Automating the detection of Turning Points: Inventory control at ComputerShop

Authors

Dr. Roger Jenkins

University of Technology, Sydney

Email: roger.jenkins@uts.edu.au

Geoff Breach

School of Management, University of Technology, Sydney

Email: geoff.breach@uts.edu.au

Abstract

Inventory control for a product catalogue of 3000 products is carried out by two managers at ComputerShop. While there is a substantial level of automation of product flow in this company, there is no analysis of inventory levels, nor of trends in demand for each product. Inventory management thus is an area that imposes a high workload on the managers and is characterized by the usual problem of inventory and demand being poorly balanced. In this paper we have studied two techniques that can be applied to detect turning points in a sequence of sales data. We have demonstrated that both of these techniques can be used to support the partial automation of decisions on inventory control.

Keywords

CUSUM, Shiryayev-Roberts , inventory.

Introduction

The ComputerShop web site describes itself as a supplier of Personal Computer (PC) hardware components to PC hardware enthusiasts. A typical ComputerShop customer is an individual who is employed as a computer technician or a person who is otherwise sufficiently skilled to be able to assemble their own PC from individual components. Our informants noted that annual turnover for ComputerShop was about $12 Million, and had been flat for some time after a period of substantial growth. Profit margin on most sales is in the order of 5% to 8%. The PC component market is very competitive, and margins are slim throughout the industry.

There are twelve people employed in the business, they operate out of a light industrial estate, have a combined warehouse and office facility, and have an overflow warehouse located elsewhere. The facility is small and fully occupied by either open plan office space or storage and goods processing areas. There is no visible public reception area with any sense of making a presentation of the company to outside agents. The general appearance of the facility is that of a business that does not expect to impress external agents with it physical appearance. This company operates via a web interface and it does not encourage a direct contact with its customers

In order to prosper in such a market place we believe that this company must have solved some problems in interesting ways. We accordingly have chosen to study the company in some detail. This paper explores one aspect of our study; the problem of managing stock levels when demand is highly variable and the cost of holding stock is very high. It should be noted that the IT system and product control activities are completely integrated. The IT system and product flow processes are seamless at ComputerShop.

In our first two meetings the problems of managing stock, and the range of products held were frequently mentioned as important issues. Company managers pointed out that ComputerShop lists a product range of between 3000 and 4000 items for sale on its web site at any given time. Stock is typically turned at least twice per month, and many items are rotated more often than that. Very few, if any, stock items stay in the ComputerShop inventory for more than twelve months. Stock control is the responsibility of one person. Stock on hand can be a very expensive luxury in this business. Across many industries the penalty for holding stock is seen as an asset problem. Supermarkets for example will see assets bound up with stock as simply related to the degree to which they can make their capital productive. ComputerShop have a very different problem. Individual lines typically have very short lives. A specific model of central processing unit (CPU) will typically have a product life of four to five months before it is superseded by a new model that is usually faster and cheaper than its predecessor. CD and DVD burners (known collectively as 'optical drives') are products that have an even shorter life cycle - a new model will supersede any given example as often as once per month.

A hard disk drive is a typical example of how technological development affected the supply chain. For some product lines ComputerShop would start to provide price reductions if the stock was held for longer than one week. One response in the industry was for a ‘drop ship’ strategy to be adopted by some retailers. In this strategy the retailer held no stock at all, a customer order triggered an order to the wholesaler for a single unit, with the customers address as destination. This obviously is not providing a sustainable advantage for retailers operating in that mode. Low inventory is of course the simplest response, and with this response is the potential for stock-outs and consequent loss of sales. One conclusion drawn from our second visit to ComputerShop was that the area of inventory management required some analysis. This area is the main basis of this report.

The core problems for ComputerShop are simple, when to reorder and how much to reorder. There are many strategies that can enable efficient operations when there is ample data on demand. The problem that we are reviewing is different to this. We are motivated by the need to simplify the problem of detecting a turning point in demand. We wish to examine some strategies that can automate the task of informing the manager that there has been a significant change in demand for a product. In order to do this we have reviewed three methods that can be used to detect turning points. Two of these methods are evaluated using one method, the Shewhart test as, a benchmark. We have then examined how these methods would perform on three different types of ComputerShop products.

Methodology

We have framed this problem similarly to the one of detecting a significant change in a process mean. We consider the rate of sales per week to be a random variable with a normal distribution drawn from a population with a mean μ, and standard deviation σ. We will explore the capability of three tests to determine a 1σ step change in the mean. The sensitivity of the test will be measured by the Average Run Length (ARL). The ARL of the test is the average number of samples following a change that is required to trigger the alarm. This is a generally used approach to measuring the performance of the CUSUM test (Montgomery 1996; Arnold and Reynolds 2001) A good test will have a low value of the ARL for a process that is out of control, but a high value for a process that is in control. Desirable values of the ARL for processes in control are often implicitly set at levels achieved using a 3σ Shewhart control chart and these values for a two sided test are 370, or 740 for the equivalent, one sided limit (Montgomery 1996). Montgomery (1996) also noted that a properly parameterized CUSUM chart will detect a 1σ shift in the mean with an ARL of 8.38 while the Shewhart test would have an ARL of 43.96.

It is possible to determine ARLs for the Shewhart and CUSUM tests analytically (Montgomery 1996), and Markov chain analysis based on the work of Brook and Evans (1972) is widely used to model the behaviour of the CUSUM test for a number of distributions. Authors will however often use Monte-Carlo simulation, particularly if the distribution of the variable is likely to be poorly described by either the Poisson or Normal distribution (Atienza Tang and Ang 2000; Chang and Fricker 1999: Khoo 2005: Koning and Does 2000, Sparks 2004). Grigg and Farewell (2004) recommend the use of simulation if there is some doubt about the validity of distributions used to model the behaviour of the process variable. In this project we expect that we will need to apply the technique in time to data that will not be adequately described by the Normal or Poisson distributing and so we will rely exclusively on Monte-Carlo simulation for ARL estimates. Our simulations have been performed using a VBA driven Microsoft Excel spreadsheet, using wherever possible the standard Microsoft Excel functions for sampling from distributions and the generation of a random number series.

Once we have characterized the performance of the three tests on synthetic timeseries we will then apply the tests to a set of timeseries data drawn from actual sales history in the company. The three techniques that will be discussed are; the standard Shewhart test, where an alarm is set when a sample value is equal to or greater than 3σ from the process mean, a CUSUM test and a Shiryayev-Roberts test. We discuss the Shewhart test briefly but provide a little more detail on the other two tests.

Shewhart test

This test emulates the process of plotting the data on a Shewhart control chart (for example as described by Montgomery, 1996) and asserting an alarm when the sample mean reaches the nominated control limit. In this work the Shewhart test is evaluated as a benchmark. It is generally accepted that the Shewhart test is not sensitive to small changes in the process mean (Montgomery, 1996) and the CUSUM and Shiryayev-Roberts tests have been proposed as more sensitive alternatives in the quality literature (see for example Montgomery, 1996; Ergashev 2004, Kenett and Pollak 1996). We evaluate the Shewhart test in its simplest form, which is the alarm is asserted only when the process mean reaches the control limit. We do not attempt to utilize information related to the sequence of data. This is consistent with our intent of using this test as a means of setting some expectations for sensitivity on the other two tests reported here.

CUSUM test

There are two forms of the CUSUM test; the tabular or the V-mask form. The tabular is preferable in the view of Montgomery (1996) and it is the form used for this project. The description outlined below follows the approach used in Montgomery (1996)

The CUSUM chart plots the evolution of two variables (C^+ and C^-) formed from results taken from a sequence of samples.

$$C_i^+ = \max[0, x_i - (\mu_0 + K) + C_{i-1}^+]$$
$$C_i^- = \max[0, (\mu_0 - K) - x_i + C_{i-1}^-]$$
$$\text{where the starting values for } C_i^+ = C_i^- = 0$$

These equations define the form that uses raw data from the process results, and where the values of K and μ_0 are expressed in process units. K is usually called the reference value (or allowance or slack value) and is often chosen to be about halfway between the target value (μ_0) and the out-of-control mean that we wish to detect quickly. It sets a window of indifference in the test such that process values that differ from the target value by less than K do not contribute to an expansion of the CUSUM value. Large values of K lead to very unresponsive CUSUM tests.

If either C^+ or C^- exceeds a decision interval *H* then the process is deemed to be out-of-control. In the context of forecasting then we will deem the level of demand to have shifted by an interesting amount. High values of *H* will also lead to unresponsive tests; low values will lead to false alarms

In our work we only operate on standardized data, and so the following equivalent forms of the equations are applicable:

where

$$z_i = \frac{x_i - \mu_0}{\sigma}$$

$$C_i^+ = \max[0, z_i - (\mu_0 + k) + C_{i-1}^+] \quad (1)$$

$$C_i^- = \max[0, (\mu_0 - k) - z_i + C_{i-1}^-] \quad (2)$$

where the starting values for $C_i^+ = C_i^- = 0$

Montgomery considers that setting h at a value of 4 to 5, and k at a value of 0.5 will generally give a CUSUM with good properties when testing for a 1σ shift in the process mean.

Shiryayev-Roberts test

The formulation of this test is based on the formulation and notation used by Ergashev (2004). That author can be referred to for a more complete history of the test.

In this test we compute the value of a statistic *R*, where *R* responds to the value of a process variable X_i and recursively to the last calculated value of *R*. This is analogous to the technique of an exponentially smoothed forecasting strategy. The calculation of *R* is modulated via a parameter m, and this value controls the sensitivity of the test. A decision interval variable *B* is set such that when R_i exceeds *B* an alarm is asserted.

Values of R are calculated using the relationship set out in equation (3):

$$R_n = (R_{n-1} + 1) \cdot \exp\left\{ mX_n - \frac{m^2}{2} \right\} \quad (3)$$

This relationship will give values of *R* that increase as the value of *X* increases. As X_i takes on negative values, when we have values less than the target or average value then the value of *R* becomes very small, and becomes an ineffective indicator of out of control values. In response we also compute the Shiryayev-Roberts value for the negative value of X_i. This allows us to track values less than the mean with the same sensitivity as those above the mean. This is shown in equation (4)

$$R_n = (R_{n-1} + 1) \cdot \exp\left\{ - mX_n - \frac{m^2}{2} \right\} \qquad (4)$$

When this test is used for a continuous stream of date the value of the decision interval (B) can be chosen to be equal to the desired ARL. For data in discrete time, this equality does not hold and approximations were used by Ergashev (2004) to set the value of B. For the work reported in our paper we have selected a value of B based on the results of our Monte-Carlo simulations presented in Table 3.

Results

Shewhart test

Results of a Monte Carlo simulation of 30,000 cycles.

In Table 1 the results of a simulation of a process using the Shewhart test to detect a 1σ shift in the process mean. The table outlines the ARL found for tests at the conventional level of 3σ, but also smaller values of the test. These values correspond to the upper control limit on a normal Shewhart control chart for the process mean. The table contains results for a one sided testing regime, so results for alarms on the low side are not reported.

Table 1 ARL for Shewhart test

	In control. Normal distribution μ=0, σ=1	Shift the mean by 1σ Normal distribution μ=1, σ=1
Limit	ARL	ARL
1.0σ	6	2
1.5σ	14	3
2.0σ	45	6
2.5σ	156	15
3.0σ	679	43

On the basis of this tabulation we can observe that a typical test where we generate an alarm with an upper control limit at 3σ then this will give us an ARL on the high side of 679 when the process is in control. For this process we observe a high side alarm with an ARL of 43 for a nominated 1 σ shift in the process mean. These are the results in the shaded cells. These results are able to be developed from a simple use of the normal probability distribution tables, they are however include to provide a basis for comparison with later results that are based on the same type of Monte-Carlo simulation methodology.

CUSUM test

Table 2 presents the results for a Monte-Carlo test for 30,000 cycles.

In this test we need to make a choice of both h and k. Montgomery (1996) suggests that a value of k = 0.5 (where K = kσ) when the process mean is expected to experience a shift of about 1σ. We have used this value for k in our simulation. Values of h of between 4 and 5 will give similar levels of performance under the base case as that noted for the Shewhart test.

Table 2 ARL for CUSUM test

H	ARL μ= 0	ARL μ= 1
1	22	3
2	67	5
3	190	7
4	498	9
5	1139	11
6	4269	13
7	6782	15

It is evident from the table that this will be achieved with substantially higher sensitivity to the shift in the process mean. Whereas the Shewhart test will signal an alarm after, on average, 43 samples, the CUSUM test will signal an alarm after about 10 samples.

Shiryayev-Roberts test

This table presents the results for a Monte-Carlo test for 30,000 cycles.

Table 3 ARL for Shiryayev-Roberts test

B	ARL μ = 0	ARL μ= 1

50	95	6
100	160	7
150	275	8
200	378	9
250	535	9
300	571	9
350	704	10
400	649	10
450	723	10
500	826	11
550	1033	11
600	797	11

In this test we have chosen a value of m=1, this is equal to the hypothesised process shift, a strategy noted by Ergashev (2004) as providing optimal performance for both the Shiryayev-Roberts and CUSUM tests. The results are indistinguishable from that of the CUSUM, and similarly, quite superior to that found for the Shewhart test. Results for the three tests are summarized in graphical form below.

Summary of results for three tests

Values were tabulated for each of the three tests for specific values of the key parameters of k and m. Values chosen for these parameters (k = 0.5 and m = σ = 1) were based on those used by Montgomery (1996) for k and Ergashev (2004) for m. It is evident from this chart that the CUSUM and Shiryayev-Roberts tests give quite similar levels of sensitivity for a similar shift in the process mean. For example if we choose settings for the parameters of each test as specified above, then for those settings that give say an ARL of 400 when the process is in control, we can expect an ARL of about 9 for the CUSUM and Shiryayev-Roberts tests when the process mean has undergone a 1σ shift. The Shewhart test would give an ARL of about 32 under these conditions.

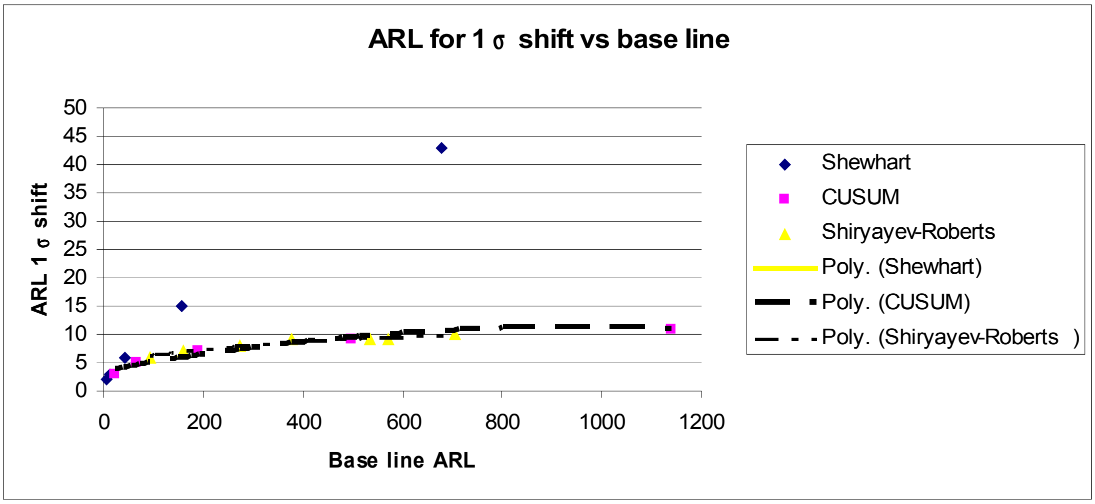

Figure 1 Impact of 1 σ shift of mean on ARL

Analysis of three ComputerShop products

In this section we exhibit the results of applying the previous three tests to three products that were selected in order to illustrate products with a different life cycle. We review firstly a product with a long life but at the end of its cycle, secondly a product with a very short life cycle, and thirdly a product at the mature phase of a relatively long life cycle. These three products illustrate a representative range of products for this company. A glossary of the notation used in the plots and tables in outlined in .

Table 4 Glossary of terms in test results tabulations

Sales	Number of units sold in this week
z	Standardized sales data
Shi	CUSUM on positive side –sales growth
Slo	CUSUM on negative side – sales decline
SR +	Shiryayev-Roberts test on positive side – sales growth
SR -	Shiryayev-Roberts test on negative side – sales decline
	Shaded cell when alarmed

Mature product: end phase

This product had a long run of significant sales up to week 64 at which time it had no further sales. The period between weeks 15 – 45 was selected as one that could be used to set management expectations for this product. Statistics for this period were

Target value	9.5
StDev	4.9
Skew	0.8

These statistics indicate that the values could reasonably be expected to have been drawn from a population with a Normal distribution, the coefficient of variation is however high. These values are used to set the target value for the standardization of the observed variable in both the CUSUM and the Shiryayev-Roberts test. Once set they are not changed over the life of the test.

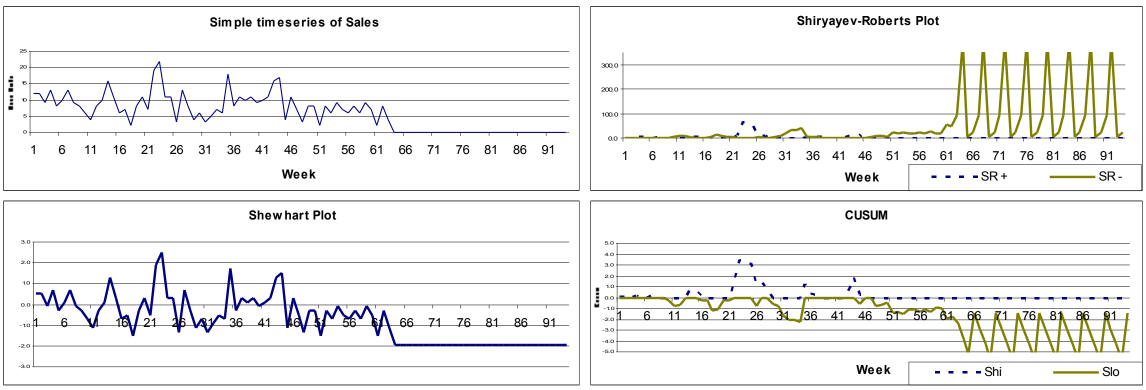

Figure 2 Test plots for Mature product: end of life

The period of interest for this product is between weeks 55 – 65. This was a period of declining sales that preceded a final sale in week 63. The Shewhart plot gave no alarms for this timeseries. The end of sales was not signalled because the target level of sales was not three times the standard deviation during the benchmark period.

The Shiryayev-Roberts plot gave an alarm in week 64, in the same week as the last sale. The CUSUM gave an alarm in week 65, one week later.

There were no false alarms during this period. The solid line plot is for the negative CUSUM and Shiryayev-Roberts values, the positive values are plotted with the dotted line.

Table 5 Test results for a mature product: End phase

Week	55	56	57	58	59	60	61	62	63	64	65
Sales	7	6	8	6	9	7	2	8	4	0	0
z	-0.5	-0.7	-0.3	-0.7	-0.1	-0.5	-1.5	-0.3	-1.1	-1.9	-1.9
Shi	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slo	-1.1	-1.3	-1.1	-1.3	-0.9	-1.0	-2.0	-1.8	-2.4	-3.8	0.0
SR +	1	0	1	0	1	1	0	1	0	0	0
SR -	18	24	21	27	19	20	58	49	92	390	0

Burst: full life cycle

This product experienced a period of intense sales between weeks 17 – 37 and then no further activity. This is a very typical product for this company.

The period between weeks 19 – 35 was selected as one that could be used to set management expectations for this product. Statistics for this period were:

Target value	10
StDev	3.9
Skew	0.6

These statistics indicate that the values could reasonably be expected to have been drawn from a population with a Normal distribution, the coefficient of variation is however high. The period of interest for this series is the peak in sales in week 34 followed by a rapid decline to zero sales by week 38.

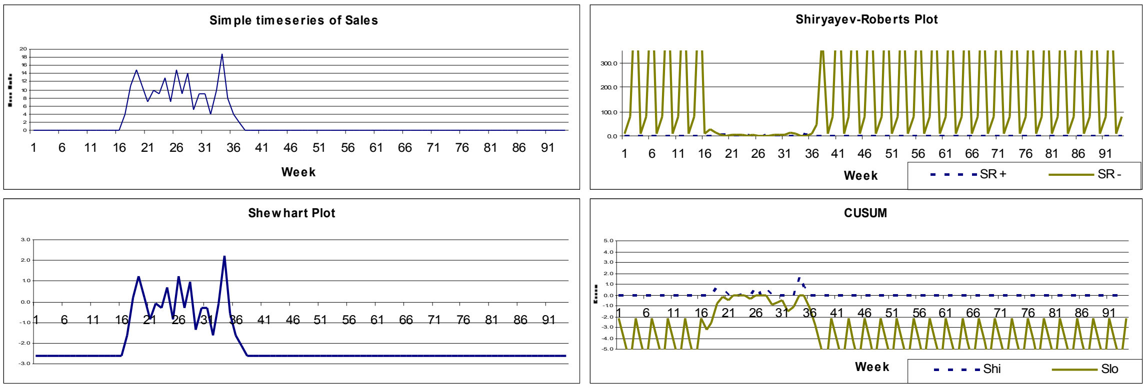

Figure 3 Test plots for Burst product: Full cycle

Again, the target value was less than three time standard deviation of demand during the benchmark period and so the Shewhart test did not detect the end period. There were no false alarms from this test. The Shiryayev-Roberts test showed no false alarms during the period and gave an alarm at 38 having accumulated a negative signal during the final decline. The CUSUM gave an alarm one week later.

Table 6 Test results for a Burst product: Full cycle

Week	29	30	31	32	33	34	35	36	37	38	39
Sales	5	9	9	4	10	19	8	4	2	0	0
z	-1.3	-0.3	-0.3	-1.6	-0.1	2.2	-0.6	-1.6	-2.1	-2.6	-2.6
Shi	0.0	0.0	0.0	0.0	0.0	1.7	0.7	0.0	0.0	0.0	0.0
Slo	-0.8	-0.7	-0.5	-1.6	-1.1	0.0	-0.1	-1.2	-2.8	-4.9	0.0
SR +	1	1	1	0	1	10	4	1	0	0	0
SR -	4	4	4	15	10	1	2	9	48	412	0

Mature product: mature phase

This product was introduced in week 20 and had experienced relatively steady sales over the period of the study. There is some suggestion of a slight decline in sales in the most recent periods, but sales could be still quite healthy. The period between weeks 24 – 33 was selected as one that could be used to set management expectations for this product. Statistics for this period were

Target value	35
StDev	4.7
Skew	0.8

These statistics indicate that the values could reasonably be expected to have been drawn from a population with a Normal distribution; the coefficient of variation is much lower than for the previous two samples. The product sold fairly steadily between weeks 20 – 95. The product was still active at the time this data was collected.

The Shewhart test has produced a number of signals during this period. Given that sales still continue, at about the benchmark level, we can deduce that all of these signals were false. There were three positive signals (at weeks 50, 66, and 68) and five negative alarms (at weeks 49, 55, 69, 72, and 93).

The Shiryayev-Roberts test has given positive alarms at week 68, and negative alarms as weeks 50 and 69. The CUSUM test has given a positive alarm at week 68 and negative alarms at week 50 and 69.

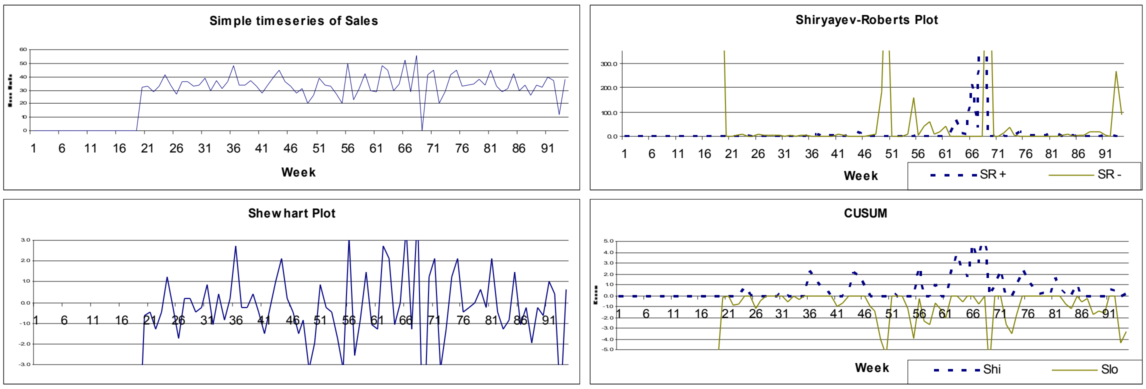

Figure 4 Test plots for Mature product: Mature phase

Inspection of the data suggests that the signals at week 68 and 69 are an artefact. Sales in week 68 clearly contain sales for 69 as well. The low sale alarm in week 50 is a response to a sequence of weeks with lower than average sales, but is a false alarm. Sales in week 51 are above average and this is followed a period of about 20 weeks of increasing sales.

Table 7 Test results for a Mature product: Mature phase

Week	48	49	50	51	64	65	66	67	68	69	70	92	93	94
Sales	31	20	26	39	30	35	52	29	56	0	41	37	12	38
z	-0.9	-3.2	-1.9	0.8	-1.1	0.0	3.6	-1.3	4.4	-7.4	1.2	0.4	-4.9	0.6
Shi	0.0	0.0	0.0	0.3	2.2	1.7	4.8	3.0	6.9	0.0	0.7	0.4	0.0	0.1
Slo	-1.4	-4.1	-5.5	0.0	-0.6	-0.1	0.0	-0.8	0.0	-6.9	0.0	0.0	-4.4	-3.3
SR +	1	0	0	2	14	9	213	36	1857	0	2	3	0	1
SR -	11	183	768	0	2	2	0	2	0	1049	0	2	264	87

Discussion and conclusions

Shiryayev-Roberts and CUSUM performed better than Shewhart. In all cases for this data set the Shewhart test failed to provide an alarm at the end of life of sales. This is not surprising given that the Shewhart test (as applied in this project) does not retain any memory of previous results. Normally the Shewhart test has a range of rules that reflect pattering in the data. The test was used in its simple form here as a benchmark. The more interesting comparison is between the performance of the CUSUM and Shiryayev-Roberts tests. Monte-Carlo simulation suggests that there should be very little difference between these two tests. For the data used here however the Shiryayev-Roberts test has given marginally superior performance to the CUSUM test. The Shiryayev-Roberts test gave a clear alarm for the end of life of the first and second products reviewed in this paper. This alarm was asserted in the same week as sales hit zero in both cases for the Shiryayev-Roberts test. The CUSUM test produced an alarm one week later in both cases.

These results support the claim by Ergashev (2004) and Kennett and Pollak (1996) that the Shiryayev-Roberts test should be considered as a useful additional approach for the detection of small shifts in a process mean. These results also support the view that we can consider the test useful in conditions where we wish to track the performance of product sales. The particular attraction of both of these techniques is that they can be implemented as a simple quantitative analysis of demand data. It is relatively easy for a manager to look at a plot of demand and draw conclusions on the state of sales for a product. This however is a difficult task when there are over 4000 SKUs that need to be reviewed! Either of these tests can provide the manager with a short list of candidates that appear to have experienced a shift in the process mean. It should be relatively easy for the manager to draw conclusions once attention has been directed at a particular SKU.

While the results support the claim that these tests are useful, it would be difficult to make a claim that one test was superior to the other based on these results. Both tests are easy to program, parameters for both tests can be readily optimized if the proposed process shift can be established. This paper has not examined the interaction of the test and its performance for unspecified shifts in the process mean and this may change the sensitivity of the tests differently. This will be the subject of further work during the course of this project as it is not possible to specify the shift in the process mean in this project.

References

Arnold, J.C. and Reynolds, Jr.M.R. 2001, "Cusum control charts with variable sample sizes and sampling intervals", *Journal of Quality Technology,* vol. **33**, no. 1, pp. 66-81.

Atienza, O.O., Tang, L.C. and Ang, B.W. 2000, "A uniformly most powerful cumulative sum scheme based on symmetry", *The Statistician,* vol. **49**, no. 2, pp. 209-217.

Brook, D. and Evans , D.A. 1972, "An approach to the probability distribution of CUSUM run length", *Biometrika,* vol. **59**, no. 3, pp. 539-549.

Chang, J.T. and Fricker, Jr.R.D. 1999, "Detecting when a monotonically increasing mean has crossed a threshold", *Journal of Quality Technology,* vol. **31**, no. 2, pp. 217-234.

Ergashev, B. 2004, "Sequential detection of US business cycle turning points: Performances of Shiryayev-Roberts, CUSUM and EWMA procedures", *Econometrics from EconWPA.* Viewed 1 July 2006, http://econpapers.repec.org/paper/wpawuwpem/0402001.htm

Grigg, O. and Farewell, V. 2004, "An overview of risk-adjusted charts", *Journal of the Royal Statistical Society,* vol. **167**, no. Part 3, pp. 523-539.

Kennett, R.S. and Pollak, M. 1996, "Data-analytic aspects of the Shiryayev-Roberts control chart: surveillance of a non-homogenous Poisson process", *Journal of Applied Statistics,* vol. **23**, no. 1, pp. 125-137.

Khoo, M.B.C. 2005, "Determining the time of a permanent shift in the process mean of CUSUM control charts", *Quality Engineering,* vol. **17**, pp. 87-93.

Koning, A.J. and Does, R.J.M. 2000, "CUSUM charts for preliminary analysis of individual observations", *Journal of Quality Technology,* vol. **32**, no. 2, pp. 122-132.

Montgomery, D.C. 1996, *Introduction to statistical quality control 3rd ed.*, John Wiley and Sons, Brisbane.

Sparks, R. 2004, "Weighted moving averages: an efficient plan for monitoring specific location shifts", *International Journal of Production Research,* vol. **42**, no. 12, pp. 2521-2528.