



Beyond Budgeting
Institute

Working Paper

BEYOND VARIANCE ANALYSIS

Part 1: WHY WE NEED A NEW WAY OF MEASURING PERFORMANCE

Steve Morlidge

March 2016

Acknowledgements

This paper is the preliminary first chapter in Steve Morlidge's next book provisionally titled 'Present Sense: A practical guide to the science of measuring performance and communicating it'.

We have agreed to 'serialise' the book in the form of Working Papers that will be added to our members' knowledge base and shared in our network. We would like to offer our appreciation to Steve Morlidge for allowing us to publish his book as it evolves in this manner. This is highly interesting and relevant for our members.

This first paper sets out the reasons why conventional approaches to performance measurement do not work, so building the case for alternative approaches that will be described in subsequent papers which are planned to be released in the coming months.

Steve has spent most of his professional career in designing and running performance management systems in Unilever. Steve co-authored 'Future Ready: Mastering business forecasting' (John Wiley, 2010) and has written many papers on business forecasting. Steve has a PhD in Management Cybernetics.

Steve is the former Chairman of the BBRT and remains a tremendous support of our network for which we are most grateful.

About Beyond Budgeting Institute and BBRT

The Beyond Budgeting Institute is at the heart of a movement that is searching for ways to build lean, adaptive and ethical enterprises that can sustain superior competitive performance. We promote a set of principles that lead to more dynamic processes and front-line accountability. Organizations that follow this approach transform their management model in line with these principles.

Our ideas are spread through the Beyond Budgeting Round Table (BBRT); a shared learning network of member organizations with a common interest in transforming their performance management models to enable sustained, superior performance. BBRT helps organisations learn from worldwide best practice studies and encourages them to share information, past successes and implementation experiences to move beyond command and control.

For more information about Beyond Budgeting and BBRT membership, please contact:

Beyond Budgeting Institute
One Kingdom Street
Paddington Central
London W2 6BD
United Kingdom
Tel: +44 20 3755 3692
Mail: info@bbrt.org
Web: www.bbrt.org

© Beyond Budgeting Association Ltd.

No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the author. BBRT member organisations may reproduce it for internal use within their organisations.

1. Introduction

This paper is the first in a series of four on the subject of variance analysis; a topic of particular interest to Beyond Budgeters since it is an integral part of the traditional management system we want to replace. Specifically, over the course of the series I will explain why and how it fails to provide a meaningful analysis of performance and what methodologies and approaches we should use instead. This first paper makes the case for going beyond variance analysis.

I hope that you will find the content not only thought provoking but also useful. It is not enough to argue that variance analysis does not work. We also have to provide an alternative, and for this to displace the traditional approach, it has to be simple and easy to implement.

These papers are a synopsis of chapters in a book I am in the process of writing. It is provisionally titled 'Present Sense: A practical guide to the science of measuring performance and communicating it'. It will be a companion volume to my earlier book 'Future Ready: Mastering business forecasting' and it will have the same aims.

I hope to provide practitioners like the readers of this paper with scientifically grounded practical methodologies and tools to facilitate the transition to an adaptive management model. At the same time, I want to reach out to those who know nothing about Beyond Budgeting (BB) or who are sceptics, and by explaining the deficiencies of traditional approaches and demonstrating that there is a better way open their minds to the possibilities.

One of the uncomfortable things about being a book author in particular is that while your work may be read by thousands of people, you rarely discover what they think about it. And before the book is published – when you have the opportunity to make improvements – the level of useful feedback you get is virtually zero. I would therefore very much appreciate receiving your comments, criticisms and suggestions and any contributions that you think might enhance the work. You can contact me at steve.morlidge@satoripartners.co.uk.

2. Variance analysis and Beyond Budgeting

Let us start by defining what we mean by variance analysis and find out from where it emerged. By 'variance analysis' I mean the systematic analysis of differences between an actual outcome for a period and the target, plan or budget for that period. A distinctive feature of variance analysis when applied to financial performance is that it often involves decomposing the difference into its component parts - volume, price, usage, and so on.

Prior to about 1900, variance analysis did not exist in a form that we would recognise. The first book on the topic appeared in 1918, and initially was applied exclusively to the analysis of costs using material standards developed by engineers for production management purposes. This was an era of rapid industrialisation when the first multidivisional companies were born, most notably in the motor industry. The complexity of these large diversified businesses required new ways to measure and manage business performance so the market was primed for the first book on Budgetary Control (1922) written by a young Chicago Business School professor called James O. McKinsey (who later went on to found a consulting firm that you might have heard of). This form of management involved setting detailed targets (budgets) for every financial component of a business exercising control through the analysis of variances from budgets. The implicit expectation behind this approach is that performance could and should be steered back to the predetermined plan.

This explains why we in the Beyond Budgeting community have a particular interest in variance analysis: It is a cornerstone of traditional budgeting. Indeed, in my view it is one of the reasons why budgeting is so hard to eradicate and why we need to find an alternative. If the only tool we have to analyse performance is a variance, then we have to have budgeting even if we do not believe in budgeting itself. Therefore, finding a better way to analyse performance removes one of the biggest obstacles to the adoption of BB.

The importance of performance measurement and the role of variance analysis in cementing the position of budgeting was recognised very early on in the life of BB, but my sense is that we have sometimes failed to give helpful practical advice to potential adopters. For example, the first BB book, written by Jeremy Hope and Robin Fraser (2003), recommends that we 'Base Controls on Effective Governance and on a Range of Relative Performance Indicators', which isn't very specific. Since then we have shared the experiences of many hundreds of early adopters who all have their own 'home grown' set of measures and metrics which is great but carries the risk of creating the perception that our advice consists of little more than 'you need to find out what works for you'. It is true that every business is different, and the fact that we do not seek to impose 'a one size fits all solution', is a strength and not a weakness of BB. But when we compare 'our' approaches with the clear cut, apparently definitive results produced by a traditional variance analysis, our approach can seem vague and unscientific.

There are, however, several other good reasons for abandoning variance analysis, even if you do not subscribe to the Beyond Budgeting view of the future.

3. Reason 1: It doesn't measure performance

The most important reason is a straightforward technical one. The simple arithmetic difference between a target and an actual is not a measure of performance of an organisation, and this applies to ANY such comparison, not just the variance analyses prepared by finance people.

One reason for this - that we are very familiar with, but the rest of the world chooses not to recognise - is the fact that it is impossible to set a target in advance that represents 'good performance'. The modern world is complex and dynamic making prediction impossible, and we know that we will never face the same set of circumstances twice. The economy changes, as do our customers and competitors, so we can do no more than hazard a guess at 'what good looks like'. The only way that we can assess performance in a truly meaningful way is to compare ourselves to peers that have faced the same set of conditions.

The second reason why a variance cannot capture performance is that any measure of reality will always contain noise: the impact of an unknowable number of random or irrelevant events, which distort and disguise 'real' performance (the signal). Therefore, even if we have a 'good' relative target we cannot assume that the difference between the actual and the target reflects performance because we do not know to what extent it is a product of chance. We know this intuitively - we instinctively discount what we think are 'small' differences - but if we do not have the ability to measure the level of noise, we have no scientific basis to distinguish between something that is safe to ignore and something that we should be acting upon.

However, the majority of the world does not use relative targets, so most of what passes as performance analysis is the result of comparing a guess with noise! This is particularly the case for accountants; the more detailed the variance analysis is, the more we rely on guesswork to set targets and the more our measures are infected by noise. The fact that 'performance information' produced in this way is meaningless is bad enough, but what makes it worse is that people might believe it and act on it!

The final reason why variances fail is that performance is the product of a pattern of behaviour over time, and so cannot be measured by a single data point. A single frame can show an image of a dancer but it takes a film to capture a performance.

Let me illustrate this with a thought experiment that involves analysing performance using numerical data. Imagine a weekend golfer - not a pro but one who sprays the ball about. Imagine that he (because it usually is) is playing an unknown (to you) golf course in unknown weather conditions. As you do not know the par for each hole, the only way you can make any judgement about his performance is from the measurements you receive - i.e. the number of shots he takes at each hole. How would you set about describing his performance? For

sure, you would not use the score on a single hole, because we don't know how long the hole is, how difficult it is, what the weather conditions are or whether the golfer suffered from bad or good luck – for example perhaps the ball hit a tree, which stopped a wayward drive going out of bounds. To understand the level of performance of our imaginary golfer we need to analyse the score from a number of holes looking for patterns. Only then (and with the help of a good comparator like other golfers scores) can we begin to answer the question, 'how good?' and 'is performance getting better or worse?'

In summary, to understand performance we need dynamic not static measures.

4. Reason 2: It is not in tune with current reality

Variance analysis is the product of an era when businesses had very little reliable data and very limited analytical power. In these circumstances, they needed to squeeze every drop of insight from the meagre supply of available data. This was only possible because that the world was relatively stable, meaning that targets were easier to set with confidence. Furthermore, the environment was less noisy, so measures were less distorted. Fast-forward a hundred years (yes it really has been that long!) and the world is very different.

Now our problem is not a shortage of data. Many managers feel like they are drowning in it and the supply is growing at an exponential rate. IBM estimates that we are generating 2.5 quintillion bytes of data each day, more than 90 per cent of which was created in the last two years (IBM, 2010). Much of this flood of data is 'out there' sitting on the Internet, but the situation is mirrored within businesses. It is difficult to measure this directly but the collapse in the cost of data storage gives us an idea of the likely scale of change. By 2012 a gigabyte of storage cost roughly 3,000,000 times less than it did in 1982 (Komorowski, 2014), and with the growth of Cloud storage over the last few years you can probably add another zero to this number. Of course having more data is not 'bad', it should be 'good'; the real problem is that we don't know what to do with it. Even if the analysis of variances were meaningful, we cannot cope with setting 3 million more targets and 3 million more pieces of information on performance. As a result of not having the ability to make sense of the data that they have collected without reference to arbitrary targets, most businesses simply ignore it.

We see the same trends in computing 'horse power'. Even thirty years ago, the majority of calculations took place manually. PC's did not exist. Today, your smartphone has much more computing power than the rocket that took Neil Armstrong to the moon. Yet, we are basically doing the same simple arithmetic operations as we did on manual calculators in the past: mainly adding, subtracting, dividing and occasionally multiplying. Technology has not only provided us with unimaginable riches in terms of data but it has also given us the ability to mine it for insights... but we do not know how to use the power at our disposal. We analyse data in the same way as before – the only difference is the speed at which we do it.

Am I suggesting that 'data analytics' will ride to our rescue and solve this problem for us? I am not. There are a number of reasons for this view, one of which is that, contrary to what many would like us to believe, technology cannot provide us with a silver bullet. The reality is that the more data we have and the greater the sophistication of our techniques, the higher is the chance of us coming up with something that is complete nonsense.

I cannot express this any better than Nate Silver did, in his book 'Signals and Noise'. Silver has become famous for his ability to analyse complex real world process, like elections, and make stunningly accurate predictions using mathematical techniques, so he is no luddite, but he is sceptical of the hype surrounding the topic:

'This is why our predictions may be more prone to failure in the era of Big Data. As there is an exponential increase in the amount of available information, there is likewise an exponential increase in the number of hypotheses to investigate. For instance, the U.S. government now publishes data on about 45,000 economic statistics. If you want to test for rela-

tionships between all combinations of two pairs of these statistics: Is there a causal relationship between the bank prime loan rate and the unemployment rate in Alabama? That gives you literally one billion hypotheses to test. But the number of meaningful relationships in the data - those that speak to causality rather than correlation and testify to how the world really works - is orders of magnitude smaller. Nor is it likely to be increasing at nearly so fast a rate as the information itself; there isn't any more truth in the world than there was before the Internet or the printing press. Most of the data is just noise, as most of the universe is filled with empty space.' (Silver, 2013,249)

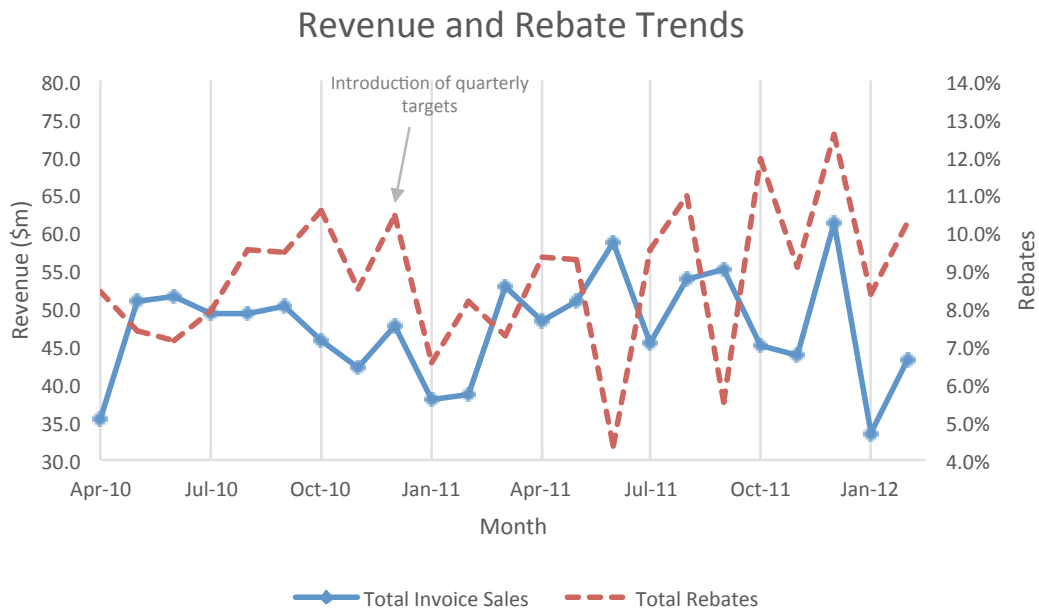
The situation is complicated enough just by virtue of the unimaginable number of mathematical combinations, but the real world is many orders more complex because its dynamics - the patterns of causation - cannot be understood through simple mathematical association. For example, a computer might notice a correlation between the sales of men's shorts and ice cream, but it cannot know whether the ice cream sales cause the sale of shorts (perhaps because the ice cream drips on legs rather than on fabric) or vice versa. It takes a human to spot that the correlation is caused by something else altogether – temperature. This is complicated further by the fact that, if we are trying to understand something like performance, we have to factor in time delays: something that we observe now might be the result of action taken one month, one quarter or even a year ago. Even worse, the key variable might be impossible to capture in the form of data at all – like 'we started using cute animals in our TV adverts'.

5. Reason 3: It drives dysfunctional behaviour

The final reason for calling time on variance analysis is that it drives the wrong behaviour. If a variance from target is perceived as 'bad', people will seek to avoid this outcome. If a fixed target is meaningless or unachievable, and if real data is rendered unpredictable by noise, it is often easier to hit the target by manipulating the system or the data than it is to change performance. In fact, it might be the ONLY way to 'hit the number'.

We are all familiar with extreme examples of this kind of behaviour from the news. Over the last decade Bernie Ebbers, Jeffrey Skilling and Bernie Madoff have been jailed for their role in large-scale frauds and the banking industry has nearly brought down the entire global economic system (without anyone being convicted!). In my view, however, the more pernicious examples of this are far more damaging precisely because they are accepted as 'normal'.

For example, the chart below shows the patterns of sales and sales discounts for a large multinational business.



After a period of steady but declining sales in quarter 2 and 3 of 2010, which investors were no doubt unhappy with, there was a push to set more aggressive quarterly targets – and hit them. This was achieved, not by increasing performance, but by manipulating it. We can see that sales start peaking at quarter end with the level of sales discounts following the same pattern. There are a number of things that appal me about this. The money wasted by paying customers to change the timing of their orders, and forcing the supply chain to accommodate unnatural patterns of demand, all done in the name of improving performance. The fact that once you go down this path, like a gambler doubling up to recoup losses, it is difficult to give it up and it often ends in tears.

But for me what is worse is that it was a blatant attempt to disguise what was really happening to performance and that everyone in the business casually accepted this exercise in deception as just being ‘part of the game’. Moreover, I see this happening everywhere and at every level where this approach to performance measurement is used.

6. The Way Forward

In summary, the conventional approach of measuring performance, which involves comparing an actual outcome for a period, is no longer fit for purpose. Targets are usually arbitrary and actuals are always infected by an unknown amount of random noise and so can never be a reliable guide of achievement. Further, the more granular the analysis, the less reliable it becomes. It also makes it difficult to uncover the underlying dynamics of performance.

Indeed, the requirement to ‘hit the numbers’ can distort patterns of behaviour in manner that is at best misleading, at worst damaging to the business. For us in the BB community it is particularly important that we are able to articulate problems with conventional approaches and provide credible alternatives to help us promote and embed the BB principles in our organisations.

But are there alternatives, and what are they?

There are three important qualities that alternatives to conventional variance analysis need:

1. They should present a dynamic perspective of performance in contrast to the static focus of traditional approaches based on the analysis of one period in isolation.
2. They should enable us to make sense of data that is infected with noise, to stop us reading meaning into randomness and help harvest the insights from rich data sources.
3. They should promote a more nuanced interpretation of performance and avoid the dysfunctional behaviour engendered by compliance to simple period-end fixed targets.

Over the next few months, I will share some approaches and techniques that address each of these in turn. These approaches aim to exploit the capabilities of modern technologies while respecting the powerful yet limited capabilities of our brains, which have not significantly changed in the last few thousand years. So, they have to be both sophisticated and intuitive - simple to use and simple to communicate.

The next working paper in this series is titled: 'From static to dynamic measures: Understanding patterns in performance'.

REFERENCES

HOPE, J.; FRASER, R. Beyond Budgeting: How Managers Can Break Free from the Performance Trap. Boston, MA: Harvard Business School Press, 2003.

IBM. What is Big Data? , 2010. www-01.ibm.com/software/data/bigdata, accessed 29/2/2016.

KOMOROWSKI, M. History of Storage Cost. 2014. <http://www.mkomo.com/cost-per-gigabyte> , accessed 29/2/2016.

MCKINSEY, J. O. Budgetary Control. New York: The Ronald Press Company, 1922.

SILVER, N. The Signal and the Noise: The Art and Science of Prediction. 2013.