www.joim.com

Journal Of Investment Management, Vol. 3, No. 3, (2005), pp. 84–91

SURVEY OF THE LITERATURE

POWER LAWS

Sanjiv Das[a] *and Jacob Sisk*[b]

We provide a brief survey of two areas in finance in which power laws may play an important role—one, in better describing the tails of return distributions; and two, in market microstructure modeling. While the existing literature in finance is not extensive, we have surveyed a collection of papers in finance, as well as in other areas, attempting to highlight cross-disciplinary connections.

1 Introduction

We examine power law distributions that arise in two areas in finance. One, in distributions of prices and returns, and two, in the microstructure of information in markets. The role of power laws in finance has been relatively less explored. In this brief article, we summarize some of the literature relating to finance from other fields, and conjecture that power laws will play an important role in our understanding of markets in the future.[1]

Frequency distributions of financial quantities abound, but are rarely identified with power laws. We have been trained to believe deeply in the sanctity of the bell curve, that the normal distribution governs most of what we see, especially quantities in finance. For almost four decades, Benoit Mandelbrot, in repeated publications, reminded us that there was a good reason to examine power laws in finance. This line of work began with an early piece (Mandelbrot, 1963) to the most recent, an entire book on the subject (Mandelbrot and Hudson, 2004). Recent research seems more accommodating of his vision, finding more often that financial variables may indeed be distributed according to power laws. Power laws encompass a family of distributions, and include more specific versions, such as Pareto's law and Zipf's law.

The basic feature of a power law distribution is that quantities of small size appear with very high frequency, and large sizes appear with very low frequency. For example, in economics (Pareto, 1896), income distribution appears to be based on a power law. Hence, the probability p of a variable x is proportional to x^{-b}, where b is the exponent of the

[a]Santa Clara University, 500 EI Camino Real, Santa Clara, CA 95053.

[b]Yahoo! Search Marketing, 74 N Pasadena Ave, Pasadena, CA 91103.

distribution. Internet networks are characterized by power laws. The degree distribution of the network is now known to be based on a power law (the degree distribution comprises the frequencies of the number of nodes each node on the network is connected to). Power law distributions are characterized by linear log-log plots, because the logarithm of the probability is equal to

$$\log(p) = a - b\log(x)$$

where a is a scaling coefficient and b, as before, is the power coefficient or tail index which quantifies the rate at which the frequency scales with the size of x. Many of the power law distributions found in nature, physics, economics, sociology, and the world at large seem to have power coefficients between 1 and 3, i.e., $b \in (1, 3)$. Fama's (1965) early work arrived at b approximately of 1.7. Fatter tails than this are often found in related empirical work. Note that the probability function p is written quite generally here, and may be specified as either the probability density function or as the cumulative distribution function. In the latter case, it is also mentioned interchangeably as Pareto's law.

Power law distributions characterize what are known as "scale-free" networks, i.e., when the law describes the degree distribution of nodes in the network graph. From the probability function presented above, it is clear that there is no peak frequency in the center of the distribution, in the manner of a normal distribution, for example. The peak frequency occurs at the minimum point of the support of the distribution. There is thus no "scale" parameter to speak of, no "normal" member of the distribution. Hence, such distributions have no scale, and are known as scale-free.

Power law distributions are related to Zipf's law (1949). Zipf's law has very much the same flavor of a power law, but for the fact that the variable x refers to the rank of the observation. For example, the size of a city s may be related to its rank r, as follows: $s = ar^{-b}$ (see Gabaix, 1999, for a detailed analysis). Another example: in English, the probability of the rth most common word is $0.1/r$ for about the first 1000 words.

Pareto's law is the instance of the power law that refers to the CDF of the probability rather than the PDF. Hence, the probability that a variable is larger than x is proportional to x^{-b}. This version of the scaling law is useful when examining the tails of distributions, and as we will see, has been successfully used in the modeling of outliers in financial data. For another example, see Adamic and Huberman (2000) who show that the distribution of visits to web sites is governed by a Pareto power law. Pareto's law has also been shown to be consistent with a realm of "winner takes all" in economics.

What is often ignored is that many more natural phenomena seem to follow power law distributions rather than regular distributions such as the normal. Given that this scale-free property is ubiquitous in nature, and because economic variables are related to physical phenomena, it is hardly surprising that many of the economic variables we examine also tend to be distributed as per a power law. In the next sections we shall examine two different settings in which the statistical edifice of power laws plays a role in finance.

2 Prices and returns

The work of Mandelbrot (1963) was followed by that of Fama (1965), who found support for the fact that equity return distributions have tails that obey a power law.

An important characteristic of this finding is that the tails are invariant to scaling and to time aggregation. Such distributions are known as "stable" and in the context of our notation, they would be called b-stable, being labeled with the power coefficient.

Hence, returns R would be such that

$$\Pr[|R| \geq x] = ax^{-b}$$

where $a > 0$ and $b > 0$. Such a distribution is also said to be "self-similar" in the tails. What this means is that if we take any cut on the tail and examine it, the tail will decay at rate b. If we then take another slice of this tail, i.e., the tail of the tail, it will look and scale exactly as its parent. And so will the tail of the tail of the tail. This does not happen with distributions such as the normal. See Figure 1, where it is clear that the log-log plot for both, normal and lognormal distributions declines at an increasing rate, and is not linear, as would be the case for a power law distribution. The tails of the normal distribution decay much faster than power law distributions. Hence, they do not depict stock returns (at least the tails) very well. In reality, we have many more large outliers than we would expect to see from distributions that experience faster tail decay than that of a power law.

So, if it appears that Mandelbrot has been right all along, then why is his wisdom so far from being generally acceptable? While this is in and of itself an issue for discussion, there may be a good reason for this lack of acceptance—new evidence shows that he may not have been altogether correct on this issue. A recent paper by Wu (2001) presents evidence that the truth lies smack in between the "normalists" and Mandelbrot. So, the tails decay a wee bit faster than would be the case if they were purely power law distributed. But they decay slower than would be for exponential tailed distributions (Malevergne *et al.*, 2003).

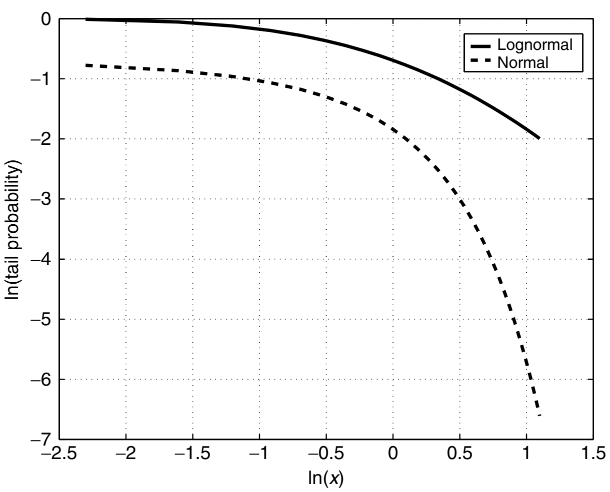

Figure 1 Log-log plot of the normal CDF. This is based on a standard normal distribution. The figure also shows the log-log plot for the lognormal distribution. This is closer to being, yet is not, perfectly linear.

Wu's idea is to model stock prices as pure jump Levy processes, but to make the jump arrival rates follow a power law, dampened by an exponential function. It speeds up the rate of decay in the tails, placing the process in between that of pure normality and pure power law. This produces exactly the desired empirical match. This model is called the DPL, standing for "dampened power law" model. The paper shows that there is a good fit to the data comprising S&P 500 returns and option prices. The model is also able to reconcile the well-known feature of index returns, i.e., that under the physical measure, as the interval between observations increases, the return distribution becomes close to normal; but under the risk-neutral distribution, as seen in option prices, no such reversion to normality appears to be seen. Hence, b-stable distributions are more apt for pricing derivative securities, even though a mere examination of the physical time series would not necessarily reveal this.

The tension between the use of a normal distribution for stock returns (lognormal in stock prices) versus a power law is an ongoing issue for modelers in finance. Curiously, similar issues come up in the field of theoretical computer science. Mitzenmacher (2004) provides an excellent historical perspective as well as an easy to read (for the almost lay person) introduction to the technical issues. He details some important features of power law distributions that we believe will be of interest to financial economists. We summarize some of his findings here:

1. Moments may not be finite. Mitzenmacher highlights the Pareto distribution of the following form: $\Pr[|R| \geq x] = [x/a]^{-b}$, with density function $p(x) = ba^b x^{-b-1}$, where for positive values of $b \leq 2$, the variance is infinite, and for $b \leq 1$, the mean too is infinite. This presents especially thorny problems in pricing securities if the risk-neutral distribution is governed by a power law. Some earlier work facing up to this problem is presented in Heston (1997) and Carr and Wu (2003). There seems to be little doubt now that tails of the return distribution are fatter than Gaussian and fatter than exponential (Malevergne *et al.*, 2003), who raise the concern that even if mean and variance are finite, there is little chance that the skewness and kurtosis will be. Whether they are completely power law is of course the open question, but the evidence seemingly is against that.
2. The shape of the lognormal distribution is very similar to that of power law distributions. As with power laws, the log-log plot of the lognormal distribution is also almost a straight line. See Figure 1 for a visual comparison of the lognormal distribution versus the normal on a log-log plot. The lognormal one approaches linearity.
3. The similarity in shape of the lognormal and power law distributions increases as the variance of the variable increases. See Figure 2 for the log-log plot of the lognormal distribution as we increase variance. Hence, high-return variance stocks appear to be closer to power law distributed than low variance stocks.

While these are only few examples of the use of power laws in looking at return distributions, sporadic bursts in research on this topic have been experienced in the past two decades or so. An excellent paper by Malevergne *et al.* (2003) lays out

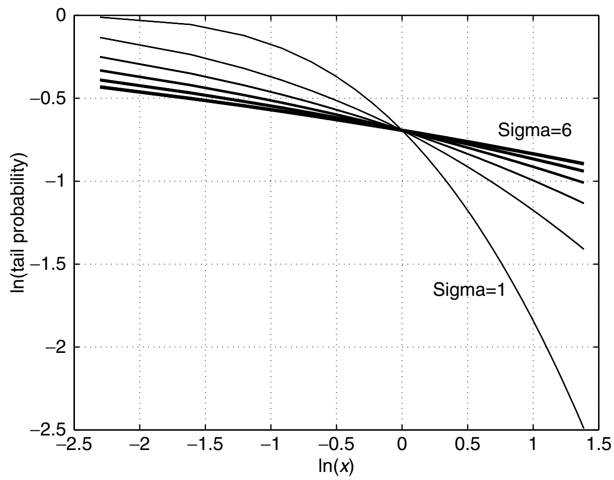

Figure 2 Log-log plot of the lognormal CDF tail. The standard deviation (`sigma`) of $\ln(x)$ is varied from 1 to 6, and the graph shows that the log-log plot becomes almost linear as the standard deviation increases. The steeply sloping line is for `sigma=1` and the flattest line is for `sigma=6`.

much of the history of exploration of these processes. This paper also explores whether stretched exponential distributions and generalized Pareto distributions fit a hundred years of stock returns better than just power law distributions with $b = 3$.

Thus, the recent trend in the literature finally seems to be moving in the direction of looking carefully at these issues. The self-similarity property brings in a link with fractals, an area that has been extensively studied. Thus, there is much that financial economists will find in terms of available tools. However, the thorny issue of unbounded moments also needs to be addressed. Even if the power coefficient is chosen so as to provide a finite mean, thus allowing pricing under the risk-neutral distribution, the variance may not be bounded. For options dealers, the notion of an implied volatility would vanish. Hence, it is clear that adopting power law distributions may pose some difficult problems that reduce its practical value. This may be why researchers and practitioners have been reticent to move to this class of processes. The future here is interestingly poised.

3 Information microstructure

Power laws also describe networks well, so that we may use them to model the information structure of markets and the sociology of investor relationships.

In a recent paper, Das and Sisk (2005) show how the discussion on stock market message boards on the internet may be used to infer the information network across stocks. By examining the commonality of discussion on message boards (as measured by the frequency of postings by one individual on multiple boards within a time interval), they are able to build up a network of information connections between stocks. Once the network of connections is established, it is then possible to subject it to analysis. In the context of power laws, they find that these networks are close to being scale-free. Some stocks have a high degree of connections and others mostly have very few. If the information network of stocks is scale-free, then it implies that information will flow very fast on the network. Quite intuitively, this classic hub and spoke network fosters rapid interchange. Every node is connected to each other one on the network with a very short number of hops, which supports the rapid flow of information. (The same intuition explains the "six degrees of separation" result in social networks.)

The idea of studying the information network might facilitate understanding the mechanism by which information enters prices, and may also be useful in explaining market herding, panics, and crashes. Because of the rapid movement of opinion the tails of the stock return distribution may be induced to be fatter. Thus, the power law (scale-free) structure of the information network may result in the presence of a power law structure in the tails of equity returns. One power law begets another!

An example is the model of Eguiluz and Zimmermann (2000), where investors who share information are pooled into clusters which evolve over time. All investors within a cluster share information and make the same (herding) trades. Trading activity disperses information across the network at a fixed rate. If this rate is below a threshold, the induced distribution of returns exhibits a power-law distribution with an exponential cutoff. Above this threshold, large returns become more likely, as do large crashes. D'Hulst and Rogers (2000) find that the cluster sizes are power law distributed with exponent $b = 5/2$. Xie *et al.* (2002) find that when the number of agents in this model is large but the probability of making a transaction is low, then there exists a high probability of the system coalescing into a single cluster.

There is other evidence of scale-free networks in the information flow through the market. Garlaschelli *et al.* (2005) model share-holding networks, a graph in which the nodes are firms, and the directed, weighted edges represent holdings in the securities of other firms. For a given firm, (a) the in-degree and (b) the sum of incoming link weights both display power law tails in an empirical study over several markets; further, (a) is a power law function of (b). They also find that the exponents of these power laws are related by a simple scaling function. It seems reasonable that people obtain new information when they have a vested interest in that information. Consequently, power laws of mutual holdings could be a major underpinning of information propagation.

Scale-free distributions may be used to model default contagion. A network model of credit correlations related to this idea is developed by Giesecke and Weber (2003). They show how the model may be used both for calibration and analysis of joint loss distributions. Such a model will also support analyzing losses economy-wide in the limit. Many of the results now widely developed in the physics and computer science literature on the limiting form of such networks will become useful in understanding market interaction and aggregation of information in finance.

Of course, the presence of a power law information network does not have purely negative implications. The reverse intuition works just as well in the context of market efficiency. Scale-free networks offer fast information aggregation even when information brokers and intermediaries have access to only limited portions of the information set (Moss *et al.*, 2000).

An important question arises at this point—what is the mechanism that leads to information structures in finance becoming scale-free or power law distributed? (Note that we are being rather loose here with the network description—what is meant is that the degree distribution of the information network will be power law distributed.) Further research is required to determine if the information network of stocks, as it grows, is becoming more power law oriented or less. To date this is not known, but may be explored by looking at the growth of the message board network from its inception, data which is indeed available today. It may also be the case that at different time scales the information structure is quite different. It is quite likely that, as markets grow and become increasingly globalized, they coalesce into power law mode, increasing the risk of a market meltdown, while concurrently becoming more efficient.

The growth of networks into their scale-free forms has been codified into the theory of *preferential attachment* of Barabasi and Albert (1999). The idea of preferential attachment is that when a new node enters the network, it tends to attach to other nodes with a probability that is proportional to the number of connections of each other node. This eventually results in the limit structure of a scale-free network. Whether the information network in markets grows in a manner consistent with preferential attachment is not known, but is certainly worth an empirical exploration.

Traces of the information transmission have never been more transparent or accessible. Internet newsgroups, mailing lists, email corpora such as Cohen (2004), closed-captioning for news media outlets are just some examples of the wealth of information available to use at a variety of different time scales. It should be possible to observe the propagation of exogenous news (or endogenous herd behavior) across investor groups an individual at a time.

Gabaix (1999) explores the mechanism of growth in detail for city sizes, and provides a useful discussion of how certain forms of multiplicative growth eventually result in power laws in the domain of economics; see Yule (1924) for what is the first attempt at this line of thinking, and Simon (1957) for the characterization of this phenomenon as the "Gibrat" principle. Mitzenmacher (2004) presents a simple analysis also showing that many different multiplicative growth models may result in variables that are approximately lognormal. Hence, this lends subtlety to the contentious debate about power law and lognormal distributions for financial variables. Both appear to be the same, yet are indeed so different.

Gabaix *et al.* (2003) provide a microstructure model that explains why many of the fluctuations we see in financial variables are power law distributed. The basis of their theory lies in the size distribution of participants in the market. When these agents (investment funds) of different sizes optimize, it result in power law distributions in returns, trading volume, and number of trades. Thus, as summarized by Mitzenmacher (2004), there are two routes to power laws: (i) preferential attachment and (ii) optimization.

Lüders *et al.* (2004) explore empirically the outcome of optimization behavior, which is posited to result in a power law relationship between prices (P) and dividends (D), i.e., $P = aD^b$. They also show that this is consistent with an economy of agents who have declining relative risk aversion. They look at six major international stock markets and find that the evidence supports the theoretically derived power law relationship.

4 Summary

Power law distributions are intricately connected to natural phenomena; in like manner, so are regular distributions such as the normal. As we have seen, the two distributions do indeed resemble each other in many ways, but are also very different, as may be evidenced in an examination of their moments.

A critical issue in the adoption of power laws is whether the moments of relevance (mean, variance at least) are bounded. Of course, even if the underlying stock (or other financial variable) has infinite mean and/or variance, but the required derivative security (a function of the underlying) has finite mean, we may still price securities as before. This logic was exploited in Carr and Wu (2003), though calibration may still be an issue. Hence, there seem to be some, but not too many critical impediments to the adoption of power laws in practical applications in finance.

Note

[1] This is by no means a comprehensive survey of power laws, nor does it aim to cover all the literature in finance. We apologize in advance for leaving out many references that were not directly related to the goals of this article. Our only purpose here is to highlight two interesting areas where power laws may have an important role to play in the development of our thinking in finance, and to analyze the linkages between these two areas. Linkages to other fields are also mentioned because there seems to be a growing connection in the recent literature.

References

Adamic, L. and Huberman, B. (2000). "The Nature of Markets in the World Wide Web." *Quarterly Journal of Electronic Commerce* **1**(1), 5–12.

Barabasi, A.-L. and Albert, R. (1999). "Emergence of Scaling in Random Networks." *Science* **286**, 509–512.

Carr, P. and Wu, L. (2003). "Finite Moment Log Stable Process and Option Pricing." *Journal of Finance* **58**(2), 753–777.

Cohen, W. (2004). "Enron Email Data Set." http://www-2.cs.cmu.edu/∼enron

Das, S. and Sisk, J. (2005). "Financial Communities." *Journal of Portfolio Management*, Summer forthcoming.

Eguíluz, V. and Zimmermann, M. (2000). "Transmission of Information and Herd Behavior: An Application to Financial Markets." *Physical Review Letters* **85**, 5659–5662.

Fama, E.F. (1965). "The Behavior of Stock Market Prices." *Journal of Business* **38**, 34105.

Gabaix, X. (1999). "Zipf's Law for Cities: An Explanation." *Quarterly Journal of Economics* **114**(3), 739–767.

Garlaschelli, D., Battiston, S., Castri, M., Servedio, V., and Caldarelli, G. (2005). "The Scale-free Topology of Market investments." *Physica A* **350**(2–4), 491–499, forthcoming.

D'Hulst, R. and Rodgers, G.J. (2000). "Exact Solution of a Model for Crowding an Information Transmission in Financial Markets." *International Journal of Theoretical and Applied Finance* **3**(4), 609–616.

Gabaix, X., Gopikrishnan, P., Plerou, V., and Eugene Stanley, H. (2003). "A Theory of Power Law Distributions in Financial Market Fluctuations." *Nature* **423**, 267–270.

Giesecke, K. and Weber, S. (2003). "Cyclical Correlations, Credit Contagion, and Portfolio Losses." Working paper, Cornell University.

Heston, S. (1997). "Option Pricing With Infinitely Divisible Distributions." Working paper, University of Maryland.

Lüders, E., Lüders-Amann, I., and Schröder, M. (2004). "The Power Law and Dividend Yields." Working paper, Universite Laval.

Malevergne, Y., Pisarenko, V., and Sornette, D. (2003). "Empirical Distributions of Log-Returns: between the Stretched Exponential and the Power Law?" Working paper, UCLA.

Mandelbrot, B.B. (1963). "The Variation of Certain Speculative Prices." *Journal of Business* **36**, 394419.

Mandelbrot, B. and Hudson, R. (2004). *The (Mis)behavior of Markets*." New York: Basic Books.

Mitzenmacher, M. (2004). "A Brief History of Generative Models for Power Law and Lognormal Distributions." *Internet Mathematics* **1**(2), 226–251.

Moss, S., Edmunds, B., and Wallis, S. (2000). "The Power Law and Critical Density in Large Multi Agent Systems." Working paper, Manchester Metropolitan University.

Pareto, V. (1896). *Cours d Economie Politique*. Geneva, Switzerland: Driz.

Simon, H. (1957). *Models of Man: Social and Rational*. Mathematical Essays on Rational Human Behavior in a Social Setting. New York: Wiley.

Wu, L. (2001). "Dampened Power Law: Reconciling the Tail Behavior of Financial Security Returns." *Journal of Business*, forthcoming.

Xie, Y., Wang, B., Quan, H., Yang, W., and Hui, P. (2002). "Finize-size Effect in the Eguíluz and Zimmermann Model of Herd Formation and Information Transmission." *Physics Review. E* **65**, 046130.

Yule, G. (1924). "A Mathematical Theory of Evolution Based on the Conclusions of Dr. J.C. Willis, F.R.S., *Philosophical Transactions B* **CCXIII**, 21–87.

Zipf, G.K. (1949). *Human Behavior and the Principle of Last Effort*. Cambridge, MA: Addison-Wesley.

Keywords: Power laws; networks; microstructure