Characterizing Probability Distributions for Farm Prices, Yields and Net Returns

by

Steven T. Buccola

Characterizing Probability Distributions for Farm Prices, Yields, and Net Returns

Steven T. Buccola

Introduction

Analysts frequently find it convenient to assume that the economic variables in which they are interested are realizations of a normal or Gaussian process. In a price forecasting model, a normally distributed dependent variable leads naturally (although not necessarily) to supposition of a normally distributed residual. The latter, in turn, permits use of t distributions for hypothesis testing and for characterization of forecast confidence intervals. A further advantage of normality is that it permits one to draw, on the basis of mean and variance statistics alone, meaningful statements about the probabilities of various events. Risk analysts often posit net return normality in order to base optimal choices on mean and variance. For example, both E-V and MOTAD frameworks assume normally distributed portfolio returns (Hazell).

The assumption that a particular variate may be Gaussian is, in practice, rarely subjected to test. Failure to test is sometimes legitimate: research time or data to investigate the normality assumption may be too costly. Or the advantages of analytic methods

---

The author is an associate professor in the Department of Agricultural and Resource Economics, Oregon State University, Corvallis, Oregon 97331.

requiring normality may be deemed so great that only the most extreme departure from normality would lead one to abandon such methods. In certain cases, the likelihood that testing would indicate extreme nonnormality may seem quite remote.

In other circumstances, however, the case for normality tests is compelling. Accuracy of t-based interval forecasts can crucially depend upon departures from normality in a model's residual. Feldstein and others, in addition, have indicated the significant errors in optimal choice that may result if normality-based decision procedures are employed in the presence of nonnormality. Conversely, the awkwardness of such distribution-free methods as stochastic dominance suggests that analysts pay a price for abandoning mean-variance approaches when they need not do so.

The purpose of the present paper is to inquire whether prices, yields, or returns on Oregon alfalfa and wheat farms can reasonably be represented as Gaussian. The analysis emphasizes data whitening procedures and the importance of employing a variety of normality tests. Irrigated alfalfa's price and net return distributions are shown to contrast sharply with those of dryland wheat. Implications for other cropping situations and for risky decision models are suggested.

Initial Considerations

The simple net return model

$$\pi = PY - C \qquad (1)$$

will be considered, where P is price in dollars per unit, Y is yield in units per acre, and C is cash cost in dollars per acre. Behavior of noncash costs raises issues beyond the scope of our study. At the time production is initiated, P, Y, and C generally are unknown and their outcomes random. The distribution of π depends upon the distribution of each of these variables as well as upon the nature and extent of their stochastic interdependence. For this reason alone, some attention to the individual densities of P, Y and C is warranted.

It is frequently argued on the basis of the central limit theorem that the latter variables are Gaussian, especially if each observation represents an average of data points from shorter time periods (Freund, p. 208). But the central limit theorem is not applicable if, in a given time interval, the major factors affecting the variable are not "numerous" or if they are not mutually independent. For example, the theorem does not apply to price averages if the shorter-period prices comprising the average are serially correlated.

There is substantial evidence of nonnormality in daily prices of many agricultural commodities. In separate studies, Houthakker, Mandelbrot, and Mann and Heifner found daily changes in the logs of a variety of cash and futures prices to be leptokurtic. That is, the distributions appeared to have "longer" or thicker tails than occur in the normal distribution. Apparently the densities usually were symmetric, implying positively skewed actual price changes.[1] Little attention has been devoted to distributions of weekly or monthly average prices.

Both Day and Bessler have provided evidence of nonnormality in field crop yields. Day found Mississippi Experiment Station cotton yields lognormally distributed, corn yields nonsignificantly skewed, and oat yields nonsignificantly or negatively skewed. The cotton yields generally were leptokurtic, while corn and oat data tended not to depart significantly from normal kurtosis. Bessler found subjective distributions of California field crop yields to be negatively skewed in some cases and positively skewed in others. Because data series corresponding to a given production technology are so short, comparable studies of farm cost distributions are not often found.

Only in special cases would knowledge of price, yield, and cost distributions be sufficient to derive analytically the distribution of net returns. Where P, Y, and C are lognormally distributed and mutually independent, revenue (PY) also is lognormal but π is not (Johnson and Kotz, p. 119). If P, Y, and C all are Gaussian, the distributions of PY and π are complex and generally nonGaussian, the extent of departure from normality depending upon statistical dependence among the three variables (Haldane). In most applied contexts, furthermore, prices, yields, and costs would not conform exactly to any theoretical probability family. Thus, it is helpful to consider historical distributions of PY and π directly.

Data Whitening

Data first were assembled to reflect historical net returns on an irrigated alfalfa farm and dryland wheat farm. State mid-month prices of baled alfalfa from 1949 to 1981 and monthly average soft white wheat prices from 1913 to 1980 were obtained from the Oregon Extension Service. Assuming the alfalfa grower sells his hay more or less continuously during the harvest season, annual alfalfa prices were estimated by averaging monthly prices from June to October.[2/] The wheat grower was assumed to sell all his wheat on the cash market in September.

Alfalfa yield data for 1949-1981 were supplied by the Medford branch of the Oregon Agricultural Experiment Station. Wheat yields for 1913-1980 were obtained from the Moro Station. As would be expected, varieties grown at the stations changed during the sample period. Alfalfa yields represent averages for Lahontan, Talent, and Vernal varieties while wheat yields are drawn successively from nine different varieties.

Per-acre cash costs were constructed by multiplying historical input prices with per-acre input requirements listed in Oregon Extension Service budgets. Items included are labor, fertilizer, pesticides, fuel, electricity, and seed. Input requirements were held fixed under the assumption they are known by the grower and that cost uncertainty is due primarily to uncertainty about input prices.[3/] Consistent series for fertilizer and pesticide inputs could be obtained only for 1965-1982. All prices and costs were deflated with the CPI to reflect constant 1980 purchasing power.

Whitening Price, Yield, and Cost Series

To permit valid normality tests a series must be random, that is, each observation drawn independently from an identical distribution. One-sample runs tests, making use of the number of runs above and below the estimated median, are useful for distinguishing serial dependence or secular shifts in the location of a distribution (Swed and Eisenhart). Runs tests were applied to all the original series just described. Results, shown at the left of table 1, suggest there typically were too few runs for the series to have been completely random. The nature of the patterns were various. Deflated prices of both commodities followed a wavelike motion with no overall trend, whereas yields had slight to moderate trends with no apparent waviness. Deflated costs of both commodities were roughly constant through 1973, after which there was a marked increase followed by a slight negative trend.

To accommodate these patterns, price first-differences were formed and linear trends removed from yield series. Post-1973 costs were segregated from earlier ones and linear trends removed from the post-1973 segments. Wheat price first-differences before 1931 were eliminated because they clearly were more volatile than in subsequent years. Similarly, cost trend residuals prior to 1974 were removed because they were less variable than residuals in the 1974-1982 period. Results of runs tests on remaining price first-differences and yield and cost trend residuals are given at the right of table 1. The tests failed to indicate, at the 95 percent confidence level, any regular

Table 1. Sample Sizes, Runs Statistics, and Means of Alfalfa and Wheat Series

	Original Series		Whitened Series		
	Sample Size	Number of Runs	Sample Size	Number of Runs[a/]	Mean (Standard Error)
Alfalfa					
Price	33	8*	32	20	-0.81 (2.58)
Yield	33	11*	33	12	0.00 (0.25)
Revenue	33	13	32	17	-4.00 (25.17)
Cost	18	2*	9	5	0.00 (2.68)
Wheat					
Price	68	14*	50	25	0.01 (0.16)
Yield	68	30	50	22	-1.13 (1.47)
Revenue	68	18*	50	28	-0.14 (6.71)
Cost	18	2*	9	4	0.00 (0.95)

a/ Runs listed are those above or below the median. An asterisk indicates presence of too few runs at the five percent error level in a two-tailed test. Critical values are given in Swed and Eisenhart.

patterns in the transformed series. As shown also at the right of table 1, expected values of transformed series were much less than their standard errors, suggesting the adjusted means were nonsignificantly different from zero.[4/]

Even after removal of secular changes in location and variance, the data may be nonstationary with respect to higher moments. Day showed that skewness and kurtosis of Mississippi field crop yields varied regularly with fertilizer level. Fertilizer data were not available for the yields used here, so this input could not be controlled. And, given the information at hand, reliable methods are not available for detecting serial changes in higher moments. Our estimates of such moments may therefore reflect their historical averages.

Whitening the Revenue Series

Mean-constant revenue series can be formulated by deducting expected revenue

$$(2) \qquad E(P_tY_t) = E(P_t)E(Y_t) + \text{cov}\,(P_t,Y_t)$$

from actual revenue P_tY_t. For simplicity in estimating (2), cov (P_t,Y_t) may be considered fixed in t. Lagged price P_{t-1} is a reasonable estimate of $E(P_t)$ since $E(P_t - P_{t-1}) \approx 0$ (table 1). A good estimate of $E(Y_t)$ is its linear trend value because detrending removed mean drift in yields. Designating the estimated yield trend by $(\hat{g} + \hat{h}t)$, (2) becomes, for both alfalfa and wheat,

$$(2)' \qquad \hat{E}(P_tY_t) = (P_{t-1})(\hat{g} + \hat{h}t) + \hat{\text{cov}}\,(P_t,Y_t).$$

To the extent (2)' is a valid model of expected revenue, the revenue deviation series

$$(3) \qquad P_tY_t - \hat{E}(P_tY_t)$$

has stationary mean of zero.

The latter hypothesis was tested by calculating (3) for both commodities and subjecting the series to runs tests. Results in table 1 show that while there was significant nonrandomness in series P_tY_t, the series transformed by (3) apparently are serially independent and location stationary. Estimates of means of both transformed series are negligibly small compared to their standard errors.

Price, Yield, Revenue, and Cost Distributions

Numerous methods are available for testing the hypothesis that a sample has been drawn from a normal parent. Birnbaum points out that the power of a test to detect departures from normality depends on the nature of the departure with which one is concerned. Kolmogorov-type tests, for example, are sensitive to the maximum distance between sample and hypothesized cumulative distribution and hence are most appropriate for detecting departures from normality in the center of the distribution. Other tests are more powerful at detecting departures in the tails. For the present purpose a mix of tests was employed.

Tests Used

To detect significant skewness, estimates of Pearson's skewness parameter $g_1 = m_3/s^3$ were used, where m_3 is the sample estimate of the third central moment and s the sample standard deviation. The distribution of g_1 is symmetrical about zero if the parent population is normal, approaching $N(0,\sqrt{6/n})$ with large sample size. Tests based on this distribution have been shown sensitive to asymmetry in the parent (Snedecor and Cochran, pp. 78-80; Gastwirth and Owens, p. 138).

Asymptotically the kurtosis parameter estimate $g_2 = (m_4/s^4) - 3$, where m_4 is the estimated fourth central moment, is at least as powerful as any other statistic in detecting nonnormal kurtosis (Geary, pp. 228-239). For small and medium samples, however, Geary and Gastwirth and Owens show that a (the ratio of mean absolute deviation to standard deviation) has relatively more power in this task. Statistic a therefore is employed for alfalfa and wheat tests while the population equivalent of g_2 is referred to in a subsequent discussion. The distribution of a is slightly asymmetric about its mean, which varies with sample size (Snedecor and Cochran, p. 80).

Finally, Monte Carlo experiments indicate Shapiro and Wilk's W is sensitive to both asymmetry and tail-length departures from normality, even in small samples (Shapiro and Wilk). Statistic W is defined as c^2/S^2, where c is the regression slope of ordered sample observations on the expected values of the normal order statistics and S^2 is the sum of squared errors about the mean. Shapiro and Wilk tabulated the distribution of W up to moderate sample sizes. They found (p. 608), for a five percent critical region at least, that W is more powerful than Kolmogorov-type or Cramer-Von-Mises-type tests in detecting a wide range of nonnormal shapes. W-tests accordingly are used in conjunction with the shape parameter tests just mentioned.

Estimates of g_1, a, and W for the transformed price, yield, revenue, and cost series are given in tables 2 and 3. Ten- and two-percent critical regions for two-tailed tests (the only significance points published for g_1 and a) are listed for comparison. Values of a falling below the critical intervals indicate significant leptokurtosis, that is, tails thicker than in the normal distribution. W-values below the critical points indicate significant nonnormality in either skewness or kurtosis. Choice between ten- and two-percent significance depends upon relative costs of a type I and type II error in this situation. The two-percent level gives greater assurance that we will not falsely reject normality, but greater risk that we will fail to reject normality when we should. In the author's judgment the ten-percent level here provides the more judicious balance of risks and the following discussion is based on ten-percent significance.

Alfalfa Test Results

Alfalfa price and yield in table 2 are nonsignificantly skewed but significantly leptokurtic. Using the rule that a variable is nonnormal if either g_1 or a indicate such, one would reject the normal hypothesis for each series. Yet in both instances W is well within the 90-percent bound, suggesting nonsignificant departures from the normal. The a estimates are not exceptionally low and thick-tailedness apparently is too weak to greatly affect the overall normality of the sample. Hence,

Table 2. Normality Tests for Alfalfa Variables

		Coefficient of Variation[a/]	Pearson's g_1	Geary's a	Shapiro-Wilk's W
Price		0.200	0.404	0.718*	0.954
Yield		0.199	-0.157	0.723*	0.956
Revenue		0.275	1.028*	0.713*	0.924*
Ln Revenue			-0.011	0.727*	0.960
Cost		0.056	0.515	0.894	0.911
Critical Regions[b/]	Ten Percent		0.645 -0.645	0.862 0.741	0.941
	Two Percent		0.958 -0.958	0.881 0.712	0.915

a/ Standard deviation of whitened series divided by mean of original series. The alfalfa price-yield correlation was -0.173.

b/ Region within which null hypothesis of normality would be rejected at indicated type I error level in a two-tail test. Regions listed here apply to price, yield, and revenue. Those for cost (n = 9) are given in Shapiro and Wilk or extrapolated from Geary.

Asterisks indicate significantly different from a normal distribution at the 10 percent error level. Pearson and Geary tests are two-tailed; Shapiro-Wilk's is one-tailed.

evidence seems best to support the null hypotheses of price and yield normality.

By contrast, the whitened revenue series is strongly positively skewed as well as leptokurtic and its W-statistic falls within the 10-percent critical region. Alfalfa price and yield were only slightly negatively correlated ($r = -0.173$), implying that the coincidence of a high price and high yield was nearly as likely as that of a high price and low yield. This resulted in a relatively long tail to the right of the revenue mode and hence a positively skewed revenue distribution. Revenue's leptokurtosis partly may be explained by thick-tailedness in both prices and yields.

Logs of revenue also were subjected to normality tests to see if revenue could adequately be characterized as lognormal. The logs indeed are symmetric and W well within the range expected of a normal distribution (table 2). Although a low <u>a</u> estimate suggests weak leptokurtosis in the logs, it would not be far wrong to say that alfalfa revenues have been lognormally distributed.

Wheat Test Results

Dryland wheat variables in table 3 present a quite different picture from those of irrigated alfalfa. September wheat price is strongly positively skewed and leptokurtic, the W-test indicating pronounced nonnormality. Interestingly, the high wheat price occurring in 1973 alone accounts for this phenomenon. When the 1973 observation was temporarily removed, all statistics shifted away from their 10-percent

Table 3. Normality Tests for Wheat Variables

		Coefficient of Variation[a/]	Pearson's g_1	Geary's a	Shapiro-Wilk's W
Price		0.231	1.979*	0.663*	0.853*
Yield		0.359	-0.042	0.803	0.980
Revenue		0.291	0.258	0.811	0.990
Cost		0.085	0.030	0.813	0.964
Critical Regions[b/]	Ten Percent		0.533 -0.533	0.849 0.751	0.955
	Two Percent		0.787 -0.787	0.866 0.728	0.938

a/, b/ See footnotes in table 2. The wheat price-yield correlation was -0.287.

critical zones. Of course, there is no basis for removing the 1973 data point.

Despite skewed wheat prices, wheat revenues in table 3 are normally distributed according to all three test statistics. There are two reasons why the product of positively skewed price and unskewed yield should result in an essentially unskewed revenue series.[5/] First, price and yield are more negatively correlated for wheat ($r = -0.287$) than they were for alfalfa. Such negative correlation reduces the frequency of especially large revenues and thus inhibits positive skewness. Second, as measured by its coefficient of variation, wheat yield is more variable than price and one expects the shape of PY to be more influenced by the shape of that factor having the greater relative variability. This is corroborated by Haldane, p. 234.

Selected combinations of alfalfa and wheat revenue also were formed to gauge the effect of farm diversification on the distribution of portfolio returns. Revenues from combinations involving more than 50 percent alfalfa were significantly nonnormal in all tests, whereas those with less than 50 percent alfalfa provided normally distributed revenues. The implication, from the present data at least, is that returns of two combined enterprises are not more Gaussian than the individual returns making up the portfolio. That is, there is no sign of the central limit theorem at work when only two enterprises are involved.

Finally, tables 2 and 3 provide no significant indication of nonnormality in cash production costs. We would not, with only 9 observations to employ, expect otherwise. As Geary remarks, only "extreme cases" of population nonnormality would ever result in rejecting the null hypothesis with such small samples. Cost coefficients of variation are, however, very low relative to those of revenue. It is shown below that relative variability of cost and revenue is an important determinant of the distribution of net returns.

Net Return Distributions

In the absence of cost series of adequate length, one cannot conduct direct statistical tests of the shapes of net return distributions. But one can make logical inferences about such shapes under certain reasonable circumstances. Anderson and Doran show that if revenue and cost are assumed independent, the second through fourth central moments of profit may be expressed as

$$\begin{aligned} \mu_{2\pi} &= \mu_{2r} + \mu_{2c} \\ \mu_{3\pi} &= \mu_{3r} - \mu_{3c} \\ \mu_{4\pi} &= \mu_{4r} + \mu_{4c} + 6\mu_{2r}\mu_{2c} \end{aligned} \tag{4}$$

where numbered subscripts indicate the moment order and π, r, c refer to profit, revenue, and cost, respectively. Substituting (4) into expressions for net return skewness ($\gamma_{1\pi}$) and kurtosis ($\gamma_{2\pi}$) gives

$$\gamma_{1\pi} = \mu_{3\pi}/\mu_{2\pi}^{1.5} = \frac{\mu_{3r} - \mu_{3c}}{(\mu_{2r} + \mu_{2c})^{1.5}}$$

(5)
$$\gamma_{2\pi} = (\mu_{4\pi}/\mu_{2\pi}^2) - 3 = \frac{\mu_{4r} + \mu_{4c} + 6\mu_{2r}\,\mu_{2c}}{(\mu_{2r} + \mu_{2c})^2} - 3.$$

These parameters may usefully be compared with gross revenue skewness (γ_{1r}) and kurtosis (γ_{2r}):

(6)
$$\gamma_{1r} = \mu_{3r}/\mu_{2r}^{1.5}$$
$$\gamma_{2r} = \mu_{4r}/\mu_{2r}^2 - 3.$$

The third and fourth central moments of cost, μ_{3c} and μ_{4c} in (5), are correlated with powers of the cost variance μ_{2c}. Hence as the latter variance falls relative to revenue variance μ_{2r}, net return skewness $\gamma_{1\pi}$ approaches revenue skewness γ_{1r} and net return kurtosis $\gamma_{2\pi}$ approaches revenue kurtosis γ_{2r}. An implication is that the shape of the net return distribution increasingly is dominated by that of the revenue distribution as cost variance falls relative to revenue variance.

To give some idea of the similarity of revenue and net return distributions in the present study, costs were assumed normally distributed ($\mu_{3c} = 0$, $\mu_{4c} = 3\mu_{2c}^2$) and independent of revenues. These are, at least for a length of run as short as a year, reasonable suppositions because farm costs react to a number of input price variables having little relation to yields or output prices. For the alfalfa data, cost variance was about 64, or just over 0.3 percent of the revenue variance figure of 18,714. The result was that $\mu_{3r}/(\mu_{2r} + \mu_{2c})^{1.5}$ was approximately equal to $\mu_{3r}/\mu_{2r}^{1.5}$, and $(\mu_{4r} + 3\mu_{2c}^{2} + 6\mu_{2r}\mu_{2c})/(\mu_{2r} + \mu_{2c})^{2}$ approximately equal to μ_{4r}/μ_{2r}^{2}, with exact equalities holding at least to the second place past the decimal. The ratio of wheat revenue variability to cost variability was similarly very high (compare their coefficients of variation in table 3). Hence, our discussion of the alfalfa and wheat revenue distributions applies about equally well to the corresponding distributions of net returns.

Conclusions

The irrigated alfalfa and dryland wheat data investigated here provide a contrast in probability distribution shapes. Alfalfa price and yield each are nonsignificantly skewed, whereas alfalfa revenue and net returns are significantly skewed to the right. Wheat price is strongly positively skewed and wheat yield symmetric, yet the corresponding revenue distribution is itself nonsignificantly skew. One expects a product of two roughly symmetric and independent variables to have a relatively long right tail, and in this sense the alfalfa results offer no surprise. A more interesting phenomenon is the failure of wheat price skewness to induce a comparable skew in wheat revenues. The failure seems largely a consequence of the symmetrical distribution of wheat yields combined with yield's superior contribution to revenue variability.

Despite their virtual symmetry, alfalfa price and yield may have significantly thicker tails than do Gaussian variates, contributing to leptokurtic alfalfa revenues. Long- or thick-tailedness in wheat prices does not, however, become manifested in the wheat revenue distribution, likely because price's leptokurtosis is mostly a reflection of its skew. Overall, alfalfa revenues and net returns studied can nearly be characterized as lognormal (with perhaps a somewhat thicker tail than the lognormal) and the corresponding wheat returns as normal. Distribution shapes of combined alfalfa-wheat net returns tend toward or away from the normal depending upon which product has the greater combination share.

The foregoing emphasizes that a number of factors can result in nonnormal returns and that indeed one might expect significant nonnormality to be more the rule than the exception. A similar conclusion may be drawn from Haldane, who has developed expressions for the third and fourth central moments of a product of two random normal deviates. Evaluations of such expressions suggest that even if price and yield are bivariate normal, only special parameter combinations of the joint distribution ever would result in approximately Gaussian revenue.

The effect of violating the normality restriction implicit in a forecasting or risk analytic procedure depends upon the procedure's robustness regarding this restriction as well as upon the degree of nonnormality present. If prices are nonnormal, one expects price forecast errors also to be nonnormal unless the exogenous information explains the nonnormality in some way. Nonnormal forecast errors, in turn, imply the standard deviation of forecast does not entirely represent the uncertainty surrounding the forecast. For example, standard deviation provides no information about asymmetry of actual price observations around projected prices.

Yield distributions different than those discussed in this paper might have been obtained from different climatic regions. Drought conditions in the Plains, for example, sometimes are severe enough to discourage wheat harvest. A sufficient number of ruined crops would skew the yield distribution to the right, producing a thick left tail. Thus, yield distributions need not be as symmetric as those shown in tables 2 and 3.

Footnotes

1. Houthakker, Mandelbrot, and Mann and Heifner investigated series of the form $\ln P_t - \ln P_{t-1}$. If the latter difference is Gaussian and (P_t, P_{t-1}) are independent, then P_t is lognormal (Johnson and Kotz, p. 119). The fact that the log difference tends to be long-tailed suggests P_t is positively skewed but with a kurtosis greater than predicted by the lognormal distribution.

2. A simple average of monthly prices misstates season average price if both the quantity and price of each cutting are different. Provided the intra-seasonal price trend and proportions sold after each cutting are relatively constant across years, the resulting bias will affect mean price only.

3. This assumption is not strictly true, as growers often cannot anticipate the amount of pesticide, for example, they will require. But input use uncertainty probably does not have great cost impact in most situations and it would be difficult to model accurately.

4. An exact confidence level cannot be established for this conclusion. Because underlying distributions may not be normal, ratios of means to standard errors of means are not necessarily t-distributed.

5. Revenue deviation series (3) is not identical to what one would obtain by simply multiplying whitened price and yield series. However, equation (3) arises from the product of price and yield in the sense that these variables enter (3) multiplicatively.

References

Anderson, J.R., and H.E. Doran. "Some Properties of Simple Functions of Random Variables." Rev. Mkt. Agr. Econ. 46 (1978): 37-47.

Bessler, D.A. "Aggregated Personalistic Beliefs on Yields of Selected Crops Estimated Using ARIMA Processes." Amer. J. Agr. Econ. 62 (1980): 666-674.

Birnbaum, Z.W. "Distribution-Free Tests of Fit for Continuous Distribution Functions." Ann. Math. Stat. 24 (1953): 1-8.

Day, R.H. "Probability Distributions of Field Crop Yields." J. Farm Econ. 47 (1965): 713-741.

Feldstein, M.S. "Mean-Variance Analysis in the Theory of Liquidity Preference and Portfolio Selection." Rev. Econ. Stud. 36 (1969): 5-12.

Freund, J.E. Mathematical Statistics, 2nd Edition. Englewood Cliffs: Prentice-Hall, Inc., 1971.

Gastwirth, J.L. and M.E.B. Owens. "On Classical Tests of Normality." Biometrika 64 (1977): 135-139.

Geary, R.C. "Testing for Normality." Biometrika 34 (1947): 209-242.

Haldane, J.B.S. "Moments of the Distributions of Powers and Products of Normal Variates." Biometrika 32 (1942): 226-242.

Hazell, P.B.R. "A Linear Alternative to Quadratic and Semivariance Programming for Farm Planning Under Uncertainty." Amer. J. Agr. Econ. 53 (1971): 53-62.

Houthakker, H.S. "Systematic and Random Elements in Short-Term Price Movements." Amer. Econ. Rev. 51 (1961): 164-172.

Johnson, N.L., and S. Kotz. Continuous Univariate Distributions -1. Boston: Houghton-Mifflin Co., 1970.

Mandelbrot, B. "The Variation of Certain Speculative Prices." J. of Business 36 (1963): 394-419.

Mann, J.S., and R.G. Heifner. The Distribution of Shortrun Commodity Price Movements. Econ. Res. Serv. Tech. Bull. No. 1536, U.S. Dept. of Agriculture, 1976.

Shapiro, S.S., and M.B. Wilk. "An Analysis of Variance Test for Normality (Complete Samples)." Biometrika 52 (1965): 591-611.

Snedecor, G.W., and W.G. Cochran. Statistical Methods, 7th Edition. Ames: Iowa State Univ. Press, 1980.

Swed, F.S., and C. Eisenhart. "Tables for Testing Randomness of Grouping in a Sequence of Alternatives." Ann. Math. Stat. 14 (1943): 66-87.

Thompson, K.J., and P.B.R. Hazell. "Reliability of Using the Mean Absolute Deviation to Derive Efficient E,V Farm Plans." Amer. J. Agr. Econ. 54 (1972): 503-06.