E Gilbert and D Strugnell

Does survivorship bias really matter? An empirical investigation into its effects on the mean reversion of share returns on the JSE (1984-2007)

ABSTRACT

This paper tests for the impact of survivorship bias by building on the work of Cubbin, Eidne, Firer and Gilbert (2006), and Bailey and Gilbert (2007). The former paper confirmed the existence of mean reversion of relative returns on the JSE, because portfolios of shares with high Price to Earnings (P/E) ratios (being those which had tended to outperform recently) underperformed significantly over five years against portfolios of shares with low P/E ratios. This definition of mean reversion is contentious, but the convention is followed. The latter paper confirmed the economic validity of this conclusion by applying liquidity constraints to portfolio formation. This tended to slightly dampen the observed effects, but confirmed the significant presence of mean reversion. In both cases, extensive efforts were made to include all delisted shares in the study to avoid the effects of survivorship bias. This paper updates both studies by extending the period for a further 21 months, and then quantifies the impact of survivorship bias by comparing the results against those of an equivalent study based on a data set of currently listed shares only. The results of our study confirm that, in general, the effects of survivorship bias are present and material. While similar patterns of mean reversion are detected on both data sets, the returns earned on portfolios selected from currently listed shares are significantly higher than the corresponding returns on portfolios selected from all shares. While survivorship bias does not necessarily affect the conclusion of the patterns of mean reversion revealed in the earlier studies, it is a potentially important issue in any empirical financial research, and effort ought to be made to avoid it.

1. INTRODUCTION

One of the most common challenges facing financial researchers in emerging markets is the lack of a clean and comprehensive data set of price and accounting data for listed firms. There may be some historical data available for currently listed companies, but historical data availability for delisted shares is always an issue.

It is well established in financial research that ignoring delisted companies when conducting historical research leads to survivorship bias in results. As Bain (1972:104) asserts, when commenting on a paper by Wagner and Lau (1971): "the use of ex-post sampling will invariably produce an upward bias in the measurement of returns on risky securities". This bias results from the use of a data set that consists of the survivors over a period, not the full set of companies that were listed over this period. As the characteristics of the survivors are likely to differ systematically from those who have delisted, the results of such a study will be biased.

However, collecting data for delisted companies is a time-consuming and expensive process, and it is legitimate to question whether it is necessary for any given piece of financial research. This paper examines the importance of survivorship bias for the results reported by Cubbin, Eidne, Firer and Gilbert (2006) and Bailey and Gilbert (2007), which found that share returns on the JSE exhibit mean reversion (in fact, their use of the term 'mean reversion' is contentious, as discussed below).

Cubbin *et al*. (2006) provided new evidence for the existence of mean reversion of relative share returns on the JSE through an adaptation of the methodology of DeBondt and Thaler (1985). They compared the average outperformance of portfolios of shares with high P/E ratios (being those which had tended to outperform recently) to portfolios comprised of shares with low P/E ratios, and showed that the high P/E portfolios underperformed significantly over a five-year holding period against portfolios of shares with low P/E ratios for shares listed over the period. Bailey and Gilbert (2007) tested the economic validity of this conclusion by applying liquidity constraints to portfolio formation. Portfolios constructed from more liquid shares tended to display dampened mean reversion results, but the presence of mean reversion was again confirmed.

Mean reversion, as interpreted by DeBondt and Thaler (1985) and as commonly understood, refers to the tendency for the relative returns of shares which have in recent history performed above the benchmark to revert to the benchmark through ensuing underperformance (on a relative basis). The results presented by Cubbin *et al.* (2006) and Bailey and Gilbert (2007) could perhaps be better described as a 'value effect', the tendency for value stocks (with low P/E ratios) to outperform their high P/E counterparts over the medium- to long-term. There will, of course, be some degree of correlation between past return and

*Respectively Department of Economics, University of Stellenbosch and School of Management Studies (Actuarial Science), University of Cape Town. The authors would like to thank Paul van Rensburg for suggesting the study; Greg Bailey for his comments on the methodology; Johan Wiid for his research contribution; I-Net Bridge for access to the data used; and especially Lesley Booysen for attempting to resolve our multitude of data-related queries. We would also like to thank two anonymous referees for their very useful comments.
Email: egilbert@sun.ac.za.

P/E ratio, but this correlation will by no means be complete. Since this study builds on the results of the two previous papers, and was prompted by responses to those papers, we continue with the use of the 'mean reversion' terminology, but readers should keep the above in mind when interpreting our findings.

Both preceding papers highlighted the fact that every effort had been made to include a comprehensive list of delisted shares, thereby eliminating survivorship bias from the results. Several respondents to presentations of these papers' results however posed the question of whether survivorship bias actually makes a difference in this instance. In essence, this question is twofold:

1. Would a different result have been reached had a universe of currently listed stocks only been used? There is no compelling *a priori* reason to suppose that the outperformance of low P/E stocks exhibits itself differently among listed and delisted stocks. And if mean reversion, as defined in Cubbin *et al.* (2006), applies whether or not delisted stocks are included, is the effort of including the latter warranted?

2. Is survivorship bias indeed a problem in equity returns-based research in South Africa generally? While well-established in the international literature, there is little empirical evidence in the South African literature to confirm that it is a problem of equal stature domestically. Our intuition would suggest that it would be an issue – there is no *a priori* reason to suspect otherwise.

This paper updates both preceding studies by extending the period under consideration for a further 21 months. It further tests for, and quantifies, the impact of survivorship bias. It does this by comparing the results of a liquidity-adjusted mean reversion study for the population of all listed shares, against those of an equivalent study based on a data set of currently listed shares only. This approach allows for a direct test for the relevance of survivorship bias in the context of mean reversion of relative returns on the JSE, as defined by Cubbin *et al.* (2006). Our secondary objective is to establish whether survivorship bias is indeed an issue worth taking seriously in South African empirical financial research generally.

Section two of this paper reviews the evidence for the importance of survivorship bias in empirical studies and previous efforts to test for mean reversion in South Africa. Section three outlines the methodology used in this study and the results are presented in Section four. Section five concludes.

2. LITERATURE REVIEW

2.1 Survivorship bias

The presence and possible effects of survivorship bias has concerned researchers in this field, particularly in the United States (US). The primary data sources used in the US for finance-related research are the COMPUSTAT database (for accounting data) and the data prepared by the University of Chicago's Centre for Research into Securities Prices (CRSP) (for price data). While the CRSP data set does not suffer from significant biases, the way the COMPUSTAT database is created is problematic.

Firstly, the accounting data for companies that are delisted are deleted from the database, causing survivorship bias. Secondly, new companies added to the dataset are included with a full history, which means that companies that do not succeed are not included (selection bias). A third problem with the COMPUSTAT database is the presence of 'look ahead' bias. This occurs when data are recorded as being available to investors at a particular time, when it actually only becomes available at a later stage. For example, the annual financial statements (or announcement of the earnings) for a company may only be made public several months after its year end. However, in the COMPUSTAT database, these earnings would have been recorded as applying to the whole financial year, i.e. they were reported immediately.

Researchers became aware of the presence and potential effects of survivorship bias in the early 1970s when looking at the question of using accounting data in explaining the cross-section of share returns. Initial tests were limited and focused on establishing whether the conclusions of previous studies were robust (Ball and Watts, 1979; Salamon and Smith, 1977). The first comprehensive examination of, and test for, the effects of this bias only happened later (McElreath and Wiggins, 1984; Banz and Breen, 1986, Davis 1994, 1996). The question of the impact of survivorship bias has also been discussed in the context of establishing the persistence of portfolio manager returns (Grindblatt and Titman, 1989; Brown, Goetzmann, Ibbotson and Ross, 1992; Carpenter and Lynch, 1999).

Ball and Watts (1979) tested for presence of survivorship bias in a study they completed in 1972 where they critically evaluated the time series properties of earnings per share (EPS) for listed companies. They randomly selected 25 shares that were in existence in 1916 and 25 shares that were in existence in 1966. They compared the characteristics of the EPS time series for these shares with 25 shares which had been in existence from 1916 to 1966. They found no significant difference in the results for these two samples and concluded that survivorship bias was unlikely to have affected their results.

McElreath and Wiggins (1984) reviewed the reasons for the delisting of the 330 firms that had left the New York Stock Exchange (NYSE) in the period 1970 to 1979. They concluded that the likely size of the bias was small due to the fact that 55% of these delisting were due to mergers. Bankruptcy and liquidations only accounted for 6% of the delisted firms. However, no further quantitative comparative analyses between this group of firms and those still listed was presented.

Banz and Breen (1986) presented the first direct test of the size and impact of survivorship bias on the COMPUSTAT database. They compared the nature of the properties of two separate populations of firms: a complete and a partially complete (currently listed only) COMPUSTAT list of firms. For the period 1974 to 1981 they independently collected accounting data on a monthly basis for all listed firms on the NYSE and the American Stock Exchange (AMEX). They then compared the differences in the returns of similar equally weighted portfolios created from this list and the COMPUSTAT dataset. The differences in returns were found to be significant at the 1% level for all portfolios, whether the raw or risk-adjusted portfolio returns were compared.

They went on to point out that this difference in returns is actually the combined effects of the survivorship and "look ahead" biases. To isolate the effects of the differences due to survivorship bias, they created a subset of firms that are included in the partial COMPUSTAT data series from their complete list. They then recreated the portfolios as explained above using their complete list and this subset of their complete list. In other words, the effects of the differences in the portfolio returns could only be due to survivorship bias. Again, the differences in returns were statistically significant.

Banz and Breen (1986) also evaluated the effects of survivorship bias on the results of studies investigating the presence of size and P/E effects on portfolio returns. They found that using the complete data series (i.e. corrected for the missing firms) leads to the rejection of earlier claims of a P/E effect on returns when size is controlled for. They concluded that survivorship bias does seem to matter.

Davis (1996) also tested for the effects of survivorship bias on the results of his earlier study (Davis, 1994). He identified, and then directly compared, the nature of the firms listed on the NYSE and the AMEX that were excluded from the COMPUSTAT data set to those in the set. He found that the nature of the excluded firms was systematically different in terms of relative size (the excluded firms are smaller) and monthly returns (the excluded firms have lower returns). While the inclusion of the shares listed in the Moody's database (but not in the COMPUSTAT database) did not change the conclusions of his previous study, it did lead to non-trivial differences in the regression results, both in terms of economic and statistical significance.

Kothari, Shanken and Sloan (1995) confirmed the importance of the effects of survivorship bias. They firstly established that the returns to the shares excluded from the COMPUSTAT database (but on the CRSP data series) were on average 9 to 10 percentage points lower than the shares included in the database. This emphasised the systematic difference of the nature of the excluded shares. They also showed that the significant book to market value (B/M) result of Fama and French (1992) could, in part, be explained by the survivorship bias in the data set they used[1].

In summary, survivorship bias matters for US data, but there have been no studies of the importance of this bias in SA. Secondly, the most persuasive way to test for its presence is to do a comparative study using two different datasets: a complete and incomplete set of share data. Our study adds to this literature by adopting this comparative approach and by using data for South Africa to test in particular for whether survivorship bias would have affected the earlier conclusions of mean reversion of relative share returns on the JSE, and in general for whether survivorship bias has a significant impact on returns.

2.2 Mean reversion

Mean reversion, in the context of share returns, was first identified by DeBondt and Thaler (1985) for the US. They showed that the returns to shares that have either out- or under-performed historically (relative to some benchmark) tend to revert, on average, to the mean, i.e. they under- or out-perform the benchmark in the future on a relative basis.

This has proved to be a robust empirical result for different time periods and countries. Cubbin *et al.* (2006) present a detailed review of the international literature and the rationales for this phenomenon. We now look at three studies that focus primarily on testing for the presence of mean reversion in SA.

Plaistowe and Knight (1986) reviewed share data for the period 1973 to 1980. They ranked shares on a market to book ratio basis, created portfolios of the top and bottom 35 shares. They then compared the average cumulative weekly returns to winner and loser portfolios over a period of a year after the portfolio formation date. They found that while the loser portfolios did not exhibit abnormal returns relative to the returns of the RDM 100 Index of industrial shares,

[1]The other reason offered for their finding that the B/M result of Fama and French (1992) was not valid was that their dataset was significantly longer (1927 to 1990) than that used by Fama and French (1941 to 1990). In other words, they believe that Fama and French's results were period specific.

the winner portfolios did exhibit significant abnormal losses

Page and Way (1992) examined the monthly returns for 204 shares traded between 1974 and 1989. Like DeBondt and Thaler (1985) they used cumulative historical returns to rank shares at the portfolio formation period. Two separate sets of portfolios were constructed. The first was comprised of non-overlapping portfolios of two years duration each; and the second of overlapping portfolios of three year duration. Two start dates per year were used (1st of July and 1st of January) for each approach to account for timing issues. For the 36-month portfolios the loser portfolios on average outperformed the winner portfolios on a relative basis by almost 15%. A similar result is achieved for the two year non-overlapping portfolios. They conclude on the basis of their analysis that there is 'clear evidence of investor over-reaction on the Johannesburg Stock Exchange' (Page and Way, 1992:43).

A key weakness of both of these studies is that their analysis is based solely on shares that were still being traded at the end of the periods under consideration – in other words their conclusions suffer from survivorship bias. As will be shown in the discussion of the data set used in this paper, the number of firms excluded from their analyses is indeed large. This paper uses a data set that includes the returns for all shares listed on the JSE All Share Index (ALSI) over the entire period under consideration.

Mean reversion was formally analysed by Bradfield and Ardington (1997), who avoided the problems of survivorship bias by basing their analysis on the JSE All-Share Index. On the basis of runs tests, autocorrelation tests and variance ratio tests, they verified that mean reversion is indeed present on the JSE based on annual returns. Importantly, they observed that this suggests that equity investment in South Africa may in fact be less risky than is often supposed: in the first instance, investors may be presumed to prefer a market which displays mean-reverting tendencies, meaning that market downturns are likely to be followed by recoveries, than one which typically trends, and in the second, extrapolations of variance of returns for long-term investors from annual estimates are likely to be overstated, due to serial dependence of returns.

3. RESEARCH METHODOLOGY

DeBondt and Thaler (1985, 1987) first examined the question of the existence of mean reversion in returns for the US. They identified 'winner' and 'loser' portfolios based on past abnormal returns. They then tracked the relative performance of these portfolios. Cubbin *et al.* (2006) applied the DeBondt and Thaler (1985) methodology to JSE share returns with one significant difference: their 'winners' represented firms with high P/E ratios and their 'losers' represented firms with low P/E ratios.

Bailey and Gilbert (2007) extended the results of Cubbin *et al.* (2006) by applying liquidity caps to 'High P/E' and 'Low P/E' portfolios in an attempt to evaluate the economic reality of the abnormal returns seemingly offered by the presence of mean reversion. They tested for effects of liquidity constraints on the presence of mean reversion in multiple portfolios by applying liquidity caps. Depending on the value of a portfolio, a share would only be considered for inclusion if its average monthly traded volume was sufficiently large. They concluded that, although dampened, mean reversion persists after application of liquidity constraints.

This paper builds on this work, by testing for the presence and effects of survivorship bias by applying the Bailey and Gilbert (2007) methodology to two separate groups of firms: the complete list of firms and the list of firms that are currently listed on the JSE.

We replicated the methodology used by Bailey and Gilbert (2007) using historic share data collected from I-Net Bridge. The key advances are: firstly, we extended the period under consideration by 21 months; and secondly, we conducted the analysis on two separate groups of shares: the currently listed shares and the complete list of shares (i.e. including the delisted shares). By comparing the difference between the returns of the high and low P/E portfolios for each set of shares, we can establish whether the conclusion of the presence of mean reversion is in any way affected by survivorship bias. As a secondary objective, by comparing the difference between returns on the current and complete share databases for each of the high and low P/E portfolios, we can establish whether survivorship bias affects returns generally; financial common sense as well as the international literature would lead us to believe that this ought to be the case, and that returns for the current set of shares will systematically outperform the complete set, but it is nevertheless worth verifying this.

Historical data for shares traded on the JSE at each month end were obtained from I-Net Bridge. Month end price, Earnings Yield (from which P/E ratio was derived) and dividend yield data were collected for all shares listed on the All Share Index (ALSI) for the JSE for the period 31 October 1984 to 30 September 2007, a 23 year period[2].

A list of delisted shares was obtained from I-Net Bridge[3] and all the data for these shares was included

[2]We excluded preference shares ('P' series) from the analysis, because they have a different dividend risk profile.

[3]We have used the old code list from I-net Bridge to identify the delisted shares. While this is the most complete and accurate list

in the dataset for this period. In total, 1631 shares were included in the analysis, of which 841 had either been delisted or were removed due to some other type of corporate action.

The data presented in Figure 1 highlight the existence and size of the difference in the two groups of shares. The dark shaded area represents the shares that are currently listed (the 'current' list) and the light shaded area indicates all the shares that were listed (the 'complete' list). The area between the two parts represents the shares that would have been excluded from the analysis if only the currently listed shares were included in the analysis. The relative variation of the two lists ranged from 0% (in the final month, by definition) to 80% (in late 1992).

While I-Net Bridge was the most complete and accurate data source that we were able to use, we still encountered significant data completeness and quality issues. The presence of the discontinuous jumps in the number of total shares in September 1991, November 1992 and September 1997 in Figure 1 suggest that the list of 'dead' shares from I-Net Bridge is still not completely accurate. Moreover, the P/E Ratio data was found to be unreliable so we used the inverse of earnings yield (EY) ratio that had been recorded more extensively and more accurately in the dataset. We also encountered obvious errors in share price and dividend yield data. Where found, these observations were deleted. The reasons for these errors have not been fully resolved and one of the project's conclusions is that we need to consolidate and clean this historical database through comparisons with other sources (e.g. McGregor BFA Net, Datastream or Reuters).

Following Bailey and Gilbert (2007), we adjusted for liquidity by applying a liquidity cap. We assumed an investor starts with a specific amount of money to be invested in the portfolio. For the sake of illustration, assume this is R1 billion. As every portfolio is made up of 35 equally weighted shares, the investment required in each share is equal to the value of the portfolio (R1 billion) divided by the number of shares in the portfolio (35). Approximately R28.5m is thus invested in each share. It is at this point that the liquidity cap comes into consideration. For the share to be considered for inclusion in the portfolio, 50% of the average monthly value of shares traded must exceed this threshold of R28,5m, suggesting that it would be possible for the portfolio manager to unwind a position in that share within a reasonable time period. This is calculated by taking the average volume of shares (which is based on the previous 12 months' volume traded) multiplied by the current share price and then dividing by two. If, for example, the average number of shares traded each month is 500 000 and the current share price is R10 then the liquidity cap for this share is 500 000 × R10 × 50% = R2,5 million.

In this example, the share would be filtered out due to its lack of liquidity for an investor of this size (R28,5m investment > R2,5m liquidity cap). However, if the investor had a portfolio of R10m, then this share would be considered available for investment (R10m/35 = R0,285m holding per share which is less than the R2,5m liquidity cap).

Using the approach outlined above, a liquidity constraint can be derived from portfolio size, share price and average volume traded. As explained above, the impact of the liquidity cap is dependent on the size of the portfolio. Seven different portfolio sizes were evaluated: R0 (i.e. no liquidity constraint), R100 000, R1 million, R10 million, R100 million and R1 billion.

These portfolio sizes reflect nominal values as at the end of the portfolio formation period of the study, which was September 2002. To ensure that portfolio sizes were comparable in real terms, the nominal size of portfolios was discounted by the monthly inflation rate for portfolios formed prior to September, 2002.

The liquidity cap was calculated using this methodology for all shares for a given portfolio size. Those shares that did not meet the liquidity cap were excluded from the rest of the portfolio construction process.

The shares that met this liquidity requirement were then ranked by their P/E ratios at the portfolio formation date and portfolios consisting of the top 35 (high P/E) and bottom 35 (low P/E) shares were constructed[4]. These shares were weighted equally and the excess monthly returns relative to the ALSI index returns were then calculated for 5 years (60 months) from the portfolio formation date. The start date was then incremented one month forward and the process repeated. Applying this process over the entire period allowed for the creation of 5184 portfolios and 311 040 monthly returns.

of 'dead' shares that we were able to obtain, the jumps in the number of firms covered by our analysis (see Figure 1) suggest that we still do not have a comprehensive list.

[4]There were some months when there were less than 70 shares in the overall portfolio (especially when liquidity constraints were introduced). The high P/E and low P/E portfolios were formed in this case by taking the list of shares available (ranked by P/E) and dividing the shares into two equally sized groups.

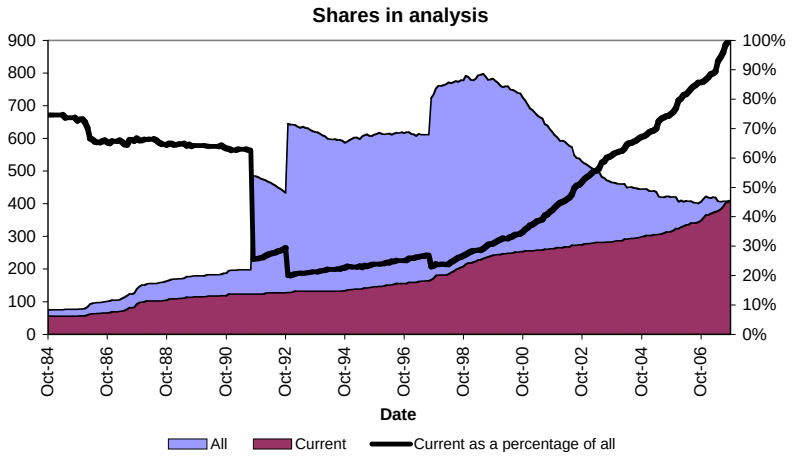

Figure 1: Number of listed and delisted firms included in the analysis

The monthly returns for each high P/E and low P/E portfolio were then combined and averaged to give the final average relative returns for each category. Thus, the returns for month one (post portfolio formation) of all the high P/E portfolios were combined to give the average high P/E portfolio return for month one. The same exercise was conducted for the low P/E portfolios to give the average returns for each month for these portfolios up to 60 months after portfolio formation.

Following the approach adopted by Cubbin *et al.* (2006) and Gilbert and Bailey (2007), our monthly portfolio returns are calculated as the sum of the capital gain (change in price/starting price) and dividend yields (DY) (with the latter converted to an equivalent effective monthly rate of return). Returns were calculated according to the following equation:

$$R_t = \frac{P_t}{P_{t-1}} + \frac{D_t}{12} - 1$$

where

R_t, P_t, D_t are respectively the return in month t and the share price and dividend yield at the end of that month.

This approach yields a total return that is a proxy for the actual monthly total return, as the dividends are not received on an equal monthly basis over the 12 months of the year. Due to lack of data on the timing of the dividend payments, this approximation was judged to be the next best solution.

One significant problem that we encountered with the implementation of this approach was the very high dividend yields associated with shares whose prices fell rapidly and approached zero. As the portfolios are equally weighted, the portfolios returns are the average of the individual component shares' returns. Thus if uncorrected, these outliers led to some portfolios having extremely high returns. To correct for this, we made all DY greater than 150% equal to zero. This minimised the effects of the real outliers on the portfolio returns[5] reported in this analysis.

4. RESULTS

The analysis was conducted for six different portfolio sizes. For reasons of brevity we report the results for all shares (i.e. no liquidity cap) and, where relevant, highlight the differences in the results for the other five portfolio sizes.

The pattern of average portfolio returns for the high and low P/E portfolios drawn from the complete and currently listed groups of shares is reported in Figure 2. These returns are the arithmetic average of portfolio returns relative to the ALSI return in the relevant month post portfolio formation.

Figure 3 shows the cumulative portfolio returns (annualised) for each of the four portfolios:

Three important observations can be drawn from the data presented in this figure. Firstly, the outperformance of the low P/E shares relative to the high P/E shares holds for both groups of portfolios. Secondly, the pattern of the returns of the portfolios drawn from currently listed shares systematically differs from those drawn from the complete list of shares. Finally, the returns of the portfolios drawn from the currently listed shares only seem to systematically exceed the returns of the portfolios drawn from the complete set of shares, with the exception of a handful of months, and certainly systematically outperform on a cumulative basis.

[5]We recognise that the use of 150% as the cutoff value is an arbitrary decision. We arrived at this number by considering the tradeoff between the effects of the cutoff level on the number of shares affected and the reality of the portfolio earning dividend yields of such high values. At levels below this value, we found that a relatively large proportion of shares were affected. Anything above this number would lead to an unrealistically large upward bias in the measured return for the portfolio containing such a share.

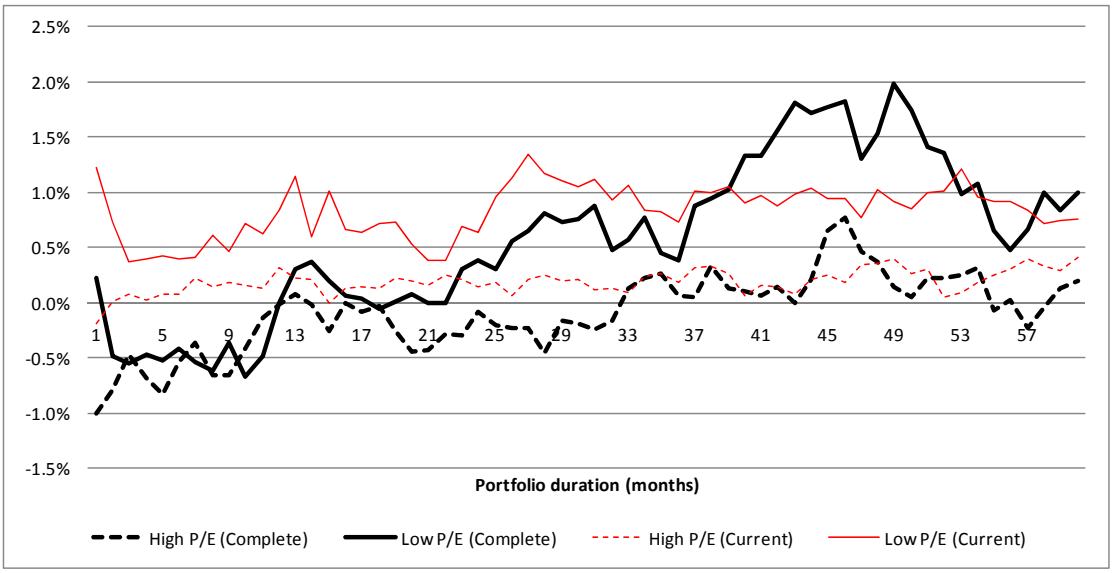

Figure 2: Returns for high and low P/E portfolios (relative to the ALSI) drawn from currently listed and complete populations – no liquidity cap

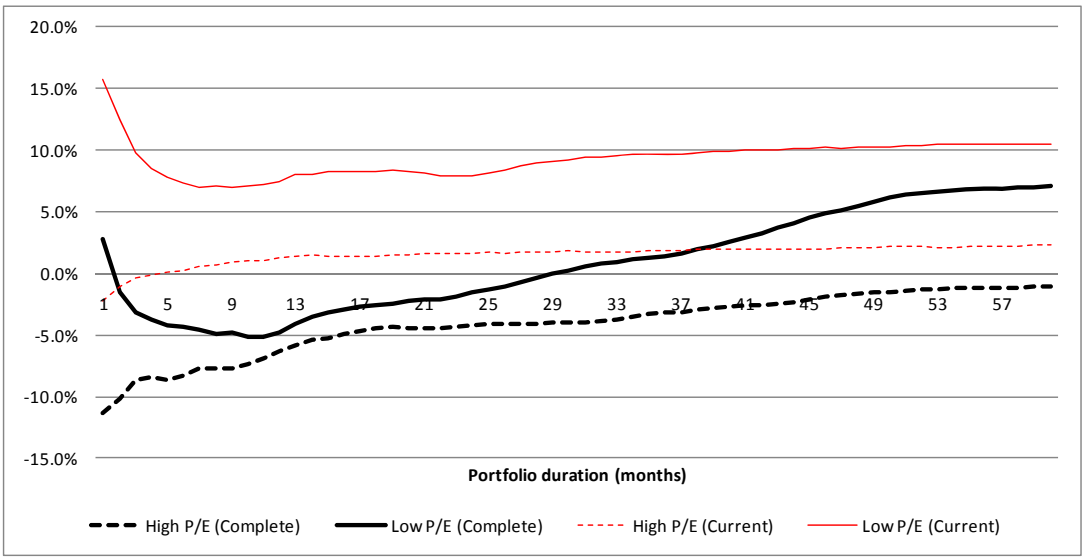

Figure 3: Cumulative annualised returns for high and low P/E portfolios (relative to the ALSI) drawn from currently listed and complete populations – no liquidity cap

The first observation suggests that mean reversion is present for portfolios drawn from both groups of shares. The second and third observations indicate that survivorship bias is present in this study. These conclusions are examined in more detail below.

The excess geometric mean returns (annualised) for each group of portfolios are summarised in Table 1. This represents the average annual portfolio return in excess of the equivalent ALSI return for the same period. The same geometric means for all five liquidity caps are presented in Figure 4.

Looking at the data in Table 1, we can see that the mean excess returns for high P/E portfolios are lower than the equivalent returns for the low P/E portfolios for both groups. Similar results hold for the differences in mean returns for the portfolios for all five liquidity caps considered. A significant limitation of our methodology is that there is a degree of non-independence between portfolio returns. We form portfolios on the basis of P/E ratio at each month-end throughout the sample period.

Table 1: Geometric mean excess returns for high and low P/E portfolios (annualised) – no liquidity cap

Portfolio	Geometric Mean Return
High P/E (Complete)	-1,05%
Low P/E (Complete)	7,08%
High P/E (Current)	2,32%
Low P/E (Current)	10,47%

Consider for example the high P/E ratio portfolio. At 31 October 1984, this consists of an equal weighting in each of the 35 shares with the highest P/E ratios at the portfolio formation date. The returns on this portfolio are then tracked for the ensuing sixty months. We then move on to 30 November 1984 and repeat the process. However, to the extent that there is some commonality between the highest P/E ratios at the two dates, there will be a degree of dependence between the first month's return for the portfolio formed on 30 November and the second month's return for the portfolio formed on 31 October. This makes precise statistical testing of the significance of differences in returns achieved in overlapping five-year periods almost impossible. The portfolio formation technique is one which is commonly used in empirical financial research, and the development of statistical testing methodology for testing the significance of returns where there is some degree of overlap, and hence dependence, seems a fruitful area for further research; for the purposes of this paper, however, we sidestep the issue by not reporting on statistical significance, following the well-ploughed furrow of, for example, DeBondt and Thaler (1985:802).

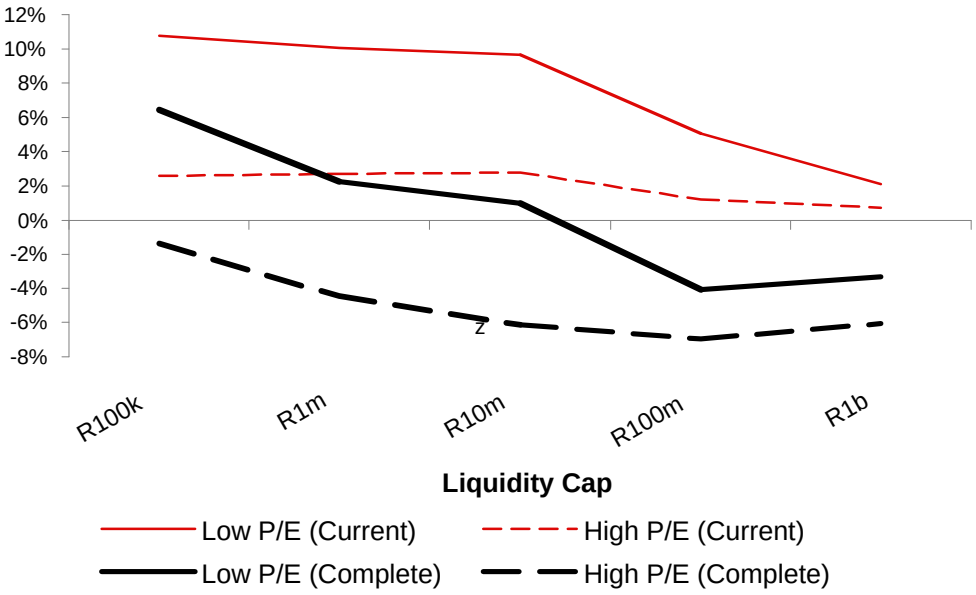

Figure 4: Impact of liquidity caps on the geometric mean excess returns of high and low P/E portfolios

The dampening effect of liquidity caps on the size of the mean reversion effect reported in Bailey and Gilbert (2007) is also confirmed here. As is illustrated in Figure 4, there is a clear apparent negative relationship between the geometric mean return for all the portfolios and the liquidity cap level, indicating that liquidity constraints do, as expected, limit portfolio managers' attempts to outperform the market.

The mean reversion result can also be clearly seen if the differences in excess returns between low and high P/E portfolio shares are plotted for the two groups of shares (Figure 5). For the complete list of shares, the average monthly excess return for the low P/E portfolio exceeds that of the high P/E portfolio for 55 out of the 60 months considered. For currently listed shares, the same result holds for all 60 months. The equivalent cumulative returns (annualised) are shown in Figure 6.

There is a great deal of similarity between the patterns of outperformance of low and high P/E portfolios illustrated in Figure 5. The differences are remarkably consistent – no matter what set of shares the portfolios are drawn from[6]. This suggests that the correction of the sample data to avoid survivorship bias was unnecessary for the research conducted by Cubbin *et al.* (2006) and Bailey and Gilbert (2007): qualitatively similar results would have been found by using a

[6]Note that no attempt was made to adjust these results for the risks of the portfolios. It is possible (though we believe unlikely) that Betas could explain this effect. Zarowin (1989, 1990), for example, finds that risk adjustments of this sort do explain the presence of mean reversion in share returns in the United States as identified by DeBondt and Thaler (1985, 1987). However, these results have been questioned (Albert and Henderson, 1995; Gropp 2004). This is an obvious area for further study, however.

database of currently listed stocks only. This is interpreted further below.

We then use the outcomes to test for the effects of survivorship bias more generally. Its impacts can be seen clearly if we plot the differences between the mean portfolio returns for high P/E portfolios drawn from the two sets of shares, and the same for low P/E portfolios. This is illustrated in Figure 6.

If there are systematic differences in the firms delisted as compared to those currently listed, then there should be a consistent difference in the returns to high (and low) P/E portfolios drawn from the two groups. In particular, we would expect there to be a positive difference between the portfolios drawn from the currently listed shares and those drawn from the complete list: a material proportion of delistings would have occurred due to failure of the listed firm, depressing returns for these shares compared to those which are still listed at the end of the study period.

This is almost exactly what we find illustrated in Figure 7. The difference between the returns for the current portfolios is positive for 50 months out of 60 for the high P/E portfolios and 43 months for the low P/E portfolios. Figure 8 shows that on a cumulative basis, the excess return is always positive.

The geometric averages of the differences in mean returns are summarised in Table 2. This represents the difference in high (and low) P/E portfolios for the current and complete lists of shares.

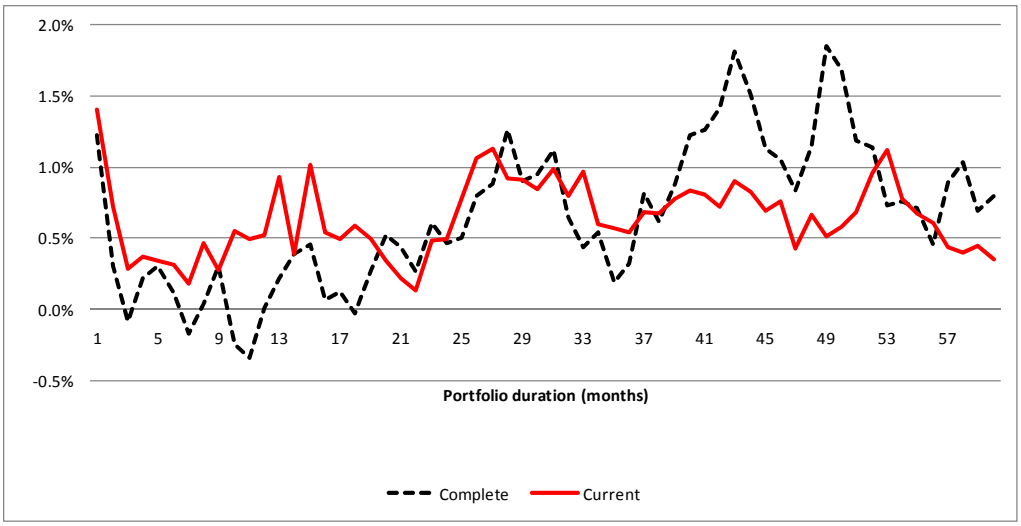

Figure 5: Difference between low P/E and high P/E portfolio monthly excess returns – no liquidity cap

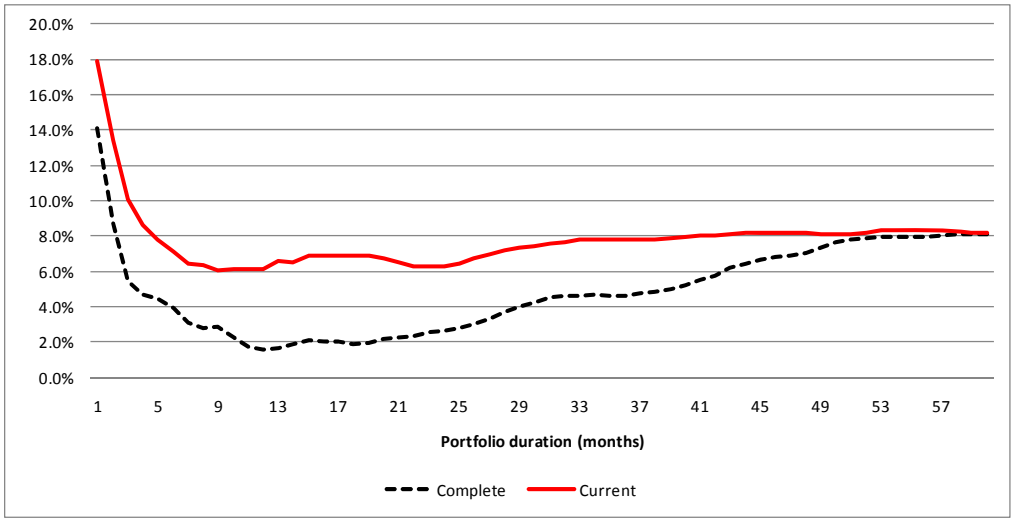

Figure 6: Difference between low P/E and high P/E portfolio cumulative annualised excess returns – no liquidity cap

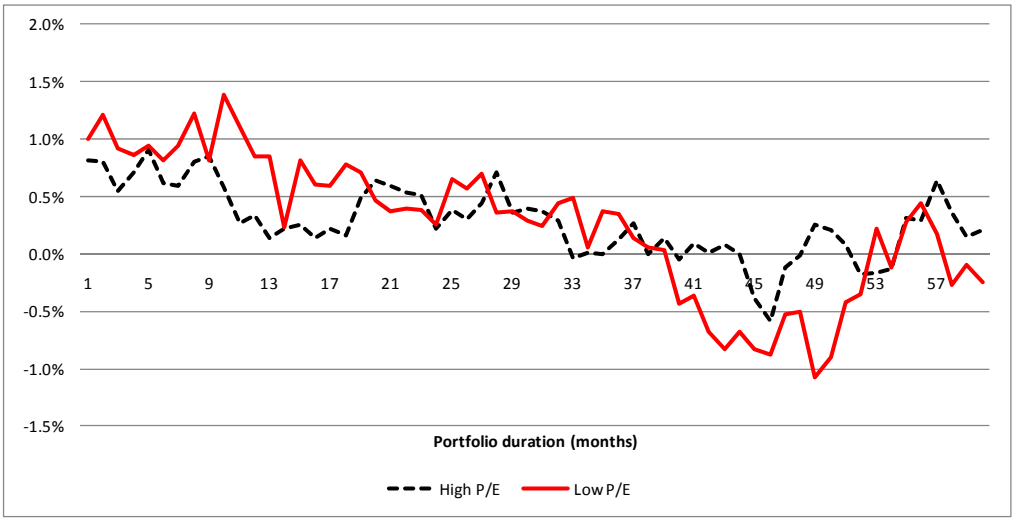

Figure 7: Differences in mean monthly portfolio returns (Current - Complete) for high and low P/E portfolios

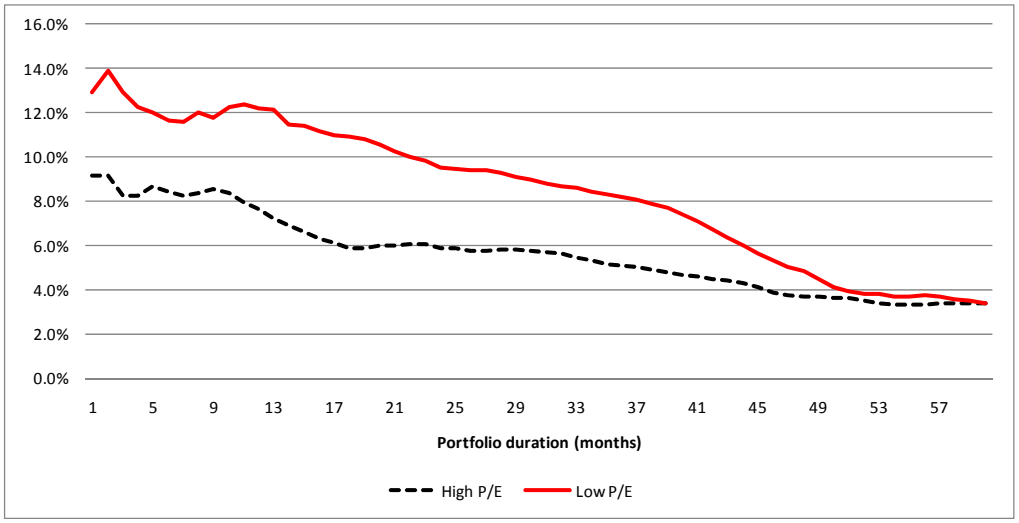

Figure 8: Differences in mean cumulative annualised returns (Current - Complete) for high and low P/E portfolios

The data presented in Table 2 shows positive differences in the portfolio returns for high (and low) P/E portfolios, although the statistical significance of these differences is not analysed for the reasons set out above. While not reported here, these results hold true for all liquidity adjusted portfolios considered.

Table 2: Geometric average difference in returns for high and low P/E portfolios for Current and Complete share lists

Portfolios compared (Current – complete)	Difference in geometric means
High P/E portfolios	3,39%
Low P/E portfolios	3,13%

The differences in the geometric average excess returns generated by the high and low P/E portfolios drawn from the two lists are reported in Figure 9 for the five different liquidity caps. This shows that survivorship bias is present, no matter what liquidity constraints are imposed.

5. CONCLUSION

The data and analysis presented in this paper lead to a clear conclusion regarding the presence of survivorship bias, but implications for researchers are nuanced.

On the one hand, the presence, and importance, of survivorship bias is clearly demonstrated in the difference of returns between the complete and listed groups of firms. The mean levels of returns from the

currently listed portfolios (both high P/E and low P/E) are higher than those generated from the portfolios of the complete data set. This is exactly the sort of bias that would be expected if the characteristics of the delisted firms are systematically different from those that are listed at each point in time. This highlights the need for researchers to take the survivorship bias problem seriously – you are likely to get incorrect answers if you do not.

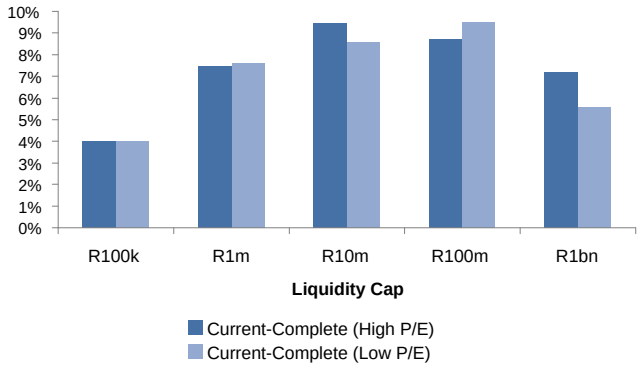

Figure 9: Differences in geometric average portfolio returns between current and complete portfolios

However, this is not necessarily the case for every avenue of financial research. It appears that the conclusions of Cubbin *et al.* (2006) and Bailey and Gilbert (2007) that mean reversion (as defined in those papers) of share returns exists on the JSE, would not have been affected had a database of currently listed stocks only been used (i.e. a data set which has not been adjusted to avoid survivorship bias).

This suggests that mean reversion of returns is a robust phenomenon as it applies to both groups of shares that we examined. Given that there is no *a priori* reason why mean reversion should not apply to all shares equally, this is perhaps not an unexpected result. However, it would in our opinion be foolhardy for researchers to rely on the absence of clear *a priori* reasons for systematic differences to justify the use of currently listed stocks for any given area of investigation. Research is, after all, often about uncovering relationships that are not immediately obvious prior to analysis, and the risk of reaching misleading conclusions suggests that all empirical work, barring studies which expressly set out to examine currently listed stocks only, ought to ensure that they are based on data sets which eliminate as far as possible any element of survivorship bias.

Our conclusions are subject to the following caveats. Firstly, we are not convinced that we have a complete data set going back in time. As illustrated in Figure 1, there are clear jumps in the number of firms covered by the data provider. This strongly suggests that the data source we used is still not complete and thus free from survivorship bias. Secondly, in the process of creating the portfolios for this analysis, we noticed anomalies in the data which raises questions regarding the quality of the entire data set. This suggests the need for a data cleaning exercise on this data set. Thirdly, when calculating portfolio returns, we have treated delisted shares as delivering a zero return in the month following their delisting. This is correct if the company is in the process of being liquidated. However, this is obviously not correct if companies are delisted for other reasons and investors receive the value of their shares[7]. Finally, we have not made any risk adjustments when comparing portfolio returns. There is some literature (Zarowin, 1989, 1990; Clare and Thomas, 1995) that suggests that this may well affect the conclusions regarding the presence of mean reversion of returns. This offers an interesting opportunity to extend this work.

In conclusion, our analysis shows that any research that excludes delisted shares is likely to be subject to survivorship bias. This may not materially affect the outcomes of the studies (as in this case), but our work suggests that including data for delisted shares is likely to have a significant effect on the results reached. Researchers should be aware of this and attempt to include such data in any empirical analysis of this sort.

REFERENCES

Albert RL and Henderson GV. 1995. Firm size, overreaction and return reversals. *Quarterly Journal of Business and Economics*, 34(4):60-80.

Bailey G and Gilbert E. 2007. The impact of liquidity on mean reversion of share returns of the JSE. *Investment Analysts Journal,* 66:1-11.

Bain JS. 1972. Diversification and risk: Comment. *Financial Analysts Journal,* 28(6):104-106.

Ball R and Watts R. 1979. Some additional evidence on survival biases. *Journal of Finance,* 34(1):197-206.

Banz RW and Breen WJ. 1986. Sample-dependent results using accounting and market data: some evidence. *Journal of Finance,* 41(4):779-793.

Bradfield, DJ and Ardington CS. 1997. A Note on the riskiness of long term investment on the JSE. *Journal for Studies in Economics and Econometrics*, 21(3):67-78.

[7]This is equivalent to assuming that the return on the delisted share is -100%. This would overstate the loss (and understate the return) if some value is received for the share on its delisting. To test for the sensitivity of our results to this assumption, we re-ran the analysis assuming a -50% return to the shares in the month following their delisting. There was not material change to our results.

Brown SJ, Goetzman W, Ibbotson RG and Ross SA. 1992. Survivorship bias in performance studies. *Review of Financial Studies,* 5(4):553-580.

Carpenter JN and Lynch AW. 1999. Survivorship bias and attrition effects in measures of performance persistence. *Journal of Financial Economics,* 54(3):337-374.

Chan LKC, Jegadeesh N and Lakonishok J. 1995. Evaluating the performance of value versus glamour stocks - the impact of selection bias. *Journal of Financial Economics,* (38)3:269-296.

Cubbin E, Eidne M, Firer C and Gilbert E. 2006. Mean reversion on the JSE. *Investment Analysts Journal,* 63:1-17.

Clare A and Thomas S. 1995. The overreaction hypothesis and the UK stockmarket. *Journal of Business Finance & Accounting,* 22(7):961-973.

Davis JL. 1994. The cross-section of realized stock returns: the pre-COMPUSTAT evidence. *Journal of Finance,* 49(5):1579-1593.

Davis JL. 1996. The cross-section of stock returns and survivorship bias: evidence from delisted stocks. *Quarterly Review of Economics and Finance,* 36(3):365-375.

DeBondt WFM and Thaler RH. 1985. Does the Stock Market Overreact? *Journal of Finance,* 40(3):793-806.

DeBondt WFM and Thaler RH. 1987. Further evidence on investor overreaction and stock market seasonality. *Journal of Finance,* 42(3):557-582.

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

Grindblatt M and Titman S. 1989. Mutual fund performance: An analysis of quarterly portfolio holdings. *Journal of Business,* 62(3):313-416.

Gropp J. 2004. Mean reversion of industry stock returns in the U.S. 1926–1998. *Journal of Empirical Finance,* 11(4):537-551.

Kothari SP, Shanken J and Sloan RG.1995. Another look at the cross-section of expected stock returns. *Journal of Finance,* 50(1):185-224.

McElreath RB and Wiggins CD. 1984. Using the COMPUSTAT tapes in financial research: problems and solutions. *Financial Analysts Journal,* 40(1):71-76.

Page MJ and Way CV. 1992. Stock market over-reaction: The South Africa evidence. *Investment Analysts Journal*, 36:35-49.

Plaistowe T and Knight RF. 1986. Premium to book value may be a contrary indicator. *Investment Analysts Journal*, 28:35-39.

Salamon GL and Smith ED. 1977. Additional evidence on the time series properties of reported earnings per share: comment. *Journal of Finance*, 32(5):1795-1801.

Wagner WH and Lau SC. 1971. The effects of diversification on risk. *Financial Analysts Journal*, 27(6):48-53.

Zarowin P. 1989. Does the stock market overreact to corporate earnings information? *Journal of Finance*, 44(5):1385-1399.

Zarowin P. 1990. Size, seasonality, and stock market overreaction. *Journal of Financial and Quantitative Analysis*, 25(1):113-125.