Research Program on Forecasting

Monitoring Processes with Changing Variances

J. Keith Ord
McDonough School of Business,
Georgetown University, Washington, DC20057.
Email: ordk@georgetown.edu

RPF Working Paper No. 2008-004
http://www.gwu.edu/~forcpgm/2006-004.pdf

July 22, 2008

RESEARCH PROGRAM ON FORECASTING
Center of Economic Research
Department of Economics
The George Washington University
Washington, DC 20052
http://www.gwu.edu/~forcpgm

Monitoring Processes with Changing Variances

Abstract

Statistical process control (SPC) has evolved beyond its classical applications in manufacturing to monitoring economic and social phenomena. This extension requires consideration of autocorrelated and possibly non-stationary time series. Less attention has been paid to the possibility that the variance of the process may also change over time. In this paper we use the innovations state space modeling framework to develop conditionally heteroscedastic models. We provide examples to show that the incorrect use of homoscedastic models may lead to erroneous decisions about the nature of the process. The framework is extended to include counts data, when we also introduce a new type of chart, the P-value chart, to accommodate the changes in distributional form from one period to the next.

Keywords: control charts, count data, GARCH, heteroscedasticity, innovations, state space, statistical process control

1. Introduction

The seminal work of Shewhart (1931) and many others on control charts focused squarely upon manufacturing processes, where production conditions might be expected to be stable over time if and when the process was in control. Typically a small sample (e.g. n= 5) was taken and the sample mean plotted on an $\overline{X}$ chart. The process was deemed to be out of control if the sample mean was more than three standard deviations from the underlying mean. The nature of such processes meant that it was often reasonable to leave a sufficient interval of time between successive sets of readings to ensure that the samples might be viewed as independent from one time period to the next. Further, if the process was deemed to be out of control, it was reasonable to assume that equipment could be reset so that the process was in control again by the time the next sample was taken. Thus, we are led to the standard formulation that the observations are independent and identically distributed at a given time and that there is independence between time periods.

Inevitably, as the popularity of such methods increased, these assumptions often became less plausible. In particular, independence over time seemed an overly strong assumption; for example, the process might deteriorate slowly so that a series of related but relatively weak signals could occur, but none strong enough to trigger an intervention. The recognition of such possibilities led to a series of ad-hoc rules designed to capture such behavior. For example, Version 15 of Minitab gives eight rules, each designed to generate approximately the same probability of a Type I error for an in-control process. These rules include:

- 1 point more than 3 standard deviations from center line (the original rule)
- 2 out of 3 points > 2 standard deviations from center line (same side; the warning area)
- 6 points in a row, all increasing or all decreasing.

Over time, it became evident that an out of control process could also lead to increased variability, without necessarily affecting the mean level, so that charts based upon the range of the sample (R-chart) or its standard deviation (S-chart) became popular. A variety of extensions and improvements have appeared, see for example Montgomery (2004); we focus upon the basic approaches in this paper.

The next major change that came about was the explicit introduction by Alwan and Roberts (1988) of time series models to describe the underlying nature of the process. Once we recognize that the process evolves over time, we can see that the assumptions made earlier cease to be valid and that erroneous decisions could be made if the time dependence is ignored. Another way of looking at this approach is to think of the usual charts as being *unconditional*, in that they rely upon the marginal distribution provided the series is stationary so that such a distribution exists. By contrast, the time series charts are defined *conditionally* upon the past values of the series. We summarize developments in this area in section 1.1.

As the use of control charts has spread beyond manufacturing processes to the study of social and economic phenomena, the emphasis has also shifted more to monitoring, rather than control. Indeed, intervention to restore “control” may be physically or politically impossible, yet we still wish to know when a process has deviated from its expected path, be it an increase in crime rates or a shift in gasoline prices. We observe that there are two distinct issues here: a process may be in or out of statistical control, and it may meet or fail to meet the goals for which the monitoring is in place. For example, crime rates may be changing in a manner that is statistically predictable, but politically unacceptable. In such circumstances, changes in volatility may occur as well as changes in the mean level, which suggests that we should consider time series models that reflect movements in the variance as well as the mean. Such processes are briefly reviewed in section 1.2.

When we consider monitoring social and economic processes, it is important to recognize that although the tools may be similar, the focus is somewhat different. In monitoring applications, we may be well aware of temporal dependence and of changes in the mean and variance over time. Our purpose is to look for unexpected shifts, and policy changes in response may be slow to take effect. For a discussion in the context of transportation indicators, see Ord and Young (2004).

A further point needs to be made when we consider monitoring social and economic processes. In order to calibrate the charts, we must either assume that the process is in statistical control during the calibration period, or that outliers can be successfully identified and adjusted. This outlier modification step must be approached with care; typically the parameter estimates may not change dramatically, but the residual variance may reduce considerably thereby narrowing the control limits. If we are overly zealous in outlier removal we may induce a "chicken little" affect whereby excessive numbers of out-of-control signals are generated in later periods.

The paper is structured as follows. In the remainder of this section we describe control charts based upon state space models, and the extension to models with changing variances. In section 2 we provide two examples of processes that are time dependent in both the mean and the variance, and illustrate how the proposed approach enables monitoring of each process in an effective manner. Then, in section 3 we consider the same question for data on counts and provide a framework for monitoring these processes, again illustrated by an example. Section 4 presents the conclusions.

1.1 Control charts based upon exponential smoothing

We may formulate the statistical model underlying simple exponential smoothing (SES) using the innovations state space approach (Snyder, 1985; Ord, Koehler & Snyder, 1997; Hyndman, Koehler, Ord and Snyder, 2008). Such models may be formulated as follows; for purposes of illustration we focus upon the *local level model*. We denote the process of interest by $\{y_t,\ t=1,2,\ldots\}$, an unobserved state variable by $\{x_t,\ t=0,1,\ldots\}$ and a random error term by $\{\varepsilon_t,\ t=1,2,\ldots\}$. For the present, we assume that the errors are independent and identically distributed with zero means and common variance σ^2, or $\varepsilon_t \sim IID(0,\sigma^2)$. We then define the *measurement* (or observation) equation and the *transition* (or state) equation respectively as:

$$\begin{aligned} y_t &= x_{t-1} + \varepsilon_t \\ x_t &= x_{t-1} + \alpha\varepsilon_t \end{aligned} \tag{1}$$

The measurement equation describes the variations about the underlying mean level (the unobservable state) whereas the transition equation updates the state in light of the latest error term; the parameter lies in the range $0 \le \alpha < 2$. Eliminating the state variable between the two equations reduces to the familiar ARIMA(0,1,1) form:

$$y_t - y_{t-1} = \varepsilon_t - \theta\varepsilon_{t-1}, \text{ where } \theta = 1-\alpha . \quad (2)$$

Although model (1) leads to (2) in a formal sense and implies that the forecasts are generated by simple exponential smoothing (SES), there is an important distinction. In the state space model we often assume that the series starts at time $t = 1$, known as the *finite start-up condition*. By contrast, ARIMA modeling typically assumes that the series extends back into the infinite past. This distinction is not a mere formality: ARIMA schemes require that (after a suitable degree of differencing) the process is stationary, whereas the state space approach does not require such an assumption. In particular, the state space model can accommodate $\alpha = 0$ which leads to the model $y_t = x_0 + \varepsilon_t$; the ARIMA model cannot reduce to this form and so requires $\alpha > 0$. Further, as we see later in the paper, one of the advantages of the state space formulation is that it provides a more straightforward way to relax the equal variances assumption.

The state space models can be extended to include seasonal components, trends and higher order lags, just like the ARIMA system. Indeed, every ARIMA model may be represented in state space form and every linear state space model can be reduced to an ARIMA scheme; for further discussion see, for example Hyndman et al. (2008, Chapter 11). In this paper we restrict attention to the simplest models since these versions often suffice for short-term monitoring; the extensions to more complex schemes are conceptually straightforward.

Alwan and Roberts (1988) showed that the SES model is often an appropriate way to check whether the process is in statistical control; in essence, we use a Shewhart chart to check the behavior of the residuals. A separate question is whether the process is behaving as desired, or on target. This question can be examined by looking at a plot of the (clearly correlated) values of the state variable over time. If the trends are not to the liking of the decision maker, an intervention is required. In this way, Alwan and Roberts clearly separate the issues of statistical control and targeted behavior, which are sometimes confused. In a later paper, Alwan and Roberts (1995) show that the failure to allow for time dependence leads to the use of misplaced control limits leading to potential errors in decision-making.

The basic model defined by equation (1) may be too limited and extensions to a broader range of state space models or their ARIMA counterparts are clearly possible. In particular, a slightly more general model is given by:

$$\begin{aligned} y_t &= \mu + \phi x_{t-1} + \varepsilon_t \\ x_t &= \phi x_{t-1} + \alpha \varepsilon_t \end{aligned} \tag{3}$$

This model corresponds to the ARIMA(1,0,1) model, with autoregressive and moving average parameters ϕ and $\phi(1-\alpha)$ and constant $\mu(1-\phi)$, reducing to equation (2) when $\phi = 1$. One reason for not going too far in the direction of increased complexity is the increase in the number of parameters to be estimated, although further extensions are sometimes required (e.g. to allow for seasonality).

For a general overview of recent developments in monitoring changes in the mean and variance, see Stoumbos, Reynolds and Woodall (2003). In this paper we consider only extensions to the Shewhart chart and do not consider cumulative sum (CUSUM) charts. Reynolds and Stoumbos (2005) examine conditions under which it may be desirable to use exponential smoothing and CUSUM charts in combination. Extensions to the multivariate case have been examined by several authors; see for example Lowry *et al*. (1992) and Pan & Jarrett (2004), but we stay within the univariate framework in this paper.

1.2 GARCH models

The focus of this paper is how to monitor processes whose variance changes over time. To do so, we must extend the models described in the previous section to accommodate such structural movements. The first formulation of this type was the ARCH Autoregressive Conditional Heteroscedastic (ARCH) model proposed by Engle (1982). Intuitively speaking, the ARCH models represent the conditional variance in a purely autoregressive way, which may be extravagant in terms of the number of parameters to be estimated. For this reason, the Generalized ARCH (or GARCH) model proposed by Bollerslev (1986) is now generally preferred. The GARCH version may be thought of as an ARMA formulation, although the details are more involved. As with models for the mean, the question of stationarity is important for ARMA models. Since the state space model may assume a finite start-up, stationarity is not necessary.

The original ARCH and GARCH models formulated changes in the variance directly in terms of the variance itself, so that conditions are required on the parameters to ensure that the estimated variance remains positive. Nelson (1991) introduced the exponential or EGARCH model which considers the logarithm of the variance, thereby avoiding the need for such conditions. We will consider both possibilities. Tsay (2005, Chapter 3) provides an excellent guide to recent extensions of these models. It will be evident from the ensuing discussion that more complex models are readily incorporated into the proposed framework.

Our discussion leads to a modification of model (3) to allow the error terms to be independent but not identically distributed, with zero means and variance at time t dependent on previous observations, denoted by V_{t-1}. The variance might then be updated according to a relationship such as either of:

$$V_t = u_{10} + u_{11}V_{t-1} + u_1(\varepsilon_t)$$
$$\ln V_t = u_{20} + u_{21}\ln V_{t-1} + u_2(\varepsilon_t) \quad (4)$$

The functions $u_i(.)$ are open to choice and the selection may well be application-specific. However, reasonable choices are $u_1(\varepsilon) = u_{12}\varepsilon^2$ and $u_2(\varepsilon) = u_{22}\ln(\varepsilon^2)$. To avoid complications when the error is (close to) zero, we could add a small positive constant in the second case, u_3 say, and use the function $\mathrm{u}_2(\varepsilon_t)=\mathrm{u}_2\ln(\varepsilon_t^2 + u_3)$. The constant variance case corresponds to $u_2 = 0$.

In addition to the variance specification, there are many variations on the basic form. For example, we may include slope and seasonal state equations in the usual way (c.f. Hyndman et al., 2008, Chapter 2) or allow the variance to depend upon the state variables used to describe the mean.

The benefit of the innovations state space approach is that we have considerable freedom in the specification of expression (4), yet parameter estimation is still straightforward and may be performed by maximum likelihood, by minimizing the sum of squared or absolute errors, or by using any other appropriate objective function.

2. Monitoring heteroscedastic processes

Following from the discussion in the previous section, we extend model (3) to include a variance function:

$$\begin{aligned} y_t &= \mu + \phi x_{t-1} + \varepsilon_t \\ x_t &= \phi x_{t-1} + \alpha\varepsilon_t \\ \ln V_t &= u_0 + u_1\ln V_{t-1} + u_2\ln(\varepsilon_t^2) \\ \varepsilon_t &\sim IN(0, V_{t-1}) \end{aligned} \quad (5)$$

The parameters may be estimated by maximum likelihood using the likelihood function:

$$\ell(x_0, v_0, \mu, \phi, \alpha, u_0, u_1, u_2) = \prod_{t=1}^{n} \frac{1}{\sqrt{2\pi v_{t-1}}} \exp\left(-\tfrac{1}{2}\frac{\varepsilon_t^{\,2}}{v_{t-1}}\right) \tag{6}$$

where x_0 and v_0 denote the start-up values for the state variables. Recall that this approach differs from the standard ARIMA version, where it is assumed that the series has an infinite past. Thus, the likelihood includes the initial values of the state variables. For a detailed discussion of the estimation issues, see Hyndman et al. (2008, Chapter 5). We now consider two examples to illustrate the importance of allowing for possibly heteroscedastic processes.

2.1 Running mileage

The first example uses a time series of the number of miles run per month by a middle-aged forecaster over the period January 1981 – March 2007. The series is plotted in Figure 1, panel (a). Model (3) was fitted over the period 1/81 to 9/96 (n = 189) and yielded the estimates $\hat{\phi} = 0.61$, $\hat{\alpha} = 1.15$ and standard error 37.8. The standardized residuals are plotted in Figure 1, panel (b). The Box-Ljung test gave a p-value of 0.643 for the first 12 lags, so that there is no indication of a seasonal pattern. This model was then used to generate one-step-ahead forecasts for the period 6/97 to 3/07, and the one-step-ahead standardized forecast errors are plotted in Figure 1, panel (c). The eight-month hiatus covered when the runner was injured and in most of those months he recorded zero mileage; the error was set to zero at 5/97 to restart the series for the one-step-ahead predictions.

It is clear from Figure 1 panel (c) that the runner's training regime changed post-injury; the average mileage is lower and the variability in the series is greatly reduced.

Because the variance has declined so much, the monitoring process is completely ineffective, with virtually all the residuals being within one standard deviation of the center line.

The results for model (5) show interesting differences. The likelihood estimates derived from (6) for the first part of the series are:

$\hat{\phi} = 0.60, \hat{\alpha} = 1.20,\ \hat{u}_0 \simeq 0, \hat{u}_1 = 0.991,\ \hat{u}_2 = 0.0095$ and $\hat{\mu} = 106$. On inspection, we observe that the estimates relating to the mean level of the process are very similar to those for the constant variance model. By contrast, Figure 2, panel (a) shows the standardized one-step-ahead forecasts based upon model (5), which presents a more reasonable picture than Figure 1 panel (c). Even so, it appears that the standard deviations may still be on the high side. In practice, the model would be periodically recalibrated when used for monitoring, so that adjustments would be incorporated more rapidly. When we refit the model over the period 6/97 to 3/07 we obtain the modified plot shown in Figure 2, panel (b), which is much more reasonable, although it clearly benefits from the wisdom of hindsight in the estimation process. Finally, Figure 2, panel (c) shows how the standard deviation of the process has declined over time, roughly by a factor of three.

2.2 Gasoline prices

The first example shows reduced volatility, which in general might be the result of process improvements (e.g. new laws) or structural changes. It is important that the monitoring system should adjust to such changes so as to avoid missing shifts in the new regime. More common perhaps, are processes whose volatility increases over time.

Gasoline prices provide a recent painful example of such a process. We consider a simple regression model for which the dependent variable is:

Y = logarithm of US retail gas prices[1] (the average price per gallon, in dollars)

and the predictor variable (lagged one month) is:

Z = logarithm of the spot price of a barrel of West Texas Intermediate (WTI) oil as traded at Cushing, Oklahoma (in dollars).

The Cushing spot price is widely used in the industry as a "marker" for pricing a number of other crude oil supplies traded in the domestic spot market at Cushing, Oklahoma. The data are monthly and cover the period January 1991 to November 2006.

We use model (5) as a local level model with $\phi = 1$ to reflect the random walk like nature of the series and extend the measurement equation to the form:

$$y_t = \beta_0 + x_{t-1} + \beta z_{t-1} + \varepsilon_t \tag{7}$$

Different versions of the model were fitted to data for the period January 1991 to December 2001; the summary results are given in Table 1. The improved fit provided by the second GARCH model is evident. The GARCH term for this model is a slight variant on the earlier model:

$$\ln V_t = 0.009 + 0.987 \ln V_{t-1} + 0.019 \ln(|\varepsilon_t| + 0.006) \tag{8}$$

Figure 4, panels (a) and (b) show the control charts based upon the constant variance and GARCH models respectively. As is to be expected, the general shapes of the two plots are similar. What is very different is the scale of the Y-axis. The constant variance model ranges from roughly +3 to -4 standard deviations with a number of warning

[1] These series are available from the US Energy Information Administration website http://www.eia.doe.gov.

signals. By contrast, the GARCH model shows no such signals. The reason for the difference is readily deduced from Figure 3 panel (c), which shows the steady increase in the standard deviation over time. The estimated standard deviation from the earlier part of the series is only 0.030, whereas the most recent values are roughly twice that. We may conclude that the series has indeed become more volatile, but that the basic model continues to describe its general movements. By contrast, as we saw in Figure 3, the average price is increasing rapidly. An effective monitoring scheme must take both the plot of levels and the plot of residuals into account, as noted by Alwan and Roberts (1988). In combination the two plots indicate that although price is clearly increasing, the process is in (GARCH-adjusted) statistical control. The impacts of potential policies should be evaluated in terms of their ability to affect both the price level and the magnitude of the fluctuations.

3. Monitoring counts

When monitoring counts, a natural starting point is the Poisson distribution; see recent discussions in Heinen (2003) and Jung, Kukuk and Liesenfeld (2006). However, it is often found that the data are over-dispersed; that is, the variance is greater than the mean. To overcome this difficulty, Harvey and Fernandes (1989) proposed the use of the negative binomial distribution (NBD). Harvey and Fernandes (1989) provided an updating procedure for parameter estimation justified by Bayesian arguments. However, Grunwald, Hazma and Hyndman (1997) showed that a weakness of this model (and others) is that the sample paths are degenerate (towards zero); nevertheless this property should not be of major concern in our application provided we focus attention upon short-term monitoring.

We opted to preserve the exponential smoothing format for the mean and then used a linear updating relationship for the variance, to maintain the similarities with the Trigg tracking function given in Cohen, Garman and Gorr (2007). That is, we updated the mean and variance using the recurrence relationships:

$$\begin{aligned} m_t &= (1-\alpha)m_{t-1} + \alpha y_{t-1} \\ v_t &= (1-\delta)v_{t-1} + \delta\varepsilon_{t-1}^2 \end{aligned} \tag{9}$$

where m_t and v_t denote the mean and variance, $\varepsilon_t = y_t - m_{t-1}$ and (α, δ) are the smoothing parameters. The corresponding NBD may be written as

$$P(Y_t = y) = \frac{\Gamma(k_t + y)}{\Gamma(k_t)y!} p_t^{k_t} q_t^y, \quad y = 0,1,\ldots \tag{10}$$

This distribution has mean and variance: $m_t = k_t q_t / p_t$ and $v_t = k_t \dfrac{q_t}{p_t^2}$.

A practical difficulty with this approach is that we may have different distributions at each point in time, which makes conventional Shewhart charts ungainly and difficult to interpret. Instead, we may plot a chart of the P-values, which can readily accommodate changing parameters or even different distributions. Since interest usually attaches to larger than expected counts, we describe only the construction of one-sided charts for the upper tail, although the lower-tail and two-tail variants are readily constructed along the same lines.

Suppose that, at time *t*, we observe the value y_t on the random variable Y_t. If the probability mass function of Y_t is $P(Y_t = y_t) = Q_t(y_t)$, we may compute the P-value as:

$$p_t = P(Y_t \geq y_t) = Q_t(y_t) + Q_t(y_t + 1) + \ldots \tag{11}$$

We then create a chart by plotting the P-values. A critical value for the chart may be established in the usual way via the Average Run Length (ARL). That is, if the desired ARL is A, the critical value is $P = A^{-1}$ when the successive tests are conditionally independent.

The P-value chart may be easier to interpret than the usual chart for the Poisson means (known as the c-chart) when the parameters change over time, although the discrete nature of the probability distribution precludes having equal probabilities for the tail areas. Nevertheless, the observed P-values seem a better framework for decision-making than the standard charts and this chart has the advantage that quite distinct distributions might be used at different time periods if deemed appropriate. For example, if the variance dropped below the mean, we may use the binomial distribution in place of the NBD.

We now illustrate the method using data on the number of murders per year in Montgomery County, Maryland, over the period 1985 – 2006. The data are given in Table 2. Cursory inspection of the data suggests there is little or no trend, so we initially fitted a Poisson model with parameter:

$$m_t = (1-\alpha)m_{t-1} + \alpha y_{t-1}. \tag{12}$$

Since we have a short series and we are interested in a pure monitoring scheme, we set $\alpha = 0.10$ and estimated the initial level of the mean by averaging the first six observations, which yielded $m_0 = 17.0$. Figure 5 shows the P-value plot for the Poisson and Table 3 gives the extreme values.

Overall, the sample variance is 45.6 whereas the mean is 20, the variance-to-mean ratio of 2.28, suggestive of over-dispersion and a need for the NBD. In view of the short series, we preset $\delta = 0.05$ and $\alpha = 0.10$ and used the mean and variance of the first six observations to initialize equations (9), yielding the values $m_0 = 17.0$ and $v_0 = 38$. As for the Poisson, Figure 5 shows the P-value plot and Table 3 gives the extreme values. It is readily seen that the NBD is much more selective when identifying extreme events.

As noted by Alwan and Roberts (1995), many SPC applications are marred by incorrect distributional assumptions. In this case, either the use of fixed parameters or the choice of the Poisson could lead to incorrect conclusions.

4. Conclusions

When we use control charts to monitor social or economic processes, temporal dependence is often a given. Further, both the mean and the variance may evolve over time in a recognizable fashion. The objective is then to model such anticipated changes so that unexpected shifts can be identified. We have used innovations state space models to describe such evolving processes and shown that constant variance models may be quite inadequate for the monitoring task; by contrast, models that allow for conditional heteroscedasticity are much more effective. Such models enable us to separate out issues of statistical control from those of underlying trends, thereby providing a decision maker with a clearer view of the underlying process.

Counts data provide a particular challenge in this context since changes in the parameters produce different distributions above and beyond shifts in the mean and variance. To accommodate such changes, we recommend using P-value charts in place of the charts used when the distributions are identical. An excessive number of small P-values may indicate model misspecification.

It is also worth noting that the P-value charts may also be used for continuous distributions. For the common case of the normal distribution, this chart and that based upon z-scores provide equivalent information, since the distributions can vary only through the mean and variance. However when the distribution is non-normal and may even be changing form over time, the P-value chart offers a flexible way of making comparisons over time. Further since P is uniformly distributed when the process is in control, a CUSUM chart could be constructed treating $-2\sum_{j=1}^{K}\ln(1-P_j)$ as chi-square with $2K$ degrees of freedom.

The paper focuses upon extensions to Shewhart charts for monitoring univariate time series. Extensions to multivariate series are issues for future research. A systematic analysis of the performance of this approach also needs to be considered, using data coded for exceptions, as in Cohen, Garman and Gorr (2008).

References

Alwan, L.C., & Roberts, H.V. (1988). Time-series modeling for statistical process control. *Journal of Business and Economic Statistics, 6,* 87 – 95.

Alwan, L.C., & Roberts, H.V. (1995). The problem of misplaced control limits. *Applied Statistics, 44,* 269 – 278.

Bollerslev, T. (1986). A conditional heteroskedastic time series model for speculative prices and rates of return. *Journal of Econometrics, 69,* 542 – 547.

Cohen, J., Garman, S. & Gorr, W. (2007). *Calibration of time series monitoring methods using expert judgment.* Working Paper, Heinz School of Public Policy and Management, Carnegie Mellon University.

Engle, R.F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica, 50,* 987 – 1007.

Grunwald, G.K.K., Hamza, K. & Hyndman, R.J. (1997). Some properties and generalizations of non-negative Bayesian time series models. *Journal of the Royal Statistical Society, series B, 59,* 615 – 626.

Harvey, A.C. & Fernandes, C. (1989). Time series models for count or qualitative observations. *Journal of Business and Economic Statistics, 7,* 407 – 422.

Heinen, A. Modeling time series count data: An autoregressive conditional Poisson model. CORE Discussion Paper 2003/62, Center of Operations research and Econometrics, Université Catholique de Louvain.

Hyndman, R.J., Koehler, A.B., Ord, J.K. & Snyder, R.D. (2008). *Forecasting with Exponential Smoothing: The State Space Approach*. New York: Springer.

Jung, R.C., Kukuk, M. & Liesenfeld, R. (2006) Time series of count data: modeling , estimation and diagnostics. *Computational Statistics & Data Analysis, 51,* 2350 – 2364.

Lowry, C., Woodall, W.H., Champ, C. & Rigdon, S. (1992). A multivariate exponentially weighted moving average control chart. *Technometrics, 34,* 34 – 53.

Montgomery, D.C. (2004). *Introduction to statistical quality control.* New York, Wiley.

Nelson, D.B. (1991). Conditional heteroskedasticity in asset returns: A New Approach. *Econometrica, 59,* 347 – 370.

Ord, J. K., Koehler, A. B. & Snyder, R. D. (1997). Estimation and prediction for a class of dynamic nonlinear statistical models. *Journal of the American Statistical Association, 92,* 1621 – 1629.

Ord, J.K. & Young, P. (2004). Estimating the impact of recent interventions on transportation indicators. *Journal of Transportation and Statistics, 7,* 69 – 85.

Pan, X. & Jarrett, J. (2004). Applying state space to SPC: Monitoring multivariate time series. *Journal of Applied Statistics, 31,* 379 – 418.

Reynolds, M.R. Jr., & Stoumbos, Z.G. (2005). Should exponentially weighted moving average and cumulative sum charts be used with Shewhart limits? *Technometrics,* 47, 409 – 424.

Stoumbos, Z.G., Reynolds, M.R. Jr., and Woodall, W.H. (2003). Control chart schemes for monitoring the mean and variance of processes subject to sustained shifts and drifts. *Handbook of statistics: statistics in industry,* C.R. Rao & R. Khattree (eds.). Amsterdam: Elsevier, pp. 53 – 71.

Shewhart, W.A. (1931). *Economic control of quality of manufactured product.* New York: Van Nostrand.

Snyder, R.D. (1985). Recursive estimation of dynamic linear models. *Journal of the Royal Statistical Society, series B, 47,* 272 – 276.

Tsay, R.S. (2005). *Analysis of financial time series*. NewYork: John Wiley & Sons. Second Edition.

Table 1: Parameter estimates for models fitted to the U.S. Gasoline Price data, 1/81 – 12/01. The AIC values are measured relative to AIC = 0 for the basic model, which is the homoscedastic version containing the local level and the lagged regression term.

Model	**α**	**β_1**	**AIC**
Basic	1.61	0.112	0.0
Basic + GARCH(U_3=0)	1.48	0.100	-29.9
Basic + GARCH with U_3	1.49	0.145	-34.0

Table 2: Number of murders per year in Montgomery County, Maryland

YEAR	Murder
1985	12
1986	8
1987	17
1988	19
1989	21
1990	25
1991	26
1992	21
1993	30
1994	34
1995	21
1996	13
1997	23
1998	13
1999	13
2000	12
2001	19
2002	32
2003	23
2004	18
2005	21
2006	19

Table 3: Extreme entries of the P-values chart for the annual data on murders in Montgomery County, Maryland

YEAR	Poisson	NBD
1993	0.0014	0.037
1994	8.7E-05	0.021
2002	0.00038	0.031

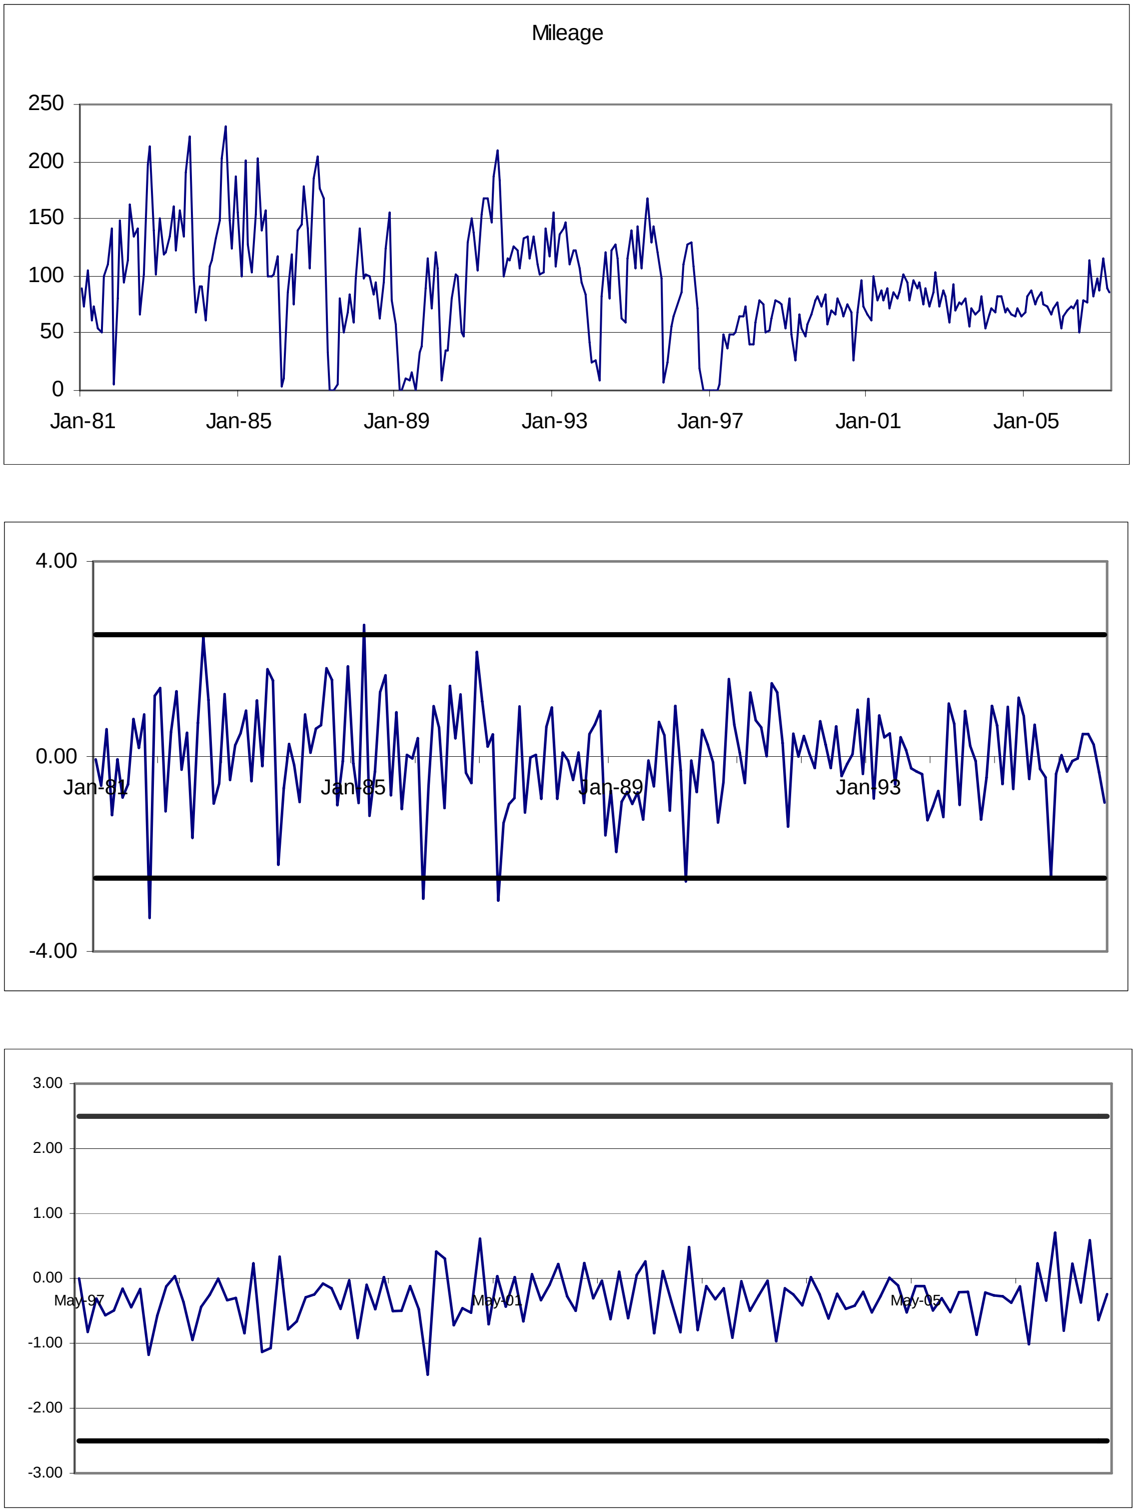

Figure 1: Plots of (a) Monthly mileage, 1/81 – 3/07; (b) standardized residuals for AR(1) model fitted over 1/81 – 9/96; (c) standardized one-step-ahead forecasts from same model for 6/97 – 3/07. The limits on the standardized charts are set at ±2.5 standard deviations.

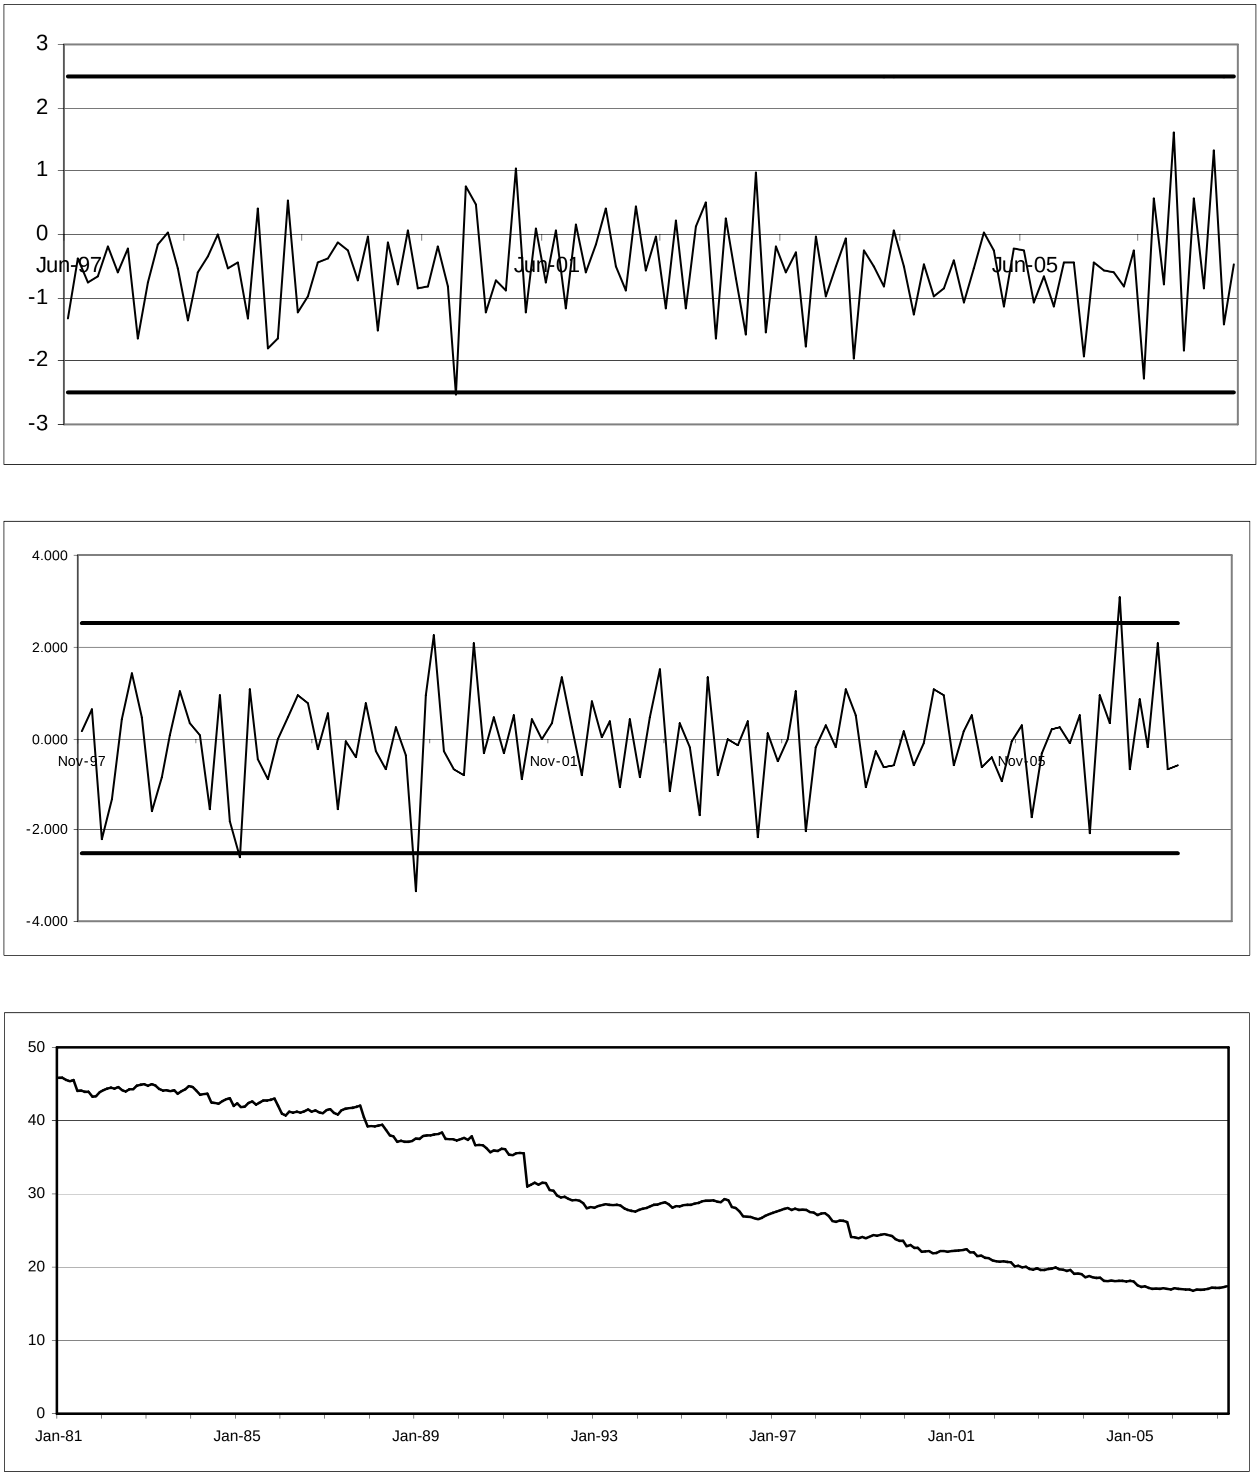

Figure 2: Plots of (a) standardized one-step-ahead forecast errors for the GARCH model for monthly mileage model for the period 6/97 – 3/07; (b) forecast errors for the same period after recalibrating the model; (c) standard deviations from the original model for the period 1/81 – 3/07. The limits on the forecast errors charts are set at ±2.5 standard deviations.

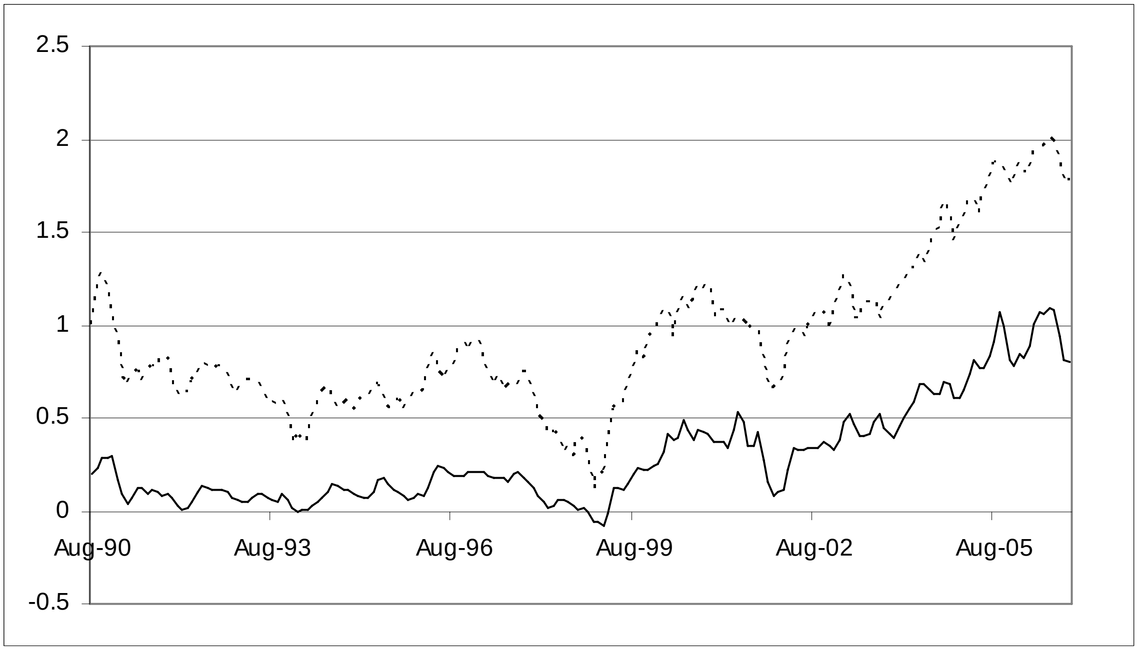

Figure 3: Plots for the log (gas price) [solid line] and log (spot price/10) [dotted line] series over the period 8/90 – 11/06. The spot price is scaled simply to place both series on the same diagram.

(a) Standardized residuals for the constant variance model

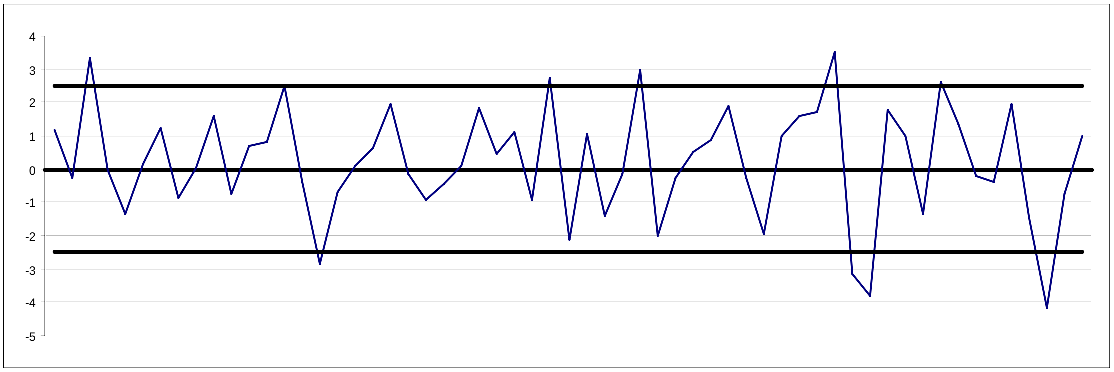

(b) Standardized residuals for the GARCH model

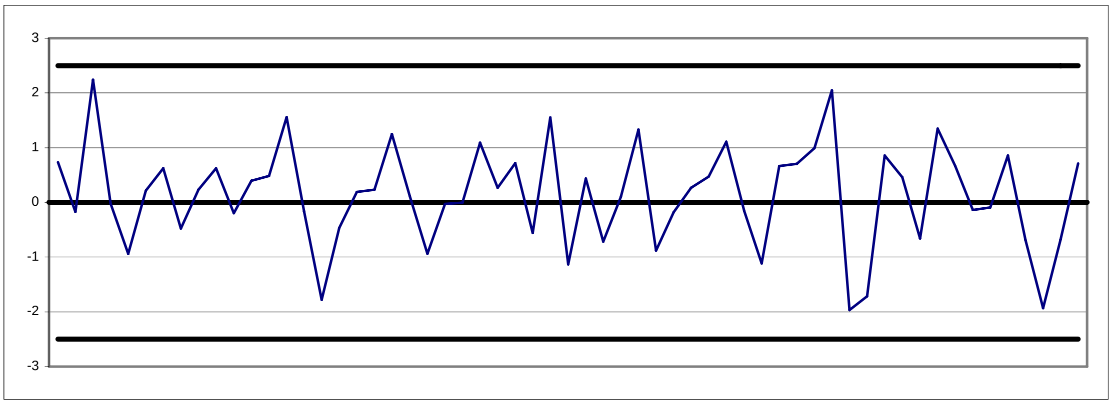

(c) Standard deviations for the GARCH model

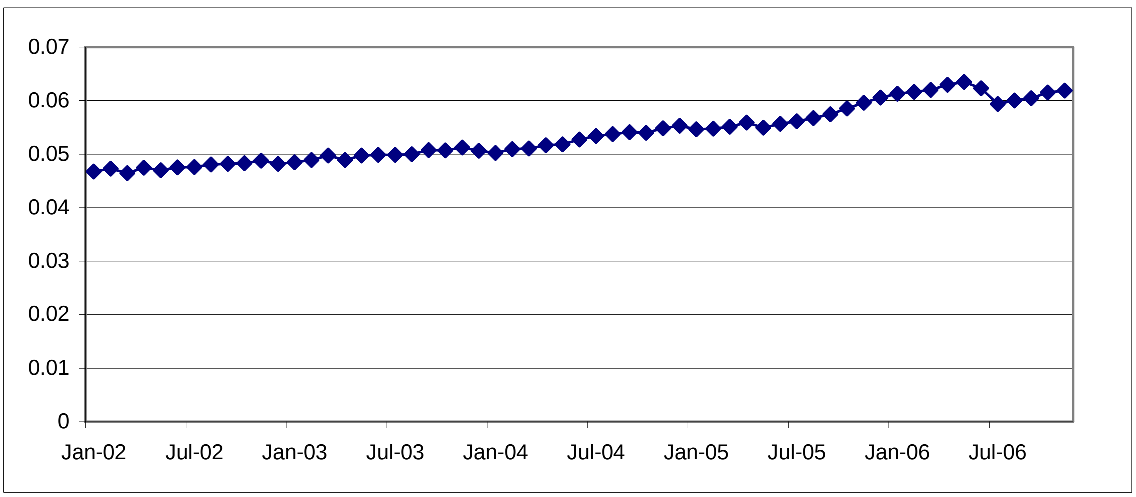

Figure 4: Plots for the log (gas price) series over the period 1/02 – 11/06. The limits on the standardized residuals charts are set at ±2.5 standard deviations.

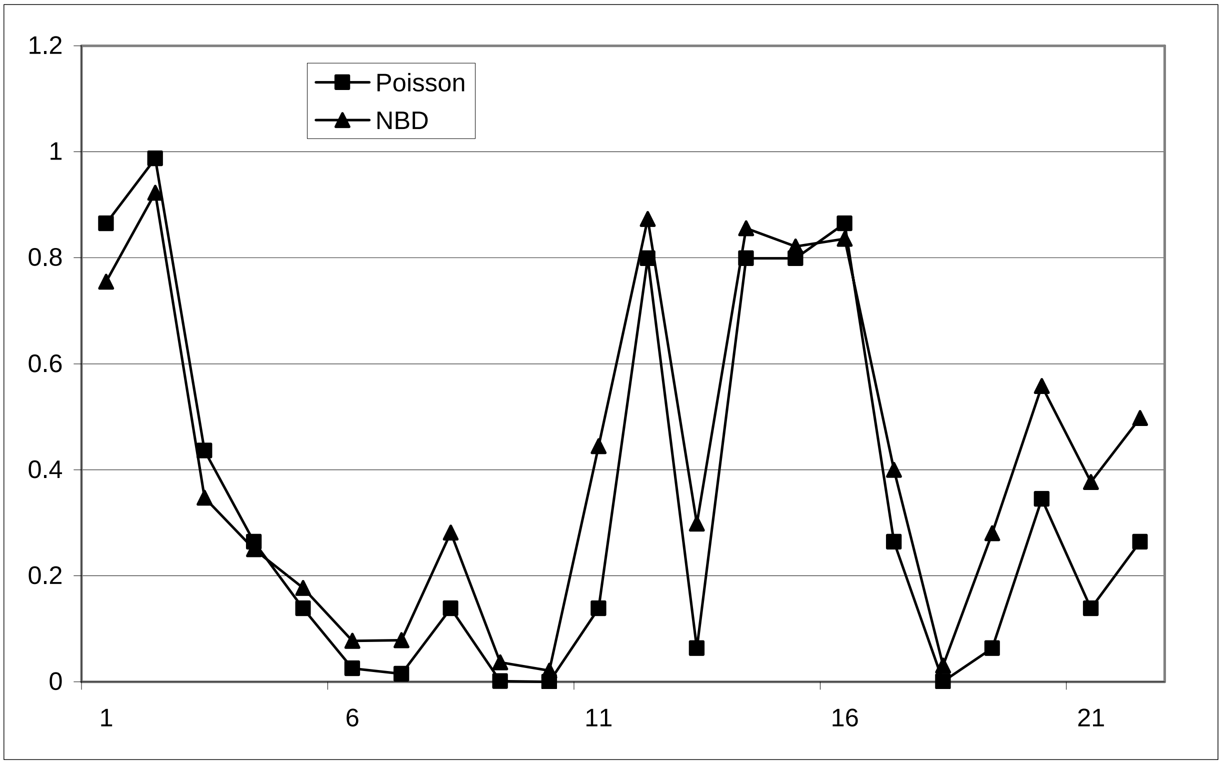

Figure 5: P-value chart for the annual murders series, based upon the Poisson (Poi) and negative binomial (NBD) distributions.