The Emerald Research Register for this journal is available at
www.emeraldinsight.com/researchregister

The current issue and full text archive of this journal is available at
www.emeraldinsight.com/1741-0401.htm

PROFESSIONAL PRACTICE

Business analytics: getting behind the numbers

Jan Emblemsvåg
Fora Form AS, Ørsta, Norway

Received August 2004
Accepted September 2004

Abstract

Purpose – To identify, discuss and provide a solution for a common problem in the mathematical analyses in business analyses, namely, paralysis by analysis.

Design/methodology/approach – The paper first discusses the scale and frequency of the paralysis by analysis problem, before discussing it in more depth before addressing a fundamental problem, which is an important root of the paralysis by analysis problem, the indiscriminate usage of central tendency measures. Finally, it discusses how variance can be turned from being a liability into an asset. The approach is conversational but examples and a case study are provided to substantiate the arguments.

Findings – The paper provides some recommendations for avoiding paralysis by analysis.

Practical implications – Basically, the paper shows by argument and example why practitioners and some researchers need to better understand the limitations and promises of mathematical analyses and to some extent how to incorporate this understanding into their work.

Originality/value – There is nothing really new in this paper, but it discusses a problem that for some reason is often ignored by practitioners and some researchers. The true value of the paper therefore lies in making practitioners, in particular, more aware of the limitations as well as the possibilities in the mathematical analyses performed in business analytics so that they can better understand what they are doing and hence get behind the numbers, as it were.

Keywords Business analysis, Mathematical analysis, Statistical methods of analysis, Analysis of variance

Paper type Viewpoint

1. Introduction

Data! Data! Data! . . . I can't make bricks without clay (Sherlock Holmes).

Studies emphasize the importance of analyzing financial data and other information – performing business analyses (BA) – to gain insight about markets, competitors, internal processes and so on is beyond doubt. For example, the IDC study in 2003 – "The Financial Impact of Business Analytics" – found a ROI (return on investment) ranging from 17 per cent to 2000 per cent with a median of 112 per cent. The average five-year ROI of analytics implementations was 431 per cent, with 63 per cent of the corporations having a payback period of two years or less (see, for example, Groh (2004)). The issue is therefore not whether we should perform BA or not, but how. Particularly, how do we avoid the common paralysis by analysis and how do we get the most insight out of the numbers? Some simple and useful avenues are discussed in

The views presented in this paper are those solely of the author.

International Journal of Productivity and Performance Management
Vol. 54 No. 1, 2005
pp. 47-58

1741-0401
DOI 10.1108/17410400510571446

Sections 2 and 3, but first a discussion of what constitutes insight – or knowledge to be more precise – is provided.

All mathematical calculations use data as input – not information, as some seem to believe. Information can only be obtained by providing attributes or relevance and purpose to data (see Joia (2000)). Thus, numbers by themselves do not constitute information, which is why information technology (IT) is not really information-oriented by default but data-oriented. Only the human mind can provide purpose and hence relevance to data, which of course can subsequently be built into data systems to provide information systems.

The next step is to turn information into knowledge, where knowledge is defined as "information combined with experience, context, interpretation and reflection" (see Davenport *et al.* (1998)). Furthermore, Kock and McQueen (1998) point out that information is "descriptive and historical, relating primarily to the past and present whereas knowledge is predictive and associative and unveils hidden facts". Thus, no BA can provide knowledge or insight directly, and often will the information derived from BA require further work. In other words, BAs are more attention directing instruments than knowledge generating instruments.

But knowledge by itself is also not enough – we must truly understand an issue to be able to act upon it wisely. Interestingly, there are not many definitions of "understanding" in the business literature. It is as if everybody assumes that once we know something we understand it by default and hence can act accordingly. Knowledge can be defined as "information combined with experience, context, interpretation and reflection" (see Davenport *et al.* (1998)). The definition of knowledge offered by Senge (1999) as "the capacity for effective action" is essentially a definition of "understanding" (see McKelvey *et al.* (1999)). Another useful definition of understanding is "to perceive the meaning of" something (see *Webster* (1989)).

With these distinctions between data, information, knowledge and understanding in mind, we realize that successful BAs require more non-analytical work than analytical work – we must during all BA efforts constantly remind ourselves what is the purpose, the context, possible alternative interpretations, experience and so on. This may sound obvious, but it is a great problem of all BA, in my opinion. For example, two economists, Deirdre McCloskey and Stephen Ziliak show that this is a persistent problem in what they refer to as economic analyses. They studied to what degree papers in the highly respected journal *American Economic Review* failed to separate statistical significance from plausible explanations of economic reality (*The Economist*, 2004a). Their findings are depressing: first, in the 1980s 70 per cent of the papers failed to distinguish between economic and statistical significance, and second, in the 1990s more than 80 per cent failed.

At the other end of the scale – the specialized financial analyses – we unfortunately find similar problems, for example in the value at risk (VAR) concept. "Its foundations are shaky in several ways", as *The Economist* (2004b) points out:

- It assumes that market-returns conform with a particular pattern, despite the pattern in financial markets are nowhere as certain as those of, say, physics.
- The problem of "fat tails" in these statistical patterns, which means that the whole calculation of statistical significance may be flawed so that for example if you believe you take 1 per cent chance of a certain loss you are in reality and unknowingly taking maybe a 5 per cent chance of the same loss.

- The relationships between various financial assets in various geographical markets calculated as covariances are assumed to continue into the future, which means that if for example shares in Japan and Argentina have shown a strong covariance over the last five years, VAR calculations will assume that to continue. This implies that the probability for relatively rare, large scale financial disasters, such as stock-market cracks, countries running into major economic problems like Argentina, will always be underestimated. The 100-year storm will come much more frequently than once every 100 years.

The largest problem of VAR is of course the illusion of certainty and control it creates, combined with the fact that not everybody really understands what a confidence level of say 1 per cent means. Well, it means that in a year you will, on average, have two days where you will exceed your loss limits – is that acceptable?

From my experience, I do not think that the art of conducting BA, which is an intermediate affair to the more general economic analyses and the more specialized financial analyses, is in any better state. In fact, in many BAs, statistical analyses are completely ignored, thus we never come as far as discussing statistical significance versus economic or business reality – we simply omit the whole issue and consequently expose the BA, ourselves and others to the risk of self-deception or ignorance.

In short, we must internalize that the BA model is just a model of reality – not reality itself, and we must consequently constantly seek to understand reality more fully via non-analytical means and bring this understanding to bear on the modeling and its statistics so that both the modeling and the statistics become more meaningful. When this process is faulty, the result is the common phenomenon of paralysis by analysis, which is discussed next.

2. Avoiding paralysis by analysis

Paralysis by analysis refers to the common phenomenon where an analysis becomes so large or unclear that we essentially no longer understand what the analysis or its output is about, which results in no action or paralysis. Many seem to think that analyses inevitably lead to paralysis by analysis, but this is an oversimplification – it occurs only when the process outlined in Section 1 faults. However, it is true that the probability increases as the number of variables – particularly output variables – in the model increases. This is because BA models with a large number of variables typically tend to address many objectives which require that the data are mathematically manipulated in several dimensions. For example, in activity-based costing (ABC) analyses we typically compute the profitability both for products and customers, yet the total cost is the same in both dimensions (product and customers) which means that all costs must be presented in two dimensions without losing track. In strategic BA models the number of dimensions increases even more as we often are interested in market, customers, products, processes, overhead and so on. As the number of objectives a BA model is to satisfy increases, it becomes increasingly important to be diligent in understanding well the reality behind the numbers, as discussed at the end of Section 1, so that clear objectives can be made for the model. From this discussion it follows that paralysis by analysis should not occur in single-objective BA models – unless the objective is very poorly defined. If paralysis by analysis still occurs, it is simply attributable to poor craftsmanship.

In order to further reduce the probability of paralysis by analysis, there are systematic ways to work during the BA modeling. Here are five important ways, that are explained subsequently: make the model modular; use precise descriptions in the model and explaining text if necessary; make self-regulating models; use control summations after every step to make sure that nothing is lost; and do not make large complex equations but break it up into smaller segments.

Making models modular has many advantages, including:

- the model becomes easier to build because we can concentrate on certain aspects of the model;
- building correct interfaces between the various dimensions of the model becomes easier;
- it becomes easier to check the logic in the model because the flow of information becomes more apparent;
- making correct assumptions becomes easier as the system boundaries of the model becomes clearer;
- the model becomes easier for other to understand and work with; and
- running Monte Carlo simulations or similar becomes easier – particularly for very large models – because parts of the model can be run sequentially and hence not run into conflict with constraints such as computer memory.

Precise descriptions and assignment of units to all variables are crucial to secure a clear delineation of the internal logic of the model. Otherwise, we may assign one interpretation to a variable and then some months later – when the first interpretation is more or less forgotten – we unconsciously assign a slightly different interpretation, and over time this may produce logical inconsistencies and ultimately model failure.

Self-regulating models are models that can be used to improve the modeling itself. This has actually very little to do with the actual model, except that it must be built logically enough to secure rational interpretations, which implies that models should not be built like black-boxes or spaghetti codes. Given a rational model, the self-regulation is generated by employing statistical sensitivity analyses that allow us to identify the most important variables, which in turn makes it possible to check the logic of these variables. Beware that non-statistical sensitivity analyses, such as the What-if technique, can be deceptive as interactions between the variables are ignored and hence the wrong conclusions can be made. The most effective approach is to use statistical sensitivity analyses in BA by employing Monte Carlo methods. For a more thorough discussion on sensitivity analyses and Monte Carlo methods consult Emblemsvåg (2003).

Using control summations and breaking up large complex equations into smaller segments are both ways to ensure that the logic in the model works as intended and to make it easier to identify logical problems.

However, the single most important way to ensure that we do not fall prey to the paralysis by analysis syndrome is to actually understand what the results actually imply, because then we can identify logical flaws, which are ones that are most difficult to identify. To do this we must both have a clear picture of reality, and we must understand the results. Understanding the results is partly a matter of having a broad understanding of the subject matter, for example it is beneficial to understand some finance in order to interpret the financial results of a model, but it is also a matter of analytical techniques, which is discussed next.

3. The curse of central tendency measures

Too often BA is based on calculating various measures of central tendency such as averages, medians and modes, while available information to also compute measures of dispersion (such as variance) is often ignored. This is a fundamental problem with many repercussions as explained and briefly exemplified in this section.

First note that there are several ways of measuring the central tendency in a set of data (see Hines and Montgomery, 1990 for more information). The four most common are:

(1) Arithmetic average – is by far the most common measure of central tendency – so common that people typically refer to it as the average – and it is simply found by adding the data values together and divide by the number of data points.

(2) Weighted average – is quite similar to the arithmetic average, except that all data are weighted by another set of data.

(3) Median – is the middle of a data set arranged in increasing order of magnitude. This measure of central tendency has the great advantage over arithmetic average that it is not easily influenced by extreme data, but it ignores the value of the data points. For example, the following set of data points, 1, 1, 3, 10, 20 will produce a median of 3 but an arithmetic average of 7. In this case, the fact that the median ignores the values of the data points seems to produce a somewhat doubtful estimate of central tendency.

(4) Mode is the data that occurs most frequently in a data set. Consequently there can by several modes in a set of data. In the previously mentioned data set, the mode is 1 – an even more doubtful estimate of central tendency than the median in that particular case.

All these measures of central tendency, however, are nothing but alternative measures of central tendency with certain strengths and weaknesses, but ultimately they share a major limitation with respect to the discussion in this paper; they cannot measure dispersion. When it comes to measures of dispersion there are only two commonly-used measures (see Hines and Montgomery, 1990 for more information):

(1) Variance – is the sum of squared of deviations between each data point and the arithmetic mean. From the variance we can calculate the standard deviation as the square-root of the variance. Variance is the most important measure of dispersion and it comes in two shapes: sample variance and population variance. The difference between the two is that sample variance is based on a limited data set (the sample) whereas the population variance is the true variance within the entire population, from which the sample originated.

(2) Sample range – is the difference between the smallest and the largest data point in a data set. The problem with this measure is that all other information is ignored and hence this measure is completely dependent on the extreme data points in a data set.

To simplify the discussion in the remainder of this paper, I confine the discussion to the usage of average (arithmetic average) when discussing central tendency measures and variance when discussing measures of dispersion since the argumentation is the same for the other measures of central tendency. The sample range measure should not be used alone – only to complement variance.

Consider the sales data of a certain Product A over a year, as shown in Table I. Some business analysts might here be tempted to only compute the averages only; average sales of $1,383 per month, average volume of 121 units per month wherefrom an average unit price of $11,42 per unit can be computed.

There is nothing wrong with computing the averages, but it is not sufficient because we miss a lot of valuable information which can be very useful. In order to understand a data set properly, we must always use both a measure of central tendency and a measure of dispersion (see Hines and Montgomery, 1990). For example, a simple plot of the monthly sales will reveal an apparently cyclic behavior (see Figure 1), which should be further investigated, depending on the purpose of the BA.

More information can be further provided by studying the percentiles of the data in Table I. The corresponding percentiles are found in Table II, from which we can start

Month	Sales ($/month)	Volume (units sold/month)
January	2,000	180
February	1,500	125
March	1,200	105
April	1,300	110
May	800	60
June	500	45
July	500	43
August	1,000	80
September	1,500	120
October	2,000	170
November	1,800	170
December	2,500	250

Table I.
Sales data of product A in Year 2

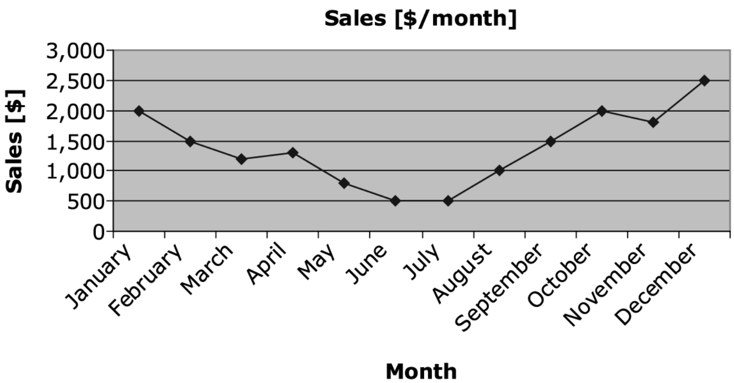

Figure 1.
Sales ($/month) of Product A

Percentiles	Sales ($/month)	Volume (units sold/month)	Average unit price ($/unit)
10%	530	46.5	10.64
20%	840	64.0	11.11
30%	1,060	87.5	11.21
40%	1,240	107.0	11.51
50%	1,400	115.0	11.70
60%	1,500	120.0	11.80
70%	1,710	155.0	12.30
80%	1,960	170.0	12.50
90%	2,000	179.0	12.50
100%	2,500	250.0	13.33

Table II. Percentiles for the data in Table I

identifying probabilities. For example, we see that the probability of the sales exceeding 2,000 units per month is only 10 per cent, and there is actually an almost 30 per cent probability that sales will fall short of 1,000 units per month. We also see that the unit price is varying by almost $3 per unit or about 25 per cent. In this fashion, we already have learned a lot more about Product A than the simple averages could ever teach us.

Of course, this is a very simple example – so simple that everybody should arguably have performed such simple analyses of variation as done here without being explicitly asked to. However, in many BA, such simple studies of the data are poorly performed or simply ignored. Just think about it; how many sales budgets actually include such discussions? How often are manufacturing costs per unit expressed by using an expected value and a band of possible costs, for example, Product A has a manufacturing cost of $10,50 ± $1,50? I believe such practices partly arise due to ignorance and partly due to misconceptions about well-known mathematical laws such as Law of Large Numbers in particular.

The Law of Large Numbers by Jacob Bernoulli (1654-1705) is often interpreted by practitioners as a few big numbers dominate the outcome – like in $100 + 1 = 101$; 100 is much larger than 1 and hence dominates. Equally erroneous are the views that it is a method for validating observations or that an increasing number of observations will increase the probability that what you see is what you will get (the Law of Averages). All it says is that "the average of a large number of throws[1] will be more likely than the average of a small number of throws to differ from the true average by less than some stated amount" (Bernstein, 1996). This implies that the relative frequency of an event of a certain probability tend to approach the true average of the event as the number of observations approaches infinity. Thus, the relative influence of each observed deviation from the true average approaches zero. In other words, the true interpretation of the Law of Large Numbers is almost the opposite of various, common misinterpretations. The Law of Large Numbers could therefore be renamed as the Law of Large Number of Observations to be clearer.

Now, let us look at a real-life case. We had performed an ABC analysis which directed attention towards the rebate system employed. The question we had to address was whether or not the rebate system works as intended, particularly in our small retailers, which was to discourage small orders and promote some economies of scale effects in ordering, manufacturing and transportation. One analyst approached the matter the standard way and computed the average sales per order, which clearly showed that there were significant differences between the various markets. This information was coupled with the ABC information and there was some agreement – small orders resulted in lower profitability.

Then, we took the same information but instead of aggregating it into market-wide averages we kept it more detailed, as illustrated by Table III. We defined the rebate intervals and summed up how many orders fitted into each rebate category. For example, for Austria we see that there is no rebate if you buy less than five units per order and 51 orders fell into this category. However, if you order between six and 11 units you get 9.0 per cent sales rebate and 22 orders fell into this category. The interpretation of the data in Table III can be done both graphically, by plotting rebates versus number of orders, or using the statistical concept of correlation. The idea is that for a rebate system to be deemed effective it should change the natural inclination to buy what you need to buy more than what you need for the time being, so that the number of small orders is reduced. Therefore, the higher the correlation between the number of orders and the rebates, the less effective is the rebate system. For Austria, we then realize that not only does the rebate system hardly impact ordering behavior but we also give away large rebates without influencing the behavior. In other words, the rebate system in Austria does not work at all as an influencer of ordering behavior. The system in the UK is somewhat better, as the small orders constitute a comparatively lower portion of the total number of orders than in Austria, but there is clearly room for improvement in both markets. In this way we see that using the data as detailed as possible allows us to gain much more information and insight from the data via the study of variance and correlation.

The lessons from these two examples are that we must always use the full amount of information data offers and particularly important is the estimation of variance, correlation, trend and probability. Averages are statistical illusions in the sense that they are supposedly what occur on average, but too often this average never occurs in reality. Often, we find that supposedly improbable events occur much more frequently than we believe.

Finally, I would just like to point out two basic cases that are described in standard textbooks on statistics, yet they are very often ignored in BAs. The first case is the importance of understanding the importance of variation when comparing the results of a model. All too often, the uncertainties surrounding the estimates or the risks are ignored and only the averages are considered. It is important to distinguish between risk and uncertainty. Roughly speaking, we may say that while uncertainty arises due

Country	Start	Stop	Rebates (%)	Number of orders	Correlation (%)
Austria	0	5	0.0	51	99.9
	6	11	−9.0	22	
	12	23	−13.0	12	
	24	199	−15.0	6	
	200	–	−17.0	0	
UK	0	5	0.0	124	79.2
	6	11	−2.0	218	
	12	25	−4.0	123	
	26	49	−6.0	5	
	50	–	−8.0	0	

Table III.
The correlation between number of orders and rebates for Austria and the UK

to complexity and hence lack of information, risk arises due to decision making and therefore concerns the future. Furthermore, risks may or may not appear with varying consequences and probabilities, while uncertainty exists all the time, albeit its magnitude may change. From this we understand that the notion used in finance that risk can be measured as volatility ignores this distinction between risk and uncertainty. For a thorough discussion on risk and uncertainty, see Emblemsvåg and Kjølstad (2002).

In Figure 2 we see that by focusing only on averages we might be tempted to conclude that the two alternatives are equally profitable. If the uncertainties and risks were included, however, we might have realized that the downside of alternative 1 is considerably larger than for alternative 2 – in fact, there is a considerable probability for loss. The upside, however, is also larger. Thus, we are now forced to consider the risks and uncertainties that are associated with each alternative. Forcing decision makers to face this reality is useful by itself.

The second case arises when the averages are different but the variation is the same as shown in Figure 3. The question here, of course, is whether or not the averages are truly different – in real life – or are they merely the result of our modeling? This is where the power of statistical methods come into play – using hypothesis testing we can decide whether or not the difference is real give a certain p value, which is the probability of rejecting a hypothesis when it in reality is true. Unfortunately however, even such methods have their weaknesses – they only work perfectly as long as the data are normally distributed and with a limited data set this is difficult to be certain about. Hence – as McCloskey and Ziliak point out – we must always make a judgment in the end, whether or not the analysis makes sense given the realities we face.

Too often, however, such judgments are not made and possibly even worse – people overreact to small differences in average values or deviations from goals, believing that the difference is real and substantial when it in reality could be entirely random.

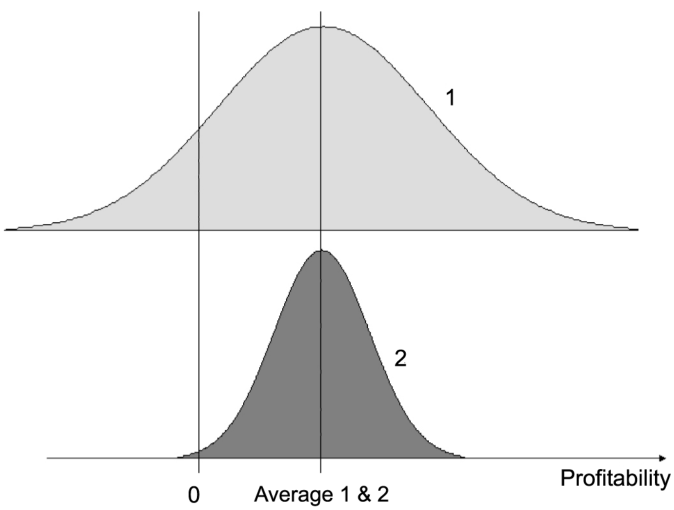

Figure 2.
Variation versus average

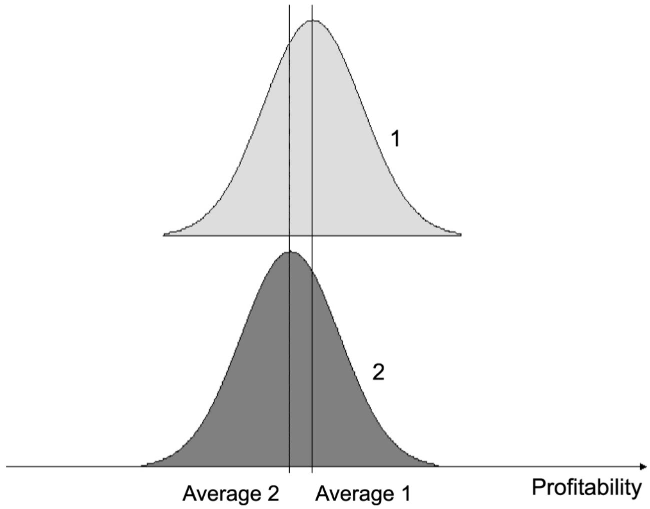

Figure 3.
Average 1 versus average 2

In reality, we find a combination of the two aforementioned cases, which further emphasizes the importance of making sound statistical and modeling judgments and reasoning thoroughly about their practical implications. When discrepancies between modeling and reality are found – a critical review must be performed. Never completely trust either a model or your perception of reality – each is flawed; the only question is which one is the least flawed and should be trusted more. Thus, excellent BA modeling is a learning process where our perception of reality is critically examined using a model and the same model is critically examined against our perceived reality and hence producing the best fit between perceived reality and model. The most important part of BA modeling is, therefore, not the answers it produces but rather the learning process it facilitates.

Such learning can be further enhanced by exploiting variance and turning it into a tool as explained next.

4. Turning variance into a potent tool

The most potent way of using variance is not through analyzing existing data but to run experiments. Here, we have an incredibly powerful statistical approach called Design of Experiments (DoE), in which analysis of variance (ANOVA) is a very important and integral part. In DoE, variance is actually added in order to exploit "... the nonlinear effects of the process parameters on the performance characteristics", (see Kacker (1996)). Similar can be said about Monte Carlo methods (see Emblemsvåg, 2003). In both cases – DoE and Monte Carlo methods – variance is thus introduced purposely and its impact on the performance characteristics measured statistically to derive as much information about the object of study.

While DoE can be extremely useful in BA, where the objective is, for example, to improve a process, it can be somewhat troublesome to apply if the process is hard to run experiments on. For example, running a DoE for the marked demand of a product is possible but highly unfeasible for many smaller corporations due to the sheer costs of running such experiments. In such cases where DoE may be too grand a task to undertake, Monte Carlo methods can provide much information, provided that a detailed and thorough mathematical model is available (see Emblemsvåg, 2003 for examples on the usage of Monte Carlo methods in large spreadsheet models).

Because both DoE and Monte Carlo methods are large subjects in themselves, they will not be explained here, as books cover these subjects well. For this paper it suffices to acknowledge that variance can be a tool in itself, but only when it is added intentionally and systematically to derive as much information as possible about the object the experiment is designed for. DoE is therefore a common tool in six sigma and other statistically-oriented quality management methodologies and it has produced impressive results in both quality and cost savings (see DeFeo, 2000; Chua, 2001). In my opinion, DoE is one of the most underestimated and ignored – yet, one of the most useful – management tools available, and that is a great loss for corporations. For example, in a survey where senior executives rated 25 "leading" management tools (see Rigby, 2001), DoE or related approaches like six sigma were not even on the list despite these tools often commanding a ROI of between 10:1 and 100:1 (see DeFeo, 2000). Thus, the single most important message to take away from this paper is that as business analysts we must pay much more attention to variance than what is commonly practiced.

5. Closure

I have tried to illustrate and argue the importance of following certain principles to avoid paralysis by analysis and of thinking in terms of variance when performing BAs, both to gain insight but also to improve the modeling itself. Data often lends itself to statistical analyses, and it is evidently foolish to not utilize the data to extract as much information as possible, but we must always critically review the model against perceived realities and vice versa. Furthermore, variance can be a tool in itself if properly used. Therefore, I believe it is well overdue for corporations to open up their eyes to variance as a powerful source of information and embrace approaches that exploit variance to the maximum and derive as much information and insight from the data as possible. When variance is exploited systematically and intentionally it becomes an asset and not a liability, as it is treated today.

Note

1. Bernoulli, like many of his contemporaries, was preoccupied by the study of chance in the context of gambling. Hence, they frequently talk about throwing dices.

References

Bernstein, P.L. (1996), *Against the Gods: The Remarkable Story of Risk*, John Wiley & Sons, New York, NY, p. 383.

Chua, R.C.H. (2001), "What you need to know about Six Sigma", *Productivity Digest*, December, pp. 37-44.

Davenport, T.H., Delong, D.W. and Beers, M.D. (1998), "Successful knowledge management projects", *Sloan Management Review*, Vol. 39 No. 2, pp. 43-57.

DeFeo, J.A. (2000), "An ROI story: the black belts of Six Sigma", *Training & Development*, July, pp. 25-7.

(The) Economist (2004a), "Signifying nothing?", Vol. 370 No. 8360, p. 63.

(The) Economist (2004b), "Living dangerously: a survey of risk", 24 January, p. 22.

Emblemsvåg, J. (2003), *Life-Cycle Costing: Using Activity-Based Costing and Monte Carlo Methods to Manage Future Costs and Risks*, John Wiley & Sons, Hoboken, NJ, p. 320.

Emblemsvåg, J. and Kjølstad, L.E. (2002), "Strategic risk analysis – a field version", *Management Decision*, Vol. 40 No. 9, pp. 842-52.

Groh, T. (2004), "Beyond ROI ... justifying a business intelligence initiative", *DM Review*, January.

Hines, W.W. and Montgomery, D.C. (1990), *Probability and Statistics in Engineering and Management Science*, John Wiley & Sons, New York, NY, p. 732.

Joia, L.A. (2000), "Measuring intangible corporate assets: linking business strategy with intellectual capital", *Journal of Intellectual Capital*, Vol. 1 No. 1, pp. 68-84.

Kacker, R.N. (1996), "Taguchi's quality philosophy", *Quality Progress*, December, pp. 21-9.

Kock, N. and McQueen, R. (1998), "Knowledge and information communication in organisations: an analysis of core, support and improvement processes", *Knowledge and Process Management*, Vol. 5 No. 1, pp. 29-40.

McKelvey, W., Mintzberg, H., Petzinger, T., Prusack, L., Senge, P., Shultz, R., Bar-Yam, Y. and Lebaron, D. (1999), "The gurus speak: complexity and organizations", *Emergence*, Vol. 1 No. 1, pp. 73-91.

Rigby, D.K. (2001), "Putting tools to the test: senior executives rate 25 top management tools", *Strategy & Leadership*, Vol. 29 No. 3, pp. 4-12.

Senge, P.M. (1999), *The Fifth Discipline: The Art & Practice of the Learning Organization*, Random House Business Books, London, p. 424.

Webster (1989), *Webster's Encyclopedic Unabridged Dictionary of the English Language*, Gramercy Books, New York, NY, p. 1854.