1

Sandpiles, Long Tails, and Complex Catastrophes

Terrorist attacks, floods, nuclear power meltdowns, economic collapses, political disruptions, and natural events like earthquakes come and go in waves – they are bursty – arriving sporadically with uneven consequences that are "long-tailed" rather than purely random. We ask why?

The modern world differs from the past mainly due to vast network connections that replace independent parts acting independently. This applies to all kinds of systems – both physical and socio-economic. Today, most systems are highly dependent on one another. The actions of one system affect other systems increasing the probability of a catastrophic cascade event conditioned on a tripping event. That is, the probability of an incident is no longer independent and represented by the normal distribution – it is highly dependent on other events. This abolishes the normal distribution as an explanation for catastrophes replacing it with a long-tailed *power law distribution*.

For example, a tsunami far across the Pacific created by an earthquake under the ocean northeast of the Japanese islands swept over the northeast coast of Japan, overwhelming the retaining wall protecting the nuclear reactor at Fukushima daiichi [reactor number one in Fukushima Prefecture]. A chain reaction that began in the Pacific Ocean propagated disaster through miles of Blue Ocean, seashore, retaining wall, and very thick cement walls into the heart of one of Japan's nuclear power plants. This was the immediate consequence of the *Great East Japan Earthquake* (also known as the *2011 Tohoku earthquake*). The consequences spread to both physical and economic systems across the globe. The connectivity led to an avalanche of consequences.

The 21st century is an age of lopsided long-tailed probability distributions rather than the classical normal distribution. Like an episodic tsunami striking the beach, modern-day events come and go in bursts – most are clustered together in time and

The Science of Resilience: Complexity, Risk Modeling, and Systems, First Edition. Ted G. Lewis.

space, with bursts separated in time, space, and consequence according to a power law. This model of reality is a form of *punctuated reality*, characterized by wave-like behavior of modern complex systems.

Scale is another property of the punctuated 21st century – both big and small events are subject to the same long-tailed fingerprint. Aerial and close-up wave-counting experiments produce long-tailed distributions, but at different scales. We say that wave-intervals *scale* because whether I measure the intervals from near or far, both measurements produce a long-tailed distribution. This curious fact is profound because it says something about the similarity of earth-shaking events versus insignificant events – earth-shaking events are a lot like everyday small events – only bigger. It is a matter of scale.

Long-tailed distributions like the ones described in this book are called *scalable*, or *self-similar*, because they are *fractals*. In simple terms, a fractal looks the same at all scales. Take a magnifying glass to a long-tailed distribution like the ones shown in this book, and you get another long-tailed distribution! They all look identical regardless of how near or far away they are observed. Ocean waves and many worldly events such as changes in stock market indexes look similar at different scales – that is, they exhibit self-similarity. The underlying patterns are the same regardless of how far or close we observe them.[1]

Many more events in modern life obey fractal or self-similar distributions than ever before in history. [This is a big claim that will take the rest of this book to justify]. Big waves (tsunamis) are just like small waves (Asilomar Beach), in terms of frequency of time intervals. The time scale (or size scale) may change, but the underlying phenomenon is the same. Big events mimic small events, and vice versa. For example, big events like the Arab Spring in Tunisia and Egypt in 2011 mimic self-similar small events like Occupy Wall Street, and the American Tea Party movement. Even small *flashmob* events that pass with little notice are fractals – disturbances operating at a relatively small scale. Similarly, political protests occur at all scales, but when aggregated into a frequency distribution, they obey a long-tailed distribution just like the ocean waves.

> *Fractal nature of nature: Observations made at different scales often belie the same underlying structural dynamics – punctuated, self-similar, and long-tailed.*[2]

---

1 Changes in a stock price obey a long-tailed distribution and so do wave intervals.
2 As I will show, the rules governing much of modernity are non-parametric, non-linear, scalable, conditional, and regulated largely by extreme statistics.

Some say history repeats itself, but I say history replicates self-similar fractals. Large-scale hurricanes, terrorist attacks, nuclear power meltdowns, financial collapses, and earthquakes are simply scaled-up small hurricanes, attacks, nuclear power hazards, financial disruptions, and earth tremors. Incidents like these happen all the time, but most of them go unnoticed (fall on the left end of the long-tailed distribution), while a small number of extreme events blow up in our face (and fall on the right end of the long-tailed distribution). Instead of repeating itself, history repeats a self-similar pattern at different scales. Mark Twain got it right when he said, “History doesn’t repeat itself, but it often rhymes.”

We put the power law to work modeling complex systems subject to faults due to connectivity of systems. The power law is a result of connectivity – structure that renders the likelihood of an incident conditional on prior incidents. Like dominos falling because an adjacent domino fell, the next incident is predicated on earlier incidents and conditions.

This chapter presents basic concepts of risk and resilience in complex systems under stress:

- Resilience is defined in terms of a system under stress and its reaction. The components of resilience are (and not limited to) design, resistance, absorption, adaptation, and recovery. The profile of an incident is represented as a v-shaped incident profile curve.
- We define systems as black boxes with inputs and outputs subject to stresses that lead to failures with *consequences*. Accordingly, we model stress and consequences as probabilistic or stochastic functions that follow a *power law* with slope q, also known as the *fractal dimension*, leading to a fundamental mathematical theory of *resilience*.
- System behaviors that demonstrate long-tailed power laws ($q < 1$) are considered *complex*. They hold the potential for *complex catastrophe*, as gauged by the slope, q, of their *exceedance probability distribution*.
- Exceedance is the probability an incident will equal or exceed a certain consequence. It can be computed two ways: (1) binned or (2) ranked. We explore both to show the advantages of each and compare their accuracy.
- These concepts are developed from their origins in the work of Per Bak and his simulations of avalanches in a hypothetical sandpile. These became known as the *BTW experiment* and form the basis of the science of resilience.
- We extend Bak’s theory to include exceedance probability, *risk*, and *MPL* – *maximum probable loss* as measures of risk and resilience.
- We find that risk and resilience depend solely on the fractal dimension, q, obtained by data analysis of recorded faults in real systems.

The Many Faces of Resilience

Intuitively, a *resilient system* is one that tolerates faults, rapidly recovers, and continues to operate, perhaps with less performance, while under stress [1]. Figure 1.1 illustrates the general idea applied to the BTW experiment. Stress is always present at some level and leads to an incident with consequences.

In contrast, Taleb [2] defines a resilient system as an *antifragile system possessing an increase in capability to thrive as a result of stressors, shocks, volatility, noise, mistakes, faults, attacks, or failures*. Taleb's definition suggests that resilience is an increase of function above "normal operating capacity" in order to compensate for stresses. Antifragility implies additional cost and capability in anticipation of stress, while resilience is a property of systems in general, regardless of their capacity for handling stress.

Resilience and resilient system design touches on at least six dimensions of systems subject to stress:

- Resilience by design – structure, redundancy, etc.
- Resistance to stress – structure, anti-fragility, etc.
- Absorbance of stress – redundancy, backup, etc.
- Adaptation – lowering vulnerability/consequences by re-design
- Recovery – time to recover, backup, restart, etc.

While we will not cover all of these aspects of systems capable of recovering from a damaging incident, we will focus on many of the attributes above, such as

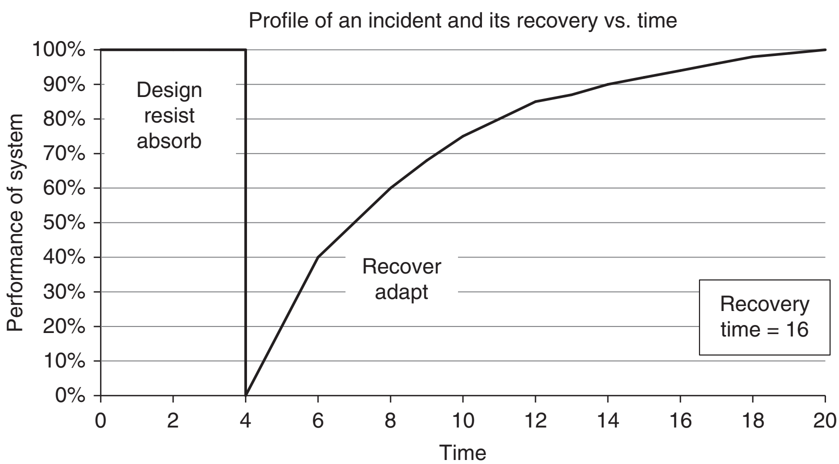

Figure 1.1 Idealized profile of an incident and its recovery over time. Resilience is quantified in terms of design, resistance, ability to absorb and adapt, and recovery time.

structure, redundancy, buying down vulnerability, and time to recover. The model proposed here assumes an unexpected incident of probabilistic magnitude will occur with a probability determined by past history, for example, a *frequentist* approach to disaster.

We take a conservative approach in which system resilience is a property that depends primarily on two factors: consequences due to faults caused by stress and recovery time following an incident. Resilience may exist in a system without additional antifragility capacity. In our definition, resilience is a byproduct of design and implementation of a system. We assume every system is designed and implemented to resist degradation due to stresses such as attacks, accidents, and aging.

The question is, what characteristics of a system lead to resilience, or anti-fragility? To answer this question, we need a quantitative model – a digital twin – of the system. Because accidents happen, our digital twin is based on probabilities. That is, it is stochastic by nature.

Figure 1.2 illustrates a simple system with inputs and outputs, stresses, and consequences. Stresses are typically stochastic anomalies such as unexpected demands, surges, attacks, or accidents. *Consequences* are the damages incurred by the stresses. Under normal operating conditions, a deterministic output is obtained from every deterministic input. For example, an electric power grid produces a known output corresponding to power generation inputs. A highway network provides transportation for a known quantity of automobiles as long as there

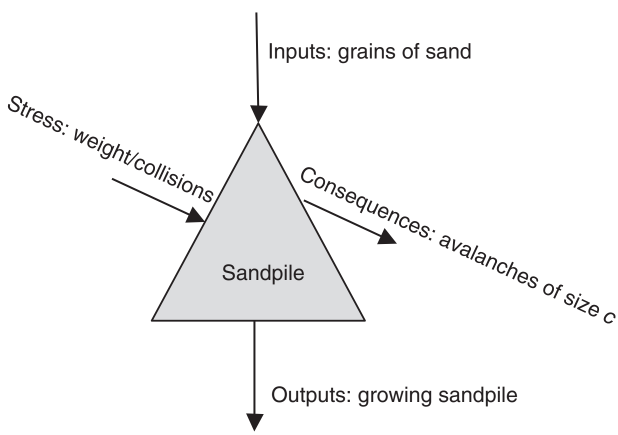

Figure 1.2 Idealized sandpile with grains of falling sand and the resulting stress they cause as inputs and growing size of sandpile and avalanches of size *c* as outputs. This forms a simple system modeled by Bak, simulated/experimentally, and validated by many others. This subsequently became known as the BTW experiment.

are no surges or accidents and traffic flows smoothly as expected. Consequences are zero as long as stresses are zero and nothing breaks or wears out.

Ultimately, something eventually breaks, or an unexpected stress has impact on the smooth operation of the system. The incident results in disruption, which is quantified in terms of consequence. Consequence is loss of productivity, time, human life, capital, or some asset that is typically quantified in terms of loss – damaged equipment, dollars or time, or both.

While a system is typically deterministic, producing known outputs given known inputs, stresses occur unexpectedly and therefore are stochastic. The probability of a stress-induced incident with consequence c is given by $\Pr(c)$. The expected loss – *risk* – is a product of consequence c and the probability of c:

$$R(c) = c\Pr(c)$$

For example, in a simple coin-tossing experiment, the probability of heads and tails is 50% each assuming the coin is fair. In a game where heads pays $1 and tails means the gambler pays $1, the expected loss per coin toss is $0.50. Each time the coin is tossed, the gambler risks losing $1.

$$\Pr(\text{tail}) = 0.5; c = \$1$$

$$R(\text{tail}) = c\Pr(\text{tail}) = \$1(0.50) = \$0.50$$

In the following, we develop a complete theory of risk and resilience from simple, straightforward building blocks such as probability and consequences. We begin in 1987 with the sandpile metaphor studied by Per Bak and associates.

Sandpile as Metaphor

In the mid-1980s, research scientist Per Bak and two colleagues, Chao Tang and Kurt Weisenfeld, performed an unusual thought experiment that became known as the BTW experiment.[3] It was extremely simple but exhibited complex behavior – the behavior of a *complex catastrophe*. Bak, Tang, and Weisenfeld stumbled onto the perfect metaphor for systems that collapse under pressure.[4]

Figure 1.2 is an idealized model of the BTW experiment that begins to *avalanche* when placed under stress. The sandpile system has inputs and outputs as shown. The researchers imagined a beachcomber pouring grains of sand on a flat stretch of beach. As the beachcomber drops grains onto the cone-shaped pile, it grows in size and eventually reaches a *critical point* – an extreme state of *self-organized criticality*. At that point a cluster of grains of sand – an avalanche – breaks away and cascades down the side of the cone.

The process repeats itself – a slow build-up of sand followed by a sudden release of a section of the cone followed by another build up and so on. The elapsed time between subsequent avalanches is called *interarrival times* and is used to predict the likelihood of the next avalanche. The sandpile process produces a series of avalanches at random times and of random sizes. In an actual physical experiment, one might record a series of size and interarrival times as measure of consequence – both size and the moment in time when the next avalanche might occur are of interest.

Consequences: $(x_1, x_2, \ldots x_n)$
Interarrival Time: $(t_1, t_2, \ldots t_k)$

Two questions intrigued Bak and associates:

1) How large are the avalanches in terms of grains of sand?
2) Is it possible to predict size and when the next landslide will occur?

BTW explored this model to understand how it behaves under stress, represented here as grains hitting other grains and dislodging them causing avalanches of variable size to occur. Avalanches are analogous to earthquakes, wildfires, and terrorist attacks, called *incidents*. Bak wanted to quantify such incidents in terms of size of the avalanches and the elapsed time between them – the interarrival time. He reasoned that the sandpile represents real-life incidents that have both consequences (size) and inevitability (interarrival time).

As it turns out, both size and interarrival time are very difficult to express formally in terms of mathematics. The best we can do is measure them and derive empirical formulas based on statistics – frequency of size and frequency of elapsed time between subsequent avalanches. The BTW experiment has been repeated thousands of times in real-life by curious researchers. Frequency of size and interarrival times are estimated using probabilistic analysis. The result is that they both obey a long-tailed distribution.

Bak initially assumed the size of avalanches (number of grains) obeyed a binomial distribution, which is the stochastic model most commonly used to describe the consequences of independent random events. A grain striking the pile hits another grain, breaking it loose, which propagates to another grain or grains, etc., ending in a cascade of grains sliding down the side of the sandpile. It is reasonable to assume cascades occur randomly without *a prior* condition.

If the avalanches are independent and uniformly likely to occur, standard probability models should apply. For example, the size of avalanches should obey a binomial, $B(k)$, or perhaps Poisson distribution and interarrival times should obey an exponential distribution, $\text{Exp}(x)$ (see Figure 1.3).

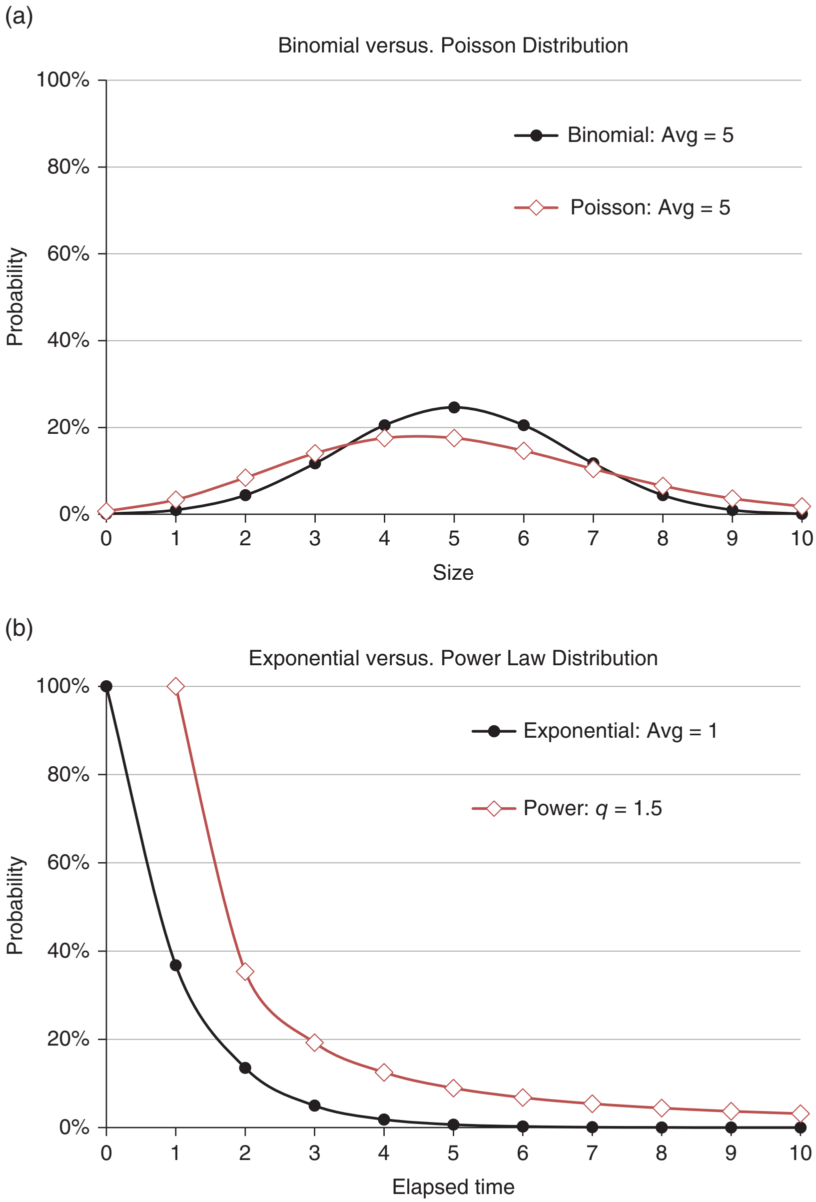

Figure 1.3 Candidate distributions for the sandpile model. (a) Binomial compared with Poisson distribution, and (b) Exponential compared with power law distribution.

$$B(k) = \binom{n}{k} p^k (1-p)^{n-k}$$

where p is the probability of success in k trials

$$\text{Exp}(t) = \lambda e^{-\lambda t}$$

where mean value of t is $\mu = 1/\lambda$

The problem is empirical data from the real world produces a different result. Instead of the expected normal, binomial, or exponential distribution, real-world distributions follow a long-tailed power law as shown in Figure 1.3b. They have many names – Zipf's law, Pareto's law, Levy walk,[3] scale-free, fractal, and others. But what really makes power laws interesting is the fact that they are *not* bell-shaped like the normal and binomial distributions. Normal and binomial distributions define uniformly random phenomena, whereas power law distributions define biased, but still random, phenomena. Normal distributions model the ideal fair world; power laws model the unfair real world.

The general form of a long-tailed distribution is given by the power law, $\Pr(x)$:

$$\text{pr}(x) = \text{A}x^{-q}$$

where A is a normalizing constant, $x > 0$, $q > 0$ is the power.

This unexpected result has shocking implications. First, $P(x)$ may lack a *mean value* and *standard deviation*. Second, $P(x)$ is lopsided with small sizes and interarrival times occurring most often and large sizes and interarrival times less often. In fact, large sizes and interarrival times are extreme because of their size and rarity. Power laws suggest something extreme is going on. One should expect the unexpected.

Power law distributions are considered nonparametric because they may not have a mean or standard deviation. A power law x – k has a mean value $x \in [1, \infty)$ only if $k > 2q > 2$, and it has a finite variance only if $q > 3$, $k > 3$. The sandpile model investigated by Bak found q approximately equal to one, leading Bak et al. to call frequencies obeying the power law $\Pr(x) = x^{-1}$, "1/f noise." "Noise" in this context means randomness.

In practice, power laws with a mean value but undefined variance implies they are capable of *black swan* behavior – *complex catastrophes*. Black swan incidents are rare with large consequences. That is, they occur at the long end of the power law corresponding to large consequences and diminishing probability of occurrence. Rare incidents with huge consequences are a sign of a complex catastrophe or potential for a complex catastrophe.

3 Levy walk is a foraging pattern where the interarrival distance or time obeys a power law.

Power laws are also interesting because of their *scale invariance,* which makes them a *self-similar fractal*. A power law remains similar to itself even when the x-axis is scaled by an amount s:

$$\Pr(sx) = A((sx)^{-q}) = As^{-q}x^{-q} = A'\Pr(x)$$

where $A' = As^{-q}$

The power q is sometimes called the *fractal dimension* and the system that exhibits scale invariance is called *scale-free*. We use the term fractal dimension to remind us that systems obeying a power law are self-similar fractals. They exhibit fractal behavior regardless of the scale. For example, large earthquakes are scaled-up small earthquakes and large wildfires are scaled-up small fires. Systems like the sandpile are self-similar because large avalanches are scaled-up small avalanches. This is another signature of complex catastrophes.

Power laws have only one parameter of interest, q. The resiliency of a system is inversely proportional to q. Small q means a longer fat tail while large q means a shorter thin tail. Recall that consequences are largest in the tail and smallest at the head of the distribution (see Figure 1.3b). Power laws are fat-tailed because $\Pr(x)$ approaches zero more slowly than $\mathrm{Exp}(x)$ as x grows without bound. This becomes more pronounced as $q > 0$ decreases.

To find q, we need only calculate the frequency of occurrence $F(x)$ versus x, and plot the logarithm of F versus logarithm of x to obtain the slope, q, of the resulting straight line:

$$\mathrm{Log}(F(x)) = \mathrm{Log}(\mathrm{Ax}^{-q}) = \mathrm{Log}(A) - q\mathrm{Log}(x)$$

Thus, q is the slope of the OLS (optimal least-squares) fit to the consequence data. This process is illustrated in Figure 1.4 for Benford's Law used to expose fraud in financial data and scale-free patterns in street addresses, stock prices, population sizes, etc. Electrical engineer and physicist Frank Benford published his empirical observations of power laws in a diverse set of data sources including statistics from baseball, atomic weights, the areas of rivers, and numbering of articles in magazines while working for General Electric in 1938.[4]

Benford's most notable law is the observation that leading digits in financial data such as accounting records obey a power law with given frequencies as shown in Figure 1.4. The number 1 appears as the leading significant digit about 30% of the time, while 9 appears as the leading significant digit less than 5% of the time. OLS fit of this data to the log–log plot of frequencies versus leading digit yields $q = 0.863$, which classifies it as extreme (because it has no mean value or standard

4 https://en.wikipedia.org/wiki/Frank_Benford

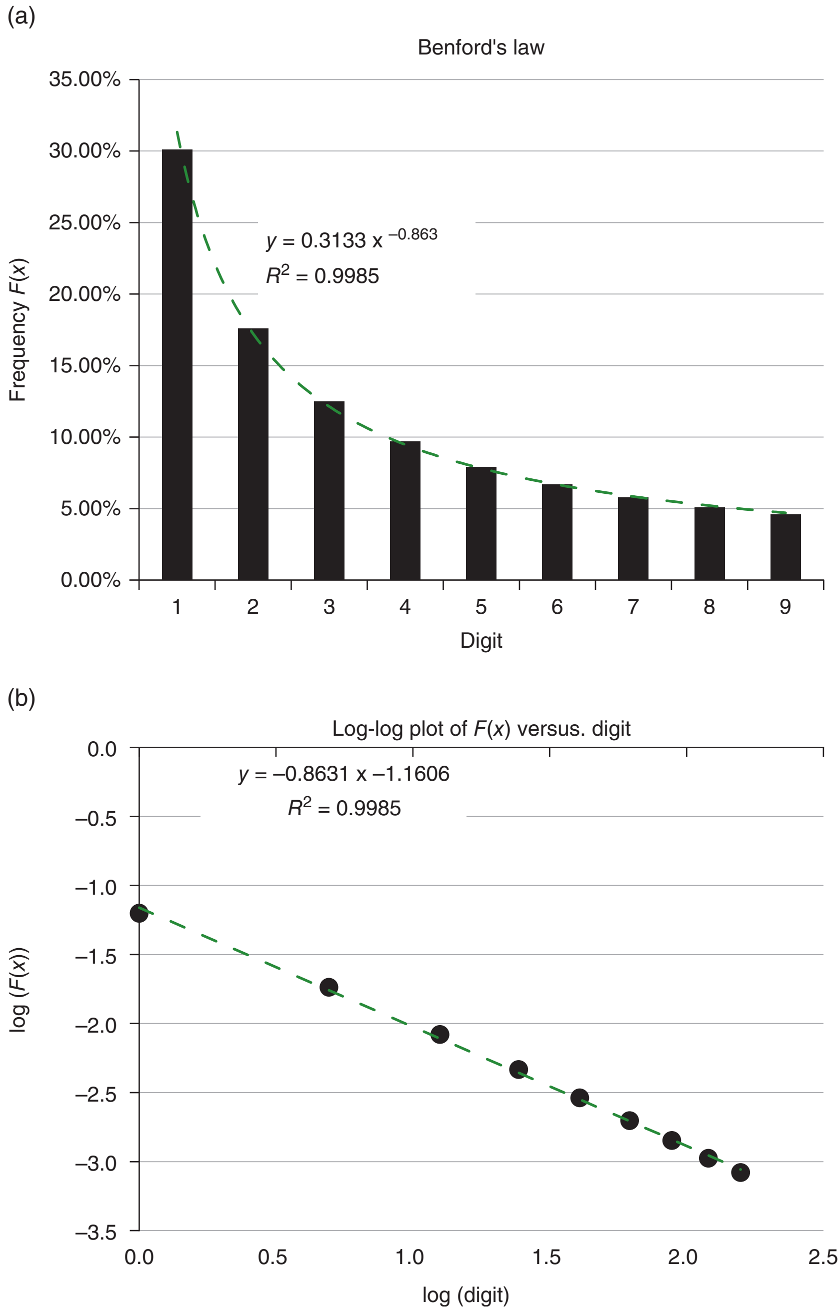

Figure 1.4 Benford's law illustrating the application of power laws to reveal financial fraud. The value of fractal dimension q is obtained by OLS fitting of the data plotted on the log–log scale as shown. (a) Frequency of leading digit obeys a power law (b) Fractal dimension q obtained from log-log plot.

deviation). Financial data that strays too far from this distribution is likely to be fraudulent.

The BTW experiment produced similar results with $q \sim 1$. This suggests that sandpiles are fractals with extreme, but highly unlikely consequences. In fact, a single parameter, q, can be construed as a general measure of risk and resilience in systems that are prone to complex catastrophes. We associate fractal dimension q with fragility because the tail of the distribution of consequences increases with declining q.

In summary, q is a measure or risk and resilience in terms of frequency of avalanches and interarrival times:

- The tail becomes longer and fatter inversely proportional with the value of q.
- Risk and resilience are inversely proportional to q.

Blocks and Springs

It is a mistake to assume sandpile avalanches are the result of independent interacting grains of sand gathering additional grains of sand as the avalanche slides down the slope of the sandpile. The sandpile system is a complex system of interacting parts (grains). Their behavior is *conditional*. That is, the behavior of a grain is conditional on the behavior of other grains. This connectivity is the root source of avalanches and also has an impact the distributions of both size and interarrival times. The more connectivity in terms of resistance and viscosity, the more likely is the distribution to be long-tailed.

Consider the simple apparatus in Figure 1.5a. Six blocks are connected by springs, which allow the blocks to slide back and forth when pushed from the left-hand side. Huang and Turcotte used a similar "slip–stick" sandpile model to explain the tectonic shifting between plates that cause earthquakes.[5] Blocks are like grains of sand, and the springs are like friction and viscosity. In addition, springs make the invisible influence of one block on the others more explicit. Clearly, movement of the inner blocks is conditional on the movement of blocks on their left.

Imagine applying an impulse to the left-most block to force it to move to the right. If the friction between block and plane is small, the impulse will push the left-most block far enough to collide with the adjacent block. If the force is even stronger, several blocks will be displaced and collide with adjacent blocks, etc. The consequence increases with the amount of *random* force applied.

Applying a random force produces an unpredictable number of collisions between adjacent blocks. Therefore, the number of blocks involved in an

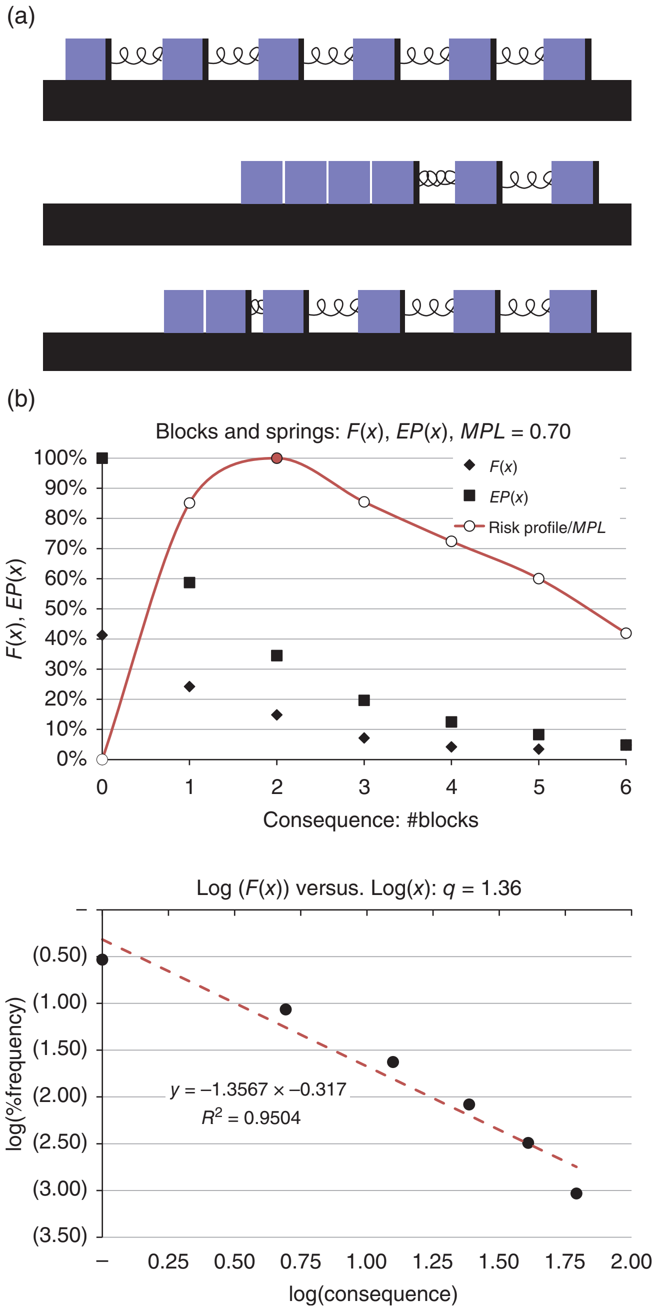

Figure 1.5 Simulation of dynamical system with blocks and springs produces a long-tailed exceedance probability from observed consequence frequencies. (a) Frequency of leading digit obeys a power law. (b) Fractal dimension q obtained from log–log plot.

"accidental collision" is also unpredictable just as the number of grains colliding in Bak's sandpile is unpredictable. Given a stochastic shock or stress, we can record the consequences and then convert them into a frequency chart similar to Benford's law.

Converting the number of times each block is moved by the adjacent block into a frequency count, we get frequency $F(x)$ versus consequence x for the simulation illustrated in Figure 1.5a. The resulting distribution, $F(x)$, obeys a power law as shown in Figure 1.5b, instead of a normal distribution, because of the dependency of the right-most blocks on the actions of the left-most blocks. The frequency of block x is conditional on the frequency of adjacent blocks.

Insurance companies use *exceedance probabilities* to calculate risk and establish premium rates because exceedance represents the worst-case scenario. That is, insurance companies want to know the *maximum probable loss* (*MPL*) due to a *consequence of size* c *or larger*. This is computed by summing over the long tail of the power law.

Specifically, let $EP(x, c)$ be the *exceedance frequency* of blocks moved. Frequency $F(x)$ of moving blocks by distance x is obtained by simply counting the number of times x occurs. Then $EP(x)$ is computed by summing frequencies that are equal to or exceed a certain consequence, c.

$$EP(x, c) = \text{sum over tail of } F(x \geq c)$$

Computing Exceedance Probability

In practice we can compute exceedance probability in two different ways. The straightforward calculation is simply the tail sum of $F(x)$ from x equals c to infinity, called the *true* or *binned* exceedance:

$$EP(x, c) = \sum_{x=c}^{x=\infty} F(x); \textit{binned} \text{ exceedance}.$$

The summation is typically done by binning the raw data to obtain $F(x)$. The x-axis is partitioned into bins, and the number of times a consequence point falls into each bin is normalized as $F(x)$. The tail sum is the sum from consequence $x = c$ to the end of the frequency distribution.

Alternatively, EP is computed as *ranked exceedance* by ranking N consequences from highest to lowest, and then assigning a rank to each:

$$EP(x) = \text{rank}(x)/(N + 1); \textit{ranked} \text{ exceedance}.$$

Here, exceedance is the expected number of moved blocks that equals or exceeds consequence c. Consequences are ranked from largest to smallest and then plotted against consequence x, yielding $EP(x)$. Note that rank(x) and $(N+1)$ produce the same sequence regardless of consequence x. Ranked exceedance is independent of the value of x. It is the value of x that determines the shape of the exceedance curve when plotting $EP(x)$ versus x.

Obviously, ranked exceedance works best when consequences cannot be binned because only one or zero numbers fall into each bin. Binned works best when the consequence data contains duplicates. The results may vary, suggesting that exceedance probability is a proxy rather than an absolute quantity. We are more interested in the shape of the exceedance probability curve than the exact values.

In real life, the exceedance probability curve typically follows a power law as shown in Figure 1.5b for the sliding blocks experiment. This curve is characterized by a single number – the fractal dimension q. Therefore, we can calculate risk – expected loss – as a function of consequence and fractal dimension, q.

Let risk and resilience be defined in terms of *maximum probable loss*, MPL, as follows. $\Re$ is the risk profile obtained by multiplying consequences by exceedance for every value of consequence x:

$$\Re = xEP(x) \approx x \cdot \mathrm{Ax}^{-q} = \mathrm{Ax}^{1-q}$$

MPL is the largest value of $\Re$.

$$\mathrm{MPL} = \max_x \Re$$

For convenience, we plot normalized risk, also known as risk profile, on the same graph with consequence and EP:

$$\Re_{\mathrm{Norm}} = \Re/\mathrm{MPL}$$

Risk profile $\Re$ and MPL are measures of resilience because they depend on fractal dimension q. In the blocks and springs example, above, the value of MPL occurs at $x = 2$, see Figure 1.5b. $MPL = 0.70$ at consequence $x = 2$ with exceedance frequency $EP(2) = 0.33$. Allowing for noise, $2EP(2) \approx 0.70$.

The shape of the risk profile is important because it suggests a way to predict *catastrophic potential*. Note that we have reduced the risk and resilience analysis down to a single factor, q. When the expression for $\Re$ is monotonically increasing, the system has catastrophic potential because $\Re$ increases without bound. Conversely, the risk profile of a resilient system decreases with an increase in consequences. Summarizing:

$$1 - q < 0 \Longrightarrow q > 1 \text{ implies resilience}$$

$$1 - q > 0 \Longrightarrow q < 1 \text{ implies fragility}$$

$$1 - q = 0 \Longrightarrow q = 1 \text{ implies a tipping point (criticality).}$$

At the critical point, $q = 1$, the system transitions from resilient to fragile or fragile to resilient. This phase transition is very noticeable because consequence increases exponentially, and normalized risk $\Re_{\text{Norm}}$ fails to decline as consequence increases. Once again, this is a signature of a system capable of catastrophic behavior under stress. In practice, the risk profile is likely to appear wavy or monotonically rising. If $\Re$ trends monotonically upward with increases in x, the system under investigation has catastrophic potential.

Two Versions of Exceedance

Exceedance may be a new idea to the reader accustomed to more traditional probability theory; hence, this section elucidates the difference between binned and ranked exceedance briefly introduced above. Ranked exceedance is preferred when consequences are distinct, separated by value, and there are few of them. Binned is preferred when consequences repeat or nearly repeat and there are many of them in the data.

We proceed by example. Consider consequences produced by a Fibonacci process Y as shown in Table 1.1a.

$$Y_{i+1} = rY_i + Y_{i-1}; i \geq 1$$
$$Y_0 = 1; Y_1 = 1; r = 0.4$$

Table 1.1 Consequences *Y* produced from a Fibonacci process.

(a) Ranked with *MPL* = 26.05

i	*Y[i]*	Rank	Ranked *EP*	Risk/*MPL*: 26.05	Risk
0	1	31	1.03	0.04	1.03
1	1	30	1.00	0.04	1.00
2	1	29	0.97	0.05	1.35
3	2	28	0.93	0.06	1.46
4	2	27	0.90	0.07	1.82
5	2	26	0.87	0.08	2.05
6	3	25	0.83	0.10	2.48
7	4	24	0.80	0.11	2.85
8	4	23	0.77	0.13	3.37
9	5	22	0.73	0.15	3.90
10	7	21	0.70	0.18	4.57
11	8	20	0.67	0.20	5.28

Table 1.1 (Continued)

(a) Ranked with *MPL* = 26.05

i	*Y[i]*	Rank	Ranked *EP*	Risk/*MPL*: 26.05	Risk
12	10	19	0.63	0.24	6.14
13	12	18	0.60	0.27	7.08
14	14	17	0.57	0.31	8.17
15	18	16	0.53	0.36	9.37
16	21	15	0.50	0.41	10.72
17	26	14	0.47	0.47	12.20
18	32	13	0.43	0.53	13.82
19	39	12	0.40	0.60	15.56
20	47	11	0.37	0.67	17.40
21	58	10	0.33	0.74	19.29
22	71	9	0.30	0.81	21.18
23	86	8	0.27	0.88	22.97
24	105	7	0.23	0.94	24.51
25	128	6	0.20	0.98	25.63
26	156	5	0.17	1.00	26.05
27	191	4	0.13	0.98	25.42
28	233	3	0.10	0.89	23.26
29	284	2	0.07	0.73	18.91
30	346	1	0.03	0.44	11.54

(b) Binned with *MPL* = 25.81

bin	*x*	*F*(*x*)	*EP*(*x*): *q* = 0.99	Risk (*x*)	Risk/*MPL*: 25.81
10	13	0.42	1.00	10.00	0.39
20	3	0.10	0.58	11.61	0.45
30	2	0.06	0.48	14.52	0.56
40	2	0.06	0.42	16.77	0.65
50	1	0.03	0.35	17.74	0.69
60	1	0.03	0.32	19.35	0.75
70	—	—	0.29	20.32	0.79
80	1	0.03	0.29	23.23	0.90
90	1	0.03	0.26	23.23	0.90
100	—	—	0.23	22.58	0.88
110	1	0.03	0.23	24.84	0.96

(*Continued*)

Table 1.1 (Continued)

(b) Binned with *MPL* = 25.81					
bin	***x***	***F*(*x*)**	***EP*(*x*): *q* = 0.99**	**Risk (*x*)**	**Risk/*MPL*: 25.81**
120	—	—	0.19	23.23	0.90
130	1	0.03	0.19	25.16	0.98
140	—	—	0.16	22.58	0.88
150	—	—	0.16	24.19	0.94
160	1	0.03	0.16	25.81	1.00
170	—	—	0.13	21.94	0.85
180	—	—	0.13	23.23	0.90
190	—	—	0.13	24.52	0.95
200	1	0.03	0.13	25.81	1.00
210	—	—	0.10	20.32	0.79
220	—	—	0.10	21.29	0.83
230	—	—	0.10	22.26	0.86
240	1	0.03	0.10	23.23	0.90
250	—	—	0.06	16.13	0.63
260	—	—	0.06	16.77	0.65
270	—	—	0.06	17.42	0.68
280	—	—	0.06	18.06	0.70
290	1	0.03	0.06	18.71	0.73
300	—	—	0.03	9.68	0.38
310	—	—	0.03	10.00	0.39
320	—	—	0.03	10.32	0.40
330	—	—	0.03	10.65	0.41
340	—	—	0.03	10.97	0.43
350	1	0.03	0.03	11.29	0.44

Figure 1.6 shows the plot of ranked *EP* versus consequence *Y* taken directly from Table 1.1. Ranked exceedance depends on the rank – in place of $F(x)$ – of the descending order of consequences. It results in a smoother exceedance curve and risk profile because of the smoothness ranking. Consider the $EP(x)$ of 0. N consequences of equal probability $\frac{1}{N+1}$ of x happening regardless of the value of x:

$$EP(x, c) = \sum_{x}^{N} \frac{1}{N+1}$$

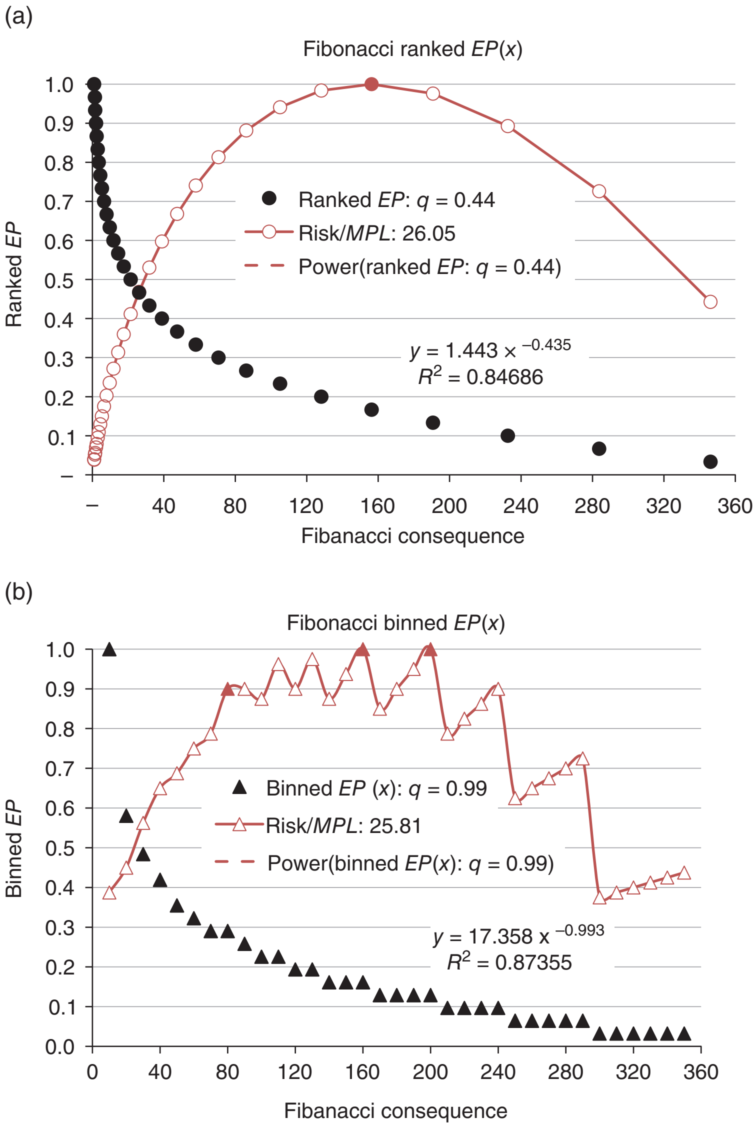

Figure 1.6 *EP* (consequence) plots the same data using two distinct definitions of *EP*. (a) Ranked: $q = 0.44$, $MPL = 26.05$. (b) Binned: $q = 0.99$, $MPL = 25.81$.

$$= \frac{1}{N+1}\sum_x^N 1 = \frac{N+1-x}{N+1} = \frac{\text{rank}(x)}{N+1}$$

A long-tailed exceedance probability curve is obtained when plotting *EP* versus *x*:

$$\frac{\text{rank}(x)}{N+1} \text{versus } x$$

In this example, ranked *EP* produces a radically lower fractal dimension, $q = 0.44$, and an *MPL* of 26.05 compared with binned exceedance (see Figure 1.6a). All curves are smooth due to the smooth curvature of rank(*x*). Keep in mind, rank(*x*) assumes all incidents occur equally likely and only the magnitude of the consequences differ!

Binned *EP* derives from the number of consequences placed in bins and is therefore dependent of the values of *x* instead of the rank of *x*. The disadvantage of this approach is that selection of bin size is rather arbitrary. Figure 1.6b shows the results of placing consequences in bins of size 10. Frequency is derived from counting the number in each bin, instead of from rank.

How do the power laws compare? Comparing binned with ranked, the power law fit is radically different with $q = 0.99$, but the values of *MPL* are similar – 25.81 versus 26.05. This, however, depends heavily on selection of bin size. A large bin size reduces the number of points in the *EP*(*x*) graph, while a small bin size decreases the precision as shown in Figure 1.6b where near-duplicate consequences together show up as stair-steps in *EP* and risk profile.

In general, use ranked when little data are available and binned when a large data set is available. Experiment with different bin sizes in order to narrow the gap between ranked and binned *EP*. Bak and associates used purely statistical methods to obtain their 1/f noise conclusion. Today, we are more interested in greater accuracy and the interpretation of the data from a complex catastrophe point of view.

Comments

The BTW experiment received the attention of many writers in many diverse fields of study. It was used as a metaphor by Al Gore in his 1992 book, *Earth in the Balance: Ecology and the Human Spirit*. Gore argued that the impending climatic disaster awaiting humankind is just like a sandpile edging ever closer to collapse.

Taleb describes a miniature Tower of Babel he built from Rio de Janeiro beach sand in his popular book *Fooled By Randomness*. But Taleb never mentions Bak's name. Nonlinear extreme events become black swans in Taleb's sequel, *The Black Swan*. Again, Taleb never mentions Bak or explains the underlying cause of black swans.

Ramo devotes an entire chapter to the sandpile metaphor, claiming, "He [Bak] could have been speaking about the Middle East, relations between the United States and China, the oil market, disease, nuclear proliferation, cyber warfare or a dozen other problems of global affairs and security." [6] Ramo's use of Bak's ideas expanded Bak's theory beyond the physical and biological sciences into the social and political realm.

Over a period of several decades, Bak's theories have gained more scientific attention and been applied to more situations than the science and technology community could have imagined back in 1987. It seems that Bak's fame was on solid ground, and his ideas would soon be used to explain why complex systems of all kinds eventually fail, recover, and fail again. Sandpile as metaphor was poised to become mainstream. But this didn't happen.

Chao Tang's description of Per Bak is insightful and may explain why the sandpile metaphor didn't revolutionize the scientific or socio-economic world. He says, "Per was a very straightforward person, unconventional, unorthodox, and antiestablishment. He liked to be controversial, and more often than not, he would embarrass the speakers in seminars and conferences by asking them simple questions and insisting to have simple answers. He had many friends, and also many enemies. To many, he may not appear to be a polite person. But he could actually be very sweet. He was a great cook – he cooked for me twice at his home (once steak and once fish), and those were among my best dining experiences."[5]

Kurt Wiesenfeld's admiration remains decades after their postdoc time together, "He was one of those larger-than-life people one often hears about and only very occasionally meets. Per was outgoing, loud, and if I had only a little interaction with him I probably wouldn't have liked him. But he wears extremely well. I will relate one story to illustrate. I visited Per at Brookhaven in the winter of 1985. At the time I was at UC Santa Cruz doing a one-year post-doc stint and was looking for another post-doc job. I saw a fair amount of loud talking/arguing that day and was amazed during private conversations that virtually everyone there made excuses (it seemed to me) for Per's aggressive manner. After working there, myself I understand it. He generated a lot of affection among those who got to know him (and, it seemed to me, a fair number of enemies among the rest). The point is that he might make a rude joke about someone's ideas but was just as likely to make the same sort of joke about himself. And that was an endearing trait."[6]

Perhaps Bak was too loud and abrasive. Perhaps he was a man ahead of his time. Or perhaps his peers recognized and then quickly forgot him like so many other pioneers of the unconventional. Bak was highly critical of the scientific establishment, and they of him. He especially disliked "big science," criticizing it for wasting money while ignoring the true geniuses of science. Bak championed the

5 Email correspondence, June 2010.
6 Email, August 2010.

underfunded lone genius. In his mind, creativity and intelligence were more important than large funding grants.

Predrag Cvitanovic, a professor of physics at the Georgia Institute of Technology, said Bak was "the most American of Danes. Danes eschew confrontation, but [Bak] was arrogant and loved to fight with his colleagues in academia. We all have stories of how we first met him, usually remembered by some outrageous statement or insult."[7] Was Bak simply too brash?

I asked Professor Tang what he thought of Bak's theories after nearly 25 years since the BTW experiment. "I think it is fair to say that the discovery of complexity leading to collapse was really a joint work of three of us, with critical inputs from Sue Coppersmith. I am of course biased, but I think the discovery opened the door to a new way of thinking about complex systems and the emergence of scaling. The original sandpile model (and the many subsequent models) is simplified and idealized. While simplified models played absolutely essential roles in many fields of physics, including statistical mechanics, it remains to be seen how general and how widely applicable it is in explaining natural phenomena. I have no doubt that it can be applied to some natural systems and the mechanism shown to be behind the observed critical behavior in these systems. But it is not clear that it is as universal and as widely applicable as Per had hoped. In any case, I think it is a great idea, and I believe it will prove to be a useful model."[8]

References

1 Lewis, T.G. (2023). The many faces of resilience. *Communications of the ACM* 66 (1): 56–61. https://doi.org/10.1145/3519262. https://cacm.acm.org/magazines/2023/1/267957-the-many-faces-of-resilience/fulltext.
2 Taleb, N.N. and Douady, R. (2013). Mathematical definition, mapping, and detection of (anti) fragility. *Quantitative Finance.* 13 (11): 1677–1689. arXiv:1208.1189.
3 Bak, P., Tang, C., and Weisenfeld, K. (1987). Self-organized criticality: an explanation of 1/f noise. *Physical Review Letters* 59: 381–384.
4 Lewis, T.G. (2011). *Bak's Sand Pile*, vol. 368. Agile Press.
5 Huang, J. and Turcotte, D. (1990). Evidence for chaotic fault interactions in the seismicity of the San Andreas fault and Nankai trough, Nature. 348: 6298. Research Library.
6 Ramo, J.C. (2009). *The Age of the Unthinkable: Why the New World Disorder Constantly Surprises Us and What We Can Do About It*, 59–60. New York, Little, Brown & Company.

7 http://en.wikipedia.org/wiki/Per_Bak
8 Email communication, June 2010.