ESSAYS IN EMPIRICAL FINANCE: FROM VENTURE CAPITAL TO THE S&P 500

By

DANIEL SEÁN FINUCANE

B.Math, The University of Waterloo, 1993

M.Math, The University of Waterloo, 1995

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

In

THE FACULTY OF GRADUATE STUDIES

Business

THE UNIVERSITY OF BRITISH COLUMBIA

December 2004

Abstract

Stocks often increase in value when they are added to the S&P 500 and lose value when they are dropped from the index. There are several hypotheses advanced for this effect. The information hypothesis is that when a firm is added this reveals positive information about the firm. This information effect would explain stock price movements associated with additions and deletions. I examine several stocks that were dropped from the S&P 500 purely because they were foreign stocks. Accordingly, there should be no negative implication regarding these stocks and the information effect should be absent. The foreign firms behave differently from the normal deletions in all time periods. Further evidence of the information effect is found when those firms which are transferred to another S&P index significantly outperform those which are deleted from S&P indices entirely. This study does, therefore, provide support for the information hypothesis.

I next look at one of the recent theories of deviation from diversification in the allocation of assets – "familiarity". I look at the specific case of the decision by defined benefit trustees to allocate assets to venture capital. One might expect that the pension funds of companies in highly venture capital industries and near venture capital clusters would be more likely to invest in venture capital, despite the high correlation with human capital and real estate. I find that pension funds do not seem to base the decision on whether or not to invest in venture capital on familiarity, as measured by industry and geography.

American corporations that sponsor pensions may have defined benefit pension plans, defined contribution (DC) plans, or both, and have been shifting toward DC plans over the past 25 years. We investigate the effect of cross-firm and time-series variation in the DC share of pension assets on corporate financial performance. Using several return measures, including return on assets, operating return on assets, and return on equity, we find that larger DC shares tend to give rise to higher returns. We interpret these results as arising in part from more efficient worker retirement and mobility decisions under DC plans.

Table of Contents

List of Tables

List of Figures

Acknowledgements

This thesis is a compilation of three distinct pieces of work and different people have helped with each section. Above all I would like to thank my committee members Jim Brander, Kai Li and Adlai Fisher for help and guidance from my proposal right through every step of the process.

For Chapter 2 I would like to thank Lisa Kramer and forum participants at UBC, Queen's School of Business, Norwegian School of Management BI, University of Exeter, University of Lancaster, the EFA 2003 Doctoral Symposium and the NFA 2003 for their helpful comments and direction.

For Chapter 3 I would like to thank Chris Veld and participants at the 2004 NFA conference. I would also like to thank the Philips Hager North Centre and the Maurice Young Centre for financial support, and Venture Economics, Price Waterhouse Coopers and Money Tree for providing data.

Chapter 4 was co-authored with Jim Brander. Kai Li also made valuable contributions to the paper. We also received helpful comments from Adlai Fisher and from Keith Head. In addition, we thank Jennifer Ng and Timothy Leung for their excellent work as research assistants. Financial support from SSHRC grant 412 98 0025 is gratefully acknowledged.

Chapter 1 - Introduction

In this empirical examination of finance I will look at very large well-established firms and very small start-up firms. In Chapter 2 I look at the effect on stock prices of firms that have been deleted from the S&P 500. The S&P 500 represents many of the largest firms in the world and I take the rare opportunity in finance to study a natural experiment to determine investor reactions in a way that would not otherwise be possible. In Chapter 3 I study the other extreme and study investment by pension funds in venture capital firms that are just getting started. Understanding how pension trustees allocate their assets to venture capital can tell us a lot about investor behaviour in general. In Chapter 4 we look not at how pension funds invest their assets, but some of the results of firms choosing to have a defined benefit plan at all. We try to learn whether the trend towards conversion from defined benefit plans to defined contribution plans has any effect on corporate performance.

Stocks often increase in value when they are added to the S&P 500 and often lose value when they are dropped from the index. There are several hypotheses advanced for this effect. The information hypothesis is that when a firm is added this reveals positive information about the firm. Conversely when a firm is dropped this reveals negative information. This information effect would explain stock price movements associated with additions and deletions. However, I am able to examine several stocks that were dropped from the S&P 500 purely because they were foreign stocks. Accordingly, there should be no negative company specific implication regarding the stocks and the information effect should be absent. On July 9, 2002 S&P announced that it was

removing seven firms from the S&P 500 index only because of their nationality. I argue that this makes the event unique, and thus it can be contrasted with other deletions to isolate and test for the information effect. It is only when we get the chance to examine a natural experiment like the removal of foreign firms studied here that one can isolate a single effect. Since these seven firms under study were removed solely to make the index officially "American", I argue that the information effect that may be present in normal deletions will not exist in these firms. All of the other effects mentioned above should, however, remain the same, thus any difference between the normal deletions and the foreign firms will be due to the information effect.

The primary goal of the next chapter is to determine whether or not there is an information effect in S&P stock deletions. I also contribute by studying in detail a unique event in recent markets where a major rebalancing has occurred for apparently exogenous reasons. This change in policy by S&P affects almost all institutional investors, indexers and mutual fund managers who are measured against the popular index.

I will examine two sets of hypotheses relating to pension funds. The first hypothesis to test is whether or not pension fund trustees use familiarity based variables to make their decision on whether or not to invest in venture capital. The primary variables of interest are geography and industry, both of which have been identified in previous literature, and both of which are identifying traits of venture capital in general and venture capital clusters in particular. This is an important issue both for trustees who have growing pension deficits to fund over a long time horizon, and to venture capitalists and entrepreneurs who require a steady long-term stream investment funds to grow.

The second hypothesis is that the choice of defined contribution plans over defined benefit plans has a positive effect on corporate performance. The literature has not, to our knowledge, identified improved corporate financial performance as an important motivation for or effect of conversion to DC pension plans. Therefore, we regard drawing attention to the possible effect of pension plan structure on corporate performance as measured by the return on assets as a valuable contribution of this paper. With the large and accelerating shift to DC plans over the last 25 years this will have ramifications across many sectors of the economy. It can also have an effect on the management of public pension plans which are backed by the taxpayer, and which are almost exclusively DB plans to date.

Chapter 2 - Distilling the Information in S&P 500 Deletions

2.1 Literature Review

It is a well accepted fact in the literature and in practice that when a company is added to the S&P 500 index, there is an immediate price increase on the day of announcement, followed by a gradual increase until the effective date. On the opposite side when companies are removed there is a corresponding decrease in price. There is, however, some disagreement on whether this effect is short term or permanent, though on balance it now seems accepted that at least some of this effect is permanent. This seeming anomaly has been examined in many ways, but until now there have been few ways of directly testing the information effect.

There are several different theories that have been put forward to explain this phenomenon.

The *indexing pressure* hypothesis[1] says that when a stock is added to the index, there is a large short-term increase in demand from index funds. This requires index funds to pay a premium to compensate sellers for selling at an earlier time than they might have wanted.

The *liquidity* hypothesis says that since the liquidity of a stock increases after addition to the index, people are willing to pay a premium for the stock.

---

[1] This is often called the *price pressure* effect. *All* of the effects that we are testing are price effects, so this term is somewhat misleading.

Finally, the *information* hypothesis says that there is new beneficial information revealed by addition to the S&P 500 index which causes the price to rise.

Of course these hypotheses are not mutually exclusive. It is likely that each one occurs to varying degrees. The problem has been the difficulty of disentangling all of these effects. Many authors have had to assume that one or more hypotheses are false in order to test another.

The indexing pressure hypothesis says that when a stock is added to the index, there is a large short-term increase in demand from index funds. Index funds have grown immensely since Vanguard started its first index fund, the Vanguard 500 Index Fund, in 1976. It is currently estimated that assets directly indexed to the S&P 500 (excluding "closet" indexers) are in excess of US$1 trillion. Since index fund managers are concerned only with tracking error, and not out-performance of the index, they are forced to buy new stocks as close to the inclusion date as possible. Consistent with this hypothesis, Harris and Gurel (1986) find that there is a short-term rise in price after inclusion which is reversed within 30 days. Interestingly the opposite of this reversal has been found in studies conducted by Beneish and Whaley (1996,1997), Dhillon and Johnson (1991), Lynch and Mendenhall (1997), and Shleifer (1986) who all find that while some of the price effect is reversed there is still a permanent price increase associated with inclusion in the index. Indexing pressure rather than price pressure seems especially apt given the finding with perhaps the only index more widely quoted than the S&P 500 – the Dow Jones Industrial Average (DJIA). This index of 30 stocks, while widely quoted, has very few assets indexed to it. Beneish and Gardner (1995)

showed that there was no short term price or volume effect associated with additions to the DJIA.

A new spin on the indexing pressure hypothesis is proposed in a recent paper that looks at a new argument to explain the price swings after addition to the S&P 500 Index – taxes. If the price is pushed upward on index inclusion in order to compensate investors for selling now instead of later, then it is reasonable to expect that if an investor also has to trigger capital gains taxes on this sale, the price pressure will be even higher. Blouin, Raedy and Shackelford (2000) examine this effect through the relationship between price appreciation on index inclusion and expected tax effect (this is a combination of recent price activity and the difference between long and short term capital gains taxes). The authors find that there is, as expected, a positive effect relating incremental taxes and price movement and estimate that this accounts for a 0.5% abnormal return, a small portion of the daily price increase observed. There are several issues with the estimation of the magnitude, some of which the authors point out. For example, much of the press about indexing concerns such mutual funds as Vanguard or from Exchange Traded Funds like Standard and Poor's Depository Receipts which track the S&P 500, which are owned by individual taxable investors. Anecdotal evidence suggests that there is at least as much indexing occurring in institutional portfolios, not even considering "closet indexers" who would also react to the changes. Many of these such as pension funds and endowment funds are tax-free, and so would not be affected by capital gains taxes[2].

[2] As the authors point out the marginal retail investor, through taxable mutual funds, would still have an effect. However, it is not clear that mutual fund managers would worry about the tax situation as much as the end investor.

One element that all earlier studies of this topic have in common is the use of closing price data. Beneish and Whaley (1996, 1997) take a new tact by using intraday data. They examine whether or not there are still profits to be made by "playing the S&P game", or trying to profit from index changes at the expense of inflexible index funds. The good news is their primary conclusion that the profits available have decreased since 1989 as more and more people become aware of the opportunity, even when allowing for the number of days between announcement and effective dates. Although it is not mentioned, this effect has to combat the opposite force of the rapid increase in indexing activity. So even though there are more people (and presumably more money) chasing the "arbitrage" profits[3], this has to combat even more index fund money. The new intraday data and methodology employed is also unique. Since all the information is released the night before, one would expect most of the price effect to occur overnight rather than be spread out through the first day or over the period between the announcement and effective dates of the index change. Unfortunately with daily closing prices it is impossible to test this. To separate the overnight effect from the rest of the first day effect Beneish and Whaley look at the price change from the close the previous night to the opening trade. One new measurement issue that is dealt with is that the S&P index is not directly tradable on the open. Each stock is opened at a different time, so that there is no official index level to compare with. To overcome this problem the authors use the S&P futures as a benchmark. To determine the profit available from buying companies added to the index they take a long position in the stock, and a

[3] Technically this is not an arbitrage profit, though it is often referred to as such. As we will see a few companies do not follow the trend except on the first day after announcement where every company decreases as expected.

corresponding short position in the S&P 500 futures contract. Although this ignores the basis of the futures contract, the strategy *is* tradable, and the basis effect is likely negligible. They find that this overnight effect is 3.06%. This return is, however, not tradable since the markets are closed[4]. The abnormal return from this first trade until the close of the first day is 3.70%. This return is available to quick traders.

A relatively new idea in the finance literature is that liquidity can move asset prices. If a stock is easily traded in large size at a reasonable price close to the prevailing bid and ask prices then an investor might rationally be willing to pay a premium for the stock over an otherwise comparable stock that was difficult or expensive to trade. Hegde and McDermott (2000) test for liquidity changes in stocks being added to the S&P 500. They document a permanent increase in liquidity measured by decreased effective spreads, increased quoted depth, increased trade frequency, as well as the well documented increase in volume. As Amihud and Mendelson (1986) show, the value of any firm is the discounted present value of future cash flows minus the discounted present value of future transaction costs. If costs – either direct costs like commissions or indirect costs like bid-ask spread or market impact – rise as a result of deletion from the index, then we would expect the price of the stock to decline in response.

The hypothesis I focus on in this paper is the information effect. As Harris and Gurel (1986) state before arguing against an information effect: "The no-information assertion is essential to any interpretation of the post-announcement price change". Does S&P have some skill in picking companies using public information that others do not have or

---

[4] Over recent years after hours trading has been a fast growing, though still illiquid option. Some of these companies are also traded on foreign markets. Both of these alternative trading methods should also be examined in future research.

do they have private information from their world leading bond rating service? There are, as in any interesting debate, arguments on both sides. Unfortunately until very recently very few papers have formally tested the thesis. There are several convincing arguments against any information effect. The most obvious is that there should be no extra returns that can be generated using publicly available information beyond the cost of trading and acquiring the information. The companies should already be fairly priced. A recent study by Denis, McConnell, Ovtchinnikov and Yu (2002) examines the relationship between analysts earnings expectations, realizations and inclusion in the S&P 500. They discover by looking at matched firms that are and are not included in the index that earnings expectations increase after inclusion in the index, as do actual earnings. There are two possibilities. One is that addition to the index actually improves either the quality of the firm, perhaps by making its name known and acceptable to more potential clients[5]. Also the extra information generated by more analyst coverage (which is confirmed by Denis et al.) and more fund managers looking at the firm will tend to increase the price regardless of the average expectations. Pruitt and Wei (1989) also show that there is more institutional ownership in the month after a stock has been added than the month before. This institutional following could lead the analysts to cover the companies more, and help to decrease monitoring and agency costs. The opposite possibility, and the one that I will assume here is that S&P has some ability or skill in picking higher quality firms to be included (or lower quality firms to be excluded). One reason to expect that they may have superior skill in picking firms, or even possibly non-public information, is that S&P is one of the biggest bond-rating firms in the world, and

[5] Shaw is an example of a Canadian company that uses index inclusion in its marketing material.

is likely to have rated the debt of any firm eligible for inclusion in the index. Information from that process may well be used for creating index products.

While this information effect is interesting in its own right, it is a key assumption used when estimating whether or not stocks have downwards sloping demand curves. In normal day to day market watching, it is not possible to tell whether the demand is perfectly elastic or not. Previous studies mentioned above have used initial and secondary equity offerings and S&P additions to conduct this test. Unfortunately they all assume that there is no information effect at all. One paper which does not require this restrictive assumption is Kaul et al. (2000). They examine another natural experiment which occurred in 1996 when the weights in the TSE 300 were recalculated after the change in the definition of float[6]. In this case the event was advertised many months in advance, so there was no actual information in the event. On the day of the re-weighting the 31 stocks which had increased weights experienced excess returns of 2.3%. This was not merely a liquidity effect since the gains were not reversed even after the volume had returned to normal.

The earliest theory proposed that is consistent with a permanent price increase is the imperfect substitutes hypothesis of Shleifer (1986) and Scholes (1972). Under the assumption of perfect frictionless markets, demand functions for stocks are perfectly elastic. So if a stock is priced $.01 above (below) fair value, investors will sell (buy) an infinite amount of it, and can then form a portfolio of assets that has the exact same risk and return characteristics to earn an arbitrage profit. Unfortunately markets are not

---

[6] Two other natural experiments using the S&P 500 are Goetzmann and Garry (1986) who look at the deletion of 7 utility companies to make room for the "Baby Bells", and Graham and Pirie (1994) who look at the rebalancing that was required after the deletion of RJR/Nabisco.

perfect and this arbitrage profit is not always available. This means that demand curves will actually slope downwards for stocks, and when a stock is added to the index, the demand curve gets shifted to the right, and thus the equilibrium price increases. Wurgler and Zhuravskaya (2002) look in detail at some of these imperfections in a way that was not possible when they were first suggested. They create portfolios of stocks that are near matches of the stocks under investigation to measure how well they might be arbitraged. Surprisingly, even for the median stock in the sample, only a quarter of the daily variation can be explained. This effect is far more notable for large stocks than small stocks. They find that price increases on inclusion are increasing in the ease of arbitrage, strongly supporting the contention that at least some of the price jump is due to a downward sloping demand curve. Even though the predominant effect of deletion is a *shift* in the demand curve up or down rather than a shift *along* the demand curve, absence of information effects are precisely the reason that S&P 500 index changes were chosen as a laboratory to test demand curve slopes.

There have been several studies done on international markets to test the same hypotheses as have been done on the S&P 500. One theory that is arguably invalid in international markets but which could occur in the S&P indices is the information effect. Most international markets use an objective rules-based index construction. In these indices there is a set review on regular dates, either quarterly or annually. There is a publicly known list of criteria which must be met to select a fixed number of stocks. In the FTSE 100, for example, reviews are conducted quarterly. Any company in the top 90 by market capitalization is automatically in, those below 110 are removed from the list, with the balance made up of firms ranked between 90 and 110. Mase (2002) finds that

there is a much smaller effect in the FTSE 100 than in previous studies of the S&P 500. None of these international studies have focused on companies that are primarily located in one country but members of an index in a second country.

2.2 Data

2.2.1 S&P 500 Index construction

Companies are only added to the S&P indices when another company needs to be removed. Firms are removed primarily due to mergers, acquisition or a restructuring that significantly alters the company and makes it no longer suitable[7]. Some firms are also removed when they no longer meet the criteria outlined below in some significant way. Although some firms may seem to be "temporarily" unsuitable, this does not automatically trigger a removal.

S&P has very specific criteria for firms to be added, though it seems that these change over time. The index committee at S&P maintains a proprietary list of companies that are considered to be suitable for inclusion. This list is updated periodically so that when a company needs to be removed, a replacement can be chosen quickly. The selection criteria for inclusion are:

1. The company must be American (this was first added at the time of the event under study in this paper).
2. Liquidity and price. The current liquidity definition states that the dollar volume traded per year must be at least 30% of the market cap. Low dollar value priced companies are not eligible.

---

[7] Bankruptcy, though rare for S&P 500 companies, would count as a restructuring. Spin-offs of subsidiaries also call for a re-evaluation, though in some cases both firms remain in the index.

3. Market capitalization. The current threshold for the S&P 500 is $3 billion. It must also have a traded float of at least 50% of the company's equity.
4. Sector balance. S&P endeavours to have its index match the performance of the economy in general as well as the stock market.
5. Operating company. Holding companies, partnerships, closed-end funds, and royalty trusts (except REITs) are not eligible.
6. Financial viability.

It is this last point specifically that those supporters of an information effect believe is most important. In their latest report on the topic, S&P states that financial viability means 4 quarters of consecutive positive net earnings. However as Beneish and Whaley (1996) suggest "the firm's financial and operating conditions are rigorously analysed to ensure that added firms will have longevity". This view would also be supported by S&P's belief "that unnecessary and excessive turnover in index membership should be avoided when possible". The index committee would thus understandably prefer companies that had good prospects to avoid need to delete the firm in the near future.

Prior to October, 1989, S&P announced any changes to the S&P index one evening, and then the changes became effective the next morning. However, in order to alleviate some of the trading imbalances they decided to pre-announce changes. The standard policy is to announce the changes 5 days in advance (the mode of the distribution charted in Figure 1). Circumstances do not allow this policy to always be followed, however. As can be seen from Figure 2-1, while approximately one quarter of the deleted companies in my sample do have a 5 day window, there is a range from 1 to 19 days.

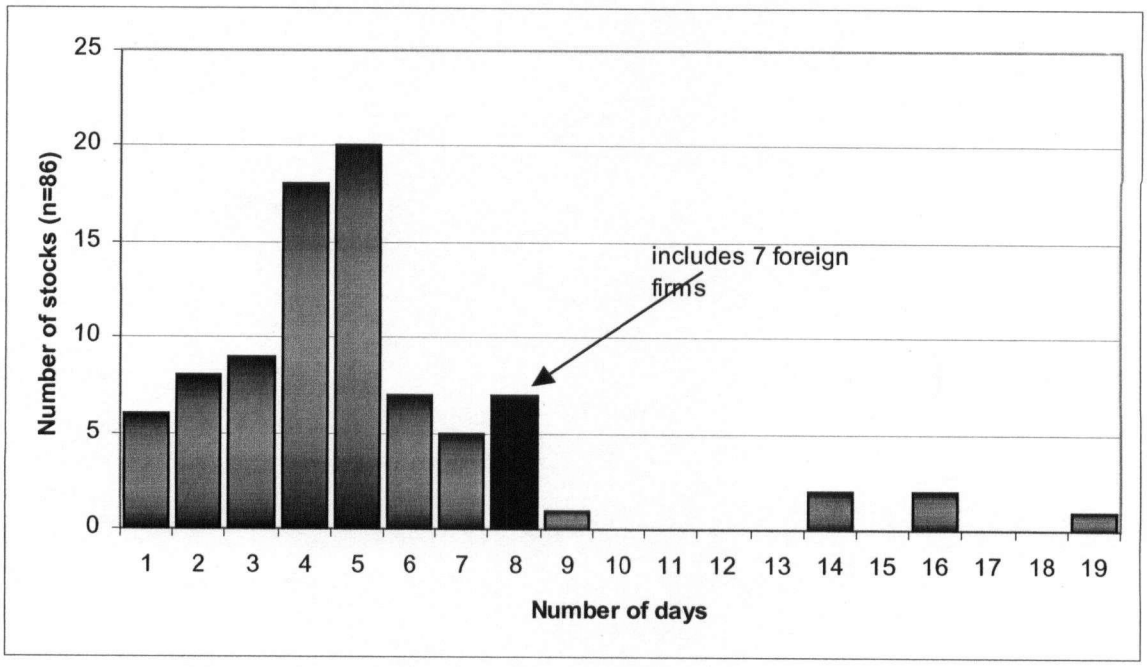

Figure 2-1 - Number of days between announcement and actual removal (effective date) of a company from the S&P 500 - Although the stated target since 1989 is to give 5 trading days notice, this is not always possible or desirable.

2.2.2 What is the event we are studying and why is it useful?

On July 19, 2002 Standard and Poor's announced that seven foreign firms would be removed from the S&P 500 index. The 5 Canadian firms were Alcan, Barrick Corp, Inco, Nortel Networks and Placer Dome. The 2 Dutch firms were Unilever and Royal Dutch Shell. S&P made the decision to ensure that the index was a proper representation of American firms and the American economy. All of the firms are still members of their respective country's S&P index, and of the S&P 1200 World index. There do exist some index funds following these indices, especially the Canadian and Dutch country indices. The demand by American index funds, and the assets benchmarked to the S&P 500 dwarfs these amounts and is likely to be the dominant force.

This event provides perhaps a unique chance to test once and for all whether or not S&P 500 inclusion provides new information to the market about the companies it chooses. Although the information hypothesis is interesting in its own right, it is also important for tests of the whether or not the demand curve for stocks is downward sloping. A fundamental assumption of many of the tests that use the S&P 500 index additions and deletions as a control event is that there is no information in the event. If there is, then these effects will generally be confounded and there will be no way of disentangling them. Our natural experiment, however, gives us exactly that opportunity. Any effects which are observed in the normal deletions but not in the foreign sample can be interpreted as information effects.

2.2.3 Sample Selection

I obtained the names and dates for all companies deleted from two primary sources. For 1990-2000 I used the data provided by J. Wurgler on his website[8]. From 1993-2002 I received dates and names from Standard and Poor's. Both of these sources provided a list of companies to begin my search. I searched the Lexis-Nexis for Standard and Poor's daily notices for announcements of any changes to the S&P 500, and cross-referenced these notices with a search of the Wall Street Journal. Any companies that were removed from the index due to merger, acquisition, spin-off or bankruptcy were removed from the sample. There is clearly confounding information in these announcements, and in most cases the company is no longer trading after the effective

[8] http://pages.stern.nyu.edu/~jwurgler/

date. There were 5 companies that entered Chapter 11 proceedings or did not have complete CRSP data within 60 trading days after being removed from the S&P 500 which I did not include in my analyses[9]. This left me with 79 companies for the normal deletions, plus the seven foreign firms that I am investigating. The list of the 86 firms along with brief reasons given for the deletion is found in the Appendix.

As indicated in section 2.2.1 above, the procedures around firm removal were changed in 1989. In order to study only a consistent regime I will only look at deletions beginning in 1990. Figure 2-2 shows the number of companies in the study by year of announcement.

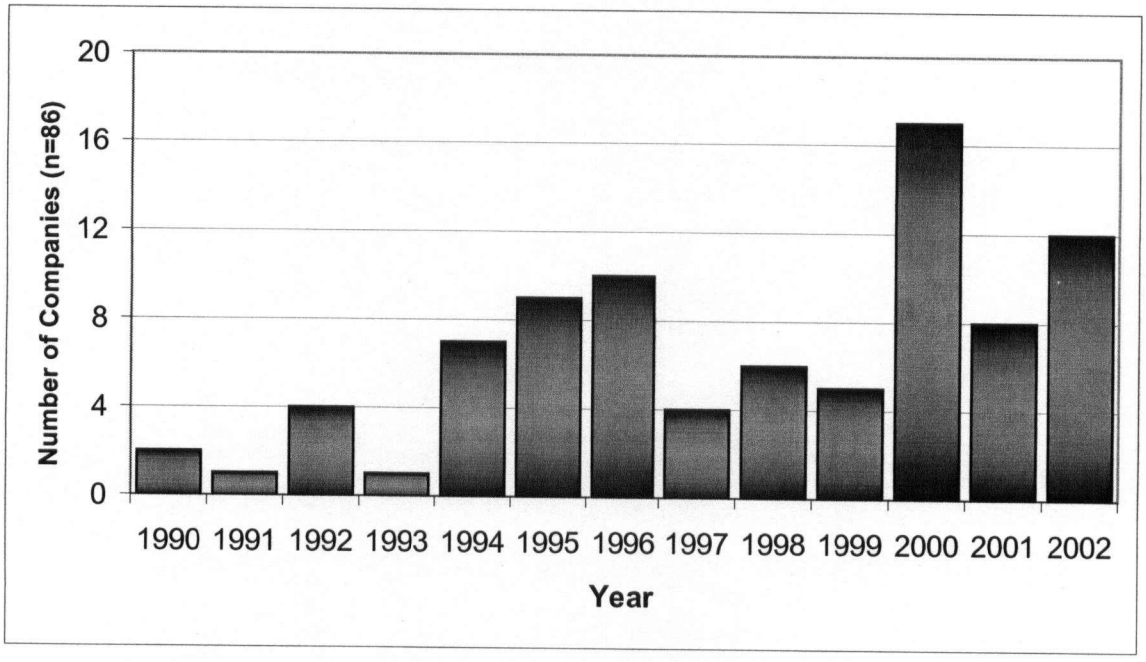

Figure 2-2 - Number of companies in total sample (normal deletions and foreign firms combined) broken out by year of announcement – There is a clear upward trend in the turnover in the S&P 500.

Price and volume information for all companies and indices was obtained from the Centre for Research in Security Prices (CRSP).

---

[9] Conoco, Conseco and Enron, which were all removed in 2002. WorldCom was suspended from trading for 3 days and Owens Corning for 1 day.

2.3 Results

2.3.1 Some summary data on the foreign firms removed

Unlike other deletions in the past (with the exception of the AT&T break-up in 1983[10]), S&P provided a great deal of information about the companies that were removed as part of a 23 page press release on the subject. They also published 2 research reports arguing that there would be no effect from the change (Dash (2002) and Dash and Ruotolo (2002)). I argue that the great trouble they went to with these companies supports the idea that unlike the normal deletions, there is no adverse information in the foreign deletions under study. Table 2.1 contains some background on the seven firms.

Table 2.1 - Background data on seven foreign firms studied - On July 9, 2002, S&P announced that it would be removing all foreign firms from the S&P 500 index in order to make it more representative of the American economy and a better benchmark for American investors. The following 7 companies were removed and form the basis for the natural experiment being studied in this paper.

Company	Royal Dutch	Unilever	Nortel Networks	Alcan Inc	Barrick Gold	Placer Dome	Inco, Ltd
Ticker	RD	UN	NT	AL	ABX	PDG	N
Country	Holland	Holland	Canada	Canada	Canada	Canada	Canada
Added to index	1957	1961	1978	1935	1993	1987	1940
Index weight	1.19%	.38%	.07%	.12%	.12%	.05%	.04%
Index Rank	15	57	286	172	180	373	381
Industry	Oil and Gas	Consumer Products	Telecom	Base Metals	Gold	Gold	Base Metals

[10] See Goetzmann and Garry (1986) for more details.

These seven foreign firms were replaced by UPS, Goldman Sachs, Prudential Securities, Electronic Arts, Principal Financial, Sungard Data and eBay. Given the concentration of the new companies in very different industries from the foreign firms, it is possible that rebalancing was also a goal of this action, or at least the timing.

2.3.2 Methodology

The basic event study methodology follows Campbell, Lo and MacKinlay (1997), with an extension described in Lynch and Mendenhall (1997) to allow for 2 event dates with different times between the two dates for each security. As mentioned before, there is no set time between announcement of removal from the index and actual removal.

There are two important dates – T_1 is the day after the announcement. The announcements are made after the close of the market, so the return the day after is the one of interest. T_2 is the date after the removal, i.e. the first day it is traded without being a member of the S&P 500 index. In our sample, the time between these two dates (T_2-T_1) ranges from 0 days to 19 days with an average of 5.1 days. The mode, as expected, is 5 days. I use the period from 200 trading days before the initial announcement to the 31 days before the announcement as an estimation window[11]. None of the previous studies suggest that any information is available prior to the announcement, or that any of the

[11] These tend to be stocks that have performed poorly, so beta will be low when markets are good. Average beta in initial sample is 0.9, though there is a large variation between companies. This should get picked up if we add a momentum factor, and in theory even if we don't, unless momentum is priced. It may be, however, that it is only the alpha that is negative, and not the beta which is low. We will also test with longer beta estimation periods in Section IV.D, but the results are qualitatively unaffected.

reaction occurs before this date in returns or in volumes. In this period I estimate the market model[12]:

$$r_i = \alpha_i + \beta_i r_m + \varepsilon_i$$

for each of the securities in the normal deletions and the foreign firms. I use the value weighted return on the S&P 500 to represent the market, given the size and nature of the companies that I am studying. The average beta in the normal deletions is 0.9, and it is 0.6 for the seven foreign firms.

Given the estimates of a and ß I can calculate the excess or abnormal return AR_i in subsequent periods as:

$$AR_i = r_i - (\hat{\alpha}_i + \hat{\beta}_i r_m)$$

This will give the excess return for a given day. To analyse the effect of the deletion event we need to combine the returns over the entire event window. The cumulative abnormal return $CAR_i(T_1,T_2)$ is defined as the sum of all excess returns over the window of interest. The average abnormal return $AAR_i(T_1,T_2)$ is defined as the $CAR_i(T_1,T_2)$ divided by the number of trading days in the period (T_2-T_1+1). To test for effects with several firms I construct two average measures – MAAR and MCAR, which are simple averages of AAR and CAR. So for N companies,

$$CAR_i(T_1,T_2) = \sum_{t=T_1}^{T_2} r_{it} - (\hat{\alpha}_{it} + \hat{\beta}_{it} r_{mt})$$

[12] We could also think of ß as a vector to follow essentially the same methodology for a Fama-French or other multi-factor model as in Brooks, Kappou and Ward (2003).

$$MCAR = \frac{1}{N}\sum_{i=1}^{N} CAR_i(T_1, T_2)$$

$$MAAR = \frac{1}{N}\sum_{i=1}^{N} \frac{1}{(T_2 - T_1)_i} CAR_i(T_1, T_2)$$

When there are the same number of days in every event window as in most event studies, MCAR and MAAR produce the same results. However when there are different numbers of days across firms – for example the window between announcement and effective date in our study - the MAAR will give greater weight to firm-days where the window is short, whereas the MCAR gives identical weight to all firm-days. In this case the T_2-T_1 would be different for each firm. MAAR is only presented in those cases where it provides useful information. Otherwise all results reported are MCARs. To determine whether any of the observed MCARs are statistically different from zero, we need to estimate the standard errors. A detailed description of this is available in Campbell, Lo and MacKinlay (1997), page 159. The important part to recognize is that there are two aspects to the estimation error. The first is the estimation error inherent in estimating the betas. This will be different for each of the stocks in the sample. To determine the daily error estimate for each firm, this must also be divided by the number of days in the window. This is the same for all firms except in the period between the announcement date and the effective date. During this period each firm will have its own denominator. These error estimates are summed up across all firms and then divided by the number of firms. The second aspect is the sampling error which I assume will be the same across all firms.

We assume, as per standard methodology that there is no correlation between the firms in the sample, or more particularly between the error estimates. When the events

are separated in time this is generally a good assumption. In the case of the foreign firms, all of the events are on the same day. I calculate a correlation matrix between the 7 firms. All of the estimated correlations are below 0.25 with the exception of Placer Dome and Barrick, two gold companies. The effect of an unaccounted for correlation would be to artificially increase the t-stats. Given the very high significance of our results it is highly unlikely that it would change the results. As a second test, the robustness check in Section 2.3.4 calculates abnormal returns using a different method. In particular since the benchmark is another gold firm in each case, any factors common to the two firms that was correlated should be properly removed.

To determine whether or not there is an information effect, I do a two-sample t-test to test for differences in the MCAR/MAAR. The information hypothesis predicts that the MCAR/MAAR should be negative for both groups of stocks, but *less* negative for the foreign sample on removal from the index. Our null hypothesis is that there is no difference between the normal deletions and foreign deletions. If we do see a difference, then we would take that as evidence to reject the null hypothesis that there is no information effect. I also examine each event window to see if the MCAR for normal deletions and the foreign firms different from zero.

2.3.3 Price results

There are several event windows that we can look at to provide some insights. A numerical summary of the results can be seen in Table 2.2. The first event window that I look at is the period leading up to the announcement. Previous authors with only one exception have not studied this period under the assumption that the information is

private to the index committee. This is unusual since in most event studies including, but not limited to takeovers, earnings announcements or restructurings the balance of evidence suggests that there is, in fact, a run-up effect to many of these events. Only Liu (2000) examined this time period, but for the Nikkei 500 index. He looked at each of the 15 days preceding the announcements and found that none of the individual days was significant. However, it is important to look at the "big picture", or the cumulative effect. I have chosen to look at the 30 day period immediately preceding the deletion announcement. Figure 2-3 shows the CARs for both the normal deletions and the foreign deletions. Given the lack of attention by previous authors, I was surprised to discover an MCAR of –11.0% (t-stat of –5.0), with a consistent downward trend for the entire period. For the foreign firms a much smaller MCAR of –4.1% (t-stat of –0.8) was observed, suggesting that this event was not foreseen by the markets. This result is not statistically significant, and there is no regular trend throughout the 30 day run-up window. So the normal deletions did indeed experience poor returns before the removal, but the foreign firms did not. What is not clear is the causality of the event. That is to say, were the companies removed *because* they did poorly in this short window, or did the market anticipate their removal and discount that possibility in advance[13].

Table 2.2 - Mean cumulative abnormal returns using the period between 200 and 31 trading days preceding the announcement date to estimate the beta of the market model - Numbers marked with a * are statistically significant at the 10% level, ** at the 5% level and *** at the 1% level. The difference between normal deletions and foreign firms is statistically significant for all event windows examined. The

[13] If we accept the second direction then this provides an interesting possibility for interpreting the results in much of the rest of the paper. If poorly performing and "small" firms in general are anticipated to be removed, that would not have been the case for the seven foreign firms. So one might argue that the larger foreign firms would have less abnormal losses in the run-up period as we observe, or less of a surprise, but then experience a greater negative return after the event to yield the same overall abnormal return. This interpretation would only strengthen the results below. Thanks to Alan Kraus for suggesting this line of reasoning.

difference between normal deletions and foreign firms is statistically significant for all event windows examined.

	Normal Deletions	Foreign Firms	Difference
30 days preceding announcement	-11.0%***	-4.1%	6.9%**
Announcement date	-8.2%***	-4.6%***	3.6%***
Announcement to effective date	-13.6%***	-10.7%***	2.9%***
Effective date only	-4.7%***	-1.5%*	3.2%***
10 days after effective date	6.9%***	-9.9%***	-16.8%***
30 days after effective date	13.7%***	-5.0%	-18.7%***
Total 30 days before to 30 days after	-10.9%***	-19.8%***	-8.9%**

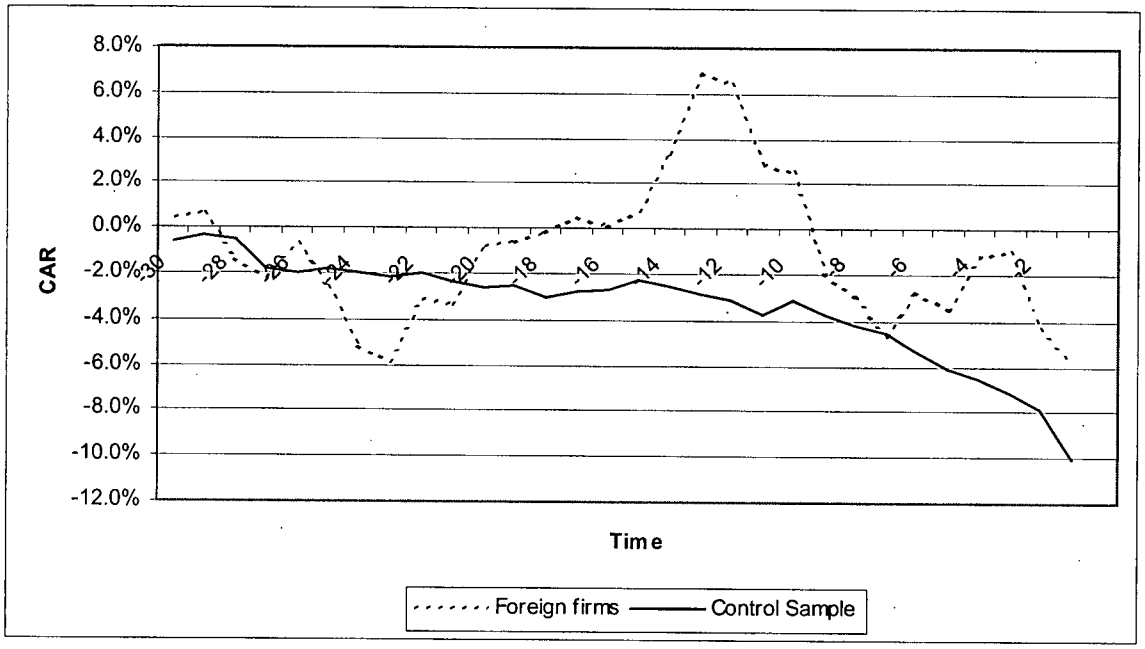

Figure 2-3 – Pre-announcement Abnormal Returns - This figure charts the Abnormal Returns of the normal deletions and the foreign stocks over the 30 trading day period leading up to the announcement of removal from the S&P 500. The anticipatory decline was surprising given that so few authors had reported results for this period before the announcement, and none were significant.

The first date that other authors have studied in detail is the announcement date (or, more precisely the first trading day after the announcement date). Under both the liquidity and the information hypotheses we should observe a negative return on this day. In both the normal deletions and the foreign firms we would expect liquidity to be a significant issue, but if there is an information effect, it should only appear for the normal deletions and not for the foreign firms. Any difference between the two would therefore represent the information effect that I am trying to detect and measure. For the 79 normal deletions I find a one day abnormal return of –8.2% (t-stat of -19.5), which is statistically different from 0. This is very similar in magnitude to the –6.26% reported by Lynch and Mendenhall (1997). For the foreign firms I find a one day abnormal return of -4.6% (t-stat of -5.2), which is also significantly different from 0. The two results are also statistically different from each other in means and medians which implies that we are able to reject the null hypothesis that there is no information effect, though it seems to be economically quite small.

The next window of interest is the effective date. If there were only index funds involved, or if only index fund managers knew that there was going to be a change or in the time before there was pre-announcement of changes, this would be the only important date to examine. This would then be effectively the same regime as pre-1989. Under both the liquidity and information hypotheses we would expect a price decline on this date. Even though most index funds spread their trades over several days, this introduces tracking error, which is the most meaningful way of measuring index manager performance. Of course it is not necessarily in the interest of shareholders or unit holders, but if liquidity is the only issue, then the prices should return to normal over

time. For the 79 normal deletions I find an abnormal return of –4.7% (t-stat of -11.3). For the foreign firms I get a one day abnormal return of only –1.5% (t-stat of -1.7). This suggests that all of the required trading was done before the effective date in the case of the deletion of the foreign firms.

The literature assumes that information effects should occur almost immediately, whereas liquidity effects can take some time just due to the large volume of trading required. Because normal deletions have different numbers of days between announcement date and effective date (not including the effective date itself) it is unclear whether AAR or CAR is more important, so I present both statistics for consideration. For example, when Microsoft was added to the index, there were 16 trading days before the effective date. This was to done to allow portfolio managers a longer time due to the very large amount of stock that needed to be traded. It is unclear whether or not all of the excess return would be spread over time, or whether there would be a greater or lesser overall effect. Looking at the total time between the announcement date and effective date the CAR for the normal deletions is –13.6% and for the foreign firms is –10.7%. Although slightly larger, this effect is very similar to the CAR of –12.69% reported by Lynch and Mendenhall (1997). The two groups are not statistically different in means from each other at the 5% significance level, but they are both statistically different from zero. This is inconsistent with the idea that the information effect is present in the normal deletions but not the foreign firms. To test whether this is possibly due to outliers I ran a Wilcoxon two sample test. The observed Z-score is 49, supporting the idea that the medians of the two distributions are different, with a median in the normal deletions of –10.5% and a median of -9.5% in the foreign firms. The MAAR for the normal deletions

is only –4.1%, while for the foreign firms it is –1.6%, which may simply represent the shorter average event window of the sample firms. It does suggest that increasing the length of the window from the current average 5 days may decrease some of the observed effects.

The last event window that I want to examine is the time after the companies have been removed from the index. We have seen that there is a strong negative reaction over the previous windows as expected. Based only on the previous information we would conclude that there is a significant information effect. What we want to see is if this reaction is reversed over time. Studies to date have offered mixed evidence. The balance of evidence currently supports that while there is some price reversal over time, there is also some permanent effect that does not get reversed. If there is new information revealed by the removal then we would not expect it to be reversed over time. Changes in price due to long-term shifts in the demand curve would also be a permanent component that would not be reversed, though that should be true for both foreign firms and normal deletions. We would, however, expect short term liquidity effects to disappear and be reversed. I look at windows 10 trading days after the effective date and 30 trading days after. These periods are overlapping, so some care needs to be taken interpreting the results, but these are similar to those used in the previous literature. Figure 2-4 graphs the CARs of both normal deletions and foreign deletions. In the normal deletions I find evidence to support that presented in previous studies. In the 10 days after the firm is removed, MCAR is 6.9% (t-stat of 5.2). After 30 days, the MCAR is 13.7% (t-stat of 6.3). Lynch and Mendenhall report a return of 4.6% in the first 5 days after the effective date, and then an additional 1.6% in the next 5 days. So while some of

the abnormal losses are reversed, there does seem to be a permanent decrease in the price of the stocks. This would be consistent with several of the hypotheses presented above including the information hypothesis and the market segmentation hypothesis. The 7.85%[14] could be considered the value of the liquidity premium that was demanded in the short run to absorb the excess supply of stock from indexers, while the approximate 13.5% difference would be the value to other sources. Surprisingly this result does not appear in the seven foreign firms examined. The MCAR over the 10 and 30 day windows is -9.9% (t-stat of -3.5) and –5.0% (t-stat of -1.0) respectively. I do not as yet have an explanation for this result.

[14] A loss of 11.27% from announcement to effective date, and then a rebound of 7.85% in the 30 days after the company has been removed from the index. 11.27%-7.85%=3.42%

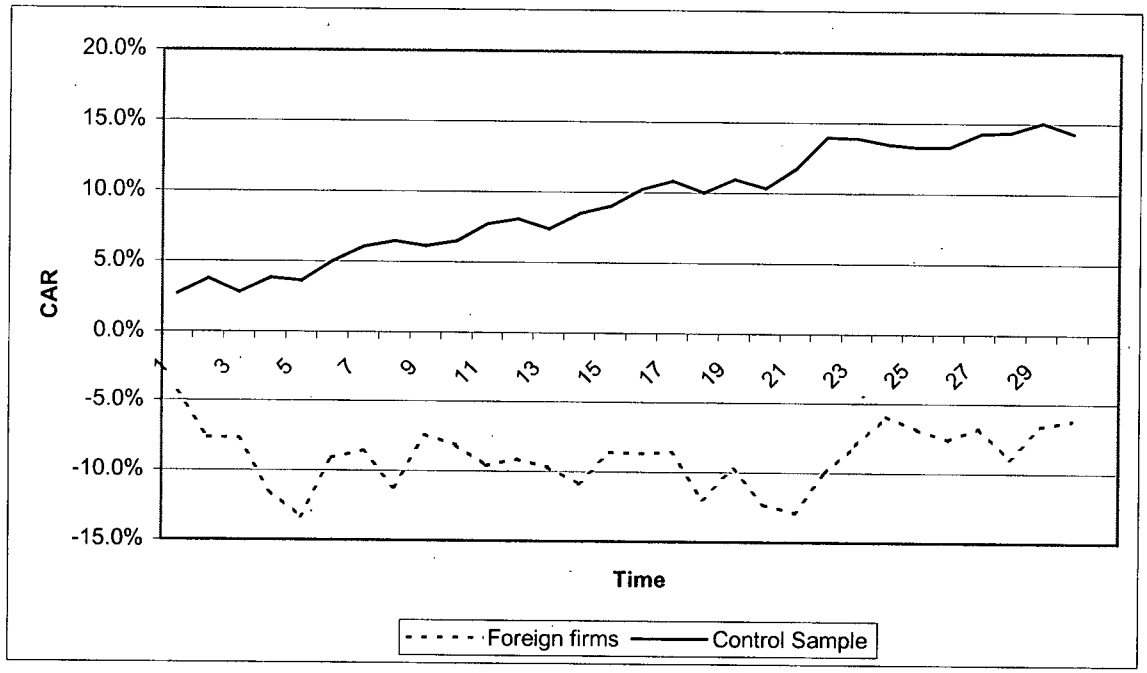

Figure 2-4 - Post-removal Cumulative Abnormal Returns from the day after the effective date until 30 days after removal - This figure charts the performance of both the normal deletions and the foreign firms for the 30 days following the effective date of removal from the S&P 500. The results from the normal deletions reaffirm previous research and reverse some but not all of the decrease on removal. A surprising result is that the foreign firms *continue to decline* after removal.

2.3.4 Robustness tests

Because the sample of deletions that I have is so small it is prudent to go beyond the standard event study methodology used in the previous section. Using MCARs with a small sample size as described above on our foreign firms has two potential drawbacks. Firstly, the seven foreign firms are considerably larger than any of the firms in the normal deletions on both an absolute basis and relative to other firms in the S&P 500 at the time of removal. Secondly, there is significant industry concentration in our small sample. Two firms are in the gold industry, two are in the basic metals industry and there is one each from the oil, high tech and food industries. With the first alternative test I will attempt to address these two problems. For each of our 86 firms I try and find a company in a similar industry and of a similar size. Starting with the entire CRSP stock universe I

sort by size within 3 digit SIC codes and choose the firm that was closest in absolute size on the day that the deletion was announced. In those cases where there are not at least five companies in the 3 digit SIC code I use the 2 digit SIC code. The abnormal return on any given day is defined as the difference between the matched firm and the firm which was deleted from the S&P 500. Calculation of CARs is then the same as previous tests. As can be seen in Table 2.3, while the point estimates are different, the qualitative results are unchanged from the standard methodology. In particular the foreign firms do not reverse the earlier price declines and have significantly worse performance over the 60 trading days surrounding the announcement and effective dates. This similarity of the two very different abnormal return methodologies provides strong evidence that the results are robust to different models.

Table 2.3 - Summary of size/industry matched abnormal returns - This represents the mean difference in performance between a firm in our sample and a firm matched by industry and market capitalization over the various windows of interest. Numbers marked with a * are statistically significant at the 10% level, ** at the 5% level and *** at the 1% level. The difference between normal deletions and foreign firms is statistically significant for all event windows examined.

	Normal Deletions	Foreign Firms	Difference
30 days preceding announcement	-10.2%***	–5.4%	4.8%***
Announcement date	-8.2%***	–5.3%***	2.9%**
Announcement to effective date	-14.8%***	–12.7%***	2.1%**
Effective date only	-5.5%***	–0.4%	5.1%***
10 days after effective date	7.6%***	-7.9%**	-15.5%***
30 days after effective date	12.8%***	–5.5%	-18.3%***
Total 30 days before to 30 days after	-12.2%***	-23.6%***	-11.4%***

The second robustness check I make is changing the estimation period for ß in the market model. In the preceding tests I used the period from 200 trading days before to 31 days before the announcement date to estimate the beta. This is the most common period

to use. One concern might be that the risk profile and thus beta of the stock might change after removal from the index. The first alternate estimation period that I use begins 31 days after the effective date and ends 90 trading days later. I also use the 60 day period preceding the announcement date. As can be seen from the results in Table 2.4 and Table 2.5, this makes little qualitative difference to the results.

Table 2.4 - Mean cumulative abnormal returns using the period between 60 and 1 trading days preceding the announcement date to estimate the beta of the market model - All differences between the normal deletions and the foreign firms are statistically different using a 2 sample t-test. The difference column is the foreign firm CAR minus the CAR of the normal deletions. Numbers marked with a * are statistically significant at the 10% level, ** at the 5% level and *** at the 1% level. The difference between normal deletions and foreign firms is statistically significant for all event windows examined.

	Normal Deletions	Foreign Firms	Difference
30 days preceding announcement	-10.6%***	–5.4%	5.2%***
Announcement date	–6.4%***	–5.3%***	1.1%**
Announcement to effective date	–14.1%***	–12.7%***	1.4%*
Effective date only	–4.7%***	–0.4%	4.3%***
10 days after effective date	6.3%***	-7.9%**	-14.2%***
30 days after effective date	13.5%***	–5.5%	-19.0%***
Total 30 days before to 30 days after	-11.2%***	-23.6%***	-12.4%***

Table 2.5 - Mean cumulative abnormal returns using the period between 31 and 90 trading days after the announcement date to estimate the beta of the market model - Numbers marked with a * are statistically significant at the 10% level, ** at the 5% level and *** at the 1% level. The difference between normal deletions and foreign firms is statistically significant for all event windows examined.

	Normal Deletions	Foreign Firms	Difference
30 days preceding announcement	-10.9%***	-2.2%	8.7%***
Announcement date	-8.3%***	-4.1%***	4.2%***
Announcement to effective date	-13.5%***	-9.1%***	4.4%***
Effective date only	-4.7%***	-1.0%	3.7%***
10 days after effective date	7.0%***	-10.3%***	-17.3%***
30 days after effective date	13.8%***	-6.2%	-20.0%***
Total 30 days before to 30 days after	-10.6%***	-17.5%***	6.9%**

Another unique aspect of this data compared to many event studies is that the group of companies being studied – the foreign deletions – all happen on the same day, whereas the normal deletions happen over a period of 12 years. One possible reason that we see smaller price declines in the foreign firms is that the effect has decreased overall through the years. To look at this I break out the abnormal returns by year of deletion. Figure 2-5 shows the average cumulative average returns of deleted stocks by year for both the overall return 30 days before the announcement date until 30 days after the announcement date and the shorter window between the announcement and effective dates.

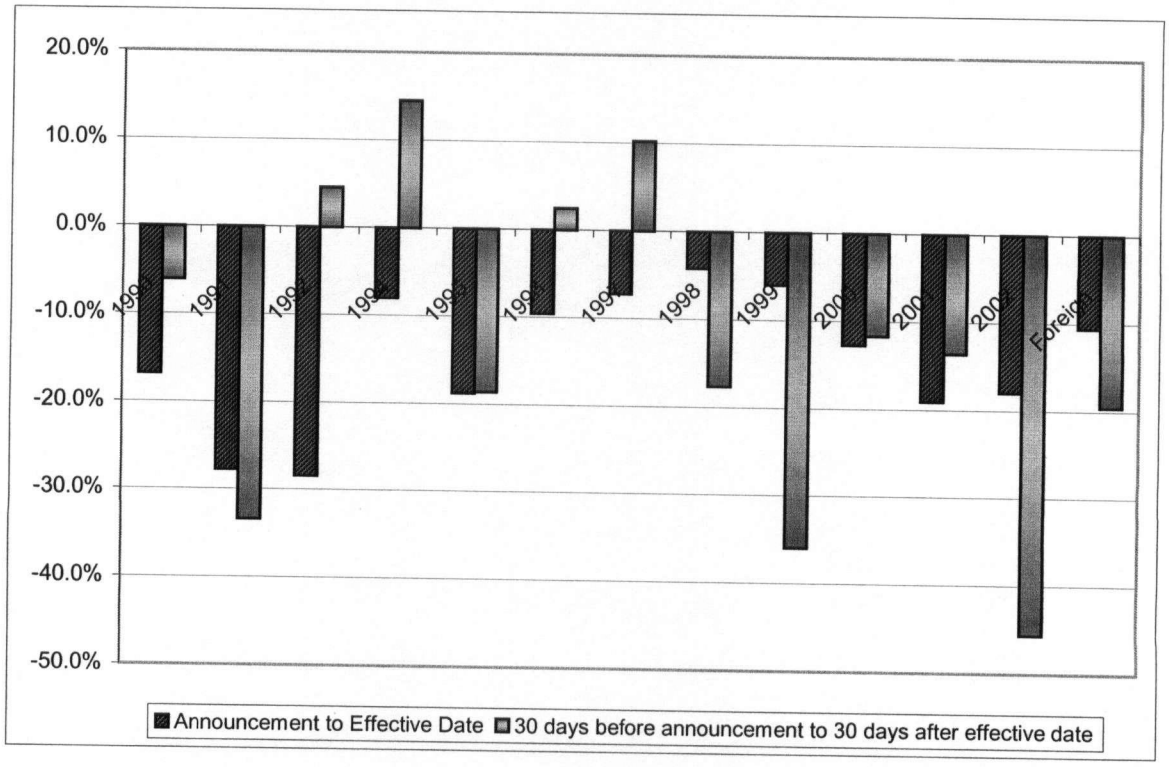

Figure 2-5 –Time Effects - Cumulative Abnormal Returns of deleted stocks by year for both the overall return 30 days before the announcement date until 30 days after the announcement date and the shorter window between the announcement and effective dates. There does not appear to be any continuous trend in either series or any of the other series not presented here.

There does not appear to be any continuous trend in either series or any of the other series not presented here. When I regress the abnormal return on the year of deletion then there is no evidence of a slope significantly different from 0. Although this differs from the result of Sui (2003), this is likely due to our small sample of deletions, whereas he concentrates on the more numerous additions. It does not appear that a shrinking price effect through time can explain the difference in performance between the normal and foreign deletions during our event window.

The final robustness check does not relate to this specific natural experiment, but enables us to make further inferences about the information hypothesis using the

deletions data. In the previous results we are limited by the very small sample size – only 79 companies in the normal deletions group and just seven foreign firms in the other group. There are several other ways to group the 86 firms. The one that I will focus on is whether or not the company remains in or is added to other S&P indices. One might think that if S&P was worried about the quality of the underlying firm then it would not add it to the S&P Mid-cap 400 or S&P Small-cap 600. If size were the only factor then it would be appropriate to put it in a smaller capitalization index. The seven foreign firms are all still part of their own national S&P indices and the S&P 1200 World Index. Of the 79 normal deletions, 7 were transferred to the S&P Mid-cap 400, 24 were transferred to the S&P Small-cap 600 and the remaining 48 were removed from all S&P indices. I hypothesize that if there is an information effect to inclusion in the S&P 500, then the abnormal returns on the firms that remain in an S&P index will be higher than on the firms that are removed from all indices. Table 2.6 summarizes the results of these tests.

Table 2.6 - Comparison between firms added to another index and those left out - Numbers marked with a * are statistically significant at the 10% level, ** at the 5% level and *** at the 1% level. There is strong evidence of an information effect since those firms that are deleted from the S&P 500 but remain in another index perform better than those which are removed from all S&P indices in all periods studied.

	Moved to 400	Moved to 600	Remained in 1200	All in indices combined	Not in any index
Number of companies	7	24	7	38	48
30 days preceding announcement	-4.7%	-0.4%	-3.9%	-1.8%**	-17.0%***
Announcement date	-4.0%***	-6.0%***	-4.5%***	-5.4%***	-9.7%***
Announcement to effective date	-8.7%***	-10.3%***	-10.6%***	-10.1%***	-15.8%***
Effective date only	-2.8%***	-2.6%***	-1.5%*	-2.4%***	-5.9%***
10 days after effective date	5.4%**	4.1%***	-9.9%***	1.7%***	8.0%***
30 days after effective date	6.3%	8.9%***	-5.0%	5.8%***	17.1%***
Total	-7.1%	-1.8%	-19.5%***	-6.1%	-15.7%***

The most interesting and striking results are those comparing the 48 firms which were not in any S&P index after deletion from the S&P 500, with the results for those which remained in one of the three S&P indices. The firms which remain in an S&P index outperform those which do not remain over all important event windows. On the announcement date the remaining firms declined an average of 5.4% while those not remaining in an index decreased by 9.7%. Between the announcement and effective dates (inclusive) the remaining firms decreased by 10.1% while those not remaining decreased by 15.8%. While those not remaining do rebound 17.1%, more than the 5.8% for the remaining firms, this is because they decreased even more. Looking from 30 days before the announcement (the lead-up period) until 30 days after the effective date those remaining firms decreased by (a non-significant) 6.1% while those not remaining decreased by 15.7%. The difference of 9.6% is significant at the 1% level. This result provides strong evidence of a long-term information effect. For those firms that remain in an index there is no significant long term effect. Although the sample is small (only 38 firms) this suggests that the information effect not only exists, but is significantly more important than the long-term liquidity effect.

2.4 Conclusions

In this paper I looked at abnormal returns surrounding the removal of seven foreign firms from the S&P 500 in July 2002 and compared them with the "typical" removal over a 12 year period. This event provided a unique opportunity to separate the information effect from the other effects that have been proposed for S&P deletions. In keeping with the findings of past literature I find an abnormal negative return between the

day of announcement and the actual day that the company is removed from the index. During this period the stock price of the normal deletions decreased by more than the seven foreign firms on the announcement date, the effective date and the time period in between. Because the seven foreign firms were removed for exogenous reasons, the difference between the returns is the expected value of information divulged by S&P. This is consistent with theories that suggest that there is information content in S&P deletions.

The surprising, and as yet unexplained result, came after the companies had been removed. In the normal deletions, I confirmed previous research which showed that some, but not all, of the losses are recaptured within the 30 trading days after the event. However with the foreign firms, not only does the loss not reverse itself, the companies' prices *continue to decline*! Even though the decline is not significantly different from zero, it is clearly a different result from what had previously been observed in all studies of S&P 500 additions and deletions. All of the above results were robust to different methods of estimating market betas and using size and industry matched firms as the benchmark asset instead of the market, and are unaffected by changes in the index effect through time.

Even stronger evidence of the information effect came from splitting the 86 firms along a different dimension. When I split the firms based on whether or not Standard and Poor's transferred the stock to a different index or removed it from all S&P indices I find some striking results. Those firms that remain in an S&P index perform significantly better (decrease in price less) than those which are not added to another index. This provides strong support for an information effect.

2.5 References

Amihud, Y. and Mendelson, H. 1986, 'Asset pricing and the bid-ask spread'. *Journal of Financial Economics* 1, 3-40

Beneish, M. and Gardner, J. 1995, 'Information costs and liquidity effects from changes in the Dow Jones Industrial Average list'. *Journal of Financial and Quantitative Analysis*, 30 (1), 135-157

Beneish, M. and Whaley, R. 1996, 'An Anatomy of the "S&P Game": The effects of changing the rules', *Journal of Finance* 51, 1909-1930

Beneish, M. and Whaley, R. 1997, 'A Scorecard from the S&P Game', *Journal of Portfolio Management* Winter 1997

Blouin, J., Raedy, J. and Shackelford, D. 2000, 'The impact of capital gains taxes on stock price reactions to S&P 500 inclusions', Working paper

Brooks, C., Kappou, K. and Ward, C. 2004, 'Gambling on the S&P 500's Gold Seal: New Evidence on the Index Effect", Working paper

Campbell, J., Lo, A. and MacKinlay, A. 1997, *The Econometrics of Financial Markets*, Princeton University Press

Dash, S. 2002, 'Price changes associated with S&P 500 Deletions', *Standard and Poor's*

Dash, S. and Ruotolo, M. 2002, 'Deletion of Canadian stocks from the S&P 500', *Standard and Poor's*

Denis, D., McConnell, J., Ovtchinnikov, A. and Yu, Y. 2003, 'S&P 500 Index Additions and Earnings Expectations', *Journal of Finance* 58(5), 1821-1840

Dhillon, U. and Johnson, H. 1991, 'Changes in the Standard and Poor's 500 list', *Journal of Business* 64:75-85

Froot, K. and Dabora, E. 1999, 'How Are Stock Prices Affected by the Location of Trade?', *Journal of Financial Economics* 53, 189-216

Goetzmann, W. and Garry, M. 1986, 'Does delisting from the S&P 500 affect stock' price? *Financial Analysts Journal*, 42, 64-68

Graham, A. and Pirie, W. 1994, 'Index fund rebalancing and market efficiency', *Journal of Economics and Finance*, 18.2, 219-229

Harris, L. and Gurel, E. 1986, 'Price and Volume effects associated with changes in the S&P 500 List: New evidence for the existence of price pressure', *Journal of Finance* 41, 815-829

Hegde, S. and McDermott, J. 2003, 'The liquidity effects of revisions to the S&P 500 index', *Journal of Financial Markets*, 6, 413-459

Kaul, A., Mehrotra, V. and Morck, R. 2000, 'Demand Curves for stocks *do* slope down: new evidence from an index weights adjustment', *Journal of Finance* 55, 893-912

Liu, S. 2000, 'Changes in the Nikkei 500: New evidence for downward sloping demand curves for stocks', *International Review of Finance*, 1:4, 245-267

Lynch, A. and Mendenhall, R. 1997, 'New evidence on stock price effects associated with changes in the S&P 500 Index', *Journal of Business* 70, 351-383

Mase, B. 2002, 'The Impact of Changes in the FTSE 100 Index', Working paper.

Pruitt, S. and Wei, J. 1989, 'Institutional ownership and changes in the S&P 500', *Journal of Finance* 44, 509-513

Scholes, M. 1972, 'The market for securities: substitution versus price pressure and the effects of information on share prices', *Journal of Business* 45, 179-211

Shleifer, A. 1986, 'Do demand curves for stocks slope down?', *Journal of Finance* 41, 579-590

Standard and Poor's Press Release, July 9, 2002, Focussing the S&P 500 on U.S. large cap stocks and the removal of non-U.S. companies in the S&P 500

Standard and Poor's Press Release, October 22, 2002, Criteria for Additions to and Deletions from a U.S. Index

Sui, L. 2003, 'The addition and deletion effects of the Standard and Poor's 500 index and its dynamic evolvement from 1990 to 2002: demand curves, market efficiency, information, volume and return', Working paper

Wurgler, J. and Zhuravskaya, E. 2002, 'Does arbitrage flatten the demand curves for stocks?' *Journal of Business*, 75(4) 583-608

Chapter 3 - Does familiarity affect the decision to invest in venture capital?

3.1 Literature Review

There is a long and distinguished history of modern portfolio theory which offers methods of allocating assets between asset classes and securities, largely based on choosing the portfolio with the lowest variance for a given expected return. "Modern Portfolio Theory" (starting with Markowitz (1952)) suggests that all investors will have a portfolio that is created to minimize risk for a given desired expected return. Most studies of asset allocation to date have focussed solely on large asset classes like stocks vs. bonds or domestic vs. foreign. With smaller or less common asset classes there have been several studies involving hedge funds and real estate[15], but to my knowledge none of the studies have concentrated on venture capital. Most of the studies (Lochoff (2002) excepted) have also considered the asset classes in an existing structure such as mean-variance without considering unique aspects of the asset class or the investor and how they fit together. There are, however, several aspects of pension funds and venture capital that may cause the underlying assumptions in these models to be false or at least misleading.

The main lesson is that funds need to be diversified. By holding multiple assets it is possible to have a portfolio with the same expected return, but lower risk, by taking advantage of the correlation structure of assets. However, despite all that we know about the advantages of diversification, in practice we do not observe it as often or as well

---

[15] See for example Lochoff (2002) for hedge funds. See Ciochetti, Brian A.; Sa-Aadu, J.; Shilling, James D. (1999), Chun, Ciochetti and Shilling (2000) or Craft (2001) for real estate.

implemented as theory would predict. I hope to examine one specific deviation from diversification which has gained considerable popularity in the academic literature over the last decade. *Familiarity* can occur for several reasons, but the underlying idea is that for reasons of cost, of comfort, of regulation or simply preferences, even sophisticated investors choose their portfolios in ways that deviate from basic diversification predictions.

I will look at the familiarity hypothesis in an area that has previously been overlooked – venture capital investment by pension funds. Pension funds do not seem to invest enough in venture capital based on a traditional mean-variance model of asset allocation (Lamm and Ghaleb-Harter (2001)). Some of the very reasons that explain why venture capital has not been looked at in more detail also explain why it is an ideal candidate for familiarity biases. Data on venture capital is very hard to come by, the management fees are very high along with a very high cost of entry, and some people believe, rightly or wrongly, that venture capital is either not appropriate for pension funds or is not "prudent" and thus illegal under ERISA. If familiarity is to be an issue with any asset class, then one would expect that venture capital would certainly be in that group.

Under the familiarity hypothesis, firms would choose to invest in companies that are geographically close to their place of business. This might not be considered "prudent" under the current diversification methods. Local firms are inexorably tied to the economy of a region. Location affects labour income, state transfers, and home prices, all of which the employees have significant exposure to outside of their pension plans. It would then be prudent of the trustees to try and diversify away this risk by investing in more distant firms. A second area where trustees might use familiarity to

sway their decisions is industry choice. It would be understandable for a trustee in high technology, for example, to lean the portfolio towards high technology firms and venture capital in particular where they are more comfortable. However, it would be prudent to invest primarily in industries and companies whose prospects are uncorrelated with those of the firm. For example if the entire technology industry is struggling, then the labour income and job prospects of employees in a high technology firm are at risk, at the exact same time as the assets in the pension plan would be at risk of a downturn. Thus under the familiarity hypothesis pension funds of firms located in high intensity venture capital areas and which are in industries that receive a lot of venture capital funds would be more likely to invest in venture capital.

3.1.1 Familiarity

"A man does not wonder at what he sees frequently, even though he be ignorant of the reason."- Cicero (De Divinatione (II, 22))

Familiarity is, at its most basic, the idea that people prefer to associate with, cheer for, make estimates of, and even invest in, things with which they are more comfortable, about which they have better information, or even about which they only think that they have better information. This manifests itself in areas as innocuous as sports betting. People are more likely not just to cheer on the home team, but to put their money down and bet on their own home team as well (Babad and Yosi (1991)). Extending to the financial markets, Strong and Xu (1999), for example, show that international investors show significant comparative optimism about their own home markets – the markets with which they are the most familiar. Thus even when there is no actual information

advantage, even a perceived information advantage or can affect behaviour. Huberman (2001) provides an excellent summary of previous research on familiarity.

There is now a significant literature which looks at familiarity as a way of explaining the way individuals and institutions choose to allocate their assets and select their securities. This began with the "home bias" puzzle, introduced by French and Poterba (1991) who analysed foreign holdings of investors in the G7 countries. Although it is generally impossible to actually view the expected returns, it is possible to infer them given the asset allocations and historical variance-covariance estimates[16]. The authors found that institutions tend to invest more in their home markets (and thus less in foreign markets) than would be predicted by standard asset allocation models. Researchers have determined that people seem to hold higher risk-return expectations for their home market than foreign markets, even among professionals. There is another possibility, of course. Unlike French and Poterba, Li (2004) infers expected returns and finds that investors would have to believe that the risk inherent in foreign investments was actually several times higher than the actual risk to account for the observed allocations.

There have been several explanations put forward to explain this apparently irrational behaviour. Foreign markets are more expensive to invest in. There are tax considerations such as dividend withholding taxes that only affect foreigners. It is more difficult to obtain information about foreign companies due to a difference in accounting standards, language and even distance. Regulation, such as the requirement in Canada to hold a maximum of 30% of pension assets in foreign assets can be circumvented, but there is a cost to that beyond the normal foreign investment costs. Kang and Stulz (1997)

---

[16] Many people mistake historical *observed* returns which can be observed for *expected* future returns which cannot be directly observed. Most asset pricing and allocation models require expected returns.

look at the cross-section of investments by foreigners in the Japanese equity market. They find that foreign investors do invest less overall than expected, but that these differences are least apparent among large companies in manufacturing industries and firms with good accounting performance, low unsystematic risk and low leverage. This is consistent with the idea that information costs are a significant barrier and can explain some of the home market bias since it is easier to get reliable information on these large companies than smaller companies.

Even within a given country familiarity can affect investors' choices. Grinblatt and Keloharju (2001) look at the behaviour of investors belonging to two separate language groups in Finland. They find that investors are more likely to hold and trade firms that are geographically close to the investor, that communicate in the same language and have CEOs of the same cultural background. More recently people have studied situations where none of these costs or barriers should exist. Within the United States, for example, research costs for retail investors are fairly homogeneous for all large capitalisation firms, the accounting standards are the same and the language of communication is always English. Huberman (2001) looked at the investment in the stocks of Regional Bell Operating Companies (RBOC) by retail investors. The author found that individuals tend to hold higher amounts of phone company stock in the company that serves their local area. The author concludes that there is a "cognitive bias" or confidence in the familiar. Zhu (2002) also looks at retail investors, but at a wider sample of stocks. He finds that retail investors are much more likely to hold firms located nearby, but that much of this effect disappears for those distant firms which have heavy advertising campaigns. The effect of distance is not limited to "unsophisticated"

retail investors. Coval and Moskowitz (1999) found that professional mutual fund managers are more likely to own companies that are located nearby than far away. This effect is more pronounced in companies which likely have higher information asymmetries such as small firms that produce non-traded goods.

3.1.2 Venture Capital

I intend to look at another instance where familiarity may appear – venture capital. This is important not only from the point of view of another data point in the familiarity literature, but also for the venture capital literature. There are several definitions used of venture capital and private equity in general. For my purpose I will think of venture capital as equity or equity-like investments[17] in early stage private companies by firms that are financial intermediaries. Although some corporations also have "venture capital arms", many of these are essentially outsourced research and development vehicles. Private equity is a broader definition that includes not only venture capital, but also mezzanine financing, leveraged buyouts, PIPES (private investments in public entities), etc.

An oft repeated concern of venture capitalists is that pension funds in particular do not invest enough of their assets in venture capital. In the United States the pension fund contribution is approximately 40% of all venture capital funding, with the rest coming from corporations, wealthy individuals and the government. There is almost no possibility for average individuals to invest in venture capital except indirectly through their pension funds. In the United States at the end of 2001 the average pension fund

[17]This includes convertible debt which is used in a large number of venture capital transactions. See Kaplan and Strömberg (2003) for example.

which held venture capital allocated 4.9% of total defined benefit assets to this asset class. As defined contribution plans (which do not currently invest in venture capital because of liquidity requirements) increase in popularity relative to defined benefit plans it will become even more important for venture capitalists to target those remaining defined benefit funds. In his 1999 address to the British Venture Capital Association, Tony Blair urged pension funds "to examine whether they and other institutional investors are being too cautious when it comes to venture capital and investing in early-stage companies". With venture capital assets are very illiquid, trading is infrequent and contracts are generally long-term. This makes pension funds a natural investor given their very long liabilities and predictable cash flows. The relationships between the general partner (the fund manager) and the limited partners (those providing the capital) are key to the success of the venture as are those between general partner and entrepreneur. Increased awareness, comfort and familiarity should be even more important for venture capital than for the public equity markets that have been studied to date. I hope to determine whether or not familiarity variables do, in fact, affect the decision to add venture capital to pension funds in the United States.

Information costs are a significant issue in all investing, but in few are they as important as in venture capital. "Good investment strategies should utilize detailed due diligence and deal screening based on intensive industry knowledge and important networks of experts" (Gompers and Lerner (2001)). In public markets like the stock market, information is released on a regular and timely basis. It is subject to strict regulation and oversight by both stock markets themselves and the Securities and Exchange Commission. There are professional analysts with brokerages and thousands

of professional investors and regular investors discovering information and forming opinions to set the price. All of this ensures that information flows swiftly into prices and that on average the price that an investor pays for a security is "fair". This is not true in the venture capital market. In this market many companies have almost no operational history to look at. The quality of the management is unknown. There are very few people who know the details of the company, and the management is only required to tell you what they want (within the bounds of fraud laws). Many companies at this stage are in very specific niches that have not been dominated by larger, older firms, so understanding the technology and the potential sales of the product requires specific expert knowledge of the market. This is not easy. It is very difficult to discover, investigate, understand and monitor such firms. One might, however, expect there to be some economies of scale to these information costs. Understanding a technological area requires a large upfront cost, but then a lower maintenance cost through time. Acquiring networks of contacts is similar.

It also takes a minimum amount of capital to get into venture capital. It is difficult to diversify a portfolio of under a million dollars in top quality projects (though the advent of funds of funds in the United States does allow smaller pension funds to get some exposure). Given these facts it is likely that large funds are more likely to invest in venture capital than smaller funds.

3.1.3 Pension funds

Pension funds in the United States have been a rapidly growing controller of funds since the early 1950s. Pension plans in the United States in 2002 controlled over

US$10 trillion in assets[18]. Three quarters of this money rests in defined benefit plans. In the beginning almost all pension funds were defined benefit plans. In a defined benefit plan the employer (usually the same as a plan sponsor), or a group of employers in a multi-employer pension plan, promises to pay the plan member a certain series of cash flows according to an agreed upon formula after retirement. Although there are several different forms that the benefit may take, most of them are based on a fixed dollar value or percentage of salary, the age of retirement and the number of years worked. The Employee Retirement Income Security Act (ERISA), introduced in 1974, governs the management of private (corporate) defined benefit pension plans. Several aspects of this act changed the way that pension plans are managed. Some of the more important features are:

- There is a requirement that the trustees act as fiduciaries with respect to the beneficiaries of the plan.
- Trustees must act solely in the interests of the beneficiaries rather than the company management or shareholders.
- The "prudent man" rule requires the trustees to manage the assets in the same way that a prudent man would manage his own assets. This has since been strengthened to what is called the "prudent expert" or "prudent investor" rule. Now the trustees are required to manage the assets as a man with experience managing money would manage his assets. One specific consequence of this is in the area of diversification.

Ali (2003) discusses the prudent investor rule in the context of hedge funds. Like venture capital he believes that too many trustees believe that it would be inappropriate

---

[18] See "Private Pension Plans, Participation, and Assets: Update", January 2003. The on-line URL for EBRI is www.ebri.org. This data is based on EBRI tabulations from U.S. Department of Labor Form 5500.

because of the high levels of perceived risk. In the US, however, it is clear that venture capital is both legal and prudent if the trustees want to invest in the asset class.

3.2 Data and Background

The dependent variable that I will study is whether or not the pension fund includes venture capital as part of its asset mix. A natural dependent variable to use would seem to be the amount of venture capital in a fund, but unfortunately is difficult to accurately measure for venture capital. Unlike most other asset classes, the decision to purchase equity in an early-stage company or to buy into a venture capital fund does not happen at the same time as the actual investment. In the United States pension funds usually sign on as limited partners to a venture capital fund. They then make a commitment to invest, say, $100 million. The venture capital fund can then draw on this money as needed to invest in companies. This can take several years to fully draw down and all of the funds may not end up being used. By this point, however, the decision to invest had already been made. Since the asset allocation data will measure how much is currently in venture capital rather than how much has been committed, the percentage can be a little misleading. I use the Pensions & Investments (P&I) survey data as the indicator of whether or not they include venture capital. There are no other sources of this data currently available. Venture capital databases which are widely available such as Venture Economics or Venture One work from a "bottom-up" perspective. They track deals at the level of individual financing or else fund company. They do not provide the information on the ultimate investor that P&I data provides us. It is unlikely that any systematic biases are introduced because of this.

The P&I study reports the 1000 largest pension funds in the United States and is produced annually. This data are collected by sending surveys to all plan sponsors. The editors then supplement returned surveys with independent research to round out the data available. The top 1000 pension funds have been reported since 1996[19]. Most results presented in this chapter are from the 2002 version of the survey. Although it is possible to analyse the data in a panel as we do in the next chapter, this is not appropriate when looking at venture capital. Venture capital partnerships are generally 10 years long with punitive provisions for early exit. Thus once a pension fund has entered the contract it will not leave the fund for at least 10 years. In 2002 the top 1000 pension funds in the United States included 618 corporate plans, 22 Multi-employer plans (also known as Taft-Hartley plans), 221 public plans and 139 Union plans based on P&I's classification. As can be seen in Table 3.1, there is a great difference in size amongst these groups as well as their propensity to invest in venture capital.

Table 3.1 - 2002 Pension plans and venture capital investment – Data from P&I.

Plan Type	Total	Average plan size ($millions)	Average size of Defined Benefit Plan ($millions)	Number with Venture Capital	Average allocation to Venture Capital
Corporate	618	$3,100	$1,709	58	1.9%
Public	221	$8,681	$8,318	62	1.5%
Union	139	$1,776	$1,641	7	1.1%
Multi-Employer	22	$11,258	$5,481	4	1.6%

[19] The 1000 funds are not the same every year, as small funds grow, big funds shrink, and companies merge or go out of business.

The rest of this paper will concentrate on the pension funds sponsored by publicly traded companies, because it is much easier to find data on them from sources like Compustat. Multi-employer and union plans do not have a single identifiable sponsor to test my hypothesis, and publicly managed plans have a single sponsor, but without suitable industry or firm characteristics. Of the top 1000 pension funds in the United States in 2002, 131 funds held venture capital in their portfolios. Of the 423 companies which have also data in Compustat, 54 held venture capital with an average allocation of 2.0%. For the purposes of testing our familiarity hypothesis I will use only American firms. I do this because it is unclear that American subsidiaries of foreign firms will behave in the same way as purely American firms. I also do not have reliable or consistent information on the venture capital markets outside of the United States. I will look at three different "familiarity" variables to see how these affect the decision to include venture capital in the pension fund – industry, geography and ownership of own company stock. This will help shed light on how institutions like pension funds make decisions to allocate their assets beyond just the more specific question of venture capital.

Data on venture capital was provided by PWC-MoneyTree. The MoneyTree Survey is a quarterly summary of all venture capital investments made in the United States. There are several different summaries of the data. The ones that we will concentrate on in this paper are the breakdown by state and the breakdown by industry. The MoneyTree survey uses 16 industry classifications (plus one "unknown/other" category). These do not match the SIC or NAICS codes so I match the pension sponsors as closely as possible to each of the 17 MoneyTree industry groups. While venture funded firms are almost always in a single niche and thus easy to categorize, the larger

firms that make up our pension sponsor sample are generally more diversified. The industry match is necessarily approximate but is unlikely to introduce any systematic biases.

The first familiarity measure that I look at is the primary industry that a firm operates in. In most corporate pension funds the trustees are made up of firm executives and staff representatives. People who work in industries that are commonly associated with venture capital such as high technology or biotechnology will likely be more familiar and comfortable with investing in firms in similar industries, including venture capital funded firms if the familiarity hypothesis holds.

The second familiarity variable, and one that has been examined in other contexts, is location. Unlike the home bias literature, all of the funds I will look at are based in the United States and most of the venture capital dollars are also invested in the United States. There is, however, a much larger informational issue in venture capital compared to public equity. Venture capital firms tend to congregate in several small locations or "clusters" around the country as do the start-ups that they fund. Although Silicon Valley is the best known and most successful, it is not the only hotbed of venture capital funding. This relationship has been studied by Lerner (1995). Venture capitalists do not just provide capital. The top venture capitalists also provide business development expertise, contacts with potential customers and suppliers, management oversight and eventually "certification" when the company goes public. To do this they rely on significant equity or equity-like ownership and board membership. Lerner looked at the board composition of venture capital backed biotechnology companies and determined that the closer a company was to the venture capitalist, the more likely that there would

be board representation by that firm: organizations with offices within 5 miles of the firm's headquarters are twice as likely to be board members as those more than 500 miles distant. Given that an individual venture capitalist can sit on as many as a dozen boards, and that the lead venture capitalist makes an average of 19 visits to the company (Gorman and Sahlman (1989)) the costs in both capital and time of visiting distant companies could be significant barriers.

Pension fund trustees who live in these areas will likely hear more about venture capital – especially the successes – and know people in the industry or related to the industry. For pension funds such as CalPERS which invest alongside their venture capital managers this would be especially important, because the trustees of CalPERS must not only act as a supplier of capital to venture capitalists, but as a venture capitalist in the broader definition of the term. For the pension fund locations I will use the state listed in the P&I pension fund survey. As Coval and Moskowitz (1999) and Goldberg, Heinkel and Levi (2003) suggest, increased personal contact can reduce moral hazard. This is an important aspect to convince normally conservative pension trustees to add venture capital to their funds. Coval and Moskowitz (1998) look at institutional investors and external managers. They find that 30% of clients of institutional money managers in Nelson's directory are located in the same city as the manager. It is likely that this would be even more true for venture capital where personal involvement with firm operations and strategy, as well as the limited partnership relationships would be much more important.

To determine the geographical "intensity" of venture capital investment I use two sources. For venture capital I use the Money Tree survey. This breaks down venture

capital investment by state. In order to standardize the amounts for each state to account for different sizes of economy, I use the Gross State Product (GSP) as provided by the Bureau of Economic Analysis (BEA). This is a similar measure to the Gross Domestic Product (GDP) at the state level.

The final familiarity variable that I use is whether or not the pension fund invests in the stock of the sponsoring firm. There is a large literature which discusses the role of own company stock in defined benefit pension plans[20]. This has all been concentrated in the defined contribution area and there has been no academic study of this in the defined benefit area. Under ERISA, defined benefit plans are allowed to hold up to 10% of the fund in the stock of the sponsoring firm, while there is no limit in defined contribution plans. Those trustees who allocate assets based more on the familiarity hypothesis would be more likely to hold own company stock, though this is clearly not the only reason for holding company stock. Although it is unusual to have one asset class as a predictor of another, it is important to realize that it is being used as a measure of familiarity rather than addressing any other reasons for holding either asset class.

Another similar variable that I look at is ownership of other "alternative" asset classes other than venture capital. I look at those asset classes that might also be considered "risky" by plan sponsors. This includes hedge funds, direct oil and gas investment, private debt, distressed debt and emerging markets equity. Having such a fine breakdown of asset classes is unique to the P&I data and compensates for the survey nature of the data.

---

[20] See Giammarino and Soumaré (2003) for a discussion

Size is a key factor that I need to control for in this study. Previous studies have shown that larger pension funds are more likely to hold venture capital than smaller funds. This is likely due to the high cost of entry. Although there are funds of funds available in the United States, it is generally only larger firms that enter the asset class. While it is not difficult to own a well-diversified large cap equity portfolio of $1 million, it is very difficult in venture capital.

3.3 Results

The first thing to note is the dominance of plan size on the decision to invest in venture capital. As can be seen in Table 3.2 there is an almost monotonic increase in the proportion of funds holding venture capital as fund size increases. Interestingly, once the decision to invest has been made, there does not seem to be any distinct trend in the *percentage* of venture capital. Because of this size needs to be included in all of the regressions in this chapter. In some cases we will also leave size out as a variable but repeat the regressions on each quintile[21].

[21] With only 375 pension funds deciles are too small to have any confidence in regressions.

Table 3.2 – Size of DB plan and venture capital investment - Size is perhaps the largest determinant of whether or not a pension fund will hold venture capital. With the exception of the fifth size decile the chance of holding venture capital increases uniformly from smallest to largest defined benefit pension plans. Once the decision to hold venture capital has been made there does not seem to be any relationship between size and amount invested in venture capital

Size Decile	Mean assets in DB plan ($millions)	Proportion of funds holding venture capital	Mean Allocation to venture capital
1 (smallest)	158	0%	0.0%
2	322	2.6%	7.5%
3	432	2.7%	0.4%
4	563	5.4%	3.4%
5	735	13.2%	2.1%
6	989	8.1%	1.3%
7	1,426	10.5%	1.3%
8	2,084	18.4%	1.5%
9	3,126	27.0%	1.7%
10 (largest)	13,660	44.7%	1.8%

The first factor that I look at is the location of the company's head office. The principal measure is geographic intensity which I define as the total amount of venture capital investment in the state in a given year divided by the Gross State Product (GSP). Those states with the highest number of dollars invested in venture capital are those that one would expect – California, Massachusetts and Texas. Looking at the top 10 states by intensity includes some surprises like Colorado, Maryland and Delaware. Those pension funds which do not invest in venture capital have an average intensity of 3.5 and those which do invest in venture capital have an average intensity of 3.2. These are not statistically different in a two-sample mean comparison test, suggesting that geography does not have an effect on the decision to invest in venture capital.

The second familiarity variable is the primary industry in which it operates. The principal measure is industry intensity which I define as the percentage of venture capital investment in the industry in a given year. Those funds which do not invest in venture capital have an average intensity of 3.0 while those which do invest in venture capital have an average intensity of 3.9. Although cursory inspection would suggest that those

industries with more venture capital investment nationally are also more likely to invest in venture capital in their pension funds, the difference is not statistically significant in a two-sample mean comparison test.

Next I look at own-company stock ownership and the ownership of other alternative assets. As mentioned above I am using both of these categorizations as a control to determine whether or not the trustees seem to use familiarity to influence their decision to allocate assets. I do not attempt to show any other link between these asset classes and venture capital. Table 3.3 shows the relationship between venture capital, own company stock and alternative assets. Those firms that choose to invest in their own company's stock are almost three times as likely to invest in venture capital than those which do not. Investing in other alternative assets seems to have only a very small correlation.

Table 3.3 - Own company stock and alternative assets

	Firms without own company stock	Firms with own company stock	Firms without alternative assets	Firms with alternative assets	Total
Firms that do not invest in venture capital	280 (90.0%)	45 (70.3%)	280 (87.2%)	45 (83.3%)	325
Firms that do invest in venture capital	31 (10.0%)	19 (29.7%)	41 (12.8%)	9 (16.7%)	50
Total	311	64	321	54	375

The next step is to look at the familiarity variables in a more general framework. Since we are modelling a probability of investing in venture capital, the most natural regression to perform is a probit regression. For the first set of regressions I will include the size of the pension fund as a control variable. Table 3.4 shows the results of these

regressions. The only variable that seems mildly significant is industry. None of the other variables are significant at the 10% level[22].

Table 3.4 - Single variable probit tests – Results of a probit model with an indicator variable equal to one if the fund holds venture capital and zero otherwise. The number of dollars in the fund is included in all of these regressions given the large significance that it has in the decision. * represents significance at the 10% level.

Coefficients (t-stat)	Geography	Industry	Own Company stock	Alternative assets
Constant	-1.2	-1.4	-1.3	-1.3
	(-10.0)	(-11.3)	(-13.6)	(-13.0)
Variable	-1.2	4.3*	2.5	.28
	(-0.5)	(1.7)	(1.48)	(1.3)
Size of fund	5.5	5.5	5.6	5.7
	(4.6)	(4.5)	(4.6)	(4.7)
Pseudo R^2	7.3%	8.2%	7.9%	7.7%

Rather than end the analysis here it is likely useful to consider which variables to include in a more complete model. Stepwise probit analysis yields the following model with a pseudo R^2 of 12.9%:

$$P(\text{venture capital}) = \underset{(-11.3)}{-1.65} + \underset{(4.4)}{5.6}*\text{fund size} + \underset{(1.8)}{4.6}*\text{industry intensity} + \underset{(3.8)}{0.8}*I(\text{own company stock})$$

In this model the ownership of company stock is now strongly significant, but the key familiarity variable – industry intensity – suggests that those firms which operate in a venture capital intensive industry like high technology are in fact more likely to include venture capital in their pension funds. One possibility is also that there is a non-linearity which would cause the effects to increase non-linearly. Rather than just include a quadratic term in the linear regression, another way of accounting for the difference in

[22] All of the significance levels are based on 2-sided tests, to recognize that one possible alternative hypothesis is that firms follow traditional asset allocation and diversify away from venture capital if it is highly correlated with human capital, real estate values, etc. Using one sided tests, while making the observed values slightly more significant, would not significantly change the conclusions.

propensity by fund size is to repeat the regression within each size quintile. This should also show any non-linearity if it exists.

Before running the regressions it is useful to see summary statistics for each of the variables before running the regressions. The summary is in Table 3.5.

Table 3.5 – Summary of data by quintile

Quantile	VC in Pension plan	Plan Assets ($millions)	Industry Intensity	Own Company Stock	Other Alternative Assets
1 (smallest)	1.3%	241	2.9%	12.0%	2.7%
2	4.0%	500	2.6%	10.7%	4.0%
3	10.7%	868	3.3%	24.0%	17.3%
4	14.7%	1,762	3.2%	17.3%	24.0%
5 (largest)	36.0%	8,463	3.5%	21.3%	54.7%

Table 3.6 - Quintile regressions – Rerunning the regressions within each size quintile suggests that the significance above is simply an artefact of the data rather than a consistent effect. *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively. Numbers in parentheses are t-stats.

Quintile	Constant	Industry	Own Company Stock	Pseudo R^2
1 (smallest)	-2.4***	5.8	n/a	1.5%
	(-3.9)	(0.4)		
2	-2.7***	26.3**	.1	25.9%
	(-4.7)	(2.0)	(.04)	
3	-1.6***	5.4	0.5	3.4%
	(-4.5)	(0.9)	(1.1)	
4	-1.0***	-7.7	0.6	5.6%
	(-3.5)	(-1.0)	(1.4)	
5 (largest)	-0.8***	5.4	1.1***	11.4%
	(-3.6)	(1.4)	(2.9)	

Table 3.6 suggests that the result above is an artefact of the data rather than a strong effect of familiarity. The industry variable is significant in only the second quintile, and own company stock is only significant in the largest 20% of pension funds. Although the point estimates of the own company stock variable do seem to increase with firm size, any non-linearity would still be statistically insignificant. As a final

confirmation the above regressions are also run on 1998 data. In 1998 there were only 9.9% of pension funds which held venture capital (compared with 13.3% in 2002) and the smallest 40% of pension funds held no venture capital at all. This overall increase is likely caused not only by the frothy IPO markets in the intervening years, but also the introduction of fund-of-funds which allowed funds to invest smaller dollar values in venture capital. The results for 1998 yield the same negative result – there does not seem to be any effect of geography, industry or alternative asset classes on the likelihood of holding venture capital[23]. Overall this suggests that familiarity does not affect the choice by pension fund trustees of whether or not to invest in venture capital counter to previous evidence on other asset classes.

3.4 Caveats

There are several caveats regarding the data that should be addressed before continuing with the conclusions. I do not intend to detract from the result obtained but to recognize that the data available in both venture capital and pension funds is very difficult to obtain, and what is available must be interpreted carefully.

The first issues are with the nature of the P&I data. The primary source of data is survey based. Unlike studies which use Form 5500 which is required by government legislation, this survey is completely voluntary. There is also no clear definition given of venture capital or private equity. The definition is left up to the plan sponsor filling out the form. In the 2002 survey, for example, there are approximately 85 pension plans that

---

[23] The ownership of own company stock in defined benefit plans was not included in the 1998 P&I survey.

claim to hold private equity, but when asked to break out the different categories claim not to hold any. The advantages of using the P&I data, outlined in the data section, outweigh this weakness.

The PWC data are very specifically venture capital only. This requires us to use venture capital instead of the broader private equity category. The geographic data that we use is based on state rather than a more specific location[24]. While there may well be a lot of venture capital investment in Silicon Valley, for example, this may not extend to southern California.

Using the firm's headquarters as a measure of geography may not be a good representation of either the trustees making the decisions or the pension beneficiaries. For a small local firm all of the trustees and employees will likely be in the same geographic area. For large corporations with operations around the country the geographic measure may be less important[25]. It is likely, however, that most company executives will be located at the head office. These people often make up a large portion of the trustees of the pension plan and will be influenced by this location. A large number of employees will also be located at or near the head office.

3.5 Conclusions

The principal question I wanted to answer was whether or not familiarity played a role in the decision to include venture capital in pension funds or not. I have looked at industry, location, ownership of own company stock and other alternative investment classes as likely familiarity variables. I found that while there seemed to be a familiarity

[24] The venture capital data is available on a county basis as well, but there is no GSP data available at such a fine level

[25] Because we only deal with US pension plans we do not need to worry about international employees.

effect related to industry, this did not stand scrutiny under different specifications of the size of the pension plan.

This result is puzzling on two opposite fronts. The familiarity literature to date has found very strong biases for investors to tilt their portfolios towards assets they are more familiar with from domestic assets instead of foreign assets to owning the local phone company instead of another phone company in the same country. This has been supported in several assets, in several countries and for several different familiarity variables. If any asset class is ex ante likely to exhibit familiarity biases one would think that it would be venture capital due to the high information costs, high perceived risk and the tendency of venture capital to cluster in certain locations and industries. One possible interpretation of my result is that the need to diversify pension plans away from the industry and geographical exposure that the beneficiaries already have overwhelms the familiarity bias of venture capital.

3.6 References

Ali, P.U. 2003, 'Hedge fund investments and the prudent investor rule' Trust Law International v19, n2: 74-89

Babad, E. and K. Yosi, 1991, 'Wishful Thinking – Against All Odds', Journal of Applied Social Psychology, 21, 1921-1938

Ciochetti, Brian A., Sa-Aadu, J., and Shilling, James D. 1999, 'Determinants of Real Estate Asset Allocations in Private and Public Pension Plans' *Journal of Real Estate Finance and Economics* v19, n3: 193-210

Chun, G. H., B. A. Ciochetti and J. D. Shilling 2000, 'Pension-Plan Real Estate Investment in an Asset-Liability Framework', *Real Estate Economics*, 28:3, 467-491.

Coval, J. and Moskowitz, T. 1999, 'Home bias at home: local equity preference in domestic portfolios', *Journal of Finance*, 54(6), 2045-2073

Coval, J. and Moskowitz, T. 1998, 'On the selection of fund managers by institutional investors', Working paper

Craft, T. M. 2001, 'The Role of Private and Public Real Estate in Pension Plan Portfolio Allocation Choices', *Journal of Real Estate Portfolio Management*, 7:1, 17-23

French, K. and Poterba, J. 1991, 'Investor diversification and international equity markets', *American Economic Review*, 81(2) 222-226

Giammarino, R. and Soumaré, I. 2003, 'Incentives and Voluntary Investment in Employer Shares', Working paper

Goldberg, M., Heinkel, R. and Levi, M. 2003, 'Foreign Direct Investment: The Human Dimension', Working paper

Gompers, P. and Lerner, J. 2001, *Money of Innovation: How Venture Capital Creates New Wealth*, Harvard Business School Press

Gorman, M. and Sahlman, W. 1989, 'What do venture capitalists do?', *Journal of Business Venturing* 4, 231-248

Grinblatt, M. and Keloharju, M. 2001, 'How Distance, Language and Culture Influence Stockholdings and Trades', *Journal of Finance*, 56:3, 1053-1073

Huberman, G. 2001, 'Familiarity Breeds Investment', *Review of Financial Studies*, 14:3, 659-680

Kang, J. and Stulz, R. 1997, 'Why is there a home bias? An analysis of foreign portfolio equity ownership in Japan.' *Journal of Financial Economics*, 46, 3-28

Kaplan, S.N., and P. Strömberg, 2003, 'Financial contracting theory meets the real world: an empirical analysis of venture capital contracts.' *Review of Economic Studies* 70, 281-315.

Lamm, R. and Ghaleb-Harter, T. 2001, 'Venture capital as an asset class: Its role in Investment Portfolios', *Journal of Venture capital*, Fall 2001, 68-79

Lerner, J. 1995, 'Venture Capital and the oversight of private firms', *Journal of Finance*, 50:1, 301-318

Li, K., 2004, "Confidence in the Familiar: An International Perspective," *Journal of Financial and Quantitative Analysis*, 39, 47-68

Lochoff, R. 2002, 'Hedge Funds and Hope: Hedge Funds Deserve Allocations', *Journal of Portfolio Management*, v28, n4: 92-99

Markowitz, H. 1952, 'Portfolio Selection', *Journal of Finance*, 7:17-91

Zhu, N. 2002, 'The local bias of individual investors', Yale ICF Working Paper No. 02-30

Chapter 4 - Pensions and Corporate Performance: Effects of the Shift from Defined Benefit to Defined Contribution Pension Plans

4.1 Literature Review

Pension plans are well known to be major participants in financial markets. As of 2002, total tax-qualified U.S. pension plans had the remarkable total of $10.1 **trillion** in assets.[26] This asset total was approximately equal in value to the entire domestic capitalization of the New York Stock Exchange (NYSE).[27] Despite the importance of pension funds as investors, corporate decisions concerning the structure of employer-provided pension plans have received relatively little attention as potential determinants of corporate performance. We argue that the determination of pension plan structure is among the most important financial decisions undertaken by firms.

Particularly worthy of interest is the gradual but cumulatively large shift that has taken place over the past 25 years from "defined benefit" (DB) pension plans to "defined contribution" (DC) plans. Whether a pension plan is a DC or DB plan is likely to have an important effect on workers' retirement decisions, job transition decisions, and human capital investment decisions. These decisions might reasonably be expected to affect a firm's performance. Furthermore, as U.S. pension plan assets continue to grow in relative importance and as the American labour force continues to age, any effects of pension

---

[26] This total and all other data in this paper exclude Social Security. Of the $10.1 trillion total, about $6.3 trillion was in employer-provided plans, $2.5 trillion was in individual retirement accounts (IRAs), and $1.3 trillion was in private insured pension accounts (i.e. accounts with insurance companies). See Employee Benefit Research Institute (EBRI) (2004a). The EBRI website at URL www.ebri.org is an good source of background pension information.

[27] The NYSE capitalization is obtainable from the NYSE website at www.nyse.com by following the market information, quick facts, and market capitalization links. As of July 31, 2003, market capitalization of domestic companies listed on the NYSE was $10.1 trillion, of which a significant share was owned by pension funds. Pension funds have large holdings of equities and also hold a range of other assets.

plan structure on corporate performance are likely to become increasingly significant in the future. We view the possible impact of the changing structure of pension plans on corporate performance as an important and interesting topic. Accordingly, the primary objective of this paper is to investigate the effect of the shift from DB to DC pension plans on corporate performance.

Over 60% of pension assets are in employer-provided pension plans. These plans must be either DB or DC plans. DB plans base retirement benefits on a formula incorporating years of service, age, and earnings. DC plans (which include so-called "401(k)" plans) establish individual employee accounts whose value is determined by employer and/or employee contributions and the investment performance of the account. Most small and medium-sized firms with pension plans have only one type of plan, but large firms commonly have both. When a firm has both types, workers may or may not have the option of self-selecting into a preferred plan or diversifying over plan types. Most pension-eligible workers have only one pension plan but a significant minority have assets in both DB and DC plans.

Both DC and DB plans must be funded in accordance with legal requirements. In the case of DC plans this is straightforward. Each period, contributions are made by the employee and/or employer to the individual worker's DC account. Managers of the pension plan make investments and sometimes provide a menu of investment options to the employee. The plan might earn positive or negative returns, but the plan is always fully funded in the sense that everything owed to the pension holder is in his or her individual account at any given time.

DB plans are more complicated. The employer must make contributions to the plan on a regular basis. The objective is that the plan should normally be "fully funded" in the sense that the actuarial expected value of plan assets should be sufficient to pay off the actuarial expected value of liabilities. However, employers are sometimes in the position of having "under-funded" plans. Instead of requiring that DB plans always be fully funded, the law allows for temporary under-funding. However, registered DB plans be insured with the Pension Benefit Guarantee Corporation (PBGC), a self-funded agency of the U.S. government[28]. If a company goes into bankruptcy and is unable to meet its DB pension commitments, the PBGC assumes the assets and liabilities of the insolvent company's DB plan, subject to certain limits. The PBGC then pays off most, but not necessarily all pension liabilities of the company's DB plans. Pension beneficiaries bear a small amount of residual risk[29]. Despite this small risk, DB plans are, in general, very safe. Even the Enron DB beneficiaries are being paid in full, unlike Enron's DC beneficiaries, whose pension funds were invested primarily in now worthless Enron stock.[30]

The shift from DB to DC pensions has been well-documented. From the data presented in EBRI (2004a) it follows that between 1985 and 2002 the DC share of private sector employer-provided pension assets rose from 34% to 54%, while the DB share fell, correspondingly, from 66% to 46%, continuing a trend that has been in place at least

---

[28] An interesting evaluation of the PBGC pension insurance program is provided by Marcus (1987).

[29] Because DB pensions are not quite fully insured, the payout under a DB pension plan is not fully independent of the performance of the sponsoring firm. If the firm goes bankrupt and turns its pension assets and liabilities over to the insurer (PBGC), some claimants may get less than their full pension. In the absence of insurance, a DB plan would be much like a long-term corporate bond, with an associated risk profile.

[30] See money.cnn.com/2004/09/13/news/midcaps/enron_pension.reut/ for details.

since the mid-1970s. Data on the relative number of private sector DB and DC plans shows a more dramatic shift. According to EBRI (2004b), the number of DB plans fell from 170,172 in 1985 to 32,351 in 2002, while the number of DC plans rose from 461,963 to 809,889. The reflects a particularly large conversion rate among small employers. Considerable attention has been focussed on this change, especially in light of highly publicized losses in DC plans at Enron and elsewhere and because of high levels of under-funding that have emerged in recent years in DB plans.

This paper examines whether the variations across firms and over time in the share of DC assets in total pension assets are associated with corresponding variation in corporate financial returns. We focus on the return on assets (ROA) as the primary measure of corporate financial performance. We also consider the operating return on assets (OROA) and the return on equity (ROE) as performance measures.

We emphasize two reasons why changes in the structure of pension plans might have a significant impact on performance. First, DB and DC pension plans have strikingly different incentives regarding retirement behaviour. More specifically, DC plans typically offer incentives to delay retirement relative to DB plans and would, therefore, lead to later retirement. If, on average, those workers who delay retirement generate net surplus for the firm, then the shift toward DC plans would tend to improve corporate performance. Conversely, if these workers are a net burden on the firm in the sense that their total compensation exceeds their value to the firm, then the shift toward DC plans might damage corporate performance.

A second reason for a possible effect of pension plans on corporate performance is that pension plans might affect job mobility decisions at early career stages because of

"vesting" requirements. Pension plans typically have a vesting period under which workers must remain with the firm for a certain period of time before they acquire the right to receive a future pension. Vesting requirements are normally much stricter under DB plans than under DC plans.[31] DB plans might therefore inhibit worker mobility relative to DC plans. In addition we would not rule out other possible influences of pension plan choice on firm performance. For example, worker effort, capital market access, and product demand might all be affected by pension plan choice in view of the fact that under-funded DB plans can play a significant role in firm insolvency. However, in this paper we emphasize worker mobility and retirement decisions.

There is a large literature dealing with pension economics and finance. Several streams in that literature are relevant to this paper. One highly relevant research stream for our analysis includes papers dealing with the effect of pension plan structure on retirement decisions. Early papers of this type include Kotlikoff and Wise (1989), Lazear (1986), and Stock and Wise (1990). In addition Friedberg and Webb (2003) is of particular interest as it demonstrates that the shift from DB toward DC pension plans is associated with later retirement.

In addition, there is a literature dealing with the relationship between pensions and productivity, usually measured by output per worker, value-added per worker, or total factor productivity. See, in particular, Ippolito (1997), Dorsey, Cornwell and Macpherson (1998), and Cornwell and Dorsey (2000). The theoretical background to this literature

---

[31] The rules associated with vesting requirements are fairly complex. They are explained in some detail in Employee Benefits and Security Administration (EBSA) (2004). An employee always has an absolute right to the employee's own contributions. Employer contributions are subject to vesting. As DB plans are normally structured in such a way as to derive principally or entirely from employer contributions, vesting is a more important constraint for DB plans than for DC plans. Over the sample period the normal DB vesting period was three to five years

consists of work by Becker (1964), Lazear (1983) and others seeking to explain possible efficiency advantages of defined benefit pension plans. The basic idea is that the tendency of DB plans to discourage early job transitions by workers and to encourage relatively early retirement in the late stages of a worker's career allows for efficient investment by the firm and the worker in the worker's firm-specific human capital. Cornwell and Dorsey (2000) undertake an empirical comparison of firms with DB and DC plans for the 1982-1991 period and emphasize their finding of a productivity premium for DB plans in large manufacturing firms, although they also report that no such effect exists in their full sample.

Within the literature on pensions there is a sub-literature that documents and seeks to explain the shift from DB toward DC pension plans, including Ippolito (1995), Kruse (1995), Mitchell and Schieber (1998), and Ostaszewski (2001), among others. In related work, Huberman, Sethi-Iyengar and Jiang (2003) estimate the determinants of participation in DC plans. One commonly given reason for the shift is an alleged administrative cost advantage for DC plans in the aftermath of the ERISA, which set the important features of the current pension environment. A second category of explanation focuses on the preference of employees for the greater transparency and portability of DC plans. A third explanation is that employers perceive less risk from defined contribution plans. This explanation may seem odd from an economist's point of view as employers are normally taken to be less risk averse than employees. Therefore, if DC plans shift risk from employers to employees they should be inefficient and would require higher wage compensation to employees than would be justified by any reduction in risk to employers. The requirement for these extra payments can only strengthen the results later in this

paper since the firm would need to perform better even after accounting for the extra wages.[32]

The fourth important class of explanation relates to "reversions", which occur when an employer terminates a defined benefit pension plan and replaces it with a defined contribution plan. When a DB plan is "over-funded" in the sense that assets exceed the actuarial requirement to fund future pension commitments, then the employer can take the surplus as income.[33] Thus one possible reason for the shift toward DC plans might arise from employers' incentives to replace DB plans with DC plans so as to realize any DB plan surplus as current income. Stone (1987), VanDerhei (1987), and Mittelstaedt (1989) find that pension reversions are more likely to be undertaken by firms in financial difficulty. Petersen (1992) presents evidence that pension reversions have the effect of transferring wealth from pension beneficiaries to firm shareholders. Ippolito and James (1992) propose an extreme form of reversion where firms are taken over by a leveraged buyout specifically to transfer money from the DB pension plan to the new equity holders.

4.2 Theoretical Structure

4.2.1 The worker's retirement decision

---

[32] In light of the greater transparency and portability of DC plans, it is possible that both employers and employees perceive DC plans as less risky.

[33] Following from the 1980 reversion of the A&P defined benefit pension plan, the courts established that the employer has the right to terminate over-funded plans and take the surplus. However, subsequent legislation and tax rulings have made such terminations less attractive. As of 2004, the tax rate on conversion surpluses is 50% but employers can reduce or avoid the tax by transferring some or all of the surplus to new pension plans or health plans.

Our theoretical context focuses on two important decisions undertaken by workers – the retirement decision at late career stages and the job mobility decision at early career stages. Consider the retirement decision first. This is a binary decision at any given time: the worker can either retire or continue working for another period. Let the value of the decision be denoted V, let continued work be represented by 0, and let retirement be represented by 1. We can write the value of continuing to work for the current period at time t, $V_t(0)$, as

$$V_t(0) = u(w_t, e_t, a) + ?E(V_{t+1}) \quad (1)$$

where u is current utility, w is this period's compensation, e is work effort, a is an attribute that varies across individuals, ? is a discount factor taking on a value between 0 and 1, and E is the expectation operator. The value of retirement is given by

$$V_t(1) = u(P_t) \quad (2)$$

where P_t is the value of pension wealth at time t. $V_t(1)$ is taken to include all benefits from retirement. The utility function in the event of retirement may include arguments other than P_t but we do not explicitly model such effects here.

The retirement decision should be viewed as irreversible. A retired person might of course take another job, but the original pre-retirement pension status cannot be restored. Thus continued work has an option value associated with it. This value should be viewed as being incorporated in $E(V_{t+1})$. In deciding whether or not to retire, the worker should compare $V_t(1)$ with $V_t(0)$. We can define the net value of retirement at time t, denoted NV_t, as the difference between the value of continuation and the value of retirement.

$$NV_t = V_t(1) - V_t(0) \quad (3)$$

The person should retire at time t if NV_t is positive.[34]

Variable a is intended to capture the idea that, even holding wage and effort fixed, some workers enjoy working more than others. One contributor to a might be health status – employment might be more enjoyable (or less costly) for workers enjoying good health. We define a such that utility is increasing in a. It follows immediately that, other things equal, the net value of retirement is decreasing in a. There is of course a distribution of w, P and other factors affecting retirement across individuals. For any given level of w, P, and other factors, workers with higher levels of a will work longer.

4.2.2 Effects of Retirement Decisions on the Firm

The value of a worker's retirement decision to the firm at time t is denoted p_t.

$$p_t = s_t(w_t, e_t, ß) + ?E(p_{t+1}) \qquad (4)$$

where s is the current surplus generated by the worker for the firm. It is the value of the worker's marginal product over and above the cost of employing the worker. This value depends on the worker's compensation, w, the worker's effort, e, and an attribute, ß, that we can think of as "productivity".

The underlying employment environment we have in mind allows for compensation to differ from the value of marginal product. Many employment contracts, particularly but not exclusively in unionised environments, imply a wage profile that may differ from the productivity profile of a given worker. Furthermore, regardless of the aggregate pattern of the wage profile, there will be individual variation. Some workers will experience less growth (or faster loss) of productivity than the average and some will

---

[34] A more detailed simulation of the worker's retirement decision in a continuous time model with stochastic marginal productivity is provided by Sundaresan and Zapatero (1997).

experience more. This individual variation in net worker contributions to the firm is captured by the variable ß. The firm would like the worker to retire at time t if p_t is negative.

4.2.3 Retirement and the Shift from Defined Benefit to Defined Contribution Plans

A defined benefit plan implies a different evolution of pension wealth, P, than a defined contribution plan. Under a DC plan, the evolution of P is easy to understand. As long as the worker keeps working and making contributions to the plan (usually supplemented by employer contributions) the plan will increase in value, subject to the effect of investment returns. Investment returns are normally also positive. The longer a worker waits before starting to make withdrawals (i.e. before retirement) the higher the value of the plan at retirement. Thus pension wealth increases in a transparent way over time and offers a clear incentive to continue working.

Determining pension wealth for a DB plan is more difficult. Pension wealth at time t can be viewed as the expected present value of payments from the pension plan. This present value is 0 until the worker has worked long enough to qualify for a future pension (the "vesting period"), when it takes a sharp jump up. The value of the DB pension then grows as age, years of service and wage levels increase.

Most DB plans have an early retirement date (ERD), when reduced benefits can be taken. Every year worked after the ERD means that the worker is giving up the opportunity to receive pension benefits for that year. This effect tends to reduce pension wealth but is usually more than offset by higher pension payments called for by plan parameters as years of work increase, up to the "full benefit retirement date" (FBRD),

when the person qualifies to begin receiving "full" pension benefits. This typically occurs in a worker's late 50's or at about age 60, depending on the specific plan and the worker's employment history. Beyond this age, pension wealth normally decreases if the person continues to work as each extra year worked represents a year of foregone pension income. In general, for workers in their early 60s and beyond, pension wealth normally decreases with further years worked under DB plans and increases with additional years worked under DC plans, inducing sharply different retirement incentives.

One important difference between DC and DB plans is that DB plans contain an important element of cross-subsidy. Workers with high levels of a (i.e. those who enjoy working) will work after the age at which pension wealth is maximized under a DC plan. In effect, they will end up receiving less than the actuarial value of the employer's contributions made on their behalf. Those who retire when pension wealth is maximized will, on the other hand, receive more than the actuarial value of the employers contributions made on their behalf.

Accordingly, we suggest that workers with high values of a will self-select into defined contribution pension plans. For large employers, this self-selection is typically available without changing jobs, as most large employers with DB plans also have DC plans available to which workers can make voluntary contributions. We also suggest that workers with high values of a will tend to have high values of ß: that a and ß are positively correlated. In order words, we suggest that workers who enjoy their work are more likely to generate positive surplus for the employer than those who do not.

Therefore, we suggest that the growth of defined contribution pension plans has differentially selected for workers who enjoy working and who tend to be more

productive to move into DC plans. Having moved into DC plans they face stronger incentives to work longer and will do so, generating positive surplus for the firm. Thus the shift from DB to DC plans should be associated with improved corporate performance.

4.2.4 Job Mobility Decisions

The theoretical implications of DC and DB plans with respect to job mobility are relatively straightforward. Because of the vesting requirements for DB plans, early job mobility is discouraged. This might be efficiency-enhancing if it encourages efficient investment by the firm in the worker's human capital. On the other hand, it might be efficiency-reducing if it simply prevents workers from moving to jobs where they would have higher marginal productivity. In principle, either effect could dominate.

Even if DC plans are efficiency-enhancing in aggregate, we must still explain how this affects individual firms. What is the consequence for a firm of favouring DC pension plans relative to DB plans? We suggest that DC plans favour better corporate performance. Our reasoning relies on self-selection, as with the retirement effect. Specifically, workers who have or anticipate having relatively high productivity would tend to self-select into DC pension plans by choosing employment with firms that have DC plans or, if they have the choice available, by selecting a DC plan in preference to a DB plan with a given employer. Such workers would also have a stronger incentive to invest in human capital acquisition. We suggest that the resulting higher productivity of workers covered primarily by DC plans provides net benefits for the firm as we would expect that the surplus associated with more efficient human capital decisions would be

shared betweens firms and workers. Furthermore, DC plans also make it easier for workers who are unhappy with their jobs to leave rather than staying simply to qualify for their pensions. We suggest that this effect is also likely to be beneficial to firms.

4.3 Data

Our analysis is based on data combined from two primary sources. One source is Compustat. The other source is Pensions and Investments (P&I), a provider of on-line and hard copy information on pensions[35]. Since 1997, P&I has provided an annual list of the top 1000 pension "sponsors". Corporations form the majority of these pension sponsors, accounting for about 620 of the top 1000 pension sponsors in any given year. However, the very largest sponsors are public entities, like the California Public Employees Retirement System (CalPERS), which has the largest pension fund, with 2003 assets of $148.8 billion. The next largest sponsors in 2003 were the Federal Retirement Thrift Investment Board, with assets of $118.8 billion and the New York State Retirement Fund, with 2003 assets of $106.8 billion. The largest corporate pension sponsors for 1997 and 2003 are shown in Table 4.1.

[35] The URL is www.pionline.com.

Table 4.1 Largest Corporate Pension Sponsors - The very largest pension plans are public sector plans, but some corporate pension sponsors also have very large pension plans. Corporate Assets are from Compustat (Item 6). Pension assets and the defined benefit and defined contribution shares in pension assets are from P&I.

Corporate Sponsor	Corporate Assets ($bill.)	Pension assets ($billion)	Defined Benefit Share (%)	Defined Contr. Share (%)
1997				
General Motors	228.9	90.6	79.1	20.9
General Electric	304.0	56.9	68.5	31.5
IBM	81.5	53.1	73.7	26.3
Lucent Tech.	24.8	48.4	74.3	25.7
Ford Motor	279.1	47.8	74.7	25.3
2003				
General Motors	448.5	89.7	79.1	20.9
General Electric	647.5	61.6	67.0	33.0
IBM	104.5	60.1	65.2	34.8
Boeing	53.0	58.2	66.3	33.7
Ford Motor	315.9	44.9	73.8	26.2

As can be seen from Table 4.1, the largest corporate sponsors have very substantial assets in both DB and DC plans. These very large firms experienced only a small shift toward DC plans over the 1997-2003 period. In addition to corporations and public employers, other organizations represented as pension sponsors include labour unions, universities, and churches. Putting together data from EBRI (2004a) and our P&I data shows that in any given year the top 1000 pension sponsors account for about 2/3 of total employer-sponsored pension assets. Figure 3.1 shows the evolution in pension assets in the top 1000 pension funds and how it is broken down by sector.[36] Public and NP includes the public sector along with universities and churches. (NP stands for "non-profit".)

[36] Pension sponsor type is missing from the data for 1998. The values for 1998 shown in Figure 3.1 are estimated by taking the average of the 1997 and 1999 values.

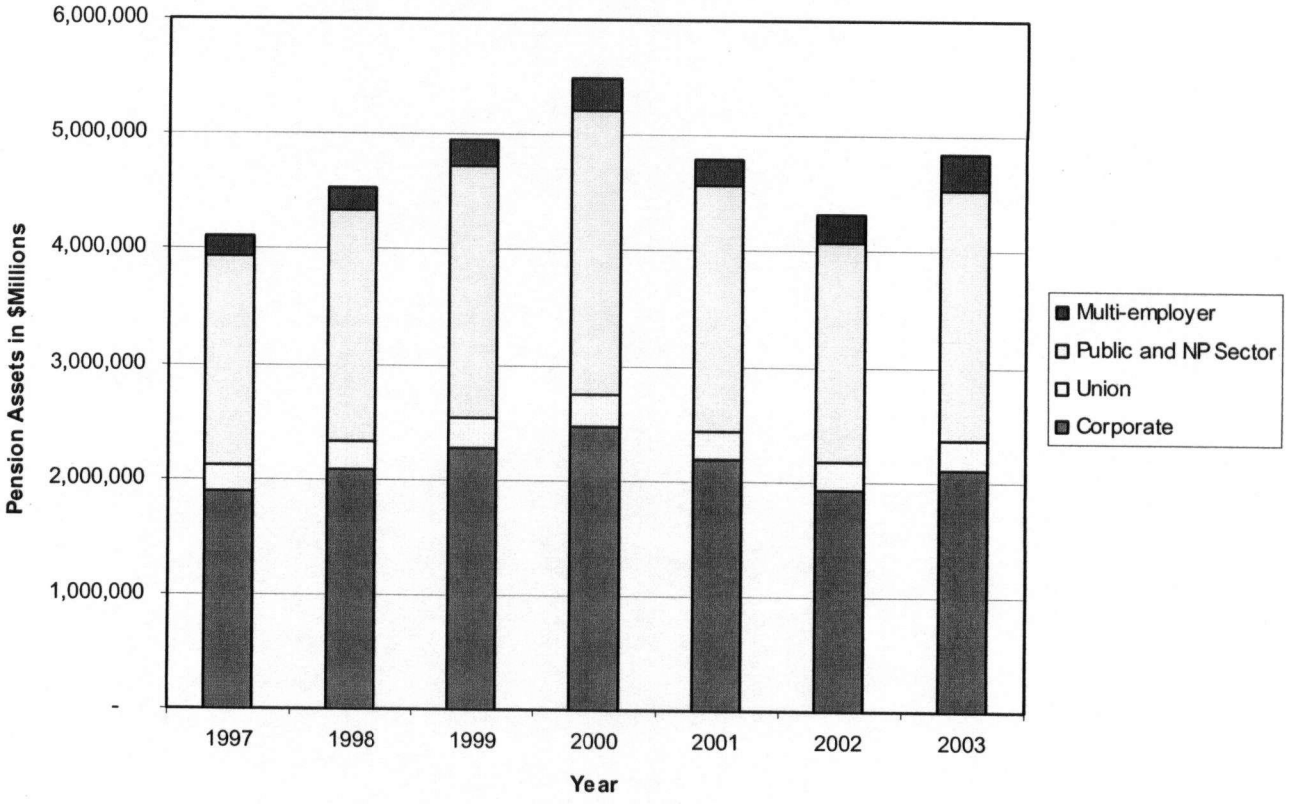

Figure 4-1 - Pension Assets for the Top 1000 Sponsors by Sponsor Type - Despite having diversified portfolios, overall pension assets followed the stock market, peaking in 2000 and recovering again in 2003.

Figure 4-1 shows that pension assets for the top 1000 sponsors peaked in 2000 and then fell significantly in 2001 and 2002 as the overall stock market fell. Figure 4-2 shows the relative importance of defined contribution plans by sponsor type. It also shows how the relative importance of the DC share has evolved over time.

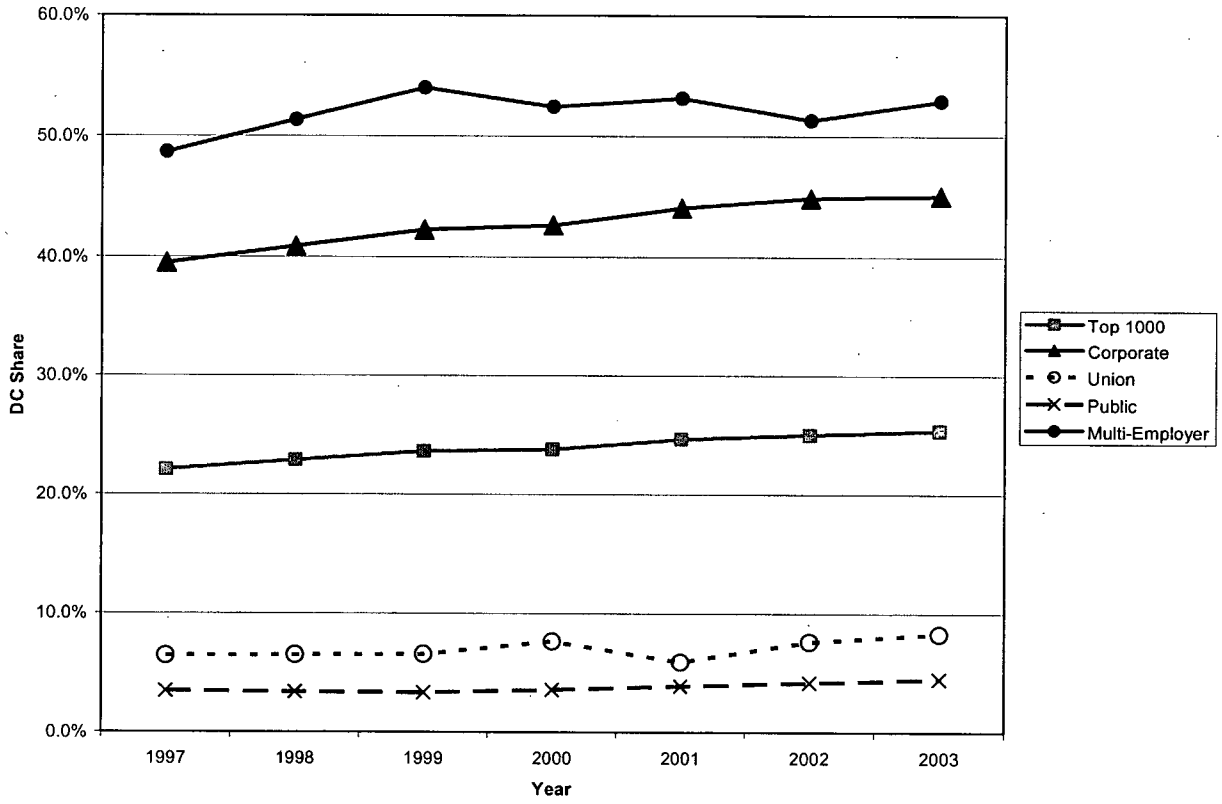

Figure 4-2 - Evolution of the DC Share in Total Pension Assets by Sponsor Type in the P&I 1000 - Although DC plans represent approximately only 20% to 25% of total pension assets, they are much more important in corporate (and multi-employer) pension plans than in public sector plans. Among corporate and multi-employer plans combined, DC assets are approaching a 50% share.

We have matched the P&I corporate pension data to Compustat data on firm performance and other characteristics of the firms. The data set used in our analysis consists of all corporate sponsors that were in the P&I top 1000 for at least one year in the 1997-2003 period and had the relevant Compustat data for that year. As a result we have a panel data set consisting of a set of large publicly traded firms tracked on an annual basis for up to 7 years. For an observation (corresponding to one firm for one year) to be "usable" it must have P&I pension data and Compustat return data for that particular firm and year. The set of usable observations varies slightly according to which

return measure is used, reflecting variations in data availability across return measures. Table 4.2 provides more detail on the structure of the panel of usable observations for the case in which the ROA is used as the return measure. We can see, for example, that there were 383 firms with usable data in 1997 and that 238 of these firms were in the data for all 7 years.

Table 4.2 - Panel Structure for Pension and ROA data– We study a total of 2879 firm-years of pension plans that have data in both the P&I Top 1000 and in Compustat. Over the 7 years covered there are 616 firms in total, of which 238 appear in all 7 years. This provides significant cross-sectional and firm-specific time series variation in the defined contribution share of pension assets and in ROA.

Year	Firms	# of years in data	# of firms
1997	383	1	86
1998	419	2	79
1999	423	3	50
2000	422	4	55
2001	407	5	49
2002	420	6	59
2003	405	7	238
Total Observations	2879	Total Firms	616

Our firm performance measures are drawn from the Compustat Industrial Annual dataset. Our primary measure of firm performance is ROA as defined by income excluding extraordinary items (Compustat item 18) divided by lagged assets (Compustat item 6). Although this is the most commonly used accounting-based return measure, Bertrand and Schoar (2003) note that managers have the capacity to undertake short-run manipulation of income. They find (as implied by their Table V on p. 1188) that the operating return on assets (OROA) is less subject to managerial fixed effects than the ROA. The OROA is defined as operating cash flow (Compustat item 308) divided by lagged assets (Compustat item 6). We find that our results are not significantly affected (and are slightly strengthened) by using OROA rather than ROA. We therefore continue

to emphasize the ROA given its widespread use as a performance measure. However, results using the OROA are reported in the paper.

When such measures are used it is important to be alert to the presence of outliers. Although both Compustat and the P&I data are of high quality, conceptual issues in the measurement and reporting of assets can lead to misleading return measures. In this case, however, outliers are not much of a problem. For our return on assets measure, the 1st percentile is about -21% and the 99th percentile is about 20%. These are plausible values for firms close to the extremes of the data. Table 4.3 provides descriptive statistics for the 1997 and 2003 cross-sections.

Table 4.3 - Descriptive Statistics - Representative summary statistics are provided for the first and last year studied. By 2003 more than 10% of corporate pension sponsors had no assets in defined benefit plans. There is a clear trend towards a higher proportion of DC assets. Pension assets and the Defined Contribution share come from P&I. ROA, corporate assets and employees come from Compustat.

1997	Pension Assets ($million)	% Defined Contr.	Return on Assets (%)*	Corp. assets ($billion)	Employees (thousands)
Minimum	432	0	-29.1	0.1	0.1
10th percentile	546	21.9	0.8	1.4	4.5
Median	1,259	41.4	5.3	5.6	21.0
90th percentile	6,596	79.7	11.3	36.3	113.9
Maximum	90,600	100	29.3	391.7	825
Mean	3,575	46.4	5.5	18.2	46.4
Standard Deviation	8,196	22.7	5.7	42.6	77.1
Observations	383	383	383	383	379
2003	Pension Assets ($million)	% Defined Contr.	Return on Assets (%)	Corp. assets ($billion)	Employees (thousands)
Minimum	557	0	-35.4	0.5	0.05
10th percentile	657	26.1	- 1.9	2.4	4.7
Median	1,657	49.3	2.9	10.7	24.5
90th percentile	8,738	100	11.0	87.4	122.6
Maximum	89,738	100	22.3	1,264.0	1,500
Mean	4,097	53.5	3.5	48.0	53.0
Standard Deviation	8,348	23.9	6.4	13.8	99.8
Observations	405	405	405	405	396

* Return on assets is defined as income excluding extraordinary items (Compustat Item 18) divided by lagged assets (Compustat Item 6).

4.4 Estimation

4.4.1 Simple OLS Regressions

The first statistical tests we report show year by year simple ordinary least squares (OLS) regressions relating the ROA to the DC share. The regression equation is given by

$$ROA_t = a_0 + a_1 DCshare_t + e_t \quad (5)$$

While these simple regressions are not to be taken as serious candidates for a "best specification", they do serve to establish clearly that there is a consistent positive relationship between the ROA and the DC share. The results are shown in Table 4.4.

Table 4.4 - Simple OLS Regressions - The return on assets is positively associated with DC Share in all years studied. In aggregate, it suggests that a firm moving from 40% assets in DC to 50% would have its ROA increased by about 0.5%, all else being equal. Dependent Variable = Return on Assets; t-ratios are shown in parentheses. *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively.

Year	1997	1998	1999	2000	2001	2002	2003	aggregate
DC share	3.24**	3.08**	8.2***	4.77***	6.27***	8.52***	6.34***	5.42***
	(2.5)	(2.1)	(5.8)	(3.5)	(3.8)	(5.6)	(4.9)	(9.9)
constant	4.0***	3.3***	1.6**	2.4 ***	-.6	-1.5*	.1	1.5***
	(6.0)	(4.2)	(2.2)	(3.3)	(.6)	(1.7)	(0.2)	(5.1)
adjusted R^2	.01	.01	.07	.03	.03	.07	.05	.03
observations	383	419	423	422	407	420	405	2879

Table 4.1 is striking. The return on assets is positively associated with the share of pension assets in defined contribution plans in each of the seven years. These effects are statistically different from zero in each year and are highly significant in aggregate. The results also suggest economic significance at high but plausible levels. ROA is measured in percentage points. DC share is measured as a fraction and is therefore between 0 and 1. Thus if a firm moved from having a 0.4 (or 40%) DC share to a 0.5 (or 50%) DC share, the aggregate regression suggests an increase in ROA by about half a percentage point from, for example, 5.0% to 5.5%.

The cross-sections (i.e. the years) are not simple replications of one another. As implied by Table 4.2 above, firms among the top 1000 pension sponsors exhibit considerable turnover. Furthermore, the return on assets shows considerable year to year variation. The share of pension assets shows more persistence, but even this variable changes significantly from year to year for many firms.

We hasten to acknowledge that there might be other factors that should be entered as control variables into a regression seeking to explain variations in ROA. We should also exploit the panel structure of the data, take advantage of other aspects of the temporal structure of the data, and take steps to assure ourselves that the DC share can be treated as an exogenous variable. Dealing with these and other issues requires more sophisticated econometrics. Nevertheless, Table 4.4 is sufficient to establish an interesting and robust fact: among large American corporations, those with relatively more important defined contribution pension plans also had better performance (as measured by return on assets) over the1997-2003 period.

4.4.2 Control Variables, Potential Endogeneity, and Heteroscedasticity

The natural next step is to incorporate control variables into our analysis. There are many possible influences on corporate returns. Omitted variables are interpreted as entering the error term. This does not cause bias in estimating the effect of the "treatment" variable (DC share) unless these omitted variables are correlated with DC share in such a way as to either exaggerate or disguise the effect of the treatment variable. One obvious effect of concern is that corporate returns vary over time with the business cycle. As there is a systematic change over time in DC share it is important to correct for

this effect. This can be handled by adding year fixed effects as regressors. A second point of concern is that ROA and DC share might both vary systematically by industry. We run the risk of attributing to DC share what is really just industry-by-industry variation. There may be specific industries where, on balance of the competing forces, it is less of an advantage to offer a DC plan over a DB plan. In general this might occur in cases where there is a lot of cost to the firm early in the career building up firm-specific human capital. Offering a DB plan might allow the firm to ensure that the employees are more likely to stay long enough to recoup this investment. This can be controlled for by introducing industry fixed effects. For industry fixed effects we use 2-digit North American Industry Classification System (NAICS) codes.[37]

There are, in addition, other control variables that might be relevant. Returns might be related to aspects of firm size, such as assets, property plant and equipment (PPE), or employment. There may also be firm specific factors which affect the choice of plan type. As Wang, Chung and Tzeng (2001) show, not just industry, but also plan age and unionization status can also affect the choice of plans and the return on assets. If there are important and stable firm specific effects influencing returns, they can be controlled for by using the first period (1997) return as a regressor. In addition, it is possible that relevant but unobserved aspects of the firm that affect ROA might be associated with lagged stock price changes. We report the effect of using lagged values of all these variables as control variables.

The endogeneity question is closely related to causality. It is difficult to infer causation with confidence, but the normal approach is to rely on the exogeneity of the

[37] The NAICS has replaced the Standard Industrial Classification (SIC) system. There are 18 two-digit industries represented in our data.

explanatory variable. To investigate the effect of x on y, we would like to exogenously vary x (holding other factors constant) and observe what happens to y. This would be strong evidence regarding any causal effect of x on y. Exogeneity can be achieved through controlled experiments, but in economics and finance it is rare that controlled experiments can by used. Instead, we normally rely on time. We can reasonably expect in most (but not all) situations that changes in x at time t might be exogenous with respect to (i.e. not caused by) changes in y at time t+1. In our case, we might look for the effect of prior changes in pension plan structure on subsequent corporate performance. Initially, we simply use lagged values for DC share as a regressor. In addition, we might expect heteroscedasticity. Accordingly, we correct for heteroscedasticity using the Huber/White/sandwich estimator as implemented in Stata 8. Results dealing with the points described above are shown in Table 4.5.

Table 4.5 - Regressions with Control Variables - To minimize any omitted variable bias we add several control variables. Year and industry fixed effects are natural inclusions to account for business cycle variation and industry-specific variation. The 1997 ROA corrects for initial differences in firm performance. Inclusion of control variables does not significantly affect the main results. Dependent Variable = return on assets in %; t-statistics are in parentheses. *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively.

Specification	1	2	3	4	5	6
DC share	5.42***	5.42***				
	(9.9)	(8.2)				
lagged DC share			6.09***	4.73***	6.18***	5.76***
			(7.9)	(5.9)	(6.4)	(5.7)
1997 return on assets				0.37***	0.32***	0.31***
				(9.0)	(7.8)	(7.48)
lagged assets						-2.13 *
						(1.9)
lagged employment						6.6***
						(4.1)
lagged PPE						20.9
						(1.3)
lagged stock price increase						.024***
						(3.5)
year fixed effects					F(5,11) = 14.4***	F(5,1911) = 12.9 ***
industry fixed effects					F(17, 2111) = 10.0***	F(17,1911) = 9.0 ***
constant	1.54***	1.54***	1.14***	-0.27	-2.3**	-1.39
	(5.1)	(4.6)	(3.0)	(0.6)	(2.4)	(1.4)
robust errors	no	yes	yes	yes	yes	yes
R^2	.03	.03	.04	0.14	.22	.25
observations	2879	2879	2231	2136	2136	1940

Table 4.5 contains a substantial amount of interesting information. First of all, the DC share has an economically and statistically significant positive effect on the return on assets, regardless of specification. Specification 1 repeats the last column of Table 4.5 showing the simple OLS regression with no control variables and no adjustments. It is present in Table 4.5 for comparative purposes. Specification 2 is the same except that robust standard errors are used to correct for heteroscedasticity. This reduces the statistical significance of the result only slightly and suggests that heteroscedasticity is not much of a problem. We continue to use this adjustment where possible. Specification 3 shows the effect of lagging DC share by one year. This should ensure strict exogeneity

of the DC share. Using lagged rather than current DC share has little effect but it is reassuring that using lagged DC share actually increases the magnitude of the coefficient on DC share slightly. Using lagged values causes us to lose a year of data. With 7 years of data in total, losing one year reduces the sample size significantly and, in this case, contributes to a marginal decline in statistical significance for DC share.

Specification 4 introduces the initial (i.e. 1997) ROA as a regressor to correct for initial (and persistent) profitability differences between firms. We expect that this captures idiosyncratic firm-specific effects. This variable is highly significant in itself and substantially increases the R^2 statistic. It reduces the magnitude and significance of the DC share effect only moderately. Specification 5 introduces year and industry fixed effects. Year fixed effects have a high level of joint significance, as indicated by the reported F statistic. If anything, correcting for year fixed effects increases the significance of the DC share. Our interpretation is that year fixed effects correct for declines in ROA due to business cycle effects that were experienced in the 2001-02 period relative to the 1997-2000 period. As the DC share was generally increasing over time, failure to adjust for business cycle effects would tend to provide a spurious weakening of the apparent positive effect of DC share on ROA. Industry fixed effects are essentially neutral. Incorporation of fixed effects for year and industry seems uncontroversial.

Finally, specification 6 incorporates several potential control variables, including lagged assets[38] (Compustat Item 6), lagged PPE (Compustat Item 8), lagged employment

---

[38] In order to get suitable scaling we use very large units for some control variables. Assets and PPE are measured in $trillion (and are therefore small fractions for each observation). Similarly, employment is measured in millions. Stock price changes are measured in dollars.

(Compustat Item 29), and the lagged change in the calendar year closing stock price (Compustat Item 24). The main result of including these various control variables is that they do not affect the magnitude or significance of the DC share very much.

Using a large group of control variables has costs and benefits. As noted earlier, control variables can correct for potential omitted variable bias. If, for example, higher employment levels for some reason cause both a higher ROA and a higher DC share, then failure to include it would tend to result in an exaggerated coefficient on the DC share. However, introducing questionable control variables can also inappropriately reduce the apparent significance of the treatment variable. For example, if a higher DC share led to higher employment but employment itself did not have an effect on ROA, then including employment would incorrectly tend to lower the estimated effect of DC share. Some researchers tend to err on the side of inclusiveness in selecting control variables. Others tend to avoid them unless they arise from a specific structural theoretical model. In our case, the battery of firm-specific control variables has little effect and we see little reason to include them. If anything, we would suggest that the initial ROA should act as a proxy for omitted firm-specific variables. On the other hand, we believe that year and industry fixed effects (which are not firm-specific effects) should be included. Accordingly, we view Specification 5 as the most relevant specification in Table 4.5. There is a strong and significant effect of DC share in all specifications.

4.4.3 Panel Data Estimation

The econometric specifications underlying Table 4.5 do not fully exploit the panel structure of the data. The natural next step is to use explicit panel data methods. Results are reported in Table 4.6.

Table 4.6 - Panel Data Estimation and First Differences – We take advantage of the panel structure of the data. Both time series variation (as shown by the "within" estimator) and cross-sectional variation (as shown by the "between" estimator) provide similar estimates of the effect of DC share. Dependent Variable = return on assets; t-statistics are in parentheses (except that z-scores are used in the random effects model); *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively.

Specification	Fixed Effects ("within" estimator)	Random Effects	"Between" estimator	First Differences (with robust errors)
DC share	4.22**	5.09***	4.54***	6.28**
	(2.37)	(6.39)	(4.9)	(2.3)
year fixed effects	F(6,2256) = 29.3***	Chi^2 (6) = 172***	F(6,608) = 2**	F(5,2207) = 12.4***
industry fixed effects				F(17, 2207) = .37
constant	3.74***	2.93***	1.79*	-1.18
	(4.2)	(5.9)	(1.89)	(1.2)
R^2	.06 (overall)	.06 (overall)	.05 (between)	.04
observations	2879	2879	2879	2231

The fixed effects estimator (sometimes called the "within" estimator) provides a fixed effect for each firm and therefore eliminates variation associated with any aspect of the firm that is time-invariant. In particular, industry effects are eliminated (i.e. rendered redundant) by the fixed effects estimator, as are initial (1997) returns. More importantly, persistent differences in DC share across firms will also be rendered redundant. Loosely speaking, fixed effects eliminate much of the cross-sectional explanatory power of the treatment variable (DC share). Any significance of the DC share variable must come from temporal variation in DC share within each firm. (If, for example, DC share were constant over time for each firm, then the fixed effects estimator would ascribe no explanatory power to DC share and would drop it from the regression.) In light of these properties we might expect the DC share to lack statistical significance when using fixed

effects. It is therefore striking that the even in the fixed effects regression the DC share is significant at better than the 5% level. It is also reassuring that the magnitude of the coefficient is very similar to the magnitude in the "between" estimation. In effect, time series variation and cross-sectional variation provide similar estimates of the effect of DC share on returns, although the statistical significance of the cross-sectional effect is greater.

The random effects estimator is reported in the second result column. We regard this specification as the preferred specification for this environment. As a computational matter, this specification uses both cross sectional and time series variation in DC share to derive estimates of its effect. At a conceptual level, we regard the random effects model as most reflective of the underlying process generating returns. Specifically, the random effects estimator posits that returns depend in part on a firm-specific effect that is random and is therefore not observed with certainty. The fixed effects estimator, in contrast, assumes that the firm-specific effect is fixed and equal to the calculated difference in intercept (from the base case) for this particular sample period. It places "too much" confidence in the idea that variation in returns is due to otherwise unobserved fixed differences in firms rather than to cross-sectional variation in DC share. In any case, the random effects specification provides an estimate of the effect of DC share that is strongly statistically significant and is similar in magnitude to the estimates in Table 4.5.

The "between" estimator, reported in the third result column, is a regression done on group means. It is comparable to a regression done on a single cross-section (i.e. a single year), but it properly incorporates all available data. It is referred to as the between estimator because it is based entirely on differences between subjects (i.e. firms) and

completely suppresses variation over time within a firm. Our panel is broad (with many firms) and short (with only 7 years). Furthermore the variable DC share evolves only relatively slowly. Accordingly, most of the variation in our data is cross-sectional. We therefore expect the "between" estimator to be statistically significant, as it is.

An alternative to using fixed effects or random effects panel data methods is to estimate a first differences model. This is reported as the fourth specification in Table 4.6. The dependent variable is the annual change in return on assets and the treatment variable is the annual change in DC share. As with the fixed effects method, this approach eliminates constant cross-sectional variation in DC share from consideration and we expect similar results to the fixed effects model. We include industry fixed effects in the regression but, as we might hope and expect when running a regression in first differences, these effects are not significant. The DC share effect is statistically significant in the first difference regression and its magnitude is similar to but slightly larger than that obtained using explicit panel data methods.

We did not use a battery of control variables in the panel data estimations or in the first difference model. Using fixed or random effects should control for firm specific omitted variables that are stable over time, and first differencing should achieve the same result.

4.4.4 Investigation of Temporal Effects Using a Balanced Panel

Although the panel data methods reported in Table 4.6 explicitly deal with the time series structure of the data, they do not take full advantage of the temporal information available in the data set. Given a 7 year time horizon and considerable

persistence in DC share, our ability to use time series information is limited. However, we can provide some useful information based on the full balanced panel of 238 firms that have return and pension data for all 7 years.

Table 4.7 - Balanced Panel - Although we are only able to use the 238 firms that are in the panel for the full 7 years, we can take full advantage of the cross-sectional and time-varying structure of this balanced panel. The DC share is still strongly significant when introduced either as the initial DC share or as the increase in DC share over the time studied. Dependent Variable: return on assets; t-statistics in parentheses. *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively

Dependent Variable	1998-2003 average ROA	1998-2003 average ROA	Change in ROA: 1997-2003	2003 ROA
1997 DC share	6.60***	5.50***	6.24**	
	(4.80)	(3.81)	(2.47)	
1997-2002 change in DC share				7.53**
				(2.28)
1997 return on assets		.32***	-.81***	.22*
		(3.81)	(7.0)	(1.7)
industry fixed effects	F(16,219) = 92 ***	F(17,218) = 1458 ***	F(16,218) = 199***	F(17,218) = 2393***
constant	0.58	-1.5	-0.95	2.1
	(0.3)	(1.0)	(0.4)	(0.9)
robust errors	yes	yes	yes	yes
R^2	.30	.42	.41	.18
observations	238	238	238	238

In the regressions reported in Table 4.7 there is one observation associated with each of the 238 firms present in the data for all 7 years. Thus the number of observations is dramatically lower than in the previously reported regressions. We might therefore expect less statistical significance. The first specification estimates the effect of initial (i.e. 1997) variation in DC share on subsequent (1998-2003) returns. This is achieved by regressing average returns for the 1998-2003 period and 1997 DC shares for the 238 firms that are in the data for all 7 years. In this regression the 1997 DC share has a strongly significant coefficient that is similar in magnitude to the coefficients associated with lagged DC shares in Table 4.6. Specification 2 is the same except that we also control for initial (1997) returns. Initial returns have a positive and significant effect, but

the explanatory power of the 1997 DC share is only slightly diminished and remains highly significant.

The third specification focuses on the effect of the initial DC share on the change in ROA. Specification 4 is essentially the mirror image of Specification 3 as we seek to estimate the effect of the change in the DC share over the longest possible period (1997-2002) on subsequent (2003) returns. The magnitude of the effect is comparable to that suggested by other regressions. In both of these regressions the DC share regressor is statistically significant at the 5% level or better and of significant economic magnitude.

4.4.5 Alternative Measures of Corporate Performance

This paper investigates the relationship between the share of pension assets in defined contribution plans and the return on assets (ROA). While the ROA is commonly used as a measure of corporate performance, there are many other such measures, including return on equity, total shareholder return, abnormal (and raw) stock market returns, total factor productivity, growth in earnings per share, etc. The first point to emphasize is that no one measure is universally "correct" while the others are "incorrect". Rather, it is important to recognize that each measure captures something different. The appropriate measure for any particular situation depends on the particular research question under consideration.

In this section we report results using two alternative return measures. One measure is the operating return on assets, which is given by operating cash flow (Compustat Item 308) divided by lagged assets (Compustat Item 6). This is similar to the standard ROA measure except that operating cash flow is thought to be less susceptible to

managerial manipulation than the numerator for ROA, which is income before extraordinary items (Compustat Item 18). The other alternative measure we use is return on equity (ROE), which is defined as income before extraordinary items (Compustat Item 18) divided by the book value of equity as measured by stockholder's equity (Compustat Item 216). We select several representative regressions for each measure, as reported in Table 4.8 and Table 4.9.

Table 4.8 - Operating Return on Assets - Using the Operating Return on Assets (OROA), which has been shown to be less susceptible to manager manipulation, instead of the ROA produces very similar results. Dependent Variable = operating return on assets in %; t-statistics are in parentheses. *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively.

Specification	1	2	3: fixed effects	4: random effects
DC share			4.07**	6.09***
			(2.42)	(6.94)
lagged DC share	7.52***	6.28***		
	(10.4)	(7.73)		
1997 return on assets		0.35***		
		(9.64)		
lagged company stock				
year fixed effects		F(5,1981)=4.9***	F(6,2121)= 12.5***	Chi^2 (6) = 76.9***
industry fixed effects		F(17,1981)= 8.4***		
constant	6.2***	5.3***	9.7***	8.3***
	(18.0)	(4.7)	(11.5)	(15.5)
robust errors	yes	yes	no	no
R^2	.06	.34	.06 overall	.06 overall
observations	2097	2006	2708	2708

The results using the operating return on assets instead of ROA as the dependent variable are similar in economic and statistical significance to the results with ROA as the dependent variable, but they are slightly stronger in general. Thus, if anything, using the OROA offers slightly stronger support for a positive effect of the DC share on returns than using ROA.

Table 4.9 - Return on Equity - Using ROE as the dependent variable leads to weaker results than ROA or OROA. This may be because of the more significant outlier problem that arises with ROE. Quantile regressions greatly reduce the influence of outliers and, in this case, restore a high level of economic and statistical significance for the DC share. Dependent Variable = return on equity in %; t-statistics are in parentheses. *, **, *** represent significance at the 10% level, 5% level, and 1% level respectively.

Specification	1	2	3	4:random effects	5:quantile regression
DC share				0.22	
				(.01)	
lagged DC share	14.0**	12.4**	18.6***		8.56***
	(2.40)	(2.04)	(2.84)		(9.12)
1997 return on assets		0.07			
		(1.03)			
lagged company stock					
year fixed effects			F(5,2207) = 2.02*	Chi^2 (6) = 8.3	F(5,2207) = 19.0***
industry fixed effects			F(17,2207) = 2.98***		F(17,2207) = 20.0***
constant	6.4*	6.2*	2.3	14.6	7.5***
	(1.8)	(1.7)	(0.6)	(1.0)	(5.4)
robust errors	yes	yes	yes	no	no
R^2	.00	.01	.04	.00 (overall)	
observations	2231	2136	2231	2878	2231

The effect of the DC share on the ROE has less statistical significance than the effect on the ROA, although the coefficient has a larger magnitude in some specifications. In various specifications the effect of the DC share is only marginally significant. In the panel data estimates it is not significant at all. Only the random effects estimation is reported, but the fixed effects and "between" estimates show a similar pattern. One concern about using ROE is that it much more prone to the influence of large outliers than is ROA. As firms can have very low levels of equity, they have very large values of ROE. For example, there are 3 observations with an ROE that exceeds 100% and 3 other observations for which ROE is less (i.e. more negative) than -100%. One way to reduce the weight of outliers is to use a quantile regression. The quantile regression used here estimates the conditional median (rather than mean) of the

dependent variable and is obtained by minimizing the sum of absolute residuals. A quantile regression is reported as specification 5. Using a quantile regression restores the magnitude and significance of the DC share effect to levels similar to those reported for ROA. This is true in specification 5 and it is also true of most of the other reported specifications.

All three measures of firm performance that we have considered rely on book or accounting values. In other contexts firm performance is often measured using stock returns based on market values rather than book values. The problem with using measures based on the stock price (including abnormal stock returns, raw stock returns, and the total return to shareholders) is that these measures identify the timing of changes in market expectations, not the timing of actual changes in performance. For example, suppose that, in 1997, some firms had a high value of the DC share. This should result in better corporate performance for those firms than for others. If this is recognized by investors, this effect will be fully reflected (i.e. fully capitalized) in the stock price in 1997 and no subsequent stock price appreciation due to this source will be observed. Thus, our data cannot be readily combined with stock price measures to investigate the effect of the DC share on investor expectations. In order to carry out a true event study based on stock prices it would be necessary, in our view, to identify previously unknown changes in pension plan structure and link these changes to daily stock prices. In this paper we are using annual data. Variations in the ROA and OROA and ROE measures provide a good broad-based indication of variations in firm performance on an annual basis. Each of these measures contains some "noise" but there is little reason to anticipate any bias in estimating the effect of the DC share on performance.

4.5 Issues for Discussion

In this analysis (as in any) there are several concerns that might arise in interpreting the results. We consider three issues of particular relevance.

4.5.1 Possible Selection Bias

A corporation enters our data set by being one of largest 1000 pension sponsors in the United States and by simultaneously having coverage in Compustat. Of the 383 such firms in the data in 1997 only 238 remain continuously in the data for the full 7 years. In total 616 different firms enter the data set for at least one year. The most important reason for exit is simply that other corporations have more rapidly growing pension assets so a particular firm might drop out of the top 1000. The second most important reason for exit is through merger or acquisition. The question is whether this process generates any selection bias that would affect the results. We believe that no significant selection bias would be expected. Inclusion in the sample is not based on any particular realization of the dependent variable (return on assets) or the primary "treatment" variable (the DC share).

It might be thought that persistently low ROA would lead to contraction of the firm and departure from the data set, inducing selection on the dependent variable. If so, this might call for using a regression technique (like the Tobit estimator) to handle censoring of the dependent variable. In fact, however, the correlation between ROA and the change in overall pension assets is virtually zero. If this effect is present it is extremely weak – certainly too weak to justify using a Tobit estimator. Furthermore, there is little reason to believe that the firms that are excluded or drop out of the data set

have a different relationship between the DC share and returns than firm that continue in the data set.

The only selection issue that seems important is that our firms are all large publicly held firms. Possibly small firms or privately held firms might be different in some relevant way. Therefore we would exercise caution in asserting that our findings necessarily apply to the entire economy. However, the firms in the data set comprise a large part of the economy in themselves and would seem to illustrate incentive effects that should apply even more broadly.

4.5.2 Defined Benefit Pension Plans and Income

Some readers might be concerned about possible bias in the results arising from the relationship between income and DB pension plan earnings. Our primary dependent variable, ROA, is defined as income excluding extraordinary items divided by assets. Income is itself affected by DB plans as the cost of funding a DB plan is a cost that is deducted in arriving at income as given by Compustat Item 18. As explained very helpfully in Kwan (2003), firms do not report actual payments to DB plans as expenses. (Recall that actual payment may fall short of the increase in net pension liability in any one year.) Instead, the Financial Accounting Standards Board (FASB) requires that income be adjusted by the net periodic pension cost (NPPC). This takes into account accrued liabilities and explicit costs of the pension plan minus the increase in actual and expected returns on plan assets. In any given year the NPPC "cost" may be either positive or negative. If the stock market is rising, then the NPPC tends to be a "negative" cost – meaning that it provides a net addition to income. If the stock market is falling the

opposite is true. There have been well publicized cases in which the NPPC provides a major part of income for a particular firm. This contrasts with the case of DC plans. With DC plans the employers' payments to the plan appear purely as a cost, just as with any other form of employee compensation. The question arises as to whether this differential treatment of DC and DB plans might bias our results in any way.

We see no concern for our analysis arising from the role of DB plans in income. The rather complicated NPPC approach is intended to approximate the true net cost of the pension plan and this cost is directly comparable to the contribution made by a firm to a DC pension plan. In both cases these are true costs comparable to other types of employee compensation. Conceptually, the NPCC approach is a correct way of determining pension plan costs to the firm and properly enters income. However, because it is only an approximation, systematic errors are possible. Nevertheless, we believe that any systematic errors in the NPCC approximation are much too small to have a noticeable effect on our results.

To the extent that the NPCC system might have had an effect, we note that stock price increases are amortized over a five year period for determining the NPCC. Accordingly, for most of our sample period (at least up until 2001) the NPCC system had a relatively favourable effect on income for firms with DB plans. Any approximation error would be in this direction. If anything, this would tend to increase the apparent relative return for firms with large DB shares and reduce the apparent return to firms with large DC shares. To the extent that this effect is present, any bias would therefore be against our primary finding that large DC shares are associated with higher returns. In

summary, taking account of the NPPC system would, if anything, increase the confidence we would have in our main results.

4.5.3 Endogeneity

It is clear that there is a correlation between the return on assets and the share of pension assets in employer-provided defined contribution (as opposed to defined benefit) pension plans. This effect comes through clearly in a regression context. It is, however, important to ask whether this effect, instead of being due to a causal influence of the DC share on the ROA, might instead be due to reverse causality. Perhaps ROA affects DC share or perhaps both variables are endogenous variables in a larger simultaneous system. We can think of these possibilities as representing possible "endogeneity" of the DC share as a regressor.

Prior theory is an important guide to answering this question. In this case we have suggested what we view as good reasons as to why pension plan structure might affect corporate performance. We can think of little theoretical rationale for reverse causality, although there are some possibilities. For example, it has been suggested that firms in financial "trouble" might terminate over-funded DB plans (if they have them) in order to obtain the surplus. This would be at most a small effect in our data.

At the empirical level we have exploited the panel structure of the data and the temporal aspect of the data to the extent possible to draw inferences about causality, including using lagged DC share as a regressor. While we cannot be definitive, it seems that changes in DC shares can be viewed as occurring prior to changes in ROA. Furthermore, when lags of various types are introduced, the estimated effects of the DC

share tend to become stronger in both magnitude and statistical significance. We have also undertaken "reverse" regressions in which we use ROA or lagged ROA as an explanatory variable for DC share. We do obtain positive and significant results for several specifications but these effects are weaker than the estimated effects of DC share on ROA. Accordingly, we feel that the evidence for an effect of the DC share on firm performance as measured by the ROA is relatively strong.

4.5.4 Choosing a Plan as an Employer

Despite the strong results presented in this paper, it is nonetheless clear that not all companies have chosen to offer only a defined contribution plan to their employees. There are several reasons that we offer to explain this, falling broadly within two categories. The first class of reasons that we do not observe 100% DC plans is that there are significant costs to changing plans. These costs include regulatory hurdles and plain inertia. People have become accustomed to the cradle to grave coverage from their employers and there is significant opposition from some groups of employees – likely those who would not self-select into DC plans. Despite this there is a growing trend to convert from DB to DC plans, and most new plans are DC only.

Secondly, the model presented in this paper suggesting that workers would self-select into firms that offer DC plans is but one of several competing factors that executives must consider in choosing a plan. There may also be solid reasons for choosing a DB plan. For example, as Black (1980) pointed out, DB plans can act as a tax shelter for corporate assets. Since any assets within a pension plan ultimately belong to the firm (though there is an offsetting liability), then the tax-free growth can be efficient.

Though this led the government to place funding limits of 120% on DB plans, it can still lead to significant tax savings. DB plans also provide a measure of financial flexibility. With a DC plan the employer must pay in a fraction of payroll every month without fail. With a DB plan there is considerable flexibility in both the amount and timing of any payments into the plan. These are some of the conflicting priorities that any executive must balance based on the specific situation of the firm.

4.6 Conclusions

The general objective of this chapter is to examine the relationship between pension plan structure and corporate performance. Higher shares of defined contribution assets in company-sponsored pensions plans are associated with higher levels of return on assets. This is true across firms in any given year and it is true over time for a representative firm. Thus both cross-sectional and time series variation are supportive of the hypothesis that the defined contribution share has a positive effect on the return on assets. These results also apply when the operating return on assets is used as the return measure. A similar (although statistically weaker) pattern emerges when return on equity is used as the performance measure.

Our theoretical expectations relate to worker incentives. Defined contribution pension plans tend to encourage later retirement than defined benefit plans. In principle, this could be either costly or helpful to the firm depending whether the worker generates net surplus for the firm in these extra working years. If anything, prior literature has suggested that encouraging earlier retirement is a positive effect of defined benefit plans. Our results suggest the opposite. Furthermore, we suggest a mechanism that appears to

have previously been ignored. Specifically, we suggest that DC plans encourage self-selection of highly productive workers into such plans as they can avoid the cross-subsidy they would otherwise provide for other workers. These more productive workers then work longer than they would under DB plans and therefore generate net benefits for the firm.

Defined contribution plans also allow easier portability as they have a less severe "vesting period" than DB plans and frequently have no vesting period at all. DC plans are therefore thought to offer fewer restrictions on worker mobility, especially in the early stages of a worker's career. Once again this could in principle either enhance or weaken firm performance. DC plans might reduce the incentive of firms to invest in the worker's human capital but they increase the incentive of workers to invest in their own human capital. DC plans also have the advantage of removing a friction that would otherwise inhibit worker transitions to jobs where they have a higher marginal product. Our analysis suggests that the positive effects of DC plans outweigh any negative effects.

The movement toward DC plans has been much slower in the public sector than in the corporate sector, perhaps because public sector organizations are less sensitive to economic incentives of the type considered here. On balance, we would suggest that more aggressive shifting toward defined contribution pension plans would be suitable for the public sector and that continued shifting toward DC plans in the corporate sector will probably continue to enhance economic performance.

4.7 References

Becker, Gary 1964, *Human Capital: A Theoretical and Empirical Analysis*, New York: Columbia University Press.

Bertrand, Marianne and Antoinette Schoar 2003, 'Managing with Style: The Effect of Managers on Firm Policies', *Quarterly Journal of Economics*, 118, pp. 1169 – 1207.

Black, Fisher 1980, 'The Tax Consequences of Long-Run Pension Policy', *Financial Analysts Journal*, 36 (1980), pp 21-28.

Bodie, Zvi, John B. Shoven, and David A. Wise 1989 eds., *Issues in Pension Economics*, Chicago: University of Chicago Press.

Cornwell, Christopher and Stuart Dorsey 2000, 'Are Pension-Providing Firms More Productive?' discussion paper.

Dorsey, Stuart, Christopher Cornwell and David Macpherson 1998, *Pensions and Productivity*, Kalamazoo, Michigan: W.E. Upjohn Institute for Employment Research.

Employee Benefit and Security Administration 2004, 'What You Should Know About Your Retirement Plan' at www.dol.gov/ebsa.

Employee Benefit Research Institute (EBRI) 2004a, 'Assets in Qualified Retirement Plans, 1985-2002: Revised', Facts from EBRI, September, as available at www.ebri.org/facts/index.htm.

Employee Benefit Research Institute (EBRI) 2004b, 'Are the Types of Retirement Plans Changing?' in FAQs About Benefits – Retirement Issues, as available at www.ebri.org/benfaq/retfaqt13.htm.

Friedberg, Leora and Anthony Webb 2003, 'Retirement and the Evolution of Pension Structure', NBER Working Paper 9999, Cambridge Mass.: NBER.

Huberman, Gur, Sheena Sethi-Iyengar, and Wei Jiang 2003, 'Defined Contribution Pension Plans: Determinants of Participation and Contribution Rates', Discussion Paper, Columbia University.

Ippolito, Richard A. and William H. James 1992, 'LBOs, Reversions and Implicit Contracts', Journal of Finance, 47, 139-167.

Ippolito, Richard A. 1995, 'Toward Explaining the Growth of Defined Contribution Pensions', Industrial Relations 34 (1), 1-20.

Ippolito, Richard A. 1997, *Pension Plans and Employee Performance*, Chicago: University of Chicago Press.

Kotlikoff, Lawrence and David A. Wise 1989, 'The Incentive Effects of Private Pension Plans' in Issues on Pension Economics, Z. Bodie, J. Shoven, and D. Wise eds., Chicago: University of Chicago Press, 283-339.

Kruse, Douglas 1995, 'Pension Substitution in the 1980s: Why the Shift toward Defined Contribution?', Industrial Relations 34(2): 218-241.

Kwan, Simon 2003, 'Pension Accounting and Reported Earnings', FRBSF Economic Letter, Federal Reserve Board of San Francisco, www.frbsf.org/publications/economics/letter/2003/el2003-19.html.

Lazear, Edward P. 1983, 'Incentive Effects of Pensions,' NBER Working Papers 1126.

Lazear, Edward P. 1986, 'Retirement from the Labor Force,' *Handbook of Labor Economics*, Elsevier Science Publishers.

Marcus, Alan J. 1987, 'Corporate Pension Policy and the Value of PBGC Insurance', in Issues in Pension Economics. Bodie, Zvi, ed. Shoven, John B., ed. Wise, David A., ed.,

National Bureau of Economic Research Project Report series Chicago and London: University of Chicago Press. 1987, 49-76.

Mitchell, Olivia S. and Sylvester J. Schieber 1998, 'Defined Contribution Pensions: New Opportunities, New Risks' in Living With Defined Contribution Pensions, O.S. Mitchell and S.J. Schieber eds., Philadelphia: University of Pennsylvania Press, 1-14.

Mittelstaedt, H. Fred 1989, 'An Empirical Analysis of the Factors Underlying the Decision to Remove Excess Assets from Overfunded Pension Plans', Journal of Accounting and Economics 11, pp. 399-418.

Ostaszewski, Krzysztof 2001, 'Macroeconomic Aspects of Private Retirement Programs', North American Actuarial Journal, 5(3): 52-6

Petersen, Mitchell 1992, 'Pension Reversions and Worker-Stockholder Wealth Transfers', The Quarterly Journal of Economics 107(3): 1033-1056

Stock, James, and David Wise 1990, 'Pensions, the Option Value of Work, and Retirement,' Econometrica, 58(5):1151-1180.

Stone, Mary 1987, 'A Financing Explanation for Overfunded Pension Plan Terminations', Journal of Accounting and Economics 11, 361-398.

Sundaresan, Suresh and Fernando Zapatero 1997, 'Valuation, Optimal Asset Allocation and Retirement Incentives of Pension Plans', Review of Financial Studies, 10(3): 631-660.

VanDerhei, Jack 1987, 'The Effect of Voluntary Terminations of Overfunded Pension Plans on Shareholder Wealth', Journal of Risk and Insurance 59, 817-838.

Wang, Jennifer Li-Ling, Chung, Ching-Fan, and Larry Yu-Ren Tzeng 2001, 'The Dynamics of Employers' Pension Choices', Working Paper

Chapter 5 - Concluding Remarks

The study of the information effect in the S&P 500 has interested finance academics for over 20 years. There have, however, been very few opportunities to test this due to the inability to separate out different effects. I take advantage of a natural experiment to isolate the information effect. I find strong evidence of an information effect, and this is robust to several different formulations of the problem. The presence of this effect casts some doubt on other studies of downward sloping demand curves using the S&P 500 as a laboratory. Other areas for future research related to the S&P 500 include a look at a cross-sectional analysis of companies that are removed. Are there company specific factors other than size that determine whether or not a company will be removed? For example some companies are removed when they are 497th largest. But why weren't the 3 smaller companies in the index removed? International markets are also an interesting and underdeveloped comparison for some hypotheses. For example it is estimated that there is $3-4 trillion indexed to the MSCI indices, and there have been some major historical rebalancings. Some work discussed above has looked at individual countries, but they are for various reasons not comparable across borders.

The two Dutch firms in particular are interesting because of the unique nature of Royal Dutch-Shell and Unilever[39]. Both firms have two different classes of shares which should trade together, but are often "mis-priced" and seem to disobey the law of one price. There is a small but significant thread of literature which studies these "siamese twins". It is an excellent laboratory to study the importance of location in trade, and the

[39] For example, see Froot and Dabora (1999).

information and liquidity transfer across markets, and this natural experiment could provide an important addition to this literature.

Although pension funds have been well studied in some branches of economics, they have been largely ignored in the finance literature, largely due to a lack of data. This is unfortunate as pension plans are growing in economic importance. They are important in finance in at least two dimensions studied above – their affect on capital markets and the effect of the pension plan on the firm. I study two different aspects of pension plans. Firstly I find that unlike studies in familiarity in many contexts from home-market bias to language barriers, that pension funds do not seem to use familiarity variables to decide whether or not to invest in venture capital. Even though venture capital has high information costs, and is concentrated in a few locations and industries – exactly the conditions that should lead to familiarity based investing – these variables seem to have no effect on the likelihood of owning venture capital. Secondly we look at the decision by the firm to have a defined benefit or a defined contribution plan. We conjecture that based on our self-selection model presented above, better employees are likely to choose firms with defined contribution plan. We find strong evidence that firms with more DC assets than DB assets have better performance as measured by ROA, OROA and ROE.

Pension plans are a more important part of the economy and the firm than they were in the past. In recent years there has been an increase in the correlation structure of stocks. Several explanations have been put forward for this but I feel that pension plans have been ignored. A defined benefit plan represents both a giant asset and a giant liability that must be properly accounted for. If most pension plans have largely similar

asset mixes then it is possible that the growth of pension assets explains the increased correlation.