Modern Wars: Fundamental Patterns and Predictions

By

Professor Michael Spagat

Department of Economics

Royal Holloway, University of London

These slides are based on work with Stijn Van Weezel, recently of Royal Holloway, and Neil Johnson, Department of Physics, University of Miami.

Presentation given at the UK Ministry of Defence, March 30, 2016

Here are two diametrically opposed positions:

1. All modern wars are the same.

2. Each modern war is unique

Of course, the truth has to lie somewhere in between.

However, I hope this talk will move you a bit toward position 1.

Some *fundamental patterns* are *rather constant across wars* – there seems to be something about how humans fight that is resistant to specifics of time, place, culture, technology, etc..

Today we will look at *event data*.

An event is a discrete violent occurrence such as a suicide bombing, battle, or air strike.

Today I will focus on just the size of each event where size refers to how many people are killed in the event.

However, we have an equally active stream of research on the timing of events where timing refers to the number of days since the last attack and until the next attack – but you will have to invite me back if you want to hear about this work.

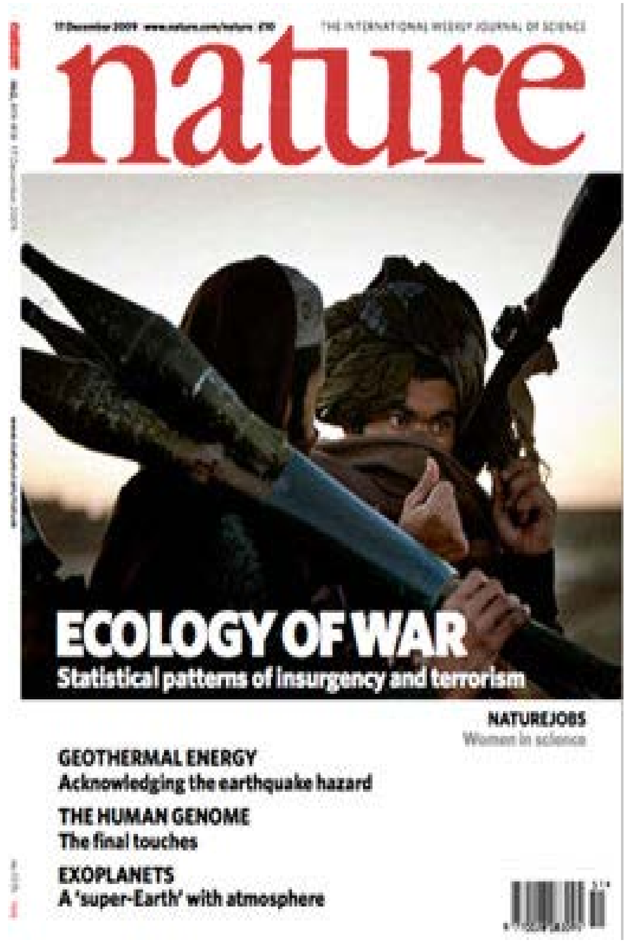

In 2009 Bohorquez et el. found remarkable regularities in the size distribution of violent events in seven different modern wars.

The graph below shows the fraction of events above all possible sizes in Afghanistan:

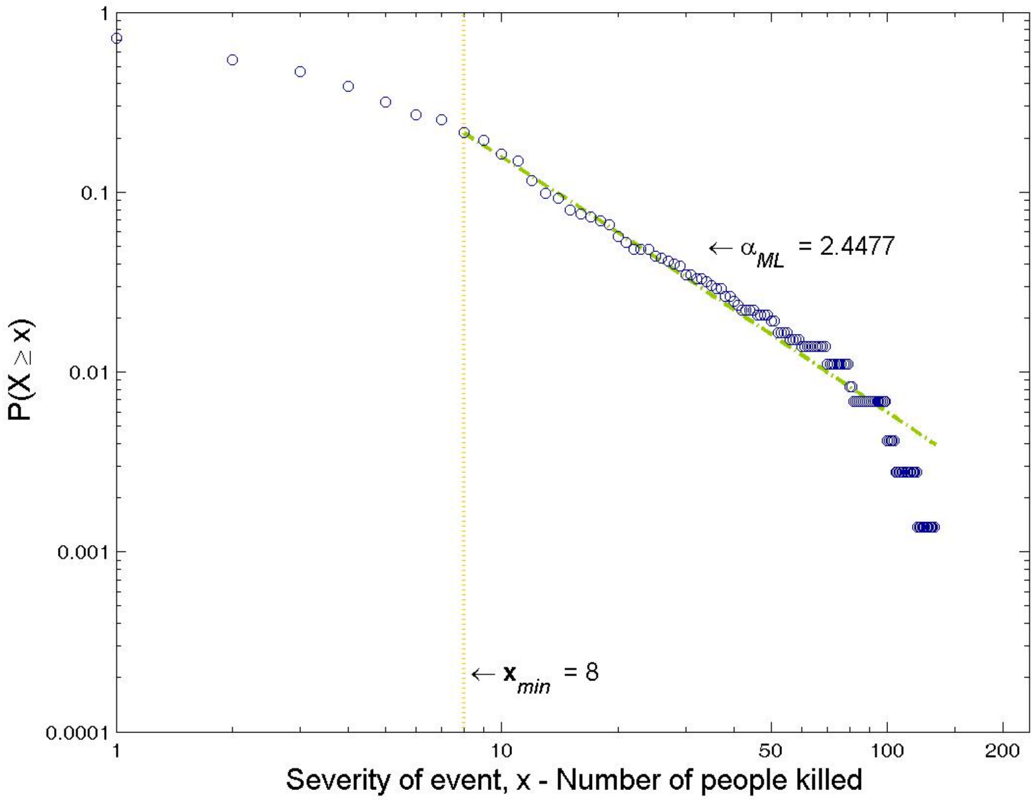

Notice that the points fit a straight line nicely for events with 8 or more people killed.

Power Law

$p(x)$ - the probability that an event will be of size x

If the distribution of event sizes fits a power law above some minimum event size then:

$$p(x) = Cx^{-\alpha} \quad \text{for} \quad x \geq x_{\min}$$

Take the logarithm of both sides and you have a straight line as in the picture on slide 6:

$$\ln(p(x)) = \ln(C) - \alpha \ln(x)$$

We obtained event data on 7 modern wars and found that all of them:

1. Fit a power law reasonably well above some minimum value.

2. Fit more or less the same power law – the estimated alphas are all around 2.5.

In addition, earlier work by Clauset, Young and Gleditsch found that the size distribution of global terrorist events fits a power law with an alpha around 2.5.

We have a family of models that can generate the sorts of event size distributions we see in the data:

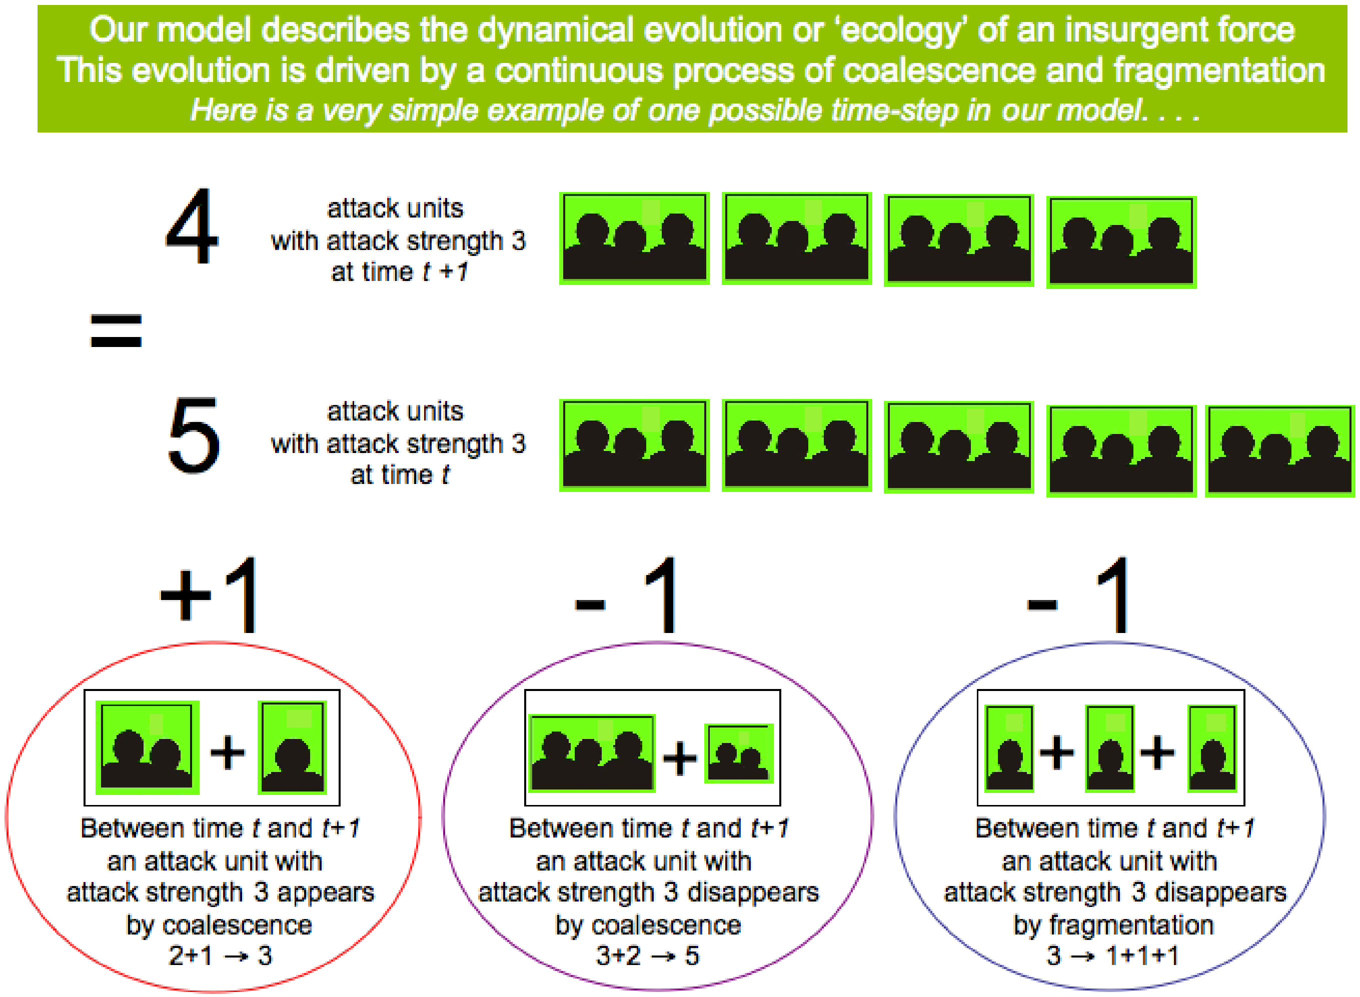

It is useful to keep three distinct things separate in your mind:

1. The *empirical data*, some of which is shown in slide 6. This is the information we have on the size distribution of real events in a real war. This record will, necessarily, be incomplete although it should be relatively complete for the bigger events.

2. *Power laws* – we use these distributions to fit the empirical data. Such fits will never be perfect and are only meant to be good above some minimum size, $x_{\min}$.

3. *Theoretical models*, of which one is sketched on slide 9. These generate size distributions of violent events that fit both the empirical data and estimated power laws reasonably well but, again, not exactly.

The rest of today's talk will be relentlessly empirical.

We will:

1. Fit power laws to conflict data.

2. Use our estimated power laws to predict future distributions of violent event sizes in future conflicts.

There are two key properties of power laws:

1. *Fat tails* – big events are much more common than they would be if the size distribution of violent events followed a normal distribution (Bell curve).

Indeed, suppose we fit a power law with alpha = 2.5 to the heights of all Americans, normalizing so that the mean of the power law equals the mean for all Americans. Then there would be about 60,000 Americans as tall as the tallest person on record, 10,000 people as tall as a giraffe and one about as big as the Empire State Building.

2. *Scale invariance* – the ratio of the number of events of size x to the number of events of size y depends only on the ratio of x to y.

Let's explore scale invariance a little more – assume that alpha equals 2.5.

The ratio of the number of events of size 10 to the number of events of size 15 is:

$$\frac{C \times 10^{-2.5}}{C \times 15^{-2.5}} \approx 2.8$$

The ratio of the number of events of size 20 to the number of events of size 30 is:

$$\frac{C \times 20^{-2.5}}{C \times 30^{-2.5}} \approx 2.8$$

So, as stated in slide 9, the only thing that matters for the ratio of the frequencies is the *ratio of the event sizes*.

This means that if we believe in alpha = 2.5 then we can make lots of quick predictions about the mixture of events of different sizes.

For example, a similar calculation shows that the ratio of events of size 10 to events of size 20 should be around 5.7. The same is true for size 15 versus size 30, etc.

We need one last technical wrinkle before moving on to predictions.

As with slide 6 we prefer to work in terms of events of a certain size *or greater*. Going over these wide ranges smooths out the data so if, for example, too many events of one size tends to be balanced by too few of a different size.

The ratio of events of size x or greater to events of size 2x or greater is around 2.8 (using calculus).

The ratio of events of size x or greater to events of size 1.5x or greater is around 1.8.

We are living in the middle of a data revolution - the [UCDP GED] database gives us event data for a large number of conflicts fought between 1989 and 2014 in Asia and Africa with more data on the way.

After eliminating a few conflicts from the database for which there are fewer than 30 events we are left with:

1. Africa – 98 conflicts with 21,239 events

2. Asia - 104 conflicts with 60,162 events

We fit a power law to each of the 202 conflicts - the next few slides summarize our findings.

This picture on the next slide gives:

1. The estimated alphas for the Africa conflicts.

2. The number of events for each conflict – more events should improve statistical reliability.

3. A p value for each conflict for a test of the hypothesis that the data for that conflict are generated by the power law we have estimated for that conflict – a low p indicates a low probability that the data are generated by the estimated power law.

Here is a summary picture for the Africa power law estimates:

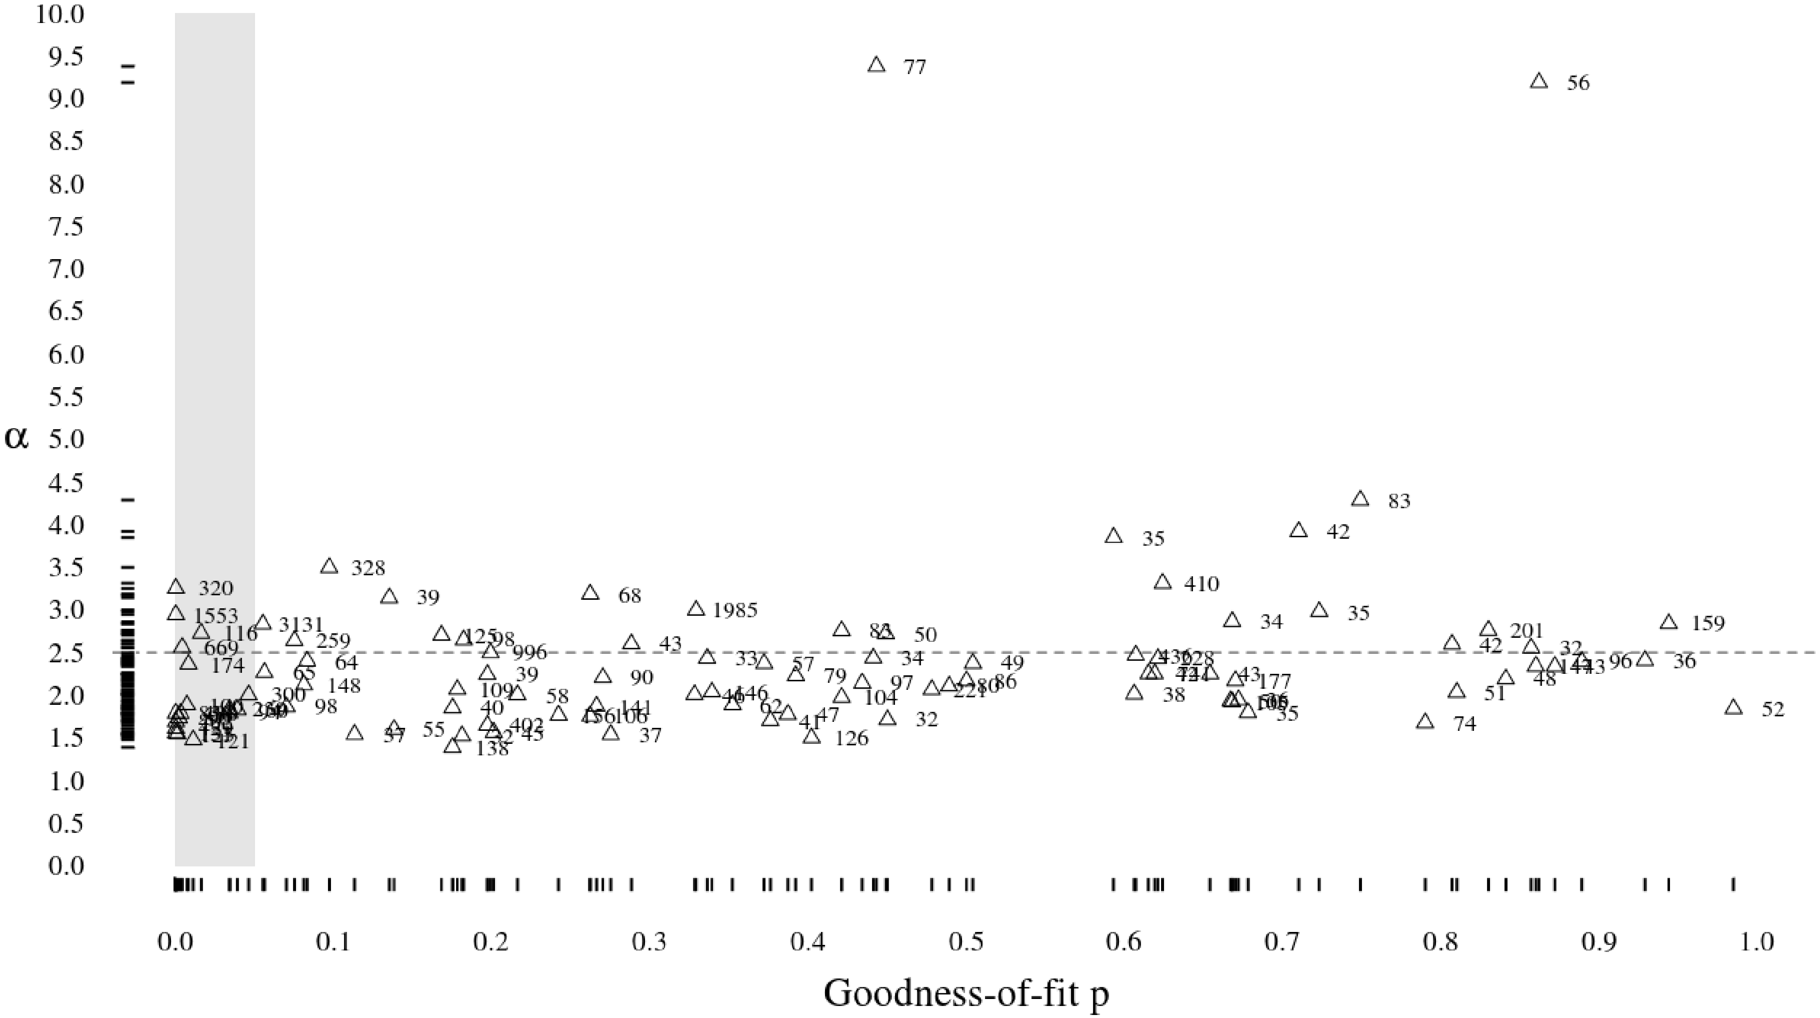

Observations:

1. The alpha values cluster around 2.5 although there are some spectacular misses.

2. The hypothesis that the data are generated from power laws fares pretty well although we sometimes reject the hypothesis at the usual significance level of 0.05. We should bear in mind that it is farfetched to imagine that the data for any conflict is literally generated by an exact power law so we should expect to reject these models with enough data.

3. If you drop the 10 highest and 10 lowest alphas then the rest of the alphas are between 1.6 and 3.1.

Here is the same picture for Asia:

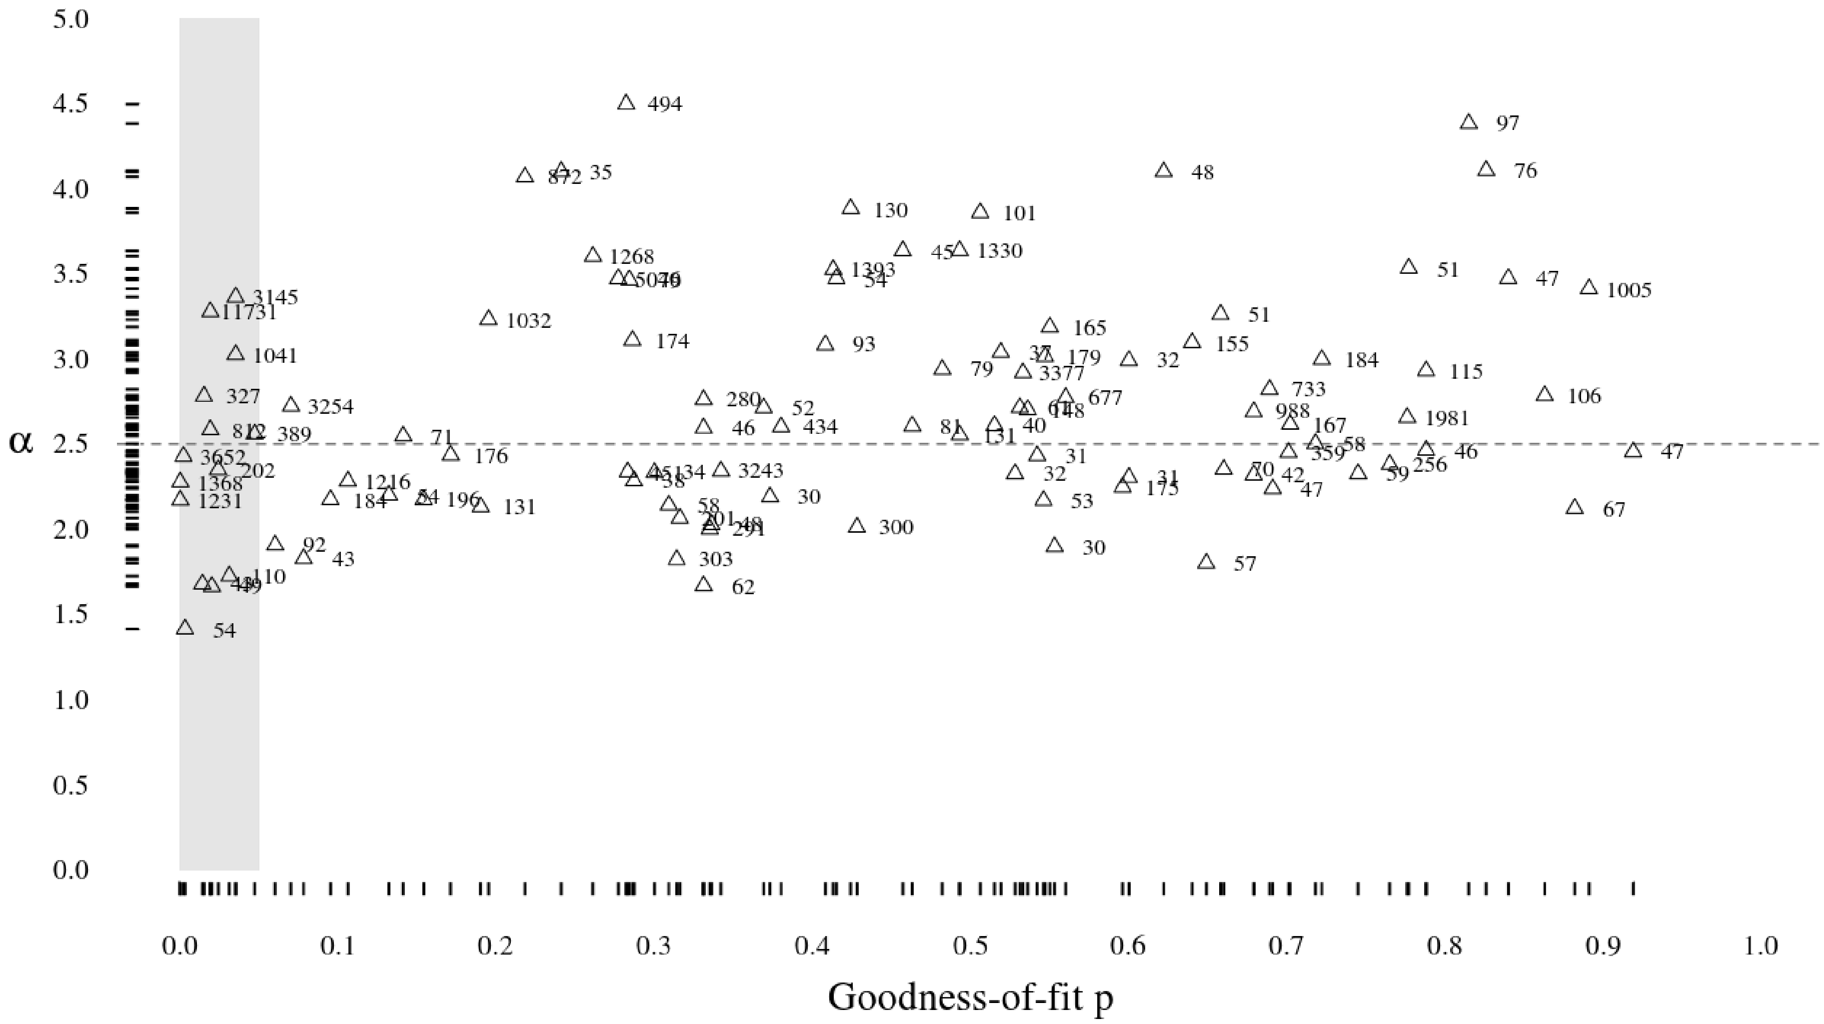

Observations:

1. The Asia picture is similar to the Africa one except that it does not have such extreme outliers.

2. If you drop the 10 highest and 10 lowest alphas then the rest of the alphas are between 2.0 and 3.6. So the alphas are a bit higher for Asia than they are for Africa but distance between the 10^{th} and 90^{th} percentiles is similar for the two continents.

This table shows how the values of alpha affect the ratios of the numbers events of various sizes.

Alpha Event Size Ratios	1.5	2.0	2.5	3.0	3.5	4.0
(>x)/(>1.5x)	1.2	1.5	1.8	2.25	2.8	3.4
(>x)/(>2x)	1.4	2.0	2.8	4.0	5.7	8.0

You can read from the table that, for example, the when alpha = 2.0 there should be twice as many events of size 10 or greater than events of size 20 or greater.

So if you use alpha = 2.5 to make predictions about the relative numbers of events of different sizes you can do reasonably well as long as the sizes of the events you are predicting are not extremely far apart.

We also can our empirical distribution of all the estimated alphas for the 202 conflicts to put a range around our point estimates using the following procedure:

1. List all of our 202 alpha estimates from smallest to largest.

2. For point estimates we use the median alpha which is 2.4. (The mean is 2.6, or 2.5 without the two outliers, but we can use 2.4 for point estimates.)

3. The alpha at percentile 2.5 is equal to 1.5

4. The alpha at percentile 97.5 is equal to 4.1.

5. We use these alphas of 1.5 and 4.1 to predict ranges for the ratios of the numbers of events of various sizes - slide 22 gives these predicted ranges almost exactly.

A more stringent exercise is to try these predictions on new data – it is much harder to make good predictions outside your sample than it is to make predictions inside your sample.

We will have to wait for truly new data to test fully out of sample predictions.

However, we can still test our prediction ability now by splitting our sample into two pieces and making predictions for one piece of the sample based on analysis of the data in the other piece of the sample.

Specifically, the procedures are:

1. Select 1/3 of the conflicts at random, i.e, select 67 out of the 202 conflicts.

2. Follow the procedures on slide 23 for these 67 conflicts.

3. Use the range of alphas for these 67 conflicts to make predictions with ranges for the ratios of event counts of various sizes for the remaining 135 conflicts.

4. Take many more random draws of 1/3 of the conflicts and record the outcomes every time.

This picture summarizes the results from this prediction exercise.

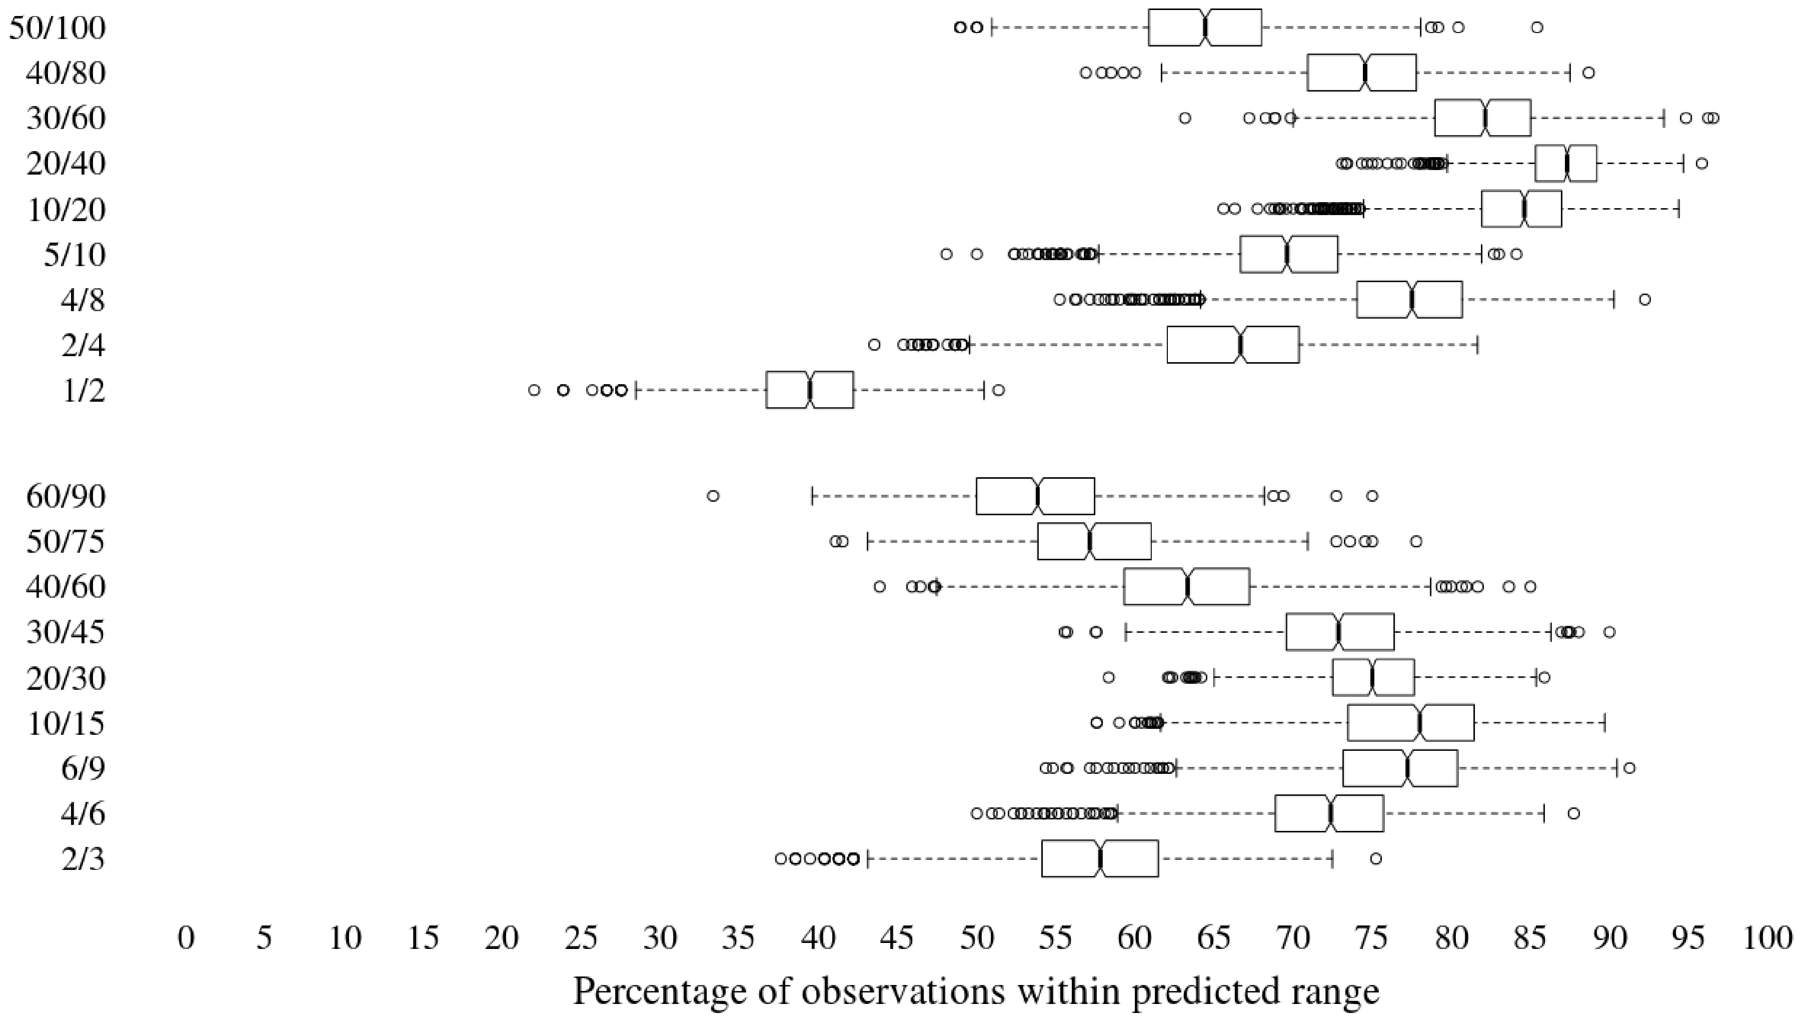

The picture on slide 26 is complicated so let's work through some bits of it. Start with the 10/15 predictions.

1. The computer picks 67 conflicts at random and uses their alpha estimates to place lower and upper bounds on the ratio of the number of events of size 10 or greater to the number of events of size 15 or greater. These bounds depend on which conflicts are picked but they will be something like 1.2 to 3.5 (see slide 22).

2. The computer then checks each of the remaining 135 conflicts (out of sample) and records a success if the ratio for that conflict is within the range established by the random draw of 67 conflicts. Otherwise, the computer records a failure.

3. The computer records the fraction of successes for that random draw and moves on to the next of 1,000 random draws.

4. From the picture on slide 26 we can see that the median success rate was just below 80%. This means that for this (median) draw the ratio of the number of events of size 10 to the number of events of size 15 fell within the predicted range for around 108 out of the 135 out-of-sample conflicts.

5. We can also read that half the draws have success rates between around 73% and 82% (if my eyeballs serve me)

6. All but a handful of draws have success rates between around 62% and 90%.

Here are some observations from scanning the whole picture.

1. The worst predictions are for the smallest events and the biggest events.

a. Power laws are only meant to apply above some $x_{\min}$ so relatively weak performance for small events is to be expected.

b. Big events are relatively rare so data are sparse at the high end and, again, relatively weak performance makes sense.

2. The best performances are for the 10/20 and 20/40 predictions. I had expected that these would be harder than x/1.5x predictions but I am guessing that the higher success rate is because the prediction ranges are much wider for the x/2x predictions than they are for the x/1.5x predictions.

3. For most of the comparisons the success rates are above 60% for at least 75% of the random draws.

Conclusions

If a new conflict started next week anywhere in the world I would be fairly confident that I could make decent predictions about the distribution of violent-event sizes in that conflict.

I would go ahead and make my predictions without any knowledge of the specific features of that conflict – presumably such features could be used to improve the predictions although right now I would not know how to use them.

Such predictions may be of some use in military planning.

But there is deeper significance to the feasibility of making such predictions - it suggests that there are fundamental features of modern wars that transcend specifics of time and place.

It would be a gross exaggeration to say that all modern wars are the same but, still, there do seem to be strong commonalities in the size distributions of violent events that ranges across wars.

Moreover, we have similar results for the *timings* of violent events….but this will have to wait for another day.