An Analysis of Brokers' and Analysts' Unpublished Forecasts of UK Stock Returns
Author(s): Elroy Dimson and Paul Marsh
Source: *The Journal of Finance*, Vol. 39, No. 5 (Dec., 1984), pp. 1257-1292
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/2327727
Accessed: 27/11/2009 07:50

THE JOURNAL OF FINANCE • VOL. XXXIX, NO. 5 • DECEMBER 1984

An Analysis of Brokers' and Analysts' Unpublished Forecasts of UK Stock Returns

ELROY DIMSON and PAUL MARSH*

ABSTRACT

This paper describes an empirical study of over 4000 specific share return forecasts made by 35 UK stockbrokers and by the internal analysts of a large UK investment institution. A comparison of forecast and realised returns reveals a small but potentially useful degree of forecasting ability. A large part of the information content of the forecasts, however, appears to be discounted in the market place within the first month. Nevertheless, an analysis of some 3000 transactions motivated by, and executed at the time of, the forecasts shows that the apparent predictive ability of the recommendations could be translated into superior performance by the fund's investment managers. Differences in forecasting ability between brokers do not appear to persist over time, but predictive accuracy can be improved by pooling simultaneous forecasts from different sources.

STOCKBROKERS' ANALYSTS ARE responsible for a significant proportion of the original research carried out in the equity market. Their findings are communicated to, and utilised by, virtually all active fund managers. Indeed many funds rely predominantly or entirely on such external research. It is therefore important to evaluate the quality of the information and forecasts provided, and how it can best be combined with internal research.

Brokers' analysts, of course, provide their clients with a wide variety of research and background information. Ultimately, however, the analysts' objective is to forecast, or contribute to a forecast, of stock returns. Fund managers need to know how good these forecasts are in order to judge whether they are obtaining value for money from their brokers. More importantly, as Treynor and Black [78] have demonstrated, they need to know the level of forecasting ability implicit in both internal and external research in order to decide how much weight to place on individual forecasts, and in order to form rational judgements on the appropriate levels of diversification and turnover for their funds. From an overall capital markets perspective, it is also important to know the extent of brokers and other analysts' forecasting ability, since this provides information on whether

* Both authors from London Business School. We would like to thank the executives of the institution which cooperated with us in this study, both for data collection and for helpful responses to earlier drafts. We benefitted from the insights of Dick Brealey, Michael Brennan, Ian Cooper, Evi Kaplanis, the Journal's referee, and participants at workshops held at the London Business School and at the European Institute for Advanced Studies in Management, Brussels. We also received useful comments from participants at the 1984 Conference on New Challenges for Management Research, Leuven, Belgium and at the 1984 Meetings of the Eastern Finance Association, Orlando, Florida, European Finance Association, Manchester, England, and International Symposium on Forecasting, London, England.

the market is efficient in the strong form sense, a question on which there has been relatively little convincing evidence, especially in the UK.

This paper describes an empirical study of more than 4000 stock return forecasts covering just over 200 of the largest UK shares. Some 80 percent of these forecasts originated from 35 UK brokers who provided them over a two-year period at the specific request of the research department of a large UK investment institution. The remaining forecasts came from the institution's own team of internal analysts. Forecasting ability is evaluated by comparing forecast returns with realised returns for individual brokers, for brokers as a group, for the internal analysts, and for the sample as a whole. In addition to evaluating "paper" performance, we also examine the actual returns achieved by the fund by analysing the performance of some 3000 transactions motivated by, and executed at the time of, the forecasts. Tests are carried out to see whether differences in brokers' forecasting ability persist over time and can be usefully exploited. We also examine the degree of dependence between forecasts for the same share which are made simultaneously by different analysts. Finally, we investigate the potential gains from using forecasts from multiple sources and the ways in which forecasts can most effectively be combined.

In Section I, we review the existing empirical evidence on brokers' and analysts' forecasting ability, including both US and UK studies. In Section II, we describe the sample, data, and methodology used in the current study. Section III contains our empirical results, while Section IV outlines our summary and conclusions.

I. Previous Studies

Strong form tests of market efficiency are concerned with whether given investors or groups have monopolistic access to (nonpublic) information relevant to price formation (see Fama [29]). Unfortunately, there is no clear cut distinction between public and private information, nor is it always clear when information becomes public, nor which investors are in the public (see Sharpe [69]). Brokers' analysts, for example, utilise both public and private information. The latter is obtained from company visits, special relationships with companies and their personnel, or may be generated as "private insights" derived from published information (see Foster [39]). When their forecasts are communicated to clients, there will be uncertainty about when (or even if) they enter the public domain. These overlaps and ambiguities make it difficult to draw a clear distinction between semistrong and strong form tests of efficiency. Nevertheless, studies of analysts' forecasting ability have always been regarded as mainstream strong form tests. Indeed, in practice, virtually all strong form tests proceed by first identifying a group with potential access to private information, and then testing for evidence of persistent abnormal returns, with little attempt to identify the precise source of the special information.

In fact, the evidence on strong form market efficiency is dominated by studies of the returns achieved by professionally managed funds, e.g., Jensen [49], Friend et al. [40], Bogle and Twardowski [9], and, in the UK, Firth [36], Ward and Saunders [80], and Cranshaw [19]. The evidence is perplexing in that professional portfolio managers are unable consistently to meet, let alone exceed, the performance of naive benchmarks. In a security market which is known to provide superior returns to certain groups of investors, including corporate insiders (Scholes [66], Finnerty [34], Pratt and DeVere [61], and Oppenheimer and Dielman [60]) and stock exchange specialists (Niederhoffer and Osborne [59], Dann et al. [20]), it is surprising that the unceasing efforts of security analysts do not appear to enhance portfolio returns. One explanation may be that the market is informationally efficient, and that analysts are therefore not in a position to produce profitable recommendations. An alternative view is that the market is not wholly efficient and that valuable information is in fact contained in analysts' research; this information does not, however, lead to consistently superior performance because of a failure by portfolio managers to translate analysts' recommendations into timely transactions of an appropriate size. Potentially favourable gross performance may thus be translated into unfavourable net performance after trading costs have been deducted (see Mains [55]). Before concluding that security analysis has no economic value, we should therefore examine the profitability of following analysts' recommendations.

Numerous studies of the performance of stock recommendations, in the US, Britain, Canada, Australia, and Hong Kong, have in fact been published. Typically, some form of residual analysis is used (see Fama et al. [30]) to identify any abnormal returns behaviour from the publication date onwards for all shares for which there have been buy, hold, or sell recommendations. The results of these studies, which are summarised in Table I, generally indicate a small degree of forecasting ability, at least before dealing costs are taken into account. The profit opportunities disclosed by these studies are, however, limited. Share purchases at the pre-publication price are on average profitable; but purchases made a day or a week after the recommendation offer no opportunity for systematic abnormal gains.

Since the value of investment advice is very rapidly eroded by publication, the economic benefits from research are likely to accrue only to those who have access to private, unpublished investment advice. Little evidence is provided on this matter by the studies summarised in Table I. We can, however, obtain some insights into the potential value of unpublished brokers' advice from Fitzgerald's [37] research. While this study detected only mildly favourable average performance in a sample of 635 recommendations from 35 British stockbrokers (see Table I), two groups of firms achieved abnormal returns. One group, a cluster of sophisticated, research-oriented firms, obtained cumulative abnormal returns during the recommendations month which rose steadily to a total of over 6 percent abnormal gain. The other group, a cluster of brokers with close ties to market makers, a strong marketing emphasis, a willingness to indulge in pre-recommendation dealings, and a belief that their dealings were relatively unethical, were rather different: the cumulative residual for their 60 recommendations first rose and then fell sharply to give a total abnormal loss of over 6 percent over the month. While Fitzgerald interpreted his results as a contravention of market efficiency, it is worth noting that some of his recommendations related to periods before the brokerage firms volunteered to be investigated, and this could have given rise to selection bias. In addition, Fitzgerald made no adjustments for thin trading bias with daily data, so that the risk and return estimates

Table I

Performance of Brokers' and Advisory Services' Recommendations

Publication	Date	Country	Number of Services	Number of Recommendations	Type of Recommendation	Percentage Return on Pre-Recommendation Price: Publication Day	After One Day	After One Week	Longer Term
Cowles [17]	1933	USA	45	7500	Published				−1.4
Cowles [18]	1944	USA	11	6904	Published				0.2
Ferber [32]	1958	USA	4	345	Published	0.5		1.1	0.7
Ruff [65]	1963	USA	1	31	Published			4.0	
Colker [14]	1963	USA	1	1054	Published				3.8
Stoffels [75]	1966	USA	3	264	Published	0.7	1.0	1.5	1.4
Money Which [57]	1968	Britain	6	90	Published				11.9
Cheney [12]	1969	USA	4	na	Published				2.0
Brealey [10]	1971	Britain	1	360	Published				4.1
Diefenbach [23]	1972	USA	24	1209	Published				2.7
Firth [35]	1972	Britain	4	1100	Published			−1.0	−3.0
Black [8]	1973	USA	1	500	Published				10.0
Logue and Tuttle [54]	1973	USA	6	304	Published				3.0
Fitzgerald [37]	1975	Britain	35	635	Unpublished	0.5	1.0	1.5	0.5
Bidwell [6]	1977	USA	11	99	Published				2.6
Shepperd [71]	1977	USA	7	1008	Published				−3.0
Davies and Canes [21]	1978	USA	8	785	Published	1.3	1.6	1.7	
Fitzgerald [38]	1978	Britain	11	467	Published	3.1	3.2	3.5	1.0
Groth et al. [42]	1979	USA	1	6014	Published			1.8	2.6
Stanley et al. [74]	1981	USA	1	4461	Published			0.7	
Finn [33]	1983	Australia	1	361	Published				2.0
Bjerring et al. [7]	1983	Canada & USA	1	92	Published	1.5		2.1	1.2
Dawson [22]	1982	Hong Kong	1	300	Published		0.4	1.2	2.9
Garmaise [41]	1982	Canada	10	6968	Unpublished				−2.0
Copeland and Mayers [16]	1982	USA	1	3840	Published	0.2		0.4	3.4
Hanna [43]/Holloway [46, 47]	1983	USA	1	1500	Published				3.6
Dimson and Marsh [25]	1984	Britain	11	862	Published			2.5	1.3
Weighted Average of All Results	1933–1984	N. America and EAFE	211	47053	Mainly Published	0.6	1.5	1.1	0.6

for the infrequently traded shares in his sample are subject to potentially large estimate errors, and conceivably, this could vary systematically with the type of broker.

In addition to studies which have looked at the performance of buy, sell, and hold recommendations, a number of researchers have looked directly at the accuracy of analysts' stock price forecasts. Most of these studies have evaluated forecasting ability by measuring the correlation coefficient between forecast returns and subsequent outcomes. The results of these investigations are summarised in Table II. One of the earliest studies, by Malkiel and Cragg [56], worked with a set of earnings predictions made by 17 investment firms (of which four were brokers) for 178 corporations for each of the years 1961–65. The authors were then able to infer the degree of over- or under-pricing for each security implicit in the analysts' forecasts, and hence obtain a set of implied forecast returns. The correlation coefficient between these implied forecasts and the subsequent returns was 0.06. In a more direct study, Hodges and Brealey [44] measured the industry forecasting ability of 25 UK institutional investors who were asked to rank 22 sector indexes on the basis of forecast returns assuming an unchanged market. Working in groups, the subjects of this study produced nine independent rankings of anticipated sector performance. The mean Spearman's rank correlation coefficient between these forecasts and the subsequent outcomes of 0.15 indicated a small, but potentially useful, degree of forecasting ability.

Ambachtsheer [1] adopted this procedure to examine some 250 forecasts of returns on US stocks made by eight analysts in a Canadian investment institution during 1981. He found a correlation coefficient between forecast and actual ranked returns of 0.16 and observed that a large part of the information content of the forecasts was discounted in the market place within the first month. In a follow-up study [2], he examined some 6800 returns forecasts made by 16 independent investment institutions during 1972–74. Of these forecasts, approximately 1000 were industry ratings while the remainder were for individual securities. For the security forecasts, Ambachtsheer found correlation coefficients of 0.17, 0.16, and 0.09 over the periods 1–3, 4–7, and 8–12 months after the forecast, respectively. Once again, this suggests that the forecasts, which in this case were cited as being for a six- to twelve-month horizon, provided greater

Table II

Estimates of Forecasting Correlations

Publication	Date	Country	Number of Organisations	Number of Recommendations	Mean Forecasting Correlation
Malkiel and Cragg [56]	1970	USA	17	178	0.06
Hodges and Brealey [44]	1972	Britain	1	198	0.15
Ambachtsheer [1]	1972	USA	1	250	0.16
Ambachtsheer [2]	1974	USA	16	6800	0.16
Ambachtsheer and Farrell [3]	1979	USA	2	1200	0.07–0.14
Garmaise [41]	1982	Canada	10	6868	0.03–0.10
Coggin and Hunter [13]	1983	USA	1	1986	0.10

explanatory power in the earlier rather than the later months. The corresponding correlation coefficients for the industry forecasts were 0.07, 0.18, and 0.11. Subsequently, Ambachtsheer and Farrell [3] used the same technique to assess the quality of two investment advisory services. They found that for a 200-stock universe over six consecutive six-month periods from September 1973 to September 1976, the average correlations between forecast ratings and actual outcomes were 0.067 and 0.135 for the Value Line and Wells Fargo stock ratings, respectively. They found that the forecast ratings from these two services were virtually independent, so that when the two sets of ratings were optimally combined, the correlation coefficient between the composite forecasts and subsequent returns was 0.152.

More recently, Garmaise [41] evaluated 6968 stock return ratings which were forecast by ten leading Canadian stockbroking firms. The average correlation between forecast return ranks and actual return ranks was computed for five successive three-month forecasts, for three six-month forecasts, and for one twelve-month forecast. The results were 0.04, 0.10, and 0.03, respectively, for these three horizons, although Garmaise reports that these figures were not significant at the 95 percent confidence level. Coggin and Hunter [13], on the other hand, found that almost 2000 forecast ranks produced by analysts at a large trust company during 1979–81 had a correlation which, at 0.10, was significantly different from zero. Finally, Farrell [31], in follow-up work to [3], reports on two other experiments where the forecast sources examined (two and four sources, respectively) proved both to have predictive ability and also to provide virtually independent forecasts. While the correlations obtained are not reported, Farrell [31] does provide information on the abnormal returns obtained from both simulated and (latterly) actual portfolios which were explicitly managed to exploit the pooled forecasts derived from these sources.

In summary, previous studies of the performance of recommendations made by brokers and advisory services indicate a small but potentially useful degree of forecasting ability. In cases where recommendations are made public, most of the information content appears to be impounded into share prices by the end of the publication day. Even for analysts' predictions which are not made public, forecasts provide their greatest explanatory power in the early months immediately after the forecast. Despite the proliferation of studies in this field, however, the existing evidence remains controversial. As Coggin and Hunter [13] point out, the small-sample studies [44, 56] are far too limited to provide reliable estimates of forecasting correlations. All the remaining studies [1, 2, 3, 31, 41] are retrospective experiments, in which an advisory organisation has collected forecast data related to the service it provides, has analysed it, and then made a decision on publication in the light of the outcome. Other organisations may have obtained neutral or perverse results which have not been disseminated, and this possibility provides an important motivation for the current study. Only a prospective experiment, in which the decision to undertake published research occurs before data are collected and analysed, can provide evidence which is not tainted by the possibility of selection bias.

Previous studies have, in the main, concentrated on establishing the existence of forecasting skills. Given the commercial, rather than academic origins of the evidence to date, it is not surprising that the existing literature provides little supplementary information on the characteristics of analysts' forecasts. One limitation common to the recommendations analysed in virtually every study is that they were either binary (buy/sell), trinary (buy/hold/sell), or at best five-point (strong or weak buy/hold/strong or weak sell) ratings. This degrades the information which could potentially be produced by the analyst, since the optimal investment in a stock depends on a scalar forecast of its specific return (see Treynor and Black [78], Rosenberg [63], Rudd and Rosenberg [64], and Sharpe [70]). Many other issues remain unresolved. For example, with the exception of Farrell [31], existing studies provide no information on whether forecasts which appear profitable in a hypothetical paper portfolio can actually be translated into profitable transactions (see Treynor [77]). Similarly, we have only snippets of evidence about the rate at which forecasts decay over time (see Ambachtsheer [1] and [2]). As Coggin and Hunter [13] point out, existing studies also tell us very little about whether there are true differences in forecasting skill between different advisors, and we have virtually no evidence on the stability of predictive ability over time. Finally, with the exception of the findings of Ambachtsheer and Farrell [3] and of Farrell [31], we know almost nothing about the degree of independence between stock price forecasts from different advisors (see Sharpe [70]), nor about the potential benefits from pooling such forecasts, even when they are not contemporaneous. The current study addresses all of these issues, providing new evidence on these important, unresolved questions.

II. Methodology, Sample, and Data

This study is designed to estimate the extent to which brokers and analysts are able to forecast the specific component of security returns. We assume that stock returns conform to the Sharpe [68], Lintner [52] capital asset pricing model (CAPM), so that:

$$R_{jt} = (1 - \beta_j)R_{ft} + \beta_j R_{mt} + u_{jt} \tag{1}$$

where R_{jt} is the return on stock j during period t, R_{mt} is the corresponding market return, R_{ft} is the risk-free rate, β_j is the beta or systematic risk of stock j, and u_{jt} is the error term, which we will henceforth refer to as the specific return on security j.

Our use of the CAPM involves assuming a specific model of equilibrium expected return, and may at first sight appear unduly restrictive, particularly in the light of Roll's [62] critique of the model's empirical testability. However, all studies of efficiency must either assume such a model, or else make equally strong implicit assumptions about the stability of the return generating process, as in, say Sharpe's [67] market model. The latter would be particularly inappropriate when estimating forecasting correlations for specific security returns. Note that, if the CAPM holds, it will not be possible to predict the specific return u_{jt} in Equation (1) from historical data. This is in direct contrast to the task of predicting the error term v_{jt} in the usual version of the market model,

$$R_{jt} = \alpha_j + \beta_j R_{mt} + v_{jt} \tag{2}$$

where α_j is projected from the observed alpha over a prior estimation period. For stocks which performed well (poorly) during the estimation period, α_j will be projected to be positive (negative); and since favourable (unfavourable) performance does not persist in weak-form efficient markets (see Fama [28] and Dryden [27]), the v_{jt} can be predicted to be negative (positive).[1] The market model would thus provide a quite inappropriate benchmark, since analysts who simply predict trend reversals will always appear to have positive forecasting skills.[2]

We assume that a key role of both internal and external research is to provide ex ante forecasts of both R_{mt} and u_{jt} in Equation (1). In any given time period t, there will thus be a single forecast for R_{mt}, and a set of u_{jt} forecasts, one for each share analysed. Forecasts of specific security returns thus dominate numerically, and also, as Treynor and Black [78] have demonstrated, in terms of economic importance. Our focus in this study is on evaluating the extent to which analysts and brokers are able to forecast the u_{jt}. Let F_{kjt} be analyst k's forecast of u_{jt} made immediately prior to the start of period t. Then we can estimate analyst k's forecasting ability by fitting the model

$$u_{jt} = \gamma_{0k} + \gamma_{1k} F_{kjt} + e_{kjt} \tag{3}$$

where γ_{0k} and γ_{1k} are parameters which can be estimated by ordinary least squares (OLS) regression, and where e_{kjt} is a zero mean error term.

The slope coefficient γ_{1k} can then be interpreted as the coefficient of forecasting ability for analyst k. (When the u_{jt} and F_{kjt} are converted to a ranking on a 1–5 scale, the slope of Equation (3) is sometimes referred to as the information coefficient or IC; see Ambachtsheer [2].) Note that when we estimate Equation (3), we are, in fact, explicitly testing whether or not the market is efficient. A value of γ_{1k} significantly different from zero would provide evidence of the existence of information which, at the time of the forecast, was not fully

[1] Equation (1) is, of course, consistent with the market model if α_j is not extrapolated from the past but is instead specified as $(1 - \beta_j)R_{ft}$ Furthermore, if the stock betas are all equal to one—and they are in fact close to this value—Equation (1) amounts to use of the market model with an α_j of zero. Consequently one possible justification for using the market model in this study would be if the expected α_j were nonzero and were also stationary over time. This is the argument that is sometimes put forward to support use of the market model in the presence of a "size effect" in security returns: it is alleged that the α_j term might capture the tendency for larger companies to underperform their smaller counterparts. However, if the sample consists of large capitalisation stocks and a capitalisation-weighted index is used, the impact of the size effect is likely to be small. In this study, our sample (described later) is confined almost entirely to stocks in the largest decile of capitalisations on the London Stock Exchange. This decile performed neutrally relative to the All Share Index during the period of our study, outperforming by 1 percent in 1980 and underperforming by 1 percent in 1981 (see Dimson and Marsh [25]). For these reasons, the use of the market model as our benchmark in this study appears to offer no compensating advantages to help offset the serious problems arising from the inherent predictability of the v_{jt} (see also footnote 2).

[2] Many technical analysts might thus appear to be good forecasters when evaluated against the market model. Similarly, if, as in the current study, the pattern of pre-forecast specific returns differs between analysts (see Section III.*B*), the market model will result in quite inappropriate rankings of forecasting ability. The use of an exclusion period (see Fama et al. [30]) might serve to mitigate some of these biases, but the choice of exclusion period itself poses serious difficulties (see Dimson and Marsh [25]). In any event, this would not prevent "cheating" by analysts who were aware that the market model was being used, and who also knew the estimation period.

discounted in share prices. As discussed in Section I, this is a strong form test, since we are analysing nonpublic forecasts, which are partially based on privileged information and on private insights.

As noted in Section I, previous studies have focussed on the correlation (ρ_k) between forecast and realised returns, since this is ultimately the variable of economic interest (see Treynor and Black [78]). Clearly, γ_{1k} is directly related to the correlation coefficient since $\gamma_{1k} = \rho_k\ \sigma_u/\sigma_f$, where σ_u^2 and σ_f^2 are variances of the realised specific returns (u_{jt}) and forecast specific returns (F_{kjt}), respectively. In this study, we will therefore present our results in terms of correlation coefficients (ρ_k), as well as reporting regression coefficients (γ_{0k} and γ_{1k}) and providing evidence on the relative dispersion of actual and forecast returns (σ_u and σ_f). Note, however, that when investigating the rate at which forecast information is incorporated into share prices, slope coefficients may be easier to interpret than correlations. When forecasts are tested against realisations on a regular basis through time, the correlation coefficient can (and usually does) fall as the horizon is extended, due to increases in the variance of realised returns (see Kendall and Stuart [51, Exercise 26.17] on attenuation errors). The slope coefficient, however, will increase as long as forecasts still have incremental predictive power for the later subperiods. A decline in correlation could thus be masking evidence of continued predictive ability. If, in addition, we monitor changes in slope coefficients, however, we obtain a direct measure of the rate of information flow.

Our forecast data consist of 4187 specific return forecasts for 206 domestic shares made during the calendar years 1980 and 1981. The data were collected by the research department of a large institution, which we will hereafter refer to as the fund. Although the fund's equity portfolio performed well over the sample period, the experiment was initiated by the authors in conjunction with executives of the fund prior to the beginning of 1980, and the forecasts therefore avoid ex post selection bias. Of the total sample, 3364 forecasts were made by 35 different UK stockbrokers. The fund collected this data systematically as part of its regular research programme, so that the securities covered represented shares in which the fund had a particular interest. While this gave rise to a small bias towards capital goods and financial stocks at the expense of other industrial groups, oils, and natural resources, the sample's industrial composition nevertheless very closely resembled that of the market.[3] The 206 shares covered accounted for some 75 percent by value of the *Financial Times*—Actuaries All Share Index, and since this index includes over 90 percent by value of all UK listed equities, the sample is representative of the stocks monitored by most institutional investors. The average beta of the shares, computed in [24] and relative to the above index, was 1.01, and the cross-sectional standard deviation of the estimated betas was 0.10.

In collecting the forecasts, the fund requested the brokers to provide forecasts of specific returns by asking them to forecast the excess return on the shares

[3] Of the 4187 forecasts, 35 percent were for capital goods shares, 24 percent for consumer goods, 4 percent for other industrial groups, 3 percent for oil and natural resources, and 34 percent for financial shares. The corresponding proportions of All Share Index constituents in these sectors were 28, 26, 11, 4, and 31 percent, respectively.

assuming an excess return on the market of zero.[4] Brokers' analysts were asked for one-year forecasts, but when analysts felt unable to forecast one-year returns, they were permitted to specify their own time horizon. In spite of this, 95 percent of the forecasts were for one year, and only 2 percent were for periods of less than six months. While 35 brokers were involved in total, for any particular share, the fund collected forecasts from only those brokers whose research departments covered the share in question, and whom the fund regarded as leaders in that sector. The 35 brokers involved in this study include all the large research-oriented London brokers together with one or two of the more prominent provincial brokers.[5] In all cases, the brokers' forecasts were prepared specifically for the fund and collected by the internal analysts at a date of the fund's choosing; they were not made publicly available. Because of the very large size of the fund's equity portfolios and the substantial commissions which it generates, we believe the brokers took this exercise seriously.

Initially, when data collection commenced in January 1980, the fund requested forecasts for the majority of companies on a quarterly basis, although this was changed to a semiannual basis in the second half of 1981. Because of the magnitude of the work programme involved in contacting a large number of brokers for a wide selection of shares, forecasts were collected on a staggered basis, although they tended to be concentrated around the end of each quarter. The number of forecasts per broker ranged from 10 to more than 300. Although most supplied around 100 forecasts, nine brokers provided less than 20, while six furnished more than 200.

[4] Clearly, there are a number of ways in which the brokers might have misinterpreted these instructions. For example, some brokers may be more accustomed to providing forecasts of returns relative to the market ($R_{jt} - R_{mt}$), rather than conditional forecasts of excess returns. Others may tend to think in terms of capital gains rather than total return forecasts (i.e., $C_{jt} - C_{mt}$, where C_{jt} and C_{mt} are the percentage capital gains on stock j and on the market index during period t). Finally, some analysts may be more used to forecasting changes in price relatives (i.e., $(P_{jt+1}/P_{mt+1})/(P_{jt}/P_{mt}) - 1 = (C_{jt} - C_{mt})/(1 + C_{mt})$, where P_{jt} and P_{jt+1} are the prices of stock j at the beginning and end of period t, and P_{mt} and P_{mt+1} are the corresponding levels of the market index). Such misunderstandings give rise to only small differences in the forecasts, and thus seem unlikely to have a material impact on our results. Furthermore, there is no evidence to suggest that any misunderstandings have taken place, since there is no correlation between the forecasts and the betas of the forecast stocks, nor is there any relationship between the forecasts and subsequent market returns. Nevertheless, the question of whether misunderstandings might potentially have affected our results is an empirical one, and so we therefore investigated the sensitivity of our results to misinterpreted instructions. To do this, we reestimated the correlation coefficient between forecasts and horizon realisations (u_{jt}) using the three alternative definitions of u_{jt} suggested above. The results were virtually identical, and in all three cases, the horizon correlation coefficient differed by less than 0.003 from that obtained assuming u_{jt} to be the realised specific return from the CAPM (Equation (4)).

[5] The brokers involved (in alphabetical order) were Buckmaster & Moore, Capel Cure Myers, James Capel, Carr Sebag, Cazenove, De Zoete and Bevan, Fielding Newson-Smith, Gilbert-Eliott, Hoare Govett, W. Greenwell, Grieveson Grant, Henderson Crosthwaite, Henry Cooke Lumsden, Kitkat & Aitken, Laing & Cruickshank, Lawrence Prust, McAnally Montgomery, L. Messel, W. N. Middleton, Montague Loebl Stanley, Panmure Gordon, Phillips & Drew, Quilter Hilton Goodison, Rowe & Pitman, Savory Milln, Scott Goff Hancock, Scrimgeour Kemp-Gee, Albert E. Sharp, Simon & Coates, Spencer Thornton, Tilney, Vickers da Costa, Vivian Grey, Williams de Broe, and Wood Mackenzie.

In addition to the 3364 brokers' forecasts, our forecast data also include 823 forecasts made by the fund's own internal analysts. During this period, the fund employed five analysts, one of whom was the manager. Each analyst specialises in a particular sector of the market. The organisational structure is such that all analysts execute trades and share a team responsibility for trading and hence, ultimately, for fund performance. The analysts are therefore responsible for monitoring stockbrokers' forecasts and using them in conjunction with their own primary research. They also analyse securities which are not followed by the brokers, and 139 of the internal forecasts were for stocks which were in this category. Thus, while the brokers' forecasts are important and are taken into account by the analysts, it is ultimately the internal team's own forecasts which underlie the fund's active dealing decisions. Finally, the fund is actively managed in the Treynor-Black [78] sense, and short sales can be executed by borrowing stock from an inventory fund; this permits the fund to take advantage of sale, as well as purchase, recommendations.

In order to evaluate all of these forecasts using Equation (3), additional data are required in the form of estimates of the realised specific returns (u_{jt}) for each security, corresponding to each forecast period. We calculated u_{jt} from Equation (1) as:

$$u_{jt} = R_{jt} - (1 - \beta_j)R_{ft} - \beta_j R_{mt} \tag{4}$$

Period t was taken as the period from the close of business on the date the forecast was made until the end of the month following the forecast horizon. Thus, for example, for a twelve-month forecast received by the fund on January 16, 1980, period t extends from the close of business on January 16, 1980 until the close of business on the last trading day of January in 1981. Security returns were calculated from the middle market closing prices recorded at the month-end on the London Share Price Database (LSPD).[6] Since the LSPD is a monthly database, the middle market closing price on the day of the forecast was collected separately. The market return (R_{mt}) was calculated as the return on the *Financial Times*-Actuaries All Share Index[7] during period t, while the risk-free rate (R_{ft}) was taken as the three-month Treasury Bill rate ruling on the date of the forecast, and appropriately adjusted for the length of period t. Finally, the β_j's were estimated from the previous five years' LSPD returns data using month-end transaction prices. We used the trade-to-trade estimation method to minimise thin trading bias in beta (see Dimson and Marsh [26]).

[6] The LSPD provides both month-end quotation prices and transaction prices for the last trade of the month (see Smithers [72, 73]). In order to minimise thin trading bias in post-recommendation returns, we therefore used the LSPD quotation prices. The latter should suffer from minimal nonsynchroneity especially in view of the high frequency with which the current sample trades; the 206 stocks had a marketability rating (mean time from last trade of the month to month-end) in the London Business School's *Risk Measurement Service* [24] of under 0.1 day.

[7] The *Financial Times*-Actuaries All Share Index is a broadly based, capitalisation-weighted arithmetic index covering 750 stocks. It is described in Bell and Greenhorn [5]. Index returns were computed by incorporating an estimate of the dividend yield based on the published yield on the Index, as in Dimson and Marsh [26]. Because the forecasts here were for large UK companies, we believe this is the best choice of a benchmark (cf. Dimson and Marsh [25]).

In Section III, we present our empirical results. First, we examine the distributional characteristics of both the forecast and actual specific returns. Next, we present evidence on the extent of brokers' and internal analysts' forecasting ability. In particular, we examine the speed with which any information contained in the forecast is incorporated into the realised returns. We then investigate whether differences in brokers' forecasting ability persist over time and can usefully be exploited. Finally, we look at the degree of dependence between simultaneous forecasts to investigate the potential gains from combining forecasts from multiple sources.

III. Empirical Results and Analysis

A. Characteristics of Forecasts and Actuals

The first two columns of Table III show the means and standard deviations of the brokers' and internal forecasts, and of the corresponding realised specific returns. The latter are measured both to the forecast horizon, and also over fixed length periods of 1, 2, 3, 6, and 12 months after the date of the forecast.[8] Looking first at the figures for the brokers, Panel A shows that the mean forecast specific return for the shares covered by our sample was 3.6 percent, indicating that the brokers made rather more purchase than sale recommendations. However, Panel A also shows that over the same time horizon the mean realised specific return for the sample was 1.0 percent, indicating that the shares for which we have forecasts only marginally outperformed the market. The brokers' forecasts were thus biased upwards by, on average, some 2.6 percent.

The second column of Panel A shows that the cross-sectional standard deviation of the brokers' forecasts (10.2 percent) was only a third that of the realised horizon returns (30.9 percent). This tendency for the forecasts to have a much narrower spread than the realised returns indicates fairly cautious forecasting behaviour on the part of the brokers. It does not correspond, for example, to the degree of "exaggeration" reported by Hodges and Brealey [44] who found that forecasts had roughly the same degree of dispersion as the subsequent realised returns. It is interesting to note that although the brokers are typically forecasting twelve months ahead, their forecasts have roughly the same dispersion as the realised specific returns in the one-month period immediately following the forecast.[9]

Panel B indicates that the internal forecasts had a very similar distribution to the brokers' forecasts. The standard deviations were virtually identical, and while the mean internal forecast was slightly higher at 4.9 percent, the mean realised return was also somewhat higher at 1.9 percent. The internal forecasts are thus biased upwards by some 3 percent or by approximately the same amount as the brokers' forecasts. The dispersion of actual returns is slightly higher (34.0

[8] These periods are fixed in length from the end of the month of the forecast. Consequently, the "one-month" period actually averages about six weeks, with a range of plus or minus two weeks.

[9] The "one-month" period is not exact in duration; see footnote 8. The narrower dispersion of forecast twelve-month returns would be rational behaviour for a forecaster who believes himself to have limited, but positive, forecasting skills.

Table III

Relationship Between Forecast and Realised Specific Returns

		Distribution of Returns[a]		Constant Term		Slope Coefficient		Correlations	
Source of Forecast	Time Since Forecast	Mean Return (%)	Standard Deviation (%)	Constant	t-Value	Slope	t-Value	Coeffic-ient	R-Squared
Panel A:	To Horizon[a]	3.6	10.2	na	na	na	na	na	na
Brokers'	1 Month	0.4	9.9	−0.093	(80.9)	0.124	(7.5)	0.128	0.016
Forecasts	2 Months	0.7	12.5	0.175	(68.1)	0.132	(6.3)	0.108	0.011
(N = 3364)	3 Months	0.4	15.3	−0.092	(62.1)	0.148	(5.7)	0.098	0.009
	6 Months	1.1	21.3	0.349	(59.6)	0.197	(5.5)	0.094	0.009
	12 Months	1.3	32.0	0.275	(54.2)	0.269	(5.0)	0.086	0.007
	To Horizon[b]	1.0	30.9	0.093	(54.1)	0.259	(5.0)	0.086	0.007
Panel B:	To Horizon[a]	4.9	10.4	na	na	na	na	na	na
Internal	1 Month	0	10.2	−0.695	(45.7)	0.134	(4.0)	0.137	0.018
Forecasts	2 Months	0.2	13.3	−0.217	(23.1)	0.089	(2.0)	0.070	0.004
(N = 823)	3 Months	0.6	15.9	0.153	(20.9)	0.096	(1.8)	0.063	0.003
	6 Months	0.2	22.4	−0.572	(24.1)	0.156	(2.1)	0.073	0.004
	12 Months	2.1	34.2	1.387	(15.0)	0.149	(1.3)	0.045	0.001
	To Horizon[b]	1.9	34.0	1.194	(13.9)	0.138	(1.2)	0.042	0.001
Panel C:	To Horizon[a]	3.9	10.3	na	na	na	na	na	na
All Forecasts	1 Month	0.3	10.0	−0.205	(92.2)	0.125	(8.4)	0.129	0.016
(N = 4187)	2 Months	0.6	12.7	0.095	(70.6)	0.122	(6.4)	0.099	0.010
	3 Months	0.5	15.4	−0.053	(65.0)	0.137	(5.9)	0.091	0.008
	6 Months	0.9	21.5	0.169	(63.4)	0.186	(5.8)	0.089	0.008
	12 Months	1.4	32.4	0.468	(55.4)	0.246	(5.0)	0.078	0.006
	To Horizon[b]	1.2	31.5	0.283	(54.7)	0.236	(5.0)	0.077	0.006

[a] The first row within each panel shows the mean and standard deviation of forecast returns. Subsequent rows show the mean and standard deviation of realised specific returns for different time periods since the date of the forecast.

[b] Most forecasts were specified over a time horizon of 12 months, but a minority spanned differing time horizons. Here returns are computed to the horizon date specified by the forecaster.

percent) than for the brokers' forecasts (30.9 percent). This may reflect some small tendency for the internal analysts to focus on slightly riskier shares, and suggests that extreme results (favourable or unfavourable) are more likely for the internal forecasts. These differences are small, however, and the overall picture is one of similarity between the distributional characteristics of forecasts and outcomes for both the brokers and internal analysts. This is perhaps hardly surprising, as the internal analysts themselves collect and review the brokers' forecasts, which are collected in the same forecasting rounds as the internal forecasts.

B. The Degree of Forecasting Ability

The remaining columns of Table III show the results of estimating the extent of forecasting ability using Equation (3). The bottom line of Panel A indicates that the brokers' predictions exhibit a small but statistically significant degree of forecasting ability when tracked to their horizon date, with a slope coefficient of 0.259 and a t-value of 5.0.[10] The correlation coefficient between forecast and realised specific returns of 0.086 is slightly lower than that found in previous studies, but still indicates a potentially useful degree of forecasting ability[11] (see Hodges and Brealey [45]).

The bottom line of Panel B in Table III shows that the internal forecasts also exhibit a positive degree of horizon forecasting ability. Here, the slope coefficient is 0.138 with a t-value of 1.2, and the correlation coefficient between forecasts and actuals is 0.042. These figures are somewhat lower than for the brokers, and at first sight this seems implausible, since the internal analysts have access to all of the brokers' forecasts. However, we have already noted that 139 of the internal forecasts were for stocks not covered by the brokers. Similarly, 306 of the brokers' forecasts were for stocks not covered by the internal analysts. Finally, of the stocks covered by both parties, there was a tendency for the analysts to produce proportionately more forecasts for those companies which were relatively neglected by the brokers. Thus, the internal forecasts were somewhat weighted towards smaller and more risky securities.

Variation in the risk characteristics of the stocks which are being forecast introduces heteroscedasticity into Equation (3). This is because the forecast error term e_{kjt} in Equation (3) is likely to have a standard deviation which, for low levels of forecasting ability, is approximately proportional to the specific risk of the stock, σ_u. To examine this possibility, we reestimated Equation (3) using

[10] Because of the lack of independence between the longer period returns, the t-values may overstate the significance of longer term forecasts. This is less of a problem when forecasts are only tracked for up to three months.

[11] We checked whether our coefficient estimates were influenced by outliers amongst the forecasts or amongst the realised returns. We did this by truncating both distributions at the ±3 standard deviation limits; the truncation was undertaken using an iterative procedure such that all nontruncated observations were within ±3 standard deviations of the truncated distribution (see Huber [48]). The estimated correlation coefficient between brokers' forecast and realised specific returns changed by −0.004 from 0.086 to 0.082; the figure for internal analysts changed by +0.001. There were similar small changes in the slope coefficients (t-values) which changed by 0.01 (−0.2) for the brokers and by +0.01 (+0.0) for the internal analysts.

weighted least squares so as to ensure an approximately homoscedastic error term. Our regression equation was:

$$u^*_{jt} = \gamma_{0k} + \gamma_{1k}F^*_{kjt} + e_{kjt} \tag{5}$$

where $u^*_{jt} = u_{jt}\ \sigma_t/\sigma_{jt}$ and $F^*_{kjt} = F_{kjt}\ \sigma_t/\sigma_{jt}$. Here, σ_t is defined as the weighted average of the σ_{jt} for all $j = 1, \cdots, 206$ stocks. Also note that, although the intercept, slope, and correlation coefficient of regression (5) have a different interpretation from the corresponding estimates using Equation (3), they provide us with an efficient method of checking whether the poorer performance of the analysts is merely a reflection of their concentration on riskier stocks.

The horizon forecasting correlation given by (5) is 0.069 for the brokers and 0.076 for the analysts, while the corresponding slope coefficients are 0.205 and 0.261, respectively. The ranking of the analysts vis-à-vis the brokers is reversed, although the differential is not significant. Evidently, the internal concentration on riskier stocks gave rise to forecasts which, with the benefit of hindsight over the sample period, proved to be less accurate. Further evidence that this was the case is given by the forecasting abilities of the brokers and internal analysts for the stocks covered by both parties. Estimated using Equation (3), the horizon forecasting correlation for this group of stocks was 0.084 for the brokers and 0.082 for the internal analysts, while the corresponding figure for stocks forecast by one party only was −0.023.[12] Although it may have been worthwhile ex ante for the internal analysts to monitor those stocks which were neglected by the brokers (and vice versa), the ex post results were perverse. For all other stocks, both internal analysts and brokers exhibited positive, potentially useful, and virtually identical degrees of forecasting ability.

It is interesting to note that while both brokers and internal analysts exhibited positive forecasting ability, the pattern of pre-forecast specific returns was quite different for the two groups. Panels A and B of Figure 1 show the cumulative specific returns over the one-year pre-forecast period for the five quintiles formed by ranking both the brokers' and the internal analysts' forecasts in ascending order by forecast return. For the brokers, the cumulative specific returns in the year leading up to the forecast were −4.7, −0.3, 0.3, 1.7, and 0.8 percent for the lowest through to the highest forecast quintiles. The corresponding figures for the internal analysts were 7.6, 4.4, 4.0, −1.4, and 3.1 percent. This may possibly indicate slight differences in the forecasting "philosophies" adopted by the brokers and the internal analysts. Conceivably, brokers may tend to forecast events which are already partially discounted by the market; forecasts by such "information-based" analysts (see Treynor [77]) may well appear to be predicting a continuation of recent relative strength or weakness. The internal analysts, on the other hand, may be "valued-based" transactors (see [77] again), whereby forecasts are to some extent triggered by the belief that a particular stock has temporarily moved away from its equilibrium value.

[12] There were in all 3742 forecasts for stocks covered by both the internal analysts and the brokers. For these forecasts, the correlation coefficients of 0.084 and 0.082 correspond to 3058 brokers' forecasts and 684 internal forecasts, respectively. The remaining 445 forecasts were for stocks covered by one party only, and for these forecasts the correlations were −0.088 for the 139 internal forecasts and 0.009 for the 306 brokers' forecasts.

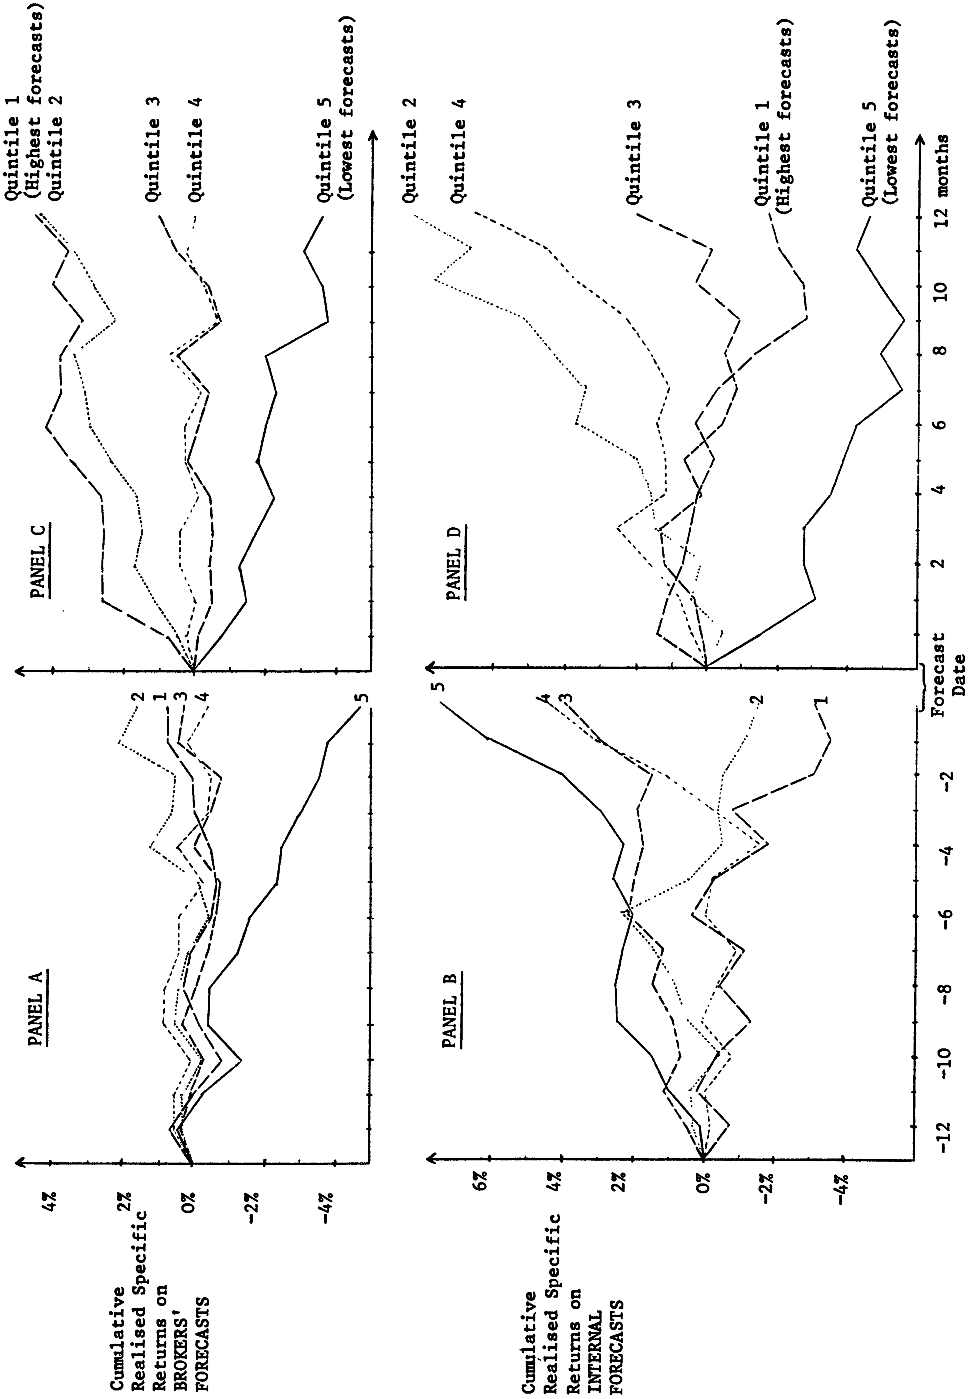

Figure 1. Performance of Forecast Quintiles for Both Brokers' and Internal Forecasts

Clearly, however, our principal concern here is with post-forecast date performance. In order to investigate the potential gains from utilising the forecasts, we therefore examined the post-forecast specific returns for each quintile, for both the brokers' and the internal analysts' forecasts. Note that for the brokers' forecasts, the average predicted returns corresponding to the lowest to the highest forecast quintiles were −10, −1, 3, 8, and 18 percent, respectively, while the corresponding figures for the internal forecasts were −8, −1, 4, 10, and 19 percent. Figure 1 (Panels C and D) shows the cumulative specific returns calculated from Equation (4) for the five different portfolios (quintiles) for both the brokers' and internal forecasts over the one-year period subsequent to the date of the forecast.

Panel C shows that for the brokers' forecasts, the performance of the five portfolios broadly conformed to the predictions, at least in terms of ranking. The cumulative specific returns over the twelve-month period after the forecast were −3.6, 0.0, 1.0, 4.4, and 4.5 percent for the lowest to the highest forecast quintiles, respectively. Clearly, therefore, as already indicated by the slope coefficients of less than unity in Table III, UK brokers have some modest forecasting ability, but at the same time, exhibit a tendency to "exaggerate" in the sense that high forecasts generally turn out to be overestimates, and low forecasts to be underestimates. This "exaggeration" tendency can be seen by contrasting the predicted range of outcomes −10 to 18, or 28 percent, with the subsequent actual range of outcomes, −3.6 to 4.5 or only 8 percent.[13]

Panel D shows a similar pattern of results for the internal forecasts with twelve-month cumulative specific returns of −4.6, 6.6, 2.0, 8.3, and −1.7 percent for the lowest through to the highest forecast quintiles, respectively. These results, while reflecting the somewhat lower overall forecasting correlation for the stocks covered by the internal analysts, nevertheless still indicate a useful degree of potential outperformance. In the next section, we examine the extent to which this potential was translated into achieved superior returns for the fund's portfolio.

C. Trading Performance and Realised Returns

We noted in Section I that previous studies of forecasting ability have focussed exclusively on the performance of paper portfolio transactions, and not on achieved returns. In practice, however, there may be significant difficulties in translating paper recommendations into timely transactions of an appropriate size. In particular, as Treynor [77] points out, the apparent profitability of paper

[13] This exaggeration tendency reflects inefficiency in our forecasts, which can be seen more clearly when we examine the mean squared error (MSE), and decompose it into its three standard components, bias, inefficiency, and random error (see Theil [76]). The total MSE for our population of percentage forecasts over the horizon period was 1055, made up of random error, 986, inefficiency, 62, and bias, 7. The inefficiency component here measures the extent of "exaggeration bias" in the forecasts, namely the tendency for negative errors to be associated with high forecasts, and vice versa. If we know the slope and constant from Equation (3), we can, of course, adjust the forecasts to remove this exaggeration bias. Indeed, unless such adjustments are made, the forecasts will actually exhibit a higher MSE than the naive forecast that all specific returns will be zero. This naive forecast would involve no inefficiency, would have a random error component of 992, and hence a total MSE of 999 for our sample of stocks.

portfolios may be illusory since, in practice, the bid-ask spread works against a trader who seeks to profit from stock selection skills. It is thus important to evaluate achieved returns based on actual trading prices before reaching any final conclusions on the value of the forecasts.

The current study provides us with a unique opportunity to evaluate not only potential but also achieved returns. Since the internal team is responsible for trading decisions, and since these, in turn, are directly motivated by its research, we can evaluate achieved as opposed to potential returns by monitoring the performance of the fund's trading record. During the period under study, a record was kept of all the fund's transactions, including the actual price at which each purchase or sale was consummated, and the transactions costs involved. While near-neutral forecasts typically generated no trading, other forecasts led to multiple trades in the same stock in the days, or even weeks, following the forecast. This pattern of trading reflected the fund's belief that they were less likely to disturb the market and alert others with a sequence of smaller trades than with a single large transaction.[14] This practice resulted in a total of 2950 transactions in the 206 stocks analysed during the two-year period under study. To evaluate the performance achieved by these transactions, we used Equation (4) to estimate the specific return, based on the actual trading price, on each transaction over the one-year period following the day of the trade. The weighted average of these specific returns was then calculated, the weights being proportional to the market value of each transaction.

The results of this exercise indicate that the transaction-value weighted specific return achieved on the portfolio of all 2950 transactions over the twelve-month period following the trade was 2.2 percent, representing some £3 million. When purchases and sales were examined separately, both sets of transactions showed positive specific returns with the sales (3.9 percent) slightly outperforming the purchases (2.0 percent). These are gross returns, however, and it is clearly of interest to compare them with the transactions costs incurred. However, over this period, the fund grew rapidly, so that a significant proportion of the transactions costs arose from the investment of new money, and was thus unavoidable. The appropriate comparison is, therefore, with avoidable transactions costs, defined as the incremental costs over and above those required simply to invest net new money. On this basis, the net specific return on the total portfolio, after deducting all avoidable transactions costs, was 1.7 percent or £2.3 million. While no direct data are available, we are nevertheless confident that this figure is significantly in excess of the internal team's salary and overhead expenses during this period. We, therefore, conclude that the apparently favourable performance of the forecasts (Figure 1) is not artefactual in nature; a trading program executed on the basis of the forecasts did, in practice, achieve superior returns.

[14] While trading decisions are motivated directly by the internal team's research and forecasts, the precise pattern and timing of the trades may well be influenced by short-term considerations regarding market conditions. The team's record may thus in part be attributable to their trading skills, as well as to their return forecasts. While it is difficult to disentangle these two effects, it is worth noting that the fund's passive portfolio, where no trades are research-motivated, but where trading skills are nevertheless exercised, shows no evidence of superior returns.

D. The Rate at Which Forecast Information Is Incorporated in Share Prices

As we noted above, previous studies have found that a major portion of the information content of forecasts is incorporated into share prices within a short period after the forecast. It is certainly of interest, therefore, to test our forecasts against realisations on a regular basis through time in order to gain some insights into the dynamics of the process of information flow. For example, analysts may state that their forecasts are for a one-year horizon, but they would seldom claim to know exactly when the private information will become public, and hence when the share price will reach its target level. Thus, although we may be dealing with what are nominally one-year forecasts, it may still be the case that information flows rapidly into the market, and that the forecast specific return is realised within, at most, only a few months of the date of the forecast.

Certainly, the results of the current study shown in Table III seem to bear this out. A comparison of the slope coefficients for one- and twelve-month periods following the forecast shows the slope increasing from 0.124 to 0.269 for the brokers' forecasts, and from 0.134 to 0.149 for the internal forecasts. This suggests that nearly half the information content of the brokers' one-year forecasts and nearly 90 percent of the information content of the internal forecasts is incorporated within share prices by the end of the first month after the forecast.[15]

To examine this in more detail, we used Equation (3) to estimate the relationship between forecast and realised specific returns in a sequence of consecutive, non-overlapping periods following the forecast date. To facilitate comparisons, we restricted our sample to only those forecasts which were specified as being for a twelve-month horizon. The periods examined and the slope coefficient estimates for both brokers' and internal forecasts are shown in Table IV. Looking first at the brokers' forecasts, the slope coefficient is positive in all subperiods examined. Clearly, however, the relationship between forecasts and actuals is at its strongest in the period up to the end of the first subsequent month. Indeed the initial two periods, which on average span the six weeks immediately following the forecast, are the only time periods examined for which the slope coefficient has a significant *t*-value. In general, the magnitude of the slope coefficient declines with the time which has elapsed since the forecast date. Note, however, that there is some evidence to indicate that the forecasts have some predictive power even beyond their nominal horizon date, although this relationship is not statistically significant for individual months.

The situation with respect to the internal forecasts is somewhat different. There is a highly significant relationship between forecasts and actuals in the period immediately following the forecast. Indeed, over the period from the forecast date until the end of the forecast month (on average, a two-week period),

[15] Note that although the slope coefficients are higher for the twelve-month than for the one-month periods, the correlation coefficients are actually lower. This is attributable to the higher variance of the twelve-month residual returns, and is precisely what we would expect if forecasts decay over time (see Section II above). Because the forecasts examined in this study have a fairly long time horizon, our twelve-month correlation coefficients may appear low when compared with the "information coefficients" which would be computed over the shorter time horizons used in other studies.

Table IV

Relationship Between Forecast and Realised Specific Returns in Consecutive Non-Overlapping Subsequent Periods

	Brokers' Forecasts			Internal Forecasts		
Period over Which Returns Were Computed	Slope Coefficient[a]	(t-Value)	Number of Observations[b]	Slope Coefficient[a]	(t-Value)	Number of Observations[b]
Forecast date to end of month	0.038	(3.6)	3178	0.124	(6.2)	782
Month 1 (first subsequent month)	0.081	(6.6)	3178	0.006	(0.2)	782
Month 2 (second subsequent month)	0.012	(1.0)	3178	−0.028	(−1.0)	782
Month 3 (third subsequent month)	0.020	(1.4)	3178	0.019	(0.7)	781
Month 4 (fourth subsequent month)	0.013	(1.0)	3178	0.013	(0.5)	781
Month 5 (fifth subsequent month)	0.001	(0.1)	3178	−0.013	(−0.5)	781
Month 6 (sixth subsequent month)	0.009	(0.7)	3169	0.026	(0.9)	778
Average results for Months 7–9	0.011	(0.8)	3175	−0.007	(−0.2)	780
Average results for Months 10–12	0.007	(0.5)	3169	0.001	(0.0)	778
Average results for Months 13–18	0.011	(0.8)	3145	−0.013	(−0.4)	760
Average results for Months 19–24	0.009	(0.6)	2958	0.028	(0.8)	731

[a] Slope coefficient (γ_{1k}) estimated from Equation (3).

[b] To facilitate comparisons, the sample used here was restricted to only those forecasts specified by the forecaster as being for a 12-month horizon.

Table V

Correlation Coefficient, z-Statistics, and Rankings of Brokers in Each Year[a]

	Overall Period			1980 Only			1981 Only		
Broker[b]	ρ	(z)	r	ρ	(z)	r	ρ	(z)	r
1	−0.02	(0.4)	14	−0.00	(0.0)	14	−0.06	(0.6)	15
2	0.25	(3.1)	2	0.38	(3.8)	1	−0.01	(0.1)	12
3	0.08	(1.1)	11	0.12	(1.3)	8	0.01	(0.1)	10
4	0.09	(1.1)	10	0.27	(2.4)	2	−0.24	(1.8)	20
5	0.12	(1.9)	7	0.25	(3.1)	3	−0.05	(0.5)	14
6	−0.09	(1.0)	17	0.10	(0.9)	10	−0.32	(2.5)	21
7	0.18	(3.1)	5	0.22	(3.0)	4	0.04	(0.5)	8
8	0.26	(4.1)	1	0.16	(2.0)	7	0.39	(4.0)	1
9	0.10	(1.3)	9	0.21	(2.0)	5	−0.05	(0.5)	13
10	0.19	(3.3)	4	0.18	(2.3)	6	0.18	(2.1)	6
11	0.15	(1.6)	6	0.11	(0.8)	9	0.21	(1.3)	5
12	−0.13	(1.1)	18	−0.23	(1.6)	20	0.04	(0.2)	9
13	0.04	(0.5)	13	−0.13	(1.0)	17	0.29	(1.8)	2
14	0.07	(0.6)	12	0.09	(0.6)	12	−0.15	(0.5)	17
15	−0.17	(1.6)	20	−0.21	(1.2)	19	−0.15	(1.0)	18
16	0.20	(1.4)	3	−0.01	(0.1)	15	0.25	(1.2)	3
17	−0.06	(0.4)	15	0.10	(0.5)	11	0.00	(0.0)	11
18	0.12	(1.3)	8	0.03	(0.3)	13	0.22	(1.6)	4
19	−0.14	(1.4)	19	−0.15	(1.2)	18	−0.13	(0.7)	16
20	−0.06	(0.4)	16	−0.10	(0.6)	16	0.11	(0.3)	7
21	−0.19	(1.7)	21	−0.25	(1.7)	21	−0.16	(1.0)	19

[a] The correlation coefficient, z-statistic, and ranking are denoted by ρ, z, and r, respectively.

[b] Brokers have been ranked in this table according to their performance in the CI Survey (see text). Those brokers with less than 50 forecasts were omitted from this table. These 14 brokers, listed in declining order of number of forecasts, had correlation coefficients (z-statistics) for the overall period as follows: 0.12(0.8), 0.10(0.6), 0.32(2.0), 0.30(1. 5), −0.08(0.3), −0.11(0.4), 0.04(0.2), −0.44(1.7), 0.80(3.5), 0.51(1.7), −0.10(0.3), 0.80(2.9), −0.03(0.1), and 0.64(1.4).

the slope and correlation coefficient shown in Table IV for the internal forecasts are considerably greater than the corresponding figures for the brokers. Thereafter, however, there is no detectable relationship between forecasts and actuals—the slope coefficient is as often negative as positive, and is never significant.

A plausible explanation for this pattern is that while the brokers' assessments will be taken into account by the internal analysts, it is ultimately the internal team's forecasts and views which determine the fund's trading decisions. Since the fund's analysts participate in trading, this helps to ensure that their views are rapidly translated into market transactions. The resultant trading activity, however, may alert others, including floor traders, to the possibility that there is private information. This in turn will cause prices to adjust very rapidly to their new equilibrium level. Once this level is achieved, we would expect to find no subsequent relationship between forecast and actual returns. Hence, in practice, what were originally regarded as, say, one-year forecasts by the internal analysts may in fact be realised in a far shorter time span due to the informational impact of the fund's own buying and selling actions. The brokers' forecasts, on the other hand, are provided privately to the fund, and thus do not in themselves necessarily generate market transactions. As Treynor [77] notes, such "paper portfolio" decisions will not have an immediate impact on security prices. Conceivably, therefore, this may help to explain why the information content of brokers' forecasts appears to be incorporated more slowly in share prices.

E. Differential Forecasting Ability

We have already noted the apparent differences in forecasting ability between brokers, as a group, and the internal analysts. It is also of interest to examine the extent to which forecasting ability varies between different brokers, and to test whether any such differences persist over time. The first column of correlation coefficients in Table V provides some evidence on this, showing the correlation coefficients between forecast and realised returns for all brokers who made more than 50 forecasts during our sample period. Clearly, there is a wide range of apparent forecasting abilities with correlation coefficient estimates ranging from −0.19 to +0.26.[16] More than a quarter of the brokers have coefficients which, on a one-tailed z-test at the 95 percent level of confidence, are significantly above zero.[17] Nevertheless, it is possible that the apparent differences between the brokers' levels of forecasting skill are a consequence of sampling error, rather than an indication of true differences in analysts' abilities.[18] To test this proposition, we estimated the variance of true correlation coefficients across analysts. Assuming that the k^{th} error variance $\sigma_{\epsilon k}^2$ is independent of analyst k's true

[16] These correlation coefficients are based on regression (3). We noted earlier that the apparent difference between the brokers, as a group, and the internal analysts is not confirmed when the more efficient regression (5) is performed. It is obviously desirable that our findings should not be highly dependent on the precise specification of the regression equation. We, therefore, reestimated all the forecasting correlations reported in this paper using Equation (5). Although the range of estimated correlations widened slightly (to −0.20 through +0.30 for the brokers represented in Table V), our results remained unchanged qualitatively and only varied slightly in quantitative terms. Following the principle of Occam's Razor, we report below only the results based on regression (3).

[17] The reason for using Fisher's z-transformation of the correlation coefficient here is that a few of the brokers (especially those with under 50 observations, which we examine below) have large estimated correlation coefficients. The z-statistic,

$$z = ½ \ln(1 + \hat{\rho})/(1 - \hat{\rho})$$

has a mean and variance of

$$E(z) = \frac{\rho}{2(n-1)}\left\{1 + \frac{5+\rho^2}{4(n-1)} + \frac{11 + 2\rho^2 + 3\rho^4}{8(n-1)^2} + \cdots\right\}$$

$$\mathrm{Var}(z) = \frac{1}{n-1}\left\{1 + \frac{4-\rho^2}{2(n-1)} + \frac{22 - 6\rho^2 - 3\rho^4}{6(n-1)^2} + \cdots\right\}$$

In this test, ρ is taken as zero, but because n is small for some brokers we refrain from using the approximate variance $1/(n-3)$. Kendall and Stuart [51, Ch. 16.33] report that the estimates given in this footnote are good approximations for n as low as 11, and give the exact distribution when $\rho = 0$.

[18] The possibility of our results being influenced by sampling error is most obvious when we consider the cases in which a broker's forecasting correlation is found to be negative. Sampling error

forecasting ability, the cross-sectional variance of the true correlation coefficients is approximately:

$$\sigma_{\rho}^{2} = \sigma_{\hat{\rho}}^{2} - \overline{\sigma_{\varepsilon}^{2}} \tag{6}$$

In Equation (6), $\sigma_{\hat{\rho}}^{2}$ is the cross-sectional variance of the estimated coefficients, where the variance is computed using a weight for each estimate $\hat{\rho}_k$ that is proportional to the degrees of freedom used in computing the estimate. The other statistic, $\overline{\sigma_{\varepsilon}^{2}}$, is the mean of all the error variances $\sigma_{\varepsilon k}^{2}$ for each of the estimates $\hat{\rho}_k$;[19] here, too, the cross-sectional mean is weighted by the degrees of freedom used in computing each estimate $\hat{\rho}_k$.[20]

Our results are shown in the bottom line of Panel A of Table VI. For the overall period, the distribution across analysts of true correlation coefficients has a standard deviation of about 0.11. A chi-square test may be used to test the null hypothesis that all of the estimated coefficients are drawn from a single population with homogeneous forecasting ability.[21] For the overall period, as well as for each subperiod, the null hypothesis may be rejected at above the 99 percent level

---

is the only reason we can conceive for finding instances of perverse forecasting ability, as witnessed by correlation coefficients of less than zero. The conventional wisdom about what constitutes a poor analyst—one who consistently recommends a purchase before the share goes down and a sale before the share goes up—is demonstrably false. Inept research results in lack of correlation between forecasts and subsequent returns, not negative correlation. Trading on the strength of such research will generate losses, not because of adverse price movements, but because of unwarranted transaction costs.

[19] We use the estimator

$$\sigma_{\varepsilon k}^{2} = (1 - \rho^{2})^{2}/n$$

and approximate ρ by $\hat{\rho}$, adjusting for degrees of freedom.

[20] In estimating the variance of true forecasting correlations, we chose to ignore three items of potentially relevant information. First, we ignored a possible set of implicit estimates of each brokers' forecasting skill. An optimal forecast for a *perfect* forecaster would, of course, have a standard deviation equal to that of the realised returns, while an optimal forecast for an analyst with *no* skill would have a standard deviation of zero (the optimal forecast would always be equal to the mean return). It is just conceivable, therefore, that the dispersion of a broker's forecasts provides some information on self-assessed forecasting skills. However, we suspect that any such self-assessments would prove to be unreliable, and the results presented in Subsection *F* below provide some support for this conjecture. We also chose to ignore the fact that, as forecasting ability (ρ, assumed known) rises above zero, an optimal forecaster's predictions will have a larger dispersion; and hence the error variance $\sigma_{\varepsilon k}^{2}$ will not be strictly independent of ρ. Finally, we ignored the argument (presented in footnote 18 above) that the distribution of true forecasting correlations should be bounded at zero. The analysis of the distribution of ρ which follows would (unfortunately) be consistent with the existence of some analysts who had truly perverse skills, but the prospective pay-off from a more sophisticated analysis would seem to be small.

[21] The test-statistic is

$$\chi_{K-1} = K\sigma_{\hat{\rho}}^{2}/\overline{\sigma_{\varepsilon}^{2}}$$

where K is the number of brokers. It is distributed chi-squared with $K - 1$ degrees of freedom if the null hypothesis is true. The test results are $\chi_{34}^{2} = 76.2$, 71.6, and 63.1 for the overall period, the first subperiod, and the second subperiod, respectively.

Table VI

Variation Between Brokers in Levels of Forecasting Skill

	Statistic	Overall Period (3364 Forecasts)	1980 Only (1965 Forecasts)	1981 Only (1399 Forecasts)
Panel A:	Mean observed value of correlation coefficient[a]	0.090	0.108	0.042
Correlation coefficients	Variance across brokers of observed coefficients[a]	0.022 (0.149^2)	0.035 (0.186^2)	0.040 (0.200^2)
	−Mean error variance of observed coefficients[a]	−0.010 (0.101^2)	−0.017 (0.130^2)	−0.022 (0.149^2)
	=Variance across brokers of true coefficients	=0.012 (0.109^2)	=0.018 (0.133^2)	=0.018 (0.133^2)
Panel B:	Mean observed value of z-statistic[a]	0.094	0.112	0.047
Fisher's z-statistics	Variance across brokers of observed z-statistics[a]	0.027 (0.163^2)	0.042 (0.205^2)	0.047 (0.217^2)
	−Mean error variance of observed z-statistics	−0.012 (0.109^2)	−0.022 (0.148^2)	−0.027 (0.165^2)
	=Variance across brokers of true z-statistics	=0.015 (0.121^2)	=0.020 (0.141^2)	=0.020 (0.141^2)

[a] All cross-sectional means and variances are weighted according to the degrees of freedom for each broker.

of confidence. In Panel B of Table VI this analysis is replicated using Fisher's z-transformation of the correlation coefficient, and the results remain unchanged.[22]

It is interesting to see how our ranking of brokers by forecasting ability compares with the rankings obtained from the Continental Illinois Annual Survey [15] of British investment analysts. Each July, Continental Illinois asks more than 200 institutional fund managers to name, in order, the three analysts who "count most" in making their investment decisions. The managers are asked to rank their choices within each of 58 sectors on the basis of depth of knowledge, quality of written material, frequency of follow-up and the success of their recommendations. The study obtains a good response rate, and in July 1981, there were 90 respondents who, between them, were responsible for 53 percent by value of total funds under management in the UK.

Although Continental Illinois produces an overall rating of stockbrokers, this is based not simply on the firm's reputation in equities, but also on a range of other sectors including fixed interest, options, overseas markets, currency, and the overall economy. To ensure appropriate comparisons, we therefore used the underlying survey results to recompute UK equity ratings, based on only the sectors covered in our research. These ratings were calculated on a weighted basis,[23] and naturally differ from the published Continental Illinois results. The brokers are listed in sequence from a "best" to "worst" rating in Table V. We then calculated the rank correlation between the UK equity ratings and our own rankings based on forecasting correlations. The rank correlation coefficient was in fact 0.27, indicating that, in spite of the differences in criteria employed and the somewhat longer period used in this study, there was nevertheless a small amount of agreement between our objective empirical results and the qualitative, subjective assessments of the market place.

F. The Stability of Forecasting Ability

It is clearly important to determine whether or not differences in brokers' forecasting ability persist over time. We therefore divided our total time period into two subperiods covering calendar years 1980 and 1981, respectively. We then recalculated the forecasting correlation coefficients for both subperiods for each broker who made more than 50 forecasts in total. The correlation coefficients and rankings for each broker in each subperiod are shown in Table V. It is clear

[22] The z-transformation is used because of the occasional high estimated correlation coefficients and low number of observations for certain brokers. In computing Panel B of Table VI, we replaced all the $\hat{\rho}_k$ by their z-transformations and replaced the $\sigma^2_{\epsilon\kappa}$ by the variance of the z-value (given in footnote 17). The test-statistics, $\chi^2_{34} = 78.3$, 64.5, and 57.7 for the overall period, the first subperiod, and the second subperiod are again significant at above the 99 percent level of confidence. Coggin and Hunter [13] argue that the z-transformation leads to overestimation of the level of forecasting skill, but it is evident that this effect is unimportant here.

[23] In the Continental Illinois Survey, individual responses are weighted, with analysts obtaining three points for a first place, two for a second, and one for a third. Sector ratings are then calculated by totalling the points obtained by each analyst. For each sector, the top three analysts are then reported, together with their percentages of the total points cast. Our UK equity ratings are calculated by summing the percentage points obtained for all analysts in each firm of stockbrokers, giving equal weight to all 35 equity sectors.

that a particular broker's position in the league table can vary significantly from one year to the next. Indeed, the correlation coefficient between the 21 forecasting correlations shown in the main body of Table V for the first and second subperiods is −0.06, suggesting that there may in fact be no relationship between the rankings in successive periods.

While Table V suggests that there is no stability in estimated forecasting correlations, Table VI provides evidence of true differences between brokers' skills. As pointed out earlier, these differences are statistically significant not only for the overall period, but also for each subperiod taken individually. While it is possible for there to be true differences between brokers that do not persist over time, this is somewhat implausible. Since we might be guilty of a Type II error in concluding that there is no relationship between individual broker's forecasting ability in one period and in the next, it is appropriate to look more closely at the stability of the forecasting correlations.

To examine the stability of forecasting ability, we computed a Bayesian predictor of the brokers' forecasting correlation coefficients. Denoting the correlation coefficient computed in the first subperiod for broker k by $\hat{\rho}_{1k}$ and the cross-sectional mean value of this statistic by $\bar{\rho}_1$, the Bayesian estimate of the forecasting correlation for broker k is:

$$\hat{\rho}'_{1k} = (\hat{\rho}_{1k}\sigma^2_{1\rho} + \bar{\rho}_1\sigma^2_{1\epsilon k})/(\sigma^2_{1\rho} + \sigma^2_{1\epsilon k}) \tag{7}$$

where the variances ($\sigma^2_{1\rho}$, $\sigma^2_{1\epsilon k}$) for the first subperiod are defined as in Equation (6).[24] As shown in Table VI, the cross-sectional variance of the true correlation coefficients σ^2_ρ was 0.133^2 in the first subperiod, while the error variances $\sigma^2_{\epsilon k}$ averaged 0.130^2 over the same period.

To test whether the Bayesian estimators $\hat{\rho}'_{1k}$ have some predictive value we would, in principle, like to run a regression of the estimated forecasting correlation in the second subperiod, $\hat{\rho}_{2k}$, on the predictors $\hat{\rho}'_{1k}$ which were computed in the prior time period. The regression would take the form:

$$\hat{\rho}_{2k} = a_0 + a_1\hat{\rho}'_{1k} + v_k$$

However, the error term v_k in this specification will be severely heteroscedastic because of the very substantial variation in the number of observations used in computing each estimate $\hat{\rho}_{2k}$. The variance in v_k will be approximately proportional to $\sigma^2_{2\epsilon k}$, the error variance for the estimate in the second subperiod $\hat{\rho}_{2k}$. Consequently, division by the square root of this variance estimate will produce a suitable specification for the regression:

$$\hat{\rho}_{2k}\sigma^{-1}_{2\epsilon k} = a_0\sigma^{-1}_{2\epsilon k} + a_1\hat{\rho}'_{1k}\sigma^{-1}_{2\epsilon k} + w_k \tag{8}$$

This specification has a reasonably homoscedastic error, and is a multiple regression with a suppressed intercept. Estimating the parameters of this relationship using the forecasting correlations for all 35 brokers produced estimated

[24] This is analogous to the Vasicek [79]-Litzenberger, Ramaswamy, and Sosin [53] estimator for beta. The correlation is, of course, a slope coefficient in a regression of realised returns on rescaled forecast returns, where the rescaling has equalised the variances of the two variables.

slope coefficients which were positive but nonsignificant. The analysis was also repeated using Fisher's z-transformation of the correlation coefficients in Equations (6)–(8), and the results were not materially different.[25]

While these results could be consistent with a very small degree of stability in analysts' forecasting ability, the estimated relationship is at best tenuous. We conclude that it is unlikely that we will be able to predict reliably which analysts will be superior on the basis of their observed forecasting ability in a prior period. Our best assumption is probably that all brokers have about the same degree of forecasting ability.[26] The best estimate of this ability is given by the forecasting correlation coefficient for brokers taken as a group.

G. Interdependence Between Forecasts

Even though differences in individual brokers' forecasting ability do not appear to persist over time, we might still be able to improve on the levels of forecasting ability shown in Table III by combining forecasts made by different brokers. The potential gains from this will depend on the extent to which different brokers' forecasts for the same share are independent, and the prevalence of simultaneous, or near simultaneous, forecasts. To investigate this, we calculated the matrix of simple correlation coefficients between forecasts made by different brokers. However, because forecasts were collected on a staggered basis, there were relatively few instances of forecasts for the same share falling on precisely the same date. We therefore included in our sample all "near simultaneous" forecasts, defined as being forecasts where there had been at least one other forecast for the same share within the previous 30 days. Brokers making less than 50 forecasts were omitted, and correlation coefficients were estimated only in situations where there were at least 20 "simultaneous" forecasts.

The majority of the correlation coefficients in fact lay in the range −0.15 to +0.30, with a mean correlation coefficient of only 0.08. This suggests that although, on average, there is a small degree of agreement between specific return forecasts made by different institutions, this association is tenuous and scarcely detectable. Indeed, although statistically, the mean coefficient is significantly different from zero, for most practical purposes it would seem a reasonable

[25] The motivation for the z-transformation was to obtain a more normal distribution of errors for use in the OLS regressions. We replaced all $\hat{\rho}_k$ and the revised means and variances in Equations (6)–(8) by their corresponding z-values. The results of regression (8) after the z-transformation were as follows (t-statistics in parentheses):

$$a_0 = 0.045\ (0.65)$$

$$a_1 = 0.035\ (0.08)$$

with a correlation coefficient of 0.22. The results using the raw correlation coefficients without transformation had very marginally higher t-statistics and explanatory power.

[26] This is the assumption which will produce (nearly) the smallest error in predicting individual brokers' levels of forecasting skill from data on the levels of skill exhibited in the past. We are not saying that all brokers have the same degree of forecasting ability, only that we cannot (on the data available to us) identify which ones are "best."

working assumption that forecasts made by different institutions are actually independent.[27]

It is particularly interesting to examine the correlation coefficient between the internal and the brokers' forecasts. Clearly, brokers did not have access to each others' predictions, so that only the internal analysts were in a position to appraise the full range of forecasts. We might therefore expect to find that the correlation coefficient between the fund's forecasts and all "simultaneous" forecasts made by other institutions was higher than the equivalent figure for a typical broker's forecasts. In fact, however, the correlation coefficient between the internal and the brokers' forecasts was only 0.11, which is only marginally higher than the typical correlation for the brokers taken as a group. This is consistent with the notion that the internal analysts exercise considerable independence rather than simply following and collating the brokers' views.

H. Pooling of "Simultaneous" Forecasts

Given that the forecasts have some information content and that "simultaneous" forecasts are virtually independent, it should be possible to improve predictive accuracy by combining simultaneous forecasts. If, as argued above, our working assumption should be that all brokers have the same degree of forecasting ability, then the optimal way of combining truly simultaneous forecasts would be to give each forecast equal weight.[28] In our study, however, each quarter's forecasts are in fact somewhat chronologically dispersed, and the evidence of other studies (Tables I and II) and of our own research (Table IV) indicates that the economic value of a return forecast diminishes rapidly with elapsed time. Indeed, if the market adjusts fully and promptly to the informational content of a forecast, the optimal pooled forecast (assuming analysts have the same forecasting ability) would give equal weight to truly simultaneous forecasts and a weight of zero to previous forecasts. A more realistic scenario, however, would be one in which the predictive value of a forecast is attenuated (but not immediately eliminated) by the passage of time.

To model this process, let us assume that the information embodied in an analyst's forecast for a stock has a stationary probability $1 - \lambda$ of being acquired by the market and impounded into that day's stock price movement, and a probability λ of not reaching the market on that day. It follows that the expected specific return conditional on the forecast is a proportion $(1 - \lambda)$ of the forecast return on the day after the forecast, a proportion $(1 - \lambda)\lambda$ on the subsequent day, a proportion $(1 - \lambda)\lambda^2$ on the third day, and so on. The expected cumulative

[27] Because the forecasts considered here are dispersed over an entire month, there is a risk of overstating the extent of independence between different brokers' forecasts. The mean correlation between timely forecasts might therefore be above 0.08. However, since brokers do not make very many truly simultaneous forecasts, the figure of 0.08 is the appropriate one to use with data of the type that are analysed in this paper. Kendall and Stuart [51, Ex. 26.7] point out that if n forecasters have the same variance, the mean of all the $n(n - 1)$ correlation coefficients must be not less than $-1/(n - 1)$, i.e., the mean correlation is bounded below by a minimum value of -0.05.

[28] See Johnston [50, Ch. 5.3] for a rigorous discussion of the impact on the forecasting correlation from combining forecasts of known statistical characteristics. Newbould and Granger [58] provide the formulae for the relative weight to be given to each forecast.

return by the ℓ^{th} day is therefore the sum of this series, i.e., $1 - \lambda^{\ell}$ multiplied by the forecast specific return.[29] The remaining proportion, λ^{ℓ}, of the forecast specific return may be expected to be realised more than ℓ days after the forecast was made.

If the rate at which information is impounded into stock prices is the same for all stocks and for all analysts, the optimal composite forecast is therefore obtained by giving a relative weight λ^{ℓ} to each ℓ-day-old forecast made in the preceding quarter. (With a quarterly cycle for collecting the input data, forecasts which are more than a quarter old receive zero weight as the broker has had the opportunity to record a revised judgement on the stock.) The optimal pooled forecast to be analysed using Equation (3) is thus an equally-weighted average of all the quarter's forecasts for the stock, where the forecasts are adjusted for the decay in their informational content up to the date of the latest forecast. This adjustment involves giving a weight λ^{ℓ} to each of the K forecasts for stock j:

$$F'_{kjt} = \textstyle\sum_{k=1}^{K} F_{kjt}\lambda^{\ell}/K \tag{9}$$

where $\ell \equiv \ell_{kjt}$ is the age in days at the start of period t of analyst k's forecast for stock j.

The characteristics of the composite forecast F'_{kjt} depend critically on the parameter λ used in Equation (9). With $\lambda = 0$ the composite forecast is, of course, identical to the raw forecast F_{kjt} if a forecast has been made on a particular day, otherwise it is zero. At the other extreme, with $\lambda = 1$ we have a consensus forecast which gives equal weight to every raw forecast made during the forecasting round. Between these limits, it can be shown that if forecasts were made daily and were included in computing the composite figure (9) regardless of their age, the average age of the data would be $d = (1 + \lambda)/(1 - \lambda)$ days; for $d > 10$, about 87 percent of the weight used in constructing the composite forecast is given to the most recent d observations (see Brown [11]).

Table VII shows the degree to which pooling of forecasts enhances the value of analysts' predictions. We start in Panel A by examining the 3364 raw forecasts produced by the brokers who were included in our study. The first column of forecasting correlation coefficients (headed $\lambda = 0$) shows how the raw forecasting ability of individual brokers varied with the intensity with which stocks were monitored. Those analysts who forecast returns for stocks which had 0 or 1–2 forecasts made by other brokers in the preceding quarter exhibited only modest success; their raw forecasting correlations were in the region of 0.05. Those analysts who forecast returns for stocks which had 3–4 or 5–6 forecasts in the preceding quarter experienced greater success; their raw correlations averaged

[29] This is the sum of a truncated geometric series:

$$\begin{aligned} S &= (1-\lambda)(1+\lambda+\lambda^2+\cdots+\lambda^{\ell-1}) \\ &= (1-\lambda)\{(1+\lambda+\lambda^2+\cdots) - \lambda^{\ell}(1+\lambda+\lambda^2+\cdots)\} \\ &= (1-\lambda)\{(1-\lambda)^{-1} - \lambda^{\ell}(1-\lambda)^{-1}\} \\ &= 1-\lambda^{\ell} \end{aligned}$$

Naturally we assume here that $0 \leq \lambda \leq 1$.

Table VII

Gains from Pooling All Forecasts Made Within the Preceding Quarter

Source of Forecasts	Number of Forecasts in Preceding Quarter	Weight Given to Latest Forecast	Correlation Coefficient for Composite Forecast with weight λ^{ℓ} Given to ℓ-Day Old Forecasts: $\lambda = 0$ (Raw Forecasts)	$\lambda = 0.90$	$\lambda = 0.95$	$\lambda = 0.98$	$\lambda = 0.99$	$\lambda = 1$ (Equal Weighting)	Sample Size
Panel A:	0	1	0.043	0.043	0.043	0.043	0.043	0.043	317
Brokers'	1–2	1	0.053	0.055	0.057	0.061	0.064	0.066	630
Forecasts	3–4	1	0.153	0.204	0.221	0.241	0.246	0.244	740
	5–6	1	0.124	0.134	0.139	0.155	0.154	0.139	832
	7–8	1	0.020	0.052	0.049	0.058	0.064	0.063	476
	≥9	1	0.045	0.029	0.028	0.024	0.029	0.043	369
	≥0	1	0.086	0.104	0.108	0.116	0.118	0.116	3364
Panel B:	0	1	0.058	0.058	0.058	0.058	0.058	0.058	393
All Forecasts	1–2	1	0.024	0.023	0.019	0.013	0.009	0.005	790
	3–4	1	0.114	0.153	0.178	0.200	0.206	0.209	803
	5–6	1	0.147	0.169	0.174	0.191	0.190	0.178	904
	7–8	1	0.026	0.090	0.111	0.145	0.156	0.160	678
	≥9	1	0.059	0.047	0.032	0.034	0.039	0.045	619
	≥0	1	0.077	0.095	0.100	0.106	0.106	0.104	4187
Panel C:									
Internal	≥1	1	0.093	0.089	0.099	0.115	0.115	0.112	522
Forecasts	≥1	0[a]	0.000	0.055	0.070	0.086	0.092	0.097	522

[a] Whereas all other composite forecasts are constructed by giving a weight of unity to current forecasts, here we omit the current (internal) forecast from the computation. This helps us check whether the internal forecasts have added predictive value to the forecasts previously provided by the brokers.

0.14. Forecasts for heavily researched stocks, those with at least seven prior forecasts in the quarter, had a correlation with realised returns of only 0.03. It is tempting to rationalise this observed pattern of forecasting correlations. Perhaps brokers' analysts only undertake detailed, information-based research if other brokers compete in analysing the same stock, while their views on under-researched stocks are of necessity based on fully discounted, public information. If this conjecture were correct, few insights could be gained from brokers on an under-researched stock, while the brokers' efforts would ensure that very heavily researched stocks were in fact efficiently priced. However, this observed pattern of forecasting correlations may very well be sample specific.

The right-hand column of forecasting correlation coefficients in Table VII (labelled $\lambda = 1$) indicates the degree of predictive ability obtained when every forecast is replaced by the simple average of all the quarter's simultaneous and prior forecasts. The brokers' forecasting correlations shown in Panel A are in all cases higher for these equally-weighted consensus forecasts than for the raw forecasts. Evidently the benefits from pooling chronologically dispersed forecasts outweigh any decay in their informational content. Taking all 3364 brokers' forecasts as a whole, the forecasting correlation coefficient is increased from 0.086, the figure for raw forecasts, to 0.116, the figure for equally-weighted consensus forecasts.

While this improvement is quite impressive, we argued above that it is unsatisfactory to give as much weight to an old forecast as to a current forecast produced by an analyst with a similar level of skill. The remainder of Panel A shows how the forecasting correlations vary as the weight λ^{ℓ} given to ℓ-day-old forecasts is increased from $\lambda = 0$ to $\lambda = 1$. When there are zero prior forecasts, the correlation coefficient is naturally unaffected by the value of λ. But it is interesting to note that a value of $\lambda = 0.95$, which gives less than one-hundredth of the weight to a one-quarter-old forecast as it does to a current forecast, is almost as effective as $\lambda = 1$ (equal weighting); the forecasting correlation falls by only 0.008. In fact, a value of $\lambda = 0.99$ appears to maximise the overall predictive value of the brokers' forecasts, although following Armstrong's [4] guidelines for practitioners, a slightly lower figure might be advocated.

Panel A evaluated the gains from pooling brokers' forecasts which, although not in the public domain, might have been available to another equally large investment institution. It appears that, by creating composite forecasts, a correlation coefficient between forecast and realised specific returns could be achieved of about 0.12. This figure is insensitive to the choice of weighting scheme for combining nonsimultaneous forecasts. Panel B turns to evaluating all the forecasts actually recorded within one specific department, focussing on both external (brokers') and internally generated forecasts. The overall impression is similar: (1) stocks which have very few/very many recorded forecasts have lower correlation coefficients, while the middle group is more predictable; and (2) there is a substantial improvement in predictive ability when forecasts are pooled, and this is not highly sensitive to the choice of weighting scheme. On average, the forecasting correlations in Panel B are marginally lower; this reflects the fact that the stocks monitored solely by the internal analysts, for which the forecasting correlations were lower, were pooled with those covered by both parties.

Finally, in Panel C of Table VII we look directly at the extent to which the internal analysts have added predictive value to the forecasts generated by the brokers. We consider all 522 internal forecasts for which there were one or more brokers' forecasts on the same stock. In the first line of Panel C, note that the internal analysts have a raw forecasting correlation (0.093) which is higher than the global figure for the brokers (0.086). The second line of Panel C shows that they have achieved as much predictive success as if they had averaged the brokers' forecasts (using the parameter value proposed above, $\lambda = 0.99$). The internal analysts therefore appear to be efficient in meeting this standard. However, we noted above that the internal forecasts are relatively independent of the brokers' forecasts, so it is possible that by pooling internal and external forecasts, they could even surpass the composite forecast based solely on the brokers' views. In fact, the first line of Panel C suggests that, for a reasonable range of values for λ, an improvement in the internal analysts' forecasting correlation to a figure of 0.11 to 0.12 would have been available if internal and external forecasts had been pooled in this way.

IV. Summary and Conclusions

We have presented an analysis of more than 4000 stock return forecasts for some 200 of the largest UK shares provided by 35 different firms of stockbrokers and the internal analysts of a large UK investment institution. We found that stock prices react rapidly to the informational content of forecasts, over half of the information value being impounded within share prices by the end of the first month after each forecast. By the time the forecasters' horizons had been reached, there was a significant correlation of 0.08 between forecast and realised specific returns. Our evidence is thus comparable with previous studies in the United States and indicates a potentially useful degree of forecasting ability.

Unlike previous research which has focussed on "paper" performance, we have also examined the returns actually achieved by the fund from following these forecasts. The 3000 trades executed by the fund during the sample period were directly motivated by the forecasts. These transactions outperformed the market by 2.2 percent in the year following the trade, representing an abnormal return of some £3 million. We therefore conclude that the apparently favourable performance of the forecasts was not artefactual in nature, but instead, had real commercial value.

Amongst the brokers there were significant differences in their levels of predictive ability, with forecasting correlation coefficients ranging between −0.19 and +0.26 for brokers recording over 50 forecasts. About half of this cross-sectional variance in brokers' forecasting abilities may be attributable to true differences between them, rather than to measurement errors in the estimated correlation coefficients. In spite of this, it proved difficult to predict the rankings of individual brokers from a knowledge of their forecasting correlations in a prior period. For practical purposes, our best assumption is probably that all brokers have about the same degree of forecasting ability.

Although it is often asserted that brokers tend to produce rather similar predictions, we found no evidence to support this. In examining all simultaneous or "near simultaneous" forecasts by different brokers for the same stock, we found that the forecasts had a mean correlation with one another of only 0.08. For most practical purposes, it is therefore reasonable to assume that share price (but not necessarily earnings) forecasts made by different institutions are virtually independent. Consequently, it should be possible to improve predictive accuracy by combining forecasts from different sources, and this is, in fact, the case. The brokers' forecasting correlation was raised to 0.12 when each forecast was pooled with simultaneous or earlier forecasts made in the same forecasting round. Internal forecasts, for which prior brokers' forecasts were available, also had their correlation raised to 0.12. These gains realised from creating a composite forecast were relatively insensitive to the choice of weights used for averaging the raw forecasts.

This study provides the first large-scale prospective study of analysts and stockbrokers' forecasting ability. It provides independent confirmation of the claims of proprietory services that forecasting correlations of around 0.10 may reasonably be expected from "good" security analysts. It also confirms that, in general, the availability of independent views from different brokers conveys considerable benefits to any institution in a position to pool forecasts from a wide selection of sources.

These benefits are not, of course, freely available to all; they may be available only to the largest institutions in exchange for acceptable volumes of fixed-commission business. If these institutions pay "excessive" commissions, there may be little net gain to the immediate recipients of the brokers' wisdom. And if brokers compete to provide clients with valuable information, the benefits of their research may unavoidably be shared with the investment community as a whole, albeit with the virtuous by-product of maintaining the efficiency of the market. Thus, while our research may be interpreted as a contradiction of strong form market efficiency, it is not necessarily at variance with the notion of an efficient market in security analysis.

REFERENCES

1. K. P. Ambachtsheer. "Portfolio Theory and the Security Analyst." *Financial Analysts Journal*, Vol. 28, No. 6 (1972).
2. ———. "Profit Potential in an 'Almost Efficient' Market." *Journal of Portfolio Management*, Vol. 1 (Fall 1974).
3. ——— and J. L. Farrell. "Can Active Management Add Value?" *Financial Analysts Journal*, Vol. 5, No. 6 (1979).
4. J. S. Armstrong. *Long Range Forecasting.* New York: Wiley, 1978.
5. D. Bell and A. Greenhorn (Eds.) *A Guide to the FT Statistics.* Financial Times Publications, 1984.
6. C. Bidwell. "How Good is Institutional Brokerage Research?" *Journal of Portfolio Management*, Vol. 3 (Winter 1977).
7. J. H. Bjerring, J. Lakonishok, and T. Vermaelen. "Stock Prices and Financial Analysts' Recommendations." *Journal of Finance*, Vol. 38, No. 1 (1983).
8. F. Black. "Yes, Virginia, There is Hope: Tests of the Value Line Ranking System," *Financial Analysts Journal*, Vol. 29, No. 5 (1973).
9. J. C. Bogle and J. M. Twardowski. "Institutional Investment Performance Compared: Banks, Investment Counsellors, Insurance Companies, and Mutual Funds." *Financial Analysts Journal*, Vol. 36, No. 1 (1980).

10. R. A. Brealey. *Security Prices in a Competitive Market.* Cambridge, Mass.: M.I.T. Press, 1971.
11. R. G. Brown. *Smoothing, Forecasting and Prediction of Discrete Time Series.* Englewood Cliffs, N.J.: Prentice Hall, 1963.
12. H. L. Cheney. "How Good are Subscription Investment Advisory Services?" *Financial Executive*, Vol. 37 (November 1969).
13. T. D. Coggin and J. W. Hunter. "Problems in Measuring the Quality of Investment Information: The Perils of the Information Coefficient." *Financial Analysts Journal*, Vol. 39, No. 3 (1983).
14. S. Colker. "An Analysis of Security Recommendations by Brokerage Houses." *Quarterly Review of Economics and Business*, Vol. 3 (Summer 1963).
15. Continental Illinois International Investment Corporation. "Ranking of U.K. Investment Analysts." Eighth Annual Survey (1981).
16. T. E. Copeland and D. Mayers. "The Value Line Enigma (1965–1978)." *Journal of Financial Economics*, Vol. 10 (1982).
17. A. Cowles. "Can Stock Market Forecasters Forecast?" *Econometrica*, Vol. 1 (July 1933).
18. ———. "Stock Market Forecasting." *Econometrica*, Vol. 12 (July–October 1944).
19. T. E. Cranshaw. "The Evaluation of Investment Performance." *Journal of Business*, Vol. 50 (1977).
20. L. Y. Dann, D. Mayers, and R. J. Rabb. "Trading Rules, Large Blocks and the Speed of Price Adjustment." *Journal of Financial Economics*, Vol. 4 (1977).
21. P. L. Davies and M. Canes. "Stock Prices and the Publication of Second-Hand Information." *Journal of Business*, Vol. 51 (1978).
22. S. Dawson. "Is the Hong Kong Market Efficient?" *Journal of Portfolio Management*, Vol. 8 (Spring 1982).
23. R. E. Diefenbach. "How Good is Institutional Brokerage Research?" *Financial Analysts Journal*, Vol. 28, No. 6 (1972).
24. E. Dimson and P. R. Marsh (Eds). *Risk Measurement Service.* London Business School, Quarterly, 1979–1983.
25. ——— and ———. "The Impact of the Small Firm Effect on Event Studies and the Performance of Published UK Stock Recommendations." Paper presented at the Annual Meetings of the Western Finance Association, Vancouver, Canada (June 1984).
26. ——— and ———. "The Stability of U.K. Risk Measures and the Problem of Thin Trading." *Journal of Finance*, Vol. 38 (1983).
27. M. M. Dryden. "A Statistical Study of UK Share Prices." *Scottish Journal of Political Economy*, Vol. 17 (1970).
28. E. F. Fama. "The Behaviour of Stock Market Prices." *Journal of Business*, Vol. 38 (1965).
29. ———. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance*, Vol. 25 (1970).
30. ———, L. Fisher, M. C. Jensen, and R. Roll. "The Adjustment of Stock Prices to New Information." *International Economic Review*, Vol. 10 (1969).
31. J. L. Farrell. "A Disciplined Stock Selection Strategy." *Interfaces*, Vol. 12 (October 1982). See also "Addendum." *Interfaces*, Vol. 13 (December 1983).
32. F. Ferber. "Short Run Effects of Stock Market Services on Stock Prices." *Journal of Finance*, Vol. 13 (1958).
33. F. J. Finn. "Internal Evaluation of Institutional Security Analysts' Research." Unpublished Working Paper, University of Queensland, Australia (1983).
34. J. E. Finnerty. "Insiders and Market Efficiency." *Journal of Finance*, Vol. 31 (1976).
35. M. A. Firth. "The Performance of Share Recommendations Made by Investment Analysts and the Effects on Market Efficiency." *Journal of Business Finance*, Vol. 4 (Summer 1972).
36. ———. "The Investment Performance of Unit Trusts in the Period 1965–75." *Journal of Money, Credit and Banking*, Vol. 9 (1977).
37. M. D. Fitzgerald. "A Proposed Characterisation of UK Brokerage Firms and their Effects on Market Prices and Returns." In E. J. Elton and M. J. Gruber (Eds.), *International Capital Markets.* New York: North-Holland, 1975.
38. ———. "Media and Investment Advisory Service Recommendations and Market Efficiency." Salomon Brothers Centre for the Study of Financial Institutions, Working Paper No. 159, New York University, December 1978.

39. G. Foster. "Briloff and the Capital Markets." *Journal of Accounting Research*, Vol. 17 (1979).
40. I. Friend, M. Blume, and J. Crockett. *Mutual Funds and Other Institutional Investors, A New Perspective*. New York: McGraw-Hill, 1970.
41. G. Garmaise. *The Analyst's Analyzer*. Toronto: Garmaise and Associates (September 1982).
42. J. C. Groth, W. G. Llewellen, G. G. Schlarbaum, and R. C. Lease. "An Analysis of Brokerage House Recommendations." *Financial Analysts Journal*, Vol. 35, No. 1 (1979). See also J. C. Groth. "Investor Objectives, Stock Recommendations and Abnormal Returns." *Quarterly Review of Economics and Business*, Vol. 18, No. 2 (1978).
43. M. Hanna. "Testing an Aggressive Investment Strategy Using Value Line Ranks: A Comment." *Journal of Finance*, Vol. 38 (1983).
44. S. D. Hodges and R. A. Brealey. "An Empirical Analysis of the Diagonal Model." In G. P. Szego and K. Shell (Eds.), *Mathematical Methods in Finance*. New York: North-Holland, 1972.
45. ——— and ———. "Portfolio Selection in a Dynamic and Uncertain World." *Financial Analysts Journal*, Vol. 29, No. 2 (1973).
46. C. Holloway. "Testing an Aggressive Investment Strategy Using Value Line Ranks: A Reply." *Journal of Finance*, Vol. 38 (1983).
47. ———. "A Note on Testing an Aggressive Investment Strategy Using Value Line Ranks." *Journal of Finance*, Vol. 36 (1981).
48. P. J. Huber. *Robust Statistics*. New York: Wiley, 1981.
49. M. Jensen. "The Performance of Mutual Funds in the Period 1945–64." *Journal of Finance*, Vol. 23 (1968).
50. J. Johnston. *Econometric Methods*. New York: McGraw-Hill, 1972.
51. M. G. Kendall and A. Stuart. *The Advanced Theory of Statistics*. Charles Griffen, 1969.
52. J. Lintner. "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *Review of Economics and Statistics*, Vol. 47 (1965).
53. R. Litzenberger, K. Ramaswamy, and H. Sosin. "On the CAPM Approach to the Estimation of a Public Utility's Cost of Equity Capital." *Journal of Finance*, Vol. 35 (May 1980).
54. D. E. Logue and D. Tuttle. "Brokerage House Investment Advice." *Financial Review*, Vol. 1 (1973).
55. N. E. Mains. "Risk, the Pricing of Capital Assets, and the Evaluation of Investment Performance: Comment." *Journal of Business*, Vol. 50 (1977).
56. B. G. Malkiel and J. G. Cragg. "Expectations and the Structure of Share Prices." *American Economic Review*, Vol. 60 (1970).
57. Money Which? "Share Tipsters." *Which?* (September 1968).
58. P. N. Newbould and C. W. J. Granger. "Experience with Forecasting Univariate Time Series and the Combination of Forecasts." *Journal of the Royal Statistical Society (Series A)*, Vol. 137 (1974).
59. V. Niederhoffer and M. F. M. Osborne. "Market Making and Reversal on the Stock Exchange." *Journal of the American Statistical Association*, Vol. 61 (1966).
60. H. Oppenheimer and T. Dielman. "On Mergers, Tender Offers, Competitive Bidding and Insider Transactions: Some Direct Evidence." *Journal of Financial Economics*, Vol. 11 (1983).
61. S. P. Pratt and C. W. DeVere. "Relationship Between Insider Trading and Rates of Return for NYSE Common Stocks, 1960–66." In J. H. Lorie and R. A. Brealey (Eds.), *Modern Developments in Investment Management*. Hinsdale, Ill.: Dryden Press, Second Edition, 1978.
62. R. Roll. "A Critique of the Asset Pricing Theory's Tests; Part 1: On Past and Potential Testability of the Theory." *Journal of Financial Economics*, Vol. 4 (1977).
63. B. Rosenberg. "Institutional Investment Management with Multiple Portfolio Managers." University of California, Berkeley, Institute of Business and Economic Research Working Paper 65 (October 1977).
64. A. Rudd and B. Rosenberg. "The Market Model in Investment Management." *Journal of Finance*, Vol. 35 (1980).
65. R. T. Ruff. "Effect of a Selection and Recommendation of a 'Stock of the Month'." *Financial Analysts Journal*, Vol. 19, No. 2 (1963).
66. M. Scholes. "The Market for Securities: Substitution versus Price Pressure and the Effects of Information on Share Prices." *Journal of Business*, Vol. 45 (1972).
67. W. F. Sharpe. "A Simplified Model for Portfolio Analysis." *Management Science*, Vol. 9 (1963).

68. ———. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, Vol. 19 (September 1964).
69. ———. "Comments" on Fama [29] *Journal of Finance*, Vol. 25 (1970).
70. ———. "Decentralised Investment Management." *Journal of Finance*, Vol. 36 (1981).
71. L. Shepperd. "How Good is Investment Advice for Individuals?" *Journal of Portfolio Management*, Vol. 3 (Winter 1977).
72. P. J. Smithers. "London Share Price Database." London Business School (February 1977).
73. ———. "LBS Financial Database Manual." London Business School (March 1980).
74. K. L. Stanley, W. G. Llewellen, and G. G. Schlarbaum. "Further Evidence on the Value of Professional Investment Research." *Journal of Financial Research*, Vol. 4 (Spring 1981).
75. J. D. Stoffels. "Stock Recommendations by Investment Advisory Services: Immediate Effects on Market Pricing." *Financial Analysts Journal*, Vol. 22, No. 2 (1966).
76. H. Theil. *Principles of Econometrics*. New York: Wiley, 1971.
77. J. L. Treynor. "Implementation of Strategy: Execution." In J. L. Maginn and D. L. Tuttle (Eds.), *Managing Investment Portfolios: A Dynamic Process*. New York: Warren, Gorham and Lamont, 1983.
78. ——— and F. Black. "How to use Security Analysis to Improve Portfolio Selection." *Journal of Business*, Vol. 46 (1973).
79. O. A. Vasicek. "A Note on Using Cross-Sectional Information in Bayesian Estimation of Security Betas." *Journal of Finance*, Vol. 28 (1973).
80. C. Ward and A. Saunders. "U.K. Unit Trust Performance 1964–74." *Journal of Business Finance and Accounting*, Vol. 3 (1976).