

Health Care Applications of Statistical Process Control: Examples Using the SAS® System

Robert N. Rodriguez
SAS Institute Inc.
Cary, North Carolina, USA

ABSTRACT

Health care applications present a new frontier for statistical process control (SPC) methods. Interest in SPC is driven by a desire to improve patient outcomes in the face of capitation, cost reduction, competition, and changing health care industry standards. This paper illustrates the use of SAS software to analyze health care data with u charts, p charts, control charts for individual measurements, analysis of means for rates and proportions, simultaneous confidence intervals for proportions, and basic forecasting methods. Each of these methods provides a graphical display that facilitates understanding of process variability.

BACKGROUND

This paper was motivated by recent conversations with a number of SAS users in health care organizations who have asked the following questions:

- "Are statistical process control (SPC) methods relevant to health care applications?"
- "How does health care SPC differ from manufacturing SPC?"
- "How can I use SAS software to make the appropriate control charts?"

The purpose of this paper is to answer the third question with a series of tutorial examples based on health care data. Although it is beyond the scope of this paper to provide comprehensive answers to the first two questions, it is useful to begin with some background concerning the adoption of continuous quality improvement (CQI) programs within the health care industry.

Continuous quality improvement is a management strategy based on measurement and feedback of statistical information for continuous improvement. Beginning in the late 1980s, this approach has been motivated by several factors, including

- capitation and cost reduction. Employers and insurance companies are forcing health care providers to accept flat annual per-patient fees. Measures of quality are required to demonstrate that quality can be maintained and improved as health care expenses are brought under control.

- competition for patients covered under managed-care plans and competition to join preferred integrated delivery networks. Hospitals joining managed-care networks can succeed in winning contracts if they can demonstrate high patient satisfaction.
- health care industry standards. The Joint Commission on Accreditation of Healthcare Organizations (JCAHO) now requires hospitals to improve organizational performance, and patient satisfaction is one of nine measures of performance. The Health Plan Employer Data and Information Set (HEDIS), developed by the National Committee for Quality Assurance, includes patient satisfaction, and it is designed to help consumers and employers compare performance of managed-care plans.

An increasing number of health care organizations are applying the quality philosophy of W. Edwards Deming, Joseph Juran, and others who led the quality revolution that took place in American manufacturing industry in the early 1980s. Deming emphasized the need for top management to assume responsibility for quality improvement. However, the theoretical foundation for his 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; refer to Neave (1990).

One example of this approach is a quality program at Humana Inc. described by Spoeri (1991), which offers physicians and administrators information that they can use for review, patient management, and quality measurement. Clinical outcome measurements are summarized in a composite results report. Hospital rates are analyzed with control charts, and facilities with significantly high rates are asked to respond. Statistical process control methods are also applied to utilization management. For instance, control charts are used to examine length of stay, charge, and cost for combinations of hospitals, departments, and physicians. This type of analysis facilitates review through feedback of data based on performance.

Statistical quality improvement has also been effective in smaller health care organizations. For example, Staker (1995) described the use of SPC to improve clinical outcomes in primary care practice at Intermountain Health Care.

There is considerable recent evidence that statistical thinking and SPC methods can play a valuable role in health care quality improvement. For historical reasons, this experience is not well represented in standard textbooks on statistical quality improvement, such as Montgomery (1991) and Wadsworth, Stephens and Godfrey (1986), leaving the impression that health care is not a good candidate for SPC applications. However, for health care professionals who are getting started with SPC, there are now several excellent resources, including Berwick (1989) and Balestracci and Barlow (1994), that motivate and illustrate the use of SPC methods. Other useful references on statistical quality improvement in health care include Al-Assaf and Schmele (1993), Benneyan (1995), Berwick (1989, 1991, 1992), Laffel and Blumenthal (1989), Longo and Bohr (1991), Plsek (1992), and VanderVeen (1992).

A number of special issues arise in health care applications of SPC. One of these is the question of what to measure. This is less problematic in manufacturing, where the variables to be controlled are often determined by engineering requirements or by experimentation. In health care applications, there is concern and debate about excessive emphasis on outcome metrics and report cards to the detriment of process understanding and improvement; refer to Benneyan and Kaminsky (1995). Another issue is the problem of how to aggregate, adjust, and present rate data, which are increasingly used to make decisions. Categorical or "attributes" data are less prevalent in manufacturing, where advances in measurement technology have resulted in greater reliance on continuous measurements of process and quality variables.

SAS EXAMPLES

This section presents a series of tutorial examples based on health care data that demonstrate the use of the SHEWHART procedure in SAS/QC® software. Each example provides SAS code that can easily be extended to handle large data sets, which are prevalent in outcome analysis and reporting applications. Furthermore, this code can be embedded behind point-and-click interfaces developed with SAS/AF® software to facilitate use of SPC methods by individuals throughout an organization. A particularly valuable tool for this type of development is the new PFD (process flow diagram) FRAME entry, which is described in SAS Institute Inc. (1995a).

For readers who are unfamiliar with SAS programming, it is worth noting that the basic displays illustrated here can be created interactively with the SQC Menu System in SAS/QC software and the Forecasting Menu System in SAS/ETS® software. These menu systems are described in SAS Institute Inc. (1995b, 1995c).

Basic $\bar{u}$ Chart

This example introduces the use of the SHEWHART procedure to construct a $\bar{u}$ chart, which is one of several control charts for count data. In manufacturing, $\bar{u}$ 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 $\bar{u}$ 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.

A health care provider uses a $\bar{u}$ chart to analyze the rate of cat scans performed each month by each of its clinics. Output 1 shows data collected for Clinic B and saved in a SAS data set named CLINICB.

Output 1. SAS Data Set CLINICB

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	2.24481
SEP94	49	25958	30	2.16317
OCT94	63	25957	31	2.23519
NOV94	64	25920	30	2.16000
DEC94	62	25907	31	2.23888
JAN95	67	26754	31	2.30382
FEB95	58	26696	28	2.07636
MAR95	89	26565	31	2.28754

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 $\bar{u}$ chart in Figure 1.

```
title 'U Chart for Cat Scans per 1,000 Members:'
      'Clinic B';

proc shewhart data=clinicb graphics;
  uchart nscanb * month /
    subgroupn = nyrsb
    tests = 1 to 4
    testmethod = standardize
    nohlabel
    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, and the option GRAPHICS specifies that the chart is to be created on a graphics device (by default, a line character chart would be created).

The UCHART statement requests a $\bar{u}$ 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 page 1233 of SAS Institute Inc. (1995d) for a summary of the available options and page 1385 for a dictionary of options. 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 the control limits. The tests are described in Chapter 41 of SAS Institute Inc. (1995d). The TESTNMETHOD=STANDARDIZE option applies a standardization method to adjust for the fact that the number of units varies from sample to sample; refer to page 1504 of SAS Institute Inc. (1995d). This option is not used when the number of units is fixed.

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 1. Basic $\bar{u}$ Chart

In Figure 1 the control limits shown are 3σ limits estimated by default from the data*; the limits vary because the number of opportunity units changes from month to month. Formulas for the limits are given on page 1241 of SAS Institute Inc. (1995d). (Alternatively, you can read pre-established control limits from a SAS data set as illustrated in the next section.) The only test signaled by the chart is Test 1, which indicates a special cause of variation leading to a rate increase in

*In general, it is recommended that at least 25 to 30 subgroup samples be used when control is being established.

March 1995.

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.

Multiple Sets of Control Limits for a $\bar{u}$ Chart

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

A health care provider uses a $\bar{u}$ 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. Output 2 shows data collected for Clinic E and saved in a SAS data set named CLINICE.

Output 2. SAS Data Set CLINICE

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	0.70889	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

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 (a change in the system is known to have occurred in September 1994 at the beginning of Phase 2).

The following statements create a $\bar{u}$ chart with a single set of default limits.

```

title 'U Chart for Office Visits per 1,000'
      'Members: Clinic E';
proc shewhart data=clinice graphics;
  uchart nvisite * month /
    subgroupn = nyrse
    cframe    = ligr
    cinfill   = yellow
    nohlable
    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 1.




Figure 2. 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=clinic graphics;
  by _phase_;
  uchart nvisite * month /
    subgroupn = nyrsa
    outlimits = vislimit
    (rename=( _phase_ = _index_ ))
  nochart;
run;
```

Output 3 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.

Output 3. SAS Data Set VISLIMIT

Control Limits for Office Visit Data									
	I	S	L	S					
	N	U	I	A	I				
	D	B	M	L	O				
	E	G	I	P	M				
	X	R	X	A	S				
1	Phase 1	NVISITE	MONTH	ESTIMATE	V	V	3	V	2302.99
2	Phase 2	NVISITE	MONTH	ESTIMATE	V	V	3	V	2623.52

The following statements combine the data and control limits for both phases in a single u chart, shown in Figure 3.

```
title 'u Chart for Office Visits'
      ' per 1,000 Members: Clinic E';
proc shewhart data=clinic
  limits=vislimit graphics;
  uchart nvisite * month /
    subgroupn = nyrsa
    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. For details, refer to "Displaying Multiple Sets of Control Limits" on page 1458 of SAS Institute Inc. (1995d).




Figure 3. u Chart with Multiple Sets of Control Limits

In Figure 3 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 3 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.

Individual Measurements Charts

This section illustrates the use of the SHEWHART procedure to construct a control chart for individual measurements in which the measure of dispersion is based on moving ranges rather than model assumptions.

A clinic uses a chart for individual measurements to analyze the number of medical/surgical days per 1,000 members per year. Output 4 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.

Output 4. SAS Data Set MEDSURG

Medical/Surgical Admissions Data			
MONTH	MSAD_E	MSAD_OTH	_PHASE_
JAN91	.	151.9	Historical
FEB91	.	136.3	Historical
MAR91	.	236.5	Historical
...	...	...	...
SEP92	.	174.0	Historical
OCT92	.	87.2	Historical
NOV92	.	260.2	Historical
DEC92	.	183.1	Historical
JAN93	618.3	269.8	New Product E
FEB93	367.4	125.6	New Product E
MAR93	116.7	65.1	New Product E
...	...	...	...
MAY94	260.0	192.6	New Product E
JUN94	151.1	179.5	New Product E
JUL94	128.0	129.6	New Product E
AUG94	203.3	180.2	New Product E
SEP94	318.0	109.9	Younger Members
OCT94	109.6	139.6	Younger Members
NOV94	124.5	167.4	Younger Members
...	...	...	...
APR95	78.7	182.4	Younger Members
MAY95	65.0	212.6	Younger Members
JUN95	112.8	137.4	Younger Members

The following step uses the IRCHART statement in the SHEWHART procedure to construct an individual measurements 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;
proc shewhart data=medsurg graphics;
  where month < '01jan93'd;
  irchart msad_oth * month /
    cframe = ligr
    split = '/'
    nohlabel;
  label msad_oth = 'Days per 1,000/Mvg Rng';
run;

```

The chart, shown in Figure 4, indicates that the utilization for the other products is a stable, predictable process.




Figure 4. 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 5.

Output 5. SAS Data Set RUNLIM

Control Limits for Medical/Surgical Days							
INDEX	_VAR	_SUBGRP	_TYPE	_LIMIT	_ALPHA		
Historical	MSAD_OTH	MONTH	ESTIMATE	2	.0026998		
New Product E	MSAD_OTH	MONTH	ESTIMATE	2	.0026998		
New Product E	MSAD_E	MONTH	ESTIMATE	2	.0026998		
Younger Members	MSAD_OTH	MONTH	ESTIMATE	2	.0026998		
Younger Members	MSAD_E	MONTH	ESTIMATE	2	.0026998		
SIGMA	_LCLI	_MEAN	_UCLI	_LCLR	_R	_UCLR	_STDDEV
3	0.00000	156.774	324.487	0	63.081	206.057	55.904
3	7.94440	149.929	291.913	0	53.404	174.446	47.328
3	0.00000	287.773	979.441	0	260.154	849.802	230.556
3	0.00000	134.593	285.900	0	56.911	185.800	50.436
3	0.00000	138.470	383.816	0	92.281	301.439	81.782

The following statements read RUNLIM to create the control chart for the rate for Product E that is shown in Figure 5. The NOLCL option suppresses the lower control limit, which is zero, and the NOCHART2 option suppresses the chart for moving ranges.

```

title 'Medical/Surgical Rate for'
      'Product E';
symbol v=dot;
proc shewhart data=medsurg
      limits=runlim graphics;
  where month >= '01jan93'd;
  irchart msad_e * month /
    nolcl
    nohlabel
    nochart2
    phaselegend
    phaselabtype = scaled
    readindex    = all
    readphase    = all
    cframe       = ligr;
label msad_e = 'Med/Surg Days per 1,000';
run;

```

Figure 5 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 5. Medical/Surgical Admissions Rates

The next statements overlay the historical control limits for the other products as dashed lines on the preceding chart. First, the limits are saved in a reference line data set named OTHERREF; refer to page 1440 of SAS Institute Inc. (1995d).

```

data otherref;
  keep _ref_ _reflab_;
  length _reflab_ $ 16;
  set runlim;
  if _index_ = 'Historical' and
     _var_   = 'MSAD_OTH';
  _ref_     = _mean_;
  _reflab_  = 'Avg Other'; output;
  _ref_     = _ucl_;
  _reflab_  = 'UCL Other'; output;
run;

```

The next statements save the rates for the other products as coordinates for symbols in an annotate data set named

OTHER. For details concerning annotate data sets, refer to Chapter 19 of SAS Institute (1990a).

```

%annomac;
data other;
  %dclanno;
  set medsurg (rename=(month=x msad_oth=y));
  if month >= '01jan93'd;
  %sequence(after);
  %system(2,2,3);
  function = 'symbol';
  size     = 2.4;
  color    = 'yellow';
  text     = 'circle';
run;

```

The following statements use the data sets OTHERREF and OTHER to construct the overlaid chart.

```

title 'Product E and Other Products';
proc shewhart data=medsurg
      limits=runlim graphics;
  where month >= '01jan93'd;
  irchart msad_e * month /
    nochart2
    nolcl
    nohlabel
    phaselegend
    phaselabtype = scaled
    readindex    = all
    readphase    = all
    vref         = otherref
    cvref        = yellow
    lvref        = 4
    llimits      = 1
    anno         = other;
label msad_e =
  'Med/Surg Days per 1,000 for Product E';
footnote j=1
  'Circles indicate rates for other products';
run;

```

The chart is shown in Figure 6. Contrary to the original expectation, it indicates that the utilization rate for the Product E is slightly lower than the rate for the other products.




Figure 6. Product E Compared with Other Products

Individual Measurements Chart for Seasonal Data

This section illustrates the use of an individual measurements chart with multiple sets of control limits that adjust for seasonal effects.

A hospital system uses a chart for individual measurements to analyze monthly variation in the number of emergency room visits per 1,000 member-years. The data are saved in a SAS data set named ERVISIT which is listed in Output 6.

Output 6. SAS Data Set ERVISIT

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
...	...	...
JUN95	S95	68.02

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 of 1990 since a new system was introduced at that point, and the average rate was known to change.

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 graphics;
  boxchart visits*month /
    npanel          = 100
    cphasebox       = black
    cphaseboxfill   = ligr
    cphasemeanconnect = black
    phasemeanconnect = black
    phasemeanconnect = black
    cphaseleg       = black
    readphase       = all
    phaselabtype    = scale
    phaselegend     = scale
    nolegend
    nolabel
    nolimits;
  label visits = 'Visits per 1000 Member Years';
run;

```

The line segments in Figure 7 connect the average of the rates within each time phase. The display reveals higher rates of emergency room visits in warm weather (May through September) and lower rates in cold weather (October through April). The overall rate is declining until October of 1994. An explanation for this effect is that the

winter of 1994/1995 was very mild, whereas the preceding winter was very cold.




Figure 7. Emergency Room Visits

The following statements use BY processing to save distinct control limits for each phase in a SAS data set named ERLIMITS.

```

proc shewhart data=ervisit;
  by _phase_ notsorted;
  irchart visits*month /
    nochart
    outlimits =
      erlimits
      rename=(_phase_=_index_);
run;

```

The next statements read the control limits from ERLIMITS and combine them in a single chart, shown in Figure 8.

```

title 'Emergency Room Visits per'
      ' 1000 Member Years';
symbol v=dot c=orange;
proc shewhart data=ervisit
      limits=erlimits graphics;
  irchart visits*month /
    readindex      = all
    readphase      = all
    phaselabtype   = scaled
    npanel         = 100
    cframe         = ( megr ligr )
    cphaseleg      = black
    llimits        = 3
    nolimitslegend
    phaselegend
    nochart2
    nolabel;
  label visits = 'Visits per 1000 Member Years';
run;

```

Figure 8 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 on "Forecasting."




Figure 8. Emergency Room Visits

Analysis of Means for Rate Data

This section illustrates the use of analysis of means (ANOM) for rate data. Since ANOM is not a commonly used SPC tool (even in manufacturing applications) a review is appropriate.

Analysis of means is a graphical and statistical method for simultaneously comparing a group of k treatment means with their grand mean at a specified significance level α . 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. The analysis of means has the same graphical presentation as a control chart except that the decision limits are computed differently, and the visual appeal of ANOM has been key to its effectiveness in industrial applications.

There are several excellent references on analysis of means. The book by Ott (1975) provides a variety of manufacturing examples. The January 1983 issue of *Journal of Quality Technology* focuses on ANOM and is a useful source of applications and computational methods. The paper by Ramig (1983) is particularly relevant here because it discusses applications to attributes data. The SUGI 13 Proceedings paper by Nelson (1988) presents a very readable motivation and overview of ANOM, and the SUGI 13 Proceedings paper by Fulewider (1988) presents a helpful tutorial on how to use the SHEWHART procedure to perform ANOM with continuous data. The application of ANOM to health care data is discussed by Balestracci and Barlow (1994) with simplified computations that facilitate the exposition of the method. The presentation here and in the next section deals with similar examples but computes the decision limits using statistically exact results due to Nelson (1983).

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 MSADMIT, which is listed in Output 7.

Output 7. SAS Data Set MSADMIT

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
1M	121	5.0141
3E	105	4.4072
1Q	124	4.3829
1K	171	4.2691
3B	88	2.8979
1C	100	2.6633
1K	112	2.3985
2C	84	2.2898
1A	69	2.2078
1T	21	2.8913
1M	130	2.8603
1O	61	2.0438
3D	66	1.8633
1J	54	1.5918
3J	38	1.3408
3G	36	1.1543
3E	26	0.8823
1G	28	0.8626
1I	25	0.5034
1L	20	0.4282
1S	7	0.2249
1P	7	0.2020
1P	2	0.1692

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 rates of admission at the $\alpha = 0.01$ level of significance. The UCHART statement in the SHEWHART procedure is used to compute the rates and display them graphically with upper and lower decision limits (UDL and LDL). A SAS macro named ANOMSIG (see the Appendix) computes the appropriate multiple (3.52) of σ for the SIGMAS= option.

```
%anomsig( 0.01 /*alpha*/, 29 /*no. of groups*/ );

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

proc shewhart data=msadmits graphics;
  uchart count95*id /
    subgroupn = myrs95
    sigmas    = %sigmult
    cframe    = blue
    cneedles   = yellow
    climits    = black
    llimits    = 1
    cinfll    = green
    lcllabel   = 'LDL'
    ucllabel   = 'UDL'
    turnhlabels
    nolegend;
  label count95 =
    'Admits per 1000 Member Years';
run;
```

The chart is shown in Figure 9. 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.




Figure 9. ANOM for Medical/Surgical Admissions Rates

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.

Despite the similarity of Figure 9 to a $\bar{x}$ 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 3σ limits that the SHEWHART procedure would compute by default for a $\bar{x}$ 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.

Analysis of Means for Proportions

A health care system uses ANOM to compare cesarean section rates for a set of medical groups. The data are saved in a SAS data set named CSECTION, which is listed in Output 8. The variable ID identifies the medical groups, the variable CSECT95 provides the number of c-sections for each group during 1995, and the variable TOTAL95 provides the total number of deliveries for each group, which serves as the "measure of opportunity" for c-sections. The variables CSECT94 and TOTAL94 provide similar counts for 1994.

Output 8. SAS Data Set CSECTION

ID	CSECT94	TOTAL94	CSECT95	TOTAL95
1A	163	1070	150	923
1K	55	369	45	298
1B	52	231	34	170
1D	19	147	18	132
3I	21	119	20	106
3M	15	96	12	105
1E	15	67	10	77
1N	6	49	19	74
1Q	12	79	7	69
3E	7	72	11	65
1R	4	47	11	49
1K	5	37	9	48
3J	7	19	7	20
1C	13	55	8	43
3B	3	36	6	43
1M	6	14	4	29
3C	6	33	5	28
1O	8	34	4	27
1J	6	24	6	22
1F	1	4	3	22
3E	2	14	4	18
1G	1	8	4	15
3D	7	29	4	13
3G	2	7	1	11
1L	2	6	2	10
1I	1	4	1	8
1P	15	81	0	3
1F	0	1	0	3
1S	1	3	1	3

The following statements perform an ANOM for the proportion of c-sections across groups at the $\alpha = 0.01$ level of significance. The PCHART statement in the SHEWHART procedure is used to compute the proportions and display them graphically with upper and lower decision limits (UDL and LDL). As in the previous section, the ANOMSIG macro is used to determine the appropriate multiple of σ for the decision limits. Here, however, the number of groups is determined from the data and passed to the macro.

```

data csection;
  set csection end=eof;
  if eof then
    call symput('ngroups',
      left( put( _n_, 4.)));
run;
%anomsig( 0.01, &ngroups );

title 'Proportion of C-Sections in 1995';
symbol v=none w=7;

proc shewhart data=csection graphics;
  pchart csect95=id /
    subgroupn = total95
    cframe    = blue
    sigmas    = &sigmult
    cneedles   = yellow
    lc1label  = 'LDL'
    uc1label  = 'UDL'
    nolegend
    turnhlabels;
  label csect95 =
    'Proportion of Cesarean Sections';
run;

```

The chart, shown in Figure 10, indicates that the variation in rates across clinics is strictly due to chance.




Figure 10. ANOM for Rate of C-Sections in 1995

In managed care reporting, it is often necessary to compare results from one year with those of the previous year. The following statements create the display shown in Figure 11 by superimposing the c-section rates for 1994 as empty yellow circles on the chart in Figure 10.

```
%annomac;
data csect94;
  set csection;
  %dclanno;
  %sequence(after);
  %system(2,2,3);
  function = 'symbol';
  xc      = id;
  y       = csect94 / total94;
  size    = 2.4;
  color   = 'yellow';
  text    = 'circle';
run;

title 'Proportion of C-Sections:'
      ' 1994 and 1995';
symbol v=dot;
proc shewhart data=csection graphics;
  pchart csect95*id /
    subgroupn=total95
    sigmas   = %sigmult
    anno     = csect94
    lcllabel = 'LDL'
    ucllabel = 'UDL'
    turnblab;
  noconnect;
  nolegend;
label csect95 = 'Proportion of Cesarean Sections';
footnote j=1 'Empty Circles Indicate 1994 Rates';
run;
```

A drawback of this display is that the decision limits apply only to the 1995 rates. For visual clarity, a better way to compare the rates is to create an ANOM chart for each year and stack the charts.




Figure 11. C-Section Rates in 1994 and 1995

The following statements create this display with a programming trick. First, the rates and decision limits are computed and saved in an OUTTABLE= data set; for details about the structure of OUTTABLE= data sets created with the PCHART statement, refer to page 1133 of SAS Institute Inc. (1995d).

```
proc shewhart data=csection;
  pchart csect94*id /
    subgroupn = total94
    sigmas    = %sigmult
    outtable  = cstab94
    (rename = ( _lclp_ = _lclx_
                _subp_ = _subx_
                _p_    = _mean_
                _uclp_ = _uclx_ ))
  nochart;

proc shewhart data=csection;
  pchart csect95*id /
    subgroupn = total95
    sigmas    = %sigmult
    outtable  = cstab95
    (renames = ( _lclp_ = _lclr_
                _subp_ = _subr_
                _p_    = _r_
                _uclp_ = _uclr_
                _exlim_ = _exlimr_ ))
    drop     = _var_ _limitn_
              _subn_ _sigmas_
  nochart;

data cstab; merge cstab94 cstab95; run;
```

Because of the way in which the variables in CSTAB have been renamed, this data set has the structure of a TABLE= input data set for the XPCHART statement, which creates the stacked display in the following step. For details concerning TABLE= data sets, refer to page 1372 of SAS Institute Inc. (1995d).

```

title 'C-Section Rates for 1994 and 1995';
symbol v=none w=7;
proc shewhart table=cstab graphics;
  xrchart csect94 *id /
    cneedles = yellow
    cinfill = green
    cframe = blue
    split = '/'
    ypct1 = 50
    xsymbol = 'Avg'
    rsymbol = 'Avg'
    lc1label = 'LDL'
    uc1label = 'UDL'
    lc1label2 = 'LDL'
    uc1label2 = 'LDL'
    turnhlabel
    nolimitsleg
    nolegend;
label _subx =
  '1994 Proportion/1995 Proportion';
run;

```

The display, shown in Figure 12, indicates that c-section rates were comparable across groups during both years.




Figure 12. Comparative ANOM

It should be pointed out that this problem can be analyzed with statistical modeling techniques such as analysis of variance, multiple comparisons, and generalized linear models. These methods lie outside the scope of this paper, but they are well supported in the SAS System by the GLM and GENMOD procedures.

Another approach, which lends itself to graphical presentation, is the computation of simultaneous Bonferroni confidence intervals for the c-section rates across medical groups. The following statements compute and display these intervals at the $\alpha = 0.01$ level.

```

data csect95;
  set csection end=eof;
  if eof then
    call symput('ngroups',
      left( put( _n_, 4. )));
run;

```

```

proc means noprint data=csect95;
  var total95;
  output out=total sum=n;
data total;
  set total;
  call symput('n',left(put(n,6.)));
data csect95;
  set csection;
  keep id p95l p95i p95x p953 p95h
      p95n p95s p95m n95;
  p95x = csect95 / total95;
  p95n = total95;
  alpha = 0.01;
  z = probit(1-(alpha/(2*ngroups)));
  p953 = p95x + z*sqrt(p95x*(1-p95x)/p95n);
  p95l = p95x - z*sqrt(p95x*(1-p95x)/p95n);
  /* assign dummy variables */
  p95h = p953;
  p95i = p95l;
  p95s = 0.001;
  p95m = p95x;
run;

```

The data set CSECT95 has the structure appropriate for input to the BOXCHART statement, which creates the box display shown in Figure 13.

```

title 'C-Sections as a Proportion of Total'
      'Deliveries in 1995';
symbol v=none;
proc shewhart history=csect95 graphics;
  boxchart p95*id ( n95 ) /
    stddevs
    xsymbol = 'Avg'
    blocklabtype = scaled
    npanel = ngroups
    cboxes = yellow
    cboxfill = orange
    turnhlabels
    nolcl
    noucl
    nolimitslegend
    nolegend;
label p95x = 'Simultaneous 99% Intervals';
run;

```




Figure 13. Simultaneous CI's for 1995 C-Section Rates

The center of each box represents the point estimate for each rate, and the edges of the box represent the upper and lower *simultaneous* confidence limits. The overlap of the boxes, *which must be viewed as an ensemble*, conveys the message that no subset of groups stands out from the rest, except for 1P and 1F (these groups had extremely low numbers of deliveries).

Bonferroni intervals are simple to compute but conservative. They are competitive with simultaneous intervals obtained using other methods, but when the number of groups is large, Bonferroni intervals are unnecessarily long. For a comprehensive discussion of Bonferroni intervals and other methods of simultaneous inference, refer to Miller (1966).

Forecasting

In the section on "Individual Measurements Chart for Seasonal Data," 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 Chapter 9 of SAS Institute Inc. (1993).

```
proc forecast data = ervisit
    interval = month
    method = winters
    seasons = month
    lead = 7
    out = outval
    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 ervisit;
    merge ervisit 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='yellow';
    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='yellow';
    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=ervisit graphics;
    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 14, 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 Chapter 19 of SAS Institute Inc. (1993). 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 14. FORECAST Analysis

The following statements create a plot of the original and seasonally adjusted series (C_1I_1).

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

title 'Emergency Room Visits';
title2 'Original and Seasonally Adjusted Data';

symbol1 i=join c=yellow v='plus';
symbol2 i=join c=red v='diamond';

proc gplot data=out;
  plot visits * date = 1
        adjusted * date = 2 /
        overlay
        legend = legend1
        haxis = axis1
        vaxis = axis2
        cframe = blue;
axis1 label=none;
axis2 minor=none
      label=('Visits per 1000 Member Years');
legend1 cborder = black
       label = none
       value = ('original' 'adjusted');
run;
```

The plot is shown in Figure 15. Adjusting for seasonal variation, the rate of emergency room visits is decreasing over time, with a slight increase late in 1994.




Figure 15. X11 Analysis

The next statements plot the final seasonal factor.

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




Figure 16. Final Seasonal Factor

The last set of statements combine the final irregular factor and the trend in a single display, shown in Figure 17.

```

data out;
  set out;
  sum = irreg + trend;

title 'Control Chart for Irregular Variation';
proc shewhart data=out graphics;
  xchart sum * date /
    cframe = ligr
    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 14 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 14. Nonetheless, both methods provide useful views, understanding, and prediction of the variation in the process.




Figure 17. Final Irregular Factor and Trend

APPENDIX: SAS Macro for Analysis of Means

The SAS macro ANOMSIG listed here computes the appropriate multiple of the standard error for ANOM with count data. Input to the macro consists of the significance level (α) and the number of groups (k). The multiple is returned as the value of a macro variable (&sigmult).

In order to use the macro in a SAS program as illustrated in this paper, you should save the macro in a file and include it in your program. For example, if you have saved the macro in a file named anomcmp.sas, you can include it as follows.

```

filename anomcmp 'c:\sugi96\anomcmp.sas';
%include anomcmp;

```

```

/******
/* NAME: ANOMSIG
/* TITLE: Macro for Multiple of Standard
/* Error for Analysis of Means With
/* Infinite Degrees of Freedom
/* REF: P. R. Nelson (1982), "Exact
/* Critical Points for the Analysis
/* of Means", Communications in
/* Statistics, A11, 699-709
/* NOTES: This macro provides the multiple
/* of standard error for analysis of
/* means (ANOM) for infinite degrees
/* of freedom. The output &sigmult
/* is the value required for the
/* SIGMAS= option in the SHEWHART
/* procedure for an ANOM involving k
/* groups and infinite degrees of
/* freedom. The input significance
/* level alpha must be 0.10, 0.05, or
/* 0.01, and k must be in the range
/* 3<=k<=60.
/******
%macro anomsig(alpha,k);
%global sigmult;
data _null_;

/* check arguments */
if &k lt 3 then do;
  put 'ERROR: k must be greater than 2.';
  abort;
end;
if &k gt 60 then do;
  put 'ERROR: Macro not valid for k > 60.';
  abort;
end;

/* build arrays to contain the halpha */
/* approximations
array h(19) h1-h19, h1=., h2=.;
if &alpha = 0.10 then do;
  array b10(4) b1-b4;
  b10(1) = 2.326409;
  b10(2) = 0.031757;
  b10(3) = -0.000493;
  b10(4) = 0.000003114;
  h3 = 2.052;
  h4 = 2.193;
  h5 = 2.289;
  h6 = 2.363;
  h7 = 2.423;
  h8 = 2.473;
  h9 = 2.516;
  h10 = 2.554;
  h11 = 2.587;
  h12 = 2.618;
  h13 = 2.645;
  h14 = 2.671;
  h15 = 2.694;
  h16 = 2.716;
  h17 = 2.735;
  h18 = 2.755;
  h19 = 2.773;
end;
else if &alpha = 0.05 then do;
  array b05(4) b1-b4;
  b05(1) = 2.597909;
  b05(2) = 0.028231;
  b05(3) = -0.000419;
  b05(4) = 0.000002551;

```

```

h3      = 2.344;
h4      = 2.468;
h5      = 2.555;
h6      = 2.621;
h7      = 2.676;
h8      = 2.721;
h9      = 2.761;
h10     = 2.796;
h11     = 2.827;
h12     = 2.855;
h13     = 2.881;
h14     = 2.904;
h15     = 2.926;
h16     = 2.946;
h17     = 2.965;
h18     = 2.983;
h19     = 2.999;
end;
else if salpha = 0.01 then do;
  array b01(4) b1-b4;
  b01{1} = 3.099409;
  b01{2} = 0.025932;
  b01{3} = -0.000397;
  b01{4} = 0.000002489;
  h3      = 2.914;
  h4      = 3.013;
  h5      = 3.084;
  h6      = 3.140;
  h7      = 3.185;
  h8      = 3.224;
  h9      = 3.258;
  h10     = 3.288;
  h11     = 3.315;
  h12     = 3.340;
  h13     = 3.362;
  h14     = 3.382;
  h15     = 3.401;
  h16     = 3.419;
  h17     = 3.435;
  h18     = 3.451;
  h19     = 3.466;
end;
else do;
  put 'ERROR: Alpha not 0.10, 0.05, or 0.01';
  abort;
end;

if ak < 20 then sigmult = h(ak);
else
  sigmult = b1      +
            b2*ak    +
            b3*ak*ak  +
            b4*ak*ak*ak;
sigmult = sigmult * sqrt((ak - 1)/ak);
call symput('sigmult',left(put(sigmult,8.3)));
%mend anomsig;

```

ACKNOWLEDGEMENTS

I am very grateful to Lynne Dancha of HealthPartners in Minneapolis for providing much of the data used in the examples, and to Brett Chapman, Martin King, and Donna Sawyer of SAS Institute for valuable assistance in the preparation of this paper.

REFERENCES

- Al-Assaf, A. F. and Schmele, J. A. (1993), *The Textbook of Total Quality in Healthcare*, Delray Beach, Florida: St. Lucie Press.
- Alwan, L. C. and Roberts, H. V. (1988), "Time Series Modeling for Statistical Process Control," *Journal of Business and Economic Statistics*, 6, 87-95.
- Balestracci, D. and Barlow, J. L. (1994), *Quality Improvement: Practical Applications for Medical Group Practice*, Englewood, Colorado: Center for Research in Ambulatory Health Care Administration.
- Benneyan, J. (1995), "Applications of Statistical Process Control (SPC) to Improve Health Care," conference presentation, *Healthcare Information and Management Systems Society*.
- Benneyan, J. and Kaminsky, F. C. (1995), "Another View on How to Measure Health Care Quality," *Quality Progress*, 28(2), 120-124.
- Berwick, D. M. (1989), "Continuous Improvement as an Ideal in Health Care," *New England Journal of Medicine*, 320, 53-56.
- Berwick, D. M. (1991), "Controlling Variation in Healthcare: A Consultation from Walter Shewhart," *Medical Care*, 29, 1212-1225.
- Berwick, D. M. (1992), "Can Quality Management Really Work in Health Care?," *Quality Progress*.
- Berwick, D. M., Godfrey, A. B., and Roessner, J. (1991), *Curing Health Care*, San Francisco: Jossey-Bass.
- Fulenwider, D. O. (1988), "Using SAS Software for the Analysis of Means," *SAS Users Group International: Proceedings of the Thirteenth Annual Conference*, 1212-1219.
- Laffel, G. and Blumenthal, D. (1989), "The Case for Using Industrial Quality Management Science in Health Care Organizations," *Journal of the American Medical Association*, 262, 2869-2873.
- Longo, D. R. and Bohr, D. (1991), *Quantitative Methods in Quality Management: A Guide for Practitioners*, Chicago: American Hospital Association.
- Miller, R. G. (1980), *Simultaneous Statistical Inference*, New York: Springer-Verlag.
- Montgomery, D. C. (1991), *Introduction to Statistical Quality Control, Second Edition*, New York: John Wiley & Sons.
- Neave, H. R. (1990), *The Deming Dimension*, Knoxville, TN: SPC Press, Inc.
- Nelson, L. S. (1974), "Factors for the Analysis of Means," *Journal of Quality Technology*, 6, 175-181.

Nelson, L. S. (1983), "Exact Critical Values for Use with the Analysis of Means," *Journal of Quality Technology*, 15, 40-44.

Nelson, P. R. (1982), "Exact Critical Values for Use with the Analysis of Means," *Communications in Statistics*, A11, 699-709.

Nelson, P. R. (1988), "Application of the Analysis of Means," *SAS Users Group International: Proceedings of the Thirteenth Annual Conference*, 225-230.

Ott, E. R. (1975), *Process Quality Control: Troubleshooting and Interpretation of Data*, New York: McGraw-Hill.

Plsek, P. E. (1992), "Introduction to Control Charts," *Quality Management in Health Care*, 1, 65-74.

Ramig, P. R. (1983), "Applications of the Analysis of Means," *Journal of Quality Technology*, 15, 19-25.

SAS Institute Inc. (1990a), *SAS/GRAPH Software: Reference, Version 6, First Edition, Volumes 1 and 2* Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1990b), *SAS/STAT User's Guide, Version 6, Fourth Edition, Volumes 1 and 2* Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1993), *SAS/ETS User's Guide, Version 6, Second Edition*, Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1995a), *SAS/AF Software: FRAME Class Dictionary, Version 6, First Edition*, Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1995b), *SAS/ETS Software: Time Series Forecasting System, Version 6, First Edition*, Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1995c), *SAS/QC Software: SQC Menu System for Quality Improvement, Version 6, Second Edition*, Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1995d), *SAS/QC Software: Usage and Reference, Version 6, Volumes 1 and 2*, Cary, NC: SAS Institute Inc.

Spoeri, R. K. (1991), "The Emerging Use of Measurement and Statistics in Health Care Quality and Productivity Improvement," presentation at the 1991 Joint Statistical Meetings, Atlanta, Georgia.

Staker, L. V. (1995), "Using Statistical Process Control to Improve Clinical Outcomes in a Primary Care Practice," presentation at the 7th Annual National Forum on Quality Improvement in Health Care.

VanderVeen, L. M. (1992), "Statistical Process Control: A Practical Application for Hospitals," *Journal for Healthcare Quality*, 14, 20-29.

Wadsworth, H. M., Stephens, K. S., and Godfrey, A. B. (1986), *Modern Method for Quality Control and Improvement*, New York: John Wiley & Sons.

Wheeler, D. J. (1993), *Understanding Variation: The Key to*

Managing Chaos, Knoxville, TN: SPC Press, Inc.

Wheeler, D. J. (1995), *Advanced Topics in Statistical Process Control: The Power of Shewhart's Chart*, Knoxville, TN: SPC Press, Inc.

SAS, SAS/AF, SAS/ETS, and SAS/QC are registered trademarks of SAS Institute Inc. in the USA and other countries.
* indicates USA registration.

Printed in the U. S. A.