

Common Cause: How Six Sigma Can Drive Better Management Reports

In many companies, performance reporting processes encourage managers to react every time a metric falls outside the desired range. But looking for a reason behind every fluctuation in results wastes resources and can camouflage the organization's true problems.

by **Forrest Breyfogle**



Forrest W. Breyfogle III is the founder and CEO of Smarter Solutions Inc. and developer of the Integrated Enterprise Excellence (IEE) business management system. You can reach him at forrest@smarter-solutions.com.

IT'S NOT UNCOMMON FOR A COMPANY to track its performance solely in terms of conformance — and non-conformance — to plan. What many organizations don't realize, though, is that this method of performance measurement can drive the wrong behaviors among employees. Nonconformance tracking can be appropriate in monitoring product quality when there is no question whether a product is satisfactory. But nonconformance tracking when comparing financial results with expectations or goals can lead to destructive behavior that actually erodes long-term business performance.

An exclusive focus on making the month's numbers often leads to the expenditure of a tremendous amount of energy on the narrow purpose of living up to promised or budgeted financial targets. Krispy Kreme shipped donuts that managers knew would be returned so that they would meet their quarterly objectives, and the Enron management team made notoriously poor decisions to meet corporate financial goals. These examples are extreme, but they are just the tip of the iceberg. Weak assessment processes lead to wrong behaviors on a much smaller scale in all kinds of companies every day.

An approach to performance measurement that I call "Smarter Six Sigma Solutions (S⁴)" or "Integrated Enterprise Excellence (IEE)" can open managers' eyes to ways in which the metrics they've chosen are driving the wrong employee behaviors, and can help them focus improvement efforts on actions that can truly impact performance.

Different Solutions for Different Problems

The biggest mistake management teams tend to make is responding in exactly the same way to all instances of nonconformance. In every realm of corporate management, there are two types of variability. One is what Dr. W. Edwards Deming, known for management practices he pioneered in post-World War II Japan, called "special cause." These problems result from short-term glitches in



a process, such as faulty product assembly by a new or temporary employee. Deming called the other type of variability “common cause”; these problems result from situations such as the predictable variability of suppliers’ raw materials or the typical variability that a process experiences due to differences in people or machines.

When an organization responds to special- and common-cause issues in the same way, it will fail to solve most of the problems that arise. Regardless of the cause of perceived nonconformance, operations staff will go into a firefighting mode whenever they are notified that a product, service, or financial result is not meeting specifications or planned objectives. They may resort to playing the blame game, and the problem may be “resolved” temporarily with Band-Aid changes or disciplinary action.

To understand how counterproductive such a response is for most variability in a company’s outcomes, consider Deming’s estimate that 94 percent of problems result from common-cause variability and only 6 percent result from special-cause variation. This means that the vast majority of nonconformance situations require modification of the system itself.

Successful executives continuously analyze their systems to identify and eliminate, or at least reduce the frequency of, common-cause nonconformance. How do they do that? By taking the right process-improvement corrective actions, which they determine based on the reports they use to assess performance relative to specification limits or goals.

How Reports Drive Bad Behavior

I’m not breaking new ground by saying that what we measure is what we get. It’s been said before. But many companies continue to use metrics that lead to the wrong kinds of behavior. Consider, for example, a call center’s measurement of duration of hold time. This metric may make good sense from a distance for an executive who wants to assess customer satisfaction. Long hold times would presumably correlate to customer dissatisfaction. However, the metric has the potential to drive the wrong behavior unless safeguards are implemented to prevent abuse.

To see how, think about an operator striving to achieve the target for this metric during an understaffed peak call period. His strategy for meeting his hold-time goal might be to periodically pick up a line that is on hold, ask the caller to hold for longer, wait for a reply, and then quickly place the caller back on hold again. It may not be bad policy to check with callers and ask whether they can continue holding, but this behavior should not simply be the result of an employee wanting his duration-of-hold-time metric to look good. A call-center business looking at hold times as a leading indicator of customer satisfaction should, instead, capture the total time elapsed from a caller’s initial connection until she speaks with the appropriate person.

Many companies, recognizing that their metrics are sending employees the wrong message and providing faulty incentives, have rethought the measures they use to gauge performance. But for most, the process stops there. That’s

When an organization responds to special-cause and common-cause issues in the same way, it will fail to solve most of the problems that arise.

Exhibit 1

Traditional Performance Metric Reporting

This is a performance report taken from the Web site of the city of Austin, Texas. The format simplifies year-to-year comparisons of data but provides no context, no way of determining whether variations between years have a common cause or a special cause. Many organizations’ reporting processes encourage managers to come up with “stories” to explain the movement of metrics in reports like this one. These stories frequently lead to firefighting efforts, in which immediate action is taken to remedy changes that actually fall within the range of predictable outcomes of the organization’s current processes.

Performance Metric	FY2001 ACTUAL	FY2002 ACTUAL	FY2003 ACTUAL	FY2003 AMENDED	FY2004 AMENDED
% customers satisfied with dispatch staff	99.99%	100%	99.99%	98%	98%
% priority-one calls dispatched to field crews within 30 minutes of receipt	99.99%	99%	99.99%	95%	95%
Labor cost per customer call taken	\$4.20	\$5.31	\$5.09	\$4.88	\$5.09
# calls taken	62,054	59,828	63,046	60,000	60,000
# priority-one calls dispatched to field crews	5,797	4,828	6,686	5,000	6,500
# work orders & component parts (segments) created in database	8,226	4,724	7,742	5,500	6,700

Charts delineated by calendar boundaries make visual comparisons between months or years easy, but they can lead to erroneous conclusions.

a shame. Just as they need to put energy into selecting the right metrics, corporate managers should think about whether the data presentation and assessment formats they use will lead to the right employee behaviors.

Exhibit 1, on page 17, exemplifies a commonly used performance measure report-out. The format has calendar boundaries that reflect only average annual results. This type of chart does not present response data as though it were a result of internal processes with inherent variability; it does not identify common-cause events as such. Nor does it reflect a desire to identify trends or to detect unusual, special-cause events on a timely basis.

Instead, the inclusion of corporate performance data in a report like this usually leads to “stories.” Someone presenting results in this format in a meeting might give an explanation for some of the upward and downward movements between quarters or years. This approach is not dissimilar to a nightly stock market report, in which analysts try to pinpoint a specific reason behind even the smallest movement in share prices. Whether we’re talking about business performance or stock prices, stories that are developed to explain changes in metrics often implicate, as causal events, occurrences that did not actually have any effect on the result. A short-term change in the performance of one metric may trigger improper corrective action, or the report may be simply viewed as a for-your-information (FYI) report that has no real value. In either case, much of the variability in metrics is likely a result of the system’s common-cause variability, and treating every change in results as a special-cause variation is inappropriate.


Charts delineated by calendar boundaries make visual comparisons between months or years easy, but they can be difficult to interpret, can lead to erroneous conclusions, and can encourage managers to come up with “stories” about the past without providing any formal system for describing what to expect in the future. Companies using this type of report must consider what it is that interests them the most: the past or the future? Reporting “stories” to explain each historical up-and-down metric movement does not provide accurate insight to what inputs collectively impact the overall response, so this form of reporting does not provide insights into what might be expected for the process in the future. If, instead, a company’s performance reporting system provided prediction statements, insight would be gained to real improvement opportunities.

Exhibit 2, below, shows an example of another popular but problematic method of presenting performance data: the profit scorecard. The horizontal axis represents the months of the year. Because the goals are annualized, the target line sits at a consistent 20 percent for every month of the year. The scorecard provides no indication of how the organization performed in the previous year, or of whether a 20 percent profit is a reasonable goal or simply a pie-in-the-sky objective.

Exhibit 2

Calendar-Based Graphs Aren’t Better

Below is an example of a graphical report that might appear on a dashboard where the company’s profit goal is 20 percent of revenue. Like the grid format in exhibit 1, this kind of simplistic graph may simplify period-over-period comparisons, but it does not help managers decide how to respond. These traditional types of reports should be replaced by process control charts that show whether the process is predictable, along with (when appropriate) a probability plot that provides a prediction statement.



If people are held accountable for achieving metrics presented in this type of format, undesirable behaviors may result. When the profit line turns red, the owner of this metric sees that she needs to take immediate action to drive the profit number back up. This could lead to Enron-like behaviors, where money is simply shifted from one area to another to make the graph look better, or it could lead to knee-jerk cost-cutting measures that significantly damage the company’s future prospects.

Many traditional performance reports reflect only fiscal-year numbers, determine trends based on a comparison to a single estimate from a previous month or year, and fail to provide predictive statistics-based statements. Traditional reporting methods don’t view the enterprise as a system of processes that result in performance numbers, nor do they account for the fact that the numbers necessarily reflect all the variability that occurs within those enterprise processes. That’s a shame, because improvements in a company’s rate of common-cause nonconformance — and so a reduction in the frequency of nonconformance in the area that accounts for the vast majority of performance variations from plan — can be brought about only through systematic improvement of these processes.

“A fault in the interpretation of observations, seen everywhere, is to suppose that every event (defect, mistake, accident) is attributable to someone (usually the nearest at hand) or is related to some special event,” Deming once noted. “Confusion between common causes and special causes leads to frustration of everyone, [to] greater variability, and to higher costs — exactly contrary to what is needed.”

Harnessing the Power of Predictability

Approaching the display of performance data from an S⁴ point of view can yield much more useful management reports. Consider how a company might monitor a manufacturing process that

has specification limits of 72 (minimum) and 78 (maximum). Managers reporting on the performance of this process in a format similar to that of exhibit 2 might be tempted to include horizontal lines on the graph at 72 and 78. Where the process's performance fell below 72 or exceeded 78, they might color the result line red to draw attention to it. But a knee-jerk reaction to these out-of-specification conditions might consume considerable resources and might not fix anything long-term from a process point of view.

Alternatively, when information is presented with a statistical upper control limit (UCL) and lower control limit (LCL), as shown in exhibit 3, at left, report recipients can make statements about common-cause and special-cause variability. The UCL and LCL numbers can be derived easily from the data. Chart out the time-series data points you've collected (represented by the first two columns of exhibit 3's data table), then find the deviation from each data point to the next — its moving range (MR) value. The UCL is the mean of your data points plus 2.66 times the average MR, where 2.66 is a three-standard-deviation constant that applies when the moving range is determined from adjacent values. The LCL is the mean of your data points minus 2.66 times the average MR.


For control charts that have no data outside the LCL and UCL and show no clear upward or downward trends, quality control professionals will state that the process is “in control.” I prefer to simply call the process predictable. What does a process control chart like the one in exhibit 3 predict? We can consider the data in the stable region (between UCL and LCL) to be a sample of not only past process performance, but also future performance — assuming nothing either positive or negative changes the system. That is to say, we can expect common-cause variability to stay within the bounds of the upper and lower control limits. We can view this data-analysis procedure as a way to get the process to talk to us — to tell us what we can expect it to deliver in its current state.

Exhibit 3

A Predictable Process

Companies would be better off analyzing their performance by plotting individual data points on a control chart like this one. The upper control limit (UCL) and lower control limit (LCL) are calculated using the mean of the data points and the mean of the moving range (MR). Once calculated, the UCL and LCL can be used to determine whether a process is predictable (i.e., “in control”), in which case reductions in the range of variance in results requires a change in the fundamental process being measured.

NO.	MEASURED VALUE	MR	NO.	MEASURED VALUE	MR	NO.	MEASURED VALUE	MR
1	76.20		18	77.60	0.60	35	77.70	2.10
2	78.20	2.00	19	77.10	0.50	36	76.90	0.80
3	74.10	4.10	20	75.20	1.90	37	76.20	0.70
4	74.10	0.00	21	76.90	1.70	38	75.10	1.10
5	75.00	0.90	22	78.30	1.40	39	76.60	1.50
6	74.50	0.50	23	72.70	5.60	40	76.60	0.00
7	75.00	0.50	24	76.30	3.60	41	75.10	1.50
8	75.00	0.00	25	78.50	2.20	42	75.40	0.30
9	71.80	3.20	26	76.00	2.50	43	73.00	2.40
10	76.70	4.90	27	76.80	0.80	44	74.60	1.60
11	77.80	1.10	28	73.20	3.60	45	76.10	1.50
12	77.10	0.70	29	78.80	5.60	46	79.30	3.20
13	75.90	1.20	30	77.60	1.20	47	75.90	3.40
14	76.30	0.40	31	75.20	2.40	48	75.70	0.20
15	75.90	0.40	32	76.80	1.60	49	77.90	2.20
16	77.50	1.60	33	73.80	3.00	50	78.00	0.10
17	77.00	0.50	34	75.60	1.80			



MEAN	76.09
MEAN MR	1.73
UCL	$76.09 + 2.66 (1.73) = 80.69$
LCL	$76.09 - 2.66 (1.73) = 71.49$

Upper and lower specification limits should not be shown on a control chart because they have nothing to do with the statement of whether a process is predictable.

Although this example shows a manufacturing process with both upper and lower control limits, the same approach could be used to display data that had only a one-sided specification, such as frequency of on-time departures for an airline, a ticketing quota for a police officer around a particular speed-limit zone, or even desired corporate revenue. The basic point is that if the process is predictable, we need to look at all the data collectively when making a statement about conformance. It's not valuable to try to explain why a single data point is up and another is down. Only if a process exhibits a special-cause condition — if one data point exceeds the UCL or LCL by a substantial margin or a series of data meets certain other statistical criteria indicating a pattern of unpredictability — should we talk about data points individually.

Whenever a process is predictable, the next obvious question is what we forecast for the future, based on the data we have. One approach to making such a prediction is to compile all the common-cause data into a dot plot. Since the process is considered predictable, the data from previous periods can be thought of as a random sampling of performance in the future.

Now we insert the specification limits — the data range that management deems acceptable for this process. A “pass/fail” dot plot like that in exhibit 4, below, highlights the number of data points that fail to conform with the process's specification limits. The 12 percent “defective” rate (i.e., six out of 50 data points) jumps out in exhibit 4. For situations like this, creating a prediction by dividing instances of nonconformance from the total data population is not the best approach. If close-to-specification data points had a slightly higher or lower response, they could have toggled from a conformance to a nonconformance state, or visa versa. And where they end up in the particular cases included in this report directly impacts the predicted failure rate.

A better approach to reporting on this information would be to create a probability plot of all the data from this control chart, as shown in exhibit 5 on page 22. In this display format, we can see that the process has a nonconformance rate of approximately 13 percent ($0.629\% + [100\% - 87.768\%] = 12.861\%$). Because we've already determined through the control chart that any nonconformance is common-cause, we can expect that nonconformance rate to continue in the future unless something either negatively or positively impacts the overall process. If this rate of nonconformance is unsatisfactory, the company needs to modify the basic process to reduce the variability of its outputs. Displaying the results in this probability-plot format shows executives that a process-improvement project, and not one-off firefighting efforts, is the way to reduce the frequency of nonconformance in this process.

Cultural Change Through Reporting

People often talk about wanting to make a cultural change in their organization. Much money is spent on workshops and consulting to address cultural change issues — and the results are frequently questionable. In contrast, changing the way managers and employees examine and react to performance data can be the impetus for real, long-lasting positive cultural change.

Again consider exhibit 2. If we had drawn horizontal lines at the specification limits and then reacted to each of the individual data points that landed beyond those specification limits, we would be reacting to common cause as though it were special cause. In a culture that responds this way to performance data, one person may receive an award for a long string of outputs that conform to production specifications, while another employee is reprimanded for missing specs a few times in his production work. This can be highly demotivating when the performance differences between the two employees result from random, common-cause variability in the company's processes. It can also lead to poor resource-allocation choices.

If a process is predictable, we need to look at all the data collectively. It's not valuable to try to explain why a single data point is up and another is down.

Exhibit 4

Dot Plot Determines Proportion of Data out of Range

Because our control chart from exhibit 3 shows variability to be common-cause, we can use a simple dot plot to describe the capability of our process. As this dot plot shows, a simple count of the number of responses outside the specification limits of 72 and 78 is six out of 50, which translates to a 12 percent nonconformance rate. However, attribute estimates are not as desirable as continuous response estimates — because, for example, if the two responses with the value 78 were slightly higher, our nonconformance percentage would be considerably larger.




Exhibit 5

Probability Plot of Process Common-Cause Variability Indicates Nonconformance Rate

Plotting performance data on a process-percentage probability plot that is not bounded by calendar month or year shows clearly the proportion of data that falls outside the specification limits.



Consider the example of an organization that monitors the frequency of accidents. One year the number of accidents causing injury increased during the month of July to 16, up from 14 in July the year before. The company issued a safety memo declaring the increase unacceptable and telling all employees they must watch a 30-minute safety video within the next month. At an average wage of \$10 per hour, payroll for the company's 1,500 employees increased by \$7,500 for August, not including time spent getting to and from the conference room where the video was shown or time spent issuing memos reminding people to attend, reviewing attendance rosters looking for no-shows, etc.

Safety is, of course, very important. But this company reacted to the increased accident rate without considering first whether the problem was with the process itself. As Deming estimated, 94 percent of nonconforming results are the direct result of the systems management has put in place. The vast majority of instances of poor performance require that the system itself be modified for improvements to occur — and those modifications require an organization to look at its systems collectively, relative to specification limits, over a long period of time. Reacting to a spike in accidents in an individual month can be counterproductive and expensive.

Organizational performance reporting using individuals control charting (i.e., plotting individual data points rather than mean data), as illustrated in exhibit 3, enables management to view a process in the same detail an airplane would when viewing the earth from flight. We do not end up reacting to small up-and-down terrain differences when flying an airplane. Similarly, in business we should not react to every small instance of lot-to-lot variability caused by differences among people and raw materials. This view from the 30,000-foot level is a no-nonsense methodology that helps organizations distinguish between common- and special-cause process variation so that they can reduce the frustration and expense of firefighting activities.

Cause — and Effect

Companies do face some performance shortfalls that result from special causes. If an individual problem is determined to be special-cause from a process point of view, the company needs to address whatever was different in that event to cause the nonconforming result (e.g., why was our customer response time yesterday exceptionally large relative to typical common-cause variability?).

Generally speaking, though, organizations can achieve more gains by continuously working to mitigate common-cause problems by improving their basic processes. Effective, long-lasting improvements to processes are not made by firefighting. They require the examination of process data over a period of stability to determine what should be done differently in the long term.

Presenting performance data in traditional management reports, with simple year-to-year comparisons of metrics, may identify results that are out of line with targets, but it does little to help executives determine how to respond to those results. How can a company fix poor performance when it doesn't know what caused that performance? Process improvement projects in Six Sigma utilize a define-measure-analyze-improve-control (DMAIC) road map to investigate the causes behind nonconforming processes using both statistical and nonstatistical techniques. Such an analysis can lead to long-lasting, sustainable improvement, and taking an S⁴/IEE approach to reporting on the analysis expands the positive impact that companies see in their top-level performance metrics. ^Bm