

Guide to

Data analysis and quality control using cusum techniques —

**Part 2: Decision rules and statistical
tests for cusum charts and tabulations**

UDC 519.244.8:658.562.012.7:519.816

NO COPYING WITHOUT BSI PERMISSION EXCEPT AS PERMITTED BY COPYRIGHT LAW



Document Sharing Hub

Cooperating organizations

The Quality Management and Statistics Standards Committee, under whose direction this British Standard was prepared, consists of representatives from the following Government department and scientific and industrial organizations:

- Confederation of British Industry
- Consumers' Association
- Electronic Engineering Association
- Institute of Cost and Management Accountants
- Institute of Quality Assurance*
- Institute of Statisticians*
- Institution of Electrical Engineers
- Institution of Production Engineers*
- Ministry of Defence*
- National Council for Quality and Reliability
- National Terotechnology Centre

The organizations marked with an asterisk in the above list, together with the following, were directly represented on the committee entrusted with the preparation of this British Standard:

- Association of Public Analysts
- British Paper and Board Industry Federation (PIF)
- Chemical Industries Association
- Department of Industry (National Engineering Laboratory)
- Economist Intelligence Unit Limited
- Electronic Components Industry Federation
- Institute of Petroleum
- Ministry of Agriculture, Fisheries and Food
- National Coal Board
- Post Office
- Royal Statistical Society
- University of Birmingham
- University of Essex

This British Standard, having been prepared under the direction of the Quality Management and Statistics Standards Committee, was published under the authority of the Executive Board and comes into effect on 29 February 1980

© BSI 10-1999

The following BSI references relate to the work on this standard:
Committee reference QMS/1
Draft for comment 77/60981 DC

ISBN 0 580 11046 X

Amendments issued since publication

Amd. No.	Date of issue	Comments
3735	September 1982	Indicated by a sideline in the margin

Contents

	Page
Cooperating organizations	Inside front cover
Foreword	iv
0 Introduction	1
0.1 Basis of cusum chart	1
0.2 Simple example of cusum chart	1
0.3 Monitoring or retrospective analysis	3
1 Scope	4
2 References	4
3 Symbols and terminology	5
4 Simple decision rules (a) Monitoring and control	6
4.1 Introduction	6
4.2 Cusum V-masks	8
4.3 Practical considerations concerning cusum masks	11
4.4 The full V-mask	11
4.5 Local decision lines	11
4.6 Average run length characteristics of basic decision rule	11
4.7 A semi-parabolic mask	14
4.8 Monitoring non-normal variables	16
5 Simple decision rules (b) Retrospective analysis	16
5.1 Introduction	16
5.2 "Span" tests	17
5.3 Construction of decision lines	17
5.4 Ad hoc tests of significance	17
5.5 Distribution-free cusum tests	18
5.6 A comparison of the four procedures	18
5.7 An example of application of the tests over a longer series	18
5.8 Automatic procedures	18
6 Examples of applications	18
6.1 Introduction	18
6.2 Setting up a monitoring procedure from a specification	18
6.3 Control of volatile matter in a synthetic resin	20
6.4 Ash-content in coal	22
6.5 Evaluation of the performance of a blending machine	24
6.6 Systematic patterns in fabric coating	26
6.7 Faults in knitted fabric	28
7 Extended cusum techniques	29
7.1 General	29
7.2 Cusums without charts	29
7.3 Application of automatic cusum procedure to package weight control	32
7.4 Binary data, or counts of events containing frequent zeros	33
7.5 An application of cusums to binary data	33
7.6 Data elements of varying size or importance	35
7.7 Estimation of standard error for weighted data	36
7.8 Application of weighted cusum to machine stoppages	38
7.9 Monitoring vehicle fuel consumption	39

	Page
Appendix A The full V-mask	44
Appendix B Local decision lines	46
Appendix C Monitoring non-normal variables	46
Appendix D Construction of decision lines for retrospective analysis	48
Appendix E Ad hoc tests of significance for retrospective analysis	49
Appendix F Distribution-free cusum tests	50
Appendix G A comparison of the four procedures	51
Appendix H An example of application of the tests over a longer series	52
Appendix J Rigorous procedure for obtaining standard errors for weighted data	57
Figure 1 — Conventional chart of data from Table 1	3
Figure 2 — Cusum chart of data from Table 1	3
Figure 3 — Cusum chart of data from Table 1, with target value 12	4
Figure 4 — Cusum chart of data from Table 1, with target value 15 but compressed cusum scale	4
Figure 5 — Some typical run length distributions	8
Figure 6 — General-purpose truncated V-mask	9
Figure 7 — Truncated V-mask applied to cusum chart: no indication of shift	10
Figure 8 — Truncated V-mask applied to cusum chart: indication of shift	10
Figure 9 — Average run length data for cusum and control charts	13
Figure 10 — Semi-parabolic cusum mask	16
Figure 11 — Application of span test to cusum chart	16
Figure 12 — Nomogram for evaluating maximum vertical height ($V_{\max}$) of cusum over span of m sample intervals	19
Figure 13 — Cusum chart for % volatile matter	21
Figure 14 — Cusum chart of ash-content data from Table 7	23
Figure 15 — Cusum chart of data from blending experiment	24
Figure 16 — Conventional chart of data from blending experiment	25
Figure 17 — Cusum charts of resin deposition data	27
Figure 18 — Cusum chart for fabric faults ($T = 7$, $\sigma = 6$)	28
Figure 19 — Alternative presentation of decision lines	29
Figure 20 — Modified cusums, S_1 and S_2	30
Figure 21 — Generalized cusum mask construction	31
Figure 22 — Cusum chart of ammunition firing trials (binary data)	34
Figure 23 — Weighted cusum chart for data of Table 12	37
Figure 24 — Weighted cusum chart of machine stoppages data	40
Figure 25 — Cusum charts for fuel use, 1 500 cc saloon car	43
Figure 26 — Full cusum V-mask	44
Figure 27 — Alternative method of constructing full V-mask	45
Figure 28 — Full V-mask applied to cusum chart ($T = 15$, $\sigma_e = 2$)	45
Figure 29 — Construction of local decision line	47
Figure 30 — Decision line constructed on cusum chart ($T = 15$, $\sigma_e = 2$)	47
Figure 31 — Test for significance of change in cusum slope, based on decision line	48
Figure 32 — Cusum chart for retrospective analysis of data in Table 20	56
Figure 33 — Typical span and decision line tests	56

	Page
Table 1 — Data for cusum plotting	2
Table 2 — Average run length data for cusum and control charts	12
Table 3 — Effect of incorrect value of σ_e on ARL	14
Table 4 — Data for construction of a semi-parabolic cusum mask	15
Table 5 — Batch % volatile matter	21
Table 6 — Average run length data for cusum scheme for % volatile matter	22
Table 7 — Data for ash-content in coal	22
Table 8 — Percentage of soluble ingredient present in discharge from a blending machine	25
Table 9 — Span tests on cusum segments	26
Table 10 — Ad hoc significance tests on cusum segments	26
Table 11 — Cusum calculations for modified procedure	32
Table 12 — Data with varying sample sizes for weighted cusum	37
Table 13 — Machine stoppages data	38
Table 14 — Fuel consumption data, 1 500 cc saloon car	39
Table 15 — Further fuel consumption details	42
Table 16 — Percentage points for cusum score (+ 1, - 1) about median (distribution-free test)	50
Table 17 — Scores and cumulatives of sequence 8 to 24	51
Table 18 — Scores and cumulatives of sequence 8 to 31	51
Table 19 — Difference $ \bar{x}_1 - \bar{x}_2 $ indicated as significant by various cusum tests (normally distributed variable)	52
Table 20 — Data and cusum calculations	53
Table 21 — Cusum span tests on data of Table 20	55
Table 22 — Tests of significance between mean levels of segments	55
Table 23 — Use of distribution-free cusum tests	57
Publications referred to	Inside back cover

Foreword

It was originally intended that this British 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 in manufacturing (e.g. management, administration, medical science, marketing, commerce) and comments on the draft standard revealed interest in a more general document. Part 1 of the standard is 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 3¹⁾ deals with applications of cusum methods to process and quality control by measurement. The techniques 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 4¹⁾ deals with control of counts of defectives.

Acknowledgement is made to the Department of Statistics and Computer Science, University College London, and Cambridge University Press for material reproduced, by permission, from *Random Normal Deviates* by H Wold published as *Tracts for computers*, No. XXV, 1954; to the authors of papers in the statistical journals whose work has been drawn upon; to the contributors of examples; to Mr P L Goldsmith for useful discussions; to ICI Fibres (Pontypool) for provision of computing facilities; 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.

A British Standard does not purport to include all the necessary provisions of a contract. Users of British Standards are responsible for their correct application.

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

Summary of pages

This document comprises a front cover, an inside front cover, pages i to iv, pages 1 to 58, an inside back cover and a back cover.

This standard has been updated (see copyright date) and may have had amendments incorporated. This will be indicated in the amendment table on the inside front cover.

¹⁾ In course of preparation.

0 Introduction

0.1 Basis of cusum chart. To avoid repetition of much of the material in Part 1, it is necessary to assume that the user of Part 2 is familiar with the basis of cusum charting as presented in Part 1. However, a brief recapitulation of the fundamentals is given before decision rules and statistical tests are detailed.

The cumulative sum chart (hereafter referred to by the generally accepted contraction “cusum chart”) is a highly informative means of graphical presentation of data which can be ordered into a logical sequence. Frequently, though not necessarily, this sequence corresponds to the order of observation on a time scale. However, examples are given where other, but equally logical, ordering is based on separation of the observations into groups having distinctive features.

The essential principle is that a *target* value, T , is subtracted from each observation. This target value is generally a constant but it may be a value that changes for each observation, e.g. a prediction from a forecasting model or a conditional target which may vary depending on the circumstances in which the observation is obtained. The cumulative sum of the deviations from target is formed, and this *cusum* (C) is plotted against the serial number of the observation.

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 (positive), the cusum slopes upwards (away from the sequence axis); conversely, when the local average is less than the target value (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.

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 slight loss in accuracy in some cases, from the chart itself.

0.2 Simple example of cusum chart. 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 were developed in Part 1.

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

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 than the conventional chart. 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 in Figure 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, in conventional terms, is 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 that 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.

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 of the cusum scale factor? Thus Figure 3 and Figure 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.

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




Figure 1 — Conventional chart of data from Table 1




Figure 2 — Cusum chart of data from Table 1

In Figure 3 and Figure 4, the change in slope around observation 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 will be detailed later).

0.3 Monitoring or retrospective analysis. Questions b) and c) were covered in detail in Part 1, but question a) was deferred to be treated in Part 2. In prescribing decision rules and statistical tests, two distinct situations should be recognized.

- a) The object of the checking is to monitor the behaviour of a series of observations against some specified or standard value such as, for example, in quality control operations. Decision rules for monitoring are presented in clause 4.
- b) The object is to examine historical data, or observations grouped in some logical manner, so as to detect any differences *between segments*; no formal standard or specification exists, the interest lying in the retrospective investigation of accumulated data. This situation is close to that of testing the significance of apparent differences between groups of observations, but it differs in that the grouping may be effected on the basis of a preliminary inspection of the cusum chart. Statistical tests for retrospective analysis are presented in clause 5.




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

1 Scope

This Part of this standard introduces simple decision rules for monitoring, control and retrospective analysis, statistical tests, examples and use of extended cusum techniques.

2 References

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

3 Symbols and terminology

3.1 General. To avoid an unduly lengthy document, it is necessary to assume some familiarity with basic statistical concepts and terminology. In the later Parts, a similar knowledge of the language of quality control will be assumed, but for this Part, 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 (real or imaginary) from which the data is a sample;
- s sample estimate of the standard deviation;
- σ standard deviation of the population;
- Σ summation, in particular as used in

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

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

$$\text{or} = \sqrt{\frac{1}{n-1} \sum x^2 - \frac{(\sum x)^2}{n}}$$

using also the abbreviated summation notation, where the range of x is assumed to extend over $i = 1, \dots, n$. σ_e will be 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 be almost exclusively the normal 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 are 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 $\mu = 0$, $\sigma = 1$.

3.2 Symbols and terms

Symbol	Term	Definition
a, A	Scale factor	$A = a\sigma_e$, and is 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
C	cusum	The cumulative sum of deviations of a variable from a target value
d	lead distance	The distance (in number of sample intervals) between a point of interest on a cusum chart and the vertex of a full V-mask

Symbol	Term	Definition
f, F	reference shift	The difference between target and reference values for a cusum tabulation. $F = f\sigma_e$ corresponds to the slope of the limb of a V- or truncated mask
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
RL	run length	The number of samples taken up to the point at which a signal occurs
ARL or L	average run length	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
h, H	decision interval	The cumulative sum of deviations from a reference value required to yield a signal. $H = h\sigma_e$ corresponds to the half-height at the datum of a truncated V-mask
i, j, r	reference value	Subscripts used to identify values of a variable, or of the cusum
K		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, and $K = T - F$ for monitoring against downward shift
m	segment length	The number of sample intervals in a cusum sequence for which a test of change in level is carried out
N	sequence length	The total number of sample intervals covered by a retrospective cusum
n	sample number	The number of individual items comprising a sample
P	proportion	The proportion of binary observations scoring 1 (“successes” in the usual Binomial distribution convention)
s	sample estimate of standard deviation	Estimated as: $\sqrt{\frac{1}{n-1} \sum (x - \bar{x})^2}$
T	target value	At a given quality level, the target or standard value from which departure in average level requires to be detected
$V_{\max}$	maximum vertical height	The maximum vertical distance of a cusum path from a chord joining the starting and finishing points of a segment
w	element size	The magnitude, in appropriate units, of a sample element
$\bar{w}$	mean element size	The arithmetic mean size of a number of sample elements
$\tilde{w}$	typical element size	A typical or notional element size adopted for charting purposes
x, y, Y	variables	Symbols denoting a random variable
$\bar{x}, \bar{y}, \bar{Y}$	mean values	Arithmetic mean of a set of random variables x, y, Y

4 Simple decision rules (a) Monitoring and control

4.1 Introduction

4.1.1 In this clause and in clause 5, simple decision rules are presented for determining whether apparent changes in cusum slope (and hence of the average level of the plotted variable) are real, or whether they may be regarded as part of the underlying variability in the data.

The present clause is concerned with monitoring and control applications: the plotting of data points in sequence as the observations are obtained, with a view to detecting as rapidly as possible any shift from a target level. Clause 5 deals with decision rules for the retrospective analysis of data already in its complete form, the object being to establish whether particular segments differ, in average level, from the preceding or succeeding segments.

4.1.2 It should be pointed out that the use of decision rules may not always be necessary, especially in cases where cusum charts are used as a means of effective presentation of series data. However, where decisions (and possibly resulting action) need to be taken, the cusum method often yields more efficient use of the available data than other traditional techniques. In this context, efficiency implies *either* that decisions about changes in level are reached more quickly; *or* that false alarms (erroneous decisions that a shift in mean has occurred when in fact there has been no change) are less frequent; *or* that both advantages may be obtained together.

4.1.3 In all applications requiring decisions to be made in the presence of variation in observed data, risks of erroneous decisions occur. In terms of classical statistical tests of significance, these risks are generally defined as:

type I error. The risk that a working (or null) hypothesis is erroneously rejected. For example, in acceptance sampling, the “producer’s risk” of erroneous rejection of a satisfactory batch corresponds to type I error, as does the risk of erroneously concluding that a process is “out of control” when in fact it is operating satisfactorily.


type II error. The risk that the working (or null) hypothesis is erroneously accepted when in fact an alternative hypothesis is true. The “consumer’s risk” in acceptance sampling, and failure to detect a real change in process conditions, correspond to type II error.

In general, these risks are measured by probabilities, and the conventional 5 % significance relates to a probability $\alpha = 0.05$ of type I error. The “power” of a test is measured by $1 - \beta$, where β is the probability of type II error at some specified departure from the working hypothesis.

4.1.4 When dealing with accumulating data, *sequential* tests of significance are often useful, permitting the appropriate decisions to be taken as soon as sufficient observations are obtained. However, the monitoring of series data raises the additional complication that it is impossible to specify in advance the point at which to begin a sequential test. The monitoring or control procedure has the dual objectives of detecting significant changes and the points at which they occur.

In order to measure the performance of such procedures, it appears more relevant to consider the average number of samples taken until a signal is indicated by the chosen decision rule. When no real change has occurred, this *average run length* should be long — perhaps hundreds of samples — as any signals will then be false alarms. When a large shift from target occurs, a short average run length is desirable so as to detect off-standard conditions rapidly. The response of a decision rule to various magnitudes of shift from target can be represented by an average run length curve (ARL curve), which provides a means of comparing decision rules in a manner analogous to that of the operating characteristic (OC curve) for comparing acceptance sampling procedures, or the power curve for comparing hypothesis tests.

It should be noted that the run length itself is subject to statistical variation. Sometimes one may be “lucky” in obtaining no false alarms over a long run, or detecting a change very quickly. Sometimes an “unfortunate” run of samples may generate false alarms, or mask a real change so that it does not yield a signal. The average run length (ARL) is used as a summary measure for comparison of schemes and decision rules, but other aspects of run length (mode or median run length, or the complete run length distribution) may sometimes deserve attention. Typical run length distributions are illustrated in Figure 5.



4.1.5 The essential features of the cusum chart relevant to the detection of changes in the underlying average level are the steepness of the cusum slope and the number of samples over which the slope persists. For monitoring and control, the object is generally to detect shifts from a target level which would yield a roughly horizontal cusum path. When the process is running at the target level (often known as an acceptable quality level or AQL), there will, because of sampling variation, inevitably be occasions when the cusum path apparently diverges from the horizontal. The decision rule should not give false alarms that indicate that such sequences represent significant departure from AQL. On the other hand, when the process moves to an unsatisfactory condition, for example towards a rejectable quality level or RQL, the decision rule should give as rapid a response as possible. The desiderata are thus:

True process condition	Cusum response
At or near AQL (target)	Long ARL (few false alarms)
Near RQL (substantial departure from target)	Short ARL (rapid detection)

The detailed specification of cusum procedures to give prescribed ARLs at particular quality levels is deferred to Part 3 of this standard. Here are presented simple rules that cover a wide range of applications. Their run length properties are related to quality levels generally defined as departures from target, measured in standard errors of the plotted variable. The standard error is assumed to be known from some preliminary evaluation, as described in clause 6 of Part 1 of this standard.

4.2 Cusum V-masks


4.2.1 General. The simplest decision rules for use in conjunction with cusum charts are embodied in V-masks. There are three slightly differing forms, but all are identical in principle and effect. Each form may be applied with any of the three scaling conventions; they will be applied to the more generally used version, as detailed in 8.2 of Part 1:1980 of this standard.

4.2.2 The truncated V-mask. Although a complete V-mask may be used, the truncated form is rather more convenient in practice. It comprises a datum point, indicated by A in Figure 6, which may be usefully identified by a small notch in the vertical section BC. Two sloping arms run from B to D and C to E, but they may be extended indefinitely beyond D or E if required. To correspond with later sections of this standard (and with Part 3) the vertical half-distances AB, AC will be termed *the decision interval*, and the lines BD, CE as *decision lines*. The slope of these decision lines, measured in scale units per sample plotting interval, will correspond to the *reference value* introduced in later sections.

The construction of the mask is detailed in Figure 6. Note that the distances $5\sigma_e$, $10\sigma_e$ may be transferred from the cusum scale if desired, or calculated from σ_e and the scale factor, if preferred.

The mask is used by placing the datum, A, over any point on the chart: this will often be the most recent point plotted, or the last point in some segment of particular interest. The AF axis is laid parallel to the sample number axis of the chart. If the path of the cusum wanders outside the sloping arms (or their extensions beyond D or E), a significant departure from the *target* value is signalled. However, if the entire cusum path remains inside the arms, no significant shift is indicated.

Figure 7 and Figure 8 show the truncated mask applied at two points on the cusum chart of the data from 0.2 (Table 1 and Figure 2). The value of σ_e is assumed to be 2.0. With the mask applied at observation 16, the cusum remains within the sloping arms, and the segment from samples 8 to 16 does not differ significantly from the target value of 15. However, with the mask applied at sample 18, the cusum path is seen to touch the *upper* decision line, indicating that, by the time the 18th value is obtained, sufficient evidence has accumulated to signal a *downward* shift from the target value.






Figure 7 — Truncated V-mask applied to cusum chart: no indication of shift




Figure 8 — Truncated V-mask applied to cusum chart: indication of shift

4.2.3 Several points are worth noting here.

- The upper decision line is *touched* at sample points 6 and 7. If the mask touches or crosses the decision line it is deemed to be a violation, yielding signals.
- The point at which the decision line is violated is on the extended arm of the mask, i.e. more than 10 samples behind the datum. The inference of a change in mean is none the less valid.
- If the mask is applied at any point before 18, no decision is signalled. If it is applied at any point between 18 and 21 (the end of the downward segment) a significant downward shift is indicated. Repeated (even unnecessary) use of the mask does not invalidate or modify the signals obtained.
- The signal occurs at sample 18. At this point, the evidence is sufficient to support a conclusion that the mean has changed, but the change is likely to have set in at some earlier stage. Inspection of the chart suggests sample 7 or 8 as the first of a sequence in which the average is displaced from the target value.

e) In a control situation, appropriate action would be indicated after plotting sample value 18 and obtaining a signal from application of the V-mask. An estimate of the new mean level may be useful in invoking the correct adjustment, and use of the method described in 13.4 of Part 1 of this standard yields:

$$C_{18} = -22, C_6 = 0;$$

$$\bar{x} = 15 + \frac{-22 - 0}{18 - 6} = 13.16$$

If action is delayed or ineffective, this estimate may be up-dated at a later stage if required.

f) Similar use of the V-mask later in this series would reveal a violation of the *lower* arm, indicating an *upward* shift by the 29th sample point.

g) The present procedure is not appropriate for testing the *difference* in means between segments 8 to 21 and 22 onward: it is designed for testing against the target value only.

4.3 Practical considerations concerning cusum masks

4.3.1 Before considering alternative forms of the preceding rule, it is necessary to consider some practical points regarding the form and construction of masks. These may be constructed from card or even paper, but a mask on transparent material (acrylic or acetate sheet) is more useful. The slope protractor (see clause 13 of Part 1:1980 of this standard) may then be incorporated on the same sheet of material, rendering estimation of local averages even simpler.

When V-masks are constructed from transparent material, removal of the truncated-V section is unnecessary, and a decision is signalled if the cusum path crosses a decision line. If a non-transparent material is used, either of two forms may be adopted as follows.

a) Cut out the truncated-V section, leaving the remainder. When the mask is used, a decision is signalled if the cusum path disappears behind the mask.

b) Remove and use the truncated-V section, discarding the remainder. A decision is then signalled if the cusum path appears outside the mask: as long as it remains hidden by the mask, no signal is indicated.

Selection of either of these forms is purely a matter of personal preference, but the transparent mask has the advantage that the entire cusum path remains visible, whether or not a decision is signalled or impending.

4.3.2 In cases where one-sided control is appropriate, a half-mask may be used if preferred. When monitoring against an upward shift, the lower portion of the mask is required, and for a downward shift, the upper half of the mask is required. However, a full mask may still be preferred on the grounds of simplicity, and any shifts in the irrelevant direction may be ignored. In many cases, evidence of such a change is of interest in directing attention to possible reasons for good process performance; a signal may then be noted, though the action generated will obviously differ from that triggered by an adverse shift.

4.4 The full V-mask. Appendix A details the construction of full V-masks.

4.5 Local decision lines. Appendix B details the construction of local decision lines.

4.6 Average run length characteristics of basic decision rule

4.6.1 It is now appropriate to consider the average run length properties of the basic decision rule formulated in 4.2 and incorporated in the various graphical methods of application by masks and decision lines. Readers familiar with control charts may find it useful to compare this ARL curve with those for two well-established control chart procedures:

- a) control chart with action line only at the 1 in 1 000 point of a target normal distribution;
- b) control chart with action line at the 1 in 1 000 point and warning line at the 1 in 40 point of a target normal distribution.

(For a description of control charts for variables, see BS 2564 or BS 600.)

The ARL data refer to observations having a normal distribution with known standard error (σ_e) which is unaffected by changes in mean level of the observations. For non-normal distributions, or where variation may be affected by changes in mean, the ARL data may be considerably affected, both for control charts and cusums.

The average run lengths refer to one-sided control of the average. Where two-sided control from a single target value is adopted (i.e. no *band* of satisfactory quality exists) some interaction between the upper and lower sides will occur. This is detailed in Part 3 of this standard, the most important effect being that the ARL at the target value is halved (there will be twice as many false alarms), but for large shifts in mean the ARL is unaffected.

The ARL data are listed in Table 2 and presented graphically in Figure 9.

Table 2 — Average run length data for cusum and control charts

Decision rules:

Cusum has $H = 5 \sigma_e$, $f = 0.5$ (slope of decision line $F = 0.5 \sigma_e$ per sample interval)

Control charts:

i) Action line at $3.09 \sigma_e$

ii) Action line at $3.09 \sigma_e$

Warning line at $1.96 \sigma_e$

i) Signal generated by single action value only

ii) Signal generated by single action value or two successive warning values

Plotted variable assumed normally distributed with standard error σ_e

Shift in process average from target value (units of σ_e)	Average run length (samples)			
	Cusum	Control Chart		Semi-parabolic cusum mask (see 4.7)
		(i)	(ii)	
0	930	1 000	640	470
0.25	140	443	256	113
0.5	38	208	126	35
0.75	17	104	52	16.5
1.0	10.5	55	26	10
1.5	5.8	18	8.9	5.3
2.0	4.1	7.3	4.2	3.4
2.5	3.2	3.6	2.5	2.3
3.0	2.6	2.2	1.75	1.74
4.0	1.9	1.2	1.2	1.21

4.6.2 Referring to Table 2 or Figure 9, note that when the process is at the target mean level, the cusum scheme gives an ARL roughly equal to that of control chart (i), and larger than that of control chart (ii): a useful feature, in that there are fewer false alarms. For shift from target in the range $0.1 \sigma_e$ to $2 \sigma_e$, signals occur (on average) more quickly with the cusum chart than with either control chart, especially in the region $0.5 \sigma_e$ to $1.5 \sigma_e$. However, *for this particular cusum decision rule*, the response to very large shifts ($> 2 \sigma_e$) is more rapid with control chart (ii), and for shifts beyond $2.8 \sigma_e$, even control chart (i) reacts more quickly than the cusum rule.

The reason for this is that the rule with $H = 5 \sigma_e$, $f = 0.5 \sigma_e$ is especially suited to detection of changes in the region of one standard error. If larger shifts are anticipated, other cusum rules may be more appropriate, giving more rapid response to large shifts whilst still maintaining much of their advantage on small or moderate changes. Full details of alternative rules will be given in Part 3, but a useful modification of the cusum mask, combining the better features of both cusum and control charts, is described in 4.7.




Figure 9 — Average run length data for cusum and control charts (see Table 2)

4.6.3 The user of the basic cusum decision rule may find it useful to have some indication of the possible effect of non-normality on the ARL characteristics. In general, skewness with the longer tail in the same sense as the direction of potential shift in one-sided control will result in shortening the ARL at target, but it will have little effect on the ARL for large shifts in mean. Conversely, if the shorter tail is in the direction of potential shift, the ARL at target level will be considerably lengthened (again with little effect on ARL for large shifts).

If the distribution is discrete rather than continuous (as may occur with coarse measurements as well as with counting of defectives or defects), the ARL near target is again shortened, with little effect on ARL for large shifts.

The effects of skewness and discreteness may be illustrated by two examples.

- For a “folded” normal distribution, as may occur in observing absolute deviations from a target quantity, the ARL at target conditions falls from 930 samples (for the normal distribution) to 370 samples.
- For a Poisson distribution with mean = 4 (so that the standard deviation is 2) a cusum scheme with $H = 10$ and $f = 0.5$ (slope of decision line = 1), has ARL 420 samples as compared with 930 samples for a normal distribution with the same mean and standard deviation.

In both cases, shifts which would yield ARL less than about 20 for a normal distribution yield very similar ARLs for the non-normal cases.

Table 3 — Effect of incorrect value of σ_e on ARL

Process average (as function of true σ_e)	Over-estimate σ_e by 10 %	Correct estimate of σ_e	Under-estimate σ_e by 10 %
T	3 000	930	410
$T + 0.5 \sigma_e$	45	38	35
$T + \sigma_e$	10	10.5	10
$T + 1.5 \sigma_e$	6	5.8	6
$T + 2 \sigma_e$	4.4	4.1	4.5

4.6.4 The importance of having a reliable estimate of σ_e has been noted in clause 6 of Part 1:1980 of this standard. Under- or over-estimation of σ_e will distort the ARL characteristics, again having little effect for large shifts but gross effects on ARL near target conditions. Thus for 10 % errors in estimation of σ_e , the distortions shown in Table 3 could occur.

4.6.5 Finally, it should be noted that the effects of non-normality, error in estimating parameters and other possible departures from explicit or implicit assumptions apply in similar degree to the use of conventional control charts, and are not a peculiarity of cusum methods.

4.7 A semi-parabolic mask

4.7.1 As was indicated in 4.6.2, other cusum decision rules may be more suitable, in some cases, than the one so far described, especially if shifts of some defined magnitude are of particular interest. In such cases, reference may be made to Part 3 of this standard or to texts dealing with cusum methods. In other cases, a decision rule which extends the more rapid detection of changes towards larger shifts in mean, whilst preserving the cusum superiority in signalling small shifts, may be useful. Such a decision rule is obtainable by a combination of conventional control chart and cusum rules, and can be embodied in a mask having a curved profile near its narrow end.

The basis of this semi-parabolic mask is the truncated-V of 4.2.2 (Figure 6). However, over the last five sample intervals at the narrow end, the mask curves in to a width (at the datum end) of $1.25 \sigma_e$ instead of $5 \sigma_e$. The data for construction of the mask are given in Table 4, and the mask appears as in Figure 10.