

ESTIMATING THE HISTORICAL AND FUTURE PROBABILITIES OF LARGE TERRORIST EVENTS¹

BY AARON CLAUSET AND RYAN WOODARD

University of Colorado, Boulder and ETH Zürich

Quantities with right-skewed distributions are ubiquitous in complex social systems, including political conflict, economics and social networks, and these systems sometimes produce extremely large events. For instance, the 9/11 terrorist events produced nearly 3000 fatalities, nearly six times more than the next largest event. But, was this enormous loss of life statistically unlikely given modern terrorism's historical record? Accurately estimating the probability of such an event is complicated by the large fluctuations in the empirical distribution's upper tail. We present a generic statistical algorithm for making such estimates, which combines semi-parametric models of tail behavior and a nonparametric bootstrap. Applied to a global database of terrorist events, we estimate the worldwide historical probability of observing at least one 9/11-sized or larger event since 1968 to be 11–35%. These results are robust to conditioning on global variations in economic development, domestic versus international events, the type of weapon used and a truncated history that stops at 1998. We then use this procedure to make a data-driven statistical forecast of at least one similar event over the next decade.

1. Introduction. The September 11th terrorist attacks were the largest such events in modern history, killing nearly 3000 people [MIPT (2008), START (2011)]. Given their severity, should these attacks be considered statistically unlikely or even outliers? What is the likelihood of another September 11th-sized or larger terrorist event, worldwide, over the next decade?

Accurate answers to such questions would shed new light both on the global trends and risks of terrorism and on the global social and political processes that generate these rare events [McMorrow (2009), Sornette (2009), Sornette and Ouillon (2012)], which depends in part on determining whether the same processes generate both rare, large events and smaller, more common events. Insights would also provide objective guidance for our long-term expectations in planning, response and insurance efforts [de Haan and Ferreira (2006), Reiss and Thomas (2007)], and for estimating the likelihood of even larger events, including mass-casualty chemical, biological, radioactive or nuclear (CBRN) events [Cameron (2000), Drennan (2007)].

Received April 2012; revised July 2012.

¹Discussed in [10.1214/13-AOAS614A](#), [10.1214/13-AOAS614B](#), [10.1214/13-AOAS614C](#), [10.1214/13-AOAS614D](#), [10.1214/13-AOAS614E](#), [10.1214/13-AOAS614F](#); rejoinder at [10.1214/13-AOAS614R](#).

Key words and phrases. Rare events, forecasting, historical probability, terrorism.

The rarity of events like 9/11 poses two technical problems: (i) we typically lack quantitative mechanism-based models with demonstrated predictive power at the global scale (which is particularly problematic for CBRN events) and (ii) the global historical record contains few large events from which to estimate mechanism-agnostic statistical models of large events alone. That is, the rarity of big events implies large fluctuations in the distribution's upper tail, precisely where we wish to have the most accuracy. These fluctuations can lead to poor out-of-sample predictive power in conflict [see Beck, King and Zeng (2000), Bueno de Mesquita (2003, 2011), King and Zeng (2001), Rustad et al. (2011), Ward, Greenhill and Bakke (2010)] and can complicate both selecting the correct model of the tail's structure and accurately estimating its parameters [Clauset, Shalizi and Newman (2009)]. Misspecification can lead to severe underestimates of the true probability of large events, for example, in classical financial risk models [Farmer and Lillo (2004), Financial Crisis Inquiry Commission (2011)].

Little research on terrorism has focused on directly modeling the number of deaths ("severity")² in individual terrorist events [McMorrow (2009)]. When deaths are considered, they are typically aggregated and used as a covariate to understand other aspects of terrorism, for example, trends over time [Enders and Sandler (2000, 2002)], the when, where, what, how and why of the resort to terrorism [Brown, Dalton and Hoyle (2004), Enders and Sandler (2006), Valenzuela et al. (2010)], differences between organizations [Asal and Rethemeyer (2008)], or the incident rates or outcomes of events [Enders and Sandler (2000), Enders, Sandler and Gaibullov (2011)]. Such efforts have used time series analysis [Enders and Sandler (2000, 2002), Enders, Sandler and Gaibullov (2011)], qualitative models or human expertise of specific scenarios, actors, targets or attacks [Wulf, Haines and Longstaff (2003)] or quantitative models based on factor analysis [Li (2005), Pape (2003)], social networks [Desmarais and Cranmer (2011), Sageman (2004)] or formal adversarial interactions [Enders and Sandler (2006), Kardes and Hall (2005), Major (1993)]. Most of this work focuses on modeling central tendencies, treats large events like 9/11 as outliers, and says little about their quantitative probability [Clauset, Young and Gleditsch (2007)] or their long-term hazard.

Here, we describe a statistical algorithm for estimating the probability of large events in complex social systems in general, and in global terrorism in particular. Making only broad-scale and long-term probabilistic estimates, our approach is related to techniques used in seismology, forestry, hydrology and natural disaster insurance to estimate the probabilities of individual rare catastrophic events [Breiman, Stone and Kooperberg (1990), de Haan and Ferreira (2006), Gumbel (1941), Gutenberg and Richter (1944), Reed and McKelvey (2002), Reiss and Thomas (2007)]. Our approach combines maximum-likelihood methods, multiple models of the distribution's tail and computational techniques to account for

²Other notions of event "size" or severity, which we do not explore here, might be the economic cost, number injured, political impact, etc. To the extent that such notions may be quantitatively measured, our algorithm could also be applied to them.

both parameter and model uncertainty. It provides a quantitative estimate of the probability, with uncertainty, of a large event. The algorithm also naturally generalizes to include certain event covariates, which can shed additional light on the probability of large events of different types.

Using this algorithm to analyze a database of 13,274 deadly terrorist events worldwide from 1968–2007, we estimate the global historical probability of at least one 9/11-sized or larger terrorist event over this period to be roughly 11–35%. Furthermore, we find the nontrivial magnitude of this historical probability to be highly robust, a direct consequence of the highly right-skewed or “heavy-tailed” structure of event sizes [Clauset, Young and Gleditsch (2007)]. Thus, an event the size or severity of the September 11th terrorist attacks, compared to the global historical record, should not be considered a statistical outlier or even statistically unlikely. Using three potential scenarios for the evolution of global terrorism over the next decade, we then estimate the worldwide future probability of a similarly large event as being not significantly different from the historical level. We close by discussing the implications for forecasting large terrorist events in particular and for complex social systems in general.

2. Estimating the probability of a large event. The problem of estimating the probability of some observed large event is a kind of tail-fitting problem, in which we estimate parameters for a distributional model using only the several largest observations. This task is distinct from estimating the distribution of maxima within a sample [de Haan and Ferreira (2006), Reiss and Thomas (2007)], and is more closely related to the peaks-over-threshold literature in hydrology, seismology, forestry, finance and insurance [Adler, Feldman and Tuqqu (1998), Breiman, Stone and Kooperberg (1990), de Haan and Ferreira (2006), Gumbel (1941), Gutenberg and Richter (1944), Reed and McKelvey (2002), Reiss and Thomas (2007), Resnick (2007)]. Here, we aim specifically to deal with several sources of uncertainty in this task: uncertainty in the location of the tail, uncertainty in the tail’s true structure, and uncertainty in the model parameters.

Our approach is based on three key insights. First, because we are interested only in rare large events, we need only model the structure of the distribution’s right or upper tail, which governs their frequency. This replaces the difficult problem of modeling both the distribution’s body and tail [de Haan and Ferreira (2006), Reiss and Thomas (2007), Resnick (2007)] with the less difficult problem of identifying a value x_{\min} above which a model of the tail alone fits well.³ That is, choose some x_{\min} and a tail model $\Pr(x|\theta, x_{\min})$ defined on $x \in [x_{\min}, \infty)$. We will revisit the problem of choosing x_{\min} below.

Second, in complex social systems, the correct tail model is typically unknown and a poor choice may lead to severe misestimates of the true probability of a large event. We control for this model uncertainty by considering multiple tail

³The notation x_{\min} should not be confused with the first order statistic, $x_{(1)} = \min_i x_i$.

models. Given these models and a common choice of x_{\min} , we use a likelihood ratio test to identify and discard the statistically implausible ones [Clauset, Shalizi and Newman (2009)]. In principle, the remaining models could be averaged to produce a single estimate with confidence intervals [Claeskens and Hjort (2008)], for example, to aid decision makers. We return to this point in more detail below.

Finally, large fluctuations in the distribution's upper tail occur precisely where we wish to have the most accuracy, leading to parameter uncertainty. Using a non-parametric bootstrap [Efron and Tibshirani (1993)] to simulate the generative process of event sizes, we incorporate the empirical data's inherent variability into the estimated parameters, weight models by their likelihood under the bootstrap distribution and construct extreme value confidence intervals [Breiman, Stone and Kooperberg (1990)].

This combination of techniques provides a statistically principled and data-driven solution for estimating the probability of observing rare events in empirical data with unknown tail structure. If such an event is observed, the algorithm provides a measure of whether its occurrence was in fact unlikely, given the overall structure of the distribution's tail. For instance, if the estimated probability is negligible (say, $p < 0.01$), the event may be judged statistically unlikely. When several tail models are plausible and agree that the probability is away from $p = 0$, the event can be judged to be statistically likely, despite the remaining uncertainty in the tail's structure.

2.1. The method. Our goal is to estimate the probability that we would observe at least ℓ "catastrophic" events of size x or greater in an empirical sample.⁴ In principle, any size x and any value ℓ may be chosen, but, in practice, we typically choose x as the largest (and thus rarest) event in the empirical data and set $\ell = 1$. To ensure that our estimate is meaningful from a historical perspective, we remove the catastrophic event(s) from the empirical sample before applying the algorithm. Here we describe the method in terms of univariate distributions, but its generalization to certain covariates is straightforward (see Appendix C.3.3).

Let $\Pr(x|\theta, x_{\min})$ denote a particular tail model with parameters θ , let $\{x_i\}$ denote the n empirical event sizes (sans the catastrophic events), and let $Y = \{y_j\}$ be a bootstrap of these data (n samples drawn from $\{x_i\}$ with replacement). To begin, we assume a fixed x_{\min} , the smallest value for which the tail model holds, and later describe the generalization to variable x_{\min} .

The fraction of empirical events with values in the tail region is $p_{\text{tail}} = \#\{x_i \geq x_{\min}\}/n$, and in each bootstrap the number is a binomial random variable with

⁴Consider events to be generated by a kind of marked point process [Last and Brandt (1995)], where marks indicate either the event's severity or that it exceeded some threshold x . Although we assume the number of marks n to be fixed, this could be relaxed to incorporate additional uncertainty into the algorithm's output.

probability p_{tail} :

$$(2.1) \quad n_{\text{tail}} \sim \text{Binomial}(n, p_{\text{tail}}).$$

The maximum likelihood estimate $\hat{\theta}$ is a deterministic function of the portion of Y above x_{min} , which we denote $\theta(Y, x_{\text{min}})$.

Given that choice, the probability under the fitted model that not one of $n'_{\text{tail}} = 1 + n_{\text{tail}}$ events is at least as big as x is

$$(2.2) \quad F(x|\theta(Y, x_{\text{min}}))^{n'_{\text{tail}}} = \left(\int_{x_{\text{min}}}^x \text{Pr}(y|\hat{\alpha}, x_{\text{min}}) dy \right)^{n'_{\text{tail}}}.$$

Thus, $1 - F(x|\theta(Y, x_{\text{min}}))^{n'_{\text{tail}}}$ is the probability that at least one event is of catastrophic size. Because the bootstrap Y is itself a random variable, to derive the marginal probability of observing at least one catastrophic event, we must integrate the conditional probability over the domain of the bootstrap distribution:

$$(2.3) \quad \begin{aligned} p(n_{\text{tail}}, \theta) &= p(n_{\text{tail}}, Y) \\ &= \int dy_1 \cdots dy_{n_{\text{tail}}} (1 - F(x; \theta(Y, x_{\text{min}}))^{n'_{\text{tail}}}) \prod_{i=1}^{n_{\text{tail}}} r(y_i | n_{\text{tail}}). \end{aligned}$$

The trailing product series here is the probability of drawing the specific sequence of values $y_1, \dots, y_{n_{\text{tail}}}$ from the fixed bootstrap distribution r . Finally, the total probability p of at least one catastrophic event is given by a binomial sum over this equation.⁵

When the correct value x_{min} is not known, it must be estimated jointly with θ on each bootstrap. Maximum likelihood cannot be used for this task, because x_{min} truncates Y . Several principled methods for automatically choosing x_{min} exist, for example, Breiman, Stone and Kooperberg (1990), Clauset, Shalizi and Newman (2009), Clauset, Young and Gleditsch (2007), Danielsson et al. (2001), Dekkers and de Haan (1993), Drees and Kaufmann (1998), Hancock and Jones (2004). So long as the choice of x_{min} is also a deterministic function of Y , the above expression still holds. Variation in x_{min} across the bootstraps, however, leads to different numbers of observations n_{tail} in the tail region. The binomial probability p_{tail} is then itself a random variable determined by Y , and n_{tail} is a random variable drawn from a mixture of these binomial distributions.

Analytically completing the above calculation can be difficult, even for simple tail models, but it is straightforward to estimate numerically via Monte Carlo:

⁵We may calculate p in either of two ways: (i) we draw n_{tail} events from a tail model alone, or (ii) we draw n events from a conditional model, in which the per-event probability is $q(x) = \text{Pr}(X \geq x | X \geq x_{\text{min}}) \text{Pr}(X \geq x_{\text{min}}) = p_{\text{tail}}(1 - F(x|\theta, x_{\text{min}}))$. When the probability of a catastrophic event is small, these calculations yield equivalent results.

1. Given n empirical sizes, generate Y by drawing y_j , $j = 1, \dots, n$, uniformly at random, with replacement, from the observed $\{x_i\}$ (sans the ℓ catastrophic events).
2. Jointly estimate the tail model's parameters θ and x_{\min} on Y , and compute $n_{\text{tail}} = \#\{y_j \geq \hat{x}_{\min}\}$ (see Appendix A).
3. Set $\rho = 1 - F(x; \hat{\theta})^{\ell + n_{\text{tail}}}$, the probability of observing at least ℓ catastrophic events under this bootstrap model.

Averaging over the bootstraps yields the estimated probability $\hat{p} = \langle \rho \rangle$ of observing at least ℓ catastrophic-sized events. The convergence of \hat{p} is guaranteed so long as the number of bootstraps (step 1) tends to infinity [Efron and Tibshirani (1993)]. Confidence intervals on \hat{p} [Breiman, Stone and Kooperberg (1990), Efron and Tibshirani (1993)] may be constructed from the distribution of the ρ values. If the tail model's c.d.f. $F(x; \theta)$ in step 3 cannot be computed analytically, it can often be constructed numerically; failing that, ρ may always be estimated by sampling directly from the fitted model.

2.2. Model comparison and model averaging. In complex social systems, we typically do not know a priori which particular tail model is correct, and the algorithm described above will give no warning of a bad choice [but see Clauset, Shalizi and Newman (2009)]. This issue is partly mitigated by estimating x_{\min} , which allows us to focus our modeling efforts on the upper tail alone. But, without additional evidence of the model's statistical plausibility, the estimate \hat{p} should be treated as provisional.

Comparing the results from multiple tail models provides a test of robustness against model misspecification, for example, agreement across models that $\hat{p} > 0.01$ strengthens the conclusion that the event is not statistically unlikely. However, wide confidence intervals and disagreements on the precise probability of a large event reflect the inherent difficulty of identifying the correct tail structure.

To select reasonable models to compare, standard model comparison approaches may be used, for example, a fully Bayesian approach [Kass and Raftery (1995)], cross-validation [Stone (1974)] or minimum description length [Grünwald (2007)]. Here, we use a goodness-of-fit test to establish the plausibility of the power-law distribution [Clauset, Shalizi and Newman (2009)] and Vuong's likelihood ratio test [Clauset, Shalizi and Newman (2009), Vuong (1989)] to compare it with alternatives. This approach has the advantage that it can fail to choose one model over another if the difference in their likelihoods is statistically insignificant, given the data.

In some circumstances, we may wish to average the resulting models to produce a single estimate with confidence intervals, for example, to aid decision makers. However, averaging poses special risks and technical problems for estimating the probability of large events. For instance, traditional approaches to averaging can obscure the inherent uncertainty in the tail's structure and can produce spuriously

precise confidence intervals [Claeskens and Hjort (2008), Hjort and Claeskens (2003)]; a Bayesian approach would be inconsistent with our existing framework; and an appropriate frequentist framework is not currently available, although one may be possible using insights from Cesa-Bianchi, Conconi and Gentile (2004).

Thus, in our application below, we elect not to average and instead we present results for each model. Even without averaging, however, several valuable insights may be drawn.

2.3. Tests of the method's accuracy. To test the accuracy of our estimation algorithm, we examine its ability to recover the true probability of a rare event from synthetic data with known structure. To generate these synthetic data, we use the power-law distribution

$$(2.4) \quad \Pr(y) \propto y^{-\alpha},$$

where $\alpha > 1$ is the “scaling” parameter and $y \geq x_{\min} > 0$. When $\alpha < 2$, this distribution exhibits infinite variance and produces extreme fluctuations in the upper tail of finite-size samples. By defining a catastrophic event x to be the largest generated event within the n synthetic values, we make the test particularly challenging because the largest value exhibits the greatest fluctuations of all. Detailed results are given in Appendix B.

We find that despite the large fluctuations generated by the power-law distribution, the algorithm performs well: the mean absolute error $\langle |\hat{p} - p| \rangle$ is small even for samples with less than 100 events, and decays like $O(n^{-1/3})$. A small absolute deviation, however, may be an enormous relative deviation, for example, if the true probability tends to zero or one. Our algorithm does not make this type of error: the mean ratio of the estimated and true probabilities $\langle \hat{p}/p \rangle$ remains close to 1 and thus the estimate is close in relative terms, being only a few percent off for $n \gtrsim 100$ events.

3. Historical probability of 9/11. Having described our statistical approach, we now use it to estimate the historical probability of observing worldwide at least one 9/11-sized or larger terrorist event.

Global databases of terrorist events show that event severities (number of deaths) are highly right-skewed or “heavy tailed” [MIPT (2008), START (2011)]. We use the RAND-MIPT database [MIPT (2008)], which contains 13,274 deadly events worldwide from 1968–2007. The power law is a statistically plausible model of this distribution’s tail, with $\hat{\alpha} = 2.4 \pm 0.1$, for $x \geq \hat{x}_{\min} = 10$ [Clauset, Shalizi and Newman (2009), Clauset, Young and Gleditsch (2007)]. A goodness-of-fit test fails to reject this model of tail event severities ($p = 0.40 \pm 0.03$ via Monte Carlo [Clauset, Shalizi and Newman (2009)]), implying that the deviations between the power-law model and the empirical data are indistinguishable from sampling noise.

This fact gives us license to treat as i.i.d. random variables the severity of these events. This treatment does force a particular and uncommon theoretical perspective on terrorism, in which a single global “process” produces events, even if the actions of individual terrorists or terrorist organizations are primarily driven by local events. This perspective has much in common with statistical physics, in which particular population-level patterns emerge from a sea of individual interactions. We discuss limitations of this perspective in Section 5.

Past work shows that this apparent power-law pattern in global terrorism is remarkably robust. Although the estimated value of α varies somewhat with time [Clauset, Young and Gleditsch (2007)], the power-law pattern itself seems to persist over the 40-year period despite large changes in the international system. It also appears to be independent of the type of weapon (explosives, firearms, arson, knives, etc.), the emergence and increasing frequency of suicide attacks, the demise of many terrorist organizations, the economic development of the target country [Clauset, Young and Gleditsch (2007)] and organizational covariates like size (number of personnel), age and experience (total number of attacks) [Clauset and Gleditsch (2009)].

Comparing the power-law tail model against log-normal and stretched exponential (Weibull) distributions, via a likelihood ratio test, yields log-likelihood ratios of $\mathcal{R} = -0.278$ ($p = 0.78$) and 0.772 ($p = 0.44$), respectively [Clauset, Shalizi and Newman (2009)]. However, neither of these values is statistically significant, as indicated by the large p -values for a test against $\mathcal{R} = 0$. Thus, while the power-law model is plausible, so too are these alternatives. This ambiguity illustrates the difficulty of correctly identifying the tail’s structure and reinforces the need to use multiple tail models in estimating the likelihood of a rare event like 9/11. Furthermore, it implies that slight visual deviations in the empirical distribution’s upper tail (see Figure 1) should not be interpreted as support either for or against any of these models. In what follows, we consider estimates derived from all three.

To apply our algorithm to this problem, we must make several choices. For consistency with past work on the frequency of severe terrorist events [Clauset, Shalizi and Newman (2009), Clauset, Young and Gleditsch (2007)], we choose x_{\min} automatically by minimizing the Kolmogorov–Smirnov goodness-of-fit statistic between the tail model and the truncated empirical data.⁶ We use the discrete power-law distribution as our tail model (which implies x_{\min} is also discrete; see Appendix A) and compare its estimates to those made using log-normal and stretched exponential models. To avoid the problem of choosing an appropriate event count distribution, we keep the number of events n fixed.

⁶Clauset, Shalizi and Newman (2009) provide a thorough motivation of this strategy. Briefly, the KS statistic will be large either when x_{\min} is too small (including nonpower-law data in the power-law fit) or when x_{\min} is too large (when sample size is reduced and legitimately power-law data thrown out), but will be small between these two cases.

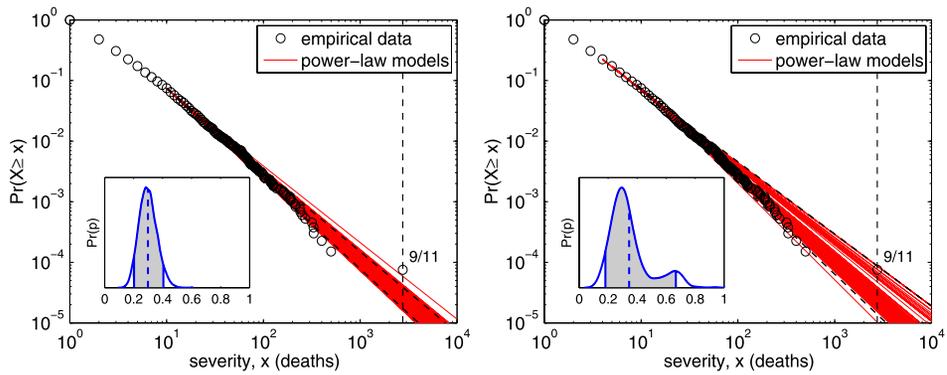


FIG. 1. Empirical severity distribution with 100 bootstrap power-law models for (a) fixed $x_{\min} = 10$ and (b) estimated x_{\min} . Overprinting illustrates the ensemble of estimated models (dashed lines show 90% CI on $\hat{\alpha}$) and the inherent uncertainty in the tail structure. Insets show the 90% confidence intervals for the estimated probability of observing at least one 9/11-sized event.

Finally, using the RAND-MIPT event data (other sources [START (2011)] yield similar results; see Appendix C.2), we define $x \geq 2749$ to be a “catastrophic” event—the reported size of the New York City 9/11 events.⁷ Removing this event from the empirical data leaves the largest event as the 14 August 2007 coordinated truck bombing in Sinjar, Iraq, which produced approximately 500 fatalities. To illustrate the robustness of our results, we consider estimates derived from fixed and variable x_{\min} and from our three tail models. We also analyze the impact of covariates like domestic versus international, the economic development of the target country and the type of weapon used.

3.1. *Uncertainty in the scaling parameter.* Let $x_{\min} = 10$ be fixed. Figure 1(a) shows 100 of the fitted bootstrap models, illustrating that by accounting for the uncertainty in α , we obtain an ensemble of tail models and thus an ensemble of probability estimates for a catastrophic-sized event. The bootstrap parameter distribution $\Pr(\hat{\alpha})$ has a mean $\langle \hat{\alpha} \rangle = 2.40$, which agrees with the maximum likelihood value $\hat{\alpha} = 2.4$ [Clauset, Shalizi and Newman (2009)].

To estimate the historical probability of 9/11, we use 10,000 bootstraps with x_{\min} fixed. Letting p denote the overall probability from the algorithm, we find $\hat{p} = 0.299$, with 90% confidence intervals of $[0.203, 0.405]$ [Figure 1(a) inset], or about a 30% chance over the 1968–2007 period.

An event that occurs with probability 0.299 over 40 years is not a certainty. However, for global terrorism, this value is uncomfortably large and implies that, given the historical record, the size of 9/11 should not be considered a statistical fluke or outlier.

⁷Official sources differ slightly on the number killed in New York City. Repeating our analyses with other reported values does not significantly change our estimates.

3.2. *Uncertainty in the tail location.* A fixed choice of x_{\min} underestimates the uncertainty in p due to the tail's unknown structure. Jointly estimating α and x_{\min} yields similar results, but with some interesting differences. Figure 1(b) shows 100 of the bootstrap models. The distribution of \hat{x}_{\min} is concentrated at $x_{\min} = 9$ or 10 (48% of samples), with an average scaling exponent of $\langle \hat{\alpha} \rangle = 2.40$. However, 15% of models choose $x_{\min} = 4$ or 5, and these produce much heavier-tailed models, with $\langle \hat{\alpha} \rangle = 2.21$.

This bimodal distribution in $\hat{\alpha}$ is caused by slight curvature in the empirical mid-to-upper tail, which may arise from aggregating multiple types of local events into a single global distribution (see Appendix C.3.3). The algorithm, however, accounts for this curvature by automatically estimating a slightly wider ensemble of models, with correspondingly greater density in the catastrophic range. As a result, the estimated probability is larger and the confidence intervals wider. Using 10,000 bootstraps, we find $\hat{p} = 0.347$, with 90% confidence intervals of [0.182, 0.669], or about a 35% chance over the 1968–2007 period.

3.3. *Alternative tail models.* Comparing these estimates with those derived using log-normal and stretched exponential tail models provides a check on their robustness, especially if the alternative models yield dramatically different estimates.

The mathematical forms of the alternatives are

$$\begin{aligned} \text{log-normal} \quad & \Pr(x) \propto x^{-1} \exp[-(\ln x - \mu)^2 / 2\sigma^2], \\ \text{stretched exp.} \quad & \Pr(x) \propto x^{\beta-1} e^{-\lambda x^\beta}, \end{aligned}$$

where we restrict each to a “tail” domain $x_{\min} \leq x < \infty$ (see Appendix A). In the stretched exponential, $\beta < 1$ produces a heavy-tailed distribution; in the log-normal, small values of μ and large values of σ yield heavy tails. Although both decay asymptotically faster than any power law, for certain parameter choices, these models can track a power law over finite ranges, which may yield only marginally lower estimates of large events.⁸

To simplify the comparison between the tail models, we fix $x_{\min} = 10$ and use 10,000 bootstraps for each fitted alternative tail model. This yields $\hat{p} = 0.112$ (CI: [0.063, 0.172]) for the log-normal and $\hat{p} = 0.187$ (CI: [0.115, 0.272]) for the stretched exponential, or roughly an 11% and 19% chance, respectively. These values are slightly lower than the estimates from the power-law model, but they too are consistently away from $p = 0$, which reinforces our conclusion that the size of 9/11 should not be considered a statistical outlier.

⁸The question of power law versus nonpower law [Clauset, Shalizi and Newman (2009)] is not always academic; for instance, macroeconomic financial models have traditionally and erroneously assumed nonpower-law tails that assign negligible probability to large events like widespread subprime loan defaults [Financial Crisis Inquiry Commission (2011)].

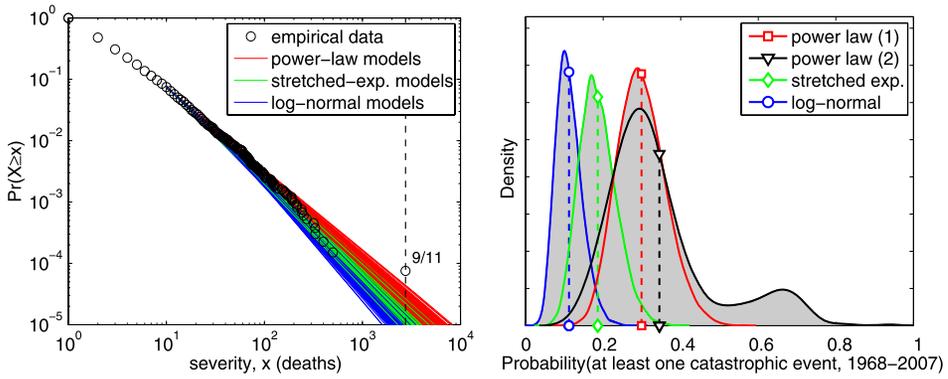


FIG. 2. (a) Empirical event severities with 100 bootstrap models for the power-law, log-normal and stretched exponential tail models, with $x_{\min} = 10$ fixed. (b) Bootstrap distributions of \hat{p} for each model, with overall estimates (Table 1) given by dashed lines.

Figure 2(a) shows the fitted ensembles for all three fixed- x_{\min} tail models, and Figure 2(b) shows the bootstrap distributions $\Pr(\hat{p})$ for these models, as well as the one with x_{\min} free. Although the bootstrap distributions for the log-normal and stretched exponential are shifted to the left relative to the two power-law models, all distributions overlap and none place significant weight below $p = 0.01$. The failure of the alternatives to disagree with the power law can be attributed to their estimated forms roughly tracking the power law’s over the empirical data’s range, which leads to similar probabilistic estimates of a catastrophic event.

3.4. *Impact of covariates.* Not all large terrorist events are of the same type, and thus our overall estimate is a function of the relative empirical frequency of different covariates and the structure of their marginal distributions. Here, we apply our procedure to the distributions associated with a few illustrative categorical event covariates to shed some additional light on the factors associated with large events. A generalization to and systematic analysis of arbitrary covariates is left for future work.

TABLE 1
Estimated per-event and worldwide historical probabilities for at least one catastrophic event over the period 1968–2007, for four tail models

Tail model	Parameters	Est. $\Pr(x \geq 2749)$ per event, $q(x)$	Est. prob. p , 1968–2007	90% CI (bootstrap)
Power law (1)	$\Pr(\hat{\alpha}, x_{\min} = 10)$	0.0000270200	0.299	[0.203, 0.405]
Power law (2)	$\Pr(\hat{\alpha}, \hat{x}_{\min})$	0.0000346345	0.347	[0.182, 0.669]
Stretched exp.	$\Pr(\hat{\beta}, \hat{\lambda}, x_{\min} = 10)$	0.0000156780	0.187	[0.115, 0.272]
log-normal	$\Pr(\hat{\mu}, \hat{\sigma}, x_{\min} = 10)$	0.0000090127	0.112	[0.063, 0.172]

For instance, international terrorist events, in which the attacker and target are from different countries, comprise 12% of the RAND-MIPT database and exhibit a much heavier-tailed distribution, with $\hat{\alpha} = 1.93 \pm 0.04$ and $\hat{x}_{\min} = 1$ (see Appendix C.3.1). This heavier tail more than compensates for their scarcity, as we estimate $\hat{p} = 0.475$ (CI: [0.309, 0.610]; Figure 6(a)) for at least one such catastrophic event from 1968–2007.⁹ A similar story emerges for events in economically developed nations, which comprise 5.3% of our data (see Appendix C.3.2). Focusing on such large events ($x \geq 10$), we estimate $\hat{p} = 0.225$ (CI: [0.037, 0.499], Figure 6(b)).

Another important event covariate is the type of weapon used. The tails of the weapon-specific distributions remain well described as power laws, but weapons like guns, knives and explosives exhibit less heavy tails (fewer large events) than unconventional weapons [Clauset, Young and Gleditsch (2007)], even as the former are significantly more common than the latter. The estimation algorithm used above can be generalized to handle categorical event covariates, and produces both marginal and total probability estimates (see Appendix C.3.3). Doing so yields an overall estimate of $\hat{p} = 0.564$ (CI: [0.338, 0.839]; Figure 7). Examining the marginal hazard rates, we see that the largest contribution comes from explosives, followed by fire and firearms.

4. Statistical forecasts. If the social and political processes that generate terrorist events worldwide are roughly stationary, our algorithm can be used to make principled statistical forecasts about the future probability of a catastrophic event. Although here we make the strong assumption of stationarity, this assumption could be relaxed using nonstationary forecasting techniques [Caires and Ferreira (2005), Clements and Hendry (1999), Shalizi et al. (2011)].

A simple forecast requires estimating the number of events n expected over the fixed forecasting horizon t . Using the RAND-MIPT data as a starting point, we calculate the number of annual deadly events worldwide n_{year} over the past 10 years. Figure 3 shows the empirical trend for deadly terrorist events worldwide from 1998–2007, illustrating a 20-fold increase in n_{year} , from a low of 180 in 1999 to a high of 3555 in 2006. Much of the increase is attributable to conflicts in Iraq and Afghanistan; excluding events from these countries significantly reduces the increase in n_{year} , with the maxima now being 857 deadly events in 2002 and 673 in 2006. However, the fraction of events that are severe ($x \geq 10$) remains constant, averaging $\langle p_{\text{tail}} \rangle = 0.082684$ (or about 8.3%) in the former case and 0.072601 (or about 7.3%) in the latter.

An estimated trend over the next decade could be obtained via fitting standard statistical models to annual data or by soliciting judgements from domain experts

⁹The implication of a larger \hat{p} for a covariate distribution, as compared to the full data set, is a smaller p for the excluded types of events. That is, a larger p for international events implies a smaller p for domestic events.

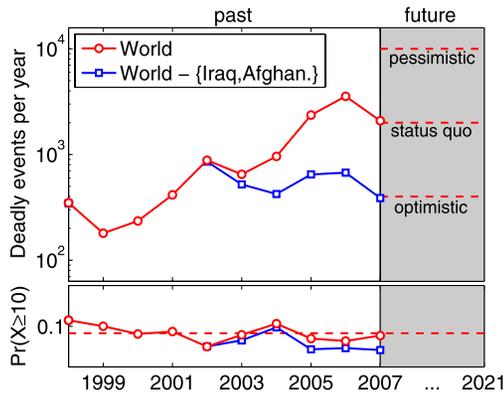


FIG. 3. (Upper) number of deadly (domestic and international) terrorist events worldwide for the 10-year period 1998–2007, and three forecast scenarios. (Lower) fraction of events that are severe, killing at least 10 individuals and its 10-year average (dashed line).

about specific conflicts. For instance, Iraq and Afghanistan may decrease their production rates of new events over the next decade, leading n_{year} to decrease unless other conflicts replace their contributions. Rather than make potentially overly specific predictions, we instead consider three rough scenarios (the future’s trajectory will presumably lay somewhere between): (i) an optimistic scenario, in which the average number of terrorist attacks worldwide per year returns to its 1998–2002 level, at about $\langle n_{\text{year}} \rangle = 400$ annual events; (ii) a status quo scenario, where it remains at the 2007 level, at about 2000 annual events; and finally (iii) a pessimistic scenario, in which it increases to about 10,000 annual events.¹⁰

A quantitative statistical forecast is then obtained by applying the estimation algorithm to the historical data (now including the 9/11 event) and then generating synthetic data with the estimated number of future events n_{tail} . For each scenario, we choose $n_{\text{decade}} = 10 \times n_{\text{year}}$ and choose n_{tail} via equation (2.1) with $p_{\text{tail}} = 0.082684$ (its historical average). Finally, we fix $x_{\text{min}} = 10$ to facilitate comparison with our alternative tail models.

Table 2 summarizes the results, using 100,000 bootstraps for each of the three tail models in the three forecast scenarios. Under the status quo scenario, all three models forecast a 19–46% chance of at least one catastrophic event worldwide in the next decade. In the optimistic scenario, with events worldwide being about 5 times less common, the models forecast a 4–12% chance. These estimates depend strongly on the overall frequency of terrorist events n_{year} . Thus, the greater the popularity of terrorism worldwide, that is, the more often terrorist attacks are launched, the greater the general likelihood that at least one will be catastrophic.

¹⁰Modeling these rough event counts via a Poisson process with rate $\lambda_{\text{scenario}}$ would refine our forecasts slightly. More detailed event production models could also be used.

TABLE 2

Forecast estimates of at least one catastrophic event worldwide over a 10-year period, using three tail models in each of three forecast scenarios

Tail model	Pr($x \geq 2749$) forecast, 2012–2021		
	“Optimistic” $n_{\text{year}} \approx 400$	“Status quo” $n_{\text{year}} \approx 2000$	“Pessimistic” $n_{\text{year}} \approx 10,000$
Power law	0.117	0.461	0.944
Stretched exp.	0.072	0.306	0.823
log-normal	0.043	0.193	0.643

Any progress in moving the general frequency of terrorism toward the more optimistic scenario is likely to reduce the overall, near-term probability of a catastrophic event.

5. Improved estimates. Our analysis places the 1968–2007 worldwide historic probability of a catastrophic event in the 11–35% range (see Table 1) and none of the alternative or covariate models provide any support for judging the size of 9/11 as statistically unlikely. The wide confidence interval illustrates the difficulty of obtaining precise estimates when accounting for model and parameter uncertainty. That being said, our calculations could be further refined to improve the overall estimates, incorporate additional sources of uncertainty or address specific questions, by relaxing portions of our i.i.d. treatment of event severities. We discuss several such possibilities here, but leave their investigation for the future.

First, our algorithm assumes a stationary event generation process, which is unlikely to be accurate in the long term. Technology, population, culture and geopolitics are believed to exhibit nonstationary dynamics and these likely play some role in event severities. Thus, techniques for statistical forecasting in nonstationary time series [Caires and Ferreira (2005), Clements and Hendry (1999), Shalizi et al. (2011)] could be used to identify subtle trends in the relevant covariates to make more accurate forecasts.

Second, our algorithm is silent regarding efforts to prevent events or mitigate their severity [Kilcullen (2010)]. However, the historical impact of these processes is implicitly present in our empirical data because only events that actually occurred were recorded. Thus, our results may be interpreted as probabilities conditioned on historical prevention or mitigation efforts. To the extent that policies have an impact on incidence and severity, more accurate estimates may be achievable by incorporating models of policy consequences or interactions between different actors. Similarly, our algorithm is silent regarding the actors responsible for events, and incorporating models of organizational capabilities, proclivities, etc. [Asal and Rethemeyer (2008), Clauset and Gleditsch (2009), Jackson et al. (2005)] may improve the estimates.

Finally, our approach is nonspatial and says little about where the event might occur. Incorporating more fine-grained spatial structure, for example, to make country-level or theatre-level estimates [Zammit-Mangion et al. (2012)] (as is now being done in seismology [Lee et al. (2011)]), or incorporating tactical information, for example, about specific CBRN attacks, may be possible. Such refinements will likely require strong assumptions about many context-specific factors [Gartzke (1999)], and it remains unclear whether accurate estimates at these scales can be made. At the worldwide level of our analysis, such contingencies appear to play a relatively small role in the global pattern, perhaps because local-level processes are roughly independent. This independence may allow large-scale general patterns to emerge from small-scale contingent chaos [Lorenz (1963), Strogatz (2001)] via a Central Limit Theorem averaging process, just as regularities in birth rates exist in populations despite high contingency for any particular conception. How far into this chaos we can venture before losing general predictive power remains unclear [Rundle et al. (2011), Ward, Greenhill and Bakke (2010)].

6. Discussion. In many complex social systems, although large events have outsized social significance, their rarity makes them difficult to study. Gaining an understanding of such systems requires determining if the same or different processes control the appearance of small, common events versus large, rare events. A critical scientific problem is estimating the true but unknown probability of such large events, and deciding whether they should be classified as statistical outliers. Accurate estimates can facilitate historical analysis, model development and statistical forecasts.

The algorithm described here provides a principled and data-driven solution for this problem that naturally incorporates several sources of uncertainty. Conveniently, the method captures the tendency of highly-skewed distributions to produce large events without reference to particular generative mechanisms or strong assumptions about the tail's structure. When properly applied, it provides an objective estimate of the historical or future probability of a rare event, for example, an event that has occurred exactly once.

Using this algorithm to test whether the size of the 9/11 terrorist events, which were nearly six times larger than the next largest event, could be an outlier, we estimated the historical probability of observing at least one 9/11-sized event somewhere in the world over the past 40 years to be 11–35%, depending on the particular choice of tail model used to fit the distribution's upper tail. These values are much larger than any reasonable definition of a statistical anomaly and thus the size of 9/11, which was nearly six times larger than the next largest event, should not be considered statistically unlikely, given the historical record of events of all sizes.

This conclusion is highly robust. Conditioning on the relative frequency of important covariates [Clauset, Young and Gleditsch (2007)], such as the degree of

economic development in the target country, whether an event is domestic or international, or the type of weapon used, we recover similar estimates, with additional nuance. Large events are probabilistically most likely to target economically developed nations, be international in character and use explosives, arson, firearms or unconventional weapons. Although chemical and biological events can also be very large [Cameron (2000)], historically they are rare at all sizes, and this outweighs the heaviness of their tail.

Furthermore, using only event data prior to 9/11 (as opposed to using all available data sans 9/11), we find a similar overall historical hazard rate. This suggests that the worldwide probability for large events has not changed dramatically over the past few decades. In considering three simple forecast scenarios for the next 10 years, we find that the probability of another large event is comparable to its historical level over the past 40 years. This risk seems unlikely to decrease significantly without a large reduction in the number of deadly terrorist events worldwide.

Of course, all such estimates are only as accurate as their underlying assumptions, and our method treats event sizes as i.i.d. random variables drawn from a stationary distribution. For complex social phenomena in general, it would be foolish to believe this assumption holds in a very strong sense, for example, at the micro-level or over extremely long time scales, and deviations will lower the method's overall accuracy. For instance, nonstationary processes may lower the global rate of large events faster than smaller events, leading to overestimates in the true probability of a large event. However, the i.i.d. assumption appears to be statistically justified at the global spatial and long-term temporal scales studied here. Identifying the causes of this apparent i.i.d. behavior at the global scale is an interesting avenue for future work.

The relatively high probability of a 9/11-sized event, both historically and in the future, suggests that the global political and social processes that generate large terrorist events may not be fundamentally different from those that generate smaller, more common events. Although the mechanism for event severities remains unclear [Clauset, Young and Gleditsch (2010)], the field of possible explanations should likely be narrowed to those that generate events of all sizes.

Independent of mechanistic questions, the global probability of another large terrorist event remains uncomfortably high, a fact that can inform our expectations [as with large natural disasters Gumbel (1941), Gutenberg and Richter (1944), Reed and McKelvey (2002)] of how many such events will occur over a long time horizon and how to appropriately anticipate or respond to them. This perspective is particularly relevant for terrorism, as classical models aimed at predicting event incidence tend to dramatically underestimate event severity [Clauset, Young and Gleditsch (2007)].

To conclude, the heavy-tailed patterns observed in the frequency of severe terrorist events suggests that some aspects of this phenomenon, and possibly of other complex social phenomena, are not nearly as contingent or unpredictable as is often assumed. That is, there may be global political and social processes that can

be effectively described without detailed reference to local conflict dynamics or the strategic trade-offs among costs, benefits and preferences of individual actors. Investigating these global patterns offers a complementary approach to the traditional rational-actor framework [Bueno de Mesquita (2003)] and a new way to understand what regularities exist, why they exist, and their implications for long-term stability.

APPENDIX A: TAIL MODELS

The functional form and normalization of the tail model should follow the type of empirical data used. For instance, if the empirical data are real-valued, the power-law tail model has the form

$$(A.1) \quad \Pr(y|\alpha, x_{\min}) = \left(\frac{\alpha - 1}{x_{\min}}\right) \left(\frac{y}{x_{\min}}\right)^{-\alpha}, \quad \alpha > 1, y \geq x_{\min} > 0.$$

Given a choice of x_{\min} , the maximum likelihood estimator for this model is

$$(A.2) \quad \hat{\alpha} = 1 + n / \sum_{i=1}^n \ln(x_i/x_{\min}).$$

The severity of a terrorist attack, however, is given by an integer. Thus, in our analysis of terrorist event severities, we use the discrete form of the power-law distribution

$$(A.3) \quad \Pr(y|\alpha, x_{\min}) = y^{-\alpha} / \zeta(\alpha, x_{\min}), \quad \alpha > 1, y \geq x_{\min} > 0,$$

where $\zeta(\hat{\alpha}, x_{\min}) = \sum_{i=x_{\min}}^{\infty} i^{-\alpha}$ is the generalized or incomplete zeta function. The MLE for the discrete power law is less straightforward, being the solution to the transcendental equation

$$(A.4) \quad \frac{\zeta'(\hat{\alpha}, x_{\min})}{\zeta(\hat{\alpha}, x_{\min})} = -\frac{1}{n} \sum_{i=1}^n x_i.$$

However, it is straightforward to directly maximize the log-likelihood function for the discrete power law in order to obtain $\hat{\alpha}$:

$$(A.5) \quad \mathcal{L}(\alpha) = -n \ln \zeta(\alpha, x_{\min}) - \alpha \sum_{i=1}^n \ln x_i.$$

Past work shows that the continuous model given by (A.3) provides a reasonably good approximation to the discrete case when x_{\min} takes moderate values. In our own experiments with this approximation, we find that when $x_{\min} \gtrsim 10$ the difference in estimated probabilities for observing one or more 9/11-sized events between using the discrete versus continuous model is at most a few percent.

Estimates of x_{\min} may be obtained using any of several existing automatic methods [Breiman, Stone and Kooperberg (1990), Danielsson et al. (2001), Dekkers

and de Haan (1993), Drees and Kaufmann (1998), Hancock and Jones (2004)]. We use the Kolmogorov–Smirnov goodness-of-fit statistic minimization (KS-minimization) technique [Clauset, Shalizi and Newman (2009), Clauset, Young and Gleditsch (2007)]. This method falls in the general class of distance minimization techniques for selecting the size of the tail [Reiss and Thomas (2007)], and was previously used to analyze event severities in global terrorism. The KS statistic [Press et al. (1992)] is the maximum distance between the CDFs of the data and the fitted model:

$$(A.6) \quad D = \max_{x \geq x_{\min}} |S(x) - P(x)|,$$

where $S(x)$ is the CDF of the data for the observations with value at least x_{\min} , and $P(x)$ is the CDF of the maximum-likelihood power-law model for the region $x \geq x_{\min}$. Our estimate \hat{x}_{\min} is then the value of x_{\min} that minimizes D . In the event of a tie between several choices for x_{\min} , we choose the smaller value, which improves the statistical power of subsequent analyses by choosing the larger effective sample size.

Our alternative tail models are the log-normal and the stretched exponential distributions, modified to include a truncating parameter x_{\min} . These distributions are normally defined on continuous variables. The structure of their extreme upper tails for $x_{\min} = 10$, however, is close to that of their discrete versions, and the continuous models are significantly easier to estimate from data. For the results presented in the main text, we used the continuous approximation of the upper tails for these models.

APPENDIX B: ESTIMATOR ACCURACY

We quantify the expected accuracy of our estimates under two experimental regimes in which the true probability of a catastrophic event can be calculated analytically.

1. Draw n values i.i.d. from a power-law distribution with $x_{\min} = 10$ and some α ; define $x = \max_i \{x_i\}$, the largest value within that sample. This step ensures that we treat the synthetic data exactly as we treated our empirical data and provides a particularly challenging test, as the largest generated value exhibits the greatest statistical fluctuations.
2. Draw $n - 1$ i.i.d. values from a power-law distribution with $x_{\min} = 10$ and some α , and then add a single value of size x whose true probability of appearing under the generative model is $p = 0.001$, that is, we contaminate the data set with a genuine outlier.

Figure 4 shows the results of both experiments, where we measure the mean absolute error (MAE) and the mean ratio between \hat{p} and the true p . Even for samples as small as $n = 40$ observations, the absolute error is fairly small and decreases with increasing sample size n . In the first experiment, the error rate decays

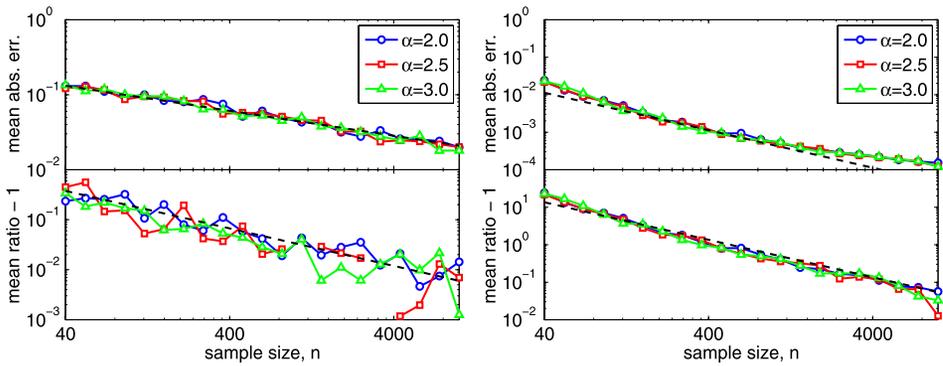


FIG. 4. The mean absolute error $\langle |\hat{p} - p| \rangle$ and mean relative error $\langle \hat{p}/p \rangle - 1$ for (a) n values drawn i.i.d. from a stationary power-law distribution with $x_{\min} = 10$ and some α , with the target size being the single largest value in the draw, and for (b) $n - 1$ values to which we add a single outlier (with true $p = 0.001$). In both experiments, both types of errors are small even for fairly small sample sizes and decay further as n increases.

like $O(n^{-1/3})$, approaching 0.01 error rates as n approaches 5000 [Figure 4(a)], while in the second it decays like $O(n^{-1})$ up to about $n = 4000$, above which the rate of decay attenuates slightly [Figure 4(b)].

Absolute deviations may mask dramatic relative errors, for example, if the true probability is very close to one or zero (as in our contaminated samples experiment). The mean ratio of \hat{p} to p would reveal such mistakes. The lower panels in Figure 4 show that this is not the case: the estimation procedure is close both in absolute and in relative terms. As the sample size increases, the estimated probability converges on the true probability. For contaminated data sets, the \hat{p}/p can be fairly large when n is very small, but for sample sizes of a few hundred observations, the method correctly estimates the relative size of the outlier’s probability.

APPENDIX C: ROBUSTNESS CHECKS

We present three checks of the robustness of our probability estimates: (i) using simple parametric models without the bootstrap, (ii) using an alternative source of terrorist event data, and (iii) using event covariates to refine the estimates. In each case, we find roughly similar-sized estimates.

C.1. Estimates using simple models. A simpler model for estimating the historical probability of a 9/11-sized or larger terrorist event assumes the following: (i) a stationary generative process for event severities worldwide, (ii) event sizes are i.i.d. random variables drawn from (iii) a power-law distribution that (iv) spans the entire range of possible severities ($x_{\min} = 1$), and (v) has a precisely-known parameter value $\alpha = 2.4$.

A version of this model was used in a 2009 Department of Defense-commissioned JASON report on “rare events” [McMorrow (2009)], which estimated the

historical probability of a catastrophic (9/11-sized or larger) terrorist event as 23% over 1968–2006. The report used a slightly erroneous estimate of the power law’s normalizing constant, a slightly different estimate of α and a smaller value of n . Here, we repeat the JASON analysis, but with more accurate input values.

Let $q(x)$ be the probability of observing a catastrophic event of size x . With event severities being i.i.d. random variables drawn from a fixed distribution $\text{Pr}(y)$, the generation of catastrophic events can be described by a continuous-time Poisson process with rate $q(x)$ [Boas (2005)]. Approximating x as a continuous variable, in a sequence of n such events, the probability \hat{p} of observing at least one of catastrophic severity is

$$(C.1) \quad \begin{aligned} \hat{p} &= 1 - [1 - q(x)]^n \\ &\approx 1 - e^{-nq(x)}. \end{aligned}$$

The rate $q(x)$ is simply the value of the complementary CDF at x , and for a power-law distribution is given by

$$(C.2) \quad \begin{aligned} q(x) &= \int_x^\infty \text{Pr}(y) \, dy \\ &= (\alpha - 1)x_{\min}^{\alpha-1} \int_x^\infty y^{-\alpha} \, dy \\ &= \left(\frac{x}{x_{\min}}\right)^{1-\alpha} \end{aligned}$$

for $x \geq x_{\min}$. Substituting $x_{\min} = 1$ and $\alpha = 2.4$ yields the per-event probability of a catastrophic event $q(2749) = 0.0000153164$.

The RAND-MIPT database records $n = 13274$ deadly events worldwide from 1968–2007; thus, substituting n and $q(x)$ into (C.1) yields a simple estimate of the probability of observing at least one catastrophic event over the same time period $\hat{p} = 1 - e^{-13274q(2749)} = 0.184$, or about 18%.

However, this calculation underestimates the true probability of a large event because the empirical distribution decays more slowly than a power law with $\alpha = 2.4$ at small values of x . Empirically 7.5% of the 13,274 fatal events have at least 10 fatalities, but a simple application of (C.2) using $x = 10$ shows that our model predicts that only 4.0% of events should be this severe. Thus, events with $x \geq 10$ occur empirically almost twice as often as expected, which leads to a significant underestimate of p .

By restricting the power-law model to the tail of the distribution, setting $x_{\min} = 10$ and noting that only $n = 994$ events had at least this severity over the 40-year period, we can make a more accurate estimate. Repeating the analysis above, we find $q(2749) = 0.0000288098$ and $\hat{p} = 0.318$, or about a 32% chance of a catas-

trophic event,¹¹ a value more in line with the estimates derived using our bootstrap-based approach in the main text.

C.2. Estimates using the Global Terrorism Database. An alternative source of global terrorism event data is the Global Terrorism Database [START (2011)], which contains 98,112 events worldwide from 1970–2007. Of these, 38,318 were deadly ($x > 0$). Some events have fractional severities due to having their total fatality count divided evenly among multiple event records; we recombined each group of fractional-severity events into a single event, yielding 38,255 deadly events over 38 years. Analyzing the GTD data thus provides a check on our results for the RAND-MIPT data.

The largest event in the GTD is 9/11, with severity 2763, and the second largest is the 13 April 1994 Rwandan massacre of Tutsi refugees, with 1180 reported fatalities. This event is absent from the RAND-MIPT data; its inclusion in the GTD highlights this data set’s broader definition of terrorism, which includes a number of genocidal or criminal events.

The best fitting power-law model obtained using the methodology of Clauset, Shalizi and Newman (2009) is $\hat{\alpha} = 2.91 \pm 0.22$ and $\hat{x}_{\min} = 39$. The $p < 0.1$ for this model may be attributable to the unusually large number of perfectly round-number severities in the data set, for example, 10, 20, 100, 200, etc., which indicates rounding effects in the reporting. (These appear in Figure 5 as small discon-

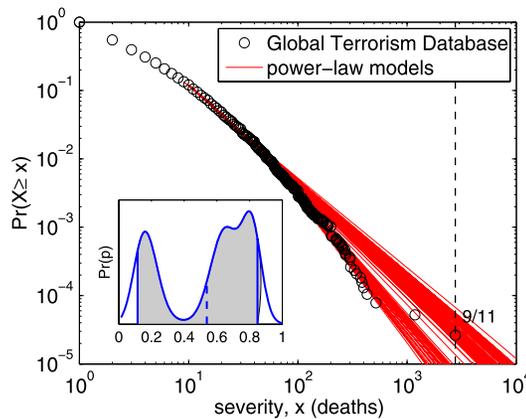


FIG. 5. Empirical distribution of event severities from the GTD [START (2011)] with 100 power-law models, fitted to bootstraps of the data. Inset shows the estimated distribution of binomial probabilities $\Pr(\hat{p})$ for one or more catastrophic events.

¹¹To make our reported per-event probabilities $q(x)$ consistent across models, we report them as $q(x) = \Pr(X \geq x | X \geq x_{\min}) \Pr(X \geq x_{\min})$, that is, the probability that a tail event is catastrophic times the probability that the event is a tail event. These values can be used with (C.1) to make rough estimates if the corresponding n is the total number of deadly events.

tinuous drops in the complementary CDF at round-number locations; true power-law distributed data have no preference for round numbers and thus their presence is a statistically significant deviation from the power-law form.)

Using the algorithm described in the main text with 10,000 bootstraps, we estimate a 38-year probability of at least one catastrophic event as $\hat{p} = 0.534$ (with 90% CI [0.115, 0.848]) or about a 53% chance. Repeating our analysis using the two alternative tail models yields only a modest decrease, as with the RAND-MIPT data.

Figure 5 shows the empirical fatality distribution along with 100 fitted power-law models, illustrating the heavy-tailed structure of the GTD severity data. Notably, the maximum likelihood estimate for α is larger here (indicating a less heavy tail) than for the RAND-MIPT data. However, the marginal distribution $\Pr(\hat{\alpha})$ is bimodal, with one mode centered on $\alpha = 2.93$ and a second larger mode centered at roughly $\alpha = 2.4$, in agreement with the RAND-MIPT data. Furthermore, the failure of the GTD-estimated \hat{p} to be dramatically lower than the one estimated using RAND-MIPT data supports our conclusion that the size of 9/11 was not statistically unlikely.

C.3. Impact of event covariates.

C.3.1. *International versus domestic, and events prior to 1998.* Events in the RAND-MIPT database with dates before 1 January 1998 are mainly international events, that is, the attacker's country of origin differed from the target's country. Subsequent to this date, both domestic and international events are included but their domestic versus international character is not indicated. Analyzing events that occurred before this breakpoint thus provides a natural robustness check for our overall estimate: (i) we can characterize the effect that domestic versus international events have on the overall estimate and (ii) we can test whether the probability estimates have changed significantly in the past decade.

The pre-1998 events comprise 12% of the RAND-MIPT database and exhibit a more heavy-tailed distribution ($\hat{\alpha} = 1.92 \pm 0.04$ and $x_{\min} = 1$). Using 10,000 bootstraps, we estimate $\hat{p} = 0.475$ (90% CI: [0.309, 0.610]) for at least one catastrophic international event over the target period. Figure 6(a) shows the empirical distribution for international events and the ensemble of fitted models, illustrating good visual agreement with the empirical distribution.

The estimate for international-only data ($\hat{p} = 0.475$) is larger than the estimate derived using the full data set ($\hat{p} = 0.347$), although these values may not be as different as they seem, due to their overlapping confidence intervals. Fundamentally, the larger estimate is caused by the heavier-tailed distribution of the international-only data. Because the full data set includes these international events, this result indicates that domestic events tend to exhibit a lighter tail, and thus generate large terrorist events with smaller probability. As a general guideline, subsets of the full data set should be analyzed with caution, as their selection is necessarily conditioned. The full data set provides the best estimate of large events of all types.

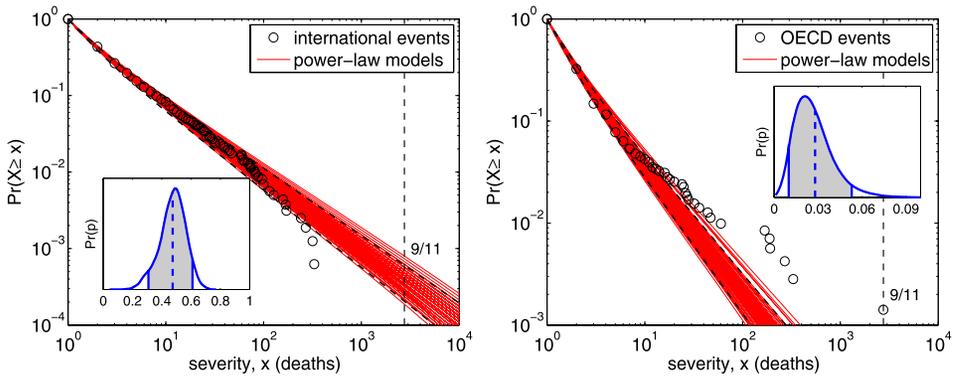


FIG. 6. Empirical distributions, with 100 power-law bootstrap models, for (a) international events (events from 1968–1997 in the RAND-MIPT database) and (b) events within the OECD nations; dashed lines show the 90% CI on $\hat{\alpha}$. Insets show the estimated distribution $\Pr(\hat{p})$ with 90% confidence intervals (shaded area) and overall estimate (dashed line).

C.3.2. *Economic development.* A similar story emerges for deadly events in economically developed nations, defined here as the member countries of the Organisation for Economic Co-operation and Development (OECD), as of the end of the period covered by the RAND-MIPT data, which are 5.3% of all deadly events. The empirical distribution [Figure 6(b)] of event severities shows unusual structure, with the upper tail ($x \geq 10$ fatalities) decaying more slowly than the lower tail. To handle this oddity, we conduct two tests.

First, we consider the entire OECD data set, estimating both α and x_{\min} . Using 10,000 bootstraps yields $\hat{p} = 0.028$ (with 90% CI [0.010, 0.053]) or roughly a 3% chance over the 40-year period, which is slightly above our $p = 0.01$ cutoff for a statistically unlikely event. Figure 6(b) shows the resulting ensemble of fitted models, illustrating that the algorithm is placing very little weight on the upper tail. Second, we apply the algorithm with a fixed $x_{\min} = 10$ in order to focus explicitly on the distribution's upper tail. In this case, 10,000 bootstraps yield $\hat{p} = 0.225$, with 90% CI as [0.037, 0.499].

C.3.3. *Type of weapon.* Finally, we consider the impact of the attack's weapon type, and we generalize the estimation algorithm to the multi-covariate case. Events are classified as (i) chemical or biological, (ii) explosives (includes remotely detonated devices), (iii) fire, arson and firebombs, (iv) firearms, (v) knives and other sharp objects, and (vi) other, unknown or unconventional. Given the empirically observed distributions over these covariates, we would like to know the probability of observing at least one catastrophic-sized event from any weapon type.

This requires generalizing our Monte Carlo algorithm: let $(x, c)_i$ denote the severity x and categorical covariate c for the i th event. Thus, denote the empirical data by $X = \{(x, c)_i\}$.

1. Generate Y by drawing $(y, c)_j$, $j = 1, \dots, n$, uniformly at random, with replacement, from the original data $\{(x, c)_i\}$ (sans the ℓ catastrophic events).
2. For each covariate type c in Y , jointly estimate $\hat{x}_{\min}^{(c)}$ and the tail-model parameters $\theta^{(c)}$, and compute $n_{\text{tail}}^{(c)} = \#\{y_j \geq \hat{x}_{\min}^{(c)}\}$.
3. For each covariate type c in Y , generate a synthetic data set by drawing $n_{\text{tail}}^{(c)}$ random deviates from the fitted tail model with parameters $\hat{\theta}^{(c)}$.
4. If any of the covariate sequences of synthetic events includes at least ℓ events of size x or greater, set $\rho = 1$; otherwise, set it to zero.

In applying this algorithm to our data, we choose $\ell = 1$ and $x = 2749$, as with our other analyses. In step 2, we again use the KS-minimization technique of [Clauset, Shalizi and Newman \(2009\)](#) to choose x_{\min} and estimate θ for a power-law tail model via maximum likelihood. Finally, as with the univariate version of the algorithm, bootstrap confidence intervals may be obtained [[Efron and Tibshirani \(1993\)](#)], both for the general hazard and the covariate-specific hazard, by repeating steps 3 and 4 many times for each bootstrap and tracking the distribution of binomial probabilities.

Using 10,000 bootstraps and drawing 1000 synthetic data sets from each bootstrap, we estimate $\hat{p} = 0.564$, with 90% confidence intervals of $[x, y]$. Again, this value is well above the cutoff for a 9/11-sized attack being statistically unlikely. Figure 7(a)–(f) shows the ensembles for each weapon-specific severity distribution. As a side effect of this calculation, we may also calculate the probability that a catastrophic event will be generated by a particular type of weapon. The following table gives these marginal probability estimates, which are greatest for explosives, fire, firearms and unconventional weapon types.

It is emphasized that these are historical estimates, based on the relative frequencies of weapon covariates in the historical RAND-MIPT data. If the future exhibits similar relative frequencies and total number of attacks, then they may also be interpreted as future hazards, but we urge strong caution in making these assumptions.

Weapon type	Historical \hat{p}	90% CI
Chem. or bio.	0.023	[0.000, 0.085]
Explosives	0.374	[0.167, 0.766]
Fire	0.137	[0.012, 0.339]
Firearms	0.118	[0.015, 0.320]
Knives	0.009	[0.001, 0.021]
Other or unknown	0.055	[0.000, 0.236]
Any	0.564	[0.338, 0.839]

(The sum of marginal probabilities exceeds that of the “any” column because in some trials, catastrophic events are generated in multiple categories.)

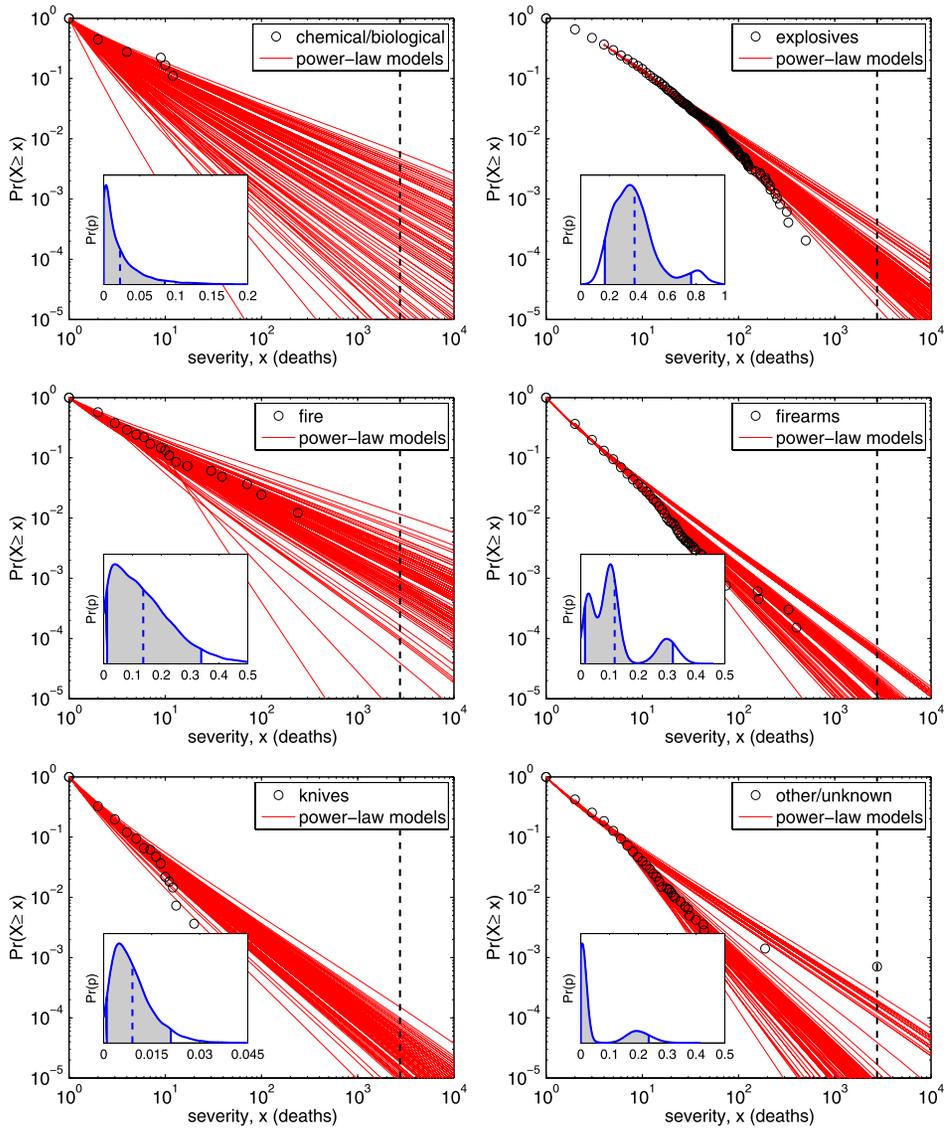


FIG. 7. Empirical distribution, with 100 power-law bootstrap models, for events using (a) chemical or biological, (b) explosives (includes remote detonation), (c) fire, arson and firebombs, (d) firearms, (e) knives or sharp objects, and (f) other, unknown or unconventional weapons. Insets: marginal distributions of estimated hazard rates $\Pr(\hat{p})$, with the region of 90% confidence shaded and the mean value indicated by the dashed line.

Acknowledgments. The authors thank Cosma Shalizi, Nils Weidmann, Julian Wücherpfennig, Kristian Skrede Gleditsch, Victor Asal, Didier Sornette, Lars-Erik Cederman and Patrick Meier for helpful conversations.

Implementations of our numerical methods are available online at <http://www.santafe.edu/~aaronc/rareevents/>.

REFERENCES

- ADLER, R. J., FELDMAN, R. E. and TUQUU, M. S., eds. (1998). *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*. Birkhäuser, Boston, MA. [MR1652283](#)
- ASAL, V. and RETHEMEYER, R. K. (2008). The nature of the beast: Organizational structures and the lethality of terrorist attacks. *Journal of Politics* **70** 437–449.
- BECK, N., KING, G. and ZENG, L. (2000). Improving quantitative studies of international conflict: A conjecture. *American Political Science Review* **94** 21–35.
- BOAS, M. L. (2005). *Mathematical Models in the Physical Sciences*. Wiley, Hoboken, NJ.
- BREIMAN, L., STONE, C. J. and KOOPERBERG, C. (1990). Robust confidence bounds for extreme upper quantiles. *J. Stat. Comput. Simul.* **37** 127–149. [MR1082452](#)
- BROWN, D., DALTON, J. and HOYLE, H. (2004). Spatial forecast methods for terrorist events in urban environments. In *Intelligence and Security Informatics. Lecture Notes in Computer Science* **3073** 426–435. Springer, Heidelberg.
- BUENO DE MESQUITA, B. (2003). Ruminations on challenges to prediction with rational choice models. *Rationality and Society* **15** 136–147.
- BUENO DE MESQUITA, B. (2011). A new model for predicting policy choices: Preliminary tests. *Conflict Management and Peace Science* **28** 64–84.
- CAIRES, S. and FERREIRA, J. A. (2005). On the non-parametric prediction of conditionally stationary sequences. *Stat. Inference Stoch. Process.* **8** 151–184. [MR2121675](#)
- CAMERON, G. (2000). WMD terrorism in the United States: The threat and possible countermeasures. *Nonproliferation Review* **7** 162–179.
- CESA-BIANCHI, N., CONCONI, A. and GENTILE, C. (2004). On the generalization ability of on-line learning algorithms. *IEEE Trans. Inform. Theory* **50** 2050–2057. [MR2097190](#)
- CLAESKENS, G. and HJORT, N. L. (2008). *Model Selection and Model Averaging*. Cambridge Univ. Press, Cambridge. [MR2431297](#)
- CLAUSET, A. and GLEDITSCH, K. S. (2012). The developmental dynamics of terrorist organizations. *PLoS ONE* **7** e48633.
- CLAUSET, A., SHALIZI, C. R. and NEWMAN, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Rev.* **51** 661–703. [MR2563829](#)
- CLAUSET, A., YOUNG, M. and GLEDITSCH, K. S. (2007). On the frequency of severe terrorist events. *Journal of Conflict Resolution* **51** 58–87.
- CLAUSET, A., YOUNG, M. and GLEDITSCH, K. S. (2010). A novel explanation of the power-law form of the frequency of severe terrorist events: Reply to saperstein. *Peace Economics, Peace Science and Public Policy* **16** article 12.
- CLEMENTS, M. P. and HENDRY, D. F. (1999). *Forecasting Non-Stationary Economic Time Series*. MIT Press, Cambridge, MA.
- DANIELSSON, J., DE HAAN, L., PENG, L. and DE VRIES, C. G. (2001). Using a bootstrap method to choose the sample fraction in tail index estimation. *J. Multivariate Anal.* **76** 226–248. [MR1821820](#)
- DE HAAN, L. and FERREIRA, A. (2006). *Extreme Value Theory: An Introduction*. Springer, New York. [MR2234156](#)
- DEKKERS, A. L. M. and DE HAAN, L. (1993). Optimal choice of sample fraction in extreme-value estimation. *J. Multivariate Anal.* **47** 173–195. [MR1247373](#)
- DESMARAIS, B. A. and CRANMER, S. J. (2011). Forecasting the locational dynamics of transnational terrorism: A network analytic approach. In *Proc. European Intelligence and Security Informatics Conference* 171–177. IEEE Computer Society, Washington, DC.

- DREES, H. and KAUFMANN, E. (1998). Selecting the optimal sample fraction in univariate extreme value estimation. *Stochastic Process. Appl.* **75** 149–172. [MR1632189](#)
- DRENNAN, M. P. (2007). Analyzing catastrophic terrorist events with applications to the food industry. In *The Economic Costs and Consequences of Terrorism* (H. W. Richardson, P. Gordon and J. E. Moore, eds.). Edward Elgar Publishing, Northampton, MA.
- EFRON, B. and TIBSHIRANI, R. J. (1993). *An Introduction to the Bootstrap*. *Monographs on Statistics and Applied Probability* **57**. Chapman & Hall, New York. [MR1270903](#)
- ENDERS, W. and SANDLER, T. (2000). Is transnational terrorism becoming more threatening? A time-series investigation. *Journal of Conflict Resolution* **44** 307–332.
- ENDERS, W. and SANDLER, T. (2002). Patterns of transnational terrorism, 1970–1999: Alternative time-series estimates. *International Studies Quarterly* **46** 145–165.
- ENDERS, W. E. and SANDLER, T. (2006). *The Political Economy of Terrorism*. Cambridge Univ. Press, Cambridge.
- ENDERS, W., SANDLER, T. and GAIBULLOEV, K. (2011). Domestic versus transnational terrorism: Data, decomposition, and dynamics. *Journal of Peace Research* **48** 319–337.
- FARMER, J. D. and LILLO, F. (2004). On the origin of power law tails in price fluctuations. *Quant. Finance* **4** 7–11.
- FINANCIAL CRISIS INQUIRY COMMISSION (2011). Financial crisis inquiry report. U.S. Government Printing Office, Washington. Available at <http://bit.ly/hEsCwo>.
- GARTZKE, E. (1999). War is in the error term. *International Organization* **53** 567–587.
- GRÜNWARD, P. D. (2007). *The Minimum Length Description Principle*. MIT Press, Cambridge, MA.
- GUMBEL, E. J. (1941). The return period of flood flows. *Ann. Math. Statist.* **12** 163–190. [MR0004457](#)
- GUTENBERG, B. and RICHTER, C. F. (1944). Frequency of earthquakes in California. *Bull. Seismol. Soc. Amer.* **34** 185–188.
- HANCOCK, M. S. and JONES, J. H. (2004). Likelihood-based inference for stochastic models of sexual network evolution. *Theoretical Population Biology* **65** 413–422.
- HJORT, N. L. and CLAESKENS, G. (2003). Frequentist model average estimators. *J. Amer. Statist. Assoc.* **98** 879–899. [MR2041481](#)
- JACKSON, B. A., BAKER, J. C., CRAGIN, K., PARACHINI, J., TRUJILLO, H. R. and CHALK, P. (2005). *Aptitude for Destruction: Organizational Learning in Terrorist Groups and Its Implications for Combating Terrorism* **1**. RAND Corporation, Arlington, VA.
- KARDES, E. and HALL, R. (2005). *Survey of Literature on Strategic Decision Making in the Presence of Adversaries*. Center for Risk and Economic Analysis of Terrorism Events, Los Angeles, CA.
- KASS, R. E. and RAFTERY, A. E. (1995). Bayes factors. *J. Amer. Statist. Assoc.* **90** 773–795.
- KILCULLEN, D. (2010). *Counterinsurgency*. Oxford Univ. Press, Oxford.
- KING, G. and ZENG, L. (2001). Explaining rare events in international relations. *International Organization* **55** 693–715.
- LAST, G. and BRANDT, A. (1995). *Marked Point Processes on the Real Line: The Dynamic Approach*. Springer, New York. [MR1353912](#)
- LEE, Y.-T., TURCOTTE, D. L., HOLLIDAY, J. R., SACHS, M. K., RUNDLE, J. B., CHEN, C.-C. and TIAMPO, K. F. (2011). Results of the Regional Earthquake Likelihood Model (RELM) test of earthquake forecasts in California. *Proc. Natl. Acad. Sci. USA* **108** 16533–16538.
- LI, Q. (2005). Does democracy promote or reduce transnational terrorist incidents? *Journal of Conflict Resolution* **49** 278–297.
- LORENZ, E. N. (1963). Deterministic nonperiodic flow. *J. Atmospheric Sci.* **20** 130–141.
- MAJOR, J. A. (1993). Advanced techniques for modeling terrorism risk. *Journal of Risk Finance* **4** 15–24.

- MCMORROW, D. (2009). Rare events. JASON Report JSR-09-108, The MITRE Corporation, McLean, VA.
- MIPT (2008). Terrorism knowledge base. National Memorial Institute for the Prevention of Terrorism, Oklahoma City, OK.
- PAPE, R. A. (2003). The strategic logic of suicide terrorism. *American Political Science Review* **97** 343–361.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T. and FLANNERY, B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge Univ. Press, Cambridge. MR1201159
- REED, W. J. and MCKELVEY, K. S. (2002). Power-law behavior and parametric models for the size-distribution of forest fires. *Ecological Modelling* **150** 239–254.
- REISS, R. D. and THOMAS, M. (2007). *Statistical Analysis of Extreme Values with Applications to Insurance, Finance, Hydrology and Other Fields*, 3rd ed. Birkhäuser, Basel. MR2334035
- RESNICK, S. I. (2007). *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer, New York. MR2271424
- RUNDLE, J. B., HOLLIDAY, J. R., YODER, M., SACHS, M. K., DONNELLAN, A., TURCOTTE, D. L., TIAMPO, K. F., KLEIN, W. and KELLOGG, L. H. (2011). Earthquake precursors: Activation or quiescence? *Geophysical Journal International* **187** 225–236.
- RUSTAD, S. C. A., BUHAUG, H., FALCH, Å. and GATES, S. (2011). All conflict is local: Modeling sub-national variation in civil conflict risk. *Conflict Management and Peace Science* **28** 14–39.
- SAGEMAN, M. (2004). *Understanding Terror Networks*. Univ. Pennsylvania Press, Philadelphia.
- SHALIZI, C. R., JACOBS, A. Z., KLINKNER, K. L. and CLAUSET, A. (2011). Adapting to non-stationarity with growing expert ensembles. Preprint. Available at <http://arxiv.org/abs/1103.0949>.
- SORNETTE, D. (2009). Dragon-kings, black swans and the prediction of crises. *International Journal of Terraspace Science and Engineering* **2** 1–18.
- SORNETTE, D. and OUIILLON, G., eds. (2012). *Eur. Phys. J. Special Topics* **205** 1–373.
- START (2011). Global terrorism database (data file). Available at www.start.umd.edu/gtd/.
- STONE, M. (1974). Cross-validators choice and assessment of statistical predictions. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **36** 111–133.
- STROGATZ, S. H. (2001). *Nonlinear Dynamics And Chaos*. Westview, Boulder, CO.
- VALENZUELA, M. L., FENG, C., REDDY, P., MOMEN, F., ROZENBLIT, J. W., EYCK, B. T. and SZIDAROVSKY, F. (2010). A non-numerical predictive model for asymmetric analysis. In *Proc. 17th IEEE International Conference and Workshops on the Engineering of Computer-Based Systems* 311–315. IEEE Computer Society, Washington, DC.
- VUONG, Q. H. (1989). Likelihood ratio tests for model selection and nonnested hypotheses. *Econometrica* **57** 307–333. MR0996939
- WARD, M. D., GREENHILL, B. D. and BAKKE, K. M. (2010). The perils of policy by p -value: Predicting civil conflicts. *Journal of Peace Research* **47** 363–375.
- WULF, W. A., HAIMES, Y. Y. and LONGSTAFF, T. A. (2003). Strategic alternative responses to risks of terrorism. *Risk Anal.* **23** 429–444.
- ZAMMIT-MANGION, A., DEWAR, M., KADIRKAMANATHAN, V. and SANGUINETTI, G. (2012). Point process modeling of the Afghan War Diary. *Proc. Natl. Acad. Sci. USA* **109** 12414–12419.

DEPARTMENT OF COMPUTER SCIENCE AND
THE BIOFRONTIERS INSTITUTE
UNIVERSITY OF COLORADO, BOULDER
BOULDER, COLORADO 80309
USA
E-MAIL: aaron.clauset@colorado.edu

DEPARTMENT OF MANAGEMENT, TECHNOLOGY
AND ECONOMICS
ETH ZÜRICH
KREUZPLATZ 5, CH-8032 ZÜRICH
SWITZERLAND
E-MAIL: rwoodard@ethz.ch