STATISTICAL DISTRIBUTION OF COITAL ACTS

by

Sutapa Das
and
Craig D. Turnbull

Department of Biostatistics
University of North Carolina

Institute of Statistics
Mimeo Series No. 2188

April 1998

STATISTICAL DISTRIBUTION OF COITAL ACTS

BY

SUTAPA DAS AND CRAIG D. TURNBULL

Department of Biostatistics
School of Public Health, CB #7400
University of North Carolina at Chapel Hill
Chapel Hill, North Carolina 27599-7400

INSTITUTE OF STATISTICS MIMEO SERIES NO. 2188

April 1998

Statistical Distribution of Coital Acts

INTRODUCTION

Quantitative data on the frequency of sexual intercourse are important for testing the effectiveness of contraceptives in preventing pregnancies and the spread of sexually transmitted diseases. This type of testing can be effectively conducted by using mathematical and statistical models. However, entering raw data into these models can be computationally cumbersome even with the current processing capabilities of modern computers, especially since most complex models consider a number of factors which may contribute to a response. For example, when examining the relative effectiveness of various contraceptive methods in preventing pregnancies, factors which can be incorporated into a simulation model include: 1) coital frequency, 2) the distribution of coital acts throughout the menstrual cycle relative to ovulation (Lachenbruch 1968), 3) cycle variability among women using different methods, 4) the capacity of women using different methods to conceive, and 5) the ability of women (or couples) to use each contraceptive method correctly (Steiner).

Given the complexity of such a model, a concise summary of relevant factors could be much more effectively incorporated into a model than could raw data. However, the use of only summary statistics, such as measures of central tendency and variance, can cause the loss of much of the information contained in original data. A more informative method would be to use probability distributions to approximate events included in models. Probability distributions can provide more statistical information by using only a few parameters than can just the mean and variance alone; in addition, probability distributions are useful for modeling purposes.

The purpose of this paper is to establish a methodology for determining a probability distribution which provides a reasonable fit to coital frequency.

LITERATURE REVIEW

In 1996, Kault analyzed the distribution of the number of sexual partners, using a Chi-square goodness-of-fit test to compare data from two Australian questionnaires (Table 1) to five probability distributions (Log-normal, Gamma, Negative Binomial, truncated Normal and Poisson). Table 2 gives the results for the overall goodness-of-fit in which data from both data sets are combined. (Kault also performs goodness of fit on each of four rows of Table 1 separately. To see the results of these analyses, refer to Kault 1996). For the overall goodness-of-fit, Kault assumes the four rows of Table 1 are independent. None of the five distributions fit the data when all categories are separate. However, the Log-normal distribution fits very well when categories 'no partner' and 'one partner' are combined and, in this situation, fits the data much better than the Gamma and Negative Binomial distributions. If the highest two 'number of partner' categories are combined the Gamma and Negative Binomial distributions perform much better than the Log-normal but they still do not provide a very good fit to the data. When both the top two and the bottom two 'number of partners' categories are combined the log-normal again gives the best fit to the data.

Table 1: This table shows the two data sets analyzed by Kault. These data sets were supplied by the National Centre for HIV Social Research, Macquarie University. The subjects are first year behavioural science students at certain universities in New South Wales who were surveyed in 1991 (data set 1) and in 1990 (data set 2).

Data set 1

	Number of partners						Total
	None	1	2-4	5-10	11-20	>20	
Female	255	143	166	43	20	4	**631**
Male	61	37	94	45	13	10	**260**

Data set 2

	Number of partners								Total
	None	1	2	3	4	5	6	7-20	>20	
Female	296	134	67	43	21	14	8	18	4	**605**
Male	92	48	44	37	25	11	7	38	6	**308**

Table 2: Overall Chi-square goodness-of-fit. Degrees of freedom for each test are equal to $(n-4k-4)$ where n is the number of categories in the data and k is the number of parameters fitted.

	Log-normal	Gamma	Negative Binomial	Truncated Normal	Poisson
All categories separate (30 categories)	97.6 (6×10^{-13})*	96.0	148.1	$>10^{7}$	9×10^{6}
Categories 0 and 1 combinded (26 categories)	15.7	82.0	98.7	$>10^{7}$	4×10^{6}
Highest two categories combined (26 categories)	91.4 (2×10^{-12})	41.2 (2×10^{-4})	43.8 (1×10^{-4})	940 ($< 10^{-16}$)	6×10^{5}
Categories 0 and 1 and highest two categories combined (22 categories)	10.8 (0.37)	27.7 (0.0023)	33.1 (3×10^{-4})	467 ($< 10^{-16}$)	2×10^{5}

*p-values are omitted for any χ^2 calculation for which the expected value in any cell is less than 5

The main purpose of Kault's analysis is to find a model that can be used to predict the likely long term course of the AIDS epidemic. Hence distributions which provide a good fit to the right hand side of the distribution of the number of sexual partners are the most beneficial, since the people on this side of the distribution will have the greatest role in the dynamics of the epidemic. Those with 0 or 1 partners will have a negligible role. In light of this, Kault suggests the Log-normal distribution as the best fitting probability distribution.

The methods and ideas presented by Kault in his analysis provided a useful guide for the analysis described in this paper.

DESCRIPTION OF DATA

The data used in this analysis are from a recent oral contraceptive-use behaviour study spanning up to three menstrual cycles for each participant. (Oakley). The data was collected from November 1994 to December 1995. The study involves 102 women recruited from two university student health clinics (n=76) and two Title X-funded clinics (n=26) in Michigan and North Carolina. One university clinic and one Title X clinic were included from each state. To qualify for the study, women had to be 18 years of age or older, eligible for the type of pill used in the study and could not

have used oral contraception in the six months prior to enrollment in the study. Monthly diary-cards were used to record frequency of sexual intercourse. Data from a total of 286 cycles were collected. All 102 women provided data for their first cycle, 95 provided data for their second cycle and 89 for their third cycle. Eighty-six women provided data for all three cycles. (See Appendix 1 for a frequency table of these data.)

It should be noted that this sample does not accurately represent the general population of women in the United States, since 77 of the 102 women were recruited from student health clinics, and all women were oral contraceptive users. This may introduce bias into the results.

METHODS

The variable of interest in the oral contraceptive-use data set is *totsex*, the number of coital acts per menstrual cycle. Since the statistical tests employed in this research assume independence between observations, the 286 cycles cannot be treated equally, as each individual woman's cycles are correlated. Therefore, analysis of *totsex* is conducted using data provided for cycle one only ($n = 102$). A new variable, *avgsex*, is created by averaging the *totsex* variable over the three cycles for women who provided data for all three cycles ($n = 86$). Since the *totsex* variable is independent between women, so is *avgsex*. The types of analyses conducted using *totsex* from cycle one are repeated using the *avgsex* variable.

The *totsex* and *avgsex* variables are modeled by using the Poisson, Negative Binomial, Gamma and Log-normal distributions. These probability distributions are used because they apply to non-negative data and they allow for indefinitely long right tail probabilities. The four distributions are described below. (Appendix 2 contains graphical representations of each distribution.)

Distributions

Poisson Probability Distribution

The Poisson probability distribution provides a model for rare events that occur in space, time, or volume. In this case, the menstrual cycle can be divided into many subintervals, each of which is so small that the probability of a woman having sexual intercourse in any subinterval more than once is essentially equal to zero. The frequency of coital acts during a menstrual cycle is equal to the total number of subintervals that contain one coital act. The formula for the Poisson distribution is:

$$p(x) = \frac{\lambda^x}{x!}e^{-\lambda}, \qquad x = 0, 1, 2, \ldots, \quad \lambda > 0$$

where x is the frequency of sexual intercourse during the menstrual cycle, $p(x)$ is the probability that the frequency of sexual intercourse equals a particular value of x, λ is the average number of coital acts, and e is a constant.

Negative Binomial Distribution

The Negative Binomial distribution involves a number of identical and independent trials or subintervals, each of which can result in either success or failure. The probability of success on each trial is p, and the probability of failure on each trial is $q = (1 - p)$. If r is any positive integer greater than 1, then the Negative Binomial distribution describes the number of the trial on which the rth success occurs. The formula for this distribution is:

$$p(x) = \binom{x-1}{r-1} p^r q^{x-r}, \quad x = r, r+1, r+2, \ldots, \quad 0 \leq p \leq 1$$

where x is the frequency of sexual intercourse and $p(x)$ is the probability that the frequency of sexual intercourse equals a particular value of x.

The two continuous distributions (Gamma and Log-normal) are used here to model coital frequency because *avgsex* is essentially a continuous variable, and it may be fit by a continuous distribution better than a discrete distribution. Since continuous distributions are often used to model discrete events, these distributions are also compared to the *totsex* variable. The use of these particular continuous distributions is based on the analysis performed by Kault.

Gamma Distribution

The formula for the Gamma distribution is:

$$f(x) = \begin{cases} \frac{x^{\alpha-1}e^{-x/\beta}}{\beta^{\alpha}\Gamma(\alpha)}, & \alpha, \beta > 0;\ 0 < x < \infty \\ 0, & elsewhere \end{cases}$$

where x is the frequency of sexual intercourse, $f(x)$ is the probability density function for x, and $\Gamma(\alpha)$ is the gamma function. To find the probability that the frequency of sexual intercourse falls within a specific interval (i.e., $Pr(a < X < b)$), the corresponding area under $f(x)$ can be found by integration.

The parameter α is called the *shape parameter* because the shape of the Gamma density changes when different values of α are used. The parameter β is called the *scale parameter* because multiplying a Gamma-distributed variable by a positive constant (i.e., changing the scale on which the measurement is made) produces a variable that has a Gamma distribution with the same α value but a different β value.

The Gamma distribution is used to model populations which possess a positively skewed frequency distribution. For example, it can be used to model the lengths of time between malfunctions for aircraft engines, the lengths of time between arrivals at a supermarket checkout line, or the lengths of time to complete a maintenance checkup for an automobile or aircraft engine.

In the special case when $\alpha = 1$, the Gamma density function is called the Exponential density function. The Exponential density function is widely used to describe events recurring at random in time. It is also useful for modeling lifetime, such as for electronic components or radioactive matter. The Gamma distribution can be used to make adjustments to the Exponential distribution in representing these lifetimes. For example, the Gamma distribution provides a useful failure-time model for a system under continuous maintenance if it initially experiences some wear or degradation, but reaches a stable state of repair as time goes on.

Log-normal Distribution

If the random variable U is Normally distributed with mean μ and variance σ^2, and if $X = e^U$, then X has a Log-normal distribution. The formula for the Log-normal distribution is:

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-(\ln x - \mu)^2/2\sigma^2}, \qquad 0 < x < \infty;\ -\infty < \mu < \infty;\ \sigma > 0.$$

The Log-normal distribution is widely used for modeling purposes. For example, it can be used to model the distribution of particle size in natural aggregates, critical dosages in drug applications, and the duration of doctors' consultations.

The parameter(s) of each of the four distributions mentioned above can take on an infinite number of values which alter the shape and/or location of the distributions in subtle or drastic ways depending on the particular distribution (See Appendix 2). The collection of these variations of each distribution is called the *family* of the distribution. Because of the wide variety of shapes and/or locations within each family, the range of parameters to be assessed can be narrowed to exclude values of

specific distributions which clearly do not fit the sample data. In the current analysis, this is done by the use of quantile-quantile plots. A range of parameters to be further assessed by the Chi-squared and Kolmogorov goodness-of-fit tests is determined for each family by using the quantile-quantile plots to reject parameter values of clearly inappropriate distributions.

Quantile-quantile Plots

A *quantile-quantile plot* depicts the theoretical quantiles (cumulative percentage points) of a theoretical probability distribution against the empirical quantiles from an observed set of sample data. If the theoretical distribution adequately fits the sample data, the theoretical and empirical quantiles should be almost equal, and the quantile-quantile plot should approximate a straight line. Here, two different procedures based on quantile-quantile plots are used to define the ranges of parameter values which undergo further testing.

Gamma and Log-normal Distributions: Normal Quantile-quantile Plots

In the first stage, *A Folio of Distributions: A Collection of Theoretical Quantile-Quantile Plots* by Edward B. Fowlkes is used to reject clearly inappropriate parameter values for the Log-normal and Gamma distribution families. This book contains theoretical *Normal quantile-quantile plots* based on a wide range of parameter values from many different families of continuous distributions. These plots show theoretical quantiles from each distribution plotted against the quantiles from the Standard Normal distribution (See Figures 4 and 5 for examples of plots based on the Log-Normal and Gamma distribution families). Because probability distributions are unique, no two Normal quantile-quantile plots are exactly the same, although plots based on distributions which are similar in shape and location depict curves which are also similar to one another.

A sample of data drawn from a population with cumulative distribution function $F(x)$, can be expected to have a Normal quantile-quantile plot very similar to the theoretical Normal quantile-quantile plot based on $F(x)$. When $F(x)$ is unknown, the sample Normal quantile-quantile plot can be compared to Normal quantile-quantile plots based on theoretical distributions which are possible candidates for $F(x)$. Because quantile-quantile plots are graphical in nature, this is a visual comparison based on the relative closeness of the shapes of the various theoretical Normal quantile-quantile plots to the plot based on the sample data. Distribution parameters producing plots which are relatively very different from the sample plot compared to plots of other parameters can be rejected as being inappropriate models. The remaining distributions can be assessed by goodness of fit tests for more precise results.

For this stage of the analysis, Normal quantile-quantile plots based on the sample data are created (See Figure 1). It should be noted that these particular quantile-quantile plots are created only for this part of the analysis in which inappropriate parameters are rejected, and are not used to assess the appropriateness of the Standard Normal distribution to the sample data. The Normal distribution is not an appropriate fit for coital frequency since it has negative values and allows for indefinitely long left tail probabilities (See Appendix 2 for a graphical representation of the Normal distribution). This distribution is used purely as a reference, through the use of Normal quantile-quantile plots, to compare the appropriateness of the Gamma and Log-normal distributions to the sample data.

Because the two plots shown in Figure 1 are similar, it is likely that the *totsex* and *avgsex* variables can be fitted by the same probability distribution. The lines added to the plots in Figure 2 show the paths that the plots would have approximated had the data been normally distributed. The shapes of these plots reveal that the data are positively skewed relative to the Normal distribution. Each plot can be divided into

Figure 1: Standard Normal quantile-quantile plots based on sample variables

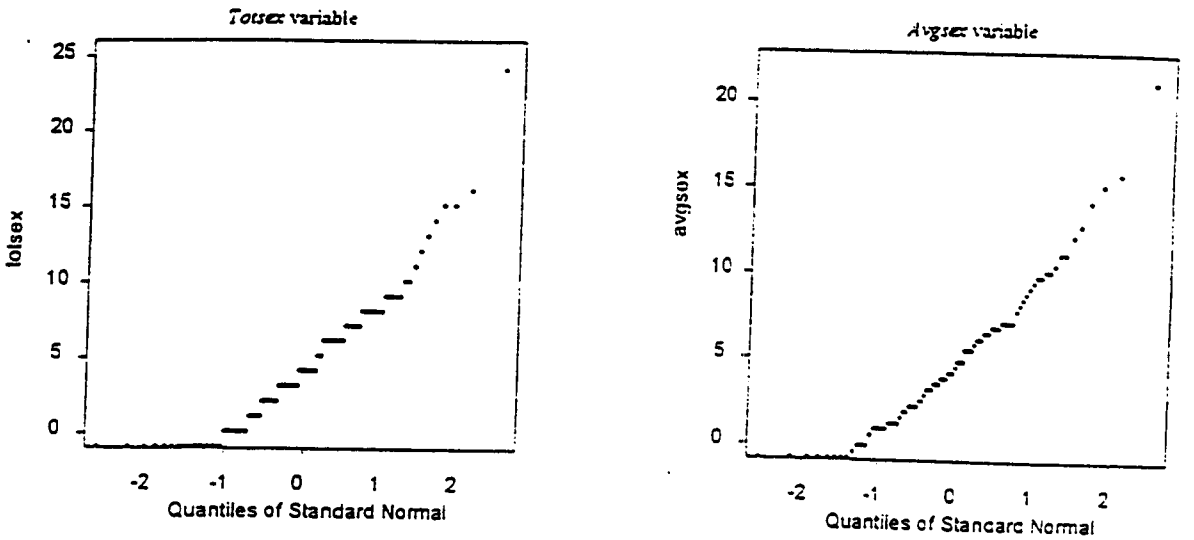

three parts as shown in Figure 3. The first part forms a horizontal line corresponding to the excess of zeros in the sample data. The second part forms an upward sloping line corresponding to the central values of the data. From looking at the graphs, these values appear to be between one and ten for *totsex* and one and seven for *avgsex*. The third part forms a more steeply upward sloping line which corresponds to the largest

Figure 2: Standard Normal quantile-quantile plots based on sample variables.
The line indicates the path the plot would follow if the data were normally distributed

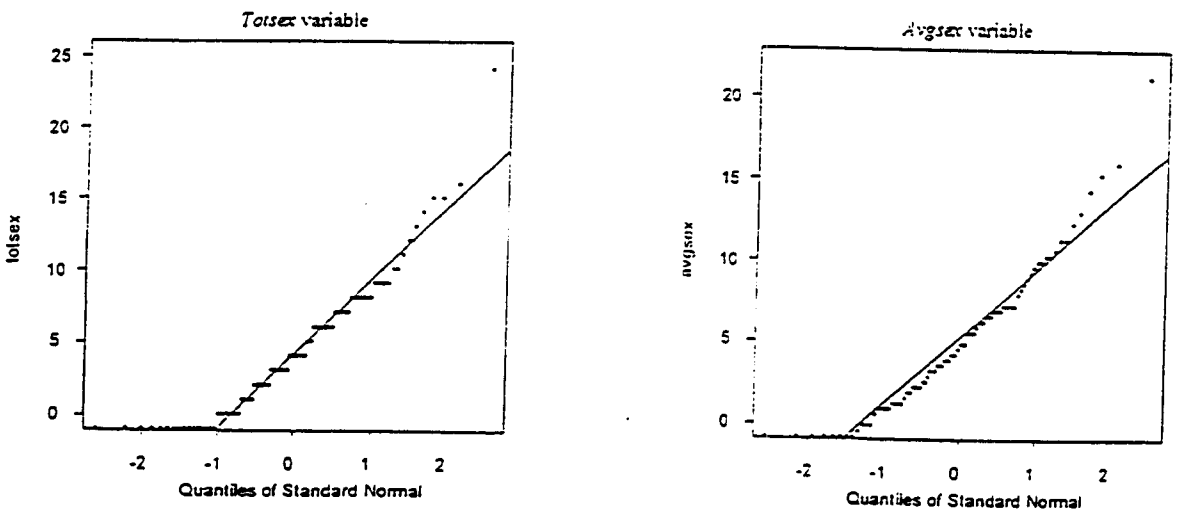

Figure 3: Standard Normal quantile-quantile plots based on sample variables divided into three regions

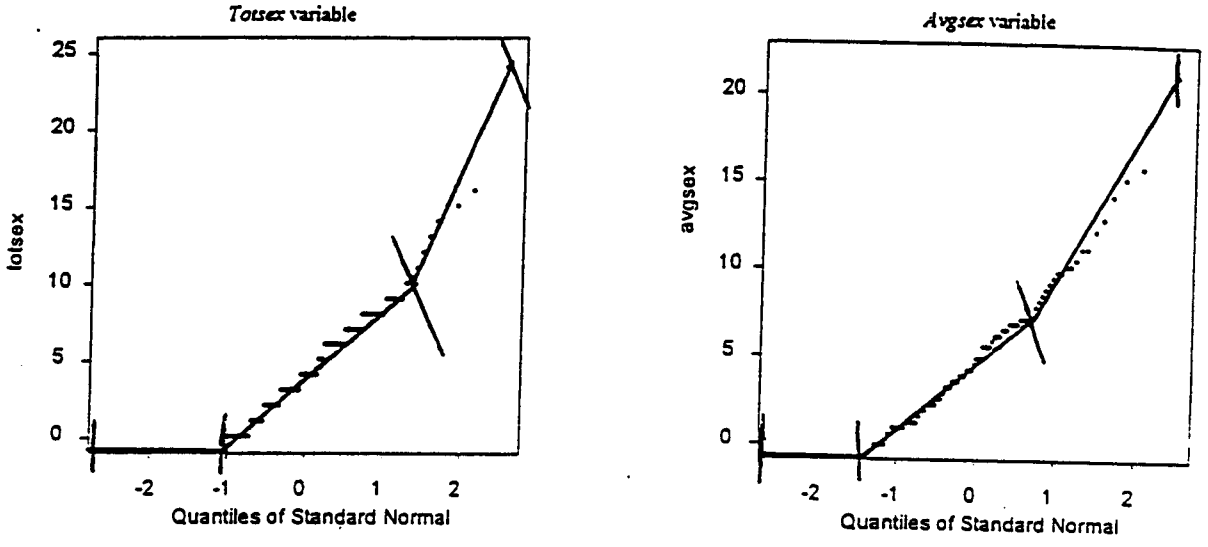

data values. When the two sample plots are compared to the plots based on theoretical distributions (See Figures 4 and 5), the relative skewness to the Normal distribution is considered, as is the relative divergence of the theoretical quantile-quantile plots from these three regions of the sample plots.

Gamma Distribution

Figure 4 shows four composite Standard Normal quantile-quantile plots based on members of the Gamma family of distributions. The scale parameter, β, is held constant at 1.0, while the shape parameter, α, varies from 0.1 to 50. Since quantile-quantile plots are invariant to changes of scale and location, it is unnecessary to compare plots based on different values of β to the sample quantile-quantile plot. This also allows the scales of the four plots in Figure 4 to vary for easier comparisons and interpretations. Cumulative probabilities for the Standard Normal distribution are marked at the top of each of the four graphs. These probabilities are not provided for the Gamma distribution since they depend on the value of α and differ for each curve.

Figure 4: Standard Normal Quantile-quantile plots based on the Gamma distribution for $\beta = 1.0$ (Fowlkes)

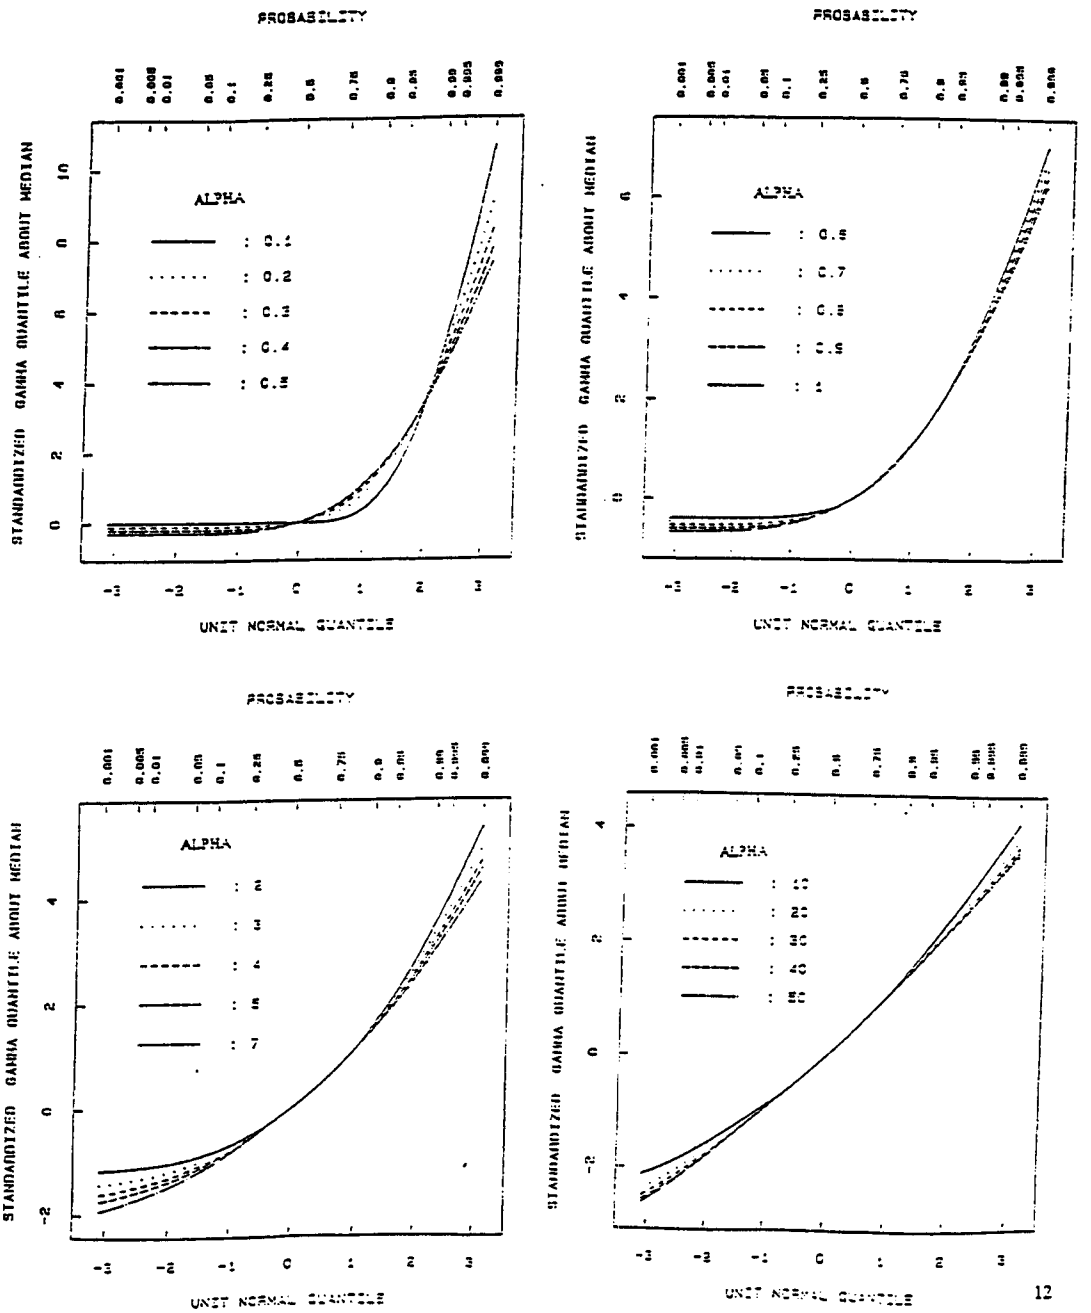

A comparison of these quantile-quantile plots reveals that as α gets larger, the shape of the Gamma distribution approaches that of the Normal distribution, and that as α gets smaller, the Gamma distribution becomes very skewed relative to the Normal distribution.

For an example of how the range of parameter values is chosen for the Gamma distribution based on Normal quantile-quantile plots, the four theoretical plots in Figure 4 are compared to the sample plots in Figure 3. The theoretical plots corresponding to values of α between 0.2 and 7.0 are of a similar skewness to the data; additionally, these curves can be divided into three regions similar to those marked on the sample plots in Figure 3. Subsequently, these parameter values cannot be rejected as possible matches to the data, and should be further tested by goodness-of-fit.

When α is 0.1, the Gamma distribution is much more skewed than the sample data, and the theoretical quantile-quantile plot is most easily divided into two regions instead of the three corresponding to the sample plot. Because of this, the Gamma distribution is unlikely to fit the sample data when α is equal to 0.1; this value is rejected from further consideration.

When α is greater than or equal to ten, the near linearity of the quantile-quantile plots indicates that the shape of the Gamma distribution is close to that of the Normal distribution. The Gamma distributions based on these parameter values are not as skewed as the sample data. Because these plots appear almost linear, they cannot be divided into regions corresponding to the sample quantile-quantile plot. The Gamma distribution is unlikely to fit the sample data when α is greater than or equal to ten, and so these values are rejected from further consideration.

LOG-NORMAL DISTRIBUTION

Figure 5 shows four composite Standard Normal quantile-quantile plots based on members of the Log-normal family of distributions. The scale parameter, e^{μ}, is held constant at 0, while the shape parameter, σ, varies from 0.1 to 2.4. A comparison of the plots reveals that as σ gets larger, the Log-normal distribution becomes very skewed relative to the Normal distribution.

For an example of how the range of parameter values is chosen for the Log-normal distribution, the four theoretical plots in Figure 5 are compared to the sample plots in Figure 3. The theoretical plots corresponding to values of σ between 0.2 and one are of a similar skewness to the data; additionally, these curves can be divided into three regions similar to those marked on the sample plots in Figure 3. Subsequently, these parameter values cannot be rejected as possible matches to the data, and should be further tested by goodness-of-fit.

When σ is 0.1, the quantile-quantile plot appears to be almost linear, and cannot be divided into regions corresponding to those of the sample quantile-quantile plot. Because of this, the Log-normal distribution is unlikely to fit the sample data when σ is equal to 0.1; this value is rejected from further consideration.

When σ is greater than one, the Log-normal distribution is much more skewed than the sample data, and the theoretical quantile-quantile plots are most easily divided into two regions instead of three. The Log-normal distribution is unlikely to fit the sample data for these values of σ, and so these values are rejected from further consideration.

Negative Binomial and Poisson Distributions: Sample Quantile-quantile Plots

The above method cannot be used to exclude inappropriate parameter values for the two discrete distributions because a comprehensive collection of theoretical Normal quantile-quantile plots for these distributions could not found. Instead, the

Figure 5: Standard Normal Quantile-quantile plots based on Log-normal distribution for $e^{\mu} = 0$ (Fowlkes)

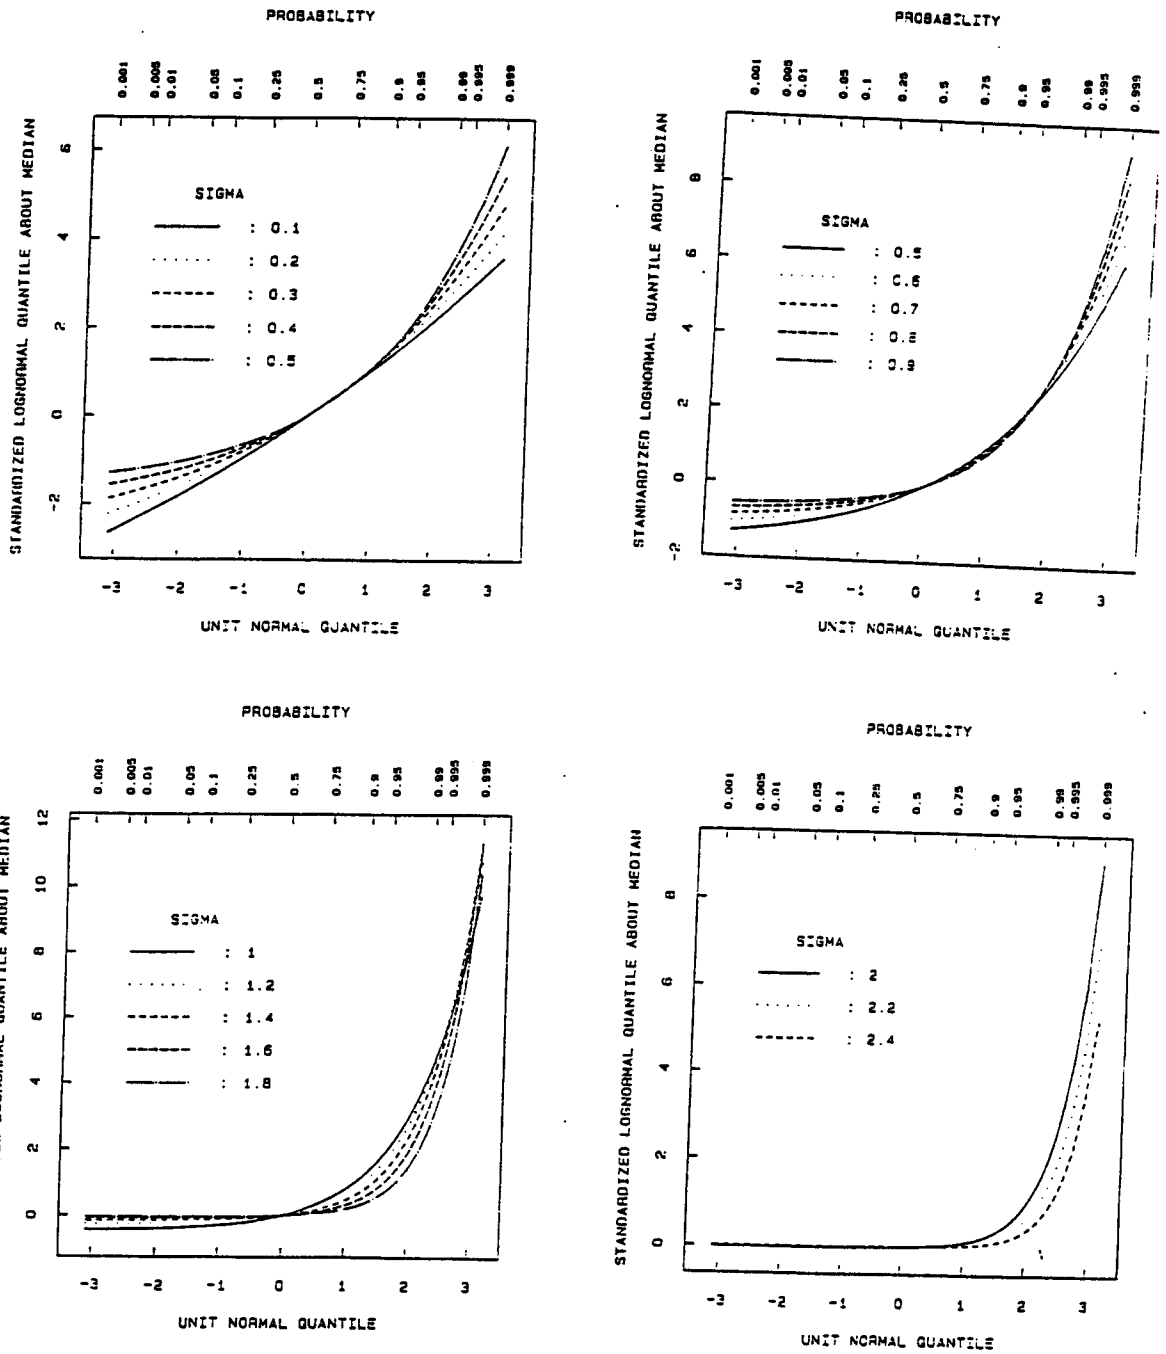

empirical quantiles from the sample data are plotted directly against the theoretical quantiles of the Poisson and Negative Binomial distributions (See Figures 6-9). For this part, if a theoretical distribution resembles the distribution of the sample data, the quantiles should be approximately equal and the quantile-quantile plot should approximate a straight line. The parameters of distributions whose quantiles clearly do not fit those of the sample data are rejected. Again, this rejection is based on a visual comparison in which parameters producing quantile-quantile plots relatively less linear than plots based on other parameter values are less likely to fit the sample data. Through this process, a range of parameter values is selected from the two distribution families for further testing using the goodness-of-fit tests.

Negative Binomial Distribution

Figures 6 and 7 show quantile-quantile plots of the sample data and different parameter values of the Negative Binomial distribution. For an example of how the ranges of values are chosen for this distribution, each plot is compared to a straight line. For both the *totsex* and *avgsex* variables, the values ($r = 20$, $p = 0.1$) can be rejected from further consideration, since these two plots could easily be divided into two regions, and therefore do not approximate a straight line. The plots for the other parameter values appear almost linear, and so these values should be further tested by goodness-of-fit.

Poisson

Figure 8 and 9 shows quantile-quantile plots of the sample data and different parameter values of the Poisson distribution. For an example of how the range of values is chosen for this distribution, each plot is again compared to a straight line. For both the *totsex* and *avgsex* variables, the values ($\lambda = 10, 20$) can be rejected from further consideration, since these plots could easily be divided into two regions, and

Figure 6: Quantile-quantile plots of totsex against Negative Binomial distribution

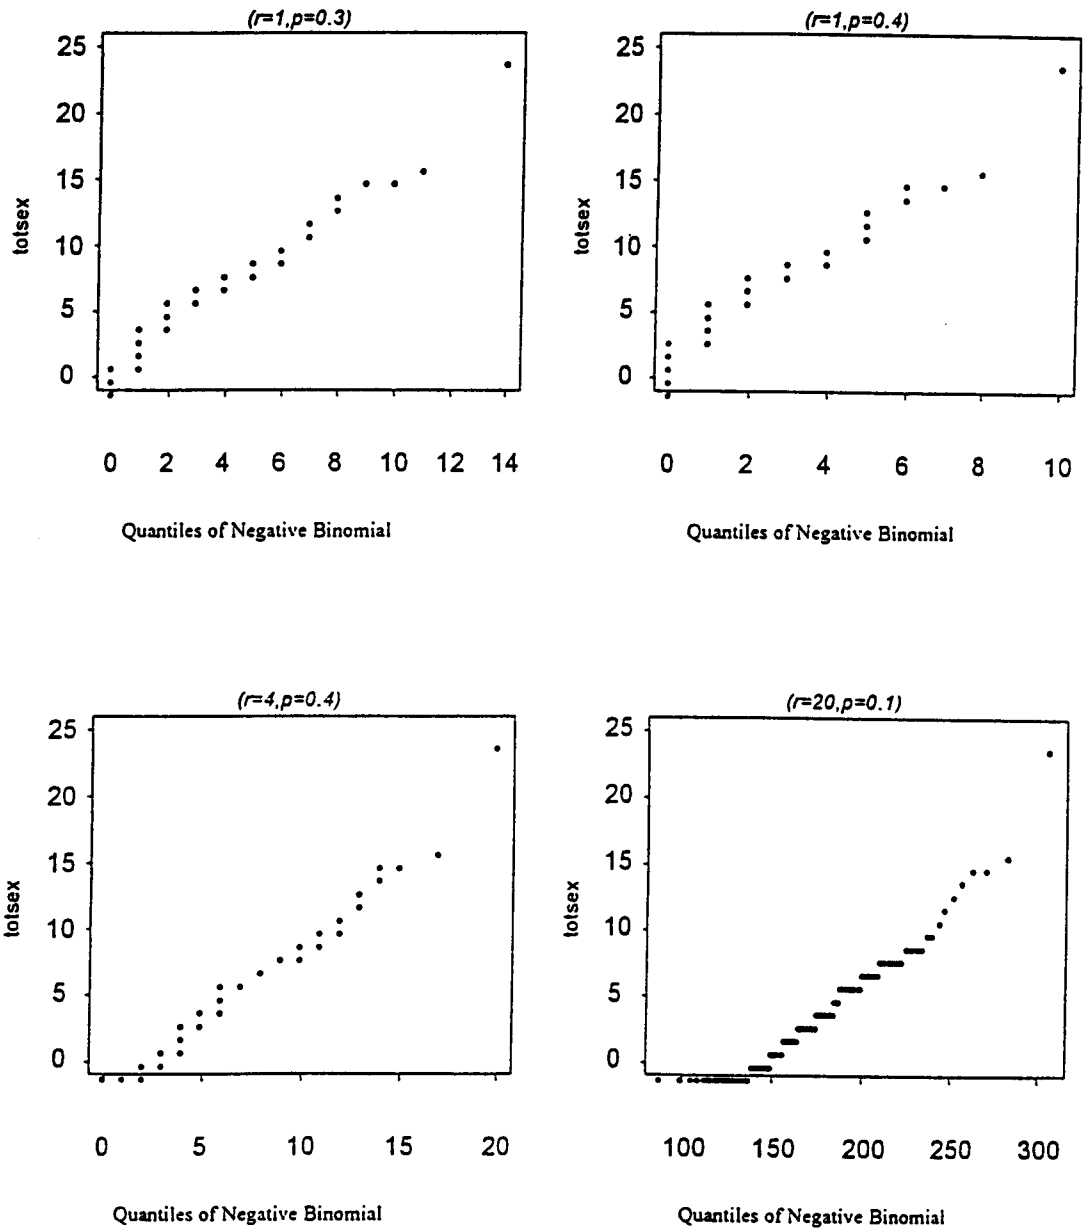

Figure 7: Quantile-quantile plots of avgsex against Negative Binomial distribution

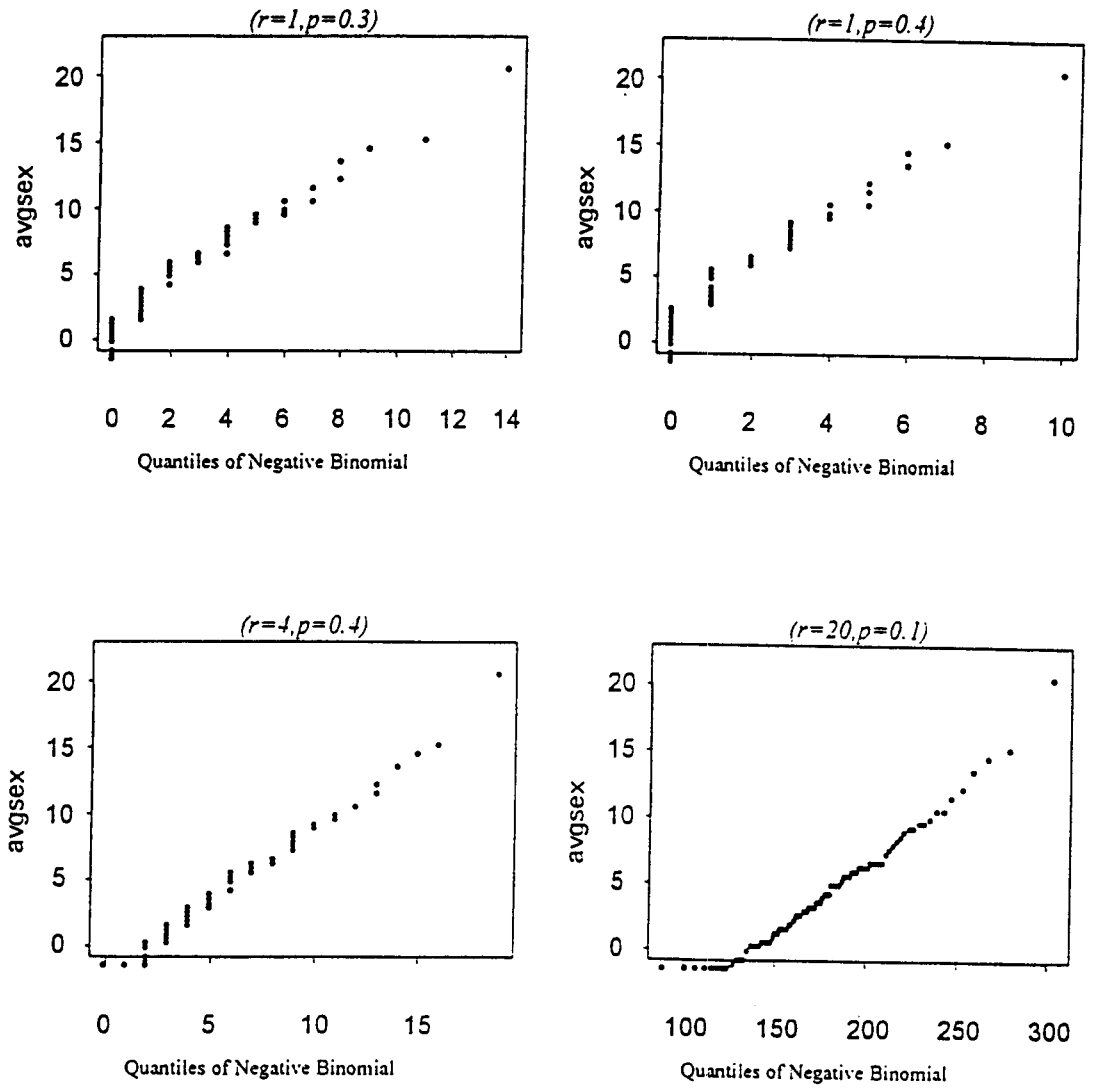

Figure 8: Quantile-quantile plots of totsex against Poisson distribution

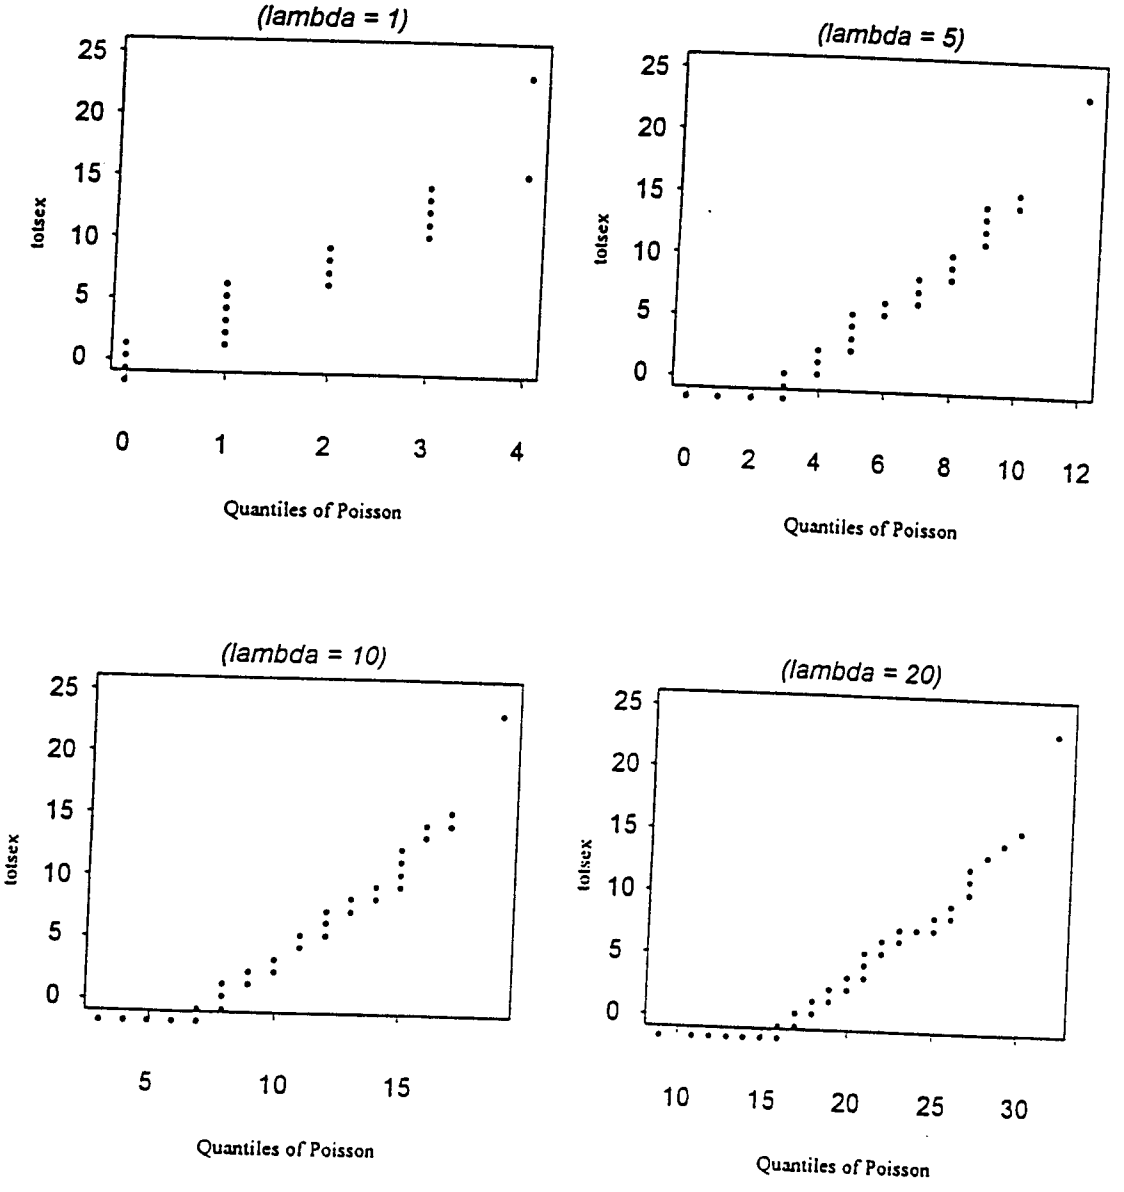

Figure 9: Quantile-quantile plots of avgsex against Poisson distribution

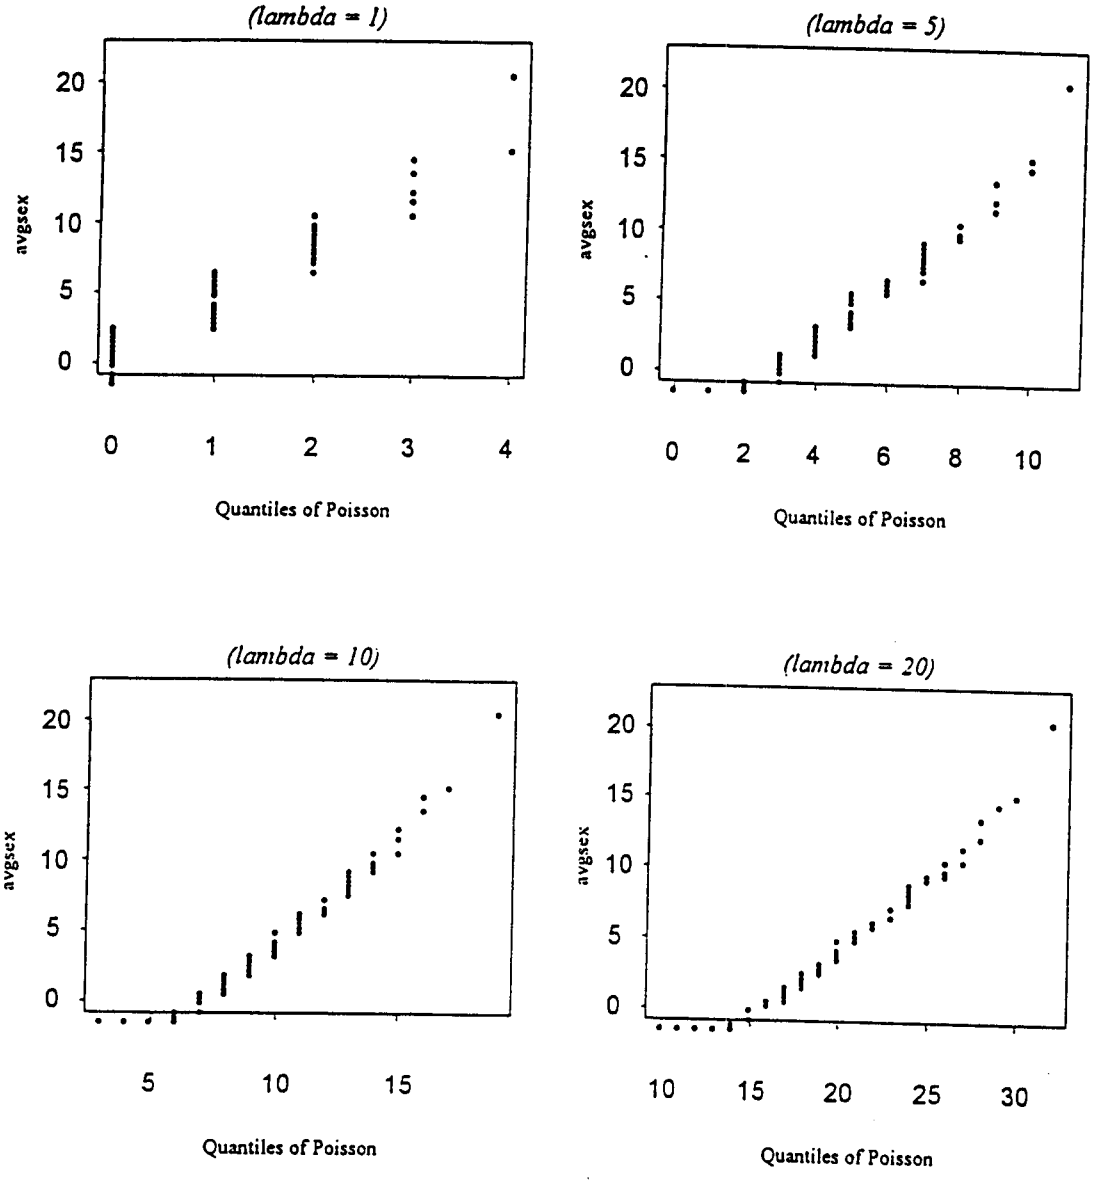

therefore do not approximate a straight line. The plots corresponding to ($\lambda = 1, 5$) appear almost linear, and so these values should be further tested by goodness-of-fit.

GOODNESS-OF-FIT TESTS

Chi-square Goodness-of-Fit Test

Goodness-of-fit tests are commonly used to assess the agreement between observed and expected results. Pearson originated the *Chi-square goodness-of-fit* test to determine the consistency of a sample of observed data to a proposed probability distribution (Koehler). To use this test, the set of all possible outcomes from the proposed distribution are divided into r intervals which are mutually exclusive and inclusive. The actual number of observations in each interval is compared to the number expected. If O_i is the sample number of observations in the i-th interval and E_i is the expected number of observations in the same interval given the particular distribution of interest, then the quantity

$$\chi^2 = \sum_{i=1}^{r} \frac{(O_i - E_i)^2}{E_i}$$

is used to determine if the fit of the proposed theoretical distribution to the data is good (i.e., adequate to describe the sample data). When the expected counts in all intervals are sufficiently large (i.e., no expected frequency is less than one and at least 80% of expected frequencies are greater than five) and the proposed distribution fits the data, this result follows the Chi-square distribution with $(r - 1)$ degrees of freedom. When the above quantity is much greater than that expected from the Chi-square distribution, a large discrepancy exists between the observed and expected counts, and the probability distribution does not fit the sample data. The *p-value* is the probability associated with the goodness-of-fit statistic which indicates the

likelihood of a particular sample being observed if the proposed distribution is correct. A small p-value (e.g., 0.01) indicates that this particular distribution does not fit the sample data well. In this analysis, a p-value greater than 0.01 indicates that the distribution fits the data.

The Chi-square goodness-of-fit test is used here to assess the appropriateness of the two discrete probability distributions (Poisson and Negative Binomial) to the coital frequency data. By nature, discrete distributions are already categorized by outcome (e.g., possible values for *totsex* = 0, 1, 2, . . .). However, the current sample size is too small to use each observed outcome as an interval since the requirements of the Chi-squared goodness-of-fit test (i.e., no expected frequency is less than one and at least 80% of expected frequencies are greater than five) are not adequately met. Therefore, the outcomes must be further categorized. During this part of the analysis, there is much debate regarding how to best create the intervals to be used for the test. One argument is that the intervals should be created according to the range of *observed* outcomes, and that these intervals should be meaningful for the data. This is difficult to do since the only meaningful interval created in this manner is that containing the outcome zero, which is the absence of coital acts during the menstrual cycle. The other intervals would be created in an arbitrary manner which satisfies the requirements of the Chi-square goodness-of-fit test.

An alternative argument suggests basing the intervals on the theoretical distribution given the specific parameter values. This method involves dividing the range of possible outcomes into a number of approximately equiprobable intervals, with all intervals wide enough to meet the requirements for the goodness of fit test. This method can be considered the default method when meaningful intervals cannot be created. The advantage of this method is that the requirements of the Chi-square goodness-of-fit test are met. It should be noted that the degrees of freedom based on

this method are $(r - 1 -$ the number of parameters), since the parameters are estimated.

After much discussion with peers and professors (Koch, MacDougall, Taff, Turnbull, Zink), and consultation with the Encyclopedia of Biostatistics (Koehler, Atkinson), intervals are created according to the second argument. Although some argue that this method may be statistically 'naive' (Turnbull), it is believed to be the more logical choice by the author, and is adequate for the needs of this research.

The set of possible outcomes for each distribution is divided into five intervals, each with probability of approximately 0.20. Because of the nature of discrete distributions, it is not always possible to have an equal probability in each interval. The expected values are compared to the number of outcomes actually observed in each interval. P-values are used to assess the goodness of fit for each distribution.

Kolmogorov Goodness-of-fit Test

The *Kolmogorov goodness-of-fit* test is used to assess the agreement of the continuous distributions (Log-normal and Gamma) to the coital frequency data. The Kolmogorov goodness-of-fit test is used to test the appropriateness of the continuous distributions because it avoids the drawbacks associated with testing continuous distributions using the Chi-square goodness-of-fit test such as having to artificially categorize the set of possible outcome values.

If $F(x)$ is the distribution of the population from which the sample data is drawn, and $G(x)$ is the proposed theoretical distribution, then the Kolmogorov parameter is the maximum absolute distance between these two distributions:

$$D = max|F(x) - G(x)|.$$

If the distributions are identical, meaning that $F(x) = G(x)$ for all possible values of x, then the Kolmogorov statistic is zero. From a graphical point of view the Kolmogorov parameter is the largest vertical distance between the two cumulative distribution functions. The Kolmogorov distance is estimated between the distributions F and G by

$$\widehat{D} = max|\widehat{F}(x) - G(x)|$$

where $\widehat{F}(x)$ is the distribution of sample data. A sufficiently large $\widehat{D}$ indicates that the proposed theoretical distribution is an inappropriate model for the sample data. The p-value associated with the Kolmogorov statistic is used to assess the adequacy of the fit.

The SPLUS statistical analysis package is used to create the quantile-quantile plots and to determine goodness-of-fit (See Appendix 3 for SPLUS programs).

RESULTS

Appendix 4 shows the results of the goodness-of-fit tests. Because many parameter values are tested, some results for distributions which clearly do not fit the data have been excluded.

The tables for the two discrete distributions show the Chi-square statistics and the corresponding p-values for the goodness-of-fit tests. Additionally, interval breaks for each test are shown, as well the observed and expected observations in each interval. The expected counts are not distributed equally among the five intervals because the intervals are only *approximately* equal to 0.20. The tables for the two continuous distributions show the Kolmogorov statistics and the corresponding p-values for the goodness-of-fit tests. P-values less than 0.0001 are recorded as '0'; the corresponding distributions clearly do not fit the data.

For the *avgsex* variable, several p-values corresponding to parameters of the Negative Binomial distribution are greater than 0.01 (See Table 3) These parameters provide a good fit to the data. The other distributions clearly do not fit this variable ($p < 0.001$). None of the distributions fit the *totsex* variable ($p<0.01$).

Table3: A summary of the results of the goodness-of-fit tests of Negative Binomial parameters which fit the data well (p>0.01). These results indicate that the Negative Binomial distribution fits the data.

r	*p*	GOF	P-value	Interval breaks	Observed values	Expected values
3	0.3	4.36	0.1130	-0.5, 3, 5, 7, 11, 100	28 15 12 22 9	22.0 16.6 14.5 19.1 13.8
5	0.1	4.61	0.0998	-0.5, 3, 4, 6, 9, 100	28 6 13 22 17	24.9 10.0 18.2 18.4 14.5
6	0.5	8.3	0.0158	-0.5, 3, 5, 6, 9, 100	28 15 4 22 17	21.8 21.2 9.7 20.3 13.0

DISCUSSION

Quantile-quantile Plots

Because of the visual nature of quantile-quantile plots, they are best used as a very general guide to the fit of various distributions to the sample data. Although goodness-of-fit results may indicate that a distribution clearly does not fit the data, it is often difficult to detect the subtle differences in plots based on distributions which appear similar in shape to the observed data. Because of this, the error associated with the quantile-quantile plot based methods is a tendency to retain too many parameter values rather than rejecting indiscriminately. This is especially true for the method involving the two discrete distributions since it is difficult to distinguish many of these plots from a straight line.

In spite of the ability of quantile-quantile plots to reject only the most unlikely candidates, they are helpful in providing a baseline from which to compare the different distributions. The Normal quantile-quantile plots are especially useful in this regard for comparing the relative skewness of the two continuous distributions

and the data to the Normal distribution, a well-known and much used distribution in statistics.

Since the method involving Normal quantile-quantile plots is more helpful than the second quantile-quantile plots based method, creating theoretical Normal quantile-quantile plots for the Negative Binomial and Poisson distributions as part of this research might have been beneficial. However, since it is more common to plot sample data directly against the proposed probability distribution, this idea is not conceived until after the analysis is completed, and so, is dropped.

Chi-square Goodness-of-Fit Intervals

As mentioned in the methods section, there is much discussion as to how best choose the intervals for the Chi-square goodness-of-fit test. After much deliberation, approximately equiprobable intervals are chosen based on the range of expected outcomes from the distribution being tested.

Some believe that intervals created in this manner are not as valid as intervals solely based on the data (Turnbull). Although this may be true according to rigorous mathematical proofs, there is nothing wrong statistically with the method used, and it is adequate for the purpose of this research. This method is also more convenient to use since the approximate probabilities corresponding to the intervals are chosen to guarantee that the requirements of the Chi-square goodness-of-fit test (i.e., large enough expected counts) are met.

Generalizability

The results of this analysis may not be generalizable to the general population of women in the United States, because the sample demographics are very different from this population. For example, 76 of the 102 women in the sample were recruited from student health clinics, 40% were younger than 20 years of age at the time of the study

and 93% had never been married. An attempt is made to compare the demographics of the sample data to those of the National Survey of Family Growth (NSFG) for 1995 (US Department of Health and Human Services). The NSFG survey involves approximately 10,000 women, and is a representive sample of the general population of women in the United States. Unfortunately, a comparison of the two samples is difficult since many demographic variables are defined differently in the two data sets. For example, the actual age of each participant is included in the NSFG data set, whereas age is dichotomized into older or younger than 20 years of age in the sample data. Despite these differences, it is still clear that the two samples appear to represent very different populations. Because of this, the distribution of coital frequency cannot be assumed to be the same for the two represented populations.

Originally, the NSFG data set was to serve as the primary sample in the analysis, since it is large and generalizable to the population of women in the United States. Unfortunately, its coital frequency variable is based on a recall of average coital frequency during the three months prior to the interview and is categorized into very broad intervals which are too imprecise to model. An attempt is made to compare the coital frequency distributions from the sample data and the NSFG data set by dividing the averaged sample coital frequency data into the categories used in the NSFG data set. This comparison is difficult since the NSFG categories do not include all possible values of *avgsex*. For example, category three includes women who have sexual intercourse once a week, which corresponds to four times per month, and category four includes women who have sexual intercourse two or three times a week, which is approximately eight to twelve times per month. This excludes values of *avgsex* which fall between four and eight times per month.

Another possible cause of bias is the nature of the study from which the data is taken, since, by definition, all participants are oral contraceptive users. Since choice of contraceptive method might be correlated with coital frequency, this study may

have produced more generalizable results if coital frequency data from women using a variety of contraceptive methods had been included. This is a particularly important point since a secondary purpose of this research is to establish a standard distribution to be used in models comparing various characteristics of different contraceptive methods.

Since goodness-of-fit tests are generally more accurate for larger samples, a larger sample size may produce results which have more validity. Another advantage of a larger sample is that the proportion of atypical observations is reduced. This effect is also achieved by increasing the number of cycles observed for each woman, as this generates values more representative of each woman's typical behaviour.

SUGGESTIONS FOR FURTHER RESEARCH

Ideally this analysis would be repeated using a larger sample which is representative of the target population. Not only would such an analysis prove or disprove the generalizability of these results to the population of women in the United States, it could also be used to compare the distribution of coital frequency of the general population to that of the more specific population of oral contraceptive users from which this sample is drawn. The analysis could also be repeated using data from women who use other contraceptive methods for further comparisons.

Because many contraceptive studies are conducted in developing nations where population control and the spread of sexually transmitted diseases are high priority issues, similar analyses based on data from various nations may be of interest, since culture and other factors may affect coital frequency. Differences in the underlying distributions of coital frequency from different areas may affect the generalizability of contraceptive studies between nations.

APPENDIX 1:
Summary of *totsex*

Frequency table of *totsex* by *cycle*

totsex	Frequency of totsex by cycle		
	Cycle 1	Cycle 2	Cycle 3
0	17	14	14
1	9	7	10
2	6	4	9
3	7	10	7
4	11	7	6
5	9	4	5
6	3	8	8
7	10	5	4
8	7	4	5
9	8	9	5
10	5	6	4
11	2	0	2
12	1	3	2
13	1	4	3
14	1	4	0
15	1	1	0
16	2	0	0
17	1	1	1
19	0	1	1
20	0	1	1
21	0	1	1
22	0	0	1
24	0	1	0
25	1	0	0
TOTAL	**102**	**95**	**89**

APPENDIX 2:
Graphical Representations of Poisson, Negative Binomial, Gamma, Log-normal and Normal Distributions (Rothschild)

NEGATIVE BINOMIAL DISTRIBUTIONS

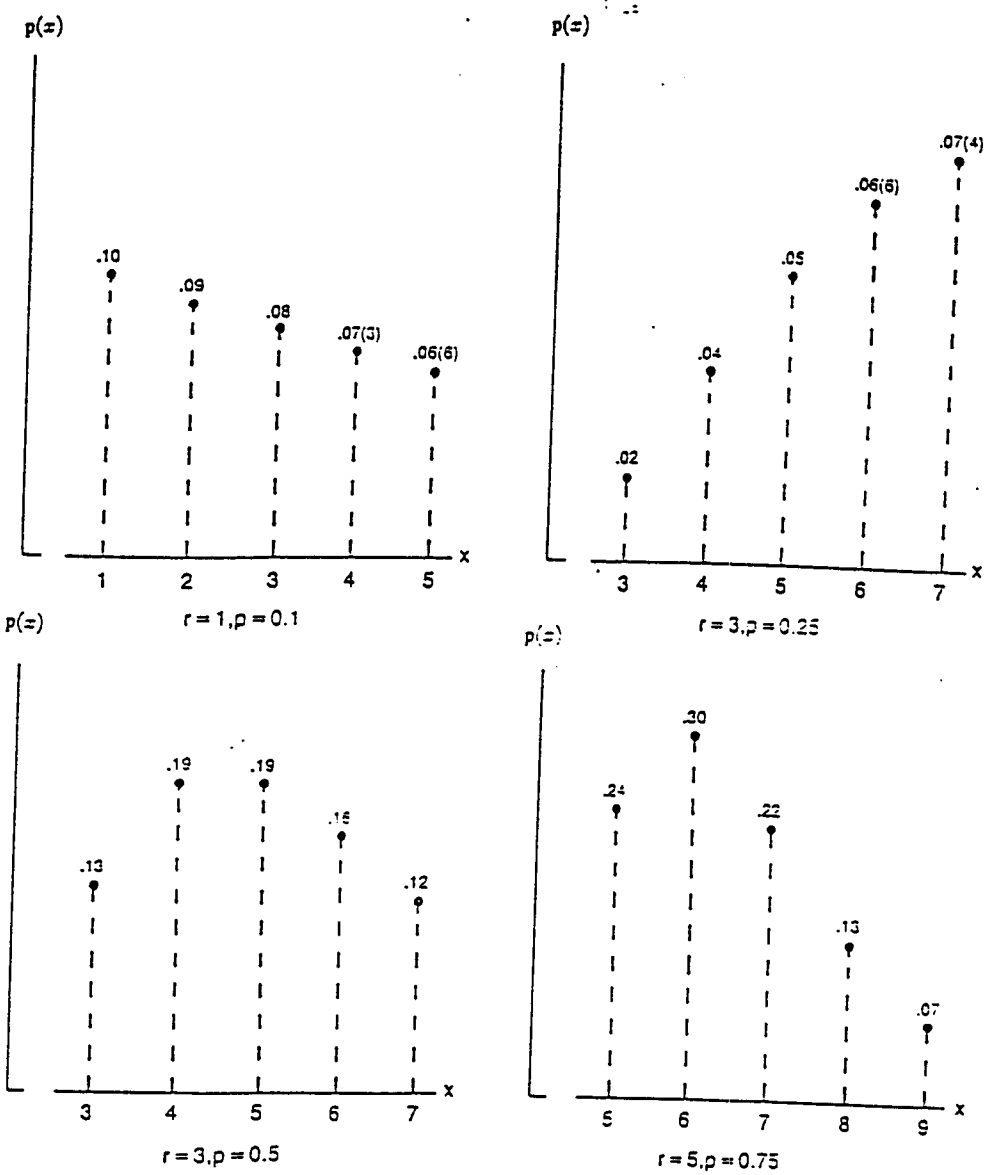

$$p(x) = \binom{x-1}{r-1} p^r q^{x-r}, \; x = r, r+1, r+2, \ldots, \quad 0 \leq p \leq 1$$

POISSON DISTRIBUTIONS

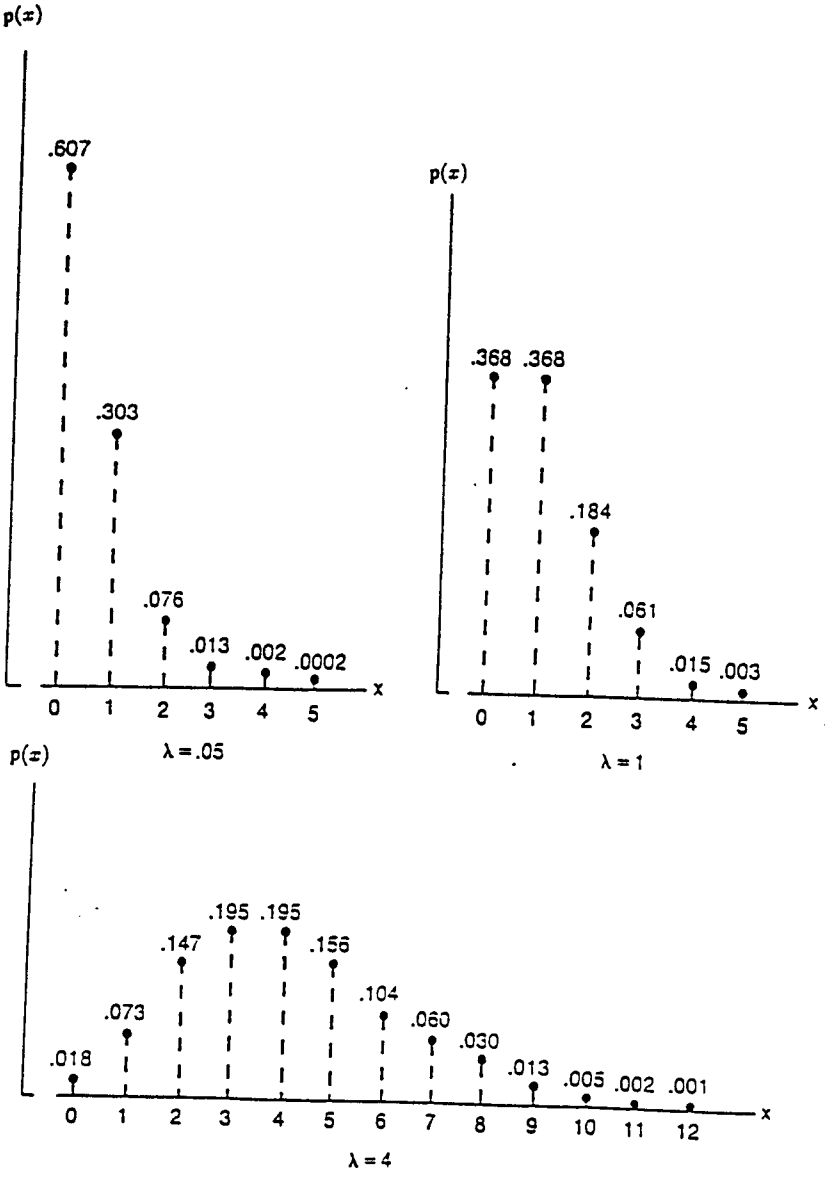

$$p(x) = \frac{\lambda^x}{x!}e^{-\lambda}, \qquad x = 0, 1, 2, \ldots, \quad \lambda > 0$$

GAMMA DISTRIBUTIONS

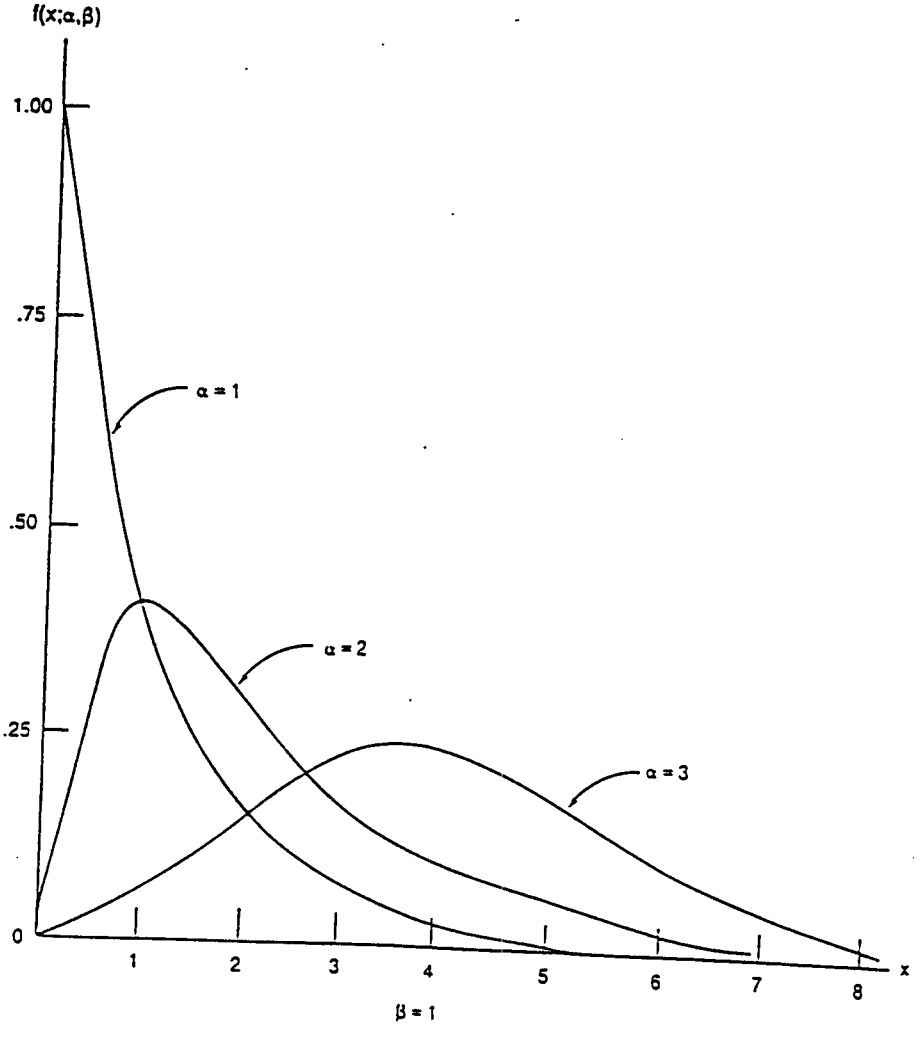

$$f(x) = \begin{cases} \frac{x^{\alpha-1}e^{-x/\beta}}{\beta^{\alpha}\Gamma(\alpha)}, & \alpha, \beta > 0;\ 0 < x < \infty \\ 0, & \textit{elsewhere} \end{cases}$$

LOGNORMAL DISTRIBUTIONS

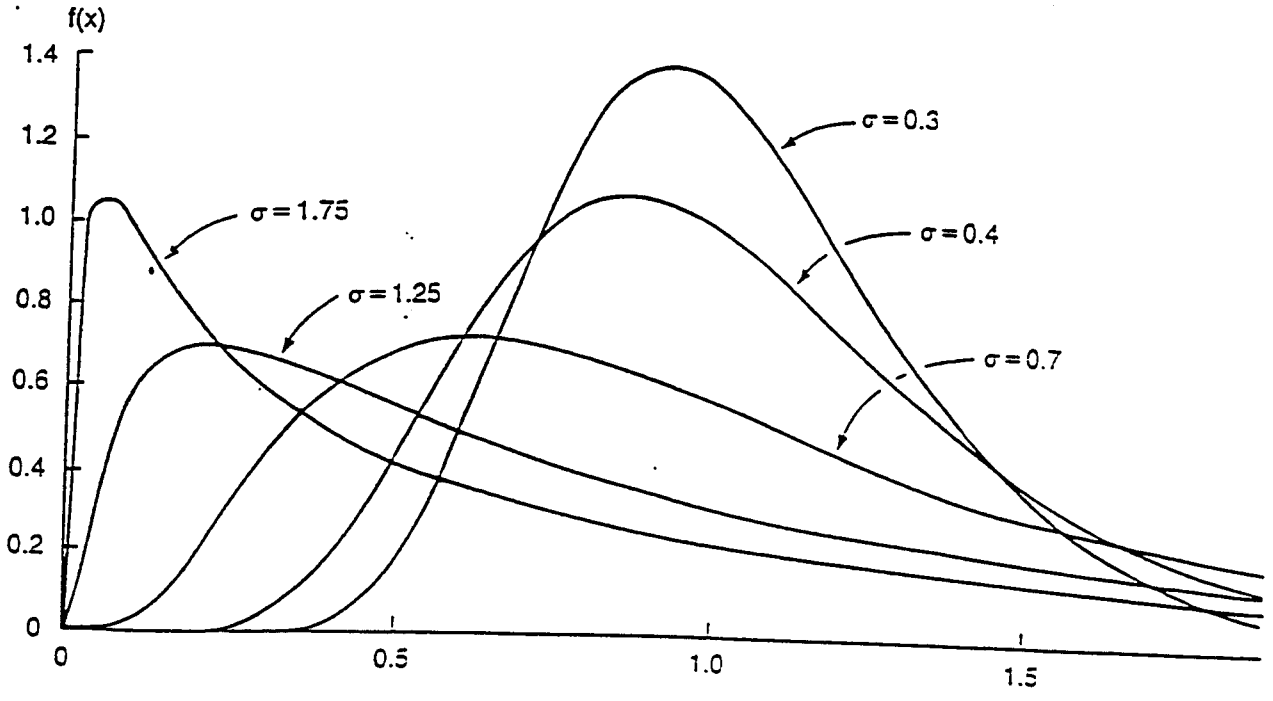

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-(\ln x - \mu)^2/2\sigma^2}, \quad 0 < x < \infty; \quad -\infty < \mu < \infty; \sigma > 0.$$

NORMAL (OR GAUSSIAN) DISTRIBUTIONS

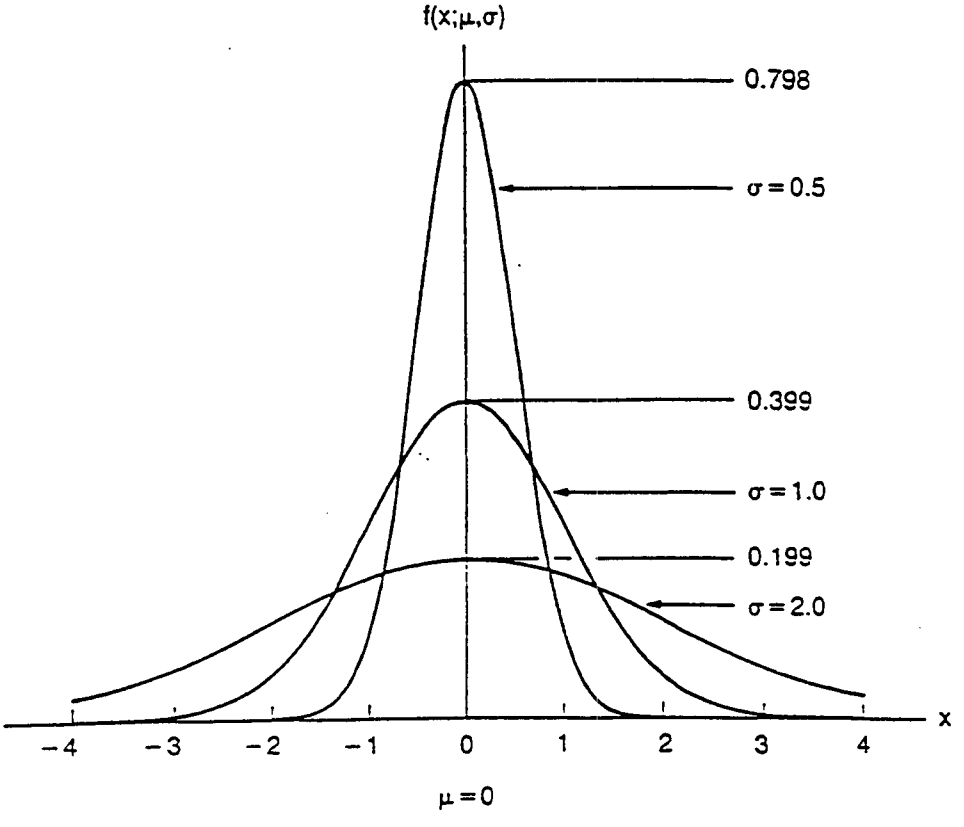

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}, \qquad -\infty < \mu < \infty;\ \sigma > 0.$$

APPENDIX 3:
S-PLUS Programs

This appendix contains examples of the S-PLUS programs used to create the quantile-quantile plots and to run the goodness of fit tests.

Quantile-quantile plots

Normal distribution

The code below creates Standard Normal quantile-quantile plots of the *totsex* and *avgsex* variables. These plots can be seen in Appendix 3.

```
mean <- 0
sd <- 1
qqnorm(totsex, mean, sd)
title("Totsex variable")
qqline(totsex)
qqnorm(avgsex, mean, sd)
title("Avgsex variable")
qqline(avgsex)
```

To create quantile-quantile plots in S-PLUS for distributions other than the Normal distribution, a simple S-PLUS function must be created for each distribution.

Negative Binomial

The code below creates a quantile-quantile plot of the Negative Binomial distribution with parameters *(r=1, p=0.3)* against the *avgsex* variable. This plot can be seen in Appendix 3.

```
size <- 1
prob <- 0.3
qqnbinom <- function(x, size, prob) {
            plot(qnbinom(ppoints(x), size, prob), sort(x))
            }
qqnbinom(avgsex, size, prob)
title("(r=1,p=0.3)")
```

Poisson

The code below creates a quantile-quantile plot of the Poisson distribution with parameter ($\lambda = 5$) against the *avgsex* variable.

```
lambda <- 5
qqpois <- function(x, lambda) {
            plot(qpois(ppoints(x), lambda), sort(x))
            }
qqpois(avgsex, lambda)
title("(lambda=5)")
```

Gamma

The code below creates a quantile-quantile plot of the Gamma distribution with parameters ($\alpha=5$, $\beta=1$) against the *avgsex* variable.

```
shape <- 5
qqgamma <- function(x, shape) {
            plot(qgamma(ppoints(x), shape), sort(x))
            }
qqgamma(avgsex, shape)
```

Log-normal

The code below creates a quantile-quantile plot of the Log-normal distribution with parameters ($\mu=5$, $\sigma=0.5$) against the *avgsex* variable.

```
mean <- 5
sd <- .5
qqlnorm <- function(x, mean, sd) {
            plot(qlnorm(ppoints(x), mean, sd), sort(x))
            }
qqlnorm(avgsex, mean, sd)
```

Chi-square goodness of fit test

Negative binomial

The code below is used to run the chi-square goodness of fit test for the Negative Binomial distribution ($r=3$, $p=0.6$) and the variable *totsex*.

The two lines below create five approximately equal intervals according to the parameters specified.

```
q <- c(0, .2, .4, .6, .8, 1)
breaks <- qnbinom(q, 3, 0.6)
```

The lines actually run the goodness of fit test and show the result, expected counts and observed counts for each interval.

```
z <- chisq.gof(totsex, cut.points = b2, dist = "negb", size=3, prob=.6)
z
z$count
z$expected
```

Poisson

The code below is used to run the chi-square goodness of fit test for the Poisson distribution ($\lambda = 6$) and the variable *totsex*.

```
q <- c(0, .2, .4, .6, .8, 1)
breaks <- qpois(q, 6)
z <- chisq.gof(totsex, cut.points = b2, dist = "pois",lambda=.6)
z
z$count
z$expected
```

Kolmogorov goodness of fit test

Gamma

The code below is used to run the Kolmogorov goodness of fit test for the Gamma distribution *(α=0.1, β=1)* and the variable *totsex*.

```
ks.gof(totsex,distribution="gamma", shape=.1)
```

Log-normal

The code below is used to run the Kolmogorov goodness of fit test for the Log-normal distribution *(μ=0, σ=0.1)* and the variable *totsex*.

```
ks.gof(totsex,distribution="lognormal", meanlog=0, sdlog=.1)
```

APPENDIX 4:
Goodness-of-fit results

This appendix contains the goodness-of-fit results and corresponding p-values for the four distributions modeled in this research. Results which indicate a good fit are highlighted. Since many parameters are tested for goodness-of-fit, some results for distributions which clearly do not fit the data have been excluded. P-values marked zero are less than 0.0001.

TOTSEX

Negative Binomial – Chi square goodness of fit

r	*p*	GOF	P-value	Interval breaks	Observed values	Expected values
1	0.1	16.9	0.0002	-0.5, 2, 4, 8, 15, 100	32 18 29 19 4	27.6 14.1 20.7 20.6 18.9
1	0.2	19.3	0.0001	-0.5, 0, 2, 4, 7, 100	17 15 18 22 30	20.4 29.4 18.8 16.3 17.1
1	0.3	89.5	0	-0.5, 0, 1, 2, 4, 100	17 9 6 18 52	30.6 21.4 15.0 17.8 17.1
2	0.1	152.9	0	-0.5, 7, 12, 18, 27, 100	72 23 6 1 0	23.0 19.4 19.7 19.7 20.3
2	0.2	15.8	0.0004	-0.5, 3, 5, 8, 12, 100	40 19 20 16 7	26.8 16.4 20.5 18.1 20.2
2	0.3	11.4	0.0033	-0.5, 1, 3, 5, 7, 100	26 14 19 13 30	22.0 26.1 20.3 13.6 20.0
2	0.4	134.7	0	-0.5, 1, 2, 3, 5, 100	26 6 8 62 0	35.9 17.6 14.1 18.2 16.2
3	0.1	328.0	0	-0.5, 13, 20, 28, 39, 100	96 5 1 0 0	21.5 20.1 20.7 19.7 19.9
3	0.2	103.8	0	-0.5, 5, 9, 12, 18, 100	59 28 8 6 1	20.7 24.3 16.4 22.4 18.2
3	0.3	12.8	0.0017	-0.5, 3, 5, 7, 11, 100	40 19 13 22 8	26.1 19.6 17.2 22.6 16.4
3	0.4	13.9	0.0010	-0.5, 2, 3, 5, 7, 100	32 8 19 13 30	18.3 11.3 23.2 19.0 30.2
4	0.1	364.8	0	-0.5, 20, 29, 38, 50, 100	101 1 0 0 0	21.9 21.3 19.7 18.9 20.2
4	0.2	210.2	0	-0.5, 8, 12, 17, 23, 100	79 16 6 0 1	21.0 20.0 23.2 19.2 18.6
4	0.3	55.7	0	-0.5, 5, 7, 10, 14, 100	59 13 20 5 5	27.6 16.3 21.8 19.4 16.8
5	0.1	360.1	0	-0.5, 27, 37, 47, 61	102 0 0 0	21.6 20.5 19.6 20.1
5	0.2	310.5	0	-0.5, 11, 16, 21, 28	94 6 1 1	20.6 21.6 20.7 20.5
5	0.3	111.7	0	-0.5, 6, 9, 12, 16, 100	62 25 8 5 2	21.5 21.0 20.0 19.4 20.2
5	0.4	11.0	0.0041	0.5, 4, 6, 8, 11, 100	33 12 17 15 8	26.1 20.5 18.3 19.0 17.0
5	0.5	14.9	0.0006	-0.5, 2, 4, 5, 7, 100	32 18 9 13 30	23.1 27.9 12.6 18.7 19.8
6	0.3	190.4	0	-0.5, 8, 11, 15, 19, 100	79 15 4 3 1	22.4 18.7 23.9 17.3 19.7
6	0.4	64.7	0	-0.5, 5, 7, 10, 13, 100	59 13 20 4 6	25.1 18.3 25.0 16.9 16.6
6	0.5	15.5	0.0004	-0.5, 3, 5, 6, 9, 100	40 19 3 25 15	25.9 25.1 11.5 24.1 15.4
6	0.6	46.0	0	-0.5, 2, 3, 4, 6, 100	32 8 10 12 40	32.2 17.1 15.4 21.3 16.1

Poisson – Chi square goodness of fit

lambda	GOF	P-value	Interval breaks	Observed values	Expected values
4	95.5469	0	-0.5, 2, 3, 4, 6, 100	32 8 10 12 40	24.3 19.9 19.9 26.6 11.3
5	40.068	0	-0.5, 3, 4, 5, 7, 100	40 10 9 13 30	27.0 17.9 17.9 25.6 13.60
6	35.1853	0	-0.5, 4, 5, 6, 8, 100	50 9 3 17 23	29.1 16.4 16.4 24.6 15.6
7	41.4252	0	-0.5, 5, 6, 8, 9, 100	59 3 17 8 15	30.7 15.2 28.5 10.3 17.3

Gamma - Kolmogorov goodness of fit

Alpha	GOF	P-value
0.5	0.6996	0
1	0.6365	0
1.5	0.5747	0
2	0.5163	0
2.5	0.4516	0
3	0.3852	0
3.5	0.3410	0
4	0.3104	0
4.5	0.2698	0
5	0.2611	0
5.5	0.2836	0
6	0.3082	0

Log-normal - Kolmogorov goodness of fit

Mu	Sigma	GOF	P-value
0	0.2	0.7448	0
0	0.5	0.6723	0
0	1	0.5503	0
1	0.2	0.5811	0
1	0.5	0.3984	0
1	1	0.2582	0
2	0.2	0.5530	0
2	0.5	0.3804	0
2	1	0.2831	0
3	0.2	0.9264	0
3	0.5	0.8204	0
3	1	0.6592	0
4	0.2	1.0000	0
4	0.5	0.9804	0
4	1	0.8706	0
5	0.2	1.0000	0
5	0.5	0.9998	0
5	1	0.9751	0

AVGSEX - Chi-square goodness of fit

Negative Binomial – Chi square goodness of fit

r	p	GOF	P-value	Interval breaks	Observed values	Expected values
1	0.1	24.8	0	-0.5, 2, 4, 8, 15, 100	21 13 33 16 3	23.3 11.9 17.5 17.4 15.9
1	0.2	33.9	0	-0.5, 0, 2, 4, 7, 100	8 13 13 21 31	17.2 24.8 15.9 13.8 14.4
1	0.3	122.1	0	-0.5, 0, 1, 2, 4, 100	8 4 9 13 52	25.8 18.1 12.6 15.0 14.5
2	0.1	110.0	0	-0.5, 7, 12, 18, 27, 100	55 25 5 1 0	19.4 16.4 16.6 16.6 17.1
2	0.2	11.5	0.0032	-0.5, 3, 5, 8, 12, 100	28 15 24 13 6	22.6 13.8 17.3 15.3 17.0
2	0.3	16.1	0.0003	-0.5, 1, 3, 5, 7, 100	12 16 15 12 31	18.6 22.0 17.1 11.5 16.9
2	0.4	78.5	0	-0.5, 1, 2, 3, 5, 100	12 9 7 15 43	30.3 14.9 11.9 15.3 13.6
3	0.1	276.9	0	-0.5, 13, 20, 28, 39, 100	81 4 1 0 0	18.1 17.0 17.5 16.6 16.7
3	0.2	63.0	0	-0.5, 5, 9, 12, 18, 100	43 26 11 5 1	17.5 20.5 13.8 18.9 15.4
3	0.3	4.36	0.1130	-0.5, 3, 5, 7, 11, 100	28 15 12 22 9	22.0 16.6 14.5 19.1 13.8
3	0.4	23.8	0	-0.5, 2, 3, 5, 7, 100	21 7 15 12 31	27.3 11.9 19.7 12.7 14.4
4	0.1	305.7	0	-0.5, 20, 29, 38, 50, 100	85 1 0 0 0	18.4 18.0 16.6 15.9 16.6
4	0.2	179.4	0	-0.5, 8, 12, 17, 23, 100	67 13 5 1 0	17.7 16.9 19.6 16.2 15.7
4	0.3	26.9	0	-0.5, 5, 7, 10, 14, 100	43 12 17 10 4	23.2 13.8 18.4 16.4 14.2
5	0.1	4.61	0.0998	-0.5, 3, 4, 6, 9, 100	28 6 13 22 17	24.9 10.0 18.2 18.4 14.5
5	0.2	48.7	0	-0.5, 2, 3, 4, 6, 100	21 7 6 13 39	29.6 13.4 11.8 16.5 14.8
5	0.3	258.5	0	-0.5, 11, 16, 21, 28, 100	77 7 1 1 0	17.4 18.3 17.4 17.3 15.7
5	0.4	71.9	0	-0.5, 6, 9, 12, 16, 100	47 22 11 4 2	18.1 17.7 16.8 16.4 17.1
5	0.5	12.0	0.0025	-0.5, 4, 6, 8, 11, 100	34 13 20 10 9	22.9 17.3 15.5 16.0 14.3
6	0.3	18.3	0.0001	-0.5, 2, 4, 5, 7, 100	21 13 9 12 31	19.5 23.5 10.6 15.7 16.7
6	0.4	31.7	0	-0.5, 5, 7, 10, 13, 100	43 12 17 9 5	21.2 15.4 21.1 14.3 14.0
6	0.5	8.3	0.0158	-0.5, 3, 5, 6, 9, 100	28 15 4 22 17	21.8 21.2 9.7 20.3 13.0
6	0.6	57.7	0	-0.5, 2, 3, 4, 6, 100	21 7 6 13 39	27.1 14.4 12.9 17.9 13.6

Poisson – Chi square goodness of fit

lambda	GOF	P-value	Interval breaks	Observed values	Expected values
4	107.9426	0	-0.5, 2, 3, 4, 6, 100	21 7 6 13 39	20.5 16.8 16.8 22.4 9.5
5	46.615	0	-0.5, 3, 4, 5, 7, 100	28 6 9 12 31	22.8 15.1 15.1 21.6 11.5
6	14.9602	0.0019	-0.5, 4, 5, 6, 8, 100	34 9 4 20 19	24.5 13.8 13.8 20.7 13.1
7	23.6786	0	-0.5, 5, 6, 8, 9, 100	43 4 20 2 17	25.9 12.8 24.0 8.7 14.6

Gamma - Kolmogorov goodness of fit

Alpha	GOF	P-value
0.5	0.7809	0
1	0.6747	0
1.5	0.6093	0
2	0.5480	0
2.5	0.4833	0
3	0.4141	0
3.5	0.3728	0
4	0.3296	0
4.5	0.2752	0
5	0.2224	0.0004
5.5	0.2141	0.0008
6	0.2417	0.0001

Log-normal - Kolmogorov goodness of fit

Mu	Sigma	GOF	P-value
0	0.2	0.8435	0
0	0.5	0.7195	0
0	1	0.5850	0
1	0.2	0.6128	0
1	0.5	0.4299	0
1	1	0.2899	0
2	0.2	0.4746	0
2	0.5	0.2932	0
2	1	0.2474	0.0001
3	0.2	0.9271	0
3	0.5	0.7811	0
3	1	0.6270	0
4	0.2	1.0000	0
4	0.5	0.9796	0
4	1	0.8707	0
5	0.2	1.0000	0
5	0.5	0.9999	0
5	1	0.9740	0

REFERENCES

Atkinson, A. C. "Goodness of Fit." Encyclopedia of Biostatistics(1998 ed.), 2, 1723-1731.

Fowlkes, E. B. (1987), A Folio of Distributions: A Collection of the Theoretical Quantile-quantile Plots, New York: Marcel Dekker, Inc.

Hastings, N. A. J. and Peacock, J. B. (1975), Statistical distributions: a handbook for students and practitioners, New York: Wiley.

Kault, D. (1996), "The Shape of the Distribution of the Number of Sexual Partners," Statistics in Medicine, 15, 221-230.

Koch, G. G., Personal Communication, Spring 1998.

Koehler, K. "Chi-Square Tests." Encyclopedia of Biostatistics(1998 ed.), 1, 608-622.

Lachenbruch, P. A. (1967), "Frequency and Timing of Intercourse: Its Relation to the Probability of Conception," In Population Studies, 21(1), 23-32.

Lachenbruch, P. A. and Glasser, J. H., (1968), "Observations on the Relationship and Timing of Intercourse and the Probability of Conception," In Population Studies, 22(3), 399-407.

MacDougall, J., Personal Communication, Spring 1998.

Martin, M. A. "Graphical Displays." Encyclopedia of Biostatistics(1998 ed.), 2, 1750-1771.

Oakley, D., Potter, L., Wong, E. and Visness, C., (1996), "Measuring Compliance Among Oral Contraceptive Users," Family Planning Perspectives, 24, 12-19.

Quade, D., (1996-7) BIOS 162: Introductory Applied Statistics, Coursepak.

Rothschild, V., (1986) Probability distributions, New York: Wiley.

SAS Institute Inc., (1990) SAS© Language Reference, Version 6, First Edition, Cary, NC: SAS Institute Inc.

SAS Institute Inc., (1990) SAS© Procedures Guide, Version 6, Third Edition, Cary, NC: SAS Institute Inc.

Statistical Sciences, Inc., (1995) S-PLUS User's Manual, Version 3.3 for Windows, Seattle: Statistical Sciences, Inc.

Steiner, M., Dominik, R., Trussell, J., Hertz-Picciotto, I. (1996), "Measuring Contraceptive Effectiveness: A Conceptual Framework," Obstetrics & Gynecology, 88(3) (Supplement).

Taff, G., Personal Communication, Spring 1998.

Turnbull, C.D., Personal Communication, Biostatistics Honours Research Course (BIOS 99), Spring 1998.

US Department of Health and Human Services. (1997), National Survey of Family Growth, Cycle 5: 1995. Hyattville, Maryland.

Wackerly, D. D., Mendenhall, W., Scheaffer, R. L. (1996), Mathematical Statistics with Applications, Belmont, California: Duxbury Press.

Wilcox, R. "Kolmogorov-Smirnov Test." Encyclopedia of Biostatistics(1998 ed.), 3, 2174-2176.

Zink, R.C., Personal Communication, Spring 1998.