

Journal of Conflict Resolution

<http://jcr.sagepub.com/>

A Generalized Aggregation-Disintegration Model for the Frequency of Severe Terrorist Attacks

Aaron Clauset and Frederik W. Wiegel

Journal of Conflict Resolution 2010 54: 179 originally published online 10 December 2009

DOI: 10.1177/0022002709352452

The online version of this article can be found at:

<http://jcr.sagepub.com/content/54/1/179>

Published by:



<http://www.sagepublications.com>

On behalf of:

Peace Science Society (International)

Additional services and information for *Journal of Conflict Resolution* can be found at:

Email Alerts: <http://jcr.sagepub.com/cgi/alerts>

Subscriptions: <http://jcr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://jcr.sagepub.com/content/54/1/179.refs.html>

>> [Version of Record - Jan 15, 2010](#)

[OnlineFirst Version of Record - Dec 10, 2009](#)

[What is This?](#)

A Generalized Aggregation-Disintegration Model for the Frequency of Severe Terrorist Attacks

Journal of Conflict Resolution

54(1) 179–197

© The Author(s) 2010

Reprints and permission: <http://www.sagepub.com/journalsPermissions.nav>

DOI: 10.1177/0022002709352452

<http://jcr.sagepub.com>



Aaron Clauset¹ and Frederik W. Wiegel²

Abstract

The authors present and analyze a model of the frequency of severe terrorist attacks, which generalizes the recently proposed model of Johnson et al. This model, which is based on the notion of self-organized criticality and which describes how terrorist cells might aggregate and disintegrate over time, predicts that the distribution of attack severities should follow a power-law form with an exponent of $\alpha = 5/2$. This prediction is in good agreement with current empirical estimates for terrorist attacks worldwide, which give $\hat{\alpha} = 2.4 \pm 0.2$ and which the authors show is independent of certain details of the model. The authors close by discussing the utility of this model for understanding terrorism and the behavior of terrorist organizations and mention several productive ways it could be extended mathematically or tested empirically.

Keywords

terrorism, severe attacks, frequency statistics, scale invariance, Richardson's Law

Richardson's Law—one of the few robust statistical regularities in studies of political conflict—states that the distribution of casualties in violent conflicts follows a power-law form, in which the probability of an event with x deaths is $p(x) \propto x^{-\alpha}$, where α is a parameter called the scaling exponent (Cederman 2003; Richardson 1948, 1960). Recent studies have used rigorous statistical methods to confirm this statistical law for wars between 1816 and 1980 (Clauset, Shalizi, and Newman 2009; Newman 2005) and have extended it to cover the severity of individual terrorist attacks

¹Santa Fe Institute, NM

²University of Amsterdam, the Netherlands

Corresponding Author:

Aaron Clauset, Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe, NM, 87501, USA

Email: aaronc@santafe.edu

worldwide from 1968 to 2008 (Clauset, Shalizi, and Newman 2009; Clauset, Young, and Gleditsch 2007).

Although Richardson's original interest in violent conflicts was inclusive of both wars and homicides, most research on the severity of conflicts has focused on wars and other large-scale events, characterizing them mainly dichotomously according to their incidence or absence (with some exceptions; see Cederman 2003 and Lacina 2006). Research on terrorism has tended to be similarly focused (Li 2005), with considerable additional attention paid to its strategic elements (Enders and Sandler 2004; Pape 2003). As a result, little is known systematically about what factors and mechanisms influence the severity of terrorist attacks. This ignorance is exacerbated in part by the extreme scarcity of systematic, quantitative data about, for instance, the recruitment, fund-raising, decision making, and structure of terrorist organizations or about the counterterrorism efforts of states. But good-quality data about terrorist attacks themselves do exist, and their systematic analysis led to the discovery that Richardson's Law includes the severity of terrorist attacks.

Power-law distributions have recently attracted a great deal of interest across the sciences and have been found to characterize the distribution of a wide variety of natural and social phenomena. Examples of power-law distributed quantities include earthquakes, floods, and forest fires (Bak, Tang, and Wiesenfeld 1987; Malamud, Morein, and Turcotte 1998; Newman 2005) as well as city sizes, citation counts for scientific papers, the number of participants in strikes, and the frequency of words in written language (Biggs 2005; Newman 2005; Simon 1955; Zipf 1949). These distributions are scientifically interesting because they depart dramatically from Central Limit Theorem assumptions of normality (or even log-normality), and extremely large or severe events are orders of magnitude more likely than would normally be expected. Furthermore, the discovery that an empirical quantity follows a power-law distribution suggests certain unusual kinds of mechanistic explanations for their origin, for example, mechanisms that rely on long-range correlations, long-term memory effects, or positive feedbacks. In the case of terrorism, an understanding of the social or political mechanisms that govern the frequency of severe terrorist attacks would have strong implications for security policies. (Readers unfamiliar with power-law distributions can refer to Appendix A for a brief primer or to reviews by Kleiber and Kotz 2003; Newman 2005; and Mitzenmacher 2004.)

For terrorist attacks, recent analyses of empirical data suggest that the distribution of event severities, that is, the number of deaths or casualties, follows such a power law and that this statistical pattern has been largely stable during the past forty years despite large changes in the global political system during the same period (Clauset, Young, and Gleditsch 2007). (For concreteness, we reproduce this result from Clauset, Young, and Gleditsch 2007 in Figure 1.) Clauset, Young, and Gleditsch (2007) further showed that the severity of terrorist attacks remains power-law distributed, although with different scaling exponents, even after controlling for the type of weapon used (e.g., firearms, explosives, etc.) or the level of economic development of the target country, but not when controlling for geographic region (e.g., North America, Europe, etc.) or tactic (e.g., hostage, assassination, suicide bombing). Other studies of the severity of such

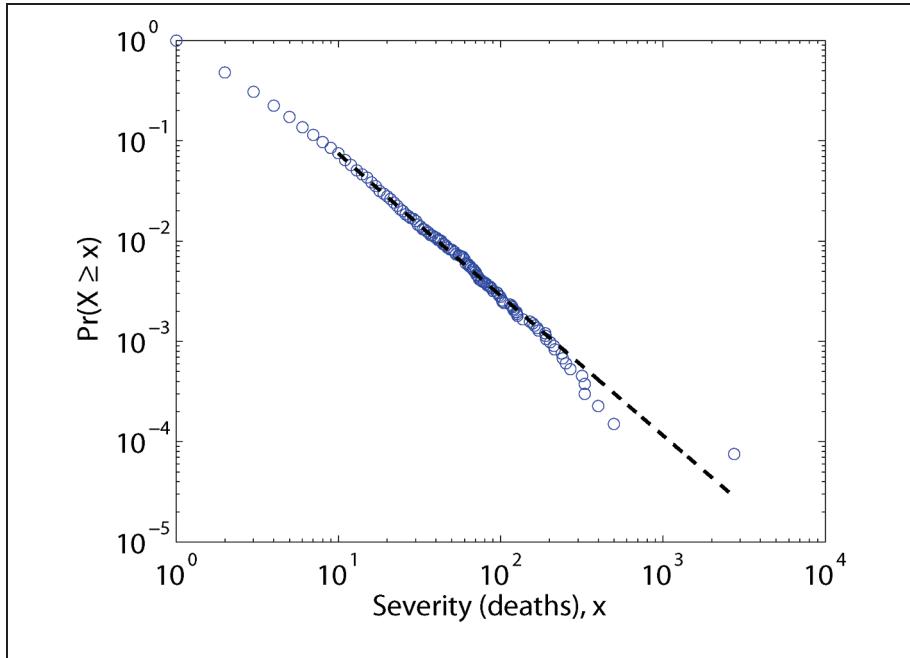


Figure I. The severities (number of deaths) for 13,274 fatal terrorist attacks worldwide from 1968 to 2008 (National Memorial Institute for the Prevention of Terrorism 2008)

NOTE: The data are plotted as a complementary cumulative distribution function $\Pr(X \geq x)$. The solid black line shows the power-law behavior of the distribution, with scaling exponent $\hat{\alpha} = 2.4 \pm 0.2$ for $x \geq 10$ (Clauset, Shalizi, and Newman 2009).

attacks go further, suggesting that the frequency and severity of events within individual conflicts, such as those in Colombia and Iraq, exhibit power-law statistics (Cioffi-Revilla and Romero 2009; Johnson et al. 2005, 2006) and that observable changes in the power law's exponent over time are indicative of real and important shifts in the underlying dynamics of the social and political generative processes.

At present, these ubiquitous power-law statistics lack a clear and well-supported explanation: what mechanisms, political or otherwise, give rise to these law-like behaviors? A scientific answer to this question may ultimately shed light, in a manner complementary to traditional studies, on the use of such tactics in violent conflicts (Li 2005; Pape 2003), the internal dynamics of terrorist organizations (Clauset and Gleditsch 2009; Cordes et al. 1985), and trends in global terrorism (Enders and Sandler 2000; U.S. Department of State 2004). It may also shed light on the connection between severity and other modalities (Clauset et al. forthcoming), for example, location and timing; suggest novel intervention strategies or policy recommendations for counterterrorism (Cronin 2003); and shed light on the connection between terrorism and other kinds of violent conflicts, such as civil and international wars (Richardson 1960; Small and Singer 1982).

To date, two explanations have been proposed for the origin of the observed power law in the frequency of severe terrorist attacks.¹ One, proposed by Clauset, Young, and Gleditsch (2007), relies on an exponential sampling mechanism in which states and terrorists compete to decide which planned events become real. In this model, terrorists invest time planning events, and the potential severity of these increases roughly exponentially with the total planning time. Through counterterrorism actions by states, along with other natural attrition factors, these potential events are then strongly sampled, with the probability of a potential event's becoming real decreasing roughly exponentially with the size of the event. That is, large events are exponentially less likely to become real than smaller events. The competition of these two exponentials produces a power-law distribution in the severity of events, with the scaling exponent α depending only on the two exponential rates.

The second mechanism, proposed by Johnson et al. (2005, 2006), is a self-organized critical model (Bak, Tang, and Wiesenfeld 1987) of the internal dynamics of a modern terrorist organization. In this model, a terrorist organization is composed of cells that merge and fall apart according to simple probabilistic rules (see below). The long-term dynamics of this aggregation-disintegration process produces a dynamic equilibrium or steady state that is characterized by a power-law distribution in the sizes of cells and, by assumption, a power-law distribution in the severity of events. In this model the scaling exponent in the steady state can be calculated exactly and is found to be $\alpha = 5/2$. This value is in good agreement with the best current empirical estimate of $\hat{\alpha} = 2.4 \pm 0.2$ (Clauset, Shalizi, and Newman 2009) for terrorist attacks worldwide from 1968 to 2008.

In this article, we mathematically study the Johnson et al. (2005, 2006) model. In particular, we generalize Johnson et al.'s specific model to a family of such models. We then analytically solve for their steady-state behavior and show that a power-law distribution is a *universal* feature² of this class of models. That is, provided the number N of radicalized individuals is large $N \gg 1$, the appearance of the power-law distribution and the value of its scaling exponent α does not depend on certain details of the model itself. Mathematically speaking, our analysis is exact in the limit $N \rightarrow \infty$. We note that our asymptotic analysis is done purely for mathematical convenience; the limit $N \rightarrow \infty$ has no social meaning, and so long as N is very large, our results should hold.

The benefits of generalizing the Johnson et al. (2005, 2006) model are twofold. First, there is the generalization itself, which extends the model in a new and important direction and demonstrates that the model's main qualitative result—the power-law distribution in event sizes—is robust to certain specific modeling assumptions. Second, by carefully describing the model's assumptions and then mathematically working out their consequences, we can more precisely identify which empirical tests are ultimately necessary to support or refute the model's assumptions and predictions. This approach defers answering the question of which mechanism produces the power-law distribution in the frequency of severe terrorist attacks; however, this seems acceptable partly because of the complexity of the model and its analysis and partly because of our currently very limited knowledge of the social and political processes that might give rise to the power-law distribution. This model-based

approach can thus highlight which empirical facts it would be useful to know and stimulate research in productive directions.

I. The Model

The model we analyze is based on five assumptions about the interaction of the terrorist cells that make up a modern terrorist organization. We make no other assumptions about the relationship between these cells and the conflict or the terrorist organization they inhabit, about the mode of attack or tactic used by an attacking cell, or that this model represents the behavior of hierarchical terrorist organizations.

Although these assumptions are straightforward to state, and allow us to mathematically analyze their consequences, they embody strong and possibly unrealistic constraints on the internal dynamics of terrorist groups that have not yet been systematically tested with empirical data. At present, however, this model is worthwhile to study mainly because it yields one prediction—a power-law distribution in the frequency and severity of events—that agrees relatively well with a wide range of empirical data (Cioffi-Revilla and Romero 2009; Clauset, Shalizi, and Newman 2009; Clauset, Young, and Gleditsch 2007; Johnson et al. 2005, 2006). By carefully exploring the behavior of this model, we can identify quantitative predictions or critical assumptions that may be tested using the available empirical data. In our concluding remarks, we discuss some of these tests and possible extensions of the model that relax some of the model’s assumptions.

The five model assumptions are the following:

1. There is a pool of N radicalized individuals who are inclined toward terrorism. We assume N to be large $N \gg 1$ and to be constant in time. This latter assumption implies that terrorists who are eliminated for any reason, for example, by counterterrorism measures, by intercell or intracell conflict, by personal preferences, or in the course of their attacks, are replaced immediately by an equal number of radicalized individuals.
2. These individuals can form cells of size $1, 2, 3, 4, \dots$. Let n_k denote the number of cells consisting of $k = 1, 2, 3, \dots$ individuals.
3. Cells grow by a process of aggregation, in which any pair of cells can merge to form a larger cell. Specifically, we assume that any pair of cells consisting of k and ℓ individuals, respectively, has a probability $A_0(k\ell)^a$ per unit time to combine into a cell of size $k + \ell$. Here $A_0 > 0$ and $a \geq 0$ are parameters of the model, and we analyze the model for general a . To be realistic when comparing with data, however, we choose $a \cong 1$ to represent the fact that the number of possible human relations between members of the two cells is $k\ell$; that is, it scales linearly with the product of the cell sizes.
4. Cells fall apart or “disintegrate” spontaneously into single individuals. Let $b(k)$ denote the probability per unit time that a given cell of k individuals will disintegrate spontaneously into k cells of size one and where $b(1) = 0$. The

explicit form of the function $b(k)$ is not needed to calculate the equilibrium distribution of cell size provided one studies the asymptotic region $N \gg 1$.

5. At any time, any cell can launch an attack. For simplicity, we assume that the attack occurs with probability (per unit time) that is independent of the cell's size, its age, the number of attacks it has previously launched, and so forth and that the severity $v(k)$ of an attack is roughly proportional to the cell's size k , that is, $v(k) \propto k$, for $1 \ll k \ll N$.

To be precise, the number of possible pairings of a k -cell with an ℓ -cell, that is, the number of potential combinations between some cell of size k and some cell of size ℓ , equals $n_k n_\ell$ for $k \neq \ell$ and $\frac{1}{2} n_k (n_k - 1)$ for $k = \ell$. However, if $N \gg 1$, we shall find that all $n_k \gg 1$; in this case, we can approximate $\frac{1}{2} n_k (n_k - 1) \cong \frac{1}{2} n_k^2$, which simplifies the mathematics considerably but does not fundamentally alter the results.

Our analysis of this model will show that the steady-state distribution of the sizes of the terrorist cells follows a power-law distribution with exponent $\alpha = 5/2$. By assumption 5, that the severity of an attack is proportional to the size of the attacking cell, this then implies that the distribution of event severities follows a power-law distribution with the same exponent.

2. The Distribution of Cell Sizes in the Steady State

From the five assumptions discussed above, we can write down the equation for how $n_k(t)$ changes with time for $k = 2, 3, \dots$:

$$\frac{dn_k}{dt} = \frac{1}{2} A_0 \sum_{i,j=1}^{\infty}' i^a j^a n_i n_j - A_0 k^a n_k \sum_{j=1}^{\infty} j^a n_j - b(k) n_k, \quad (1)$$

where \sum' denotes a summation over all natural numbers i and j such that

$$i + j = k. \quad (2)$$

The equation for dn_1/dt is not needed in our analysis. In words, the first term on the right-hand side of equation 1 represents the increase of the number of cells of size k because of the aggregation of two smaller cells, the second term measures the decrease of this number because such a cell can itself merge with another cell, and the third term represents the loss of these cells because of spontaneous disintegration.

As we are interested mainly in the steady-state behavior of this model, we denote $\lim_{t \rightarrow \infty} n_k(t)$ by n_k^* , where $*$ is not an exponent but a label that denotes the attached variable being in its steady-state limit. Equation 1 now simplifies to

$$\frac{1}{2} A_0 \sum_{i,j}' i^a j^a n_i^* n_j^* = A_0 k^a n_k^* \sum_j j^a n_j^* + b(k) n_k^*, \quad (3)$$

for $k = 2, 3, \dots$. As a technical detail, we point out that the term with $j = k$ in the second summation in the right-hand sides of equations 1 and 3 comes from the fact that the

number of pairs k, k equals $\frac{1}{2}n_k^2$ (see section 1), but as each combination of two such cells leads to the decrease of n_k by two, the loss term is proportional to $2 \cdot \frac{1}{2}n_k^2 = n_k^2$.

A simple way of solving the set of equations given in equation 3 is by introducing the generating functions (Wilf 2006):

$$f(z) \equiv \sum_{k=1}^{\infty} k^a n_k^* z^k \quad (4)$$

$$g(z) \equiv \sum_{k=1}^{\infty} b(k) n_k^* z^k. \quad (5)$$

That is, we multiply equation 3 by z^k and then sum over k from 2 to ∞ . This reduces our system of equations to

$$\frac{1}{2} A_0 f(z) f(z) = A_0 f(1) \{f(z) - n_1^* z\} + g(z), \quad (6)$$

where we used the fact that $b(1) = 0$ because a cell of one individual cannot disintegrate into single individuals. (Readers unfamiliar with generating functions can refer to Appendix B for a brief primer, and to Wilf 2006 for a more thorough introduction.)

Although the solution of equation 6 is difficult for general z and N , it is much simpler in our case where z is fixed and the limit $N \rightarrow \infty$ is studied. For $N \gg 1$, the equilibrium frequencies n_k^* will be proportional to N (for k smaller than some cutoff k_0 , which we need not calculate explicitly; see Appendix C). Hence the leading orders of magnitude (in N) of the various terms in equation 6 are

$$f(z) \sim N \quad (7)$$

$$g(z) \sim N \quad (8)$$

$$\frac{1}{2} A_0 f(z) f(z) \sim N^2 \quad (9)$$

$$A_0 f(1) \{f(z) - n_1^* z\} \sim N^2. \quad (10)$$

This means that for z fixed and $N \gg 1$, equation 6 can be replaced by

$$\frac{1}{2} f^2(z) - f(1)f(z) + f(1)n_1^* z = 0, \quad (11)$$

which has the solution

$$f(z) = f(1) - \sqrt{f^2(1) - 2f(1)n_1^* z}. \quad (12)$$

Substituting $z = 1$ shows

$$f(1) = 2n_1^*, \quad (13)$$

and gives

$$f(z) = 2n_1^* \left\{ 1 - \sqrt{1-z} \right\}. \quad (14)$$

The definition of $f(z)$ given in equation 4 shows that the term $k^a n_k^*$ can now be found as the coefficient of z^k in the power series expansion of equation 14. For small values of k , these coefficients can be calculated by hand from the series

$$f(z) = 2n_1^* \left(\frac{1}{2}z + \frac{1}{2} \cdot \frac{1}{4}z^2 + \frac{1}{2} \cdot \frac{1}{4} \cdot \frac{3}{6}z^3 + \frac{1}{2} \cdot \frac{1}{4} \cdot \frac{3}{6} \cdot \frac{5}{8}z^4 + \dots \right). \quad (15)$$

For example, the first four terms are

$$2^a n_2^* = \frac{1}{4} n_1^*, \quad (16)$$

$$3^a n_3^* = \frac{1}{8} n_1^*, \quad (17)$$

$$4^a n_4^* = \frac{5}{64} n_1^*, \quad (18)$$

$$5^a n_5^* = \frac{7}{128} n_1^*. \quad (19)$$

To obtain the coefficients for $k \gg 1$, one can use Cauchy's theorem, which gives the contour integral

$$k^a n_k^* = i \frac{n_1^*}{\pi} \oint_C z^{-k-1} \sqrt{1-z} dz, \quad (20)$$

where i denotes the imaginary number and where the contour C encircles the origin of the complex z -plane once in the counterclockwise direction. This contour can be deformed into a contour C' , which encircles the branch cut $1 \leq z < \infty$ once in the clockwise direction. For z near to the branch point at $z = 1$, it is convenient to first write

$$z = 1 + \zeta \quad (21)$$

$$z^{-k-1} \cong e^{-(k+1)\zeta}. \quad (22)$$

When ζ has a small positive imaginary part, one can write $\sqrt{-\zeta} = -i\sqrt{|\zeta|}$; when ζ has a small negative imaginary part, one writes $\sqrt{-\zeta} = +i\sqrt{|\zeta|}$. Hence, we find the asymptotic result

$$k^a n_k^* \cong \frac{2}{\pi} n_1^* \int_0^\infty \sqrt{\zeta} e^{-(k+1)\zeta} d\zeta = \frac{1}{\sqrt{\pi}} n_1^* (k+1)^{-3/2}, \quad (23)$$

for $k \gg 1$. (An alternative approach to this result would express $\{1 - \sqrt{1-z}\}$ as a ratio of Γ -functions and use asymptotic analysis.) For k as small as 5, the last equation gives reasonably close approximations of the true values, for example, for $5^a n_5^*$ the value of 0.038 whereas the exact value (from equation 19) is 0.055.

This analysis thus shows that the number of cells consisting of k terrorists, at equilibrium, is given by the power law

$$n_k^* \cong \frac{1}{\sqrt{\pi}} n_1^* k^{-a-3/2}, \quad (24)$$

for $k \gg 1$. Hence, because of model assumption 5, that the severity of an event is proportional to the size of the attacking cell, the probability p_k that a terrorist attack will claim k victims will also have a power-law distribution is

$$p_k \propto k^{-\alpha}, \quad (25)$$

for $k \gg 1$, with an exponent

$$\alpha = a + 3/2. \quad (26)$$

As mentioned before, we assume that $a \cong 1$ (see section 1), which leads to the prediction

$$\alpha = 5/2. \quad (27)$$

In fact, for $a = 1$ and $b(k) \propto k$, this model can be solved exactly, that is, with no approximations, and doing so recovers the results of Johnson et al. (2005, 2006).

The value in equation 27 is in good agreement with recent estimates from empirical data (Clauset, Shalizi, and Newman 2009; Clauset, Young, and Gleditsch 2007), which give $\hat{\alpha} = 2.4 \pm 0.2$ for terrorist attacks worldwide since 1968.

3. Concluding Remarks

Thus, we find that the class of dynamical models studied here produces a steady state in which the number of terrorist cells of size k , and by assumption the severity of their attacks, follows a power-law distribution. This feature implies that the dynamics of this model system are characterized by self-organized criticality (Bak, Tang, and Wiesenfeld 1987). Furthermore, we find that the scaling exponent of this distribution $\alpha = 5/2$ is (for $N \gg 1$) independent of the manner in which terrorist cells disintegrate (represented by the function $b(k)$). That is, whether cells tend to disintegrate due to internal conflict, external efforts, some combination of these, or other factors does not change the fundamental character of the frequency-severity distribution of attacks. In this sense, the statistical properties predicted by these models show a form of universality.

However, other statistical properties of the model should depend on the function $b(k)$ in a crucial way. For the record, we give three such properties.

- The explicit determination of n_1^* , the number of lone terrorists, is a function of N , the number of radicalized individuals.
- A terrorist cell will grow in the course of time by combining occasionally with a smaller cell. As a result, the size of a particular cell will be time dependent. For $a = 1$ in particular, we find that the size of a terrorist cell increases exponentially with time. Similarly, each cell of size $k > 2$ has a probability to disintegrate, which will also be time dependent.
- The previous problem is especially interesting if one starts with a single, radicalized individual. The theory presented here makes it possible to calculate the speed with which such an individual cycles through cells of various sizes, in the steady state.

From a policy perspective, an important question for this model concerns the difficulty of inducing qualitative changes in the steady-state behavior via realistic interventions. For instance, the independence of the model system's behavior from the particular manner in which cells disintegrate suggests that efforts focused mainly on breaking up terrorist cells may not produce long-term changes in the severity of terrorist attacks unless they are paired with additional interventions, such as reducing the pool of radicalized individuals by other means. On the other hand, the aggregation process, that is, the manner in which terrorist cells can achieve coordinated behavior, is a clear target, and its frustration may have a strong influence on the frequency of severe attacks. We leave for future work the articulation of specific intervention strategies based on this model.

Because many questions remain about the accuracy of this model for understanding modern terrorism and its utility for counterterrorism efforts, we remain modest about its long-term value. First, there is the question of the dependence of the central prediction—the power-law distribution in the frequency of severe attacks—on the particular assumptions we have described here. Already we have shown that the power-law prediction does not depend on the function $b(k)$, and it may be that other model assumptions can also be eliminated, relaxed, or made more realistic while this behavior is preserved (see, for instance, Ruszczycki et al. 2008).

For example, in most conflicts, the number of radicalized individuals N is unlikely to remain constant and may not vary slowly relative to the replacement of individuals lost from counterterrorism activities and so forth or relative to the aggregation-disintegration dynamics. Changes in N should thus induce perturbations to the model's steady-state behavior. Furthermore, empirical research may show that cells do not launch attacks with probability independent of their size. If larger cells launched attacks more frequently than smaller cells, it could be possible to adjust the aggregation dynamics so as to produce correspondingly fewer of these large cells, thus leaving the qualitative behavior of the model unchanged.

Existing analyses have focused on the steady-state behavior, but real organizations may exhibit a transient period of non-power-law behavior during which they

self-organize to the critical state. The character and duration of this transient behavior depends on the initial distribution of cell sizes, but for reasonable initial conditions, it is unknown what specific behavior we should expect. Finally, the strategic utility of terrorist attacks is widely accepted (Clauset et al. forthcoming; Enders and Sandler 2004; Kydd and Walter 2002; Pape 2003). However, the model assumes that attacks are largely stochastic in nature, and it is unknown whether these two perspectives can be reconciled. Research on this model would benefit greatly from mathematical generalizations that move us toward discovering the most general version that still produces the power-law distribution.

Second, although the model correctly predicts the distribution of event severities, this agreement is a relatively indirect test of the model's accuracy, and a stronger test would consider the accuracy of the model's specific assumptions or its predicted dynamics. Tests along these lines may also point out the most useful mathematical generalizations. For the record, we describe a number of ways the model can be tested.

Anecdotal evidence, including post hoc analyses of severe events like the September 11 attacks (Sageman 2004), suggests that extremely severe attacks often require significantly more resources and manpower than small-severity attacks (see Clauset, Young, and Gleditsch 2007 for additional discussion), but it is unknown whether this is a systematic relationship and whether the precise form of assumption 5, that is, $v(k) \propto k$, is sufficiently accurate. Without access to data about the internal dynamics of terrorist organizations, a direct test of this assumption seems impossible. However, research on determining which factors correlate with the severity of terrorist attacks may indirectly address this question (for instance, see Asal and Rethemeyer 2008; Clauset and Gleditsch 2009; Clauset, Young, and Gleditsch 2007; Harrison 2006).

Furthermore, the assumption that cells initiate attacks independently of their age or history may prove to be overly simplistic, and systematic correlations could produce deviations from the expected power-law form. That being said, recent work finds no significant deviations from a power-law distribution for attacks worldwide that killed at least ten individuals (Clauset, Young, and Gleditsch 2007), and it remains to be seen whether other kinds of systematic correlations exist. The model defined here also predicts that the severity of attacks by individual terrorist organizations should follow a power law. Johnson et al. (2005, 2006) previously analyzed the conflict in Colombia, which is largely defined by the actions of the Revolutionary Armed Forces of Colombia and found evidence supporting this fact. However, a more systematic study of individual organizations is needed to fully vet this hypothesis.

Deviations, however, may not mean that the entire model is incorrect. Non-power-law behavior could be indicative of the aforementioned transient, noncritical behavior. In addition, the model assumes that terrorist cells interact with other cells only within the same organization—for example, Taliban fighters do not aggregate with Revolutionary Armed Forces of Colombia fighters—however, some evidence suggests that cells sometimes do interact across organizational boundaries, for instance, between organizations involved in the same conflict (Araj 2008; Pedahzur and Perlinger 2006), between allied organizations (Sageman 2004), or when fighters from different conflicts are jailed together (McKeown 2001). Thus, in some cases, the set of cells that

constitutes a “group” in the sense of the model may not correspond to a single identifiable terrorist organization; instead, a group may be a somewhat amorphous set of cells, spread over multiple organizations. Thus, the set of events by which to test the power-law hypothesis may not always break cleanly at organizational boundaries. The extent to which the network of organizational alliances worldwide structures and constrains the set of possible interactions between cells is largely unknown but likely plays an important role in the global dynamics of terrorism.

This discussion points to a more critical test of the accuracy of the model: validating the aggregation-disintegration dynamics themselves. As described above, without detailed data about the internal organizational dynamics or about the actions of many individual fighters, this part of the model seems difficult to test directly. Conventional wisdom suggests that aggregation-disintegration dynamics are unrealistic, as interactions between cells could pose security risks to the larger organization or to ongoing operations. However, recent analyses of organizations involved in the “global jihad” indicate that interactions, including aggregations, do indeed occur with some frequency and that such interactions may be critical to the execution of particularly severe attacks (Fouda and Fielding 2003; Sageman 2004). But it remains unclear how often such aggregations occur, how widespread they are, and how necessary they are to the execution of large attacks. Taking this anecdotal evidence at face value, it still remains unclear whether they occur frequently enough to allow an organization or set of cells to converge on the critical state—exhibiting the power-law distribution in cell sizes—in a timely fashion.

Ideally, all of these assumptions and predictions will be tested with empirical data to determine just how realistic, and thus how useful, this model is. Because of the scarcity of systematic, quantitative data about terrorism, some of these assumptions may prove impossible to test directly. On the other hand, by focusing on the model’s testable predictions, it may be possible to test the model indirectly using available data. These empirical tests, along with the mathematical tests of the dependence of the power-law result on the model’s particular assumptions, are promising avenues for future work on Richardson’s Law as applied to the severity of terrorist events.

Appendix

A. Power-Law Distributions

Some readers may be unfamiliar with power-law distributions (sometimes also called “Zipf’s law” or “Pareto distributions” after two early researchers who championed their study; Pareto 1896; Zipf 1949), and this section is to serve as a brief, and somewhat informal, primer on the topic. What distinguishes a power-law distribution from the more familiar Normal distribution is its heavy tail. That is, in a power law, there is a nontrivial amount of weight far from the distribution’s center. This feature, in turn, implies that events orders of magnitude larger (or smaller) than the mean are relatively common. The latter point is particularly true when compared to a Normal distribution, where there is essentially no weight far from the mean.

Although there are many distributions that exhibit heavy tails, the power law is special and exhibits a straight line with slope α on doubly logarithmic axes. (Note that some data being straight on log-log axes is a necessary but not a sufficient condition of being power-law distributed.) This behavior is termed *scale invariance* because the power law admits the following property: multiplying its argument by some factor k results in a change in the corresponding frequency that is independent of the function's argument. For example, if $p(x) = Cx^{-\alpha}$, then

$$p(k \cdot x) = Ck^{-\alpha}x^{-\alpha} = k^{-\alpha}p(x),$$

for every value x . For this reason, the exponent α is called the “scaling exponent” (for historical reasons, $\alpha - 1$ is sometimes called the “Pareto exponent”), and the distribution is said to “scale.” This property also implies that there is no qualitative difference between large and small events.

Power-law-distributed quantities are not uncommon, and many characterize the distribution of familiar quantities. For instance, consider the populations of the 600 largest cities in the United States (from the 2,000 census). Among these, the average population is only $\bar{x} = 165,719$, and metropolises like New York City and Los Angeles seem to be outliers relative to this size. One clue that city sizes are not well explained by a Normal distribution is that the sample standard deviation $\sigma = 410,730$ is significantly larger than the sample mean. Indeed, if we modeled the data in this way, we would expect to see 1.8 times fewer cities at least as large as Albuquerque (population 448,607) than we actually do. Furthermore, because it is more than a dozen standard deviations above the mean, we would never expect to see a city as large as New York City (population 8,008,278), and the largest we would expect would be Indianapolis (population 781,870).

Figure 2 shows the empirical data for these 600 cities, plotted on doubly logarithmic axes as a complementary cumulative distribution function $Pr(X \geq x)$ (the standard way of visualizing this kind of data). The scaling behavior of these empirical data is clear, and the corresponding power-law model (black line) a reasonably good fit. In contrast, a truncated normal model is a terrible fit. These notions of goodness of fit can be made precise using an appropriately defined significance test, such as the one described by Clauset, Shalizi, and Newman (2009).

As a more whimsical second example, consider a world where the heights of Americans were distributed as a power law, with approximately the same average as the true distribution (which is convincingly Normal when certain exogenous factors are controlled). In this case, we would expect nearly 60,000 individuals to be as tall as the tallest adult male on record, at 2.72 meters. Furthermore, we would expect ridiculous facts such as 10,000 individuals being as tall as an adult male giraffe, 1 individual as tall as the Empire State Building (381 meters), and 180 million diminutive individuals standing a mere 17 centimeters tall. In fact, this same analogy was recently used to describe the counterintuitive nature of the extreme inequality in the wealth distribution in the United States (Crook 2006), whose upper tail is often said to follow a power law.

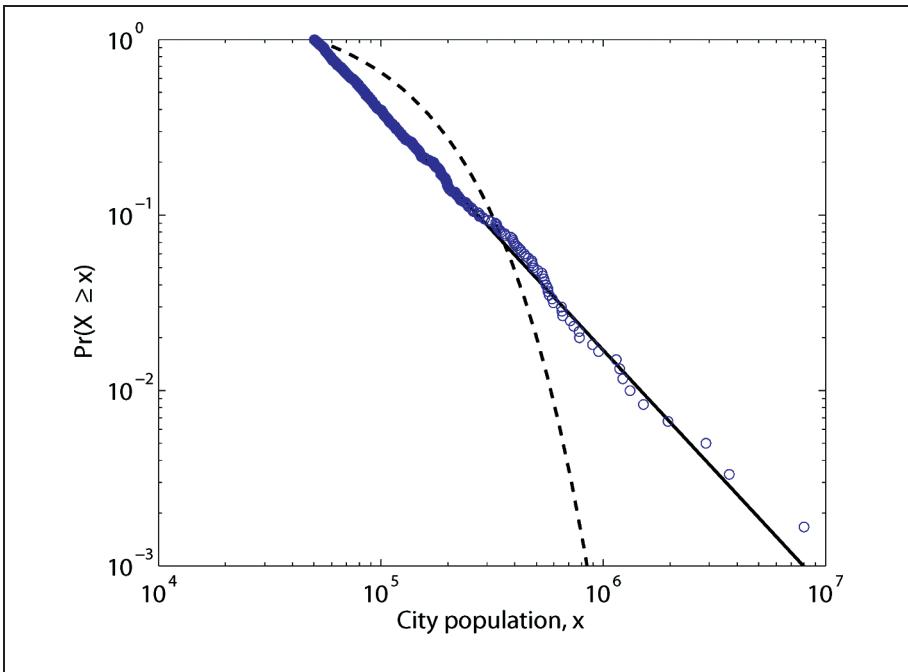


Figure 2. The sizes of the 600 largest cities in the United States, that is, those with population $x \geq 50,000$, based on data from the 2000 census

NOTE: The data are plotted as a complementary cumulative distribution function $\Pr(X \geq x)$. The solid black line shows the power-law behavior that the distribution closely follows, with scaling exponent $\alpha = 2.36 \pm 0.06$, while the dashed black line shows a truncated normal distribution with the same sample mean.

Although much more can be said about power laws, we hope that the curious reader will take away a few basic facts from this brief introduction. First, heavy-tailed distributions do not conform to our expectations of a linear, or normally distributed, world. As such, the average value of a power law is not representative of the entire distribution, and events orders of magnitude larger than the mean are, in fact, relatively common. Second, the scaling property of power laws implies that at least statistically, there is no qualitative difference between small, medium, and extremely large events, as they are all succinctly described by a very simple statistical relationship. Readers who would like more information about power laws should refer to the extensive reviews by Kleiber and Kotz (2003), Newman (2005), and Mitzenmacher (2004).

B. Generating Functions

Generating functions are a mathematical tool for representing and doing calculations with infinite sequences. Suppose you have two infinite sequences: (c_0, c_1, c_2, \dots) and (d_0, d_1, d_2, \dots) . Their generating functions are defined by

$$F(z) \equiv \sum_{k=0}^{\infty} c_k z^k, \quad (28)$$

$$G(z) \equiv \sum_{k=0}^{\infty} d_k z^k. \quad (29)$$

Both are analytic functions of the complex variable z . Their product $H(z) = F(z)G(z)$ is a power series

$$H(z) = \sum_{k=0}^{\infty} h_k z^k \quad (30)$$

with coefficients that are sums of products of the c_k and d_k :

$$h_k = \sum_{m,n=0}^{\infty} {}' c_m d_n, \quad (31)$$

where again \sum' denotes a summation over all natural numbers m and n such that $m + n = k$. This property was used in section 2.

It is often easier to calculate a generating function than to work explicitly with the sequence of the expansion coefficients. Once the function is known explicitly, the coefficients can be calculated from Cauchy's theorem

$$h_k = \frac{1}{2\pi i} \oint_C H(z) \frac{dz}{z^{k+1}}, \quad (32)$$

where C encircles the origin of the complex z -plane once in the counterclockwise direction.

Readers who would like more information about generating functions and their use in mathematical analysis should refer to the textbook by Wilf (2006).

C. The Cutoff k_0 and the Value of n_1^*

The full equation for $n_1(t)$ follows from the model assumptions in section 1. It has the form

$$\frac{dn_1}{dt} = \sum_{k=2}^{k_0} k b(k) n_k - A_0 n_1 \sum_{\ell=1}^{\infty} \ell^a n_{\ell}, \quad (33)$$

which gives for the stationary state the equation

$$\sum_{k=2}^{k_0} k b(k) n_k^* = A_0 n_1^* \sum_{\ell=1}^{\infty} \ell^a n_{\ell}^*. \quad (34)$$

This equation connects the cutoff k_0 with n_1^* . The right-hand side equals $A_0 n_1^* f(1)$, where equation 4 was used. Using equation 13, one can rewrite this as

$$\sum_{k=2}^{k_0} k b(k) n_k^* = 2A_0 (n_1^*)^2. \quad (35)$$

The value of n_1^* can then be calculated from the relation

$$N = n_1^* + \sum_{k=2}^{\infty} k n_k^*, \quad (36)$$

which expresses the fact that the total number of radicalized individuals should equal N . For the case $a = 1$, the definition in equation 4 shows that one can rewrite equation 36 in the form

$$\sum_{k=1}^{\infty} k n_k^* = N = f(1). \quad (37)$$

Combining this expression with equation 13 gives $N = 2n_1^*$, so one finds

$$n_1^* = \frac{1}{2} N, \quad (38)$$

that is, half the number of these individuals are singletons and half that number are part of larger cells.

To now calculate the cutoff k_0 (for $k > k_0$, we assume $n_k^* = 0$), one rewrites equation 35 in the form

$$\sum_{k=2}^{k_0} k b(k) n_k^* = \frac{1}{2} A_0 N^2. \quad (39)$$

As an example for explicit calculation, we take the case $a = 1$ and

$$b(k) = B_0 k^b, \quad (40)$$

for $\frac{1}{2} < b < \frac{3}{2}$, where the exponent b is some number in the vicinity of unity. Equation 24 now gives equation 39 the form

$$\sum_{k=2}^{k_0} k b(k) n_k^* \cong \frac{B_0 N}{2\sqrt{\pi}} \sum_{k=2}^{k_0} k^{b-3/2}, \quad (41)$$

where a small error is neglected, which is due to the fact that we used the $k \gg 1$ asymptotic expression for n_k^* for all $k \geq 2$. The series on the right of equation 41 can be approximated by an integral, which gives

$$\sum_{k=2}^{k_0} k^{b-3/2} \cong \int_2^{k_0} k^{b-3/2} dk \cong \left(\frac{1}{b - \frac{1}{2}} \right) k_0^{b-1/2}, \quad (42)$$

for $k_0 \gg 1$. With these results, equation 39 takes the form

$$\frac{B_0}{2\sqrt{\pi}} \left(\frac{1}{b - \frac{1}{2}} \right) k_0^{b-1/2} = \frac{1}{2} A_0 N, \quad (43)$$

which gives an explicit value for the cutoff:

$$k_0 = \left[\frac{A_0}{B_0} \left(b - \frac{1}{2} \right) \sqrt{\pi N} \right]^{1/(b-\frac{1}{2})}. \quad (44)$$

The essential feature of this result is that $k_0 \gg 1$ when $N \gg 1$. At the cutoff, the value of $n_{k_0}^*$ is proportional to a negative power of N :

$$n_{k_0}^* \propto N^{1 - \frac{5}{2}(b-\frac{1}{2})^{-1}}, \quad (45)$$

where one uses equations 24, 38, and 44. Hence for $k > k_0$, all numbers $n_k^* \ll 1$ and are therefore irrelevant. These features of the cutoff show that its existence is a mathematical artifact only, with no consequences for the distribution of cell sizes for realistic values of k .

Authors' Note

The authors thank Kristian S. Gleditsch, Christopher K. Butler, Libby Wood, Lars-Erik Cederman, Sidney Redner, and Neil F. Johnson for helpful conversations and two anonymous referees for comments on an earlier draft of this article.

Acknowledgments

This work was supported in part by the Santa Fe Institute.

Declaration of Conflicting Interests

The authors declared no conflicts of interest with respect to the authorship and/or publication of this article.

Funding

The authors received no financial support for the research and/or authorship of this article.

Notes

1. We note that a wide variety of mechanisms can produce power-law distributions. Most of these processes, however, are not well suited for explaining the severity of terrorist attacks

(see Clauset, Young, and Gleditsch 2007 for some discussion). As such, we focus our attention on the two mechanisms that have been proposed, both of which have some empirical support.

2. Here, *universality* denotes the robustness of certain qualitative features of a mathematical model to certain specific modeling assumptions. This usage is distinct from, and should not be confused with, the less technical usage of the same term to denote a natural or social phenomenon that appears to be independent of certain contingent or contextual details.

References

- Araj, B. 2008. Harsh state repression as a cause of suicide bombing: The case of the Palestinian-Israeli conflict. *Studies in Conflict & Terrorism* 31 (4): 284-303.
- Asal, V., and R. K. Rethemeyer. 2008. The nature of the beast: Organizational structures and the lethality of terrorist attacks. *Journal of Politics* 70 (2): 437-49.
- Bak, P., C. Tang, and K. Wiesenfeld. 1987. Self-organized criticality: An explanation of 1/f noise. *Physical Review Letters* 59 (4): 381-84.
- Biggs, M. 2005. Strikes as forest fires: Chicago and Paris in the late 19th century. *American Journal of Sociology* 111 (1): 1684-1714.
- Cederman, L.-E. 2003. Modeling the size of wars: From billiard balls to sandpiles. *American Political Science Review* 97 (1): 135-50.
- Cioffi-Revilla, C., and P. P. Romero. 2009. Modeling uncertainty in adversary behavior: Attacks in Diyala province, Iraq, 2002-2006. *Studies in Conflict & Terrorism* 32 (3): 253-76.
- Clauset, A., and K. S. Gleditsch. 2009. Developmental dynamics of terrorist organizations. Preprint. <http://arxiv.org/abs/0906.3287> (accessed June 17, 2009).
- Clauset, A., L. Heger, M. Young, and K. S. Gleditsch. Forthcoming. The strategic calculus of terrorism: Substitution and competition in the Israel-Palestine conflict. *Cooperation & Conflict*.
- 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., M. Young, and K. S. Gleditsch. 2007. On the frequency of severe terrorist events. *Journal of Conflict Resolution* 51 (1): 58-87.
- Cordes, B., B. M. Jenkins, K. Kellen, G. Bass, D. Relles, W. Sater, M. Juncosa, W. Fowler, and G. Petty. 1985. *A conceptual framework for analyzing terrorist groups*. Arlington, VA: RAND.
- Cronin, A. K. 2003. Behind the curve. *International Security* 27 (3): 30-58.
- Crook, C. 2006. The height of inequality. *Atlantic Monthly* 298 (2): 36-37.
- Enders, W., and T. Sandler. 2000. Is transnational terrorism becoming more threatening? A time-series investigation. *Journal of Conflict Resolution* 44 (3): 307-32.
- Enders, W., and T. Sandler. 2004. What do we know about the substitution effect in transnational terrorism? In *Researching terrorism trends, achievements, failures*, ed. A. Silke. London: Frank Cass.
- Fouda, Y., and N. Fielding. 2003. Masterminds of terror: *The truth behind the most devastating terrorist attack the world has ever seen*. New York: Arcade.
- Harrison, M. 2006. Bombers and by standers in suicide attacks in Israel. *Studies in Conflict & Terrorism* 29:187-206.
- 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. Preprint. <http://arxiv.org/physics/0506213> (accessed June 29, 2005).

- Johnson, N. F., M. Spagat, J. A. Restrepo, O. Becerra, J. C. Bohorquez, N. Suarez, E. M. Restrepo, and R. Zarama. 2006. Universal patterns underlying on going wars and terrorism. Preprint. <http://arxiv.org/physics/0605035> (accessed May 3, 2006).
- Kleiber, C., and S. Kotz. 2003. *Statistical size distributions in economics and actuarial sciences*. Hoboken, NJ: John Wiley.
- Kydd, A., and B. Walter. 2002. Sabotaging the peace: The politics of extremist violence. *International Organization* 56 (2): 263-96.
- Lacina, B. 2006. Explaining the severity of civil wars. *Journal of Conflict Resolution* 50 (2): 276-89.
- Li, Q. 2005. Does democracy promote or reduce transnational terrorist incidents? *Journal of Conflict Resolution* 49 (2): 278-97.
- Malamud, B. D., G. Morein, and D. L. Turcotte. 1998. Forest fires: An example of self-organized critical behavior. *Science* 281:1840-41.
- McKeown, L. 2001. *Out of time: Irish Republican prisoners Long Kesh 1972-2000*. Belfast, UK: Beyond the Pale.
- Mitzenmacher, M. 2004. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics* 1 (2): 226-51.
- National Memorial Institute for the Prevention of Terrorism. 2008. Terrorism knowledge base. <http://www.tkb.org> (accessed January 29, 2008).
- Newman, M. E. J. 2005. Powerlaws, Pareto distributions and Zipf's law. *Contemporary Physics* 46 (5): 323-51.
- Pape, R. A. 2003. The strategic logic of suicide terrorism. *American Political Science Review* 97 (3): 343-61.
- Pareto, V. 1896. *Cours d'Economie Politique*. Geneva, Switzerland: Droz.
- Pedahzur, A., and A. Perlinger. 2006. The changing nature of suicide attacks: A social network perspective. *Social Forces* 84 (4): 1987-2008.
- Richardson, L. F. 1948. Variation of the frequency of fatal quarrels with magnitude. *Journal of the American Statistical Association* 43:523-46.
- Richardson, L. F. 1960. *Statistics of deadly quarrels*. Pittsburgh, PA: Boxwood Press.
- Ruszczycki, B., B. Burnett, Z. Zhao, and N. F. Johnson. 2008. Relating the microscopic rules in coalescence-fragmentation models to the emergent cluster-size distribution. Preprint. <http://arxiv.org/abs/0808.0032> (accessed June 8, 2009).
- Sageman, M. 2004. *Understanding terror networks*. Philadelphia: University of Pennsylvania Press.
- Simon, H. A. 1955. On a class of skew distribution functions. *Biometrika* 42 (34): 425-40.
- Small, M., and J. D. Singer. 1982. *Resort to arms: International and civil wars, 1816-1980*. Beverly Hills, CA: Sage.
- U.S. Department of State. 2004. *Patterns of global terrorism, 2003*. Washington, DC: U.S. Department of State. <http://www.state.gov/documents/organization/31912.pdf> (accessed January 22, 2005).
- Wilf, H. 2006. *Generating functionology*. 3d ed. Wellesley, MA: A. K. Peters.
- Zipf, G. K. 1949. *Human behavior and the principle of least effort*. Cambridge, MA: Addison-Wesley.