The Developmental Dynamics of Terrorist Organizations

Aaron Clauset
Kristian Skrede Gleditsch

SFI WORKING PAPER: 2009-07-021

SANTA FE INSTITUTE

The developmental dynamics of terrorist organizations

Aaron Clauset[1] and Kristian Skrede Gleditsch[2,3]

[1]*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM, 87501, USA*
[2]*Department of Government, University of Essex, Wivenhoe Park, Colchester CO4 3SQ UK*
[3]*Centre for the Study of Civil War, Oslo, Norway*

Traditional studies of terrorist group behavior [1] focus on questions of political motivation, strategic choices, organizational structure, and material support [2–7], but say little about the basic laws that govern how the frequency and severity (number of deaths) [8] of their attacks change over time. Here we study 3,143 fatal attacks carried out worldwide from 1968–2008 by 381 terrorist groups [9], and show that the frequency of a group's attacks accelerates along a universal trajectory, in which the time between attacks decreases according to a power law in the group's total experience; in contrast, attack severity is independent of organizational experience and organizational size. We show that the acceleration can be explained by organizational growth, and suggest that terrorist organizations may be best understood as firms whose primary product is political violence. These results are independent of many commonly studied social and political factors, suggesting a fundamental law for the dynamics of terrorism and a new approach to understanding political conflicts.

High-quality empirical data on terrorist groups, such as their recruitment, fundraising, decision making, and organizational structure, are scarce, and the available sources are not typically amenable to scientific analysis [10]. However, good-quality data on the frequency and severity of their attacks do exist, and their systematic analysis can shed new light on which facets of terrorist group behavior are predictable and which are inherently contingent. Each record in our worldwide database of 3,143 fatal attacks [9] includes its calendar date t, its severity x, and the name of the associated organization, if known (see Supplementary Information).

For each group, we quantify the changes in the frequency and severity of a group's attacks over its lifetime using a *development curve*. This curve maps a group's behavior onto a common quantitative scale and facilitates the direct comparison of different groups at similar points in their life histories. To construct this, we plot the behavioral variable, such as the time (days) between consecutive attacks Δt or the severity of an attack x, as a function of the group's maturity or experience k, indexed here by the cumulative number of fatal attacks (Fig. 1). Combining the developmental curves of many groups produces an aggregate picture of their behavioral dynamics, and allows us to extract the typical developmental trajectory of a terrorist group.

Constructing a combined development curve for the 381 organizations in our database, we find that the time between consecutive attacks Δt changes in a highly regular way (Fig. 2a), while the severity of these attacks x is independent of organizational experience (Fig. 2c).

Empirically, the time between attacks decreases quickly as a group gains experience. For example, the mean delay between the first and second fatal attacks is almost six months, $\langle \Delta t \rangle = 168.6 \pm 0.6$ days, while after 13 attacks, the mean delay is only $\langle \Delta t \rangle = 27 \pm 1$ days. More generally, the envelope or distribution of delays $p(\Delta t, k)$ can be characterized as a truncated log-normal distribution with constant variance σ^2 and a characteristic delay between attacks μ that decreases systematically with experience k. Mathematically:

$$p(\log \Delta t, \log k) \propto \exp\left[\frac{-(\log \Delta t + \beta \log k - \mu)^2}{2\sigma^2}\right] \quad , \quad (1)$$

where β controls the trajectory of the distribution toward the natural cutoff at $\Delta t = 1$ day. For small k, i.e., during a group's early development, this model predicts a mean delay between attacks that decays like a power law $\Delta t \approx \mu k^{-\beta}$; however, as k increases, this trend is attenuated as the mean delay asymptotes to $\Delta t = 1$ (see Supplementary Information). Under this model, $\beta = 1$ would indicate a simple linear feedback between a group's attack rate and its experience. However, we find $\hat{\beta} = 1.10 \pm 0.02$, indicating a faster-than-linear feedback between the accumulation of experience and the rate of future attacks.

This model successfully predicts that the distributions of normalized delays $\Delta t\, k^{\hat{\beta}}$ will collapse onto a single log-normal distribution with parameters $\hat{\mu}$ and $\hat{\sigma}$ (Fig. 2b). However, individual attacks cannot be considered fully independent ($p = 0.00 \pm 0.03$; see Supplementary Information), indicating that significant temporal or inter-group correlations may exist in the timing of a group's future attacks [11] (Fig. 1).

In contrast, the severity of an attack x is independent of group experience k (Fig. 2c; $r = -0.024$, t-test, $p = 0.17$), as illustrated by the collapse of the severity distributions $p(x \,|\, k)$ onto a single invariant heavy-tailed distribution [8, 12] (Fig. 2d; see Supplementary Information). For example, the mean severity of a group's first attack $\langle x \rangle = 6.7 \pm 0.9$ is only slightly larger than the mean severity of all attacks by very experienced groups ($k > 100$) $\langle x \rangle = 5.1 \pm 0.6$. Thus, contrary to common assumptions, young and old groups are equally likely to produce extremely severe events. Older groups, however, remain significantly more lethal overall [13] because they attack much more frequently than small groups, not because their individual attacks are more deadly.

Here, we consider four explanations for the acceleration in the frequency of attacks: (i) organizational learning [10], (ii) organizational growth, (iii) sampling artifacts, and (iv) averaging artifacts [14]. Organizational learning—commonly studied in manufacturing [15, 16] and called "learning by doing"—implies that terrorist groups are born clumsy and increase their attack rate primarily because their existing

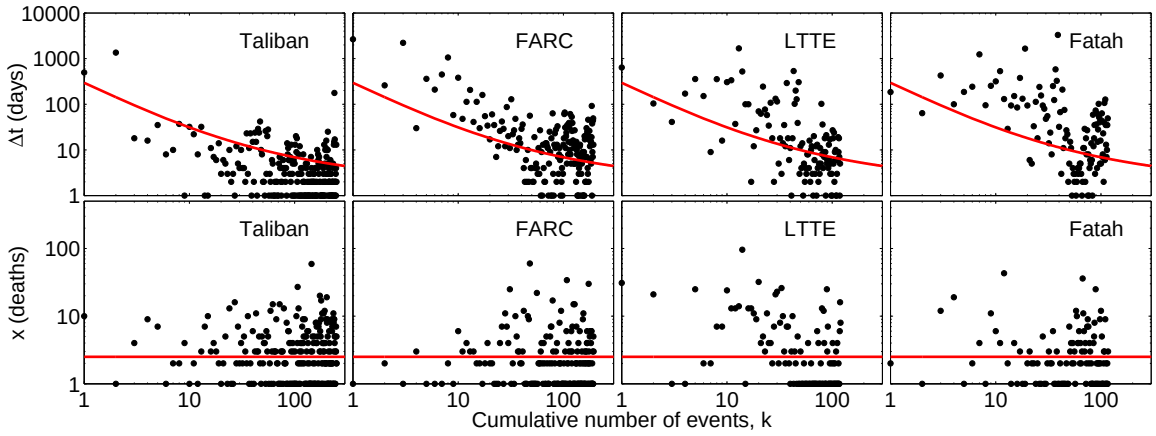

FIG. 1: **Individual development curves.** Time until the next fatal event Δt, and the severity of the event x, as a function of individual group experience k, for the four most experienced groups (Taliban, Revolutionary Armed Forces of Colombia (FARC), Liberation Tigers of Tamil Eelam (LTTE), and al Fatah), along with population-level curves derived from analyzing all groups together. Similar results hold for less experienced groups.

members learn to be more efficient, e.g., better planning, coördination, and execution. In contrast, organizational growth implies that groups are born small and increase their attack rate primarily by recruiting new, replaceable, relatively independent members, e.g., adding new terrorist cells. Straightforward tests of the data can eliminate both artifactual explanations (see Supplementary Information), indicating that the acceleration is real, even at the level of an individual group (Fig. 1).

The relative importance of organizational learning and organizational growth cannot be estimated using frequency and severity data alone. Untangling their effects requires data both on event planning and execution, and on a group's size and recruitment at various points in its life history. To our knowledge, systematic data on event planning and execution do not exist. The best available data on group sizes, taken from an expert survey [13], are coarse (roughly order of magnitude) estimates of the maximum size achieved by each of the 381 groups over the 1998–2005 period; of these 161 conducted at least one fatal attack, and 80 conducted at least two.

The growth hypothesis predicts that a group's maximum size will be inversely related to the minimum delay between its attacks over the 1998–2005 period. An analysis of variance indicates that the average minimum delays in the four size categories are significantly different (Fig. 3a; n-way ANOVA, $F = 9.98$, $p < 0.000013$), and further that larger organizational size is a highly significant predictor of increased attack frequency ($r = -0.49$, t-test, $p < 10^{-5}$). In contrast, size, like experience (Fig. 2c), is not a significant predictor of attack severity (Fig. 3b; see Supplementary Information).

Although operational, organizational, and political circumstances vary widely across terrorist groups, the systematic nature of our results suggests several general conclusions. The strong dependence of attack frequency on experience (Fig. 2a) suggests that the timing of events is governed primarily by organizationally internal factors, like growth and learning, related to group development, e.g., recruitment, personnel turnover, and internal coördination. Our analysis of group sizes lends significant support to the growth hypothesis, but without additional data, we cannot eliminate the possibility that these groups are also learning. Even so, our results implicate these internal factors as leverage points for decreasing the incidence rate of future attacks. In contrast, internal factors seem to play a marginal role in the severity of any particular attack (Figs. 2c and 3b), implying that the lethality of larger and more mature groups [13] is explained by their more frequent activity rather than more deadly activity. Further, we note that curtailing the frequency of a group's attacks, perhaps by limiting growth, would reduce the *cumulative risk* of very severe attacks.

Learning and growth may also constrain other behaviors, such as fundraising, resource availability, political motivation [6, 17], strategic choices in timing [18] and targets [4], and the use of tactics like suicide bombs [5]. However, most groups never achieve a high level of experience ($k > 100$), and it is unclear to what degree this leaky pipeline is caused by group death, e.g., from counter-terrorism activities or internal conflicts [19, 20], versus shifts away from violence. Similarly, it is unclear whether or how these dynamics change when a group does become highly experienced. Once a group is large enough to execute almost daily attacks, it may be more like a social institution, and thus face different constraints and incentives than smaller, still developing groups. It is also unclear what mechanism explains the power-law distribution in the severity of terrorist attacks [8, 12], and its independence of organizational experience and size. Two possibilities are (i) the advantages and disadvantages of youth (small size and the element of surprise vs. poor resources and clumsy attacks) balance those of maturity (more resources and better planning vs. risk aversion and hostile attention from governments), yielding apparent independence; and (ii) severity is inherently random, governed by contingent details associated with the

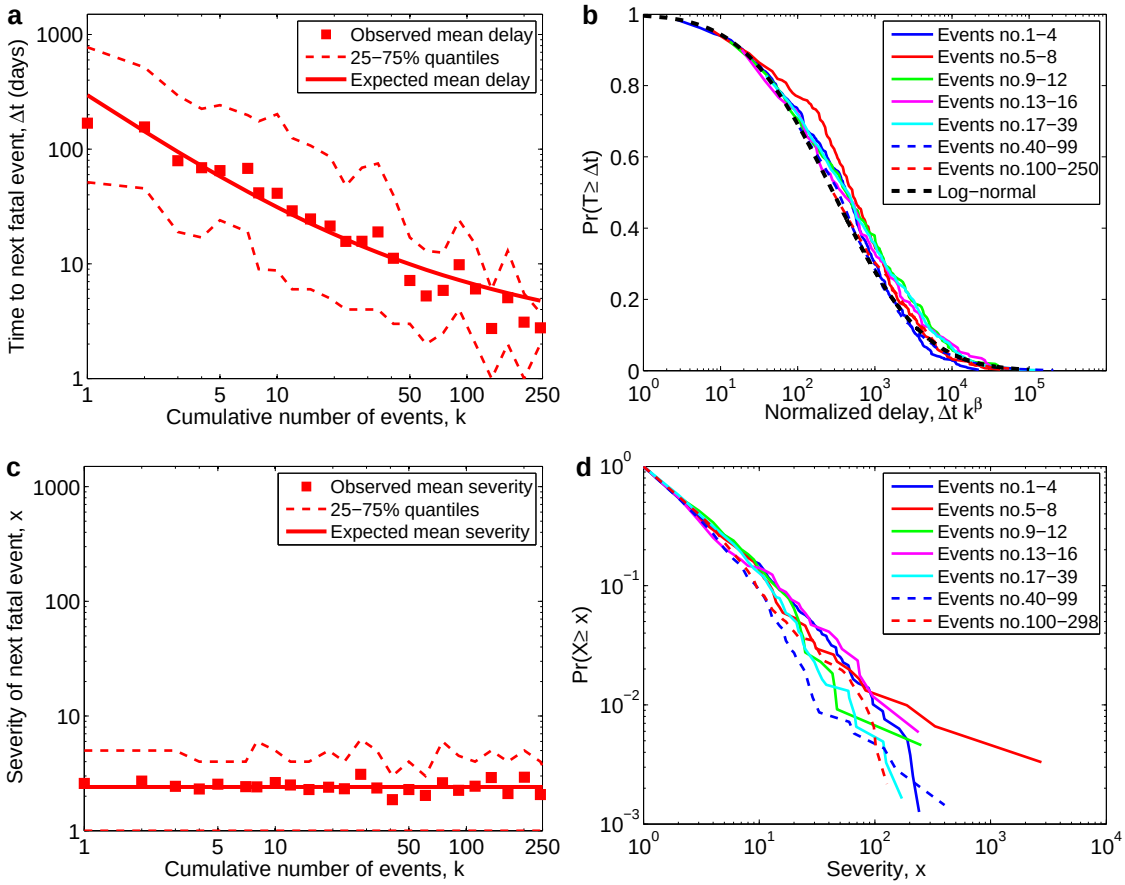

FIG. 2: **Combined development curves.** **a**, The mean delay $\langle \log \Delta t \rangle$ between attacks by a terrorist group, with 25^{th} and 75^{th} percentile isoclines, as a function of group experience k. The solid line shows the expected mean delay, from the model described in the text. **b**, The distributions of normalized delays $p(\Delta t\, k^{\beta})$, showing the predicted data collapse onto an underlying log-normal distribution. **c**, The mean severity $\langle \log x \rangle$ of fatal attacks by a terrorist group, with 25^{th} and 75^{th} percentile isoclines, as a function of group experience k. The solid line (with slope zero) shows the expected mean delay, from a simple regression model. **d**, The distributions of event severities $\Pr(X \geq x)$, showing the data collapse onto an underlying heavy-tailed distribution.

particular attack, the particular group, etc.

The development curves shown here are similar in form to production curves found in manufacturing [15], in which per-item production costs tends to decrease like a power law in the cumulative number of items produced. In this light, terrorist groups may best be understood as firms whose principal product is political violence, and whose production rates depend largely on organizational growth and the availability of low-skill labor. Thus, studies of event-driven, non-terrorist organizations, e.g., some non-profits, political activism groups, and commercial firms, could shed additional light on the dynamics of terrorist groups. A better understanding of how these groups are like, and unlike, non-terrorist human social groups may indicate which counter-terrorism policies, e.g., limiting growth by stifling recruitment or by forcing organizational turnover, are likely to be effective.

Finally, we note that our results are independent of many commonly studied factors, including geographic location, time period, ideological motivations (religious, separatist, reactionary, etc.), and political context. Some aspects of terrorism are thus not nearly as contingent as is widely assumed, and quantifying dynamical patterns in political conflict can serve the broader goal of understanding what regularities exist, why they exist, and what they imply for long-term social and political stability.

Acknowledgments The authors thanks L.-E. Cederman, K. Drakos, B. Karrer, J. McNerney, J. H. Miller, M. E. J. Newman, A. Ruggeri, D. Sornette, C. R. Shalizi, and M. Young for helpful conversations. This work was supported in part by the Santa Fe Institute, the Economic and Social Research Council (RES-062-23-0259), and the Research Council of Norway (180441/V10).

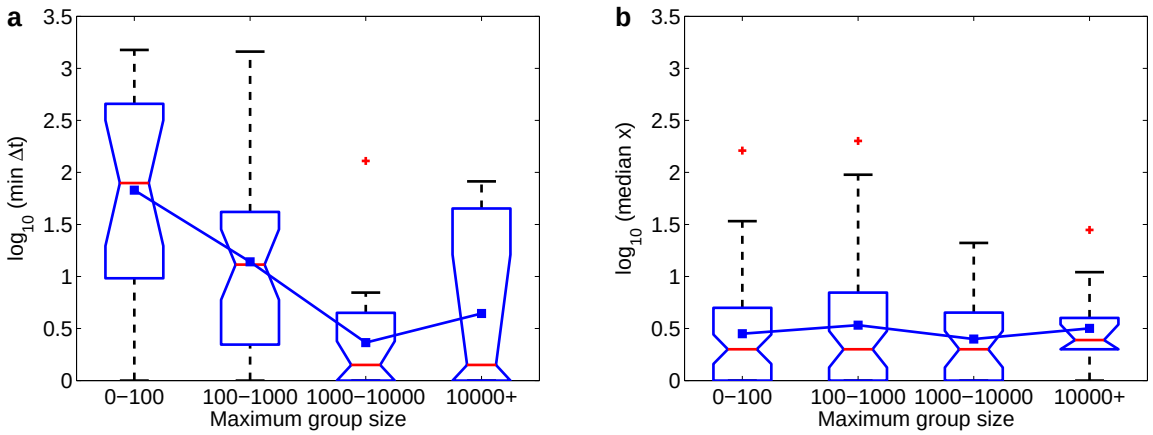

FIG. 3: **Frequency and severity versus group size** s. Box-plots of the distributions of a group's **a**, minimum delay $\log(\min \Delta t)$ and **b**, median attack severity $\log(\text{median } x)$ for attacks over the period 1998–2005, within each of four estimated maximum size categories [13]. For convenience, we connect the means of each category, which are significantly different in the case of delays (n-way ANOVA, $F = 9.98$, $p < 0.000013$), but indistinguishable in the case of severities (n-way ANOVA, $F = 0.59$, $p = 0.62$).

[1] We use a relatively standard definition of *terrorism*: a violent act by a non-state actor intended to create fear for political purposes. Other definitions exist, but tend to be similar; variations here do not impact our results significantly.

[2] Cordes, B. *et al. A Conceptual Framework for Analyzing Terrorist Groups* (RAND Corporation, Arlington, 1985).

[3] Hoffman, B. *Inside Terrorism* (Columbia University Press, New York, 1998).

[4] Enders, W. & Sandler, T. What do we know about the substitution effect in transnational terrorism? In Silke, A. (ed.) *Researching Terrorism Trends, Achievments, Failures* (Frank Cass, London, 2004).

[5] Pape, R. A. The strategic logic of suicide terrorism. *American Political Science Review* **97**, 343–361 (2003).

[6] Sageman, M. *Understanding Terror Networks* (University of Pennsylvania Press, Philadelpha, 2004).

[7] Li, Q. Does democracy promote or reduce transnational terrorist incidents? *Journal of Conflict Resolution* **49**, 278–297 (2005).

[8] Clauset, A., Young, M. & Gleditsch, K. S. On the frequency of severe terrorist events. *Journal of Conflict Resolution* **51**, 58–87 (2007).

[9] National Memorial Institute for the Prevention of Terrorism. Terrorism Knowledge Base (2008). `http://www.tkb.org` (access date January 29, 2008).

[10] Jackson, B. A. *et al. Aptitude for Destruction: Organizational Learning in Terrorist Groups and Its Implications for Combating Terrorism*, vol. 1 (RAND Corporation, Arlington, 2005).

[11] Clauset, A., Heger, L., Young, M. & Gleditsch, K. S. The strategic calculus of terrorism: Substitution and competition in the Israel-Palestine conflict (2009). *Cooperation & Conflict*, to appear.

[12] Clauset, A., Shalizi, C. R. & Newman, M. E. J. Power-law distributions in empirical data (2009). *SIAM Review*, to appear.

[13] Asal, V. & Rethemeyer, R. K. The nature of the beast: Organizational structures and the lethality of terrorist attacks. *Journal of Politics* **70**, 437–449 (2008).

[14] Gallistel, C. R., Fairhurst, S. & Balsam, P. The learning curve: Implications of a quantitative analysis. *Proc. Natl. Acad. Sci. USA* **101**, 13124–13131 (2004).

[15] Dutton, J. M. & Thomas, A. Treating progress functions as a managerial opportunity. *Academy of Management Review* **9**, 235–247 (1984).

[16] Argote, L. Group and organizational learning curves: Individual, system and environmental components. *British Journal of Social Psychology* **32**, 31–51 (1993).

[17] Krueger, A. B. *What Makes a Terrorist: Economics and the Roots of Terrorism* (Princeton University Press, Princeton NJ, 2007).

[18] Kydd, A. & Walter, B. Sabotaging the peace: The politics of extremist violence. *International Organization* **56**, 263–296 (2002).

[19] Cronin, A. K. How al-Qaeda ends. *International Security* **31**, 7–48 (2006).

[20] Jones, S. G. & Libicki, M. C. *How Terrorist Groups End: Lessons for Countering al Qa'ida* (RAND Corporation, Arlington, 2008).

Supplementary Information

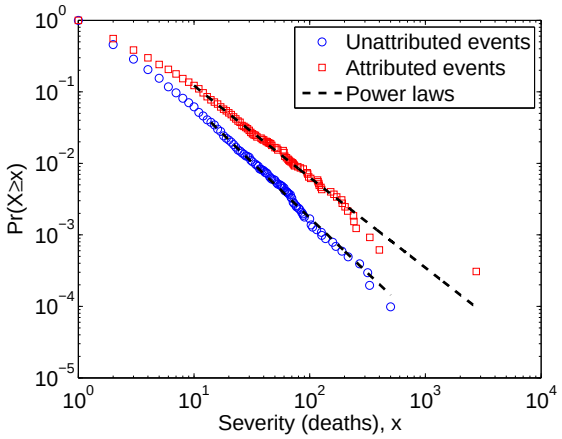

FIG. S1: The severity distributions for attributed and unattributed events, along with maximum likelihood power-law models for their tails. See Table S1 for details.

Appendix A: Terrorism Data

1. Events

Many organizations track terrorist events worldwide, but few provide their data in a form amenable to scientic analysis. The most popular source of information on terrorist events in the political science literature is the ITERATE data set [S1], which focuses exclusively on transnational terrorist events, i.e., events involving actors from at least two countries. For the analysis of terrorist groups, however, domestic events, i.e., events involving actors from only one country, are equally relevant. Thus, we use the data contained in the National Memorial Institute for the Prevention of Terrorism [S2] database, which largely overlaps with the ITERATE data but also includes fully domestic terrorist events since at least 1998.

The MIPT database combines the RAND Terrorism Chronology 1968-1997, the RAND-MIPT Terrorism Incident database (1998-2008), the Terrorism Indictment database (University of Arkansas and University of Oklahoma), and DFI International's research on terrorist organizations. Since we last collected event data [S2], MIPT mandate was changed by the U.S. Department of Homeland Security such that maintaining and making publicly available this data was no longer a priority. Our understanding is that the MIPT data will eventually be merged with the Global Terrorism Database (GTD) managed by the START program, a U.S. Department of Homeland Security Center of Excellence, at the University of Maryland, College Park.

Records in the MIPT database are derived by a team of specialized scholars from original sources, largely reports in newspapers worldwide, accessible via the LexusNexis service and other means. This human element presents the possibility that some records, particularly those corresponding to nonlethal events, could be missing or corrupted by poor reporting [S3]. In order to mitigate the effect of such biases, we focus our analysis on the portion of the database that is most trustworthy [S4], namely those events that resulted in at least one fatality, since these events are typically reported more reliably in the media [S5].

The MIPT database contains 35,688 recorded terrorist attacks worldwide from January 1968 to January 2008; these events occurred in almost 6,000 cities in over 180 nations. 13,274 (37.2%) of the events resulted in at least one fatality while 10,085 (28.3%) were attributed to one of 910 organizations. The intersection of these criteria, i.e., events that were both fatal and attributed, includes 3,143 events attributed to 381 organizations. (106 such events were attributed to more than one organization; in our analysis, we give all groups credit for such an event.) This set of organizations includes most common ideological motivations, such as national-separatist, reactionary, religious, and revolutionary ideologies [S6], and spans a wide variety of political contexts.

In some cases, groups carried out multiple attacks on the same day. For our analysis of the delay between consecutive attacks, these few cases present a difficulty by inducing a delay variable of $\Delta t = 0$. For simplicity, we combine such concurrent attacks into a single record with a severity equal to the sum of the fatalities of its component attacks.

2. Organizations

There are very few sources of systematic data on terrorist organizations worldwide. The START program at the University of Maryland currently hosts some data under their "Terrorist Organization Profiles" (TOPs) program; this data was originally developed by Detica, Inc., a British defense contractor, and was accessible from the MIPT from c.1998 to March 2008. Estimates of group sizes in this database are relatively few, and the methodology used to derive them opaque.

Jones and Libicki, working at the RAND Corporation, com-

	n	$\langle x \rangle$	σ	$x_{\max}$	$\hat{x}_{\min}$	$\hat{\alpha}$	p
Attributed	3,143	6.6	52.1	2749	10 ± 3	2.24 ± 0.12	0.78 ± 0.03
Unattributed	10,131	3.5	10.5	500	14 ± 3	2.55 ± 0.12	0.19 ± 0.03
All events	13,274	4.2	51.3	2749	10 ± 3	2.39 ± 0.09	0.39 ± 0.03

TABLE S1: Summary statistics, along with parameters and significance values for power-law tail models, for fatal terrorist events attributed and unattributed to some terrorist organization.

piled a database of information on 649 terrorist groups. This database includes estimates of the group's peak size over its entire lifetime [S7]. Unfortunately, the breadth of time over which the size estimate holds makes it difficult to tightly relate size with behavior.

Finally, Asal and Rethemeyer extracted data from the MIPT list of organizations and augmented it with data drawn from public sources and estimates from a panel of experts at the Monterey Institute of International Studies (MIIS) [S8]. This data provides coarse (roughly order of magnitude) estimates of maximum size over the 1998–2005 period. Given the care taken in constructing the Asal and Rethemeyer data, it appears to be the most accurate data on terrorist group sizes currently available. Of the groups they consider, 161 conducted at least one fatal attack over the 1998–2005 period, and 80 conducted at least two. For our analyses of group size (Appendix E), we consider the former set for questions of event severity, and the latter set for questions of event frequency.

3. Attributed vs. unattributed events

Because a large fraction of the MIPT events are not attributed to any particular terrorist organization, an important question is whether the unattributed events differ from the attributed events in important ways. For example, are attributed events systematically more severe than unattributed events? This might be the case if there were a systematic bias or incentive for attacking organizations to be associated with their more severe events, perhaps to gain greater media attention. We can test for such a bias by separately considering the severity distributions for attributed and unattributed events in the MIPT database. Fig. S1 shows these distributions, along with their corresponding power-law fits; Table S1 summarizes these data and models.

The results show that the attributed- and unattributed-event severity distributions are extremely similar. Their averages and standard deviations differ primarily because these are high-variance distributions, and the presence of a few very large attributed events drives outsized differences in these summary statistics. There are slight differences in the bodies (small x, see Fig. S1), and the upper tails are almost identical [S9]: both are plausibly distributed according to power laws and span largely indistinguishable ranges of severity, but have slightly different scaling exponents (Table S1). Although we find some evidence for a very weak relationship between severity and attribution, our results imply that the severity of attributed and unattributed events do not differ in qualitatively important ways.

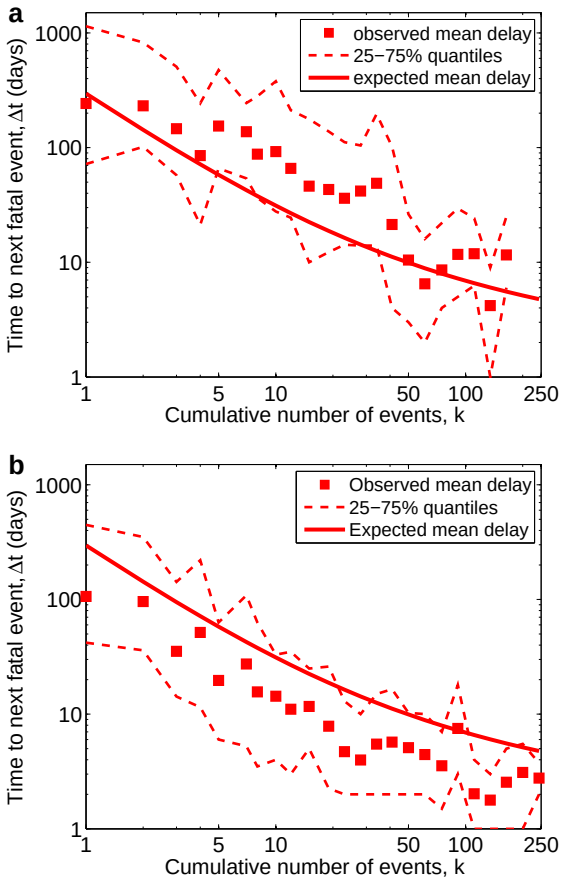

FIG. S2: The attack frequency development curves for groups whose first attack was **a**, in 1968–1997 or **b**, in 1998–2008, illustrating that the development curve progression for attack frequency is not merely a transnational or domestic phenomenon, and is robust to the change in the definition of the MIPT database in 1998. The expected mean delay shown here is taken from the main text, and thus serves as a reference point.

There is no clear theoretical reason to expect that the delay between events Δt should be systematically related to attribution. However, because the development curve analysis described in Appendix B omits all of the events carried out by a group but which are not attributed to it, the empirical delay between consecutive attacks is at best an upper-bound on the true delay. (Unfortunately, it is unknown how the observed delay between attributed events relates to the true delay between consecutive events.) This implies that our estimate of the ac-

celeration of attack rates (characterized by β) may overestimate its true value. The systematic patterns shown in Fig. S3 and Fig. S7 suggest that, unless they are highly pathologically distributed as a function of k, incorporating correctly labeled unattributed events is unlikely to fundamentally change our results.

4. Domestic vs. transnational events

The MIPT data has an important bias, due to its particular maintenance history. From 1968–1997, the database was maintained by the RAND Corporation as part of its project on transnational terrorism. As a result, almost no domestic terrorist attacks are included before 1998, after which the scope of the database was significantly expanded (in part due to the Oklahoma City bombing in 1995) to include purely domestic events worldwide. Thus, an important check on the generality of our results is to ask whether the attack frequency development curves depend on mixing data across the critical 1998 date. To test for this sensitivity, we construct combined development curves from events by groups whose first attack was before 1 January 1998 (mainly transnational groups), and separately from events by groups whose first attack was on or after 1 January 1998 (transnational and domestic groups). Groups in the MIPT database are not coded as being transnational or domestic, and adding such data is a highly non-trivial task for a database of this size.

Figure S2 shows that the development curve phenomenon is robust to this division of data. That being said, one difference is worth noting. The development curve for the 1968–1997 data shows a much slower overall rate of acceleration than does the 1998–2008 curve. The origin of this difference may be related to the rise of religiously motivated terrorism in the 1990s and beyond; however, sorting out its cause is beyond the scope of both this test and this study. For our purposes, it suffices to demonstrate that the development curve's overall form is robust to the MIPT's coverage bias.

Appendix B: Attack frequency development curve model

This model comes from the observation that the distributions of delays $p(\Delta t, k)$ for different values of k collapse onto a single universal curve (Fig. S3b) of the form

$$p(\log \Delta t, \log k) = C \exp\left[\frac{-(\log \Delta t + \beta \log k - \mu)^2}{2\sigma^2}\right] \tag{B1}$$

$$C = \frac{\sqrt{2/\pi}}{\sigma\left(1 - \mathrm{Erf}\left[\frac{\beta \log k - \mu}{\sigma\sqrt{2}}\right]\right)} \ ,$$

where Erf is the error function and C is the normalization constant. In words, the logarithm of the delay is normally distributed $\mathcal{N}(\mu, \sigma)$ (or equivalently, the delay is log-normally distributed), but has a natural lower cutoff at $\Delta t = 1$ day

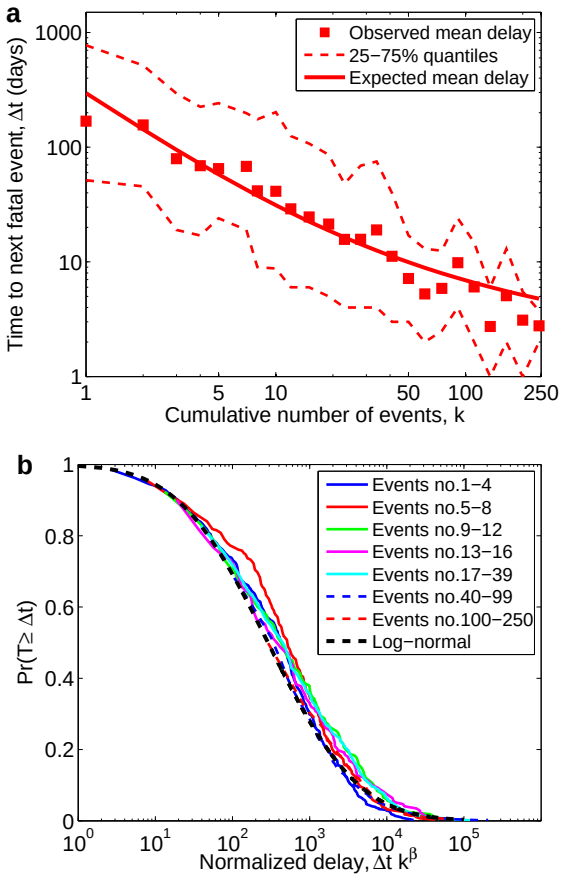

FIG. S3: **a**, The mean delay $\langle \log \Delta t \rangle$ between attacks by a terrorist group, with 25^{th} and 75^{th} percentile isoclines, as a function of group experience k. The solid line shows the expected mean delay, from the model described in the text. **b**, The distributions of normalized delays $p(\Delta t\, k^{\hat{\beta}})$, showing the predicted data collapse onto an underlying log-normal distribution.

(which comes from the resolution of the recorded data) and a μ parameter that decreases systematically as k increases. The parameter μ denotes the characteristic delay between attacks, and in particular the delay between the first and second attacks, while σ^2 denotes the variance in the expected delay. Due to the breadth of the log-normal distribution, a non-trivial value of σ implies a wide degree of variability in the waiting time, even given the characteristic delay μ. By integrating, we can derive the expected mean delay from Eq. (B1), which has the form

$$\mathrm{E}[\log \Delta t] = \mu - \beta \log k + \left(\frac{\exp\left[\frac{-(\beta \log k - \mu)^2}{2\sigma^2}\right]\sqrt{2/\pi}}{\sigma^{-1}\left(1 - \mathrm{Erf}\left[\frac{\beta \log k - \mu}{\sigma\sqrt{2}}\right]\right)}\right) . \tag{B2}$$

For small values of k, the expected delay is dominated by the first two terms, and thus decays according to a power-law $\Delta t \approx \mu k^{-\beta}$, where μ represents the initial rate of attack for a group. At larger values of k, the expected delay is dominated

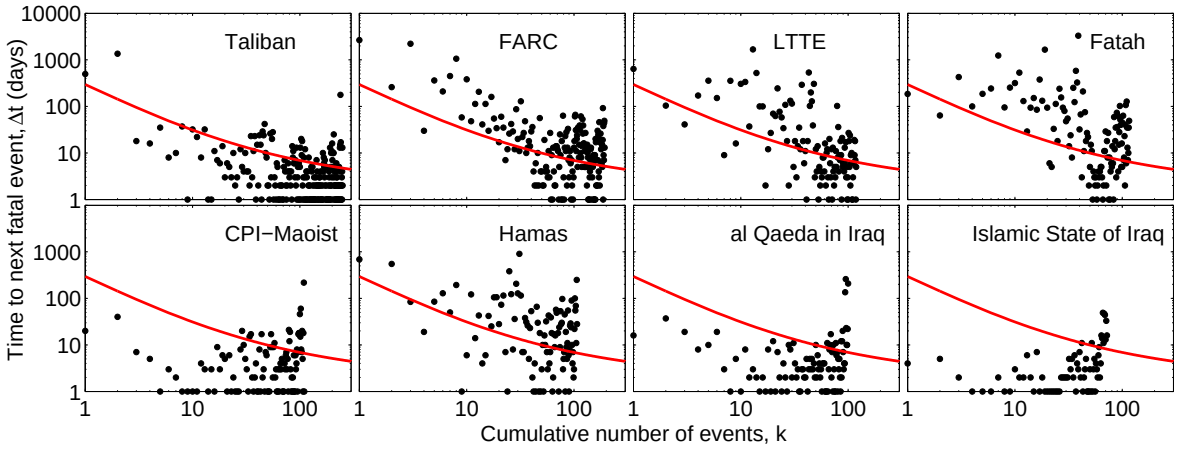

FIG. S4: Individual attack frequency development curves for the eight most experienced groups, along with the expected mean delay curve shown in Fig. S3 (i.e., the combined curve for all groups). Notably, there is no evidence of the threshold behavior predicted by the averaging hypothesis, or the universally low initial delay predicted by the sampling hypothesis; similar results hold for less experienced groups.

by the third term in Eq. (B2), which makes the expected delay to approach $\Delta t = 1$ more slowly than a power law.

The model's parameters μ, σ and β can be estimated directly from the data using standard numerical procedures (the simplex or Nelder-Mead method [S10], in this case) to maximize the likelihood of the data. Standard error estimates for the parameters can then be numerically estimated using the method of Fisher information to approximate the width of the log-likelihood function in the vicinity of the maximum [S11]. Applying these methods to our data yields $\hat{\mu} = 5.67 \pm 0.05$, $\hat{\sigma} = 2.12 \pm 0.02$, and $\hat{\beta} = 1.10 \pm 0.02$. The estimated value of β has particular significance: $\beta = 1$ indicates a simple linear feedback between the gain in experience and the change in attack rate, with deviations above and below indicating super- and sub-linear feedback. The estimated $\hat{\beta} > 1$, thus indicates super-linear feedback, which implies an explosive acceleration in the attack rate in finite time. However, this explosive growth is strongly attenuated by the third term in Eq. (B2), which prevents the actual attack rate from reaching extremely high levels.

Although desirable, estimates of individual group trajectories are difficult to obtain using this procedure, due to a severe $O(k_*^{-1/4})$ finite-size bias in the estimation of μ (where k_* denotes the total number of attacks by the group), and slightly less severe biases in estimating σ and β. On the other hand, combining data from many groups yields a less severe $O(n^{-1/2})$ bias in each parameter (where n is the total number of delays combined). For the combined development curve, we have $n \approx 2500$, and the finite-size bias is negligible (less than 0.1 for each parameter).

Finally, we point out that very few groups (e.g., Hamas, Fatah, LTTE, FARC, etc.) manage to become highly experienced ($k \gtrsim 100$). The sparsity of the data for large k means that the fit in this region is primarily controlled by the delays at much smaller values of k, where the vast majority of data lay (Fig. S3a). This explains the slight misfit to the delays for highly experienced groups, relative to the empirical observations, and highlights the fact that the behavior of inexperienced groups is predictive of the behavior of more mature groups.

Appendix C: Tests of Attack Delay Development Curve

1. Averaging and sampling artifacts

As described in the main text, there are several possible explanations for the observed universal development curve for the frequency of terrorist attacks by a group. These are (i) organizational learning [S12], (ii) organizational growth, (iii) sampling artifacts, and (iv) averaging artifacts [S13]. Here we show that the sampling and averaging explanations can be eliminated using the available empirical data.

In the sampling scenario, groups are born with some unique, but fixed, attack rate μ_i. As each group progresses through its lifetime, groups with slower attack rates (larger μ_i) die out first. This implies that the maximum experience k_* achieved by a group should be inversely proportional to its attack rate μ_i. The net result of sampling groups in this way is to leave progressively fewer groups with large μ_i values at larger experiences k, which gives the illusion of a smooth trend toward faster attack rates.

In the averaging scenario, groups can exhibit one of two attack frequencies: daily attacks (fast attackers) or attacks at some rate μ (slow attackers). Initially, all groups are slow attackers; however, at a group-specific experience threshold $k_{\circ,i}$, they switch behaviors and become fast attackers. The individual development curve for such a group would be a step function. By combining many such step functions, each with a different step location $k_{\circ,i}$, i.e., by averaging across the different thresholds, the combined development curve show a smooth trend toward faster attack rates even though no indi-

vidual group behaves that way. (As an aside: this averaging across thresholds model has been used to explain the apparently progressive learning curves for certain behaviors in animals [S13] and is thought to underly some evolved regulatory behaviors among eusocial insects [S14].)

It is instructive to again use the factory analogy. In the case of selection artifacts, different factories have different intrinsic production rates, but slower factories go out of business sooner, leaving a progressively larger fraction of fast factories at later times. Thus, although no factory exhibits a progressive trend toward faster production, as time passes, a larger fraction of the slow factories have died off, and the average rate of the remaining factories progressively increases. In the case of averaging artifacts, all factories begin with slow production rate, but at a randomly chosen time in the future, each switches from slow to fast production (due to a sudden insight or a sudden windfall profit). Thus, although no individual factory exhibits a progressive trend toward faster production, as time passes, a larger fraction of factories have switched to the faster rate, and the average rate of all factories progressively increases.

Both the sampling and averaging explanations can be tested using the available event frequency data. The sampling explanation predicts an inverse relationship between a group's characteristic attack rate μ_i and its maturity $k_{*,i}$ (total number of attacks), i.e., groups with larger μ_i (slower attack rate) die out more quickly and thus should exhibit a smaller number of total events. Unfortunately, the finite-size bias mentioned in Appendix B prevents us from testing this prediction directly, by first estimating each group's μ_i via maximum likelihood and then testing if μ_i varies inversely with $k_{*,i}$. Instead, we must use a more indirect test. Here, we use the first delay variable for a group $\Delta t_i^{(1)}$ as a proxy for the group's attack rate μ_i. Regressing each group's initial delay against its total experience yields no significant relationship ($r = -0.057$, t-test, $p = 0.46$), and similarly for the logarithm of delay and experience ($r = -0.114$, t-test, $p = 0.14$). (Notably, the empirical $k_{*,i}$ underestimates the total number of attacks by a group for all groups still active today; however, this particular kind of bias is unlikely to transform an inverse relationship between $\Delta t_i^{(1)}$ and $k_{*,i}$ into the null relationship we observe here.)

A strong test of the sampling explanation, which also tests the averaging hypothesis, is to plot the individual attack frequency development curves for the groups with the most fatal attacks. Fig. S4 shows these curves for the eight most experienced groups in our database. Notably, none of these curves exhibits the predicted switching behavior (from low to high frequency) predicted by the averaging hypothesis, and very few (Islamic State of Iraq, and possibly al Qaeda in Iraq) exhibit extremely short initial delays; instead, each generally exhibits a progressive increase in attack rate, and similar results hold for less experienced groups.

Thus, although the combined development curve may not represent precisely the development of any individual organization—and indeed, we see that individual groups do indeed deviate from the population trend (Fig. S4)—its general shape is an accurate description of the tendencies of each group.

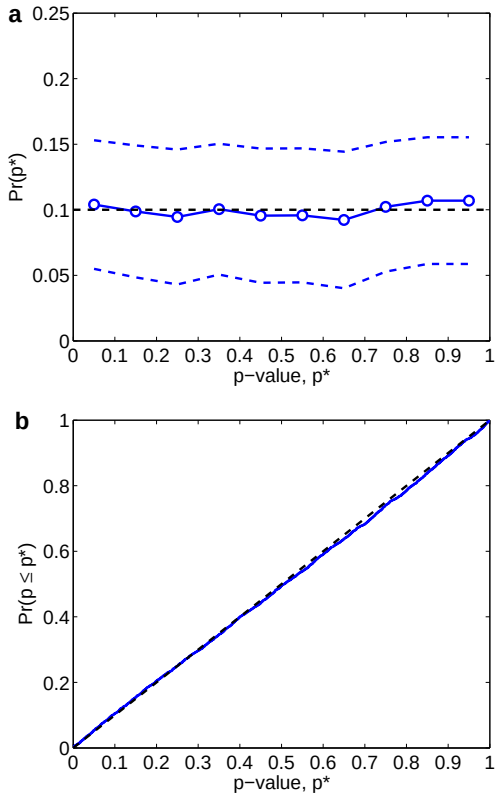

FIG. S5: The numerically estimated distribution of p-values (4000 reps.) for data drawn iid from the generative model described in Eq. (B1), using the procedure described in Appendix C 2, shown as **a**, the histogram with standard error estimates and **b**, the cumulative distribution function. The dashed line shows the null-model of a uniform distribution, and the estimated distribution's uniformity indicates that the Monte Carlo estimation of the p-value is unbiased.

2. Calculation of statistical significance; p-value

We numerically estimate the statistical significance of the model given in Eq. (B1) using a semiparametric Monte Carlo procedure. Under this procedure, we use the empirical experience data $\{k_i\}$, but generate new delay variables $\{\Delta t_i\}$ by drawing them iid from the estimated model, conditioned on their corresponding k_i value (see Appendix B). As the test statistic, we use the Kolmogorov-Smirnov distance [S15] between the empirical data and the truncated log-normal distribution, averaged over experience variables:

$$D^* = \frac{1}{k_{\max}} \sum_{k=1}^{k_{\max}} \max_i |S_k(i) - P_k(i)| \quad , \qquad \text{(C1)}$$

where $S_k(i)$ is the empirical distribution function for experience k and $P_k(i)$ is the corresponding theoretical cumulative distribution function. We average the KS distances, rather than, say, take their maximum, because the number of delays n_k observed at each value of k is roughly inversely proportional to k, and when n_k is small, the KS distance will be

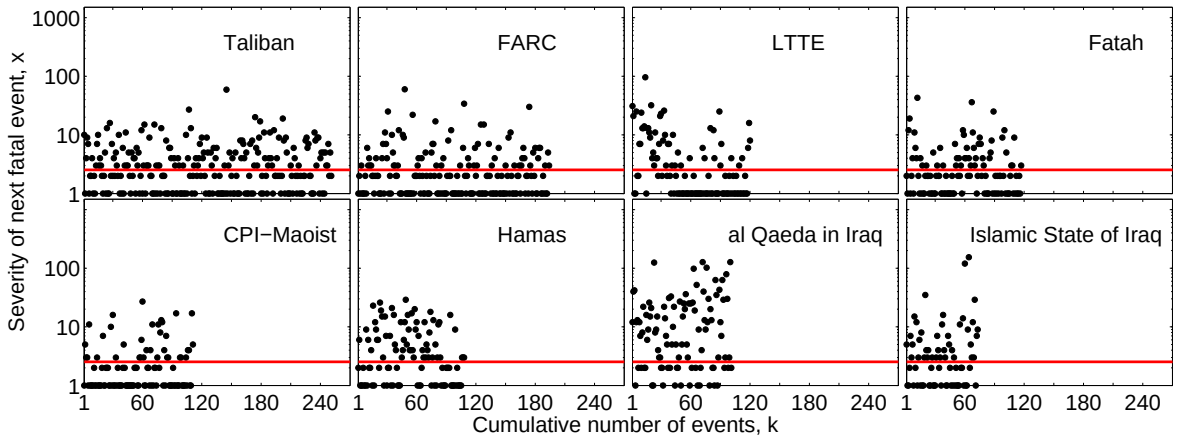

FIG. S6: Individual attack severity development curves for the eight most experienced groups, along with the expected mean severity curve shown in Fig. S7a.

large, regardless of how well (or poorly) the model fits the data. The average KS distance avoids this problem.

To calculate the p-value, we use the following procedure:

1. estimate μ, σ and β for the empirical $\{\Delta t_i\}$ data using maximum likelihood ;
2. compute the empirical value of the test statistic D^* (Eq. (C1)), relative to this model ;
3. for $j = 1 \ldots N$ repetitions (for $N \gg 1$), do the following:
   (a) for each event delay k_i, draw a new delay variable Δt_i independently from $p(\Delta t \,|\, k_i)$ (the model estimated in step 1) ;
   (b) estimate model parameters μ_j, σ_j, and β_j from the data generated in step 3a ;
   (c) compute the test statistic D_j^* for this same data (from step 3a relative to its own estimated model, with parameters μ_j, σ_j, and β_j, from step 3b) ;
4. the p-value is defined as the fraction of the test statistics D_j^* that are at least as large as the empirical test statistic, i.e., $D_j^* \geq D^*$.

To check that this choice of test statistic yields an unbiased p-value, we conducted a numerical experiment to measure the distribution of p-values for data truly drawn iid from the model. To be unbiased, i.e., to have the correct interpretation as a probability that the data were in fact drawn from the null model, this distribution should be uniform on the unit interval. Fig. S5 shows the results of this test, confirming that the test statistic D^* is unbiased.

When applied to the empirical delays, we estimate $p = 0.00 \pm 0.03$, indicating that the null model can be rejected. Because the normalized delay distributions $p(k^{\hat{\beta}}\,\Delta t)$ collapse onto the underlying log-normal distribution, however, this rejection implicates the assumption of independence as being invalid, as opposed to the general form of the model, i.e., there are significant correlations in the timing of attacks.

Appendix D: Analysis of attack severity development curve

Fig. S6 shows the individual attack severity development curves for the same eight groups as in Fig. S4, along with the population-level trend described below. Most notably, these individual development curves largely reflect the combined development curve result that experience and severity are independent, even at the level of individual groups.

Fig. S7a repeats the attack severity developmental curve from the main text. The slope of a simple trend line (regressing $\log k$ against $\log x$) is not distinguishable from zero ($r = -0.0238$, t-test, $p = 0.17$), and the solid line shows the zero-slope null hypothesis. Further, if we consider the distribution of severities at a particular value of k, we see that they do not seem to depend on k (Fig. S7b), and that they resemble power-laws distributions [S4]. Taking this similarity at face value, we can test whether the severity distribution $p(x \,|\, k)$ depends strongly on the group's experience k. Fig. S8 (upper) shows the results of fitting each of the distributions $p(x \,|\, k)$ for the first 40 values of k with a power-law distribution with $x_{\min} = 1$ [S9].

The estimated power-law exponents are very stable: the correlation between $\hat{\alpha}_k$ and k is negligible ($r = 0.0274$, t-test, $p = 0.87$), and the average estimated value $\langle \hat{\alpha}_k \rangle = 1.74$ is indistinguishable from the estimated value for all the data together (i.e., ignoring k), $\hat{\alpha}_{\rm all} = 1.74 \pm 0.01$. Thus, the distribution of attack severities appears to be independent of organizational experience, i.e., the probability that an attack will cause x fatalities does not depend on the number of attacks k a group has done in the past.

This result has several implications for understanding the origin of the power-law form of the severities distribution for terrorist attacks worldwide [S4, S9]. In so far as terror-

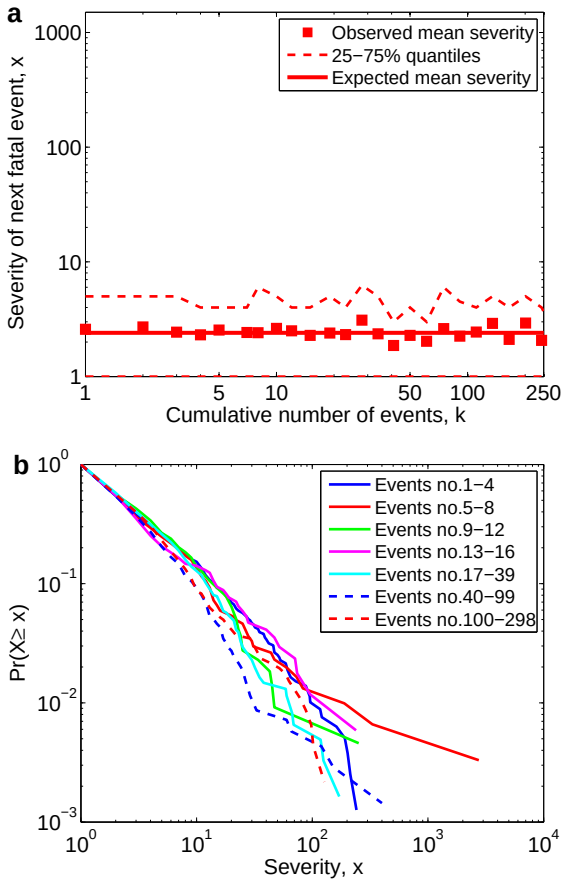

FIG. S7: **a**, The mean severity $\langle \log x \rangle$ of fatal attacks by a terrorist group, with 25^{th} and 75^{th} percentile isoclines, as a function of group experience k. The solid line (with slope zero) shows the expected mean delay, from a simple regression model. **b**, The distributions of event severities $\Pr(X \geq x)$, showing the data collapse onto an underlying heavy-tailed distribution.

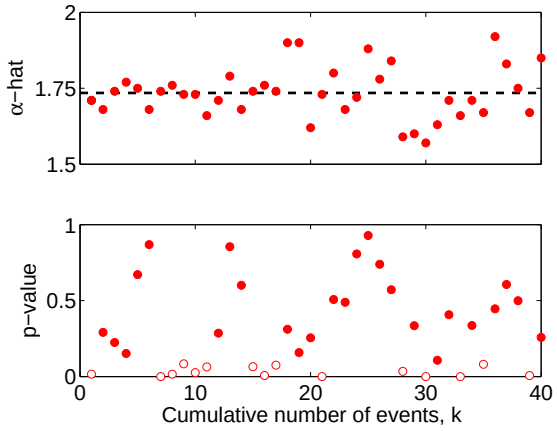

FIG. S8: Power-law analysis of event severities, for $k \leq 40$: (upper) the maximum likelihood exponent $\hat{\alpha}$ for a power-law distribution with $x_{\min} = 1$, and (lower), the corresponding p-value of the fit. The slope of the trend line in the upper panel is not statistically different from zero.

ist attacks are characteristic of the more general category of asymmetric or guerrilla warfare, a model proposed by Johnson et al. suggests that the severity power law comes from self-organized critical behavior in the internal dynamics of terrorist organizations [S16, S17]. Their model assumes that all groups are composed of a number of cells, and that these cells merge (pairwise) or fall apart (into individuals) according to a Markov process. Under these dynamics, the steady-state distribution of cell sizes can be shown to follow a power law under relatively general conditions [S18, S19], and by assuming that cells attack independently, at roughly equal rates, and induce fatalities in proportion to their size, this model yields a power-law distribution in the severity of attacks.

The independence of event severity x and organizational experience k, however, suggests that this explanation requires additional assumptions to explain terrorist attack severities. For instance, we might assume that groups are born with a power-law distribution of cell sizes, as the model otherwise predicts an initial transient period of non-power-law behavior while the cells self-organize away from their initial state toward their steady state. This transient period should appear in our data as non-power-law behavior in the severity of events at small k, with the distribution at $k = 1$ reflecting the distribution of initial cell sizes.

However, Figs. S7a,b and S8 show that even for $k = 1$, we see power-law-like behavior. This suggests that (i) terrorist groups are not internally self-organized critical, (ii) groups converge to their steady-state distribution of cell sizes before making any attacks, or (iii) other assumptions (e.g., about the behavior of cells) conspire in a complicated way to nevertheless produce a power-law distributions for small k. Common sense and historical evidence suggest that the second possibility is probably not the case: most terrorist groups start out small and do not wait very long before beginning their attacks [S20, S21].

An important caveat to our stability analysis is that more of the 40 distributions tested above have $p < 0.1$ than we would expect from an iid situation (37% vs. 10%; Fig. S7a, lower). This implies that there is more structure (or more variance) in the severity data than the simple iid power-law hypothesis would lead us to expect. That is, there are likely significant correlations in the severity of subsequent attacks by the same group. One place this structure could be hiding is in the extreme right tail.

Appendix E: Additional analyses of group size

On the question of how the frequency of attack varies with organizational size, in the main text, we argue that the minimum delay $\min \Delta t$ is the appropriate dependent variable to consider because the size estimate corresponds to the maximum size over the 1998–2005 period. For completeness, we additionally considered whether the mean and median delays

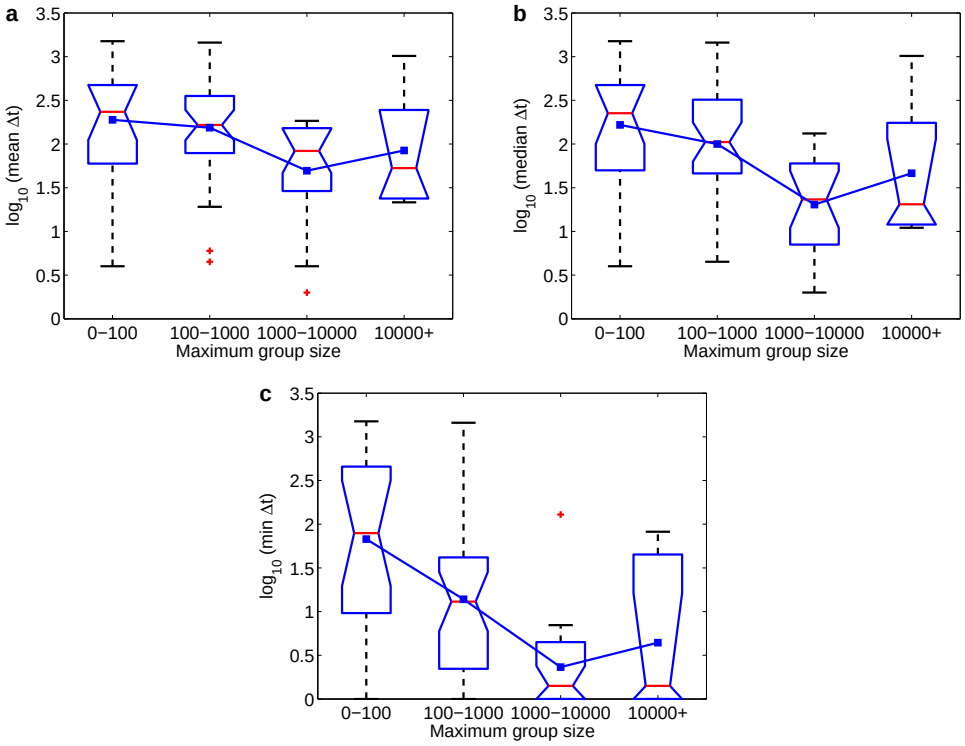

FIG. S9: Boxplots showing the distributions of **a**, mean, **b**, median and **c**, minimum delays between consecutive fatal attacks by groups in the 1998–2005 period. In all cases, the trend shows that larger size is indicative of a faster attack rate.

vary with group size, and found similar results (Fig. S9). In all cases, the categorical means were significantly different (n-way ANOVA, $F = 3.78$ and $p = 0.0139$ for mean delay; $F = 7.44$ and $p = 0.0002$ for median delay), and the trends point in the same direction as for the minimum delay. That is, in all cases, larger organizational size is indicative of greater attack frequency.

Similarly, we considered how the mean, median and maximum severity of attacks varies with organizational size. Here, the mean and median severities are not related to organizational size (n-way ANOVA, $F = 1.14$ and $p = 0.3352$ for mean severity; $F = 0.59$ and $p = 0.6219$ for median severity). Only in the case of maximum severity do we find a significant relationship (n-way ANOVA, $F = 4.53$ and $p = 0.0045$); however, such a relationship is expected, given that larger groups attack much more frequently. That is, consider a situation in which the severity x of an attack is an iid random variable drawn from some distribution P. It is well known from extreme value theory that the expected maximum observed severity will increase monotonically with the number of draws from P [S22]. Thus, because larger groups attack more often, i.e., they have many more chances to produce a severe attack, a positive relationship between the maximum severity and group size is expected.

[S1] Mickolus, E., Sandler, T., Murdock, J. & Fleming, P. International terrorism: Attributes of terrorist events 1968-2003 (ITERATE) (2004). Dunn Loring, VA: Vinyard Software.

[S2] National Memorial Institute for the Prevention of Terrorism. Terrorism Knowledge Base (2008). `http://www.tkb.org` (access date January 29, 2008).

[S3] Danzger, M. H. Validating conflict data. *American Sociological Review* **40**, 570–584 (1975).

[S4] Clauset, A., Young, M. & Gleditsch, K. S. On the frequency of severe terrorist events. *Journal of Conflict Resolution* **51**, 58–87 (2007).

[S5] Snyder, D. & Kelly, W. R. Conflict intensity, media sensitivity and the validity of newpaper data. *American Sociological Review* **42**, 105–123 (1977).

[S6] Miller, G. D. Confronting Terrorisms: Group motivation and successful state policies. *Terrorism and Political Violence* **19**, 331–350 (2007).

[S7] Jones, S. G. & Libicki, M. C. *How Terrorist Groups End: Lessons for Countering al Qa'ida* (RAND Corporation, Arlington, 2008).

[S8] Asal, V. & Rethemeyer, R. K. The nature of the beast: Organizational structures and the lethality of terrorist attacks. *Journal of Politics* **70**, 437–449 (2008).

[S9] Clauset, A., Shalizi, C. R. & Newman, M. E. J. Power-law distributions in empirical data (2009). *SIAM Review*, to appear.

[S10] Nelder, J. & Mead, R. A simplex method for function mini-

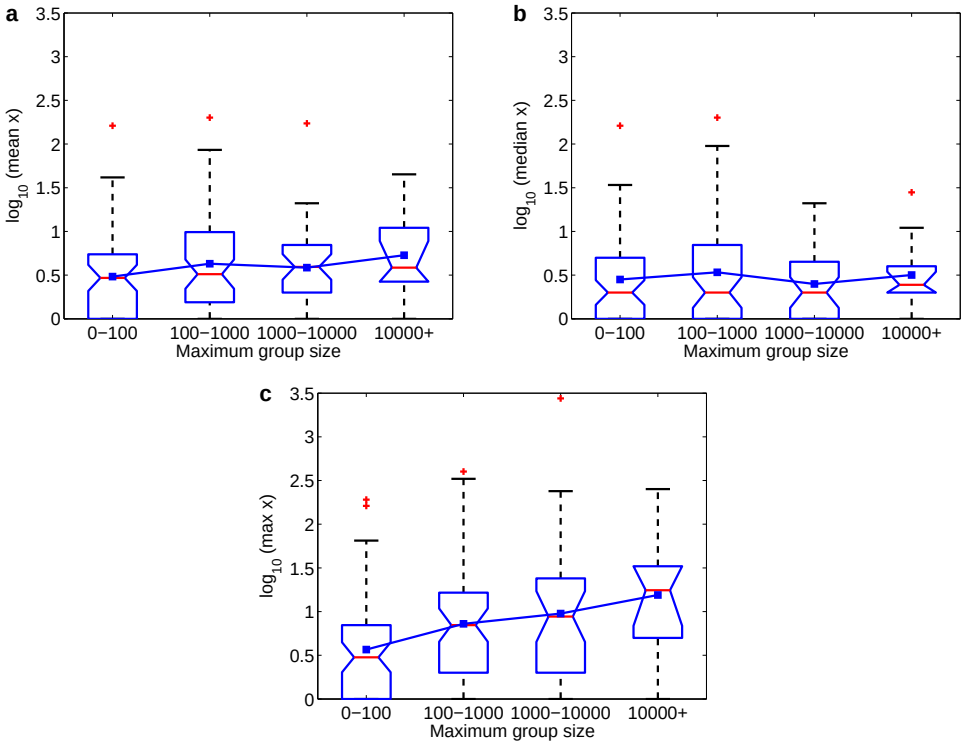

FIG. S10: Boxplots showing the distributions of **a**, mean, **b**, median and **c**, maximum severity of fatal attacks by groups in the 1998–2005 period. For the mean and median severity, larger organizational size is not indicative of more severe attacks; only in the case of maximum severity is a significant relationship found, which is to be expected for other reasons (see text).

mization. *Computer Journal* 308–313 (1965).

[S11] Barndorff-Nielsen, O. E. & Cox, D. R. *Inference and Asymptotics* (Chapman and Hall, London, 1995).

[S12] Argote, L., Insko, C. A., Yovetich, N. & Romero, A. A. Group learning curves: The effects of turnover and task complexity on group performance. *Journal of Applied Social Psychology* **25**, 512–529 (1995).

[S13] Gallistel, C. R., Fairhurst, S. & Balsam, P. The learning curve: Implications of a quantitative analysis. *Proc. Natl. Acad. Sci. USA* **101**, 13124–13131 (2004).

[S14] Jones, J. C., Myerscough, M. R., Graham, S. & Oldroyd, B. P. Honey bee nest thermoregulation: Diversity promotes stability. *Science* **305**, 402–404 (2004).

[S15] Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press, Cambridge, UK, 1992).

[S16] Johnson, N. F. *et al.* From old wars to new wars and global terrorism (2005). E-print, `arxiv:physics/0506213`.

[S17] Johnson, N. F. *et al.* Universal patterns underlying ongoing wars and terrorism (2006). E-print, `arxiv:physics/0605035`.

[S18] Ruszczycki, B., Burnett, B., Zhao, Z. & Johnson, N. F. Relating the microscopic rules in coalescence-fragmentation models to the emergent cluster-size distribution (2008). E-print, `arxiv:0808.0032`.

[S19] Clauset, A. & Wiegel, F. W. A generalized fission-fusion model for the frequency of severe terrorist attacks (2009). E-print, `arxiv:0902.0724`.

[S20] Hoffman, B. *Inside Terrorism* (Columbia University Press, New York, 1998).

[S21] Sageman, M. *Understanding Terror Networks* (University of Pennsylvania Press, Philadelpha, 2004).

[S22] de Hann, L. & Ferreira, A. *Extreme Value Theory: An Introduction* (Springer-Verlag, New York, 2006).