REJOINDER OF "ESTIMATING THE HISTORICAL AND FUTURE PROBABILITIES OF LARGE TERRORIST EVENTS" BY AARON CLAUSET AND RYAN WOODARD

BY AARON CLAUSET AND RYAN WOODARD

University of Colorado, Boulder and ETH Zürich

We are pleased that our article has stimulated such thoughtful discussion, and we appreciate the discussants' interest in exploring the methodological and practical points of our study. The discussants make a number of excellent points worthy of future investigation and highlight the difficulty of making accurate statistical estimates of the likelihood of rare events in general, and of large terrorist events in particular. Our study is certainly not the final word on these topics and we look forward to future developments in these areas. In our rejoinder, we focus on selected points that will clarify the context of our study and open questions, including the choice of statistical models that are consistent with reasonable mechanistic models for the data under study, and the value of simple models in controlling uncertainty in complex social systems.

Not all tail models are equal. Two key motivations in our use of the power-law or simple Pareto tail model were (i) its previous use in modeling terrorist event severities, and (ii) its status as the only tail model with published mechanisms for the frequencies of large terrorist events. Although the debate is ongoing as to which mechanism, if any, is the correct explanation for the observed heavy-tailed pattern in event severities [see Clauset and Gleditsch (2012) for discussion], these mechanisms provide an important theoretical grounding for any statistical modeling of terrorism's upper tail. *Without such mechanisms, there is little theoretical justification for favoring one particular tail model over another to estimate extreme event probabilities.* Thus, we believe some amount of priority should be given to estimates derived from distributions like the power law, which have articulated and plausible underlying mechanisms for terrorist event severities. However, the surest way to reduce our ultimate uncertainty as to the likelihood of future large events is to identify and test alternative mechanisms for the heavy-tailed pattern in terrorist event severities, and we look forward to new work in that direction.

Disagreement among tail models. The statistical framework we presented is entirely general and can thus be used in conjunction with (i) any well-defined, automatic method for identifying the upper tail region, and (ii) any well-defined probabilistic model of an upper tail. (Although we modeled severities as i.i.d. random variables, this is not a requirement, and a clear understanding of the statistical

correlations among terrorist events at the global scale could, in principle, be incorporated into a more detailed model.) Given these choices, data-driven estimates are then produced.

Even without regard to their theoretical motivation, not all tail models are reasonable choices in this framework. Any model can be fitted to the data, but if it is a poor fit, we are under no obligation to trust its results. How then should we decide which models are good fits? The models used in our analysis (power law, log-normal and stretched exponential) were all previously demonstrated, under a combination of standard hypothesis tests and likelihood ratio tests, to meet this criteria [Clauset, Shalizi and Newman (2009)]. Of course, more flexible models, like the generalized Pareto distribution (GPD), the tapered Pareto or a piecewise Pareto, may also provide reasonable fits, as several discussants suggested. The difficulty, however, is how to interpret or reconcile models that produce conflicting estimates for the likelihood of a large event, and how to choose among models with different levels of flexibility.

Modern statistics does not offer clear answers to these questions because of the role played by $x_{\min}$, the smallest value for which the tail model holds. Unlike traditional model parameters, changing $x_{\min}$ changes the sample size, which confounds changes in statistical power with the usual bias-variance trade-off. The result is an additional risk in overfitting, particularly with flexible models like those suggested by the discussants, as larger values of $x_{\min}$ are considered. Simple models, like the power-law distribution, would seem to offer some protection against this risk. (We note that this problem of identifying reasonable tail models is ubiquitous in complex social systems with heavy-tailed variables, including financial markets [Farmer and Lillo (2004), Financial Crisis Inquiry Commission (2011)].) The method we used to choose $x_{\min}$ has reasonable properties and performs well in practice [Clauset, Shalizi and Newman (2009)]. New research on how best to choose $x_{\min}$ for out-of-sample forecasts and how to rigorously compare models estimated from differently sized samples would help resolve some of these difficulties.

Conclusions. The statistical modeling and forecasting of terrorist event severity is a relatively new endeavor that combines interesting methodological problems and tricky forms of uncertainty in an application area with arguably genuine practical benefit. There are many interesting and important questions worthy of study, and the discussants have identified a number of them. Better estimates and a deeper understanding of the pattern of large terrorist events and the mechanisms that produce them can inform our expectations (as with large natural disasters) of how many such events will occur over a long time horizon and how to appropriately anticipate or respond to them.

In closing, we note that the relatively smooth distribution of the sizes of terrorist events worldwide presents a puzzle. Given the highly contingent nature of individual events and individual conflicts, 9/11 being an outstanding example, how can the global distribution be so regular? This striking pattern suggests both that

accurate estimates of the probability of large events may be derived from modeling the relative frequency of much less severe events and that some aspects of global terrorism may not be as contingent or unpredictable as is often assumed. Understanding the origin of this global-level pattern, and the mechanisms by which local-level dynamics give rise to it, is itself an important research direction with real implications for understanding the fundamentals of violent political conflict. We look forward to new insights in these directions.

REFERENCES

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

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.

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