
Shrunk earnings predictions are better predictions

MANFRED KEIL, GARY SMITH§ and MARGARET H. SMITH*§

Claremont McKenna College, Claremont, CA 91711, USA and §Pomona College, Claremont, CA 91711, USA

Analysts' earnings forecasts are not perfectly correlated with actual earnings. One statistical consequence is that the most optimistic and most pessimistic forecasts are usually too optimistic and too pessimistic. The forecasts' accuracy can be improved by shrinking them towards the mean. Insufficient appreciation of this statistical principle may partly explain the success of contrarian investment strategies, in particular why stocks with the most optimistic earnings forecasts underperform those with the most pessimistic forecasts.

I. INTRODUCTION

Predicted earnings are not perfectly correlated with actual earnings. One consequence is that the largest predicted earnings growth rates are more likely to be excessively optimistic predictions than to be overly pessimistic. Similarly, the lowest predicted growth rates are more likely to be too pessimistic than too optimistic. If so, the adjustment of stock prices when earnings turn out to be closer to the mean than was predicted may partly explain the success of contrarian strategies.

This paper's objective is to see whether the statistical principle of regression towards the mean can be used to improve analysts' forecasts of earnings growth rates. Section II gives a brief overview of the literature on contrarian strategies. Sections III and IV focus on regression towards the mean and its implications. Section V presents the forecasting model, Sections VI and VII apply it to analysts' forecasts, and Section VIII looks at the performance of portfolios based on earnings predictions.

II. CONTRARIAN STRATEGIES

There is considerable evidence of abnormal returns from 'value' strategies that select stocks with low ratios of price to dividends (O'Higgins and Downes, 1992; McQueen *et al.* 1997), earnings (Nicholson, 1960, 1968; Basu, 1977; Jaffe

et al. 1989), book value (Rosenberg *et al.* 1985; Fama and French, 1992), and cash flow (Chan *et al.* 1991). Believers in market efficiency (such as Chan, 1988; Fama and French, 1992) argue that these abnormal returns must be a compensation for the riskiness of value strategies; sceptics (for example, Lakonishok *et al.* 1994) argue that systematic pricing mistakes (such as the incautious extrapolation of earnings growth rates or a failure to distinguish between a good company and a good stock) create opportunities for contrarian investors.

Other evidence of successful contrarian strategies is provided by Debondt and Thaler (1985 and 1987), who found that portfolios of poorly performing 'loser' stocks outperformed portfolios of previous winners by substantial margins, even though the winner portfolios were riskier. Similarly, Fama and French (1988) and Poterba and Summers (1988) conclude that stock returns are mean-reverting over long horizons. Bauman and Downen (1988), La Porta (1996), and DeChow and Sloan (1997) found negative relationships between predicted earnings growth rates and subsequent stock returns.

III. REGRESSION TOWARD THE MEAN

Regression to the mean is often observed in sequential data, but it actually occurs in a much wider range of contexts (Schmittlein, 1989). For instance, any bivariate

*Corresponding author. E-mail: msmith@pomona.edu.

normal variables with equal variances and a correlation between 0 and 1 exhibit regression to the mean (Maddala, 1992). Suppose, for example, that height and weight are bivariate normal and that each has been standardized to have mean zero and standard deviation one. Because height and weight are imperfectly correlated, the tallest person is usually not the heaviest and the heaviest person is usually not the tallest. Height and weight regress to the mean relative to each other.

The same principle applies to corporate earnings. Suppose that the distributions across firms of 1997 and 1998 earnings growth rates are bivariate normal and have each been standardized to have zero mean and zero standard deviation. The company with the highest growth rate in 1997 usually does not have the highest growth rate in 1998, and vice versa. Or suppose that predicted and actual 1998 earnings growth rates are bivariate normal and have each been standardized. The company with the highest predicted growth rate usually does not have the highest actual growth rate, and vice versa.

The educational testing literature provides a framework for explaining this statistical phenomenon. A person's observed test scores fluctuate about the unobserved latent trait measured by the test. This latent trait (the 'true score') can be interpreted as the expected value of a person's test score, with the difference between a person's test score and true score called the 'error score' (Lord and Novick, 1968). Among a group of test takers, those who score the highest are likely to have had positive error scores: it is possible, but unusual, for someone to score below his or her true score and still have the highest score on a test. Since a score that is high relative to the group is also likely to be high relative to that person's true score, this person's score on another test is likely to regress towards the mean.

This framework is directly applicable to a company's earnings. Actual earnings and predicted earnings both deviate from the probabilistic expected value of a company's earnings ('true earnings'). Actual or predicted earnings that are high relative to a group of companies are also likely to be high relative to that company's true earnings. It is possible, but unlikely, that the most profitable company in 1998 had a negative error score that year, with earnings below its expected value. It is possible, but unlikely, that the company predicted to be the most profitable in 1998 had a negative error score that year, with the prediction below the expected value of earnings.

One can consequently anticipate regression towards the mean when comparing consecutive earnings data or when comparing predicted and actual earnings. Freeman and Tse (1992) and Fama and French (2000) investigate the first question and find that successive earnings regress to the mean, although Fama and French attribute this regression to competitive forces rather than the purely statistical explanation that the error scores of companies with relatively high earnings are more likely to be positive than

negative. Here, we investigate the second question: regression to the mean in comparing predicted and actual earnings.

IV. ARE WE AWARE OF REGRESSION TO THE MEAN?

There is well-established evidence that regression to the mean is a pervasive but subtle statistical principle that is often misunderstood or insufficiently appreciated. Kahneman and Tversky (1973) note that people are often surprised when regression occurs and invent fanciful theories to explain it. If pilots who excel in a training session do not do as well in the next session, it is because the flight instructors praised them for doing well. If bright wives have duller husbands, it is because smart women prefer to marry men who are not as smart.

In the stock market, Keynes (1936) observed that 'day-to-day fluctuations in the profits of existing investments, which are obviously of an ephemeral and nonsignificant character, tend to have an altogether excessive, and even absurd, influence on the market.' Lakonishok, Shliefer, and Vishny (1994) and La Porta (1996) provide formal evidence. The regression to the mean explanation is that investors do not fully appreciate the extent to which profit fluctuations are random variation about true earnings.

Regression towards the mean should not be confused with the fallacious law of averages, which states that an unusual run of successes must be balanced by a run of failures; for example, the incorrect belief that a short-run surplus of heads in coin flips must be balanced by an offsetting future deficit. With corporate earnings, the fallacious law of averages implies that companies with above-average earnings growth rates are due to have below-average growth rates. The correct principle of regression towards the mean implies that those companies with the highest growth rates will, on average, continue having above-average growth rates, but not as high as previously since their high growth rates were more likely affected by good luck than bad.

One regression-towards-the-mean fallacy is to misinterpret the temporary nature of extreme observations as evidence that the standard deviation is shrinking. In the 1930s, Horace Secrist, a statistics professor at the Northwestern University, wrote a book with the provocative title *The Triumph of Mediocrity in Business*. Secrist had found that businesses with exceptional profits in any given year tend to have smaller profits the following year, while firms with very low profits generally do somewhat better the next year. From this evidence he concluded that strong companies were getting weaker, and the weak stronger, so that soon all would be mediocre. The president of the American Statistical Association wrote an enthusiastic review of this book (King 1934); another statistician pointed out that

Secrist had been fooled by regression towards the mean (Hotelling 1933; see also Friedman, 1992). In any given year, companies with exceptional profits relative to other companies are likely to have experienced good fortune.

V. A MODEL OF REGRESSION TOWARDS THE MEAN

Let the analysts' forecast growth rate f and the actual growth rate g both depend on a company's true earnings growth rate μ and the usual independent error terms:

$$f = \mu + \varepsilon \tag{1}$$

and

$$g = \mu + \nu \tag{2}$$

Because one is concerned here with idiosyncratic risks that affect individual companies, rather than macrorisks that affect all companies, all variables are measured as deviations from their respective means across companies.

A firm's true earnings μ is the expected value of its earning growth rate. If ε is independent of μ then the variance of the forecast f across companies is equal to the variance of μ plus the variance of ε , and is therefore larger than the variance of true earnings across companies.

If one estimates the relationship between actual and predicted earnings

$$g = \beta f + \omega$$

the least-squares slope is:

$$b = \frac{\sum fg}{\sum f^2} = \frac{\sum (\mu + \varepsilon)(\mu + \nu)}{\sum (\mu + \varepsilon)^2} \tag{3}$$

Taking probability limits of both sides of Equation 3:

$$\text{plim } b = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\varepsilon^2} \tag{4}$$

The slope coefficient is expected to be less than 1 (which implies regression to the mean) even though the underlying population parameter is 1. The larger the variation in the error score ε relative to the variation in true earnings μ , the further the slope coefficient b is from 1. This is the classical errors-in-variables result.

Thus the least-squares predicted deviation of a firm's growth rate from the mean growth rate is a fraction of the deviation of the analyst forecast from its mean:

$$\hat{g} = bf \tag{5}$$

For a Bayesian interpretation, recognizing that the error term represents the cumulative effects of a great many omitted variables and appealing to the central limit theorem, it is assumed that the error term ε is normally distributed with mean 0 and standard deviation σ_ε . A convenient conjugate prior for a firm's true earnings is provided by a

normal distribution with mean μ_0 and standard deviation σ_0 . The mean of the posterior distribution for μ is partway between the forecast and our prior mean:

$$\hat{\mu} = \frac{\sigma_0^2}{\sigma_0^2 + \sigma_\varepsilon^2} f + \frac{\sigma_\varepsilon^2}{\sigma_0^2 + \sigma_\varepsilon^2} \mu_0 \tag{6}$$

If there was no information about a firm's true earnings other than the analysts' forecast, one might set the prior mean equal to 0 and set the prior standard deviation equal to the standard deviation of true earnings. If so, Equation 6 becomes equivalent to Equations 4 and 5:

$$\begin{aligned} \hat{\mu} &= \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\varepsilon^2} f \\ &= bf \end{aligned}$$

If μ were observable, one could use Equation 2 to make unbiased predictions of actual earnings growth. Because μ is not observed, one would want to use the analysts' forecasts to predict actual earnings growth. If one had the forecast for only one company, one might use that forecast as is. However, looking at this forecast in relation to the forecasts for other companies, one should take into account the statistical argument that those forecasts that are optimistic relative to other companies are probably also optimistic relative to this company's own prospects. It would be unusual if analysts were unduly pessimistic about a company and it still had one of the highest predicted growth rates.

The model suggests that the accuracy of earnings predictions may be enhanced by shrinking the analysts' forecasts towards the mean. By how much? The appropriate shrinkage depends on the correlation between forecast and actual growth rates. If forecast and actual growth rates were perfectly correlated, one would not shrink the forecasts at all. A firm predicted to be one percentage point above the mean will, on average, be one percentage point above the mean, and one can use analyst forecasts with no shrinkage. If forecast and actual growth rates were uncorrelated, the forecast would be useless in predicting earnings. One would shrink each forecast to the mean completely, thereby making no effort to predict which companies will have above- or below-average growth rates.

If the data are standardized to have not only zero means, but also standard deviations of one, the least squares equation for predicting actual earnings growth from the analysts' forecasts is

$$\hat{g} = rf \tag{7}$$

where r is the correlation between g and f . If the forecast growth rate is f standard deviations above its mean, the actual growth rate is predicted to be rf standard deviations above its mean.

The correlation between the current forecast and future actual earnings are not observed. However, the argument suggests that the forecasts f are adjusted for, say 1992,

by calculating the correlation r between forecast and actual earnings growth rates in 1991 and then rf used to predict the 1992 values of the actual earnings growth rate g .

VI. DATA

Earnings growth rates are worked with since the value of a company's stock depends critically on its growth rate. In addition, comparisons of earnings per share across firms are muddled by the fact that earnings per share depend on the arbitrary number of shares outstanding. Split-adjusted earnings growth rates do not depend on the number of shares and are a much more useful metric for comparing firms.

Each year, earnings growth rates are standardized across firms to have a zero mean and unitary standard deviation. The regression-towards-the-mean phenomenon applies to relative values. Thus the concern here is not with predicting the average growth rate, but rather the growth rates of individual companies relative to the mean. The lesson of regression to the mean is that those firms whose growth rates are predicted to be far from the mean will probably have growth rates closer to the mean. Thus standardized earnings growth rates are analysed.

The current-year and next-year analysts' forecasts of earnings per share are used from the beginning of the First Call historical database in 1990 through 2000. First Call tabulates earnings estimates from more than 100 security analysts and its database is updated whenever an analyst issues a new or revised earnings estimate. The median of the analysts' forecasts are used as of 30 April of each year for those companies with December fiscal years that had predicted earnings reported by at least five analysts. The study is restricted to companies with December fiscal years so that all earnings are affected by the same macroeconomic surprises. Companies with December fiscal years are required to file 10-K reports by 31 March; 30 April forecasts are used to ensure that earnings for the preceding fiscal year were available to analysts. Stocks are looked at followed by at least five analysts as these companies tend to be highly visible and closely scrutinized. Any systematic inaccuracies in analyst forecasts cannot be explained away as careless guesses about unimportant companies. The median prediction is used to reduce the influence of outliers.

The First Call database is comprised mainly of the large, financially secure companies that are of interest to the brokerage firms and institutional investors that subscribe to their service. The analysis of short-term forecasts and the restriction to companies followed by at least five analysts further ensures that there is little survivorship bias. Since all forecasts are entered contemporaneously, there are no backfilling issues.

For each company the actual fiscal-year earnings per share A_t and the forecast earnings per share A_{t-k} are

used where $k=1$ or 2 depending on whether the forecast was made in April of the current or previous year. The forecast percentage change f_t in earnings per share is

$$f_t = 100 \left(\frac{F_{t,t-k}}{A_{t-k}} - 1 \right)$$

Companies with nonpositive earnings in the base year are excluded as percentage changes are meaningless. Companies are also excluded whose earnings are predicted to increase or decrease by more than 100%, as these are presumably special situations or rebounds from special situations and the large values might influence the results unduly.

The actual annual percentage change in earnings per share is

$$g_t = 100 \left(\frac{A_t}{A_{t-k}} - 1 \right)$$

In order to avoid problems with the square roots of negative growth rates, we do not annualize the growth rates.

The means and standard deviations of f_t and g_t across firms each year are used to calculate the standardized forecast and actual percentage changes in earnings per share.

VII. RESULTS

The model suggests that analysts' forecasts can be improved by shrinking each forecast towards the mean; that is, predicting a company's g from rf rather than f . To see whether this is so, the forecasts are adjusted for regression to the mean by estimating the correlation between the forecast and actual percentage changes for the year preceding the forecast, always using forecasts with the appropriate horizon (either 1 or 2 years) that had been made for the most recent fiscal year whose results are known at the time of the current forecast. For example, to adjust the analysts' forecasts made in April 1997 of fiscal 1998 earnings, the correlation between actual 1996 earnings and the analysts forecasts made in April 1995 of fiscal 1996 earnings are estimated.

The correlation coefficient is estimated by least median squares, a robust estimation procedure (that is also equivalent to minimizing the median of the absolute values of the residuals). The adjusted forecasts are then calculated by multiplying this correlation coefficient times the standardized forecast percentage change in earnings per share. Forecasting accuracy is measured in two ways. First, the number of firms are tabulated each year for which the adjusted or unadjusted forecasts are closer to the actual values. Second, the mean absolute error (MAE) is calculated for the adjusted and unadjusted forecasts. Tables 1 and 2 show the results.

The number of companies covered increased rapidly in the early 1990s as the First Call database grew. A comparison of the mean forecast and actual growth rates each year shows

Table 1. Current-year forecasts of standardized percentage change in earnings

	Number of companies	Mean No. of analysts	Mean forecast	Mean actual	Correlation	More accurate		MAE	
						f	rf	f	rf
1991	88	8.0	3.15	-41.86	0.68				
1992	407	9.4	14.11	-5.90	0.24	70	337	0.83	0.71
1993	601	9.7	17.15	3.65	0.12	153	448	0.84	0.61
1994	673	10.0	15.84	8.95	0.38	268	405	0.79	0.71
1995	740	9.7	15.87	4.05	0.18	231	509	0.80	0.64
1996	920	9.9	16.87	4.32	0.15	175	745	0.79	0.51
1997	964	9.5	16.74	-1.46	0.10	245	719	0.83	0.58
1998	1039	10.0	16.82	-6.07	0.25	417	622	0.80	0.73
1999	967	10.4	12.92	4.20	0.40	379	588	0.78	0.67
2000	868	10.7	18.28	4.04	0.24	208	660	0.80	0.63
Total	7267	10.0	15.99	1.23	0.19	2146	5033	0.80	0.64

Table 2. Next-year forecasts of standardized percentage change in earnings

	Number of companies	Mean No. of analysts	Mean forecast	Mean actual	Correlation	More accurate		MAE	
						f	rf	f	rf
1992	48	6.8	29.03	-2.53	0.05				
1993	280	7.8	39.22	11.09	0.04				
1994	365	8.4	42.24	30.66	0.41	186	179	0.80	0.79
1995	395	8.6	38.43	31.65	0.38	146	249	0.83	0.72
1996	457	8.5	40.28	15.60	0.05	114	343	0.90	0.72
1997	648	8.9	36.93	13.57	0.25	173	475	0.88	0.75
1998	715	8.7	38.45	2.35	0.10	241	474	0.90	0.77
1999	795	9.5	42.41	-6.14	0.00	195	600	0.92	0.72
2000	741	9.6	37.37	16.42	0.26	205	532	0.84	0.67
Total	4444	8.9	39.20	11.60	0.12	1264	2852	0.87	0.73

that analysts tend to be too optimistic, as several other studies have documented (for example, Dreman and Berry, 1995; Easterwood and Nutt, 1999). Overall, the average median forecast is 15.99% and 39.20% for the current year and next year respectively, compared with average actual values of 1.23% and 11.60%. The question addressed here, however, is not whether the forecasts should be uniformly adjusted downward, but rather whether the forecasts should be compacted by making the relatively optimistic predictions less optimistic and making the relatively pessimistic predictions less pessimistic.

Whether gauged by the number of more accurate predictions or by the mean absolute errors, the adjusted forecasts are more accurate than the unadjusted forecasts in every year. Overall, the adjusted forecasts are more accurate for 70% of the predictions and reduce the mean absolute error by about 20%. One can use the binomial distribution to test the null hypothesis that each method is equally likely to give a more accurate prediction. For the current-year predictions, with the adjusted model more accurate 5033 of 7179 times, the two-sided p value is 2.7×10^{-254} , for the next-year predictions, with the adjusted model more accurate 2852 of 4116 times, the two-sided p value is 4.3×10^{-135} .

Interestingly, if one ignores the analysts' forecasts and use $f=0$ to predict relative growth rates, the forecasts are more accurate than the analysts for 65% of the current-year forecasts and 67% of the next-year forecasts. This result is reminiscent of the higgledy-piggledy research (Little, 1967; Lintner and Glauber, 1967) that indicated that annual earnings growth rates across firms are independent of historical growth rates. The wrinkle here is that analysts' predictions of relative performance, which are presumably based on much more than historical growth rates, are less accurate than the prediction that all firms will grow at the same rate.

On the other hand, the analysts' forecasts do contain useful information in that forecasts that have been shrunk toward zero are more accurate than $f=0$ in 53% of the current-year cases (two-sided $p=0.00000006$) and 52% of the next-year cases (two-sided $p=0.0036$).

VIII. PORTFOLIO RETURNS

If analysts are, on average, excessively optimistic about the companies that they predict will have the largest earnings increases and are overly pessimistic about the com-

Table 3. Five portfolios based on current-year earnings forecasts

	Portfolio				
	1	2	3	4	5
Number of stocks	1321	1326	1306	1319	1319
Average analyst f	-1.245	-0.314	-0.069	0.253	1.378
Average actual g	-0.417	-0.022	0.060	0.079	0.301
Standard deviation	0.044	0.035	0.039	0.043	0.055
Beta coefficient with S&P 500	0.757	0.620	0.753	0.842	1.117
Arithmetic average return	0.177	0.156	0.151	0.146	0.114
Geometric average return	0.170	0.147	0.142	0.134	0.102

Table 4. Five portfolios based on next-year earnings forecasts

	Portfolio				
	1	2	3	4	5
Number of stocks	808	810	808	808	807
Average analyst f	-1.105	-0.442	-0.183	0.210	1.523
Average actual g	-0.367	-0.062	0.045	0.084	0.301
Standard deviation	0.038	0.039	0.042	0.047	0.054
Beta coefficient with S&P 500	0.671	0.772	0.889	0.956	1.023
Arithmetic average return	0.366	0.345	0.319	0.270	0.233
Geometric average return	0.355	0.320	0.297	0.258	0.222

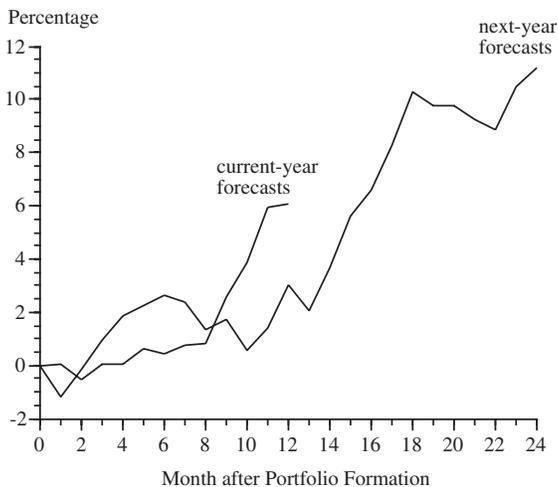


Fig. 1. Cumulative difference between Portfolio 1 and Portfolio 5 returns after formation

panies predicted to have the smallest increases, stock prices may be too high for the former and too low for the latter – mistakes that will be corrected when earnings regress to the mean relative to these forecasts. If so, stocks with relatively small earnings growth predictions may outperform stocks with relatively large predictions.

To test this strategy, the stocks in the First Call database that had monthly returns in the CRSP data base were identified. Portfolios were formed on 30 April of each year based on the analysts' predicted earnings growth rates for the current fiscal year. Portfolio 1 consisted of the 20% of the stocks with the lowest predicted growth rates, Portfolio 5 contained the 20% with the highest growth rates. Equal dollar investments were made in each stock in each portfolio; if the stock was taken private or involved in a merger during the next 12 months, it is assumed that the proceeds were reinvested in the remaining stocks in the portfolio. At the end of 12 months, the portfolio returns were calculated and new portfolios were formed.

Just as one looks at relative earnings, one looks at relative stock returns. One is not trying to predict the direction of the stock market, but rather how a portfolio of stocks with pessimistic earnings forecasts does relative to a portfolio of stocks with optimistic forecasts.

Table 3 shows the results. As expected, the standardized actual growth rates are closer to zero than are the standardized forecast growth rates. In addition, the portfolios with low forecast earnings growth rates outperform the portfolios with high forecast growth rates by substantial and statistically significant margins (ANOVA $p = 0.018$).

Table 4 shows similar results using the analysts' year-ahead predictions. Five portfolios were formed on 30 April of each year based on the analysts' predicted earnings growth rates for the next fiscal year. All portfolios were held for 24 months, until 30 April following the fiscal year for the earnings predictions. Standardized actual earnings growth rates are closer to zero than are the standardized forecast growth rates. The low-forecast portfolios do better than the high-forecast portfolios by substantial and statistically significant margins ($p = 0.00003$).

To explore the timing of these return differentials, Fig. 1 shows the monthly cumulative difference between the performances of Portfolio 1 and 5, for portfolios based on current-year and next-year earnings predictions. The portfolios based on the current-year forecasts show little difference in cumulative returns through the eighth month after formation; the cumulative difference then spurts upward, from 0.9 percentage points at the end of December to 6.1 percentage points at the end of April, the period during which the fiscal-year earnings are reported. For portfolios based on next-year forecasts, the cumulative difference is only 2.0 percentage points 13 months after formation and then rises to 11.2 percentage points by 30 April following the fiscal year for the earnings predictions. An 11.2 percentage-point difference over two years is roughly consistent with a 6.1 percentage-point difference over one year.

It is hard to imagine that this excess return is some kind of risk premium. Growth stocks not only have relatively uncertain cash flows but, because of their long durations, are also more sensitive to changes in required rates of

return. Tables 3 and 4 show that the monthly returns for Portfolio 5 have a higher standard deviation than does Portfolio 1 and more systematic risk (as measured by beta in relation to the S&P 500).

The timing patterns documented in Fig. 1 further indicate that the cumulative return differences between Portfolios 1 and 5 are not a risk premium. In particular, if it were a risk premium why would the return differential be concentrated in the second half of each holding period? A more plausible explanation is that stock prices adjust as investors learn that earnings will be closer to the mean than was predicted.

IX. SUMMARY

The empirical success of contrarian investment strategies may be due to insufficient awareness of regression to the mean—the tendency of extreme observations to exaggerate differences in the latent traits that generate the data. The logic of regression to the mean applies both to successive earnings and to the relationship between predicted and actual earnings. The latter is investigated here. Persuasive evidence is found that relative earnings forecasts are systematically too extreme – too optimistic for companies predicted to do well and too pessimistic for those predicted to do poorly. The accuracy of these forecasts can be improved consistently and substantially by shrinking them towards the mean. It is also found that, for both current-year and next-year forecasts, portfolios of stocks with relatively optimistic earnings predictions underperform portfolios of stocks with relatively pessimistic predictions. Most tellingly, the return differentials are concentrated in the second half of the holding period, when investors are learning that earnings will be closer to the mean than was predicted.

REFERENCES

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*, **32**, 663–82.

Bauman, W. S. and Downen, R. (1988) Growth projections and common stock returns, *Financial Analysts Journal*, **44**, 9–80.

Chan, K. (1988) On the contrarian investment strategy, *Journal of Business*, **61**, 147–63.

Chan, L., Hamao, Y. and Lakonishok, J. (1991) Fundamentals and stock returns in Japan, *Journal of Finance*, **46**, 1739–64.

DeBondt, W. F. M. and Thaler, R. (1985) Does the stock market overreact?, *Journal of Finance*, **40**, 793–805.

DeBondt, W. F. M. and Thaler, R. (1987) Further evidence on investor overreaction and stock market seasonality, *Journal of Finance*, **42**, 557–80.

DeChow, P. and Sloan, R. (1997) Returns to contrarian investment strategies: tests of naive expectations hypotheses, *Journal of Financial Economics*, **43**, 3–27.

Dreman, D. N. and Berry, M. A. (1995) Analysts forecasting errors and their implications for security analysis, *Financial Analysts Journal*, **51**, 30–40.

Fama, E. F. and French, K. R. (1988) Permanent and temporary components of stock prices, *Journal of Political Economy*, **96**, 246–73.

Fama, E. F. and French, K. R. (1992) The cross-section of expected stock returns, *Journal of Finance*, **47**, 427–65.

Fama, E. F. and French, K. R. (2000) Forecasting profitability and earnings, *Journal of Business*, **73**, 161–75.

Freeman, R. and Tse, S. (1992) A nonlinear model of security price responses to unexpected earnings, *Journal of Accounting Research*, **30**, 185–209.

Hotelling, H. (1933) Review of ‘The Triumph of Mediocrity in Business’, *Journal of the American Statistical Association*, **28**, 463–5.

Jaffe, J., Keim, D. and Westerfield, R. (1989) Earning yields, market values, and stock returns, *Journal of Finance*, **44**, 135–48.

Kahneman, D. and Tversky, A. (1973) On the psychology of prediction, *Psychological Review*, **80**, 237–51.

Keynes, J. M. (1936) *The General Theory of Employment, Interest and Money*, Macmillan, New York.

King, W. I. (1934) Review of ‘The Triumph of Mediocrity in Business’, *Journal of Political Economy*, **42**, 398–400.

Lakonishok, J., Shliefel, A. and Vishny, R. W. (1994) Contrarian investment, extrapolation, and risk, *Journal of Finance*, **49**, 1541–78.

La Porta, R. (1996) Expectations and the cross-section of stock returns, *Journal of Finance*, **49**, 1715–42.

Lintner, J. and Glauber, R. (1978) Higgledy, piggledy growth in America, in *Modern Developments in Investment Management*, 2nd edn (Eds) J. Lorie and R. Brealey, Dryden, Hinsdale.

Little, I. M. D. (1966) *Higgledy Piggledy Growth Again*, Basil Blackwell, Oxford.

Lord, F. M. and Novick, M. R. (1968) *Statistical Theory of Mental Test Scores*, Addison-Wesley, Reading, MA.

Maddala, G. S. (1992) *Introduction to Econometrics*, 2nd edn, Macmillan, New York.

McQueen, G., Shields K. and Thorley, S. R. (1997) Does the ‘Dow-10 investment strategy’ beat the Dow statistically and economically?, *Financial Analysts Journal*, **53**, 66–72.

Nicholson, S. F. (1960) Price-earnings ratios, *Financial Analysts Journal*, **16**, 43–5.

Nicholson, S. F. (1968) Price ratios, *Financial Analysts Journal*, **24**, 105–9.

Poterba, J. M. and Summers, L. (1988) Mean reversion in stock returns: evidence and implications, *Journal of Financial Economics*, **22**, 27–59.

O’Higgins, M. B. and Downes, J. (1992) *Beating the Dow*, Harper Collins, New York.

Rosenberg, B., Reid, K. and Lanstein, R. (1985) Persuasive evidence of market inefficiency, *Journal of Portfolio Management*, **11**, 9–17.

Schmittlein, D. C. (1989) Surprising inferences from unsurprising observations: do conditional expectations really regress to the mean?, *The American Statistician*, **43**, 176–83.

Copyright of Applied Financial Economics is the property of Routledge, Ltd. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.