THE LAW OF GENIUS AND HOME RUNS REFUTED

JOHN DINARDO and JASON WINFREE*

In a lively, provocative article, DeVany claims inter alia that the size distribution of home runs follows a continuous "power law" distribution which is nested in a larger class of "stable" statistical distributions characterized by an infinite variance. He uses this putative fact about the size distribution of home runs to argue that concern about the use of steroids to enhance home run ability is necessarily misplaced. In this article, we show that the initial claim is false and argue that the subsequent claim about the potential importance of steroid use does not follow from the first. We also show that the method used to establish that the size distribution of home runs is characterized by an infinite variance is unreliable and will find evidence "consistent" with infinite variance in all but the most trivial of data sets generated by processes with finite variance. Despite a large and growing literature that spans several fields and uses methods and arguments similar to DeVany's, we argue that mere inspection of the unconditional distribution of some human phenomenon is unlikely to yield much insight. (*JEL* C16, L83)

I. INTRODUCTION

"Empirical regularities in biology, as in other fields, can be extremely interesting. In particular, such regularities may suggest the operation of fundamental laws. Unfortunately, apparent regularities sometimes cannot stand up under close scrutiny" (Solow, Costello, and Ward 2003).

A lively, provocative article by DeVany (2007) in *Economic Inquiry* argues that:

- "the statistical law of home run hitting is the same as the laws of human accomplishment developed by Lotka . . ., Pareto . . ., Price . . ., and Murray . . .,"
- "there is no evidence that steroid use has altered home run" hitting,
- "the greatest accomplishments in [science, art, and music] all follow the same universal law of genius," and
- "the stable Paretian model developed here will be of use to economists studying extreme accomplishments in other areas,"

which apparently follows from his claim that the size distribution of annual home run production has a finite mean but infinite variance and follows a "power law distribution." DeVany's argument is not unique: it is part of a large and growing literature where claims of the ubiquity of power laws are legion.[1]

DeVany takes the additional step of connecting this statistical analysis to an argument about the effect of steroids on home run

*Data and programs can be obtained by e-mailing jwinfree@umich.edu. We would like to thank Benjamin Keys, Thomas Buchmueller, Ron Mittelhammer, the editor, and two anonymous reviewers for helpful comments. No steroids were consumed during the production of this article.

DiNardo: Professor, School of Public Policy and Department of Economics, University of Michigan and NBER, Ann Arbor, MI 48109-3091. Phone 734-647-7843, Fax 734-763-9181, E-mail jdinardo@umich.edu

Winfree: Assistant Professor, Program in Sport Management, University of Michigan, Ann Arbor, MI 48109-2013. Phone 734-647-5424, Fax 734-647-2808, E-mail jwinfree@umich.edu

1. A recent survey by Newman (2005) cites evidence that a diverse number of things allegedly "follow power law" distributions including "city populations, the sizes of earthquakes, moon craters, solar flares, computer files, wars, the frequency of use of words in any human language, the frequency of occurrence of personal names in most cultures, the number of papers scientists write, the number of citations received by papers, the number of hits on web pages, the sales of books, music recordings and almost every other branded commodity, the numbers of species in biological taxa, and people's annual incomes."

ABBREVIATIONS

CCDF: Complementary Cumulative Density Function
CLE: Central Limit Theorem
MLE: Maximum Likelihood Estimates
SOC: Self-Organized Criticality

Economic Inquiry
(ISSN 0095-2583)
Vol. 48, No. 1, January 2010, 51–64

doi:10.1111/j.1465-7295.2008.00176.x
Online Early publication November 10, 2008

hitting by major league ballplayers: "steroid advocates have to argue that the new records are not consistent with the law of home runs *and*, that the law itself has changed as a result of steroid use."

Our purpose of this article is to suggest that the above should be met with a fair amount of skepticism.

First, we try to provide some background on previous attempts to identify the existence of universal laws and provide the intellectual context for DeVany's claims.

Second, we show that DeVany's claims follow from a flawed statistical inference procedure. His procedure, with probability 1, would find evidence consistent with "infinite variance" for virtually any nontrivial data set. To do so, we first analyze the size distribution of a quantity, which could not follow a power law distribution and show that using DeVany's inference procedure, we would be led to the same (incorrect) claim. We also discuss the important distinction—elided by DeVany and others writing in related literatures—between an unconditional distribution and a conditional distribution.

Third, we observe that the size distribution of home runs *cannot* follow a power law distribution and show that the posited class of distributions provide an inadequate approximation to the data, *at best*.

Fourth, while concurring with DeVany's implicit criticism that "steroid advocates" who rely on recent "trends" to substantiate their views have not made their case, we suggest that the problem is that the question is ill posed. The level and distribution of total home runs in any given year is minimally a function of hundreds of things: the quality of pitching, the weather, the introduction of new ball parks, the number of games played, the distribution of baseball talent across the teams, and so forth. To claim that only one "cause" is responsible for a trend involves some (possibly unstated) assumption about the myriad of other factors. Indeed, what is sauce for the goose is sauce for the gander: those seeking to support *or* deny the claim that increased use of steroids have led to increased home run hitting will have to employ considerably more "shoe leather" than mere statistical analysis of the unconditional distribution of home runs per player or time trends in home run hitting.

We conclude by observing that neither examination of time trends in annual home run production nor examination of the unconditional distribution of home runs will settle the dispute between "steroid advocates" and "steroid opponents" and that more convincing evidence will have to be sought elsewhere.

II. DOES A POWER LAW IMPLY "SELF-ORGANIZING CRITICALITY" AND SO FORTH?

We are not the first to argue that claims about universal laws should be met with some skepticism. Indeed, our criticisms are depressingly familiar.[2]

> The stringency with which the goodness of a fitted model should be assessed depends to a degree on the claims that are being made about the model. The claim that a model is correct, as opposed merely to providing a useful approximation, should be subjected to particularly close scrutiny. Such claims have been made about the power law model for size-frequency data without adequate scrutiny. (Solow, Costello, and Ward 2003)

Claims about the ubiquity of statistical distributions have a long history. A classic example is from Feller (1940).

> The logistic distribution function ... may serve as a warning. An unbelievably huge literature tried to establish a transcendental "law of logistic growth"; measured in appropriate units, practically all growth processes were supposed to be represented by a function of [a particular distributional form] Lengthy tables, complete with chi-square tests, supported this thesis for human population, for bacterial colonies, development of railroads, etc. Both height *and* weight of plants and animals were found to follow the logistic law even though it is theoretically clear that these two variables cannot be subject to the same distribution. Laboratory experiments on bacteria showed that not even systematic disturbances can produce other results. Population theory relied on logistic extrapolations (even though they were demonstrably unreliable). The only trouble with the theory is that not only the logistic distribution but ... other distributions can be fitted to the *same material with the same or better goodness of fit.* In this competition the logistic distribution plays no distinguished role whatever; most contradictory theoretical models can be supported by the same observational material.
>
> Theories of this nature are short-lived because they open no new ways, and new confirmations of the same old thing soon grow boring. But the naive reasoning as such has not been superseded by common sense, and so it may be useful to have an explicit demonstration of how misleading a mere goodness of fit can be. Feller (1940) as cited in Brock (1999).

2. See Keller (2005) for a useful review of some of the history.

Brock (1999), cited by DeVany, cites Feller to warn economists and others against making precisely the types of claims DeVany makes:

> I will make the general argument here that, while useful, these "regularities" or "transcendental laws" must be handled with care because . . .most of them are "unconditional objects" i.e. they only give properties of stationary distributions, e.g., "invariant measures," and, hence, can not say much about the dynamics of the stochastic process which generated them. To put it another way, they have little power to discriminate across broad classes of stochastic processes.

Even active researchers in the area have begun to observe "that research into power laws . . . suffers from glaring deficiencies" (Mitzenmacher 2006). Nonetheless, a long history of researchers making extravagant claims about phenomenon that derive from the resemblance of their size distribution to some statistical distribution has not slowed down the making of the claims. Feller's (1940) rejection of "universal models of growth," Solow, Costello, and Ward's (2003) rejection of power laws in biology, Miller and Miller and Chomsky's (1963) rejection of the usefulness of Zipf's law of word length (Zipf 1932) are a few examples of prior (apparently failed) attempts to raise the level of discourse and raise the quality of attempts to "validate" or subject such theorizing to "severe testing" (Mayo 1996).[3]

Our argument is complicated by at least two issues:

1. DeVany argues that "steroid advocates" are wrong. Unfortunately, he cites no one actually making the claims he attributes to such advocates.

2. DeVany makes claims about the size distribution of home runs and refers vaguely to notions of "self-organized criticality" (SOC) without spelling out the implications of such notions for hypotheses about the effect of steroids on home run hitting.[4]

An important concern, which we address in Sandpiles, SOC, and Home Runs? section, revolves about (2). What is the law of genius? How would we know if some phenomenon was subject to such a law? Indeed, what does it mean to say, as DeVany does, that home runs are "more like the movies . . . or . . . earthquakes . . . than dry cleaning?"

A useful introduction to "complexity theory" for economists can be found in Krugman (1996). And though we cannot recapitulate the logic entirely, we sketch the notion of SOC which we believe is key to understanding the implicit argument DeVany makes. Only then is it possible to understand why some might find it plausible to assert that "the law of home runs" might look something like "the law of earthquakes" and why such an assertion might lead some to suggest that "steroids don't matter."

Sandpiles, SOC, and Home Runs?

To place both our arguments and DeVany's in context, it would be most helpful to provide a comprehensive review of some of the arguments made by students of "self-organizing" or "complex" systems which lie at the heart of some of DeVany's analysis. We cannot obviously do that here.[5]

3. It is routinely claimed that the putative fact that size distribution of word lengths follows Zipf's implies something important about language, for example, Li (1992) observes that "probably few people pay attention to a comment by Miller in his preface to Zipf's book [(Miller 1965)] . . ., that randomly generated texts, which are perhaps the least interesting sequences and unrelated to any other scaling behaviors, also exhibit Zipf's law." See also Perline (1996) for an enlightening discussion.

4. We would hasten to add that this omission may be for no other reason than editorial constraints as DeVany cites some of the relevant literature.

5. Krugman (1996) provides a sober yet optimistic discussion of this approach from an economist's perspective. For an enthusiastic appraisal and simple introduction, see Bak (1996) or Bak and Chen (1991). Krugman (1996) identifies three components of the complex system:

> "1. Complicated feedback systems often have surprising properties.
> 2. Emergence—[situations in which] large interacting ensembles of individuals [or neurons, magnetic dipoles, . . .] exhibit collective behavior very different from [what one might have] expected by simply scaling up the behavior of the individual units.
> 3. Self-organizing systems: systems that, even when they start from an almost homogeneous or almost random state, spontaneously form large scale patterns."

As Krugman observes, these components, especially the first two, are not unique to complex systems. The standard general equilibrium model, for example, can be described as displaying complex feedback (everything depends on everything else). As to "emergence," it is possible to view the Pareto optimality as "emergent" behavior generated by self-interested agents. Despite having some of the features associated with complex systems, neither of these would usually be viewed as examples of "complex systems." (We do not mean to suggest that complex systems have not been developed or used by economists. An example of a classic model exhibiting all three components [and generally considered to be an example of this approach] is Schelling's famous model of segregation, Schelling 1969, 1978.)

Instead, we think we can convey much of the implicit logic at the core with a short description of the canonical example of a system displaying self-organizing criticality—Bak's sandpile (Bak 1996; Bak, Tang, and Wiesenfeld 1988; Bretz et al. 1992; Nagel 1992; Winslow 1997).

Tesfatsion (2007) provides a nice intuitive explanation, which covers most of the important points:

> When you first start building a sand pile on a tabletop of finite size, the system is weakly interactive. Sand grains drizzled from above onto the center of the sand pile have little effect on sand grains toward the edges. However, as you keep drizzling sand grains onto the center, a small number at a time, eventually the slope of the sand pile "self organizes" to a critical state where breakdowns of all different sizes are possible in response to further drizzlings of sand grains and the sand pile cannot grow any larger in a sustainable way. Bak refers to this critical state as a state of self-organized criticality (SOC), since the sand grains on the surface of the sand pile have self-organized to a point where they are just barely stable.
>
> What does it mean to say that "breakdowns of all different sizes" can happen at the SOC state?
>
> Starting in this SOC state, the addition of one more grain can result in an "avalanche" or "sand slide," i.e., a cascade of sand down the edges of the sand pile and (possibly) off the edges of the table. The size of this avalanche can range from one grain to catastrophic collapses involving large portions of the sand pile. The size distribution of these avalanches follows a power law over any specified period of time T. That is, the frequency of a given size of avalanche is inversely proportional to some power of its size, so that big avalanches are rare and small avalanches are frequent. For example, over 24 hours you might observe 1 avalanche involving 1,000 sand grains, 10 avalanches involving 100 sand grains, and 100 avalanches involving 10 sand grains ….
>
> At the SOC state, then, the sand grains at the center must somehow be capable of transmitting disturbances to sand grains at the edges, implying that the system has become strongly interactive. The dynamics of the sandpile thus transit from being purely local to being global in nature as more and more grains of sand are added to the sandpile (Tesfatsion 2007).

Stipulating to this being an accurate description of avalanches in sandpiles[6] and stipulating to the ubiquity of such SOC in diverse fields and situations, some of the leaders in this field have drawn some rather wide-ranging implications for science or social science.

> If this picture is correct for the real world, then we must accept instability and catastrophes as inevitable in biology, history, and economics. Because the outcome is contingent upon specific minor events in the past, we must also abandon any idea of detailed long-term determinism or predictability. *Large catastrophic events occur as a consequence of the same dynamics that produces ordinary events. This observation runs counter to the usual way of thinking about large events, which …looks for specific reasons (for instance, a falling meteorite causing the extinction of dinosaurs) to explain large, catastrophic events.* Bak (1996, p. 32, emphasis added)

To put it yet a different way, the sandpile forms, experiences avalanches, and so forth as a consequence of a *single* causal process. Great catastrophes arise from the *identical* mechanism as the periods of noncatastrophes.

We think DeVany means to make a similar argument regarding the production of home runs: home runs are the "catastrophe" in a SOC process. Applying Bak's and DeVany's logic to home run production, we might be led to conclude that the process that produces a year with few home runs for an individual batter can be identical to the process that produces a year with an extremely large number of home runs. Moreover, a further hunt for causes for extreme events might be unwarranted.

As we discuss in detail below, this seems an unwise inferential leap. Even in the case of sandpiles, the fact that avalanches *can* arise from the same causes that generate periods of low avalanche activity does not necessarily imply that other causes are not or cannot be at work. We conjecture, for example, that the introduction of a typical 3 yr old with a plastic shovel into a sandpile laboratory might predictably lead to avalanches even in a system that until that time exhibited SOC. At a minimum, we doubt that many parents would accept without question a 3 yr old's denial of involvement with the sandpile avalanche on the grounds that he or she could not have caused the avalanche since the sandpile exhibited SOC—especially if the 3 yr old is observed in the vicinity of the avalanche with sand all over his or her clothes.

III. A POWERLESS POWER LAW TEST

The bulk of the statistical analysis in DeVany is in section 5 "The Distribution of Home Runs" and section 6 "The Law of Home Runs." The core of the statistical argument and upon

6. While such a process is rather easy to generate in a computer simulation (Winslow 1997), actual practice is quite different. In laboratory experiments with sandpiles, the sand and setup require a fair amount of tweaking to behave in the idealized way described above (Bretz et al. 1992; Nagel 1992).

which the subsequent statistical analysis rests is that the *unconditional* distribution of home runs hit in a year follows a so-called "stable distribution."[7] In particular, the claim is made that the distribution of home runs is characterized by a subset of this class of stable distributions in which the variance of home runs is infinite. Consequently, DeVany infers that "this makes it a 'wild' statistical distribution, far different from the normal (Gaussian) distribution that people are tempted to use in their reasoning about home runs and most other things. Things are not so orderly in home runs; they are rather more like the movies ... or earthquakes ... than dry cleaning."

How does DeVany establish that the size distribution of home runs follows a power law? The method is simple. Fit the data to a "stable" distribution and check whether the estimated parameters are consistent with a stable distribution with infinite variance. If so, conclude that the data are generated from a "wild" statistical distribution and follow the "universal law of genius."

> The exponent is a measure of the probability weight in the upper and lower tails of the distribution; it has a range of $0 < \alpha \leq 2$ and the variance of the stable distribution is infinite when $\alpha < 2$. The basin of attraction is characterized by the tail weight of the distribution (α). This remarkable feature tells us that the weight assigned to extreme events is the key distinguishing property of a stable probability distribution ....[8] The tails of a stable distribution are Paretian and moments of order ≥ 2 do not exist when $\alpha < 2$. This is typical of many extraordinary accomplishments, as seen in the works of Lotka, Pareto, and Murray. Its mean need not exist for values of $\alpha < 1$. When $\alpha = 2$, the stable distribution is the normal distribution with a finite variance. The parameter α is called the tail weight because it describes how rapidly the upper tail of the distribution decays with larger outcomes of the random variable; smaller implies a less rapid decay of probability.

FIGURE 1
Does *X* Follow A Power Law ?

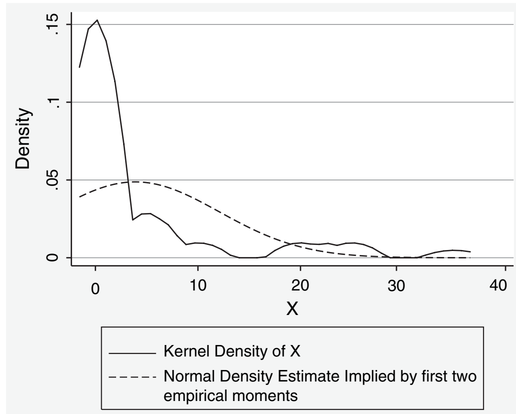

Put more simply, DeVany's procedure is estimate the four parameter stable distribution. If the estimated value of $\alpha < 2$ conclude that the distribution of home runs has infinite variance.

Does X Follow the Law of Genius?

To make clear why this analysis is problematic, we perform a similar analysis on a different random variable, which we call *X* for the moment. Following DeVany, we display a smoothed histogram of *X* (using the conventional "Silverman rule-of-thumb bandwidth") and compare it with the normal distribution implied by the empirical mean and variance of *X* in Figure 1.

As is true with size distribution of home runs, the size distribution of *X* is decidedly nonnormal. As with the home run data, the upper tail is poorly fit by the normal distribution. Table 1 repeats the more formal analysis in DeVany (2007). The table displays our estimates of the four parameters of the stable distribution by maximum likelihood using the same program as DeVany (2007) but using data *X*.[9] We display our results for *X* along side DeVany's results using the same data on individual home run hitting [DeVany (2007); Table 1].[10] While the distribution of *X* and the distribution of home runs are not

7. "Stable distributions are a rich class of probability distributions that allow skewness and heavy tails and have many intriguing mathematical properties" (Nolan 2007). One difficult aspect of these distributions is that, except in a few special cases, there exists no closed-form expression for the probability density and distribution functions.

8. The stable distribution has a total of four parameters. For the other three parameters, "... the skewness coefficient—$1 \leq \beta \leq 1$ is a measure of the asymmetry of the distribution. Stable distributions need not be symmetric; they may be skewed more in their upper tail than in their lower tail. The scale parameter γ must be positive. It expands or contracts the distribution in a non-linear way about the location parameter δ which is the center of the distribution" (DeVany 2007). Following DeVany, we limit our discussion to just the one parameter, α.

9. See Rimmer and Nolan (2005) for details.

10. Our estimates of the four parameters are identical to those estimated by DeVany (2007), although our calculated value of the maximized log likelihood function is somewhat larger than reported in the article.

TABLE 1
X versus the Home Run Data: Fitted to the "Stable" Distribution

	Index α	β	Scale	Location
DeVany's data	1.6422	1.00	6.219	12.30
X	1.07657	0.966409	1.28782	11.484

Notes: MLE estimates of the four-parameter stable distribution. The estimates in the first row replicate DeVany (2007) using home run data from 1950 to 2004. The estimates in the second row use data on variable "*X*"; see text for details.

identical, they both have the properties which are "consistent" with the random variable *X* possessing an infinite variance, namely that the estimated value of α (one of the four parameters of the stable distribution) is less than 2.

Does *X* follow a power law? No. We defined *X* as the number of mentions (times five) of the word "normal" or "normality" on a page of the web draft of DeVany (2007).[11] Surely, *X* does not possess infinite variance: presumably, the number of words that *Economic Inquiry* will allow to be printed on a page is finite; an author who proposed to submit an article including nothing but the words normal and normality would stand a low chance of having the article included in the journal.

Why is DeVany's procedure flawed? Most simply, observing that the estimated value of α is less than 2 can only be construed as evidence for an infinite variance conditional on the data actually following the stable distribution. Among *stable* distributions (Nolan 2007), those consistent with finite variances only occur on the boundary of the parameter space—when $\alpha = 2$. Since $\alpha \leq 2$—by definition—DeVany's procedure will always provide evidence for an infinite variance unless it reaches the boundary. Putting aside the considerable difficulties in maximum likelihood estimation when the true value of the parameter lies on the boundary of the parameter space, even a variable "just shy of normality"—that is, $\alpha = 1.\overline{9}$—is consistent with an infinite variance. More importantly, if the data are not from the stable distribution—say, uniformly distributed, exponentially distributed, and so forth—such a procedure will almost surely result in an estimated value of $\alpha < 2$.

IV. OTHER PROBLEMS WITH THE ANALYSIS

There are other significant problems with the analysis in DeVany (and in much of literature which purports to have found evidence for the workings of "power laws"):

1. There is a failure to distinguish between conditional and unconditional distributions. If the number of at bats, for example, were allowed to follow a power law, the relationship between home run hitting and at bats could be nonstochastic, deterministic, and purely mechanical and the unconditional distribution would follow a power law. Home run hitting in such a situation would be more like dry-cleaning than "genius" despite the fact that the size distribution of home runs followed a power law.[12] Mere inspection of the variance of the unconditional distribution of total home runs, in general, tell us *nothing* about whether steroids matter.

2. Like much of the literature, DeVany does not contemplate the possibility that the observed size distribution of home runs is a mixture of many different—individual—(nonpower law) statistical distributions. Hence, estimating the parameters of a single (falsely imposed) statistical distribution cannot, in general, be adequate for reliable inference about the potential existence of a "fundamental" law.[13,14] Perline (2005) shows that data are often cited as

11. The data we used were as follows:

Page	2	3	6	7	8	9	10	11	13	16	20	22	41	The other 32 pages
Number of mentions	1	1	5	7	5	1	4	2	1	1	2	4	1	0

We multiplied the number of mentions by five. The data were collected using the (undated) web draft which was created on June 14, 2006. For the kernel density estimate we used an Epanechnikov kernel and a bandwidth of 1.558. The normal density estimate used the sample mean of *X*, which was 3.89 and had a sample standard deviation of 8.18.

12. In our previous draft, we generated a toy example in which steroids improved performance and the unconditional distribution of home runs had infinite variance. In this example, when we conditioned on the number of at bats, the variance of home runs was either very finite or zero (DiNardo and Winfree 2007).

13. It is possible, however, that aggregation of objects following their own power law could itself produce another power law. See, for example, Gabaix (1999).

14. In our previous draft, for example, we observed that the assumption that total home runs are independently and identical distributed as a power law was clearly violated. In such a world, we would also expect that the individual with, say, the maximum home runs in a season would be essentially chosen at random from all players. Today's home run leader might be next season's zero home run hitter. A focus on a single unconditional distribution would, in general, ignore such difficulties.

following a power law, but a more careful look illustrates that the data are often a mixture of different distributions.

A. Power Law as "Law" and "Approximation"

In Section III, we documented the difficulties with DeVany's inference procedure as well as the more general problem of reasoning about the existence of a law from the unconditional distribution from a quantity. Thus far, we have argued that the inference *procedure* was faulty. Nonetheless, it remains possible that the *inference* drawn from such a procedure might be correct: a broken clock is still correct twice a day, as the adage goes.

Unfortunately, such is not the case here. The problem is so grave that it is considerable work to even contemplate a situation in which it might be reasonable to characterize the distribution of home runs as following a "power law" and hence having a tail that is "subject to bursts or avalanches." That is, the distribution of home runs *cannot* follow the distribution that DeVany posits and even if it could, he is not licensed to draw the inferences that he does about the nature of home run production.

We agree with criticism of related work on SOC that a serious problem with this literature is the unwillingness to put the proposition that an outcome follows a power law to even a minimally severe test. In a discussion, Solow, Costello, and Ward (2003) suggest that the problem with much of the power law literature in biology is the failure to evaluate the power of the power law against an explicit *alternative*.[15] Indeed, DeVany, following a tradition in the "complex studies" literature, considers no alternative to a distribution with infinite variance (except the stable normal distribution).

15. Specifically, they considered data from Yule (1925), an early proponent of a power law hypothesis. Yule is better known perhaps for his work in Economics where he documented a positive correlation between the degree of pauperism in a district and the generosity of provision of food for the poor; this was used to argue that there was a *causal* relationship between the generosity of such relief and the degree of pauperism in Yule (1899). See Freedman (1999) for a discussion. Yule used data representing the frequencies of genera of different sizes for snakes, lizards, and two Coleopterans (Chrysomelidae and Cerambycinae). When Solow, Costello, and Ward (2003) examined four of Yule's cases, they were able to reject the discrete power law distribution proposed by Yule (1925) versus a discrete nonparametric alternative in three of the four cases.

While we wholeheartedly concur with this critical judgment (and therefore compare a power law to other distributions), we wish to emphasize that in the *present case* such an analysis is superfluous: there are other even more insurmountable obstacles.

B. Why Home Runs are Immediately Inconsistent with a Stable Distribution

The most immediate problem is that the size distribution of home runs is *immediately* inconsistent with the posited distribution in DeVany (2007) even before approaching a systematic analysis of the data:

1. The number of home runs is bounded below by 0. Indeed, Figure 4 of DeVany displays only *part* of the estimated probability density function, that part where the number of home runs is greater than or equal to zero. The estimated power law distribution, if it were to be taken literally, predicts that 11% of baseball players would have a negative quantity of home runs. We think it safe to assume that negative home runs do not exist.
2. The number of home runs by a given player is discrete, not continuous as posited by the class of stable functions DeVany chose to estimate. No one will ever hit 1.2 home runs in a season. Somewhat surprisingly, DeVany makes a related observation regarding team production of home runs when he dismisses the "home runs per game statistic."[16]
3. If we are willing to assume that the number of games, at bats, and so forth in a given year is bounded from below by 0 and above by some arbitrarily large value $\overline{M}$ then it immediately follows that for any discrete distribution, the variance is bounded by $\frac{\overline{M}^2}{4}$, which is the variance of the Bernoulli distribution with equal-sized mass points at 0 and $\overline{M}$. Indeed, there is a long literature on establishing bounds for the variance of distributions with finite domain—for example, Muilwijk (1966), Gray and Odell (1967), Jacobson (1969)—where the bounds can be tightened

16. From DeVany (2007, p. 22): "If you think for a moment about the constraints of a ball game, it becomes obvious that home runs per game cannot be a well-behaved statistic that can be used to make sharp comparisons. The number of home runs in a game is an integer, not a continuous variable. The number of league games is an integer too. Dividing these numbers will give rational numbers, but they will not be distributed normally and will have strong modes at a few typical values."

TABLE 2
Maximum Likelihood Estimates of the Size Distribution of Home Runs per Player in Major League Baseball—1950–2004[a]

Distribution	Stable[b]	Stable[c]	Stable[d]	Negative Binomial	Discrete Power Law
Index (α)	1.6422	1.64221	1.64221		1.378
β	1.00	1	1		
Scale	6.219	6.21928	6.21928		
Location	12.30	12.3041	12.3041		
r				1.506172	
p				0.1141677	
Log likelihood	−39,294.2	−43,812.4	−43,217.8	−41,780.7	−47,552.8
Number of observations	11,992	11,992	11,992	11,992	11,552
Includes 0 Home Runs	Y	Y	Y	Y	N

[a]Version 5.3 of the data were obtained at http://baseball1.com/content/view/57/82/. Following DeVany (2007), we drop observations in the year 2005 or persons with less than 200 at bats. Therefore, all player-years with at least 200 at bats from 1959 to 2004 were in the sample. We note that the data also include multiple observations from some players in the same year if they played for multiple teams or had multiple "stints." This also implies that a player's home run total is only for a specific team for that year and not necessarily the entire season.

[b]Estimates reported in DeVany (2007).

[c]Estimates from using the Sloglikelihood command to calculate the maximum likelihood value in *Mathematica* (Rimmer and Nolan 2005).

[d]The maximized value of the log likelihood function is calculated by adding the log of the probability distribution function at each home run value observed in the data.

under various conditions (assumptions about symmetry, unimodality, etc.).[17]

4. There is no single description of a power law. Indeed, in the case of discrete variables, it is common to define a power law as a probability mass function (Newman 2005) such that:

$$f(x) \propto x^{-\alpha}. \quad (1)$$

This is of course problematic if x can take the value of 0. One may choose the expedient of focusing on observations above which exceed some threshold (and above 0) in the discrete case and describing the results as consistent with "the upper tail following a power law"[18] (even if the distribution above some threshold follows some other nonpower law distribution[19]), but an estimation procedure that allows one an extra degree of freedom to choose this threshold after looking at the data are obviously not going to be very powerful.[20]

C. Fitting Unconditional Distributions

Despite the substantial caveats we have enumerated, we present several different attempts at fitting single distributions to home run data in Table 2.

In the first four specifications, we consider the data *including* zeros. In the fifth, we conduct an analysis excluding data on individuals who hit no home runs ("excluding zeros").

17. N.B. The existence of bounds somewhere in the data generation process is not necessarily inconsistent with some version of a power law. For example, a random walk model of *growth* with a (lower) barrier could produce a size distribution consistent with Zipf's laws. See, for example, Gabaix (1999). More descriptively, accurate models would have to allow for the "birth" and "death" of new ballplayers.

18. The "Hill estimator" (Hill 1975) is one popular way to assess "upper tails." Consider the case when the upper tail of the distribution of some random variable x follows: $1 - F(x) = x^{-\alpha}L(x)$, where $L(x)$ is constant above some threshold. The Hill estimator of α uses only information from the highest k-order statistics from a sample of size $n - \xi_{n:n}, \ldots, \xi_{n-k:n}$. The estimator of α is given by: $H_k^{(n)} \equiv \sum_{i=1}^{l} k^{-1}\log_{i=1}^{k}\log(\xi_{n-i+1:n}) - \log(\xi_{n-k:n})$ where $1 \leq k < n$ (Haeusler and Teugels 1985). We are not aware of any attempt to evaluate the properties of this estimator when a researcher gets to choose k.

19. See Nolan (2007).

20. Perhaps *obvious* is not the correct word. See the useful discussion in Perline (2005) for a demonstration how a judicious choice of a lower truncation point can transform data generated by the most mundane of nonpower law distributions into data whose upper tail seems to follow a power law distribution. We also concur in his judgment that "Shoehorning the data into one- or two-parameter models, such as the Pareto or Yule or the lognormal, while simultaneously excluding some inconvenient portion of the distribution, has too long been the norm. Many of the examples of inverse power laws proposed through the years are probably FIPLs (False Inverse Power Laws) best represented by finite mixtures of distributions."

In the first column of the table, we present DeVany's estimates of the stable distribution. In the next two columns, we reproduce our estimates using two variants of the same *Mathematica* program used by DeVany to generate his results. In the next column, we report the maximum likelihood estimates of the two parameter negative binomial distribution.

Next, we repeat the exercise with a sample that excludes all individuals with zero home runs and present the results of fitting an appropriate version of a *discrete* power law.[21]

We draw several conclusions from this statistical analysis:

1. With the exception of the maximized value of the likelihood function, our estimates of the parameters are essentially identical to DeVany's estimates.[22]

2. Despite having four parameters, the stable distribution does a poor job of "fitting" the data. The negative binomial distribution, with only two parameters, for example, results in a higher value of the maximized log likelihood. If you were to believe that the stable distribution or the negative binomial distribution were the only two hypotheses to be considered, considered them equally likely (and were willing to overlook the negative and fractional home run predictions of the stable distribution) the "weight of the evidence" (Good 1981; Peirce 1878) would *still* be against the power law distribution.[23] Of course, if you were to allow other possibilities you would certainly reject the stable distribution and quite possibly the negative binomial distribution. We also illustrate this point in Figure 2 with a graph of the estimated stable distribution, negative binomial distribution, and the histogram of the data. Clearly, the negative binomial distribution estimates the actual distribution better than the stable distribution. It does not mistakenly predict negative home runs.

FIGURE 2

Empirical and Fitted MLE Estimates of Probability Density Function

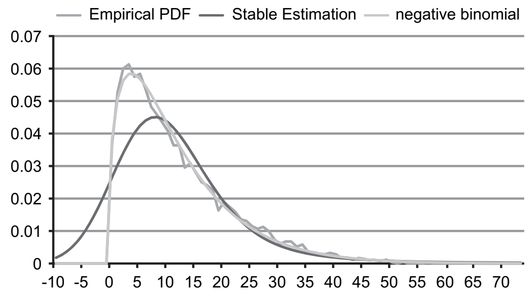

3. The situation looks no better when we focus just on the positive observations. As before, the weight of the evidence is against the power law distribution.

We would like to stress that the problem is not unique to DeVany:

> While the arguments found in the statistics literature concerning the use of scaling distributions for modeling high variability/infinite variance phenomena have hardly changed since Mandelbrot's attempts in the 1960s to bring scaling distributions into mainstream statistics, discovering and explaining strict power law relationships has become a minor industry in the complex science literature. Unfortunately, a closer look at the fascination within the complex science community with power law relationships reveals a very cavalier attitude toward inferring power law relationships or strict power law distributions from measurements. (Willinger et al. 2004)

Figure 3, taken directly from DeVany (2007)[24], is a case in point. As he describes

21. In a previous draft, following Solow, Costello, and Ward (2003), we also estimated the parameters of an alternative class of distributions that has declining tails. Specifically, we fit the size distribution of home runs subject only to the constraint that frequency with which individuals hit a specific number of home runs is nonincreasing in the number of home runs. That is, if p_k is the probability of a player hitting k home runs and n is the highest number of home runs that can be hit in a season,

$$(2) \qquad p_n < p_{n-1} - \ldots < p_k < p_{k-1} < p_{k-2} < \ldots < p_0.$$

This is a more severe test than the negative binomial since the class of nonpower law alternatives implicitly considered is larger. It, indeed, did fit the data better than the discrete power law distribution.

22. We corresponded briefly with Professor DeVany on the subject. We have not been able to determine the source of the discrepancy in the estimate of the maximized value of the log likelihood function, but it may be a typographical error or a different version of Mathematica given the almost exact correspondence between his and our estimates of the parameters of the distribution and the fact that we appear to be using the exact same data set (judging by the number of observations and sample means DeVany reports).

23. When comparing only two statistical hypotheses, the difference in the value of the log likelihood function can be interpreted as the (Bayesian) posterior log odds ratio if the initial probabilities attached to the two possibilities were .5.

24. It is labeled as Figure 5 in his article.

FIGURE 3
Plot of Fitted Stable Distribution Estimates and Empirical Cumulative Density Function from DeVany (2007)

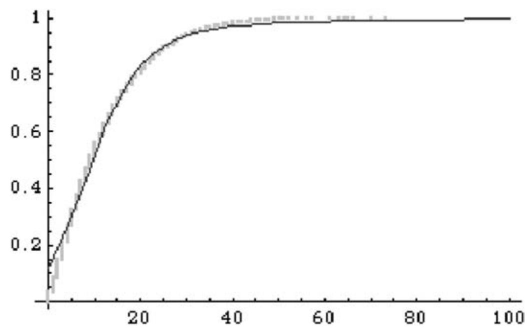

it, this figure displays a "remarkable" fit of the "cumulative theoretical and empirical distributions." One need not cavil about the definition of remarkable to demonstrate that with a more appropriate metric of "fit," the home run data are not well approximated by a power law.

The problem with DeVany's figure is, as Willinger et al. (2004) demonstrates, that such a display is quite powerless; with such a plot, it is difficult to distinguish power law from non-power law data or discriminate among power laws (i.e., different values of α).[25] Even if we stipulate to "ignoring the zeroes," it is easy to generate a more powerful visual test of the proposition.

One aspect of "self-similarity"—as this property is referred to in the complex systems literature[26]—is that the definition in Equation 1 implies the "complementary cumulative density function" (CCDF) is *linearly* related to size:

$$(3) \quad \log(1 - P(x \leq x_0) \approx a - \alpha \log(x_0),$$

where $(1 - P(x \leq x_0)) \equiv P(x > x_0)$ is the CCDF or one minus the cumulative probability of hitting at least x_0 home runs. The approximation becomes exact as $x_0 \to \infty$. This property suggests a useful visual display to assess the fit of the data to a power law: one merely plots the natural logarithm of the CCDF against the log of size. This particular display highlights the fit (or lack of fit) in the *tails* of the distribution and makes it relatively easy to distinguish the fit of the tail to different choices of α. Often, researchers use a rank-size plot, or quantile-quantile plot, to examine the fit of a power law distribution. Although the rank-size plot is a useful tool in some settings, one advantage of the log CCDF is that it is more powerful at detecting the goodness of fit (or lack thereof) in the tail of the distribution.

As Figure 4 demonstrates, the power law provides a poor approximation globally and in the tails of the distribution. The most appropriate power law—the simple discrete version of the power law—gives the *worst* fit to the data, globally and in the all-important tail. The "inappropriate" power law (the continuous stable version) gives a slightly better fit but fits quite poorly in the tail. The negative binomial distribution—which is as well behaved as it is possible for a distribution to be—seems more deserving of the moniker remarkable than the power law distributions in terms of quality of fit.

Figure 4 also helps explain why it is much easier to find a power law if one is allowed to characterize *part* of the distribution *that one chooses after the fact* as being a power law: it is easy to convince oneself that even a very convex shape is linear if one can systematically ignore part or most of the curve.[27]

There is another, informal, yet instructive way to evaluate how well the continuous stable distribution works as "the law of genius." Under the hypothesis that the fitted continuous stable law distribution is correct, we can use estimates from the cumulative density function to generate predictions for the number of genius home run hitters we should have expected to see over the period from 1959 to 2004. We can also do the same with our "dry-cleaning" distribution, the negative binomial distribution.

For example, according to the negative binomial distribution, the expected number of players that would have hit 100 or more home runs is 0.23; the expected number who

25. A frequently employed test in this literature employs variations of the *Q-Q* plot, which are also problematic. See Willinger et al. (2004).

26. DeVany discusses this briefly in section 10 and on page 11: "[if the distribution is from the stable distribution] this implies that any way you look at the process you should [see] that the distribution has the same shape."

27. The fit of the continuous stable distribution fit to the nonzero observations is no better than that produced by using all the observations as DeVany does and to avoid clutter, it is dropped from Figure 4.

FIGURE 4
Log-Log Plot of Complementary Cumulative Distribution Function of Home Runs

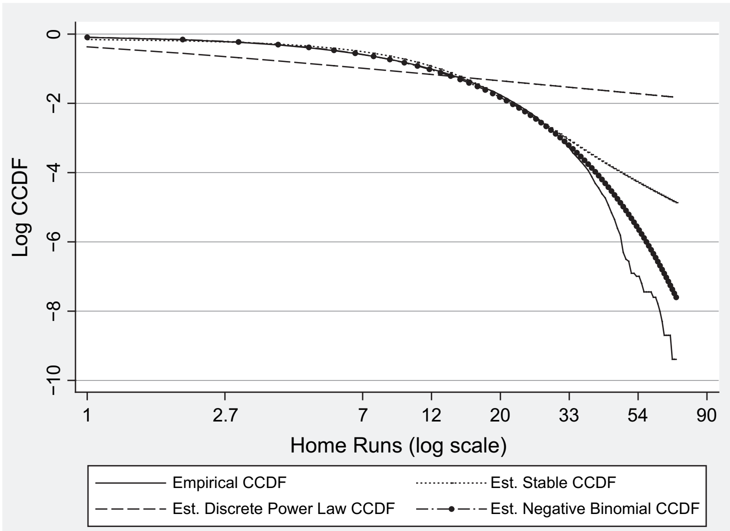

would hit more than 1,000 according to the same estimates is essentially 0 (the all-time record for home runs in a season is 73).

This is arguably a sign of bad fit for the negative binomial distribution. However bad the fit, the continuous stable law fits remarkably worse; our estimates from that distribution suggest that there should have been more than 48 players to hit 100 home runs or more. Again, according to that same distribution, we would expect .88 players to hit 1,000 home runs or more. Worse yet is the estimated discrete power law distribution: by that distribution, we would have expected to see 1,709 players hit 100 or more home runs and 716 players would have to hit 1,000 or more home runs. This would be unthinkable to most baseball fans, especially since record for *at bats* is 716. (An important distinction between the continuous and the discrete versions of the power law distributions being discussed is that the former has more parameters. It is not surprising it fits better, even for discrete data.)

We hasten to add that although the news is unremittingly bad for the power law distribution that we and DeVany have estimated, we do not mean to suggest that we believe that any of our alternative distributional choices are realistic or even particularly useful. Moreover, even if some "fix" of the sample or estimation procedure were to lead to a proper statistical test that could not reject some subset of the data from following a power law, none of DeVany's other inferences about steriods, and so forth would be warranted. In the context of this type of problem, the whole idea of fitting a parametric model of the size distribution of home runs seems like a really bad idea (except perhaps as a "quick and dirty" way to communicate some features of the data). Like any human endeavor (and much else), home run hitting is a process so ill-understood that it would be a miracle if any simple parametric model (such as the stable distribution) were able to characterize it.[28]

Apropos of why one would expect some outcome to be distributed as a power law or some other class of distributions, it is also

28. Indeed, one of the serious problems with the power law hypothesis is that it would be difficult to learn about without enormous amounts of data. The wild distributions discussed by DeVany take their character from the extreme tails of the distribution. Such phenomenon are consequently "rare" and therefore quite difficult to learn about. Heyde and Kou (2004), for example, observe that there are good reasons to doubt simple comparisons of likelihood in this context. In part, this is a problem because of the importance of correctly characterizing the tails of the distribution. A sharp ability to discriminate between a tail following a power law distribution and a tail following an exponential distribution generally requires enormous amounts of data, at a minimum (Heyde and Kou 2004).

important to remember that the ubiquity of the normal distribution in statistical analysis does *not* arise because the characteristics of the *objects of study* are distributed normally—rather they often follow because we are studying systems that can be well approximated by "chance set ups" (Hacking 1965)—the randomized controlled trial is the canonical example of such a set up—and the *sample means* of such a process can be shown, by some variant of the Central Limit Theorem,[29] to be approximately normal even when the outcomes under consideration are *not* distributed normally, as long as the outcomes have finite variance.[30] As we discussed in Section IV, assessing whether an outcome has a finite variance can often be established merely by demonstrating that outcome is bounded.

V. CONCLUDING REMARKS

As we have argued thus far, mere inspection of the size distribution of a random variable is insufficient to draw any conclusions about the *process* generating the data. While some distributions might allow for a rough approximation of the data, and this may be sufficient for some purposes, an *approximation* is *not* adequate for the purpose of drawing some sort of "causal" inference. That is, it may be fair to say that Zipf's power law—$P(\text{size} > S) \propto \frac{1}{S}$—provides a rough approximation to the size distribution of cities (Gabaix 1999), but quite another (inappropriate) matter to infer anything about the mechanism of city growth directly from that fact. To take one example from economics, Gabaix (1999) demonstrates that the mechanisms that could induce a Zipf's law for cities could be very different and result in very different inferences: "[although] the models [might be] mathematically similar, they [may be] economically completely different."[31]

Though not the focus of this article, we do believe that there may be a link between steroids and home runs. Indeed, numerous players have admitted and/or tested positive for performance-enhancing drug use. Moreover, many other players have been implicated with steroid use with varying degrees of evidence. For example, the Mitchell report, sponsored by Major League Baseball, linked 88 players to performance-enhancing drugs. Guilt was not proved by the Mitchell report, but it made clear that steroid use was not rare in Major League Baseball. Many of these same players have had unprecedented seasons; we doubt this is mere coincidence. Casual observation would also suggest that over the past 10 or 15 yr, there has been an increase of players hitting a large number of home runs at surprisingly older ages.

That is, we are sympathetic to the idea that, for certain ballplayers, it is possible that judicious use of steroids may contribute to some (possibly temporary) increase in home run hitting. It is important to stress, however, we have *not* made this case here. Our most important point is that proof of such a claim would take a great deal more work than mere casual inspection of statistics on home run hitting than we have engaged in here.[32] A higher standard of evidence is needed to establish or refute such a claim.

As we have demonstrated, *none* of the statistical analysis provided in DeVany (2007) speaks to the claim that the causal impact of the judicious use of steroids on home run hitting is zero. Inferring the existence of fundamental causal laws—that is, the law of genius—from the statistical distribution of some outcome is difficult, at best.

The view that aspects of the human condition or human behavior could be summarized by autonomous statistical laws has a long and not entirely distinguished history. It is ironic, given the aspersions cast on the normal distribution in DeVany (2007), that Galton's explorations into the normal distribution were in part motivated by a quest similar to DeVany's—to explain the "exceptional" and

29. In fact, the notion of Lévy stable or stable distribution is so named since it is a CLT of sorts for variables with infinite variance (Gnedenko 1943, Gnedenko and Kolmogorov 1954).

30. Alternatively, if the log likelihood of the data generation process is approximately quadratic with a constant Hessian, it can be shown that the maximum likelihood estimator of a quantity is approximately normal (Geyer 2005; LeCam 1986; LeCam and Yang 2000).

31. Indeed, Gabaix (1999) states simply that "economic models [for describing the size distribution of cites] have been inadequate." See also Krugman (1996).

32. N.B. It is certainly possible that even if some players could benefit from judicious use of steroids, it could also be true that some people would not benefit. In the terminology of the "treatment effect" literature, there might well exist "treatment effect heterogeneity." The effect of steroids on a typical major league baseball player might be positive, while the effect of steroids on the home run production of the authors of this article might well be zero or negative.

FIGURE 5
Galton's Quincunx from *Natural Inheritance*

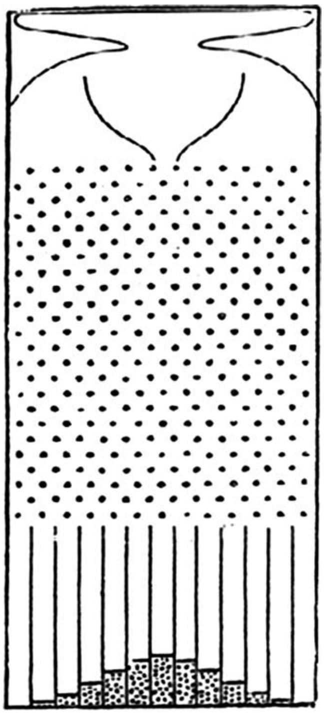

"human genius."[33] Galton worried about breeding mediocrity. Others took the existence of apparently stable (i.e., nonchanging) distributions as vitiating free will.[34] Indeed, using different language, Galton (1892) was among the first to use simulation to display an "emergent" system. Galton's famous quincunx was a vertical board with equally spaced pegs and a hole at the top in which marbles could be placed. The marbles entered the top of the device and were allowed to fall randomly[35] to reach the bins at the bottom. A figure from his book (Galton 1894) is displayed in Figure 5. The normal distribution that resulted was described as "order out of chaos."[36]

Today, we think of it as a useful mechanical model of the normal distribution as the limiting distribution of the binomial and few would attribute any "deeper" rational for this behavior.

We believe it is fair to say that there has been no convincing evidence of the existence of any causal laws regarding any aspect of the human condition regulated by the normal distribution (or any other distribution) since such ideas were proposed in the nineteenth century. The class of stable distributions investigated by DeVany (2007) may prove to be an exception, although we think it quite unlikely. If, nonetheless, economists are to take up DeVany's suggestion that the "stable Paretian model developed here will be of use to economists studying extreme accomplishments in other areas," we can only hope such claims will be subject to far more rigorous scrutiny than they have up to this point. Until then, we think it is wise to treat such claims with great skepticism.

33. See the discussion, especially chapter 21, in Hacking (1990). It is no accident, for example, that one of Galton's most significant efforts was entitled "Hereditary Genius: An Inquiry into its Laws and Consequences" (Galton 1882).

34. See Hacking (1990) and DiNardo (2007) for discussion and citations. Hacking nicely summarizes one example of this view, which arose during the nineteenth century: "A problem of statistical fatalism arose. If it were a law that each year so many people must kill themselves in a given region, then apparently the population is not free to refrain from suicide."

35. A sufficient condition for the marbles to be distributed normally at the bottom of the quincunx is that when a marble arrives at any peg, each marble has the *same* probability of heading left *or* right.

36. Galton's description of the normal distribution ("Law of Frequency of Error") echoes language used to describe SOC. From Galton (1894): "*Order in Apparent Chaos*—I know of scarcely anything so apt to impress the imagination as the wonderful form of cosmic order expressed by the 'Law of Frequency of Error.' The law would have been personified by the Greeks and deified, if they had known of it. It reigns with serenity and in complete self-effacement amidst the wildest confusion. The huger the mob, and the greater the apparent anarchy, the more perfect is its sway. It is the supreme law of Unreason. Whenever a large sample of chaotic elements are taken in hand and marshaled in the order of their magnitude, an unsuspected and most beautiful form of regularity proves to have been latent all along" (p. 66).

REFERENCES

Bak, P. *How Nature Works: The Science of Self-Organized Criticality*. New York: Springer-Verlag, 1996.

Bak, P., and K. Chen. "Self-Organized Criticality." *Scientific American*, 1991, 264(1), 26–33.

Bak, P., C. Tang, and K. Wiesenfeld. "Self-Organized Criticality." *Physical Review A*, 1988, 38(1), 364–74.

Bretz, M., J. B. Cunningham, P. L. Kurczynski, and F. Nori. "Imaging of Avalanches in Granular Materials." *Physical Review Letters*, 1992, 69, 2431–34.

Brock, W. A. "Scaling in Economics: A Reader's Guide." Manuscript, Department of Economics, University of Wisconsin, Madison, 1999.

DeVany, A. Forthcoming. "Steroids, Home Runs and the Law of Genius." *Economic Inquiry*, 2007.

DiNardo, J. "Interesting Questions in Freakonomics." *Journal of Economic Literature*, 2007, 45, 973–1000.

DiNardo, J. and J. Winfree. "The Law of Genius and Home Runs Refuted." Unpublished draft, University of Michigan. 2007. Accessed August 23, 2007. http://www-personal.umich.edu/~jdinardo/lawsofgenius.pdf.

Feller, W. "On the Logistic Law of Growth and Its Empirical Verifications in Biology." *Acta Biotheoretca*, 1940, 5, 51–66.

Freedman, D. A. "From Association to Causation: Some Remarks on the History of Statistics." *Statistical Science*, 1999, 14, 243–58.

Gabaix, X. "Zipf's Law for Cities: An Explanation." *Quarterly Journal of Economics*, 1999, 114, 739–67.

Galton, F. *Hereditary Genius: An Inquiry into its Laws and Consequences*. 2nd ed. London: Macmillan, 1892. Accessed May 28, 2007. http://galton.org/books/hereditary-genius/. Web site maintained by Gavan Tredoux.

———. *Natural Inheritance*. Macmillan and Company, 1894.

Geyer, C. J. "Le Cam Made Simple: Asymptotics of Maximum Likelihood without the LLN or CLT or Sample Size Going to Infinity." Technical Report 643, School of Statistics, University of Minnesota, 2005.

Gnedenko, B. V. "Sur la distribution limite du terme maximum d'une serie aleatoire." *Annals of Mathematics (Second Series)*, 1943, 44, 423–53.

Gnedenko, B. V. and A. N. Kolmogorov. *Limit Distributions for Sums of Independent Random Variables*. New York: Addison-Wesley, 1954.

Good, I. J. "An Error By Peirce Concerning Weight of Evidence." *Journal of Statistical Computation and Simulation*, 1981, 13, 155–57.

Gray, H. L. and P. L. Odell. "On Least Favorable Density Functions." SIAM Review, October 1967, 9, 715–20.

Hacking, I. *The Logic of Statistical Inference*. Cambridge: Cambridge University Press, 1965.

———. *The Taming of Chance Number 17. In 'Ideas in Context.'* Cambridge, England: Cambridge University Press, 1990.

Haeusler, E., and J. L. Teugels. "On Asymptotic Normality of Hill's Estimator for the Exponent of Regular Variation." *Annals of Statistics*, 1985, 13, 743–56.

Heyde, C. C., and S. G. Kou. "On the Controversy Over Tailweight of Distributions." *Operations Research Letters*, 2004, 32, 399–408.

Hill, B. M. "A Simple General Approach to Inference About the Tail of a Distribution." *Annals of Statistics*, 1975, 3, 1163–74.

Jacobson, H. I. "The Maximum Variance of Restricted Unimodal Distributions." *Annals of Mathematical Statistics*, 1969, 40, 1746–52.

Keller, E. F. "Revisiting 'Scale-Free' Networks." *Bioessays*, 2005, 27, 1060–68.

Krugman, P. *The Self-Organizing Economy*. Cambridge, MA: Blackwell, 1996.

LeCam, L. *Asymptotic Methods in Statistical Decision Theory*. New York: Springer-Verlab, 1986.

LeCam, L., and G. L. Yang. *Asymptotics in Statistics: Some Basic Concepts*. 2nd ed. New York: Springer-Verlag, 2000.

Li, W. *Random Texts Exhibit Zipf's-Law-Like Word Frequency Distribution*. IEEE Transactions on Information Theory, 1992, 38, 1842–45.

Mayo, D. G. *Error and the Growth of Experimental Knowledge Science and Its Conceptual Foundations*. Chicago: University of Chicago Press, 1996.

Miller, G. A. "Introduction," in *The Psycho-biology of Language: An Introduction to Dynamic Philology*, by George Kinglsey Zipf. Cambridge, MA: MIT Press, 1965.

Miller, G. A, and N. Chomsky. "Finitary Models of Language Users," in *Handbook of Mathematical Psychology*, Vol. 2, edited by R. D. Luce, R. R. Bush, and E. Galanter. New York: Wiley and Sons, 1963, 419–91.

Mitzenmacher, M. "Editorial: The Future of Power Law Research." *Internet Mathematics*, 2006, 2, 525–34.

Muilwijk, J. "Note on a Theorem of M.N. Murthy and V.K. Sethi." *Sankhya, Series B*, 1966, 28 (Pt 1, 2), 183.

Nagel, S. R. "Instabilities in a Sandpile." *Reviews of Modern Physics*, 1992, 64, 321–25.

Newman, M. E. J. "Power Laws, Pareto Distributions and Zipf's Law." *Contemporary Physics*, 2005, 46, 321–53.

Nolan, J. P. *Stable Distributions—Models for Heavy Tailed Data*. Boston: Birkhäuser, 2007. Accessed 20 August 2008. http://academic2.american.edu/~jpnolan/stable/chap1.pdf

Peirce, C. S. "The Probability of Induction." *Popular Science Monthly*, 1878, 12, 705–18.

Perline, R. "Zipf's Law, the Central Limit Theorem, and the Random Division of the Unit Interval." *Physical Review E*, 1996, 54(1), 220–23.

———. "Strong, Weak and False Inverse Power Laws." *Statistical Science*, 2005, 20, 68–88.

Rimmer, R. H., and J. P. Nolan. "Stable Distributions in Mathematica." *Mathematica Journal*, 2005, 9, 776–89.

Schelling, T. C. "Models of Segregation." *American Economic Review*, 1969, 59, 488–93.

———. *Micromotives and Macrobehavior*. New York: W.W. Norton & Company, 1978.

Solow, A. R., C. J. Costello, and M. Ward. "Testing the Power Law Model for Discrete Size Data." *American Naturalist*, 2003, 162, 685–89.

Tesfatsion, L. "Introductory Notes on Complex Adaptive Systems and Agent-Based Computational Economics." Technical Report, Department of Economics, Iowa State University, 2007. Accessed 8 August 2008. http://www.econ.iastate.edu/classes/econ308/tesfatsion/bat1a.htm.

Willinger, W., D. Alderson, J. C. Doyle, and L. Li. "More 'Normal' than Normal: Scaling Distributions and Complex Systems," in *Proceedings of the Winter 2004 Simulation Conference*, Vol. 1, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith and B. A. Peters. Piscataway, NJ: IEEE Press Piscataway, 2004, 5–8.

Winslow, N. "Introduction to Self-Organized Criticality and Earthquakes." Discussion Paper, Department of Geological Sciences, University of Michigan, 1997.

Yule, G. U. "A Mathematical Theory of Evolution Based on the Conclusions of Dr. J.C. Willis." *Philosophical Transactions of the Royal Society of London, Series B* (Containing Papers of a Biological Character), 1925, 213, 21–87.

———. "An Investigation into the Causes of Changes in Pauperism in England, Chiefly During the Last Two Intercensal Decades (Part I.)." *Journal of the Royal Statistical Society*, 1899, 62, 249–95.

Zipf, G. K. *Selective Studies and the Principle of Relative Frequency in Language*. Cambridge, MA: MIT Press, 1932.