Choice among methods of estimating share yield

The search for the growth component in the discounted cash flow model.

David A. Gordon, Myron J. Gordon, and Lawrence I. Gould

The yield at which a share of stock is selling, also called its expected return or required return, is an important statistic in finance. Firms use it in choosing among investment opportunities and financing alternatives, and investors use it in making portfolio decisions. Nevertheless, the yield at which a share is selling is a difficult quantity to measure, which has limited its use in the practice of finance. This paper develops and tests a basis for choice among alternative methods of estimating a share's yield.

A share's yield, like a bond's yield, is the discount rate that equates its expected future payments with its current price. A bond's yield is easy to measure under the common practice of ignoring default risk, as the future payments are then known with certainty. The future payments on a share, however, are dividends and market price, and these payments are uncertain.

The common practice is to represent these future dividend payments with estimates of two numbers: One is the coming dividend, and the other is a growth rate. The latter can be an estimate of the long-run growth rate in the dividend or of the growth rate in price over the coming period. In the latter case, the estimate is called the expected holding-period return (EHPR); in the former case, it is called the discounted cash flow yield (DCFY).[1] In either case, the estimate of a share's yield reduces to the sum of its dividend yield and a future growth rate, with the latter inferred in some way from historical data.

There is a wide variety of acceptable methods for using historical data to estimate future growth. This variation in method is illustrated in the testimony of expert witnesses before public utility commissions on the fair return for a public utility. In these cases, the estimates and the methods used are a matter of public record. Some idea of the various methods can be found in Morin (1984) and Kolbe, Read, and Hall (1984). The performance of alternative estimating methods has been examined in Gordon (1974), Kolbe, Read, and Hall (1984), Brigham, Shome, and Vinson (1985), and Harris (1986).

We have derived our basis for comparing the accuracy of alternative methods for estimating the DCFY on a share from the generally accepted propositions that yield should vary according to risk, and that beta is the best estimate of risk. Hence, the DCFY should vary among shares with beta, and, between two methods for estimating growth, the superior method is the one for which the variation in yield among shares is explained better by the variation in beta among the shares.

First we present simple, plausible, and objective measurement rules for implementing four popular and/or attractive methods for estimating the DCFY. We then describe how sample statistics may be used to judge the accuracy of each method. We also describe how the CAPM model has been used to estimate share yield and explain why we do not compare it with the various DCFY methods. The following section carries out the comparison with samples of utility and industrial shares, and the last section pre-

DAVID A. GORDON is in charge of transaction finance at Scotia McLeod, a subsidiary of the Bank of Nova Scotia in Toronto. MYRON J. GORDON is Professor of Finance at the Faculty of Management at the University of Toronto (Ontario M5S 1V4). LAWRENCE I. GOULD is Professor and Head of Accounting and Finance at the University of Manitoba in Winnipeg (Manitoba R3T 2N2).

sents the conclusions that may be drawn from the findings.

ALTERNATIVE MEASUREMENT RULES FOR A SHARE'S YIELD

Under the DCF method or model for estimating the expected return on a stock, the yield for the jth stock is:

$$DCFY_{jt} = DYD_{jt} + GR_{jt}, \quad (1)$$

where:

$DCFY_{jt}$ = DCF yield on the jth stock at time t,

DYD_{jt} = dividend yield on the jth stock at time t, and

GR_{jt} = long-run growth rate in the dividend on the jth stock that investors expect at time t.

In what follows, we omit the time and firm subscripts on the variables when they are not required. Also, DCFY will refer to the unknown true yield on a share.

The difficult problem in arriving at the DCFY is estimation of the long-run growth rate that investors expect. Four estimates of that quantity are:

EGR = rate of growth in earnings per share over a prior time period, usually the last five years;

DGR = rate of growth in dividend per share over a prior time period, usually the last five years;

FRG = consensus among security analyst forecasts of the growth rate in earnings, over the next five years; and

BRG = an average over the prior five years of the product of the retention rate b and rate of return on common equity r on a stock.

The estimate of share yield that incorporates each of these estimates of growth is denoted KEGR, KDGR, KFRG, and KBRG, respectively.

A case can be made for each of the four methods for estimating growth. KEGR, KDGR, and KBRG have been widely used in public utility testimony and in research on stock valuation models. The rationale for KEGR is the belief that the past growth rate in earnings is the best predictor of future growth in earnings and dividends. The rationale for KDGR is that the future growth rate in dividends is the statistic we want to estimate, and the past dividend record is free of the noise in past earnings.[2] The rationale for KBRG is that all variables will grow at this rate if the firm earns r and retains b. Furthermore, as Gordon and Gould (1980) show, KEGR and KDGR will be biased in one direction or another if r and b have changed over the last five years. As for KFRG, security analysts are professionals employed to forecast future performance; their forecasts are widely accepted by investors. The IBES collection of forecast growth rates of security analysts compiled by Lynch, Jones, and Ryan has increased the popularity of this estimate.

As stated earlier, we may also take the yield on a share as the sum of the dividend yield and the expected rate of growth in price over the coming period. This estimate of a share's yield is widely used in testing the CAPM, with the average HPR over the prior five years commonly used in such empirical work. On the other hand, this estimate of a share's yield varies so widely among firms and over time as to be patently in error as an estimate of share yield.[3]

BASIS OF COMPARISON

To compare the accuracy of the four estimates of the DCFY stated above, we regress the data under each estimate on beta for a sample of shares. If KEGR is the estimate,

$$KEGR_j = \alpha_0 + \alpha_1 BETA_j + \epsilon_j. \quad (2)$$

The rationale for this expression lies in the risk premium theory of share yield, where the share yield is equal to the interest rate plus a risk premium that varies with the share's relative risk. Hence, if BETA is an error-free index of relative risk, α_0 is equal to the interest rate, and α_1 is the risk premium on the market portfolio or standard share.[4]

The higher the correlation between KEGR and BETA, assuming that α_1 is positive, the greater the confidence we may have in KEGR as an estimate of DCFY. We cannot rely solely on the correlation, though, in selecting among the methods for estimating DCFY. Errors in KEGR as a basis for estimating the DCFY on the jth share have random and systematic components. The former is ϵ_j, and its average value can be taken as the root mean square error of the regression (MSE). The larger the root MSE of the regression, the less attractive KEGR is as an estimate of share yield, because the error makes the problem of choice between $KEGR_j$ and $KEGR_j - \epsilon_j$ more acute. (That problem will be discussed shortly.)

The systematic error is the difference between the unknown true yield on the jth share, $DCFY_j$, and the value predicted by Equation (2). There is no obvious measure of the systematic error, as we do not know $DCFY_j$, but sample values of α_0 may provide information on its average value. The difference between α_0 and the interest rate is an indicator of systematic error, because the difference is zero under the risk premium theory. Error in the measurement of BETA biases α_0 upward, but, with the same BETA for each share used in all four regressions, differences in α_0 are indicators of systematic error.[5]

In addition to regression statistics, the sample mean and standard deviation of KEGR is a source of information on its accuracy as a method for the estimation of DCFY. If the mean departs radically from the long-term bond rate, or if the standard deviation indicates an unreasonable range of variation among shares, the accuracy of the method is open to question. Also, the sample mean may be a source of information on the systematic error for a method of estimation. Hence, sample values for the mean, standard deviation, correlation, root MSE, and constant term all contribute to a judgment on a method's accuracy for estimating the DCFY on a share. Unfortunately, there is no simple criterion for choice among the alternatives.

Once a conclusion is reached on the most accurate method for estimating DCFY — say, KEGR — we then have the problem of choice between $KEGR_j$ and $KEGR_j - \epsilon_j$ for the jth share. If the random error in $KEGR_j$ is due to error in its measurement for the jth share, we simply use the value predicted by Equation (2), which is $KEGR_j - \epsilon_j$. On the other hand, KEGR and DCFY may vary among shares with other (omitted) variables as well as BETA, in which case ϵ_j is also due to the omitted variables, and $KEGR_j$ may be the better estimate of DCFY. Unfortunately, we have no basis for choice among these two hypotheses, and the smaller the root MSE the less troublesome the problem of choice between them.

A more favorable tax treatment of capital gains over dividends should make investors prefer capital gains to dividends. As Brennan (1973) has shown, the yield investors require on a share would then vary with the excess of its dividend yield over the interest rate. To recognize this, Equation (2) becomes

$$KEGR_j = \alpha_0 + \alpha_1 BETA_j + \alpha_2 DMI_j + \epsilon_j, \qquad (3)$$

with DMI_j the excess of the dividend yield over the interest rate for the jth firm. Although the tax effect should make α_2 positive, its information in DMI on share risk would tend to make α_2 negative. That is, dividend yield varies inversely with expected growth, and we would find α_2 negative insofar as growth is risky. To the extent that these two influences of the dividend yield offset each other, α_2 will tend toward zero.

The CAPM theory of how expected return varies among shares has been proposed as an alternative to the DCF model for measuring yield. Its value for the jth stock is

$$EHPR_j = INTR + BETA_j[EHPR_m - INTR], \qquad (4)$$

where:

$EHPR_j$ = expected holding-period return on the jth share,

INTR = one-period risk-free interest rate,

$EHPR_m$ = expected holding-period return on the market portfolio.

There is an important difference between this CAPM model of share yield and the DCF model represented by Equation (1). The latter is merely an instrument for measuring share yield: There is nothing in the DCF model that explains the variation in yield among shares. The CAPM, on the other hand, is a theory on why and how yield varies among shares, but one must go outside of the theory to estimate the variables on the right-hand side of Equation (4). Given rules for estimating the variables, EHPR and BETA, empirical work then provides a joint test of the theory and the estimating rules, such as we are carrying out here.[6]

The CAPM nonetheless has been used to estimate share yield in testimony before regulatory commissions by assigning numbers to each of the quantities on the right-hand side of Equation (4). For INTR, a long-term bond yield is sometimes used instead of a one-period rate. BETA is estimated by conventional methods.

The big problem is the expected return on the market portfolio. Here the practice has been to use the average realized risk premium over a period of about fifty years as the estimate of $EHPR_m - INTR$ in Equation (4). Although the implicit assumption is that the risk premium is a constant over time, we would expect the premium to change from one period to the next for various reasons, among them changes in the interest rate, the risk premium on the market portfolio, and the relative taxation of interest and share income. Hence, this estimate of share yield is more or less in error at any particular time, but we have no way of estimating this error and comparing the method with the others.

COMPARATIVE PERFORMANCE

We carried out our empirical work with a sample of 75 large electric and gas utility firms and a sample of 244 firms that includes 169 industrial firms drawn from the S&P 400. We obtained share yield under the four methods for estimating it as of the start of the year for the years 1984, 1985, and 1986.

For the explanatory variables, BETA for each share on each date was obtained by regressing the monthly HPRs for the share on the monthly HPRs for the S&P 500 over the prior five years. DMI for a share is its dividend yield less the interest rate on the one-month Treasury bill at the start of each year. EGR and DGR are the growth rates in earnings and in dividends per share, respectively, over the prior five years as reported on the Value Line Tape. BRG is a weighted

average of the retention growth rates over the prior five years,[7] and FRG is the average of forecast growth rates in earnings over the next five years reported by IBES. The corresponding estimates of share yield were obtained by adding the dividend yield at the start of each year to the estimate of growth.

Table 1 presents the statistics that we obtained with KBRG and KFRG as the estimates of DCFY for the sample of utility shares and of all shares. The means of KBRG for the utility shares seems reasonable, with the interest rate on ten-year government bonds the standard of comparison, the latter being 11.67%, 10.43%, and 9.19% at the start of 1984, 1985, and 1986, respectively.[8] The standard deviations for KBRG are small enough to make its range of variation well within the bounds of reason. The lower means for all shares reveal that the means for industrial shares are below the means for utility shares.[9] This casts doubt on the accuracy of KBRG as a basis for estimating the DCFY on industrial shares, because industrials are riskier than utility shares.

The beta model explains none of the variation in KBRG among utility shares, but the two-factor model is a substantial improvement. The DMI coefficient, α_2, is positive and significant in every year, meaning that the unfavorable tax effect of a high dividend yield dominates the favorable risk effect. The coefficient on BETA is positive and significant in two of the three years. The only disturbing feature of the data is the sharp fall in R^2 and the corresponding rise in the root MSE relative to the standard deviation of KBRG as we go from 1984 to 1986.

The KBRG statistics for all shares are substantially inferior to the utility share statistics. This forces the unhappy conclusion that, for industrial shares, BETA is a poor measure of risk, or KBRG is a poor measure of DCFY, or both.

The KFRG statistics for the utility sample are superior to the KBRG statistics. The means are reasonable under the two criteria of being above the interest rate and moving with it. The range of variation of KFRG suggested by its standard deviations seems reasonable. The statistics for the beta model are a slight improvement on the corresponding statistics for KBRG. Furthermore, the two-factor model does a good job of explaining the variation in KFRG among

TABLE 1

Sample and Regression Statistics for KBRG and KFRG, Utility Shares and All Shares, 1984, 1985, and 1986

	KBRG			KFRG		
	1984	1985	1986	1984	1985	1986
			UTILITY SHARES (75)			
Mean	14.84	14.38	12.93	15.64	14.56	12.93
Standard Deviation	2.51	1.87	1.80	2.26	1.43	1.42
Beta Model α_0	14.26	13.96	13.05	15.14	13.48	12.74
α_1	1.44	1.21	−0.28	1.25	3.09	0.42
t-statistic	(0.97)	(1.12)	(0.19)	(0.93)	(4.14)	(0.37)
Root MSE	2.52	1.87	1.81	2.26	1.29	1.43
R^2	0.013	0.017	0.001	0.012	0.190	0.002
Two-Factor Model α_0	12.45	12.75	12.42	13.30	12.46	11.97
α_1	3.45	2.11	0.11	3.28	3.85	0.89
t-statistic	(3.13)	(2.19)	(0.08)	(3.83)	(6.33)	(0.88)
α_2	0.68	0.45	0.34	0.68	0.38	0.41
t-statistic	(8.22)	(4.88)	(2.81)	(10.73)	(6.52)	(4.65)
Root MSE	1.82	1.63	1.73	1.41	1.03	1.26
R^2	0.491	0.262	0.100	0.620	0.491	0.232
			ALL SHARES (244)			
Mean	12.98	13.19	11.86	16.17	15.87	14.31
Standard Deviation	3.86	3.21	3.52	2.60	2.32	2.30
Beta Model α_0	15.00	14.71	13.90	15.56	14.50	12.57
α_1	−2.47	−1.91	−2.40	0.74	1.72	2.05
t-statistic	(4.23)	(4.15)	(4.25)	(1.83)	(5.29)	(5.70)
Root MSE	3.73	3.10	3.40	2.59	2.20	2.16
R^2	0.069	0.066	0.069	0.014	0.104	0.118
Two-Factor Model α_0	14.34	14.42	13.95	15.40	14.61	12.75
α_1	0.09	−1.18	−2.51	1.37	1.44	1.61
t-statistic	(0.13)	(2.04)	(3.45)	(2.69)	(3.52)	(3.49)
α_2	0.48	0.17	−0.02	0.12	−0.06	−0.10
t-statistic	(6.04)	(2.09)	(0.24)	(2.01)	(1.12)	(1.53)
Root MSE	3.49	3.08	3.41	2.57	2.20	2.15
R^2	0.191	0.083	0.070	0.030	0.108	0.127

utility shares. The R^2s are higher here than for KBRG in every year. Finally, α_2 is positive and significant in every year, and α_1 is not significant only in 1986.

The implicit means of KFRG for the industrial shares seem high but not beyond reason. On the other hand, the regression statistics for the all-shares sample are not good, which leads to the same unhappy conclusion for industrial shares as we reached for KBRG.

Table 2 presents the statistics that we obtained using KEGR and KDGR as estimates of the DCFY on the shares in our samples. Comparison of the regression statistics with those in Table 1 reveals that KEGR and KDGR, particularly the former, fall short by a wide margin of the performance of KBRG and KFRG as estimates of the DCFY on a share.

CONCLUSION

We have compared the accuracy of four methods for estimating the growth component of the discounted cash flow yield on a share: past growth rate in earnings (KEGR), past growth rate in dividends (KDGR), past retention growth rate (KBRG), and forecasts of growth by security analysts (KFRG). Criteria for the comparison were the reasonableness of sample means and standard deviations and the success of beta and dividend yield in explaining the variation in DCF yield among shares. For our sample of utility shares, KFRG performed well, with KBRG, KDGR, and KEGR following in that order, and with KEGR a distant fourth. If we had used past growth in price, it would have been an even more distant fifth. Nevertheless, none of the four estimates of growth performed well under the criteria for a sample that included industrial shares.

Before closing, we have three observations to make. First, the superior performance by KFRG should come as no surprise. All four estimates of growth rely upon past data, but in the case of KFRG a larger body of past data is used, filtered through a group of security analysts who adjust for abnormalities that are not considered relevant for future growth. We assume this is done by any analyst who develops retention growth estimates of yield for a firm. If we had done this for all seventy-five firms in our utility sample, it is likely that the correlations

TABLE 2

Sample and Regression Statistics for KEGR and KDGR, Utility Shares and All Shares, 1984, 1985, and 1986

	KEGR			KDGR		
	1984	1985	1986	1984	1985	1986
			UTILITY SHARES (75)			
Mean	16.16	0.32	14.91	16.49	15.76	14.13
Standard Deviation	3.31	3.47	4.66	3.12	2.41	2.21
Beta Model α_0	15.45	16.18	9.51	15.75	14.53	12.30
α_1	1.75	0.40	−7.87	1.83	3.53	3.99
t-statistic	(0.89)	(0.20)	(2.16)	(0.99)	(2.64)	(2.32)
Root MSE	3.32	3.49	4.55	3.12	2.32	2.15
R^2	0.010	0.001	0.050	0.013	0.087	0.069
Two-Factor Model α_0	14.20	15.83	18.76	14.10	13.56	12.64
α_1	3.13	0.66	−8.03	3.65	4.25	3.78
t-statistic	(1.66)	(0.32)	(2.18)	(2.23)	(3.26)	(2.20)
α_2	0.47	0.13	−0.13	0.61	0.35	−0.18
t-statistic	(3.32)	(0.66)	(0.42)	(5.02)	(2.86)	(1.21)
Root MSE	3.11	3.50	4.58	2.70	2.21	2.14
R^2	0.142	0.007	0.063	0.269	0.180	0.087
			ALL SHARES (244)			
Mean	11.14	9.42	7.88	15.08	13.63	11.35
Standard Deviation	10.67	11.67	11.45	6.08	6.30	6.71
Beta Model α_0	15.96	18.28	19.55	15.15	0.04	15.39
α_1	−5.90	−11.16	−13.70	−0.09	−1.78	−4.74
t-statistic	(3.62)	(7.07)	(8.10)	(0.09)	(1.92)	(4.41)
Root MSE	10.41	10.65	10.18	6.09	6.27	6.47
R^2	0.051	0.171	0.213	0.000	0.015	0.074
Two-Factor Model α_0	14.84	18.01	19.91	14.31	14.11	14.79
α_1	−1.56	−10.49	−14.62	3.17	0.63	−3.25
t-statistic	(0.77)	(5.27)	(6.72)	(2.73)	(0.55)	(2.36)
α_2	0.81	0.15	−0.21	0.61	0.55	0.34
t-statistic	(3.51)	(0.55)	(0.67)	(4.57)	(3.47)	(1.72)
Root MSE	10.18	10.67	10.19	5.86	6.13	6.45
R^2	0.097	0.172	0.215	0.060	0.062	0.035

would have been as good or better than those obtained with the analyst forecasts of growth.

Second, we examined shares and not portfolios, because our objective is to estimate the DCFY for shares and not for portfolios. As common practice in testing the CAPM has been to execute tests on portfolios instead of shares, we classified our population of shares into ten portfolios on the basis of their beta values. Regression statistics were substantially unchanged, except that correlations increased dramatically.

Finally, we must acknowledge that we have no basis for estimating the expected HPR or DCF yield for industrial shares with any confidence. Theories on financial decision-making in industrial corporations that rely on that statistic have a weak empirical foundation.

[1] The EHPR is a one-period return, while the DCFY is a yield to maturity measure. The two may differ in actuality because of measurement problems, but they also may differ in theory. That is, they may differ in the same way that interest rates on bonds of different maturities may differ. See Gordon and Gould (1984a). This source of difference between EHPR and DCFY will be ignored here.

[2] A widely accepted hypothesis is that dividends contain information on earnings, because management sets the dividend to pay out a stable fraction of normal or permanent earnings.

[3] Over a five-year period, there may even be a negative rate of growth in price for a large number of firms. Furthermore, this negative growth rate may be larger in absolute value than the dividend yield, which leads to the conclusion that investors are holding such shares to earn a negative return. The frequency of negative rates of growth in price is reduced as the prior time period used in its calculation increases in length. As that takes place, however, the estimate of the expected return for a firm approaches a constant or a constant plus the dividend yield. The expected return on a share is one statistic for which it is an error to assume that expectations are on average realized.

[4] Equation (2) is similar to the CAPM according to Sharpe, Lintner, and Mossin. They arrived at this expression under very rigorous assumptions. The heuristic risk premium model is adequate for our purposes.

[5] It may be thought that Theil's (1966) decomposition of the difference between the actual and predicted values of a variable can be used here, but in fact that decomposition applies to a different problem. It assumes that the observed (actual) past values of a variable are free of error, and it decomposes the error in a model that is employed to explain the past values. The purpose of Theil's decomposition is to cast light on the possible error in using the model to predict future values of the dependent variable. Our problem is to determine which set of observed values is closest to the true values, with the risk premium theory of share yield and BETA as the source of information on the true values. Theil's method would be appropriate for decomposing the difference between the actual and predicted values of the *realized* holding-period return on a share. The actual values here can be observed without error.

[6] There is an enormous volume of empirical work devoted to discovering whether the theory is true, but this empirical work does not provide useful estimates of the EHPR on a share. To test the truth of Equation (4), the practice has been to regress EHPR on BETA for a sample of firms with the average realized HPR over the prior five or so years used as an estimate of the EHPR. Because of the large error in the realized HPR over a prior time period, as noted earlier, neither the actual values of the dependent variable nor the values predicted by the model are usable as estimates of share yield. See Fama and MacBeth (1973) and Friend, Westerfield, and Granito (1978).

[7] BRG for a year is earnings less dividend divided by the end-of-year book value. The estimate of the expected value as of the start of 1986 is 0.3BRG85 + 0.25BRG84 + 0.20BRG83 + 0.15BRG83 + 0.10BRG82. If any value of BRG was negative, it was set equal to zero.

[8] We expect the yields on shares to be above the risk-free interest rate, but with a high enough interest rate the more favorable tax treatment of shares can reduce the yield below the interest rate. Interest rates were not that high in these years. See Gordon and Gould (1984b).

[9] The statistics reported for all shares and for utility shares were also obtained for industrial shares. All methods of estimation performed so poorly for industrial shares, however, as to suggest no confidence can be placed in any of them. To save space, we do not present statistics for the industrial shares. Whatever we want to know about them can be deduced by comparing the data for all shares and utility shares.

REFERENCES

Brennan, M.J. "Taxes, Market Valuation and Corporate Financial Policy." *National Tax Journal*, 23 (1973), pp. 417-427.

Brigham, E., D. Shome, and S. Vinson. "The Risk Premium Approach to Measuring a Utility's Cost of Equity." *Financial Management*, Spring 1985, pp. 33-45.

Fama, E., and J.D. MacBeth. "Risk, Return and Equilibrium: Empirical Tests." *Journal of Political Economy*, 81 (May 1973), pp. 607-636.

Friend, I., R. Westerfield, and M. Granito. "New Evidence on the Capital Asset Pricing Model." *Journal of Finance*, 33 (June 1978), pp. 903-917.

Gordon, M.J. *The Cost of Capital to a Public Utility*. East Lansing, Michigan: Michigan State University, 1974.

Gordon, M.J., and L.I. Gould. "Comparison of the DCF and HPR Measures of the Yield on Common Shares." *Financial Management*, Winter 1984a, pp. 40-47.

——. "The Nominal Yield and Risk Premium on the TSE-300, 1956-1982." *Canadian Journal of Administrative Sciences*, 1 (1984b), pp. 50-60.

——. Testimony Before the Federal Communications Commission in the Matter of American Telephone and Telegraph Company. FCC Docket No. 79-63, April 1980.

Harris, R.S. "Using Analysts' Growth Forecasts to Estimate Shareholder Required Rates of Return." *Financial Management*, Spring 1986, pp. 58-67.

Kolbe, A.L., J.A. Read, and G.R. Hall. *The Cost of Capital: Estimating the Rate of Return for Public Utilities*. Cambridge, MA: MIT Press, 1984.

Morin, R.A. *Utilities' Cost of Capital*. Arlington, VA: Public Utilities Reports, Inc., 1984.

Theil, H. *Applied Economic Forecasting*. Chicago: North Holland, 1966.