Extreme value modelling of SARS-CoV-2 community transmission using discrete Generalised Pareto distributions

Abdelaati Daouia, Gilles Stupfler, Antoine Usseglio-Carleve

Extreme value modelling of SARS-CoV-2 community transmission using discrete Generalised Pareto distributions

Abdelaati Daouia[a], Gilles Stupfler[b] & Antoine Usseglio-Carleve[c]

[a] University of Toulouse Capitole, TSE - Decision Mathematics and Statistics, France (ORCID: 0000-0003-2621-8860)

[b] Univ Rennes, Ensai, CNRS, CREST - UMR 9194, F-35000 Rennes, France (ORCID: 0000-0003-2497-9412)

[c] Avignon Université, Laboratoire de Mathématiques d'Avignon EA 2151, 84000 Avignon, France (ORCID: 0000-0002-8148-3758)

Abstract. Superspreading has been suggested to be a major driver of overall transmission in the case of SARS-CoV-2. It is therefore important to statistically investigate the tail features of superspreading events (SSEs) to better understand virus propagation and control. Our extreme value analysis of different sources of secondary case data indicates that SSEs associated with SARS-CoV-2 may be fat-tailed, although substantially less so than predicted recently in the literature, but also less important relative to SSEs associated with SARS-CoV. The results caution against pooling data from both coronaviruses. This could provide policy- and decision-makers with a more reliable assessment of the tail exposure to SARS-CoV-2 contamination. Going further, we consider the broader problem of large community transmission. We study the tail behaviour of SARS-CoV-2 cluster cases documented both in official reports and in the media. Our results suggest that the observed cluster sizes have been fat-tailed in the vast majority of surveyed countries. We also give estimates and confidence intervals of the extreme potential risk for those countries. A key component of our methodology is up-to-date discrete Generalised Pareto models which allow for maximum-likelihood based inference of data with a high degree of discreteness.

Keywords. COVID-19, Superspreading, Cluster size, Secondary cases, Extreme value theory, Discrete extremes.

Introduction

Superspreading events (SSEs) have been recognised as a significant source of disease transmission for respiratory coronaviruses such as SARS-CoV and SARS-CoV-2 [1, 2]. SSEs may be defined as outbreaks in which a given individual (the index case)

infects a number of people (secondary cases) well above a certain measure, such as the average or median number of infections. The number of secondary cases resulting directly from an index case can be viewed as a random variable, say Z, defining the so-called offspring distribution. For both coronaviruses, events having triggered more than 6 secondary cases have been suggested to constitute SSEs [3]. Data on such SSEs that was curated and reported in [3] in the early stages of the COVID-19 pandemic is necessarily scarce: it consists mainly of 15 SSEs associated with SARS-CoV and 45 SSEs associated with SARS-CoV-2, each represented by a number of secondary cases Z_i resulting from a single given index case in Europe, Asia or North America. The natural framework for the analysis of SSEs, and more generally of atypical observations far away from the mean, is extreme value theory. Following this framework, it was argued in [3] that SSEs are fat-tailed, although this was done by pooling the 60 available SSEs from SARS-CoV and SARS-CoV-2. A careful investigation of these SARS-CoV and SARS-CoV-2 datasets reveals that the two largest observations in the pooled data are SARS-CoV SSEs; given the small sample size, one may wonder whether the reported estimate of tail heaviness is representative of the tail behaviour of SARS-CoV-2 SSEs.

This constitutes the motivation for this work, whose overarching goals are to show how to conduct a principled extreme value analysis of community transmission parameters, and to carry out such an analysis in the example of SARS-CoV-2. By focusing directly on the raw SARS-CoV-2 data considered in [3], we provide evidence of a lighter upper tail for SSEs with significantly less tail exposure than predicted in their study. We arrive at the same conclusion by making use of a more recent and much larger publicly available surveillance and contact-tracing database containing the number of secondary cases Z_i for 88,527 index cases in the Indian states of Andhra Pradesh and Tamil Nadu [4]. We also analyse two other South Korean contact-tracing datasets, one collected in the first half of 2020 [3], the other during the summer of 2021 when the Delta variant of SARS-CoV-2 was responsible for the majority of positive cases [5]. The fat-tailedness of the secondary cases distribution is found to be rather clear in the 2021 sample of data, while the analysis of the 2020 data is less conclusive. In all these samples of data we find point estimates of the extreme value index suggesting that the secondary cases distribution has a finite third moment, which stands in contrast with the earlier finding of [3] of a distribution with an infinite variance.

In addition to that, we consider the broader problem of large community transmission, as it represents the other fundamental source of pandemic risk. Large infection clusters, along with SSEs, have been argued to play an important role in the transmission of SARS-CoV-2 [2]. In a similar spirit to [2], we define a cluster of SARS-CoV-2 cases in our analysis as a local outbreak involving a minimum of two cases, including confirmed close contacts with epidemiological linkage over a limited period of time. We consider two databases constructed from government reports [6, 7, 8, 9] and media sources [10], comprising 15 samples of SARS-CoV-2 cluster sizes recorded in 11 countries and 4 US states. Our results show that 13 of these 15 countries and states have fat-tailed cluster size distributions, thus facilitating the process of inferring their risk category in terms of large community transmission. This allows us to better understand the drivers of superspreading and

cluster formation in the ongoing COVID-19 pandemic. The recent theory of discrete extremes [11, 12, 13, 14] is our basic tool to address the highly discrete nature of SARS-CoV-2 secondary transmission data and cluster sizes. Its use constitutes our main statistical contribution to the study of the transmission of the SARS-CoV-2 virus. As we illustrate throughout the paper, estimating and inferring the extreme value index and extreme percentiles of the underlying discrete distributions with this methodology is much easier and accurate than with classical extreme value methods such as the Hill and Generalised Pareto maximum likelihood estimators, which heavily rely on the continuous data assumption.

The structure of the paper is as follows. We first describe the methods employed throughout our study, including the discrete Generalised Pareto Distribution fitted to exceedances over a high threshold by means of the maximum likelihood estimator. We then analyse our datasets, first on SARS-CoV-2 secondary case numbers and then on cluster sizes, using these methods. A Discussion section gathers and contrasts these findings and concludes with additional comments about the scope, limitations and robustness of our results, as well as ideas for further work.

Methods

We use several methods from extreme value theory, which constitutes the correct mathematical framework for the analysis of high observations from a random phenomenon [15]. We are particularly interested in methods that can describe so-called heavy-tailed random variables, which infrequently but regularly generate very high values and therefore appear to be relevant in the analysis of SARS-CoV-2 transmission. A random variable X is heavy-tailed (or fat-tailed) if and only if its distribution function $\mathbb{P}(X \leq x)$ can be expressed as $\mathbb{P}(X \leq x) = 1 - x^{-1/\xi}\ell(x)$, where ℓ satisfies $\ell(tx)/\ell(t) \to 1$ as $t \to \infty$ for any positive real number x. Informally, the tail behaviour of X is controlled by the extreme value index $\xi > 0$, which must be estimated to get a precise understanding of tail heaviness. A standard estimator in this context is the Hill estimator [16]. For a dataset $Z_1, \ldots, Z_n$, the Hill estimator at threshold u is defined as

$$\hat{\xi}_u^H = \frac{1}{\sum_{i=1}^n \mathbb{1}_{\{Z_i > u\}}} \sum_{i=1}^n \log\left(\frac{Z_i}{u}\right) \mathbb{1}_{\{Z_i > u\}}.$$

It is of course crucial, before using the Hill estimator, to ascertain whether the distribution of the data points indeed has a heavy tail. A common diagnostic method is the mean excess plot, which estimates the values of the mean excess function $E(u) = \mathbb{E}[Z - u | Z > u]$ as function of u. A natural estimate of $E(u)$ is given, for each threshold u, by its empirical counterpart

$$\hat{E}(u) = \frac{\sum_{i=1}^n Z_i \mathbb{1}_{\{Z_i > u\}}}{\sum_{i=1}^n \mathbb{1}_{\{Z_i > u\}}} - u.$$

A heavy-tailed distribution will typically have mean excess plots exhibiting a linear upward drift for large values of u, see for example Section 1.2.2 in [17]. It has, however, been observed in the extreme value literature [18] that the mean excess function

very often exhibits a non-linear behaviour at the right end of the mean excess plot, due to very high variability of the estimate of $E(u)$ when u is close to the highest Z_i. As a consequence, good statistical practice recommends to confirm a diagnostic of a heavy tail using other extreme value tools. One such general approach, which does not presuppose that the data is heavy-tailed, consists in using the Generalised Pareto maximum likelihood estimator, defined as, according to Section 5.3.2 in [17]:

$$\left(\hat{\xi}_u^{GP}, \hat{\sigma}_u^{GP}\right) = \underset{(\xi,\sigma)\in(-1/2,\infty)\times(0,\infty)}{\arg\min} \sum_{i=1}^{n} \left[-\log\sigma - \left(\frac{1}{\gamma}+1\right) \log\left(1+\xi\frac{Z_i-u}{\sigma}\right)\right] \mathbb{1}_{\{Z_i>u\}}.$$

The Generalised Pareto maximum likelihood estimators are valid even when the underlying distribution is not heavy-tailed, which has made them very popular in the natural sciences [19].

However, both the Hill and Generalised Pareto estimators of ξ suffer from jagged sample paths when the data points Z_i come from a distribution with a high degree of discreteness. This behavior makes it extremely difficult to choose an accurate estimate of ξ, which renders the two methods highly unsatisfactory. The essential reason behind this phenomenon is that both estimators are built under the – generally incorrect – assumption that the data points come from a pure (Generalised) Pareto distribution, which is continuous, and as such, they cannot be expected to handle a substantial degree of discreteness. We exemplify this phenomenon in Fig. 1: notice, in the top panels, the stark difference in stability and smoothness of sample paths between a Hill plot for continuous data Z_i and its counterpart for data rounded to the nearest integer up. The bottom panels show that the Hill estimator for discrete data tends to be strongly biased and much more so than the Hill estimator for continuous data.

A statistically principled alternative is to employ proper discrete models to construct an estimator of the extreme value index. This was pursued by [13], which used so-called D-GPD (for Discrete-Generalised Pareto Distribution) models to introduce the maximum likelihood estimators

$$\left(\hat{\xi}_u, \hat{\sigma}_u\right) = \underset{(\xi,\sigma)\in\mathbb{R}\times(0,\infty)}{\arg\min} \sum_{i=1}^{n} \log\left(\left(1+\xi\frac{Z_i-u}{\sigma}\right)^{-1/\xi} - \left(1+\xi\frac{Z_i-u+1}{\sigma}\right)^{-1/\xi} \right) \mathbb{1}_{\{Z_i\geq u\}}.$$

When $\xi = 0$, the convention we adopt is that $(1+\xi z)^{-1/\xi} = \exp(-z)$, for any $z \in \mathbb{R}$. These maximum likelihood estimators of the extreme value index ξ and scale parameter σ of the D-GPD model are readily obtained through the R maximisation routine `optim`. Using the classical theory of maximum likelihood estimators, confidence intervals for ξ may be derived from $\hat{\xi}_u$ by estimating the total Fisher information matrix $I(\xi,\sigma)$ using a finite difference method and then deducing the following $100\alpha\%$-confidence interval for ξ:

$$\left[\hat{\xi}_u + \sqrt{\left(\hat{I}(\xi,\sigma)^{-1}\right)_{1,1}}\,\Phi^{-1}\left(\frac{1-\alpha}{2}\right),\ \hat{\xi}_u + \sqrt{\left(\hat{I}(\xi,\sigma)^{-1}\right)_{1,1}}\,\Phi^{-1}\left(\frac{1+\alpha}{2}\right)\right],$$

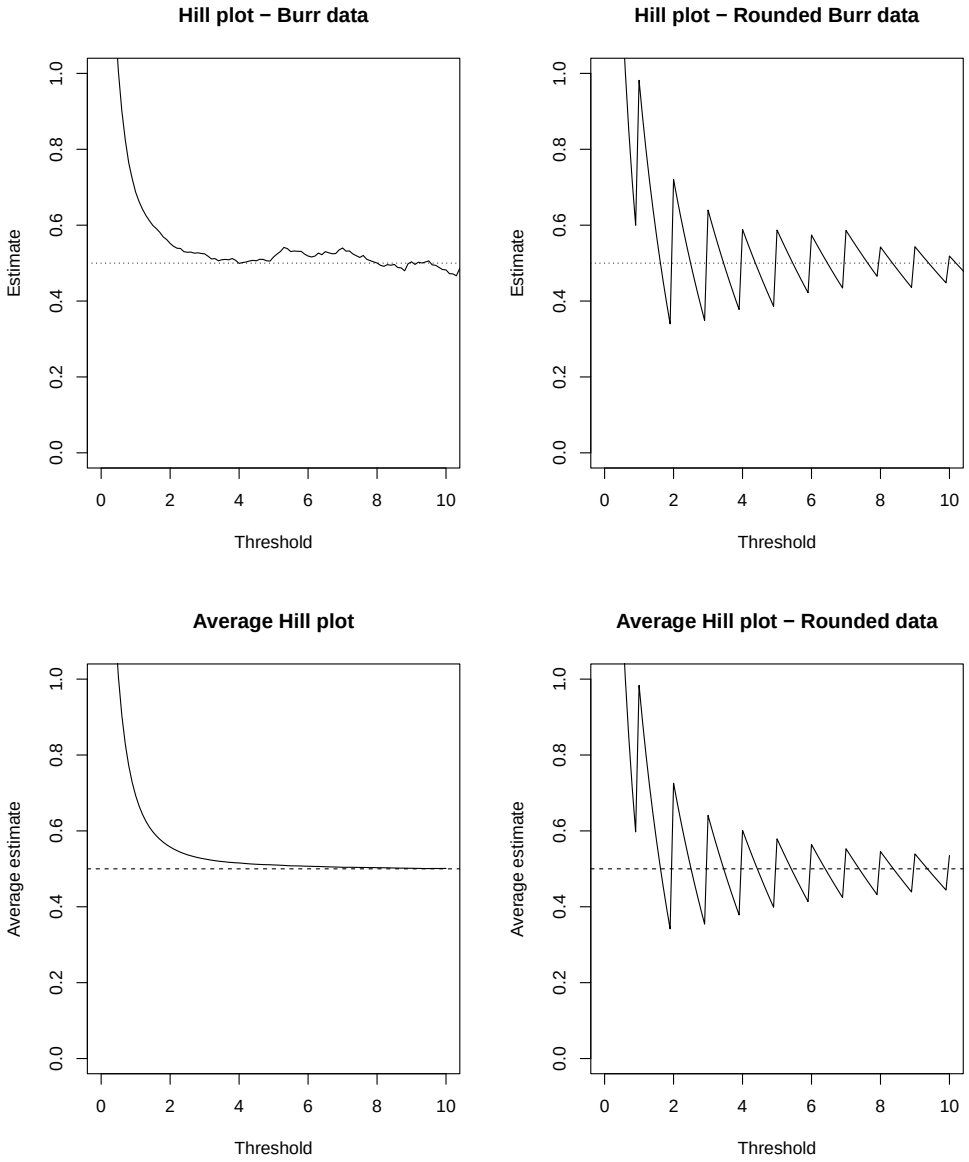

Figure 1: Top panels: Hill plots as functions of the threshold value u, for $n = 10{,}000$ simulated data points Z_i from the Burr distribution with probability density function $f(x) = \xi^{-1} x^{-\rho/\xi - 1}(1 + x^{-\rho/\xi})^{1/\rho - 1}$ (for $x > 0$) with $\xi = 1/2$ and $\rho = -1$ in the left panel, and for the data $\lceil Z_i \rceil$ (*i.e.* the smallest integer larger than or equal to Z_i) in the right panel. Bottom panels: Averaged Hill plots when this experiment is repeated $N = 1{,}000$ times.

where Φ denotes the standard normal distribution function and Φ^{-1} its inverse (quantile function). Modelling $Z-u$ conditional on $Z \geq u$ by a D-GPD distribution with parameter estimates $(\hat{\xi}_u, \hat{\sigma}_u)$ suggests the following estimate of the 100αth percentile of Z adapted from [12, Formula (5) p.41]:

$$\hat{q}_\alpha = \left\lceil \frac{\hat{\sigma}_u}{\hat{\xi}_u} \left(\left(\frac{n(1-\alpha)}{\sum_{i=1}^n \mathbb{1}_{\{Z_i \geq u\}}} \right)^{-\hat{\xi}_u} - 1 \right) + u - 1 \right\rceil,$$

for α large enough. Here, $\lceil \cdot \rceil$ denotes the ceiling function, that is, $\lceil x \rceil$ denotes the smallest integer larger than or equal to x. Estimating this quantile by plugging in the aforementioned estimates of ξ and σ makes it possible to infer extreme quantile levels and therefore get precise information on the tail behaviour of a distribution with a large degree of discreteness.

For comparison purposes, we will contrast the resulting extreme quantile estimates with those provided by the (conditioned) negative binomial distribution. Recall that the probability mass function of the negative binomial distribution (with parameters $r > 0$ and $p \in (0,1)$) conditional on $Z > u$, is given by

$$\mathbb{P}_{p,r,u}(Z = k) = \frac{\frac{\Gamma(k+r)}{k!\,\Gamma(r)} p^r (1-p)^k}{1 - \sum_{i=0}^u} \frac{\Gamma(i+r)}{i!\,\Gamma(r)} p^r (1-p)^i, \text{ for all } k > u.$$

Here Γ denotes Euler's Gamma function. With a dataset $z_1, \ldots, z_n$, the parameter estimators are therefore obtained as the maximum log-likelihood solution

$$\underset{(p,r)\in(0,1)\times(0,\infty)}{\arg\max} \sum_{i=1}^n \log \mathbb{P}_{p,r,u}(Z = z_i).$$

Ever since the seminal work of [1], the negative binomial distribution has been widely used to describe the number of secondary cases resulting from an index case of SARS-CoV. As suggested in [3, 21], this model has exponentially decreasing probability mass functions and thus cannot be expected to accurately represent tail heaviness in SARS-CoV-2 transmission data. We provide below further evidence for this claim, and for the suitability of D-GPD maximum likelihood estimates in the context of discrete data, through several datasets gathering numbers of SARS-CoV-2 secondary cases and cluster sizes in different settings.

Data and results

Analysis of secondary case data. Our first two datasets were reported in [3]. They consist of 15 SSEs associated with SARS-CoV (Dataset S1) and 45 SSEs associated with SARS-CoV-2 (Dataset S2), each resulting in more than 6 secondary cases, along with month of occurrence and location of the superspreading event, and its setting. We refer to [3] for further details about the construction of these datasets. Pooling the 15 SSEs associated with SARS-CoV and 45 SSEs associated with SARS-CoV-2 into a single sample and making use of a Generalised Pareto approximation, [3] has suggested that the distribution of the number of secondary

cases Z belongs to the Fréchet maximum domain of attraction [20], that is, the set of Pareto-type distributions, with tail index ξ between 0.5 and 1 (the estimate provided in [3, Fig. 1 E] is $\hat{\xi} \approx 0.6$). The index ξ tunes the tail heaviness of the distribution, with higher positive values indicating a heavier upper tail: moments of order higher than or equal to $1/\xi$ do not exist. An estimate of ξ around 0.6 means that the second moment of Z does not exist, reflecting the outsized contribution of SSEs to overall transmission. Most importantly perhaps, these findings on the tail heaviness of Z invalidate the conventional assumption that Z follows a negative binomial distribution for either coronavirus, whereas this assumption was widely adopted in the literature on disease transmission ever since the influential work [1] on SARS-CoV, and it is still widely employed for SARS-CoV-2, see *e.g.* [5, 22, 23].

Based on our statistical analysis of these datasets, summarised in Fig. 2, one may however argue that the method of [3] is inappropriate for examining the tail behaviour of their particular 60 SSEs. The sparsity of data on SSEs is addressed by combining the 15 and 45 observations associated with SARS-CoV and SARS-CoV-2 into a single sample, whereas the two datasets correspond to completely different distributions (Fig. 2 (a)) and should not be pooled accordingly. This is apparent from either a Kolmogorov-Smirnov test, with $p-$value 0.015, or the more common approach making the questionable assumption that Z follows a negative binomial distribution. The conditional (given $Z > 6$) negative binomial fit of the probability mass function to the Z_i (by construction larger than 6), calculated as described in the last paragraph of the Methods section (Fig. 2 (b)), already suggests that the upper tail of Z for SARS-CoV appreciably dominates that for SARS-CoV-2. In other words, even a naive analysis of the SSE distributions, using the classical negative binomial distribution and not accounting for the heavy tail in the data, indicates that the SSEs for SARS-CoV and those for SARS-CoV-2 exhibit different statistical behaviour. This is confirmed by a proper extreme value analysis of the data (Fig. 2 (c)): the ξ estimates obtained from the Hill estimator in the special case of SARS-CoV-2 vary between 0.35 and 0.45, and as such differ substantially from the various competing estimates found to vary between 0.5 and 1 in [3]. Even the 90% confidence intervals of ξ for SARS-CoV-2 (dashed red lines in Fig. 2 (c)) only partially contain the estimated tail index plot for SARS-CoV (solid blue line), reflecting a net difference between the two heavy-tailed distributions of secondary cases associated with SARS-CoV and SARS-CoV-2. This conclusion is corroborated by the mean excess function estimates (Fig. 2 (d)), which similarly indicate the relevance of separating the analysis for each coronavirus. This suggests that although SARS-CoV and SARS-CoV-2 belong to the same family of respiratory diseases, superspreading events are larger in scale for SARS-CoV in comparison to SARS-CoV-2. For all these reasons, pooling the data before applying extreme value tools can lead to misleading conclusions on the propagation of the SARS-CoV-2 virus.

Yet, the low sample size of this SSE dataset puts a question mark over the quality of the statistical analysis. Trustworthy extreme value inference may require a larger sample size, of the order of at least several thousands. This is why we also analysed a much larger Indian secondary case dataset of size $n = 88{,}527$ (Database S3). This comprehensive surveillance and contact-tracing database was collected in 2020 by the public health authorities of the two Indian states of Andhra Pradesh and Tamil

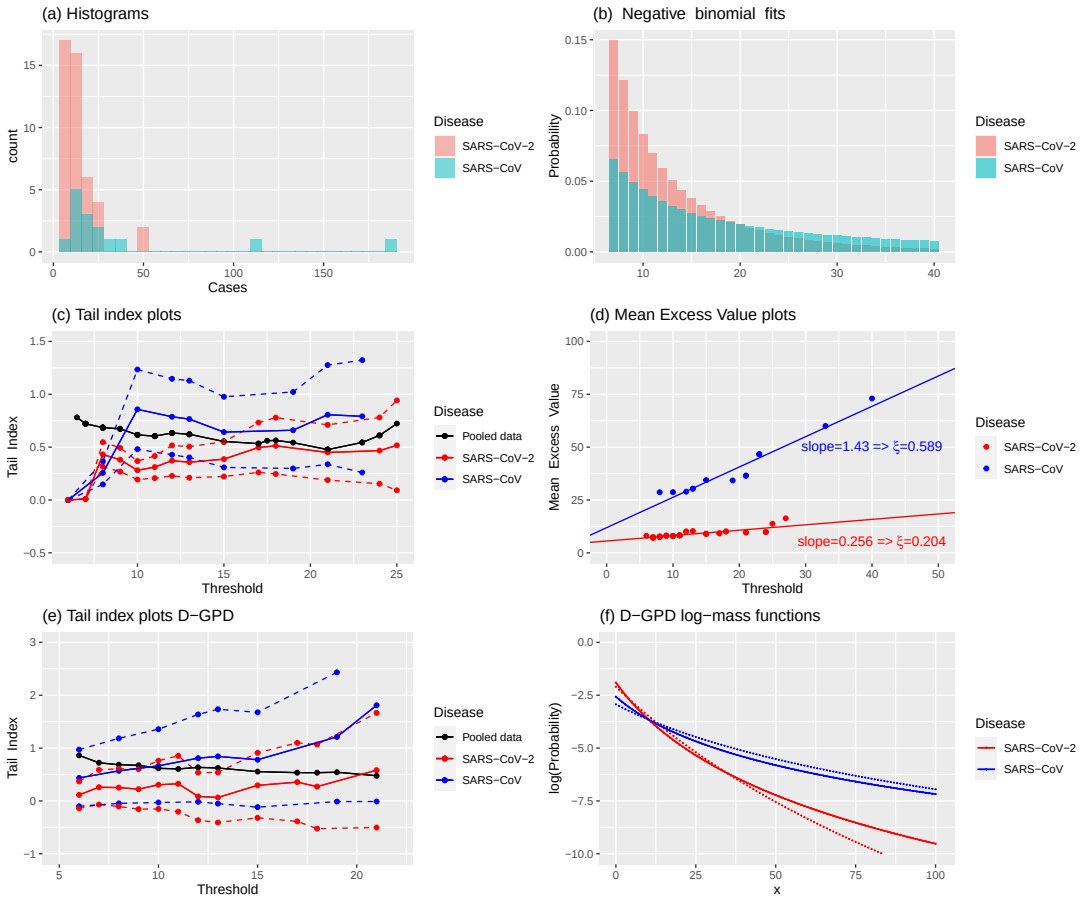

Figure 2: Secondary case data from [3] (Datasets S1 and S2). (a) Histogram of the number of secondary cases for SARS-CoV (blue, $n = 15$) and SARS-CoV-2 (red, $n = 45$) SSEs. (b) Fitted probability mass function, conditional on $Z > 6$, of the negative binomial distribution for SARS-CoV (blue) and SARS-CoV-2 (red) SSEs. (c) Hill estimates of ξ for SSEs associated with SARS-CoV (solid blue), SARS-CoV-2 (solid red), and the pooled data (solid black), obtained from the exceedance values $Z_i - u$ given $Z_i \geq u$, as function of the threshold u, along with the resulting 90% confidence intervals for SARS-CoV (dashed blue) and SARS-CoV-2 (dashed red) SSEs. (d) Mean excess plots of SARS-CoV (blue) and SARS-CoV-2 (red) SSEs, quantified by the average of the exceedances $Z_i - u$ given $Z_i \geq u$, as function of u. (e) Discrete GPD maximum likelihood estimates of ξ for SARS-CoV (solid blue) and SARS-CoV-2 (solid red) SSEs, calculated from the exceedances $Z_i - u$ given $Z_i \geq u$, as function of u, along with their corresponding 90% confidence intervals (dashed lines), and the Hill plot produced by combining SARS-CoV and SARS-CoV-2 SSEs. (f) Logarithm of the probability mass functions $\mathbb{P}_{\sigma,\xi}(X = x)$ of the D-GPD fits to the exceedance values $Z_i - u$ given $Z_i \geq u$, for the thresholds $u = 6$ (dotted lines) and $u = 10$ (solid lines), for SARS-CoV (blue) and SARS-CoV-2 (red).

Nadu, whose residents total about 10% of India's population. It was studied for instance in [4] and [21], and we refer to the latter for more information about the database's construction and contents. Results are reported in Fig. 3. Although the barplot of this data (Fig. 3 (a)) gives evidence of a considerable right skewness and its summary extreme value analysis (Fig. 3 (b)) suggests a heavy right tail, it should be noted that since the Z_i range from 0 to 39 with a sample size of 88,527, the data is necessarily highly discrete with a large number of tied observations (see Table 1).

Z	0	1	2	3	4	5	6	7	8	9	10	11
Count	62,540	17,493	4,885	1,730	802	444	267	149	67	44	29	22

Z	12	13	14	15	16	17	18	19	21	22	23	25	28	31	37	39
Count	14	16	3	3	4	4	1	1	2	1	1	1	1	1	1	1

Table 1: Secondary case data (Database S3) for SARS-CoV-2 from Andhra Pradesh and Tamil Nadu (India).

Ignoring the discrete nature of the Z_i by modelling their tail behaviour with the (Generalised) Pareto distribution is inappropriate as this typically results in unreliable tail index estimates and confidence intervals [13]. This becomes obvious here by superimposing both the classical Hill and continuous Generalised Pareto maximum likelihood estimators of the extreme value index, as functions of a varying threshold u in Fig. 3 (c). Clearly, both plots are so volatile and jagged that it is hard to identify any stable region and therefore a reasonable point estimate of ξ cannot easily be determined. We address this limitation by applying the recent theory of discrete extremes developed in [11, 13] and based on the discrete Generalised Pareto distribution (D-GPD). The D-GPD, first employed by [12] to model road accidents and more recently in [14] to model hospital congestion, has been shown to outperform the continuous GPD when there are a large number of tied observations: see the simulated Poisson and discrete Inverse-Gamma examples in Section 3.1 of [13], which respectively show that the GPD provides poor fits and poor tail estimates when the data is highly discrete, while the D-GPD distribution performs well. Its closed-form survival and probability mass functions allow for an exact likelihood-based inference. Using the D-GPD distribution to fit exceedances $Z_i - u$ above the threshold u (rather than trying to fit the whole of the distribution, as [21] did using a discrete Pareto distribution) results in a much smoother and stable fit (Fig. 3 (c)), and leads to an estimate of ξ around 0.24 with the 90% confidence intervals overwhelmingly suggesting an estimate greater than 0, thus confirming the heavy-tailed nature of SARS-CoV-2 SSEs (Fig. 3 (d)) in this sample. Interestingly, revisiting the small SARS-CoV-2 SSE dataset (Dataset S2) of size 45 using the D-GPD maximum likelihood estimation method (Fig. 2 (e)) results in an estimate of around 0.25, in agreement with the results from the Indian secondary case data. This suggests that the distribution of SARS-CoV-2 SSEs has a finite third moment and possibly even a fourth moment. These results are different from those obtained for the SARS-CoV SSEs. The latter rather point towards a distribution with infinite variance and thus a much heavier right tail. This is confirmed by considering the fitted D-GPD probability mass functions for secondary cases (Fig. 2 (f)) that decrease much more rapidly for SARS-CoV-2 than for SARS-CoV.

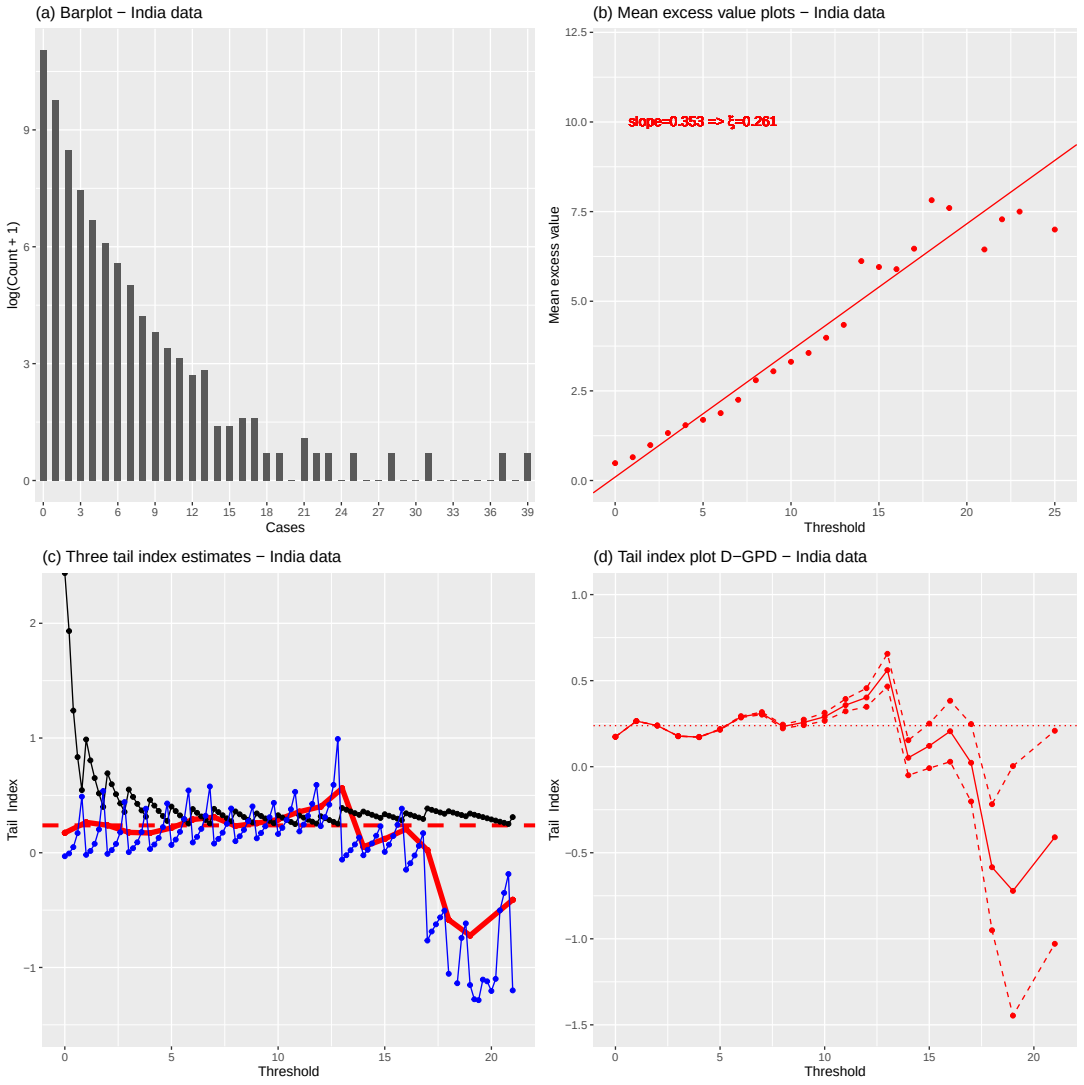

Figure 3: Secondary case data (Database S3) for SARS-CoV-2 from Andhra Pradesh and Tamil Nadu (India). (a) Bar plot of the $\log(Z_i+1)$ ($n = 88{,}527$). (b) Mean excess plots of secondary cases. (c) Hill (solid black), continuous GPD maximum likelihood (solid blue) and discrete GPD maximum likelihood (solid bold red) estimates of ξ. (d) Discrete GPD maximum likelihood estimates of ξ (solid red) and their associated 90% confidence intervals (dashed red). In panels (c) and (d), the averaged discrete GPD estimate $\hat{\xi} = 0.239$ over the stable region $u \in [0, 10]$ is indicated with the horizontal red line.

To examine the extreme value behaviour of the SARS-CoV-2 offspring distribution in different conditions, we turn to the analysis of two contact-tracing datasets in South Korea, a country which has a similar population density to the Indian state of Tamil Nadu, but did not resort to any full lockdown and has one of the largest and best-organised epidemic control programmes in the world. The first dataset was collected in the first half of 2020 (Database S4), while the second was collected during the fourth community epidemic in the summer of 2021 (Database S5) in the context of the assessment of transmission dynamics for the Delta variant of SARS-CoV-2. The first dataset, which consists of $n = 5{,}165$ numbers of SARS-CoV-2 secondary cases Z_i, was analysed in [3]. See Table 2.

Z	0	1	2	3	4	5	6	7	8	9
Count	4,558	364	114	62	27	7	7	4	4	1

Z	10	11	12	15	17	18	21	24	27	51
Count	2	3	1	2	2	1	2	2	1	1

Table 2: Secondary case data (Database S4) for SARS-CoV-2 collected in South Korea in the first half of 2020.

We revisit the estimation of, and inference about, the underlying extreme value index by comparing the D-GPD estimates with the classical GPD and Hill estimates. Results are displayed in Fig. 4. A least squares fit to the first part of the mean excess plot (Fig. 4 (b)) suggests a linearly increasing fit to the mean excess function with a slope of around 0.85, but this ignores the flat or even slightly linearly decreasing right-hand part of the data cloud. This throws the assumption that the offspring distribution is heavy-tailed in doubt, although the barplot of the data (Fig. 4 (a)) would tentatively back the heavy tail assumption. The Hill estimator, which presupposes that the data is heavy-tailed and graphed as a black line in Fig. 4 (c), does not exhibit any stable region which would allow to produce a reasonable point estimate. In such scenarios, best practice in extreme value theory requires calculating alternative extreme value estimators whose consistency does not rest upon the heavy tail assumption (unlike the Hill estimator), such as the general GPD and D-GPD estimators. These are also represented in Fig. 4 (c). Clearly, the paths of these two estimates follow a similar trajectory which is very different from that of the Hill plot. They point towards substantially lower estimates of ξ, and even though the estimates are overall larger than 0, the validity of the heavy tail assumption $\xi > 0$ is not obvious for this dataset. Fig. 4 (d) further supports this observation: in the (somewhat) stable region around the threshold $u = 10$, the 90% confidence interval produced through maximum likelihood theory contains the value 0. Our conclusion from the analysis of this dataset is that the distribution of the number of secondary cases is either fat-tailed but with a low tail index, or perhaps even light-tailed. As a consequence, our finding is qualitatively different from that of [3], since we do not obtain ξ estimates similar to those found by merging Datasets S1 and S2.

The second South Korean contact-tracing dataset comprises $n = 33{,}903$ SARS-CoV-2 numbers of secondary cases Z_i (Database S5) detected between 25th July 2021 and 15th August 2021. It was initially explored in [5], where it was highlighted that the Delta variant accounted for the majority of those cases. We therefore inves-

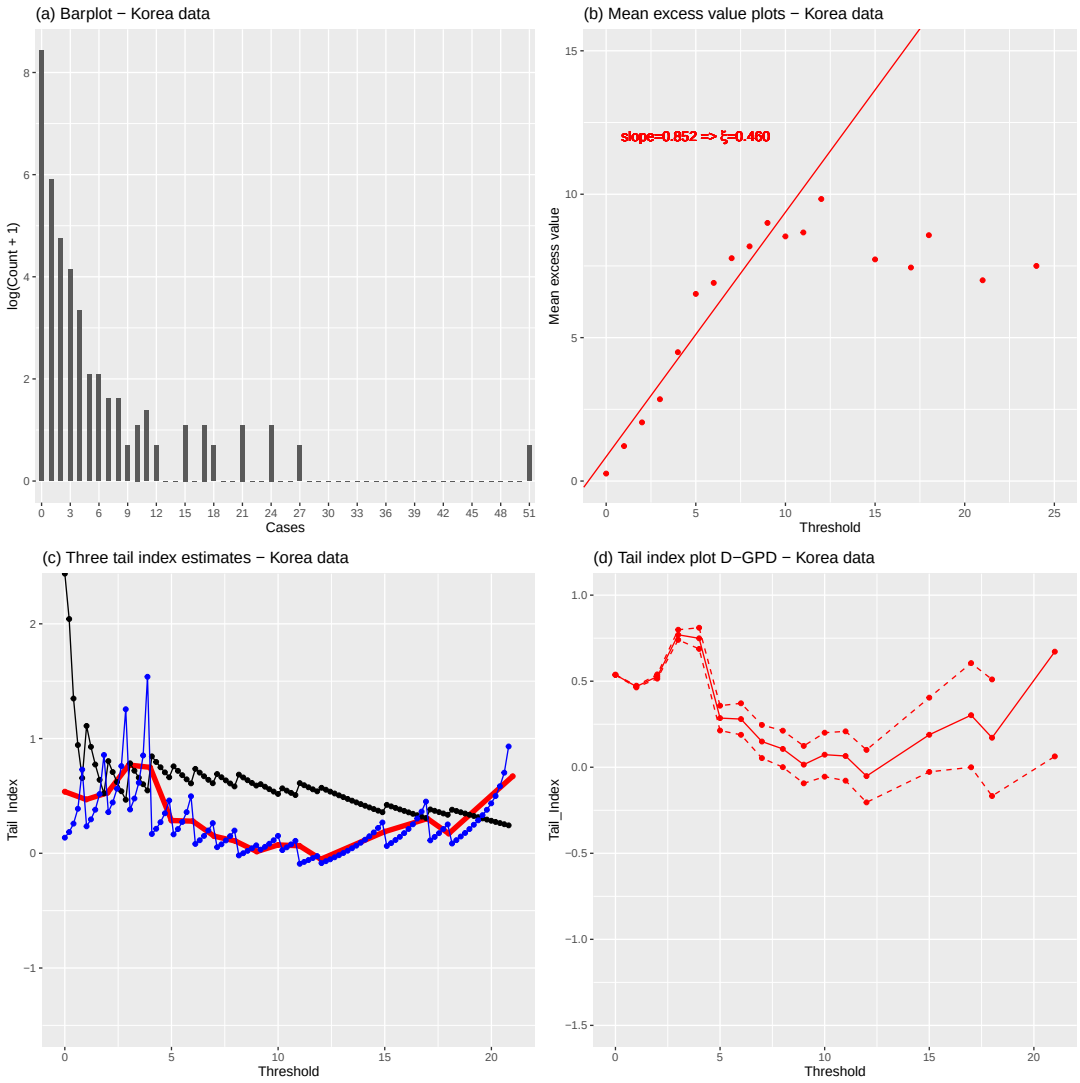

Figure 4: Secondary case data (Database S4) for SARS-CoV-2 from South Korea (first half of 2020). (a) Bar plot of the $\log(Z_i + 1)$ (n = 5,165). (b) Mean excess plots of secondary cases. (c) Hill (solid black), continuous GPD maximum likelihood (solid blue) and discrete GPD maximum likelihood (solid bold red) estimates of ξ. (d) Discrete GPD maximum likelihood estimates of ξ (solid red) and their associated 90% confidence intervals (dashed red).

tigate this dataset to ascertain whether the tail behaviour of SSEs is substantially different for the Delta variant. The data is presented in Table 3 below. The results we obtain for this dataset are displayed in Fig. 5. The barplot of the data in Fig. 5 (a) again backs the assumption of a heavy tail, but here, the mean excess plot in Fig. 5 (b) suggests a more convincing linearly increasing fit to the mean excess function with a slope of around 0.3. The Hill estimator and both continuous and discrete GPD maximum likelihood estimators, represented in Fig. 5 (c), appear to support the fat tail assumption of the offspring distribution which is mainly dominated here by the Delta variant. Once again, the D-GPD estimate has a much smoother and more stable sample path, with a stable zone over $u \in [1, 10]$ indicating a point estimate of around 0.21. The 90% confidence interval of the D-GPD estimate over that region, provided in Fig. 5 (d), does not contain 0 and offers further justification of the assumption that the offspring distribution is heavy-tailed in this dataset, in contrast to the 2020 South Korea data where the validity of this conclusion is much less clear.

Z	0	1	2	3	4	5	6	7	8	9	10
Count	29,193	2,154	1,121	594	332	207	113	53	53	21	21

Z	11	12	13	14	15	16	17	18	19	21	22	24	32
Count	6	8	5	3	3	2	2	3	3	1	2	2	1

Table 3: Secondary case data (Database S5) for SARS-CoV-2 collected in South Korea from 25th July 2021 to 15th August 2021.

Analysis of cluster size data. We broaden our analysis by examining whether SARS-CoV-2 cluster sizes are fat-tailed. We consider a database of 15 samples of cluster sizes recorded in 11 countries and 4 US states. We define a cluster as a local outbreak involving a minimum of two cases, including confirmed close contacts with epidemiological linkage over a limited period of time. The number of reported clusters per country or state varies from 29 (France) to 4,769 (Colorado, USA). The database is constructed from government reports [6, 7, 8, 9] (Database S6) and media sources [10] (Database S7). The median cluster sizes were 5 (Database S6) and 33 (Database S7), and the largest clusters had sizes 1,761 (Database S6, in a Colorado prison) and 7,000 (Database S7, in an Italian football stadium). We denote by Y_i the number of SARS-CoV-2 cases in cluster i. The ξ estimates from each sample of cluster sizes allow to infer the risk category of the corresponding country/state in terms of local community transmission.

Figs. 6 and 7 display the D-GPD maximum likelihood estimates of ξ as functions of the cluster size u. A common practice for selecting a suitable pointwise estimate of ξ is to pick out a sufficiently high threshold u corresponding to a stable region of the plot [15], as indicated by the vertical dashed lines in Figs. 6 and 7. The final selected estimates are reported in Table 4, where 13 out of the 15 countries or states appear to have fat-tailed cluster size distributions (confirmed at the 90% confidence level except for China). The analysis for California and UK & Ireland was inconclusive. For the California dataset, this is possibly due to a strong degree of heterogeneity (see the histogram in the bottom left panel of Fig. 7). A stratified study of the

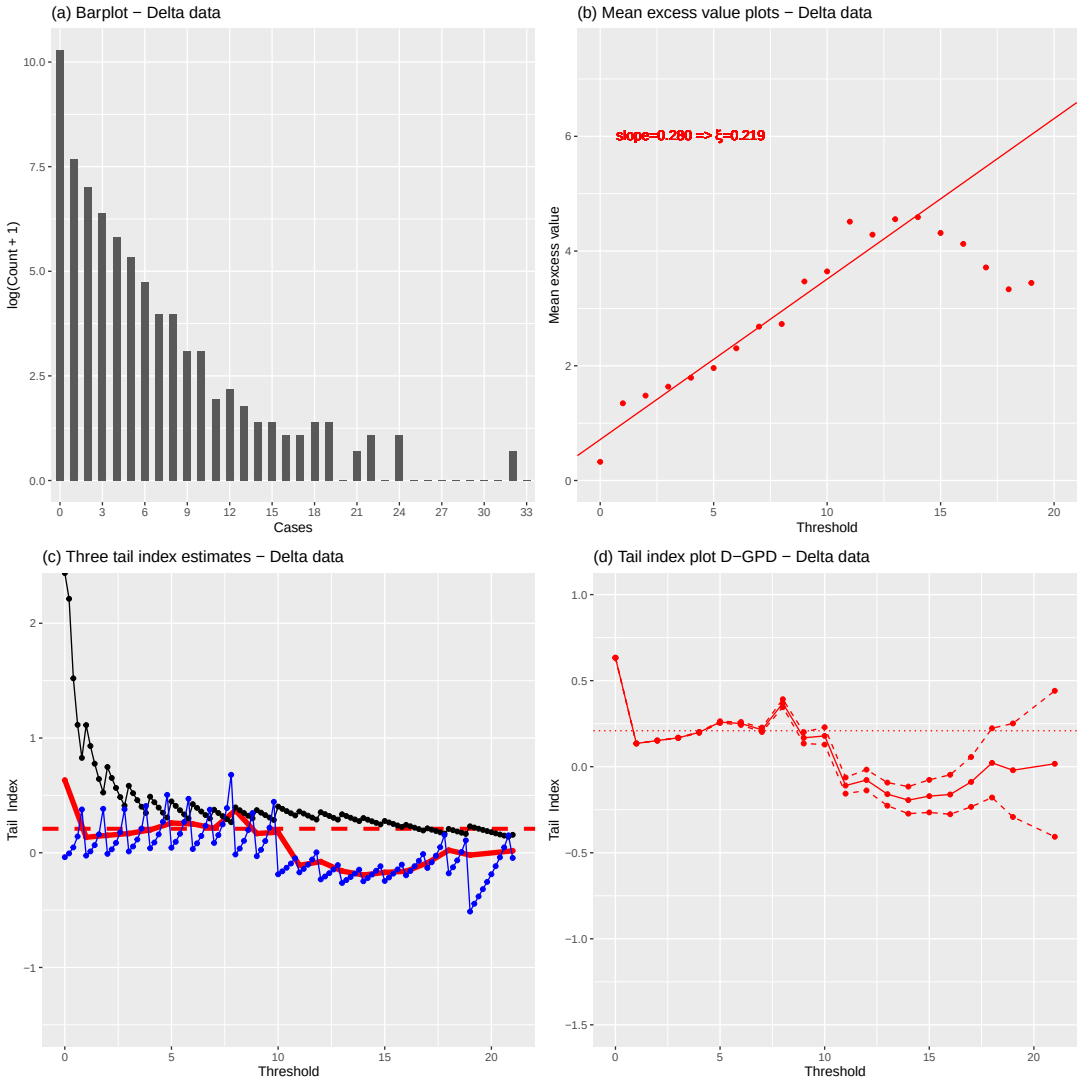

Figure 5: Secondary case data (Database S5) for SARS-CoV-2 from South Korea (July-August 2021). (a) Bar plot of the $\log(Z_i + 1)$ ($n = 33{,}903$). (b) Mean excess plots of secondary cases. (c) Hill (solid black), continuous GPD maximum likelihood (solid blue) and discrete GPD maximum likelihood (solid bold red) estimates of ξ. (d) Discrete GPD maximum likelihood estimates of ξ (solid red) and their associated 90% confidence intervals (dashed red). In panels (c) and (d), the averaged discrete GPD estimate $\hat{\xi} = 0.209$ over the stable region $u \in [1, 10]$ is indicated with the horizontal red line.

Californian data might be more conclusive. For the UK & Ireland dataset, the fact that the sample is so small (30 clusters) in two countries with highly developed healthcare and contact tracing systems is suspicious and may suggest reporting issues.

Using the D-GPD model, one can gain further insight into large cluster sizes by providing extrapolated estimates of extreme percentiles q_α potentially beyond the sample maximum, through the estimate $\hat{q}_\alpha$ described in the Methods section. Estimated 95th and 99th percentiles are given in Table 4. One may also match the estimated percentiles with actual observations to get a sense of what would constitute a conducive environment for the formation of large SARS-CoV-2 clusters. For example, the estimated 95th percentile of 120 cases in Kerala is close to two clusters of 113 cases (nursing home) and 132 cases (local transmission) already observed in Kerala. Likewise, the estimate $\hat{q}_{0.95} = 272$ cases in Canada is fairly close to a cluster of 324 cases in Canadian nursing homes. In Oregon, the estimated 99th percentile $\hat{q}_{0.99} = 124$ cases is in the vicinity of a cluster of 134 cases in a care home setting. In Colorado, the estimate $\hat{q}_{0.99} = 140$ cases is close to a cluster of 134 cases in a nursing home. All of these clusters bar one (the local transmission cluster in Kerala) correspond to indoor environments where social distancing is difficult to practice.

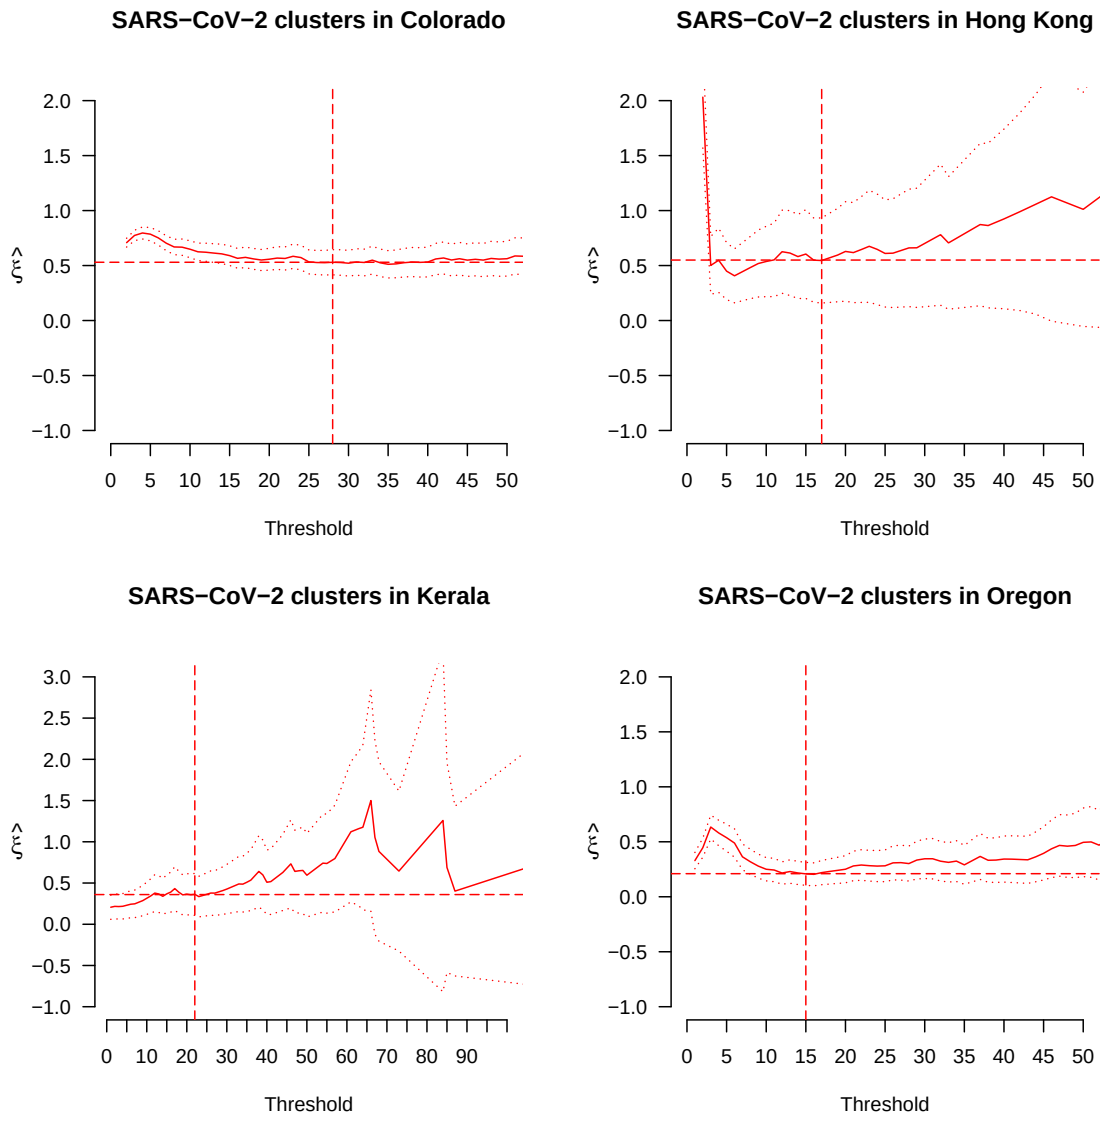

Figure 6: Analysis of cluster cases, for the four countries/states where the source is official data (Database S6). Plots of discrete GPD maximum likelihood estimates of ξ (solid lines), along with their 90% confidence intervals (dotted lines) and the final selected estimates (horizontal dashed lines) and thresholds (vertical dashed lines).

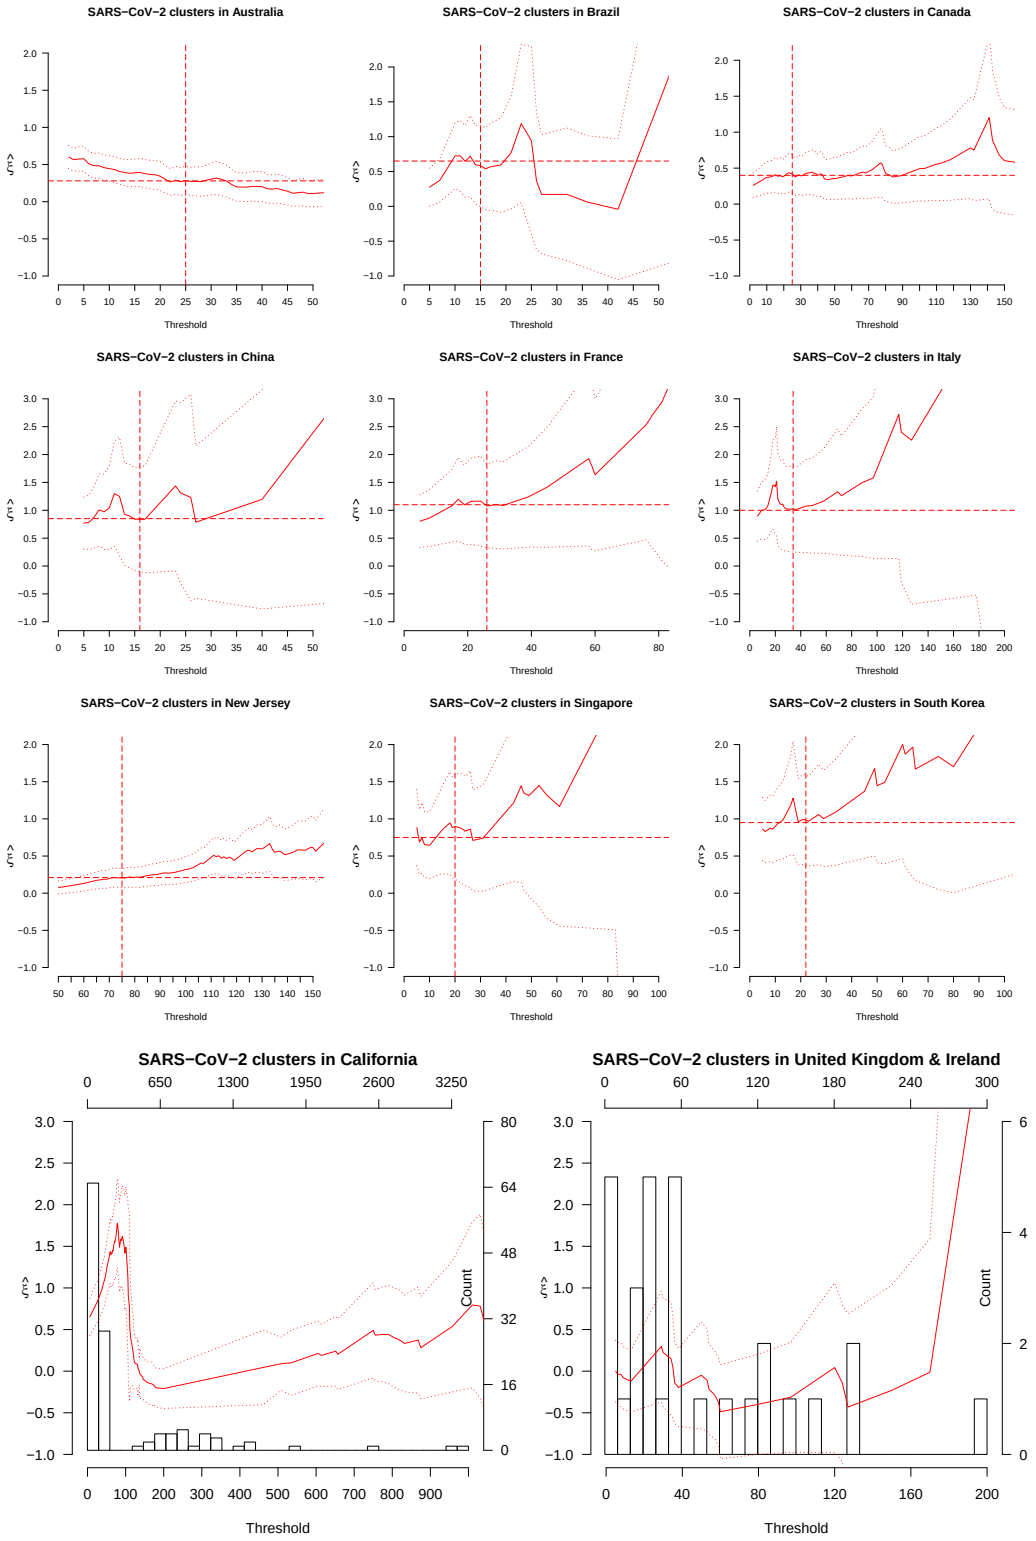

Figure 7: Analysis of cluster cases as in Fig. 6, with the results obtained from the data whose sources were the media (Database S7). The top 9 plots refer to those countries and states for which the extreme value analysis was conclusive. The bottom 2 plots refer to those for which the extreme value analysis was inconclusive.

Database S6

Location	n	$\hat{\xi}$ [90% CI]	u (n_u)	$\hat{q}_{0.95}$	$\hat{q}_{0.99}$	Max. Y_i (Setting)
Colorado, USA	4,769	0.53 [0.41, 0.64]	27 (474)	48	140	1,761 (Prison)
Hong Kong	54	0.55 [0.16, 0.93]	17 (34)	119	310	732 (Dancing)
Kerala, India	113	0.36 [0.11, 0.62]	22 (60)	120	255	580 (Unknown)
Oregon, USA	795	0.21 [0.10, 0.31]	15 (254)	64	124	639 (Prison)

Database S7

Location	n	$\hat{\xi}$ [90% CI]	u (n_u)	$\hat{q}_{0.95}$	$\hat{q}_{0.99}$	Max. Y_i (Setting)
Australia	355	0.28 [0.09, 0.48]	25 (145)	157	326	662 (Cruise ship)
Brazil	42	0.58 [0.00, 1.16]	15 (22)	82	220	191 (Hospital)
Canada	100	0.42 [0.15, 0.69]	25 (74)	272	624	1,500 (Meat processing plant)
China	34	0.84 [−0.12, 1.80]	16 (10)	99	401	368 (Market)
France	29	1.08 [0.32, 1.83]	26 (17)	443	2,530	2,500 (Religious gathering)
Italy	41	1.02 [0.25, 1.79]	34 (15)	378	2,013	7,000 (Stadium)
New Jersey, USA	183	0.20 [0.08, 0.33]	75 (157)	299	496	1,042 (Prison)
Singapore	45	0.90 [0.19, 1.61]	20 (21)	156	661	797 (Worker housing)
South Korea	45	0.98 [0.37, 1.59]	22 (24)	324	1,616	5,016 (Religious gathering)

Table 4: Final results for SARS-CoV-2 cluster sizes by country (first column), the corresponding sample size n (second column), D-GPD maximum likelihood ξ estimate and 90% confidence interval (third column), selected cluster size threshold u and associated number n_u of exceedance values $Y_i - u$ given $Y_i \geq u$ upon which the ξ estimate is calculated (fourth column), D-GPD maximum likelihood 95% and 99% percentile estimates of cluster size (fifth and sixth columns), and the sample maximum (last column). The top table corresponds to data from official sources (Database S6), and the bottom table to data from media sources (Database S7). The results reported in the latter table only concern the 9 countries and states for which the extreme value analysis was conclusive.

Discussion

In summary, we have investigated four datasets of secondary case numbers Z_i for SARS-CoV-2 as a way to estimate and infer the extreme value index of the related underlying offspring distribution. Motivated by the highly discrete nature of such data, we used the Discrete GPD (D-GPD) maximum likelihood estimation method which produces smoother and more stable plots of the associated D-GPD estimator than the classical continuous GPD and Hill estimators. We first provided evidence that the small SSE dataset (Dataset S2) compiled by [3] during the early phase of the COVID-19 pandemic was fat-tailed, thus confirming their findings, although we show in various ways that this dataset should not be pooled with their 15 SSEs associated with SARS-CoV (Dataset S1), since they correspond to substantially different distributions. On the other hand, as accurate extreme value inference requires a large sample size in general, we also analysed an Indian secondary case dataset of size 88,527 collected in 2020 (Database S3), which contains a very large number of tied observations. The D-GPD estimate of the tail index is around 0.24, which is in full agreement with the estimate of around 0.25 found by revisiting the small SSE dataset of size 45 from [3]. The distribution of SARS-CoV-2 SSEs therefore appears to have at least a finite third moment, whereas that of SARS-CoV SSEs is found to have a much heavier upper tail with infinite variance and therefore stronger superspreading effect. In an effort to account for the quality of implemented control programmes as well as the nature of the variant under study, we used two extra South Korean contact-tracing datasets. For the first dataset (Database S4), collected in the first half of 2020 and used in [3], we cannot disprove that the distribution of the number of secondary cases is light-tailed. By contrast, for the second South Korean dataset (Database S5) collected during the summer of 2021, in which the majority of cases correspond to the Delta variant of SARS-CoV-2 [5], we obtained a D-GPD estimate, $\hat{\xi} \approx 0.21$ clearly suggesting a heavier upper tail for the Delta variant and therefore more pronounced superspreading potential in South Korea relative to the first half of 2020.

We broaden our analysis by providing evidence that SARS-CoV-2 cluster sizes are typically fat-tailed, based on 15 samples from 11 countries and 4 US states. We infer the risk exposure and risk category of each country and state by making use of D-GPD maximum likelihood estimates of both the extreme value index and extreme percentiles, along with their associated confidence intervals. For the sake of simplicity, we used a straightforward threshold selection rule, which is to spot a stability region in the estimates (as a function of the threshold value) and choose an estimate whose value is representative of those reached in this region. This practice, colloquially known as "eyeballing", is standard in applied extreme value analysis: see for example the discussion in p.77 of Chapter 4 in [24]. It applies reasonably well to the D-GPD sample paths, because they are overall much smoother and more stable than the standard Hill and GPD maximum likelihood sample paths, which are not designed to handle the discreteness of the data. The development of more elaborate statistical techniques for the choice of threshold in discrete GPD maximum likelihood estimation, such as methods based on asymptotic MSE minimisation or the bootstrap in the spirit of the approaches outlined in Section 5.4 of [25] for Hill

estimation, is an open question which is beyond the scope of this paper.

A limitation of our study lies in the quality of the data, as it is not obvious whether all SSEs or clusters over a given time period were available, or whether cluster sizes were correctly recorded. To check robustness against missing data, we have reproduced part of our analysis of cluster data by removing 10% of observations at random in each sample containing at least 100 data points, and replicating this experiment 10,000 times. Robustness against poor recording was checked by multiplying each observation Y_i by an independent normal variate W_i having mean $\mu = 1$ and standard deviation $\sigma = 0.05$, and then reproducing our analysis of cluster data on the $Y_i' = W_i Y_i$, this experiment being again replicated 10,000 times. There is indeed some variation in the resulting estimates of ξ (Figs. 8 and 9), but this does not affect our conclusion on the fat-tailed behaviour of the data, except in rare situations when almost all the large values in the data go missing. This highlights the importance of accurate data reporting as a prerequisite to such analyses.

It should be noted that, in classical epidemiological models, accurate estimation of the basic reproduction number R_0 is of crucial importance as it informs the extent of restrictions on social interactions and other control measures that should be imposed to terminate the spread of an epidemic. The range of R_0 for SARS-CoV-2 has been revised in [26] to 4.7-11.4, which is considerably higher than most early estimates. This might explain why moderate restrictions that were implemented in some nations, e.g. France, Italy, Spain, the UK, Australia and New Zealand, turned out to be insufficient and replaced by nationwide or statewide lockdowns and/or border closures. It should be clear that our results are, by construction, robust to misspecified estimates of the expected number of secondary cases R_0 since they solely rely on extreme values of numbers of secondary cases.

Our approach can be viewed as a proof of concept that transmission data from a respiratory disease should not be pooled with data from a similar disease, since similar R_0 numbers or parameters of average transmission do not, in general, induce similar parameters of large community transmission. As such, preparing proactive control measures actually requires a fine assessment of how unequal the distributions of SSEs associated with different SARS-CoV-2 variants are. [27] conclude that the reproductive number of the Delta variant is far higher than that of the historical SARS-CoV-2 virus. Similarly, [28] estimate that the effective reproduction number of the Omicron variant is more than 3 times that of the Delta variant in Denmark. Our analysis of secondary case data did not, strictly speaking, allow one to conclude statistically that SSEs linked to the Delta variant had a different extreme value index from those linked to the original strains of SARS-CoV-2. However, in the contact-tracing data recorded in South Korea, we did find a heavy tail in the offspring distribution when the Delta variant made the majority of cases, as opposed to when it did not. This tentative finding of a heavier tail in the data linked to the Delta variant is coherent with the higher reproductive number of the Delta variant found in [27]. The question of estimating parameters of large community transmission for the Omicron variant remains open, as we could not find a dataset whose sample size would enable us to draw statistically principled conclusions about the tail behaviour of Omicron-related SSEs.

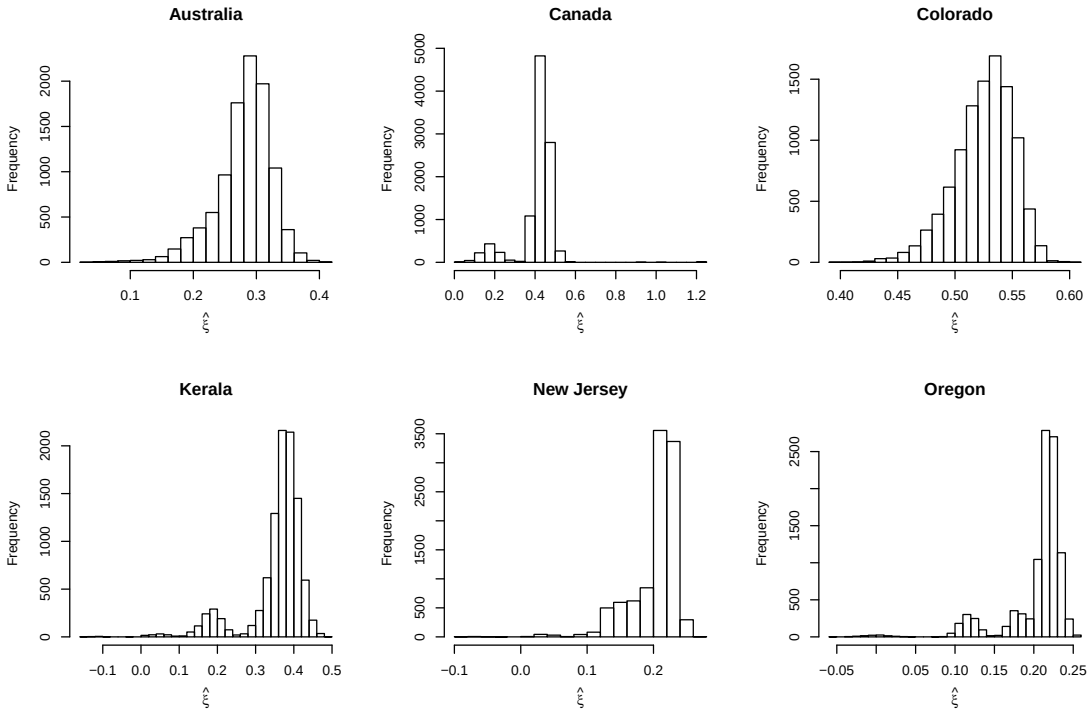

Figure 8: Robustness check (with respect to data omission) for the analysis of cluster cases (Databases S6 and S7). Histograms of the 10,000 estimates of ξ obtained by omitting at random 10% of the data. This was done only for the six samples containing at least 100 data points.

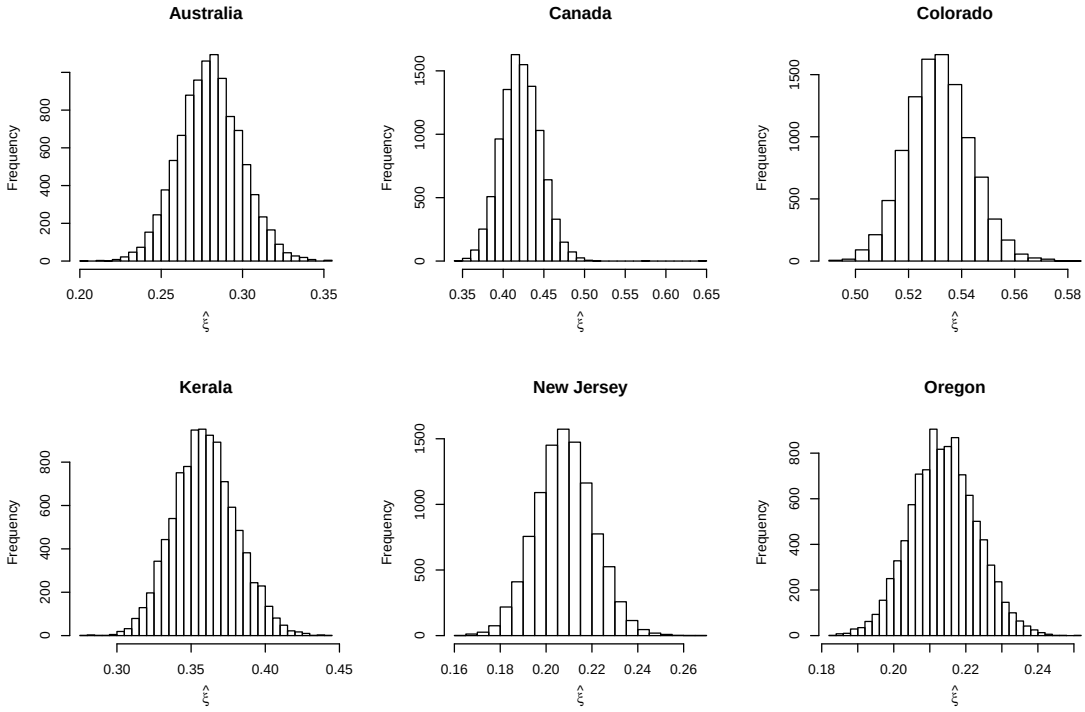

Figure 9: Robustness check (with respect to poor recording of the data) for the analysis of cluster cases (Databases S6 and S7). Histograms of the 10,000 estimates of ξ obtained by multiplying each data point by a random draw from the normal distribution with mean $\mu = 1$ and standard deviation $\sigma = 0.05$. This was done only for the six samples containing at least 100 data points.

Ethics. This article does not present research with ethical considerations.

Data accessibility. All datasets and the R code used for their statistical analysis are available at https://github.com/AntoineUC/SARS-CoV-2-codes. Datasets S1 and S2 can be found in the file `sse.R`. Database S3 can be found in the file `traceDatSaved.Rdata`. Databases S4 and S5 can be found in the files `sse_korea_2020.txt` and `sse_korea_2021_period_2.txt`, respectively. Databases S6 (apart from the Colorado data) and S7 can be found in the file `clusters.Rdata`, while the Colorado cluster size data can be found in `colorado.txt`.

Authors' contributions. A.U.C. undertook data curation and wrote the code for the statistical analysis and visualisation of the results. All three authors participated in the statistical analysis of the data and in drafting and revising the manuscript.

Competing interests. The authors declare no competing interests.

Funding. This research was supported by the French National Research Agency (grant numbers ANR-19-CE40-0013, ANR-17-EURE-0010). G. Stupfler also acknowledges support from an AXA Research Fund Award on 'Mitigating risk in the wake of the COVID-19 pandemic'.

Acknowledgements. The authors acknowledge an anonymous Associate Editor and two anonymous reviewers for their very helpful comments that led to a greatly improved version of this paper.

References

[1] J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, W. M. Getz, Superspreading and the effect of individual variation on disease emergence. *Nature* **438**, 355-359 (2005). (doi:10.1038/nature04153)

[2] D. C. Adam, P. Wu, J. Y. Wong, E. H. Y. Lau, T. K. Tsang, S. Cauchemez, G. M. Leung, B. J. Cowling, Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong. *Nat. Med.* **26**, 1714-1719 (2020). (doi:10.1038/s41591-020-1092-0)

[3] F. Wong, J. J. Collins, Evidence that coronavirus superspreading is fat-tailed. *PNAS* **117**, 29416-29418 (2020). (doi:10.1073/pnas.2018490117)

[4] R. Laxminarayan, B. Wahl, S. R. Dudala, K. Gopal, C. Mohan B, S. Neelima, K. S. Jawahar Reddy, J. Radhakrishnan, J. A. Lewnard, Epidemiology and transmission dynamics of COVID-19 in two Indian states. *Science* **370**, 691-697 (2020). (doi:10.1126/science.abd7672)

[5] S. Ryu, D. Kim, J.-S. Lim, S. T. Ali, B. J. Cowling, Serial interval and transmission dynamics during SARS-CoV-2 Delta variant predominance, South Korea. *Emerg. Infect. Dis.* **28**, 407-410 (2022). (doi:10.3201/eid2802.211774)

[6] State of Colorado, https://covid19.colorado.gov/covid19-outbreak-data, last updated on 2 June 2021 (resolved outbreaks only). Accessed 27th September 2021.

[7] Government of Hong Kong, https://www.chp.gov.hk/files/pdf/local_situation_covid19_en.pdf, last updated on 6 September 2021. Accessed 6th September 2021.

[8] Government of Kerala, https://covid19jagratha.kerala.nic.in/home/clusterList. Accessed 21st July 2021.

[9] State of Oregon, https://www.oregon.gov/oha/covid19/Documents/DataReports/Weekly-Outbreak-COVID-19-Report-2021-08-25-FINAL.pdf?utm_medium=email&utm_source=govdelivery. Accessed 25 August 2021.

[10] K. Swinkels, SARS-CoV-2 Superspreading Events Database, https://kmswinkels.medium.com/covid-19-superspreading-events-database-4c0a7aa2342b. Accessed 21st July 2021.

[11] T. Shimura, Discretization of distributions in the maximum domain of attraction. *Extremes*, **15**, 299-317 (2012). (doi:10.1007/s10687-011-0137-7)

[12] F. Prieto, E. Gómez-Déniz, J. M. Sarabia, Modelling road accident blackspots data with the discrete generalized Pareto distribution. *Accident Analysis & Prevention* **49**, 71:38 (2014). (doi:10.1016/j.aap.2014.05.005)

[13] A. Hitz, R. Davis, G. Samorodnitsky, Discrete Extremes. arXiv [Preprint] (2017). https://arxiv.org/abs/1707.05033 (doi:10.48550/arXiv.1707.05033)

[14] S. Ranjbar, E. Cantoni, V. Chavez-Demoulin, G. Marra, R. Radice, K. Jaton, Modelling the extremes of seasonal viruses and hospital congestion: the example of flu in a Swiss hospital. *J. Roy. Stat. Ser. C*, to appear (2022). (doi:10.1111/rssc.12559)

[15] L. de Haan, A. Ferreira, *Extreme Value Theory: An Introduction*, Springer-Verlag, New York (2006). (doi:10.1007/0-387-34471-3)

[16] B. M. Hill, A simple general approach to inference about the tail of a distribution. *Ann. Statist.* **3**, 1163-1174 (1975). (doi:10.1214/aos/1176343247)

[17] J. Beirlant, Y. Goegebeur, J. Segers, J. Teugels, *Statistics of Extremes: Theory and Applications*, John Wiley & Sons, Chichester (2004).

[18] S. Ghosh, S. Resnick, A discussion on mean excess plots. *Stoch. Proc. Appl.* **120**, 1492–1517 (2010). (doi:10.1016/j.spa.2010.04.002)

[19] S. Coles, *An Introduction to Statistical Modeling of Extreme Values*, Springer-Verlag, London (2004). (doi:10.1007/978-1-4471-3675-0)

[20] P. Cirillo, N. N. Taleb, Tail risk of contagious diseases. *Nat. Phys.* **16**, 606-613 (2020). (doi:10.1038/s41567-020-0921-x)

[21] C. Kremer, A. Torneri, S. Boesmans, H. Meuwissen, S. Verdonschot, K. Vanden Driessche, C. L. Althaus, C. Faes, N. Hens, Quantifying superspreading for COVID-19 using Poisson mixture distributions. *Sci. Rep.* **11**, 14107 (2021). (doi:10.1038/s41598-021-93578-x)

[22] N. Islam, Q. Bukhari, Y. Jameel, S. Shabnam, A. M. Erzurumluoglu, M. A. Siddique, J. M. Massaro, R. B. D'Agostino, COVID-19 and climatic factors: A global analysis. *Environmental Research* **193**, 110355 (2021). (doi:10.1016/j.envres.2020.110355)

[23] H. Hwang, J.-S. Lim, S.-A. Song, C. Achangwa, W. Sim, G. Kim, S. Ryu, Transmission dynamics of the Delta variant of SARS-CoV-2 infections in South Korea. *J. Infect. Dis.* **225**, 793–799 (2022). (doi:10.1093/infdis/jiab586)

[24] M. Jacob, C. Neves, D. Vukadinović Greetham, *Extreme Value Statistics*, in: Forecasting and Assessing Risk of Individual Electricity Peaks. Mathematics of Planet Earth, Springer, Cham (2020). (doi:10.1007/978-3-030-28669-9)

[25] M. I. Gomes, A. Guillou, Extreme value theory and statistics of univariate extremes: A review. *Int. Stat. Rev.* **83**, 263-292 (2015). (doi:10.1111/insr.12058)

[26] M. Kochańczyk, F. Grabowski, T. Lipniacki, Super-spreading events initiated the exponential growth phase of COVID-19 with $\mathcal{R}_0$ higher than initially estimated. *Royal Society Open Science* **7**, 200786 (2020). (doi:10.1098/rsos.200786)

[27] Y. Liu, J. Rocklöv, The reproductive number of the Delta variant of SARS-CoV-2 is far higher compared to the ancestral SARS-CoV-2 virus. *J. Travel Med.* **28**, taab124 (2021). (doi:10.1093/jtm/taab124)

[28] K. Ito, C. Piantham, H. Nishiura, Relative instantaneous reproduction number of Omicron SARS-CoV-2 variant with respect to the Delta variant in Denmark. *J. Med. Virol.* **94**, 2265-2268 (2022). (doi:10.1002/jmv.27560)