DRAFT: DO NOT QUOTE, CITE, OR CIRCULATE WITHOUT AUTHOR'S PERMISSION

Miles Townes mdtownes@gwmail.gwu.edu

Estimating Richardson's Power-law

Policymakers and citizens alike often wonder as to the timing of the next major war. Beneath this wondering lies fear – of being caught unprepared, of being attacked, of being killed - so there is no small relief in any effort that suggests a pattern or predictability to international conflict.

One of the more striking empirical patterns of conflict, described by L. F. Richardson (1948), holds that the frequency of "fatal quarrels" decreases as the lethality of the incidents increases; obvious as this seems, the relationship is suprisingly regular. 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. Richardson determined the slope of this line "by the principle of Maximum Likelihood" (1948; 532); his graph of the line is shown here as Figure 1 (the straight line portion between magnitudes 3 and 8 – ie. 10^3 and 10^8 casualties - is circled; the power-law does not seem to hold at less than magnitude 3).

Despite its apparent importance, relatively little effort has been made to test or reinterpret Richardson's law – although it is occasionally noted in the literature (e.g. Singer and Small 1972; Zinnes 1976). More recently, Lars-Erik Cederman (2003) has replicated and modified Richardson's result, as have Cioffi-Revilla and Midlarsky (2004). Both studies, however, use OLS regression to estimate the parameters of the power-law (see equation 1, below).

Meanwhile, a number of scholars outside of political science have developed new methods to identify and assess power-laws – most notably, Newman 2005 and Clauset, Shalizi, & Newman 2007. The latter study is particularly useful: it suggests better methods to estimate, test, simulate, and interpret power law distributions, based on the method of maximum likelihood. In this paper, I will

Miles Townes mdtownes@gwmail.gwu.edu

apply those methods to revisit Richardson's law and its implications for violent conflict. I begin with a demonstration of the accuracy gained via MLE over OLS. I then estimate the parameters for Richardson's power-law, and test whether a power-law distribution in fact offers the best fit. Finally, I test for variations in the parameters of Richardson's power-law across various subcategories of war. I demonstrate that Richardson's power-law result persists, but also expand that result to offer interpretation and extension that suggests Richardson's result is not simply a mathematical curiosity, but a fundamental characteristic of systems of conflict.

Figure 1. Richardson's graph of the relationship between frequency and magnitude of deadly quarrels (with the power-law region circled in grey dashed line).

Miles Townes mdtownes@gwmail.gwu.edu

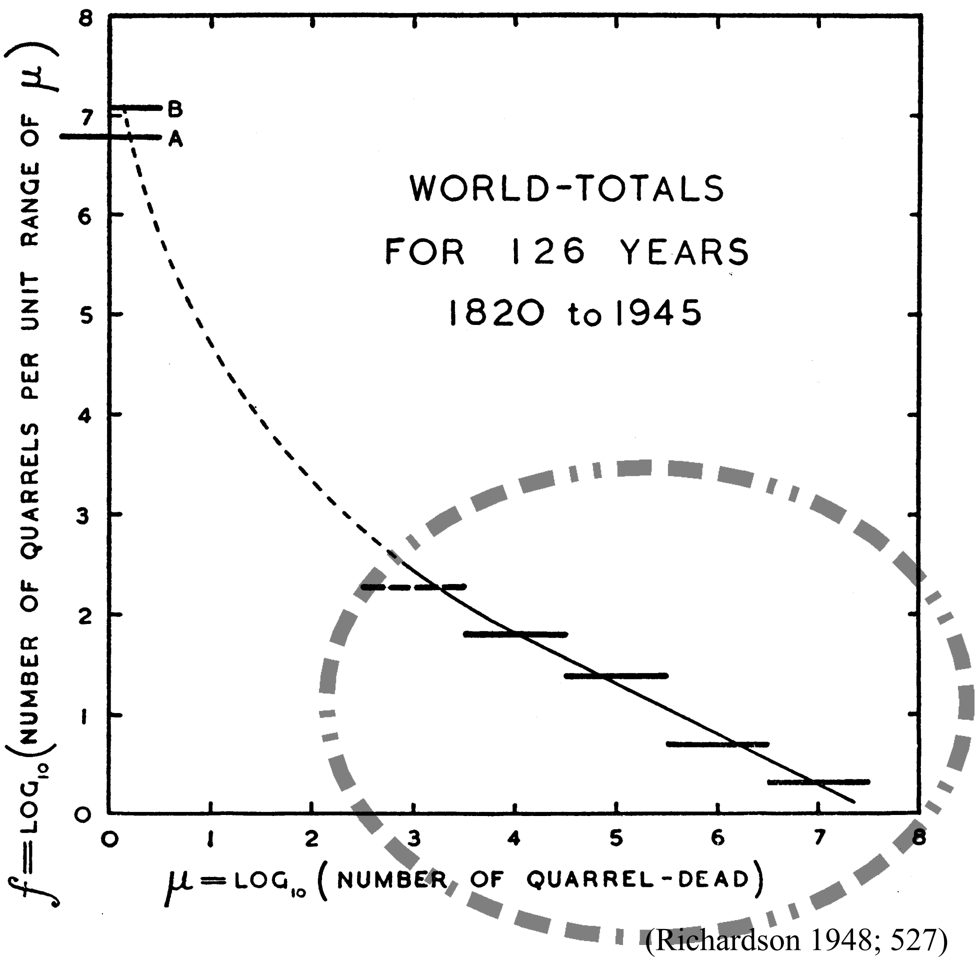

(Richardson 1948; 527)

I. Data Sources

I use the Correlates of War data on conflict; specifically the datasets for Interstate, Extra-state, and Intra-state wars (Sarkees 2000). The three datasets cover roughly the same time span: 1823-1990 for Interstate wars, 1816-1997 for Intra-state wars, and 1817-1975 for Extra-state wars. The last case is anomalous because it refers to colonial wars, a form of conflict that ceased to exist when the last major colonies gained independence. (In this respect, diplomats and policymakers can claim to have completely eliminated war – of a sort.) There are 79 observations in Interstate wars, 108 in Extra-state wars, and 213 in Intra-state wars.

Miles Townes mdtownes@gwmail.gwu.edu

The datasets combine to a total 400 observations, with no missing values for deaths; in this study, I refer to deaths as the variable of interest, "magnitude". There are 19 observations of conflicts less than 1000 deaths, which I drop from analysis to leave 318 observations in the total dataset. Interstate wars has 1 such observation, Extra-state 10 such observations, and Intrastate wars has 8. This leaves 78 observations for Interstate wars, 98 observations for Extra-state wars, and 205 observations for Intra-state wars. The decision to drop observations less than 1000 is based on two considerations: First, COW coding protocols define "war" as any conflict resulting in more than 1000 casualties. Observations with fewer than 1000 deaths are therefore anomalous. Second, the power-law estimation procedure requires a lower bound, x_{min}, which can be determined through analytical methods but may also be assumed for theoretical or practical reasons. In this case, I take on x_{min} as 1000, following the same assumptions as COW coding practices; in so doing I lose less than 5% of available data. (Note that there are a large number of observations coded at 1000 exactly; this may be evidence of a form of censoring among the data coders – that is, violent conflicts with unknown but significant casualties are coded at 1000 casualties by default. This is not the problem for power-law estimation that it is for OLS models based on a different distributional form, for reasons somewhat beyond the scope of this paper.)

Richardson (1948) assumed that all "deadly quarrels" were of a piece – that they were the product of a process intrinsic to human society, regardless of whether the incident was called a war, civil war, riot, or murder. Though I might use this assumption to combine the three datasets into one, it is instead preferable to test whether they admit such combination; IR theory suggests that there are in fact important differences between civil, colonial, and international wars. The Kolmogorov-Smirnov (KS) test for two-samples permits me to do so. The KS statistic measure the maximum difference in frequencies between two distributions – in STATA , this comes with a p-value of significance. The null hypothesis is that two samples are drawn from the same underlying distribution; a high KS statistic (and small p-value) suggests that we can reject the null hypothesis. The results of pairwise KS tests of the three datasets are shown in Table 1. As the table demonstrates, Interstate and Intra-state (that is, international and civil) wars pass the test; they appear to be drawn from the same distribution. Extra-state (colonial) wars fail the null hypothesis for both Interstate and Intra-state conflicts. Although I suspect this has more to do with data collection and coding than with empirical fact, I nonetheless proceed as though Extra-state Wars were distinct from the other conflict types. I create a fourth dataset – Int-wars – merging both Inter- and Intra-state Wars into a single whole. Where possible in my estimation and testing, I will test all four resulting datasets: Interstate Wars, Intra-state Wars, Extra-State Wars, and Int-Wars.

Table 1. Pairwise two-sample KS tests for distribution, COW datasets

	Intra-state Wars	Extra-state Wars
Interstate Wars	.1054 p = .556	.2928 p = .001
Intra-state Wars	--	.2270 p = .002

II. Methods

A power-law is formed when a quantity x is drawn from a probability distribution

$$p(x) = Cx^{-\alpha} \quad \text{(1) (Newman 2005;1 eq. 1)}$$

This relationship can also be described as $ln\ p(x) = -\boldsymbol{\alpha}\ ln\ x + ln\ c$, which forms a straight line on a plot

Miles Townes mdtownes@gwmail.gwu.edu

drawn on log-log scales (in OLS regression, this is the model estimated). Thus the power-law region of Richardson's plot (Fig. 1) begins and roughly magnitude 3 and continues through magnitude 7 – a straight-line with decreasing slope. (Incidentally, it does not matter whether the histogram is drawn in base-10 logarithm or natural logarithm, except that the axes increments are different; the line is still straight either way for a power-law distribution).

The parameter of interest in this distribution is $\boldsymbol{\alpha}$ – the slope of the power law. Clauset et al (2007) provide estimators for both discrete and continuous data; in this case, the data are discrete, so I will use their approximation

$$\hat{\alpha} \simeq 1 + n \left[\sum_{i=1}^{n} \ln \frac{x_i}{x_{min} - 1/2} \right]^{-1}$$ (2) (Clauset el al 2007; 5 eq. 22)

The authors suggest that this offer "quite good results" when x_{min} > ~6 (Clauset et al, 2007: 6); given that my value for x_{min} is 1000, I am reasonably confident in the approximate estimator. That my x_{min} is chosen a priori also permits straightforward calculation of the standard deviation by the equation

$$\hat{\sigma} = (\alpha - 1) \frac{n}{(n-1)\sqrt{n-2}}$$ (3) (Clauset et al 2007;

24)The choice of x_{min} and the estimate of alpha allows me to compute a cumulative distribution:

$$P(x) = \int_x^{\infty} p(x')\,dx' = \left(\frac{x}{x_{min}} \right)^{-\alpha+1}$$ (4) (Clauset et al 2007; 2, eq. 7):

These are the basic tools necessary to describe and interpret the power-law distributions in my dataset. I implement them in STATA 10, using hand calculations as necessary to verify accuracy. (STATA .do files are available on my web page; see Resource Note for more information.)

OLS vs. MLE

Given that equation 1 has a straightforward linear form, many students of power-law distributions are tempted to use ordinary least squares (OLS) regression to estimate the relevant parameters. However, power-law distributions violate key assumptions of OLS regression, with the consequence that OLS introduces biases into the estimates.[1] A full discussion of this problem is beyond the scope of this paper (and would be redundant to the discussion in Clauset et al, 2007 – see especially Appendix A). Instead, I present a brief example to compare OLS versus MLE methods for estimating the parameters of Richardson's Law.

Among those studies using OLS to replicate Richardson's Law, L-E Cederman's (2003) is perhaps most notable (see also Cioffi-Revilla and Midlarsky 2004). Although this paper is not intended primarily as critique of Cederman, it is worth engaging his treatment of power-law estimation as example of "how not to", as Cederman also uses COW data for Inter-state Wars. In particular, it might be instructive to replicate his result by OLS and compare that to an MLE estimation. However, although the two methods would give different results, it would not then be clear which result was more accurate. A better approach would be to use data with a known distribution, and compare the subsequent results.

To that end, I generated a simulated dataset with a known power-law distribution, using methods suggested by Clauset et al (2007). The function for generating the data is:

$$x = \left\lfloor r(x_{min} - 1/2)(1 - r)^{-1/(\alpha-1)} + 1/2 \right\rfloor \quad \text{(5) (Clauset et al, 2007; p. 40 D.6)}$$

where r is a random uniform variable distributed between 0 and 1. This data generation process allows me not only to compare MLE to OLS regression, but also to test the accuracy of my estimator. In STATA,

1 In particular, the assumptions that errors are independent and homoskedastic are compromised, resulting in biased estimates for the variance. Another significant problem is that the power-law is a type of probability distribution, and thus constrained to values between 0 and 1; there is little available literature on procedures to impose these constraints in OLS regression. Finally, OLS is highly sensitive to outliers – which are a specific problem in the COW datasets.

Miles Townes mdtownes@gwmail.gwu.edu

I produced data with $\alpha = 1.5$ over 1000 observations. The MLE estimation procedure allows the estimate to be produced with no further effort. OLS, however, requires a dependent variable; in this case, the rank-frequency was determined for each observation as $P(M \geq m) = (N - n)/N$, where N is the number of observations in the sample and n the rank of the observation. Thus the line being estimated is

$$\log P(M > m) = \log C - \alpha \log m \qquad (q)$$

(Note that log in this case is not specific to base-10 or natural logarithms.) This is the same procedure used in Cederman 2003. The result and comparison to MLE is shown as Table 1.

Table 1. Comparison of simulated data to MLE and OLS estimates

	simulated	MLE	OLS
estimate for α	-1.500	-1.5003	-.4955
standard error	--	.0158	.0006
95% C.I.	--	(-1.469, -1.531)	(-.4967, -.4943)

(note: negative signs have been added to the simulated and MLE output for convenience; both procedures assume α is negative and do not require or report it as such.)

As the table shows, OLS regression produces a very different estimate from the MLE. There are two problems that account for the difference: first, OLS does not in fact estimate α, but instead the slope of the distribution function, $-(\alpha - 1)$; the OLS value of α can be obtained simply by subtracting 1 from the estimate, in this case -1.4955. However, even correcting for the unit difference in the estimates, note that the OLS estimate is still biased.[2] Though the difference appears small, perhaps even negligible, it bears remembering that these parameters are used exponents on data that include large numbers – e.g. 10^6 and greater. Furthermore, the 95% confidence interval for OLS (-1.4967, -1.4943)

2 Further research shows that the source of bias is entirely due to the presence of outliers, and that robust regression (using the Huber objective function) is much more accurate than OLS regression.

excludes the value of α = 1.5, per the simulated data. (Note that the slight error in the MLE estimates is due primarily to variation generated by the data simulation process; this error generally is less than 1% for sufficiently large simulations.)

With this result in mind, I now turn to the specific estimates Cederman provides for Richardson's Law. Table 2 shows a comparison of Cederman's reported result, my OLS replication of same, my MLE replication in log10, and my replication in natural log. As shown, Cederman's result differs from the MLE estimates. There is an additional reason for this, beyond the inaccuracy of the OLS method: namely that Cederman uses log 10, as opposed to natural logarithms. This primarily affects his standard errors and confidence intervals. In any case, MLE is clearly preferable to OLS for the present purpose.

Table 2. Comparison of Cederman's method with MLE

	Cederman	OLS replication	MLE ($\log_{10}$)	MLE (ln)
estimate for α	-.41	-.41	-1.40	-1.40
standard error	--	.0062	.0449	.0682
95% C.I.	--	(-.4195, -.3951)	(-1.3114, -1.488)	(-1.256. -1.533)

III. Estimating Richardson's Law

Using the four datasets and syntax implementing the discrete approximation, I estimated the values for α as shown in Table 3. The STATA implementation includes standard errors and confidence intervals, but hand calculation demonstrates that these are not equivalent to the results provided by equation 3, above. Although the differences are slight – .0458 from hand calculation vs .0447 from STATA – Table 3 reports the results from hand calculation in the hope that more accuracy now will lead to better interpretation later. (I intend to improve my STATA implementation so that it returns correct

Miles Townes mdtownes@gwmail.gwu.edu

standard errors.)

For this implementation, STATA does not provide a goodness-of-fit estimate or a p-value for that fit; Clauset et al (2007) suggest the KS statistic for the real versus expected distribution (derived from the estimate of α and equation 5 above), which is also shown in Table 3. However, Clauset et al caution that this does not provide an appropriate p-value for the goodness-of-fit:

> ...there is no known formula for calculating the p-value, but we can still calculate it numerically by the following Monte Carlo procedure. We generate a large number of synthetic data sets drawn from the power-law distribution that best fits the observed data, fit each one individually to the power-law model using the methods of Section III, and calculated the KS statistic for each one relative to its own best-fit model, and then simply count what fraction of the time the resulting statistic is larger than the value D observed for the true data. This fraction is our p-value. (p. 11)

The results of this Monte Carlo procedure are included in Table 3.

As Table 3 shows, each of the datasets is well-fitted to a power-law, although Intra-state wars significantly more so than the other COW datasets. (Note that the apparent difference between this estimate for Interstate Wars and that from Table 2 is due to the dropped observation – the Falklands War, at 910 deaths – which changed the estimate of α from 1.39963 in Table 2 to 1.39438 in Table 3.) The plots of the four datasets – the log of magnitude against the log of frequency, by equation 5 above – are shown as Figure 2. The plots make clear the difference between α for Extra-state Wars and the other datasets.

Miles Townes mdtownes@gwmail.gwu.edu

Table 3. Comparison of datasets using MLE estimation[1] for α

	Interstate Wars	Extra-state Wars	Intra-state Wars	Inter- & Intra-state Wars
estimate for α	1.39	1.76	1.47	1.45
standard error[2]	.0458	.0763	.0332	.0267
95% C.I.	(1.305, 1.484)	(1.602, 1.928)	(1.405, 1.535)	(1.394, 1.473)
KS statistic	.9783	.9866	.9720	.9853
p-value[3]	.8548	.7668	.9992	.992
n	78	98	205	283

[1] all estimations were performed using the natural log variant of the estimator, for all obs > 1000

[2] standard errors and C.I.s are calculated by hand; STATA's native result is slightly inaccurate. .It is not clear what method the program uses to calculate standard errors.

[3]the null hypothesis is that the data are distributed according to a power-law; high values (such as those here) suggest the distribution does follow a power-law.

Figure 2. Power-law Plots for All Datasets

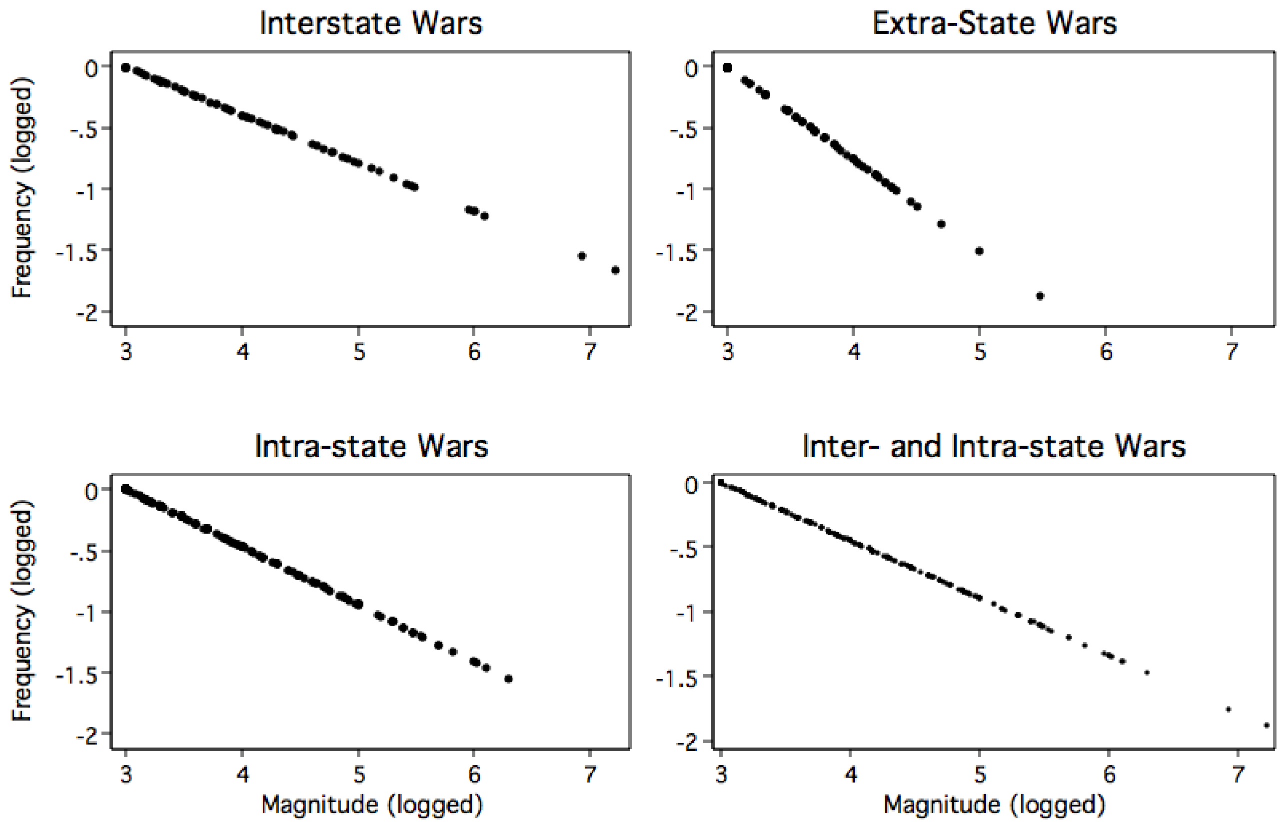

IV. Interpretation

Now satisfied that I can confirm Richardson's Law, I turn to the question of interpretation: what does it mean that conflict is distributed according to a power-law? The most obvious consequence is the suggestion of systematic relationship between the smallest and the largest conflicts – even across the theoretic categories of civil and international war. Where the security literature tends to treat large conflicts as epiphenomenal, Richardson's law suggests instead that such conflicts are merely the very tail end of a highly regular power-law distribution – that is, the data show nearly exact linearity in the log-log scales; large wars are rare, but regular nonetheless. With correct parameters for the distribution, I can now quantify that regularity. For example, World War II created more than 10,000,000 fatalities – so by order of magnitude, 10^7. Putting this value into the CDF (eq. 5) along with the estimate for $\boldsymbol{\alpha}$

Miles Townes mdtownes@gwmail.gwu.edu

generated for Inter-state Wars gives

$$\Pr (X > 10^\wedge 7) = \left(\frac{10^7}{10^3}\right)^{-0.39} = .0275$$

The same procedure with ***α*** from Inter- and Intra-state Wars combined gives

$$\Pr (X > 10^\wedge 7) = \left(\frac{10^7}{10^3}\right)^{-0.45} = .0158$$

These values translate to roughly 2-3 wars per hundred and 1-2 wars per hundred, respectively for each dataset, with casualties in or above the tens of millions range. (The difference illustrates the importance of accuracy in estimating ***α*** – a ~5% difference in the original estimates translates to a ~75% difference between the two predictions.) The values are somewhat higher than the observed frequency of wars of that magnitude – namely 1 in 78 or 1 in 283. This shows that World War II was not a *sui generis* event that might never happen again – rather we should expect wars of this size and larger more frequently than we have seen.

It is also instructive to know what the distributions suggest for wars of heretofore unseen lethality – for example, a conflict a billion or more fatalities. By the same method – using the more accurate, more optimistic value of ***α*** (from the combined dataset) – for a magnitude of 10^9 deaths

$$\Pr (X > 10^9) = \left(\frac{10^9}{10^3}\right)^{-0.45} = .002$$

That is, 2 wars per thousand, with a billion or more casualties each; by our 95% confidence interval, the range is from roughly 1.5 wars per thousand to 4.3 wars per thousand. (Incidentally, the 95% c.i. for Interstate Wars gives a range of 1.2 to 15 wars per thousand.) In more concrete terms, I might ask whether I can expect to see a war of such lethality in my lifetime. Assuming I live to my 80s and given the rate of wars in the combined dataset – 283 wars over 181 years, so approximately 1.5 wars per year

– the next 50 years should yield roughly 75 wars. Multiplying this by the estimated frequencies of a magnitude 10^9 conflict, per the combined dataset confidence interval, gives me a range from .113 to .323 – that is, an 11% to 32% chance of witnessing a war with a billion or more casualities.[3]

Simulation offers another way to interpret these results. Assuming that the data offer only a partial – 181 year – segment of a larger distribution, simulation can provide a glimpse of the larger picture of that distribution. As mentioned, the simulation process is not entirely accurate, but produces a reasonable sketch of the distribution. Figure 3 shows the results of a simulated data run of 283 conflicts added to a plot of the real data for the combined Inter- and Intra-state Wars dataset. The simulated data are added after 1997, the year indicated by the dashed reference line. It should be noted that error intrinsic to the simulation process gives the simulated data an $\hat{\alpha} = 1.441$, which differs from but is nonetheless comparable to the estimated $\hat{\alpha} = 1.446$ for the real data.

Figure 3. Real and simulated data for $\boldsymbol{\alpha} = 1.44\sim$

3 The facts that such a war will likely be nuclear, and that the first millions of casualties will be nearly instantaneous, and that I might be among the first to die in such a war, are all irrelevant to the question at hand.

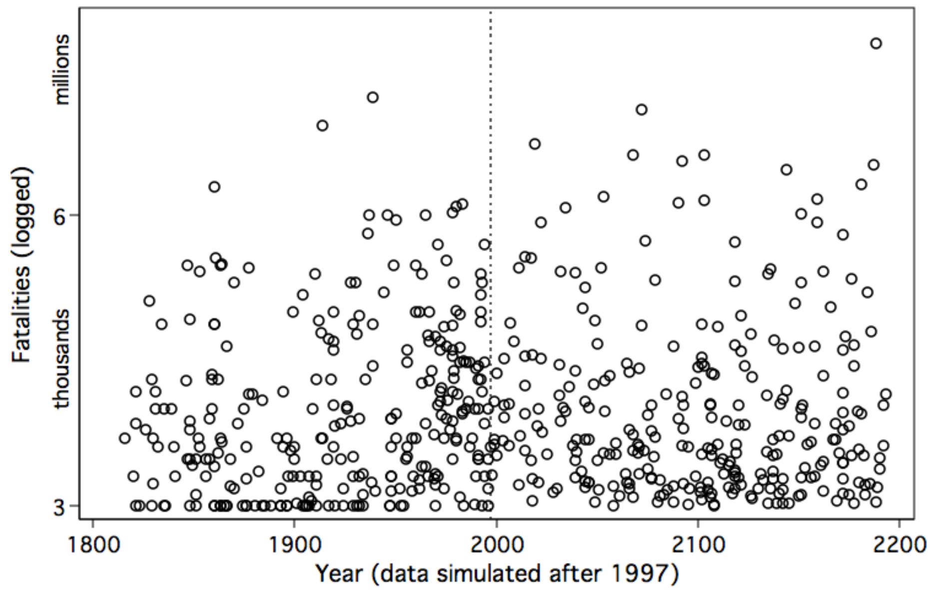

(The values for year were simulated by taking the mean of the poisson-distributed differences of start years in the real data, and generating a similar randomly-generated poisson distribution which gave the increment for year from one simulated observation to the next; thus the real data cover 181 years whereas the simulated data cover 196 simulated years, due to the random-generation process.)

The graph is illustrative but potentially misleading. First, it is only one possible simulation of innumerably many. I include this one because it is fairly conservative in scope, with a maximum conflict magnitude comparable to that of the real data; other simulations of the same parameters generated maximum magnitudes of 10^9, 10^{10} – even 10^{12}. So Figure 3 should not be taken as strictly indicative of the shape of conflicts to come, but merely suggestive of the implications of the power-law distribution. Second, the power-law distribution says nothing about the specific timing of events, only their relative frequency; so the year-to-year timing of events for the simulated portion of the plot is essentially meaningless. In fact, the simulation holds constant both α and the mean of the Poisson

Miles Townes mdtownes@gwmail.gwu.edu

distribution of wars in time – and assumption that may not be justified. Nonetheless, the graph at least helps to visualize the long-run implications of the power-law distribution.

Another more reliable way to visualize the implications of the power-law distribution is to generate a simulated dataset with a large number of observations. As the size of the dataset increases, the simulated frequencies of events of various magnitudes should approach the expected frequency given by calculation above. Figure 4 shows distribution – both as randomly generated and ordered from smallest to largest – of 10,000 simulated events from a power-law distribution with α = 1.44~. There are 22 events with magnitude greater than 10^9 and 179 with magnitude greater than 10^7; this comports with the expected frequencies of .002 and .0158 – although once again, there is some degree of error.

A final way to interpret the data also involves simulation: it is the answer to the question, what does it mean α is one value or another – what consequence does α imply? Figure 5 shows the results of 1000 simulated datasets for each value of α between 1 and 2.5 (in .05 increments), in which 250 wars were created and the total casualties calculated for that dataset. The black linear plot shows the median value for total casualties from the 1000 simulated datasets, and the gray range shows the area between the minimum and maximum values for total casualties for each value of α. The top values of the range are likely spurious, and should be taken as indication of how sensitive exponential distributions are to even slight error. Because of occasional extreme outliers, the mean of each simulation run was often highly skewed; thus the use of medians in this plot.

Miles Townes mdtownes@gwmail.gwu.edu

Figure 4. Scatterplots of 10,000 simulated events from power-law distribution of $\boldsymbol{\alpha} = 1.44$~

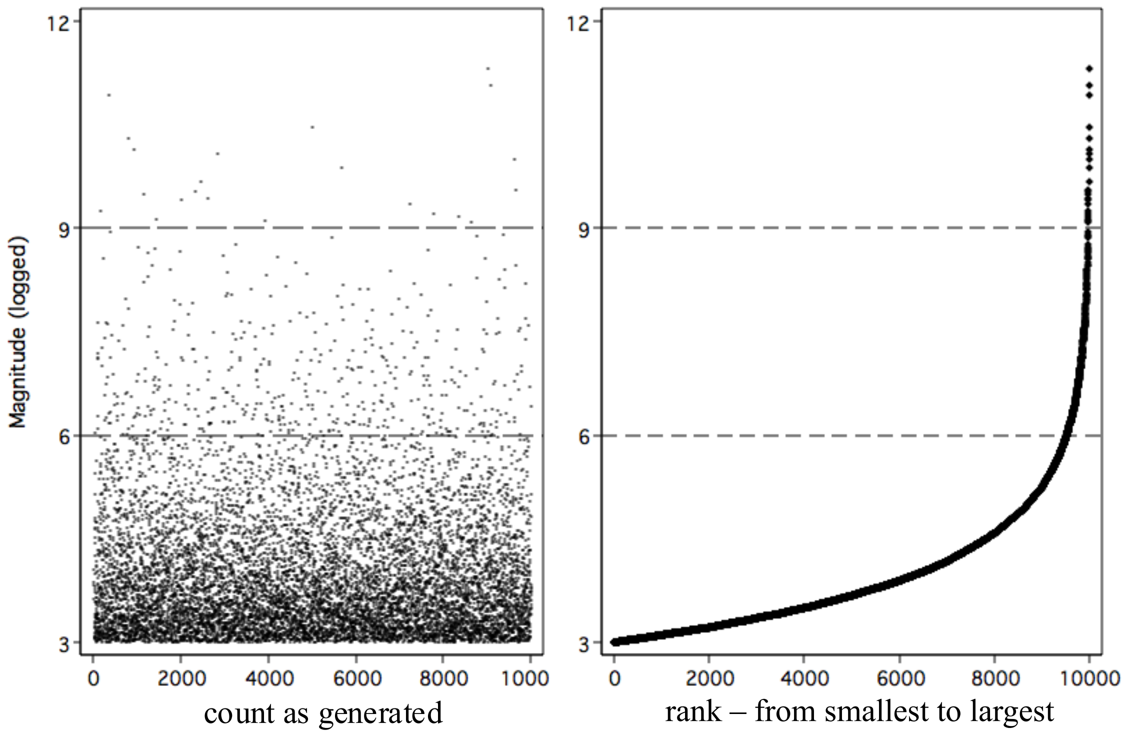

count

Figure 5. Median and Range for Total Casualties for Simulated Values of $\boldsymbol{\alpha}$

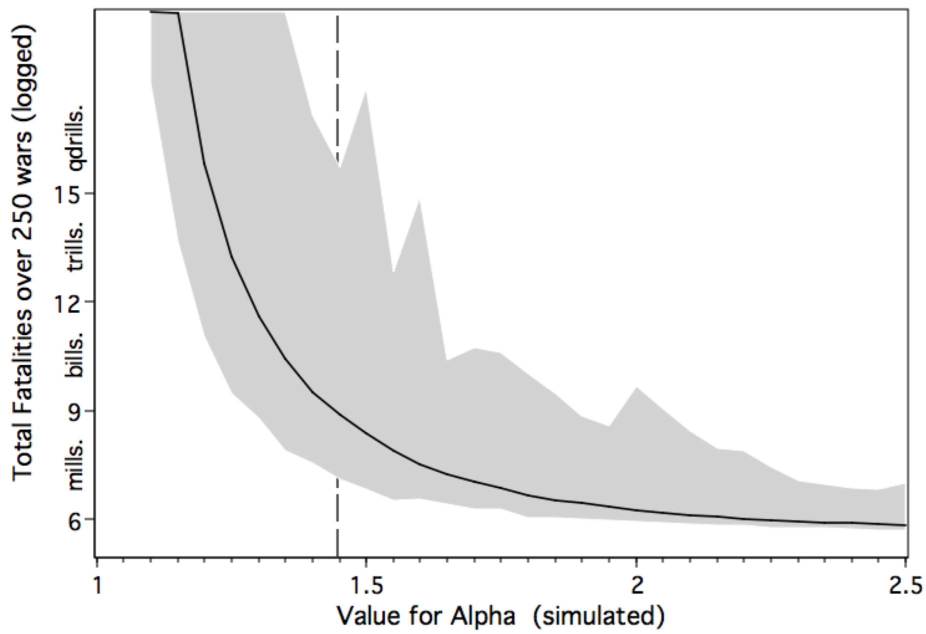

(In retrospect, it would have been more straightforward to simulate 283 wars, but I did not realize when I chose 250 how labor-intensive this graph would be. I intend to modify that process in the near future to make it more efficient.)

The reference line in Figure 5 marks 1.447 – the estimated value of α for the combined dataset. The most recent 250 real wars from that dataset generated a total of 44.5 million battle deaths; this gives $\log_{10} = 7.65$, a value between the median of 8.89 and minimum of 7.12 for simulated datasets at $\alpha = 1.45$. Notice that *decreasing* values of α are *increasingly* dangerous, in terms of potential total casualties, asymptotically approaching 1 (at which point the power-law distribution diverges; recall that in OLS, the quantity being estimated is not α, but $-(\alpha - 1)$ which can be less than 1). The increasing range not only suggests more casualties on average, but greater volatility in the system; the casualties for $\alpha = 1.45$ range from tens of millions at least to *quadrillions* at most .

So it is that Richardson's Law has profound implications for the future of deadly conflict in the

Miles Townes mdtownes@gwmail.gwu.edu

world. The result may not allow us to predict the next war with any accuracy, but we can at least say something useful about the next hundred or so wars: specifically, we should expect that the wars of the future should be no less deadly – and perhaps substantially more so – than those history has already recorded. This assumes, of course, that the shape of the distribution is invariate – that α = 1.44~ is indeed a "law". I now turn my attention to that assumption.

V. Richardson's Law: Statistic, or State of Nature?

If the KS statistic is a meaningful test of two different distributions (the literature is sparse and technical) then it offers a tool by which to discover any changes or differences in the power-law distribution – namely, whether the α parameter varies across categories of wars. Given that, as shown above, different values of α suggest different levels of danger (i.e. expected total casualties from a given series of wars), an increase in α for a given category of wars suggests that these wars are less dangerous, for whatever reason. As above, Extra-state wars have an α higher than that for other wars, suggesting they are somewhat less dangerous than those other wars. But even within the combined dataset, there are different possible categories by which wars might be analyzed. These are explicitly post-theoretical categories, and so here I relax the assumption that these wars are all of the same type. Instead, I use theoretical claims about the nature of war to divide conflicts into subpopulations. In particular, the categories I want to test are

- Century: are wars in the 20th century more dangerous than wars in the 19th?
- Nuclear weapons: are wars in the nuclear age more or less dangerous than wars prior?
- Geography: are wars in Africa more or less dangerous than wars in Europe?

This is a two-step process. First, I separate out the categories from the combined dataset of Inter- and

Miles Townes mdtownes@gwmail.gwu.edu

Intra-state wars, and then perform the KS test. If that test shows a significant difference, I then estimate the α parameter for each category and compare. Table 4 shows the results of the KS tests; of the three categorical divides, Europe/Africa is entirely insignificant. This is further demonstrated by the power-law estimation: α is 1.34 European wars and 1.38 for African wars, with large standard errors due to small sample sizes (N = 65 total). That the two should be so close is surprising, but this is partly a consequence of insufficient data. Furthermore, a number of wars that occurred in Africa are excluded from consideration because they are part of the Extra-state wars dataset, and not the Inter- and Intra-state wars under consideration here.

There is, however, a significant difference between pre- and post-nuclear wars, and a less significant difference between wars in the 19th and 20th centuries. For these categories, the process of estimating α, standard error, and 95% confidence intervals provides even more evidence of a substantive difference between the subcategories. The results of those procedures are shown in Table 5.

Table 4. Two-sample KS tests for difference of distribution across categories of wars

	Pre/Post-nuclear	Europe vs. Africa (20th Century)	19th/20th Centuries
KS statistic	.2051	.0826	.1964
p-value	.006	1.000	.015

Table 5. Estimates for α, interval, and significance by theoretic category of war

	Century		Nuclear	
	19th	20th	Pre-	Post-
α	1.574	1.399	1.511	1.383
std. error	.0059	.0022	.0032	.0083
95% CI	(1.558, 1.586)	(1.388, 1.415)	(1.499, 1.523)	(1.379, 1.387)

Miles Townes mdtownes@gwmail.gwu.edu

KS statistic	.9873	.9794	.9930	.9357
p-value	.7748	.9908	.7012	1.00
N	98	185	161	122

There are clear and substantive differences between the α estimates for 19th Century and 20th Century Wars. There are also clear and substantive differences between the α estimates for pre- and post-nuclear wars. In both cases, the later value is lower than the earlier value, suggesting that conflicts as a system are becoming *more deadly*. In particular, the distribution of wars in the post-nuclear era (i.e. since 1945) show the lowest value for α – 1.38 – from any dataset or subset tested here. This distribution excludes World War II, the deadliest war on record, which is counted as pre-nuclear. As evidence of how much more deadly an $\alpha = 1.38$ distribution is, consider that by the CDF the expected frequency $(\Pr X > 10^9) = .005$ – that is, a billion casualty conflict once every two hundred wars. Moreover, by the simulation process used to generate Figure 5 above, the minimum of simulated total casualties for α = 1.40 was 34.7 million; the median was 3.3 billion. Not only are wars as a systemic phenomenon becoming more dangerous, but there is no evidence of a declining frequency of conflict: 65% of wars occurred in the 20th century (53% of the combined dataset's timeframe) and 43% since just 1945 (28% of the timeframe).

Apart from their appalling implications, these results also demonstrate the potential for meaningful tests of systems of conflict. Because (or at least, when) wars are power-law distributed, standard statistical tests of relative fatalities tend to be less than helpful. For example, the mean of fatalities in African wars in the 20th century was 65,844; for European wars (not counting WWII) it was 923,676 – an order of magnitude higher. This might lead the researcher to conclude or assume some fundamental difference between the two continents as conflict systems, when in fact the present analysis does not support such a claim. However, this section – the sub-category tests – are intended to be suggestive only, and not a definitive study. It is entirely possible that the Africa-Europe non-difference might be challenged, perhaps refuted, by more systematic application of the present methods.

VI. Richardson's α

The dominant approach to the analysis of conflict is to focus on special cases – great-power wars, major civil wars, World Wars – and from there to generalize outwards toward an understanding of war as a systematic feature of human society. Richardson's approach was the opposite: to study violence as a systemic phenomenon, and from there work his way inwards to special circumstances.

In replicating the power-law with current data, I have shown that Richardson's power-law result still stands as one of few strong empirical regularities in international relations, that OLS is inappropriate to its estimation, and that maximum likelihood methods offer the most precise, most reliable estimates for the parameters of power-law distributions. I have used these methods to provide substantive interpretation of the power-law distribution of fatal conflict, and suggested methods by which various systems or eras of conflict might be compared and analyzed via those methods. There is, it should be clear, a great deal of work still remaining towards our understanding of violent conflict as a systemic phenomenon; it is also clear that we possess many of the tools necessary to that work.

Prior studies of Richardson's work have assumed that his power-law result is fixed over time and space, and so they seek only to replicate and refine that result (or worse, merely take it as given). They have, perhaps, taken the "law" in power-law too literally. This has been true even as advances in computing power provided researchers with tools beyond Richardson's imagining – especially the ability to fit, test, and simulate power-law distributions quickly and repeatedly. Given the power of these tools, and the fact that the study of war never lacks for scientific advance, that study would benefit from the adoption of Richardson's α – estimated by MLE – as a primary descriptive statistic of systems of conflict. That is, the α parameter of the power-law distribution offers the conflict researcher a more fundamental and concise measure of systems of conflict than do conventional statistical tools. The next step – beyond the scope of this paper – is to discover what factors account for variance (or invariance) in Richardson's α, specifically to determine whether human or institutional activity can affect α for the better. The present research suggests not, but that implication begs to be disproved.

Resource note: a draft of this paper, links to relevant web pages, and copies of the STATA 10 .do files used in this research are available at the following URL:

http://home.gwu.edu/~mdtownes/MLE

References

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

Miles Townes mdtownes@gwmail.gwu.edu

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/

Gould, W., J. Pitblado, W. Sribney (2006). *Maximum Likelihood Estimation with STATA*(3rd ed.). College Station, TX; STATA Press.

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.

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