Modern Wars: Fundamental Patterns and Predictions

Neil Johnson

Department of Physics, University of Miami

Michael Spagat

Department of Economics, Royal Holloway University of London

Stijn van Weezel

School of Economics, University College Dublin

Introduction

Bohorquez et al. (2009) studies the distributions of event sizes within nine modern wars with event size defined by the number of people killed, finding that each distribution is well approximated by a power law with the power coefficients clustering around 2.5. Bohorquez et al. (2009) also presented a theoretical model that is driven by the coalescence and fragmentation of groups within warring organizations and that generates power-law distributions of violent events with the power coefficients clustering around 2.5. Johnson et al. (2013) showed that newly released event data for a number of further conflicts generally conformed to the same patterns and surveyed a number of extensions and elaborations of the family of coalescence-fragmentation models, thus confirming the robustness of the tendency toward power laws with coefficients near 2.5. Johnson et al. (2016) demonstrates that online ISIS communities display coalescence and fragmentation behaviours that are consistent with the theoretical models. A separate line of research ((Clauset and Woodard, 2013a) and (Clauset and Woodard, 2013b)) finds that power laws with coefficients around 2.5 fit well the size distribution of terrorist events. Clauset and Woodard (2013a) and Clauset and Woodard (2013b) use this established pattern to predict the probability of a terrorist attack comparable in scale to the 9/11 one.

The present paper has two main objectives. First, we exploit new event data on armed conflict and terrorism to explore the empirical patterns in the size distributions of violent events in both contexts. For war we use the new version of the data used in Johnson et al. (2013), enabling us to extend the reach of the analysis to quite a few additional conflicts, particularly in Asia. Our empirical work in terrorism innovates by operating at the organization level, enabling us to demonstrate that the size distributions of violent events perpetrated by individual terrorist organizations greatly resemble the distributions we find for armed conflicts. This finding deepens the link already identified in Bohorquez et al. (2009) between terrorism and insurgency. The second objective of our paper is to explore the potential for exploiting the regularities uncovered by the previous conflict literature to make useful predictions about the mixtures of event sizes in future wars. We base our predictions on the observed pattern, further supported by theory, that event size distributions in modern wars tend to follow power laws with coefficients near alpha. We use the range of power-law coefficients observed in dozens of modern conflicts to quantify the bands we place around our predictions. We keep a scorecard on the success rates of our predictions and conclude that it is possible to make useful predictions about the distribution of event sizes in future conflicts.

Data

Data on conflict events is taken from the Georeferenced Event Dataset (GED) from the Uppsala Conflict Data Programme (Sundberg and Melander, 2013). This is the most comprehensive and accurate georeferenced dataset on conflict available (Eck, 2012; Weidmann, 2013, 2015) that systematically collects information on the number of people killed in each event. The GED contains detailed information on the location, timing, and severity of conflict events, along with information on the warring parties generating these events. The dataset covers various forms of warfare such as conflicts between governments and rebel groups, non-state based conflicts (also known as communal violence), and violence perpetrated by the state or insurgency groups against civilians. We use the data covering all conflicts in Africa and Asia between 1989-2014 in the most recent version (v.4) of the dataset.[1] The GED coding rules only include conflicts that have reached a total fatality threshold of 25-battle related deaths in a year thus excluding some low intensity conflicts. However, this restriction should hardly matter for us since it excludes only minor conflicts that, anyway, may have been excluded for the other reason

[1] Although the GED offers a global dataset, conflicts in Asia and Africa are covered better than those in Europe and Latin America which only go back to 2005, thereby missing the Yugoslav Wars and much of the conflict in Colombia. Also note that the Syrian civil war is currently not included in the GED data.

of not having enough events to allow us to reliably fit a power law to the size distribution the violent events.

We include only true single events in our analysis, removing a small number of fatality counts that are not broken down to the event level. We also drop conflicts with fewer than 30 events. We are left with 98 African conflicts with 21,239 events and 104 Asian conflicts with 60,162 events.[2]

We offer a parallel analysis of terrorist incidents since there is some evidence suggesting that conflicts with terrorist organizations are similar to insurgency-counterinsurgency situations (Bohorquez et al. (2009)). For this work we use the Global Terrorism Database (GTD), which is provided by the National Consortium for the Study of Terrorism and Responses to Terrorism (START). A novel feature of the GTD dataset is that it includes both domestic and trans/inter-national terrorist incidents, in contrast to the other datasets on terrorism. The GTD is updated annually and provides the most comprehensive dataset on terrorist events that is publicly available. The GTD covers the period from 1970 to 2015 and includes detailed information on incident times, locations, fatality counts and, when identifiable, the perpetrating group or individual. We include only events with at least one fatality, that are definitely acts of terrorism according to the coding and that are attributed to a known organization. Finally, we use only events occurring after 1997 because the GTD coding procedures changed that year (see their codebook for more information). This leaves us with 13,859 terrorist attacks carried out by 57 groups between 1998-2015.

Results

We follow the same approach as Johnson et al. (2013) and fit model $Ms^{-\alpha}$ to the data using maximum likelihood estimation.[3] Let s denote the number of fatalities in an event. The power law coefficient is α and M is a normalisation factor such that the cumulative probability distribution sums to unity. Figure 1 plots the estimated α values for the African and Asian conflicts against the p-values of bootstrapped tests of the hypotheses that the data for particular conflicts are generated by the fitted power laws for these conflicts.[4] There are certainly some anomalous observations with α far from 2.5 or very low p values suggesting that the power-law hypothesis should be rejected. Low p values are not necessarily and important issue since, presumably,

[2]Note that for Afghanistan we split the UCDP's state-based data into two separate conflicts so that the fight after the beginning of Operation Enduring Freedom is treated as separate conflict.

[3]We use the "poweRlaw" package in R (Gillespie, 2014) which, in turn, is based on the work by Clauset et al. (2009).

[4]The reported p-values are based on bootstrap resampling using 1000 iterations.

no distribution of violent conflict events is, literally, generated by a pure power law so we would expect to always reject the power-law hypothesis with enough data even when it is useful to model the event-generating process as a power law. Estimated α's far from 2.5 are more of an issue. These results could come from data problems, e.g., not having enough data or some serious flaws in the data-gathering processes for particular conflicts, or it could be that these conflicts are fundamentally different from most other conflicts in the modern world.

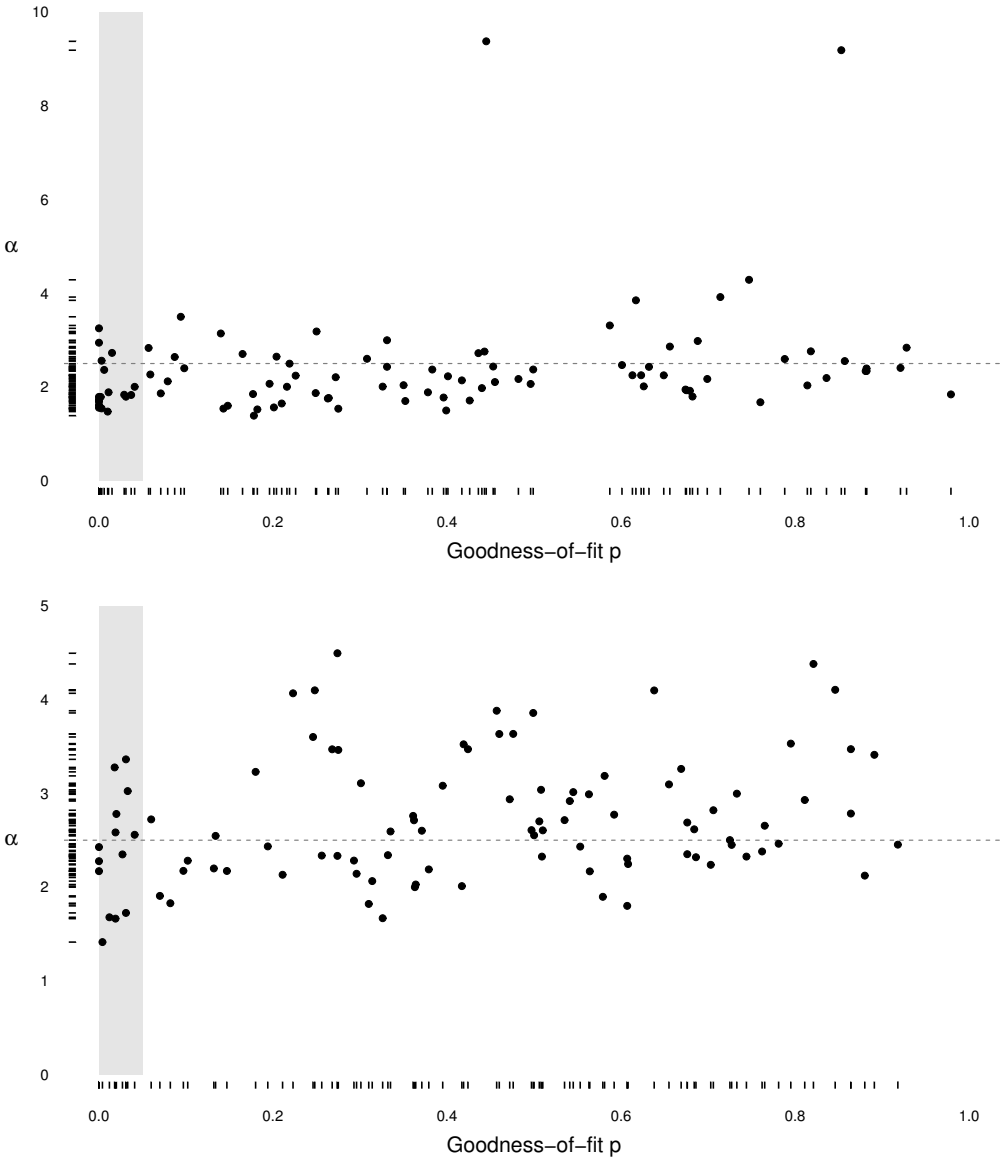

Figure 1: The estimated α parameter plotted against p-value as a measure for the goodness of fit for African (*top*) and Asian (*bottom*) conflicts.

Figure 2 provides the same p versus α information we gave for Figure 1 but for terrorist groups using the GTD data. It shows that the the distributions of violent events generated by terrorist organizations are also generally well fit by power laws with α that cluster around 2.5. Thus, there appears to be a close connection between terrorism and insurgency, at least regarding the processes for generating violent events.

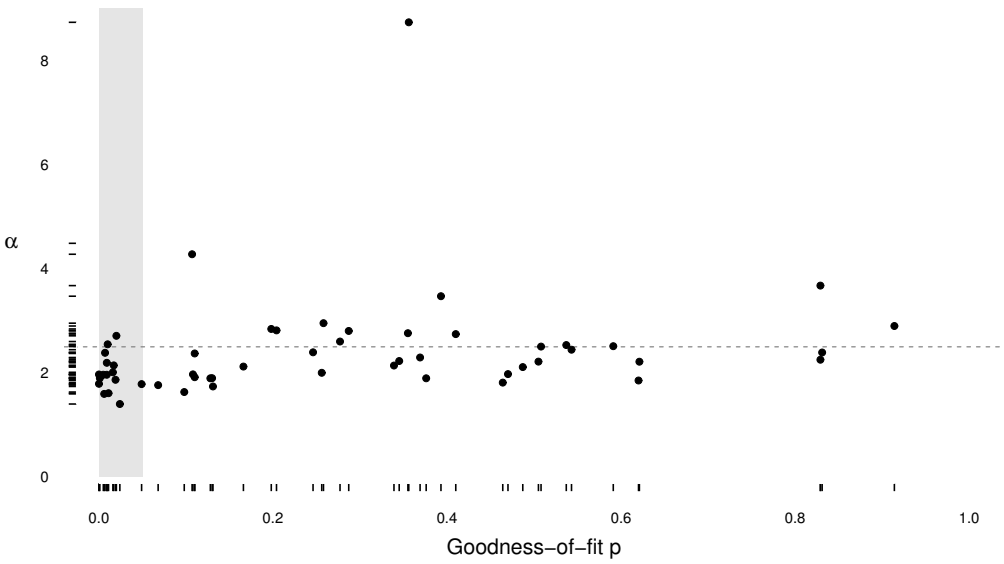

Figure 2: The estimated α parameter plotted against p-value as a measure for the goodness of fit for terrorism.

Figure 1 is consistent with previous empirical research finding, supported by theory, that power laws with α values clustering around 2.5 generally fit the empirical distributions of event sizes in modern wars (Bohorquez et al. (2009) and Johnson et al. (2013)). We exploit this general pattern in conflict dynamics to generate out-of-sample predictions for the expected ratios of event counts for a number of pairs of size ranges. We calculate the successes and failures of these predictions.

Specifically, we follow these procedures.

1. Randomly split the sample into two parts, using the selected one third of the conflicts to generate out-of-sample predictions for the remaining two thirds of the conflicts.

2. Fit power laws to the selected one third.

3. Order all the α estimates from smallest to largest and calculate the range running from percentile 2.5 to percentile 97.5.

4. Use the lower and upper bounds of this range to make predictions for the lower and upper bounds of the ratios of event-size counts for various ratios of event size ranges. For example, if the lower bound for α is 2.0 then the upper bound for the ratio of the number of events of size S or greater to the number of events of size $2S$ or greater is 2 while if the upper bound for α is 3.5 then the lower bound for the same ratio of event-size ranges is about $5.7S$. The corresponding figures for S and $1.5S$ are $1.5S$ and $2.8S$ respectively.

5. Check these predictions against the data for the two thirds of conflicts that were kept out of the randomly selected sample. Although we could check a near-endless list of predictions we confine ourselves to just the ratios for which we multiply the event size by either 1.5| or 2.0.

6. Start over, taking another draw of 1/3 of the conflicts and testing another set of out-of-sample predictions.

We repeat this procedures $1,000$ times.

Figure 3 displays the results of this simulation exercise. For most event-size ratios the success rates are above 60% for at least 75% of the $1,000$ draws. The best prediction performance is for the event-size ratios of 10/20 and 20/40 for which the median success rates are in the 80's and even the worst runs tend to score well above 60%. Unsurprisingly, the worst prediction performances are when the events are either very small or very large. The power laws are not even meant to apply below some level s_{min} so one could argue that, if anything, it is surprising that the predictions work as well as they do for the small events. For the large events the data are sparse so prediction is particularly difficult.

Figure 4 shows that out-of-sample predictions works almost as well as in-sample prediction for this scheme. The solid curve give the success rates when we use all data. We repeat the exercise using only in-sample data as well and compare the performance of the model with out-of-sample cross validation, as shown in figure 4 where the grey-shaded area indicates the 50% interval of the out-of-sample results.

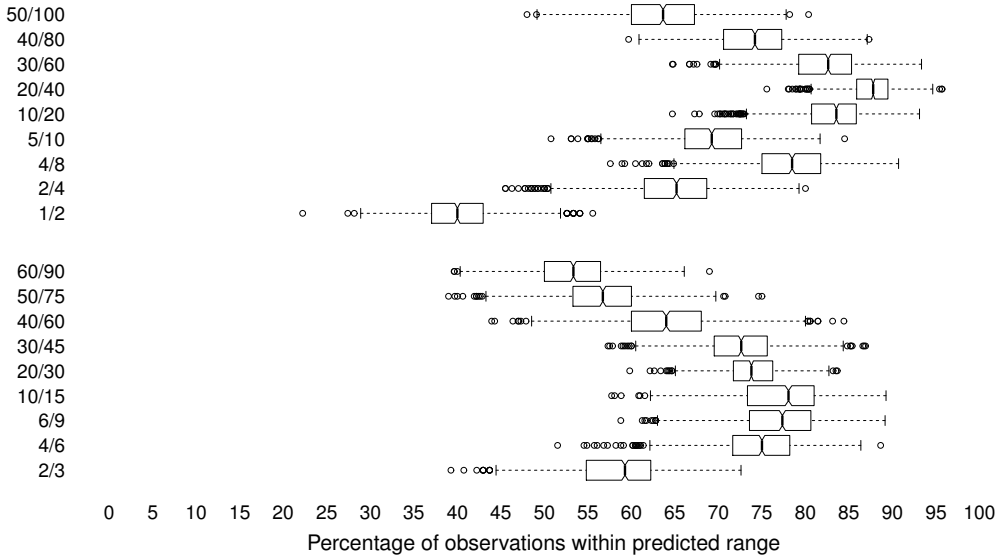

Figure 3: Boxplots illustrating the distribution of the percentage of out-of-sample observations that fall within the predicted range of fatality ratio A/B.

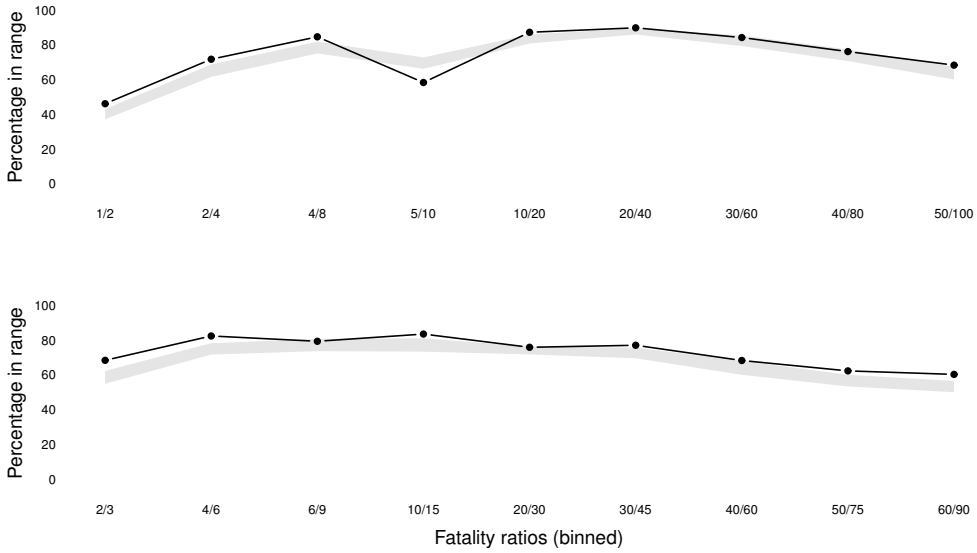

Figure 4: Number of fatality ratios within the theoretically predicted range (2-6) for in-sample data. The grey-shaded area indicates the 50% interval of the the out-of-sample results.

References

Bohorquez, J. C., S. Gourley, A. R. Dixon, M. Spagat, and N. F. Johnson (2009, dec). Common ecology quantifies human insurgency. *Nature 462*(7275), 911–4.

Clauset, A., C. Shalizi, and M. Newman (2009). Power-law distributions in empirical data. *SIAM Review 51*, 661–703.

Clauset, A. and R. Woodard (2013a). Estimating the Historical and Future Probabilities of Large Terrorist Events. *Annals of Applied Statistics 7*(4), 1838–1865.

Clauset, A. and R. Woodard (2013b). Rejoind of "Estimating the Historical and Future Probabilities of Large Terrorist Events". *Annals of Applied Statistics 7*(4), 1838–1865.

Eck, K. (2012, mar). In data we trust? A comparison of UCDP GED and ACLED conflict events datasets. *Cooperation and Conflict 47*(1), 124–141.

Gillespie, C. S. (2014). Fitting heavy tailed distributions: the poweRlaw package. (July), 1–16.

Johnson, N., M. Zheng, Y. Vorobyeva, A. Gabriel, H. Qi, N. Velasquez, P. Manrique, D. Johnson, E. Restrepo, C. Song, and S. Wuchty (2016). New online ecology of adversarial aggregates: ISIS and beyond. *Science (New York, N.Y.) 352*(6292), 1459–1463.

Johnson, N. F., P. Medina, G. Zhao, D. S. Messinger, J. Horgan, P. Gill, J. C. Bohorquez, W. Mattson, D. Gangi, H. Qi, P. Manrique, N. Velasquez, A. Morgenstern, E. Restrepo, N. Johnson, M. Spagat, and R. Zarama (2013). Simple mathematical law benchmarks human confrontations. *Scientific reports 3*, 3463.

Sundberg, R. and E. Melander (2013). Introducing the UCDP Georeferenced Event Dataset. *Journal of Peace Research 50*(4), 523–532.

Weidmann, N. B. (2013). The Higher the Better? The Limits of Analytical Resolution in Conflict Event Datasets. *Cooperation and Conflict 48*(4), 567–576.

Weidmann, N. B. (2015). On the Accuracy of Media-based Conflict Event Data. *Journal of Conflict Resolution 59*(6), 1129–1149.