Beyond Budgeting Institute

Working Paper

BEYOND VARIANCE ANALYSIS

Part 4: GOOD OR BAD?
TARGETS, PERFORMANCE MEASUREMENT AND PERFORMANCE MANAGEMENT

Steve Morlidge

November 2016

Acknowledgements

This paper is the preliminary fourth and last 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 'serialised' the book in the form of Working Papers that have been published, 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 has evolved in this manner. This is highly interesting and relevant for our members.

The Beyond Budgeting community has a particular interest in variance analysis: It is a cornerstone of traditional budgeting. Indeed, 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 a significant obstacle to the adoption of Beyond Budgeting. Over this series of papers Steve has made the case for using alternative analytical approaches to measuring and making sense of performance and demonstrated how they address the weaknesses of the traditional approach.

Steve Morlidge 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. We help organizations learn from worldwide best practice studies and encourage them to share information and 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

Introduction

This is the final paper in a series of four on the subject of variance analysis.

When I talk about Beyond Budgeting and our disdain for traditional approaches to targets and their use in assessing performance I am often asked 'how can we know how we are doing if we don't have targets?' Overcoming this objection is therefore key if we are to succeed in convincing people that there is a viable alternative to traditional budgeting and all of the dysfunctional practices and behaviour associated with it.

In the second paper in this series I argued that traditional approaches to analysing performance by comparing a single data point to a fix target results in both producers and consumers of management information lose sight of business trends which are a critical dimension of performance and which provide the context for detailed analysis. In the last paper we explored the impact of noise on our ability to make sense of performance, particularly in the context of the need to exploit the large data sets that businesses now collect. I showed how simple statistical approaches can help us filter out noise to expose the relatively small number of data points worthy of investigation and avoid the common trap of over interpretation – seeing patterns in data that do not exist.

The extensive use of trend analysis and the need to explicitly incorporate a probabilistic component to our assessment of performance exposes a gap in our toolkit that I aim to fill in the first part of this paper in which I describe a method of tracking performance against a moving comparator (such as a target) whilst allowing for noise.

However, having the ability to track a moving target in a probabilistic manner only solves one half of problem with traditional variance analysis. The second part of this paper picks up the other more fundamental part of the problem – the arbitrary nature of most targets. In particular I want to answer the question: 'do we need targets at all?' Traditionalists cannot conceive of management without targets but the followers of W. Edwards Deming and other quality gurus argue that all forms of arbitrary quotas are damaging and unnecessary. Is there a middle way or do we have to make a simple binary choice between targets or no targets? And when we refer to something as a 'target', what do we actually mean?

But before we tackle the tricky and contentious issue of targets and target setting let us explore an approach to tracking performance that addresses the technical flaws of conventional variance analysis.

How to track performance in a dynamic and noisy environment

In my second paper we explored the power of using trend-based approaches to measuring and analysing performance. This dynamic view of performance – when it is presented graphically - that is intuitive and easy to interpret enables us to answer the critical question 'is performance getting better or worse'? Although this clearly helps us grasp whether performance is heading in the right 'direction' it doesn't help us determine whether we are going 'fast enough'. Revenue or margins might be growing, or costs declining but is the rate at which they are changing excellent, satisfactory or disappointing? In short, by tracking trends we know whether performance is broadly consistent with the purpose of the organisation or not, but we do not know how well we are fulfilling it.

To answer this question we need a comparator, which for the sake of argument we will call a moving target, and a means of assessing the gap between it and the actual performance. Moreover, as reality unfolds, we might need to change our target, so our method has to be flexible enough to accommodate this. We have also seen how fixed targets can promote dysfunctional 'peaking' patterns driven by the need to 'hit the number' at period end. Finding some way to unambiguously measure performance in a dynamic way, against a moving target that dispenses with period end peaking would clearly be very useful.

In the third paper we also explored the impact of noise on the measurement process and why we have to take account of multiple data points, not just one, and express our judgement in probabilistic terms. The solution I

presented for helping to filter out the impact of noise - control charts - uses very stringent standards for statistical significance and assumes that the 'signal' we are interested in is stable over the short term. This works well where we have a well-defined process that is not very volatile or if we have large amounts of noisy data that we do not have time to analyse in detail that we need to quickly 'scan' for abnormal events. The issue with using control charts to track critical business variables like revenue, costs and profit that are heavily influenced by environmental factors or events is that they are inherently unstable. Also control charts use the criterion of 3 standard deviations to detect an abnormal event, which might require 8 or more data points, and what business can afford to wait 8 months to respond to a change in a key variable like total revenue? So there are circumstances where the precision that control charts provide may not be appropriate.

Finally, BB implementers may need a target-tracking methodology to reassure sceptical colleagues that abandoning variance analysis doesn't mean losing all sense of control, in terms that they can understand – performance against target.

In summary, managing a complex and dynamic system that is open to its environment means that we need some point of reference, but it is equally clear that we need a way of comparing it to actual performance that is probabilistic.

So we need a statistically based process of measuring performance against a dynamic target or benchmark. Ideally this should be simple to administer and as simple and intuitive as possible. Fortunately there is a tool that meets this brief: it is called a 'tracking signal'.

Tracking signal

The clue to what a tracking signal does is in the title; it tracks the difference, over time, between actual outcomes and a comparator and raises an alarm when the difference goes beyond defined limits.

It does this by taking account of the average variance over time rather than a single data point and uses statistical criteria rather than judgement or opinion to signal significant deviations between the actual and the comparator. One attraction of this approach for Beyond Budgeters is that the comparator (i.e. the target) doesn't need to be fixed; it can vary over time, thereby helping us break out of the straightjacket of annual target setting. The approach I will now describe has been in use for half a century so we can be confident that it is robust. Better still it is an application of a simple method that we have already encountered in the context of trend measurement – the smoothed average.

Before we dive into the math let me illustrate the idea behind the tracking signal with a simple example.

Take this series of differences between an actual and a target (where the sign represents whether the actual is above or below the target).

+4, -4, +4, -4

The average net difference between the target and actual is the sum of these is zero and the average absolute difference between the two (i.e. ignoring the sign) is 4. The ratio between the net and the absolute difference gives us a 'tracking score', which in this case is zero (0/4). A tracking score of zero means that the actual is (on average) bang in line with the target and that the variances can be assumed to be random and so safely ignored.

But what happens when the actual is not tracking the target? If, for example, all the values in the sequence above were positive rather than being an equal mix of positive or negative, the tracking score would be 4/4, i.e. +1.0. If, on the other hand, all the differences were negative the score would be -4/4, i.e. -1.0. A score of plus or minus 1.0 means that you have 100% statistical confidence that performance is 'off track' – either too high or too low, respectively.

You don't need a fancy technique to work this out in the examples I have just used. But what if the actual is off target but in a less clear-cut way? If there were three positive differences and only one negative the tracking score would be 0.5. So it is 'off target', but is this difference significant from a statistical point of view? Would we be justified taking action to close this gap or would we be simply reacting to noise?

To answer this question we need to make one simple adjustment to our method. Rather than using a simple average to measure the net and absolute deviation we need to use a smoothed average, in the way described in my second paper[1]. If we do this we can compare the tracking signal to published statistically based control limits, as shown in the table below (based on a smoothing parameter of 0.2):

Tracking Score	Confidence Level
0.92	99%
0.81	98%
0.74	95%
0.66	90%
0.54	80%

This means that if you have a Tracking Score of plus or minus +/- 0.92 then there is a 1 in 100 (1%) chance of a false positive signal (Type 1 error). In other words there is a very low probability that the 'off track' signal is the result of pure chance, so action is justified. But at 0.54 there is a 1 in 5 (20%) chance of a false positive so we should either ignore the tracking signal or seek out other evidence to justify intervention.

In practice I suggest you set the lower warning control limit at the 90% confidence level (a tracking signal greater than plus or minus 0.66). This means that using monthly data on average you will get roughly one false 'warning' signal a year (a 1:10 chance), which I judge to be inconclusive but worth the effort of investigating. I recommend you set the higher limit at 95% (above 0.74). Since the chance of a false positive signal is only 1 in 20 I'm inclined to treat this as strong evidence of a problem (with the actual or the target) unless I find evidence to the contrary (such as a large outlier in the data series).

Below is an example showing how the Tracking Score can be used to track performance and generate alarms.

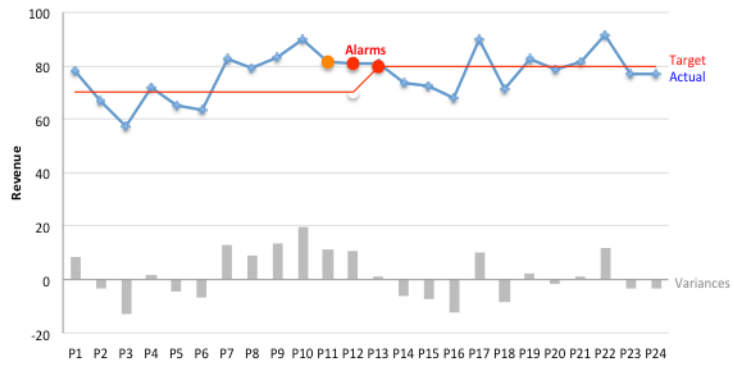

This chart plots actual revenue against a rate-based target. The difference between the two (the variance) is shown at the bottom of the chart. Alarms generated by the Tracking Signal are marked in amber and red. Note that they 'switch off' as soon as the target is adjusted.

[1] Steve Morlidge: Beyond Variance Analysis - Part 2: From static to dynamic measures: Understanding the patterns in Performance, April 2016

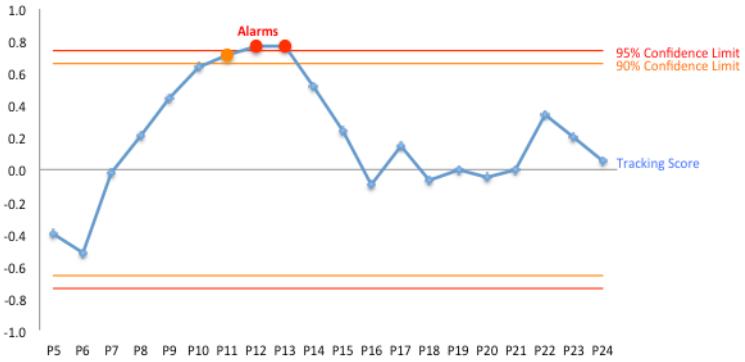

Here the calculated Tracking Signal is plotted along with confidence limits. Note that the Tracking Score does not need to be reset when the target is changed as it automatically adjusts to the new situation.

As you can see, a rate-based target has been set for revenue against which actual performance is monitored using a Tracking Signal. This results in an amber alarm being generated in Period 10 (when the 90% confidence limit is breached) and red alarms 11 and 12. In this example, management launch an in depth analysis in response to these alarms and decide that the improvement is real and should sustainable and so increased the performance benchmark in P13. As a result the variable comes back into line with the target and the alarm 'switches itself off'.

Below is an example of the use of this technique with a moving target and actuals (using a Moving Annual Total).

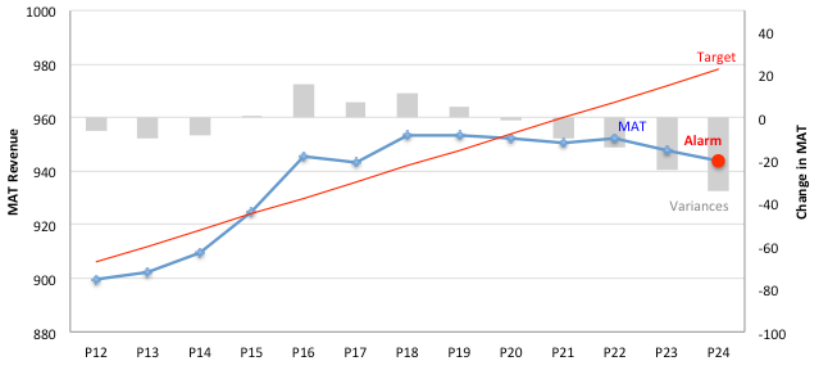

Here the MAT is compared to a trended revenue target. Variances are plotted on a secondary axis.

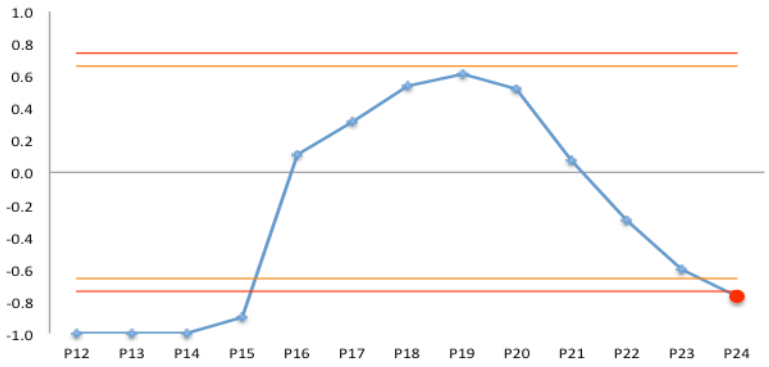

A Tracking Signal alarms is generated in Period 24. Note that no alarms are generated in the first four periods because smoothed averages cannot generate statically significant results until a certain amount of data has been collected.

The tracking signal allows us to continuously monitor performance based on sound statistical criteria in a simple and intuitive manner. This has many practical benefits, for example:

- It enables us to move away from fixed, period end targets
- It enables us to regulate the performance of the business continuously, since targets can be changed as and when required rather than as part of an annual set piece exercise.
- It can be used to monitor the behaviour of any kind of variable against a reference value. For example it can be used to track independent variables (drivers) or the accuracy of forecasts.
- The sensitivity of the system can be managed by changing the level at which alarms are triggered
- It provides a scientifically robust approach for exception reporting to replace the arbitrary criteria used by most current 'traffic lighting' systems.

Do we need targets at all?

The tracking signal gives us a much more sophisticated tool to use to manage and track the performance of a business, but given the 'charge sheet' we have against fixed targets of any kind, and the fact that we have 'target free' methods of measuring performance such as trend based measures and control charts, we still have to answer the question 'do we need targets at all?'

My sense is that the answer is not a simple binary 'yes' or 'no'. It is clear that the traditional approach of creating an arbitrary fixed target for almost every measure is unnecessary and damaging, but it is difficult to conceive an organisation operating without any form of quantitative limit or guidelines. But we definitely do need to be smarter in how we apply them.

I will illustrate this with an example that will be familiar to many readers.

A case study

In the world of business we have come to see target setting as natural and straightforward, but it is anything but. To illustrate why we are wrong to be so relaxed about target setting I will use an example of target setting in sport, specifically English football. Just like business, football is a team game and shares an obsession with measuring and motivation, so it is a good model of the challenge of target setting in organisations.

For the benefit of readers not familiar with football here are a few simple facts.

Every football team is made up of 11 players, each of whom has a defined role, just like the employees in a business organisation. This might involve scoring goals (attackers) stopping the opposition scoring goals (the goal keeper), block the path to goal (defenders) or acting as a link between attack and defence (midfielders). Within each of these groups individual players may have specific additional duties but everyone is expected to contribute to the team whenever they can. So, when necessary, an attacker might take on some defensive duties or vice versa. The English Premier League has 20 clubs who compete to win the league and to avoid relegation (3 teams per season). The position of each team in the league table is determined by the number of points that they accumulate based on a system of 3 points for a win and 1 for a draw.

What kind of targets – if any – should be used to measure the performance of a team and motivate its players?

Clearly, the league position is the most appropriate measure of performance since the ultimate goal is to win and there is a level playing field since over a season every team plays every other, so performance is truly comparable. This is what we call a relative target.

Why wouldn't you use a fixed quota of points as a target instead of the relative target of the league position?

Perhaps we should also phase it over the year – perhaps by quarter of the season like businesses typically do?

The reason why a points based target might not make sense is that although there is a relationship between the number of points and performance, the number of points is not a perfect predictor of the league position – which is the ultimate goal. So since the inception of the Premier League in 1992, a club has won the league with as little as 75 points and has failed to do so with 89. At the other extreme, a club has survived relegation with 34 points but suffered the drop with as many as 42. So while it might make sense for a team that has ambitions to survive to have rough aim to securing 40 points it doesn't guarantee that this will be enough and so should never become the only target given to the team.

In the same way, having an ambition to achieve 10 points in the first quarter of the season might be a good performance milestone but it would make a poor target since, in the first quarter, simply by chance, your team might have had easy matches against struggling teams weakened through injury. Even at the end of the season when the league position is known fans and commentators will take notice of contextual factors like the injury list, or how well the team performed in other competitions or whether the final league position was an improvement on the previous year in assessing the performance of a team.

In summary, the only true measure of performance is relative performance made by comparing the output (numbers of points achieved) to an equivalent organisation facing the same conditions. Other metrics can be used as a guide, but they should only be used as a proxy for the real aim, or to achieve a subordinate purpose (e.g. to improve an aspect of performance such as their defensive record), not as the goal.

When it comes to measuring and motivating individual members of the football team things get even trickier.

Some attributes of individual performance have always been measured, such as the number of goals scored by each player, but these days many other things are quantified as well, such as the number of passes or tackles made, how many were successful and the distance each player runs in a game.

Clearly there is a relationship between each of these things and the performance of the team. In business parlance we might call them performance drivers, KPI's or leading indicators. If you would always like your strikers to score more goals, and the more ground your midfielders cover in the game the better…all other things being equal.

But, in a complex system like football, all other things are never equal, which means that performance on one measure does not automatically translate to performance on another. For example, in the 2015/6 season Leicester City won the Premier League despite having fewer touches of the ball than any other team in the competition.

But what would happen if – as is often done in business – a football manager decides his or her players should be given targets on each of these measures in order to manage their performance. After all, it makes sense that is every individual achieves their targets, the team as a whole will achieve it's goal … doesn't it?

If the answer to this rhetorical question isn't obvious to you imagine what would happen if a full back (whose job involves sprinting up and down the edge of the pitch to support both attack and defence) discovered that he was 600 metres short of his target five minutes before the end of the game. The likelihood is that he would spend the rest of the game running up and down the touchline irrespective of what was happening in the game – particularly if he was incentivised to hit his target! And what about the midfielder who made easy passes all the time to fulfil his or her quota of 'successfully completed' passes or had achieved the annual quota of 6 goals halfway through the season? Would fans be happy if he stopped trying to score any more goals then on? Finally, what if the main striker was on a bonus for each goal he scored? Would this increase the performance of the team or would team mates be jealous and stop helping him[2]? And in all of these cases how would we decide in advance what the right target for each measure should be?

[2] When I used this example at a conference I was speaking at in Brazil the translator we used told me that this has actually happened in the team he supported. Nobody passed the ball to the striker!

The key message is that context is critical when it comes to assessing performance or determining the right thing to do, particularly in a complex environment. Excessive focus on targets crowds out context. Hitting the target become the purpose rather than an aid to achieving the purpose … the means become the end.

In the context of football it is easy to see how the automatic application of fixed targets to measures would promote stupid and suboptimal behaviour and provide us with very little useful information about the contribution of different parts of the football organisation to the performance of the whole. And yet, in business, the ritual practice of attaching an arbitrary comparator to almost every dimension of a business that can be measured is not only treated as normal but is often promoted as good practice.

But, on the other hand if there were were no quantitative reference points for a football team or constraints placed upon the actions of its players actions we can be fairly sure the consequences would be not be good. So what is the right thing to do?

It is clear that we need to think more deeply about targets and target setting than we normally do.

A good place to start is by asking ourselves the question: 'what is a target?'

What is a target?

'Target' is one of those words that is used so liberally we seldom stop to consider what we mean when we use it. My guess is that if ten people in a room were asked to define it they would come up with ten different answers. Moreover, the problem is further confused by the ambiguous and inconsistent use of language. Words such as target, budget, objective, goals, benchmarks, guidelines, reference point and comparator are sometimes used to describe the same phenomenon.

I don't think there is much to be gained by trying to finesse the way words are used…not at least until we have a clear sense in our heads of what 'it' is, whatever term we use to describe it.

Instead let's answer the question by taking a step back and think the problem through – starting with performance measures. A performance measure is a value of a variable that we believe is relevant to the purpose of the organisation – its ability to survive and thrive. This could be an input 'flow' variable (e.g. cost, time taken), an output 'flow' variable (e.g. revenue or customer service) or a 'stock' variable (e.g. work in progress). We could also measure a KPI (Key Performance Indicator), which is often a compound measure, such as costs expressed as a percentage of income.

By definition, if we believe a performance measure captures something important we should care what the measure tells us. At the most basic level we must know whether, all other things being equal, bigger values are better than smaller values or less variation is better than more variation. And because resources are always finite we cannot avoid making choices. If there were no constraints on action at all an organisation would run out of money or fail to meet the needs of investors, which would threaten its viability. We must, therefore, have a hope or an aspiration for the value of a performance measure. Following this line of logic, a target is a measurable expression of this aspiration, such as 'more than', less than' or 'between'.

Conventionally, however, the word 'target' is usually applied in a very restricted form; typically describing a fixed constraint applied to output variables (performance) and anchored at a point in time. Terms like 'budget' tend to be used to describe constraints applied to input variables, which are also fixed at a point in time, whereas guidelines convey a sense of a range, perhaps one that does not have a fixed 'maturity date'. But whatever the term used they all have the same purpose, to constrain or influence behaviour by acting as a reference point for the value of a variable.

Expressed in these terms it is fairly clear that targets – whatever word is used to describe them - are needed in some form. The question therefore is not should we have targets or not, but how many should there be, what form should they take and how should a difference between actual outcomes and target be interpreted?

What form can targets take?

Traditionally we have been brainwashed into thinking of targets needing to take a very particular form. For example, I was trained that targets had to be 'SMART' – an acronym standing for Specific, Measurable, Attainable, Relevant and Time bound. But targets can take many different forms as shown in the chart below.

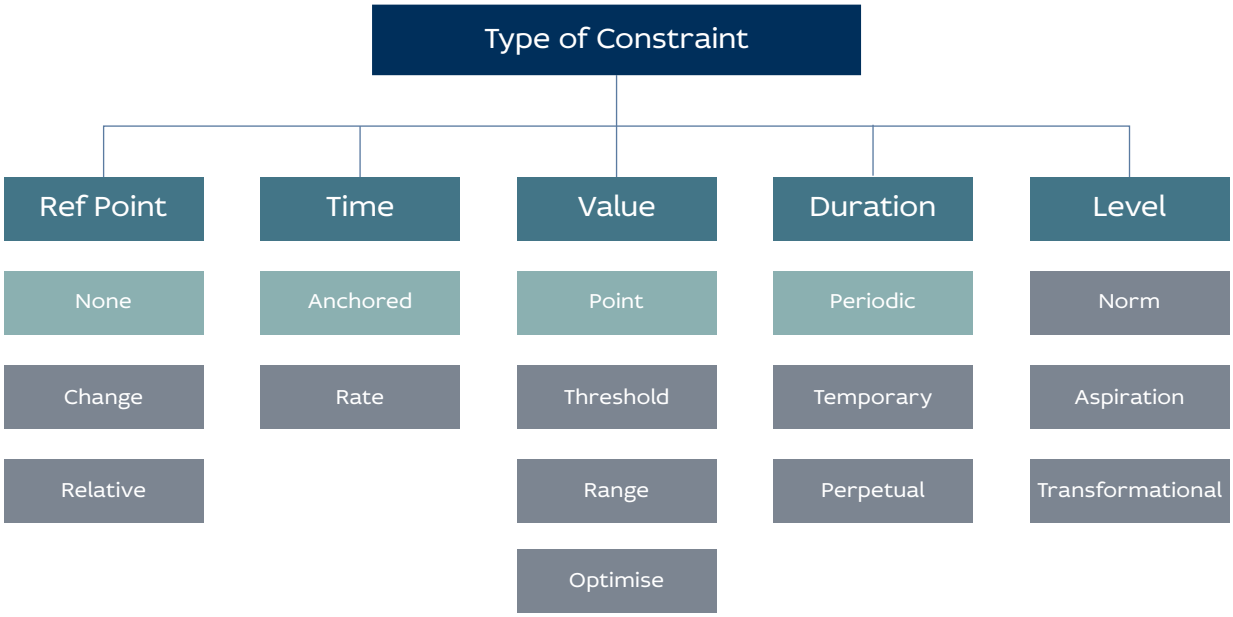

It is clear that targets can be expressed in many different ways. So how do we decide what is the right approach to use, and what implications does our choice have for how we measure performance? Let us explore these questions by considering each of the target setting dimension shown above in turn.

Reference Point

Conventional SMART targets have no external reference point: they are either imposed or negotiated. Either way the result is an arbitrary value; which begs the important question: what is the 'right' level for a target?

The future is unpredictable because it has never happened before so we have no firm basis on which to determine, in advance, what represents good performance. Even though it has been crossed millions of times in history a sailor in a transatlantic race can only judge his or her performance against that of other competitors in the same race facing the same weather and sea conditions. A football team can only judge their performance with reference to other teams and this cannot be predicted. And in business the performance of an organisation is ultimately judged (by shareholders at least) against that of its peers - which also cannot be predicted.

The simple logic for setting targets relative to peers is compelling and for this reason BB organisations like Handelsbanken embed this principal right through their performance management system. The bank as a whole has one simple goal – to consistently beat the average return on equity of its peer banks. And within Handelsbanken the performance of individual branches is measured by their cost to income ratio relative to their internal peers. In this way the target is always set at the 'right level' because the average performance always reflects achievable performance in the prevailing economic conditions. And performance is continuously driven upwards without the targets ever being changed because no-one want to be below average.

In practice, however, it's not always easy to follow Handelsbanken's example. For example, it may be impossible to define sensible internal or external peer groups against which performance can be compared and in the case of the latter getting hold of comparable performance information can be difficult. In these circumstances using targets based on a rate of change (i.e. relative to the past) could be a good alternative to a 'pure' relative target since it is not as dependent on having reliable up to date information on peer group performance – a reasonable understanding of long term trends in a business or industry might allow us to set aspirations at the right level. So, a rough sense of historic trends might be sufficient to set a rate based target for reducing costs at say 3% year on year (i.e. improving relative to the past). Or holding the KPI for administrative costs at around the industry norm of 5% of revenue might be another good (relative target) alternative to a fixed (absolute) annual budget.

Time

Traditionally targets are anchored in a point in time and expressed in absolute terms e.g 'profit should be $1m in the year to December'. But this isn't perfectly aligned with the purpose of a business, which is to continuously generate profit. Arguably, then, it would be better to express the target as a rate. This would also make manipulating the system to 'hit the number' impossible to sustain. As a result performance data would be more reliable and effort would be focussed on managing the business rather than managing the numbers.

Value

Most conventional targets are expressed as an exact value but it is difficult to justify this because the existence of noise makes it very unlikely that we would ever hit a number exactly without manipulating the system – which is the reason why period end peaking is an endemic in many publically quoted companies. Also, expressing a target in such precise terms may not reflect the behaviour we might want to encourage. For example, if we set a revenue target of $100m we don't want people to stop if they can sell more, just as if we set a cost budget of $100m we wouldn't want all this money this to be spent if we didn't need to. But this is exactly what we often see happening because targets are set in such a crude manner.

There are many other ways in which targets can be set that would more accurately reflect the objectives of the target setting process and less likely to generate dysfunctional behaviour. For example, rather than using a 'point target', an objective could be stated as 'more than' or 'less than' or 'between'. For example, in the human body it is very important to maintain equilibrium, which it does by controlling variables within ranges. For example body temperature is 'targeted' at 36.1°c to 37.2°c, but our heart rate is allowed to fluctuate more widely, generally between 60 and 100bpm in normal conditions, and often very much wider. In the same way organisational stability is a prerequisite for measured rational decision making in businesses. So, rather than setting a point targets for customer service, for example, it would be much more sensible to set it as a range since higher or lower levels of performance are costly.

This last point illustrates that targets should never be set in isolation. For example, it would make sense to set a customer service target as a range for the reasons explained above but to optimise cost within that constraint – in other words cost should be 'as low as possible' or 'lower than last period' with no lower limit.

Duration

Another benefit of expressing a target as a rate or in relative terms is that they do not need to be periodically renewed, thereby eliminating the bureaucracy and disruption associated with the annual process. Dispensing with annual process also means that targets can be changed more frequently, thereby enhancing agility. Targets could also be expressed as milestones, which lapse or are renewed not at a defined date but by the achievement of an objective. For example, it might make sense to set a target for the amount spent on a project or activity within its defined life rather than to bundle it with many other projects in an annual budget. An improvement goal that is achieved within three months should be reviewed then rather than waiting for the annual target setting exercise.

Level

Finally targets can be set at different levels. They can be 'stretch' targets that are difficult to achieve, in which case failing to meet them should be interpreted as an incentive to try harder rather than as a failure. On the other hand they can be set at the minimum acceptable level in which case variances would have a very different meaning. Or they could be norms that help orientate or guide behaviour, in which case it might not be appropriate to use them to analyse performance at all. And it is possible for one variable to have multiple targets representing different levels of achievement like the bronze, silver and gold medals awarded at the Olympic Games, all of which represent success but at different levels of achievement.

In summary, traditionally targets have been expressed in one very constricted way – as a fixed, 'point in time' absolute value. This is partly because we have had no other way of measuring performance other than crude variance analysis. Adopting a set of alternative measurement methodologies, frees the target setting process from this artificial constraint, enabling to use them in more intelligent and subtle ways that are better aligned with corporate purpose.

Conclusion

At its simplest a target represents a boundary. Boundaries help guide decision making, coordinate actions, constrain or encourage behaviour as well as serving as a reference point for assessing and interpreting performance. Organisations need boundaries, because without them they would not be organised. So targets are necessary, but they are not inherently good or bad. Like any tool they can be used or misused.

Over recent years we have seen an increase in the use of targets in business, partly because we have more data than ever before but also because they appear to simplify the complex task of management, reducing it to a simple arithmetic operation. But in 'solving' this problem we have created a whole host more. In business today there are too many arbitrary targets that are too tightly specified and locked into inflexible and bureaucratic planning process. The sway that targets hold on the corporate imagination means that those things that cannot be easily measured and so cannot be targeted become neglected. And the linking of incentives to those things that can become the primary focus of attention distorting collective behaviour in sometimes grotesque ways. Hitting the target has become the purpose of work rather than a means to achieving a purpose.

What has this got to do with performance reporting?

There is an old adage that if the only tool that you have got is a hammer everything looks like a nail. I believe that the crude and limited manner in which we traditionally measure and analyse 'performance' has contributed to this state of affairs. We need to focus less on points in time and more on trends. We need to acknowledge the existence of noise. This means that we need to adopt a statistical perspective on data and cease to treat the simple arithmetic difference between two numbers as potentially meaningful. And tools like tracking signals give us the ability to do this in a simple and intuitively attractive way.

But it would be a mistake to use tracking signals as a sticking plaster and carry on using targets in the same way as before. We need to question the assumptions upon which targets are based and adopt a more sceptical stance rather than treating them as 'god given' and immune to criticism. We need to analyse the performance of variables that are not or cannot be targeted. And we should help fight the corporate vice of setting targets in order to help analyse or to 'drive' performance since there is a good chance that it will have exactly the opposite effect to that intended.

In summary, we don't need more SMART targets to help measure and manage performance; we need to use targets in a smarter way, and the Beyond Budgeting model provides a framework that targets the root cause of the problem – the way that we have been taught to think about management.

The end result will be a more nuanced, thoughtful and insightful analysis of the state of organisational affairs, beyond variance analysis.

The aim of this series of working papers is to share insights that might be of interest to the Beyond Budgeting community; particularly those that might help make the journey out of the land of budgeting easier. The other purpose is to expose material and arguments that I plan to use in a forthcoming book to friends who I hope will help me do the best job I can of communicating it to a more sceptical audience. 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