Chapter 2

Low hanging fruit syndrome (a.k.a. picking the population of least impact…)

PRESENTING SYMPTOMS

- Successful changes are completed yet on average not much overall impact is achieved
- Project teams demonstrate improvement on their individual charts, yet there is a sense of disappointment by the sponsor that the overall goal wasn't met or the underlying target didn't really shift
- Initial round of projects got great results, and then when asked to do a second round on the same topic the results were disappointing, despite a lot of hard work by everyone
- Someone always seems to get good results on their improvement projects but you don't
- There are concerns about sustainability (proven or otherwise) of results

DESCRIPTION

Anita Roddick, social entrepreneur and business woman who used to own the Body Shop business is credited with saying, "If you've ever spent the night in a tent with a mosquito then you know that small things really do matter."

The **population of impact** is the small, *purposely and analytically identified*, target group where if you place your efforts for change, will have the greatest overall benefit.

I'm not a detail person and I'm definitely not a statistician and it was probably for this reason that it took me many years of working on improvement projects before I really understood what the population of impact really meant. It is not about broadly identifying the target group of people to work with. It is not even a little bit of analysis checking who might be the most interesting group to work with. It is definitely not the group who volunteers to work on the project, though that is often the most prevalent group both joining the project and often the group targeted to work on.

What I have learnt is that if we want to make an *impact*, and the key word here is impact, on improving outcomes for patients in healthcare, then we need to do an inordinate amount of diagnostic work upfront, before the action starts, to identify the population of impact.

There's nothing new in this. It is good old fashioned quality improvement work, though in the haste of many of the projects I have seen in action, this does tend to get forgotten.

Let's look at an example of what happens when we go off beam and then track back to the underlying behaviours. In a large primary care group practice they have identified that 500 of the total 1000 of their patients who have diabetes have reported in a survey that they are unhappy with the care they are receiving and these same patients also appear not to be receiving the type of care that current evidence suggests would be appropriate. So the leadership team asks the local quality improvement co-ordinator to run some projects in each of their 10 primary care practices to improve the diabetic care for these patients.

Each of the 10 practices enthusiastically sets out on implementing the best evidence into practice using a variety of techniques. In a matter of months they each can demonstrate that for a sample of their patients they have implemented a variety of new initiatives ranging from group visits, newsletters, in one they have implemented a new clinic, in another medication reviews and the like. They hold a review session where everyone shows a poster with charted improvement results. It is likely these very specific improvement results are focused narrowly on areas on which they have chosen to work, may only reflect a very specific clinical focus, few may show other outcomes (balancing measures) and in my experience very few report on unintended consequences generated in the system.

However, despite these reported successes, on a repeat patient survey 350 patients still reported dissatisfaction and the overall clinical results showed that many patients were still not receiving the test results and care they needed. So what is going on? Despite all this hard work and enthusiastic activity at the practice level, something is missing. Impact.

If, before starting the action phase of the project, the quality improvement co-ordinator, working with the primary care practices, had spent a few weeks, sometimes it is months, doing some detailed analysis understanding the target population - to find the population of impact, she may have identified the unique group to benefit *directly* from the proposed changes.

If you have the choice between working with 89 diabetic patients on whom you can make a 7% improvement, or 13 patients on whom you can make a 56% improvement, which group will you decide to work with? In some primary care practises they will identify that a small number of patients account for the majority of their poor results. If they work to improve the care for these few identified patients they may make a large impact on the overall. This is often the case with medication spend. Rushing to test a

group visit model with a few patients for whom there are little or small issues reported may help you test out a new way of working, however, it may have little impact on those patients who require solutions to their problems.

Now, I hear you say this is not as easy a decision as it sounds. It does depend, of course on what you are trying to achieve, so you will need to refer back to your aims (which you will have agreed before you started…). However, I urge you to consider how these decisions multiply up over a number of primary care practices. For a leadership team who have spread as their goal this is an important consideration. There is a real dilemma here between the *number* of patients and the extent and depth of the total *impact*. The tension between these two also drives the type and timing of solution that gets designed and implemented, and then of course the **knock on consequence of its ability to be spread to others.**

When investigating the lack of impact in some programs and projects my conclusion is that the cause is a combination of the lack of prior analysis and then is closely linked to the use of Plan-Do-Study-Act (PDSA) cycles.

While this model of starting off by testing small cycles of change is incredibly useful as a means of discovering what works and also reducing the risks, thus getting people engaged in the process, it does encourage the bypassing of the population of impact focus unless the project leader and team are very disciplined.

You'll have figured by now that this main issue is the Pareto Effect; that the minority of causes or inputs, leads to the majority of effects or outputs. So fancy that. This may mean that we can work on a small part of our system and get the maximum benefit. What it also means is that we might have been working on large parts of our systems and not getting much benefit at all. Why would we do that? Because often it is an easy way

forward to get a quick fix, to appear to be taking action - and that is why we call this the low hanging fruit syndrome.

Picking the apples, plums and pears from the orchard in our back garden is always fraught with difficulty and temptation. The problem with crawling along the lower branches is they eventually break off under your own weight if you "go out on a limb". And the best fruit always seems at the top of the tree – just out of reach, where you have to take risks to get at it. In the end, each year we have to balance our actions to maximise the amount of produce from each tree while minimising the wastage, sharing our fruit with the birds and bees, avoiding any injury to ourselves and not causing any damage to the trees. It's a balancing act.

As with all balancing acts, when you know a bit how the system works you can manipulate it. When it comes to projects and the spread of good practice I have come across some "cheat" methods. For instance I know how to get good results for projects. All I have to do, or recommend, is to go for the low hanging fruit (or keep quiet about the population of impact – the difficult target groups). One way to look good (in the short term) is to go for these easy groups and get great results, though I know that when they all get added up, across a community (like in a spread program) the overall impact (long term for the organisation and for patient care) may not be significant.

Now there is another facet to this issue which demonstrates both the importance of the Pareto effect as well as some other technicality around the way we have come to use numbers in many of our improvement projects. If we remember, the whole point of what we're looking at here is to understand why good ideas don't spread. Well, let's look at a waiting time example.

Say a hospital aims to reduce the length of time a patient has to wait for an elective knee surgery from 90 days to 45 days. Currently it has 100 patients waiting and the variation for the past year is very stable at no more than plus or minus 10 in any one month. So, the leadership team sets the improvement team a target of 80% of patients to meet the 45 day target 95% of the time.

Immediately the improvement project leader will be spending a great deal of time working out how to make this calculation work and how to optimise the denominator and nominator of this target to his benefit. Many project leaders and teams become experts at contextualising targets and ensuring redesign work is highly specific to their environment. The successful change they do eventually make may not be the same ones that may work in another hospital to achieve the same target for precisely this reason. It would be much easier and more positive for the improvement team to be working in whole numbers and to be designing a solution to look at the whole 100 patients for the 45 day target. That way they will also be encouraged to seek within that group which of the 100 patients, because they will fall into population of impact categories, will best be served most appropriately. By doing detailed Pareto analysis upfront on their whole, they will learn to assess their system imbalances and also discover more information and knowledge about the predictability in their system. This is nearly always not what is expected, quite often counter-intuitive and can best be discovered by data analysis. This is the crux of good Pareto analysis that works with numbers and not percentages and averages.

One other advantage of working with actual numbers is it helps to engage the team members and wider group. Aims and goals like the one above may make sense to managers, however, to many clinical staff who deal with "whole people" percentages and averages have little direct meaning. To be able to count the number of people whose lives have been

saved or improved is very rewarding. Or to count the amount of medication errors reported, or the amount of money saved this week on this ward by not throwing away drugs – all contributes to sense of purpose and meaning underlying the improvement work.

What has this got to do with spreading good practice? I've seen great examples like the few mentioned above where there has been spontaneous adoption. People like to see what others are doing and if it appears to be working well and it is being talked about then they'll have a go too. And that has got to be an easy way to get an idea spread around an organisation. I have very few examples of people telling me how they have spontaneously adopted something like reducing errors by 95%, or saving 80% of our medicines budget each week. In comparison, I do have examples of spontaneous and rapid adoption of good practice where the aims have been expressed in terms of saving £xx/$xx medication wasted each week on the ward, or y less administrative hours spent per week due to reduced medication errors on the ward.

So the challenge of the population of impact is to find that small group on which to test out our changes, right from the beginning, knowing this is the group where we intend to make the most impact, not where we expect to have the most successful change. Successful change does not necessarily mean useful improvement nor does it mean easy spread, as we have already seen from pilotitis.

SUGGESTED TREATMENT

Depending on the topic on which you are working the recovery from this condition is fairly straightforward once you have recognised you have a problem and are committed to resolve it.

Do the analysis.

Depending on what you are trying to achieve, work out, in detail, using the actual numbers, which population will have the most impact. You may have to cut your data a number of ways. Chart it using a Pareto chart (if don't know how to do a Pareto chart – do an Internet search to find out – basically it looks like a histogram except it is discrete data rather than data over time). It's also important that you chart your data so you can see the *pictures* that it throws up. There is nothing more intimidating or disheartening than rows of figures. Make sure that whichever way you cut your data you are working with the large chunks of population and if you're not, then you understand why you aren't.

Spread is about impacting the population, about moving an idea that works, across a large number of people. Work out whether and how you are making an impact for that large population. If necessary make the appropriate changes to your project work and make them now.

On a more dynamic level, PDSA cycles are a great change process, especially for the start of a project. However, an even better methodology is the broader Improvement Model (please see further reading for more details) which has the addition of three questions to PDSAs

1. What are you trying to achieve?
2. How will you know you've made a difference?
3. What changes can you make?

As a minimum, ensure teams are using the Improvement Model as this will keep a proper focus. If it were me, I would enhance the second question and rewrite it as "How will you know you've made an impact?" as that will force not just the measurement approach, but also some discussion at team meetings about the concept of impact.

PREVENTING OCCURRENCES

Define the aims for your improvement work clearly at the outset, with impact in mind. You'll be surprised but I do find large programs that have no clear answer to the question "What do we want to achieve?" and that also means they have no way to measure their progress. We'll see later that even spread programs need measurement, so even at the basic improvement project level there needs to be this starting point.

Carry out a *population of impact analysis* before starting any improvement project and specifically one that has a spread intention. From small actions, large impact – both depth and breadth can occur. This is important both at a strategic and operational level. It can be done at a systems level, organisational level as well as very specific project level; with each level contributing to the impact whole.

A common and very effective improvement spread methodology is the Breakthrough Collaborative designed by the Institute for Healthcare Improvement (Cambridge, USA www.ihi.org), and I recommend it – please see further reading for more details. One useful step I have learned is to add in a diagnostic step before the first Learning Session as part of the prework phase to ensure sufficient analysis is done to discover the population of impact. In a complex topic this can take some extra weeks as teams develop the detailed assessments, data, pareto analyses and histograms.

PROGNOSIS

If caught in time and appropriate diagnostic action taken the outlook is very good.

Institute for Healthcare Improvement www.ihi.org for more information and training on the Breakthrough Collaborative Series.

Koch, Richard., The 80-20 Principle, Nicholas Brearley Publishing, London, 2003; Excellent introduction covering the broad principles in an easy read with some useful examples.

Langley, G., Nolan, K., Nolan, T., Norman ., Provost L., The Improvement Guide, Jossey-Bass, 1999; Covers the essential basics with examples from outside healthcare.

Nelson, G., Batalden, P. & Ryer, . Clinical Improvement Action Guide, Joint Commission, 1998; This is the standard and most comprehensive text for any improvement team in healthcare**.** It is a most practical workbook, written by practitioners who know their stuff.