Companies are Seldom as Good or as Bad as They Seem at the Time

Gary Smith, Pomona College

Abstract

Professional forecasts of corporate earnings are not correlated perfectly with actual earnings. Those companies that forecasters are optimistic about usually do not do as well as predicted, while companies given pessimistic predictions typically do better than predicted. Insufficient appreciation of this statistical principle may help explain the success of contrarian investment strategies, in particular why stocks with relatively optimistic earnings forecasts underperform those with pessimistic forecasts.

key words: earnings forecasts, regression to the mean, contrarian strategies

Companies are Seldom as Good or as Bad as They Seem at the Time

Joseph Stiglitz has explored widely and deeply the implications of imperfect information. Sometimes, the information is asymmetric; sometimes, it is just imperfect. The most obvious implications are that even completely rational people may make decisions they later regret and that those who are more informed may take advantage of the less informed. More subtly, and perhaps more importantly, actions often reveal information—which, at times, provides an incentive for information-revealing action and, at other times, inaction. Companies pay dividends to demonstrate that their profits are not illusory, guarantee products to demonstrate confidence in their durability, and are reluctant to reward valuable employees publicly lest they be targeted by other firms.

To focus on the profound implications of relaxing the ingrained, yet wildly implausible, classical assumption that everyone is informed perfectly, Stiglitz typically maintains the similarly implausible classical assumption that everyone makes completely rational decisions based on the information they do possess. The current paper looks at one implication of the fact that stock-market investors are not only imperfectly informed, but also imperfectly rational. The result documents yet another reason to question the efficient-market claim that "given the available information, actual [stock] prices at every point in time represent very good estimates of intrinsic values" (Fama, 1965, p. 34).

There is a useful distinction between investors possessing information and processing information. Stiglitz's work often explores situations in which two sides to a transaction possess different information; for example, when bidding for oil leases, pricing medical insurance, and governing corporations.

Possessing valuable information in the stock market is knowing something about a company that others do not know. Processing information valuably is thinking more clearly about things that are well known; for example, not letting large losses affect one's investment decisions (Smith, Levere, and Kurtzman, 2009).

This paper examines a more subtle possible source of error in processing information. Professional corporate earnings forecasts are widely publicized, yet it is apparently not widely recognized that because of the imperfect correlation between predicted and actual earnings, the most extreme predictions are likely to be too extreme. The subsequent adjustment of stock prices when earnings turn out to be less extreme than predicted may partly explain the success of contrarian investment strategies.

1. Contrarian Strategies

"Value" investment strategies identify stocks that are out of favor, as gauged by their low market prices relative to dividends, earnings, book value, or cash flow (early research includes O'Higgins and Downes, 1992; Basu 1977; Rosenberg, Reid, and Lanstein, 1985; Chan, Hamao, and Lakonishok, 1991). Fama and French (1992) argue that because the stock market is efficient, the high returns from value strategies must be risk premia—even if we cannot pinpoint why value stocks are risky. An obvious alternative to this circular argument is that investor foibles in processing information (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 (Lakonishok, Shliefer, and Vishny, 1994).

The success of contrarian strategies is also documented by the mean-reversion literature. For example, Debondt and Thaler (1985 and 1987) found that portfolios of poorly performing "loser" stocks outperformed portfolios of previous winners by substantial margins, even though the winner portfolios were riskier.

One way to tie together the empirical success of value strategies and the mean reversion of stock prices is the possibility that investors have an insufficient appreciation of the statistical principle of regression to the mean, both with respect to fluctuations in actual earnings and, also, in the imperfect relationship between predicted and actual earnings.

2. Regression to the Mean

Regression to the mean commonly occurs when observed data are an imperfect measure of an underlying trait. For example, because observed heights are an imperfect estimate of the genetic influences that parents pass on to their children, parents whose heights are far from the mean tend to have children whose heights are closer to the mean (Galton, 1889). Similarly, because test scores are an imperfect measure of mastery, students whose scores are far from the mean on one test tend to score closer to the mean on a second test of the same material (Lord and Novick, 1968).

In business, Secrist (1993) found that the most and least successful companies (as measured, for example, by the ratio of profits to assets) tended to become more nearly average as time passed. Secrist (p. 24) speculated that firms with "superior judgment, merchandizing sense, and honesty" were undermined by competition from "unfit" firms, causing businesses to converge to mediocrity. The initial reviews (for example, Riegel 1933, Elder 1934, King 1934) of Secrist's book, *The Triumph of Mediocrity in Business*, were glowing; then Hotelling (1933) pointed out that Secrist had been fooled by regression toward the mean.

In any given year, the most successful companies are more likely to have had good luck than bad, and to have done well not only relative to other companies, but also relative to their own "ability." The opposite is true of the least successful companies. This is why the subsequent performance of the top and bottom companies is usually closer to the average company. At the same time, their places at the extremes are taken by other companies experiencing fortune or misfortune. These up-and-down fluctuations are part of the natural ebb and flow of life and do not mean that all companies will soon be mediocre. As Hotelling put it, Secrist's calculations "really prove nothing more than that the ratios in question have a tendency to wander about." (p. 164)

3. Are We Aware of Regression to the Mean?

Secrist is not alone. An anonymous reviewer for a paper (Schall and Smith, 2000) on regression to the mean in baseball wrote that,

> There are few statistical facts more interesting than regression to the mean for two reasons. First, people encounter it almost every day of their lives. Second, almost nobody understands it. The coupling of these two reasons makes regression to the mean one of the most fundamental sources of error in human judgment.

People often do not anticipate statistical regression and, when it occurs, they search for causal explanations. 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 (Kahneman and Tversky, 1973). If students who do poorly on one test do better subsequently, it was because they were given special attention (Thorndike, 1942).

Mean reversion in stock returns may be due to this insufficient appreciation of regression to the mean in earnings, cash flow, and other fundamental determinants of stock prices. Keynes (1936, p. 138) 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."

For simplicity, suppose that a company's return on assets (ROA) fluctuates randomly about a constant expected value. The firm's ROA will consequently exhibit regression to the mean (extreme values tend to be followed by values that are closer to the mean) and the *change* in ROA will exhibit mean reversion (ROA increases tend to be followed by ROE decreases, and vice versa).

Now suppose that investors do not fully appreciate regression to the mean. They see an increase in earnings and bid up the price of the stock. When earnings regress to the mean, the stock price takes a hit. Thus, a price increase (and high return) tends to be followed by a price decrease (and low return). This argument that regression to the mean in earnings causes a mean reversion is prices explains the conclusion of Lakonishok, Shliefer, and Vishny (1994):

> First, a variety of investment strategies that involve buying out-of-favor (value) stocks have outperformed glamour strategies over the April 1968 to April 1990 period. Second, a likely reason that these value strategies have worked so well relative to the glamour strategies is the fact that the actual future growth rates of earnings, cash flow, [and sales] etc. of glamour stocks relative to value stocks turned out to be much lower than they were in the past, or as the multiples on those stocks indicate the market expected them to be.

We can apply the same logic to earnings forecasts if we think of the forecasts as predictors of actual earnings in the same way that parents' heights are predictors of their children's heights, test scores are predictors of future test scores, and current earnings are predictors of future earnings. The most optimistic predictions are more likely to be overly optimistic than to be excessively pessimistic; so, the companies with the most optimistic forecasts probably won't do as well as predicted—nor are the companies with the most pessimistic forecasts likely to do as poorly as predicted.

4. A Model of Regression Towards the Mean

Let μ_i be the statistical expected value of company *i*'s earnings growth rate and Y_i be the median of the expert forecasts of this growth rate. Even if the forecasts are unbiased with Y_i equal to μ_i plus a random error term,

$$Y_i = \mu_i + \varepsilon_i \tag{1}$$

the variance of the forecasts across firms is larger than the variance of the actual expected values. Extreme forecasts tend to be too extreme.

If μ_i were known, we could use it to make unbiased predictions of actual earnings growth. Because μ_i is not observable, we use the analysts' forecasts. If we had the forecast for only one company, we might use that forecast as is. However, looking at the forecast for one company in relation to the forecasts for other companies, we 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.

This line of reasoning suggests that the accuracy of earnings predictions for any one company may be enhanced by shrinking the prediction toward the average prediction for all companies. By how much?

In the educational testing literature, μ_i is the expected value of a person's test score (fittingly labeled this person's "ability") and Y_i is this person's observed score on a single test that is intended to measure this ability.

Truman Kelley (1947) showed that the optimal estimate of a person's ability is a weighted average of this person's score and the average score of the group the person came from:

$$\hat{\mu}_i = rY_i + (1-r)\overline{Y} \tag{2}$$

The term r is the test's reliability, which measures the extent to which performances are consistent. If a group of students take two comparable tests, reliability is the correlation between their scores on these two tests.

If test scores were completely random, like guessing the answers to questions written in a language the test-takers don't understand, the reliability would be zero, and the best estimate of a person's ability is the average score of the group. At the other extreme, a perfectly reliable test would be one where some people do better than others, but any single person gets the same score, test after test. Now, the best estimate of a person's ability is the test score, regardless of the group score.

In between these extremes, Kelley's equation is a weighted average of the individual score and the group score, with the weight given to the individual score increasing as the test's reliability increases. Kelley's equation not only recognizes that performances tend to regress to the mean, but it tells us how much regression to expect.

Although Kelley derived the equation named after him by using the thoroughly conventional approach to statistics that prevailed in the 1940s (and for many decades afterward), Kelley's equation can also be derived using Bayesian reasoning. If the error term in Equation 1 is normally distributed and our prior distribution for a person's ability is normally distributed with a mean equal to the average score, the mean of the posterior distribution for ability is given by Equation 2.

Applying Kelley's equation to corporate earnings forecasts, the optimal estimate of the expected value of a company's earnings growth is a weighted average of the analysts' forecast for this company and the average forecast for all companies, with the weights depending on the correlation between actual and predicted earnings growth across companies.

If forecast and actual growth rates were perfectly correlated, we would not shrink the forecasts at all. If forecast and actual growth rates were uncorrelated, the forecasts would be useless in predicting earnings and we would shrink each forecast to the mean completely, making no effort to predict which companies will have above-average or below-average growth rates.

Analysts forecasts need not be unbiased. Several studies have documented that analysts tend to be too optimistic (for example, Dreman and Berry, 1995; Easterwood and Nutt, 1999). I will consequently work with standardized actual and predicted growth rates, each normalized to have a mean of zero and standard deviation of one.

Kelley's equation now simplifies to

$$\hat{Z}_i = rZ_i \tag{3}$$

where Z_i is the standardized expert forecast and $\hat{Z}_i$ is the adjusted forecast, taking regression to the mean into account.

5. Data

Some companies have higher earnings per share than other companies simply because they have fewer shares outstanding. I consequently analyze the predicted and actual growth rates of earnings per share, both normalized across firms each year to have a mean of zero and standard deviation of one. Regression to the mean applies to growth rates across firms, and implies that firms whose growth rates are predicted to be far from the mean will probably have growth rates closer to the mean, regardless of whether the average growth rate is high or low.

All forecasts were taken from the IBES data base, which records the forecast, the forecaster, and the time and date when the forecast is announced. Since all forecasts are entered contemporaneously, there are no backfilling issues. The IBES data base maintains a record of every forecast made by each forecaster so that it is possible to identify a forecaster's most recent forecast as of any specified date.

I used the median of the analysts' forecasts for all companies with December fiscal years so that all earnings results are buffeted by the same macroeconomic factors. Companies with December fiscal years are required to file 10-K reports by March 31; I used the most recent forecasts as of April 30 to ensure that earnings for the preceding fiscal year were available to analysts. I looked at both the current-year and year-ahead forecasts. For April 30, 2005, for example, the current-year forecasts are for fiscal 2005 and the year-ahead forecasts are for fiscal 2006.

I restricted my analysis to companies with forecasts from at least 25 analysts to ensure that these companies are prominent, highly visible, and closely scrutinized. Any systematic inaccuracies in analyst forecasts cannot be explained away as careless guesses about unimportant companies. I used the median prediction to reduce the influence of outliers. I also excluded companies that did not have stock return data in the CRSP data base because I wanted to track the stock performance of the companies I analyzed.

I excluded situations in which percentage changes were meaningless because base-year earnings were negative. I also excluded situations in which earnings are predicted to increase or decrease by more than 100 percent, as these are presumably special situations or rebounds from special situations and the large values might influence the results unduly.

These various restrictions limited the years studied to 1992 through 2014, with an average of 54 companies per year for the current-year forecasts and 38 companies per year for the year-ahead forecasts.

6. Results

Kelley's equation was used to adjust the forecasts for regression to the mean. The correlation between forecast and actual earnings was estimated using the most recent data available at the time of the forecast. For example, for the current-year forecasts on April 30, 2005, I estimated the correlation between current-year forecasts on April 30, 2004, and actual 2004 earnings—both of which were available on April 30, 2005. For the year-ahead forecasts on April 30, 2005, I estimated the correlation between year-ahead forecasts on April 30, 2003, and actual 2004 earnings.

Forecasting accuracy was measured in two ways: by the number of adjusted or unadjusted forecasts that were closer to the actual values, and by the mean absolute error (MAE) for the adjusted and unadjusted forecasts. Tables 1 and 2 show the results.

Overall, the average median forecast earnings growth is 10.57 percent and 17.09 percent for the current year and year ahead respectively, compared with average actual values of 9.86 percent and 12.99 percent. 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 the relatively pessimistic predictions less pessimistic.

It is striking that, over the entire period, the correlation between predicted and actual growth rates across firms was 0.57 for the current-year forecasts and 0.05 for the year-ahead forecasts. Remember, the current-year prediction are as of April 30, four months into the fiscal year being forecast. The year-ahead forecasts are evidently little better than dart throws.

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. Overall, the adjusted forecasts are more accurate for 58 percent of the current-year predictions and 69 percent of the year-ahead forecasts. The shrunken forecasts reduce the mean absolute errors by about 10 percent and 35 percent, respectively. We can use the binomial distribution to test the null hypothesis that each method is equally likely to give a more accurate prediction. The two-sided p value is 2.0×10^{-9} for the current-year predictions and 4.7×10^{-29} for the year-ahead predictions.

7. Portfolio Returns

The previous section showed that analysts are, on average, excessively optimistic about the companies that they predict will have relatively large earnings increases and overly pessimistic about the companies predicted to have relatively small increases. If investors pay attention to analysts (or make similar predictions themselves), stock prices may be too high for companies with relatively optimistic forecasts and too low for the companies with relatively pessimistic forecasts—mistakes that will be corrected when earnings regress to the mean relative to these forecasts. If so, stocks with relatively pessimistic earnings predictions may outperform stocks with relatively optimistic predictions.

To test this conjecture, four portfolios were formed on April 30 of each year based on the analysts' predicted earnings growth for the current fiscal year. Portfolio 1 consisted of the quartile of stocks with the highest predicted growth, Portfolio 4 contained the quartile with the lowest growth. Equal dollar amounts were invested in each stock. If the stock was taken private or involved in a merger during the next 12 months, 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.

For the year-ahead predictions, four portfolios were formed on April 30 of each year based on the forecast quartiles, but now the stock returns were tracked over the next 24 months. I also looked at two annual strategies. In the first strategy, the portfolios were formed on April 30 of each year and held for one year. They were consequently liquidated a year ahead of the actual earnings announcement. In the second strategy, the portfolios were formed one year after the prediction and held for the one year.

For example, the year-ahead forecasts on April 30, 2005, are for fiscal year 2006. In the first annual strategy, the portfolio is formed on April 30, 2005, based on the year-ahead forecasts at that time, and held until April 30, 2006, at which time the proceeds are invested based on the year-ahead forecasts on April 30, 2006, which are for fiscal year 2007. In the second annual strategy, the portfolio is formed on April 30, 2006, based on the year-ahead forecasts made on April 30, 2005, and held until April 30, 2007, at which time the proceeds are invested based on the year-ahead forecasts on April 30, 2006. These two strategies are easily implemented in practice and can help determine whether performance differences based on year-ahead forecasts occur during the first or second 12 months after the forecasts are made.

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

Table 3 shows the results for the current-year forecasts. As expected, the actual growth rates are closer to zero than are the predicted growth rates. In addition, the portfolio with lowest forecast earnings growth rates outperformed the portfolio with highest forecast growth rates, on average, by about 2 percentage points a year, though the matched-pair two-sided p value is 0.60. The difference in geometric returns (12.05% versus 8.56%), means that over the complete period 1992 through 2014, the most-pessimistic portfolio would have grown to more than twice the value of the most-optimistic portfolio. A $10,000 investment in the most-optimistic stocks would have grown to $66,000, while a $10,000 investment in the most-pessimistic stocks would have grown to $137,000.

Table 4 shows the results for the year-ahead forecasts. The actual growth rates are much closer to zero than are the predicted growth rates, reflecting the near-zero correlation between predicted and actual growth rates. In addition, the average two-year returns are substantially higher for the most-pessimistic portfolio than for the most-optimistic portfolio, and the matched-pair two-sided p value is 0.0264.

Table 5 breaks the year-ahead performance into the first and second years following the predictions, as explained above. The most-pessimistic portfolios do better than the most optimistic portfolios in both years, though the differences are more pronounced during the second year. The matched-pair two-sided p values are 0.23 for the first-year portfolios and 0.03

for the second-year portfolios. For the first-year strategy, a $10,000 investment in the most-optimistic stocks would have grown to $46,000, while a $10,000 investment in the most-pessimistic stocks would have grown to $179,000. For the second-year strategy, a $10,000 investment in the most-optimistic stocks would have grown to $40,000, while a $10,000 investment in the most-pessimistic stocks would have grown to $221,000.

8. A Risk Premium?

It is unlikely that the superior performance of the pessimistic portfolios reflects some kind of risk premium. Growth stocks are not only likely to have relatively uncertain cash flows but, because of their long durations, are also more sensitive to changes in required rates of return. Tables 3, 4, and 5 show that the annual returns for the most optimistic portfolios have higher standard deviations and more systematic risk (as measured by beta in relation to the S&P 500).

The timing patterns documented in Table 5 further indicate that the cumulative return differences between Portfolios 1 and 4 are not a risk premium. If it were a risk premium why would the return differential be larger in the second year of the 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.

Nonetheless, I gauged the riskiness of the portfolios by estimating the Fama-French (1993) three-factor model using daily percentage-return data from Ken French's web site (2015).

$$R = \alpha + \beta_1 MKT + \beta_2 SMB + \beta_3 HML + \varepsilon$$

where

R = return on portfolio minus the return on Treasury bills

MKT = the value-weighted return on all NYSE, AMEX, and NASDAQ stocks (from

CRSP) minus the return on Treasury bills

SMB = average return on three small-stock portfolios minus the average return on three large-stock portfolios (size factor)

HML = the average return on two value portfolios minus the average return on two growth portfolios (value factor)

This model reflects the historical evidence that stock returns are affected by common macro factors; small stocks tend to outperform large stocks (Banz 1981; Reinganum 1981); and value stocks tend to outperform growth stocks (Rosenberg , Reid and Lanstein, 1985).

Chan (1988) and Fama and French (1992) argue that any systematic differences in returns attributable to these factors must represent risks that matter to investors who must be rewarded for bearing these risks. Others, including Lakonishok, Shliefer, and Vishny (1994), interpret the above-average returns as evidence of investor errors—for example, value stocks generally outperform growth stocks because investors overreact to news, causing stocks to be mispriced temporarily. Either way, the question here is whether the relatively strong performance of the pessimistic portfolios can be explained by these three factors.

The estimates are shown in Table 6. The alphas are negative for the most optimistic portfolios and positive for the most pessimistic portfolios, although the only value that is statistically significant at the 5% level is for the year-ahead pessimistic portfolio. A daily excess return of 0.02% is roughly 5% on an annual basis.

There is no consistent pattern for the size (SMB) factor. The coefficients of the value factor tend to be negative for the optimistic portfolios and positive for the pessimistic portfolios, indicating (not suprisingly) that analysts tend to be more optimistic about growth stocks than value stocks.

9. Summary

Warren Buffett's (2008) succinct summary of a contrarian investment strategy is, "Be fearful when others are greedy and be greedy when others are fearful." His aphorism is generally thought to apply to market bubbles and panics. The evidence presented here suggests a similar mindset for individual stocks, even during normal times: "Companies are seldom as good or as bad as they seem at the time."

The evidence is persuasive that 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 toward the mean. These are not fly-by-night companies. They are prominent, widely followed, and closely scrutinized firms.

In addition, portfolios of stocks with relatively optimistic earnings predictions underperform portfolios of stocks with relatively pessimistic predictions. Most tellingly, the return differentials for the year-ahead portfolios are concentrated in the second year of the holding period, when investors are learning that earnings will be closer to the mean than was predicted.

It is not just corporate earnings and it is not just stock prices. Whenever there is imperfect information—whenever there is uncertainty—people may make flawed decisions based on an insufficient appreciations of regression to the mean.

References

Banz, R. 1981, The relationship between return and market value of common stocks, Journal of *Financial Economics*, 9 (1), 3–18.

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 (3), 663–682.

Buffett, W. 2008, Buy American. I am, *The New York Times*, October 16, 2008.

Chan, K. 1988, On the contrarian investment strategy, *Journal of Business*, 61 (2), 147–163.

Chan, L., Hamao, Y. and J. Lakonishok, 1991, Fundamentals and stock returns in Japan, *Journal of Finance*, 46 (5), 1739–1764.

DeBondt, W. F. M. and R. Thaler, 1985, Does the stock market overreact?, *Journal of Finance*, 40 (3), 793–805.

DeBondt, W. F. M. and R. Thaler, 1987, Further evidence on investor overreaction and stock market seasonality, *Journal of Finance*, 42 (3), 557–580.

Dreman, D. N. and M. A. Berry, 1995, Analysts forecasting errors and their implications for security analysis, *Financial Analysts Journal*, 51 (3), 30–40.

Elder, R. F. 1934, Review of The Triumph of Mediocrity in Business by Horace Secrist. *American Economic Review*, 24 (1), 121–122.

Fama, E. F. 1965. The behavior of stock market prices, *Journal of Business*, 38 (1), 34-105.

Fama, E. F. and K. R. French, 1992, The cross-section of expected stock returns, *Journal of Finance*, 47 (2), 427–465.

Fama, E. F. and K. R. French, 1993, Common risk factors in the returns on bonds and stocks, *Journal of Financial Economics*, 33 (1), 3-56.

French, K. R., Data Library; Retrieved July 14, 2015, from http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Galton, F. 1889, *Natural Inheritance*, New York: Macmillan.

Hotelling, H. 1933, Review of The Triumph of Mediocrity in Business, *Journal of the American Statistical Association*, 28 (184), 463–465.

Kahneman, D. and A. Tversky, 1973, On the psychology of prediction, *Psychological Review*, 80 (4), 237–251.

Kelley, T. L. 1947. *Fundamentals of Statistics*, Cambridge, MA: Harvard University.

Keynes, J. M. 1936, *The General Theory of Employment, Interest and Money,* New York: Macmillan, 153–154.

King, W. I. 1934, Review of The Triumph of Mediocrity in Business by Secrist H., *Journal of Political Economy*, 42 (3), 398–400.

Lakonishok, J., Shliefer, A. and R. W. Vishny, 1994, Contrarian investment, extrapolation, and risk, *Journal of Finance*, 49 (5), 1541–1578.

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

O'Higgins, M. B. and J. Downes, 1992, *Beating the Dow*, New York: Harper Collins.

Reinganum, M. 1981, Misspecification of capital asset pricing: Empirical anomalies based on earnings' yields and market values, *Journal of Financial Economics*, 9 (1), 19–46.

Riegel, R. 1933, Review of The Triumph of Mediocrity in Business by Horace Secrist. *Annals of the American Academy of Political and Social Science*, 170 (Nov.), 178–179.

Rosenberg, B., Reid, K. and R. Lanstein, 1985, Persuasive evidence of market inefficiency, *Journal of Portfolio Management*, 11 (3), 9–16.

Schall, T. and G. Smith, 2000, Baseball players regress toward the mean, *The American Statistician*, 54 (4), 231-235.

Secrist, H. 1933. *The Triumph of Mediocrity in Business*, Evanston, Ill: Northwestern University.

Smith, G., Levere, M. and R. Kurtzman, 2009, Poker player behavior after big wins and big losses, *Management Science*, 55 (9), 1547-1555.

Thorndike, R. L. 1942. Regression fallacies in the matched group experiment, *Psychometrika*, 7 (2), 85-102.

Table 1 Current-Year Forecasts of Standardized Percentage Changes in Earnings, 1992-2014

Mean No. of Companies	Mean No. of Analysts	Mean Forecast, %	Mean Actual, %	Correlation	More Accurate		MAE	
					Z	rZ	Z	rZ
53.9	29.9	10.57	9.86	0.57	516	724	0.48	0.43

Table 2 Year-Ahead Forecasts of Standardized Percentage Changes in Earnings

Mean No. of Companies	Mean No. of Analysts	Mean Forecast, %	Mean Actual, %	Correlation	More Accurate		MAE	
					Z	rZ	Z	rZ
38.1	29.2	17.09	12.99	0.05	276	601	0.90	0.58

Table 3 Four Portfolios Based on Current-Year Earnings Forecasts

	Portfolio			
	1	2	3	4
Average Number of Stocks	13.39	13.26	13.78	13.48
Average Forecast Z	1.20	0.23	-0.24	-1.17
Average Actual Z	0.82	0.13	-0.17	-0.77
Annual Return				
Mean, %	11.84	11.39	14.27	14.00
Standard Deviation, %	27.30	24.53	19.93	19.89
Beta Coefficient	0.94	1.25	0.91	0.88
Geometric Mean, %	8.56	8.58	12.46	12.05

Table 4 Four Portfolios Based on Year-Ahead Earnings Forecasts

	Portfolio			
	1	2	3	4
Average Number of Stocks	9.35	9.48	9.65	9.65
Average Forecast, Z	1.33	0.07	-0.34	-1.02
Average Actual, Z	0.10	0.03	-0.05	-0.08
Two-Year Return				
Mean, %	23.57	30.48	32.27	35.99
Standard Deviation, %	41.85	40.85	35.23	39.21
Beta Coefficient	1.18	1.13	1.06	1.01

Table 5 Year-Ahead Portfolios, first-year and second-year annual returns

	Portfolio			
	1	2	3	4
First Year				
Mean, %	11.09	14.16	14.48	15.25
Standard Deviation, %	28.33	24.34	18.54	20.44
Beta Coefficient	1.36	0.99	0.81	0.89
Geometric Mean, %	6.90	11.56	13.01	13.37
Second Year				
Mean, %	9.63	14.83	15.28	17.02
Standard Deviation, %	25.82	22.58	23.30	20.97
Beta Coefficient	1.01	1.06	1.01	0.91
Geometric Mean, %	7.87	12.26	12.05	14.80

Table 6 Fama-French Factor Model

		Intercept	R_M - R_f	SMB	HML	R^2
Current-Year	Portfolio 1	-0.01	1.21	0.09	-0.10	0.65
		[0.07]	[99.82]	[3.70]	[4.26]	
	Portfolio 2	0.00	1.11	-0.04	-0.28	0.73
		[0.18]	[119.72]	[2.18]	[15.37]	
	Portfolio 3	0.01	0.99	-0.10	0.22	0.68
		[1.08]	[109.03]	[5.41]	[11.99]	
	Portfolio 4	0.01	1.02	0.11	0.28	0.59
		[0.68]	[90.87]	[4.89]	[12.82]	
Year-Ahead	Portfolio 1	-0.01	1.18	0.05	0.04	0.63
		[0.61]	[135.83]	[2.70]	[2.37]	
	Portfolio 2	0.01	1.07	-0.03	-0.07	0.65
		[1.35]	[139.79]	[2.25]	[4.62]	
	Portfolio 3	0.01	0.96	0.22	0.61	0.61
		[1.71]	[129.41]	[11.25]	[15.16]	
	Portfolio 4	0.02	1.01	-0.09	0.07	0.61
		[2.24]	[128.26]	[5.83]	[4.77]	

[ ]: t-values