FREQUENCIES

Published by
Ekaros Analytical Inc.

The Journal of Size Law Applications
Jan 2002

Vol 1 no 2

Zipf Ranking, Benford's Law, and Ratio Analysis

Ratio analysis, as a term, is extremely broad. The most commonly understood meaning of financial ratio analysis has to do with basic business measures like profit as percentage of revenue, the debt-to-equity ratio, revenue earned per full time employee, and the like. Many more specialized and abstract types of financial ratio analysis have also been proposed, such as intellectual capital ratios or risk ratios. All such techniques rest on a fact which is seldom considered explicitly: how the underlying data are distributed.

In this issue's companion essay by David Coderre, we examine ratio analysis as an audit technique, focusing on detecting anomalies in data sets such as payroll or accounts payables transaction files. This is a particularly worthwhile application because it can pinpoint specific items for investigation (rather than whole categories). However, for this or any other ratio analysis technique to be useful, we require the presence of an underlying size law, or in some cases more than one, governing the quantities involved. Without a clear prior expectation of what we will find *in general*, we cannot assign a meaning to any particular ratio. This is the subject of the present essay.

In the history of financial ratio analysis, one discovery stands out as critically altering our expectations, all across the range of techniques: a paper by E. B. Deakin in 1976 for *Accounting Review*, demonstrating that the data underlying financial ratios almost invariably has a strongly skewed, non-normal, logarithmic distribution. (Work by Eugene Fama in the 1960s had already shown this was true of stock markets; Deakin's paper generalized the finding.)

This was a major discovery because *normative* conclusions (whether a ratio is good or bad) tend to require a norm, or average expected value. For example, we may wish to attempt to forecast future profitability by examining return-on-assets, return-on-equity, or debt-to-assets ratios. But if the actual values for assets, equity, and debt are all logarithmically distributed, then it can be virtually meaningless to speak of an 'average' for any of these measures—and thus a normal level of return on assets or return on equity.

In a normal or bell-shaped distribution curve, the mean, median, and mode are the same value, and we tend to refer to this value somewhat loosely as the 'average'. But the mean value of a skewed, non-normal distribution, such as shown in Figure 1, is completely different from the median value, and the median is very different from the mode. Most items in the data set—some-

Figure 1. Normal (bell-shaped) versus logarithmic distributions

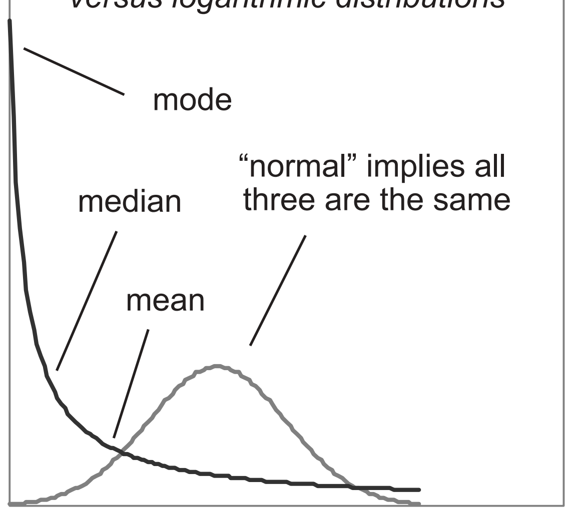

times an overwhelming majority—are *below* the mean but *above* the mode.

Objective financial ratio analysis is hampered by the fact that, depending on the value we choose as our norm, we can arbitrarily force the majority of items in our sample to be below or above it. Such analysis suffers even more from another, related fact. When dealing with real-world data, a normal distribution implies that we can trust future values to cluster in the same way as historical ones, around the same mode/median/mean value. Non-normal distributions, however, tend to be *nonstationary*, meaning that the mean, median, and mode all tend to change with time. Statistical tests and techniques that would otherwise be applicable can be rendered useless by this fact. Even very sophisticated tools, such as neural nets or options-pricing equations, will yield misleading results when handling non-normal, nonstationary distributions.

Numerous investigators since 1976 have attempted to re-normalize business data sets, using indexing schemes or multiple-variable approaches. The most that is typically obtained from these efforts is that the distributions become less dramatically skewed. Few if any become normal.

As an example of the perverse effect of non-normal distributions, consider the frequently quoted statistic that most mutual fund managers (the figure varies from year to year, but as many as 75 to 95 percent) produce results that are below the overall market average. If we are expecting a bell-shaped performance curve, with a single well-defined norm, this fact seems shocking and suspicious: what are all those managers doing? But if we are accustomed to financial performance data following a logarithmic trend, then there is no shock. *Of course* most individual funds perform below the mean: that's the (non-normal) norm! The same goes for the quarterly output of salespeople, or patents earned per researcher, or countless other performance measures.

One case that perfectly illustrates the potential disasters implicit in a normalizing approach to forecasting and assessing non-normal financial results is the infamous collapse of Long-Term Capital Management (LTCM) in 1998.

LTCM was founded in 1993 as a 'hedge fund' for large institutional investors. It specialized in global arbitrage, making profits from, for example, the differences between bond rates in Italy and Germany. Among its founders were two soon-to-be Nobel laureates in economics, Myron S. Scholes and Robert K. Merton. Two decades earlier, Mr. Scholes, Mr. Merton, and colleague Fischer Black had created an elegant model for pricing options—essentially quantifying the risk that a stock or other security will move up or down, and assigning a price for a given amount of potential movement. Roger Lowenstein, in *When Genius Failed*, explains how this normalizing model became the foundation for LTCM's worldwide trading.

> [Such hedges are] based on the assumption that the volatility of stocks is, over time, consistent. The stock market, for instance, typically varies by about 15 to 20 percent per year. Now and then the market might be more volatile, but it will always revert to form—or so the mathematicians in Greenwich believed. It was guided by the unseen law of large numbers, which assured the world of a normal distribution . . . of quiet trading days and market crashes.

The driving assumption that emerged from the normalizing model was that over time, differences in bid and ask rates (or 'spreads') will tend to narrow. A large, well-capitalized fund that can afford to wait will profit hugely by betting against temporarily large spreads.

The management of LTCM was not entirely naive about models. Wall Street already knew, in large part because of investigators like Fama and Deakin, that a normalizing model did not match reality. Nonetheless, traders everywhere used such models (and still do today), because no competing theory offered a quantifiable, systematic alternative, and even a doubtful model tends to offer better results than guessing.

In its first few years, LTCM profited hugely, returning as much as 71 percent per year to its investors. By late 1997, it was the largest hedge fund in the world, with more than $100 billion in assets, and was also very heavily leveraged in terms of debt. If volatility and the size of bid-ask spreads on one market exceeded the levels in the model, LTCM was supposedly 'hedged' by having similar bets on other markets. But if markets all around the world were ever to go into crisis simultaneously—so that bonds in Russia and stocks in Japan all fell at once, and fell dramatically, driving the gap between bid and ask prices higher everywhere—then, despite the model's showing odds of trillions to one against such an event, LTCM's complex web of trades could be overwhelmed by losses greater than its reserves in a matter of days.

In August and September 1998, after currency turmoil in Asia spread to Russia and then the rest of the world, the markets went against the models. LTCM's narrow equity reserve of $4.5 billion disappeared. A massive bailout orchestrated by the New York Federal Reserve Bank was required to rescue the fund from bankruptcy—otherwise, many believed, as much as $1 trillion in a global network of related investments might have been in peril. The historical rate of market volatility had proved to be an inadequate and misleading guide to what would happen in future—and LTCM was far from being alone in its losses. Economic forecasters around the world were stunned by the 'Asian flu' and its aftermath.

Coping with non-normal and nonstationary financial data, or 'fat tails' as they are known, is clearly much more than an academic exercise. There are literally billions of dollars to be made by whoever best solves the problem posed by such distributions.

Zipf ranking requires a highest value, while Benford's Law assumes distributions that are uniform and unbounded. To perform effective ratio analysis we need to consider both.

Our objective in this series of articles is to lay a foundation for better handling of such problems, by bringing together several different laws of non-normal distribution and showing how they can be understood as aspects of a single universal law.

Zipf Ranking and Benford's Law

In the late 1930s, George Kingsley Zipf (1902-1950), a Harvard linguist, published the first study of what came to be known as Zipf's Law—a consistent ranking pattern that shows up in the frequency of occurrence of different words, populations of cities within a country, and the sizes of competing companies (to name just three examples).

Zipf's Law is very easy to state. If the average quantity obtained for the highest ranked item is Q, then the average quantity that will be obtained for the Nth ranked item will be Q/N. The largest city in a given country will average twice the size of the second largest, three times the size of the third-largest, and so on. In English, the most commonly used word is "the," which appears on average just twice as often as the second-ranked word, "of". The third-ranked word, "to," appears about one-third as often as "the," and so on. The law is usually written this way:

$$\mathbf{I \times Qi = C}$$

where I is the rank, Qi is the quantity of the ith-ranked item, and C is a constant specific to the data set.

Zipf's Law has gained increasing respect among theoreticians in recent years as reflecting a deeper underlying order in nature. A number of papers have been written attempting to explain why Zipf's Law works the way it does, with only partial success.

One property of Zipf's Law that will be important for our discussion has to do with the number of items found in successively larger value ranges. For example, consider a series of 1,000 items that obey the law. The lowest ranked item is therefore 1/1,000th as large as the highest. If we call the value of the lowest ranked item X, then from X to just under 2X we will find, on average, 500 items (because the 500th ranked item should have value exactly 2X). Furthermore, from 2X to 4X we will find 250 items; from 4X to 8X, we will find 125; and so on. For each doubling of the range, the number of additional items found is exactly *half* the number in the previous range.

We turn now to Benford's Law. In 1938, Frank Benford, an engineer and physicist at General Electric, published a paper regarding the digit frequencies of 'anomalous numbers,' meaning numbers from everyday life such as street addresses, populations, physical constants, and baseball statistics. The first digit '1' occurs far more often than other first digits, about 30.1 percent of the time, and there is a smooth progression down to first-digit '9', which occurs 4.6 percent of the time (see Table 1 on the next page for a list of values). Although Benford's Law is stated in terms of digit frequencies, it also assumes an underlying order in terms of the quantities the digits represent: in effect, it assumes that the data set approximates a uniform geometric sequence. (See Issue 1 or *Digital Analysis Using Benford's Law*, by Mark Nigrini, for more on Benford's Law.)

Benford's Law is very similar in form and implications to Zipf's Law. Are the two related in some way? And are there perhaps practical applications that arise from knowing both laws, that cannot be found using one?

There are strong similarities. Both are essentially logarithmic size laws, and both apply to a very wide variety of phenomena. For certain kinds of data, such as the sizes of companies, both laws have seemed to apply at once, depending on circumstances. The two are in fact separate aspects of a third, broader law—a point that will be developed in this series of articles. By clarifying what makes them distinct and how they work together, we will gain a powerful new understanding of ratio analysis.

Table 1: Digit Frequencies for Zipf Sequences

Initial Digit	1000	1500	2000	3000	4000	5000	6000	7000	8000	9000	Total Count	Percent	Benford
1	56	84	111	167	222	278	233	189	144	100	1584	31.7%	30.1%
2	186	27	39	55	75	94	111	130	150	166	1033	20.7%	17.6%
3	92	140	17	29	36	45	56	65	72	84	636	12.7%	12.5%
4	55	82	111	16	23	27	33	38	45	50	480	9.6%	9.7%
5	39	56	75	11	16	20	22	27	30	33	329	6.6%	7.9%
6	26	40	54	81	10	13	18	17	21	25	305	6.1%	6.7%
7	19	31	38	59	7	10	11	16	15	17	223	4.5%	5.8%
8	15	23	31	46	63	8	10	11	15	13	235	4.7%	5.1%
9	12	17	24	36	48	5	6	7	8	12	175	3.5%	4.6%

Zipf's Law is distinct from Benford's Law in that it specifies a highest value for the set. It describes what we should expect to find starting at the *upper limit* of a distribution. Zipf's Law is also a ranking law, which implies that the items in the set are directly related to one another. By contrast Benford's Law, as originally derived, treats data sets as unbounded, and does not require items in the set to be directly related. It turns out that in many practical applications, a 'Benford set' (meaning a data set that obeys Benford's Law) can be broken down into several distinct Zipf sets, each one obeying or closely approaching Zipf's Law.

Here is a concrete illustration. Table 1 summarizes the digit frequencies found for ten Zipf sequences of 500 numbers each, based on highest values of 1000, 1500, 2000, 3000, and so on up to 9000. These are compared to the expected digit frequencies based on Benford. It should be apparent from these that a single Zipf sequence will not closely follow Benford's Law except occasionally by accident. In half of these sequences, the most common initial digit is not 1, but 2, 3, or 4.

This lack of conformity is to be expected. Benford's original paper included frequencies for twenty different data sets, many of which did not follow the digit-frequency law very well on their own. Data sets based strictly on Zipf's Law are therefore not particularly unusual in deviating. However, Benford found that by combining different sets of data into a single large table, he could obtain closer and closer conformity. The same occurs here: if we combine all ten data sets, the digit-frequency pattern for the group conforms much more closely to Benford—to within a few percent. And this is not actually a case of combining different *kinds of data*; here we are combining sequences that may be very similar, but have different *starting values*.

The case is equally compelling when we look at real-world data. We will base this discussion on the *Fortune* listings of the 500 largest public corporations in America, from the April 2001 issue. Right away we notice that Zipf's Law does not seem to apply to the top five, or the top ten (shown in Table 2 on next page).

To conform closely to Zipf's Law, Wal-Mart would have needed to earn more like $105 billion, or half what Exxon Mobil did, and Ford's revenues would have needed to be around $53 billion, or one-fourth. Verizon's earnings are three times too large for tenth position. The top earners are much too close together.

Exxon Mobil	$210.4 billion
Wal-Mart	$193.3 billion
General Motors	$184.6 billion
Ford Motor	$180.6 billion
General Electric	$129.9 billion
Citigroup	$111.8 billion
Enron	$100.8 billion
IBM	$88.4 billion
AT&T	$66.0 billion
Verizon Communications	$64.7 billion

Table 2: Sales for top 10 corporations

The rest of the list also diverges from an exact Zipf ranking: #500 on the list is Qualcomm, with $3.2 billion in earnings. To fit properly with Zipf's Law, Qualcomm ought to have earned $210/500 = $0.4 billion. This is a substantial difference, about eight times too large. The entire list is more closely spaced than we expect in theory.

The same is true if we consider the top 500 in terms of employees. Wal-Mart held #1 in this category in 2001 with 1,244,000 employees. The next four were General Motors with 386,000, McDonald's with 364,000, United Parcel Service with 359,000, and Ford Motors with 346,000. Although the drop from #1 to #2 is large, for #4 and all subsequent rankings the real value is larger than we would expect based on the ranking law. In 500th place we find a three-way tie with 12,700 employees each—much larger than the value of 2,488 we would expect based on Zipf's Law.

This effect has been observed elsewhere. For example, although Zipf's Law does apply to city populations in most countries, it fits less well to data from the former USSR or mainland China, where freedom to move from place to place was, or still is, limited. It also fits less well if cities from several countries are grouped together, for similar reasons. Just as with Fortune 500 companies, the lower-ranked cities in such cases are larger (closer to the #1-ranked value) than we would expect under strict application of the law.

Rank	Actual	Zipf
1st	100.0 %	100.0 %
2nd	60.1 %	50.0 %
3rd	38.2 %	33.3 %
4th	28.1 %	25.0 %
5th	21.5 %	20.0 %
6th	17.8 %	16.7 %
7th	14.2 %	14.3 %
8th	11.4 %	12.5 %
9th	10.2 %	11.1 %
10th	9.1 %	10.0 %

Table 3: Actual and expected rankings within industries

We could also say that the top-ranked cities are *smaller* than expected: given freedom of movement, more people might opt to move from remote Chelyabinsk to Moscow or New York, making Moscow or New York larger, Chelyabinsk smaller, and the entire list more widely spaced. In general, when a data set is made up of subsets that are unrelated or weakly related (and thus are not ranked in any practical sense), the values diverge from Zipf ranking by being *too close together*.

In addition to ranking the largest U.S. companies by overall revenue, *Fortune* also breaks down its list into 69 industry categories. Here we can more plausibly test for compliance to Zipf's Law, taking averages across the various industries: if the second-largest company in each industry category averages half the size of the largest, and if the third-largest averages one-third the size of the largest, and so on, we can say that Zipf's Law applies within industries. In general, the firms in an industry category do directly relate to one another, because they are competing for the same business; so we expect to find good agreement.

Table 3 shows the results. Summing the revenues of all the top-ranked firms in 69 categories, and comparing them to the total revenues of the second-ranked firms, we get a ratio of 1.66 to 1. The ratios for each successive position, down through the tenth-ranked firm, are broadly consistent with Zipf's Law.

The ratio between #1 and #2 is still not 2 to 1, however. Several explanations come to mind. First, several of the categories in the *Fortune* lists do not qualify as distinct industries, and the companies within such categories are not competing with one another. For example, one *Fortune* category is 'specialty retailers,' which includes Bed Bath & Beyond, Barnes & Noble, and Petsmart. The revenues earned by Petsmart are not in any obvious sense taken away from Bed Bath & Beyond, or from Barnes & Noble.

Second, companies ranked within a distinct industry may not necessarily be strict rivals, or even at arm's length from one another. Owens Corning, ranking second in 'building materials & glass,' was created as a joint venture in the 1930s by Owens-Illinois (currently ranking first in the category) and Corning. A similar situation prevails for 'beverages,' where Coca-Cola ranks first, Pepsico ranks second, Coca-Cola Enterprises ranks third, Pepsi Bottling ranks fifth, and PepsiAmericas ranks sixth.

Finally, there are many firms in the Fortune 1000 that sprawl across multiple categories. They compete, simultaneously, with dozens of rivals in different markets. The entire 'diversified financials' category, led by General Electric, fits this description.

All of these arguments tend to make an individual industry category more closely spaced than Zipf ranking predicts. However, the opposite tendency is also a concern. *Fortune* rankings do not include privately held or foreign companies, only publicly traded corporations that are located in the U.S. Some of the largest companies in America are privately held, as for example the statistical software company SAS, with revenues exceeding $1 billion per year. Subsidiaries of European or Asian corporations also account for an substantial proportion of U.S. business each year. Randomly excluding major competitors in a given market will distort the results by leaving gaps in the ranking. For example, the lowest-ranked companies on a list of the 100 largest Canadian software firms average around 1/300th the size of the top-ranked firm—much smaller than expected—in part because these firms compete with U.S. firms that are not listed.

Because the *Fortune* data is not designed to show market share in particular markets, it makes sense to examine more specialized data sets as well. One such public data set is the Top 500 ranking published every six months at www.top500.org, showing the world's 500 most powerful supercomputers, with their respective manufacturers. This is a global market, and an extremely competitive one, and so it ought to conform very closely to the ideal Zipf values. The state of the art in supercomputing is advancing so quickly that major changes in market position occur regularly. However, if we take an average across the past nine years, the market-share ratio of the leader to its nearest rival (in computers sold, not dollars) is 1.92 to 1. The ratio from #1 to #3 is 2.75 to 1. The ratios for rankings 4-10 also tend to follow Zipf closely.

These findings are not new, so we will not belabor the reader with endless examples. Zipf found the same trends in the 1940s. However, it is important to be clear about the details, as wrong assumptions can lead us to wildly different conclusions from the same data. In a 1976 essay, strategic consultant Bruce Henderson discerned a similar pattern in market share, calling it 'the rule of three and four':

> A stable competitive market never has more than three significant competitors, the largest of which has no more than four times the market share of the smallest. The conditions which create this rule are:
>
> A ratio of 2 to 1 in market share between any two competitors seems to be the equilibrium point at which it is neither practical nor advantageous for either competitor to increase or decrease share. This is an empirical observation.
>
> Any competitor with less than one quarter the share of the largest competitor cannot be an effective competitor. This too is empirical but is predictable from experience curve relationships.

Notice that Henderson, who was apparently unaware of Zipf's earlier work, thought in terms of a geometric progression that is significantly steeper than what we find in practice, expecting #3 to hold one-half the share of #2, and so on down. Thus the sixth-ranked competitor in a market, per Henderson, should get around 1/32nd of the market share of the leader. In fact, in a worst-case situation involving rivals competing head-to-head for a single integrated market, we find the sixth-ranked competitor tends to hold 1/6th of the leader's share . . . and in situations where markets are fragmented and competition is not necessarily direct, lower-ranked competitors will hold even larger shares. (A fuller discussion of the impact of Henderson's work can be found starting on page 19.)

Benford's Law and 'range bins'

Hav ing shown that market share rankings do approach closely to Zipf's Law (provided that we properly specify the market), we need to confirm that when we group them into a larger data set, they also obey Benford's Law. Here we face a minor technical hurdle and a more serious theoretical one.

The minor hurdle is that we cannot simply compile the digit frequencies for the revenue figures for the *Fortune* 500 or *Fortune* 1000, because these figures are biased by the arbitrary cutoff point. If we take the 500, the lowest value is $3.2 billion, and the resulting distribution will be biased with too many threes. The 1000 largest are similarly suspect, given their cutoff at $1.2 billion. The result in each case is known as a small subset, and the theory behind digital analysis holds that such subsets will not conform. A properly constructed data set for this purpose would have as its lower value the smallest public company with revenue above some round multiple of 10, such as $1 billion or $100 million. Instead of exactly 500 or 1,000 companies, we would then have a set of some arbitrary number like 1,343, but all else being equal we would expect good conformity to Benford. Details of this method of analysis can be found in Mark Nigrini's *Digital Analysis Using Benford's Law*.

However, we can determine conformity to Benford by several indirect means. First, the totals for profits, assets, and employees are not so tightly restrained, and their relative frequency values are just slightly outside the normal limits of statistical variation for a data set of this size.

There is another indirect means to test the revenue figures, but this is where our major theoretical problem (and insight) arises. We saw that in a set obeying Zipf's Law, each doubling of the value range adds fewer and fewer data items—50 percent fewer each time. If there are 16 companies in a given market earning between $1 million and $2 million per year, then given Zipf ranking there will be 8 between $2 million and $4 million, and 4 between $4 million and $8 million, and so on. We can call this a 'range bin' analysis, in which we count the items as they fall into a series of fixed ranges.

However, for a data set to perfectly satisfy the assumptions behind Benford's Law, there cannot be any such decline. We expect to find the same number of additional items within each larger bin—not 50 percent as many, but 100 percent each time. This condition of uniform geometric spacing is a key assumption behind the law.

Ironically, even though Benford's Law is observed occurring in many diverse data sets, in testing for this condition we are asking for something unreasonable on the face of it. Such a data set is impossible in reality because it will never reach a limit; it will become infinitely large.

A data set might approximate this ideal by having the same number of items in several successive range bins,

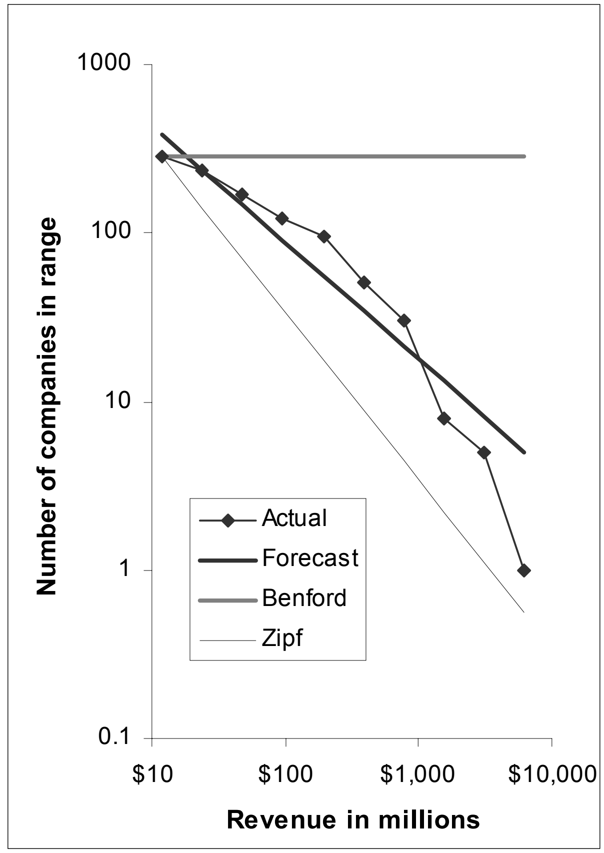

Figure 2. Actual distribution of Fortune 1000 revenues compared with Benford and Zipf (grouped by range bins of size 2.0)

and then simply running out of data. But in practice, this doesn't happen. Instead, most real data sets show a moderate decline with each doubling of the range.

For example, if a company has 1,000 customers who bought between $1,000 and $2,000 worth of merchandise during the year, then the number who bought between $2,000 and $4,000 will typically not be 1,000 (per Benford) or 500 (per Zipf), but some number between the two, such as 600 or 700. The number for the bin from $4,000 and $8,000 will then fall between 350 and 500, i.e., 60 to 70 percent of 60 to 70 percent of 1,000.

In Figure 2, based on *Fortune* data from 1981, the actual drop varies from 53 to 83 percent, well above Zipf. As we approach the highest value, the distribution drops off slightly more quickly. These final few data points represent only about 1 percent of the sample, a dozen companies. The trend of the line otherwise remains consistently higher than the predicted Zipf values, and much lower than those of an ideal Benford set.

The data set comes close to conforming to Benford's Law, but tends to deviate in the direction of having slightly too many '1's and not enough '9's—that is, '1' shows up more than the value of 30.1 percent that Benford arrived at in 1938. This reflects the fact that there is a rough geometric spacing between items, but the spacing getting larger as the items get larger. For very large and diverse data sets this tends not to matter, but the deviation from theory often shows up in smaller sets—for example, with the benefit of hindsight we can see evidence of it in Benford's original data.

Range bins and the 'ontic' distribution

How can all this help to produce better financial ratio analysis? Part of the story will be revealed next issue, but here are two clear demonstrations. The distribution of revenues among the 500 does not follow Benford's Law or Zipf's Law closely enough to use either one as a forecasting guide. Each of these laws turns out to be only a weak approximation of the actual facts. However, using the range-bin method we can identify a very strong regularity in their distribution, one that carries over into other calculated results (such as ratios) that use the revenue values.

In Figure 3, the employee counts for 1981 are shown.

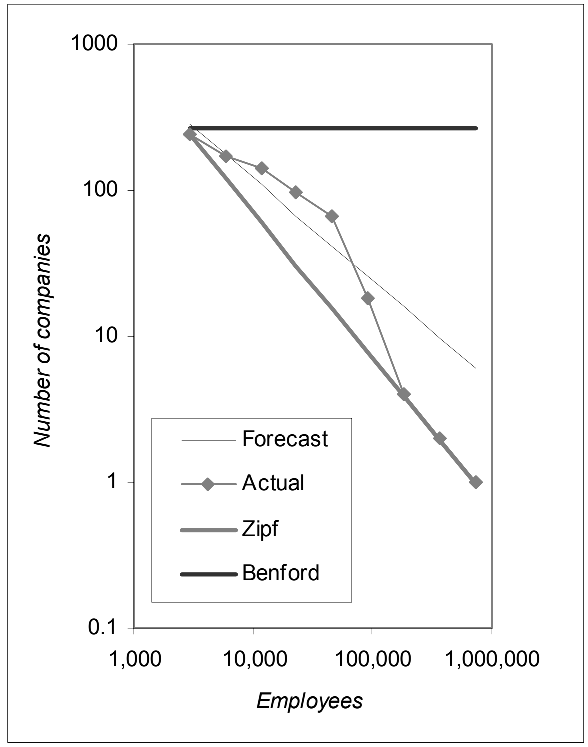

Figure 3. Employee data for the Fortune 1000 shows the same slope as revenue data

The line marked "forecast" has precisely the same slope as the forecast line in the revenue counts for 1981. This same line is also the best fit for employee and revenue

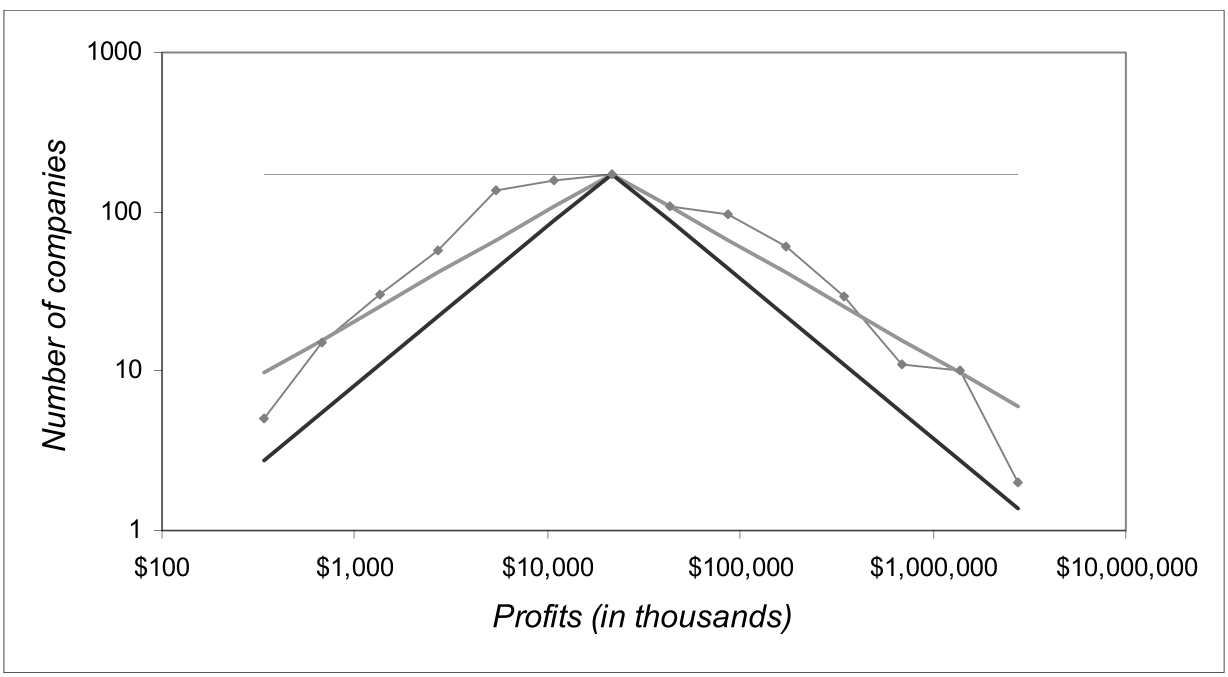

Figure 4. Distribution of Fortune 1000 profit totals showing triangular structure

data from other years. In each case, the fit breaks down for the largest 5 to 10 items, which tend to revert to a Zipf-type distribution; but for 99 percent of the data set the fit is very good. The same slope shows up in data sets unrelated to the *Fortune* companies as well.

We can call this the 'ontic' distribution for short. The slope of the ontic distribution is such that each time the range is doubled, the number of items in the new range segment is approximately 62 percent of the previous value. This is a key characteristic of the distribution: a constant average slope of about 62 percent, regardless of the particular quantities being measured.

Now consider what happens if we subtract cost values (with a 62 percent slope) from revenue values (also with a 62 percent slope) to obtain profits. Historically this has been judged to result in a highly skewed distribution, without a clear norm. Attempting to normalize this result by dividing profits into revenues has not worked. In fact, the record of Wall Street analysts in forecasting profits closely resembles the chronic underperformance by mutual fund management: most analyst forecasts are wrong. One study by David Dreman of the years 1973-96 found that consensus projections of quarterly earnings were off by a median amount of 42 percent, despite being issued as little as two weeks before the company itself reported! Large error rates have been found in every industry, for large companies and small ones, in booms and recessions.

However, despite this disappointing record, profit forecasts still form an essential part of how Wall Street does business. Therefore, any potential for improvement in the methodology must belooked at.

In Figure 4, we see that the profit figures form a kind of triangle. The right side of the triangle has the classic 62 percent slope. The left side of the triangle has a 62 percent slope as well . . . only pointing the opposite way! We have symmetry, admittedly on a log-log graph—and along with symmetry comes a new way of arriving at a consistent mode, mean, and median value. This is a stunning regularity for what is generally acknowledged to be non-normal data.

Wherever we can find regularities, we have a basis for making predictions and building models. There are more regularities to explore as well as many implications. We will continue this study next issue.

How to Think and Win in a Nonlinear Universe

The 80/20 Principle: The Secret to Success by Achieving More with Less

The Power Laws: The Science of Success

by Richard Koch

In 1998, Richard Koch wrote a very readable single-volume introduction, in non-mathematical language, to the idea of root-cause Pareto analysis as it applies to work and life—that is, the 80/20 principle. Three years later he followed up with a broader study of fundamental laws in general, first published in the UK as *The Power Laws*, and now available in North America under the title *The Natural Laws of Business*. Together these books constitute perhaps the best treatment of size laws and nonlinearity ever written for business.

Pareto analysis is now a century old, and is widely known among MBA students and engineers, but the 80/20 principle as Mr. Koch describes it is not so much about size laws or distributions or even mathematics, as it is a methodology for identifying where the easiest path to success can be found.

The same approach is evident in *The Power Laws*, which touches on a total of 93 different laws, ranging from Newton and Darwin to Mandelbrot and Zipf. Mr. Koch notes in his introduction that he and his researcher actually looked at more than a thousand theories and principles before settling on these particular ones.

Those already familiar with these laws may find the books a little thin, as there are no theorems or proofs in either one. *The 80/20 Principle* includes several worked examples of Pareto analysis, while *The Power Laws* avoids explicit mathematics entirely.

However, for readers who want to apply nonlinear principles outside of a narrowly technical sphere, believing that the nonlinear character of the universe can actually work for us rather than against us, these books are a splendid effort and long overdue.

The first paragraph of the earlier book sets out Mr. Koch's thesis, and we can easily read 'the 80/20 principle' as standing in for power laws in general:

> The 80/20 Principle can and should be used by every intelligent person in their daily life, by every organization, and by every social grouping and form of society. ...This book, the first ever on the 80/20 Principle, is written from a burning conviction, validated in personal and business experience, that this principle is one of the best ways of dealing with and transcending the pressures of modern life.

To this I feel like replying: Amen, brother! The stakes involved are indeed high, and there is a clear need in this field for books written by visionaries, rather than solely by technicians. Whether Mr. Koch succeeds in his loftier aims or not—and reviews have cited both books as among the most influential of the decade—he deserves respectful attention for having attempted them.

About *The 80/20 Principle*

The book is divided into four parts. The "Overture" offers a compact , 40-page summary of what the 80/20 principle is, what we know about its fundamental causes, and a simple set of math tools (mainly bar graphs and frequency distribution charts) to deal with it.

In tracing the history of the principle, Mr. Koch offers some worthwhile clarification: although the principle is clearly attributable to him, Pareto himself never referred to his discovery as "the 80/20 principle". Pareto's original formula was considerably more abstract, expressed in terms of logarithms, and it was not until a generation later that the simpler summarization became common—that, for example, "about 80 percent of all property is held by about 20 percent of the population." Mr. Koch then cites Zipf's principle of least effort and Juran's rule of the vital few as having clarified and extended Pareto's original work. He also mentions the early contributions of Edward Deming and IBM.

However, he believes that the biggest breakthrough in terms of explaining such distribution patterns has only come with the advent of chaos theory, arguing that "[b]oth chaos theory and the 80/20 Principle assert (with a great deal of empirical backing) that the universe is unbalanced. They both say that the world is not linear; cause and effect are rarely linked in an equal way."

Mr. Koch does not leave us with only the generality that chaos theory and the 80/20 principle are related. A few pages later he notes in what ways chaos theory and the 80/20 Principle differ:

> These observations regarding sensitive dependence on initial conditions do not exactly illustrate the 80/20 Principle. The examples given involve change over time, whereas the 80/20 Principle involves a static breakdown of causes at any one time. Yet there is an important link between the two. Both phenomena help to show how the universe abhors balance. In the former case we see a natural flight away from a 50/50 split of competing phenomena. A 51/49 split is inherently unstable and tends to gravitate towards a 95/5, 99/1 or even 100/0 split. Equality ends in dominance; that is one of the messages of chaos theory. The 80/20 Principle's message is different yet complementary. It tells us that, at any one point, a majority of any phenomenon will be explained or caused by a minority of the actors participating in the phenomenon.

In arguing for chaos theory as the cause of the distributions we observe, Mr. Koch offers in passing that "no one else appears to have made the link," but on this point he is wrong, and the error points up a minor shortcoming of the book. My brief Internet search turned up numerous references to Pareto and chaos, including at least one physics paper from 1984 that explicitly examines the 1/f Zipf-Pareto distribution law in terms of "intermittent chaos". References to Pareto and chaos have already become noticeable in papers on "fat tails" and their many implications in finance, and date from well before this book first appeared in 1998. However, while *The 80/20 Principl*e occasionally falls short in terms of citing prior work, *The Power Laws* manages to give a diverse range of figures like Mandelbrot at least a brief mention.

The next section of the book, "Corporate Success Needn't Be A Mystery," describes how to apply the 80/20 principle in business. It includes several case studies, and here if nowhere else we do get mathematical details. In one case study, Mr. Koch describes how he assisted Electronic Instruments Inc. (a pseudonym) in identifying fifteen distinct competitive segments, and generating revenue and profit figures for each. These segments varied enormously in relative profitability.

The top six represented 26.2 percent of revenue, but almost 83 percent of profits; the bottom three represented 16.7 percent of revenue, and subtracted 31 percent of profits. This left a mediocre middle group of six segments that contributed 57 percent of revenues and 47 percent of profits.

After considering such questions as the maturity of each market, the strength of the competition, and whether the relevant production lines needed additional investment, Electronic Instruments Inc. arrived at a strategic plan. Two loss-making segments were closed or sold out, and the third received new investment to make it more competitive. Meanwhile, one of the mediocre middle group was put in "harvest" mode, continuing to make sales but canceling any further investment. The most profitable segments were aggressively expanded.

The charts associated with each case study lay out a classic Pareto-style root cause analysis, so that even non-mathematicians should find it easy to identify problem areas. In the case of a publishing firm, it is quickly apparent that authors are to blame for nearly 80 percent of typesetting overruns.

Cause	**Number**
Authors late with corrections	45
Authors late with original manuscript	37
Authors make too many corrections	34
Figures need correction	13
Book longer than planned	6
Proofreader late	3
Index compiler late	3
Permissions received late	2
Typesetter's computer fault	1
Typesetter's correction errors	1
Schedule changed by editor	1
Schedule changed by marketing	1
Schedule changed by printer	1
Fire at typesetter's	1
Legal dispute with typesetter	1
Total	150

So far this is familiar territory. I suspect that most MBA students would be capable of carrying out such a basic analysis already, if asked to, and the same goes for many engineers. However, the number who fully appreciate the importance of the law is probably much smaller—and this is Mr. Koch's most important point.

There are several interesting side arguments in this section that deserve illustration using actual case studies, but never quite get it. For example, Mr. Koch observes that in the 1970s, the Fortune 500 firms collectively held about 60 percent of the U.S. gross national product, and that by the 1990s, that share had fallen to 40 percent. He asks, "Why does it happen that in practice, as opposed to theory, the advantages of scale and market share fail to translate into higher profitability? Why is it that firms often see their sales mushroom yet their returns on sales and capital actually fall, rather than rise as the theory would predict?"

He then answers by saying, "The problem is not extra scale, but extra complexity. Additional scale, without additional complexity, will always give lower unit costs. . . . Yet additional scale is rarely just more of the same. Even if the customer is the same, the extra volume usually comes from adapting an existing product, providing a new product, and/or adding more service." Complexity in effect aggravates the 80/20 tendency, lengthening the low-value "tail" of the distribution and lowering the average return.

There is a wealth of potential insight contained in this idea, which seems to me a potential elaboration on the Wright-Henderson Law, but Mr. Koch's treatment amounts to part of one chapter, about a half-dozen pages in all. Only one brief real-world example and one survey of German companies are offered to document the claim. Mysteriously, neither Theodore Wright nor Bruce Henderson are mentioned by name, although Mr. Koch has credited Henderson in past books as well as in *The Power Laws*.

This particular side argument goes to the heart of a growing debate about how far the "experience curve" principle can actually be extended. (For more on Wright-Henderson and the many connections between strategy consultants and nonlinear thinking, see the editorial starting on page 19.)

Though well written, this entire section is still in the nature of a preliminary; we have not yet arrived at the climax of the argument. In the midst of case studies and statistics, the author specifically warns several times against embracing specific Pareto-related techniques such as root cause analysis, benchmarking or reengineering, without first resolving to think in nonlinear terms about the business as a whole.

Mr. Koch very effectively reminds us that:

- 80 percent of the profits made by all industries are made by 20 percent of industries.
- 80 percent of the profits made in any industry are made by 20 percent of firms.
- 80 percent of any firm's profits come from 20 percent of its customers.
- 80 percent of the value perceived by customers relates to 20 percent of what an organization does.
- 80 percent of the benefit from any product or service can be provided at 20 percent of the cost.

Therefore, as Mr. Koch puts it, "Unless you have used the 80/20 Principle to redirect your strategy, you can be pretty sure that the strategy is badly flawed." There is no point in making incremental improvements to individual product lines or processes without first making sense (in these terms) of the overall business context—and your own personal context as well.

This is the theme of the third and longest section, "Work Less, Earn and Enjoy More". Here Mr. Koch can again claim to have come up with some original ideas. Unfortunately, as more than one reviewer has noted, the 80/20 principle does seem to apply, so that the longest section of the book is also the thinnest in terms of specific content. There are five topics: time, relationships, effort versus reward, money, and happiness.

Mr. Koch's most interesting argument is about time. He insists that contrary to "time management evangelists," we do not actually suffer from a shortage of time. Most of our really significant gains are made with a very small investment of time. This suggests that our *real* problem is to identify opportunities to make such gains; finding the time to pursue the opportunities is relatively easy. We should accordingly "declare a time revolution" and focus on what really matters.

At first glance this seems to rest on straightforward Pareto analysis. A factory suffering from machine tool breakdowns can resolve most of its problems by identifying the 3 or 4 most common causes of breakdown, and eliminating them; it would take much longer to identify and eliminate *every* cause of breakdown. At some point, diminishing returns will set in and we will be better off putting our energies into something else. The same is surely true of everyday life—pursuing happiness has much in common with pursuing profit.

However, I found myself wondering what Mr. Koch would make of the "all or nothing" situations that commonly arise in both business and everyday life. No doubt the space shuttle would be 80 percent ready after we had completed the easiest 20 percent of the pre-launch checklist—but the remaining 80 percent is still unavoidable, if the shuttle is to launch successfully. The pattern identified by Pareto does indeed hold, but it cannot necessarily help us shorten the task if none of the items are optional. I am sure many parents would cite child care as a form of all-or-nothing responsibility that similarly defeats Pareto analysis—there are plenty of easy tasks in caring for an infant, and some much more difficult ones, but knowing which are which does not relieve us from performing them all.

Mr. Koch does answer this concern indirectly. He offers advice for those who doubt whether they can give up their low-value activities, because of pressure from family or colleagues: "Try your new policy and see what happens. Since there is little value in the activities you want to displace, people may not actually notice. Even if they do notice, they may not care enough to force you to do them if they can see that this would take major effort on their part."

This advice forms part of a major philosophical theme in Mr. Koch's work, something that I have not seen addressed in any other reviews of the book. The 80/20 principle is easy to outwit (or exploit) on a piecemeal, ad hoc basis, by discarding processes, products, customers and relationships that offer little value and seeking out those that offer more. It is not so easy to overcome the principle on a global, social level. We are unlikely ever to see a world in which 80 percent of firms have high profits, and 20 percent low, or an income and property distribution in which the rich outnumber the poor—or even a world in which 80 percent of mutual fund managers perform above the mean.

Therefore a successful 80/20 life strategy will consist of two distinct types of policy. First, we should seek to isolate 80-percent opportunities in which the low-value element of the distribution can simply be discarded, leaving no identifiable "losers," as for example in Pareto's classic machine-shop case. However, we should also forthrightly accept the fact that not everyone's work can have equal value, and strive to be among the 20 percent of winners. When low-value tasks form an unavoidable part of some process, we should do our best to delegate them. Let the other guy's company service the unprofitable markets! In his defense of capitalism and the morality of self-interest, Mr. Koch comes across as a forthright libertarian—and in his final section, "Crescendo," he has some further interesting things to say about Pareto's reputation as fascism's equivalent of Karl Marx.

In turning to relationships, Mr. Koch proposes a thought experiment. Make a list of your top 20 personal relationships, he says, and assign each one a percentage value, so that their importance to you adds up to 100. The resulting distribution will nearly always tail off sharply, such that the top 3 to 5 entries make up most of the value.

This is interesting, because it is a rare instance of an investigator proposing to apply Zipf or Pareto-type laws to strictly mental variables—not to measurements such as time spent, but on the subjective *perceived value* of that relationship. If Zipf's Law or some other existing law applies to such variables, the implications are enormous.

The final section, "Cresendo," takes up Pareto and the 80/20 principle at the global or social level. The subtitle here is "Progress Regained," and the message is unflinchingly optimistic. Pareto's work on the incomes of elites became closely associated with both conservatism and fascism, and to some extent any investigation of economic distributions became tainted by that association. In making the 80/20 principle a central element of our moral philosophy, Mr. Koch asks, are we conceding that the 80/20 principle is inherently fascist or conservative? He then proceeds to answer his own question: No, because the deeper implications of the principle are not conservative at all. "At its heart, then, the 80/20 Principle is not just descriptive, and it does not glorify what it describes. It is *prescriptive*; it observes a failure to reach anything like an optimal state of affairs, and it points the way to great improvements in the status quo."

The 80/20 principle is easy to outwit or exploit on a piecemeal, ad hoc basis . . . but we are unlikely ever to see a world in which 80 percent of firms have high profits, and 20 percent low.

There are three elements to the argument here. First, that the absolute level of wealth in a society is ultimately more important than the relative distribution. Standards of living have skyrocketed upward such that the lower middle class in most countries have better health, housing and education than even the wealthiest 1 percent did two centuries ago. Focusing on the really productive 20 percent of activities will ensure that total wealth continues to rise. Given Pareto's thesis that the overall distribution does not get any worse but stays relatively constant, we are all better off in the end.

Second, that state interventions inevitably tend to fall in the low-value end of the performance range. The 80/20 principle flourishes where individuals are free to seek out opportunities with high returns, and in that sense an 80/20 philosophy is inherently biased in favor of capitalism. Mr. Koch endorses greater school choice, deregulation, lower taxation, and similar policies on these grounds.

The third point, concerning Pareto himself, is more implicit than explicit. I take Mr. Koch as arguing that Pareto was fatally misunderstood by the authoritarian right and equally by their critics, when he said that history depends on elites. It is true that history is made by elites, but primarily in the sense of performance, *not* authority. "Progress always comes from a small minority of people and organized resources who demonstrate that previously accepted ceilings of performance can become floors for everyone." Thus Thomas Edison and Leonardo da Vinci belong to "elites" in this sense, while Mussolini, a self-styled authoritarian elitist, actually serves as figurehead for a vast, inefficient tail of statist regulation.

The "Crescendo" is in many ways the most audacious and difficult part of the book, because it seeks to reconcile Pareto the engineer with Pareto the sociologist, and to reunite two historically very distinct bodies of thought by restating their common underlying principle: the 80/20 principle should be understood as *a search for excellence emerging from chaos*.

Mr. Koch invites us to re-imagine Vilfredo Pareto, both as a visionary of a free and dynamic society, and at the same time as inventor of the key mathematical tools needed to function in that society. He acknowledges that we "can only explain the 80/20 principle by positing some deeper meaning or cause that lurks behind it." He also acknowledges that despite his best efforts, "Pareto's sociology failed to find a persuasive key."

Thus the challenge Mr. Koch presents is simple: we have had a century of scientific progress since Pareto's first attempts, and a century of failed experiments in social engineering. Can what we now know about chaos theory, markets, and the limits of political authority help us to discover the key that eluded Pareto?

About *The Power Laws*

To summarize Mr. Koch's more recent book is in some ways a forlorn task, because the book itself is a fast-paced summary. There is no completely satisfactory way to address 93 different laws in just over 300 pages, but the result in this case is both readable and relevant.

The opening chapters on natural selection, genetics, survival by differentiation, evolutionary psychology, and game theory are of somewhat less interest here because they are not focused on size laws, but they are well organized and clear. Nonlinearity appears as a theme throughout the book. For example, perhaps because Mr. Koch views it as organizational learning and thus biological in nature, the Wright-Henderson experience curve actually gets its first mention in the chapter on natural selection.

Interestingly, no mention is made of Michael Rothschild or *Bionomics*, despite some striking similarities in viewpoint. The lack of mathematical detail is less objectionable in dealing with biological analogies, as outside of game theory and some helpful nonlinear principles thrown in, there simply isn't much available.

The middle third of the book relates business to the laws of physics. Here I perceive a flaw: in his eagerness to cover as much science as possible, Mr. Koch occasionally brings in laws or principles whose relevance to business seems strained. He begins with a chapter on Newton, exploring the analogy of "corporate gravity", followed by a chapter on relativity, which Mr. Koch turns into an extended metaphor on (a) how to manage our time, and (b) the difference between perception and reality. In introducing the topic, Mr. Koch remarks, "Einstein's theories are difficult enough to apply in science; can they really be applied in business? In a strict sense, the honest answer is 'no'."

Undeterred, he cites a study by the Boston Consulting Group, which reads in part, "Typically, less than 10 percent of the total time devoted to any work in an organization is truly value-added." What follows echoes a lengthier argument from *The 80/20 Principle*; but the earlier effort, free of the burden of the relativity metaphor, is better focused. The same problem occurs in the chapter on quantum mechanics. I suspect that had the book been organized around just two themes—biological analogies in business, plus nonlinear mathematics—it would also have reflected the state of advanced business thinking much better.

The final third of the book is entitled, "The Non-Linear Laws". It includes chapters on chaos, the 80/20 principle, and tipping points, concluding with a chapter on entropy and the law of unintended consequences. It keeps to the thesis established in *The 80/20 Principle*, that waste, failure and mediocrity outnumbers efficiency, success, and greatness, and the latter emerges in highly unexpected ways. Mr. Koch observes that of all the laws presented in the book, the 80/20 principle is "the most universal".

Where the book really shines is in Mr. Koch's occasional efforts to reconcile different nonlinear laws and synthesize a single, overarching understanding. For example, he explains the relationship between Moore's Law and the Wright-Henderson experience curve—something few writers ever do, despite endless journalistic references to Moore's Law. Gordon Moore, founder of Intel, observed in 1965 that the computing power of individual silicon chips was doubling every year. The reconciliation with Wright-Henderson lies in showing why the rate of increase has been slowing. The doubling period had increased to two years by 1975, and in 1999 Moore had observed that doublings now occurred every four or five years. This is consistent with the experience curve, which predicts that gains in productivity become progressively smaller as experience grows.

We can only guess what Mr. Koch might turn his attention to next. Whatever it may be, readers are well advised to pay attention.

Variance & Ratio Analysis

by David G. Coderre

In the same way that financial ratios that give indications of the relative health of a company, fraud ratios point to possible symptoms of waste, abuse, and fraud. In this article we will examine the techniques of ratio analysis for fraud auditors.

Ratio analysis is possibly the most powerful identifier of potential fraud, even more powerful than its better known cousin—digit-frequency analysis using Benford's Law. While both examine the data for patterns to highlight possible fraudulent transactions, ratio analysis not only highlights anomalies in large groups of records, but also pinpoints the specific transactions that are unusual.

Digital analysis using Benford's Law will identify a set of records whose amounts begin with the same 'x' digits more often than we expect. This often signals some kind of unusual human intervention, such as a data entry error, or a wasteful duplication of some activity, or (in some cases) an ongoing fraud in which numbers are being invented.

For example, we might find in a very large corporate accounts-payable file that '306' occurred 39,721 times, instead of the expected number of times (23,392). The problem remains, however, to determine which of the many transactions starting with '306' are inappropriate, and which are correct. Of course, we can carry out additional analyses of the '306' records, such as taking subtotals for each contracting officer and each vendor, and this will often reveal a trend worth investigating. The drawback of such techniques is that they rely on volume to work—a false or incorrect transaction that appears exactly once will not draw attention to itself in terms of digit frequencies. In contrast, ratio analysis can identify the exact transactions that are unusual and, possibly, fraudulent. No additional analyses are required, aside from applying a filter.

Ratio analysis identifies potential frauds by computing the variance in a set of transactions and then calculating the ratios for selected numeric fields.

Three commonly employed ratios are:

- the ratio of the highest value to the lowest value (maximum/minimum);
- the ratio of the highest value to the next highest (maximum/2nd highest); and
- the ratio of one numeric field to another, such as the current year to the previous year or one operational area to another (field1/field2).

Maximum/minimum ratio

For example, users concerned about prices paid for a stock-numbered item could calculate the ratio of the maximum unit price to the minimum unit price for each stock number. If the ratio is close to 1, then they can be sure that there is not much variance between the highest and lowest prices paid. However, if the ratio is large, this could be an indication that too much was paid for the product in question. Even if thousands of products were purchased, the ratio analysis will focus attention on the few products with a significant variance in the unit price. In addition, it will identify the specific unit price, for the identified product, that is abnormal.

Product No	Max	Min	Ratio
Product 1	235	127	1.85
Product 2	289	285	1.01

Product 1 has a large difference in the unit price between the minimum and maximum (ratio of 1.85); whereas Product 2 has a smaller variance in the unit prices (ratio of 1.01). Audit should review the transactions for the unit prices of $235 and $127 for Product 1 to ensure proper payments were made, or in the case of accounts receivable, that payments were received. The auditor may want to determine why abnormally high unit prices were paid/charged for a particular product. Unexplained variance in unit prices paid/charged for items could be a symptom of fraud. In the accounts payable area it could be indicative of collusion between the purchaser and the vendor—involving kickbacks for excessive payments. You could be paying too much for a product, or receiving a product of inferior quality.

The range of analyses and, hence, the types of fraud that can be discovered, are not limited to unit prices or even to accounts payable or receivable.

Errors in accounts payable (case 1)

An audit of the A/P section was underway and ratio analysis was employed to highlight any unusual payments made to vendor. The A/P section handles close to one million invoices a year and produced checks totaling over $10.3B. Ratio analysis was used to examine pattern of payments made to vendors.

A large ratio (Max/Max2) indicates that the maximum value is significantly larger than the second highest value. Auditors and fraud investigators would be interested in these unusual transactions as they represent a deviation from the norm. Unexplained deviations could be symptoms of problems, for example, high ratios suggest payments made to the wrong vendor.

ABC Corp was paid 314 times, with 313 payments falling in the range $123.93 to $5,441.62. Then there was a single payment of $110,345.21, more than 20 times the next highest payment. A review determined that the purchasing manager, who was retiring in two weeks, had made a major purchase with a company that had offered him a job.

XYZ Ltd. had 762 payments in the range $749.34 to $3,461.97 and one of $23,489.92. The largest payment - 6.79 times larger than the next highest - should have gone to XZY Ltd, but was paid to XYZ Ltd in error. When XZY Ltd. complained that they had not been paid, the accounts payable clerk checked all payment to XYZ Ltd and did not see the payment—since it was made to XZY Ltd—and a rush check was cut for them.

Vendor	Max	Max2	Max/Max2	Min	Max/Min	Cnt
ABC Corp	$110,345.21	$5,441.62	20.28	$123.93	890.38	314
XYZ Ltd	$ 23,489.92	$3,461.97	6.79	$749.34	31.35	763
ZZZ Co.	$299,345.00	$4,281.03	69.92	$442.28	676.82	140

Table 1. Ratio analysis results for a series of vendors

Detecting payroll fraud (case 2)

The auditors were engaged in a routine review of overtime payments made over the last year. The audit typically sought to ensure that prior approval was received and the appropriate forms were completed to ensure prompt and accurate payment. Further, given the thousands of overtime records, sampling was traditionally used.

This year, the auditors decided that the review of overtime payments would be an ideal candidate for variance/ratio analysis. The ratio of highest to second highest (max/max2) and highest to lowest (max/min) was calculated for every job category. The most significant results were as shown below.

Job Title	Max	Max2	Max/Max2	Min	Max/Min	Cnt
Payroll Clk	$948.36	$241.62	3.92	$126.83	7.47	12
Clerk Lvl 2	$800.30	$321.99	2.49	$49.34	16.22	77
Mgr Lvl 1	$883.07	$883.07	1.00	$883.07	1.00	1

A detailed review of the payments determined that a payroll clerk had fraudulently manipulated the system to pay herself almost 4 times the next highest overtime payment. In the second case, a Clerk Lvl 2 managed to receive almost 2.5 times the next highest overtime payment. The final record, a Manager Level 1 received an overtime payment, even though they were not entitled to overtime - hence only 1 overtime payment had been made to employees in the Mgr Lvl 1 category.

The payment to ZZZ Co for $299,345.00 was different again. It was a mistake in data entry. The amount should have been $2,993.45 which would fit nicely in the range $442.28 to $4,281.03.

It should be noted that not all of the above transactions were fraudulent—but identification of the errors and subsequent recovery was still significant. In the A/R area it could be a sign of collusion between a customer and the salesman—involving kickbacks for overcharging for items.

Corruption in contracting (case 3)

Johnathan, one of the contracting officers, had devised a great scheme in which he won and so did the companies who were willing to do business under his conditions. Companies who were not willing to provide him with a little extra would not get the contract.

The auditors decided to use digital analysis as part of their review of the contracting section. One of the analyses calculated the total contract amount by supplier for each of the past two years. A ratio of current year to previous year was calculated and the statistics command was used to look at the minimum, maximum, average and highest and lowest 5 ratios. While the average was close to 1.0 the highest and lowest 5 values showed that some companies had significant decreases in business, while others had experienced significant increases in business.

The auditors reviewed the detailed records for all companies that had a ratio of less than 0.7 or more than 1.30. The records were extracted to a file and totals were calculated by contracting officer. For companies that had seen an increase in business, the results revealed that Johnathan had raised many of the contracts. In comparison, Johnathan had raised no contracts with the companies that had seen a decrease in business. The auditors learned of Johnathan's kickback scheme when they interviewed salesmen from the companies that had ratios less than 0.7. Interviews with salesmen from the firms that had increased sales by 1.30 or more added credence to the fraud accusations. Both groups of salesmen said that they were told they would onlyget business if they paid Johnathan a kickback.

Field1/ field2 ratio

Analyses based on individual transactions, such as the max/max2 technique, are particularly attractive because

Inventory fraud (case 4)

The auditors were concerned about control weakness in the shipping department. Since there was no automated link between the inventory system, the contracting system and the billing system, the auditors usually selected a sample of shipping records and verified these to the contracts and the billing department. This year they decided to run a ratio analysis on the quantity shipped by Vendor. The results were interesting.

Company	Max	Max2	Max/Max2	Min	Max/Min	Cnt	Shipping Clerk
ABC Ltd	130	12	10.83	10	13.10	48	Fradlen
Geo Tech	305	30	10.17	30	10.17	22	Fradlen
X-Tech	200	20	10.00	10	20.00	17	Fradlen

One of the shipping clerks, Fradlen, was shipping more than the ordered quantities. He would then call the companies and have them ship back the excess. The fraud occurred when he gave them a different warehouse as the 'ship to' location - his house.

The analysis also pointed out another interesting fraud scheme. A salesman who consistently had extra items shipped to customer - more than they had ordered - perpetrated this one. He hoped that the customers would simply pay for the extra goods rather than go to the trouble of returning the unwanted goods. If not, the items would be duly returned to the company. The fraud occurred because the shipping quantities were used to calculate the salesmen's quotas - from which their bonuses were derived. A certain amount of returns were expected so no formal reconciling was done to determine actual returns. The salesman took advantage of the customers who paid for more than they ordered, and abused the system by not reporting his higher than normal returns.

they pinpoint suspect transactions. However, the more traditional kind of financial ratio analysis, involving aggregate totals, also has a place in investigations.

This 'field1/field2' analysis compares the data from two years, or two different departments, operating areas, and so on. For example, the ratio of current year's purchases to last year's purchases for each supplier can find symptoms of possible fraud, such as kickbacks in the contracting section. If the total purchases from a supplier has gone from $100,000 to $400,000, for a ratio of 4.0, or from $400,000 to $100,000, for a ratio of 0.25, audit might beinterested in performing additional analysis to determine why there was such a large increase or decrease in purchases. In either case, further analysis, such as computing totals by contracting officer, may identify kickbacks—where one firm benefited by paying the kickback and the other firm lost business because they refused to pay the kickback.

Setting realistic thresholds

For a large file, for example an accounts payable system involving hundreds or thousands of vendors, one challenge that arises is setting an appropriate threshold. Shall we investigate all max/max2 results greater than 2.0? Can we say in theory what number to use?

If each set of payments to an individual vendor followed Zipf's Law (see the article starting on page 1), such that the average value of max/max2 was 2.0, choosing 2.0 as a cutoff would lead to our investigating every result that was even marginally greater than average—close to half of the vendors on file. This is likely to be too many to be efficient. But in fact, payments to individual vendors more typically follow an 'ontic' distribution, and are more closely spaced in size, so setting a threshold of 2.0 will usually mean investigating relatively few items.

In practice, we do not yet have any general laws that we can apply to set a ratio threshold. Auditors should be guided by what they find in their own data, from year to year and file to file. In fact, a good policy would be to compute the max/max2 ratios for all vendors, save the results to a file, and then stratify the file to determine what number of vendors typically exceed a ratio of 1.25, or 1.50, and so on.

This parallels a finding many have made when performing digit-frequency analyses using Benford's Law. Because of statutory requirements, or unique practices, some industries have data that never conforms to the expected digit-frequency values laid down by Benford. However, the resulting digit frequency values unique to that business are generally consistent from year to year. A 'custom Benford' analysis using last year's data as an expected baseline can still detect deviation. The same is true when we are setting our thresholds for investigation in ratio analysis.

Conclusions

The use of the fraud ratios can highlight hidden fraud by focusing audit attention of the few transactions—representing payments to vendors, credit card purchases, or whatever they may be—that are unusual based upon the patterns in the data. The beauty of variance analysis is that the ratios are calculated based on your data—highlighting anomalies in the data compared, not to theoretical or industry benchmarks, but your data. The ratios highlight anomalies in your data, as compared to your data. Further, even if there are millions of transactions, the anomalies will stand out.

For example, consider an A/P system that pays 1 million invoices every year to 100,000 vendors. The standard audit practice of selecting a sample of 30 transaction is not going to reveal anomalies. But a variance analysis by vendor (note, each vendor only has an average of 10 transactions) will highlight, not only the total number and dollars by vendor, but also the highest, 2nd highest and lowest payments. The ratios max/max2 and max/min will focus on vendors with a large variance in their payments, finding the few anomalies among the million payments. Finally, running the ratios for several years can define trends in the data—making it easier to spot possible frauds as they occur.

While I have made a case for the power of ratio analysis over Benford's Law, I should point out that different types of fraud will require different types of analyses. No one analysis is always going to find the fraud—you need to ensure your fraud toolbox is loaded with many tools and techniques. So consider carefully the risks your company is exposed to, and the potential frauds—then choose the appropriate tool for the job.

Editorial

Developing the Paradigm: The Legacy of Boston Consulting Group

The "vital few" engaged in bringing discoveries to the attention of the world have often had interesting connections to one another. This is certainly true in the case of four men once associated in various ways with Boston Consulting Group: Bruce Henderson, Michael Rothschild, Frederick Reichheld, and Richard Koch.

Their story actually begins with Theodore P. Wright, 1895-1970, and an article he wrote, "Factors Affecting the Cost of Airplanes," for the *Journal of the Aeronautical Sciences* in 1936. Wright described a method of estimating labor costs on long production runs of aircraft, using what he called a "learning curve". Wright found that each time the total number of units produced was doubled, the average cost fell to 80 percent of the previous level, in a classic 'log-log' relationship. Wright and his brief paper both came to have enormous impact on the planning of aircraft production during WW II and on the modern art of cost estimating.

Bruce Henderson and business strategy

Bruce D. Henderson, 1915-1992, founder in 1963 of Boston Consulting Group (BCG), and known by some as "the originator of modern corporate strategy", took Wright's discovery further. By the early 1960s the 'learning curve' effect had also been observed for Liberty ships and several other wartime items, and the outline of a general theory was emerging. In a 1966 study for Texas Instruments, Henderson established that Wright's observed productivity gain of 20 percent per doubling of production also applied to electronics, so that the cost of calculators (for example) could be expected to fall dramatically as production volumes grew. BCG and an array of academics went on to show that the curve applied to an incredible range of products, from motorcycles to toilet paper to crushed limestone.

Henderson argued that this predictable effect could and should have impact beyond the realm of factory cost-estimating. Because costs will continually fall with increasing volume—and any competitor's costs will fall in the same way—corporations should ultimately be left with only two practical strategies: aggressive pursuit of maximum market share, or if this proves impossible, divestiture. Staying ahead on the experience curve should lead to consistently lower costs than any competitor, creating a perpetual "cash cow" in the profit differential. Falling behind ought to mean fighting an ever-widening cost gap leading eventually to zero or negative profit.

In a 1973 essay entitled "Failure to Compete," Henderson argued that based on the experience curve, "The dominant producer in every business should increase his market share steadily . . . Failure of an industry to concentrate is failure to compete and a failure of the national economy to optimize productivity and reduce inflation." Henderson insisted that a "benign monopoly" was the inevitable fate of every well-organized industry, and that conventional economic thought (which thought in terms of static, rather than perpetually decreasing costs) was bankrupt.

However, closer scrutiny of the evidence showed that unlimited concentration did not actually happen in *any* industry. Henderson grappled for some years with the puzzling question of why market share does not increase indefinitely for the leader, but seems to stabilize in a consistent ranking pattern— in fact, closely following the ranking law discovered by Zipf several decades earlier, although Henderson did not credit Zipf in the paper and may never have been aware of this connection. In "The Rule of Three and Four" (1976), Henderson modified his argument to cover the observed facts: "For the leader, the opportunity diminishes as the share difference widens. A price reduction costs more and the potential gain is less. The 2 to 1 limit is approximate, but it seems to fit."

This somewhat *ad hoc* set of observations offered a foundation for successful strategy for many BCG clients in the 1970s and 1980s, and the Henderson approach became widely accepted in business, as well as (more slowly) in academia. The strategic useof the experience curve, the concept of cash cows, and the expand-or-harvest decision matrix, all still hold a powerful grip on managerial thinking. Boston Consulting Group remains an influential billion-dollar enterprise.

Michael Rothschild and bionomics

In the late 1980s, Henderson turned increasingly to biological analogies as a means of understanding the limi tations of economic thought. In the months before his death he became an enthusiastic supporter of the thesisof *Bionomics* (1990), written by former BCG colleague Michael Rothschild.

Mr. Rothschild in turn credited Henderson (here, in an article for *Upside* magazine) with nothing short of an intellectual revolution:

> By substituting new knowledge for labor and materials, experience-driven innovation keeps pushing costs down. As Henderson put it, when a firm is properly managed, its "product costs will go down forever." . . . . Without his “experience curve,” there is no final and fully satisfying explanation for falling costs, rising incomes, and the phenomenon of economic growth. More than anyone else, he made it both possible and necessary for economic thinkers to break free of the static, zero-sum mentality that has gripped the “dismal science” for 200 years.

As Mr. Rothschild observed, the major political conflicts of the 20th century were rooted in an outdated view of the universe. “Both Marxism and Western economics were established before Darwin. To date, neither side has ever seriously considered evolutionary biology as a metaphorical paradigm. Consequently, neither side has yet resolved the central dilemma of how the economy changes.”

The approach of *Bionomics* was more historical and philosophical than immediately practical, devoting more than half its length to cellular biology, analogies between organisms and organizations, early theories of evolution, and the failures of Marxism. Although Mr. Rothschild made the Henderson experience curve the centerpiece of the work, he offered little that was new on how engineers, managers or marketers could actually use it.

Seeking to turn the intellectual revolution into something more tangible, Mr. Rothschild founded the non-profit Bionomics Institute. For several years, word of mouth as well as favorable reviews (many by free market and libertarian publications) linked the bionomics thesis to the growing excitement over the New Economy and the seemingly infinite potential of the Internet.

However, in 1996 interest seemed to have peaked, and Mr. Rothschild moved on to found Maxager Technologies, a software company specializing in “enterprise profit management”. *Bionomics* is mentioned in the media less often now than formerly, and the Institute has been “in hibernation” for the past five years.

Frederick Reichheld and loyalty economics

While bionomics appears to have stalled, the success of BCG has led to much more than just one spinoff. William Bain, a former BCG employee, created Bain & Company in the early 1970s. A global consulting firm in the billion-dollar range, Bain has developed a distinctive ‘loyalty’ practice, with important implications for the future of the experience curve.

The founder of loyalty consulting at Bain was Frederick Reichheld, who joined the firm in 1977 and became its first Bain Fellow in 1999. Mr. Reichheld is the author of two widely sold books on loyalty—*The Loyalty Effect* (1996), and *Loyalty Rules!* (2001).

Although it may seem to be quite distinct from the earlier work of Bruce Henderson, and even opposed to it at points, Mr. Reichheld’s approach can best be understood as a much-needed expansion and elaboration. Where Henderson focused on productivity growth over a product lifespan, Mr. Reichheld has examined productivity growth and profit growth from a variety of angles, each time discovering the same underlying pattern—the same logarithmic distribution.

Just as a long production run will bring costs down further than a short one, a long-term customer tends to be more profitable than a short-term one—and so does a long-term employee, and so does a long-term investor. All such relationships follow a form of logarithmic learning or experience curve. In one of his key insights, Mr. Reichheld took issue in *The Loyalty Effect* with what he called the “misuse” of learning curves.

Because early work on learning curves focused on manufacturing, enthusiasts seeking to exploit the concept tended to do so as well. But a long-term employee in a service position can also produce 2 or 3 times as much as a new hire doing the same job, something a single-minded pursuit of manufacturing productivity can obscure. Says Mr. Reichheld,

> Eventually, these companies were ambushed by competitors who improved customer value through performance and quality innovations instead of cost and price reductions. This is what happened to Henry Ford’s Model T, which was available at wonderfully low cost, but in only one model and one color.

The real key, according to Mr. Reichheld, is not to maximize production volumes, but rather to cultivate a more subtle variable—loyalty—which can sometimes be measured in dollars or units of output, but is best measured in units of time, or else in the quality of the relationship itself. Mr. Reichheld wrote in 1996,

> We have spent some ten years studying loyalty leaders and their business systems, and what we have learned has radically altered our view of business economics, and led us to develop a very different model . . . What drives this model is not profit but creation of value for the customer, a process that lies at the core of all successful enterprises. . . . The physics that governs the interrelationships and energy states of a business system's elementary particles—its customers, employees, and investors—we call *the forces of loyalty*. Because of the linkages between loyalty, value, and profits, these forces are measurable in cash flow terms.

Here we see a bold re-statement of the meaning of economics, very much in the spirit of the Henderson revolution, but focusing on "the forces of loyalty" rather than a biological metaphor or a narrow measure of factory output. A company with the right kind of loyal employee will attract the right kind of loyal customers, and its profitable growth will then bring in equally loyal investors—a virtuous cycle of value creation.

As Mr. Reichheld notes in the introduction of *The Loyalty Effect*, many of his colleagues doubted on hearing his title whether any good could come of emphasizing loyalty in the chaotic world of 1990s capitalism. He has gone on to demonstrate, with a wealth of specific data, that concrete, measurable, reliable results can be obtained using loyalty principles. If our goal is a broad intellectual system for exploiting logarithmic distribution laws, he is arguably the first to lay out such a system and to make it work in real life.

However, Mr. Reichheld's thesis still suffers from a serious drawback: a subtle theoretical gap that critics have not yet succeeded in addressing, but that stands out when the work is seen in its historical context. When Bruce Henderson's paradigm of the "benign monopoly" met the reality of stable market shares, it was profoundly shaken: his explanation for why a market invariably stops concentrating was not nearly as lucid or compelling as his argument for why it should continue. The bionomics approach celebrated the power of the experience curve but did not make any headway in accounting for its limits. The same problem arises with the 'forces of loyalty' argument.

This problem is well illustrated by Chapter 7 of *The Loyalty Effect*, which is devoted to the lessons that can be learned from failure. Chapters 1-6 lay out the principles of loyalty economics, and examine what they mean for productivity or retention of customers, employees, and investors. Then, in this pivotal chapter, Mr. Reichheld begins by dismissing a key thesis of bionomics, that bankruptcy for companies is comparable to the life cycle of species:

> It's tempting to accept this high corporate mortality rate as the wholesome result of natural selection in the economic sphere. When companies fail to adapt to changing environments, they die—and a good thing, too. Otherwise we'd be burdened with the kind of useless dinosaur industries that weighed down the Soviet Union. But companies are not species, and the Darwinian analogy doesn't really work. Genes can't learn; it's only the relatively glacial process of successful random mutation that produces adaptive change in species. Companies have free will. They can learn from their failures, and they can change—or they can decline to learn and refuse to change—the choice is theirs.

The balance of the chapter is devoted to the psychology behind ignoring failures, and the merits of root-cause analysis in dealing with employee or customer defections. Mr. Reichheld's argument, that examining failure may actually be more effective than trying to imitate success, is certainly very relevant and useful. In arguing the merits of root-cause analysis he specifically cites Pareto, which suggests that we are still focused on size-law principles to some degree. However, it is at this point that we fail to come to grips with the critical question of what ultimately limits the loyalty factor or the experience curve—where these things stop working. Is it possible to forecast such failures as readily as we forecast gains in productivity? If not, why not?

The title of Chapter 7 is "In Search of Failure," parodying the 1982 classic by Tom Peters, *In Search of Excellence*. The theme of the chapter is that business failure is essentially a choice rather than an accident, and that companies that examine their failures will survive and prosper. Mr. Reichheld seems to be arguing that declining performance is not something that we can forecast in mathematical terms, but instead lies in an organizational refusal to learn.

Mr.Reichheld notes that of the companies profiled by Tom Peters, two-thirds have since underperformed the market, and only one-fifth would meet the same standard of excellence today. He explains their steep decline in terms of systems theory:

> When a system is working well . . . it's impossible to explain why. Its success rests on a long chain of subtle interactions, and it's not easy to determine which links in the chain are most important. Even if the critical links were identifiable, their relative importance would shift as the world around the system changed.

Sooner or later, Mr. Reichheld says, something does change and the system no longer works well. From the peaks of excellence we can look forward to an almost inevitable decline. So we are wiser to focus on understanding our failures, where we have room to move up.

This is as close as we come to a theoretical explanation for the limits we perceive to the experience curve, or an attempt to systematically forecast failures. While what Mr. Reichheld says about systems may be true, it does not go far enough. It raises, or ought to raise, the disturbing doubt that everything we have been told so far about experience curves and loyalty principles is partly the result of selective inattention: we are seeing the success stories celebrated, and contrary evidence has been set aside, without serious discussion.

There is in fact a great deal of empirical evidence to show that organizations that stick with a standardized product or methodology can actually perform less and less well over long spans of time—and also that despite the experience curve, many companies do actually *decline* in overall productiveness as they grow in size. These trends are no less real than the never-ending productivity growth first pointed to by Wright and Henderson, and they deserve equal recognition. These trends are alluded to in Mr. Reichheld's books, in Richard Koch's books, and elsewhere in the literature. However, to date no one has succeeded in producing an integrated theory that predicts when performance will *decline* logarithmically, as opposed to continually forecasting that it will improve.

As an example: several times Mr. Reichheld discusses the importance of keeping teams as small as possible, in many cases between 5 and 10 people. He notes that this principle has long been recognized by military organizations. His examples make it clear that small teams are more efficient, more profitable, and more accountable. What is not so clear is how we are to reconcile this with loyalty principles. For example, if we are managing employee turnover well so that as we add employees, we do not suffer a loss of experience (and thus per-employee productivity), then why would a larger team be less effective than a small one? Is there something about having a greater number of people doing the same work that is inherently inefficient? If so, what explains the tremendous economies of scale that we so often observe?

This is where we can examine the ragged edge of the experience curve paradigm, and realize that in some ways it has not advanced at all since 1936. Mr. Reichheld cannot be faulted as being naive or as withholding facts from the reader. Instead, there is a conflict between what he can tell us from a theoretical perspective, and what he simply knows anecdotally. Many important insights in Mr. Reichheld's work are offered as asides, and are not necessarily integrated with the overall loyalty theme. When Mr. Reichheld's professional experience and Bain's research both point one way, and the loyalty concept seemingly points elsewhere, there is not much that he can do except to tell the story as he knows it.

Mr. Reichheld's works have been hailed as "putting loyalty economics on the map", and clearly represent a major step forward from where we stood in the 1960s. Yet it is also very clear that with this new loyalty model we are still far from really understanding the logarithmic laws that surround us.

Richard Koch and the 80/20 principle

Richard Koch is a former BCG consultant and Bain partner who left in 1983 to co-found another consulting firm, The LEK Partnership. He later took over the ailing Filofax company and returned it to profitability. In the 1990s he advised venture capital groups and Internet enterprises in the UK and South Africa. He has said he has two passions, ideas and horse racing. His restless professional journey stands in interesting contrast to that of Mr. Reichheld, who has stayed with one firm and one focus for the past 25 years.

Mr. Koch has written eleven different books on business, the most recent two on the power of logarithmic size laws (or as he calls it, the 80/20 principle). His approach, as described in the separate review in this issue, is the most sweeping to date, taking the principle beyond business or engineering into our mental life, habits, and relationships.

Mr. Koch seems to pay more attention than his predecessors to what we might call the underlying perversity of the universe, meaning the apparent tendency for weak performances to predominate as opposed to strong ones. His chosen explanatory framework is not experience, biology, or loyalty, but rather chaos, and this makes his contribution, in my view, the most important to date.

This is the reason for our separate article-length review of his two most recent books. The writings of Bruce Henderson, Frederick Reichheld, and Michael Rothschild all deserve serious study, beyond what we have given them here—but I would argue that Mr. Koch's approach shows the greatest promise of advancing our understanding of the mathematics of failure.

Challenges & opportunities

The story of BCG and its legacy has many lessons to offer. First, as I said at the outset, it underlines the tendency for knowledge to spread outward from a very small minority. In the terminology commonly used by marketers, these men were in part "early adopters" of a powerful new technology, and in part pioneers in their own right. They applied the original insight of Theodore Wright to an ever-growing range of problems, and changed our understanding of business.

Second, and more trivially, it is regrettable that books written for business readers by professional consultants do not consistently reference the work of competitors, and skim over contemporary academic work. For example, Mr. Reichheld's two books do not mention Bruce Henderson, Michael Rothschild, or Richard Koch (or for that matter Theodore Wright). The story of BCG, its origin, and its legacy in these four men, is an important chapter in the history of ideas. Yet to my knowledge, it has never been told before. It is impossible to perceive these intellectual connections without careful background research.

Third, there is a tremendous need now to put the experience curve, the Pareto 80/20 principle, the loyalty paradigm, Zipf's Law, Benford's Law, and our whole understanding of size laws, onto a more rigorous foundation. This is crucial for many reasons.

Our early intuitions about cause and effect, and the limits of applicability to these laws, have proven wrong several times over. These laws keep turning out to have wider application than we suspected, but also to run up against baffling limits such as the market share paradox. And at the same time that popular enthusiasm and business interest in these laws has been peaking, the original Wright curve has been coming under increased academic scrutiny, because the early explanatory metaphor of "learning," later modified to "experience," has been found not to fit the facts as well as was once believed.

Research teams reviewing classic productivity-growth studies from the 1940s through the 1970s have found that investments in physical capital, acceptance of lower quality product, and many other factors are needed to achieve productivity growth, and that true learning by doing (sometimes called the Horndal effect, after an iron works in Sweden in the 1930s) is governed by employee turnover and "forgetting," and therefore varies widely in how effective it can be.

Some of these points are already dealt with in the work of Mr. Reichheld, but his arguments are vulnerable as well. The research findings of the Bain loyalty team are much newer, and have not been subjected to anything like the decades of independent scrutiny received by Wright's original thesis. When the loyalty approach is examined more closely, the ultimate causes of the loyalty effect are likely to prove elusive, even as the effect itself remains demonstrable.

One of our goals at *Frequencies* is to help lay out the rigorous foundation needed to deal with all these paradoxes. We begin in this issue with an examination of Zipf's Law, Benford's Law, and a deeper understanding of their relevance to ratio analysis. In Issue 3 we will continue with an additional major article on forecasting using Zipf's Law, and we will review the history of Wright-Henderson and other productivity curves in detail, showing how they are applied practically. Starting with Issue 4, we will present a radical re-interpretation of Wright-Henderson and the loyalty paradigm, to explain how *declining* performance over time and *improving* performance over time can actually be explained by a single unified theory, resting on Pareto, Benford, and Zipf—namely, the ontic distribution.

These are exciting times and exciting topics. We hope you will stay with us for the entire journey.

Contributors wanted

We are actively soliciting new material. If you want to write for *Frequencies*, contact us at:

frequencies@ekaros.ca

How to subscribe

Frequencies is published quarterly in Vancouver, Canada. A two-year subscription is $275.00 US for eight issues, overseas mailing charges not included. The unused portion of your subscription will be refunded if for any reason you are not satisfied.

Please be sure to provide your full name, mailing address, e-mail and telephone number. Send your cheque in U.S. funds or international money order to:

Frequencies
c/o Ekaros Analytical Inc.
125A - 1030 Denman St., #408
Vancouver, B.C. V6G 2M6

Anyone needing to receive an invoice before making payment should mark their subscription request accordingly. E-mailed subscription requests should be directed to:

frequencies@ekaros.ca

Frequencies staff

Editor and publisher Dean M. Brooks is an engineer and technical writer with 15 years' experience in defense, aerospace, and software. He was the editor for Mark Nigrini's groundbreaking *Digital Analysis Using Benford's Law,* published by Global Audit Publications in May 2000.

Contributing writers in this issue

David G. Coderre is the author of three books: *CAATTs and other BEASTs for Auditors*; *Fraud Detection: Using Data Analysis Techniques to Detect Fraud*; and the *Fraud Toolkit for ACL*. He has 18 years' experience in the audit profession and frequently lectures on audit-related topics. He can be reached at dcoderre@cyberus.ca.

In this Issue

Coming Next Issue

Competitive Analysis Using Zipf's Law
by Dean Brooks

The Experience Curve & Increasing Returns: An Historical Review
by Dean Brooks

See the editorial, p. 23, for details of these upcoming articles.

Upcoming Events

August 12, 2002

Dean Brooks will speak at the national convention of the American Statistical Association, in New York City, on the topic:

"Naked-Eye Quantum Mechanics: Practical Applications of Benford's Law for Integer Quantities"

> ABSTRACT: Benford's Law (1938) predicts that digit frequencies for many scientific, engineering, and business data sets will follow P(dd)=log(1+1/dd) for digits dd. This law has been used by auditors since 1989 to detect errors and fraud in data sets. Benford also postulated a separate law for integer quantities. This little-known variant of the law is shown to be substantially correct, despite an error by Benford in its derivation. The integer variant is then shown to be extraordinarily common in everyday life, correctly predicting the distribution of footnotes per page in textbooks, sizes of groups walking in public parks and visiting restaurants, fatality counts in air crashes, repeat visits to service businesses, and purchasequantities for goods.

For more details contact us at: *info@ontics.com*