

Chapter 49

Specialized Control Charts

Chapter Table of Contents

OVERVIEW	1755
AUTOCORRELATION IN PROCESS DATA	1756
Diagnosing and Modeling Autocorrelation	1757
Strategies for Handling Autocorrelation	1758
MULTIPLE COMPONENTS OF VARIATION	1763
Preliminary Examination of Variation	1763
Determining the Components of Variation	1766
SHORT RUN PROCESS CONTROL	1770
Analyzing the Difference from Nominal	1770
Testing for Constant Variances	1776
Standardizing Differences from Nominal	1777
NONNORMAL PROCESS DATA	1779
Creating a Preliminary Individual Measurements Chart	1779
Calculating Probability Limits	1780
MULTIVARIATE CONTROL CHARTS	1783
Calculating the Chart Statistic	1783
Examining the Principal Component Contributions	1786

Chapter 49

Specialized Control Charts

Overview

Although the Shewhart chart serves well as the fundamental tool for statistical process control (SPC) applications, its assumptions are challenged by many modern manufacturing environments. For example, when standard control limits are used in applications where the process is sampled frequently, autocorrelation in the measurements can result in too many out-of-control signals. This chapter also considers process control applications involving multiple components of variation, short production runs, nonnormal process data, and multivariate process data.

These questions are subjects of current research and debate. It is not the goal of this chapter to provide definitive solutions but rather to illustrate some basic approaches that have been proposed and indicate how they can be implemented with short SAS programs. The sections in this chapter use the SHEWHART procedure in conjunction with various SAS procedures for statistical modeling, as summarized by the following table:

Process Control Application	Modeling Procedure
Diagnosing and modeling autocorrelation in process data	ARIMA
Developing control limits for processes involving multiple components of variation	MIXED
Establishing control with short production runs and checking for constant variance	GLM
Developing control limits for nonnormal individual measurements	CAPABILITY
Creating control charts for multivariate process data	PRINCOMP

Autocorrelation in Process Data

See SHWARIEW
in the SAS/QC
Sample Library

Autocorrelation has long been recognized as a natural phenomenon in process industries, where parameters such as temperature and pressure vary slowly relative to the rate at which they are measured. Only in recent years has autocorrelation become an issue in SPC applications, particularly in parts industries, where autocorrelation is viewed as a problem that can undermine the interpretation of Shewhart charts. One reason for this concern is that, as automated data collection becomes prevalent in parts industries, processes are sampled more frequently and it is possible to recognize autocorrelation that was previously undetected. Another reason, noted by Box and Kramer (1992), is that the distinction between parts and process industries is becoming blurred in areas such as computer chip manufacturing. For two other discussions of this issue, refer to Schneider and Pruett (1994) and Woodall (1993).

The standard Shewhart analysis of individual measurements assumes that the process operates with a constant mean μ , and that x_t (the measurement at time t) can be represented as $x_t = \mu + \epsilon_t$, where ϵ_t is a random displacement or error from the process mean μ . Typically, the errors are assumed to be statistically independent in the derivation of the control limits displayed at three standard deviations above and below the central line, which represents an estimate for μ .

When Shewhart charts are constructed from autocorrelated measurements, the result can be too many false signals, making the control limits seem too tight. This situation is illustrated in Figure 49.1, which displays an individual measurement and moving range chart for 100 observations of a chemical process.

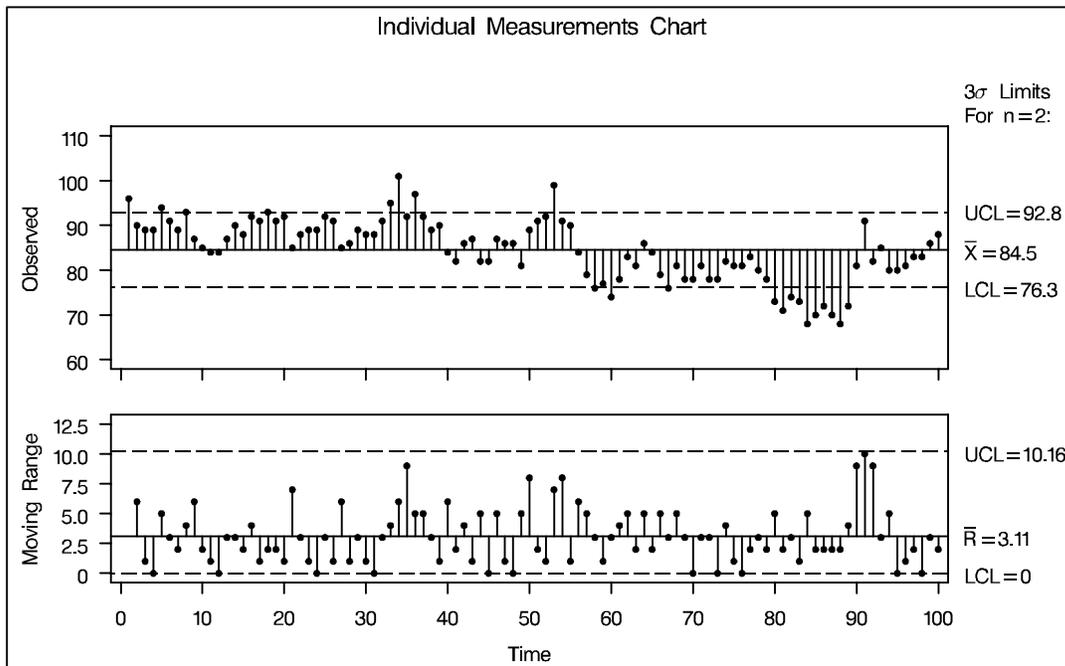


Figure 49.1. Conventional Shewhart Chart

The measurements are saved in a SAS data set named CHEMICAL.* The chart in Figure 49.1 is created with the following statements:

```

symbol value=dot;
title 'Individual Measurements Chart';
proc shewhart data=chemical;
    irchart xt*t / cneedles = black
                npanelpos = 100
                split      = '//';
    label xt = 'Observed/Moving Range'
          t  = 'Time';
run;

```

Diagnosing and Modeling Autocorrelation

You can diagnose autocorrelation with an autocorrelation plot created with the ARIMA procedure.

```

proc arima data=chemical;
    identify var = xt;
run;

```

Refer to *SAS/ETS User's Guide* for details on the ARIMA procedure. The plot, shown in Figure 49.2, indicates that the data are highly autocorrelated with a lag 1 autocorrelation of 0.83.

Autocorrelations																								
Lag	Covariance	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
0	48.348400	1.00000																						
1	40.141884	0.83026									.													
2	34.732168	0.71837									.													
3	29.950852	0.61948									.													
4	24.739536	0.51169									.													
5	20.594420	0.42596									.													
6	18.427704	0.38114									.													

Partial Autocorrelations																							
Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
1	0.83026										.												
2	0.09346									.				**	.								
3	0.00385									.				.									
4	-0.07340									.	*			.									
5	-0.00278									.				.									
6	0.09013									.				**	.								

Figure 49.2. Autocorrelation Plots for Chemical Data

The partial autocorrelation plot in Figure 49.2 suggests that the data can be modeled with a first-order autoregressive model, commonly referred to as an AR(1) model.

*The measurements are patterned after the values plotted in Figure 1 of Montgomery and Mas-trangelo (1991).

$$\tilde{x}_t \equiv x_t - \mu = \phi_0 + \phi_1 \tilde{x}_{t-1} + \epsilon_t$$

You can fit this model with the ARIMA procedure. The results in Figure 49.3 show that the equation of the fitted model is $\tilde{x}_t = 13.05 + 0.847\tilde{x}_{t-1}$.

```
proc arima data=chemical;
  identify var=xt;
  estimate p=1 method=ml;
run;
```

Maximum Likelihood Estimation					
Parameter	Estimate	Standard Error	t Value	Approx Pr > t	Lag
MU	85.28375	2.32973	36.61	<.0001	0
AR1,1	0.84694	0.05221	16.22	<.0001	1
Constant Estimate			13.05329		

Figure 49.3. Fitted AR(1) Model

Strategies for Handling Autocorrelation

There is considerable disagreement on how to handle autocorrelation in process data. Consider the following three views:

- At one extreme, Wheeler (1991b) argues that the usual control limits are contaminated “only when the autocorrelation becomes excessive (say 0.80 or larger).” He concludes that “one need not be overly concerned about the effects of autocorrelation upon the control chart.”
- At the opposite extreme, automatic process control (APC), also referred to as engineering process control, views autocorrelation as a phenomenon to be exploited. In contrast to SPC, which assumes that the process remains on target unless an unexpected but removable cause occurs, APC assumes that the process is changing dynamically due to known causes that cannot be eliminated. Instead of avoiding “overcontrol” and “tampering,” which have a negative connotation in the SPC framework, APC advocates continuous tuning of the process to achieve minimum variance control. Descriptions of this approach and discussion of the differences between APC and SPC are provided by a number of authors, including Box and Kramer (1992), MacGregor (1987, 1990), MacGregor, Hunter, and Harris (1988), and Montgomery and others (1994).
- A third strategy advocates removing autocorrelation from the data and constructing a Shewhart chart (or an EWMA chart or a cusum chart) for the residuals; refer, for example, to Alwan and Roberts (1988).

An example of the last approach is presented in the remainder of this section simply to demonstrate the use of the ARIMA procedure in conjunction with the SHEWHART

procedure. The ARIMA procedure models the autocorrelation and saves the residuals in an output data set; the SHEWHART procedure creates a control chart using the residuals as input data.

In the chemical data example, the residuals can be computed as forecast errors and saved in an output SAS data set with the FORECAST statement in the ARIMA procedure.

```
proc arima data=chemical;
  identify var=xt;
  estimate p=1 method=ml;
  forecast out=results id=t;
run;
```

The output data set (named RESULTS) saves the one-step-ahead forecasts as a variable named FORECAST, and it also contains the original variables XT and T. You can create a Shewhart chart for the residuals by using the data set RESULTS as input to the SHEWHART procedure.

```
title 'Residual Analysis Using AR(1) Model';
proc shewhart data=results(firstobs=4 obs=100);
  xchart xt*t / npanelpos = 100
              split      = '/'
              trendvar   = forecast
              xsymbol    = xbar
              ypct1      = 40
              vref2      = 70 to 100 by 10
              lvref      = 2
              nolegend;
  label xt = 'Residual/Forecast'
        t = 'Time';
run;
```

The chart is shown in Figure 49.4. Specifying TRENDVAR=FORECAST plots the values of FORECAST in the lower chart and plots the residuals ($XT - FORECAST$) together with their 3σ limits in the upper chart.*

Various other methods can be applied with this data. For example, Montgomery and Mastrangelo (1991) suggest fitting an exponentially weighted moving average (EWMA) model and using this model as the basis for a display that they refer to as an *EWMA central line control chart*.

Before presenting the statements for creating this display, it is helpful to review some terminology. The EWMA *statistic* plotted on a conventional EWMA control chart is defined as

$$z_t = \lambda x_t + (1 - \lambda)z_{t-1}$$

*The upper chart in Figure 49.4 resembles Figure 2 of Montgomery and Mastrangelo (1991), who conclude that the process is in control.

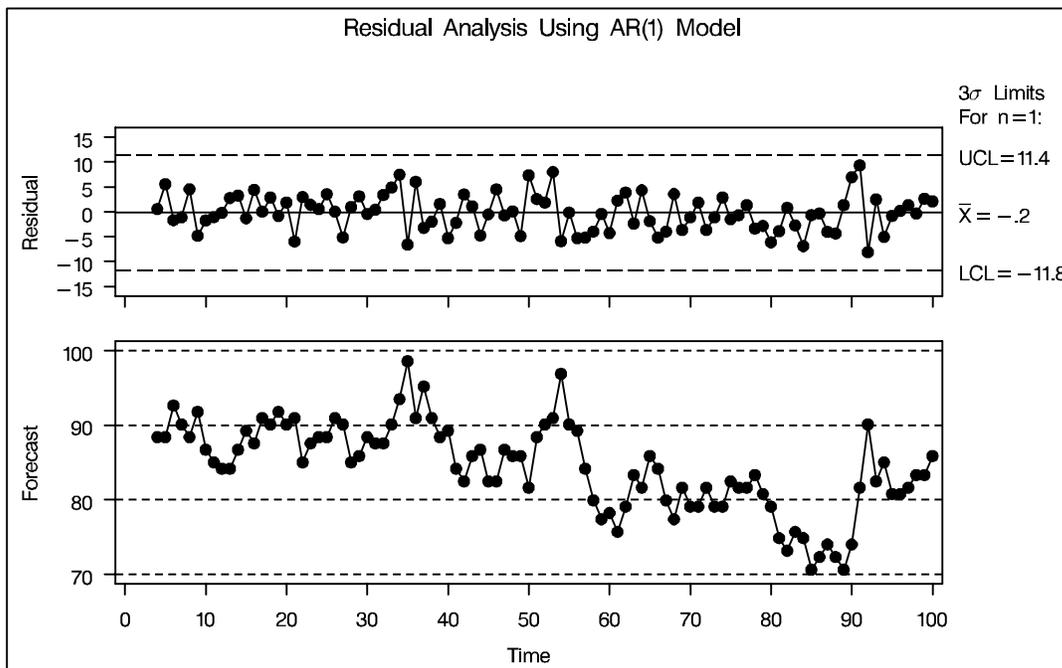


Figure 49.4. Residuals from AR(1) Model

The EWMA chart (which you can construct with the MACONTROL procedure) is based on the assumption that the observations x_t are independent. However, in the context of autocorrelated process data (and more generally in time series analysis), the EWMA statistic z_t plays a different role:* it is the optimal one-step-ahead forecast for a process that can be modeled by an ARIMA(0,1,1) model

$$x_t = x_{t-1} + \epsilon_t - \theta\epsilon_{t-1}$$

provided that the weight parameter λ is chosen as $\lambda = 1 - \theta$. This statistic is also a good predictor when the process can be described by a subset of ARIMA models for which the process is “positively autocorrelated and the process mean does not drift too quickly.”†

You can fit an ARIMA(0,1,1) model to the chemical data with the following statements. A summary of the fitted model is shown in Figure 49.5.

```
proc arima data=chemical;
  identify var=xt(1);
  estimate q=1 method=ml noint;
  forecast out=ewma id=t;
run;
```

*For a discussion of these roles, refer to Hunter (1986).

†Refer to Montgomery and Mastrangelo (1991) and the discussion that follows their paper.

Maximum Likelihood Estimation					
Parameter	Estimate	Standard Error	t Value	Approx Pr > t	Lag
MA1,1	0.15041	0.10021	1.50	0.1334	1
	Variance Estimate		14.97024		
	Std Error Estimate		3.86914		
	AIC		549.868		
	SBC		552.4631		
	Number of Residuals		99		

Figure 49.5. Fitted ARIMA(0,1,1) Model

The forecast values and their standard errors (variables FORECAST and STD), together with the original measurements, are saved in a data set named EWMA. The EWMA central line control chart plots the forecasts from the ARIMA(0,1,1) model as the central “line,” and it uses the standard errors of prediction to determine upper and lower control limits. You can construct this chart, shown in Figure 49.6,* with the following statements:

```

data ewma;
  set ewma(firstobs=2 obs=100);
run;

data ewmatab;
  length _var_ $ 8 ;
  set ewma (rename=(forecast=_mean_ xt=_subx_));
  _var_    = 'xt';
  _sigmas_ = 3;
  _limitn_ = 1;
  _lclx_   = _mean_ - 3 * std;
  _uclx_   = _mean_ + 3 * std;
  _subn_   = 1;

  title 'EWMA Center Line Control Chart';
  proc shewhart table=ewmatab;
    xchart xt*t / npanelpos = 100
                xsymbol   = 'Center'
                cinfill   = ligr
                llimits   = 1
                nolegend;
    label _subx_ = 'Observed'
          t     = 'Time' ;
  run;

```

Note that EWMA is read by the SHEWHART procedure as a TABLE= input data set, which has a special structure intended for applications in which both the statistics to be plotted and their control limits are pre-computed. The variables in a TABLE= data

*Figure 49.6 is similar to Figure 5 of Montgomery and Mastrangelo (1991).

Part 9. The CAPABILITY Procedure

set have reserved names beginning and ending with the underscore character; for this reason, FORECAST and XT are temporarily renamed as `_MEAN_` and `_SUBX_`, respectively. For more information on TABLE= data sets, see “Input Data Sets” in the chapter for the chart statement in which you are interested.

Again, the conclusion is that the process is in control. While Figure 49.4 and Figure 49.6 are not the only displays that can be considered for analyzing the chemical data, their construction illustrates the conjunctive use of the ARIMA and SHEWHART procedures in process control applications involving autocorrelated data.

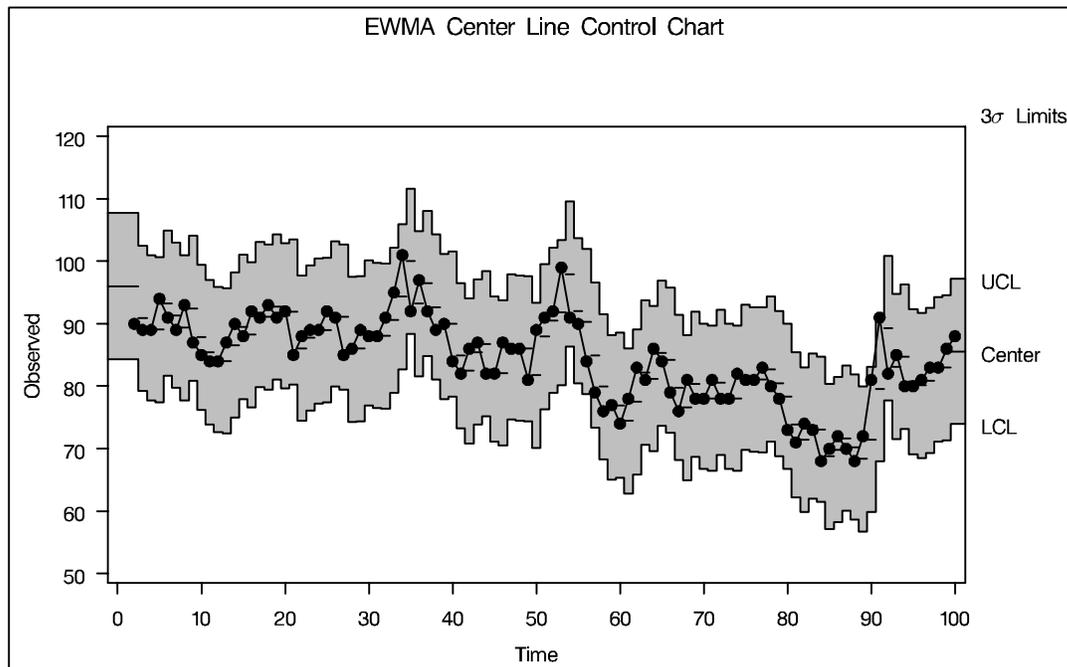


Figure 49.6. EWMA Center Line Chart

Multiple Components of Variation

In the preceding section, the excessive variation in the conventional Shewhart chart in Figure 49.1 is the result of positive autocorrelation in the data. The variation is “excessive” not because it is due to special causes of variation, but because the Shewhart model is inappropriate. This section considers another form of departure from the Shewhart model; here, measurements are *independent* from one subgroup sample to the next, but there are multiple components of variation for each measurement. This is illustrated with an example involving two components.*

See SHWMULTC
in the SAS/QC
Sample Library

A company that manufactures polyethylene film monitors the statistical control of an extrusion process that produces a continuous sheet of film. At periodic intervals of time, samples are taken at four locations (referred to as lanes) along a cross section of the sheet, and a test measurement is made of each sample. The test values are saved in a SAS data set named FILM. A partial listing of FILM is shown in Figure 49.7.

sample	lane	testval
1	A	93
1	B	87
1	C	92
1	D	78
2	A	87
.	.	.
.	.	.
.	.	.
56	D	75

Figure 49.7. Polyethylene Sheet Measurements in the Data Set FILM

Preliminary Examination of Variation

As a preliminary step in the analysis, the data are sorted by lane and visually screened for outliers (test values greater than 130) with box plots created as follows:

```
proc sort data=film;
  by lane;

title 'Outlier Analysis';
proc shewhart data=film;
  boxchart testval*lane / boxstyle = schematicid
                        idsymbol = dot
                        cboxfill = megr
                        vref      = 130
                        vreflab  = 'Outlier Cutoff'
                        hoffset  = 5
                        stddevs
                        nolegend
                        nolimits ;

  id sample;
run;
```

*Also refer to Chapter 5 of Wheeler and Chambers (1986) for an explanation of the effects of subgrouping and sources of variation on control charts.

Specifying `BOXSTYLE=SCHEMATICID` requests schematic box plots with outliers identified by the value of the ID variable `SAMPLE`. The `STDDEVS` option specifies that the estimate of the process standard deviation is to be based on subgroup standard deviations. Although this estimate is not needed here because control limits are not displayed, it is recommended that you specify the `STDDEVS` option whenever you are working with subgroup sample sizes greater than ten. The `NOLEGEND` and `NOLIMITS` options suppress the subgroup sample size legend and control limits for lane means that are displayed by default. The display is shown in Figure 49.8.

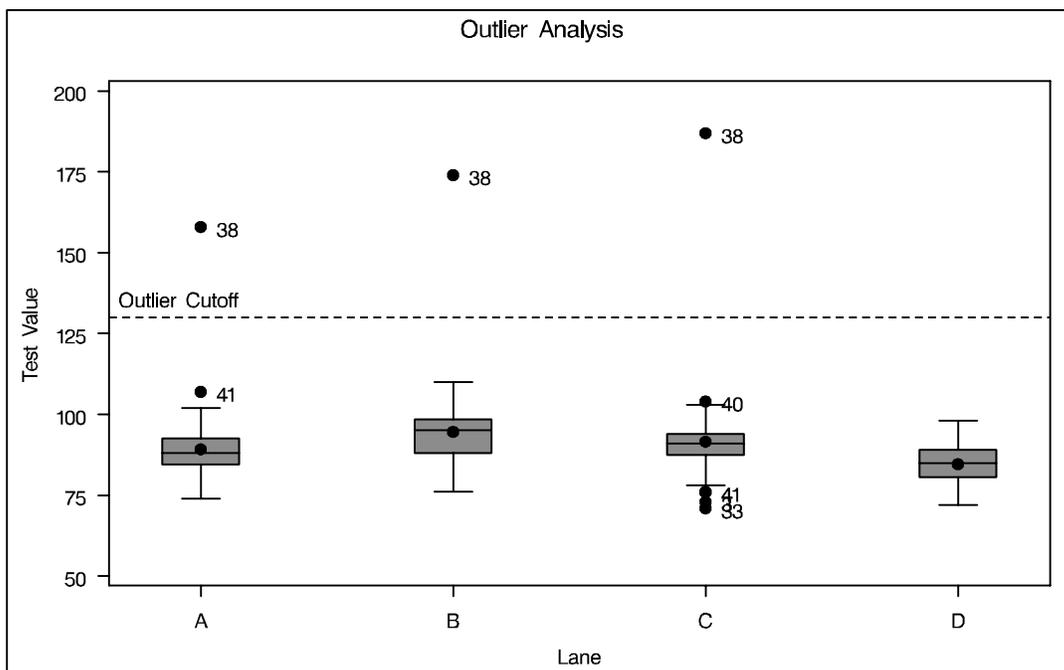


Figure 49.8. Outlier Analysis for the Data Set FILM

Figure 49.9 shows similarly created box plots for the data in FILM after the outliers have been removed.*

```

data film;
  set film;
  if testval < 130;

  title 'Variation Within Lane';
  proc shewhart data=film;
    boxchart testval*lane / boxstyle = schematicid
      idsymbol = dot
      cboxfill = megr
      hoffset = 5
      stddevs
      nolegend
      nolimits ;

    id sample;
  run;

```

*For the remainder of this section, unless otherwise indicated, it is assumed that the outliers are deleted from the data set FILM.

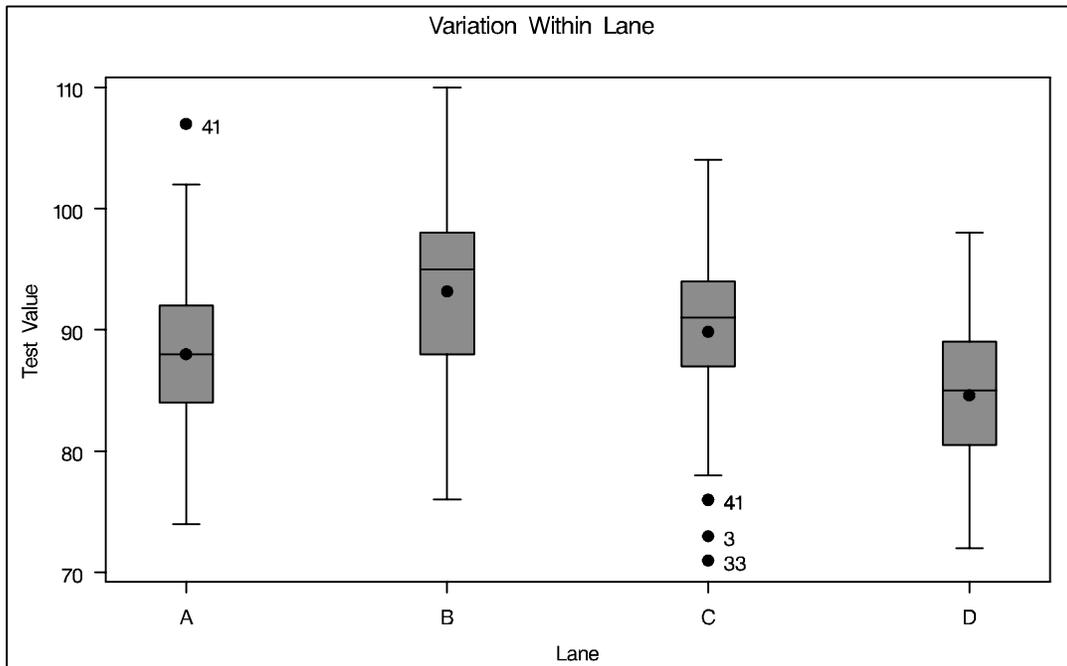


Figure 49.9. The Data Set FILM Without Outliers

Since you have no additional information about the process, you may want to create a conventional \bar{X} and R chart for the test values grouped by the variable SAMPLE. This is a straightforward application of the XRCHART statement in the SHEWHART procedure.

```
proc sort data=film;
  by sample;

  symbol value=dot;
  title 'Shewhart Chart for Means and Ranges';
  proc shewhart data=film;
    xrchart testval*sample /
      split      = '/'
      npanelpos = 60
      limitn     = 4
      coutfill  = megr
      nolegend
      alln;
    label testval='Average Test Value/Range';
  run;
```

The \bar{X} and R chart is displayed in Figure 49.10. Ordinarily, the out-of-control points in Figure 49.10 would indicate that the process is not in statistical control. In this situation, however, the process is known to be quite stable, and the data have been screened for outliers. Thus, the control limits seem to be inappropriate for the data.

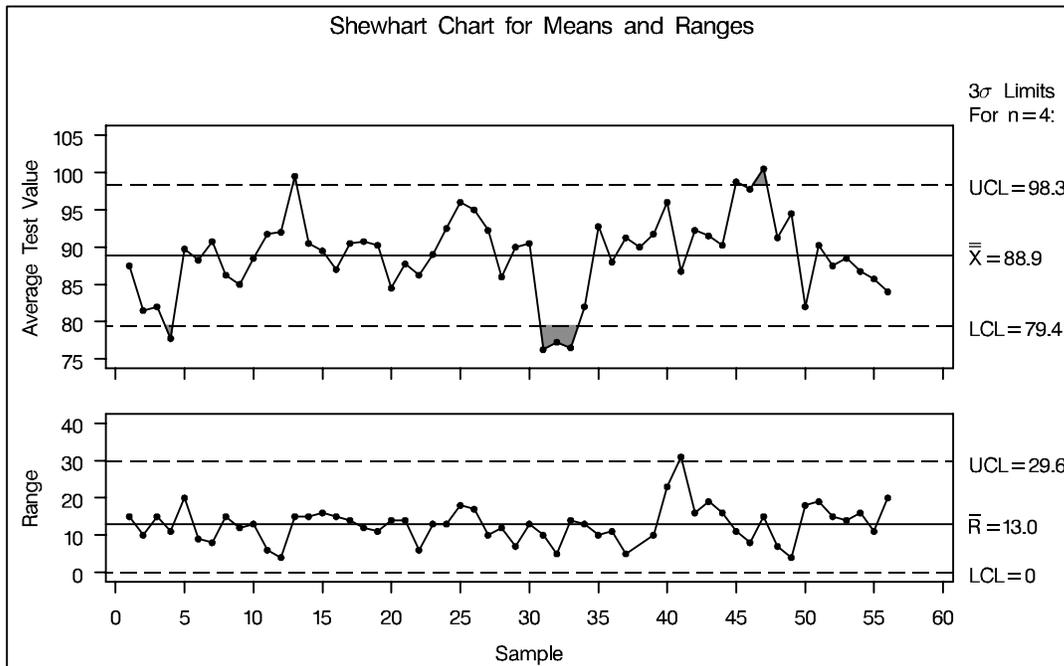


Figure 49.10. Conventional \bar{X} and R Chart

Determining the Components of Variation

The standard Shewhart analysis assumes that sampling variation, also referred to as *within-group* variation, is the only source of variation. Writing x_{ij} for the j^{th} measurement within the i^{th} subgroup, you can express the model for the conventional \bar{X} and R chart as

$$x_{ij} = \mu + \sigma_W \epsilon_{ij}$$

for $i = 1, 2, \dots, k$ and $j = 1, 2, \dots, n$. The random variables ϵ_{ij} are assumed to be independent with zero mean and unit variance, and σ_W^2 is the within-subgroup variance. The parameter μ denotes the process mean.

In a process such as film manufacturing, this model is not adequate because there is additional variation due to changes in temperature, pressure, raw material, and other factors. Instead, a useful model is

$$x_{ij} = \mu + \sigma_B \omega_i + \sigma_W \epsilon_{ij}$$

where σ_B^2 is the between-subgroup variance, the random variables ω_i are independent with zero mean and unit variance, and the random variables ω_i are independent of the random variables ϵ_{ij} .*

*This notation is used in Chapter 3 of Wetherill and Brown (1991), which discusses this issue.

To plot the subgroup averages $\bar{x}_i \equiv \sum \frac{x_{ij}}{n}$ on a control chart, you need expressions for the expectation and variance of \bar{x}_i . These are

$$E(\bar{x}_i) = \mu$$

$$\text{Var}(\bar{x}_i) = \sigma_B^2 + \frac{\sigma_W^2}{n}$$

Thus, the central line should be located at $\hat{\mu}$, and 3σ limits should be located at

$$\hat{\mu} \pm 3\sqrt{\hat{\sigma}_B^2 + \frac{\hat{\sigma}_W^2}{n}}$$

where $\hat{\sigma}_B^2$ and $\hat{\sigma}_W^2$ denote estimates of the variance components. You can use a variety of SAS procedures for fitting linear models to estimate the variance components. The following statements show how this can be done with the MIXED procedure:

```
proc mixed data=film;
  class sample;
  model testval = / s;
  random sample;
  make 'solutionf' out=sf;
  make 'covparms' out=cp;
run;
```

The results are shown in Figure 49.11. Note that the parameter estimates are $\hat{\sigma}_B^2 = 19.25$, $\hat{\sigma}_W^2 = 39.68$, and $\hat{\mu} = 88.90$.

Covariance Parameter Estimates					
		Cov Parm	Estimate		
		sample	19.2526		
		Residual	39.6825		
Solution for Fixed Effects					
Effect	Estimate	Standard Error	DF	t Value	Pr > t
Intercept	88.8963	0.7250	55	122.61	<.0001

Figure 49.11. Partial Output from the MIXED Procedure

The following statements merge the output data sets from the MIXED procedure into a SAS data set named NEWLIM that contains the appropriately derived control limit parameters for the average test value:

```

data cp;
  set cp sf;
  keep est;

proc transpose data=cp out=newlim;

data newlim;
  set newlim;
  drop _name_ _label_ col1-col3;
  length _var_ _subgrp_ _type_ $8;
  _var_ = 'testval';
  _subgrp_ = 'sample';
  _type_ = 'estimate';
  _limitn_ = 4;
  _mean_ = col3;
  _stddev_ = sqrt(4*col1 + col2);
  output;
run;

```

Here, the variable `_LIMITN_` is assigned the value of n , the variable `_MEAN_` is assigned the value of $\hat{\mu}$, and the variable `_STDDEV_` is assigned the value of

$$\hat{\sigma}_{\text{adj}} \equiv \sqrt{4\hat{\sigma}_B^2 + \hat{\sigma}_W^2}$$

In the following statements, the SHEWHART procedure reads these parameter estimates and displays the \bar{X} and R chart shown in Figure 49.12:

```

title 'Control Chart With Adjusted Limits';
proc shewhart data=film limits=newlim;
  xrchart testval*sample / npanelpos = 60;
run;

```

The control limits for the \bar{X} chart are displayed as $\hat{\mu} \pm \frac{3}{\sqrt{n}}\hat{\sigma}_{\text{adj}}$. Note that the chart in Figure 49.12 correctly indicates that the variation in the process is due to common causes.

You can use a similar set of statements to display the derived control limits in NEWLIM on an \bar{X} and R chart for the original data (including outliers), as shown in Figure 49.13.

A simple alternative to the chart in Figure 49.12 is an “individual measurements” chart for the subgroup means. The advantage of the variance components approach is that it yields separate estimates of the components due to lane and sample, as well as a number of hypothesis tests (these require assumptions of normality). In applying this method, however, you should be careful to use data that represent the process in a state of statistical control.

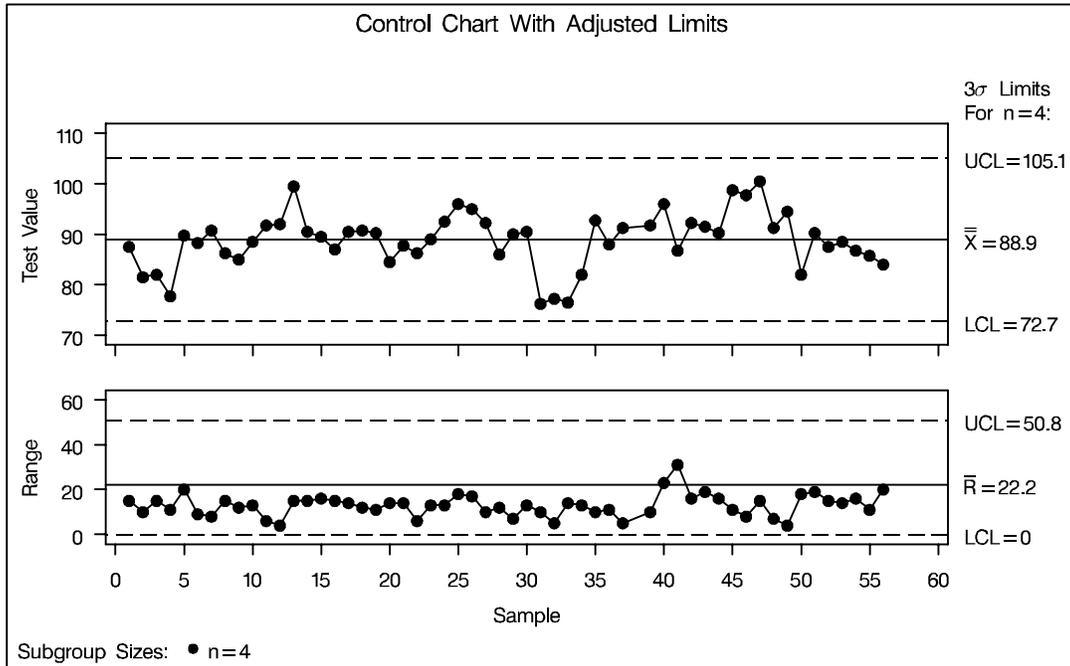


Figure 49.12. \bar{X} and R Chart with Derived Control Limits

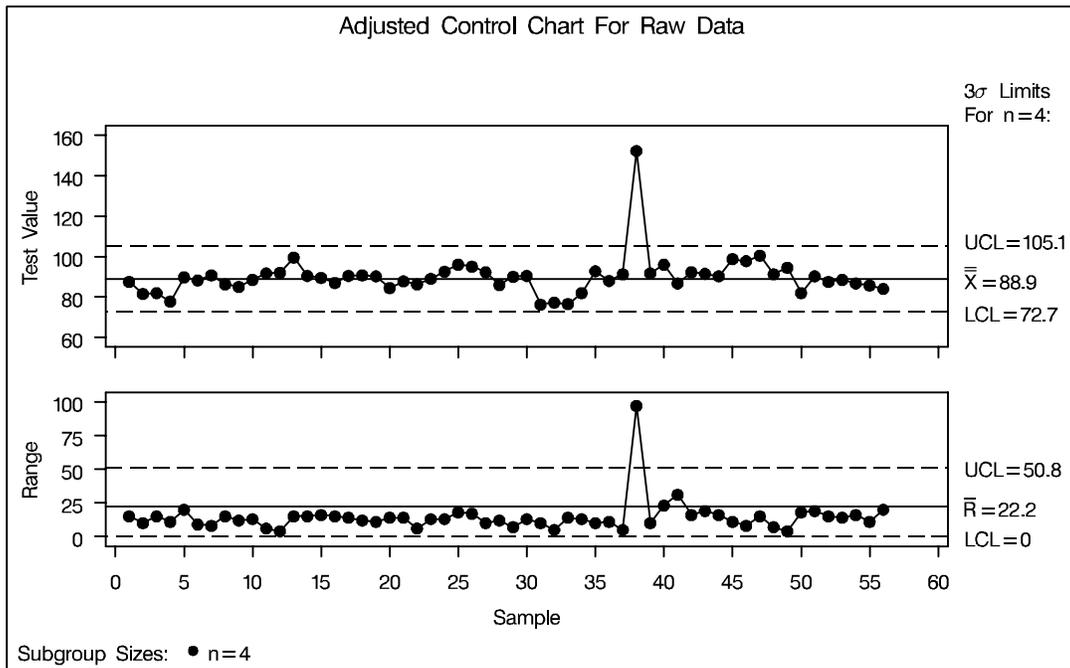


Figure 49.13. \bar{X} and R Chart with Derived Control Limits for Raw Data

Short Run Process Control

See SHWSRUN
in the SAS/QC
Sample Library

When conventional Shewhart charts are used to establish statistical control, the initial control limits are typically based on 25 to 30 subgroup samples. Often, however, this amount of data is not available in manufacturing situations where product changeover occurs frequently or production runs are limited.

A variety of methods have been introduced for analyzing data from a process that is alternating between short runs of multiple products. The methods commonly used in the United States are variations of two basic approaches:*

- the *difference from nominal* approach. A product-specific nominal value is subtracted from each measured value, and the differences (together with appropriate control limits) are charted. Here it is assumed that the nominal value represents the central location of the process (ideally estimated with historical data) and that the process variability is constant across products.
- the *standardization* approach. Each measured value is standardized with a product-specific nominal and standard deviation values. This approach is followed when the process variability is not constant across products.

These approaches are highlighted in this section because of their popularity, but two alternatives that are technically more sophisticated are worth noting.

- Hillier (1969) provided a method for modifying the usual control limits for \bar{X} and R charts in startup situations where fewer than 25 subgroup samples are available for estimating the process mean μ and standard deviation σ ; also refer to Quesenberry (1993).
- Quesenberry (1991a, 1991b) introduced the so-called *Q chart* for short (or long) production runs, which standardizes and normalizes the data using probability integral transformations.

SAS examples illustrating these alternatives are provided in the SAS/QC sample library and are described by Rodriguez and Bynum (1992).

Analyzing the Difference from Nominal

The following example[†] is adapted from an application in aircraft component manufacturing. A metal extrusion process is used to make three slightly different models of the same component. The three product types (labeled M1, M2, and M3) are produced in small quantities because the process is expensive and time-consuming.

Figure 49.14 shows the structure of a SAS data set named OLD, which contains the diameter measurements for various short runs. Samples 1 to 30 are to be used to estimate the process standard deviation σ for the differences from nominal.

*For a review of related methods, refer to Al-Salti and Statham (1994).

[†]Refer to Chapter 1 of Wheeler (1991a) for a similar example.

sample	prodtype	diameter
1	M3	13.99
2	M3	14.69
3	M3	13.86
4	M3	14.32
5	M3	13.23
.	.	.
.	.	.
.	.	.
30	M3	14.35

Figure 49.14. Diameter Measurements in the Data Set OLD

In short run applications involving many product types, it is common practice to maintain a database for the nominal values for the product types. Here, the nominal values are saved in a SAS data set named NOMVAL, which is listed in Figure 49.15.

prodtype	nominal
M1	15.0
M2	15.5
M3	14.8
M4	15.2

Figure 49.15. Nominal Values for Product Types in the Data Set NOMVAL

To compute the differences from nominal, you must merge the data with the nominal values. You can do this with the following SAS statements. Note that an IN= variable is used in the MERGE statement to allow for the fact that NOMVAL includes nominal values for product types that are not represented in OLD. Figure 49.16 lists the merged data set OLD.

```
proc sort data=old;
  by prodtype;
data old;
  format diff 5.2 ;
  merge nomval old(in = a);
  by prodtype;
  if a;
  diff = diameter - nominal;
proc sort data=old;
  by sample;
run;
```

sample	prodtype	diameter	nominal	diff
1	M3	13.99	14.8	-0.81
2	M3	14.69	14.8	-0.11
3	M3	13.86	14.8	-0.94
4	M3	14.32	14.8	-0.48
5	M3	13.23	14.8	-1.57
.
.
.
30	M3	14.35	14.8	-0.45

Figure 49.16. Data Merged with Nominal Values

Part 9. The CAPABILITY Procedure

Assume that the variability in the process is constant across product types. To estimate the common process standard deviation σ , you first estimate σ for each product type based on the average of the moving ranges of the differences from nominal. You can do this in several steps, the first of which is to sort the data and compute the average moving range with the SHEWHART procedure.

```
proc sort data=old;
  by prodtype;

proc shewhart data=old;
  irchart diff*sample /
  nochart
  outlimits=baselim;
  by prodtype;
run;
```

The purpose of this procedure step is simply to save the average moving range for each product type in the OUTLIMITS= data set BASELIM, which is listed in Figure 49.17 (note that PRODTYPE is specified as a BY variable).

Control Limits By Product Type						
prodtype	_VAR_	_SUBGRP_	_TYPE_	_LIMITN_	_ALPHA_	_SIGMAS_
M1	diff	sample	ESTIMATE	2	.002699796	3
M2	diff	sample	ESTIMATE	2	.002699796	3
M3	diff	sample	ESTIMATE	2	.002699796	3
LCLI	_MEAN_	_UCLI_	_LCLR_	_R_	_UCLR_	_STDDEV_
-3.13258	0.13000	3.39258	0	1.22714	4.00850	1.08753
-1.77795	-0.06500	1.64795	0	0.64429	2.10458	0.57098
-3.22641	-0.19143	2.84356	0	1.14154	3.72887	1.01166

Figure 49.17. Values of \bar{R} by Product Type

To obtain a combined estimate of σ , you can use the MEANS procedure to average the average ranges in BASELIM and then divide by the unbiasing constant d_2 .

```
proc means data=baselim noprint;
  var _r_;
  output out=difflim (keep=_r_) mean=_r_;

data difflim;
  set difflim;
  drop _r_;
  length _var_ _subgrp_ $ 8;
  _var_ = 'diff';
  _subgrp_ = 'sample';
  _mean_ = 0.0;
  _stddev_ = _r_ / d2(2);
  _limitn_ = 2;
  _sigmas_ = 3;
run;
```

The data set DIFFLIM is structured for subsequent use by the SHEWHART procedure as an input LIMITS= data set. The variables in a LIMITS= data set provide pre-computed control limits or—as in this case—the parameters from which control limits are to be computed. These variables have reserved names that begin and end with the underscore character. Here, the variable `_STDDEV_` saves the estimate of σ , and the variable `_MEAN_` saves the mean of the differences from nominal. Recall that this mean is zero, since the nominal values are assumed to represent the process mean for each product type. The identifier variables `_VAR_` and `_SUBGRP_` record the names of the process and subgroup variables (these variables are critical in applications involving many product types). The variable `_LIMITN_` is assigned a value of 2 to specify moving ranges of two consecutive measurements, and the variable `_SIGMAS_` is assigned a value of 3 to specify 3σ limits. The data set DIFFLIM is listed in Figure 49.18.

Control Limit Parameters For Differences					
<code>_var_</code>	<code>_subgrp_</code>	<code>_mean_</code>	<code>_stddev_</code>	<code>_limitn_</code>	<code>_sigmas_</code>
diff	sample	0	0.89006	2	3

Figure 49.18. Estimates of Mean and Standard Deviation

Now that the control limit parameters are saved in DIFFLIM, diameters for an additional 30 parts (samples 31 to 60) are measured and saved in a SAS data set named NEW. You can construct short run control charts for this data by merging the measurements in NEW with the corresponding nominal values in NOMVAL, computing the differences from nominal, and then constructing the short run individual measurements and moving range charts.

```
proc sort data=new;
  by prodtype;

data new;
  format diff 5.2 ;
  merge nomval new(in = a);
  by prodtype;
  if a;
  diff = diameter - nominal;
  label sample = 'Sample Number'
        prodtype = 'Model';

proc sort data=new;
  by sample;

symbol1 value=dot    color=black;
symbol2 value=plus   color=black;
symbol3 value=circle color=black;
title 'Chart for Difference from Nominal';
proc shewhart data=new limits=difflim;
  irchart diff*sample=prodtype / split='//';
  label diff = 'Difference/Moving Range';
run;
```

Part 9. The CAPABILITY Procedure

The chart is displayed in Figure 49.19. Note that the product types are identified with symbol markers as requested by specifying PRODTYPE as a *symbol-variable*.

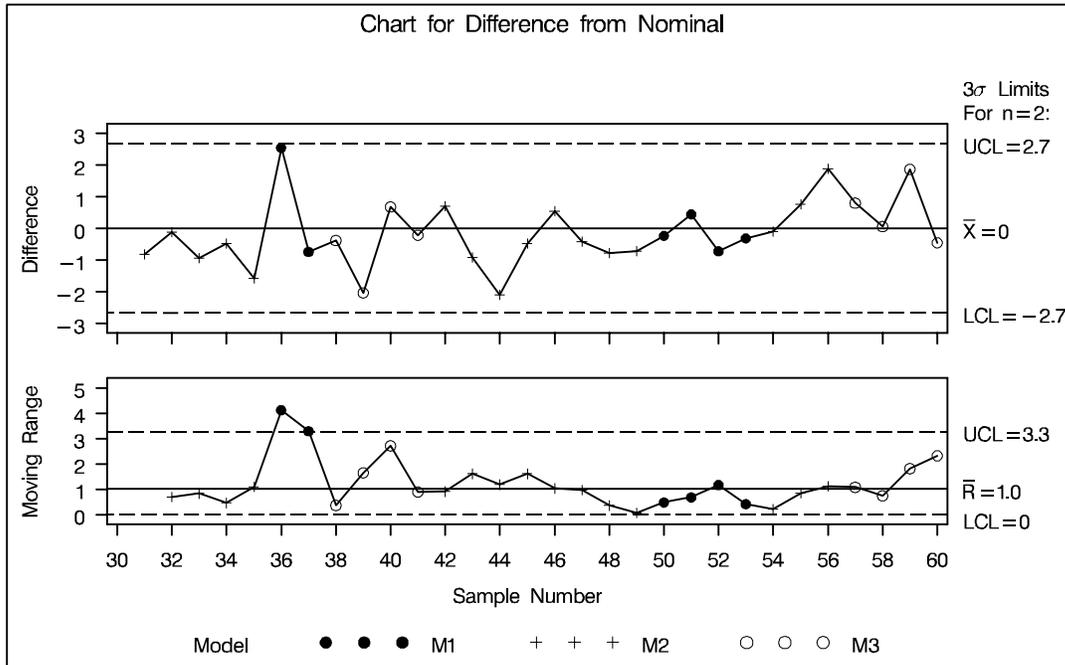


Figure 49.19. Short Run Control Chart

You can also identify the product types with a legend by specifying PRODTYPE as a *_PHASE_* variable.

```

symbol v=dot c=yellow;
title 'Chart for Difference from Nominal';
proc shewhart data=new (rename=(prodtype=_phase_))
  limits=difflim;
  irchart diff*sample /
    readphases=all
    phaseref
    phasebreak
    phaselegend
    split='//';
  label diff = 'Difference/Moving Range';
run;

```

The display is shown in Figure 49.20. Note that the PHASEBREAK option is used to suppress the connection of adjacent points in different phases (product types).

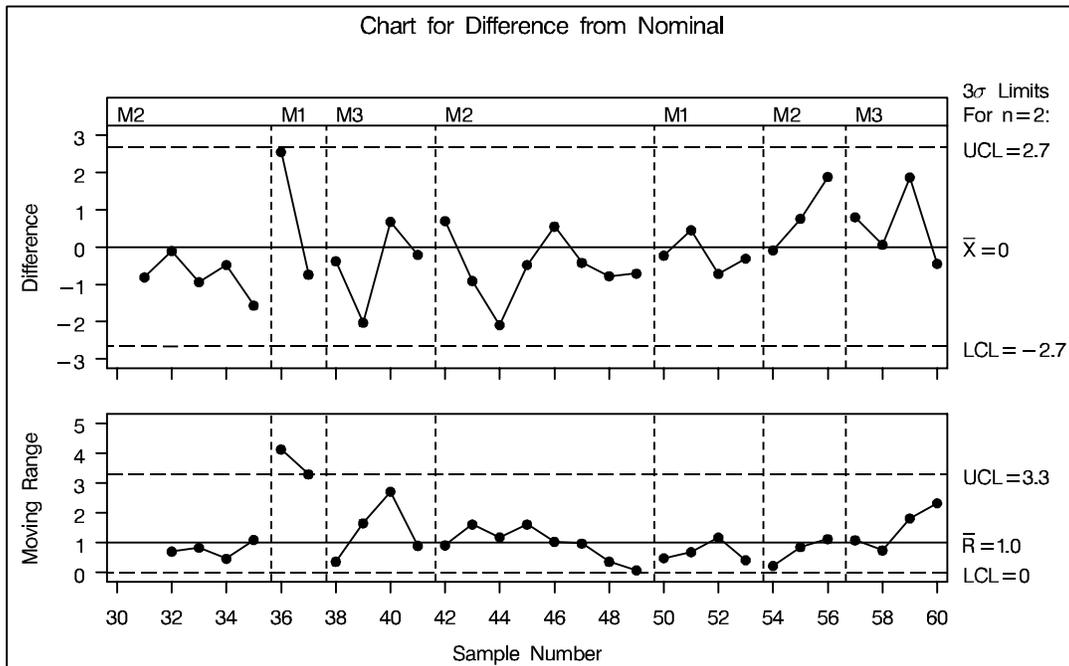


Figure 49.20. Identification of Product Types

In some applications, it may be useful to replace the moving range chart with a plot of the nominal values. You can do this with the `TRENDVAR=` option in the `XCHART` statement* provided that you reset the value of `_LIMITN_` to 1 to specify a subgroup sample of size one.

```

data difflim;
  set difflim;
  _var_ = 'diameter';
  _limitn_ = 1;

title 'Differences and Nominal Values';
proc shewhart data=new limits=difflim;
  xchart diameter*sample (prodtype) /
    nolimitslegend
    nolegend
    split = '/'
    blockpos      = 3
    blocklabtype  = scaled
    blocklabelpos = left
    xsymbol       = xbar
    trendvar      = nominal;
  label diameter = 'Difference/Nominal'
    prodtype     = 'Product';
run;

```

*The `TRENDVAR=` option is not available in the `IRCHART` statement.

The display is shown in Figure 49.22. Note that you identify the product types by specifying PRODTYPE as a *block variable* enclosed in parentheses after the subgroup variable SAMPLE. The BLOCKLABTYPE= option specifies that values of the block variable are to be scaled (if necessary) to fit the space available in the block legend. The BLOCKLABELPOS= option specifies that the label of the block variable is to be displayed to the left of the block legend.

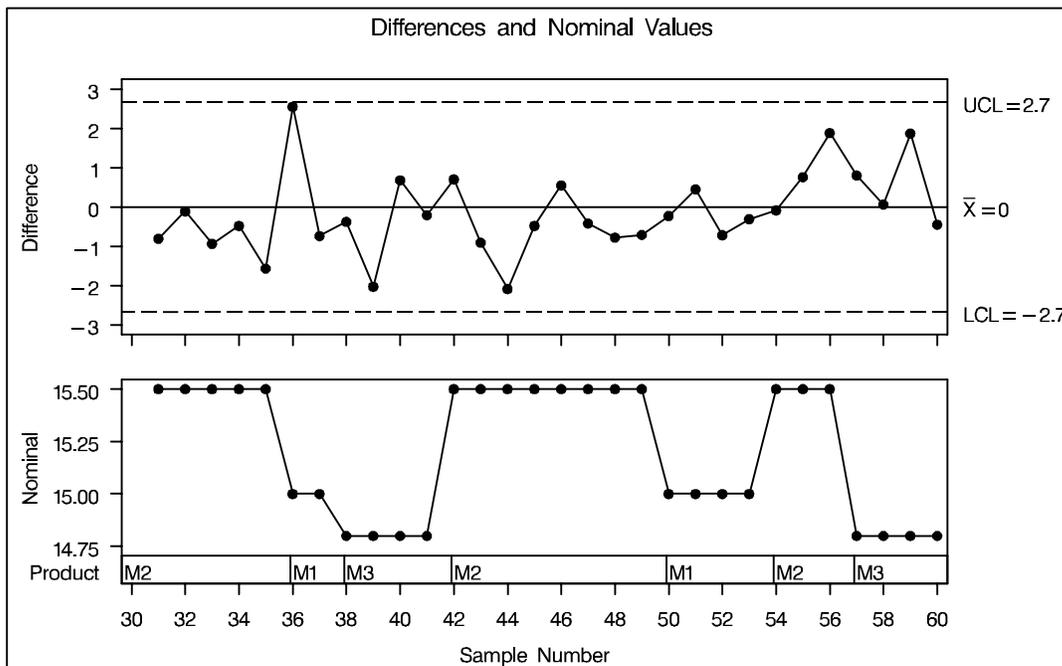


Figure 49.21. Short Run Control Chart with Nominal Values

Testing for Constant Variances

The difference-from-nominal chart should be accompanied by a test that checks whether the variances for each product type are identical (homogeneous). Levene's test of homogeneity is particularly appropriate for short run applications because it is robust to departures from normality; refer to Snedecor and Cochran (1980). You can implement Levene's method by using the GLM procedure to construct a one-way analysis of variance for the absolute deviations of the diameters from averages within product types.

```
proc sort data=old;
  by prodtype;

proc means data=old noprint;
  var diameter;
  by prodtype;
  output out=oldmean (keep=prodtype diammean) mean=diammean;

data old;
  merge old oldmean;
  by prodtype;
  absdev = abs( diameter - diammean );
```

```

proc means data=old noprint;
  var absdev;
  by prodtype;
  output out=stats n=n mean=mean css=css std=std;

title 'Test for Constant Variance';
proc glm data=old outstat=glmout ;
  class prodtype;
  model absdev = prodtype;
run;

```

A partial listing of the results is displayed in Figure 49.22. The large p -value (0.3386) indicates that the data do not reject the hypothesis of homogeneity.

The GLM Procedure					
Dependent Variable: absdev					
Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	2	1.02901063	0.51450532	1.13	0.3373
Error	27	12.27381243	0.45458565		
Corrected Total	29	13.30282306			

Figure 49.22. Levene's Test of Variance Homogeneity

Standardizing Differences from Nominal

When the variances across product types are *not* constant, various authors recommend standardizing the differences from nominal and displaying them on a common chart with control limits at ± 3 .

To illustrate this method, assume that the hypothesis of homogeneity is rejected for the differences in OLD. Then you can use the product-specific estimates of σ in BASELIM to standardize the differences from nominal in NEW and create the standardized chart as follows:

```

proc sort data=new;
  by prodtype;

data new;
  keep sample prodtype z diff diameter nominal _stddev_;
  label sample = 'Sample Number';
  format diff 5.2 ;
  merge baselim new(in = a);
  by prodtype;
  if a;
  z = (diameter - nominal) / _stddev_ ;

proc sort data=new;
  by sample;

```

Part 9. The CAPABILITY Procedure

```

title 'Standardized Chart';
proc shewhart data=new;
  irchart z*sample (prodtype) /
    blocklabtype = scaled
    mu0          = 0
    sigma0       = 1
    split        = '//';
  label prodtype = 'Product Classification'
        z        = 'Standardized Difference/Moving Range';
run;

```

Note that the options MU0= and SIGMA= specify that the control limits for the standardized differences from nominal are to be based on the parameters $\mu = 0$ and $\sigma = 1$. The chart is displayed in Figure 49.23.

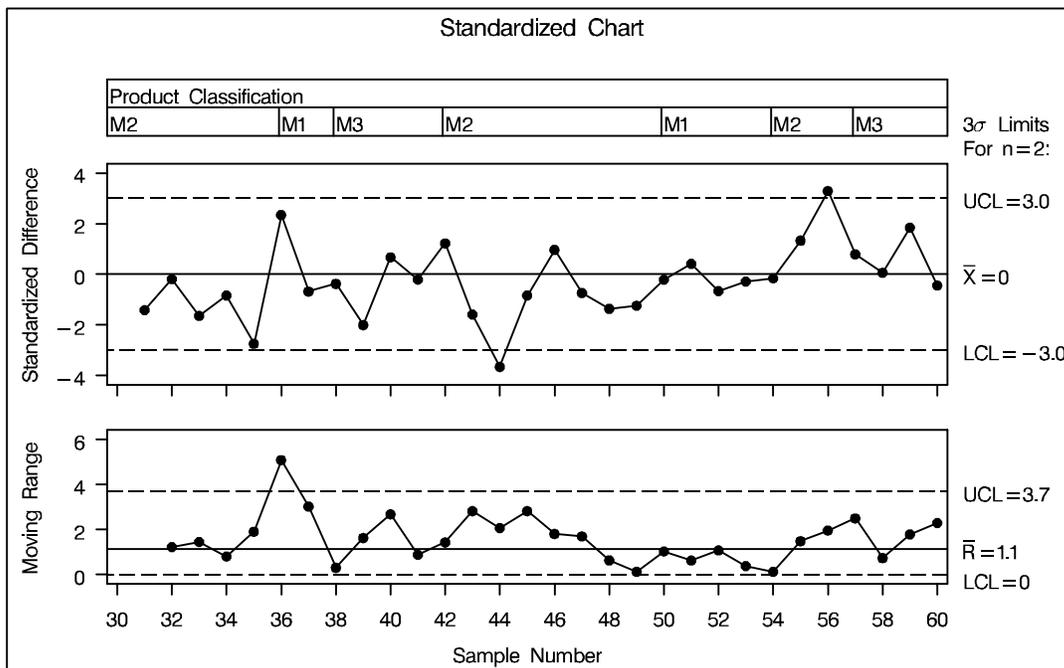


Figure 49.23. Standardized Difference Chart

Nonnormal Process Data

A number of authors have pointed out that Shewhart charts for subgroup means work well whether the measurements are normally distributed or not.* On the other hand, the interpretation of standard control charts for individual measurements (X charts) is affected by departures from normality.

See SHWNONN
in the SAS/QC
Sample Library

In situations involving a large number of measurements, it may be possible to subgroup the data and construct an \bar{X} chart instead of an X chart. However, the measurements should not be subgrouped arbitrarily for this purpose.† If subgrouping is not possible, two alternatives are to transform the data to normality (preferably with a simple transformation such as the log transformation) or modify the usual limits based on a suitable model for the data distribution.

The second of these alternatives is illustrated here with data from a study conducted by a service center. The time taken by staff members to answer the phone was measured, and the delays were saved as values of a variable named TIME in a SAS data set named CALLS. A partial listing of CALLS is shown in Figure 49.24.

recnum	time
1	3.233
2	3.110
3	3.136
.	.
.	.
50	2.871

Figure 49.24. Answering Times from the Data Set CALLS

Creating a Preliminary Individual Measurements Chart

As a first step, the delays were analyzed using an X chart created with the following statements. The chart is displayed in Figure 49.25.

```

title 'Standard Analysis of Individual Delays';
proc shewhart data=calls;
  irchart time * recnum /
    rtmplot = schematic
    outlimits = delaylim
    cboxfill = grey
    nochart2;
  label recnum = 'Record Number'
        time = 'Delay (minutes)' ;
run;

```

You may be inclined to conclude that the 4th point signals a special cause of variation. However, the box plot in the right margin (requested with the RTMPLOT= option) indicates that the distribution of delays is skewed. Thus, the reason that the measure-

*Refer to Schilling and Nelson (1976) and Wheeler (1991b).

†Refer to Wheeler and Chambers (1986) for a discussion of subgrouping.

ments are grouped well within the control limits is that the limits are incorrect and not that the process is too good for the limits.

Note: This example assumes the process is in statistical control; otherwise, the box plot could not be interpreted as a representation of the process distribution. You can check the assumption of normality with goodness-of-fit tests by using the CAPABILITY procedure, as shown in the statements that follow.

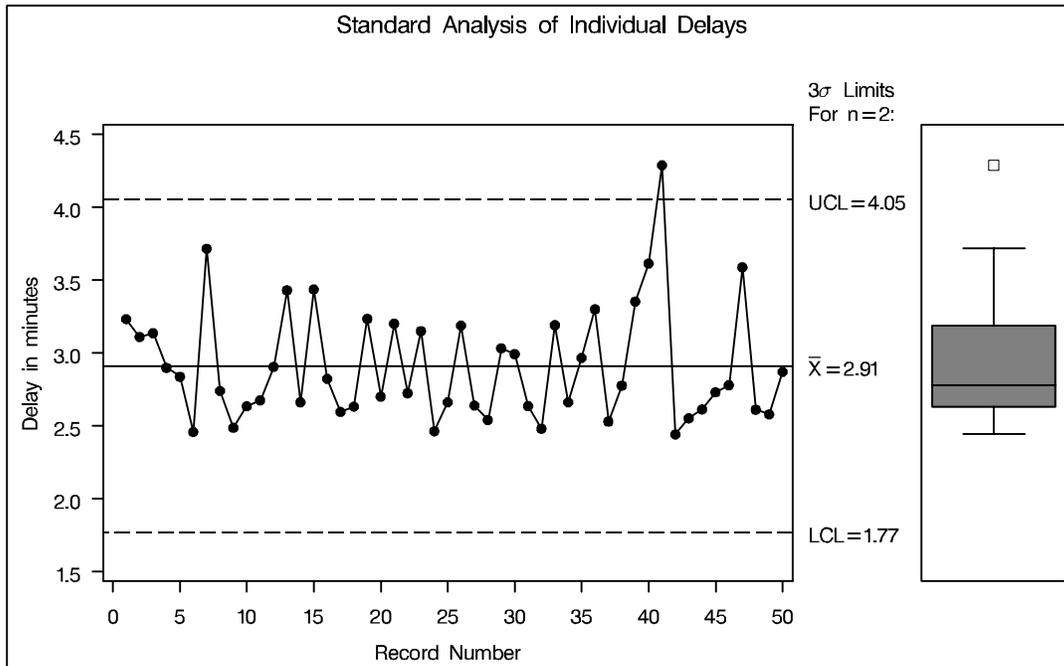


Figure 49.25. Standard Control Limits for Delays

Calculating Probability Limits

The OUTLIMITS= option saves the control limits from the chart in Figure 49.25 in a SAS data set named DELAYLIM, which is listed in Figure 49.26.

—	S	—	L	—	S	—	—	—	S
—	U	—	I	—	A	I	—	—	T
—	B	T	M	—	L	G	L	M	U
V	G	Y	I	—	P	M	C	E	C
A	R	P	T	—	H	A	L	A	L
R	P	E	N	—	A	S	I	N	I
—	—	—	—	—	—	—	—	—	—
time	recnum	ESTIMATE	2	.002699796	3	1.77008	2.91038	4.05068	0.38010

Figure 49.26. Control Limits for Standard Chart from the Data Set CALLS. The control limits can be replaced with the corresponding percentiles from a fitted lognormal distribution. The equation for the lognormal density function is

$$f(x) = \frac{1}{x\sqrt{2\pi\sigma}} \exp\left(-\frac{(\log(x)-\zeta)^2}{2\sigma^2}\right) \quad x > 0$$

where σ denotes the shape parameter and ζ denotes the scale parameter.

The following statements use the CAPABILITY procedure to fit a lognormal model and superimpose the fitted density on a histogram of the data, shown in Figure 49.27:

```

title 'Lognormal Fit for Delay Distribution';
proc capability data=calls noprint;
  histogram time /
    lognormal(threshold=2.3 color=black w=2)
    cfill = grey
    outfit = lnfit
    nolegend ;
  inset n = 'Number of Calls'
    lognormal( sigma = 'Shape' (4.2)
              zeta = 'Scale' (5.2)
              theta ) / pos = ne;
run;

```

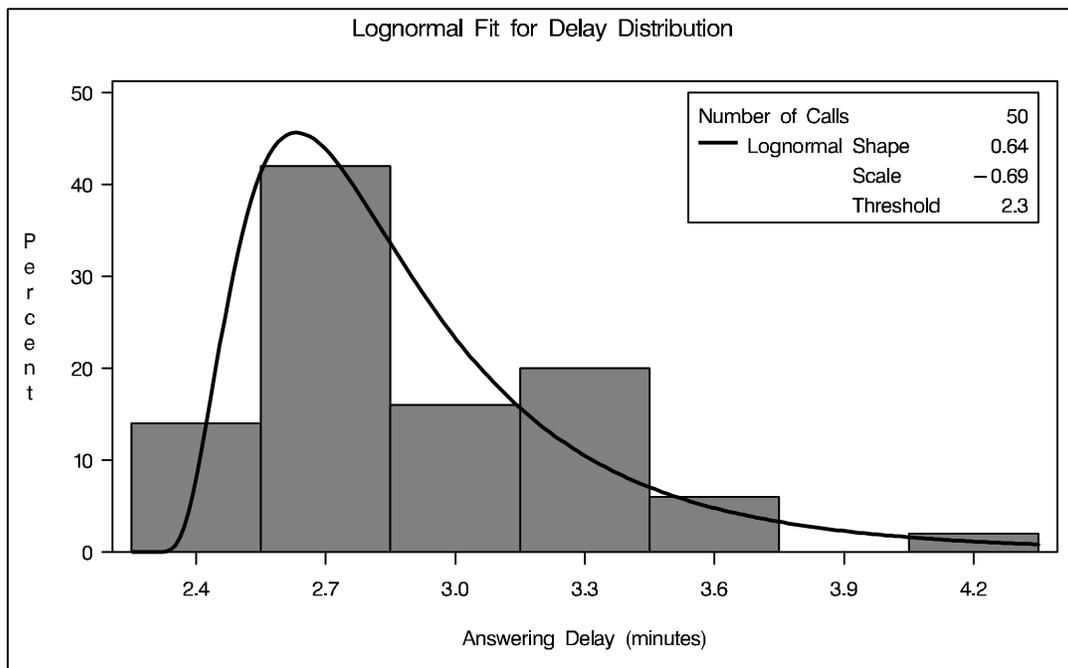


Figure 49.27. Distribution of Delays

Parameters of the fitted distribution and results of goodness-of-fit tests are saved in the data set LNFIT, which is listed in Figure 49.28. The large p -values for the goodness-of-fit tests are evidence that the lognormal model provides a good fit.

VAR	_CURVE_	_LOCATN_	_SCALE_	_SHAPE1_	_MIDPTN_
time	LNORMAL	2.3	-0.68910	0.64110	4.2
ADASQ	_ADP_	_CVMWSQ_	_CVMP_	_KSD_	_KSP_
0.34854	0.47465	0.058737	0.40952	0.092223	0.15

Figure 49.28. Parameters of Fitted Lognormal Model in the Data Set LNFIT

Part 9. The CAPABILITY Procedure

The following statements replace the control limits in DELAYLIM with limits computed from percentiles of the fitted lognormal model. The $100\alpha^{\text{th}}$ percentile of the lognormal distribution is $P_\alpha = \exp(\sigma\Phi^{-1}(\alpha) + \zeta)$, where Φ^{-1} denotes the inverse standard normal cumulative distribution function. The SHEWHART procedure constructs an X chart with the modified limits, displayed in Figure 49.29.

```

data delaylim;
  merge delaylim lnfit;
  drop _sigmas_ ;
  _lcli_ = _locatn_ + exp(_scale_+probit(0.5*_alpha_)*_shape1_);
  _ucli_ = _locatn_ + exp(_scale_+probit(1-0.5*_alpha_)*_shape1_);
  _mean_ = _locatn_ + exp(_scale_+0.5*_shape1_*_shape1_);

title 'Lognormal Control Limits for Delays';
proc shewhart data=calls limits=delaylim;
  irchart time*recnum /
    rtmplot = schematic
    cboxfill = grey
    nochart2 ;
  label recnum = 'Record Number'
        time = 'Delay (minutes)' ;
run;

```

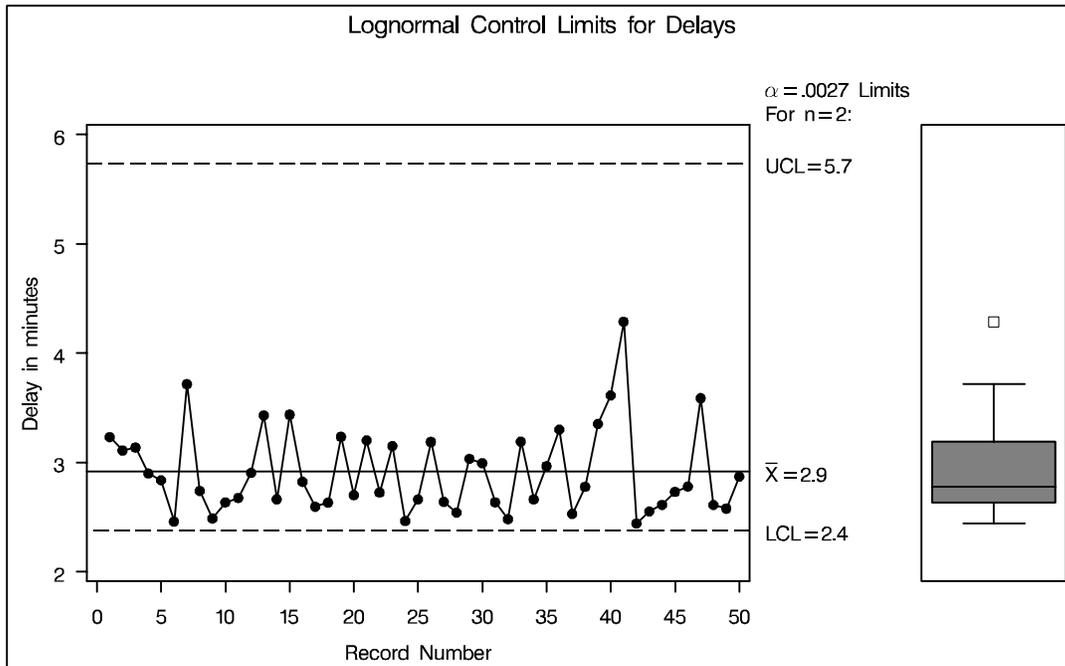


Figure 49.29. Adjusted Control Limits for Delays

Clearly the process is in control, and the control limits (particularly the lower limit) are appropriate for the data. The particular probability level $\alpha = 0.0027$ associated with these limits is somewhat immaterial, and other values of α such as 0.001 or 0.01 could be specified with the ALPHA= option in the original IRCHART statement.

Multivariate Control Charts

In many industrial applications, the output of a process characterized by p variables that are measured simultaneously. Independent variables can be charted individually, but if the variables are correlated, a multivariate chart is needed to determine whether the process is in control.

See SHWT2
in the SAS/QC
Sample Library

Many types of multivariate control charts have been proposed; refer to Alt (1985) for an overview. Denote the i^{th} measurement on the j^{th} variable as X_{ij} for $i = 1, 2, \dots, n$, where n is the number of measurements, and $j = 1, 2, \dots, p$. Standard practice is to construct a chart for a statistic T_i^2 of the form

$$T_i^2 = (\mathbf{X}_i - \bar{\mathbf{X}}_n)' \mathbf{S}_n^{-1} (\mathbf{X}_i - \bar{\mathbf{X}}_n)$$

where

$$\bar{X}_j = \frac{1}{n} \sum_{i=1}^n X_{ij}, \quad \mathbf{X}_i = \begin{bmatrix} X_{i1} \\ X_{i2} \\ \vdots \\ X_{ip} \end{bmatrix}, \quad \bar{\mathbf{X}}_n = \begin{bmatrix} \bar{X}_1 \\ \bar{X}_2 \\ \vdots \\ \bar{X}_p \end{bmatrix}$$

and

$$\mathbf{S}_n = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{X}_i - \bar{\mathbf{X}}_n)(\mathbf{X}_i - \bar{\mathbf{X}}_n)'$$

It is assumed that \mathbf{X}_i has a p -dimensional multivariate normal distribution with mean vector $\mu = (\mu_1 \mu_2 \cdots \mu_p)'$ and covariance matrix Σ for $i = 1, 2, \dots, n$. Depending on the assumptions made about the parameters, a χ^2 , Hotelling T^2 , or beta distribution is used for T_i^2 , and the percentiles of this distribution yield the control limits for the multivariate chart.

In this example, a multivariate control chart is constructed using a beta distribution for T_i^2 . The beta distribution is appropriate when the data are individual measurements (rather than subgrouped measurements) and when μ and Σ are estimated from the data being charted. In other words, this example illustrates a start-up phase chart where the control limits are determined from the data being charted.

Calculating the Chart Statistic

In this situation, it was shown by Gnanadesikan and Kettenring (1972), using a result of Wilks (1962), that T_i^2 is exactly distributed as a multiple of a variable with a beta distribution. Specifically,

$$T_i^2 \sim \frac{(n-1)^2}{n} B\left(\frac{p}{2}, \frac{n-p-1}{2}\right)$$

Tracy, Young, and Mason (1992) used this result to derive initial control limits for a multivariate chart based on three quality measures from a chemical process in the

Part 9. The CAPABILITY Procedure

start-up phase: percent of impurities, temperature, and concentration. The remainder of this section describes the construction of a multivariate control chart using their data, which are given here by the data set STARTUP.

```
data startup;
  input sample impure temp conc;
  label sample = 'Sample Number'
        impure = 'Impurities'
        temp   = 'Temperature'
        conc   = 'Concentration' ;
  datalines;
1  14.92  85.77  42.26
2  16.90  83.77  43.44
3  17.38  84.46  42.74
4  16.90  86.27  43.60
5  16.92  85.23  43.18
6  16.71  83.81  43.72
7  17.07  86.08  43.33
8  16.93  85.85  43.41
9  16.71  85.73  43.28
10 16.88  86.27  42.59
11 16.73  83.46  44.00
12 17.07  85.81  42.78
13 17.60  85.92  43.11
14 16.90  84.23  43.48
;
```

In preparation for the computation of the control limits, the sample size is calculated and parameter variables are defined.

```
proc means data=startup noprint ;
  var impure temp conc;
  output out=means n=n;

data startup;
  if _n_ = 1 then set means;
  set startup;
  p       = 3;
  _subn_  = 1;
  _limitn_ = 1;
```

Next, the PRINCOMP procedure is used to compute the principal components of the variables and save them in an output data set named PRIN.

```
proc princomp data=startup out=prin outstat=scores std cov;
  var impure temp conc;
run;
```

The following statements compute T_i^2 and its exact control limits, using the fact that T_i^2 is the sum of squares of the principal components.* Note that these statements

*Refer to Jackson (1980).

create several special SAS variables so that the data set PRIN can subsequently be read as a TABLE= input data set by the SHEWHART procedure. These special variables begin and end with an underscore character. The data set PRIN is listed in Figure 49.30.

```

data prin (rename=(tsquare=_subx_));
  length _var_ $ 8 ;
  drop prin1 prin2 prin3 _type_ _freq_;
  set prin;
  comp1 = prin1*prin1;
  comp2 = prin2*prin2;
  comp3 = prin3*prin3;
  tsquare = comp1 + comp2 + comp3;
  _var_ = 'tsquare';
  _alpha_ = 0.05;
  _lclx_ = ((n-1)*(n-1)/n)*betainv(_alpha_/2, p/2, (n-p-1)/2);
  _mean_ = ((n-1)*(n-1)/n)*betainv(0.5, p/2, (n-p-1)/2);
  _uclx_ = ((n-1)*(n-1)/n)*betainv(1-_alpha_/2, p/2, (n-p-1)/2);
  label tsquare = 'T Squared'
        comp1 = 'Comp 1'
        comp2 = 'Comp 2'
        comp3 = 'Comp 3';
run;

```

T2 Chart For Chemical Example									
var	n	sample	impure	temp	conc	p	_subn_	_limitn_	comp1
tsquare	14	1	14.92	85.77	42.26	3	1	1	0.79603
tsquare	14	2	16.90	83.77	43.44	3	1	1	1.84804
tsquare	14	3	17.38	84.46	42.74	3	1	1	0.33397
tsquare	14	4	16.90	86.27	43.60	3	1	1	0.77286
tsquare	14	5	16.92	85.23	43.18	3	1	1	0.00147
tsquare	14	6	16.71	83.81	43.72	3	1	1	1.91534
tsquare	14	7	17.07	86.08	43.33	3	1	1	0.58596
tsquare	14	8	16.93	85.85	43.41	3	1	1	0.29543
tsquare	14	9	16.71	85.73	43.28	3	1	1	0.23166
tsquare	14	10	16.88	86.27	42.59	3	1	1	1.30518
tsquare	14	11	16.73	83.46	44.00	3	1	1	3.15791
tsquare	14	12	17.07	85.81	42.78	3	1	1	0.43819
tsquare	14	13	17.60	85.92	43.11	3	1	1	0.41494
tsquare	14	14	16.90	84.23	43.48	3	1	1	0.90302
comp2	comp3	_subx_	_alpha_	_lclx_	_mean_	_uclx_			
10.1137	0.01606	10.9257	0.05	0.24604	2.44144	7.13966			
0.0162	0.17681	2.0410	0.05	0.24604	2.44144	7.13966			
0.1538	5.09491	5.5827	0.05	0.24604	2.44144	7.13966			
0.3289	2.76215	3.8640	0.05	0.24604	2.44144	7.13966			
0.0165	0.01919	0.0372	0.05	0.24604	2.44144	7.13966			
0.0645	0.27362	2.2534	0.05	0.24604	2.44144	7.13966			
0.4079	0.44146	1.4354	0.05	0.24604	2.44144	7.13966			
0.1729	0.73939	1.2077	0.05	0.24604	2.44144	7.13966			
0.0001	0.44483	0.6766	0.05	0.24604	2.44144	7.13966			
0.0004	0.86364	2.1692	0.05	0.24604	2.44144	7.13966			
0.0274	0.98639	4.1717	0.05	0.24604	2.44144	7.13966			
0.0823	0.87976	1.4003	0.05	0.24604	2.44144	7.13966			
1.6153	0.30167	2.3320	0.05	0.24604	2.44144	7.13966			
0.0001	0.00010	0.9032	0.05	0.24604	2.44144	7.13966			

Figure 49.30. The Data Set PRIN

You can now use the data set PRIN as input to the SHEWHART procedure to create the multivariate control chart displayed in Figure 49.31.

```

symbol value=dot;
title 'T' m=(+0,+0.5) '2'
      m=(+0,-0.5) ' Chart For Chemical Example';
proc shewhart table=prin;
  xchart tsquare*sample /
    xsymbol = mu
    nolegend ;
run;

```

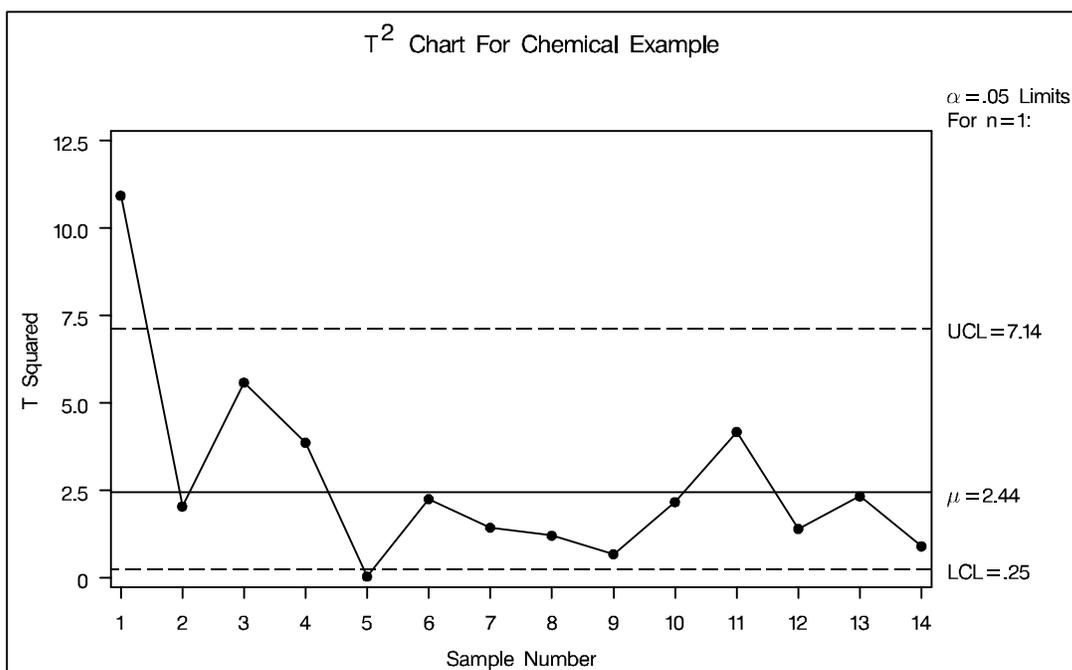


Figure 49.31. Multivariate Control Chart for Chemical Process

The methods used in this example easily generalize to other types of multivariate control charts. You can create charts using the χ^2 and F distributions by using the appropriate CINV or FINV function in place of the BETAINV function in the statements on page 1785. For details, refer to Alt (1985), Jackson (1980, 1991), and Ryan (1989).

Examining the Principal Component Contributions

You can use the *star options* in the SHEWHART procedure to superimpose points on the chart with stars whose vertices represent standardized values of the squares of the three principal components used to determine T_i^2 .

```

symbol value=none;
proc shewhart table=prin;
  xchart tsquare*sample /
    starvertices = (comp1 comp2 comp3)
    startype     = wedge
    cstars       = black
    starlegend   = none
    starlabel    = first
    staroutradius = 4
    npanelpos    = 14
    xsymbol      = mu
    nolegend ;
run;

```

The chart is displayed in Figure 49.32. In situations where the principal components have a physical interpretation, the star chart can be a helpful diagnostic for determining the relative contributions of the different components.

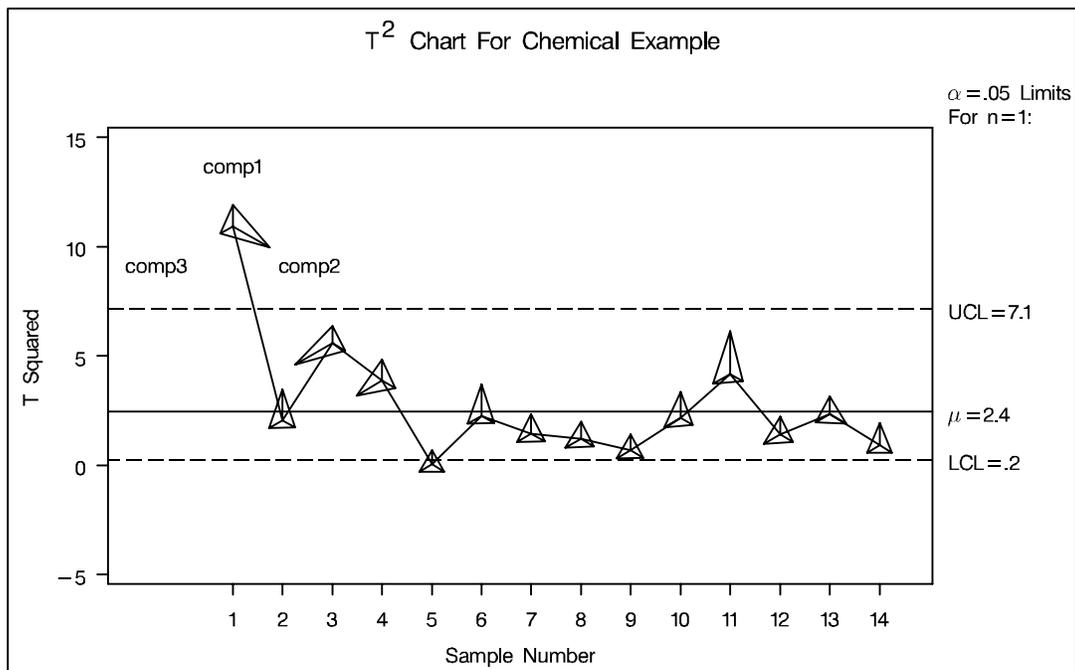


Figure 49.32. Multivariate Control Chart Displaying Principal Components

For more information about star charts, see “Displaying Auxiliary Data with Stars” on page 1701, or consult the entries for the STARVERTICES= and related options in Chapter 46, “Dictionary of Options.”

Principal components are not the only approach that can be used to interpret multivariate control charts. This problem has recently been studied by a number of authors, including Doganaksoy and others (1991), Hawkins (1991, 1993), and Mason and others (1993).

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *SAS/QC[®] User's Guide, Version 8*, Cary, NC: SAS Institute Inc., 1999. 1994 pp.

SAS/QC[®] User's Guide, Version 8

Copyright © 1999 SAS Institute Inc., Cary, NC, USA.

ISBN 1-58025-493-4

All rights reserved. Printed in the United States of America. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, by any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of the software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, October 1999

SAS[®] and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute in the USA and other countries. [®] indicates USA registration.

IBM[®], ACF/VTAM[®], AIX[®], APPN[®], MVS/ESA[®], OS/2[®], OS/390[®], VM/ESA[®], and VTAM[®] are registered trademarks or trademarks of International Business Machines Corporation. [®] indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

The Institute is a private company devoted to the support and further development of its software and related services.