War size distribution: Empirical regularities behind conflicts

Rafael, González-Val

Universidad de Zaragoza Institut d'Economia de Barcelona (IEB)

12 January 2015

War size distribution: Empirical regularities behind conflicts

Rafael González-Val

Universidad de Zaragoza & Instituto de Economía de Barcelona

***Abstract*:** This paper analyses the statistical distribution of war sizes. Using the method recently proposed by Clauset, Shalizi, and Newman (2009), we find moderate support for a Pareto-type distribution (power law), using data from different sources (COW and UCDP) and periods. A power law is a plausible model for the size distribution of a pool of all wars and a sample of wars in many years, although the log-normal distribution is a plausible alternative model that we cannot reject. The random growth of conflicts could generate both types of distribution. We study the growth rates of battle deaths and random growth cannot be rejected for most of the distribution, although the results also reveal a clear decreasing pattern; the growth of deaths declines faster if the number of initial deaths is greater.

***Keywords*:** war size distribution, battle deaths, power law, Pareto distribution.

***JEL*:** D74, F51, N40.

1. Introduction

In one of the first analyses of the statistics of war, Richardson (1948) studied the variation of the frequency of fatal quarrels with magnitude. He collected a data set for violent incidents (wars and homicides), measured by the number of victims, from 1820 to 1945, and his calculations revealed that the relationship between magnitude (size) and frequency (number) of both wars and small crime incidents could be satisfactorily fitted by a straight decreasing line with a negative slope, suggesting a power law function. This striking empirical regularity could have important implications, but it has remained almost unexplored from either a theoretical or an empirical point of view for many years.

Only a few papers follow Richardson's approach (Roberts and Turcotte 1998; Cederman 2003; Clauset, Young, and Gleditsch 2007), and they also find evidence of power law behaviour. Roberts and Turcotte (1998) find a power law dependence of number on intensity, taking into consideration several alternative measures of the intensity of a war in terms of battle deaths, using Levy's (1983) data set of 119 wars from 1500 to 1974 and Small and Singer's (1982) data set of 118 wars from 1816 to 1980. Cederman (2003) finds strong support for a power law distribution, using interstate war data from 1820 to 1997 from the Correlates of War Project. Based on this empirical evidence, he also proposes an agent-based model of war and state formation that exhibits the same kind of power law regularities. Clauset, Young, and Gleditsch (2007) extend Richardson's analysis to study the frequency and severity of terrorist attacks worldwide since 1968, also finding a linear relationship between the frequency and the severity of these deadly incidents.

The results of these studies are similar to Richardson's original result. However, as Levy and Morgan (1984) point out, all these studies focus on the distribution of all wars rather than on the wars occurring in a given period, although the frequency of wars in a given period is also assumed to be inversely related to their seriousness. Levy and Morgan (1984) try to address this latter point by calculating Pearson correlation indexes between frequency and intensity, finding a negative correlation. They use Levy's (1983) data set for wars between 1500 and 1974, aggregating wars in 25-year periods.

Finally, there is another strand of related literature. All the studies previously mentioned use between-conflict data, but other papers (Bohorquez et al. 2009; Johnson et al. 2011) focus on within-conflict incidents (attacks). Surprisingly, these studies conclude that the size distribution or timing of within-conflict events is also power law distributed. Bohorquez et al. (2009) show that the sizes and timing of 54,679 violent events reported as part of nine diverse insurgent conflicts exhibit remarkable similarities. In all cases, the authors cannot reject the hypothesis that the size distribution of the events follows a power law, but they can reject log-normality. They build on this empirical evidence to propose a unified theoretical model of human insurgency that reproduces these features, explaining conflict-specific variations quantitatively in terms of the underlying rules of engagement. Johnson et al. (2011) uncover a similar dynamic pattern using data about fatal attacks by insurgent groups in both Afghanistan and Iraq and by terrorist groups operating worldwide. They estimate the escalation rate and the timing of fatal attacks, finding that the average number of fatalities per fatal attack is fairly constant in a conflict. Furthermore, when they calculate the progress curve, they obtain a straight line, which is best fitted by a power law.

This paper contributes to the literature in several ways. First, in the spirit of Richardson (1948), we estimate the distribution of a pool of all wars. Second, using yearly data, we estimate the war size distribution by year from 1989 to 2010 to determine whether differences exist between the overall distribution of all wars and the year-by-year distribution (Clauset, Young, and Gleditsch (2007) carry out a similar analysis for terrorist attacks by year). Using Clauset, Shalizi, and Newman's (2009) methodology, we find that the power law is a plausible fit in most of the cases, but there is a plausible alternative model, the log-normal distribution. Finally, we study the behaviour of the growth rates for those conflicts that last longer than one period. We empirically test random growth, which could generate both types of distribution, Pareto and log-normal.

The paper is organised as follows. Section 2 introduces the databases that we use. Section 3 contains the statistical analysis of war size distribution and its evolution over time, and Section 4 concludes.

2. Data

We measure war size using the number of recorded battle deaths, that is, the battle-related combatant fatalities. The data come from two international data sets: the

Correlates of War (COW) (Version 4.0) (2010) Project and the Uppsala Conflict Data Program (UCDP/PRIO) Armed Conflict Dataset (Version 5) (2011).

We consider wars in which the government of a state was involved in one form or another. The COW Project distinguishes three kinds of state wars: interstate (between/among states), intra-state (within states), and extra-state (between/among a state(s) and a non-state entity). According to the COW war typology, a war must have sustained combat, involve organised armed forces, and result in a minimum of 1,000 battle-related combatant fatalities within a 12-month period; for a state to be considered a war participant, the minimum requirement is that it has either to commit 1,000 troops to the war or to suffer 100 battle-related deaths. This requisite condition was established by Small and Singer (1982). Interstate wars are those in which a territorial state is engaged in a war with another state. Intra-state wars are wars that predominantly take place within the recognized territory of a state; they include civil, regional, and intercommunal wars. Finally, extra-state wars are those in which a state is engaged in a war with a political entity that is not a state, outside the borders of the state. Extra-state wars are of two general types: colonial and imperial. The COW data cover 95 different interstate wars from 1823 to 2003, 190 intra-state wars from 1818 to 2007, and 162 extra-state wars from 1816 to 2004.[1] Thus, the COW data set covers all the conflicts over a long period and enables us to estimate the size distribution of a large pool of modern wars.

The UCDP/PRIO Armed Conflict Dataset is a joint project between the Uppsala Conflict Data Program at the Department of Peace and Conflict Research, Uppsala University, and the Centre for the Study of Civil War at the International Peace Research Institute in Oslo (PRIO). The UCDP defines conflict as 'a contested incompatibility that concerns government and/or territory where the use of armed force between two parties, of which at least one is the government of a state, results in at least 25 battle-related deaths'.[2] There are two important differences between the UCDP and the COW data. First, the UCDP data set includes four different types of conflict: extrasystemic, interstate, internal, and internationalised internal. Second, the UCDP data

[1] More information about war classifications and the lists of interstate, intra-state, and extra-state wars included in the database can be found in Sarkees and Wayman (2010).
[2] More information about the UCDP/PRIO Armed Conflict Dataset can be found in Gleditsch et al. (2002). The data set is available to download from http://www.pcr.uu.se/research/ucdp/datasets/.

set contains information about conflicts by year from 1989 to 2010. Thus, we can estimate the year-by-year size distribution.

The data presented by the UCDP are based on information taken from a selection of publicly available sources, printed as well as electronic. The sources include news agencies, journals, research reports, and documents from international and multinational organisations and NGOs. Global, regional, and country-specific sources are used for all countries. The basic source for the collection of general news reports is the Factiva news database (previously known as Reuters Business Briefing), which contains over 8,000 sources. There is not usually much information available on the exact number of deaths in a conflict, and the media coverage varies considerably from country to country. However, the fatality estimates given by the UCDP are based on publicly accessible sources.

The project uses automated events data search software that makes it possible to retrieve all the reports containing information about individuals who have been killed or injured. Each news report is then read by UCDP staff, and every event that contains information about individuals who have been killed is coded manually into an events data set. Ideally, these individual figures are corroborated by two or more independent sources. These fatalities are later aggregated into low, high, and best estimates for every calendar year. The lack of available information means that it is possible that there are more fatalities than the UCDP high estimate, but it is very unlikely that there are fewer than the UCDP best estimate. Here we use the best estimate figure in all cases.

Table 1 shows the sample sizes for each year and the descriptive statistics. There is a decrease in the number of ongoing armed conflicts over time, and this decrease is especially marked in the last few years (the average number of wars by year from 1989 to 2000 is 43.8, while in the 2001–2010 period it is 33.3). Moreover, the conflicts in the last few years have been less intense: the average number of battle deaths per war also decreases over time.

Roberts and Turcotte (1998) suggest that a pool of wars from different periods (like the COW data set) can be criticised because the global population changes substantially over a long time period. The same number of battle deaths would not represent the same war intensity if there had been a huge change in the world population. Some authors try to correct for this by using relative measures of size: Levy

(1983) defines the intensity of a war as the number of battle deaths divided by the population of Europe in millions at the time of the war, because estimates of the total world population may not be reliable for early periods. In this paper, we also define a relative measure of size as the ratio of battle deaths to the sum of the populations (in thousands) of the combatant countries of the conflict in the year of the start of the conflict.[3] Population data are also taken from the COW Project.[4] This ratio represents the number of deaths per thousand inhabitants in the countries involved in the war. However, note that this normalisation is not necessary when all the conflicts take place in the same period.

3. Results

3.1 War size distribution

Let S denote the war size (measured by the recorded battle deaths); if this is distributed according to a power law, also known as a Pareto distribution, the density function is $p(S) = \frac{a-1}{\underline{S}}\left(\frac{S}{\underline{S}}\right)^{-a}$ $\forall S \geq \underline{S}$ and the complementary cumulative density function $P(S)$ is $P(S) = \left(\frac{S}{\underline{S}}\right)^{-a+1}$ $\forall S \geq \underline{S}$, where $a > 0$ is the Pareto exponent (or the scaling parameter) and $\underline{S}$ is the number of battle deaths in the war at the truncation point, which is the lower bound to the power law behaviour. It is easy to obtain the expression $R = A \cdot S^{-a}$, which relates the empirically observed rank R (1 for the largest conflict, 2 for the second largest, and so on) to the war size. As Clauset and Wiegel (2010) point out, one of the properties of the power law is that there is no qualitative difference between large and small events; multiplying the argument (S) by some factor λ results in a change in the corresponding frequency that is independent of the argument.

This expression is applied to the study of very varied phenomena, such as the distribution of the number of times that different words appear in a book (Zipf 1949), the intensity of earthquakes (Kagan 1997), the losses caused by floods (Pisarenko

[3] The author thanks one anonymous referee for this suggestion.

[4] The COW Project includes a fourth category of war, wars between or among non-state entities. We exclude these wars (62 observations) from our analysis because in these cases it is not possible to quantify the populations involved on either side of the conflict (or even the population of the region in which the combat occurred, since COW only distinguishes six major areas), and thus no relative measure of size can be calculated.

1998), forest fires (Roberts and Turcotte 1998), city size distribution (Soo 2005), and country size distribution (Rose 2006).

Taking natural logarithms, we obtain the linear specification that is usually estimated:

$$\ln R = \ln A - a \ln S + u\,, \tag{1}$$

where u represents a standard random error ($E(u) = 0$ and $Var(u) = \sigma^2$) and $\ln A$ is a constant. The greater the coefficient $\hat{a}$, the more homogeneous are the war sizes. Similarly, a small coefficient (less than 1) indicates a heavy-tailed distribution. However, this regression analysis, which is commonly used in the literature, presents some drawbacks that have been recently highlighted by Clauset, Shalizi, and Newman (2009); of these, the main one is that the estimates of the Pareto exponent are subject to systematic and potentially large errors.[5]

Therefore, to estimate the power laws, we use the innovative method proposed by Clauset, Shalizi, and Newman (2009). This has been used to fit power laws to different data sets; Clauset, Shalizi, and Newman (2009) apply it to find moderate support for the power tail behaviour of the intensity of wars from 1816 to 1980, measured as the number of battle deaths per 10,000 of the combined populations of the warring nations (data sets from Roberts and Turcotte 1998 and Small and Singer 1982), and the behaviour of the severity of terrorist attacks worldwide from February 1968 to June 2006, measured as the number of deaths directly resulting from the attacks (data from Clauset, Young, and Gleditsch 2007). They also use this method with other data sets from many very different fields (e.g., the human populations of US cities in the 2000 US Census, the intensity of earthquakes occurring in California between 1910 and 1992, or the number of 'hits' received by websites from America Online Internet service customers in a single day). In a recent work, Brzezinski (2014) uses this methodology to study the power law behaviour of the upper tails of wealth distributions, using data on the wealth of the richest persons taken from the 'rich lists' produced by business magazines.

[5] Preliminary results obtained from the OLS estimation of Eq. (1) indicate that the power law provides a very good fit to the real behaviour of the whole distribution (all the observations) for our pool of COW wars (using deaths and relative deaths) and the yearly UCDP data set. The estimated R^2 is greater than 0.9 in all cases, and the estimated Pareto exponent is always less than 1, indicating that the distribution is heavy-tailed; this means that the average war loss is controlled by the largest conflicts. However, as indicated in the main text, these OLS results are not robust (Clauset, Shalizi, and Newman 2009).

The maximum likelihood (ML) estimator of the Pareto exponent is:

$$\hat{a} = 1 + n\left(\sum_{i=1}^{n} \ln \frac{S_i}{\underline{S}}\right), \ \forall S_i \geq \underline{S}.$$

The ML estimator is more efficient than the usual OLS line regression if the underlying stochastic process is really a Pareto distribution (Gabaix and Ioannides 2004; Goldstein, Morris, and Yen 2004). Clauset, Shalizi, and Newman (2009) propose an iterative method to estimate the adequate truncation point ($\underline{S}$). The exponent a is estimated for each $S_i \geq \underline{S}$ using the ML estimator (bootstrapped standard errors are calculated with 1,000 replications), and then the Kolmogorov–Smirnov (KS) statistic is computed for the data and the fitted model. The $\underline{S}$ lower bound that is finally chosen corresponds to the value of S_i for which the KS statistic is the smallest.[6]

Figure 1 shows the results for the COW data, covering all state (inter-, intra-, and extra-state) wars from 1816 to 2007. The data, plotted as a complementary cumulative distribution function (CCDF), are fitted by a power law, and its exponent is estimated using the ML estimator. For illustrative purposes, a log-normal distribution is also fitted to the data by maximum likelihood (the blue dotted line). The optimal lower bound for both distributions is estimated using Clauset, Shalizi, and Newman's (2009) method. The black line shows the power law behaviour of the upper tail distribution. The first graph shows the battle deaths' distribution, with an estimated Pareto exponent of 1.74 for $deaths \geq 9{,}540$, and the second displays the relative deaths, with a scaling parameter of 1.90 for $relative\ deaths \geq 0.60$. The power law appears to provide a good description of the behaviour of the distribution. In contrast, the fit of the log-normal distribution is poor, especially for the highest observations. Nevertheless, visual methods can lead to inaccurate conclusions (González-Val, Ramos, and Sanz-Gracia 2013), especially at the upper tail, because of large fluctuations in the empirical distribution (Clauset and Woodard 2013), so next we conduct statistical tests on the goodness of fit.

Clauset, Shalizi, and Newman (2009) propose several goodness-of-fit tests. In the same way as Brzezinski (2014), we use a semi-parametric bootstrap approach. The

[6] The power laws and the statistical tests are estimated using the poweRlaw R package developed by Colin S. Gillespie (based on the R code of Laurent Dubroca and Cosma Shalizi and the Matlab code of Aaron Clauset) and the Stata codes developed by Michal Brzezinski, which are all freely available on their webpages.

procedure is based on the iterative calculation of the KS statistic for 1,000 bootstrap data set replications. This method samples from observed data and checks how often the resulting synthetic distributions fit the actual data as poorly as the ML-estimated power law. Thus, the null hypothesis is the power law behaviour of the original sample for $S_i \geq \underline{S}$. Nevertheless, this test has an unusual interpretation because, regardless of the true distribution from which our data were drawn, we can always fit a power law. Clauset, Shalizi, and Newman (2009) recommend the conservative choice that the power law is ruled out if the p-value is below 0.1: 'that is, it is ruled out if there is a probability of 1 in 10 or less that we would merely by chance get data that agree as poorly with the model as the data we have.' Therefore, this procedure only allows us to conclude whether the power law is a plausible fit to the data. Table 2 shows the results of the tests; the p-values of the test for both COW samples, deaths and relative deaths, are higher than 0.1, confirming that the power law is a plausible approximation to the real behaviour of the data. This evidence confirms Cederman's (2003) results and the original result of Richardson (1948).

Finally, we also compare the linear power law fit with the fit provided by another non-linear distribution, the log-normal distribution, using Vuong's model selection test to compare the power law with the log-normal.[7] The test is based on the normalised log-likelihood ratio; the null hypothesis is that the two distributions are equally far from the true distribution, while the alternative is that one of the test distributions is closer to the true distribution. High p-values indicate that one model cannot be favoured over the other, and this is the conclusion reached with the COW data – see Table 2. Overall, using Clauset, Shalizi, and Newman's (2009) terminology, we obtain moderate support for the power law behaviour of our pool of wars: the power law is a plausible fit but there is a plausible alternative as well.

Remember that this is the distribution of a pool of all wars over a long period. Next, we use the yearly UCDP data set to estimate the war size distribution by year from 1989 to 2010. We fit a power law for each period of our yearly sample of wars; Figure 2 displays the results for two representative years (1998 and 2007) of the two

[7] In Figures 1 and 2, the lower bound for both distributions (log-normal and power law) is calculated using Clauset, Shalizi, and Newman's (2009) method. The lower bounds can be different, but to compare the distributions the threshold must be the same for both distributions, so to run the test we use the same lower bound, the estimated value corresponding to the power law.

possible cases.[8] In 1998, the distribution seems clearly non-linear and the power law fit is poor, while in 2007, the power law provides a good fit to the real behaviour of the distribution. The latter one is the predominant case, because the power law is rejected in only 8 of the 22 years considered. Figure 3 summarizes the results of the estimates by year, showing the estimated Pareto exponent and the results of the goodness-of-fit test at the 10% significance level (the p-values are reported in Table 2); most of the rejections are located in the first periods of our sample. Moreover, the results of Vuong's model selection test (Table 2) indicate that the fit provided by the power law is not significantly better than the log-normal fit in any year. However, these year-by-year results must be interpreted with caution; these sample sizes by year are small and Clauset, Shalizi, and Newman (2009) warn that their method can break down in small samples, for which a lack of statistical power can make low p-values inherently unlikely.

Although in some years the standard error of the scaling parameter is high because the number of observations above the estimated truncation point is low, the estimated values fluctuate between 2 and 2.5. These values are similar to those obtained by Clauset, Young, and Gleditsch (2007) and Clauset, Shalizi, and Newman (2009) in their analysis of terrorist attacks. Clauset, Young, and Gleditsch (2007) develop a theoretical model to explain this power law pattern.[9] Their model is a variation of the Reed and Hughes (2002) mechanism of competing exponentials, which yields to a power law distribution for the observed severities. The scaling parameter depends on the growth rate for attacks and the hazard rate imposed on events by states, and, making some assumptions (equal rates with a slight advantage to states due to their longevity and large resource base), the model generates $a \approx 2.5$. Clauset and Wiegel (2010) provide an alternative theoretical explanation, generalising the model of Johnson et al. (2005). This model, which is based on the notion of self-organised criticality and which describes how terrorist cells might aggregate and disintegrate over time, also predicts that the distribution of attack severities should follow a power law form with an exponent of 2.5.

3.2 Growth analysis

[8] The results for all the years are available from the author upon request.

[9] Saperstein (2010) and Clauset, Young, and Gleditsch (2010) discuss the implications of Clauset, Young, and Gleditsch's (2007) model.

The above results show what we consider to be a snapshot of the size distribution of wars from 1989 to 2010. For each year, we obtain the estimated coefficients of the Pareto exponent and conduct a goodness-of-fit test that indicates the plausibility of the power law model in many cases. The literature that studies the distribution of financial assets (Gabaix et al. 2006) and of firm (Sutton 1997) and city (Gabaix 1999) sizes usually concludes that this kind of Pareto-type distribution is generated by a random growth process. Moreover, there is another plausible alternative model that we could not reject in the previous empirical analysis, the log-normal distribution. Random growth rates of conflicts could generate both types of distribution – log-normal and Pareto – if there is a lower bound to the distribution (which can be very low) (see Gabaix 1999).[10] The hypothesis usually tested is that the growth of the variable is independent of its initial size (the underlying growth model is a multiplicative process).[11] To check whether this is true for war sizes, we carry out a dynamic analysis of growth rates using two different non-parametric tools. The UCDP data set enables us to calculate the yearly growth rates of battle deaths for conflicts that lasted for more than one year. We define g_i as the growth rate $\left(\ln S_{it} - \ln S_{it-1}\right)$ and normalise it (by subtracting the contemporary mean and dividing by the standard deviation in the relevant year), where S_{it} is the *i*th war's size (battle deaths).[12] We build a pool with all the growth rates between two consecutive years; there are 639 battle deaths–growth rate pairs in the period 1989–2010.

First, we study how the distribution of growth rates is related to the distribution of initial battle deaths (Ioannides and Overman 2004). Figure 4 shows the stochastic kernel estimation of the distribution of normalised growth rates, conditional on the distribution of initial battle deaths at the same date. To make the interpretation easier, the contour plot is also shown. The plot reveals a slight negative relationship between the two distributions, although there is a great deal of variance. However, most of the observations are concentrated in two peaks of density; the higher one corresponds to conflicts with a small number of deaths (below 5 on the logarithmic scale, i.e. fewer than 150 casualties) and the lower one to the less numerous group of conflicts with a

[10] Only a small change from the log-normal generative process yields a generative process with a power law distribution, that is, a bounded minimum that acts as a lower reflective barrier to the multiplicative model (Gabaix 1999, Mitzenmacher 2004).
[11] In the firm and city size literature this hypothesis is called 'Gibrat's law'.
[12] The growth rates need to be normalised because we are considering growth rates from different periods jointly in a pool.

high number of battle deaths (7 on the logarithmic scale, which means around 1,100 casualties). Note that the conditional distribution of growth rates is equal to zero for both types of war, indicating that both distributions are independent for most of the observations.

To gain a clearer view of the relationship between growth and initial battle deaths, we also perform a non-parametric analysis using kernel regressions (Ioannides and Overman 2003). This consists of taking the following specification:

$$g_i = m(s_i) + \varepsilon_i,$$

where g_i is the normalised growth rate and s_i the logarithm of the *i*th war's number of initial battle deaths. Instead of making assumptions about the functional relationship m, $\hat{m}(s)$ is estimated as a local mean around point s and is smoothed using a kernel, which is a symmetrical, weighted, and continuous function in s.

To estimate $\hat{m}(s)$, the Nadaraya–Watson method is used, as it appears in Härdle (1990, Chapter 3), based on the following expression:

$$\hat{m}(s) = \frac{n^{-1}\sum_{i=1}^{n} K_h(s - s_i) g_i}{n^{-1}\sum_{i=1}^{n} K_h(s - s_i)},$$

where K_h denotes the dependence of the kernel K (in this case an Epanechnikov) on the bandwidth h. We use the bandwidth $h = 0.5$.[13] As the growth rates are normalised, if the growth was independent of the initial number of deaths, the non-parametric estimate would be a straight line on the zero value and values different from zero would involve deviations from the mean.

The results are shown in Figure 5. The graph also includes the bootstrapped 95% confidence bands (calculated from 500 random samples with replacement). The estimates confirm the negative relationship between size and growth observed in Figure 4, although we cannot reject the premise that the growth is different from zero (random growth) for most of the distribution. Nevertheless, the decreasing pattern is clear: the greater the number of initial deaths, the lower the growth rate. This points to a certain degree of convergence (mean reversion) across wars, which we can interpret as

[13] The results using Silverman's optimal kernel bandwidth were similar.

evidence of the 'explosive' behaviour of conflicts, because the greater the number of initial deaths, the faster the decline in the growth of deaths.

Gabaix and Ioannides (2004) explain how random growth can be compatible with a degree of convergence in the evolution of growth rates, by putting forward what they call deviations from random growth that do not affect the distribution. We can adapt their theoretical framework to war growth. We start from:

$$\ln S_{it} - \ln S_{it-1} = \mu(X_{it}, t) + \varepsilon_{it}, \tag{2}$$

where X_{it} is a possibly time-varying vector of the characteristics of war i; $\mu(X_{it}, t)$ is the expectation of war i's growth rate as a function of the specific conflict characteristics at time t; and ε_{it} is white noise. In the simplest specification, ε_{it} is independently and identically distributed over time (this means that ε_{it} has a zero mean and a constant variance that is uncorrelated with ε_{is} for $t \neq s$) and $\mu(X_{it}, t)$ is constant.

Gabaix and Ioannides (2004) consider two types of deviation, relaxing both assumptions. We are interested in the consequences of relaxing the assumption of an i.i.d. ε_{it}, assuming constant $\mu(X_{it}, t) = \mu$. The following stochastic structure for ε_{it} is assumed: $\varepsilon_{it} = b_{it} + \eta_{it} - \eta_{it-1}$, where b_{it} is i.i.d. and η_{it} follows a stationary process. Replacing in (2), we obtain:

$$\ln S_{it} - \ln S_{i0} = \mu t + \sum_{s=1}^{t} b_{is} + \eta_{it} - \eta_{i0}.$$

The term $\sum_{s=1}^{t} b_{is}$ gives a unit root in the growth process (hence random growth), while the term η_{it} can have any stationarity. According to Gabaix and Ioannides (2004), this means that we can obtain a Pareto-type distribution even if the war growth process contains a mean reversion component, as long as it contains a non-zero unit root component.

4. Conclusions

Richardson's (1948) seminal study established a negative relationship between the frequency and the severity of wars, introducing a new empirical regularity. The aim of this paper is to provide robust evidence for or against Richardson's claim.

First, we estimate the distribution of a pool of all wars using COW state (inter-, intra-, and extra-state) war data from 1816 to 2007. Our estimates confirm Cederman's (2003) results and the original result of Richardson (1948); the power law provides a plausible fit to the real behaviour of the distribution. Second, using UCDP yearly data, we estimate the war size distribution by year from 1989 to 2010, finding that a power law cannot be rejected in most of the periods. Furthermore, the estimated values fluctuate around 2.5, a value similar to that of other studies that have analysed terrorist attacks. If we add that some studies conclude that the size distribution and timing of within-conflict events is also power law distributed (Bohorquez et al. 2009; Johnson et al. 2011), all this evidence points to a universal pattern across and within war sizes. Nevertheless, the log-normal distribution is a plausible alternative model that we cannot reject in any case.

Finally, a study of the growth rates of battle deaths reveals that random growth cannot be rejected for most of the distribution, which could explain the resulting Pareto (power law) or log-normal size distribution. Nevertheless, a clear decreasing pattern is also observed: the greater the number of initial deaths, the faster the decline in the growth of deaths, although this mean reversion behaviour can be compatible with random growth.

References

Bohorquez, J. C., S. Gourley, A. R. Dixon, M. Spagat, and N. F. Johnson. 2009. "Common Ecology Quantifies Human Insurgency." *Nature* 462: 911–914.

Brzezinski, M. 2014. "Do wealth distributions follow power laws? Evidence from 'rich lists'. " *Physica A* 406: 155–162.

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

Clauset, A., C. R. Shalizi, and M. E. J. Newman. 2009. "Power-law Distributions in Empirical Data." *SIAM Review* 51(4): 661–703.

Clauset, A., and F. W. Wiegel. 2010. "A Generalized Aggregation-Disintegration Model for the Frequency of Severe Terrorist Attacks." *Journal of Conflict Resolution* 54(1): 179–197.

Clauset, A., M. Young, and K. S. Gleditsch. 2007. "On the Frequency of Severe Terrorist Events." *The Journal of Conflict Resolution* 51(1): 58–87.

Clauset, A., M. Young, and K. S. Gleditsch. 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(1), Article 12.

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

Correlates of War (Version 4.0). 2010. http://www.correlatesofwar.org/

Gabaix, X. 1999. "Zipf's Law for Cities: An Explanation." *Quarterly Journal of Economics* 114(3): 739–767.

Gabaix, X., P. Gopikrishnan, V. Plerou, and H. E. Stanley. 2006. "Institutional Investors and Stock Market Volatility." *The Quarterly Journal of Economics* 121(2): 461–504.

Gabaix, X., and Y. M. Ioannides. 2004. "The Evolution of City Size Distributions." In *Handbook of Urban and Regional Economics*, Vol. 4, edited by J. V. Henderson and J. F. Thisse, 2341–2378. Amsterdam: Elsevier Science.

Gleditsch, N. P., P. Wallensteen, M. Eriksson, M. Sollenberg, and H. Strand. 2002. "Armed Conflict 1946–2002: A New Dataset." *Journal of Peace Research* 39(5): 615–637.

Goldstein, M. L., S. A. Morris, and G. G. Yen. 2004. "Problems with Fitting to the Power-law Distribution." *The European Physical Journal B – Condensed Matter* 41(2): 255–258.

González-Val, R., A. Ramos, and F. Sanz-Gracia. 2013. "The Accuracy of Graphs to Describe Size Distributions." *Applied Economics Letters* 20(17): 1580–1585.

Härdle, W. 1990. *Applied nonparametric regression*. Econometric Society Monographs. Cambridge: Cambridge University Press.

Ioannides, Y. M., and H. G. Overman. 2003. "Zipf's Law for Cities: An Empirical Examination." *Regional Science and Urban Economics* 33: 127–137.

Ioannides, Y. M., and H. G. Overman. 2004. "Spatial Evolution of the US Urban System." *Journal of Economic Geography* 4(2): 131–156.

Johnson, N., S. Carran, J. Botner, K. Fontaine, N. Laxague, P. Nuetzel, J. Turnley, and B. Tivnan. 2011. "Pattern in Escalations in Insurgent and Terrorist Activity." *Science* 333: 81–84.

Johnson, N. F., M. Spagat, J. Restrepo, J. Bohorquez, N. Suarez, E. Restrepo, and R. Zarama. 2005. "From Old Wars to New Wars and Global Terrorism." arXiv:physics/0506213 [*physics.soc-ph*].

Kagan, Y. Y. 1997. "Earthquake Size Distribution and Earthquake Insurance." *Communications in Statistics. Stochastic Models* 13(4): 775–797.

Levy, J. S. 1983. *War in the Modern Great Power System, 1495–1975*. Lexington: University Press of Kentucky.

Levy, J. S., and T. C. Morgan. 1984. "The Frequency and Seriousness of War." *Journal of Conflict Resolution* 28(4): 731–749.

Mitzenmacher, M. 2004. "A Brief History of Generative Models for Power Law and Lognormal Distributions." *Internet Mathematics* 1(2): 226–251.

Pisarenko, V. F. 1998. "Non-linear Growth of Cumulative Flood Losses with Time." *Hydrological Processes* 12(3): 461–470.

Reed, W. J., and B. D. Hughes. 2002. "From Gene Families and Genera to Incomes and Internet File Sizes: Why Power Laws are so Common in Nature." *Physical Review E*, 66: 067103.

Richardson, L. F. 1948. "Variation of the Frequency of Fatal Quarrels with Magnitude." *Journal of the American Statistical Association* 43(244): 523–546.

Roberts, D. C., and D. L. Turcotte. 1998. "Fractality and Self-organized Criticality of Wars." *Fractals* 6(4): 351–357.

Rose, A. K. 2006. "Cities and Countries." *Journal of Money, Credit and Banking* 38(8): 2225–2245.

Saperstein, A. M. 2010. "A Comment on the Power Law Relation Between Frequency and Severity of Terrorist Attacks." *Peace Economics, Peace Science and Public Policy* 16(1), Article 7.

Sarkees, M. R., and F. W. Wayman. 2010. *Resort to War: A Data Guide to Inter-State, Extra-State, Intra-state, and Non-State Wars, 1816–2007*. Washington, DC: CQ Press.

Small, M., and J. D. Singer. 1982. *Resort to Arms*. Beverly Hills: Saga Publications.

Soo, K. T. 2005. "Zipf's Law for Cities: A Cross-country Investigation." *Regional Science and Urban Economics* 35: 239–263.

Sutton, J. 1997. "Gibrat's Legacy." *Journal of Economic Literature* 35(1): 40–59.

Uppsala Conflict Data Program (UCDP/PRIO) Armed Conflict Dataset (Version 5). 2011. http://www.pcr.uu.se/research/ucdp/datasets/

Zipf, G. 1949. *Human Behaviour and the Principle of Least Effort*. Cambridge, MA: Addison-Wesley.

Table 1. Armed conflict battle deaths: descriptive statistics by year

Year	Observations	Mean Size	Standard Deviation	Minimum	Maximum	Max. Location
1989	43	1,256.651	3,023.588	25	18,403	Ethiopia
1990	50	1,631.6	5,057.416	25	30,633	Ethiopia
1991	51	1,372.471	3,436.919	25	21,790	Iraq, Kuwait
1992	53	676.2453	1,142.743	25	4,989	Bosnia-Herzegovina
1993	45	852.6889	1,955.79	25	12,054	Angola
1994	47	727.0213	1,505.68	25	8,829	Afghanistan
1995	41	698.7318	1,249.098	25	5,061	Afghanistan
1996	41	591.0732	955.7285	25	3,533	Turkey
1997	39	927.3075	1,948.249	25	10,033	Congo
1998	40	881.8	1,297.505	25	4,891	Sudan
1999	39	2,035.283	7,521.621	25	47,192	Eritrea, Ethiopia
2000	37	2,016.649	8,161.813	25	50,000	Eritrea, Ethiopia
2001	36	603.6111	800.9718	25	3,407	Sudan
2002	32	551.4063	787.1231	25	3,947	Nepal
2003	30	697.5001	1,512.132	25	8,202	Australia, Iraq, United Kingdom, United States of America
2004	32	566.6875	891.6652	25	3,499	Iraq
2005	32	358.0313	533.4645	25	2,364	Iraq
2006	33	527.2122	853.8419	25	3,656	Iraq
2007	35	487.7714	1,049.312	25	5,828	Afghanistan
2008	37	738.6217	1,586.588	25	8,413	Sri Lanka
2009	36	858.3056	1,805.982	25	8,162	Sri Lanka
2010	30	640.6	1,425.88	25	6,374	Afghanistan

Source: UCDP Battle-related deaths dataset v5 (2011), available at: www.pcr.uu.se/research/ucdp/datasets/

Table 2. Power law fit

Data	Lower bound $\underline{S}$	Pareto exponent $\hat{a}$	Standard error	Power law test p-value	Power law vs. log-normal p-value
COW pool 1816–2007, deaths	9540	1.74	0.11	0.18	0.49
COW pool 1816–2007, relative deaths	0.60	1.90	0.12	0.58	0.26
UCDP yearly data:					
1989	25	1.42	0.07	0.00	0.30
1990	1058	2.11	0.30	0.81	0.67
1991	25	1.45	0.06	0.01	0.46
1992	25	1.48	0.07	0.02	0.17
1993	53	1.54	0.09	0.22	0.18
1994	935	2.48	0.43	0.58	0.59
1995	97	1.65	0.12	0.03	0.24
1996	25	1.52	0.08	0.09	0.32
1997	25	1.49	0.08	0.22	0.26
1998	25	1.45	0.07	0.00	0.54
1999	1403	2.26	0.40	0.98	0.65
2000	1010	2.14	0.35	0.92	0.66
2001	72	1.59	0.11	0.00	0.57
2002	1086	3.54	1.04	0.96	0.66
2003	478	2.27	0.37	0.71	0.68
2004	25	1.50	0.09	0.01	0.31
2005	1046	4.71	1.66	0.95	0.69
2006	191	1.99	0.22	0.67	0.50
2007	69	1.74	0.14	0.56	0.31
2008	295	1.91	0.23	0.93	0.46
2009	353	2.05	0.26	0.53	0.70
2010	175	1.83	0.21	0.95	0.50

Note: The lower bound and the Pareto exponent are estimated using Clauset, Shalizi, and Newman's (2009) methodology. The power law test is a goodness-of-fit test. H_0 is that there is power law behaviour for $S_i \geq \underline{S}$. The power law vs. log-normal test is Vuong's model selection test, based on the normalized log-likelihood ratio: H_0 is that both distributions are equally far from the true distribution while H_A is that one of the test distributions is closer to the true distribution.

Figure 1. The intensity of wars from 1816 to 2007, 447 observations

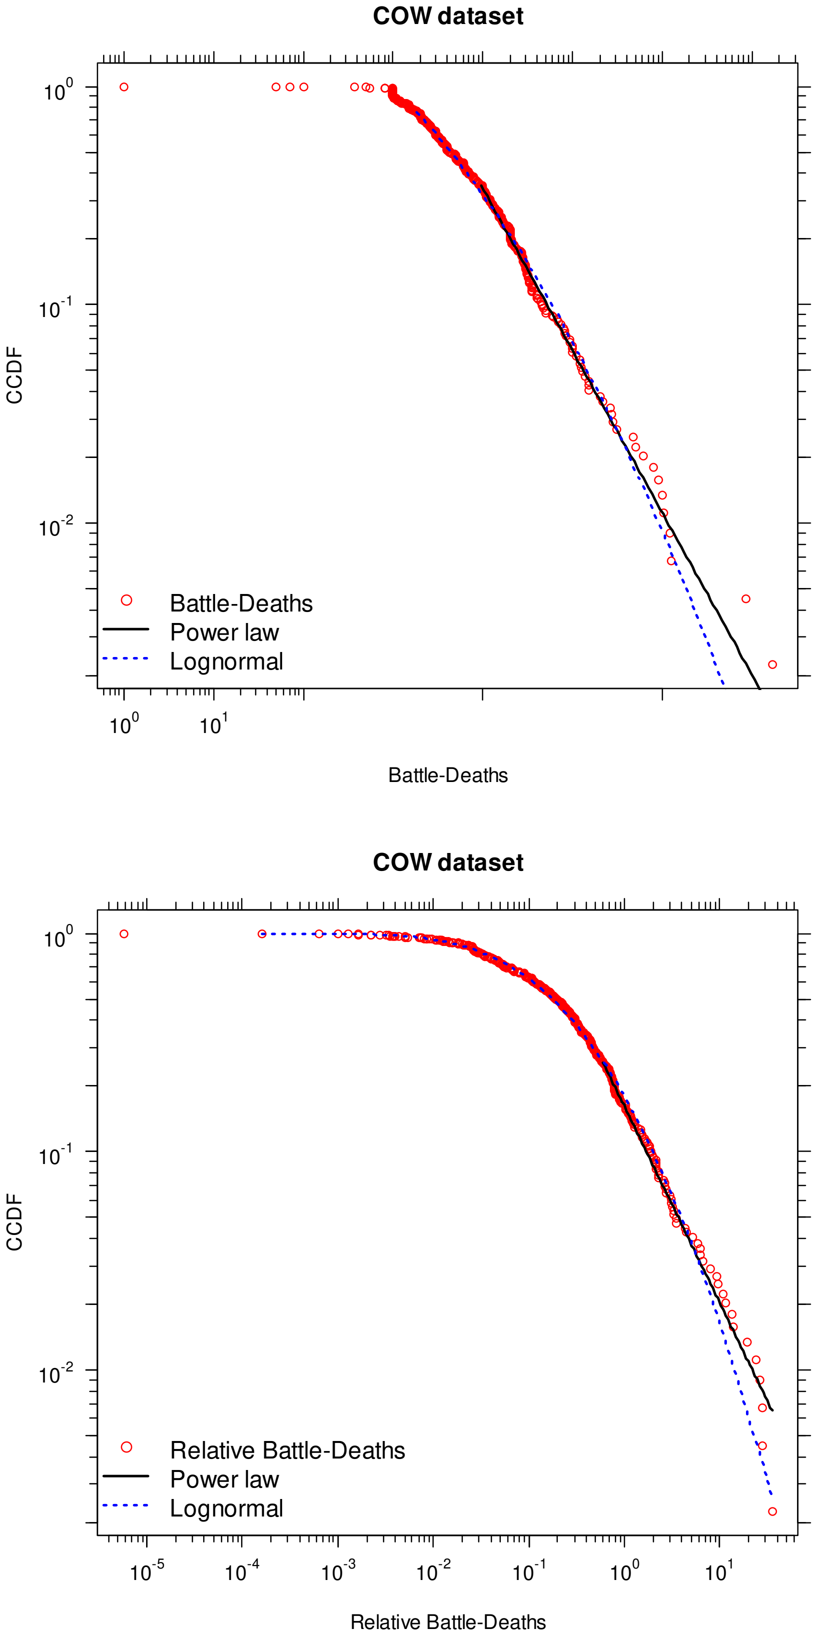

Note: COW inter-, intra-, and extra-state war data (v4.0). The data are plotted as a complementary cumulative distribution function (CCDF).

Figure 2. War size distribution in 1998 and 2007

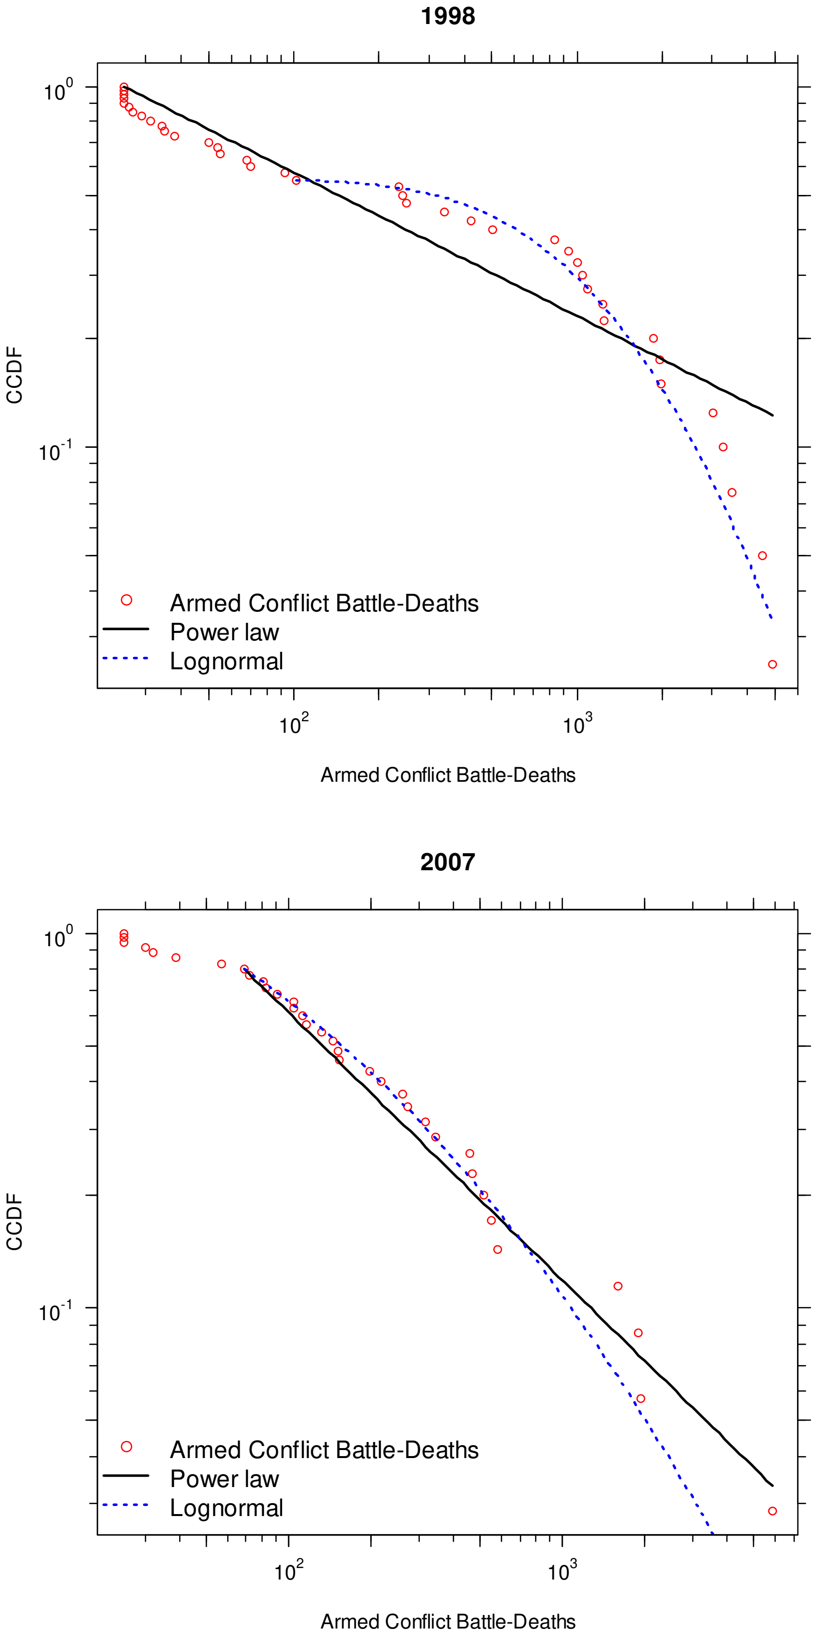

Note: UCDP Battle-related deaths dataset v5 (2011). The data are plotted as a complementary cumulative distribution function (CCDF).

Figure 3. Power law fit, UCDP yearly data

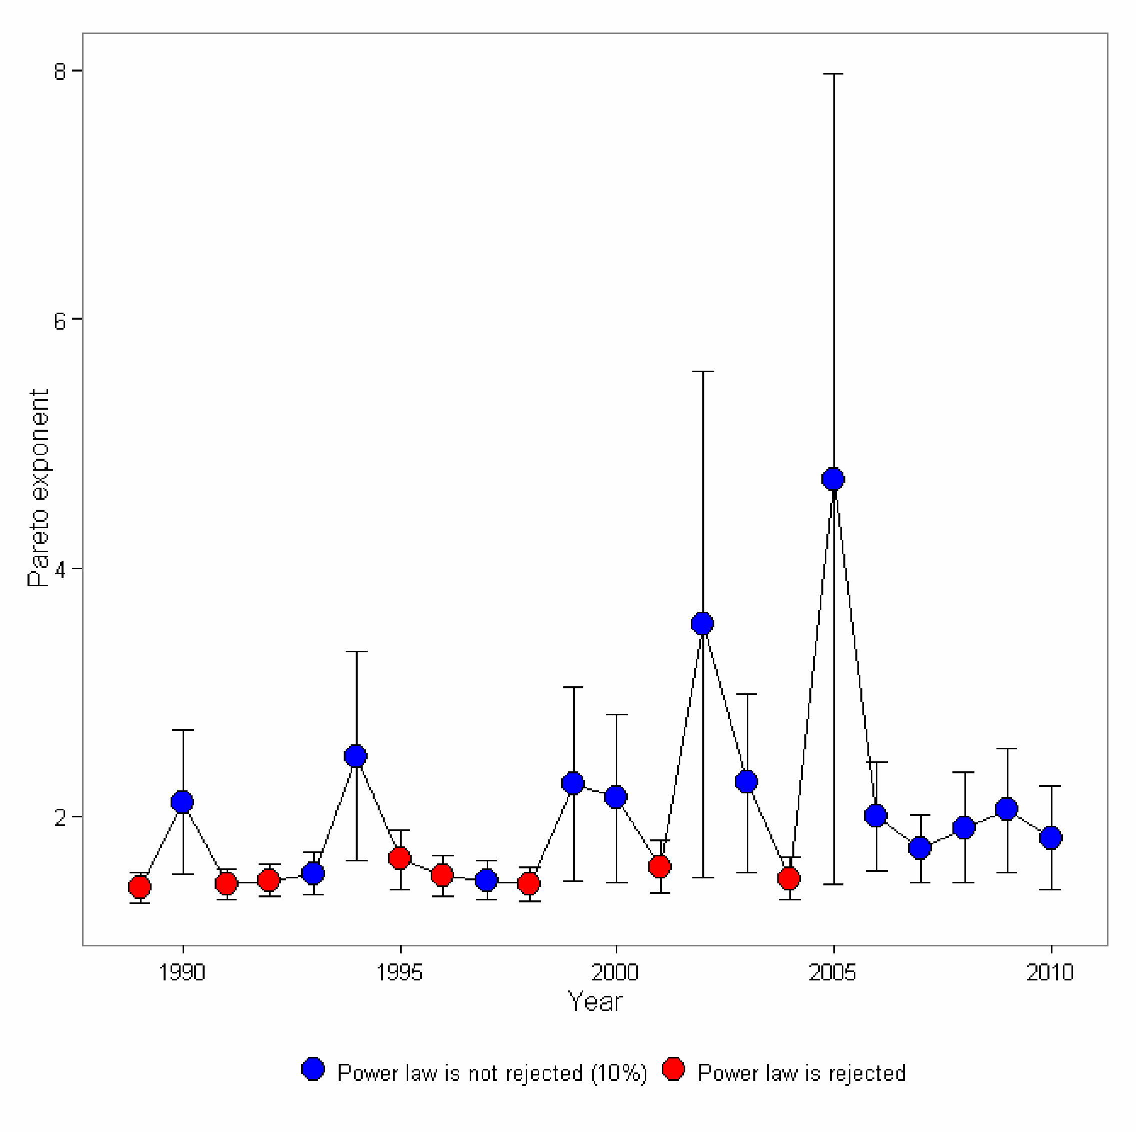

Notes: UCDP Battle-related deaths dataset v5 (2011). The Pareto exponent is estimated using Clauset, Shalizi, and Newman's (2009) methodology. The graph also shows the results of the power law goodness-of-fit test for the 10% significance level.

Figure 4. Stochastic kernel, battle deaths to growth rates

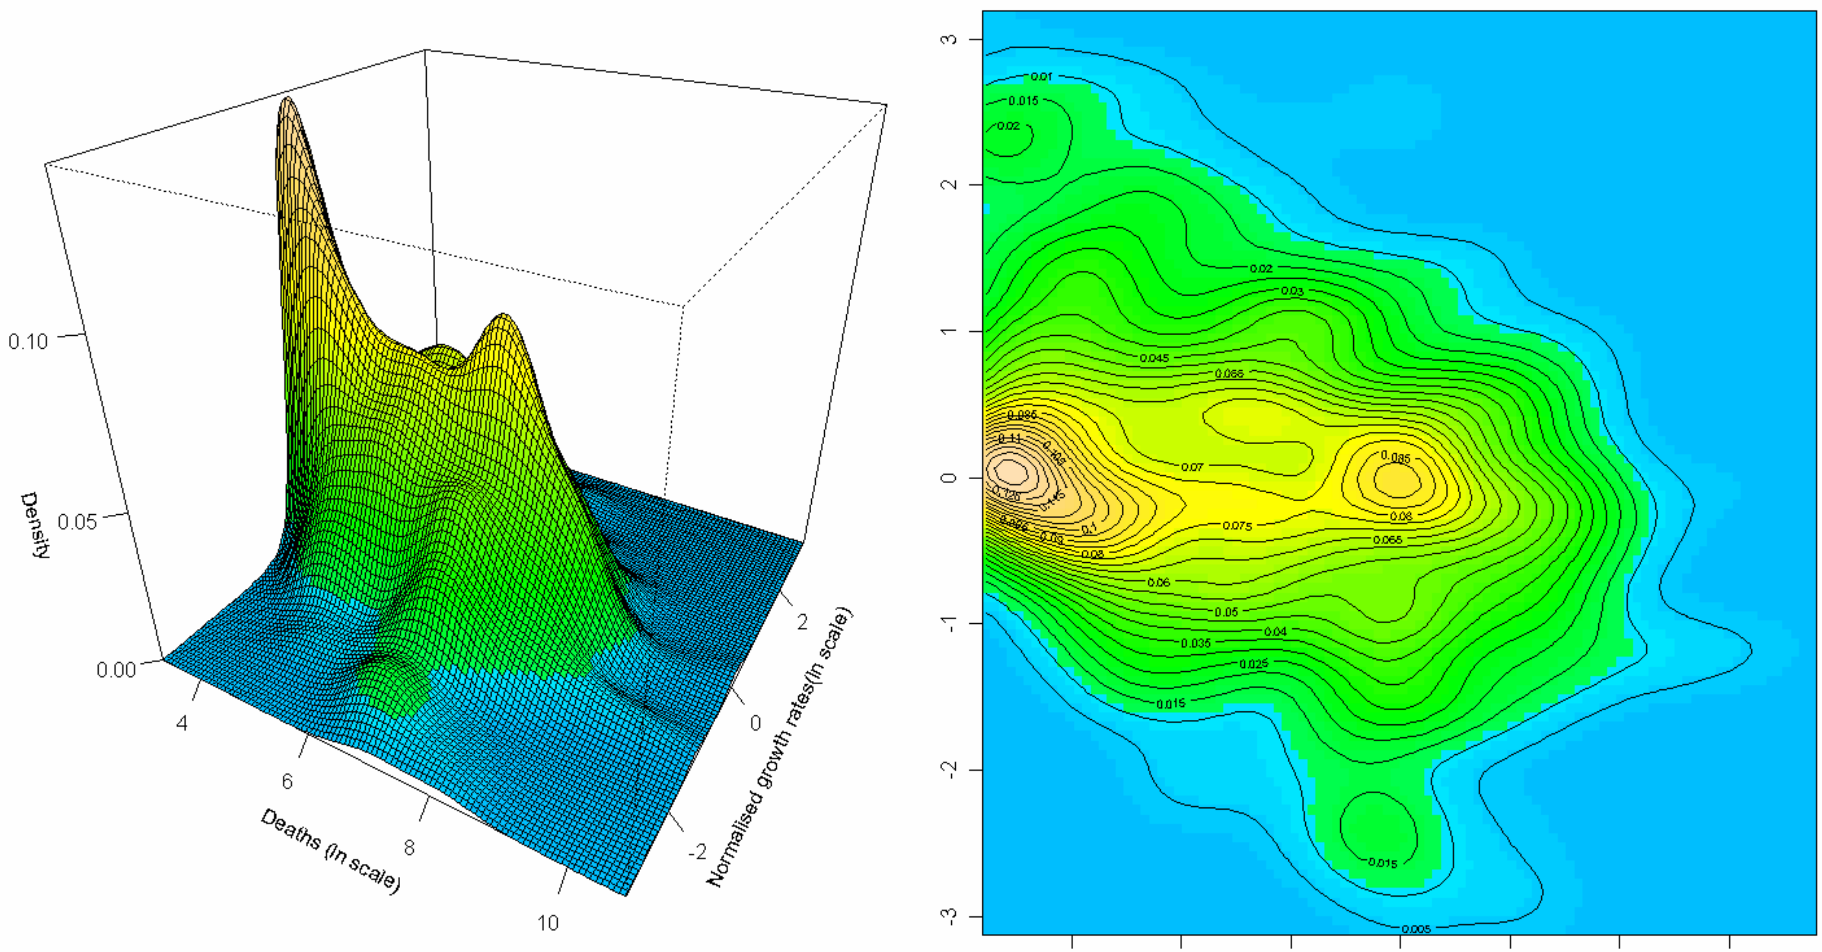

Note: UCDP Battle-related deaths dataset v5 (2011), 639 observations.

Figure 5. Kernel estimate of growth (bandwidth 0.5), 639 observations

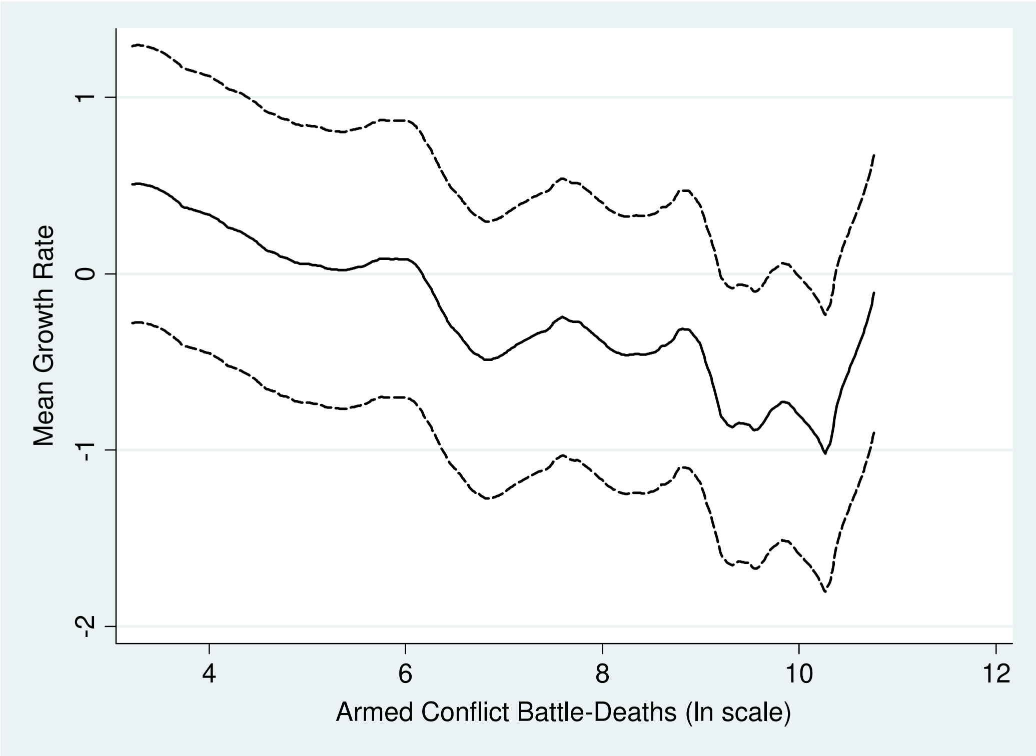

Note: UCDP Battle-related deaths dataset v5 (2011).