PT PERSPECTIVES ON TERRORISM

Research Note

Power-Law Online Terrorism and Extremism

Stephane J. Baele*

Volume XIX, Issue 2
June 2025

ISSN: 2334-3745
DOI: 10.19165/FGGW5148

Abstract
This Research Note draws attention to the ubiquity of heavy-tailed, power-law types of data distribution in all major dimensions of digital extremism and terrorism, and discusses the important implications stemming from this pattern as well as the mechanisms underpinning it. Highly unequal distributions with a very long lower tail and a sharply rising upper tail indeed characterize linguistic content, individual behaviours, and overarching structures of extremist/terrorist digital ecosystems – a striking feature that has largely escaped terrorism and extremism scholars' scrutiny as well as Counter Violent Extremism (CVE) and Counter-Terrorism (CT) practitioners, despite its scientific and practical significance.

Keywords
Online extremism; power-law; heavy-tailed distributions; terrorism.

**Corresponding author: Stephane Baele, ISPOLE, UCLouvain (Belgium). Email: stephane.baele@uclouvain.be*

Introduction

In 1949, a singular book drew attention to what appeared to be a quasi-ubiquitous statistical regularity in human activity. In *Human Behavior and the Principle of Least Effort*, George Zipf (1949) showed that plotted ranked distributions of seemingly *any* linguistic item in written and spoken language[1] displays the same pattern: a few occurrences have very high frequencies while most items are only found rarely (hence constituting a 'heavy tail' in the distribution graph), in contrast to normal bell-shaped distributions.[2] In other words, all these distributions resemble power-laws (where graphs follow the general form $y=cx^a$), revealing, according to the author, "a matter of fundamental regularity" surely emerging from "an underlying governing principle." In the final section of his book, Zipf looked beyond linguistics, and revealed that this regularity holds in most domains of human activity, from the size of cities to income;[3] dozens of almost identical graphs led him to claim that human societies invariably distributed power, capabilities, and prestige highly unequally. A staunch positivist, he presented his book as a demonstration that scientifically observing human interactions in the same, detached way we study objects and animals, similarly leads to the discovery of general laws governing the 'human ecology.'

Seventy-five years after Zipf's volume, power-law types of distributions have indeed been shown to characterise datasets pertaining to pretty much all domains of human – and even non-human – activity, from US or Chinese family names to Tour de France performances, and from moon craters to the evolution of species. As Reed and Hugues summarise, "distributions with power-law behaviour in one or both tails are ubiquitous in physics, biology, geography, economics, insurance, lexicography, internet ecology, etc."[4] Newman similarly writes that "one can, without stretching the interpretation of the data unreasonably, claim that power-law distributions have been observed in language, demography, commerce, information and computer sciences, geology, physics and astronomy, and this on its own is an extraordinary statement."[5]

All this may seem quite removed from the considerations and preoccupations keeping awake counter violent extremism (CVE) and counter-terrorism (CT) scholars and practitioners. Yet as this Research Note seeks to show, heavy-tailed power-law types of distributions also abound in the realm of digital terrorism and extremism, and this observation has key implications when it comes to both understanding dynamics of radicalisation and violence and designing countermeasures. Indeed knowing that the distribution of a certain phenomenon is akin to a power-law (rather than, say, a bell-shaped) not only helps us specify the nature of this reality, but also unlocks sound strategies to address it. We will see, for example, that the vast majority of members in any given extremist forum barely contribute to discussions, and therefore that broad-brush strategies targeting (non-existent) posters with 'average' activity would prove inefficient.

As such, the present work highlights, builds on, and complements, both at once the literature on power-laws on the Internet and computer systems[6] and the few analyses of conflict and terrorism studying heavy-tailed distributions.[7] It does this by showing that power-law types of distributions are found across the three major dimensions of online terrorism/extremism: 1) the linguistic content of what is being said in extremist/terrorist online spaces, 2) the behaviour of the users within these spaces, and 3) the structures and hierarchies connecting these spaces (as well as their users) together.

I proceed in three steps. First, I put forward two methodological considerations that together frame how (not) to understand the present endeavour; second, I provide empirical evidence for heavy tailed, power-law type of distributions across the three key dimensions of online extremism; third, I discuss three meanings and implications of this recurring pattern for knowledge about, and measures against, digital extremism and terrorism. It is hoped that this empirical snapshot and discussion will trigger a wider interest in a widespread yet still poorly understood phenomenon, and unlock new, innovative avenues for CT-CVE efforts.

Two methodological considerations: Datasets and fitness to power-law

Before delving into the data, two brief methodological considerations are warranted. First, the graphs presented below should not be understood as particularly exceptional or unique, cherry-picked instances of power-law types of distributions found in a single dataset. Rather, they are typical examples among many similar cases encountered throughout more than a decade of research on digital terrorism and extremism using large datasets and corpora.[8] As such, the empirical section below merely displays the tip of the iceberg: it provides typical illustrations of original data systematically collected from niche to very popular digital spaces and communication channels used by terrorists and extremists (positioned across the ideological and linguistic spectrums). While the examples come from the far-right and incel digital ecosystems, similar data is found in other ideological online spaces such as jihadist forums.

Second, this research note does not claim that all presented distributions perfectly fit power-laws, which is why I write "heavy-tailed" or "power-law *types*" of distributions. Following Zipf's practice, a long-lasting first wave of research has used count and rank distribution graphs to evidence cases of power-laws, usually replotted with logarithmic x and y axes to quickly evaluate closeness to power-law: in "log-log" graphs, power-law distributions indeed appear as straight lines, in contrast to other functions similarly displaying heavy tails (such as some exponential, Weibull, or polynomial variants). For instance, Reed and Hughes observe that when it comes to the frequency of US family names "the closeness of the points to a straight line (corresponding to power-law behavior) is impressive," a pattern they find in a range of other contemporary phenomena.[9] However, a more recent strand of research has lamented that the classic literature tends to "rely solely on qualitative inspection of a log-log plot to detect power-law behaviour', which 'is a not a reliable approach"[10] as it is not discriminative enough to proclaim the existence of a genuine empirical power-law – and hence to clarify the exact dynamics generating phenomenon captured by the distribution. Some scholars have even shown that using "careful statistical tests to data that were reported to exhibit power-law relationships" results in 'negligible' evidence;[11] Clauset, Shalizi and Newman's systematic testing demonstrated that "there is only one case – the distribution of the frequencies of occurrence of words in English text – in which the power law appears to be truly convincing."[12] For Lima-Mendez and van Helden, who examined power-law claims in biological networks, this 'lack of consistency between theoretical models and data' mean that the ubiquity of power-law distributions is no more than a 'myth'.[13] However, this type of criticism has itself been toned down by scholars like Zhang and colleagues, who have suggested that sampling and measurement are important factors explaining imperfect fit with power-laws[14] – a claim which raises profound epistemological and ontological issues (and which is itself derided by Lima-Mendez and van Helden).[15]

This paper sidesteps this debate by taking a pragmatic stance. I contend that the extent to which the data presented below fits a power law is irrelevant for the mainly propaedeutic purpose of this Research Note, which merely seeks to raise awareness of a phenomenon and trigger reflections about its implications. At this stage, we simply aim at displaying a striking regularity in online extremism and terrorism datasets (that they resemble power-law distributions), and to infer, at a descriptive level only, insights about why and how these distributions matter. It is therefore hoped that this Research Note will trigger more thorough articles studying with more depth and accuracy the complex dynamics causing the distributions displayed here (for example, with fitness tests against power-law and other types of heavy-tailed distributions). At the level of the present study, what Stumpf and Porter wrote is therefore applicable: “knowledge of whether or not a distribution is heavy-tailed is far more important than whether it can be fit using a power law.”[16] I, therefore, stick to the field’s older common rule of thumb that “a candidate power law should exhibit an approximately linear relationship on a log-log plot over at least two orders of magnitude.”[17] At this Research Note stage, this level of analysis is already sufficient to highlight important dynamics and situations characterising digital extremism and terrorism.

Empirical evidence

Language in extremist and terrorist online spaces

To perpetuate Zipf’s lineage, I first investigate the linguistic content of extremist digital spaces linked to terrorist activity. Extracting and plotting words or n-grams frequencies within a range of extremist forums, blogs, websites, and social media accounts (on both mainstream platforms like Facebook and “alt-tech” ones like Telegram or Gab), as well as within entire coherent groupings of such spaces (for example, the “incelosphere”,[18] or the constellation of French extreme-right websites), typically reveals power-law types of exponent graph lines with extremely long low tails, regardless of ideology or language. For example, Figure 1 below plots side-by-side the 1,000 most frequently occurring bi-grams in original posts made to the infamous “political” (i.e., white supremacist) chatroom of the *8chan* image-board,[19] on the one hand, and in the content scraped from the website of the banned French far-right movement *Les Identitaires*, on the other hand.[20] Despite differences in language and size (the first corpus contains more than 8 million words, against approximately 150,000 for the second), the two panels are remarkably similar and clearly resemble power-law types of distributions (especially in their log-log variant). Notice how the three most frequent bigrams on *8chan* (‘United States’, ‘Donald Trump’, and ‘social media’) and *Les Identitaires* (‘France’, ‘même’, and ‘Français’) occupy very similar positions on their respective graphs despite occurring at widely different counts and having different meaning. Comparable graphs are obtained with a variety of other corpora, in line with the original literature on power-laws in language (including Zipf’s book).

Figure 1. Distribution of most frequently occurring bigrams on the Les Identitaires website (left, range) and the 8chan/pol image-board (right, blue)

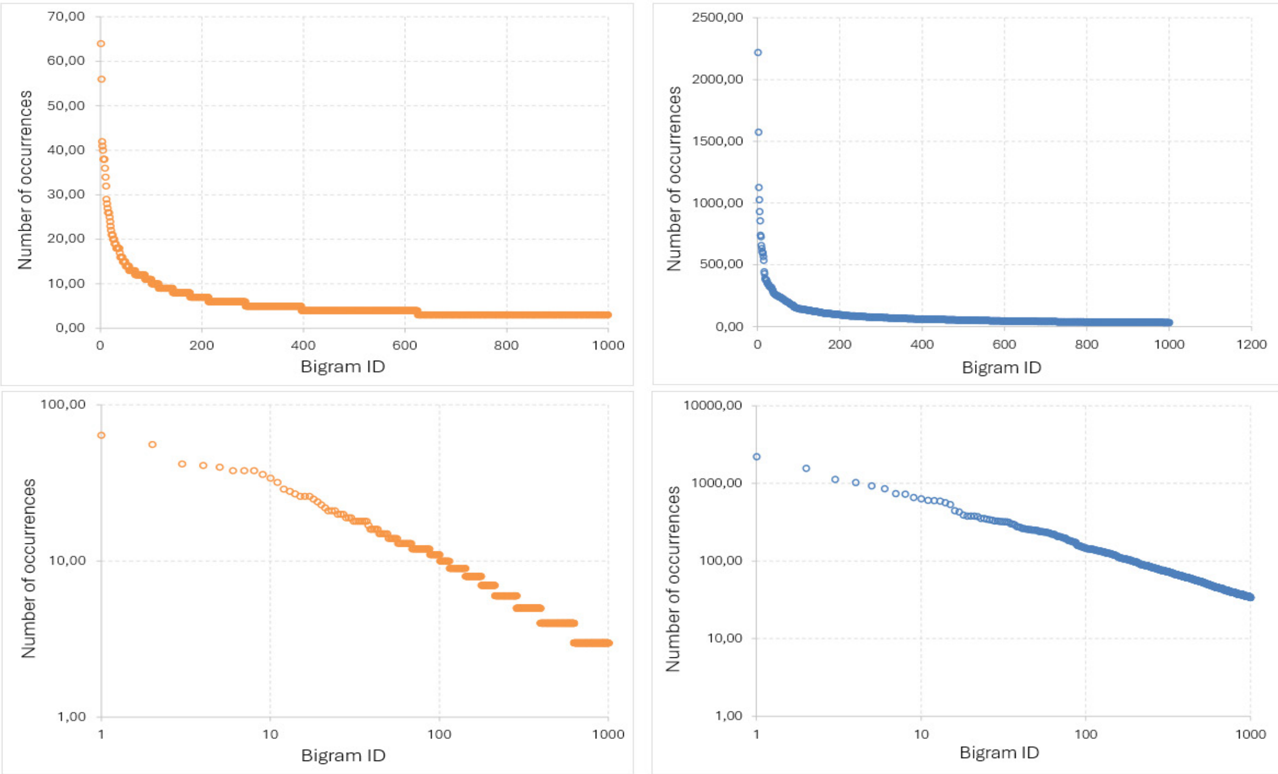

Users' behaviour in extremist online spaces

I now turn to the behaviour of users within these spaces, such as individual posting volume or propensity to outlink or react (for example, like, share, reply) to others' contributions. Initial evidence for heavy-tailed distributions for extremist digital behaviours already exists. In a recent paper coalescing anecdotal observations on the 'super-posters' phenomenon,[21] Baele and colleagues' application of the Jenkins-Fisher algorithm to individual posting volume data from seventeen extremist forums (including jihadist, far-right, and incel ones) produced invariable results. All of them were mathematically characterised by a 4-categories structure of posting activity topped by a small clique of "hyper-posters" 'responsible for a disproportionate number of the interactions'.[22] Plotting this activity in rankings, the authors incidentally remarked, in passing, that 'posting in all the forums follows a power-type curve where the most active members occupy the extreme end of a continuum', an observation that 'holds in a remarkably stable way across all forums'.[23] Here we highlight this contribution and confirm its observations.

Figure 2 below displays two illustrative examples of posting activity from two extremist forums of different sizes and expressing dissimilar ideologies linked to acts of terrorism. The left panel ranks members of the *Incels.net* forum[24] according to their posting volumes; one person posted no less than 15,325 times and the second most active member posted 11,547 times; the heavy tail is populated by 598 individuals who have posted only once and 324 twice. The right panel represents users' posting volumes on the small white supremacist *Chimpout.com* forum; only 167 members were active during data collection, with one individual posting 3,141 times and the second most active person 2,287 times; at the other end, 46 people posted only once and sixteen twice. While slightly curved lines on the log-log plots suggest that other functions than power-laws might better fit these distributions, these graphs nonetheless demonstrate strikingly similar appearances with two far outliers followed by a less engaged minority and a lethargic mass of users.

Figure 2. Posting behaviour on Incels.net (left, orange) and Chimpout.com (right, blue)

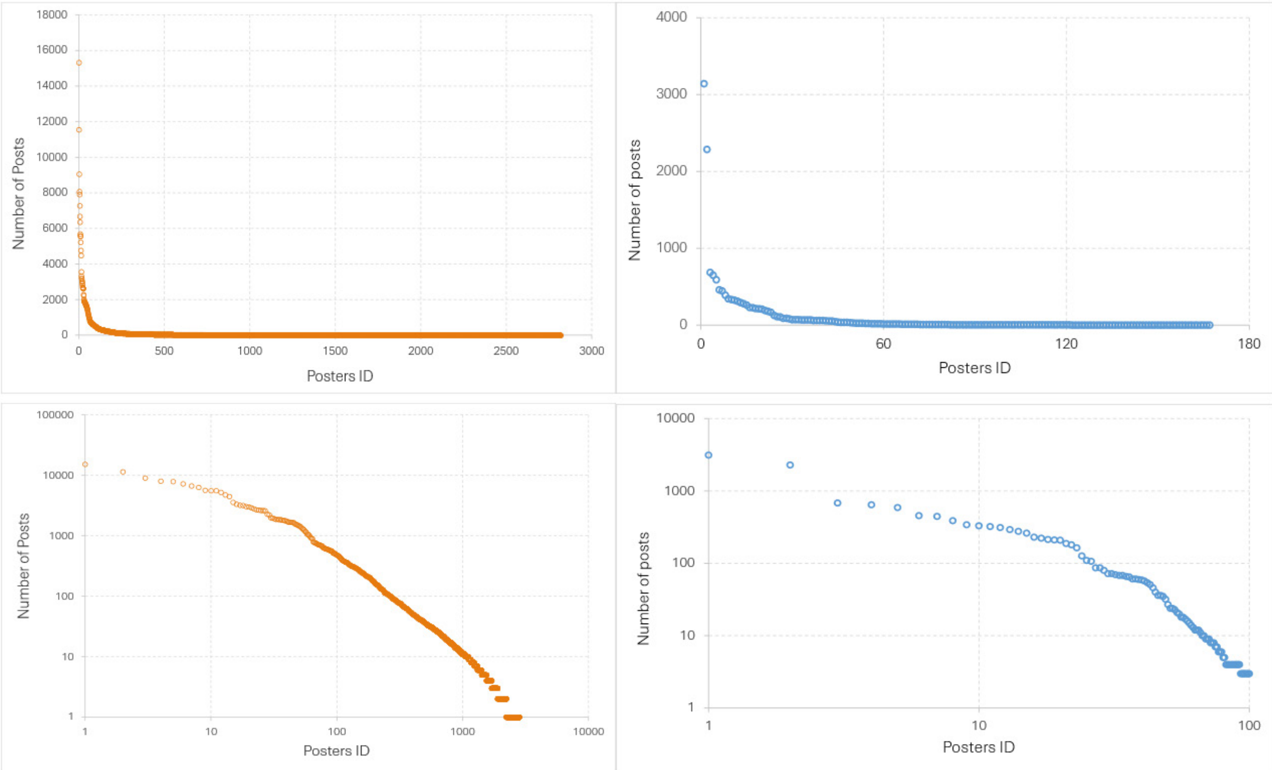

Influence and hierarchy in extremist digital ecosystems

Finally, I turn to the last key dimension of online terrorism and extremism: the structures and hierarchies within, and across, these spaces. Typically, the literature approximates the importance, status, or influence of an online space or user (for example, an Instagram account, a blog, a member of a forum) with simple network statistics such as in-degree (number of incoming connections), centrality, or algorithms like PageRank or HITS. These statistics, which help identify and rank the most popular or strategic nodes within a broader interconnected ecosystem, can be computed at several levels (plotting the connections between multiple *online spaces* and even *communities*, as in Bovet and Grindrod's network of British far-right Telegram channels;[25] or analysing interactions between *individual users* within a single group/channel/forum as in Baele and colleagues' aforementioned study of "hyper-posters"),[26] and can be complemented by direct popularity measures (for example, number of visits to a website, average number of 'likes' and other reactions to an account or poster's interventions).

Regardless of these measurement choices, datasets tend to render similarly heavy-tailed distributions. Figure 3 below offers a good example, showing in-degree rankings for all nodes contained in two totally unrelated networks. The left panel displays data from a vast network representing the German far-right Telegram ecosystem (11,643 nodes, each corresponding to an individual channel) where the three most influential channels have in-degrees of 300, 274, and 269 respectively (at the other end, no less than 2,687 channels have an in-degree of 1, and 1,711 of 2). The right panel plots numbers from the much smaller network of French far-right websites (286 nodes, each corresponding to an extremist website) where the top three sites have in-degrees of 40, 35, and 35 respectively (the heavy tail comprises 52 sites having in-degrees of 1, and 21 of 2). Heavy-tailed distributions in online extremist networks mean that these are highly unequal ecosystems whereby a small number of major sites/individuals dominate large numbers of unpopular or less active sites/individuals. Open-source data from other researchers brings about comparable graphs representing a similarly hierarchical reality. For instance, plotting the numbers contained in Tables 6 and 7 from Bovet and Grindrod's abovementioned study results in heavy-tailed distributions for the most outlinked domains in UK far-right Telegram communities,[27] and Conway and colleagues recently claimed that digital platforms most used by terrorist groups at different points in time resemble power-laws.[28] This pattern is somewhat unsurprising given that power-law types of phenomena have been highlighted and studied with particular intensity within network science and digital network analysis;[29] Zhang and colleagues for example claim that power-law structures directly reflect 'universal behaviours in the evolution of real-world networks'.[30] In line with Lima-Mendez and van Helden's critique of such claims as over-simplifying 'myths' (see methods section above), we nonetheless see that the end of the heavy tails of the above distributions depart from the straight line characterising the other end – a sign that the mechanisms underpinning the emergence of these networks are probably not those behind power-laws and thus warrant further examination.

Figure 3. *Network influence (as approximated by nodes' in-degree) of online spaces in the German far-right Telegram (left, orange) and the French far-right websites (right, blue) ecosystems*

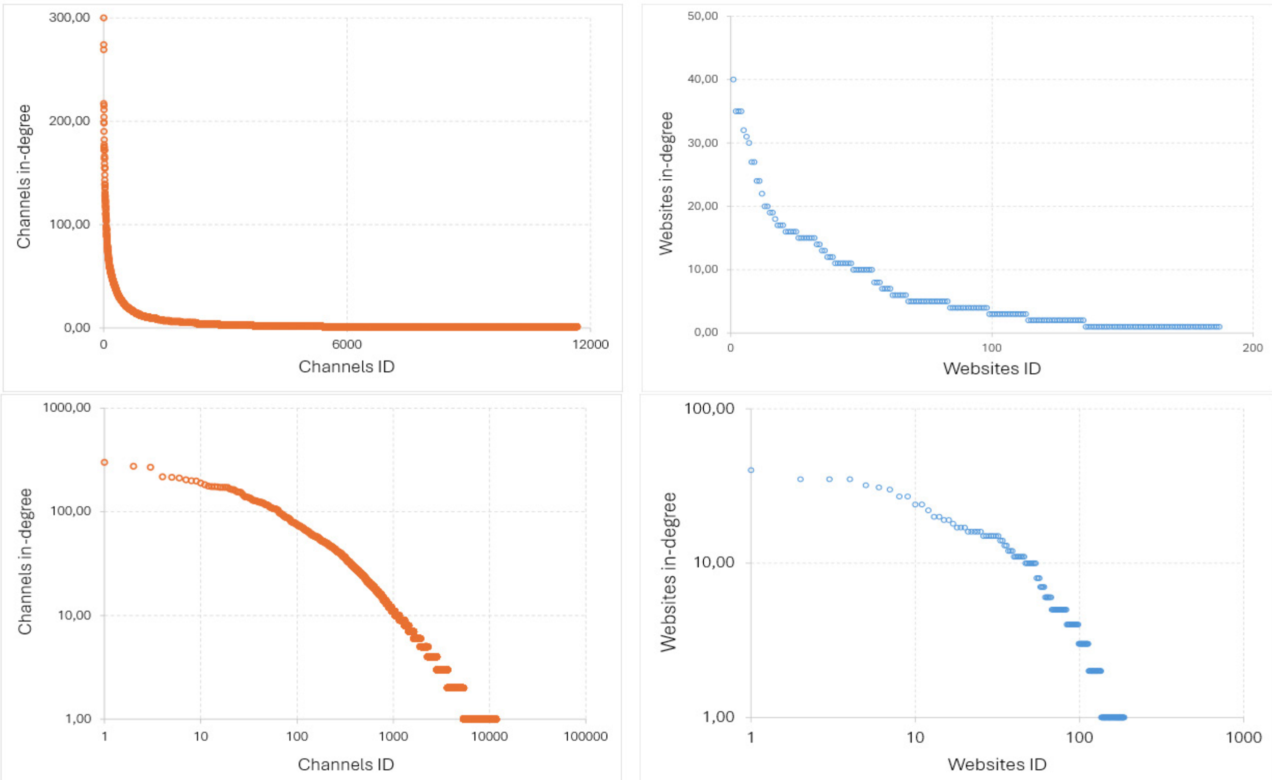

Power-law digital terrorism and extremism: So what?

Having provided examples of the ubiquity of heavy-tailed, power-law types of distributions in extremist and terrorist digital ecosystems, this concluding section opens up a discussion on the scientific and practical implications of this phenomenon. There are numerous factors involved in the dynamics producing these distributions, which makes their accurate understanding a complex task; however, staying at the descriptive level allows us to put forward three insights – which could guide further in-depth studies – that are already critical for CT/CVE practitioners.

First, power-law distributions often reflect "Yule" processes, colloquially called "rich-get-richer", whereby the bigger members of a growing system tend to attract increasingly more traction than their less initially endowed or latecomer counterparts.[31] For example, "the probability of a city gaining a new member is proportional to the number already there; the probability of a paper getting a new citation is proportional to the number it already has."[32] This 'preferential attachment' mechanism characterises much of the web, where more connected sites are more likely to obtain even more connections in the future,[33] and is thus highly likely to underpin some of the heavy-tailed distributions put forward above. Extremist sites that initially enjoyed some success or were set up earlier within an emerging system (for example, a new social media platform, an emerging ideology) are more likely to be found and engaged with by newcomers, extremist influencers who reach a critical mass of followers at one particular point in the development of the broader ideological ecosystem surrounding them are likely to get even more, etc. In their study of the far-right online ecosystem, Baele, Brace and Coan argue that this type of evolution process occurs 'organically' but can also be strategically fuelled by practices such as fake followers or artificial webs of multiple and interconnected accounts/sites tactically created by individuals who are keenly aware of the competitive nature of online extremism and eager to reach the threshold leading to snowballing influence.[34] This echoes Zipf's initial suggestion that power-laws are due to a fundamental law of 'biosocial' competition and cooperation in human societies, which results in highly unequal hierarchies.

Further research should therefore study the emergence and evolution of extremist ecosystems to test whether their development indeed presents Yule processes and truly follows a power-law logic. Recognising the reality of Yule processes in extremist ecosystems indeed offers pathways for clinical CVE interventions attuned to diachronic evolution dynamics, not only by pointing with accuracy (and as early as possible) to the most important and influential spaces or individuals who ought to be dealt with (with rippling effects on others), but also by offering direction (in terms of particular words to avoid or promote) for linguistic counter-narrative interventions aimed at changing discourse, themes, and hate levels within extremist spaces. Carthy and colleagues recently lamented that most empirical evaluations of counter-narratives effectiveness rested on weak evidence and pointed to small effects;[35] re-thinking counter-narratives programmes around statistical lexical data holds promise in designing more efficient messages. The aim of efficient CT/CVE operations cognizant of heavy-tailed distributions thus becomes to systematically target the rare yet ultra-important items, ideally *as they start to emerge as outliers*, for example by using power-law statistics as a forecasting element for identifying then stifling rising extremist accounts while monitoring the others to prevent them from reaching a critical threshold, or as a tool to guide more pertinent "inoculation" counter-narratives[36] based on increasingly popular discourses and words. These interventions have therefore a particularly important place in CT/CVE work on emerging ecosystems and new platforms, as they can help derail (for individuals or online spaces deemed problematic/dangerous), or promote (for those deemed more acceptable or less extreme) snowballing influence in early stages.

Second, some have argued that power-law distributions reflect efficiency equilibriums, i.e., systems where a few items (words, tasks, etc.) are used/done very often because they are comparatively more useful/important than most others, which are thus rarely used/done. Barabási for example argued that humans typically queue tasks based on perceived priority: typically, "most tasks [are] rapidly executed, whereas a few experience very long waiting times," resulting in a power-law distribution of tasks performed across time.[37] This logic applies to the words we use (almost never, for the vast majority) or the acquaintances we spend time with.

Further study should investigate whether such process underpins unequal activity levels in extremist spaces; here the hypothesis would be, in line with classic work on involvement in terrorism and extremism,[38] that only a small minority of individuals sustain a high level of involvement in these spaces because, contrarily to most people, they think the gains (be it in terms of social, symbolic, or economic capital, or in psychological, emotional, and well-being terms) offset the costs (on the same dimensions). The goal would be here to empirically find whether this kind of "rational choice" logic is the most important mechanism producing highly unequal patterns of participation in extremist digital milieux (as opposed to, or in complement to, for example, increasing degrees of radicalisation translating into growing involvement). A similar process is probably at play behind some of the linguistic distributions shown above. Since Rosch's work, cognitive science has indeed shown how basic-level categories dominate everyday language because they are the most efficient ones (not too specific, not too broad) to convey ideas.[39] Each milieu therefore has its own basic-level categories, which are the most useful ones geared at the tasks at hand (highly specialised words may well be basic-level categories in technical professional milieux such as hospitals or corporate financial law firms). The linguistic distributions put forward above may thus reflect the fact that the most frequent words and bigrams occur at such a high rate because they are the most useful, efficient, and direct ones to express, signal, and demonstrate belief in a given ideology or worldview (for example, "Donald Trump", "social media", "United States", and "white race" for the *8chan* board). Again, these observations indicate clearer and more pertinent pathways for CT/CVE interventions than intuition-based programmes. For example, regularly posting large volumes of incoherent 'garbage' into an extremist forum content is a power-law inspired intervention that would make it less rewarding for the vast majority of its (less active, or newer) members, thereby lowering participation or even cutting the 'heavy tail' altogether. Another type of power-law inspired intervention would aim at gradually shifting the platform's ideology, with multiple posts using moderately occurring words that are less extremist in themselves than the three or four most frequent ones (to keep the *8chan* example, trying to replace "white race" in the four most frequent terms with an open alternative like "human race", while keeping the three others).

Finally, some of these distributions might more worryingly reflect an exponentially growing extremist ecosystem – as mentioned above, exponentially growing networks observed ("killed") at one random point in time indeed display power-law distributions at one of their tails.[40] As such, the power-law appearance of data from an extremist ecosystem might evidence the exponential character of a particular stage of its development, with critical tipping points where stable systems become hugely more popular. Future research should study the chronology of particular extremist digital milieux with this hypothesis in mind, to locate exactly when – and if at all – an exponential growth (or stabilisation) happens. One such hypothesis, for alt-tech platforms, would be that the introduction of more restrictive terms of use by mainstream platforms or the takedown of a rival one triggered a moment of exponential growth. Such critical moments (for example, when *Parler* was shut down, or when the Islamic State's Telegram activity was hit by a sustained Europol-coordinated effort)[41] should concentrate CVE practitioners' efforts so as to prevent the growth of new ecosystems or the rising popularity of particular online spaces and communication channels. Attuning to heavy-tailed distribution is therefore not only a tool for designing multidimensional intervention procedures that consider indirect implications of particular actions (for example, banning/jailing an influencer from a social media platform, changing the terms of use of a platform, etc.), but also potentially a *forecasting* method.

In sum, appraising digital terrorism and extremism from the perspective of heavy-tailed, power-law types of distribution not only sheds light on important underlying structures and dynamics, but also indicates original avenues for more effective CVE interventions and even potential forecasting based on predictable statistical structures. Incidentally, it also leads us to replace the study of extremist and terrorist digital ecosystems within the broader social sciences, in the sense that most of the behaviours and developments found there are typical of all other realms of human activity. Extremism and terrorism studies tend to have a bias towards the search for highly unusual patterns, practices, 'markers' and indicators of radicalism, extremism and violence risk, whereas the key to a range of efficient counter-measures may well hold in precisely the opposite: non-specific, typical, commonly human patterns of "human behaviour" as Zipf would write.

Funding and acknowledgements

The data used in this work comes from three main funded projects: two CREST research awards (the "ConCel" project, ESRC ES/V002775/1; and the "MineChans" project, ESRC ES/N009614/1) and a NORFACE Joint Research Program "Democratic Governance in a Turbulent Age" consortium grant (the "ExID" project, EU Horizon 2020 grant agreement No. 822166). The author thanks Travis Coan, Olivier Sterck, Filippo Dionigi, Oscar Baele, and John Carbone for fruitful discussions on previous iterations of this paper, and is grateful to Lewys Brace for the data collection during the MineChans and ConCel projects and the initial discussions on power-law types of phenomena when writing the earlier "super-posters" article.

Stephane J. Baele is Professor of International Relations at the University of Louvain (Belgium), and Honorary Associate Professor of Security and Political Violence at the University of Exeter (UK). His multidisciplinary work, which includes ISIS Propaganda *(OUP) and a range of scientific articles, focuses on extremist and violent political actors' communications.*

Endnotes

1 Examples given include the words written in literary works such as *Beowulf*, Joyce's *Ulysses*, and Homer's *Iliad*; but also articles length in American newspapers, frequency of German stem forms and Chinese characters, or progress in children's acquisition of language. To clarify the notions of power-law and normal distributions to readers unfamiliar with statistics, I offer an illustrative example in **Appendix 1**.

2 George Zipf, *Human Behavior and the Principle of Least Effort* (Cambridge, MA: Addison-Wesley, 1949).

3 As Zipf himself noted, power-law distributions in economic systems (as, for instance, with income and wealth) are known since the end of the 19th Century with Pareto (power laws are therefore sometimes known as Pareto distributions)

4 William Reed and Barry Hughes, "From Gene Families and Genera to Incomes and Internet File Sizes: Why Power-Laws Are So Common in Nature." *Physical Review E* 66, no.6 Pt 2 (2002): 067103/1-067103/4.

5 Mark Newman, "Power Laws, Pareto Distributions and Zipf's Law." *Contemporary Physics* 46, no.5 (2005): 330.

6 Aniket Mahanti, Niklas Carlsson, Anirban Mahanti, Martin Arlitt, and Carey Williamson, "A Tale of the Tails: Power-Laws in Internet Measurements." *IEEE Network* 27, no.1 (2013): 59-64.

7 For example: on wars, see Lars-Erik Cederman, "Modeling the Size of Wars: From Billiard Balls to Sandpiles." *American Political Science Review* 97, no.1 (2003): 135-150; on terrorist attacks, see Aaron Clauset, Maxwell Young, and Kristian Skrede Gleditsch, "On the Frequency of Severe Terrorist Events." *Journal of Conflict Resolution* 51, no.1 (2007): 58-87; on insurgency escalation, see Neil Johnson, Spencer Carran, Joel Botner, Kyle Fontaine, Nathan Laxague, Philip Nuetzel, Jessica Turnley, and Brian Tivnan, "Pattern in Escalations in Insurgent and Terrorist Activity." *Science* 333, no.6038 (2011): 81-84.

8 See acknowledgements for information on the projects that yielded the data.

9 Reed and Hughes, "From Gene Families and Genera to Income and Internet File Sizes".

10 Mark Newman, "Power Law Distribution." *Significance* 14, no.4 (2017): 10-11.

11 Michael Stumpf and Mason Porter, "Critical Truths About Power Laws." *Science* 335 (2012): 665.

12 Aaron Clauset, Cosma Shalizi, and Mark Newman, "Power-law Distributions in Empirical Data." *SIAM Review* 51 (2009): 661-703. The authors provide instructions for implementing these tests on empirical data in several computer programmes and coding languages; see https://aaronclauset.github.io/powerlaws/.

13 Gipsi Lima-Mendez and Jacques van Helden, "The Powerful Law of the Power Law and Other Myths in Network Biology." *Molecular Biosystems* 5 (2019): 1485.

14 Xiaojun Zhang, Zheng He, Liwei Zhang, Lez Rayman-Bacchus, Yue Xiao, and Shuhui Shen, "The Truth About Power Laws: Theory and Reality." *arXiv:2105.10372* (2021).

15 Lima-Mendez and van Helden, "The Powerful Law", 1485.

16 Stumpf and Porter, "Critical Truths", 666.

17 Stumpf and Porter, "Critical Truths", 666.

18 By this we mean the constellation of digital spaces hosting misogynistic "incel" discussions linked to terrorist attacks such as Alek Minassian's in Toronto and Scott Beierle's in Tallahassee (both in 2018).

19 This is the chatroom where shooters like the 2019 Christchurch terrorist Brenton Tarrant interacted and announced their actions. On the role of this board on Tarrant's shooting and subsequent attacks, read Stephane Baele, Lewys Brace, and Travis Coan, "The 'Tarrant Effect': What Impact Did Far-right Attacks Have on the *8chan* Forum?" *Behavioral Sciences of Terrorism and Political Aggression* 15 (2020): 1-23.

20 Both sites are now defunct.

21 From, among others, Benjamin Ducol, "Uncovering the French-Speaking Jihadisphere: An exploratory Analysis", *Media, War & Conflict* 5, no.1 (2012): 51-70; Bennett Kleinberg, Isabelle van der Vegt, and Paul Gill, "The Temporal Evolution of a Far-Right Forum." *Journal of Computational Social Science* 4 (2021): 1-23; and Ryan Scrivens, Thomas Wojciechowski, Joshua Freilich, Steven Chermak, and Richard Frank, "Comparing the Online Posting Behaviors of Violent and Non-Violent Right-Wing Extremists". *Terrorism & Political Violence* 35, no.1 (2021): 192-209.

22 Stephane Baele, Lewys Brace, Travis Coan, and Elahe Naserian, "Super- (and Hyper-) Posters on Extremist Forums." *Journal of Policing, Intelligence & Counter Terrorism* 18, no.3 (2022): 243-281.

23 This pattern has nothing to do with extremism: it is also present in online spaces such as money-saving forums (see Todd Graham and Scott Wright, "Discursive Equality and Everyday Talk Online: The Impact of Superparticipants." *Journal of Computer-Mediated Communication* 19(2014): 625-642), political parties' discussion boards (Raphael Kies, *Promises and Limits of Web-Deliberation* (Basingstoke: Palgrave, 2010)), or bulletin boards (Xia-Meng Si and Yun Liu, "Power-Law Distribution of Human Behaviors in Internet Forums." *2010 International Symposium on Intelligence Information Processing and Trusted Computing*: 286-289). Beyond posting counts, other digital behaviors display heavy-tailed structures, such as contributions length (Pawel Sobkowicz, Mike Thelwall, Kevan Buckley, Georgios Paltoglou,

and Antoni Sobkowicz, “Lognormal Distributions of User Post Lengths in Internet Discussions?” *EPJ Data Science* 2, no.2 (2013): 1-20) or posting rhythm across time (Yukie Sano, Kenta Yamada, Hayafumi Watanabe, Hideki Takayasu, and Misako Takayasu, “Empirical Analysis of Collective Human Behavior for Extraordinary Events in the Blogosphere”. *Physical Review E* 87, no.012805 (2013): 1-8.

24 This forum was an important component of the “incelosophere” – see note 18 above.

25 Alexandre Bovet and Peter Grindrod, “Organization and Evolution of the UK Far-Right Network on Telegram.” *Applied Network Science* 7, no.76 (2022).

26 Baele, Brace, Coan and Naserian, “Super- (and Hyper-) Posters”.

27 Bovet and Grindrod, “Organization and Evolution of the UK Far-Right network on Telegram”.

28 Maura Conway, Sean McCafferty, and Stuart Macdonald, “Online Terrorist Content Hosting.” Presentation at *Terrorism & Social Media (TASM) Conference*, 19 June 2024.

29 For example Albert-László Barabási, “The Origin of Bursts and Heavy Tails in Human Dynamics.” *Nature* 435 (2005): 207-211; Albert-László Barabási and Réka Albert, “Emergence of Scaling in Random Networks.” *Science* 286 (1999): 509; Alex Fabrikant, Elias Koutsoupias, and Christos Papadimitriou, “Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet.” *Automata, Languages and Programming (ICALP)* 2002: 110-122.

30 Zhang et al., “The Truth About Power Laws”.

31 Fabrikant, Koutsoupias and Papadimitriou, “Heuristically Optimized Trade-Offs”.

32 Newman, “Power laws”.

33 Michael Mitzenmacher, “Editorial: The Future of Power Law Research.” *Internet Mathematics* 2, no.4 (2005): 525-534.

34 Stephane Baele, Lewys Brace, and Coan “Uncovering the Far-Right Online Ecosystem: An Analytical Framework and Research Agenda”. *Studies in Conflict & Terrorism* 46, no.9 (2020): 1599-1623.

35 Sarah Carthy, Colm Doody, Katie Cox, Denis O’Hara, and Kiran Sarma, “Counter-Narratives for the Prevention of Violent Radicalization: A Systematic review of Targeted Interventions.” *Campbell Systematic Reviews* (2020).

36 See Kurt Braddock, “Vaccinating Against Hate: Using Attitudinal Inoculation to Confer Resistance to Persuasion by Extremist Propaganda.” *Terrorism & Political Violence* 34, no.2 (2022): 240-262.

37 Barabási, “The Origin of Bursts and Heavy Tails”.

38 See for example Sophia Moskalenko and Clark McCauley, “Measuring Political Mobilization: The Distinction Between Activism and Radicalism.” *Terrorism & Political Violence* 21, no.2 (2009): 239-260; James Khalil and Lorne Dawson, “Understanding Involvement in Terrorism and Violent Extremism: Theoretical Integration Through the ABC Model.” *Terrorism & Political Violence* 36, no.6 (2024): 757-772; or Arie Kruglanski, Ewa Szumowska, Catalina Kopetz, Robert Vallerand, and Antonio Pierro, “On the Psychology of Extremism: How Motivational Imbalance Breeds Intemperance.” *Psychological Review* 128, no.2 (2021): 264-289.

39 See for example, Eleanor Rosch, Carolyn Mervis, Wayne Gray, David Johnson, and Penny Boyes-Braem, “Basic Objects in Natural Categories.” *Cognitive Psychology* 8, no.3 (1976): 382-439.

40 Reed and Hughes, “From Gene Families and Genera to Income and Internet File Sizes”.

41 I take these particular examples because they have already been commented by experts. See for example , Saharsh Agarwal, Uttara Ananthakrishnan, and Catherine Tucker, “Content Moderation at the Infrastructure Layer.” *SSRN Preprint*, https://ssrn.com/abstract=4232871; Amarnath Amarasingam, Shiraz Maher, and Charlie Winter, “How Telegram Disruption Impacts Jihadist Platform Migration.” *CREST Reports* (2021), https://www.crestresearch.ac.uk/resources/how-telegram-disruption-impacts-jihadist-platform-migration/.

Appendix

Among the distributions often (imperfectly) found in when plotting social data, one is regularly claimed to be particularly frequent: the normal, or Gaussian, "bell-shaped" distribution (as in standardized IQ testing for instance). Our Research Note emphasises another one: the heavy-tailed power-law distribution. To illustrate the differences between these two types of distributions to those not familiar with statistics, we present below two variants of a situation which, whilst hypothetical, will be familiar to many readers.

Imagine a university programme (say, a BA in terrorism studies) at a very large institution. This programme has a yearly enrolment cohort of 1,000 students. At the end of his/her first year, each of these students is given a mark between 0 and 100 (with 0 being the worst mark and 100 a perfect mark). Students ending their first year in 2024 have been ranked in Figure 1 below: we can see that the distribution of how many students obtain each mark resembles a bell and is therefore close to a "normal" distribution. Very few students obtain either very low or very high marks, whilst large numbers of students obtain marks around the high 50s / low 60s. Students ending their first year in 2025 have been ranked in Figure 2 below, and their performances present a very different distribution indeed, close to a power-law. It is characterized by a long tail made of very few students populating a wide marking range (from 0 to around 70) followed by increasing numbers of students as marks get higher, with the last two data points standing out (the opposite would have been possible too, with one or two students getting extremely high marks and the vast majority getting really poor grades).

We can quickly see that these two distributions evidence two starkly different realities – an average 2024 first-year cohort versus a brilliant 2025 one. But we also understand that evidencing these two realities should enable the programme director to design very different pedagogies and support strategies to best teach his/her students. The argument I make in the Research Note is similar: by documenting power-law types of distributions in the field of online terrorism and extremism, practitioners can design sound, data-driven counter-extremism/terrorism strategies that would have been different had they faced normal or other forms of distributions.

Figure 1. Students' marking ranks resembling a normal distribution

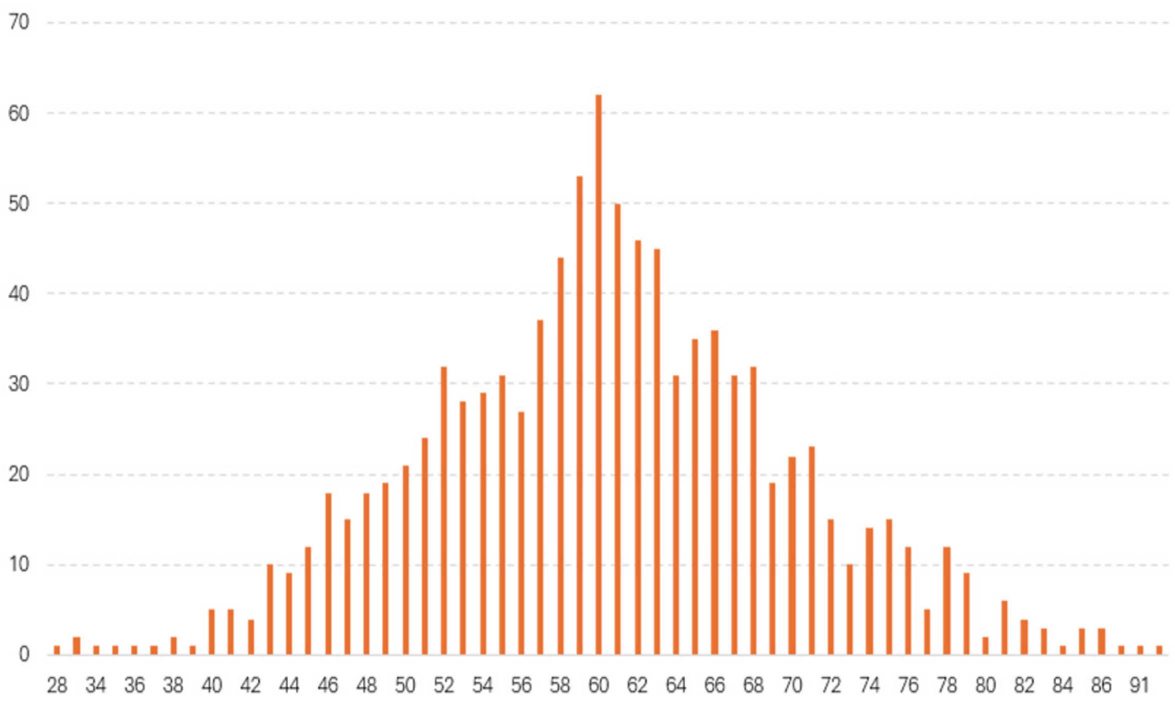

Figure 2. Students' marking ranks resembling a power-law distribution

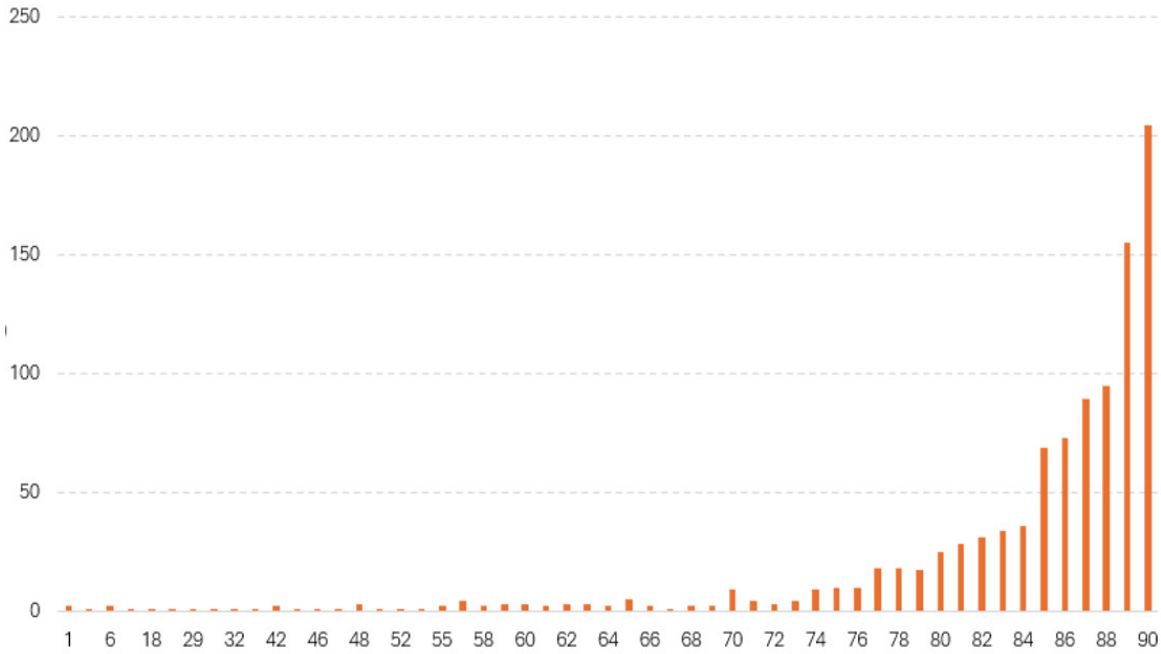

About

Perspectives on Terrorism

Established in 2007, *Perspectives on Terrorism* (PT) is a quarterly, peer-reviewed, and open-access academic journal. PT is a publication of the International Centre for Counter-Terrorism (ICCT), in partnership with the Institute of Security and Global Affairs (ISGA) at Leiden University, and the Handa Centre for the Study of Terrorism and Political Violence (CSTPV) at the University of St Andrews.

Copyright and Licensing

Contact

E: pt.editor@icct.nl
W: pt.icct.nl