,000$). For an observed number of events N_{obs}, we bisect the ordered set $\mathbf{N}$ to find the smallest i such that $N_i \geq N_{\text{obs}}$, and subsequently compute the corresponding p value. Depending on the total number of background events within each mainshock region, the empirical distribution may not be robust or representative of the underlying "true" distribution. Nonetheless, it is useful to compare the results from the empirical distributions with those assuming a parameterized (gamma or Poisson) distribution.

3. Results

The IET distributions and the corresponding fit with the gamma distribution for all 46 events is given in supporting information Figure S1. The values of γ and μ obtained from the fitting procedure, and their
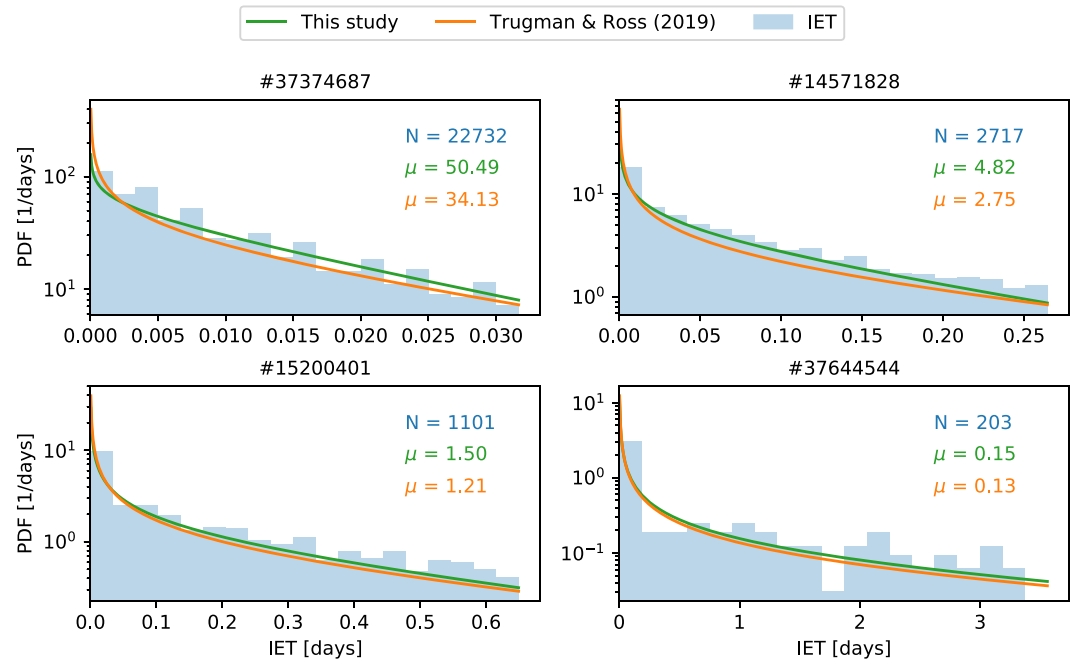
Figure 1. Interevent time (IET) distributions of four selected mainshock events (indicated in each panel's title), along with the best fit gamma distributions (this study and Trugman & Ross, 2019). The total number of events N and the inferred background rate μ (events per day) are as indicated in each panel.

comparison with those of Trugman and Ross (2019), are reported in Table S1. Examples of four selected events are given in Figure 1. Throughout the remainder of this work, the same four selected events are considered. In most cases, the quality of the fit is good if the number of events in the catalog is sufficient to produce robust statistics, indicating that the IET statistics are indeed captured reasonably well by a gamma distribution. The seismicity rate parameters μ inferred in this study are overall similar to those estimated by T&R, though our μ values are systematically higher. The minor discrepancies between our estimates of μ and those of T&R possibly arise from a difference in the selection procedure of the background seismicity (or foreshocks). Using the maximum likelihood approach, the solution of equation (2) in terms of γ is unique, and therefore the differences between the estimates of μ in the present study and T&R must result from a difference in the population statistics of τ. One possible origin could lie in the preprocessing of the QTM catalog, as T&R accounted for (occasional) station outages (p.c. D. Trugman), which was not done in this study. Furthermore, we exclude values of $\tau = 0$, that is, events that occur simultaneously in the catalog, since such cases are not permitted by the gamma distribution.

Based on random sampling of the best fit gamma distributions, distributions of the 20-day event counts (N) are obtained (see Figure 2). As discussed in section 2.2, N itself approximately follows a gamma distribution, which is in strong contrast to the Poisson distribution assumed by T&R. As can be clearly seen in Figure 2, the probability density given by a Poisson distribution rapidly decays toward zero with increasing N, often well before the median value of N is reached. This has tremendous implications for the resulting p value estimates, as the integral over the tail of the Poissonian N-distribution is effectively zero for any range of observed N.

Somewhat surprisingly, even though a gamma distribution describes the IET well, the resulting distribution of N obtained in this way does not match the empirical distribution of N. As is clearly seen in Figure 2 (and Figure S2), the empirical distributions are often much broader and more uniform than predicted based on the IETs. The mismatch between the two sampled distributions is unlikely to be attributable to the number of events and statistical robustness of the empirical distribution, since event #37374687 comprises close to 25,000 background events. Nonetheless, the heavy tails of the empirical distributions further disqualify a Poisson distribution to describe the 20-day event counts.

Figure 2. Distributions of the number of events (N) observed in a 20-day time window of four selected mainshock event (indicated in each panel's title), based on random sampling of the IET gamma and empirical distributions. N itself follows a gamma distribution, as indicated by the green line in each panel. The Poisson distribution as assumed by Trugman and Ross (2019) is indicated by the orange line.

To highlight this, we plot the probability curves (i.e., survival functions) of observing at least N events in a 20-day window of the four selected mainshock events (Figure 3). While the survival functions based on the gamma and empirical distributions plotted in Figure 2 remain well above the significance threshold in three out of four cases, the survival function based on the Poisson distribution (as adopted by T&R) sharply drops to zero in all selected cases, suggesting that all the observed earthquake sequences cannot be attributed to the background seismicity (i.e., $p < 0.01$ in all cases). In our analysis based on the gamma distribution, only 15 out of 46 mainshocks (32.6%) are characterized by elevated seismicity rates that are statistically significant ($p < 0.01$). This estimate decreases to 10 out of 46 (21.7%) when interrogating the empirical distributions.

4. Discussion

4.1. Temporal Fluctuations of Background Seismicity

While the analysis of the seismicity over 20 days prior to each mainshock has indicated that 33% of all mainshocks exhibit significantly elevated seismicity rates (compared to the average background rate), it is at this point not known whether this enhanced seismicity is uniquely associated to the mainshock event, or whether similar excursions from the background seismicity occur also in the absence of mainshock events. These excursions are expected to occur either in the form of swarm activity, or as aftershocks of smaller ($M_w < 4$) earthquakes. To this end, we compute the p value at a given moment in time by sampling the gamma distribution (as described in section 2.2) over a 20-day sliding window. This sliding window traverses the full duration of 380 days prior to the mainshock with strides of 1 day, so that the seismicity rate over this period is continuously evaluated relative to the (average) background seismicity rate. As before, we consider a p value less than 0.01 to define a statistically significant increase in seismicity. The result of this analysis for four selected mainshocks is presented in Figure 4 (the analysis of all 46 mainshocks is given in Figure S3).

The variability in the background seismicity rate can be quantified by computing the fraction of time-windows for which $p < 0.01$ for all mainshock events combined. Doing so gives a fraction of 15.7%, that is, at any given moment in time there is a 15.7% probability of observing an elevated seismicity rate. When comparing this value with the estimation that 33% of all mainshocks are preceded by significantly elevated
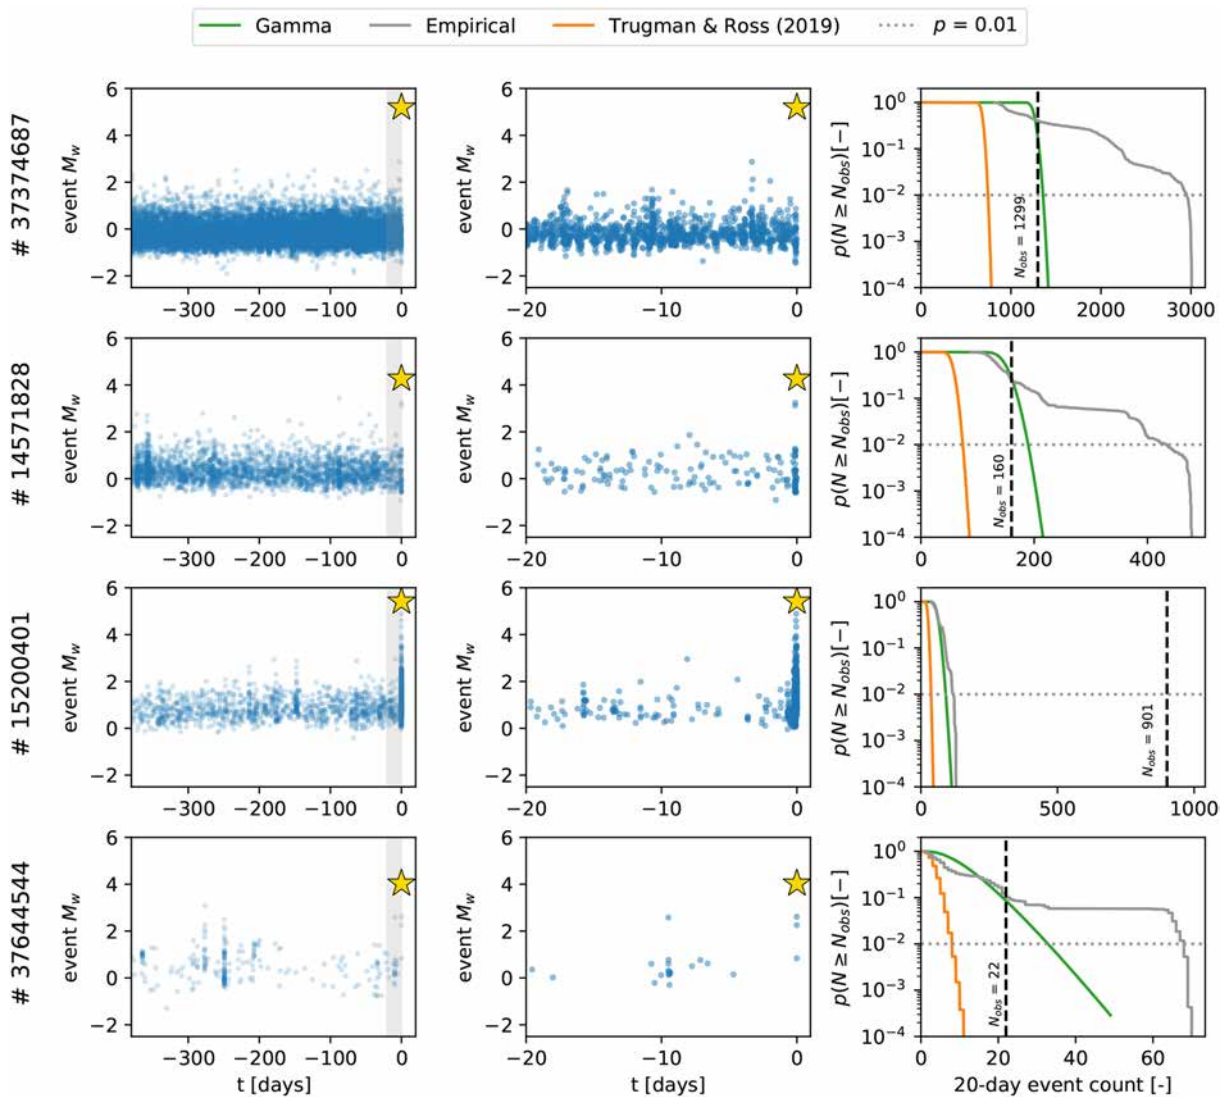
Figure 3. Synthesis of p value estimates of four selected mainshock events. (left and middle panels) Temporal record of earthquake occurrences from 380d (left) and 20d (middle) prior to the mainshock. The mainshock is indicated by the yellow star. (right panels) Probability curves of observing at least N events in a 20-day window, compared with the number of events observed 20d prior to the mainshock (N_{obs}). The present analysis based on random sampling of the gamma distribution, the empirical distribution, and the Poisson survival function based on the rate parameter μ estimated by Trugman and Ross (2019) are indicated by the green, gray, and orange lines, respectively. The significance threshold of 0.01 is represented by the horizontal dotted line.

seismicity rates, we conclude that about half of these "foreshock sequences" would have been expected purely on the basis of a fluctuating background rate. This leaves only a small number of mainshocks (about 8) in the catalog that are truly preceded by a foreshock sequence associated with the nucleation process of the earthquake.

In a way, the analysis of the background fluctuations gives a measure of the "false positive" rate, as none of the time-windows that exhibit a significantly elevated seismicity rate are associated with a mainshock (according to the definition of T&R). From conducting a brief sensitivity analysis (see Text S1; Agnew & Jones, 1991; Michael & Jones, 1998), we find that the false positive rate trades-off almost perfectly with the detection rate, so that, after correction for the background fluctuations, the percentage of mainshocks with statistically significant foreshock sequences falls in the range of 10% to 20%.

4.2. Statistical Significance of Foreshock Activity

When relaxing the assumption of Poissonian behavior, and when considering the temporal fluctuations in the background seismicity rate, only a minority of mainshock events is preceded by foreshock sequences.

Figure 4. Estimated p values for the 380 days prior to four selected mainshocks computed within a sliding 20-day window, based on the gamma and empirical distributions (respectively indicated by the green and gray lines). For reference, the significance threshold of 0.01 is indicated by the dotted line.

This is in strong contrast to the 70% estimated by T&R, who further alluded to the possibility of detecting more foreshock sequences in catalogs of higher completeness. The sample of mainshock events is small (46 in total), so it can be questioned whether or not the results presented here are significant considering the variation between mainshocks. However, a first-order sensitivity analysis (see Text S1) suggests that the results are robust with respect to certain arbitrary choices. Moreover, our findings are within the (wide) range of estimates given by previous studies (e.g., Abercrombie & Mori, 1996; Chen & Shearer, 2016; Jones & Molnar, 1976; Reasenberg, 1999), providing some confidence in the approach adopted in this study.

The quantitative differences between this study and that of T&R can also be seen qualitatively by considering the seismicity over the full 380d preceding the mainshock. For instance, events #37374687 and #37644544 (top and bottom rows in Figure 3) seem to exhibit seismicity rates over the 20 days immediately prior to the mainshock that are not unusually high considering the preceding 380 days. It is therefore not intuitive to imagine a p value that is practically zero, as given by the Poissonian survival function. On the other hand, event #14571828 (second row in Figure 3) does seem to exhibit a short burst of seismic activity on the same day as the mainshock event. Unfortunately, owing to the extent of the time window over which the p value is considered (20 days), this clear burst of activity is of insufficient duration to exceed the p value threshold. Therefore, to decide whether or not a particular mainshock exhibits significant foreshock activity, visual inspection of the catalog is recommended, rather than to rely purely on the computed p values. Alternatively, other seismological metrics such as a changing Gutenberg-Richter b value could be considered to alleviate the limitations of a fixed temporal window (Gulia & Wiemer, 2019).

5. Conclusions

In this study, we reassessed the significance of foreshock activity in Southern California, following the analysis of Trugman and Ross (2019). While the characterization of the background seismicity rate based on the interevent time (IET) method is valid, the subsequent assumption that all earthquakes in the catalog are statistically independent (and may therefore be described by a Poisson distribution) is overly restrictive. Consequently, the number of events expected to be observed in a 20-day time interval directly prior to a mainshock event is severely underestimated, and the number of mainshocks exhibiting statistically significant foreshock activity overestimated. Based on random sampling approaches that do not invoke the assumption of Poissonian behavior, we estimate that only 33% (15 out of 46) of all mainshocks are preceded by elevated

seismicity rates, while about half of that fraction is a priori anticipated based on the ubiquitous fluctuations in the background seismicity rate. In other words, we expect that only about 15% of all mainshocks exhibit a foreshock sequence uniquely associated with the earthquake preparation process.

Appendix A: Proof of Equation (6)

Suppose X and Y are independent random variables following a gamma distribution, with probability density functions (PDFs) $\rho_X(x; \gamma_x, \mu)$ and $\rho_Y(y; \gamma_y, \mu)$, respectively. The PDF of the sum of these variables ($Z = X + Y$) is found by convolving ρ_X with ρ_Y, that is,

$$\rho_Z(z) = \int_{-\infty}^{+\infty} \rho_X(x)\rho_Y(z-x)\mathrm{d}x \tag{A1}$$

For $x, y \geq 0$ (and consequently $z - x \geq 0$), the PDFs are nonzero only over the range $[0, z]$, which defines the upper and lower limits of the integration. Substituting the definition of the gamma distribution PDF (see equation (1)), we get

$$\begin{aligned}\rho_Z(z) &= \int_0^z \frac{x^{\gamma_x-1}e^{-\mu x}}{\mu^{-\gamma_x}\Gamma(\gamma_x)} \frac{(z-x)^{\gamma_y-1}e^{-\mu(z-x)}}{\mu^{-\gamma_y}\Gamma(\gamma_y)}\mathrm{d}x \\ &= \frac{e^{-\mu z}}{\mu^{-[\gamma_x+\gamma_y]}} \int_0^z \frac{x^{\gamma_x-1}(z-x)^{\gamma_y-1}}{\Gamma(\gamma_x)\Gamma(\gamma_y)}\mathrm{d}x\end{aligned} \tag{A2}$$

A change of variables $x = zt$ gives:

$$\rho_Z(z) = \frac{e^{-\mu z}}{\mu^{-[\gamma_x+\gamma_y]}} z^{\gamma_x+\gamma_y-1} \int_0^1 \frac{t^{\gamma_x-1}(1-t)^{\gamma_y-1}}{\Gamma(\gamma_x)\Gamma(\gamma_y)}\mathrm{d}t \tag{A3}$$

Finally, we recognize the emergence of the beta distribution, defined as

$$B(t; \alpha, \beta) = \Gamma(\alpha+\beta)\frac{t^{\alpha-1}(1-t)^{\beta-1}}{\Gamma(\alpha)\Gamma(\beta)} \tag{A4}$$

which, upon substitution, gives

$$\begin{aligned}\rho_Z(z) &= \frac{e^{-\mu z}}{\mu^{-[\gamma_x+\gamma_y]}} \frac{z^{\gamma_x+\gamma_y-1}}{\Gamma(\gamma_x+\gamma_y)} \int_0^1 B(t; \gamma_x, \gamma_y)\mathrm{d}t \\ &= \frac{e^{-\mu z}}{\mu^{-[\gamma_x+\gamma_y]}} \frac{z^{\gamma_x+\gamma_y-1}}{\Gamma(\gamma_x+\gamma_y)} \\ &= \frac{z^{\gamma_z-1}e^{-\mu z}}{\mu^{-\gamma_z}\Gamma(\gamma_z)}\end{aligned} \tag{A5}$$

In other words, the sum of two gamma distributed variables produces a new gamma distributed variable, with shape parameter $\gamma_z = \gamma_x + \gamma_y$. Taking $\gamma_x = \gamma_y = \gamma$, the sum over N gamma distributed variables yields a gamma distributed variable with shape parameter $N\gamma$.

References

Abercrombie, R. E., & Mori, J. (1996). Occurrence patterns of foreshocks to large earthquakes in the western United States. *Nature*, *381*(6580), 303–307. https://doi.org/10.1038/381303a0

Agnew, D. C., & Jones, L. M. (1991). Prediction probabilities from foreshocks. *Journal of Geophysical Research*, *96*(B7), 11,959–11,971. https://doi.org/10.1029/91JB00191

Bouchon, M., Durand, V., Marsan, D., Karabulut, H., & Schmittbuhl, J. (2013). The long precursory phase of most large interplate earthquakes. *Nature Geoscience*, *6*(4), 299–302. https://doi.org/10.1038/ngeo1770

Chen, X., & Shearer, P. M. (2016). Analysis of foreshock sequences in California and implications for earthquake triggering. *Pure and Applied Geophysics*, *173*(1), 133–152. https://doi.org/10.1007/s00024-015-1103-0

Dodge, D. A., Beroza, G. C., & Ellsworth, W. L. (1996). Detailed observations of California foreshock sequences: Implications for the earthquake initiation process. *Journal of Geophysical Research*, *101*(B10), 22,371–22,392. https://doi.org/10.1029/96JB02269

Gulia, L., & Wiemer, S. (2019). Real-time discrimination of earthquake foreshocks and aftershocks. *Nature*, *574*(7777), 193–199. https://doi.org/10.1038/s41586-019-1606-4

Hainzl, S., Scherbaum, F., & Beauval, C. (2006). Estimating background activity based on interevent-time distribution. *Bulletin of the Seismological Society of America*, *96*(1), 313–320. https://doi.org/10.1785/0120050053

Jones, L., & Molnar, P. (1976). Frequency of foreshocks. *Nature*, *262*(5570), 677–679. https://doi.org/10.1038/262677a0

Jones, L. M., & Molnar, P. (1979). Some characteristics of foreshocks and their possible relationship to earthquake prediction and premonitory slip on faults. *Journal of Geophysical Research*, *84*(B7), 3596–3608. https://doi.org/10.1029/JB084iB07p03596

Marsan, D., & Enescu, B. (2012). Modeling the foreshock sequence prior to the 2011, M_W 9.0 Tohoku, Japan, earthquake. *Journal of Geophysical Research*, *117*, B06316. https://doi.org/10.1029/2011JB009039

Michael, A. J., & Jones, L. M. (1998). Seismicity alert probabilities at Parkfield, California, revisited. *Bulletin of the Seismological Society of America*, *88*(1), 117–130.

Reasenberg, P. (1985). Second-order moment of central California seismicity, 1969–1982. *Journal of Geophysical Research*, *90*(B7), 5479–5495. https://doi.org/10.1029/JB090iB07p05479

Reasenberg, P. A. (1999). Foreshock occurrence before large earthquakes. *Journal of Geophysical Research*, *104*(B3), 4755–4768. https://doi.org/10.1029/1998JB900089

Roeloffs, E. A. (2006). Evidence for aseismic deformation rate changes prior to earthquakes. *Annual Review of Earth and Planetary Sciences*, *34*(1), 591–627. https://doi.org/10.1146/annurev.earth.34.031405.124947

Ross, Z. E., Trugman, D. T., Hauksson, E., & Shearer, P. M. (2019). Searching for hidden earthquakes in Southern California. *Science*, *364*(6442), 767–771. https://doi.org/10.1126/science.aaw6888

Trugman, D. T., & Ross, Z. E. (2019). Pervasive foreshock activity across Southern California. *Geophysical Research Letters*, *46*, 8772–8781. https://doi.org/10.1029/2019GL083725