

CONFIRMED
MAR 1992

British Standard Guide to
**Data analysis and quality control
using cusum techniques**

Part 4. Cusums for counted/attributes data

Guide pour l'analyse des données et le contrôle de la qualité par les techniques cumulatives
Partie 4. Sommes cumulées des données de comptage (attributs)

Leitfaden für die Datenanalyse und Qualitätslenkung mit Cusum-Verfahren
Teil 4. Cusum-Verfahren für zählbare und Attribut-Merkmale

Contents

	Page		Page
Foreword	1	14. ARL data for schemes with $K = 0.1$, Poisson variable	23
Cooperating organizations	Back cover	15. ARL data for schemes with $K = 0.25$, Poisson variable	23
Guide		16. Critical values of D for testing binomial or Poisson distribution assumptions	25
0. Introduction	2	17. Observed and expected frequencies for dustbin data	26
0.1 Basis of cusum chart	2	18. Observed and expected frequencies for fabric fault data	27
0.2 Simple example of cusum chart	2	19. Cusum calculations for dustbin data	30
0.3 Monitoring of retrospective analysis	3	20. Cusum calculations for fabric fault data	30
0.4 Scaling the chart	4	21. Cusum calculations for rearranged data of 8.6 (and table 20)	31
1. Scope	5	22. Comparison of alternative schemes for $P_1 = 0.005$	32
2. Terminology and symbols	5	Figures	
3. Decision rules: general	6	1. Conventional chart of data from table 1	3
4. Decision rules for counted data	9	2. Cusum chart of data from table 1	3
5. Cusum procedures for Poisson variables	10	3. Cusum chart of data from table 1, with target value 12	4
6. Cusum procedures for binomial variables	14	4. Cusum chart of data from table 1, with target value 15 but compressed cusum scale	4
7. Cusums for low rates of occurrence of events	20	5. Form and simple construction for cusum mask for counted data	7
8. Testing the validity of Poisson or binomial distribution assumptions	24	6. Cusum and conventional charts of Poisson data	8
9. Examples of the application of cusum methods to counted data	27	7. Cusum and conventional charts of binomial data	8
Appendix		8. Cusum chart for data of table 2, using $H = 8$ and $K = 6$	14
A. Basic decision rules and practical points about cusum masks	33	9. Cusum chart for data of table 3, using $H = 3$ and $K = 1.5$	14
Tables		10. Cusum charts for data of tables 20 and 21	28
1. Data for cusum plotting	2	11. Conventional charts for data of tables 20 and 21	28
2. Cusum data for Poisson variable with $m_a = 4$	7	12. Cusum chart for data of table 19	29
3. Cusum data for binomial variable $n = 20$, $\dot{p} = 0.03$	9	13. General-purpose truncated V-mask	33
4. Comparison of ARL curves for three discrete distributions	11	14. Truncated V-mask applied to cusum chart: no indication of shift	34
5. C1 and C2 schemes for Poisson variables	11	15. Truncated V-mask applied to cusum chart: indication of shift	34
6. ARL characteristics of cusum schemes for Poisson variables	12		
7. Cusum tabulation for data of table 2 using $H = 8$ and $K = 6$	13		
8. Cusum tabulation for data of table 2 using $H = 3$ and $K = 1.5$	13		
9. C1 and C2 schemes for binomial variables	16		
10. ARL data for schemes for binomial variables	17		
11. ARL's for schemes with $H = 1$, binary events	21		
12. Examples of schemes for weapons trial application (7.4)	21		
13. ARL's for schemes with $H = 1$, Poisson variables	22		

Compliance with a British Standard does not of itself confer immunity from legal obligations.

Foreword

This British Standard was prepared under the direction of the Quality Management and Statistics Standards Committee. It was originally intended that this standard should form part of an overall revision of British Standards relating to statistical quality control techniques. However, cusum methods have found applications in fields beyond that of quality control, and comments on the draft standard revealed interest in a more general document. Part 1 of the standard represents an attempt to meet this demand, and is intended as an introduction to the principles of cusum procedures, and charting.

Part 2 details decision rules, significance tests for retrospective analysis of accumulated data, and contains a variety of applications of cusum charting, some concerned mainly with data presentation, some with monitoring in quality or budgetary control, and others with retrospective analysis.

Part 3 deals with applications of cusum methods to process and quality control by measurement. The tech-

niques cover control of both average level and variability, and may be implemented using charts, tabulations or computers. They are especially appropriate when location of change points, and estimation of the corrections required to bring the process on to target, are desirable features of the control system. Part 4 deals with control of counts of defectives and general applications involving data based on counting rather than measurement.

Acknowledgement is made to the authors of papers in the statistical journals whose work has been drawn upon, to the contributors of examples of applications, to Mr P L Goldsmith for useful discussions, to ICI Fibres (Pontypool) for provision of computing facilities, to Mr D A Budgett for assistance with the mathematics of clause 7 of this Part, and to Dr Derek Bissell (a representative of the Institute of Statisticians) who was primarily responsible for the development work on this standard, carried out on behalf of BSI Committee QMS/1/4.

British Standard Guide to

Data analysis and quality control using cusum techniques

Part 4. Cusums for counted/attributes data

0. Introduction

0.1 Basis of cusum chart

0.1.1 To avoid repetition of material in the earlier Parts, it is necessary to assume that the reader is familiar with the basis of cusum charting as presented in Part 1, and with the concepts underlying decision rules as set out in Part 2. However, for convenience a brief recapitulation of the essential features of cusum charts and decision rules is given.

0.1.2 The cusum chart is a highly informative means of graphical presentation of data which can be ordered into a logical sequence. For process or quality control purposes, this sequence will generally correspond to the order of observation on a time scale (but examples of other bases for ordering are given in Part 2).

The essential principle is that a *target* value, T , is subtracted from each observation. This target value is usually a constant (examples where the target value is not constant are given in Part 2). In this Part, the target value(s) will be related to a proportion of binary events (e.g. defectives/non-defectives, successes/failures) or to a level of events (defects, faults, breakdowns) per unit of observation.

The cumulative sum of the deviations from target is formed, and this *cusum*, C , is plotted against the serial number of the observations.

0.1.3 Unlike conventional plotting (where the average level of the values in the series corresponds to height above the serial number axis), the cusum method of plotting results in the representation of average by the local *slope* of the chart. When the average corresponds to the target value, the path of the cusum lies roughly parallel to the sequence axis. When the local average of the series is greater than the target value, the cusum slopes upwards (away from the sequence axis); conversely, when the local average is negative, the cusum slopes downwards. The greater the discrepancy between the local average and the target value, the steeper the slope of the cusum path.

0.1.4 The result of plotting in the cusum mode is that changes in average level over different subdivisions of the sequence of observations are clearly indicated by changes in slope of the chart. The local averages in such subdivisions can be readily calculated, either from the numerical values of the cusum (from which the chart is plotted) or directly, albeit with some small loss in accuracy in some cases, from the chart itself.

0.2 Simple example of cusum chart

0.2.1 The above principles are best appreciated from a simple example. The calculations and plotting procedure will, at this point, be developed without mathematical

symbolism; appropriate algebraic formulae are presented later.

It is supposed that the data given in table 1 have been obtained, over a time sequence in the order shown, and that a target value of 15 is appropriate.

Table 1. Data for cusum plotting

Observation number	Observed value	Deviation from target (15)	Cumulative deviations
1	12	-3	- 3
2	17	+2	- 1
3	14	-1	- 2
4	14	-1	- 3
5	17	+2	- 1
6	16	+1	0
7	14	-1	- 1
8	11	-4	- 5
9	13	-2	- 7
10	14	-1	- 8
11	15	0	- 8
12	11	-4	-12
13	14	-1	-13
14	16	+1	-12
15	13	-2	-14
16	14	-1	-15
17	11	-4	-19
18	12	-3	-22
19	13	-2	-24
20	16	+1	-23
21	12	-3	-26
22	18	+3	-23
23	18	+3	-20
24	17	+2	-18
25	20	+5	-13
26	15	0	-13
27	14	-1	-14
28	18	+3	-11
29	20	+5	- 6
30	16	+1	- 5
31	18	+3	- 2
32	14	-1	- 3
33	16	+1	- 2

For a conventional chart, as in figure 1, the observed values are plotted against their corresponding observation numbers. There is some indication that the last dozen values appear to be clustered around a different mean level from the first 20 or so.

Plotting in the cusum mode gives a much clearer display. The cusum (column 4 of table 1) is plotted against the observation number on conventional axes, using the y

('vertical') axis for the cusum and the x ('horizontal') axis for the observation number, as shown in figure 2.

0.2.2 The chart clearly separates into three segments. From observations numbered 1 to 7 (inclusive) the cusum path is generally parallel to the observation number axis, i.e., the path is, in conventional terms, roughly horizontal. From observations 8 to 21 inclusive the path is recognizably downward (despite local irregularities such as at observations 14, 20). From observations 22 to 33, the path is upward (again with local irregularities which do not blur the overall pattern).

Thus it could be inferred that:

- (a) observations 1 to 7 constitute a sample from a 'population' whose mean is at or near the target value (15);
- (b) observations 8 to 21 appear to have been sampled from a population whose mean is below 15;
- (c) observations 22 onward appear to come from a population whose mean is greater than 15.

0.2.3 There are now a number of questions that might be asked.

- (a) In the light of the underlying variability (as indicated, for example, by the minor irregularities in the cusum path) can it be concluded that the changes in slope represent real shifts in average rather than merely lucky or unlucky runs of samples from a stable population?
- (b) If the changes are real, how should the data be used to estimate local averages?
- (c) To what extent might the inferences or estimates be affected by the choice of the target value or the

cusum scale factor? Thus figures 3 and 4 show the same series plotted firstly with the same cusum scale but with a target of 12, and secondly with a target of 15 but a compressed cusum scale.

In both figures 3 and 4, the change in slope around observations number 8 is less apparent, though the change around number 21 is still visible, but it is less easy to 'pinpoint' in figure 4. Thus the choice of target value and scale factor need careful thought, to avoid either the suppression of useful information or, conversely, the exaggeration of spurious effects. It is also clear from figure 3 that use of an inappropriate target value may result in the chart running off the upper or lower edge of the graph paper (although this problem may also be overcome by replotting from a new zero at any point in the sequence; this procedure is detailed in clause 14 of Part 1 : 1980).

0.3 Monitoring or retrospective analysis

0.3.1 Questions (b) and (c) of 0.2.3 are covered in detail in Part 1, and question (a) in Part 2, where two distinct situations are recognized.

- (a) The object of the checking is to monitor the behaviour of a series of observations against some specified or standard value. This situation is the one that applies especially to process and quality control. Basic decision rules, based on a truncated mask, are given in clause 4 of Part 2 : 1980, and are summarized in appendix A. Decision rules and masks appropriate to counted data are developed in this Part.
- (b) The object is to examine historical data retrospectively, so as to detect any differences between segments. No formal standard or specification need




Figure 1. Conventional chart of data from table 1




Figure 2. Cusum chart of data from table 1




Figure 3. Cusum chart of data from table 1, with target value 12




Figure 4. Cusum chart of data from table 1, with target value 15 but compressed cusum scale

exist. The methods described in clause 5 of Part 2 : 1980 may apply in the field of process or quality control, especially in checking data from feasibility or process capability studies, in order to detect any anomalous behaviour of the system under observation. This use of cusums is described in clause 10 of Part 3 : 1981.

0.4 Scaling the chart

0.4.1 The relationship between the cusum and sample number scales will determine the appearance of the chart. Too large a cusum scale will produce a chart with a saw-tooth appearance on which small changes will be difficult to detect, while large shifts in average will quickly send the cusum path off the top or bottom of the graph sheet. Too small a scale will effectively damp out minor variations, but may also suppress real changes in mean level.

The most widely accepted convention is that a distance on the vertical (cusum) scale equal to one sample interval should represent approximately $2\sigma_e$ units of the plotted variable (σ_e , the standard error, is the subject of 0.4.2).

This convention is most easily applied as follows.

Decide on a convenient interval for the sample number axis (the i -scale). This may be, for example, 5 mm (or 0.2 in) for desk use, or 1 cm, 2 cm (or 0.5 in, 1 in) for wall display. Mark off this interval on the cusum scale, and let it represent about $2\sigma_e$ units of the values to be plotted. If $2\sigma_e$ is an inconvenient value, an alternative between $1.5\sigma_e$ and $3\sigma_e$ may be preferred for ease of plotting.

Further details of plotting conventions are given in clause 8 of Part 1 : 1980.

0.4.2 To establish the scale factor (and to implement the decision rules) an estimate of σ_e , the standard error of the plotted points, is required. This subject is considered in clause 7 and appendix B of Part 1 : 1980. In many applications involving counted observations, the distribution model (e.g. binomial or Poisson) will provide the required value of σ_e ($\sqrt{npq}$ for the binomial situation, using the conventional notation, or $\sqrt{m}$ for Poisson observations with mean level m per observation).

However, it is useful to check that the assumed standard error is realistic, and a simple method is to form the successive differences, $y_j - y_{j+1}$, between observations over a sequence where the process or system appears to operate at a satisfactory and stable level. The standard error may then be estimated as:

$$\sigma_e = \sqrt{\frac{1}{2(g-1)} \sum_{j=1}^{g-1} (y_j - y_{j+1})^2} \quad (1)$$

where g is the number of observations in the sequence.

Note that g has been adopted for the sequence length in this Part, to avoid confusion with m (as used in Parts 1 to 3) in view of the use of m as the parameter of a Poisson distribution.

Note also that the method of clause 7 of Part 1 : 1980 and 0.6 of Part 3 : 1981, based on *absolute* successive differences, is not recommended for counted data, especially where the mean number of events per observation is below 10.

0.4.3 The $2\sigma_e$ rule for scaling is applicable provided that σ_e itself is not less than 1.0. For smaller values, as may occur in monitoring for low rates of occurrence of events, the modification described in 7.4.2 of Part 2 : 1980 is advisable. This is as follows.

(a) Calculate the average number of observations (e.g. samples) required to yield one 'success' or event. Frequently this will be $1/p$ or $1/np$ (for binomial situations) or $1/m$ (for Poisson situations) where p , np or m are less than 1.0.

(b) Round this average number *up* to a convenient integer for plotting and adopt it as the horizontal interval for the cusum chart. For small rates of occurrence of events, the intervals between individual sample points on the chart may be very small, a reminder that the individual observations contain little information in isolation.

(c) Mark off the vertical (cusum) scale in intervals of the same length as the horizontal scale, and label these marks as consecutive *even* integers upward and downward from zero.

(d) The cusum is now plotted as:

$C = \Sigma(x - p)$ for binary data, where p is the proportion of units scoring 1;

or $C = \Sigma(x - np)$ for binomial data comprising samples of size n with proportion p of defectives (or other attribute of interest);

or $C = \Sigma(x - m)$ for counts of Poisson type events.

1. Scope

This Part of BS 5703 details cusum methods applicable to binary data (sequences of trials where the outcomes are scored as 0 or 1) or counts of attributes or events. Particular applications are in the fields of quality control by counts of defectives or defects, or by gauging. Proving trials where results are classified as success/failure (the less likely outcome conventionally being scored as 1, the more likely as 0) are also amenable to the methods described herein.

NOTE. The titles of the publications referred to in this standard are listed on the inside back cover.

2. Terminology and symbols

2.1 General. To avoid an unduly lengthy standard, it is necessary to assume some familiarity with basic statistical concepts and terminology. The essential requirements are an understanding of the following:

- $\bar{x}$ arithmetic mean of a set of data;
- μ arithmetic mean or expected value of the population from which the data is a sample;
- s sample estimate of the standard deviation;
- σ standard deviation of the population;
- n, p parameters (sample size, and proportion of items having an attribute of interest) of a binomial distribution;
- m parameter (rate of occurrence, per unit of observation, of an event of interest) of a Poisson distribution;

NOTE. λ and a are frequently used alternatives to m in other statistical texts.

Σ summation, in particular as used in

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2}$$

or $\sqrt{\frac{1}{n-1} \left\{ \sum x^2 - \frac{(\sum x)^2}{n} \right\}}$,

using also the abbreviated summation notation, where the range of x is assumed to extend over $i = 1, \dots, n$.

σ_e is used to denote the standard error of a sample statistic, the conventional measure of variation of a variable calculated from a sample; the standard error of the mean is the most familiar example.

Some notions from elementary probability and distribution theory (independence, correlation) are used in the standard. Where decision rules are discussed, the distribution of the observed variable will be specified (in this Part, this will often be a binomial or Poisson distribution).

An account of this basic material may be found in most elementary statistical text books. The only other requirement of the reader is an appreciation of the terms used in plotting graphs on two axes. The convention adopted is that the sequence in which the data occur is referred to the x-axis (abscissa, or conventional 'horizontal' scale), and the cumulative sum is referred to the y-axis (ordinate, or 'vertical' scale).

Careful thought has been given to the notation adopted for cusum parameters. In the statistical journals, an informal convention has evolved using the symbols H (or h), K (or k), and θ . There seems no reason to depart from the use of H and K , but θ has been adopted as a standard symbol (for another purpose) in BS 5532.

In cusum terms, the required symbol represents an angle; ϕ (phi) has therefore been adopted, leading to f to represent its tangent.

The upper case forms H, K, F are used in conjunction with charts referred to an observed variable, and the lower case (h, k, f) to a standardized variable, e.g. one having a standard Normal distribution $\mu = 0, \sigma = 1$. In this Part, the upper case forms are generally appropriate to counted data.

2.2 Symbols and terms

Symbol	Term	Definition
A	scale factor	The number of units of the plotted variable, measured on the cusum axis, corresponding to one sample interval on the sequence axis of the cusum chart.
AQL	acceptable quality level	The limiting value of the process average regarded as acceptable for process control purposes. NOTE. This definition is compatible with those given in BS 4778 and BS 5532, but extends its use to cover measurable characteristics in a manner consistent with current industrial practice.
ARL or L	average run length	At a given quality level, the average number of samples taken up to the point at which a signal occurs. This is generally related to a stated quality level and the symbol ' L ' then carries an appropriate subscript, as for example L_a or L_r for average run length at AQL or RQL respectively.
C	cusum	The cumulative sum of deviations of a variable from a target value.
F	reference shift	The difference between target and reference values for a cusum tabulation. F corresponds to the slope of the limb of a truncated V-mask.
g	segment length	The number of sample intervals in a cusum sequence for which a test of change in level is carried out. NOTE. m was used for segment length in Part 3, but g is adopted in this Part to avoid confusion with the Poisson parameter.
H	decision interval	The cumulative sum of deviations from a reference value required to yield a signal. H corresponds to the half-height at the base of a truncated V-mask.
i, j, r		Subscripts used to identify values of a variable, or of a cusum.
K	reference value	The value from which differences are calculated for a cusum tabulation, and the qualifying value for initiating a cusum sequence. $K = T + F$ for monitoring against upward shift.
m	Poisson parameter	The rate of occurrence of an event of interest.
N	sequence length	The total number of sample intervals covered by a retrospective cusum.
n	sample size	The number of individual items comprising a sample.
p	proportion	The proportion of binary observations scoring 1 ('successes' in the usual binomial distribution convention).
RL	run length	At a given quality level, the number of samples taken up to the point at which a signal occurs.
RQL	rejectable quality level	In the quality control sense, a level of quality regarded as unsatisfactory, which should be detected quickly by a quality control procedure.
T	target value	The target or standard value from which departure in average level requires to be detected.
x, y, Y	variables	Symbols denoting a random variable.
$\bar{x}, \bar{y}, \bar{Y}$	mean values	Arithmetic mean of a set of observations of random variables x, y, Y .
ϕ		Slope of the limb of the cusum mask, such that $\tan \phi = F$.

3. Decision rules: general

3.1 Attention is drawn to appendix A of this Part and clause 4 of Part 2 : 1980, for an account of the basis of cusum decision rules, and the concept of measuring effectiveness of a monitoring procedure using average run length (or other features of the run length distribution).

Subclause 7.2 of Part 2 : 1980 and, in more detail, clause 8 of Part 3 : 1981 should also be consulted for details of the tabulation method of operating a cusum procedure, and for the relationship between the decision parameters for such tabulations and the masks used in connection with charts. However, a brief resumé of the main points is presented in this clause.

3.2 Decision rules are necessary where some action is to be taken on the occurrence of a 'signal' from a cusum (or other) control procedure. The rules need to be objective so that decisions do not differ according to conscious or subconscious biases of the individual assessing the data. Risks of erroneous decision need to be considered,

and are often measured by an average run length (ARL) criterion. The two types of risks are as follows.

(a) The risk that a signal from the control procedure is, in fact, a false alarm when the process or system is operating satisfactorily. Such false alarms should be infrequent, giving a long ARL, here interpreted as an average number of observations between false alarms.

(b) The risk that out-of-control situation is not detected for some time after its onset. Here, the ARL should be short, and may be interpreted as the average number of observations between the onset of the change and its detection by the control procedure.

3.3 The features of the cusum chart that give an indication of a possible change in the level of the monitored variable are the steepness of the cusum slope and the number of observations over which this slope has persisted. Cusum masks of the type indicated in figure 13 take these two features into account. For counted data, their parameters (decision interval and slope of the decision line) require some modification.

For application to tabulation procedures for counted data, the cusum is recorded as $\Sigma(x - K)$, where $K = T + F$ and F is the slope of an appropriate cusum mask. A signal occurs if the cusum tabulation exceeds H , the decision interval. Negative cusums are ignored, cumulation starting again from zero as soon as any observation exceeds K . In this Part, for the reasons given in A.2.2, attention is limited to procedures for upward shifts, so that the situation where $K = T - F$ (described in Parts 2 and 3) is not covered.

3.4 Similarly, only one half of a cusum mask is required, this being the *lower* half so that it is suitable for detecting *upward* shifts in the level of the monitored variable. A general form of mask for counted data is shown in figure 5.




Figure 5. Form and simple construction for cusum mask for counted data

3.5 For the various applications described in later clauses, the average run length properties of the decision rules are referred to H , $K (= T + F)$ and a measure of the departure of the process or system from its target performance. The target level itself may relate to an acceptable proportion (p_a) of items having the attribute of interest, or to an acceptable rate of occurrence, m_a , of the event being monitored. Often p_a will refer to a proportion defective and m_a to a rate of occurrence of defects, but these situations are far from exhaustive.

3.6 At this stage, some examples will clarify the initial steps in setting up and scaling a cusum chart. One example will require the use of the general scaling convention introduced in Parts 1 to 3 and re-stated in 0.4.1 of this Part. The other will require the use of the special scaling rules given in 0.4.3 of this Part.

Example 1

Consider the case of an (assumed) Poisson variable with target mean rate of occurrence of events (e.g. faults, defects, breakdowns) per observation of about 4. Suppose the following sequence of 25 observations was obtained:

4, 4, 3, 4, 5, 2, 5, 1, 6, 5, 0, 3, 3,
7, 8, 2, 4, 1, 3, 6, 3, 2, 2, 4, 6.

For a Poisson distribution with $m = 4$, a variance of 4 is also anticipated, giving a standard error for the observations of $\sqrt{4} = 2$. Using equation (1), form the 24 successive differences, 0, 1, 1, 1, 3, 3, 4 etc., square, add the squares and divide by $2(25 - 1)$, i.e. by 48. Finally, take the square root of this result. This gives:

$$\sigma_e = \sqrt{\left(\frac{194}{48}\right)} = \sqrt{4.04} = 2.01$$

which is close to the anticipated value of 2.

Deciding on a plotting interval between observations of, say, 5 mm, use a scale of $2\sigma_e = 4$ cusum units per 5 mm.

Using a target value of $T = 4$, form the cusum as follows in table 2.

Table 2. Cusum data for Poisson variable with $m_a = 4$

Observation No.	x	$x - T$	Cusum $= \Sigma(x - T)$
1	4	0	0
2	4	0	0
3	3	-1	-1
4	4	0	-1
5	5	1	0
6	2	-2	-2
7	5	1	-1
8	1	-3	-4
9	6	2	-2
10	5	1	-1
11	0	-4	-5
12	3	-1	-6
13	3	-1	-7
14	7	3	-4
15	8	4	0
16	2	-2	-2
17	4	0	-2
18	1	-3	-5
19	3	-1	-6
20	6	2	-4
21	3	-1	-5
22	2	-2	-7
23	2	-2	-9
24	4	0	-9
25	6	2	-7

Finally, plotting cusum against sample number, the chart shown in figure 6 is obtained; a conventional chart is given below for comparison. The cusum remains substantially 'horizontal', indicating a mean level close to the target. To calculate the actual mean level over any sequence of interest, note that the *exact* arithmetic mean is given by:

$$\bar{x}_{j+1,j} = T + \frac{C_j - C_i}{j - i}$$

Thus, for the mean level over, say, observations 16 to 25 inclusive:

$$i + 1 = 16 \text{ (hence } i = 15)$$

$$j = 25$$

$$C_i = \text{cusum at observation } 15 = 0$$


$$C_j = \text{cusum at observation } 25 = -7$$

$$\text{then } \bar{x}_{15,25} = 4 + \frac{-7 - 0}{25 - 15} = 3.3, \text{ the correct mean value}$$

for the final 10 observations.


The application of decision rules to this example is deferred to clause 5.

Cusum
 $\Sigma(x - T)$



(a) Cusum chart with $T = 4$

No. of
 events, x



(b) Conventional chart


Figure 6. Cusum and conventional charts of Poisson data

Example 2

Consider a binomial application with $n = 20$, $p = 0.03$ under target conditions. The anticipated mean number of items (per sample of size 20) having the attribute of interest is $np = 0.6$, and the standard error is $\sqrt{np(1-p)} = 0.76$ approximately. Because this is less than 1, the scaling rule of 0.4.3 is appropriate.


Again taking 5 mm as a horizontal plotting interval, it needs now be taken to represent the number of samples required to generate (on average) one 'event'. This is given by $1/np = 1.67$ samples, which may conveniently be rounded to 2. The cusum scale is then marked off in even integers at each 5 mm on the vertical axis, yielding the chart of figure 7.

Cusum
 $\Sigma(x - T)$



(a) Cusum chart with $T = 0.6$

No. of
 events, x



(b) Conventional chart (using same horizontal scale for comparison)

Figure 7. Cusum and conventional charts of binomial data

Suppose that the sequence of 40 samples given in table 3 is obtained, yielding the attribute counts as shown. With $T = 0.6$, form the cusum, $\Sigma(x - 0.6)$ in this case, and plot the cusum against the sample number as in figure 7. Again a conventional chart is shown for comparison.

Table 3. Cusum data for binomial variable
 $n = 20, p = 0.03$

Sample no. i	Attribute count, x_i	$x_i - T$ ($T = 0.6$)	Cusum ($x_i - T$)
1	1	0.4	0.4
2	0	-0.6	-0.2
3	0	-0.6	-0.8
4	0	-0.6	-1.4
5	1	0.4	-1.0
6	0	-0.6	-1.6
7	0	-0.6	-2.2
8	1	0.4	-1.8
9	0	-0.6	-2.4
10	0	-0.6	-3.0
11	1	0.4	-2.6
12	1	0.4	-2.2
13	0	-0.6	-2.8
14	0	-0.6	-3.4
15	0	-0.6	-4.0
16	2	1.4	-2.6
17	0	-0.6	-3.2
18	0	-0.6	-3.8
19	1	0.4	-3.4
20	1	0.4	-3.0
21	1	0.4	-2.6
22	0	-0.6	-3.2
23	1	0.4	-2.8
24	0	-0.6	-3.4
25	1	0.4	-3.0
26	1	0.4	-2.6
27	0	-0.6	-3.2
28	1	0.4	-2.8
29	0	-0.6	-3.4
30	0	-0.6	-4.0
31	0	-0.6	-4.6
32	0	-0.6	-5.2
33	3	2.4	-2.8
34	1	0.4	-2.4
35	0	-0.6	-3.0
36	0	-0.6	-3.6
37	1	0.4	-3.2
38	0	-0.6	-3.8
39	1	0.4	-3.4
40	0	-0.6	-4.0

The cusum chart gives a much clearer indication of the mean level of the sequence; for example, it appears that a lower average applied over samples 1 to 15 than over samples 16 to 40. The cusum values may be used to estimate the mean levels in these two segments, thus:

First segment: $i + 1 = 1$ (hence $i = 0$)

$$j = 15$$

$$C_0 = 0$$

$$C_{15} = -4.0$$

$$\text{with } T = 0.6, \text{ we have } \bar{x}_{1,15} = 0.6 + \frac{-4 - 0}{15 - 0} = 0.33;$$

Second segment: $i + 1 = 16$ (hence $i = 15$, the final point in the previous sequence)

$$j = 40$$

$$C_{15} = -4.0$$

$$C_{40} = -4.0$$

$$\text{and } \bar{x}_{16,40} = 0.6 + \frac{-4.0 - (-4.0)}{40 - 15} = 0.6 \text{ exactly, as implied}$$

by the fact that with no change in cusum between C_{15} and C_{40} , the 'process' was operating at precisely its target value.

It would, of course, be useful to check that the low mean level over the first segment differs significantly from the target level, and an appropriate test of significance is described in appendix E of Part 2 : 1980. Decision rules for upward shifts from target are detailed later in this Part.

4. Decision rules for counted data

4.1 Clauses 5 and 6 present details of decision rules appropriate to Poisson and binomial variables. For each of a range of target conditions, two decision rules are given. The target conditions are specified by the mean rates of occurrence of events for Poisson variables, or as a proportion of values having some particular attribute for binomial variables.

One of the decision rules yields an average run length at the target level of at least 1000 observations, and the other an ARL (at the same target levels) of at least 200 observations. Retaining the notation used in Part 3, these are labelled C1 and C2 schemes respectively. Where conventional control schemes are shown for comparison, these are denoted S1 and S2 respectively (again following the convention established in Part 3).

4.2 In fact, a wide variety of decision rules could be devised having approximately the same ARL's at the target level. They often differ in their speed of reaction to shifts (from the target conditions) in mean rate or proportion of the occurrences under scrutiny. The principles adopted in selecting parameters for the C1 and C2 procedures are two-fold:

- (a) to give the greatest sensitivity to shifts yielding an ARL in the region of 5 to 10 observations;
- (b) to keep the parameters simple, generally as integers but adopting fractional steps of not less than $\frac{1}{4}$ for small rates (or proportions) of occurrences. Some special cases arise with very small rates or proportions, as detailed in clause 7.

4.3 Again following the notation of earlier parts, the parameters comprising the decision rule are denoted by H (decision interval) and K (reference value). The reference value K is appropriate to tabulation or computerized procedures. The corresponding decision line slope for cusum charts is given by $F = K - m_a$, where m_a is the target rate of occurrence for the chosen cusum procedure.

Thus for charts, cumulation is based on $\Sigma(x - m_a)$, and a signal occurs when the cusum path cuts the *lower* limb of a mask based on H and F , or cuts a decision line (likewise based on H, F) constructed *above* the suspected change point. For tabulation, the cusum becomes $\Sigma(x - K)$, reverting to zero whenever a negative cusum results from this cumulation, and yielding a signal if the cusum equals or exceeds H .

4.4 In Part 3, when dealing with measured variables, it is necessary to obtain H and F as $h\sigma_a$ and $f\sigma_a$, using the standard error (σ_a) of the monitored variable. In most applications covered by this Part, H and K are obtained directly from the tables, and F as $K - m_a$. Only for large rates of occurrence, where the counts of events or attributes can be assumed to be approximately Normally distributed, is the use of $h\sigma_a$ and $f\sigma_a$ (and hence $K = T \pm f\sigma_a$) necessary. However, as pointed out in **0.4.2**, it is always advisable to check that the observed standard error is consistent with the assumed distribution model. Clause 8 indicates suitable methods for such assessment.

4.5 The tables in clause 5 and 6 provide values of H and K appropriate to target levels m_a (for Poisson variables) or p_a (for binomial variables). In the latter case, the expected number of binomial events per sample of an item is also conveniently represented by $m_a = np_a$.

The general procedure for use of these tables is to choose an AQL (m_a or p_a) from the tables, and to use the appropriate pair of cusum parameters H and K in one of two ways.

- (a) For tabulation, form $\Sigma(x - K)$ and test against H in the usual manner.
- (b) For charting, plot $\Sigma(x - m_a)$ and use a mask with decision interval H and slope F as in figure 5, where $F = K - m_a$.

In some applications, where a target value (say m_t) falls between a pair of listed m_a values, it may be desirable to plot the cusum $\Sigma(x - m_t)$ rather than $\Sigma(x - m_a)$. This may be useful in order to obtain a 'horizontal' cusum path when the rate of occurrence of events is at or close to m_t . In this case the procedure is as follows.

- (c) Plot $\Sigma(x - m_t)$, and use a mask with decision interval H , but with slope equal to $K - m_t$ units per sample interval.

5. Cusum procedures for Poisson variables

5.1 Control of rates of occurrence of defects is most often exercised by counting the number of such defects in a specified quantity of the medium in which they occur. This quantity may, for example, be defined as a number of manufactured items subject to defects, as a length or area of product such as cable or sheet, or a volume of liquid or gas in which the 'defects' are particles of other material. Such control methods are not confined to industrial quality monitoring, but may be applied in any situation where counts of events or

incidents provide the data. The methods are thus applicable to monitoring accident rates, machine breakdowns, clerical performance in respect of error rates, blood cell counts, etc.

5.2 The generally assumed distribution model in such cases is the Poisson distribution. The validity of the Poisson model hinges on the independence of the events, and their occurrence at an average rate which (in the absence of assignable causes) is assumed to be stable. It is prudent to check that these are reasonable assumptions in any particular application, and if possible, to test them using one of the methods outlined in clause 8.

5.3 The Poisson distribution is also frequently adopted as a computationally simple approximation to the binomial distribution in cases where the proportion parameter (p) of the latter is not greater than 0.1. The most common application is in dealing with 'proportion defective' situations, where the proportion of defectives is generally substantially below 10 % under target conditions.

The Poisson approximation may also be adopted in using cusum procedures for binomial variables and (as occurs with the OC curve for sampling plans) results in a procedure that is never less discriminating than a true binomial scheme. In cusum applications, this means that the ARL curve is, for a given pair of (H, F) parameters, steeper for a binomial variable than for a true Poisson variable having the same mean rate of occurrence of events. The most noticeable feature is the longer ARL (for the binomial situation) at the target level, though there is also a slightly more rapid response of the cusum to large shifts in rate of occurrence where the true distribution is binomial rather than Poisson.

It also needs to be noted that where there is non-independence between events, or where the rate of their occurrence varies randomly in time or over the sampled medium, this is likely to result in departure from the Poisson distribution. This departure is often most apparent as an excess of variance over mean level, and the effect is to shorten ARL at target conditions whilst slightly slowing the response to large changes in the event rate.

5.4 As an example of these effects, table 4 shows the ARL data for three types of distribution over a range of rates of occurrence of events. The three types are as follows.

- (a) Binomial, with mean rate of occurrence of 'events' equal to np . The value of n is fixed at 10, while p is varied so as to match np to the m parameter of the Poisson distributions in (b).
- (b) Poisson, with 'event rate', m , equal to 1.0 at target, with hypothetical shifts up to $m = 5.0$.
- (c) Negative binomial with values of m (mean rate of occurrence of events) matching those of the Poisson distributions in (b) and with $k = 10$ (the symbol k , although similar to that for the cusum parameter K , is used in view of its widespread use as the 'exponent' or 'shape' parameter of Gamma and negative binomial distributions).

For each mean rate of occurrence, m , in table 4, the binomial distribution with $n = 10$ has smaller standard deviation than the Poisson, and the negative binomial