chapter 1

Why Do We Need To Change The Way We Do Stuff?

I explain why traditional approaches – especially variance analysis – fail as a way of making sense and communicating performance and how technology has sometimes become part of the problem rather than the solution.

INTRODUCTION

It was 1996 – the year that the Spice Girls split up and Dolly the sheep was cloned. I was the Financial Controller for a large unit of a highly respected multinational company, and I had just made my own personal breakthrough – the final piece of the jigsaw was in place.

Performance reporting at the touch of a button; burned straight onto acetate. The wonders of modern technology!

Whatever next?

And I couldn't wait for my next board presentation where I could show off my new baby.

But when I came to unveil my new innovation to the senior team I was swamped by a tidal wave of indifference. Deflated by my disappointment, I cornered the Sales Director as he grabbed a dose of nicotine and quizzed him about it. I knew him well and liked him. He was charismatic and funny but blunt; a quality that he chose to share with me at this moment.

'What a load of old b*******' was the gist of his opinion on my efforts.

'OK, so volume in Tesco was up by half a million against budget in March and prices were down by £300k. So what?' He shrugged and walked off. He clearly didn't, as I had hoped, feel obliged here or in the board meeting to provide a blow by blow analysis of how this had come about or mount a defence of his sales team. He didn't even seem to think it was worth explaining his analysis of my presentation, probably because I had just demonstrated what an idiot I was, so it wasn't worth the effort of stating the obvious.

As I reflect back on this experience now, I wonder how is it possible to be so wrong on so many levels? How could an intelligent, experienced professional be so unconsciously incompetent?

Before I explain why I came to appreciate the extent of my ignorance and stupidity conclusion, I need to describe what I had done, and how I came to make the mistake. This might make some readers from a conventional finance background uncomfortable because I'm sure that most of you have done – or are doing – the same kind of thing.

Let's start with some background.

In 1992 when I was given the job of Financial Controller in a new business, I found my department in the throes of a meltdown. Thirty of the original complement of 120 people had just been 'released', along with my predecessor, and the people left were unsurprisingly not the most motivated bunch I have ever met. The financial accounts, which ran on an old IBM System 360 computer (a design dating back to 1964), were months out of date and the bank account hadn't been reconciled for the last two quarters. The asset accountant, who had 30 years' experience under his belt and was charged with the task of trying to find the reason for a multimillion-pound hole in his accounts (using an old manual calculator he was fond of), had just revealed that he was number dyslexic. This meant he couldn't tell the difference between the numbers 6 and 9, which went some way to explaining the hole. And my 'colleague' who headed up the management accounts part of my empire had already decided that I was his mortal enemy and was doing everything possible to help me fail in my new job.

You don't have to know anything about accountancy to appreciate that this was not a good place to be.

Four years down the line we had turned everything around and I was justifiably proud of what my team had achieved, and over confident as a result. We had implemented an integrated ERP system (we were SAP's first customer for its new client–server software in the UK). We had introduced fast close processes and integrated the management and financial accounts to produce a single 'version of the truth', complete with forecasts within a day after the bookmonth end. This meant that the senior team had a beautiful and colourful report pack on their desk within five days of the end of the period. And if they had any kind of query, we could drill down from the report data to individual invoices and tell them who had made the entry, when and the name of their pet dog (I made up the last bit – SAP doesn't have this as a field in their master records). As far as modern, 'best practice' finance processes were concerned, we had therefore ticked every box.

So, what went wrong when I presented my slick set of acetates to the board?

My first big mistake was to assume that the difference between a target and an actual for a period provided useful information about performance. It didn't and it usually doesn't. This might sound an outrageous thing to say given that most business performance reports work on this premise, and not just in finance.

Let me explain what I mean by imagining what my Sales Director might have said, had he taken the trouble.

> Look, we both know that the budget number was just a 'made up number' produced over a year ago. We had no idea, and we never will, what a 'good' number for sales would be in Tesco in March because we don't know how fast the market is going to grow, and how much share they are likely to have of it. Nor do we have any idea what our competitors will do or whether we will want to defend our share of Tesco's market, if they attack us, or increase our take if we spot a vulnerability. And you are kidding yourself if you think there was any science behind the target for the year. I negotiated it with the CEO – we just haggled and agreed on a number. And remember both our bonuses are tied to the number we agree to – so we were actually negotiating my pay. He wanted a high number and I wanted a lower one, so we met in the middle. And then your accountants took the number we agreed on and spread it across the months and played around with it so that the total business had a profit number for the first quarter that the guys in head office would wear.

What the imaginary sales director is telling me here is that whenever there is uncertainty or volatility in the business environment any target that is set in advance can never be more than a 'good guess' or a statement of aspiration. It is not a realistic estimate of an achievable level of performance.

Also, because it doesn't take account of the actual business context, it might drive the wrong behaviour. To use a military analogy, a commander of an army might want to advance on all fronts but if the enemy is throwing all of their forces at your position, you might be doing very well just to hold it.

Finally, target setting is often highly politicized, which makes them an even more unreliable guide to performance. In such circumstances, targets are really just a tool used by participants to try to get people to do what they want (pay me more money, pull out all the stops) rather than a foundation for the rational analysis of performance.

My imaginary conversation didn't end there.

> Also, Steve, you have to remember that we had a big promotion at the beginning of the month, so sales we would expect to make in March were booked in February because we had to stock up Tesco stores in advance. But the cost of securing the deal was booked to March because that is when we ran the promotion, so your conclusions about March performance was false.
>
> And, even in a normal month, we know that a lorry full of product breaking down or being turned away from the warehouse on a Saturday night at the end of the month can make a significant difference to the numbers – and that's before I make allowance for somebody's Aunt Nellie being on holiday.
>
> So that's why I don't take any notice of your numbers, and why I would be nuts to use them to take any decisions.

To translate: in any single period, there are a host of things that can distort 'reality' as recorded in the books. Some of these are knowable and could allow for them if we are prepared to put in the effort and make some educated guesses. But some of the things that distort the numbers we can never know about. Some will affect the timing of when we record a value; others are just random events that distort the picture. In other words, any measurements we make, even if they are indisputable statements of fact, will contain an unknowable level of 'noise'. And the more detail we go into the more this noise will obscure what is really going on. To paraphrase, we will just see trees – not the wood. When we assess performance, we need to be able to discount the noise so we can focus our attention on the signal. Only then can we begin to extract meaning and decide the best course of action.

In summary, my mistake was that I had compared a single data point (containing an unknown amount of random noise) with a target (which was a politically motivated guess made 12 months before) and assumed that the difference was meaningful. When you put it in those terms it doesn't sound too clever, does it?

But with the benefit of hindsight, I can see that I made even more errors.

For example, the 'information' that I had presented to the board looked something like Table 1.1.

Table 1.1:
An example of the sort of table I presented to the Board

A typical traditional performance report, in this case reporting on sales compared to budget.

Sales (£k)

	Q3				**Variance v Budget**			
Product Groups	**Budget**	**April Estimate**	**July Estimate**	**Actual**	**April Estimate**	**July Estimate**	**Actual**	**%**
Product A1	£ 184.9	£ 185.9	£ 171.7	£ 176.1	£ 1.0	-£ 13.2	-£ 8.8	-4.8%
Product A2	£ 214.3	£ 226.7	£ 214.1	£ 214.7	£ 12.4	-£ 0.2	£ 0.4	0.2%
Product A3	£ 123.3	£ 117.7	£ 115.0	£ 110.4	-£ 5.6	-£ 8.3	-£ 12.9	-10.5%
Product Group A	£ 522.4	£ 530.2	£ 500.7	£ 501.1	£ 7.8	-£ 21.7	-£ 21.3	-4.1%
Product B1	£ 156.8	£ 161.2	£ 156.2	£ 151.9	£ 4.4	-£ 0.6	-£ 4.9	-3.1%
Product B2	£ 162.2	£ 175.5	£ 163.8	£ 163.8	£ 13.3	£ 1.6	£ 1.6	1.0%
Product Group B	£ 319.0	£ 336.7	£ 320.0	£ 315.8	£ 17.7	£ 1.0	-£ 3.2	-1.0%
Product C1	£ 195.7	£ 200.5	£ 201.0	£ 193.1	£ 4.8	£ 5.3	-£ 2.6	-1.3%
Product CX	£ -	£ -	-£ 13.5	£ -	£ -	-£ 13.5	£ -	
Product Group C	£ 195.7	£ 200.5	£ 187.5	£ 193.1	£ 4.8	-£ 8.2	-£ 2.6	-1.3%
Total Product Groups	**£ 1,037.1**	**£ 1,067.4**	**£ 1,008.2**	**£ 1,010.0**	**£ 30.3**	**-£ 28.9**	**-£ 27.1**	**-2.6%**

Tables of numbers are still the data analysts' default mode of presentation in business. It is not difficult to see why.

You can cram a lot of 'information' into a small space, so can use them to answer many of your audience's questions. And they are easy to generate. You don't have to put much thought into it.

But, on the downside, it can be hard work trying to make sense of the data in tables.

Because of the shortcomings of targets, it might be difficult to answer basic questions like, 'Is this good performance?' And it is also almost impossible to work out whether things are getting better or getting worse.

If the numbers for previous periods are not displayed on the page (which they often aren't because it takes up too much space) the reader has to somehow reconstruct them in their head based on what they were able to remember from previous periods. Even if history was presented in the table, it is very difficult to determine trends from a limited subset of raw numbers; it requires a lot of cognitive effort and it is easy for two people to arrive at different 'answers'. To my mind it is like throwing a handful of jigsaw pieces on a table and saying to your audience 'there you go, now go work out what the picture is'.

Finally, it is also very difficult for the audience of decision makers to answer the question, 'Has something significant happened in the period that I should pay attention to?' The 'something significant' might be a 'one-off' problem with a process that requires remedial action or something that may provide early warning of a change in performance, good or bad.

If the recipients of reports like this can't make sense of them, or – perhaps worse – come to different interpretations of the facts, then the next steps in the decision-making process will be difficult, fraught and misguided.

The point I am making here is not that 'tables are bad' but that information professionals need to pay more attention to what information their audience really needs and how best to present it to ensure that they collectively get the right message, quickly and easily.

Once your audience has a clear and consistent understanding of the reality of performance, they can then do their job – bringing together their disparate knowledge, experience and skills to understand the reasons why things happened and to work out what to do for the best. And, because we human beings have very advanced visual pattern-recognition skills, we need a good understanding of how to effectively present information in a graphical form. Lists and tables of numbers just won't work.

This story illustrates the two themes in this book.

First, in order to do a better job, information professionals like the 'me' of 20 years ago need to improve our methods of communicating information – in this case reporting on performance to teams of decision makers.

This is the 'Present' of the book's title. It is a kind of 'art', and analysts are not recruited for their artistic skills. But the good news is that while there isn't a 'right' answer there are many practices that are clearly wrong, and by learning to avoid these even the most artistically challenged amongst us can quickly improve our work. I can vouch for that because I have successfully applied these ideas in my own business, and I'm no artist.

Second, we need to do a better job of extracting meaning from data. We have to do more than merely process data. We must develop an ability to understand trends and spot other meaningful patterns in performance by learning how to separate signals from noise. This accounts for the 'Sense' in the title of the book. The good news here is that, while this requires developing a bigger repertoire of analytical skills, most of the basic requirements can be fulfilled using a handful of very simple tools that every competent information analyst can easily and quickly master. They are tried and tested approaches that have been used in other walks of business life for many years that I wish the 'me' of 20 years ago had known about.

My objective in writing this book is to package this knowledge so that every practising professional can apply it in their own business as soon as they put the book down. This is why I have chosen to describe and demonstrate these approaches using the tool that every reader will have access to – Microsoft Excel – rather than reference any one of the Business Intelligence (BI) or visualization tools out there on the market. To the frustration of every software vendor in this field, Excel is still the tool of choice for analysing and reporting on performance, but also using common tools enables you to quickly try it for themselves. You don't have to 'take my word for it'. And by experimenting with these ideas yourself you can produce something you can immediately benefit from and become a better-informed purchaser of software when and if you decide to exploit these ideas and insights in a more sustainable way.

You have now heard my confessional, but don't worry about me – or beat yourself up if you find yourself in a similar position. It's not our fault. The sad fact is that none of us have ever been taught how to do things any other way, which is quite remarkable considering how far and how fast IT has advanced over the last few decades.

In my case I qualified as a management accountant in the early 1980s, and then, as now, performance reporting was regarded as the central pillar of the profession. Also, then as now, if you picked up the examination syllabus and looked for performance reporting the key technique that you are expected to have mastered is 'variance analysis'.

To better understand why this is, and why there is a need to change, let us get a handle on exactly what has changed in the 'information space' in the 30 or so years since I started work.

SOME HISTORY

Back in 1980 the computer hadn't made it out of the payroll office of most businesses. In my first 'real job' I generally added up numbers by hand, but if the list was too long or I didn't have the time I took them to one of the 'girls' (usually a formidable lady of advancing years – my mother-in-law used to be one!) operating a comptometer – a device like a manual calculating machine that you operated using all your fingers simultaneously. After a year or so I got my first electronic calculator but only after I submitted a capital proposal and had it duly recorded as an asset in the balance sheet!

Where computerized data did exist, there was not very much of it, and it was usually highly unreliable. In those days all computer systems were bespoke; customized, full of bugs and interfaces with other systems that were notoriously prone to failure and error.

In the days where data was in short supply and unreliable and where the world moved at a slower pace, analysing variances made some kind of sense. It allowed us to squeeze a lot of information out of a small amount of data and provided a mechanism to highlight problems in the data. Back then we didn't worry about the conceptual flaws in the target-setting process, our inability to understand the dynamics of performance and separate signals from noise, because *these were not the problems we then had*. And in the days when any information published within an organization was produced on a manual typewriter there was no alternative to tables of numbers.

Fast forward 30 years and the data problems we used to suffer have largely disappeared. Integrated ERP systems now provide users with copious amounts of reliable 'internal' data about our organizations, stored in huge data warehouses, and in more recent times this has been supplemented by a torrent of external data about the marketplace, our customers and competitors.

It's difficult to quantify the scale of this change, so I will use the cost of data storage as a proxy measure for the quantity of data available to businesses. In 1981,it cost roughly $700,000 to purchase a gigabyte of hard disk storage capacity. By 2014 a gigabyte cost about 3 cents – roughly a 25,000,000-fold plus reduction![1] And this number doubles every 14 months, following the same trajectory as Moore's Law, which describes the rate of change in computing processing power.[2]

Cloud technology and the mediation of so many transactions through the web has supercharged the trend towards ever greater data storage, so my guess is that we could probably add at least an extra zero to this number. IBM estimates that we are generating 2.5 quintillion bytes of data each day, more than 90% of the data that has *ever* existed was created in the last two years. You can provide your own adjective to describe the scale of this change, so long as it means 'something really big'.

The year 1981 was also when IBM sold its first ever PC. It cost the equivalent of $15,000 but ran something like 500 times slower than today's equivalent machine.

But it is not only the scale and quality of the data about the world and our technological ability to process it that are different. The world that we are seeking to understand is itself changing at an increasing pace – largely as the result of the same technological processes.

The military use the acronym VUCA to describe the volatile, uncertain, complex and ambiguous world we now inhabit. But, while the military has evolved

1. To get a sense of the scale of this change and its implications for the way that we process data, imagine that in 1981 we had one orange and that we used a manual juice extractor to squeeze all the information we possibly could get out of this single fruit. But by 2015 we have 25 million oranges, which would fit in about 200 large articulated lorries. If all these lorries were lined up it would create a traffic jam 4 kilometres long. So, the good news is that we have a lot more oranges than we used to have. The bad news is that we are still using the 'one orange at a time' manual juice extractor.

2. Source: http://www.mkomo.com/cost-per-gigabyte

to meet these new challenges and to exploit the huge volumes of intelligence and cheap computing power available to them, performance management professionals have barely changed their approach at all. Variance analysis is over 100 years old but in 2016 it is still promoted by professional accounting bodies as the 'gold standard' approach to performance analysis.

The only significant change in performance management practice over the last 30 years is the increased attention given to non-financial measures, or KPIs. This shift followed the 1987 publication of *Relevance Lost: The Rise and Fall of Management Accounting* by Bob Kaplan and Tom Johnson, in which they criticized traditional finance performance measures as 'too little, too late and too aggregated'. After *Relevance Lost* Kaplan teamed up with David Norton to write *The Balanced Scorecard*, which described a technique that kicked off what has been called 'the performance measurement revolution'.

The increased focus on, and use of, non-financial information that the BSC brought about was undoubtedly right and valuable. The battle to introduce different perspectives into traditional financially dominated performance reporting has now largely been won. But it has done nothing to address the other problem that was beginning to emerge at about the time Kaplan and Norton's book hit the shelves. If anything, it has made matters worse. As I mentioned earlier, when I started regularly speaking at business conferences the complaint I heard most frequently was 'we don't have enough data'. But we are no longer data impoverished – in fact, we are now overloaded with data, and adding new metrics makes this problem more acute. So, I now hear 'we have too much data'. For the first time in our history mankind is facing problems of data abundance and evolution has not equipped us to deal with it.

THE BALANCED SCORECARD

The Balanced Scorecard (BSC) is the best-known example of a performance management framework that attempts to integrate financial and non-financial metrics. It has four perspectives that contain different attributes of performance that are taken to be causally linked. So, financial performance is the result of certain aspects of the customer experience that are in turn determined by the performance of important (internal) business processes. Finally, according to this framework, business processes are made more effective and efficient though a process of learning.

The example of a BSC for a regional airline shown in Figure 1.1 illustrates how one is constructed. First, the critical attributes in each perspective are identified and the causal links between them plotted on what Kaplan and Norton call a strategy map. Then key measures for each of these attributes are defined and targets set. Improvement is achieved through initiatives, which, if successful, will be reflected in the level of financial performance.

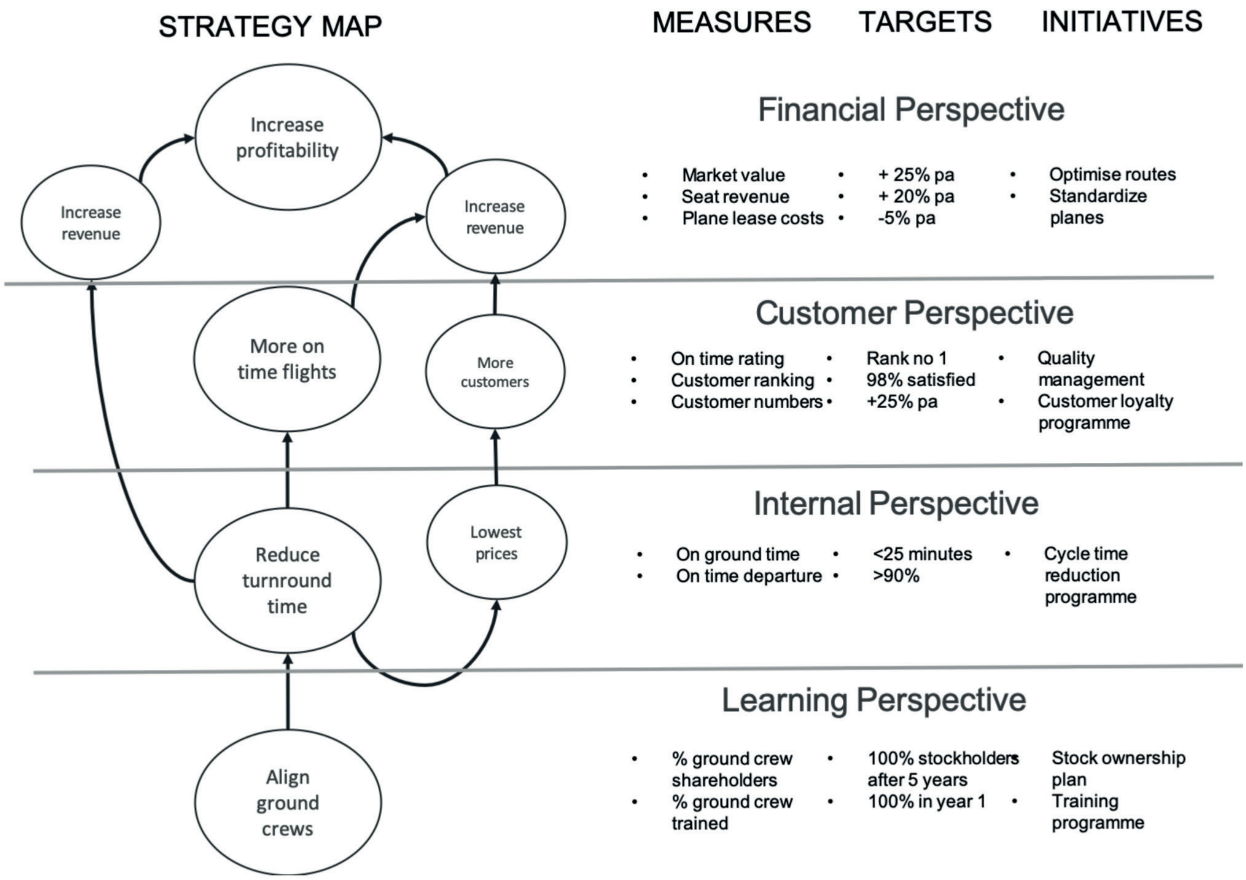

Figure 1.1:
A Balanced Scorecard

An example of what a BSC could look like for an airline company.

I see some parallels with the issue we face with food in the developed world.

In the distant past our ancestors didn't have access to enough calories so our bodies evolved to gorge on energy intensive foodstuffs like sugars and fats whenever they became available, which was infrequently. But things are different now. In many countries these 'unhealthy' foodstuffs are no longer in short supply but still our bodies crave more, because we can't turn off genetically programmed mechanisms. As a result, the developed world has an obesity crisis.

I think we now have the same problem with data – we think we have too much because we don't know what to do with it. Our organizations have therefore become bloated and slow moving and suffer from a range of weight related diseases – but the craving for more doesn't go away. We somehow feel that by consuming just a little bit more our hunger will be assuaged – but it never is.

Is the problem really that we have too much data? Or is it that the methods we use to process it are too crude to exploit these newfound riches? Do we need to cut back on our intake or should we upgrade our metabolic processes so that we can consume more of what we 'eat'?

I'm sure that a lot of data is captured 'just because we can' with no clear idea of how it can be used, but my belief is that we should focus more on the latter. Our real problem is not that we have too much data – it is our impoverished ability to make sense of it.

The BSC helped usher in an era when the range and nature of things that we measured expanded enormously, but the issue of how to analyse the resulting data has never been addressed. It was taken for granted that the way to analyse non-financial numbers was to compare them to a target at the end of a financial period. Where did the target come from? Nobody said: which left unchallenged the questionable assumption that any difference between the measure and the target, or any change in the measure, is meaningful and therefore grounds upon which action should be taken.

It is a similar story when it comes to communicating performance information. While everyone now seems to agree that graphics and visualization are 'a good thing' that we should do more of, there is little structured guidance about how to use graphical techniques to present and communicate performance information in an effective way.

So, while the amount of reliable data that businesses have at their disposal has expanded exponentially, as has the computational power potentially available to analyse and communicate it, the techniques that we use evolved in an era of data scarcity and have not changed at all.

The other thing that hasn't changed is the bandwidth of the human brain. We are no cleverer than we were 100 or 1,000 years ago. If anything, the challenge facing communicators has become bigger due to the fact the capacity available for performance information has shrunk because of the increase in the demands on our attention from other parts of our densely connected digitized world. And the intolerance for anything that is difficult to digest has increased.

Unsurprisingly the call now is for 'simplification'.

The 's' word is often used to support moves to focus attention on a narrower range of metrics – the 'vital few' – or to use high-level aggregates or averages. Because the human brain is finite and the demands on it are potentially infinite some form of selection is inevitable. But, if this is done unscientifically, we risk throwing away or systematically ignoring information that has been carefully and expensively collected. To paraphrase Stafford Beer, the great systems thinker, ignorance is the ultimate form of simplification.

The urge to simplify is the driver of probably the only 'innovation' in data analysis or presentation that has made it into mainstream use in the last 30 years. RAG charts, sometimes called traffic lights, use red, amber and green colour coded icons or numbers to signify whether something is good, bad or indifferent. In principle the idea is great, but if there is no scientific rationale behind the classification of values (which in my experience there rarely is) all we do is hard code our ignorance and make it more visually attractive. Attractive, that is, to everyone except that 10% of the male population who are red–green colour blind. And any member of the audience coming from parts of the world where red signifies something 'good' rather than 'bad' is likely to be totally confused.

There is no doubt that in order to extract sense from our superabundant data we need to find some way to filter out what is irrelevant or distracting so that we can focus our limited attention on that small proportion of our data sets that contain actionable insights. We then need a way of communicating this quickly and effectively in a manner that preserves this meaning. The problem is that the tools we traditionally use – especially in Finance – are incapable of performing this role.

I started this chapter with a story about a turning point in my career when I first realized that what I had been doing was completely inadequate. If you already 'get it' or are impatient to get to the meat of the book you might want to skip the next section because I am going to drill down into the problem some more. But the assumption that performance can only be understood by comparison with targets is so deeply embedded that I expect many readers will need more persuasion before they are prepared to leave behind the comforting assumptions of a working lifetime.

Simple 'actual to target' pairwise comparisons are ubiquitous in organizational life, but I suspect that accountants will be more reluctant to acknowledge their limitations. The edifice of financial performance management is built on what we call variance analysis – which is an 'actual against target' comparison on drugs.

Variance analysis takes the difference between target and actual and breaks it down into its constituent parts thereby – so the theory goes – exposing the reason for the deviation. Except that it doesn't. It just creates the illusion of insight.

Understand how variance analysis works and you understand the reason why we finance guys struggle to extract and communicate information in a way that other members of the organizational community find meaningful and useful.

VARIANCE ANALYSIS: RIGOROUS BUT WRONG

It is easy to forget that the commonplace 'furniture' of working life, barely noticeable because we are so familiar with it, hasn't always been there. Variance analysis is one such piece of social technology. And it wasn't 'discovered', it was invented to help people deal with a particular set of problems at a certain time. Only when we appreciate the specific problems that it was designed to fix can we begin to recognize its shortcomings when faced with a different set of problems of the sort that we now have.

Prior to about 1900 variance analysis didn't exist in its current form. The first book on the technique was published in 1918 and initially it was used to analyse costs using targets based on the material standards developed by engineers for production control purposes. Prior to the 20th century most businesses were run by owners rather than professional managers or by engineers focussed on manufacturing operations. However, the 1920s was an era of rapid industrialization, when the first multidivisional companies were born, initially in the motor industry – the 'high-tech' sector of the time. The complexity of these large diversified businesses and the separation of ownership from day-today management meant that new ways of measuring and managing performance were needed.

The market was therefore primed for the idea of exercising control through budget-based analysis and a young Chicago Business School professor called James O McKinsey duly obliged by publishing the first book on the topic in 1922. When he went on to found the world's first consulting company budgetary control became the world's original consulting 'product'.

McKinsey and his peers advocated managing by setting detailed targets (budgets) for every financial component of a business and subsequently controlling performance by analysing variances from budgets. The expectation was that performance could be steered back to the predetermined plan in the same way that a manufacturing plant could be directed to fulfil its production quota.

Given the rudimentary state of business record keeping and computation technology at the time this approach was a considerable advance on the alternative – an amalgam of financial accounts and production records. Although the process of creating a budget is often time consuming and tedious the process of analysing performance by referring to budgets – once they have been set – is straightforward, so it is easy to see why the idea caught on. So simple and obvious in fact that it can blind us to its shortcomings.

I will use a very simple sales example, in Table 1.2, to illustrate these shortcomings.

Table 1.2:
Simple Sales Variances

A simple variance analysis example. But how meaningful is it?

Revenue	**January**	**February**	**March**
Actual	25	40	45
Budget	30	35	48
Variance	**-5**	**5**	**-3**
%	**-20%**	**13%**	**-7%**

As you can see, in Table 1.2 sales revenue is lower in January than in the other months in the first quarter, but the fact that the budget is lower for the month might suggest that this is unavoidable – perhaps the result of a post-Christmas dip. Revenue in March is marginally higher than in February, but the budget was set even higher. So, as a result, the variance against budget in March was negative. That suggests that performance is better in January and February than in March.

But is that right? How sure can we be of this? Are these variances instead the result of a flawed process of target setting or perhaps just a manifestation of random fluctuations in revenue around a stable long-term trend? The uncomfortable truth is that with a simple pairwise comparison of a single period value and a target set using an unknown process we simply can't answer any of the basic 'big' questions we have about performance with any confidence.

- Is this good performance or bad?
- Is it getting better or worse?
- Has anything significant happened that requires action to be taken?

But accountants don't stop here. They need more detail, and they like tidy arithmetical processes.

If you pick up any textbook on budgetary control, you will probably find somewhere in the section on performance analysis a picture like the one in Figure 1.2.

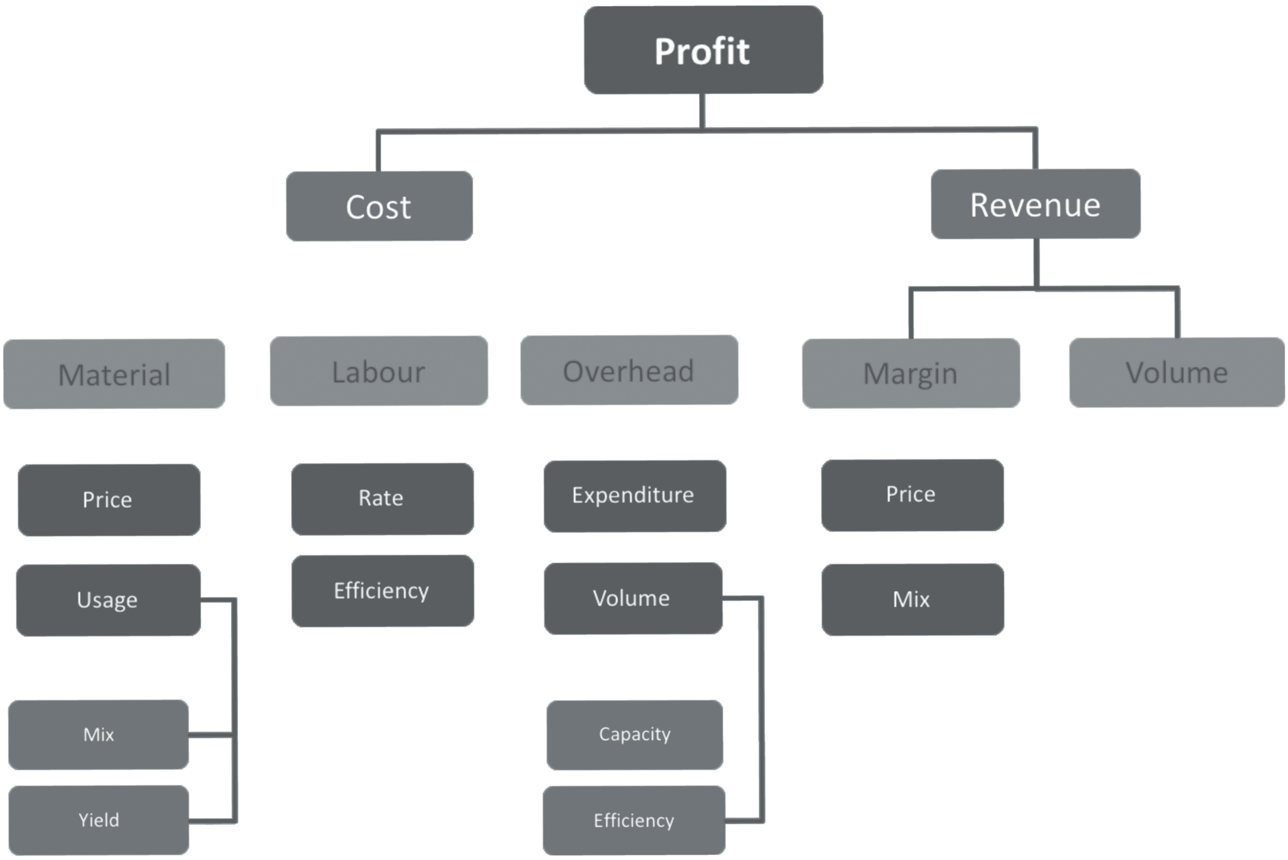

Figure 1.2:
Variance tree

This demonstrates how a high-level profit variance can be broken down into the various contributory factors to provide a comprehensive breakdown of the reasons for the deviation from target or plan – assuming that every element of the plan is reasonable and every measure is free from noise, which is unlikely ever to be the case.

You can understand the appeal of the variance tree. It is logical and structured. There appears to be no escape from the forensic glare of the controllers in finance. But its apparent rigour is also its weakness.

One of the reasons why accountants like budgets is that they are arithmetically coherent – they add up neatly. Because budgets are built from the bottom up, we can drill down from a high-level variance to lower-level ones to uncover the 'root cause' – or so they would like to believe. In Table 1.3, for example, I have broken down the revenue variance into two constituent parts: price and volume.

Revenue	January	February	March
Actual	25	40	45
Budget	30	35	48
Variance	**-5**	**5**	**-3**
%	**-20%**	**13%**	**-7%**
Breakdown			
Actual			
Price/unit	5	5	5
Volume (units)	5	8	9
Budget			
Price/unit	5	5	6
Volume (units)	6	7	8
Variances			
Price	**0**	**0**	**-9**
Volume	**-5**	**5**	**6**
Variance %			
Price	**0%**	**0%**	**-20%**
Volume	**-20%**	**13%**	**13%**

Table 1.3:
Revenue variance broken down into price and volume

A worked example showing how 'comprehensive' variance analyses can be misleading.

Now our variance analysis in Table 1.3 is telling us that the volume performance in March was better than the other two months and that the negative revenue variance in March is the result of a price rise that we assumed for March that didn't take place.

Does this now make the performance in March good – a continuation of a favourable upward trend in sales? Or is it bad because we failed to get the price increase through? Or was the price increase a guess made six months ago when the budget was set or a 'plug' figure somebody used to make the total budget come back to the numbers that head office wanted to see?

Again, we just don't know. More detail generates more questions not more answers. And the more granular the data gets the more noise swamps any signals that the data may contain. We become overwhelmed by data about individual events and so lose sight of the patterns that convey meaning and insight, in the same way that we lose a sense of the picture on a TV screen when we focus on individual pixels. Too many trees but no wood. All content, but no context.

To be clear, I am not arguing that variance analysis can never work – just that it is wrong to use it to try to understand the performance of a business as a whole rather than, say, to use it to control product costs.[3] Why?

3. I don't mean to imply that targets work perfectly in a production environment either! Indeed, one of the approaches to performance management that avoids the use of arbitrary targets originated in manufacturing, as we will discover later in the book.

1. To produce a meaningful variance analysis, you need meaningful targets to start with. The manager of a manufacturing operation can do this because she can refer to a product specification that defines exactly what materials should be and in what quantities. But the manager of a sales force can only speculate about the behaviour of his customers or competitors, which makes for unreliable targets.
2. It also assumes that every piece of data is a pristine manifestation of the truth, unsullied by chance or random factors. Production processes are explicitly designed to minimize variation. Raw materials have to conform to specifications and production processes are tightly defined. Other business processes are unavoidably more open to the outside world that cannot be controlled in this way, so randomness and noise have a much bigger impact on the data. As a result, no single data point represents 'the truth'.
3. Finally, conventional variance analysis assumes history is largely irrelevant. It is understood that all the information you need to interpret performance and to determine the appropriate course of action is contained in a single variance number. But life is not like that – it is a continuum. What you measure today is a product of actions taken in the past and an exclusive focus on a single period obscures the play of cause and effect and the direction of travel of the business – good or bad. A single piece of data is a snapshot in time, like a frame in a film. To understand performance – the behaviour of a system in time – we need to stitch the snapshots together. Variances do not give us the information we need to make decisions because every image in our performance movie is distorted by noise and the myriad of false and arbitrary assumptions made in setting targets and is presented to us one frame at a time.

The bottom line is that while variance analysis might be simple and provide us with neatly packaged answers, we should treat the results with a great deal of caution. Instead of simplifying the task of analysing performance, variance analysis makes it more complex, and more reliant on personal interpretation.

And, the more granular the level of analysis the bigger the problems become. More detailed analysis requires that more targets be set, which makes the process more bureaucratic and the targets more arbitrary. Greater detail also means that noise has a bigger impact on the data. So, instead of providing clarity, variance analysis creates irrelevant and misleading information. And by supressing the time dimension it makes it more difficult to see data in its historical context – the 'bigger picture'.

Finally, the obsession with targets and variances makes the whole process of performance reporting laborious and time consuming, but it also can make the business less responsive. As we have seen with the March price increase in the example in Table 1.3, targets that might have made sense when they were set quickly become out of date but the interconnected nature of the budgeting system makes changing targets difficult and disorientating because single changes cascade through the whole network.

In conclusion, while simple pairwise comparisons of actuals and targets to analyse performance met an important need in the formative years of professional management when data was scarce and communication difficult, it leaves us poorly equipped to tackle the challenges and opportunities faced by 21st-century managers. And while full-blown variance analysis of the kind shown here is a less common feature of financial performance management than it used to be, the target-driven performance management culture is in rude health right throughout organizations, despite its manifest weaknesses.

VARIANCES ARE NOT ALL BAD

My criticisms of variance analysis are directed at its use to analyse performance. This shouldn't be taken to mean that I believe all forms of comparison are invalid. Quite the contrary.

A number in isolation signifies nothing. It has to be compared to something else to mean something – the issue is what should it be compared to and what conclusions can be drawn from this? Indeed, Part 2 of this book will go on to advocate methods based on comparing numbers from the ***same data series*** to:

- Expose patterns of behaviour
- Help spot changes in the level and nature of performance
- Quantify the level of noise and so by exception detect signals buried in the data

Also, in my previous book (***Future Ready***) I strongly advocated systematically comparing actuals with forecasts, not to judge performance but as a means of testing and improving the models and assumptions on which the forecasts were based. Variances between actual and forecast are a reflection of the reliability of the forecast ***not*** the quality of performance.

Finally, I have devoted a lot of space to decrying the way in which targets are set and how they are compared to actual data, but there are ways to do both of these things that avoid most of the problems associated with simple variance analysis, as we will discover later.

Variance analyses have gripped the corporate imagination because they are simple to calculate and to understand. They are also a seductive tool for senior executives since they offer the prospect of being able to direct performance without having to get involved in the management of the business: you just set targets once a year and administer systems of rewards and punishments to encourage compliance. You can see why that might be popular!

But perhaps another reason why variance analysis took hold in the 1920s was that the results could easily be communicated using the technology available at that time – paper and typewriters. Tables of numbers are easy to produce and compress a lot of information into a relatively small space whereas charts had to be drawn manually and take up a lot of real estate on the page.

Tables are an efficient way of cramming pieces of numerical data, like variances, onto a single sheet of paper, and all other things being equal the less space that you use the better. But does that make them an effective device for communicating information about performance?

Let me answer this question using another example. Table 1.4 gives a summary of the simulated performance of a simple one-product company XYZ Ltd that I created for this book, using variances shown in a conventional table format.

	Quarter 4				Year to Date				2017	
	Actual	Budget	Variance	v Last Year	Actual	Budget	Variance	v Last Year	Budget	v Last Year
Turnover	245	262	-16	-4%	944	970	-26	5%	989	5%
COS	131	139	8	3%	506	527	20	-4%	539	-7%
GP	114	123	-9	-4%	438	443	-6	5%	450	3%
Gross Margin %	47%	47%	0%	0%	46%	46%	1%	0%	45%	-1%
A&P	36	37	1	2%	141	141	-0	-2%	140	1%
Overheads	39	40	1	-5%	139	148	9	-4%	136	2%
Profit	38	46	-7	-14%	157	154	3	11%	174	11%

Table 1.4:
Performance of XYZ Ltd

Note: COS = Cost of Sales, GP = Gross Profit and A&P = Advertising and Promotional expenditure.

How easy is it to make sense of the analysis of performance in this table. Put yourself in the position of the recipient of this report and use it to analyse the performance of XYZ Ltd. Try to answer these questions:

1. Is this good or bad performance?
2. Is it getting better or worse?
3. Has anything significant just happened?
4. How credible is the budget for next year?

How did you find it?

If you are used to analysing numerical information you will probably find it easy to come up with some answers, but you might be less sure in them than you would have been before you read my criticisms of variance analysis. And I'm confident that if you share this table with someone you will find that they come to slightly different conclusions than your own.

The other thing that I'm sure that you will have noticed is that trying to make sense of all this is hard work. You will have had to really concentrate in order to answer the questions I posed, and you might have noticed your eyes jumping around the table comparing one number with another as you tried to build up a picture of what was going on. If you weren't aware of this take another look at the table and retrace the path that your eyes took.

Any lack of confidence you might have in your conclusions and the sense of effort you experienced is not because the task I set you is inherently difficult. My 'toy' company is hugely less complex than anything that you deal with in real life. These uncomfortable sensations are a sure sign that I chose a very poor way to communicate information. It simply isn't a good fit with the way that our brains work.

The human brain has evolved over millions of years to efficiently assimilate information from our natural environment. This is why our visual perception is so much better developed than any other of our senses and why it is particularly good at spotting patterns and movement – provided information is presented in a way that appeals to our eyes.

In contrast, the symbolic systems we use in the West for communicating numbers emerged only 500 years ago when Leonardo Fibonacci introduced Arabic numbering systems into Europe.[4] We are born with highly-developed visual circuitry, but we have to laboriously train our brain to use numbers, which is why so much of our formal education is devoted to it, and – even then – some 'well-educated' people fail to acquire more than the most basic level of numeracy. But even if you are highly numerically literate you will find it much more effortful to assimilate the numbers in the tables in this chapter than if the same data was presented visually. And while two people may come to different conclusions about the meaning of a set of numbers there is much less chance of them perceiving the same shape in a different way, for the same reason: it is 'easy' and 'more natural'.

4. In fact, the Arabs imported the idea from India. It should really be called the Hindu numbering system.

And yet, despite our unsatisfactory personal experiences with tables and the ubiquity of computing power at our fingertips, performance reporting in business – particularly that produced by finance people – is still hugely reliant on tabular presentations and on decks of paper.

In summary, thanks to technological advances our capability to collect data, to analyse it and to communicate information to an audience has increased enormously over the last few decades. But our chronic addiction to approaches created to solve the problems faced by the first professional managers nearly a century ago severely limits our ability to exploit this potential. This is wasteful and it imprisons our minds. Our world is colourful, rich and full of life and ambiguity, but the 'pictures' of it that we create in our head are no better than bad caricatures.

The problem is clear. The challenge is to work out what to do differently.

THE BOTTOM LINE

To summarize, there are four main problems associated with traditional approaches that we need to overcome if we are to make the process of understanding and communicating performance fit for the modern world.

Too static

Performance is continuous but conventional approaches present us with a series of snapshots separated in time. As a result, trend information is blurred, and patterns of cause and effect are hard to detect.

Failure to distinguish between signals and noise

Data carries signals but is also unavoidably infected with noise. Conventional approaches do not enable us to distinguish between the two. Because we do not recognize the existence of noise, we falsely assume that *any difference between two numbers is meaningful*. And by acting on corrupted information we often amplify the noise, hence making it even more difficult to distinguish fact from fiction.

Confuse variance with performance

Comparing 'point' targets to actual outcomes is an unreliable guide to real performance because of the existence of noise and because targets are usually set without any sense of context many months in advance, often by an arbitrary or politically-driven process.

Do not take account of the brain's processing limitations and strengths

Our capacity to collect data is unlimited but the bandwidth of our brains is fixed and subject to ever-increasing demands. Conventional approaches to analysing and communicating performance, using numbers alone, are difficult for brains to process. This is inefficient and increases the risk of confusion and misinterpretation.

WHY TECHNOLOGY WON'T SAVE US

However, you might be thinking, what about the things that we hear so much about these days when we go to conferences or read business blogs? The media is full of stories about Big Data, data analytics, data visualization, dashboards and more recently AI. Surely, they will have solved these problems?

It goes without saying that better software tools are clearly *good things*. But it would be wrong to assume that there is a technological silver bullet out there that will solve all our problems for the same reason that buying a Stradivarius is unlikely to improve your fiddle playing. Tools carry potentiality; they don't deliver any results on their own. It takes hours of practice to play a violin and as far as data analysis is concerned most of us are still in kindergarten. Also, overplaying the technological card blinds us to the beautiful music we can make right now on our desktop machines, if we only knew how. But there are other more fundamental reasons why I'm sceptical about some of the technological hype out there.

Big Data is a new word that describes a phenomenon that is not new. But it has come to assume a much more prominent and important place in our everyday lives for good reason. We had lots of data back in the 1990s but what we are faced with now is a whole new ball game:

1. Volume. There is a lot more data than there ever has been, and it is growing at an exponential rate. For example, IBM estimates that we are generating 2.5 quintillion bytes of data each day, more than 90% of which was created in the last two years.

2. Velocity. The data is available much more quickly particularly if it is sourced from the Internet or smart devices like our phones or the sensors in our car.
3. Variety. It is available in many forms. Traditionally data was highly structured – classified and organized – but today a lot of data is unstructured, particularly if it takes the form of text (e.g. tweets), sounds or images.

BIG DATA ISN'T NEW – THE FUTURE HAS ALREADY HAPPENED

If you think that I am making too much of this Big Data stuff – that it is either not going to happen or that somehow technology will sort this out for us when it does – or that I am exaggerating when I say that the risk is that we will simply ignore this, let me tell you a story.

I can't remember exactly when, perhaps around the end of the 1990s, the retailers that my old company sold to started sharing their EPOS (electronic point of sale) data with us. Think of it – every day (or even more frequently) we had records of how much we had sold and in which store.

Think of it.

How useful would it be to be able to track the sales of our products on a daily basis? We could spot emerging trends almost instantly and work out how well our interventions (promotions/advertising/product innovations) were working. With the retailers' support we could carry out trials in limited areas to find out what worked and what didn't before pouring huge sums of money in. In theory, we could dispense with an enormous amount of expensive market research, which was necessarily less reliable because it was based on consumer intentions not their actual behaviour. And we could check that our customer did what they promised to do when we handed over cash to pay for prominent displays in their stores.

But what actually happened?

Every day the files landed with a big electronic thud on the company's servers and just sat there, neglected and unloved.

Why? Because, although the potential benefits were obvious, no one had any idea how to extract meaning from such a huge volume of fast-changing, noisy and messy data. So we just ignored it.

The future has already arrived and is waiting for us to hop on the bus. Get on or get left behind.

Just having data is of course of little value, and this is where data analytics come in. To convert this data into information we need to structure the data and then look for patterns in it. In theory these patterns represent information that we can do something with.

Unfortunately, it is not quite that simple.

The first problem we face is our friend noise. The more granular and unstructured data is the more noise it contains. And as anyone who has ever looked up at a cloud and seen a rabbit (or Elvis) knows, it is possible for us to detect patterns in noise and, as it turns out, computers are prone to this as well. This was recognized a long time ago, by Johnny von Neumann the brilliant mathematician. 'With four parameters I can fit an elephant,' he said, 'and with five I can make him wiggle his trunk.' What he meant was that the more data we have and the greater the sophistication of our techniques the better the chance of coming up with something that is complete nonsense.

I can't express this any better than Nate Silver did, in his book *The Signal and the Noise: The Art and Science of Prediction*. Silver is famous for his ability to analyse complex real-world process, like elections, and make stunningly accurate predictions based on the application of mathematical technique, so he is no luddite.

> "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 relationships 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**." (my emphasis)

Silver refers to the complexity that is an inevitable consequence of scale – the unimaginable number of mathematical combinations. But, even if it were not complex, it is easy to demonstrate why the real world cannot be understood through simple mathematical association alone. 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 bare legs rather than on fabric) or vice versa. And it takes a human being to spot that both are caused by something else altogether, which might not appear in the data set at all – temperature.[5] The difficulty of spotting causal patterns in data is complicated further when we have to factor in the time dimension as well. Something that we observe now might be the result of actions taken one month, one quarter or even a year ago, and the key event might even not have been captured as 'data' at all – like 'that's when we started using cute animals in our TV adverts'.

5. Arguably we only know for sure that a relationship is causal rather than purely a correlation through action – when we do something and get the response we expect. Analysis alone simply provides us with a plausible hypothesis to test.

I could go on, but you should have got the message by now. This technological stuff might be great for letting you know that 'people who order this item also bought…' but, because this 'insight' is the product of a relatively simple correlation, I wouldn't use it to do anything complicated like choosing what meal to cook for my in-laws this evening.

Likewise, there has been an explosion of interest in data visualization – the use of graphics to help us understand complex phenomena. This is undoubtedly a good thing since it exploits the visual pattern-seeking power of our brains and the processing capabilities of computers. But there is a downside to this potential. Fancy software is not cheap and not only is it complex to use but achieving good results requires a degree of design skill that it is unreasonable to expect from the average analyst who probably can't tell the difference between a Rembrandt and a regression analysis. Ultimately, only brains can understand how brains work and it takes a skilled brain to work out how best to design something that generates the right response in another individual's brain.

Dashboards go some way to solving this problem for us. They are specifically designed to help us make sense of performance data and come with a pre-packaged set of design templates and graphical tools. But they have a couple of downsides.

Firstly, their ability to convey insights is limited by the ability to extract meaning from data. So, if you do not know how to separate noise from signal you could be doing the equivalent of saying 'look at the rabbit' rather than 'it's a cloud, just ignore it'. Communicating meaningless noise in a compelling fashion is not what we should be aiming for.

There is a second, more profound reason that I think can easily be overlooked in our understandable excitement with what this kind of technology can do.

Performance reporting is embedded in a social process we call decision-making. Computerized dashboards are, however, designed for individual use. They are tools for personal productivity and enlightenment. This is not a bad thing, but because everyone will use them differently, they do not help create the shared collective consciousness necessary for effective organizational action.

This was brought home to me forcibly recently when I got a request from a user of the forecast performance reporting dashboard that my own software company sells. 'You know these charts you have on the top right-hand corner of the home page?' he asked me. 'Is there any easy way to copy a bunch of these for different accounts into PowerPoint?' My first reaction was unprintable, along the lines of 'That's not the point you dummy! We built this thing so that you could do everything on screen: zoom in and out, up and down, slice and dice – anything you want. You can't do that with PowerPoint you…*****!!!'

On reflection, however, I realized that I was the dummy. Our client was reporting on forecast performance to his company's leadership team, made up of people from other functions, such as sales, all of who had an input into the forecasting process. Quite properly, he wanted to use a specific subset of the charts for our tool to influence the behaviour of his colleagues in a very particular way. He had a specific message that he wanted to communicate so it was important that the target audience saw the same charts at the same time along with the same contextual commentary that he – as an expert – was able to supply.

The point I had failed to realize is that dashboards are designed to help an *individual* make sense of things for herself, but that they are unsuited for the task of reporting, communicating a distilled and crafted message to a *group* of decision makers. Reporting is a social process not a personal one, so we need a tool that can broadcast information to an audience not help individuals to generate it for themselves.

The implication of this insight is that we should not think of information professionals as being like plumbers who design and maintain data pipework (although this is important). They have responsibility for content as well. They are not authors of fiction or public relations (PR) professionals spinning a story for their clients. Their – your – role is closer to that of a serious journalist, responsible for presenting a clear, concise and balanced view of the world, without disguising its complexity and ambiguity.

In this role technology is your friend but it's not a silver bullet that will make all your problems disappear. You need to learn how it can best serve you by desk learning and careful thought. But there is no substitute for the experience that comes from experimentation with simple tools like Excel that you already know how to use. While it is unlikely to be the final destination, experimenting with spreadsheets will make us better-informed consumers when it comes to buying and using more sophisticated tools.

THE SOLUTIONS AND HOW TECHNOLOGY *CAN* HELP US

So far, I've focussed on problems. From now on it's about solutions.

I have tried to ensure that the solutions I recommend are conceptually and scientifically robust but also practical and easy to apply. And rather than using rhetoric or spin stories based on artfully edited case studies to win you round I will show you how to do a better job and encourage you to try things out for yourself.

The solutions I will demonstrate mirror the problems I have defined.

Dynamic – to expose trends

Performance is a pattern of behaviour that cannot be captured by focussing on individual data points. We need to measure and analyse performance in a way that exposes trends and trajectories.

Filtered – to deal with noise

We must to be able to separate the wheat (the signal) from the chaff (the noise) to focus our limited attention, knowledge and intellect on the right things and avoid being distracted by – or worse reacting to – randomness. Simple arithmetic can't so this for us. We need to view data through a probabilistic lens and use statistical filters to help us extract insights from big noisy data sets.

Reframe the meaning of performance – to reflect its complex multi dimension

We need nuanced targets or comparators that are less arbitrary than the targets produced by traditional processes like annual budgeting and better reflect the actual business context and performance potential. They have to help us track performance *over* time, rather than being anchored on points *in* time, and help us filter out the influence of noise.

Communicate more effectively – to exploit the capability of our brains

It is imperative for us to develop a better understanding of how to communicate meaning effectively to an audience of decision makers with differing knowledge, experiences and capabilities. This needs to be based on an understanding of how the brain processes information and a grasp of good design principles that exploit its capabilities.

I have argued that we should not look to technology as a silver bullet that will solve all our problems for us. But it is equally clear to me that we will not be successful in achieving any of these goals without the help of technology. However, we first need to learn how to best exploit it.

All of the techniques and approaches I describe in the following pages can be deployed using paper and personal calculators if you want to adopt a 'back to nature' approach to performance measurement. This would be perverse, so I expect and encourage you to experiment using desktop productivity software such as spreadsheets. Excel will, however, only get you so far. To implement methods that I recommend at scale and at speed in real life you will most likely need specialized software – but having experimented using simpler, more forgiving technology you will have developed a good grasp of what will work best in your organization and what you want from a specialized software solution.

I have also explained that I do not believe that dashboards are the solution to the problems of communicating meaning – primarily because they are designed to support personal understanding and enquiry not the social process of decision-making. Enormous progress has, however, been made in the last few years in understanding and codifying good practice in dashboard design and graphical communication, led by pioneers like Edward Tufte and Stephen Few. And these principles can be applied directly to the design of performance reports even if they continue to be produced in Excel or PowerPoint and distributed on paper. Again, I encourage you to practice with these ideas and experience for yourself the impact they can make. And what you learn can then be applied to customizing the dashboard software you have already got or are minded to buy.

That's enough of the preamble to this book – let's start ambling.

Learning	Reason
Why we shouldn't confuse data with truth.	All data contains noise, as well as a signal. It requires more than one data point and a scientifically robust inference process to enable us to detect the difference between the two.
How targets can be unhelpful as a guide to performance.	Fixed targets assume that you can know in advance what represents 'good performance', which is often difficult in a fast-moving world. Also, the process of target setting is often political and, particularly at a granular level, somewhat arbitrary.
Why traditional 'actual to target' comparisons are an unreliable guide to performance.	Comparing a single data point infected with an unknown level of noise with a target of unknown validity is more likely to confuse and misdirect actions than enlighten or promote wise interventions.
Why tables of numbers are not effective at communicating performance information.	The human brain struggles to assimilate information presented in a numerical form and tables present data in a way that requires a lot of effort to isolate relevant information.
How traditional approaches fail to guide intervention by decision makers.	Because numbers either represent single data points or are highly summarized it is difficult to distinguish between what is 'normal' to be ignored and what is abnormal, requiring an intervention.
Why these long-standing weaknesses have now become a major problem for organizations.	The volume of data, the pace of change and the shortened attention span of decision makers makes the failure of traditional methods to assess and communicate performance information effectively a major problem.
Why attempts to simplify using tools like RAG 'traffic light' charts are dangerous.	Attempts to classify performance without any scientific rational for making distinctions is unhelpful at best and positively misleading at worst.
Why pairwise comparisons between actuals and forecasts are helpful.	Deviations between actuals and forecasts should not be used to judge performance but they are essential to ensure that forecasts are a reliable guide to the future.
How Big Data makes things worse not better.	The ratio of signals to noise declines as the amount of data increases. So, while the potential for greater insight exists it is more difficult to extract meaning.
Why sophisticated mathematical techniques will not provide the solution.	Mathematical techniques rely on correlation, so at best they can do no more than spot patterns that may or may not be the result of causal relationships.