Power-law Analysis of Structural Factors in International War

Miles D. Townes

CSS 625 Complexity in Social Science
6 May 2009

Abstract

Most attempts to explain Richardson's power-law result for international conflict focus on complex mechanisms to generate his result. The present study instead tests whether the power-law is a deterministic consequence of the structure of international system. By testing potential structural variables for power-law distributions, this study shows that at least one such variable is not distributed according to a power-law. Regression analysis of other variables suggests that Richardson's law is not a trivial consequence of those factors, either. Meanwhile, this study confirms Richardson's result for fatalities from international conflict, and shows that duration of conflict likely follows a similar power-law. The study also draws on the methodological literature and simulation techniques to demonstrate deficiencies with the procedures commonly used to analyze and test data for power-law distributions in small samples, and that robust regression provides more accurate results in these conditions. The study also demonstrates that appropriate test of goodness-of-fit remain a challenge to power-law analyses; much work remains to be done, both in regards to Richardson's law and to power-law analysis more generally. Future attempts to explain Richardson's law should focus on ideational and institutional dimensions of the international system.

1. Introduction

L. F. Richardson (1948) describes the relationship between the magnitude of a conflict and its frequency, specifically that the frequency of "fatal quarrels" decreases as the lethality of the incidents increases. In more technical terms, Richardson discovered an inverse **power-law** distribution between the size of violent incidents and their frequency: when the logs of frequency are regressed on the logs of size, the relationship is a straight line with negative slope.

With growing interest in **complexity** theory and its potential application to questions of social science, several recent authors have attempted to replicate or explain Richardson's power-law. These include Lars-Erik Cederman (2003), Cioffi-Revilla and Midlarsky (2004), Brunk (2000, 2001), Roberts and Turcotte (1998) (also Min et al, forthcoming). Most of these studies posit a variation on **self-organized criticality** as the basic mechanism of the power law, and these explanations are all "complex" in the technical sense of the word.

To date no study has looked at simpler potential explanations of Richardson's law. This is an important question: whether international conflict demonstrates **emergent** complexity, or is simply an artifact of underlying complexity deeper in the international system. More concretely, it may be the case that the power law is directly traceable to material facts about the international system, such as the distribution of army sizes, or national wealth. This paper explores some of these possibilities by looking at structural variables of conflict, to test whether they follow a power-law. If these variables do follow a power-law, it can be argued that Richardson's law is a trivial consequence of structural inputs.

Instead, I take as my research hypothesis that the structural variables do not follow a power-law. The two variables of interest are army size and population size. I also test duration, though this is not a structural input; if duration is also power-law distributed, it provides further evidence of the emergent complexity of conflict. My results replicate Richardson's power-law for casualties and also suggest that war duration follows a power-law as well; I demonstrate that army size does not follow a power-law, but that population size might. However, I also demonstrate that population size is not correlated with casualties, when other structural variables are taken into account. The results confirm my hypothesis, suggesting that the consequences of war, measured by fatalities, is evidence of emergent complexity in the international system.

2. Method of Analysis

For my data I use the COW dataset on international wars (Sarkees 2000). The data include a number of parameters, only three of which I consider “structural” for present purposes: army size, population size, number of participant countries. I would also have liked to include a measure of economic strength – GDP – but such data were not available when I conducted my analysis. I cleaned the data so that army size, population, and deaths all measured individual persons. Participants measure countries, and so are not directly comparable to individual person, on which more later. I also used COW data on war duration, measured in days. There were 79 international wars in the dataset, with missing values for army size in two observations and duration for one observation.

From this data, the following analysis depends on two key assumptions:

1. A variable which might directly generate Richardson's law will itself be power-law distributed.
2. Such a variable will correlate to the power-law of fatalities observation-by-observation.

The second assumption is important: it is entirely possible that data for a variable to be power-law distributed, but for that distribution to be unrelated to Richardson's law. However, power-law analysis is not the best tool to test for this possibility. Instead, regression across the variables will provide a better indication of correlations. So a brief regression analysis is the first step in this study –

not to test hypotheses, but to inform the hypotheses that are subject to power-law analysis.

As discussed above, this analysis begins with a linear model of the variables under consideration.

$$deaths = \beta_0 + \beta_1 population + \beta_2 armysize + \beta_3 duration + \beta_4 participants \quad (2.1)$$

Because testing for power-law distribution implies that the log-log distribution of data is more relevant than the level-level distribution, I also test a logged model for the data:

$$lndeaths = \beta_0 + \beta_1 lnpop + \beta_2 lnarmy + \beta_3 lnduration + \beta_4 lnpart \quad (2.2)$$

The literature on estimating the power-law itself is somewhat divergent. Many authors use OLS (for example, Cederman 2003), which Clauset et al (2007) and Alfarano and Lux (forthcoming) argue is not an accurate. Both of the latter studies recommend **MLE** methods based on the Hill (1975) estimator of the **Pareto index** α; Clauset et al offer a version of the estimator for discrete data, which I use in this analysis.

$$\hat{\alpha} \simeq 1 + n \left[\sum_{i=1}^{n} \ln \frac{x_i}{x_{min} - 1/2} \right]^{-1}$$

(2.3) (Clauset el al 2007; 5 eq. 22)

In addition to the potential for greater accuracy, this estimator also has the convenient property that it does not require an a priori calculation of the **complementary cumulative distribution function.** It does, however, require the choice of an x_{min} – the minimum value for which the power-law holds. Because the distribution for *deaths* holds a power-law across its range, x_{min} for each variable is assumed to be the smallest observation of that variable. (These values are indicated in the following section.)

However, Long (1997, 54) notes that MLE methods are valid only asymptotically, and should not be used for sample sizes smaller than 100. Given the sample size here – 79, at most – the the question remains: what is the best method for estimating the power-law parameter? To answer this question, it helps to examine the objections to OLS raised by Clauset et al:

> ...the estimates of the slope are subject to systematic and potentially large errors [...], but there are a number of other serious problems as well. First, errors are hard to estimate because they are not well-described by the usual regression formulas, which are based on assumptions that do not apply in this case. For continuous data, this problem can be exacerbated by the choice of binning scheme used to construct the histogram, which introduces an additional set of free parameters. Second, a fit to a power-law distribution can account for a large fraction of the variance even when the fitted data do not follow a power law, and hence high values of r^2 cannot be taken as evidence in favor of the power-law form. Third, the fits extracted by regression methods usually do not satisfy basic requirements on probability distributions, such as normalization, and hence cannot be correct. (2007, 31)*

Working backwards through these objections, the last is obviated by correct construction of the CCDF; an appropriately constructed CCDF will satisfy the basic requirements on probability distributions. Next, concerning r^2, Clauset et al (2007) provide alternative methods for testing the power-law distribution which do not depend on the method used to estimate the slope parameter. The same is true of error estimation: once the parameter is accurately estimated, the errors are easily calculated from that result. Indeed, the key issue seems to be OLS's sensitivity to error – especially outliers; this prevents OLS from returning accurate estimates of the parameter, and the rest of the problems follow from that deficit. However, this is problem is easily addressed through robust regression methods. At least three methods of robust regression – ie regression insensitive to outliers - are native to STATA 10 (for details, see the Stata manuals and online help files): `rreg`, `qreg`, `bsqreg`. The method `rreg` – hereafter referred to as robust regression – implements a method based on Huber iteration (see Belsley et al 1980, 233-234); `qreg` and `bsqreg` both implement versions of quantile regression.

As a sketch comparison of these methods, I first constructed a synthetic dataset with a known, fixed alpha parameter of 1.5, a minimum of x = 1000, a sample size of 100, and a corresponding CCDF. I use the following formula to generate the synthetic data:

* Given their concern about inaccurate standard errors, the authors of Clauset et al. (2007) suggest instead calculating standard errors directly from the estimate of α, choice of x_{min}, and the sample size (see equations 3.2 and 3.6, pp. 6-7). Because the estimates of error 1) follow trivially from parameters already under consideration and 2) are unnecessary to the present analysis, I do not report standard errors except in Tables 3.1 and 3.2

$$x = x_{min}(1-p)^{-1/(\alpha-1)}$$ (2.4) (following methods from Clauset et al 2007)

To generate a random dataset with the desired α, *p* can be replaced by a random number from a uniform distribution between 0 and 1. This method results in a dataset with a small amount of error in the parameter. Alternatively, *p* can be replaced by a frequencies from a fixed distribution – in this case *i/N*; there was no stochastic element in the generation of these data, so the dataset was free of error, allowing for accurate comparison of the estimates for alpha by the various methods available. Under these circumstances, OLS performed far better than MLE; all linear regression methods identified the correct value of the parameter, *b* = -.5 (where the Pareto index α = |b - 1|), but MLE returned α = 1.5166. This demonstrates the problems with MLE in analyzing small sample sizes.

To test the sensitivity of the estimators to outliers, I increased the magnitude of the largest observation from 10^7 to 10^8 – representing a single, large outlying error in the dataset. Under these circumstances, MLE outperformed OLS (see table). However, all three of the robust regression techniques proved better than MLE; all three correctly identified the underlying parameter as -.5 (ie 1.5). Similar test (not reported here) with N=400 and N=1000 show that MLE becomes more accurate for larger sample sizes. A decisive Monte Carlo simulation of the small sample size problem is beyond the scope of this paper, but for present purposes I assume similar results would be obtained from more extensive comparison.

Table 2.1 Estimation procedures with synthetic, outlier data*

Method	`reg`	`rreg`	`qreg`	`bsqreg`	MLE
α = 1.5, N=100	1.47103	1.5	1.5	1.5	1.5105743

*(largest observation increased by an order of magnitude)

This suggests the robust regression procedure native to Stata 10 is the best procedure for estimating α. An estimate of α is not, however, indicative of a power-law; the estimates supply no useful information about goodness-of-fit to the power-law distribution. The **Anderson-Darling** test has been suggested as a method to test goodness-of-fit, but the test for power-laws is not widely available in statistical packages and there are no useable tables of critical values to assess results compiled by hand.* Moreover, I found no instances in the literature of the Anderson-Darling test used to assess the goodness-of-fit for a power-law distribution. The test is a potential improvement on the Kolmogorov-Smirnov test, but some amount of work remains to be done before it is useable in power-law analysis.

The Kolmogorov-Smirnov (KS) test is still a useful indicator of goodness of fit; Clauset at al propose a version of KS test using Monte Carlo simulation to estimate the p-value of the empirical result. Having previously implemented these test in Stata 10, I will use it again here. This tests only whether or not the power-law distribution fits the data; it does not test whether the power-law distribution is the best fit to the data. In some cases, different distributions will provide better fits to data that nonetheless pass the KS test (see Clauset et al 2007 for more details); the test only indicates whether we can reject the hypothesis that the data fits the power-law distribution. (This is, of course, an inversion of the usual null-hypothesis, so the critical p-value will be 1-.05 = .95.). The results of these estimations and tests are reported below.

3. Results and Findings

As discussed above, the first step in this analysis is regression of the models in Eq. 2.1 and 2.2. The estimates (by robust regression) are reported in Tables 3.1 and 3.2. One surprise from these results is that neither armysize nor population are significant at the p=.05 level – not in level-level, nor in log-log. One possibility is that model is misspecified, that *armysize* and *population* are highly correlated. A test confirms the correlation (not reported - see Figure 2.1 for a scatterplot); to correct for

* I did not abandon the Anderson-Darling test for lack of trying. I wrote code for Stata to implement it, and verified my results through spreadsheet calculations. The problem was that I could not find a frame of reference by which to interpret my results and assign critical values.

this, I drop *armysize* from Eq. 2.1 and again perform the regression; estimates for this model are shown as Table 3.3. Again, *population* is no longer significant – nor is *participants*; the implication of this result I will defer for now.

Table 3.1 Robust regression of *deaths* on *population*, *armysize*, *duration*, and *participants*

	Coefficient	Std. Error	P > |t|	95% CI
Population	-.000016	7.39^-6	.031	-.000031 -1.52e-06
Duration	17.596	.5542	.000	16.48953 18.70233
Army Size	.0054	.0016	.001	.0026045 .008884
# Participants	2294.207	546.71	.000	1202.973 3385.442

Table 3.2 Robust regression of logged (ln) variables

	Coefficient	Std. Error	P > |t|	95% CI
Population	.0430	.2692	.874	-.4937871 .5797911
Duration	.7991	.1133	.000	.5730913 1.025059
Army Size	.2348	.2513	.354	-.2666617 .7354136
# Participants	.6632	.3349	.052	-.0046648 1.330974

Figure 3.1 Scatterplot of *lnarmy* and *lnpopsize*

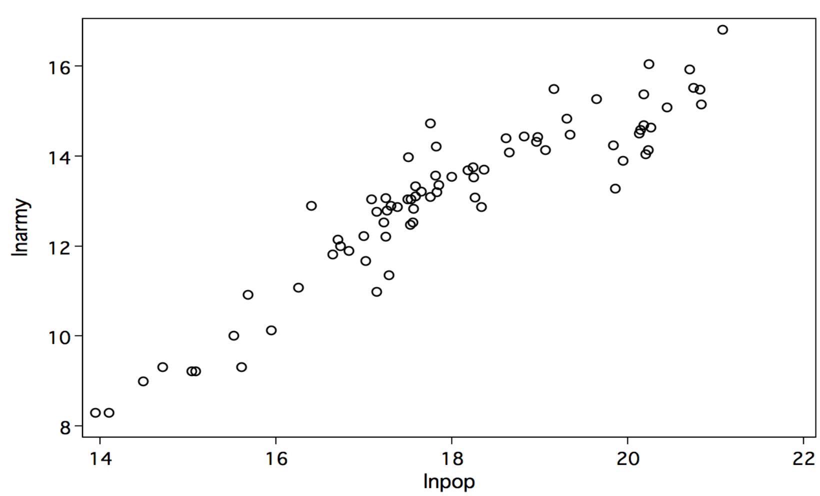

Table 3.3 Level-level model, with *armysize* removed

	Coefficient	Std. Error	P > |t|	95% CI
Population	4.17e-06	3.09e-06	0.182	-2.00e-06 .0000103
Duration	17.94588	.3958178	.000	17.15665 18.73512
# Participants	-218.6371	401.1436	0.587	-1018.495 581.2206

Based on these results, I predict that *duration* will follow a power-law distribution, but that *army* and *population* will not (*participants* covers a range from 2 to ~30, insufficient for a power law; it is not tested). The first step is to estimate the alpha parameter; the results of these estimates by OLS, robust regression, quantile regression, and MLE are reported in Table 3.4. To better visualize the results, a table of scatterplots of logged data appears as Figure 3.2.

Using the estimates of α from robust regression (rreg), I next perform a KS test on each of the distributions, including a p-value estimation via the Monte Carlo simulation. I include a synthetic data set in the process as means of comparison. The results are shown as Table 3.5. For *armysize* and *duration*, the Hill estimator failed to converge – suggesting these data are not power-law distributed. However, the fact that the estimator did converge for the other two variables does not in itself indicate that *deaths* and *populations* are power-law distributed.

The KS test provides strong evidence that *army* size is not power-law distributed; however, it fails to reject any of the other datasets. That *duration* has the best (i.e. lowest) p-value of any of the distributions except the synthetic data is unexpected; based on the Hill estimator result, *duration* seemed unlikely to pass the KS test. A visual inspection of the scatterplot nonetheless confirms that *duration* appears to be power-law distributed, and least for a large portion of the right-side tail.

Though these results appear to disconfirm the research hypothesis, there a few problems remaining. First, if size of the population brought to bear on a conflict is only weakly correlated with

Table 3.4 Estimates for α from COW Inter-state conflicts data

	reg	rreg	qreg	MLE
Deaths	1.4073	1.3802	1.3796	1.3996
Population	1.4808	1.4863	1.4664	1.0902
Army Size	1.4065	1.3076	1.3976	---
Duration	1.515	1.4979	1.538	---

Figure 3.2 Scatterplots of logged data for *deaths*, *armysize*, *duration*, and *population* with prediction lines derived from robust regression estimates of *b*

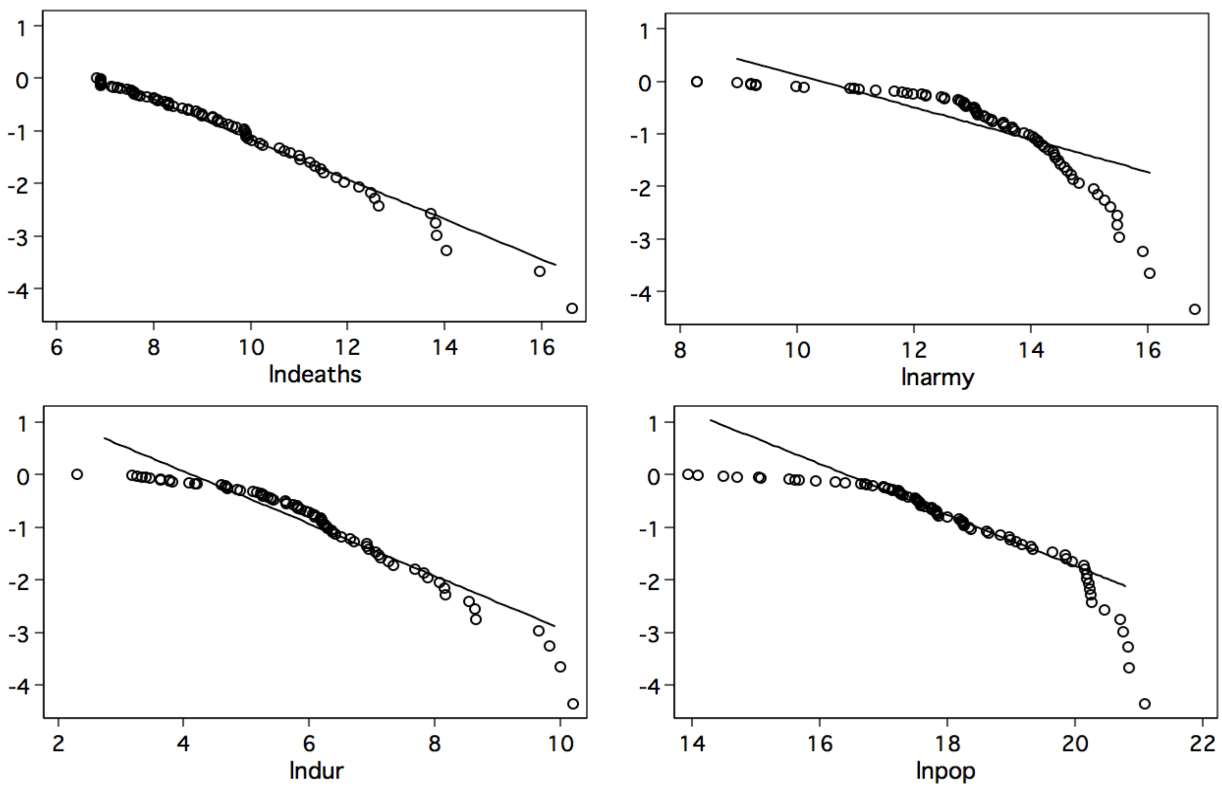

Table 3.5 Kolmogorov-Smirnov tests for power-law distribution

	$\hat{\alpha}$	xmin	n	KS	p-value*	reject power-law? (ie. P>.95?)
Synthetic data	1.5	1000	100	.99	.6504	No
Deaths	1.3802	1000	79	.9752	.91	No
Population	1.4863	1000000	79	.9708	.9388	No
Army Size	1.3076	4000	77	.9524	.994	Yes
Duration	1.4979	10	78	.9804	.7824	No

*Calculated by procedure detailed in Clauset et al 2007.

the consequent fatalities; even in a model (not reported) stripped of all variables except *deaths*, *duration*, and *population*, the coefficient for the last of these is neither substantively or statistically significant. So it may well be that the data for population sizes are power-law distributed, but the fact that these observations appear to be uncorrelated with those of fatalities makes it implausible that population size is a structural factor in conflict severity. Second, if duration itself is another measure of conflict severity is in fact power-law distributed, this too supports the hypothesis that conflict severity is an emergent property of international conflict. This suggests two further models in need of testing:

$$deaths = \beta_0 + \beta_1 pop + \beta_2 army + \beta_3 participants \quad (3.1)$$

$$duration = \beta_0 + \beta_1 pop + \beta_2 army + \beta_3 participants \quad (3.2)$$

The results of these models are reported in Tables 3.6 and 3.7; both are significant at the $p > .05$ level. The two sets of results are, however, somewhat contradictory if duration and deaths are both understood as measures of war severity. In the first model, population and army size are significant factors in severity of the conflict, where the number of participant nations is not. In the second model, these factors are reversed. The reasonable inference is that while both duration and deaths measure conflict severity, they are not equivalent measures. However, these results do not lead me to reject the hypotheses that conflict severity is not a trivial result of structural inputs; in fact, these results raise more questions than they answer.

Table 3.6 Robust regression of structural variables on *deaths*

Deaths	coef	Std. error	P>|t|	95% CI
Population	-.0000187	-6.95e-06	.009	-.0000326, -4.83e-06
Army Size	.0054	.001425	.000	.0025598, .0082415
Participant (nations)	840.55	564.8148	.062	-285.6624, 1966.755

Table 3.7 Robust regression of structural variables on *duration*

	coef	Std. error	P>|t|	95% CI
Population	-8.73e-08	2.55e-07	.733	-5.96e-07 4.21e-07
Army Size	2.20e-06	.0000478	.963	-.000093 .0000974
Participant (nations)	46.43	20.24	.025	77.38034 347.1769

4. Discussion

The results in the previous section present a complex – for lack of a better word – picture of international conflict as a function of structural inputs. Of the two variables tested as structural inputs to Richardson's law, only population could be considered to follow a power-law distribution. However, the analysis offers contradictory evidence of correlation between the remaining variables. It remains implausible that Richardson's law is a trivial consequence of underlying structural parameters in the international system.

4.1 Discussion of Findings

Of the variables tested, army size is definitively not power-law distributed; this is the only clear conclusion available in these results. However, the fact that army size is not power-law distributed may be due in part to the fact that "army size" as measured by the Correlates of War project only indicates the pre-war size of the army, not the size of forces brought to bear over the course of the conflict. The sizes of armies one or two years into a major conflict are likely substantially larger than the standing forces prior to the outbreak of war, and probably more closely proportional to the overall population size. However, this does not meant that population size is the most structurally important factor determining the number of casualties in a conflict. When the OLS model controls for duration and number of participant nations, population size is not a significantly correlated to the number of deaths resulting from conflict. In fact, robust regression suggests that only duration is significantly correlated to the level of casualties – but duration is itself better understood as a consequence, not an input.

Regarding methodology, these results demonstrate quite clearly the deficiencies of the various methods typically used in the literature to estimate power-law parameters. OLS regression is too sensitive to outliers; MLE is less sensitive, but not valid for small sample sizes. That robust regression techniques correctly estimate the Pareto index for small sample sizes is perhaps the most substantively useful finding of the present study.

4.2 Discussion of Broader Implications

Richardson's power-law remains a fascinating and challenging puzzle to the study of international conflict; despite worthy efforts, it still has not been "solved" definitively. Meanwhile, the fact that war severity assumes a power-law distribution hints at deep and pervasive complexity in the international system. The question remains: from whence does this complexity emerge?

In this analysis I have suggested (but, granted, not proven) that complexity is not directly consequential to structural and material facts of the international system. The types of data we have about the material dimensions of international conflict do not provide an obvious or satisfying explanation for Richardson's law. We might move closer towards that explanation by improving our data, or our analytic techniques (on which more in the next subsection), but it is more likely that complexity in international relations is not the result of material facts about the international system.

Richardson's original result is evidence of this, as are several recent replications or extensions of his work (including several by the present author, unpublished). In these studies the power-law result is shown to extend to sub-state conflicts, including civil wars, that cannot be assumed to be manifestations of the international system. Richardson's assumption was that all anthropogenic violence is the result of a similar, undescribed process – from one-off murders up to World Wars. However, there are good reasons to think that violent conflict between political entities is a very different phenomenon from interpersonal violence, including Richardson's own results: he was unable to show that interpersonal violence followed a power law, and his result only pertains to incidents with casualty levels greater than 10^3 - that is, the same minimum threshold set by COW for their dataset on international conflict. If Richardson's power-law is a solely a phenomenon of political violence, then its explanation might lie in politics: namely, the institutional and ideational nature of political interactions. This is non-material dimension of the international system, and one for which there are few adequate measures. Complexity in international politics may be emergent from the intellectual and emotional activities that make up political interaction – not the material resources available to support those activities. The preliminary regressions in Tables 3.1, 3.2, 3.3, 3.7 support this – there is some significance to the number of nations involved in conflict, apart from the number of people in those nations. If nation-states only mattered as containers of populations, I would expect the opposite result. Instead, the present results show a need for complexity theorists to engage politics as politics, not merely agglomerations of material facts.

4.3 Implications for Future Research

That politics is more than is more than material facts does not mean there is no need for improvement or extension on the present study. Indeed, this work suggests several further steps for research on Richardson's power-law:

1. Conduct a Monte Carlo simulation of robust regression versus MLE techniques, to ensure that the estimation methods produce the best possible estimate of the Pareto index, α. The

accuracy of $\hat{\alpha}$ is important both for interpretation, but also for calculating standard errors and confidence intervals.

2. Add data on other structural parameters to the model, especially GDP or some other indicator of economic capacity. Economic capacity is relevant because we know it has played a critical role in major wars; the ability to sustain a modern army in combat requires significant economic mobilization, revenue, and manufacturing capacity. That such data is not included here is a serious deficit, and prevents the conclusions from being more definitive.

3. Analyze fatalities and duration as two dependent variables in a simultaneous-equations model. These two variables are probably consequences of a single process, and should be analyzed as such. Specifically, this would entail a recursive model of conflict severity, in which conflict duration is both a dependent variable and an independent variable affecting the number of fatalities in a conflict. Such an analysis is entirely feasible, but beyond the scope of this study.

4. Define and validate a method for applying the Anderson-Darling test to the present data, in order to provide a more sensitive test of goodness-of-fit. As these results show, the Kolmogorov-Smirnov test is too sensitive to type II error; it fails to reject distributions that probably do not in fact follow a power-law distribution. The Anderson-Darling test is not a replacement for the KS test, but rather a complement.

5. Test the data for *deaths*, *duration*, and *population* against other possible functional distributions, to determine whether the power-law distributions offer the best fit to the data. The present tests only indicate whether or not the data might follow a power-law; they do not, however, indicate whether other distribution functions provide a better fit. I suspect that population size in particular is better fit to a different distribution. Methods exist to test these other distributions, but they are computationally intensive.

5. Summary

Recent attempts to explain Richardon's power-law for “deadly quarrels” have focused on complex mechanisms to generate his result. The present study, however, sought to test whether that result is in fact more deterministic given the structure of the international system. Testing other major structural variables in the international system for power-law distributions led me to reject only one variable, army size, as definitely not following a power-law. However, this study also confirmed Richardson's result for fatalities from international conflict, and showed that duration of conflict likely follows a similar power-law. Taken as a whole, the results suggest that complexity in international conflict is not the consequence – at least, not a trivial consequence – of physical and structural aspects of the international system. Future attempts to explain Richardson's law should focus on ideational and institutional dimensions of the international system.

In the course of this study, I have also demonstrated some deficiencies with the methodologies used to analyze and test data for power-law distributions. OLS and MLE are both problematic for small sample size; robust regression provides more accurate estimates of alpha. Meanwhile, goodness-of-fit remains a challenge to power-law analyses; much work remains to be done, both in regards to Richardson's law and to power-law analysis more generally.

Note on methods: many of the procedures used here – including the MLE estimation algorithm, the KS test and p-value procedure, and the synthetic data generation procedure - are available as Stata .do files at http://home.gwu.edu/~mdtownes/MLE/ .

Bibliography

Belsley, D.A., E. Kuh, and R.E. Welsch. (1990). Regression Diagnostics: Identifying Influential Data and sources of collinearity. (NY; Wiley)

Brunk, G.G. 2001. "Self-organized criticality: a new theory of political behaviour and some of its implications". *British Journal of Political Science*, 31 (2001); 427-445.

Brunk, G.G. 2002. "Why do societies collapse?: A theory based on self-organized criticality". *Journal of Theoretical Politics*, 14:2 (2002); 195-230.

Cederman, L-E. (2003). "Modeling the Size of Wars: From Billiard Balls to Sandpiles". *American Political Science Review*, 97:1 (February 2003), pp. 135-150.

Cioffi-Revilla, C. and M. Midlarsky (2003). "Highest Magnitude Warfare: Power Laws, Scaling, and Fractals in the Most Letha International and Civil Wars". In Diehl, Paul, ed. (2004). *The Scourge of War*. (Ann Arbor; U. Mich. Press).

Clauset, A., C. R. Shalizi, and M.E.J. Newman (2007). "Power-law distribution in empirical data". arXiv:0706.1062v1 [physics.data-an] 7 Jun 2007.
-see also Clauset's website, http://www.santafe.edu/~aaronc/powerlaws/

Levy, J. and C. Morgan. 1984. "The frequency and seriousness of war: an inverse relationship". *Journal of Conflict Resolution*, 28:4 (December 1984); 731-749.

Long, J.S. (1997). Regression models for categorical and limited dependent variables. (Thousand Oaks, CA: Sage)

Newman, M.E.J. (2005). "Power laws, Pareto distributions and Zipf's law". *Contemporary Physics*, 46:5 (Sept-Oct 2005), pp. 323-351.

Richardson, L.F. (1948). "Variation of the Frequency of Fatal Quarrels with Magnitude". *American Statistical Association Journal,* (December 1948), pp. 523-546.

Roberts, D.C. and D.L. Turcotte. 1998. "Fractality and self-organized criticality of wars". *Fractals*, 6:4 (1998); 351-357.

Sarkees, M. R. (2000). "The Correlates of War Data on War: An Update to 1997," *Conflict Management and Peace Science*, 18/1: 123-144.