

DISOBEYING POWER-LAWS

PERILS FOR THEORY AND METHOD

G. CHRISTOPHER CRAWFORD

Abstract: The “norm of normality” is a myth that organization design scholars should believe only at their peril. In contrast to the normal (bell-shaped) distribution with independent observations and linear relationships assumed by Gaussian statistics, research shows that nearly every input and outcome in organizational domains is power-law (Pareto) distributed. These highly skewed distributions exhibit unstable means, unlimited variance, underlying interdependence, and extreme outcomes that disproportionately influence the entire system, making Gaussian methods and assumptions largely invalid. By developing more focused research designs and using methods that assume interdependence and potentially nonlinear relationships, organization design scholars can develop theories that more closely depict empirical reality and provide more useful insights to practitioners and other stakeholders.

Keywords: Power-law distributions; Gaussian statistics; Pareto; nonlinear statistical methods; theory building

As myriad studies in nearly every area of business and science indicate, the norm of normality is a myth. Instead, we exist in a world where power-law (i.e., Pareto) distributions are ubiquitous. In contrast to the traditionally assumed normal (Gaussian) distribution of organizational outcomes, where events are completely independent and identically distributed, power laws identify the fundamental interconnectedness and interdependence of events (Andriani & McKelvey, 2009). As shown in Figure 1, power-law distributions are highly skewed, with long, fat tails (a downward-sloping straight line when plotted on log-log axes) that identify outliers (i.e., extreme events). When graphed on regular scales, the power-law distribution looks like Figure 1a; log-log axes are shown in Figure 1b. These distributions are interesting because of the various outcomes contained in their elongated tail, represented by the “Paretian world” area to the right of the shaded region in Figure 1b.




Fig. 1. Power-Law Distributions: Normal and Log-Log
Source: Figure 1b is from Boisot and McKelvey (2010).

The shaded Gaussian region on the left of Figure 1b represents the vast majority of organizations. There is an obvious contrast in size and potential scope of influence as firms move down the slope into the Pareto region. Though infrequent, an outcome in the long

tail of a power-law distribution – the large circle at the bottom right of Figure 1b – is of disproportionate influence on the entire system. To illustrate, the small circles in the upper left might represent the 17 million Mom & Pop retail stores in the United States, while the largest circle in the bottom right might represent Walmart. Accordingly, this distribution represents something qualitatively different that must be taken into account. The distinctiveness of the distribution and its complex effects make it relevant to both practitioners and scholars. Empirical studies have discovered power-law patterns in nearly every aspect of the internal and external contingencies explored by organization design scholars. For example, such distributions have been found in U.S. firm size, overall industry structure, competitive performance advantages, industry sector and firm growth rates, firm survival and exit, network structure, market share prices, product innovations and technological breakthroughs, entrepreneurial growth expectations, new venture performance, and individual performance – and in many others in physical, natural, biological, and social systems.

Power-law distributions emerge as a result of tension and connectivity dynamics among agents in a system. However, it is often difficult to see – let alone understand – these patterns if they are viewed only at one level. Although power-law distributions are pervasive in domains relevant to organization design, the importance of these unique statistical signatures is seldom explored.

POWER-LAW EFFECTS ON EMPIRICAL OBSERVATIONS

Traditional (Gaussian) statistical analyses are not applicable to firms in all regions of the distribution when power laws are present (Boisot & McKelvey, 2010). Though empirical reality continually displays evidence of skewed outcomes, scholars continue to use Gaussian statistical techniques that assume normal distributions of outcomes, linear relationships among variables, stable means, finite variance, and independence of events. Indeed, using Gaussian assumptions and methods to explain power-law phenomena can lead to inaccurate conclusions, under-specified theoretical models, and misleading normative recommendations, all of which reduce the credibility of scholarly research (O’Boyle & Aguinis, 2012). The prevalence of power laws has implications for both research design and statistical analysis.

Research Design Implications

Starbuck and Nystrom (1981) suggested decades ago that useful prescriptions for organization design do not come from conducting indiscriminant empirical studies. Instead, they argued, “truly innovative designs have to originate in deviant cases or in fantasies rather than in statistical norms (1981: 8).” Here, “deviant” cases are those outside “normal” (i.e., in the tail of the power-law distribution) because of their size, rarity, and potential scope of influence on the environment. Several aspects of power-law distributions will highlight their applicability to the empirical study of organization design. As the Gaussian and Paretian regions of Figure 1 suggest, there are distinct differences in the inputs and outcomes for firms in each region. First, these distributions exhibit data that are both linear and nonlinear. While assumptions of independence and additive relationships can aid in understanding linear relationships, these assumptions are wholly violated in nonlinear relationships. This also suggests that studying extreme, nonlinear outcomes – the companies that organizational scholars frequently use as deviant case examples and seek to explain such as Apple, Facebook, and Enron – is implicitly hindered by the use of traditional statistical methods. Second, once firms are large enough to be in the Pareto region of the distribution, their activities become multiplicative and nonlinear, where they have the potential to influence the outcomes of other firms in the sector.

When we use Gaussian methods to study firms in a population, we are not measuring performance as much as we are constraining it. Few phenomena adhere to a power law over all values. Instead, the power law most often applies for values greater than some minimum – this is the tipping point of the distribution, where the tail begins. As an example, the graph in Figure 2 shows an analysis of annual revenue of the Inc. 5000 Fastest-Growing Companies in the United States. Using the `plfit.m` script from Clauset, Shalizi, and Newman (2009) in MATLAB to construct a semi-parametric bootstrap maximum likelihood estimation of fit with a power-law model, the graph shows the data tip from linear to nonlinear when firms

reach \$158M in revenue. At this size, firms have a much greater potential to influence the surrounding environment (i.e., to have co-evolutionary effects). Therefore, it is important to design studies that include entire populations of interest – or at least large random samples thereof in order to maximize data collection in all regions of the distribution.




Fig. 2. Maximum Likelihood Estimation of Power-Law Fit: Inc. 5000 Fastest-Growing Private U.S. Companies

Note: The entire dotted line represents the maximum likelihood estimate of the power-law tail's slope, calculated as 1.79. The tipping point of this distribution, where nonlinear and co-evolutionary effects begin, is \$158M. Kolmogorov-Smirnov goodness of fit is 0.036 (≤ 0.10 is desirable).

Statistical Implications

If viewed through the lens of traditional statistics, entities in the tail of the distribution are outliers – they are either viewed as random or a hindrance to obtaining statistical significance. However, rather than being deleted or transformed, these outliers often should be the ones that really matter to organization design scholars. The statistical average does not help scholars understand the true dynamics of the environments in which those firms exist, and it does not provide much instructive relevance for managers. What matters are the extremes! When data are skewed, especially where agent connectivity and interdependence are prevalent, it is likely that power-law dynamics are influencing the distribution. As an example, a normal distribution has a skewness of 0, and data are considered skewed if that number is above 3 (Greene, 2007). Interestingly, in Figure 2, the skewness of these data is 36. Thus, as skew increases, the more the distribution has the potential to exhibit power-law characteristics. All of this becomes very problematic for researchers using Gaussian techniques that assume normal distributions with stable means and finite variance, a problem that has been extensively documented (Andriani & McKelvey, 2009; O'Boyle & Aguinis, 2012; Simon, 1955).

In power-law distributions, the mean is unstable and variance is nearly infinite; therefore, no single observation can represent the average of the system. Revisit Figure 1b: Whereas the tip of the downward-pointing bracket is implied as the mean of the distribution, it is probably closer to the median. In highly skewed distributions, extreme values on the right often pull the mean beyond the lower bound of the power-law tail. According to O'Boyle and Aguinis (2012), this suggests that nearly 70 percent of the population is performing below average. As

in Figure 2, the mean of the distribution is \$73M while the median is only \$11M. This implies that any explanatory or predictive theory built or tested using linear statistical methodology on the decidedly non-normal inputs and outcomes at every level in the domain must include a discussion about how the violation of Gaussian assumptions affects the analysis.

The most important variance in the data is buried by traditional robustness techniques. To wit, a common rebuttal from econometric scholars is, “Just transform the data to reduce the influence of outliers.” Doing so undoubtedly increases the probability of a significant finding (and may increase the probability of a favorable review by editors). However, such a transformation of data obfuscates the actual effect of the differences that firms experience. In a study of the retail industry, for example, transforming the data does not reduce the real magnitude of Walmart’s influence on a corner Mom & Pop store. Thus, deleting or manipulating outliers to achieve statistical significance is misguided. Schoonhoven (1981) suggests that assuming linear relationships is “misplaced” and that testing for nonlinear effects should be mandatory in all empirical analyses.

It would behoove scholars to more thoroughly understand their data prior to conducting Gaussian analyses. Two questions organization design scholars can ask as they develop their research projects are: “Can the performance of an outlier influence the outcomes of others in the distribution?” and “Can one node accurately represent the average of the population?” If the answer to the first question is “yes” or the second question “no”, then quantitative nonlinear techniques like Poisson processes, Bayesian neural networks, historical extreme event analysis, or deep structure analysis may produce more accurate descriptions of the system outputs (O’Boyle & Aguinis, 2012). Similarly, techniques using non- or semi-parametric distributions, or agent-based models, where probabilistic interactions among firms can be simulated in a virtual environment, could be used to more accurately reflect empirical reality.

POWER-LAW EFFECTS ON THEORY BUILDING

When it is likely that empirical observations are influenced by power-law effects, theory building efforts should reflect their presence. Power laws are called “scale-free” distributions because they look the same regardless of the scale used to measure them. In these distributions, the relationships among the size of the events are fractal – they have self-similar behavioral patterns and physical characteristics, where the small appears similar to the big, and individual sub-parts look the same as the whole (West & Deering, 1995). Here, theory development is demanding because it requires knowledge about the whole system and about the underlying emergence of all the sub-systems.

Researchers must investigate causality from both the bottom up and the top down. Simon (1968) and others suggest that the emergence of power-law distributions is driven from the bottom up by the simple rules of the agents (e.g., workers, firms, teams) as they interact within a system. From the bottom up, rules represent an agent’s recursive decision-making heuristics for achieving desired outcomes. These heuristics influence an agent’s habitual behavioral strategies for interacting with the environment. Over time, as successful strategies are given positive feedback from the environment, power-law patterns of outcomes emerge. As agents accumulate resources (e.g., revenue, employees) and become large enough to be in the tail of a distribution in a local environment, there is an increased potential for co-evolutionary effects on the global environment, one level of analysis higher. From the top down, rules (e.g., corporate goals, cultural context, regulatory restrictions) and inputs (e.g., quality and quantity of both competitors and resources in the environment) impose tension on the sub-systems. Together, the rules and inputs throughout the entire system require theory that explains this co-evolving causality. Thus, power laws may be generated in a process as shown in Figure 3.

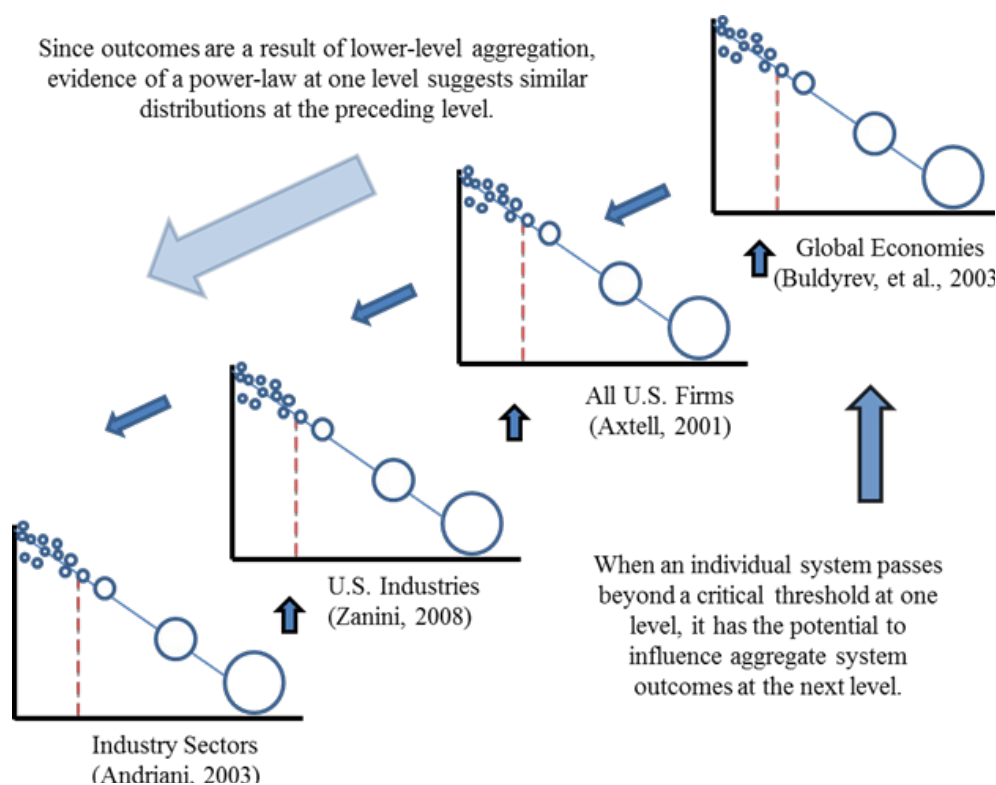


Fig. 3. Cycle of Bottom-up Emergence and Co-evolutionary Causality

When power-laws are present, theories to explain them are scale-free. In this case, an explanation for outcomes at one level of analysis should also explain inputs and outcomes at preceding levels. A scale-free hypothesis can pose as a simple theory that provides the best scalable explanation for the empirical regularities found in the data – one that seeks inference toward the best explanation by examining a mass of data and suggesting a plausible explanation for the patterns. Such a hypothesis is a simple, parsimonious, plausible, and falsifiable theory (Popper, 1963) for the emergence of power-law distributions and for the extreme outcomes therein.

In contrast to the probabilistic determinism of traditional statistics, a scale-free hypothesis needs to identify inputs and rules that could connect to produce an extreme outcome. A scale-free theory provides plausible anticipation rather than prediction. What is “plausible”? As Simon (1968:449) explains, “It is not inconsistent with everyday knowledge. At the moment they [i.e., the simple set of generative mechanisms] are introduced, they are already known (or strongly suspected) to be not far from the truth.” This may create push-back, though, from reviewers who have experience and comfort with highly refined econometric models that supposedly control for alternative configurations and provide robust statistical significance. Thus, scholars can improve the efficacy of their theories and methods by integrating power-law logic into their research designs to facilitate community-wide acceptance.

INCORPORATING POWER-LAW REASONING INTO RESEARCH DESIGNS

Scale-free theory building efforts need to integrate disparate empirical knowledge sources. The more the knowledge is representative of the entire population – whether from random-digit dialing, meta-analysis, or experimental studies – the better. Subsequent empirical testing will require nonlinear methods to discover generative mechanisms. For example, Crawford (2012) hypothesized that a new venture founder’s resource endowments and expectations for future growth generated the highly skewed distribution of outcomes in entrepreneurship. Analyzing organizational forms in three representative samples, at three different stages, at multiple time periods – starting with nascent pre-organizing expectations, to emergent

outcomes, to active start-up firms, to hyper-growth companies ($N = 11,000$) – the study used semi- and non-parametric bootstrap simulations to find statistically identical power-law slopes for nascent expected growth and for firm growth rate in all three samples. This suggests that expectations for growth influence the distribution of outcomes throughout the domain. Future theory-building efforts in this subject area should account for (and collect data on) growth expectations – from both the top down and the bottom up – at multiple levels to limit unobserved variable bias (Carroll & Harrison, 1998).

Scholars need to study the causal dynamics that underlie the emergence of power-law distributions by linking methodologies to the points on the distribution where they are most useful. Scholars may hypothesize that power laws come from “one generative mechanism” as suggested by Boisot and McKelvey (2010) or from “a simple set of mechanisms” as posited by Simon (1968). Depending on the research question of interest, scholarly investigation to explain power laws should be organized around one definition of generative mechanisms or the other. Interest in one mechanism should investigate the agent rules for interaction at multiple levels of analysis and points in time, searching for a common dynamic at the beginning of the focal process that continues to the final outcomes of interest – this will be indicated by a universal slope, α , of the power-law tail. Interest in a set of mechanisms should identify power laws at multiple units and levels of analysis, focusing on both top-down and bottom-up rules and inputs like the type discussed in the previous section – for each distribution, significance can be tested with Komolgorov-Smirnov statistics or p-values. Similarly, studying extreme outcomes in the tail of the distribution with multiple qualitative field studies or a hermeneutics study can provide insight and theoretical grounding.

Boisot and McKelvey (2010:428) maintain that “research must engage with the power-law distribution as a whole, without privileging one particular region at the expense of another.” This may not necessarily be required, however. Higher-level (i.e., mid-range) theory building must account for the emergence of extreme outcomes in the domain, but instead of amassing all outcomes together, a lower-level theory may focus on one particular region of the distribution. It is especially important to specifically state any theoretical assumptions and boundary conditions that apply to different treatments of the data in different regions of the distribution. Either way, scholars must start with the interdependence of observations as the null hypothesis unless proven otherwise. Gaussian methods and assumptions should only be used if the null is rejected.

CONCLUSION

Andriani and McKelvey (2000:16) say that “no statistical finding should be accepted into organization science if it gains significance via some assumption-device by which extreme events and (nearly) infinite variance are ignored,” and O’Boyle and Aguinis (2012) assert that all existing theories of individual and organizational performance that have been tested using Gaussian techniques must be revisited. Both of these assertions, in my view, are too restrictive. However, scholars who ignore or disobey power laws in their empirical and theoretical studies of organization design do so at the peril of invalidity. Whereas it is of primary importance to account for power laws (and their generative mechanisms) while building and testing theories of organization design, the most rigorous, robust, and practically relevant theories will be crafted abductively. Here, stylized empirical facts from multiple sources – inductive field studies and deductive data analyses – can be integrated with computational simulation models of the generative process to develop explanations not only of what is but of what might be. This will provide scholars with the ability to develop prescient theory that can both explain the past and foretell the future (Corley & Gioia, 2011).

REFERENCES

- Andriani P, McKelvey B. 2009. From Gaussian to Paretian thinking: causes and implications of power laws in organizations. *Organization Science* 20: 1053-1071.
- Boisot M, McKelvey B. 2010. Integrating modernist and postmodernist perspectives on organizations: a complexity science bridge. *Academy of Management Review* 35(3): 415-433.

- Carroll GR, Harrison JR. 1998. Organizational demography and culture: insights from a formal model and simulation. *Administrative Science Quarterly* 43(3): 637-667.
- Corley KG, Gioia DA. 2011. Building theory about theory building: what constitutes a theoretical contribution? *Academy of Management Review* 36(1): 12-32.
- Crawford GC. 2012. Emerging scalability and extreme outcomes in new ventures: power-law analyses of three studies. In L. A. Toombs (Ed.), *Proceedings of the Seventy-Second Annual Meeting of the Academy of Management*: ISSN 1543-8643.
- Greene WH. 2007. *Econometric Analysis*. 6th ed. Prentice-Hall, Englewood Cliffs, NJ.
- O'Boyle EH, Aguinis H. 2012. The best and the rest: revisiting the norm of normality of individual performance. *Personnel Psychology* 65: 79-119.
- Popper K. 1963. *Conjectures and Refutations*. Routledge and Keagan Paul, London, UK.
- Schoonhoven CB. 1981. Problems with contingency theory: testing assumptions hidden within the language of contingency "theory". *Administrative Science Quarterly* 26(3): 349-377.
- Simon HA. 1955. On a class of skew distribution functions. *Biometrika* 42(3/4): 425-440.
- Simon HA. 1968. On judging the plausibility of theories. In B. Van Rootselaar and J.F. Staal (Eds.), *Logic, Methodology, and Philosophy of Science III*. North-Holland, Amsterdam.
- Starbuck WH, Nystrom PC. 1981. Why the world needs organizational design. *Journal of General Management* 6(3): 3-16.
- West BJ, Deering B. 1995. *The Lure of Modern Science: Fractal Thinking*. World Scientific, Singapore.

G. CHRISTOPHER CRAWFORD

Ph.D. Candidate
Department of Management and Entrepreneurship
University of Louisville
E-mail: christopher.crawford@louisville.edu