

## WHO SHOULD WE FEAR MORE: BIOHACKERS, DISGRUNTLED POSTDOCS, OR BAD GOVERNMENTS? A SIMPLE RISK CHAIN MODEL OF BIORISK

---

Anders Sandberg and Cassidy Nelson

---

The biological risk landscape continues to evolve as developments in synthetic biology and biotechnology offer increasingly powerful tools to a widening pool of actors, including those who may consider carrying out a deliberate biological attack. However, it remains unclear whether it is the relatively large numbers of low-resourced actors or the small handful of high-powered actors who pose a greater biosecurity risk. To answer this question, this paper introduces a simple risk chain model of biorisk, from actor intent to a biological event, where the actor can successfully pass through each of  $N$  steps. Assuming that actor success probability at each independent step is sigmoidally distributed and actor power follows a power-law distribution, if a biorisk event were to occur, this model shows that the expected perpetrator would likely be highly powered, despite lower-powered actors being far more numerous. However, as the number of necessary steps leading to a biological release scenario decreases, lower-powered actors can quickly overtake more powerful actors as the likely source of a given event. If steps in the risk chain are of unequal difficulty, this model shows that actors are primarily limited by the most difficult step. These results have implications for biosecurity risk assessment and health security strengthening initiatives and highlight the need to consider actor power and ensure that the steps leading to a biorisk event are sufficiently difficult and not easily bypassed.

**Keywords:** Public health preparedness/response, Risk assessment, Biotechnology, Biological event

THE RISE OF SYNTHETIC BIOLOGY and biotechnology promises many benefits to the world: improved medicine, greener farming and industry, enhanced production of biofuels and nanomaterials, and new organisms with designer functions. Universities and large corporations are not the only ones exploring opportunities made possible by the growing bioeconomy; other active participants include a vital community of “biohackers” and small to medium enterprises. Despite positive advancements being sought by these actors, any technology able to significantly change the world also has the potential to be misused. It is, therefore, crucial to find ways to systematically assess the risks associated with the expansion and democratization of biotechnology, particularly as it relates to the field of health security.

Synthetic biology is “small tech” that reduces reliance on large expensive facilities and easily tracked resources. Much of its power comes from sharing information. Once a genetic sequence or methodology has been published online, it is nearly impossible to control its further dissemination and use. The risks of synthetic biology are multifold: living organisms are self-replicating and can be robust and invasive and new tools such as gene drives have the ability to potentially affect entire species in the wild. Past pandemics demonstrate that biological threats can kill tens of millions of people, and an engineered pathogen pandemic could pose an even greater threat than what arises in nature.<sup>1</sup> Biowarfare and bioterrorism have demonstrated that biotechnology can be used maliciously, and if a contagious

---

Anders Sandberg, PhD, is a Senior Research Fellow; Cassidy Nelson, MBBS, MPH, is a Research Scholar; both are at the Future of Humanity Institute, University of Oxford, Oxford, UK.

agent was deliberately released, it would be hard to accurately control or contain. Even nonmalicious intentions can cause damage, such as ecological devastation from an introduced invasive species or the nonzero risk of accidental escape of gain-of-function pathogens such as influenza from a research laboratory.<sup>2</sup>

### *The Problem of Biorisk*

These concerns have led to questions about the safety of the proliferation of biotechnology. As this technology becomes more widespread and more capable, the greater the likelihood it could end up in the hands of dangerous actors. At the same time, others claim many fears about bioterrorism are overblown, citing the role tacit knowledge and institutions play in preventing or enabling actors to succeed at weaponizing biology.<sup>3</sup> In 2015, when experts were asked who would be the most likely actor behind a biological weapon attack if one were to occur in the next 10 years, they offered a wide range of opinions on the risks posed by both state and nonstate actors.<sup>4</sup> This lack of consensus underscores the persistent uncertainty about from where the majority of biorisk arises.

This paper seeks to analyze one aspect of this process: should we be most concerned about the few large, powerful actors with significant resources, or the more numerous small actors? In particular, the first time a sufficiently large biorisk event occurs, will it come from a relatively low-skilled, low-resourced actor (eg, biohacker), a better skilled but potentially low-resourced actor (eg, a disgruntled postdoc), or an actor with potentially substantial skills and resources (eg, a government program)?

## METHODS AND RESULTS

### *The Risk Chain*

A simple model of biorisk is the risk chain is shown in Figure 1. Given an actor has a bad intention (which could be irresponsible, misguided, or malicious), they may consider using biology to cause harm. This motivates them to develop or capitalize on a biological idea, which is then converted into biological data, such as a pathogen genome.

These data are transformed into a live biological artefact as it is processed through a biosystem, cultured, and tested for effect, before successfully being dispersed into the target environment. The consequences of this work ensue.

The availability of new technologies makes transitions between steps easier or enables some to be bypassed altogether. Contemporary examples of this for each step are shown in the risk chain (Figure 1). Powerful actors, such as nation states, have larger conceptual toolboxes than weak actors, making them more likely to know that biology could be weaponized. Developments in the field of genomics and breakthroughs in DNA sequencing make converting an idea into a biological data, such as a genetic sequence, easier. Standardized laboratory parts and advancements in synthetic biology can streamline converting biological data into a live biological artefact, such as a viral pathogen. Improved understanding and parameterization of biological processes during their design make successful biosystem use more feasible, reproducible, and reliable. Improvements in lab automation and *in silico* computer simulations for testing pathogenicity may reduce or completely eliminate the need for culturing or animal model testing altogether. The spread of information hazards, such as detailed instructions about bioweapon dispersal techniques, could reduce the difficulty in dissemination of a constructed bioweapon.

### *A Simple Model of Biorisk*

Assume a risk chain containing  $N$  steps. In order to cause a disaster, an actor needs to succeed with all  $N$  steps. Skilled or well-equipped actors are more powerful because they have a higher probability of succeeding than actors with fewer resources. We can model this by assuming a success probability that is a logistic or sigmoid function of actor “power”  $x$  as follows:

$$p(x) = 0.5 + 0.5 \tanh(g(x - x_0))$$

where  $x_0$  represents the difficulty of the task and  $g$  represents the sharpness of the transition, with a higher value indicating that the probability quickly shifts from near zero to near one as the actor passes  $x_0$ . This function provides a robust basis for the model, given that we can reasonably




Figure 1. Biorisk chain model. Starting with an actor with an insight into the nefarious misuse of biology (left), progression is required through each step for a biorisk event with recognized consequences to occur (far right).

assume agent power will increase probability of success, with marginal returns greater for weak actors and diminished for powerful actors. The exact form of the curve does not matter significantly for the model.

If each step occurs independently and with the same difficulty, the probability of complete success is given by:

$$P(\text{success}|x) = p(x)^N$$

This is essentially just a sharper sigmoid with inflection point higher than  $x_0$ , as demonstrated in Figure 2. The location of the inflection point is given by  $1/2 = p(x)^N$ , leading to an explicit (if unwieldy) expression:

$$x = x_0 + (1/g) \tanh^{-1} \left( 2^{1-1/N} - 1 \right)$$

For this expression, it is evident that as the number of steps,  $N$ , increases the power,  $x$  grows logarithmically.

Let the distribution of actor powers be  $f(x)$ . The probability that a successful actor has power  $x$  is given by the following relationship:

$$\begin{aligned} P(X=x|\text{success}) &= \frac{P(\text{success}|X=x)P(X=x)}{P(\text{success})} \\ &= Kp(x)^N f(x) \end{aligned}$$

where  $K$  is a normalization factor. The overall effect is a cutoff that filters out actors that are weak, as they have little chance of success.

There are fewer actors with great power and many more actors with low power. A simple model, motivated by the

typical distribution skews of organization sizes, wealth, casualties in war, frequency of terrorist events, and a wide range of other empirical phenomena, is a power-law distribution,<sup>5,6</sup> given by:

$$P(x \leq X \leq x + dx) = Cx^{-\alpha} dx$$

where  $\alpha > 1$  determines the skewness. Low values of  $\alpha$  have a broad distribution with proportionally more powerful actors than distributions with high values of  $\alpha$ . In this case, the expected distribution of successful actors will be  $Kp(x)^N x^{-\alpha}$ .

### An Example with 6 Steps

To provide a concrete example, let us set  $N=6$ ,  $\alpha=2$ ,  $x_0=10$ , and  $g=0.1$ . The choice of  $x_0$  is arbitrary since power-laws are scale free. The choice of  $g$ , which indicates actors are 1 order of magnitude weaker than the threshold power still have some chance of success, is motivated by demonstrated projects such as OpenPCR (Chai Inc) and RepRap that produce tools substantially cheaper than typical commercial tools.<sup>7</sup> This example produces a distribution of successful actor powers as shown in Figure 3.

If a biorisk event occurs, the expected power of the actor causing it in this scenario will be large, above 200. Because of the  $x^{-2}$  heavy tail, the expectation diverges at very high power; in practice, an upper cutoff occurs at some point due to limitations to realistic power acquisition. This paper is not concerned with extreme tail behavior for power, as it does not change the conclusions of the model—these issues will be ignored henceforth.




Figure 2. Blue: Probability of success when attempting a step of difficulty 10 for varying levels of actor power based on the sigmoid model as a function of actor power. Red: Probability of success passing through a 6-step task chain. Green: Threshold difficulty.




Figure 3. The probability distribution of the power of actors successful in passing all steps in an  $N=6$  risk chain.

The median successful actor power is far smaller, 38. This is due to the heavy tail of actors: the presence of a few very powerful actors tends to lead bias estimates upward. However, even the median is notably larger (3.8 times) than the task difficulty  $x_0$ ; this suggests the number of steps in the risk chain favors more powerful actors since actors just at the threshold of ability are likely to fail at least once. To be dangerous, an actor needs to be able to complete each step with a high probability. This has important implications for biosecurity risk mitigation strategies, as it implies that it may be effective to make just 1 step in the chain very difficult for sufficient risk reduction.

Changes to the sharpness of the transition  $g$  have a simple effect: high values make the threshold effect sharper, leading to an essentially truncated  $f(x)$  distribution. Low values reduce the threshold and make the number of actors the salient factor. Changing  $f(x)$  to a lognormal distribution or modifying  $\alpha$  has little qualitative effect: the heavy tail of a few powerful actors and a large weak group, not the details of the curve, gives this model its behavior. In fact, light-tailed exponential distributions of actor power also give the same results.

### Chain Length

If the number of steps in the chain is reduced, weak actors have higher chances of success and, since there are many more of them, this can outweigh their low individual success rates. As demonstrated in Figure 4, as the number of steps  $N$  declines, the distribution becomes bimodal with a growing peak close to zero power and another above  $x_0$ ; eventually, the low-power peak dominates.

Hence, if the risk chain is short, we should expect most biorisk from the numerous low-resource amateurs rather than very skilled and well-resourced major groups. Obviously, a short risk chain also increases the baseline probability of an eventual biorisk. Because of the logarithmic dependence of the effective cutoff on  $N$ , early reductions in step number have a less overall effect than later ones. This could lead to a false sense of complacency as technology advances at first do not appreciably enable low-power actors to do dangerous things and, therefore, a “warning shot” event is unlikely to occur, followed by an accelerated increase in risk from just a few more improvement steps compared to the past.

As the steps in a risk chain increase, both the median and mean actor power rise as demonstrated in Figure 5. However, because the distribution skew also grows as the risk chain length increases, the median and mean actor power diverge with greater growth to the latter for each increase in the  $N$  step length.

### Unequal Step Difficulties

What if the difficulties of the steps differ from each other? In this case, the  $p(x)^N$  term is replaced by a product of individual difficulties:  $\prod_i p(x, x_i)$  where  $x_i$  is the difficulty of step  $i$ . It turns out that the overall effect is that the hardest step dominates: actors below the  $\max(x_i)$  have little chance of success unless the spread of difficulties is narrow. For example, if we assume a log-uniform distribution of difficulties in some range  $[a, b]$  then the maximum difficulty will be distributed as a log-beta




Figure 4. The probability distribution of successful actor powers for risk chains of different lengths. The blue curve corresponds to  $N=6$  as seen in Figure 3.

distribution with  $\alpha = N$ ,  $\beta = 1$ : it will tend to approach the upper limit as  $N$  increases.

The probability of success for actors of different power with increasing step difficulty is shown in Figure 6. Note that difficulty 0.1 and 0.5 are easier than the other steps; the light orange (difficult step = 10) represents the equal case and the others are harder steps. Difficulty pushes up the mean and median actors and drives down their overall

probability of success. However, the smaller curves here are due to the fact that the x-scale is logarithmic, so the apparent areas under the curve are not the same. Making 1 step easier in a long risk chain does not have a big effect, except by reducing the effective length of the chain; reducing step difficulty has a larger effect in shorter chains.

The overall effect is to render the easier steps irrelevant compared to the hardest step, reducing the effective length




Figure 5. For a successful actor, mean (orange) and median (blue) power as a function of risk chain length.






Figure 6. Probability of success for actors of different power given steps of different difficulty.

of the chain. The cutoff at  $\max(x_i)$  becomes softer and smaller than the  $p(x)^N$  cutoff, implying the expected and especially median power of a successful actor will be less than in the equal difficulty case where  $x_0 = \max(x_i)$ .

It is worth noting that this is equivalent to if actor power is regarded as multivariate (eg, an actor may be skilled at DNA synthesis but less good at dispersal) but correlated. In this case, actors can then be classified by their weakest skill, and if this is below the difficulty of the task, they are then considered unlikely to be a threat.

## DISCUSSION

This simple risk chain model suggests that the first actor to cause a biorisk event of note is likely to be several times more skilled or resourced than what is needed to actually perform the operation but will originate from the slightly larger pool of powerful actors rather than the highest-powered elite.

To this end, we may be rationally more worried in the near future about a disgruntled highly skilled researcher at a well-equipped university or corporate laboratory, of which there are many, than about biohackers (too weak) or nation state bioweapon programs (too few). The domain of concern should be larger-populated groups that can pass through the entire risk-chain without extreme effort.

This is an important insight for policy prioritization. It highlights that we should be more concerned about an “insider” threat, as an actor will have to spend considerable time inside our scientific institutions to gain the necessary “power.” However, a notable exception for “outsider” threats is where a powerful actor lends the necessary power to other actors who would otherwise fail to realize their goal, such as through theft of created

pathogens, transfer of technology, or release of information hazards<sup>8</sup> that reduces the number of steps or their difficulty in the risk chain. This suggests we also may not be currently paying sufficient attention to threats posed by state-sponsored actors, given they may impose both direct and indirect risks through enabling weaker actors.

Furthermore, we generally consider measures taken to combat biosecurity incidents as being sufficient in most professional settings to reduce the risk from misguided low-powered actors. For example, greater focus has been recently placed on the need to strengthen oversight of biohackers to reduce careless or malicious activities.<sup>9</sup> The findings in this paper suggest that resources would be better allocated to further strengthening outreach on potential risks in more traditional settings, with a particular focus on ensuring individual skilled scientists and engineers are familiar with national and institutional biosecurity policies and risk management practices.

While efforts to address nefarious actors often focus on biohackers,<sup>10</sup> notable resources are also invested in dealing with the threats posed by states through international instruments such as the Biological Weapons Convention and the 1925 Geneva Protocol. When international efforts, such as United Nations Security Council Resolution 1540, do focus on biothreats from nonstate actors, they often use a top-down approach, potentially limiting their impact on the types of actors that pose the greatest risk. Greater focus and resources should be invested in the work done by professional scientific and engineering communities to raise awareness about, and implement measures to manage, the threat posed by the deliberate misuse of biology.

This paper shows that as technology lowers the step threshold, the number of actors who can cause harm increases. However, a more pertinent problem may be the simplification of procedures that allow actors to reduce the

difficulty of steps in the risk chain or eliminate 1 or more steps entirely by, for example, copying proven mechanisms developed by others in online methodologies or through streamlined automation of laboratory procedures. Technology transfer from advanced groups can enable weaker actors, as has been previously seen with the development and subsequent proliferation of nuclear weapons.<sup>11–13</sup> The potential for the difficulty threshold to be lowered is compounded by information hazards in biotechnology, of which there are many published examples.<sup>14</sup>

These issues not only increase the overall risk but also increase the likelihood that small groups or individuals—who are inherently harder to monitor—can trigger biorisk events. It might be tempting to overestimate the ease with which this happens, given the tacit knowledge still required,<sup>3</sup> but it would be unwise to assume that the current situation and capability levels required will still exist in the coming decades, given the rapid pace of biotechnological innovation. Commercial and academic drivers focus time and resources specifically on changing the current difficulty status of steps that bottleneck synthetic biology developments.

At first glance, this suggests there is currently a window to refocus efforts on preventing deliberate misuse to encompass a wider range of actors. This opportunity, while encouraging, will likely prove challenging, as scaling current safety and security arrangements to fully cover such an expanded range of relevant actors (in terms of numbers, types, and distribution) may be difficult. Even more significant impacts may occur as lowered thresholds could challenge the very regimes we use. For example, current technology control relies on a very limited number of actors being able to develop key capabilities or having no ready way to bypass certain steps. This assumption may not hold true in the longer term as technological developments evolve and knowledge transfer continue. This means we should use the current window to develop safety and security approaches tailored specifically for a more distributed and democratized biology.

Not all of our efforts should be focused on the more numerous, medium-power actors. In the past, the deadliest human-caused events have been wars, democides, and famines caused by states or other large groups. The deadliest biological weapon attacks have historically been governmentally organized. For example, the Japanese military's use of biological weapons in Manchuria during WWII is believed to have led to several hundred thousand plague deaths,<sup>15</sup> and the spread of smallpox in the new world, which killed millions, was deliberately deployed for military reasons in several instances.<sup>16,17</sup> In comparison, the largest biological terrorist attack to date, the Bhagwan Shree Rajneesh attack in Oregon in 1984, caused 751 injuries and resulted in no deaths.<sup>18</sup>

Actor power may have little to do with the relative effects of a biorisk compared to their ability to create one, since once a biological agent is released it may propagate, repli-

cate, and mutate on its own in unpredictable ways. The number of historical chemical/bioweapon fatalities looks like they could have been power-law distributed with a fairly low  $\alpha$ , indicating their power distribution is less skewed, although the data is very limited.<sup>6</sup> Nevertheless, powerful actors are more likely to be able to access, test, optimize, disperse, and integrate pathogens into their doctrines successfully and in an undisturbed fashion. This adds extra weight to monitoring larger groups despite their relative rarity and lower probability of being originators of biorisks.

It should be noted that our simple model formalism is agnostic on the valence of the consequences. It could just as well be used as a model of actors trying to achieve a good outcome in some domain. Lowered thresholds and fewer steps mean that more groups have a higher chance of succeeding, and that the relatively rare powerful groups may become less important for overall progress—they would instead have a reason to focus on far harder problems where they have the decisive comparative advantage.

This model has various limitations, and further research could help estimate the number of crucial steps and distribution size of different malicious actors through a risk assessment process. Estimates of the difficulties of performing different biorisk events—releasing a known pathogen, constructing one *de novo*, gain-of-function enhancement of a pathogen followed by accidental release, ecological sabotage using gene drives, and so on—might be useful for calibrating the model but could pose an information hazard if published, making actors aware of the step(s) on which to concentrate. However, broadly understanding the difficulty spectrum would also enable health security initiatives to focus on the groups most likely to produce certain risks. Future studies could examine technology transfer in more detail, in particular the role practical knowledge (*metis*) and technical knowledge (*techne*) play in enabling lower-powered actors.

One of the advantages of this model is that the risk chain could also be turned into a risk tree or risk network to handle more complex scenarios, such as different types of biorisk events or ranges of adversaries able to detect and mitigate some activities. In particular, the formal risk assessment is best treated as a mixed model of several such scenarios. Future research could then consider a portfolio optimization approach that considers risk mitigation across multiple options and possibilities.

## CONCLUSION

It is in the interest of all nations and the international community as a whole to prevent deliberate biological attacks from occurring. As the fields of synthetic biology and biotechnology accelerate toward new understanding, capabilities, and wider accessibility, assessment and understanding of the steps required in the chain of actions leading

to deliberate events is crucial for health security preparedness. Risk mitigation activities should consider the relative abundance and power of malicious actors and prioritize preventing steps in the risk chain from being eliminated or easily bypassed. Ensuring that the actions required to carry out an attack remain sufficiently difficult will help safeguard humanity against the threats posed by the deliberate misuse of biology.

## ACKNOWLEDGMENTS

We would like to thank Lord Martin Rees, who inspired this model during a Clinical Sequencing Evidence-Generating consortium meeting in Berlin 2014. We also thank Luc Henry and Thomas Landrain for their stimulating discussion about the views of the do-it-yourself community on biorisk. We also thank Piers Millett who reviewed and commented on an early draft of this manuscript.

## CONFLICTS OF INTEREST AND FUNDING

The authors declare no competing conflicts of interest. For this research and manuscript preparation, the authors received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No 669751).

## REFERENCES

1. Schoch-Spana M, Cicero A, Adalja A, et al. Global catastrophic biological risks: toward a working definition. *Health Secur.* 2017;15(4):323-328.
2. Merler S, Ajelli M, Fumanelli L, Vespignani A. Containing the accidental laboratory escape of potential pandemic influenza viruses. *BMC Med.* 2013;11:252.
3. Jefferson C, Lentzos F, Marris C. Synthetic biology and biosecurity: challenging the "myths." *Front Public Health.* 2014;2:115.
4. Boddie C, Watson M, Ackerman G, Gronvall GK. Assessing the bioweapons threat. *Science.* 2015;349(6250):792-793.
5. Clauset A, Shalizi CR, Newman MEJ. Power-law distributions in empirical data. *SIAM Rev.* 2009;51(4):661-703.
6. Clauset A, Young M, Gleditsch KS. On the frequency of severe terrorist events. *J Conflict Resol.* 2007;51(1):58-88.
7. Pearce JM. Building research equipment with free, open-source hardware. *Science.* 2012;337(6100):1303-1304.
8. Esvelt KM. Inoculating science against potential pandemics and information hazards. *PLoS Pathog.* 2018;14(10):e1007286.
9. Blazeski G. The need for government oversight over do-it-yourself biohacking, the wild west of synthetic biology. *Law School Student Scholarship.* 2014;411. Accessed May 29, 2020. [https://scholarship.shu.edu/student\\_scholarship/411](https://scholarship.shu.edu/student_scholarship/411)
10. Ahteensuu M. Synthetic biology, genome editing, and the risk of bioterrorism. *Sci Eng Ethics.* 2017;23:1541-1561.
11. Autio E, Laamanen T. Measurement and evaluation of technology transfer: review of technology transfer mechanisms and indicators. *Int J Technol Manage.* 1995;10(7-8):643-664.
12. Kroenig M. *Exporting the Bomb: Technology Transfer and the Spread of Nuclear Weapons.* Ithaca, NY: Cornell University Press; 2010.
13. Lui Y, Lui J. Analysis of Soviet technology transfer in the development of China's nuclear weapons. *Comp Technol Transfer Soc.* 2009;7(1):66-110.
14. Lewis G, Millett P, Sandberg A, Snyder-Beattie A, Gronvall G. Information hazards in biotechnology. *Risk Anal.* 2019;39(5):975-981.
15. Harris S. Japanese biological warfare experiments and other atrocities in Manchuria, 1932-1945, and the subsequent United States cover up: a preliminary assessment. *Crime, Law Soc Change.* 1991;15(3):171-199.
16. Stearn EW, Stearn AE. *The Effect of Smallpox on the Destiny of the Amerindian.* Boston, MA: Bruce Humphries; 1945.
17. Henderson DA, Inglesby TV, Bartlett TV, et al. Smallpox as a biological weapon: medical and public health management. *JAMA.* 1999;281(22):2127-2137.
18. Crupi RS, Asnis DS, Lee CC, Santucci T, Marino MJ, Flanz BJ. Meeting the challenge of bioterrorism: lessons learned from West Nile virus and anthrax. *American J Emerg Med.* 2003;21(1):77-79.

Manuscript received September 20, 2019;  
revision returned March 3, 2020;  
accepted for publication April 29, 2020.

Address correspondence to:  
Cassidy Nelson  
Research Scholar  
Future of Humanity Institute  
University of Oxford  
16-17 St. Ebbe's Street  
Oxford OX1 1PT, UK

Email: cassidy.nelson@gtc.ox.ac.uk

(Appendix follows →)



## APPENDIX

### APPENDIX: ACTOR GROUP SIZES

There exists economic literature about business sizes, studying the stylized fact of why they are lognormal or power-law distributed.<sup>1,2</sup> In the latter paper, the authors found an exponent  $\alpha \approx 2$  for US firms. However, for smaller groups the fit might be more lognormal or even shift to normal and uniform on the smallest end. Similarly, it is well known that individual wealth, especially in the extreme tail, has a power-law form.

National power (as measured by gross domestic product, military budget, research budget, population, etc.) is also clearly a highly skewed distribution. However, it might not be a full power-law.<sup>3</sup>

Scientific team sizes appear to be a mixture of Poisson-distributed core teams and a power-law tail of very large extended teams. The typical team is small (6.7 people in astronomy), but the power-law tail has become more extreme over time as large-scale collaborations have become more common.<sup>4</sup>

Estimates for terrorist groups are very hard to come by. Looking at the Big, Allied, and Dangerous dataset<sup>5</sup> gives a size distribution of terrorist organisation as seen in Figure A1. It should be noted that the lowest category also contains cases where size is entirely unknown. In any case, this data seems to suggest that a heavy tail distribution of group sizes is possible.

For a joint risk model, these different distributions could be combined. However, an accurate weighing (or even definition) of “power” may not be possible since the different types of group may have very different risk profiles: terrorist groups are motivated to cause deliberate harm, research teams

and corporations may house malicious individuals using their resources for deliberate harm but also contribute through accidental risk, governments produce biorisks through bio-warfare research and accidental risks. To weigh together their contributions to overall biorisk more quantification of both the likelihood of malintent, what pathways of biorisk are available, and how size/wealth maps onto problem-solving capacity. This is an interesting, if major, research project.

For the purposes of this paper, it is enough to conclude that for nearly any kind of actor of concern and definition of what capacities are needed to pass through the risk chain the distribution of actors will be highly skewed.

### REFERENCES

1. Simon HA, Bonini CP. The size distribution of business firms. *Am Econ Rev.* 1958;48(4):607-617.
2. Axtell RL. Zipf distribution of US firm sizes. *Science.* 2001; 293(5536):1818-1820.
3. Skipper RK. Zipf's Law and its correlation to the GDP of nations. *The University of Maryland McNair Scholars Undergraduate Research Journal.* 2011;3:217-226.
4. Milojević S. Principles of scientific research team formation and evolution. *Proc Natl Acad Sci.* 2014;111(11):3984-3989.
5. Asal V, Rethemeyer RK, Anderson I. Big allied and dangerous (BAAD) database 1 – lethality data, 1998-2005. V3, UNF:5:2Z77QCNIImKUu2OVS6hqccw== [fileUNF]. Harvard Data Universe website. <https://doi.org/10.7910/DVN/GPEUFH>




Figure A1. Fraction of terrorist groups in different size intervals, based on Asal et al.<sup>5</sup>

**This article has been cited by:**

1. John T. O'Brien, Cassidy Nelson. 2020. Assessing the Risks Posed by the Convergence of Artificial Intelligence and Biotechnology. *Health Security* **18**:3, 219-227. [[Abstract](#)] [[Full Text](#)] [[PDF](#)] [[PDF Plus](#)]