Journal of Accounting and Economics 1 (1979) 165–185. © North-Holland Publishing Company

THE INFORMATION CONTENT OF FINANCIAL ANALYSTS' FORECASTS OF EARNINGS

Some Evidence on Semi-Strong Inefficiency

Dan GIVOLY and Josef LAKONISHOK

Leon Recanati Graduate School of Business Administration, Tel-Aviv University, Tel-Aviv, Israel

Received October 1978, revised version received August 1979

The paper assesses the information content of revisions in financial analysts' forecasts of earnings by analyzing the relation between the direction of these revisions and stock price behavior. Abnormal returns during the months surrounding the revisions in analysts' forecasts are computed and evaluated. The results strongly indicate that information on revisions in forecasts of earnings per share is valuable to investors. It is also suggested that market reaction to the disclosure of analysts' forecasts is relatively slow and gives rise to potential abnormal returns to investors who act upon this type of publicly available information.

1. Introduction

An extensive body of literature has examined the information content of earnings and the efficiency of the market with respect to their disclosure. The evidence shows that earnings possess information value: a knowledge of the contents of the forthcoming earnings announcement yields an abnormal return [see for example Ball and Brown (1968), Brown and Kennelly (1972) and Foster (1977)]. In addition, the observed stock price reaction to an earnings announcement continues over several weeks or months after the announcement [see, for example, Beaver (1968, 1975), Jones and Litzenberger (1970), Foster (1977), Latané and Jones (1977) and Watts (1978)]. This finding casts some doubt on the validity of the hypothesis of the semi-strong efficiency of the market.

Recently, however, Ball (1978) has suggested that the observed 'inefficiency' might be due to omitted variables or to other specification errors in formulating the equilibrium model of returns. Ball has also recommended procedures which might mitigate the effect of these weaknesses. These procedures were utilized in a later research by Watts (1978) who still found statistically significant abnormal returns after the public release of quarterly reports.

The observed association between stock price movements and the content of earnings announcement might reflect, in part, the continuous efforts by investors to correctly forecast future earnings. The keen interest of investors in future earnings and the weight they assign to them is manifested by, among other things, the number of brokerage houses that produce earnings forecasts on a regular basis and by the attention devoted by the financial community to the issue of the disclosure of management earnings forecasts.

Financial analysts' forecasts (FAF) have recently received an increased attention in accounting literature: Barefield and Comiskey (1975), Basi, Carey and Twark (1976) and Crichfield, Dyckman and Lakonishok (1978) evaluated the accuracy and some other statistical properties of these forecasts. Gonedes, Dopuch and Penman (1976) used them as a proxy for management forecasts in an attempt to evaluate empirically the desirability of mandatory disclosure of the latter.

The purpose of this study is to measure the information content of revisions in FAF. The pattern by which the market changes its expectations in the periods following the revisions in FAF will also be examined and some tentative conclusions concerning the efficiency of the market to revisions in FAF will be drawn.

The methodology of the study incorporates the steps suggested by Ball (1978), thus minimizes the biases in measuring the abnormal returns.

2. Research design

The methodology of the study involves the examination of the association between revisions in financial analysts forecast of earnings per share (EPS) and stock price movements.

The response of stock prices to changes (revisions) in financial analysts' forecasts is measured by the abnormal return in the months surrounding the revision month. Existence of abnormal returns during that period is consistent with the hypothesis that revisions in FAF have information content to investors. Furthermore, if the market is efficient with respect to the release of revisions in FAF, stock price changes associated with that information would coincide with the revisions, i.e., no abnormal returns would be expected after the public release of the revisions.

It should be noted, however, that observed abnormal returns could be due to shortcomings of the equilibrium model used to estimate them. Ball (1978) points to two possible causes for the failure of the model to properly describe the process by which equilibrium expected returns are determined. One is the fact that the variable whose information content is being tested acts as a proxy for variables which determine equilibrium expected returns and which are not included in the equilibrium model. The second cause is errors in measuring the market portfolio. Since abnormal returns are observed in the periods following the release of FAF revisions, reference is also made to the question of whether those abnormal returns are indicative of market inefficiency or are due to deficiencies in the equilibrium model.

Another potential problem in interpreting the results is the fact that stock price movements observed in the period of revisions of FAF might be caused by events other than the release of the revisions. In particular, stock price movements in that period might be triggered by the announcement of quarterly reports. In fact, the data reveal some concentration of revisions in months in which quarterly reports are customarily announced. To assess the bias that could be introduced by attributing the effect of quarterly report announcements to FAF revisions, further analysis was conducted under which revisions which occur during actual announcements months were excluded.

The abnormal returns in the study are computed separately for upward and downward revisions. Numerous variations of this measure are employed, each pertaining to a different set of two parameters: one is the period over which the abnormal returns are compounded (the holding period), and the other is the magnitude of the revision.

Denoting the month the FAF is revised as month 0 and the surrounding months according to their position relative to the revision month (i.e., by -1, $+1$, etc.), the following holding periods were used:

(1) Months -1, 0, 1 and 2, $[-1,2]$ in notation form.
(2) Months 0, 1, and 2, $[0,2]$ in notation form.
(3) Months 1 and 2, $[1,2]$ in notation form.
(4) to (7) are holding periods of one month each, corresponding to -1, 0, 1, and 2 respectively; $[-1]$, $[0]$, $[1]$, and $[2]$ in notation form.

If there is any market response to the revisions in earnings forecasts, it is likely to be more pronounced in the months immediately surrounding the revision (i.e., over the above holding periods). Nevertheless, months which are farther away from the revision were also examined and the findings are reported with the rest of the results.

Eleven size groups of revisions are defined ranging in increments of 1% from 'greater than 0%' to 'greater than 10%'.

Clearly, the abnormal returns during periods $[-1,2]$ and $[-1]$ can be achieved only by investors who know the direction of the forthcoming revision. Strictly speaking, the same is true also for periods $[0,2]$ and $[0]$ since the exact date of the revision is some time during the revision month.

Any abnormal returns observed at the post revision period (holding periods $[1,2]$, $[1]$, and $[2]$) indicate a gradual and slow dissemination of information contained in FAF revisions and imply that investors might (depending on their transaction costs) profitably act upon this information.

To compute the abnormal return, the normal rate of return is defined according to the familiar market model,

$$E(\tilde{R}_{it}|R_{mt}) = \alpha_i + \beta_i R_{mt},$$

where $\tilde{R}_{it}$ denotes the rate of return of security i for period t, α_i and β_i are parameters and R_{mt} is the actual market rate of return for period t.

This study uses monthly rates of return (adjusted for capitalization) and employs monthly compounding. The market rate of return is represented by the equally weighted Fisher Index (composed of all securities listed on the New York Stock Exchange). Since the test period (the period during which abnormal returns are being measured) should be completely divorced from the estimation period, $\hat{\alpha}_i$ and $\hat{\beta}_i$ used for month t are estimated by data of prior years. Specifically, the parameters for a given year are estimated from the four years (48 months) preceding that year. The monthly abnormal returns are measured by the difference

$$\hat{\varepsilon}_{it} = R_{it} - (\hat{\alpha}_i + \hat{\beta}_i R_{mt}),$$

where R_{it} and R_{mt} are the observed values of the respective rates of return and $\hat{\alpha}_i$ and $\hat{\beta}_i$ are estimated from a regression equation.

The total abnormal return for each of the seven types of holding periods is computed for each revision. The results are presented in terms of the abnormal return produced, on average, during any type of holding period over all revisions. The statistical significance of the observed abnormal returns is tested and the test statistics are provided. Breakdown of the results into years and industries is also shown.

To illustrate, the computations relating to holding period $[-1,2]$ are as follows:

Let us denote

r as the revision index for a given security,
t as the chronological month index,
$t(r)$ as the month of revision r,
i as the security index,
N as the number of revisions in the sample,
$\hat{\varepsilon}_{it}$ as the abnormal return for security i in month t.

Let $A_{i,r}[-1,2]$ be the abnormal return of security i during holding period $[-1,2]$ around revision r,

$$A_{i,r}[-1,2] = \prod_{t=t(r)-1}^{t(r)+2} (1+\hat{\varepsilon}_{it}) - 1.$$

Let $A[-1,2]$ be the abnormal return produced, on average, during holding period $[-1,2]$,

$$A[-1,2]=\frac{1}{N}\sum_{i}\sum_{r\in i} A_{i,r}[-1,2].$$

As explained, the results are presented in terms of cross sectional averages of abnormal returns over revisions. The abnormal returns are computed in two ways. One is according to the procedure described above; the other involves a standardization of the abnormal returns which might help the analysis and ease the interpretation of the results. The following discussion clarifies the nature of this standardization.

It is expected that for a random sample the average abnormal returns will not be significantly different from zero. However, in non-random samples cross-section average abnormal returns might be observed. Whenever the information content of a set of signals is investigated, as is done by this study, there is a potential danger of misinterpretation: the abnormal returns could be unduly attributed to the signals. Since this study's sample is not large compared to the number of firms in the population and as well is not representative of all industries (for a description of the sample, see below), some periods might yield abnormal returns that are significantly different from zero without implying anything regarding the value of the information that was produced during those periods. To circumvent this potential distortion, the residuals from the market line, $\hat{\varepsilon}_{it}$, were also standardized with respect to their contemporaneous cross-sectional average, as follows:

$$\hat{\varepsilon}^{s}_{it}=\hat{\varepsilon}_{it}-\bar{\hat{\varepsilon}}_{t},$$

where $\hat{\varepsilon}^{s}_{it}$ is the standardized abnormal return of stock i at month t, and

$$\bar{\hat{\varepsilon}}_{t}=\frac{1}{n}\sum_{i=1}^{n}\hat{\varepsilon}_{it}$$

(n is the number of firms in the sample).

By performing this transformation, the cross-sectional average abnormal return in each month becomes zero, thus making easier the interpretation of the results in the following sense: $\hat{\varepsilon}^{s}_{it}$ becomes a measure of the excess return, during t, of holding share i over a buy-and-hold strategy for the sample.

The abnormal return obtained through this transformation for a given holding period surrounding the revision could be interpreted as the average rate of return over that period from a strategy under which a certain dollar amount of all the shares in the sample is sold short and an equal dollar amount of shares with revisions is bought.[1] If the cash proceeds from the short sale are collected at the time of the transaction no investment outlay is required. The resulting portfolio will practically have no systematic risk.[2]

The standardization procedure might potentially remove industry effects contained in the residual returns. However, as is evident from the results the differences between the standardized and the unstandardized abnormal returns are minor.

3. Sample and data

The need to collect and process manually a large portion of the data posed a limitation on the sample size. In order to enable some inter-industry comparisons, the sample consists of three industries (Standard Industrial Two-Digit Classification): Chemicals and Allied Products (Industry 28), Petroleum Refining and Related Industries (Industry 29), and Transportation Equipment (Industry 37). These industries were chosen since they include a relatively large number of firms which could potentially satisfy the following criteria:

(1) Availability of earnings forecasts for all years in the survey period.
(2) New York Stock Exchange listing. This criterion was introduced since stock price data were derived from the monthly CRSP tape which contains only NYSE stocks.
(3) Fiscal year ending December 31.

The final sample consisted of 49 companies. For each company, actual earnings, EPS forecasts, and monthly stock returns were collected for the eight years 1967 to 1974.

Forecasts of EPS were collected from Standard and Poor's Earnings Forecaster. A weekly publication which first appeared in 1967, the Earnings Forecaster lists in each issue the outstanding EPS forecasts for about 1500 companies. The forecasts are those made by S&P and by about fifty other

[1]This interpretation is offered by Watts (1978, p. 131).

[2]This statement assumes that the β of the sample and the β of the securities with forecast revisions are equal. Since separate portfolios are constructed for securities with upward revisions and for securities with downward revisions, the statement must further assume that the β of the shares with upward revisions is equal to the β of the shares with downward revisions. For these assumptions to hold it is necessary that (1) the frequency of the number of revisions per company is independent of its β and that (2) the sign of the revision over time is not autocorrelated. No significant correlation was found between the β of a security and the number of its revisions. The validity of the second condition has also been confirmed and the results of the statistical tests are presented in the next section. Moreover as is evident from table 1 below, there is practically no difference between the β of the shares with upward revisions and the β of the shares with downward revisions.

security analysts and brokerage houses who agreed to submit their forecasts, upon release, for publication.

Typically, three to five forecasters are actively engaged in forecasting the earnings of a given company. As many as fifteen different forecasts might simultaneously be available for companies with a widely traded stock. This situation poses a difficulty in identifying 'the' forecast of next year's EPS. One solution is to regard the average or the median forecast as the relevant forecast. This approach has some weaknesses: First, changes in the value of this forecast might occur whenever a new forecaster joins or an old forecaster drops from the initial group of forecasters. Secondly, contemporaneous revisions by separate analysts make it difficult to identify the information content of the change in this measure as 'good' or 'bad' news.

Finally, changes in expectations under the above definition will invariably be gradual: The data reveal that revisions of various forecasters do generally move together but that there is also some lag between the first revision and the 'followers', which makes the average revision change slowly and gradually over time.[3] Yet, if the market is efficient, stock prices might be primarily affected by the release of the first revision. Thus, the observed relationship between the average forecast and stock price movements will only partially reflect the true association between the content of FAF revisions and stock prices.

The above shortcomings of the average forecast measure led to the selection of the revisions produced by the most active forecaster for each company (the one with the greatest number of revisions) as the representative of the group of forecasters. The most active forecaster is likely to be the first to respond to new information. At the same time, he is probably the one who specializes in the stock and as such, the most watched and followed by investors; if this is not the case the selection of the most active forecasters might bias downward the association between stock prices and revisions in earnings. The most active forecaster was chosen based on publicly available information at the time the forecast was issued. The forecaster selected each year was the one who was most active in the previous year. The most active forecaster for 1967, the first year for which data is available, was selected based on the experience of the first two months which were excluded from the computation of abnormal returns. For companies with no forecast during the first two months, the first forecaster to make a forecast was selected as the most active.

[3]The significance of the observed gradual movement was tested. For each company the average forecast was computed each month and the signs of the changes in the average forecast between consecutive months was tested for serial dependence using runs test [see Siegel (1956, p. 57)]. In over 90% of the firms a significant positive serial dependence (5% significance level) was found.

To ascertain that the EPS forecast figures and the actual figures were compatible in terms of the dilution definition and the treatment of extraordinary items, actual EPS figures were collected from the Earnings Forecaster. (To assure accuracy, the earnings' figures were compared with those recorded in the Compustat tape.) Adjustments were occasionally called for.[4]

The dates of the actual announcement of the annual and the quarterly reports were collected from the Wall Street Journal. The annual announcement date is the date of the announcement of the audited statements or of the release of the preliminary earnings, whichever is earlier.

4. Empirical results

Tables 1 and 2 describe some general characteristics of FAF. Note that only revisions produced by the most active forecasters are presented and analyzed. Table 1 presents the number of revisions by year, size, and direction. The number of revisions above 10% is 234, of those above 5% is 584 and the number of all (above 0%) revisions is 1,420.

The revisions are quite evenly distributed over the years. The average number of all revisions (above 0%) per year is 178 with a standard deviation of 29.3. There is almost an equal number of upward and downward revisions (693 against 727 for all revisions).

Table 1

Cumulative distribution of the number of financial analysts' earnings forecast revisions and average β by size of revision, direction (up/down) and year.

		Year									
Size (%)	Direction[a]	67	68	69	70	71	72	73	74	All years	Average β
	T	13	18	19	22	27	17	41	77	234	1.01
Above 10	U	2	8	3	7	13	1	36	67	137	1.02
	D	11	10	16	15	14	16	5	10	97	0.99
	T	42	59	55	76	61	49	95	147	584	0.97
Above 5	U	7	27	9	15	27	16	84	124	309	0.98
	D	35	32	46	61	34	33	11	23	275	0.96
Above 0	T	160	195	201	155	146	150	185	228	1420	0.95
(all re-	U	55	97	47	39	55	70	150	180	693	0.94
visions)	D	105	98	154	116	91	80	35	48	727	0.95

[a]'T' denotes total revisions, 'U' denotes upward revisions, and 'D' denotes downward revisions.

[4]In many instances the published forecasts changed not because of a revision but as a result of a change in the definition of EPS (e.g., from fully diluted to primary EPS). In those instances all EPS figures (forecasted and actual) were adjusted to conform with one definition.

Table 2

Relative frequency of revisions in financial analysts' forecasts of earnings by month of release.[a]

	Jan.	Feb.	Mar.	Apr.	May	June	July	Aug.	Sept.	Oct.	Nov.	Dec.
Percentage of all revisions	10.4	8.6	8.9	7.9	10.7	5.7	9.9	7.0	4.7	8.8	10.8	6.6

[a]All companies in the sample have fiscal years ending December 31.

The average β's presented in the last column of table 1 were computed across revisions. The β for a given revision was that applicable, at the time of the revision, to the share for which the revision was made. The results indicate that the β's for large revisions are somewhat greater than those for small revisions. There is practically no difference between β's of upward revisions and β's of downward revisions.

Table 2 shows the frequency of revisions by months. The distribution apparently reflects the pattern of information arrival. Annual and quarterly earnings are undoubtedly a prominent input for FAF; indeed, there is some concentration of revisions in the months in which the annual and quarterly reports are usually released. Still, a significant number of revisions are made in other months. These revisions presumably reflect the arrival of non-accounting information such as GNP, interest rate and inflation rate information, and events specific to the firm.

The main results of the study are presented in the form of abnormal returns in months surrounding the revision month. As pointed out by Ball (1978), an experimental pitfall exists if the variables used to estimate the information content of earnings are highly autocorrelated across time, and are therefore more likely to be associated with variables which explain abnormal return. In such cases the result might well be an overstatement of abnormal returns leading to erroneous conclusions. The earnings variable in this study is the sign (direction) of the revision in the earnings forecast. Therefore, serial dependence between consecutive revisions could be tested by using the run test.

For samples in which the number of positive signs or of negative signs exceeds 20, the number of runs is well approximated by the normal distribution [see Siegel (1956, p. 57)]. Since this condition does not hold for all firms in the sample, sole reliance on the firms' Z values would be inappropriate as a test for randomness in the sample. The distribution of the Z values across firms presented in table 3, indicates the extent of dependence in the sign of the revision. The median Z value in the sample is -0.58 for all revisions and -0.26 for revisions of over 5%, indicating a positive dependence. However, this dependence is small. The last column of table 3 provides the frequency with which the hypothesis of independence is rejected

Table 3

Distribution of the Z values of runs in the sign of financial analysts' earnings forecast revisions (upward = positive, downward = negative) and frequency of rejection of the hypothesis of serial independence between the sign of consecutive revisions.

Revision size	Mean	Percentile 0.10	0.25	0.50	0.75	0.90	Percentage of rejection of the hypothesis of independence, at the 5% level
Revisions above 0%	−0.75	−2.02	−1.57	−0.58	1.15	1.42	13.5
Revisions above 5%	−0.56	−1.55	−1.28	−0.26	0.82	1.09	5.8

at the 5% significance level. The computations show that the null hypothesis of independence cannot be rejected for most firms at the 5% significance level (for cases with less than 20 observations the computations are based on the exact sampling distributions). The rejection rate is 13.5% when all revisions are considered and 5.8% when only revisions of over 5% are considered. It seems reasonable to conclude, therefore, that the sign of the revision is not highly autocorrelated across time.[5] This implies that the earnings variable used here (the sign of the revision) does not proxy for possible omitted variables in the market model and that therefore the reported returns in the periods surrounding the revisions in FAF are not likely to be overstated by the 'proxy-effect' described by Ball. (A possibility still exists, however, that the returns are overstated due to a more complex proxy-effect.)

Tables 4, 5, and 6 present results for revisions over 5%. (The results for other revision-size groups portray basically the same phenomena and will be commented upon later.)

The main results are summarized in table 4. The table presents the abnormal returns which are generated over each of the seven alternative holding periods described above. During each of the periods around upward and downward revisions, average abnormal returns are recorded for the entire sample.

The table is divided into three panels. The first panel presents the

[5]A bias due to model misspecification could still arise: suppose that the sign of the revisions of a particular firm over time is indeed generated by a random process but that the probability of a positive sign to a given revision is 0.99. In this case the sign of the revision might very likely act as a proxy for omitted variables (which are presumably related to the firm's characteristics). To ascertain that our data are free of this potential bias, the null hypothesis of an equal probability for a positive and negative sign was tested for each firm using the binomial test. The null hypothesis was rejected (at the 5% significance level) only in 7.7% of the firms for all revisions (above 0%) and in 9.6% of the firms for revisions exceeding 5%.

standardized abnormal returns for *all* revisions (over 5%). The second panel shows these abnormal returns only for revisions which were released in months with no quarterly earnings announcements. The third panel presents the unstandardized abnormal returns for all revisions. Panels (1) and (3) are based on 584 revisions, while Panel (2) is based on 385 revisions.

If the generating process of the returns is correctly specified, then the existence of abnormal returns in the period surrounding the release of FAF revision might serve as an indication that this event has information content or reflects, at least in part, contemporaneous information conveyed to the market by other sources. If the market is efficient, abnormal returns should be zero in the months after the release of the revision in FAF.

Two main findings emerge from table 4: One is the existence of abnormal returns in the months surrounding FAF revisions. In particular, positive abnormal returns are observed in periods surrounding upward revisions, and negative abnormal returns are recorded in periods surrounding downward revisions. This finding suggests that FAF revisions convey or reflect information.

The second finding is that abnormal returns prevail well after the release of FAF revisions – indicating inefficiency of the market with respect to these revisions.

To test whether each of the average abnormal returns presented in table 4 is significantly different from zero, a t-test was conducted. To compute the mean abnormal returns a cross-section as well as time-series pooling of observations was performed. Since the t-test employed assumes that the observations are independent, the validity of this assumption should be considered. In the present sample, dependence between the observations could take different forms. The most common type of dependence which might exist in the data is perhaps the one created by pooling across firms. Concentration of revisions of a particular industry in a given period might reflect a common underlying industry factor that triggered the revisions and hence might indicate possible dependence of the abnormal returns. To assess the extent of this form of dependence, the distribution of revisions between industries was examined for each month. Using the binomial test it was found that in only 6 out of 81 months with revisions, the number of revisions produced by any one industry was significantly above or below the number of revisions expected for that industry given its relative frequency of revisions over all months (at the 5% significance level). It should be further noted that the total number of revisions in these 6 months was small (28 revisions out of 584). The danger of revision clustering is considerably reduced when only large (over 10%) revisions are considered. Yet, as will be reported below, the main results for the large revisions are similar to those for the other revisions.

Clearly, cross-sectional dependence for an industry could still exist even in

Table 4

Average abnormal returns per holding period by direction of financial analysts' earnings forecast revisions, for revisions over 5% (percentage),[a]

Panel	Direction (upward/ downward)	Monthly holding periods							Monthly return on buy-and-hold
		[−1, 2]	[0, 2]	[1, 2]	[−1]	[0]	[1]	[2]	
(1) All revisions (standardized)	U	4.7	3.1	2.7	1.1	0.9	1.7	1.0	0.6
	D	−3.8	−1.9	−1.0	−1.6	−1.2	−0.4	−0.6	
(2) No earnings announcements (standardized)	U	5.3	4.1	3.4	1.4	0.7	2.5	0.9	0.6
	D	−3.2	−2.6	−1.1	−1.1	−0.9	−0.9	−0.5	
(3) All revisions (unstandardized)	U	4.9	3.3	2.8	1.1	0.9	1.7	1.0	0.6
	D	−3.9	−2.0	−1.0	−1.6	−1.2	−0.4	−0.6	

[a]All values are different from zero for a two-tailed test at the 5% significance level.

the absence of clustering of revisions of that industry in certain periods. The potential dependence exists whenever more than one revision of the same industry is made during one period. Although about 50% of the revisions of every industry occur in months with a single revision of that industry, some cross-sectional dependence of this type could still be present.

Pooling abnormal returns across years might raise a question on the validity of another assumption necessary for the use of the t-test – that concerning the stationarity of this variable's distribution over time. To test whether this assumption holds the eight-year period (1967–1974) was divided into two subperiods of four years each (1967–1970, 1971–1974). The equality of the means of the first and the second subperiods was tested under the assumption of equal variances (the t-test). Then, assuming that the mean abnormal returns in the two subperiods are equal, the equality of the variances of the first and the second subperiod was tested (the F-test). The hypothesis that the variances of the abnormal return are equal in the two subperiods could not be rejected for all holding periods for both upward and downward revisions, at the 5% significance level. The hypothesis of equal means could not be rejected for revisions in both directions and for all holding periods, at the 5% significance level, except for holding period [1] of downward revisions.

The null hypothesis that the average abnormal return is equal to zero was rejected for all values reported in table 4 (at the 5% significance level, two-tail test). Given the potential dependence between observations, the t-tests results may be biased and should therefore be regarded with some caution. (It should be noted, however, that a substantial bias must be present for the results to become insignificant: the t-values for most holding periods are above 3.0.)

Computed, but not presented, are the abnormal returns for holding periods during months farther from the revision month. The abnormal return for months preceding month -1 and following month $+2$ are small and insignificant. (Their absolute values do not exceed 0.2%.)

The last column of table 4 shows the (geometric) average of the monthly return yielded by a buy-and-hold policy for the entire sample over the eight-year period. The buy-and-hold policy assumes that at the beginning of the period, an equal investment is made in the 49 stocks and retained until the end of the period. The comparison between these 'normal' returns and the monthly abnormal returns produced during months -1, 0, 1, and 2, underlines the materiality of the abnormal returns in the months surrounding the public release of FAF. For instance, an investor who acts upon publicly available information can obtain over a two-month period, [1,2], an abnormal return of 2.7% to 3.4%, which represents an increase of at least 225% over his 'normal' monthly return (return of 2.7% against the 'normal' return over two-month period which is 1.2%).

The magnitude of the abnormal returns compared with the normal monthly returns rules out the possibility that the results are due to non-stationarity of the securities' betas or to errors in their measurement. Even a deviation of 100% of the calculated beta from the true beta would not produce abnormal returns of such magnitude.

As explained, Panel (2) presents the results only for revisions which did not occur during months of quarterly announcements. The motivation behind this separate analysis was the possibility that the abnormal returns could be due in part to the effect of quarterly earnings announcements. The comparison between Panels (1) and (2) shows that in most cases the abnormal returns tend to be somewhat higher for revisions which occur in months without earnings announcements. This suggests that actual earnings announcement may not be more powerful than other news signals in inducing analysts to change their future valuation of the company, which implies, in turn, that analysts are quite capable of predicting future earnings. Nonetheless, the difference between the two Panels appears to be statistically insignificant. The null hypothesis that revisions which occur in the months of earnings announcements produce the same abnormal returns as the remaining revisions was tested. For the purpose of the test the mean abnormal returns relating to revisions which occur in the months of quarterly announcements and the mean abnormal returns relating to revisions which occur in the other months were computed for each holding period. The t-test for the difference between the means was used. The null hypothesis could not be rejected for any of the holding periods at the 5% significance level.[6]

Panel (3) exhibits the results for all revisions in terms of the unstandardized abnormal returns. As is evident from the comparison between this Panel and Panel (1), the effect of the standardization of the returns on the results is relatively small.

Since the results in the three Panels are basically similar, the rest of the description of the findings and their analysis will be done in terms of the standardized returns which appear in Panel 1.

Tables 5 and 6 present results which re-inforce the main thrust of the findings reported above. Table 5 presents the abnormal returns which are generated over the seven alternative holding periods for each industry. The results for each industry are similar to the results for the total sample. Most of the abnormal returns are significantly different from zero at the 5% significance level. (Note that the potential bias in the t-test which was discussed earlier might be present also in the results of table 5.)

[6]Here, again, a bias could be introduced due to a potential cross-sectional dependence. However, since such dependence would result in a downward bias in the estimate of the standard deviation (assuming positive cross-sectional dependence), the conclusion of no difference is still correct.

Table 5

Average abnormal returns per holding period by industry and direction of financial analysts' earnings forecast revisions, for revisions over 5% (percentage).

	Direction (upward/ downward)	Holding period [−1,2]	[0,2]	[1,2]	[−1]	[0]	[1]	[2]	Monthly return on buy-and-hold
Chemicals and allied products	U	5.0	2.9	2.1	2.2	0.9	1.0	1.1	0.7
	D	−3.4	−1.6[a]	−1.2	−1.9[a]	−0.4	−0.7	−0.4	
Petroleum refining and related industries	U	2.3	2.3	2.4	0.1[a]	0.1[a]	1.8	0.8	0.8
	D	−4.5	−3.2	−1.9	−1.9	−1.0	−0.7	−1.1	
Transportation equipment	U	12.7	10.3	4.9	2.2	5.4	3.1	1.7	−0.3
	D	−3.1	−2.3	1.2[a]	−0.6	−3.2	0.8[a]	0.3[a]	

[a]All values except those marked by superscript 'a' are different from zero under a two-tailed test at the 5% significance level.

Table 6 provides a breakdown of the results by years. For most years, and during each of the holding periods, the 'correct' sign of the abnormal return is recorded, i.e., positive abnormal returns are observed for upward revisions and negative abnormal returns are observed for downward revisions. For instance, abnormal returns during holding period [1, 2] following upward revisions are positive in all eight years. A null hypothesis that the probability of positive abnormal return is 0.5% can therefore be rejected (using a binomial test) for this combination of holding period and direction of revisions, at the 0.4% significance level $[(1/2)^8]$. Fourteen combinations of holding period and direction of revision (7 × 2) exist. Therefore fourteen such

Table 6

Average abnormal returns per holding period by year and direction of financial analysts' earnings forecast revisions, all revisions over 5% (percentage).

Year	Direction (upward/ downward)	Holding period [−1, 2]	[0, 2]	[1, 2]	[−1]	[0]	[1]	[2]
1967	U	−3.5	−2.0	0.6	−1.4	−2.3	−2.5	3.1
	D	−3.7	−1.1	1.5	−2.3	−2.5	−0.6	2.2
1968	U	3.6	1.7	0.1	1.9	1.7	−0.4	0.5
	D	0.5	−0.3	0.9	0.7	−1.1	2.4	−1.5
1969	U	8.1	4.2	3.7	3.1	0.4	2.9	0.9
	D	−4.7	−3.2	−2.3	−1.6	−1.0	−1.0	−1.2
1970	U	9.4	8.0	6.1	1.9	2.1	2.0	3.8
	D	−2.3	−0.8	−0.2	−1.4	−0.5	0.3	−0.5
1971	U	1.8	2.9	2.6	−1.0	0.3	1.1	1.3
	D	−0.9	−0.8	−0.1	−0.3	−1.0	−0.1	−0.1
1972	U	3.3	2.2	2.1	1.0	0.3	2.4	−0.2
	D	−7.8	−5.8	−3.7	−2.3	−2.0	−1.5	−2.4
1973	U	5.2	3.8	3.0	1.3	0.8	1.8	1.3
	D	−5.4	−3.4	−4.5	−2.1	0.6	−2.3	−1.6
1974	U	5.1	4.0	2.9	1.1	1.1	2.3	0.6
	D	−10.4	5.0	−2.6	−5.0	−1.5	−2.7	0.2
Simple yearly average	Upward	4.1	3.1	2.6	1.0	0.6	1.2	1.4
standard deviation	only	1.4	1.0	0.7	0.5	0.5	0.6	0.5
t-value		2.9	3.1	3.7	2.0	1.2	2.0	2.8

tests could be conducted. For all fourteen combinations a 'correct' sign of the abnormal return is recorded in at least six out the eight years, and for nine of these combinations, a 'correct' sign of the abnormal return is found in at least seven years. This means that the above null hypothesis can be rejected (using a binomial test) for all fourteen combinations at the 15% significance level, (which is the probability of obtaining at least six successes out of eight trials under the null hypothesis) and for nine out of fourteen combinations at the 3.5% significance level (the probability of obtaining at least seven successes).

A parametric test (t-test) could also be applied to test the hypothesis that the average annual abnormal return is zero. The t values for the upward revisions are presented in the table. For six of the seven holding periods, the t values are 2 or above. Furthermore, as evident from the table, the elimination of one year, 1967, from the sample, would increase the average abnormal return considerably. Recall that 1967 was the first publication year of the Earnings Forecaster and the general quality of the data could possibly be inferior to that of later years.

The two tests for table 6 yield results similar to those of tables 4 and 5. Since these tests are less subject to cross-sectional dependence, they allow to draw more affirmative conclusions. The results reported for the total sample, for each of the three industries, for revisions of both directions and for all eight years, are very similar. This makes it possible to generalize the conclusions beyond the framework of a 49-company sample.

All the results reported so far relate to revisions over 5%. Table 7 provides a breakdown of the main findings by revision size. Presented are the average

Table 7

Average abnormal returns for selected holding periods by size and direction of financial analysts' earnings forecast revisions (percentage).

Size of revision	Direction (upward/ downward)	Holding period [−1,2]	[0,2]	[1,2]
0%	U	3.9	2.5	1.9
	D	−3.0	−1.8	−1.1[a]
5%	U	4.7	3.6	2.7
	D	−3.8	−2.2	−1.0
10%	U	5.6	4.6	3.4
	D	−6.1	−4.3	−2.7

[a]All values except those marked by superscript 'a' are different from zero under a two-tailed test at the 5% significance level.

abnormal returns for all revisions (i.e., those exceeding 0%) and for those exceeding 5% and 10%. For all size groups, positive (negative) abnormal returns are observed for upward (downward) revisions. In addition, there is an increase in the absolute magnitude of the abnormal returns as the size of the revision increases. This finding seems consistent with expectations. The reason that the reported results are for 5% rather than for 10% is that for some years there are only a few revisions which exceed 10%.

5. Implications for investment policy

The existence of abnormal returns in periods [1] and [2] can be utilized by investors. An investor might hold a portfolio which consists of companies that have recently had an upward revision of their earnings. To implement this policy, the investor would, at the end of each month, search and add to his portfolio stocks which have just had an upward revision. These stocks will be held for two consecutive months and then sold.

The large number of securities for which earnings forecasts are made and the high frequency of (upward) revisions reasonably assure the investor that at almost any given time a well-diversified portfolio could be constructed from stocks with recent upward revisions.[7]

The trading policy described above, when applied to over-5% revisions, produces in this sample a portfolio with a β of 0.95 which yields an abnormal return of 2.7% over a two-month period, or 17.3% on annual basis (compounded). In comparison, the average annual return from a buy-and-hold policy, representing a β of 0.94 is only 7.4%. However, most investors incur transaction costs and cannot enjoy the full benefit offered by this trading policy. Still, an investor subject to 1% transaction cost in each direction, could achieve an abnormal return of 0.7% over a two-month period, or 4.3% annually. This abnormal return represents a 58% improvement over the performance of a buy-and-hold policy. Similarly, an investor acting upon upward revision of 10% or more could earn an annual abnormal return of 8.7% after transaction costs, which represents a 118% improvement over the return from a buy-and-hold policy.

There are reasons to believe that the trading policy described above might not fully exploit the opportunity faced by investors to gain from the (publicly

[7]Consider the following rough estimates: Assume that the frequency of upward revisions for about 1500 companies followed by the Earnings Forecaster is similar to the frequency of upward revisions for the companies in this sample (309 upward revisions during 8 years, or 0.06 revisions per company per month). An investor would have a monthly selection of 90 companies, on average (0.06×1500). Assuming a binomial distribution, the probability that this selection be reduced in a given month to, say, 70 companies is slim (less than 5%). This wide selection reasonably guarantees the possibility of constructing a well-diversified portfolio at any given time.

available) information on FAF. First, the data consist of monthly rates of return. As a result month 0 was excluded from the holding period; incorporation of daily returns would have further increased the abnormal returns, reflecting the opportunities that exist within month 0.

Second, no use was made of the information on downward revisions. The results indicate that such revisions are (preceded and) followed by periods with negative abnormal return. One way to utilize this information is to go short during these periods. However, the profitability of such a policy is doubtful since its yield must outperform the expected positive return on the market. Another way to use information on downward revisions is to buy shares with upward revisions and to sell short an equal dollar amount of shares with downward revisions. As explained in the Research Design section, such a portfolio could be designed to have a β of 0. If, in addition, the cash proceeds from the short sale are received upon transaction, the strategy does not involve an investment outlay. Acting upon revisions of over 5% for holding period [1,2] would yield, under this strategy, an excess return of about 3.7% (the combined abnormal return from the short and long positions, in table 7). Given the similarity between the β's of upward and downward revisions, the β of the portfolio produced by this strategy would be very close to zero. For most investors, who incur the full transaction cost, such a strategy (which involves two 'round-trips', or at least 4% transaction cost) is unprofitable. However, when applied to revisions of over 10%, the strategy yields an excess return of 6.1% which probably more than offsets the cost of transaction for many investors.

Third, as was shown, concentrating on revisions in months without earnings announcements could increase the abnormal returns even further.

Finally, the *incremental* transaction costs of the proposed trading policies are probably lower than the level implied above. Due to lack of synchronization between income and consumption, many investors find it necessary to temporarily reduce or increase their portfolio size. Transaction costs are therefore incurred anyway in the frequent process of expanding or contracting a portfolio. Selection of the stocks to be added or eliminated guided by the recent occurrence of upward or downward revision could improve the yield performance without incurring additional transaction costs.

6. Discussion and conclusions

The results of the study indicate that FAF revisions convey information to the stock market or reflect variables which determine stock prices. Significant abnormal returns begin to form as early as two months prior to the release of the revision. In an efficient market, no abnormal returns should be observed in the periods following the revision. Yet, the findings show that the market does not respond instantaneously to FAF revisions: significant abnormal returns are observed during the two months following the month of the revision.

Not only are the reported abnormal returns significant, but they are of a considerable magnitude as well. Holding a stock during four months surrounding an upward revision of over 5% results, on average, in an abnormal return of 4.7%, representing a 195% improvement over a buy-and-hold policy. Furthermore, a substantial portion of this abnormal return, 2.7%, is observed in the two months following the revision month. This implies that an investor acting on publicly available information and incurring the full transaction cost could still earn an abnormal return of 0.7% during this two-month period (outperforming a buy-and-hold policy by 58%).

The observed abnormal returns reported in the study might have been the result of a failure of the underlying equilibrium model to properly specify the process which determines equilibrium returns. Ball, who has analyzed this possibility, outlined conditions which reduce the likelihood that the earnings variable is a proxy for omitted variables and under which therefore the bias toward overestimating abnormal returns is minimized. The experimental design of this study meets these conditions. In particular, the variable under investigation (the sign of the revision) is not highly autocorrelated across time and the experiment to test market efficiency is predictive. That is, abnormal returns are measured over periods commencing some time after the revision became public to investors. Furthermore, the fact that abnormal returns tend to disappear two or three months after the revision (for both upward and downward revisions) and the persistence of the findings in each of the three industries and in almost every year) reduce the likelihood that the abnormal returns are due to the 'proxy-effect' described by Ball.

Acknowledgement

The authors are extremely grateful to the referee and to Ross Watts and Jerold Zimmerman for their constructive criticism and helpful suggestions. This research was supported by the Israel Institute of Business Research, Tel-Aviv University.

References

Ball, R., 1978, Anomalies in relationships between securities' yields and yield-surrogates, Journal of Financial Economics, Sept., 103–126.

Ball, R. and P. Brown, 1968, An empirical evaluation of accounting income numbers, Journal of Accounting Research, Autumn, 159–178.

Barefield, R. and E. Comiskey, 1975, The accuracy of analysts forecasts of earnings per share, Journal of Business Research, July, 241–252.

Basi, B., K. Carey and R. Twark, 1976, A comparison of the accuracy of corporate and security analysts' forecasts of earnings, The Accounting Review, April, 244–254.

Basu, S., 1977, Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis, Journal of Finance, June, 663–683.

Beaver, W.H., 1968, The information content of annual earnings announcements, Empirical research in accounting: Selected studies, 1968, Supplement to the Journal of Accounting Research, 67–92.

Beaver, W.H., 1975, The information content of the magnitude of unexpected earnings, Unpublished manuscript (Stanford University, Stanford, CA).

Brown, P. and J.W. Kennelly, 1972, The informational content of quarterly earnings: An extension and some further evidence, Journal of Business, July, 403–415.

Crichfield, T., T. Dyckman and J. Lakonishok, 1978, An evaluation of security analysts' forecasts, The Accounting Review, July, 651–668.

Foster, G., 1977, Quarterly accounting data: Time series properties and predictive ability results, The Accounting Review, Jan., 1–21.

Gonedes, N.J., N. Dopuch and S.M. Penman, 1976, Disclosure rules, information production, and capital market equilibrium: The case of forecast disclosure rules, Journal of Accounting Research, Spring, 89–137.

Jones, C. and R. Litzenberger, 1970, Quarterly earnings reports and intermediate stock price trends, Journal of Finance, March, 143–148.

Latané, H. and C. Jones, 1977, Standardized unexpected earnings – A progress report, Journal of Finance, Dec., 1457–1465.

Siegel, S., 1956, Nonparametric statistics (McGraw-Hill, New York).

Watts, R.L., 1978, Systematic 'abnormal' returns after quarterly earnings announcements, Journal of Financial Economics, June–Sept., 127–150.