

Paper 002-29

**SAS® SPM Solution for Healthcare: Quality Improvement for Providers  
Using Statistical Process Control**Robert N. Rodriguez and Susan B. Lewellen  
SAS Institute Inc., Cary, NC**ABSTRACT**

Across the country, significant issues threaten the ability of hospitals to meet their commitments to their communities. These issues include the problems of retaining qualified staff and identifying incompetent staff, increasing costs of staff and supplies, and public pressure. Institute of Medicine studies show that over half of medical deaths in hospitals are preventable, and statewide data reveal variability in hospital quality.

The healthcare industry generates large amounts of patient-specific data. However, few hospitals can use the data to identify unusual variability in staff and physician performance, cost of care, and preventable incidents that affect the outcome of a patient's care. SAS SPM for Healthcare provides the ability to access multiple data sources and create analysis-ready data; refer to SAS Institute Inc. (2004a). As illustrated in this paper, statistical process control (SPC) can then be used to identify variability due to special causes and focus further study to reduce variability. These techniques lead to improvements in quality of care, reduction of costs, opportunities to grow market share, and negotiation of better third-party payment.

This paper provides examples that explain the use of SAS statistical software to analyze health care data with  $u$  charts,  $p$  charts, control charts for individual measurements, methods for discovering trends over time, basic forecasting methods, comparative histograms, analysis of means for rates and proportions, and model-based adjustments of mortality rates.

**BACKGROUND**

Quality improvement programs have long been a part of the work of a hospital. However, since 2000, when the publication of *To Err Is Human* (Institute of Medicine 2000) exposed the impact of hospital errors on patients, providers and the public have come to understand more of the nature of hospital errors. The report states that the hospital industry is 10 years behind the aviation industry in the prevention of errors. While the hospital industry is not in agreement as to the extent of the errors, it is in agreement that serious patient safety issues exist and are difficult to assess due to multiple sources of patient care data, as well as limited use of information technology for recording, storing, accessing, and analyzing patient care data.

In *Crossing the Quality Chasm* (Institute of Medicine 2001), the follow-up report from the Institute of Medicine, the authors list six areas in which healthcare systems function at levels that are far lower than they should be:

- Avoiding unintended injuries to patients
- Providing evidence-based services where needed
- Ensuring that patients' values are respected
- Reducing harmful delays that affect both patients and providers
- Avoiding waste of materials and time
- Providing a consistent level of care to all patients

Examples of these functions include the proper administration of medications, prompt access to care, and protection from infections.

In an effort to improve performance in these dimensions, many hospitals have turned to the quality management philosophies of W. Edwards Deming, Joseph Juran, and others, which have been successful in the manufacturing sector. The premise of the Deming approach is that in order to make continual improvements, top management must measure, understand, and act upon the variability in business processes. Wheeler and Polling (1998) characterize this approach as follows: “Instead of focusing on outcomes, such as expenses and profits, this better way focuses on the processes and systems that generate the outcomes. Rather than trying to directly manipulate the results, it works to improve the system that causes the results. Rather than distorting . . . the data, it seeks to use the data to understand the system as a basis for improving the system.”

The analytical cornerstone of the Deming approach is “statistical thinking,” which starts with the recognition that all processes are subject to variability and that improvement comes about through understanding and reduction of variability. *Statistical process control* (SPC) is the use of statistical methods to distinguish and eliminate special (sporadic) causes of variation and subsequently reduce common causes of variation, which are inherent in the system.

Without careful use of SPC, organizations waste valuable time by overreacting to common cause variation rather than eliminating special causes and bringing processes into a state of stability (statistical control). When a process is stable, the variation is predictable rather than chaotic, making it possible to implement improvements that reduce common cause variation and change the mean level of the process. Lilford et al. (2003) states, “Process measurement enables a whole service to improve because most providers lie in the middle of any distribution and more gain can be achieved by shifting the means than by truncating the tail.” Because intervention and change can be assessed on a meaningful scale, sustained improvement becomes a realizable goal.

In their 1990 book *Curing Health Care*, Berwick, Godfrey, and Roessner (1990) provide excellent examples of the use of the basic tools of quality improvement and especially the use of SPC for improving hospital processes. McFadden, Towell, and Stock (2004) describe the need for statistical analysis of data in their model for controlling and managing hospital errors.

This paper provides a series of tutorial examples that use hospital data to illustrate the effectiveness of SPC as a management tool for identifying problems and making informed decisions. The paper is written with two audiences in mind:

- hospital executives, especially CMOs, Directors of Quality Management, Directors of Decision Support, and others who make decisions based on analytical reports about hospital quality. The scenarios and graphical displays in the examples should be of particular interest to these readers.
- hospital quality analysts and information technology specialists who prepare reports. The technical details of the examples should help SAS programmers get started using appropriate SAS procedures for creating charts.

Although SPC is often referred to as “the art of making control charts,” making the *right* chart is only the beginning—and it is ultimately management response to a chart that determines the effectiveness of SPC.

## EXAMPLES

The examples that follow are presented in increasing order of analytical complexity, beginning with simple control charts and progressing to statistical models for prediction. Common to all the examples is the need for making meaningful comparisons—either across time or across groups, and so careful adjustments of data, rates, and decision limits play an important role. A second theme is the creation of graphical displays that provide rich visual content in presentations of analytical information.

The basic examples emphasize the use of the SHEWHART and ANOM procedures in SAS/QC® software for statistical quality improvement. The SHEWHART procedure creates control charts, which are the fundamental SPC method for deciding whether a process is in control and for monitoring an in-control process. The ANOM procedure, newly added in SAS/QC 9.1, provides a simple but highly interpretable display for

comparing responses across groups and identifying any that are significantly different. Other examples illustrate the use of statistical modeling procedures in SAS/STAT® and SAS/ETS® software.

An excellent reference on SAS Web tools, many of which can be used to deliver SPC content in healthcare solutions, is the recent book by Pratter (2003).

### EXAMPLE 1: BASIC *U* CHART FOR RATE OF CAT SCANS

This example introduces the use of the SHEWHART procedure to construct a *u* chart, which is one of several control charts for count data. In manufacturing, *u* charts are typically used to analyze the number of defects per inspection unit in samples that contain arbitrary numbers of units. However, in general, the event that is counted need not be a “defect.” A *u* chart is applicable when the counts can be scaled by some measure of opportunity for the event to occur, and when the counts can be modeled statistically by the Poisson distribution. The SHEWHART syntax for this example is described in detail since it extends to other types of control charts that can be constructed with the procedure.

A health care provider uses a *u* chart to analyze the rate of CAT scans performed each month by each of its clinics. Figure 1 shows data collected for Clinic B and saved in a SAS data set named CLINICB.

CAT Scan Data for Clinic B				
month	nscanb	mmsb	days	nyrsb
JAN94	50	26838	31	2.31105
FEB94	44	26903	28	2.09246
MAR94	71	26895	31	2.31596
APR94	53	26289	30	2.19075
MAY94	53	26149	31	2.25172
JUN94	40	26185	30	2.18208
JUL94	41	26142	31	2.25112
AUG94	57	26092	31	2.24681
SEP94	49	25958	30	2.16317
OCT94	63	25957	31	2.23519
NOV94	64	25920	30	2.16000
DEC94	62	25907	31	2.23088
JAN95	67	26754	31	2.30382
FEB95	58	26696	28	2.07636
MAR95	89	26565	31	2.28754

Figure 1. SAS Data Set CLINICB

The variable NSCANB is the number of CAT scans performed each month, and the variable MMSB is the number of members enrolled each month (in units of “member months”). The variable DAYS is the number of days in each month. The following SAS statements compute the variable NYRSB, which converts MMSB to units of “thousand members per year.”

```
data clinicb; set clinicb;
  nyrsb = mmsb * ( days / 30 ) / 12000;
run;
```

Note that NYRSB provides the “measure of opportunity,” which corresponds to the number of inspection units in manufacturing applications.

The following statements create the *u* chart in Figure 2.

```
title 'U Chart for CAT Scans per 1,000 Members: Clinic B';
proc shewhart data=clinicb;
  uchart nscanb * month / subgroupn = nyrsb tests = 1 to 4 nohlabel
          testnmethod = standardize nolegend;
  label nscanb = 'Rate per 1,000 Member-Years';
run;
```

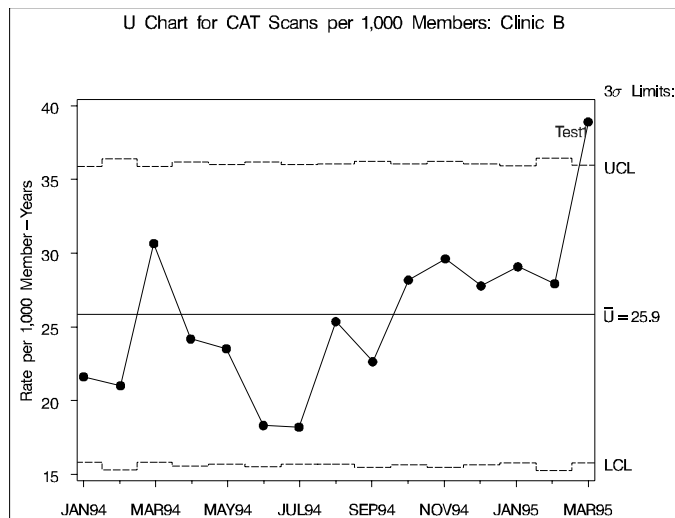
The PROC SHEWHART statement invokes the SHEWHART procedure. The DATA= option specifies the input data set.

The UCHART statement requests a  $u$  chart. After the keyword UCHART, you specify the *process* or count variable to analyze (in this case, NSCANB), followed by an asterisk and the *subgroup-variable* that identifies the sample (in this case, MONTH).

The SUBGROUPN= option specifies the number of “opportunity” units per sample. You can use this option to specify a fixed number of units or (as in this case) a variable whose values provide the number of units for each sample.

You can specify options for analysis and graphical presentation after the slash (/) in the UCHART statement. Refer to SAS Institute Inc. (2004c) for details on syntax and statistical methods. The TESTS= option requests tests for special causes, also referred to as runs tests, pattern tests, and Western Electric rules. For example, Test 1 flags points outside of the control limits. The TESTNMETHOD=STANDARDIZE option applies a standardization method to adjust for the fact that the number of units varies from sample to sample.

The NOHLABEL option suppresses the label for the horizontal axis (which is unnecessary since MONTH has a datetime format), and the NOLEGEND option suppresses the default sample size legend. The LABEL statement assigns a temporary label to the variable NSCANB that is displayed on the vertical axis.



**Figure 2.** Basic  $u$  Chart

In Figure 2, the control limits shown are  $3\sigma$  limits estimated by default from the data; the limits vary because the number of opportunity units changes from month to month. The increase in the rate of CAT scans for March 94 is interpreted as common cause variation since it lies within the control limits, whereas the increase for March 95 should be investigated.

You can use the SHEWHART procedure to create a wide variety of control charts. Each of the standard chart types is created with a different chart statement (for instance, you use the PCHART statement to create  $p$  charts). Once you have learned the basic syntax for a particular chart statement, you can use the same syntax for all the other chart statements.

#### EXAMPLE 2: CONTROL LIMITS FOR A $U$ CHART WITH A KNOWN SHIFT IN RATE

This example illustrates the construction of a  $u$  chart in situations where the process rate is known to have shifted, requiring the use of multiple sets of control limits.

A health care provider uses a  $u$  chart to report the rate of office visits performed each month by each of its clinics. The rate is computed by dividing the number of visits by the membership expressed in thousand-

member years. Figure 3 shows data collected for Clinic E and saved in a SAS data set named CLINICE.

Office Visit Data for Clinic E					
month	_phase_	nvisite	nyrse	days	mmse
JAN94	Phase 1	1421	0.66099	31	7676
FEB94	Phase 1	1303	0.59718	28	7678
MAR94	Phase 1	1569	0.66219	31	7690
APR94	Phase 1	1576	0.64608	30	7753
MAY94	Phase 1	1567	0.66779	31	7755
JUN94	Phase 1	1450	0.65575	30	7869
JUL94	Phase 1	1532	0.68105	31	7909
AUG94	Phase 1	1694	0.68820	31	7992
SEP94	Phase 2	1721	0.66717	30	8006
OCT94	Phase 2	1762	0.69612	31	8084
NOV94	Phase 2	1853	0.68233	30	8188
DEC94	Phase 2	1770	0.70809	31	8223
JAN95	Phase 2	2024	0.78215	31	9083
FEB95	Phase 2	1975	0.70684	28	9088
MAR95	Phase 2	2097	0.78947	31	9168

**Figure 3.** SAS Data Set CLINICE

The variable NVISITE is the number of visits each month, and the variable MMSE is the number of members enrolled each month (in units of “member months”). The variable DAYS is the number of days in each month. The variable NYRSE expresses MMSE in units of thousand members per year. The variable \_PHASE\_ separates the data into two time phases since a change in the system is known to have occurred in September 1994 at the beginning of Phase 2.


The following statements create a *u* chart with a single set of default limits.

```

title 'U Chart for Office Visits per 1,000 Members: Clinic E';
proc shewhart data=clinice;
  uchart nvisite * month / subgroupn = nyrse   cframe = ligr  nohlabel
        cinfill = yellow  nolegend;
  label nvisite = 'Rate per 1,000 Member-Years';
run;

```

The CFRAME= option specifies the color for the plot area, and the CINFILL= option specifies the color for the area between the limits. The chart is shown in Figure 4.



**Figure 4.** *u* Chart with Single Set of Limits

The default control limits are clearly inappropriate because they do not allow for the shift in the average rate that occurred in September 1994.

The following statements use BY processing to compute distinct sets of control limits from the data in each phase and save the control limit information in a SAS data set named VISLIMIT. The NOCHART option is specified to suppress the display of separate control charts for each phase.

```
proc shewhart data=clinice; by _phase_;
    uchart nvisite * month / subgroupn = nyrse nochart
        outlimits = vislimit (rename=( _phase_ = _index_ ));
run;
```

Figure 5 shows a listing of VISLIMIT. Note that the values of the lower and upper control limit variables `_LCLU_` and `_UCLU_` are equal to the special missing value V; this indicates that these limits are varying. The variable `_INDEX_` identifies the control limits in the same way that the variable `_PHASE_` identifies the time phases in the data.

Control Limits for Office Visit Data									
<code>_index_</code>	<code>_VAR_</code>	<code>_SUBGRP_</code>	<code>_TYPE_</code>	<code>_LIMITN_</code>	<code>_ALPHA_</code>	<code>_SIGMAS_</code>	<code>_LCLU_</code>	<code>_U_</code>	<code>_UCLU_</code>
Phase 1	nvisite	month	ESTIMATE	V	V	3	V	2302.99	V
Phase 2	nvisite	month	ESTIMATE	V	V	3	V	2623.52	V

Figure 5. SAS Data Set VISLIMIT

The following statements combine the data and control limits for both phases in a single *u* chart, shown in Figure 6.

```
title 'U Chart for Office Visits per 1,000 Members: Clinic E';
proc shewhart data=clinice limits=vislimit;
    uchart nvisite * month / subgroupn = nyrse
        cframe = ligr
        cinfill = yellow
        readindex = all
        readphase = all
        nohlabel nolegend
        phaselegend nolimitslegend;
    label nvisite = 'Rate per 1,000 Member-Years';
run;
```

The READINDEX= and READPHASE= options match the control limits in VISLIMIT with observations in CLINICE by the values of the variables `_INDEX_` and `_PHASE_`, respectively.

In Figure 6, no points are out of control, indicating that the variation is due to common causes after adjusting for the shift in September 1994.

Note that both sets of control limits in Figure 6 were estimated from the data with which they are displayed. You can, however, apply pre-established control limits from a LIMITS= data set to new data.

In applications involving count data, control charts for individual measurements can sometimes be used in place of *u* charts and *c* charts, which are based on a Poisson model, as well as *p* charts and *np* charts, which are based on a binomial model. Wheeler (1995) makes the point that charts based on a theoretical model “allow one to detect departures from the theoretical model,” but they require verification of the assumptions required by the model. On the other hand, charts for individual measurements often provide reasonably approximate empirical control limits, as illustrated in the next example.




Figure 6.  $\bar{u}$  Chart with Multiple Sets of Control Limits

### EXAMPLE 3: INDIVIDUAL MEASUREMENT CHART FOR ADJUSTED UTILIZATION RATES

A clinic uses a chart for individual measurements to analyze the number of medical/surgical days per 1,000 members per year. Figure 7 shows a partial listing of a SAS data set named MEDSURG that contains this information. The variable MSAD\_E provides the medical/surgical utilization rate for Product E, a new benefits plan that was introduced in January 1993, and the variable MSAD\_OTH provides the rate for all other products. The variable \_PHASE\_ breaks the data into time phases. It was originally expected that the rate for Product E would start out equal to that of the other products and would increase over time.

Medical/Surgical Admissions Data			
month	msad_e	msad_oth	_phase_
JAN91	.	151.936	Historical
FEB91	.	136.286	Historical
MAR91	.	236.516	Historical
.	.	.	.
NOV92	.	260.200	Historical
DEC92	.	183.097	Historical
JAN93	618.290	269.807	New Product E
FEB93	367.393	125.571	New Product E
.	.	.	.
JUL94	128.032	129.581	New Product E
AUG94	203.323	180.194	New Product E
SEP94	318.000	109.900	Younger Members
OCT94	109.645	139.645	Younger Members
.	.	.	.
APR95	78.700	102.400	Younger Members
MAY95	65.033	212.613	Younger Members
JUN95	112.800	137.400	Younger Members

Figure 7. SAS Data Set MEDSURG (Partial Listing)

The following step uses the IRCHART statement in the SHEWHART procedure to construct an individual measurement and moving range chart for the historical rate of the other products prior to the introduction of Product E.

```

title 'Historical Medical/Surgical Rate of Other Products';
symbol v=dot h=2.4pct;
proc shewhart data=medsurg; where month < '01jan93'd;
  irchart msad_oth * month / cframe = yellow npanel = 100
    split = '/' nohlabel;
  label msad_oth = 'Days per 1,000/Mvg Rng';
run;

```

The chart, shown in Figure 8, indicates that the utilization for the other products is a stable, predictable process.




Figure 8. Historical Utilization Rates

Now, consider a comparison between the other products and the new product. Begin by computing control limits for the rates for each product and for each of the time phases.

```
proc shewhart data=medsurg; by _phase_ notsorted;
  irchart (msad_oth msad_e) * month / nochart
          outlimits = runlim (rename=( _phase_ = _index_ ));
data runlim; set runlim;
  _lcli_ = max( _lcli_, 0 );
run;
```

The control limits are saved in the SAS data set RUNLIM, which is listed in Figure 9.

_index_	_VAR_	_SUBGRP_	_TYPE_	_LIMITN_	_ALPHA_	_SIGMAS_
Historical	msad_oth	month	ESTIMATE	2	.002699796	3
New Product E	msad_oth	month	ESTIMATE	2	.002699796	3
New Product E	msad_e	month	ESTIMATE	2	.002699796	3
Younger Members	msad_oth	month	ESTIMATE	2	.002699796	3
Younger Members	msad_e	month	ESTIMATE	2	.002699796	3
_LCLI_	_MEAN_	_UCLI_	_LCLR_	_R_	_UCLR_	_STDDEV_
0.00000	156.774	324.487	0	63.081	206.057	55.904
7.94440	149.929	291.913	0	53.404	174.446	47.328
0.00000	287.773	979.441	0	260.154	849.802	230.556
0.00000	134.593	285.900	0	56.911	185.900	50.436
0.00000	138.470	383.816	0	92.281	301.439	81.782

Figure 9. SAS Data Set RUNLIM

The following statements read RUNLIM to create the control chart for the rate for Product E that is shown in Figure 10. The NOLCL option suppresses the lower control limit, which is zero. The NOCHART2 option suppresses the chart for moving ranges, which are needed for the computations but add little or no information to the display.



```

title 'Medical/Surgical Rate for Product E';
symbol v=dot h=2.4 pct;
proc shewhart data=medsurg limits=runlim; where month >= '01jan93'd;
    irchart msad_e * month / nolcl          nohlabel
                                nochart2    phaselegend
                                phaselabtype = scaled
                                readindex    = all
                                readphase    = all
                                cframe       = yellow
                                npanel      = 100;

    label msad_e = 'Med/Surg Days per 1,000';
run;

```

Figure 10 reveals that the rate for Product E dropped in October 1994. Subsequent investigation showed that a large number of younger and healthier members began using the product at this point. Prior to this time the membership was small and varied, which accounts for the high variability in the rate during the introductory phase.




Figure 10. Medical/Surgical Admissions Rates

#### EXAMPLE 4: TRENDS IN EMERGENCY DEPARTMENT UTILIZATION

A hospital system tracks the utilization of the emergency department over time. Figure 11 shows a partial listing of census data. The variable TOTAL records the total number of hours per 24-hour day spent by patients in the department.

Emergency Department Visits Data	
day	hours
01JAN2002	435
02JAN2002	417
03JAN2002	252
04JAN2002	352
05JAN2002	445
06JAN2002	388
07JAN2002	415
08JAN2002	376
09JAN2002	346
10JAN2002	442

Figure 11. SAS Data Set BYDAY (First 10 Observations)


A simple—but very useful—preliminary step in analyzing utilization rates is to plot them over time. A scatterplot smoothing method such as LOESS can help to visualize trends when the data are noisy. The following statements use the LOESS procedure in SAS/STAT software to compute a smooth fit for the average number of hours per day along with 95% confidence limits for the average.

```
proc loess data=byday;
  ods output OutputStatistics = EDfit;
  model hours = day / smooth=.3 direct alpha=.05 residual all;
run;

title1 "Emergency Department Visits";
title2 "LOESS Fit with 95% Confidence Bands for Average Hours per Day";

proc gplot data=EDFit;
  plot (DepVar Pred LowerCL UpperCL)*day /
    overlay vaxis=axis1 haxis=axis2 cframe=ligr;
  format day monyy.;
  label DepVar="Hours per Day";
  axis1 label=(angle=90 rotate=0) minor=none;
  axis2 label=none minor=none;
run;
```

The scatterplot is displayed in Figure 12.



**Figure 12.** LOESS Smooth of Visits per Day

This display reveals that the average number of hours increased from 500 per day to 640 per day during 2002 and reached a new level of 800 per day during 2003. The average number dropped during the second half of 2003.

A similar trend is evident in the number of hours per month, which is provided by the variable TOTAL in the data set BYMONTH listed in Figure 13. For simplicity, assume that the values of TOTAL have been adjusted to a 30-day month.

When presented with the display in Figure 12, administrators concluded that the increase during 2002 resulted from the closing of a nearby clinic. An investigation of the downturn at the end of 2003 by information technology staff concluded that it is an artifact of a lag in the reporting system (incomplete data rather than a real effect).

Emergency Visits Data by Month				
month	_phase_	hours	ndays	
JAN02	Old Capacity	13.375	744	
...	...	...	...	
AUG02	Old Capacity	15.922	744	
SEP02	Clinic Close	19.489	720	
...	...	...	...	
DEC02	Clinic Close	18.568	744	
JAN03	New Capacity	23.086	744	
...	...	...	...	
AUG03	New Capacity	21.814	744	
SEP03	Lag	18.657	720	
...	...	...	...	
NOV03	Lag	14.999	720	

Figure 13. SAS Data Set BYMONTH (Partial Listing)

After taking this information into account, hospital analysts decide to create a control chart for TOTAL for those periods of time when utilization was determined to be stable. The following statements illustrate how to construct a chart with control limits for selected periods.

```
proc shewhart data=bymonth;
  by _phase_;
  irchart hours*month / outlimits=limits(rename=( _phase_=_index_)) nochart;
run;

proc sort data=bymonth;
  by month _phase_;
run;

data limits;
  set limits;
  if _index_ in('Clinic Closed' 'Lag') then do;
    _mean_ = . ; _lcli_ = . ; _ucli_ = . ;
  end;
run;

title "Emergency Department Utilization";
proc shewhart data=bymonth limits=limits;
  label hours = "Hours per Month (thousands)";
  irchart hours*month / readindex = all    cframe = (ligr white)
                    phaselegend         nochart2
                    phaseref            nohlabel
                    nolimitslegend      phaselabtype = scaled
                    wlimits = 3         llimits = 1
                    cphaseleg = black;
run;
```

The control chart is shown in Figure 14.




Figure 14. Control Limits for Periods of Process Stability

#### EXAMPLE 5: INDIVIDUAL MEASUREMENTS CHART FOR SEASONAL EFFECTS IN UTILIZATION

This section illustrates the use of an individual measurements chart with multiple sets of control limits that adjust for seasonal effects. A partial listing of the data is shown in Figure 15.

A hospital system located in Minnesota uses a chart for individual measurements to analyze monthly variation in the number of emergency room visits per 1,000 member-years.

Emergency Room Visits per 1000 Member Years		
month	_phase_	visits
JAN90	90	92.58
FEB90	90	82.77
MAR90	90	81.26
APR90	90	82.66
MAY90	90	94.97
JUN90	90	100.63
JUL90	90	108.43
AUG90	90	82.88
SEP90	90	91.33
OCT90	W91	74.68
NOV90	W91	75.40
DEC90	W91	78.92
JAN91	W91	74.32
FEB91	W91	80.28
MAR91	W91	79.75

Figure 15. SAS Data Set ERVISIT (Partial Listing)

The variable VISITS provides the rate of emergency visits, and the variable \_PHASE\_ groups the monthly observations into seasonal time phases. Seasonal grouping was not done prior to October 1990 since a new system was introduced at that point, and the average rate was known to have changed.

The following statements create a preliminary display of the data that highlights the seasonal structure of the rates with boxes that enclose the points for each time phase.

```

title 'Emergency Room Visits per 1000 Member Years';
symbol v=dot;

proc shewhart data=ervisit limits=erlimits;
    boxchart visits*month / nochart2          nohlabel
                                nolimits      nolegend

```



```

cframe      = ( yellow ligr )
cphaseleg   = black
cconnect    = black
llimits     = 3
readindex   = all
readphase   = all
phaselabtype = scaled
nolimitslegend phaselegend;
label visits = 'Visits per 1000 Member Years';
run;

```

Figure 17 shows that after adjusting for seasonality, the remaining variability in the rates can be attributed to common causes. It is natural to consider how statistical methods might be used to predict the future behavior of the system, and this is discussed in the section in the next example.




Figure 17. Emergency Room Visits

#### EXAMPLE 6: FORECASTING EMERGENCY ROOM VISITS

In Example 5 multiple sets of control limits were used to adjust for a seasonal effect in the rate of emergency room visits. This section describes the use of two different time series models to analyze the data.

First, the FORECAST procedure with the Winters method is used to generate forecasts and confidence limits for the rate of emergency room visits; for details, refer to SAS Institute Inc. (2004b).

```

proc forecast data      = ervisit2
              interval = month
              method    = winters
              seasons   = month
              lead      = 7
              out       = outval
              outest    = est
              outfull   outresid;

  id date;
  var visits;
run;

```

Next, the forecasts are merged with the original data.

```

data forecast; keep date forecast;
  set outval (rename=(visits=forecast));
  if _type_='FORECAST';
data ervisit2; merge ervisit2 forecast; run;

```

For subsequent display, the forecasts and confidence limits are saved as coordinates in an Annotate data set.

```

%annomac;
data forecast;
  %dclanno; %system(2,2,4);
  line=3; size=1; color='black';
  drop _type_ _lead_;
  set outval (rename = (visits=y date=x));
  if _type_='FORECAST' and x>='01jul95'd;
  if x='01jul95'd then function='MOVE';
  else          function='DRAW';

data 195;
  %dclanno; %system(2,2,4);
  line=1; size=1; color='black';
  drop _type_ _lead_;
  set outval (rename = (visits=y date=x));
  if _type_='L95' and x>='01jul95'd;
  if x='01jul95'd then function='MOVE';
  else          function='DRAW';

data u95;
  %dclanno; %system(2,2,4);
  line=1; size=1; color='black';
  drop _type_ _lead_;
  set outval (rename = (visits=y date=x));
  if _type_='U95' and x>='01jul95'd;
  if x='01jul95'd then function='MOVE';
  else          function='DRAW';
data annotate;
  set forecast 195 u95;
  when='A';
run;

```

Finally, the XCHART statement in the SHEWHART procedure is used to display the forecast values and the confidence intervals. A plot of the residuals (the differences between the observed rates and the forecasted rates) is aligned above the forecast plot, and control limits for individual measurements based on moving ranges are displayed for the residuals.

```

title 'Observed and Forecasted Emergency Room Visits';
proc shewhart data=ervisit2;
  xchart visits * date / cframe = ligr
                    cconnect = black
                    npanel = 100
                    trendvar = forecast
                    split = '/'
                    anno2 = annotate
                    ypct1 = 50
                    nolegend nohlabel;
  label visits = 'Residual/Visits per 1000 Years';
run;

```

The display, shown in Figure 18, shows that after adjusting for seasonal and trend effects, only common cause variation is evident in the rate of visits. The forecast plot indicates a drop in the rate of visits at the end of 1995. Refer to Alwan and Roberts (1988) for discussion of a similar approach to dealing with time series effects in SPC.

You can also use the X11 procedure to seasonally adjust the emergency room data; for details, refer to SAS Institute Inc. (2004b). The X11 procedure models the observed rate at time  $t$  as  $O_t = S_t C_t D_t I_t$ . Here,  $C_t$ , the long-term trend cycle component, has the same scale as the data  $O_t$ , and  $S_t$  (the seasonal or intrayear component),  $D_t$  (the trading-day component), and  $I_t$  (the residual component) vary around 100 percent.




Figure 18. FORECAST Analysis

The following statements create a plot of the original and seasonally adjusted series ( $C_t I_t$ ).

```
proc x11 data=ervisit2 noprint;
  monthly date=date; var visits;
  output out=out b1 = visits d10 = seasonal d11 = adjusted
          d12 = trend d13 = irreg;
run;

title h=4 pct 'Emergency Room Visits';
title2 h=4 pct 'Original and Seasonally Adjusted Data';

symbol1 i=join h=2.4 pct c=black v='plus';
symbol2 i=join h=2.4 pct c=black v='diamond';

proc gplot data=out;
  plot visits * date = 1
        adjusted * date = 2 / overlay frame legend = legend1 haxis = axis1
        vaxis = axis2 vminor = 0;

  axis1 label=none;
  axis2 label=(a=90 r=0 'Visits per 1000 Member Years');
  legend1 cborder=black label=none value=('original' 'adjusted');
run;
```

The plot is shown in Figure 19. Adjusting for seasonal variation, the rate of emergency room visits is decreasing over time, with a slight increase late in 1994.



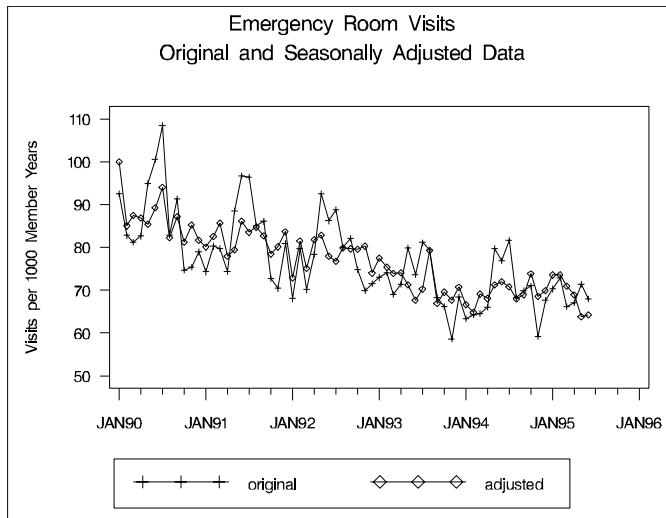


Figure 19. X11 Analysis

The next statements plot the final seasonal factor, as shown in Figure 20.

```

title 'Final Seasonal Series';
symbol i=join;
proc gplot data=out;
  plot seasonal * date / cframe = yellow vaxis = axis1
                        haxis = axis2 vminor = 0;
  axis2 label=none;
  axis1 label=(a=90 r=0 'Seasonal Factor');
run;

```

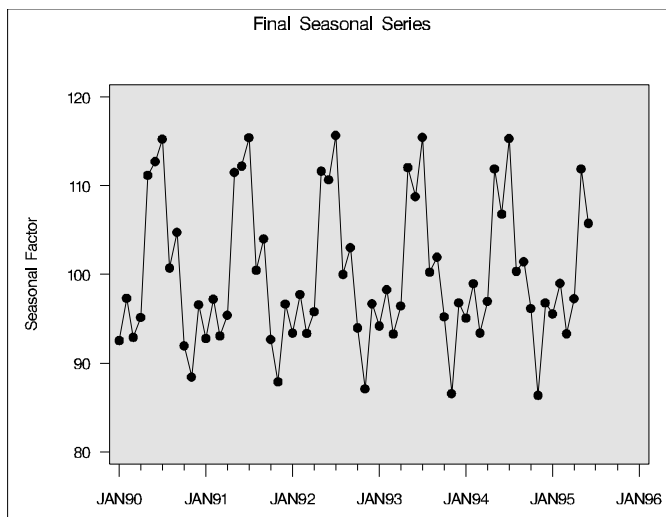


Figure 20. Final Seasonal Factor

The last set of statements combine the final irregular factor and the trend in a single display, shown in Figure 21.

```

data out; set out; sum = irreg + trend; run;
symbol v=none;
title 'Control Chart for Irregular Variation';
proc shewhart data=out;
  xchart sum * date / cframe = yellow npanel = 100 trendvar = trend
                    split = '/' nohlabel nolegend;
  label sum = 'Irregular Series/Trend';
run;

```

Note that the irregular factor is not the same as the residual displayed in Figure 18 since these values were computed using two different time series models. Likewise, the final trend is not the same as the forecast displayed in Figure 18. Nonetheless, both methods provide useful views, understanding, and prediction of the variation in the process.




Figure 21. Final Irregular Factor and Trend

### EXAMPLE 7: ANALYSIS OF MEANS FOR ADMISSIONS RATES

Examples 7 and 8 illustrate the use of analysis of means (ANOM) for rate data.

Analysis of means is a graphical and statistical method for simultaneously comparing averages, rates, or proportion for a set of groups or individuals with their overall mean at a specified significance level  $\alpha$ . Statistically, this method can be thought of as an alternative to analysis of variance for a fixed effects model. Analysis of means can also be thought of as an extension to the Shewhart chart because it considers a group of sample means instead of one mean at a time in order to determine whether any of the sample means differ too much from the overall mean.

In practice, the analysis of means has the same graphical presentation as a control chart except that the decision limits are computed differently. The visual simplicity and content of ANOM are key to its effectiveness in a variety of applications. An up-to-date presentation of analysis of means in the context of designed experiments is the recent textbook by Nelson, Coffin, and Copeland (2003).

A health care system uses ANOM to compare medical/surgical admissions rates for a group of clinics. The data are saved in a SAS data set named MSADMITS, which is listed in Figure 22.

Medical/Surgical Admissions Data		
id	count95	myrs95
1A	1882	58.1003
1K	600	18.7263
1B	438	12.8933
1D	318	6.8545
3M	183	6.3708
3I	220	6.1274
1N	121	5.0141
3H	105	4.4072
1Q	124	4.3829
1E	171	4.2691
3B	88	2.8979
1C	100	2.6633
1H	112	2.3985
3C	84	2.2898
1R	69	2.2078
1T	21	2.0913
1M	130	2.0603
1O	61	2.0438
3D	66	1.8633
1J	54	1.5918
3J	30	1.3408
3G	36	1.1543
3E	26	0.8823
1G	28	0.8626
1I	25	0.5034
1L	20	0.4282
1S	7	0.2269
1F	7	0.2020
1P	2	0.1692

Figure 22. SAS Data Set MSADMITS

The variable ID identifies the clinics, the variable COUNT95 provides the number of admissions during 1995, and the variable MYRS95 provides the number of 1,000 member-years, which serves as the “measure of opportunity” for admissions.

The following statements perform an analysis of means for the admission rates at the  $\alpha = 0.01$  level of significance. The UCHART statement in the ANOM procedure is used to compute the rates and display them graphically with upper and lower decision limits (UDL and LDL).

```

title 'Analysis of Medical/Surgical Admissions';
symbol v=none w=7;

proc anom data=msadmits;
    uchart count95*id / groupn      = myrs95
                    alpha         = 0.01
                    cneedles      = black
                    climits       = black
                    llimits       = 1
                    cinfill       = yellow
                    turnhlabels   = nolegend
                    haxis         = axis1;
    axis1 value = ( h=2.2 pct ) ;
    label count95 = 'Admits per 1000 Member Years';
run;

```

The chart is shown in Figure 23. The needles (requested with the CNEEDLES= option) emphasizes deviations from the overall mean, and the limits UDL and LDL apply to the rates *taken as a group*.

The chart answers the question, “Do any of the clinics differ significantly from the system average in their rates of admission?” The answer is that Clinics 1D and 1M have higher rates that cannot be attributed to chance variation alone. Likewise, Clinic 1T has a lower rate of admission. This answer would be the same regardless of how the clinics were ordered from left to right on the chart. The reason that the decision limits

flare out monotonically from left to right is that the clinics happen to be displayed in decreasing order of MYRS95, and the width of the limits is inversely related to the square root of MYRS95.

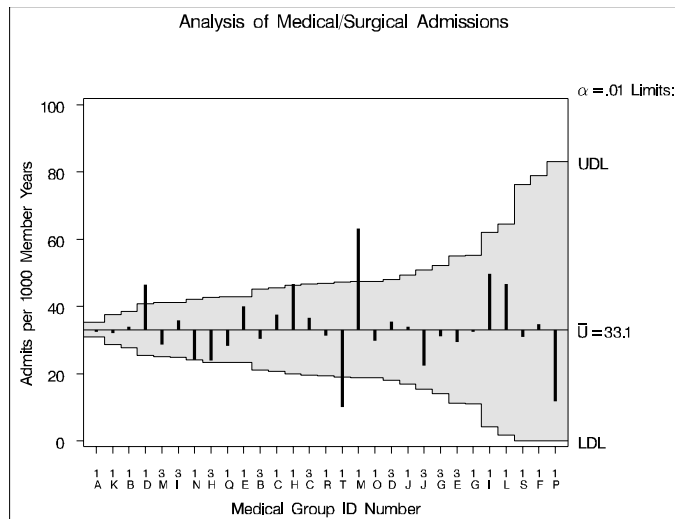


Figure 23. ANOM for Medical/Surgical Admissions Rates

Despite the similarity of Figure 23 to a  $u$  chart, it is important to understand the differences between ANOM and control charting:

- Analysis of means assumes that the system is statistically predictable, whereas a major reason for using a control chart is to bring the system into a state of statistical control; refer to Chapter 1 of Wheeler (1995).
- The decision limits UDL and LDL are not the same as the  $3\sigma$  limits that the SHEWHART procedure would compute by default for a  $u$  chart. The reason is that control limits are applied to the rates *taken one at a time*, whereas the decision limits are applied to the rates *taken as a group*.
- Runs tests, which you could request with the TESTS= option for a control chart, are not applicable in ANOM since the data are not sequential in time.

#### EXAMPLE 8: ANALYSIS OF MEANS FOR PROPORTIONS OF C-SECTIONS AND VBACS

A hospital uses ANOM to compare cesarean section rates for a set of physicians. The data are saved in a SAS data set named BYPHYS (not shown here). The variable PHYSID identifies the physicians, the variable CSECTS provides the number of c-sections in one year, and the variable TOTAL2 provides the total number of deliveries, which serves as the “measure of opportunity” for c-sections.

C-Section/VBACS data				
physid	csects	total	vbacs	total2
BK	64	176	8	38
CE	23	87	4	12
CT	8	25	.	.
FE	40	137	1	18
FT	0	59	.	.
HE	37	160	4	16
JO	46	181	7	29
LT	15	53	.	.
NR	30	157	6	21
OL	35	97	3	13

Figure 24. SAS Data Set CSECTVBAC (First 10 Observations)

The following statements perform an ANOM comparing the proportions of c-sections across physicians at the  $\alpha = 0.01$  level of significance. The PCHART statement in the ANOM procedure is used to compute the proportions and display them graphically with upper and lower decision limits (UDL and LDL).

```

title "C-Section Rates";
proc anom data=csects;
  label physid = "Physician ID";
  label csects = "Proportion of Cesarean Sections";
  pchart csects*physid / groupn = total nolegend cframe = ligr
        turnhlabels cinfill = yellow cneedles = black wneedles = 7;
run;

```

The chart, shown in Figure 25, shows that three physicians were significantly different because they performed no c-sections, while the rate for one physician (BK) was significantly higher than the average.

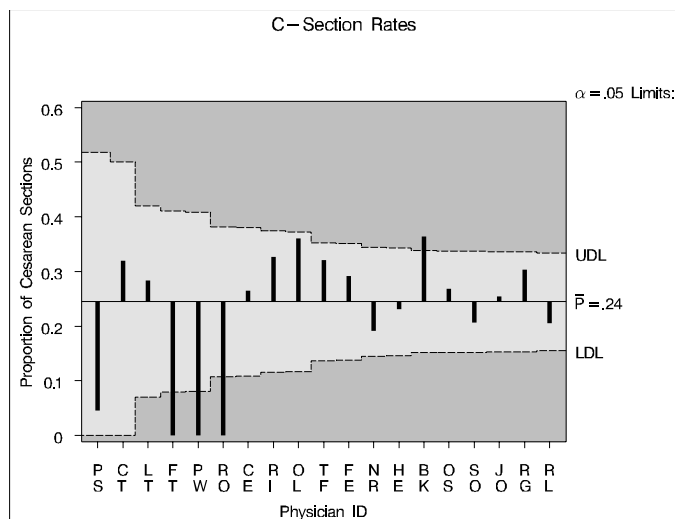


Figure 25. ANOM for Proportion of C-Sections

ANOM is also used to compare VBACS (vaginal birth after cesarean section) rates for the same group of physicians. A hospital uses ANOM to compare cesarean section rates for a set of physicians. The data are saved in a SAS data set named CSECTVBAC, which is listed in Figure 24. The variable PHYSID identifies the physicians, the variable VBACS provides the number of VBACS in one year, and the variable TOTAL2 provides the total number of previous c-sections.

The following statements perform an ANOM comparing the proportions of VBACS across physicians at the  $\alpha = 0.01$  level of significance.

```

title "VBAC Rates";
proc anom data=vbacs;
  label physid = "Physician ID";
  label vbacs = "Proportion of VBACS";
  pchart vbacs*physid / groupn = total2 nolegend cframe = ligr
        cinfill = yellow cneedles = black wneedles = 7;
run;

```

The chart, shown in Figure 26, indicates that the variation in VBAC rates is due to chance.




Figure 26. ANOM for Proportion of VBACS

### EXAMPLE 9: EXPLORATORY ANALYSIS OF LENGTH OF STAY IN HOSPITAL

A hospital administrator is concerned about the possibility that patients with congestive heart failure (CHF) who are readmitted to the hospital may tend to have a longer length of stay if their first length of stay was too short. A partial listing of the data for length of stay is shown in Figure 27.

Distribution of Length of Stay for Readmitted Patients		
los	prevlos	prevtype
4	0	Less than 3 days
3	3	More than 3 days
7	4	More than 3 days
1	5	More than 3 days
2	3	More than 3 days
2	1	Less than 3 days
6	2	Less than 3 days
6	1	Less than 3 days
6	5	More than 3 days
2	6	More than 3 days

Figure 27. SAS Data Set LOSReadmit (First 10 Observations)

The variable LOS records the length of stay for readmitted patients, and the variable PREVLOS records the previous length of stay (PREVTYPE is derived from PREVLOS). Other patient-specific variables are available but are not shown in Figure 27.

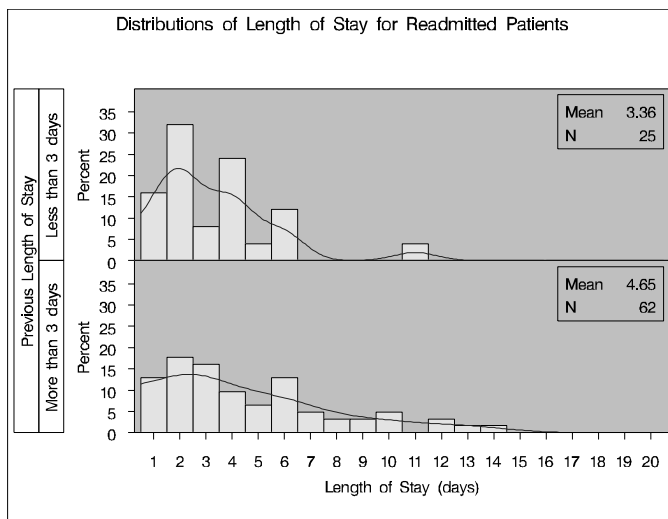
A variety of statistical methods, including regression and multivariate correlation measures, can be used to analyze the data. Here, histograms (a basic SPC tool) are used to compare the distributions of length of stay for patients whose previous stay was less than three days with those whose previous stay was more than three days. The following statements use the CAPABILITY procedure in SAS/QC software to create a comparative histogram display.

```

title "Distributions of Length of Stay for Readmitted Patients";
proc capability data=LOSReadmit noprint; var los;
  label los = "Length of Stay (days)" prevtype = "Previous Length of Stay";
  comphistogram los / class = prevtype midpoints = 1 to 20 by 1 maxnbin = 20
    kernel(lower=0) cframe=ligr cfill=yellow;
  inset mean (5.2) n / position=ne;
run;

```

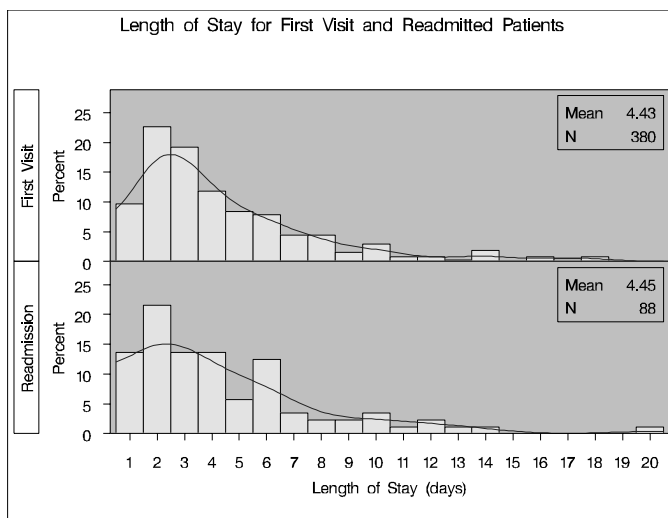
The display is shown in Figure 28. The distributions are smoothed with a nonparametric kernel density estimate as an aid to visualization.



**Figure 28.** Comparative Histogram for Length of Stay

This display shows that the average length of stay is longer (3.4 days) for patients whose previous stay was less than three days than the average (4.7 days) for patients whose previous stay was more than three days. Furthermore, the distribution for the second group has a longer tail. This might be expected if the original length of stay is largely an indicator of severity of illness, and if it is not confounded with other factors. To further explore this assumption, comparisons should be made by splitting the data by levels of other patient variables that might be relevant.

Another question was raised about the difference in length of stay for CHF patients admitted for the first time and for patients who were readmitted. A comparative histogram for these two distributions, presented in Figure 29, shows that the two distributions are very similar, and the averages are remarkably close.



**Figure 29.** Length of Stay Distributions for Original Stay and Readmission

**EXAMPLE 10: RISK-ADJUSTED MORTALITY RATES**

A hospital wants to compare congestive heart failure mortality rates for patients assigned to their physicians after accounting for the severity of the patient's illness at admission and any other relevant factors. The hospital decides to base the comparison on six months of mortality data for its physicians and also has access to statewide mortality data.

One widely used approach is to model the probability of a patient's death and then for each physician divide the observed number of deaths by the expected number to obtain a physician risk-adjustment factor. The product of this adjustment factor with the overall observed mortality rate for the hospital is the physician's risk-adjusted mortality rate. Other approaches are discussed at the end of this example.

To model the probability of death, the hospital creates its own risk-of-mortality score based on the state data. Such scores are typically based on a logistic regression model

$$\text{logit}(p_{ij}) = \log\left(\frac{p_{ij}}{1 - p_{ij}}\right) = \alpha + x_{j1}\beta_1 + \cdots + x_{jm}\beta_m$$

where  $p_{ij}$  is the probability of death for patient  $j$  assigned to physician  $i$ ,  $j = 1 \dots n_i$ , and  $x_1, \dots, x_m$  are descriptive covariates in the model. The score for patient  $j$  is  $x_j\beta = \alpha + x_{j1}\beta_1 + \cdots + x_{jm}\beta_m$ , and the patient's probability of mortality is computed as

$$p_{ij} = \frac{e^{x_j\beta}}{1 + e^{x_j\beta}}$$

The databases contain one observation for each patient with variables (RACE), age (AGE), a physician identifier (PHYSID), and an externally determined risk-of-mortality value (APRRROM) which takes values from 1 for patients who were admitted with a low risk of mortality to 4 for patients who were admitted with severe risk of death. The variable MORTFLAG indicates whether the patient died while under care.

The following statements fit a logistic regression model to the state data and apply the resulting model to the hospital's data. Since the APRROM variable is an ordered categorical variable, it is specified in the CLASS statement with the ORDINAL parameterization where the lowest level is the baseline value. The EVENT= option specifies that the logit of the probability of death is modeled. The SCORE statement applies the state model to the HOSPITAL data and writes the predicted probabilities to the FITOUT data set. Physicians with fewer than 10 patients in the data were not included. The model was validated and goodness of fit was assessed, but these steps are not shown here.

```
proc logistic data=state;
  class aprrom(param=ordinal) race / param=ref;
  model mortflag(event='1') = aprrom race age;
  score data=hospital out=fitout;
run;
```

In the following program, the physician risk-adjustment factor is computed as the ratio of the number of observed deaths (OBSDEATH) to the expected number of deaths (EXPDEATH) for each physician, and the adjusted mortality rate is the product of this factor with the hospital's mortality rate. Confidence intervals (LOWER, UPPER) are computed as described in Hosmer and Lemeshow (1995).

```
data fitout; set fitout; by physid;
  retain obsdeath expdeath var numPatients totDeath 0;
  if first.physid then do;
```



```

        obsdeath=0; expdeath=0; var=0; numPatients=0;
    end;
    obsdeath + mortflag;
    expdeath + P_1;
    var + P_1 * ( 1 - P_1 );
    numPatients + 1;
    alpha=0.05;
    z = quantile('normal',1 - ( alpha / 2 ) );
    if last.physid then do;
        amr = obsdeath / expdeath * &meanDeath;
        lower = ( obsdeath - z * sqrt( var ) ) / expdeath * &meanDeath;
        if lower < 0 then lower = 0;
        upper = ( obsdeath + z * sqrt( var ) ) / expdeath * &meanDeath;
        output;
    end;
    keep physid numPatients lower amr upper;
run;

```

The following statements create a display of the adjusted mortality rate and a 95% confidence interval for each physician.

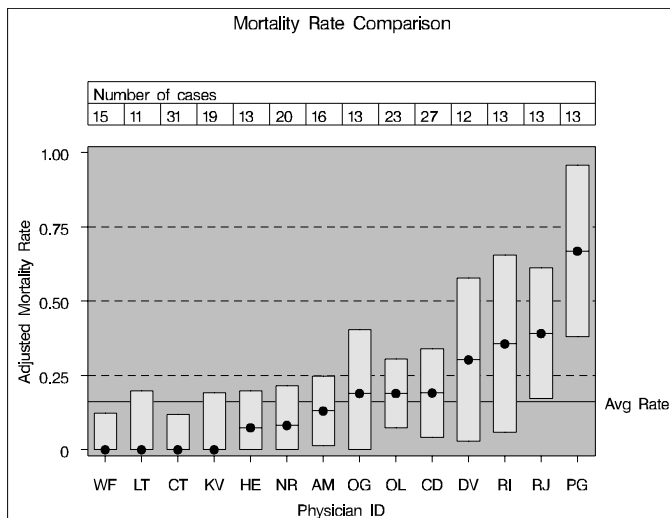
```

data adjustedmr;
    set fitout;
    if numpatients > 9;
    rename amr = amrx numPatients = amrn lower = amrl upper = amr3;
    alpha = .05;
    amrh = upper;
    amrl = lower;
    amrs = .001;
    amrm = amr;
run;

symbol v=dot;
title1 "Mortality Rate Comparison";
proc shewhart history=adjustedmr;
    label amrx = "Adjusted Mortality Rate";
    label physid = "Physician ID";
    label amrn = "Number of cases";
    boxchart amr*physid(amrn) / vref=.25 .5 .75
        cframe=ligr cboxfill=yellow cboxes=black blockrep
        nolcl noucl nolegend nolimitslegend xsymbol="Avg Rate";
run;

```

The display is shown in Figure 30. The horizontal reference line indicates the overall mortality rate for the hospital. An interval that lies completely above the reference line (as for physicians PJ and PG) indicates a rate that is significantly higher than the average. Note, however, that the intervals cannot be used to make simultaneous comparisons of physicians as in the analysis of means.



**Figure 30.** Adjusted Mortality Rates and 95% Confidence Intervals

This example is intended to illustrate the use of one possible model-based adjustment to mortality rates, and it does not reflect the various statistical approaches that have been proposed in the medical literature. Here, the method used to compute confidence intervals assumes that the covariance matrix from the state model is unknown. Hosmer and Lemeshow (1995) also show how to construct intervals when the covariance is known (you can obtain the covariance matrix by specifying the OUTEST= and COVOUT options in the LOGISTIC procedure) and describe bootstrap methods to construct the intervals.

DeLong et al. (1997) propose including the physicians as random intercept effects, fitting the model with the NLMIXED procedure or the GLIMMIX macro, then using the parameter estimates for each physician to compute odds ratios and confidence intervals. DeLong et al. (1997) also discuss several other risk-adjustment methods. Bayesian hierarchical modeling is discussed in Burgess et al. (2000) and the references therein. Bronskill et al. (2002) consider longitudinal modeling of provider performance over time. The multiple-comparison problem can be approached by using a Bonferroni-type adjustment.

## ACKNOWLEDGMENTS

We are grateful to Kim Price and Aubrey Wooldridge of Centra Health in Lynchburg, Virginia, for providing the data in Examples 4, 8, 9, and 10. We are grateful to Lynne Danha of HealthPartners in Minneapolis for providing the data in the other examples, which were originally discussed by Rodriguez (1996). The examples are intended to illustrate statistical methods and SAS programming techniques. While the examples are based on actual data, the results do not necessarily represent actual practice at the organizations that provided the data. We are also grateful to Virginia Clark, Michael Crotty, David DeNardis, Bob Derr, and Bucky Ransdell of SAS Institute, who provided valuable assistance in the preparation of this paper.

## REFERENCES

- Alwan, L. C. and Roberts, H. V. (1988), "Time Series Modeling for Statistical Process Control," *Journal of Business and Economic Statistics*, 6, 87–95.
- Berwick, D. M., Godfrey, A. B., and Roessner, J. (1990), *Curing Health Care, New Strategies for Quality Improvement*, San Francisco, CA: Jossey-Bass.
- Bronskill, S. E., Normand, S. T., Landrum, M. B., and Rosenheck, R. A. (2002), "Longitudinal Profiles of Health Care Providers," *Statistics in Medicine*, 21, 1067–1088.
- Burgess, J. F. J., Christiansen, C. L., Michalak, S. E., and Morris, C. N. (2000), "Medical Profiling: Improving Standards and Risk Adjustments Using Hierarchical Models," *Journal of Health Economics*, 19, 291–309.

- DeLong, E. R., Peterson, E. D., Muhlbaier, L. H., Hackett, S., and Mark, D. B. (1997), "Comparing Risk-Adjustment Methods for Provider Profiling," *Statistics in Medicine*, 16, 2645–2664.
- Hosmer, D. W. and Lemeshow, S. (1995), "Confidence Interval Estimates of an Index of Quality Performance Based on Logistic Regression Models," *Statistics in Medicine*, 14, 2161–2172.
- Institute of Medicine (2000), *To Err Is Human, Building a Safer Health System*, Washington, D.C.: National Academy Press.
- Institute of Medicine (2001), *Crossing the Quality Chasm: A New Health System for the 21st Century*, Washington, D.C.: National Academy Press.
- Lilford, R. J., Mohammed, M. A., Braunholtz, D., and Hofer, T. P. (2003), "The Measurement of Active Errors: Methodological Issues," *Quality and Safety in Healthcare; Patient Safety Methodology*, 12 supplement 2, ii8–ii11.
- McFadden, K. L., Towell, E. R., and Stock, G. N. (2004), "Critical Success Factors for Controlling and Managing Hospital Errors," *Quality Management Journal*, 11, 61–74.
- Nelson, P. R., Coffin, M., and Copeland, K. A. F. (2003), *Introductory Statistics for Engineering Experimentation*, New York: Elsevier Academic Press.
- Pratter, F. E. (2003), *Web Development with SAS by Example*, Cary, NC: SAS Institute Inc.
- Rodriguez, R. N. (1996), "Health Care Applications of Statistical Process Control: Examples Using the SAS System," *SAS Users Group International: Proceedings of the Twenty-First Annual Conference*.
- SAS Institute Inc. (2004a), *SAS Solutions for Health Care*, <http://www.sas.com/industry/healthcare/spm/index.html>: last accessed Feb. 6, 2004.
- SAS Institute Inc. (2004b), *SAS/ETS 9.1 User's Guide*, Cary, NC: SAS Institute Inc.
- SAS Institute Inc. (2004c), *SAS/QC 9.1 User's Guide*, Cary, NC: SAS Institute Inc.
- Wheeler, D. J. (1995), *Advanced Topics in Statistical Process Control*, Knoxville, TN: SPC Press, Inc.
- Wheeler, D. J. and Polling, S. R. (1998), *Building Continual Improvement: A Guide for Business*, Knoxville, TN: SPC Press, Inc.

## CONTACT INFORMATION

Robert N. Rodriguez  
SAS Institute Inc.  
SAS Campus Drive  
Cary, NC 27513  
(919) 531-7650  
[Bob.Rodriguez@sas.com](mailto:Bob.Rodriguez@sas.com)

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

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