SYSTEMATIC RISK CHARACTERISTICS OF CORPORATE EQUITY

Geoffrey Shuetrim

Research Discussion Paper
9802

February 1998

Economic Research Department

Reserve Bank of Australia

The views expressed herein are those of the author and do not necessarily reflect those of the Reserve Bank of Australia. The paper has benefited from discussions with Alexandra Heath, Marc Henry, Louise Keely, Philip Lowe, Henry Overman, Gleb Sandmann and David Webb and I am grateful to the Financial Markets Group for facilitating this research which forms part of my Ph.D. thesis at the London School of Economics and Political Science.

Abstract

This paper finds evidence that firms may manipulate their systematic risk. This contrasts with previously held views that changes in estimates of systematic risk were an artefact of the estimators used. The central finding is that firms take actions which result in their equity betas adjusting toward unity, where equity betas are a common measure of systematic risk. This convergence phenomenon appears to result in older and larger firms having equity betas that are closer to unity than smaller and younger firms. The relationship between equity beta convergence and firm size is reconciled with the well documented negative correlation between equity betas and firm size. Also, greater deviations of systematic risk from the market average are found to be associated with a higher probability of being delisted. Having refuted the hypothesis that observed changes in systematic risk are an artefact of the estimation process, some implications for asset-market efficiency are explored.

JEL Classification Numbers: G10, G11, G12

Keywords: systematic risk, equity betas, convergence

Table of Contents

SYSTEMATIC RISK CHARACTERISTICS OF CORPORATE EQUITY

Geoffrey Shuetrim

1. Empirical Objectives

What can be inferred about the behaviour of publicly listed corporations from the behaviour of their security returns? This paper develops a body of evidence characterising the degree of co-movement between equity returns of individual firms and the return on the entire equity market. The focus is on how the degree of co-movement changes through time and on how it is related to observed characteristics of the firms. The evidence strongly suggests that the degree of co-movement associated with a given firm's equity converges, through time, to the market average degree of co-movement. Previously, this convergence phenomena was thought to have been a statistical artefact of the estimation techniques. Robust evidence in this paper refutes this hypothesis, suggesting instead that the driving force behind convergence is the preferences of the managers or owners who control firms.

The degree of co-movement between a firm's equity return and the return on the equity market is commonly referred to as systematic risk. In finance parlance, systematic risk is that component of risk which cannot be diversified away by investing across a wide variety of assets. In this sense, the systematic risk of an individual asset return is that part of equity return volatility driven by economy-wide shocks rather than idiosyncratic or asset-specific shocks. Generally, the systematic component of equity risk is estimated using some normalisation of the covariance between the return on a firm's equity and the return on the market, however, broadly the market is defined. A firm's idiosyncratic equity risk is then defined as the residual variation in the firm's equity return.

In this paper, systematic risk is estimated using monthly equity returns on all stocks listed on the New York Stock Exchange (NYSE) from January 1926 to December 1992. The data are obtained from the Center for Research in Security Prices (CRSP) database. Systematic risk is estimated for slightly fewer than

4 000 firms. In estimating systematic risk, the ‘market’ return is measured as the NYSE weighted-average equity return. This corresponds to the interpretation of systematic risk as a measure of the degree of co-movement between the rate of return on the individual firm and the rate of return on the entire equity market.

The evidence is presented in two stages. First, the law of motion governing the systematic risk is estimated for various sub-samples. Second, the relationship between systematic risk and both the age and size of firms is characterised. The methodology for estimating both the law of motion and the relationships between systematic risk and observed features of the firms is similar to that used by Quah (1996) to examine macroeconomic convergence across nations. The methodology characterises the entire distribution of observations rather than focusing on the representative firm. This type of analysis is shown to deliver greater insight, especially in the analysis of the link between firm size and systematic risk.

Aside from providing robust evidence that the convergence of systematic risk is driven by the behaviour of firms rather than by flaws in the estimation techniques, this paper develops several interesting relationships between the systematic risk of equity returns and characteristics of firms. For example, firms with extremely high systematic risk and firms with extremely low systematic risk both have a relatively high probability of being restructured compared to firms with systematic risk that is closer to the market average. This paper also shows that larger firms and older firms tend to have systematic risk exposure that is closer to the market average than do smaller and more recently listed firms. Together, the wealth of empirical regularities are strongly suggestive that systematic risk is manipulated within firms. The apparent relationships between systematic risk and firm characteristics reinforces this message because they would not arise if the observed behaviour of systematic risk was driven by estimation techniques alone.

The remainder of this paper is structured as follows. Section 2 examines the older literature associated with tests of market efficiency and asset-pricing models. This examination is necessary to understand many of the issues addressed by an alternative methodology, presented in Section 4, for analysing systematic risk convergence. Section 3 describes the techniques used to estimate systematic risk for each month and firm in the dataset. It describes the data used and documents some of the more elementary features of the systematic risk estimates. Section 4 presents the key findings, beginning with a thorough characterisation of the time-series behaviour of systematic risk. The systematic risk estimates are then related to other properties of equity returns and, importantly, to the size and the age of firms. Finally, Section 5 summarises the findings and ties them into the conclusion that systematic risk is being manipulated by firms.

2. Is Systematic Risk Convergence a Statistical Artefact?

To date, attention has focused on the behaviour of a particular estimator of systematic risk, popularised in the late 1960s and early 1970s by empirical studies of the Capital Asset-pricing Model (CAPM). Before describing the early evidence of systematic risk manipulation, it is necessary to discuss how the CAPM as developed by Treynor (1961), Sharpe (1964), Lintner (1969), Mossin (1969), and Black (1972) is fundamentally related to measures of systematic risk.

The CAPM is a single-period model of asset returns, based on the mean-variance optimisation of Markowitz (1959) and the equilibrium assumption that markets clear. Mean-variance analysis delivers two-fund separation such that the expected return on any minimum-variance portfolio can be expressed as a linear combination of the expected returns on any two other distinct minimum-variance portfolios. In the zero-beta CAPM of Lintner and Black, the two minimum-variance portfolios of interest are the market portfolio and the zero-beta portfolio (the portfolio with zero correlation with the market portfolio). Alternatively, in the standard CAPM, the two minimum-variance portfolios of interest are the market portfolio and the risk free security. The market clearing condition implies that the return on the market portfolio is efficient in the sense that it is the portfolio with the minimum variance, given its expected return.

The CAPM defines a useful empirical measure of an asset's systematic risk. This measure of systematic risk is the coefficient, b_i, on the market portfolio's expected return in the equation defining the equilibrium relationship between excess returns on the firm and excess returns on the market:

$$E\left(r_{it} - r_{zt}\right) = b_{it} E\left(r_{mt} - r_{zt}\right) \tag{1}$$

where r_{it} is the return on asset *i*, r_{zt} is the return on the risk free or zero-beta portfolio and r_{mt} is the return on the market portfolio.

To estimate the relationship between returns (with its CAPM interpretation), several substantial steps must be taken. First, Equation (1) must be converted from its *ex ante* form by replacing the expected returns with observed data. This conversion implicitly assumes that the rate of return on any asset is a 'fair game' so that, over many realisations, the expected return will equal the average return. In other words, expectations are not biased. With normally distributed returns and independent expectation errors, this fair game assumption delivers an *ex post* equation:

$$r_{it} - r_{zt} = b_{it}\left(r_{mt} - r_{zt}\right) + e_{it} \tag{2}$$

It is also necessary to identify and measure the return on the market portfolio so that it can be used as a regressor on the left-hand side of Equation (2). In this regard, r_{mt} is often approximated by an average return on the entire equity market. Then, in a population regression, b_{it} is estimated by:

$$\hat{b}_{it} = \frac{Cov\left(r_{it}, \tilde{r}_{mt}\right)}{Var\left(\tilde{r}_{mt}\right)} \tag{3}$$

where $\tilde{r}_{mt}$ is the proxy for r_{mt}. Clearly $\hat{b}_{it}$ is a positive linear transformation of the covariance between firm *i*'s return and the proxy for the market return. In this way it can be interpreted as a measure of systematic risk.

As Roll (1977) points out, the market portfolio includes all possible assets. Specifically, it is not sufficient to use an average of returns within the *equity* market alone. Roll shows that the only valid way to test the CAPM is to test whether the true market portfolio is mean-variance efficient. Because the true market portfolio is impossible to construct, Roll concludes that standard CAPM testing strategies, using proxies for the market portfolio, are uninformative.

However, this caveat on the literature testing the CAPM does not interfere with the usage of estimated betas in Equation (2) as measures of systematic risk. Interpreting the betas in Equation (2) as measures of systematic risk does not depend on a market-clearing condition or mean-variance efficiency of the approximation for the market return. Thus, approximating the market portfolio with the value-weighted average return on the equity market does not interfere with the interpretation of the betas, estimated from dynamic versions of Equation (2), as measures of market-wide or systematic risk.

Given that the CAPM yields a useful measure of systematic risk, one would expect variations in systematic risk to be evident in studies that attempt to estimate Equation (2). During the period of intensive testing of the CAPM in the 1970s and early 1980s it was widely observed that estimated equity betas evolve through time to eventually exhibit systematic risk characteristics that are similar to those of the entire equity market.[1] More precisely, in sequential sub-samples, estimates of beta tend towards unity.[2] Black, Jensen and Scholes (1972), Blume and Friend (1973), and Fama and MacBeth (1973) refer to this phenomenon as beta convergence. These studies all test the CAPM in two stages because of a suggestion by Blume (1970) that measurement error in estimates of equity betas for individual securities is ameliorated to some extent by using the equity returns on *portfolios* formed from groups of firms in the sample. First, equity betas are estimated for all stocks individually. The stocks are ranked by these initial beta estimates and grouped into a number of portfolios. The ranking process is intended to retain variation in equity betas across portfolios of firms.

In the second stage, a subsequent sample period is used to compute betas for each portfolio. These portfolio betas are then used in cross-section regressions explaining excess rates of return using equity betas and a number of additional explanatory variables that should not be significant under the assumptions of the CAPM model. Typically, these and other studies have found that there is a consistent tendency for the second-stage beta estimates to be less extreme than the first-stage beta estimates. The first stage beta estimates of portfolios can be shown to be the average of the first stage betas estimates of the firms comprising the portfolio.

The most popular explanation for this convergence phenomenon is that it is driven by measurement error. This explanation has Bayesian foundations. Fama and

---

[1] Contributors to this literature include Friend and Blume (1970), Black, Jensen and Scholes (1972), Miller and Scholes (1972), Blume and Friend (1973), Blume and Husic (1973), Fama and MacBeth (1973), Basu (1977), Reinganum (1981), Litzenberger and Ramaswamy (1979), Banz (1981), Gibbons (1982), Stambaugh (1982) and Shanken (1985).

[2] The beta for the market return is unity by definition.

MacBeth (1973, p. 615) explain that forming portfolios of securities that have been ranked by their estimated betas 'causes bunching of positive and negative sampling errors within portfolios'. The intuition is that, again from Fama and MacBeth, 'in a cross section of $\hat{b}_i$, high $\hat{b}_i$ tend to be above the corresponding true b_i and low observed $\hat{b}_i$ tend to be below the true b_i'. A firm is grouped into a low beta portfolio either because it had a low beta *or* because its beta estimate had a negative measurement error. Thus negative measurement errors are bunched in the low beta portfolios. Similarly, the positive measurement errors tend to be bunched in the high beta portfolios. When new betas are estimated for the portfolios in subsequent time spans, the measurement errors within each portfolio have zero expected value and so a convergence of extreme portfolios towards the market beta should be observed.

However, other authors have discovered a similar convergence phenomenon when estimating the equity betas of individual firms rather than portfolios. Klemosky and Martin (1975) show that naive, no-change forecasts of a firm's beta often have twice the mean-square forecasting error of methodologies that explicitly adjust OLS beta estimates towards the market beta. This paper and other studies focus on the betas of individual firms suggesting that portfolio formation is not the only reason for the convergence witnessed in CAPM tests. This finding has provoked a more focused study of the convergence phenomenon. Is beta convergence a statistical artefact or a behavioural phenomenon?

The statistical artefact argument has been made rigorous in Blume (1975) and Vasicek (1973) using the Bayesian concept of prior distributions. Both argue that in the cross section we tend to observe equity betas that are normally distributed around unity and concentrated between zero and two. Using this prior information to form a Baysian estimator of equity betas should eliminate the measurement error bias because it is a weighted average of the classical estimator and the prior expected value of the equity beta.[3] The weight placed on the classical estimator depends on the information content in the data sample. With such short data samples being used to estimate equity betas in tests of the CAPM, the sample likelihood function does not dominate the prior information and so the prior information will adjust estimates of equity betas toward unity. This adjustment

[3] The weighted average property of the Baysian estimator requires the loss function to be quadratic.

should eliminate the measurement errors that have affected previous tests of the CAPM.

Using portfolios of equity securities, Blume (1975, p. 785) tests whether the beta convergence observed in the CAPM tests can be entirely explained by the omission of prior information and finds that the convergence tendency of adjusted estimates of equity betas remains significant at the five per cent level. The test is of the null hypothesis that a cross-section regression of adjusted betas in period one on actual betas in period two yields a slope coefficient equal to zero. Blume concludes that '*a major reason for the observed regression [to unity] is real non-stationarities in the underlying values of beta and that so-called 'order bias' is not of dominant importance*'.[4]

Although this paper reaches similar conclusions, it does so using a substantially different methodology. Beta estimates are obtained for individual firms rather than for portfolios preventing measurement errors from being aggregated as in previous studies. While betas for individual firms have large standard errors compared to the betas of portfolios, this should not cause *convergence*. This paper is also differentiated from previous work by the fact that the estimation of sequential betas is done using the Kalman Filter which optimally updates the next beta based on the current beta and the next observation on equity returns.

This approach, of formally modelling the variation in equity betas avoids the contradiction, inherent in previous studies, of estimating betas under the assumption that they are fixed in sample and then making inferences about their movements through time. Obtaining a time series of equity betas for individual firms facilitates characterisation of the law of motion governing the adjustments from the current beta to the beta in the next period. It is the features of this law of motion, rather than the types of tests conducted by Blume that form the basis of the conclusion that equity betas tend to converge toward unity.

Despite these substantial differences, Blume's Bayesian approach to incorporating prior information is retained by applying a tight prior to the initial estimate of each firm's equity beta. The initial beta estimate is then updated using the Kalman filter so that the prior information gets carried through to estimates of beta throughout the

---

[4] Order bias is that bias in systematic risk estimates arising from measurement error.

entire time dimension for each firm. This use of prior information prevents measurement error from being related to the extremeness of beta estimates.

This paper also goes beyond the previous studies of equity beta convergence by characterising the relationships between the distribution of equity betas across firms and observed characteristics of firms. The measurement error view of beta convergence cannot explain the increased probability of delisting for firms with extreme betas or the relationship between the concentration of betas and the size and age of firms. By demonstrating that such relationships are a robust feature of the data, the view that beta convergence is a statistical artefact is more strongly refuted.

Remaining is the alternative hypothesis that beta convergence is behavioural. This theory is strongly supported by Blume (1975) who finds that the order bias arising from measurement error in the context of portfolios of securities explains less than half of the observed convergence in equity betas. He concludes that firms with extreme betas seem to take on investment projects with less extreme risk characteristics. This conclusion ignores the many channels through which firms may adjust the risk characteristics of their equity returns by manipulation of their financial structures. Equity issues, leveraged buy-outs and equity carve-outs are examples of such manipulations that will influence the systematic risk profile of the income stream generated by equity holdings. Nonetheless, the essential point that beta convergence is not merely a statistical artefact remains clear. The next two sections investigate Blume's results in a more sophisticated manner, avoiding the many difficulties plaguing beta estimation using portfolios of securities.

3. Characterising Systematic Risk

Analysis of the systematic risk associated with equity returns begins by obtaining minimum mean-square estimates of equity betas for each observation, indexed by *i* for the firm and *t* for the month. The model is a generalisation of the CAPM type relationship used to discern the systematic risk component in equity returns, $b_{it}r_{mt}$, from the idiosyncratic component, e_{it}:

$$r_{it} = g_t + b_{it}r_{mt} + e_{it} \tag{4}$$

where r_{it} is the equity return for firm *i*, r_{mt} is the value-weighted average return on the NYSE, including all distributions of income such as dividends and bonuses.

How is it possible to estimate the separate equity betas for each observation as is necessary to capture the cross-section and time dimensions of equity betas? The answer is that a substantial structure needs to be placed upon the model defining equity betas. Specifically, the coefficients, g_t and b_{it}, are assumed to adjust through time according to the transition equations:

$$
\begin{aligned}
g_t &= g_{t-1} \\
b_{it} &= b_{it-1} + h_{it}
\end{aligned}
\qquad (5)
$$

The intercept coefficient, g_t, is time invariant in the sense that it is not affected by a shock in each period whereas the slope coefficient, b_{it}, is a random walk process, adjusting by the shock, h_{it}, in each period. It is assumed that the idiosyncratic shocks to returns are normally distributed with mean zero and variance, s^2_{ie}. Likewise, the shocks that change the beta coefficient from period to period, h_{it}, are assumed to be normally distributed with mean zero and variance s^2_{ih}. Further, Both shocks are assumed to be independently and identically distributed through time and are independent of each other.

From an econometric perspective, Equations (4) and (5) define a state space form that can be estimated using the Kalman filter to extract a sequence of betas for a given vector of hyper-parameters (see Harvey 1989). Equation (4) is the measurement equation and Equation (5) is the transition equation. Constrained maximum likelihood methods are then applied to the log-likelihood function formed from the resulting prediction-error decomposition. The model is an extension of the random walk with noise model in Harvey (1989, p. 37) allowing for an explanatory variable with a time-varying parameter. By imposing this structure on the betas, the problem of estimating betas for every time period is reduced to one of estimating only s^2_{ie} and s^2_{ih}. When initialising the Kalman Filter, it is necessary to specify the current ‘state’ of the system. This paper assumes that the initial intercept is zero with a diffuse prior while the beta coefficient is unity with a prior variance of 0.25 which approximately matches the cross-sectional variance of betas around unity observed by Vasicek (1973).

It should be noted that the random walk transition equation for b_{it} prevents the beta extraction model from imposing any form of mean reverting behaviour. Rather, the equity betas have been estimated under the null hypothesis that betas do not converge. This strengthens the conclusions, drawn in the next section, that betas do exhibit convergence to unity.

A generalisation of the model would allow the intercept to vary over time as well as the equity beta coefficient. This generalisation is rejected by the data which cannot distinguish between shocks to the intercept and the idiosyncratic shocks to returns in the measurement equation, e_{it}. Stationary processes for the equity betas were also explored. However, in the majority of cases, the auto-regressive coefficient on the lagged slope coefficient in the state equation approaches unity when maximising the log-likelihood function. While the density of betas appears to be stationary, the samples used for firms are too short and high frequency to capture their mean reversion.

Given estimates of the hyper-parameters and applying the fixed-interval smoother of Jazwinski (1970, pp. 216–217) to condition on the full set of observations, $t=1\ldots T$, yields estimates of $\hat{b}_{it}$ for all observations. The smoothing process is required to ensure that the precision with which betas are estimated is independent of their placement in the time dimension of the sample.

3.1 The Data

The time-varying parameter model described above is estimated for every firm with consecutive data for one year or more. This reduces the number of firms in the sample from 4 343 to 3 992. It also means that inferences from this study are only applicable to firms that survive the first year after their initial public offering (IPO). Using firms with only twelve months of data could give beta estimates that are very imprecise given that 1 004 firms in the sample have between 12 and 60 observations. However, the imprecision in their beta estimates should not lead to evidence of beta convergence. While the imprecision will overstate the mobility of equity betas, it will not bias the results toward beta convergence because of the tight initial prior around unity. The tight prior will tend to force firms with uninformative data to have equity betas that start close to unity and diverge through time.

To implement the beta estimation process, the market return is proxied by the value-weighted average return on the NYSE. In light of Roll's (1977) critique, using the value-weighted average return on the NYSE as a proxy for the market limits the implications of this work for the CAPM. However, this paper does not attempt to advance the literature on CAPM testing. Rather, the estimates of the equity betas are interpreted as measures of covariation between the return on an individual firms' equity and the returns on the equity market. To the extent that these equity betas are related to features of firms, the use of the value-weighted average return on the NYSE does not flaw the analysis.

There are several reasons why using the value-weighted average return on the NYSE is preferable to using more complex measures of market performance. First, managers and investors can easily compare firm performance to that of the value-weighted average return on the NYSE. This makes direct endogenous responses of firms to their equity betas more plausible. More complex characterisations of the market portfolio (e.g. incorporating fixed income assets, real estate and even non-marketable assets) which should be used to test the CAPM, may not be as relevant to the investigation of firm behaviour precisely because they are not easily observed by firms. Alternative covariation benchmarks like that used in Breeden's (1979) consumption CAPM, are also less relevant, despite their sophistication, because of the difficulties in adjusting consumption data to obtain a reasonable measure of consumption flow. A second point in favour of using the value-weighted average return on the NYSE is that Fama and French (1992) report that it yields similar results to using the broader valued weighted equity return on NYSE, AMEX and NASDAQ listed firms.

One way in which this paper differs from the literature testing the CAPM is that betas are constructed using raw, rather than excess, returns. Before presenting these betas, it is worth emphasising how small an effect this has for each firm. Figures 1 and 2 graph the standard errors for the measurement equation and transition equation respectively. They show the estimates computed using raw returns on the horizontal axis against those computed using excess returns on the vertical axis.

Figure 1: s_{ie} Comparison

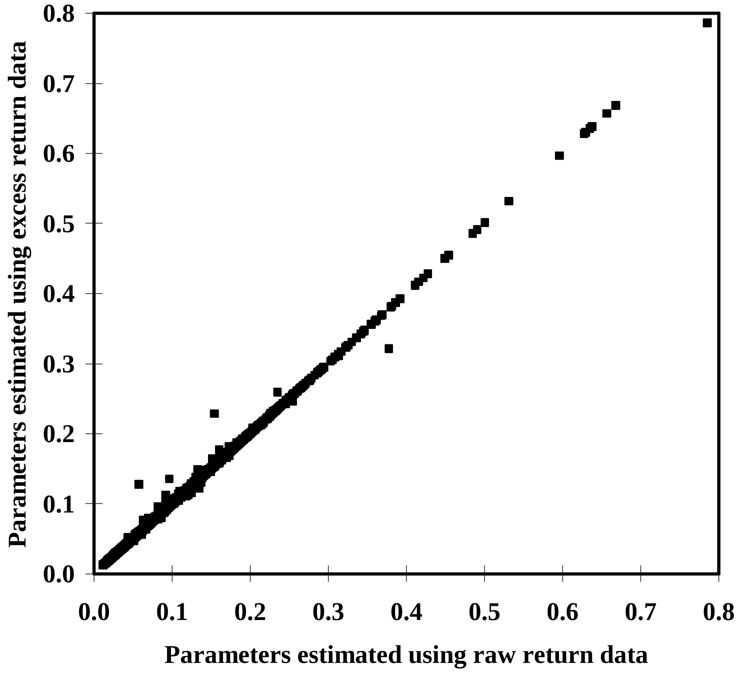

Figure 2: s_{ih} Comparison

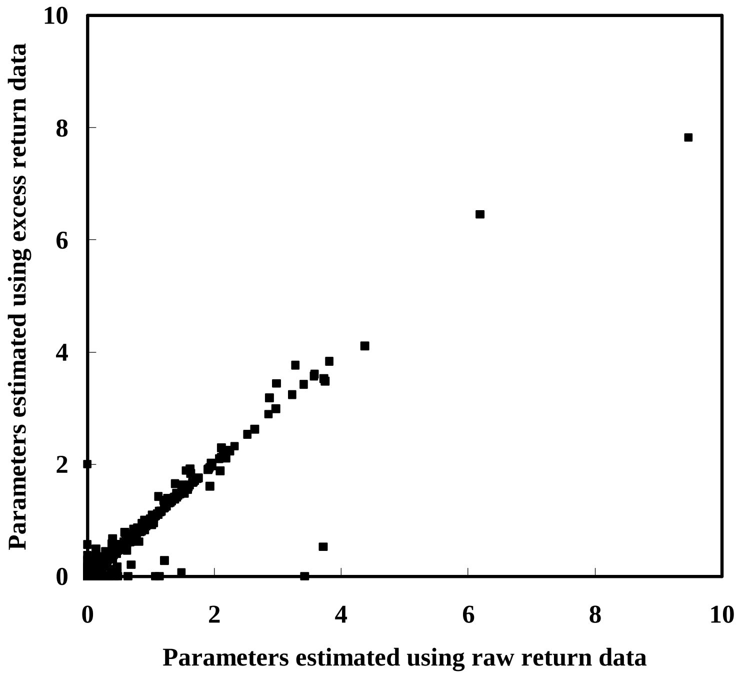

The fact that both of these figures approximately form a 45° line suggests that the choice of using raw or excess returns is not going to influence results significantly. This suggestion was verified by explicitly generating the reported results using excess returns. Raw returns have been made the focus of attention to clarify the interpretation of equity betas as measures of co-movement between firm performance and the performance of the entire equity market. Estimating equity betas using raw returns also de-emphasises any contribution being made to the asset-pricing literature.

3.2 Estimates of the Equity Beta Models

Although reporting the maximum likelihood estimates for each of the firms is uninformative, it is possible to estimate density functions showing how each of the hyper parameters, s_{ie} and s_{ih}, is distributed across firms. These estimated density functions are shown in Figures 3 to 5. To interpret the density functions, it is important to know how they have been constructed.

The density functions are estimated by pooling across firms and using non-parametric kernel-density estimation as discussed in Silverman (1986). Intuitively, the density estimate is a smoothed histogram wherein each observation in the histogram is replaced by the kernel function. The kernel function is simply a continuous, differentiable function that integrates to unity. In other words, it is a density function itself. In this paper, the standard normal distribution is used as the kernel function. The individual kernel functions, one for each observation, are integrated to obtain the estimate of the population density from which the sample has been drawn. By replacing each observation with the kernel function, this density estimate is continuous, smooth and it integrates to unity.

Formally the kernel density estimate at x of random variable X is given by:

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{x - X_i}{h} \right) \tag{6}$$

where, X_i is the ith realisation in the sample, h is the window width, n is the number of observations in the sample and $K(\cdot)$ is the kernel function.

The smoothness of the density estimate depends upon the choice of the smoothing parameter, h, referred to as the window width. This parameter is a scalar value for univariate density estimation. It defines the extent to which the probability mass associated with each observation is smoothed out over the support of the density. The larger the window width the more smoothing occurs in the estimation procedure because each observation is spread over a wider region of the support. In most cases greater smoothing reduces the variance of the density estimate while increasing the bias.

When constructing the density estimates, a subjective approach to window width selection is adopted. The subjective approach is recommended by Silverman (1986) in situations where interest focuses on the shape of the density rather than on applying more formal non-parametric inference techniques. Experimentation suggests that the information content of the univariate density estimates is unaffected over a wide range of window widths.

The estimated density functions in Figures 3 to 5 are suggestive of the range of parameter values obtained across the sample of firms. For example, Figure 3 indicates that for most firms, the standard deviation of idiosyncratic shocks to equity returns is less than 0.15. Likewise, from Figure 4, the shocks to equity betas for most firms have a standard deviation below 0.025. The density functions provide a feel for the location and dispersion of the hyper-parameter estimates. It is clear from Figure 5 that the majority of firms have very low signal-to-noise ratios, s_{ih} / s_{ie}, implying that betas generally adjust very slowly compared to the overall volatility of equity returns. However, the estimated signal-to-noise density has a substantial upper tail with a few extreme firms having signal-to-noise ratios above one. The hyper parameters in these cases are usually very imprecisely estimated.

Figure 3: Density of s_{ie}

The window width is 0.01

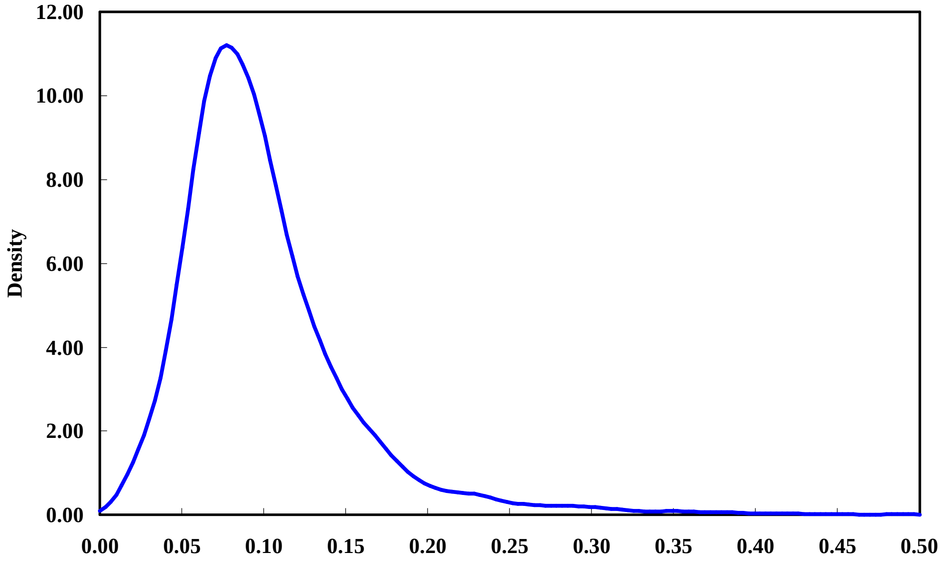

Figure 4: Density of s_{ih}

The window width is 0.01

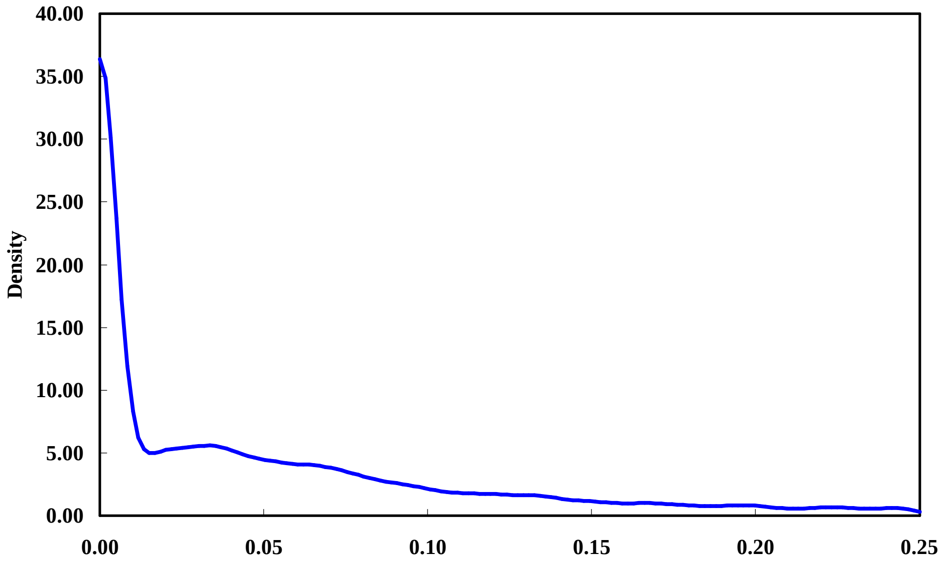

Figure 5: Density of s_{ih} / s_{ie}

The window width is 0.01

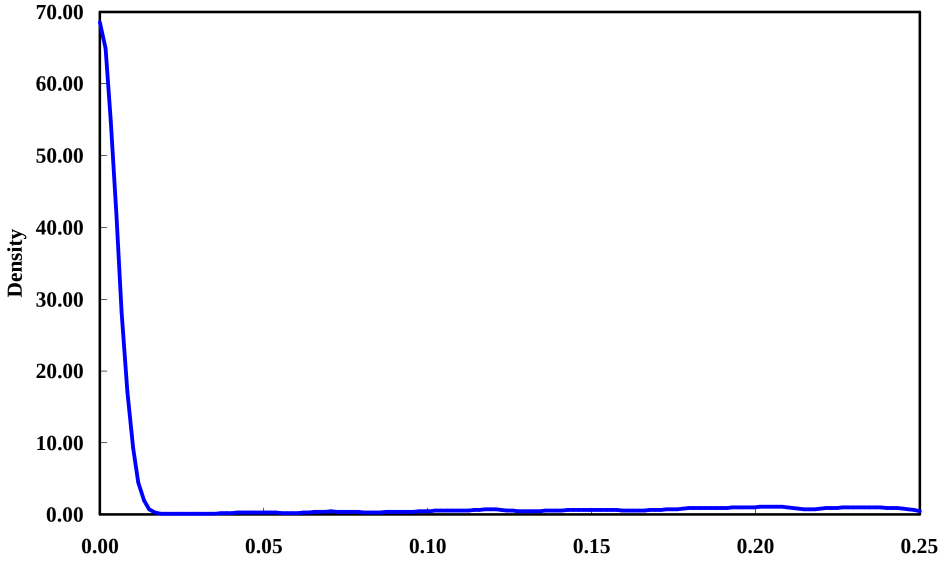

3.3 Equity Betas and Other Characteristics of Equity Returns

Additional information about firms' equity returns is provided in the form of densities of equity betas and the standard errors of estimated equity betas. The standard errors of the equity beta estimates are computed directly from the sequence of state variance-covariance matrices obtained for each firm. The data has been arbitrarily broken up into the sub-samples: 1926–50, 1951–75 and 1976–92. A visual comparison of the densities in Figure 6 and in Figure 7 is an informal means of examining the assumption that it is valid to pool observations across the entire time dimension.

The density functions generated for different time periods are fairly homogenous, confirming the validity of pooling across time. Stability of the densities through time was formally tested using the Kolmogorov (1933) and Smirnov (1939) test for which critical values are tabulated in most textbooks on non-parametric testing methods. These tests contradict the visual message that the densities are similar by powerfully rejecting the null hypothesis that the densities are stable through time even at the one per cent level in a two-sided test. The importance of the rejection of the null hypothesis is difficult to assess because it is driven by the extremely high power afforded by the large dataset. The rejection of the null hypothesis is a common phenomenon with large datasets and is an implication of the fact that null hypotheses are statements of equality that are almost never going to be true.

Figure 6: Equity Beta

The window width is 0.1

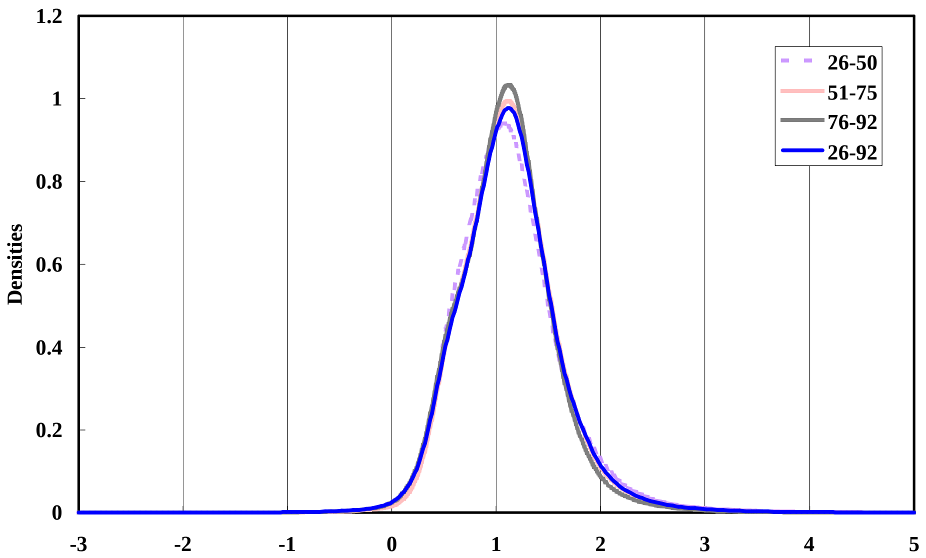

Figure 7: Standard Errors of Estimated Equity Return Betas

The window width is 0.02

Equity returns are measured in per cent per month

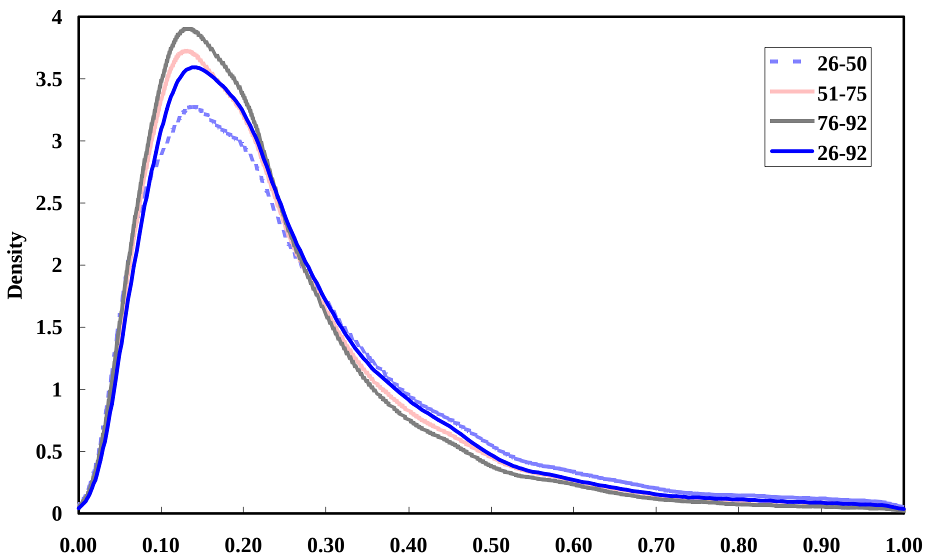

Figure 6 shows that the vast majority of betas lie between zero and two, though extreme observations are observed on both sides of this range. The concentration of beta estimates around unity is consistent with the statistics on portfolio betas discussed in Fama and French (1992). It is also consistent with the cross-sectional densities described in Vasicek (1973). These consistencies support the validity of the adopted beta estimation methodology.

Figure 7 shows the precision with which betas of individual firms have been estimated. With standard errors averaging 0.25 for the entire sample and having a standard error themselves of 0.13, it is clear that the betas are imprecisely estimated. It should be noted that this lack of precision in individual beta estimates will affect the conclusions to be drawn about beta mobility. This is because, although individual firm's betas may be estimated with substantial error, these measurement errors will be highly correlated through time for given firms. In the extreme case where the measurement error is constant through time, for a given firm, the measurement errors will have no impact on the representations of beta mobility in the next section. As the serial correlation of beta measurement errors declines, the estimated laws of motion for betas will overstate beta mobility. However, it will not lead to an overstatement of the extent to which betas have a *tendency to converge* to unity given the tight prior around unity imposed on the beta estimation procedure.

4. The Stylised Facts

The next analytical stage is intended to estimate the relationships between equity return behaviour and both previous equity return behaviour and observed firm characteristics. This section presents non-parametric representations of these relationships. Bivariate and conditional densities are estimated, relating equity betas both to their own lagged values and to observed characteristics of the firms. These densities are estimated by pooling across both firms and time. A key benefit of using estimated densities is that many of the econometric difficulties encountered in trying to model the 'representative' firm or portfolio can be avoided. Banz (1981) describes and addresses many of these difficulties within a linear regression structure.

4.1 Computing the Law of Motion for Beta

The law of motion for equity betas is characterised as a first-order Markov process with the current equity beta defining the current state. This law of motion describes the density from which the next period's equity beta is drawn, conditioning on the value of the equity beta today. The entire law of motion is then just the full set of conditional densities, one for each possible value of the current equity beta.

The conditional densities are obtained by using the standard kernel density estimation techniques described in Silverman (1986). First, the bivariate density of (b_{it}, b_{it+k}) is estimated where k is the number of months taken for a single transition from one value of beta to another. The marginal density of b_{it} can then be obtained by integrating out b_{it+k} from the joint density. Next, the densities of b_{it+k} conditional on a particular value of b_{it} can be computed by dividing the joint density at (b_{it}, b_{it+k}) by the value of the marginal density at the chosen value of b_{it}. Obtaining this conditional density of the next period's equity beta, for each possible value of the current equity beta, yields the desired law of motion.[5]

A substantial complication arises in estimating the bivariate densities for the current and future equity betas because of the possibility that firms get delisted. Before presenting the estimated laws of motion, the method for handling firms that get delisted must be understood. This complication is discussed below.

4.2 Delisting: The Absorbing State

Firms are delisted for a variety of reasons. These reasons include:

1. merger with another company;
2. share issue exchanged for share issue trading elsewhere;
3. liquidation of the firm; and
4. being dropped from the exchange for a variety of reasons, generally with a high probability of management upheaval.

---

[5] As in the univariate case, the window width needs to be selected when estimating the bivariate density. In the most general case for bivariate densities, the window width is defined by a 2×2 smoothing matrix. For most applications, only the diagonal elements of this smoothing matrix need be non-zero. Following standard practice, the smoothing matrix is diagonal for the density estimates presented below. The diagonal elements of the smoothing matrix determine the amount of smoothing across the two dimensions of the data. Again, the window widths have been chosen subjectively following the recommendations of Silverman when estimating densities for purposes of visual inspection.

Figure 8 shows the number of observed delisting in each category.

Figure 8: Number of Occurrences of Each Type of Delisting

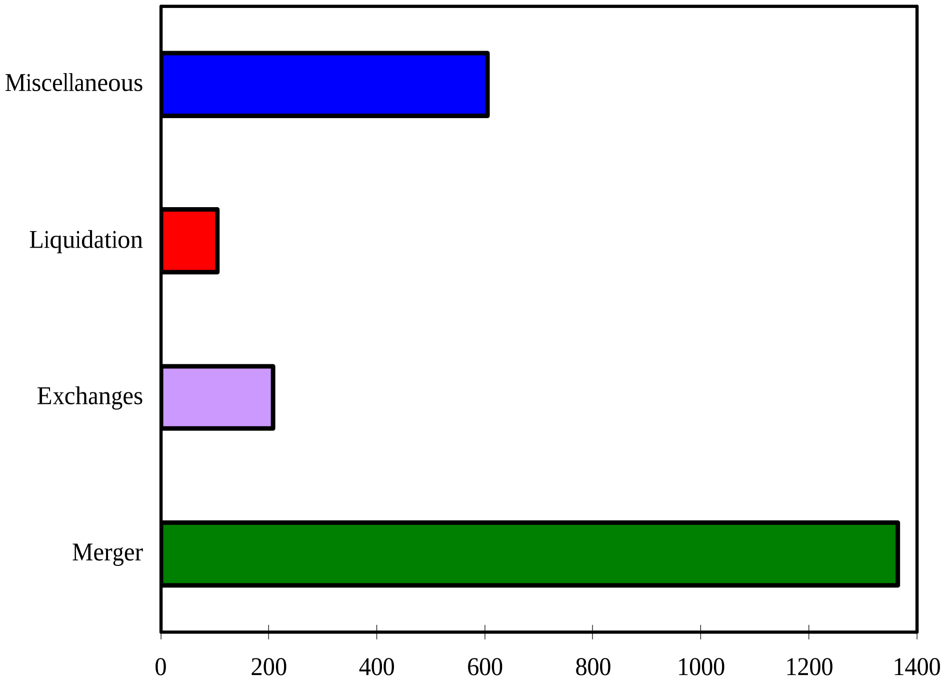

In all cases, the delisting marks the end of equity return observations from which betas can be estimated. This means that, given the current beta, the firm can either have an equity beta in the next period or be delisted. The possibility of being delisted adds an extra element to the distribution of possible outcomes for the firm in the next period. Because this extra element is discrete, special care must be taken in estimating the laws of motion for equity betas because the kernel density estimation techniques are only useful for distributions with continuous support.

The potential that firms get delisted means that the bivariate density must be constructed from two components. First, the joint density is estimated using all of the observations for which firms are not delisted. Second, the density of current equity betas is estimated for all firms that are delisted in the next period. The mass in each these two densities is then scaled by the number of observations used in their construction relative to the total number of observations available to ensure that full joint density, taking into account the possibility of being delisted, integrates to unity.

4.3 The Laws of Motion

The one-month-ahead law of motion, for firms that are not delisted, is shown in Figure 9. It is represented using a contour plot where the contours are numbered to indicate their height (three dimensional surfaces fail to represent these functions informatively because of their extreme slope in some regions). In most figures, the shape of the function will be a diagonal peak running from lower left to upper right. The one-month-ahead law of motion, for firms that are delisted in the next period has been graphed separately in Figure 10. Figures 9 and 10 together describe the full law of motion.

Figure 9: One-month-ahead Law of Motion

The window width is 0.1

Conditioning on firms that are not delisted in the next period

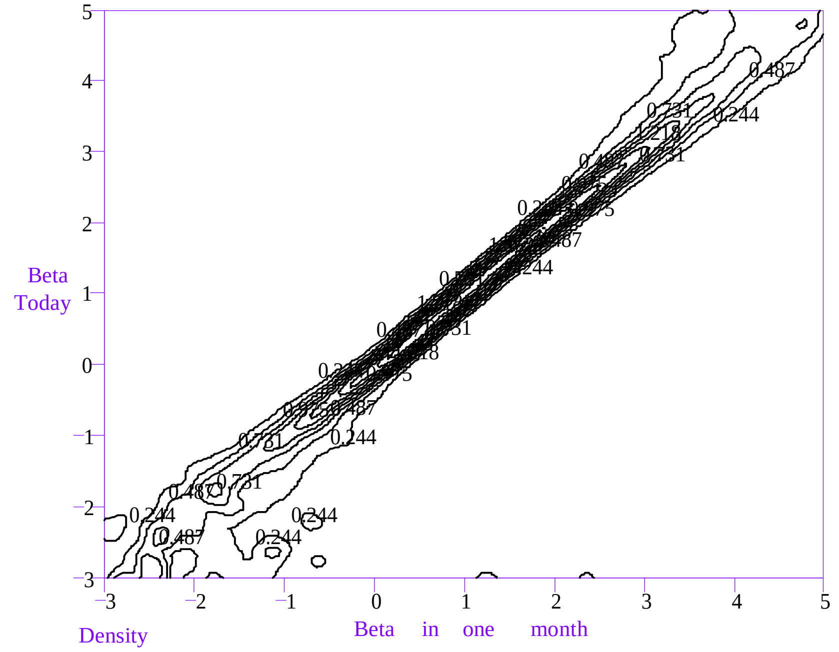

If the evolution of equity betas is not dependent upon the current beta then any horizontal cross-section taken through Figure 9 for a given current value of beta would look identical to each other. A vertical ridge running down Figure 9 with a peak on unity would be an extreme example of convergence, with the beta in one month's time being unrelated to the beta today. A ridge along the 45° line from lower left to the upper right would indicate no convergence.

Figure 10: One-month-ahead Probability of Being Delisted

The window width is 0.1

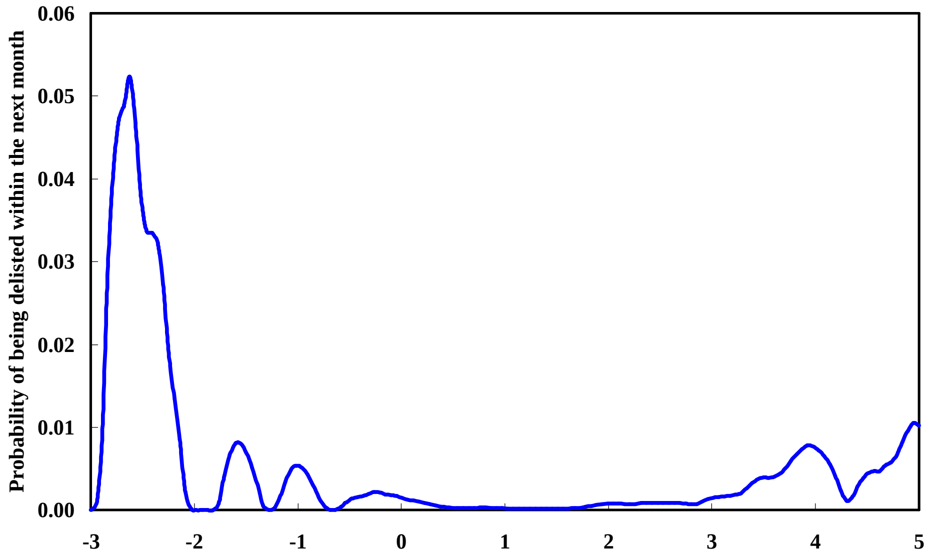

The law of motion depicted in Figure 9 is much closer to the latter characterisation with the main feature being a ridge along the 45^0 line. However, the contour lines do exhibit some skewness in the conditional densities of equity betas next month for firms with extreme equity betas today.

For example, consider a firm with a beta of 3 today. Then, reading horizontally across the contour plot from 3 on the vertical axis suggests that the probability density for that firm's beta next month is negatively skewed, toward unity. This can be inferred from the fact that the contour lines are wider apart on the left of the ridge running diagonally across Figure 9 in the vicinity of the beta equals 3 horizontal cross-section. This evidence is consistent with beta convergence toward unity. This is because extreme betas today imply a greater probability of less extreme betas next month compared to the probability that the beta next month will be more extreme.

Figure 10, showing the probability of being delisted within the next month conditioned on today's equity beta, underpins the claim that firms with more extreme betas have a higher probability of being delisted in the near future. Although the risk associated with betas below unity is greater, the U-shape of this component of the law of motion is clear even within the more typical beta range, 0–2.

Assuming that delistings are detrimental to management because of the potential loss of incumbency, Blume's (1975, p. 794) conclusion that '*part of this observed regression tendency represented real nonstationarities in the betas of individual securities*' may well have a substantive underpinning. Managers may be driving the beta of their firm toward unity in an effort to reduce the probability of delisting. The evidence in this paper goes further than Blume who was unable to identify potential driving forces behind the convergence of systematic risk characteristics. Figure 10 embodies a motive for this observed convergence behaviour.

Note also that, even though the risk of delisting is small, this is the risk of being delisted over a very short time horizon. The U-shaped relationship between betas and the risk of delisting is made more remarkable by recognising that most firms rapidly shift their beta back towards unity. By reacting to the risk of being delisted, firms increase the expected time to delisting, conditional upon the current beta. If firms did not react endogenously to their betas, the fraction of firms being delisted from the extreme beta states would be considerably greater.

The skewness of next period's conditional equity beta density for extreme current equity betas manifests itself as a pair of kinks in the ridge running along the 45° line in the laws of motion. The more sharply kinked the ridge becomes for extreme current betas, the more powerful the tendency toward convergence.

The kinks indicating convergence tendencies are much more prominent in the one and five-year-ahead laws of motion shown in Figures 11 and 13. These components of the laws of motion have been estimated directly from the data rather than by iterating forward the one-month law of motion to determine the conditional densities of future betas after undergoing the transitions implied by the one-month law of motion. This has been done because of a bias in the estimation of the laws of motion for betas currently in the vicinity of unity. The mobility of betas near unity is very low, as indicated by the fact that the ridge in the law of motion is nearly degenerate between 0.5 and 1.5. The bias arises from the fact that a fixed window width is used when estimating the bivariate density functions for current and future equity betas. There is a trade-off between using a window width wide enough to give an informative characterisation of the relatively diffuse densities for extreme betas and using a window width narrow enough to accurately represent the near degenerate density for betas near to unity. A consequence of this bias is that iteratively applying the laws of motion gives a misleadingly rapid adjustment rate.

In both of these laws of motion, there continues to be high persistence for betas near unity and very low persistence for more extreme betas. Outside the 0–2 range for beta today, there is almost no persistence evident over a 5-year time horizon. This is powerful evidence of beta convergence, despite the lack of a formal statistical test against the null hypothesis that the conditional beta densities of extreme betas are not skewed toward unity.

Figure 11: One-year-ahead Law of Motion

The window width is 0.1

Conditioning on firms that are not delisted in the next period

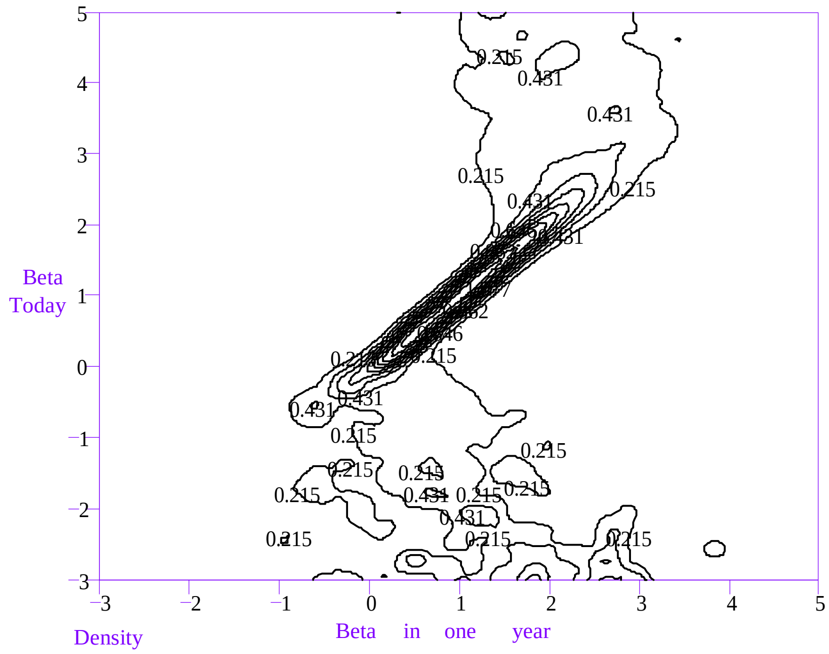

Finally, note that Figures 12 and 14 confirm the U-shaped relationship between equity betas and the probability of being delisted in the next period.

Figure 12: One-year-ahead Probability of Being Delisted

The window width is 0.1

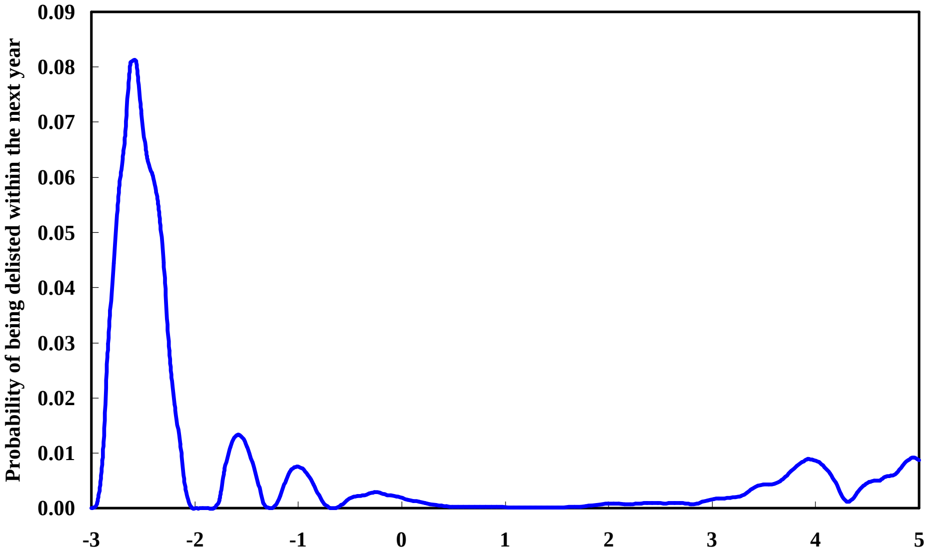

Figure 13: Five-year-ahead Law of Motion

The window width is 0.1

Conditioning on firms that are not delisted in the next period

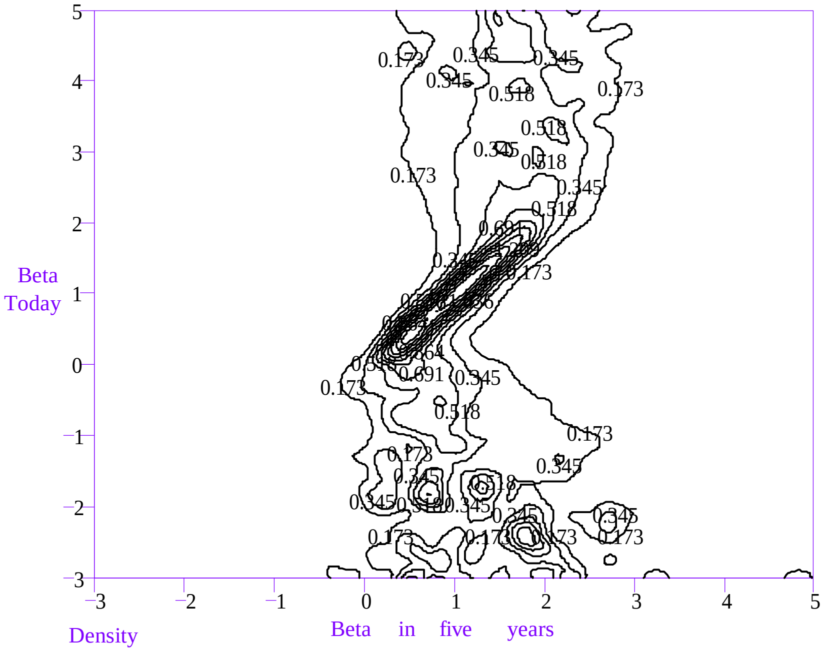

Figure 14: Five-year-ahead Probability of Being Delisted

The window width is 0.1

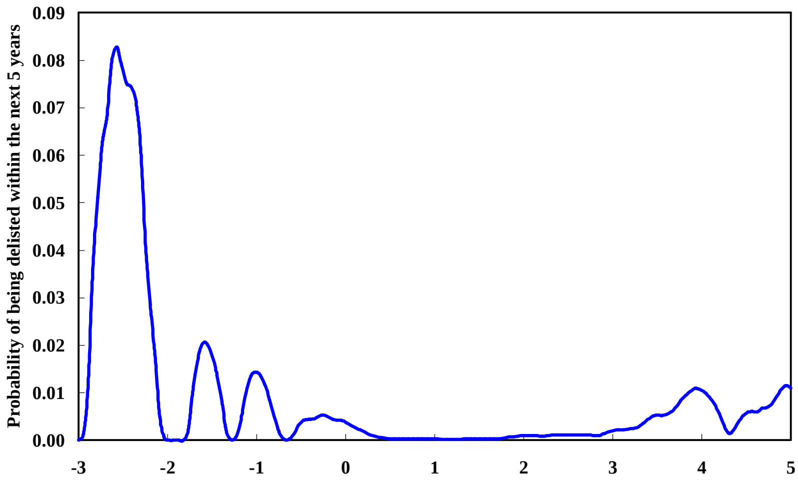

4.4 The Ergodic Beta Density

As a secondary check on the density dynamics described above, it is possible to estimate the ergodic density of equity betas across firms from the law of motion. The ergodic density can be interpreted as the long-run cross-sectional density of betas that would obtain if the estimated law of motion were the true law of motion. It is constructed by iteratively applying the one-step-ahead law of motion to obtain the infinite step ahead law of motion. Under certain regularity conditions on the estimated law of motion, this iterative process converges such that the conditional density of the infinite step ahead equity beta is independent of the current equity beta. This conditional density is the ergodic density. If this ergodic density is similar to the sample density of betas, then the ergodic density provides corollary evidence that the density dynamics, described in the previous subsection, are a reasonable characterisation of publicly listed firms' equity returns.

To generate the ergodic density, assumptions must be made regarding the betas of newly listed firms because the laws of motion include a probability that firms get delisted. According to the estimated law of motion, some mass of firms becomes delisted in each period. Without introducing new firms to offset those that are delisted, the probability mass of listed firms would dwindle to zero and the ergodic density would just indicate that all firms end up being delisted which is uninformative. To replace the mass of firms that get delisted with every iteration of the law of motion several assumptions need to be made. Specifically, it is assumed that:

- new firms exactly replace old firms so that the total number of firms is unaltered; and

- the betas of new firms are drawn from the empirical density of equity betas for newly listed firms shown below.

Note that the probability density is slightly tighter around unity for newly listed firms (Figure 15). This is an artefact of the tight prior around unity imposed on the beta state vector at time zero in the estimation process. This result is reversed if a diffuse prior is used. Imposing the artificially tight prior strengthens the empirical evidence for beta convergence, given that betas are held artificially close to unity at the beginning of the sample. Without beta convergence being a feature of the data, this mild bias imposed on the initial betas would generate a finding of beta divergence. Also note that the density of betas for firms just prior to delisting (also shown in Figure 15) is more diffuse than the equity beta density estimated using all available data. This is the feature in the data driving the U-shaped probabilities of delisting conditioned on beta, shown in the laws of motion.

Figure 15: Beta Densities

The window widths are 0.1

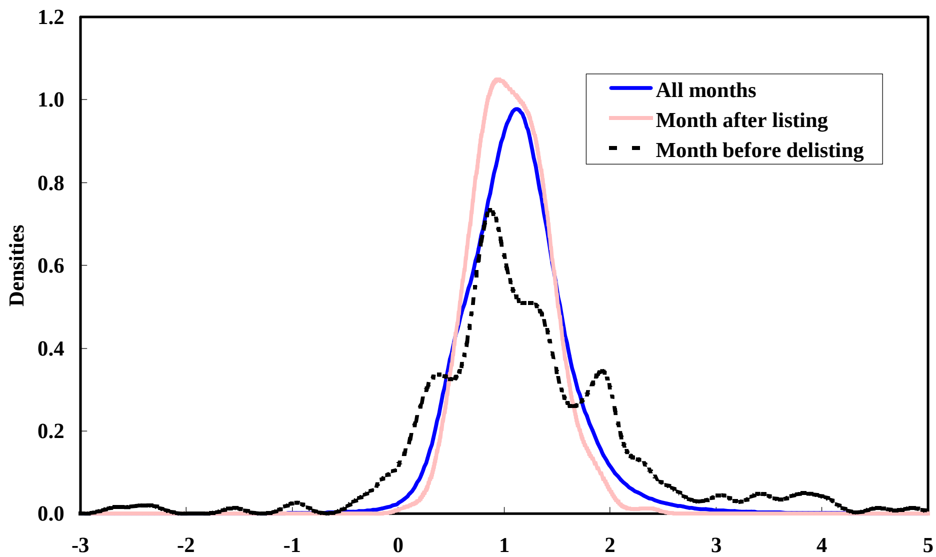

Thus, the ergodic density is estimated by iterative application of the one-month-ahead law of motion, augmented by the equity beta density of newly listed firms. This procedure yields the estimated ergodic beta density in Figure 16.

Figure 16: Ergodic and Actual Beta Densities

The window widths are 0.1

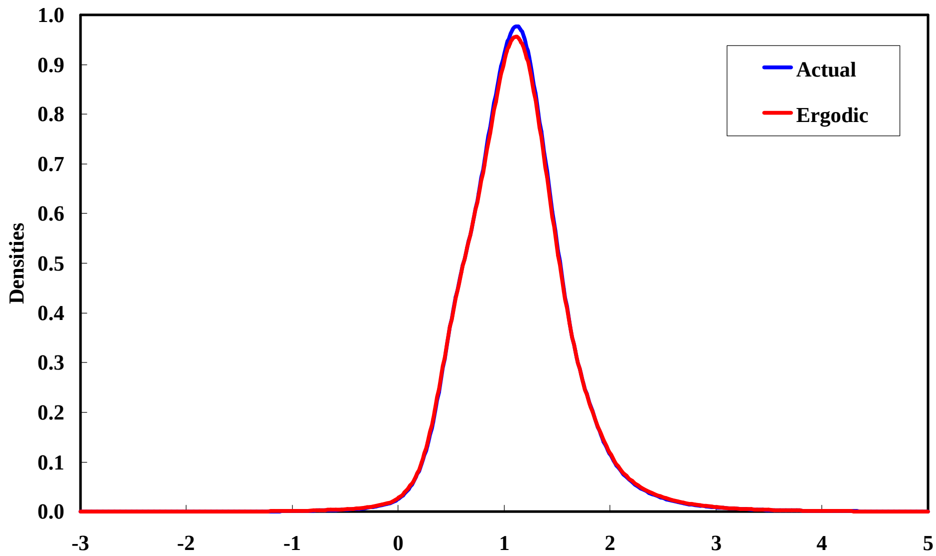

This estimated steady-state suggests that the density of betas is unimodal and highly concentrated around unity. This pattern is very similar to the prior density suggested by Vasicek (1973) and it closely matches the empirical density of equity betas obtained by Fama and French (1992) using quite different estimation techniques. These similarities represent corollary evidence that the equity beta estimation methods used in this paper are not capturing a substantially different aspect of equity return behaviour to that captured in the previous literature. It is also confirmatory to observe that the limiting density predicted by the estimated laws of motion is almost exactly the same as the observed cross-section density. It is clear that the forces for convergence do not collapse the cross-section density to a mass point at unity. Instead, the convergence tendency is offset by shocks to existing firms and by the listing of new firms.

What can be inferred from the estimated laws of motion for equity return betas? Clearly, the equity beta convergence found by Blume (1975) is strongly supported. The laws of motion also suggest a quite robust relationship between betas and the risk of being delisted in the near future. The remainder of this section explores some more obvious relationships between equity betas and characteristics of firms that further emphasise that equity beta convergence is not a statistical artefact caused by measurement error.

4.5 Firm Size and Betas

In the same way that the density of next period's beta can be computed, conditional on the current beta, densities of current betas can be estimated, conditioning on firm size. As in Banz (1981), firm size is measured as the market capitalisation of a firm in period *t* relative to the average market capitalisation of all NYSE listed firms in period *t*.

$$S_{it} = \frac{C_{it}}{\frac{1}{N_t}\sum_{j=1}^{N} C_{jt}} \tag{7}$$

where S_{it} is firm *i*'s size in period *t* and C_{it} is the firm's market capitalisation and N_t is the total number of firms in existence in period *t*. This normalisation of market capitalisation means that firm sizes are directly comparable across the time dimension, unlike the logarithmic transformation to market capitalisation applied in Chan and Chen (1988) and Fama and French (1992).

Figure 17 shows that larger firms have a tighter density of equity betas around unity. The key feature in Figure 17 is the contrast between the conditional equity beta density for large firms and the conditional equity beta density for small firms. This contrast is very clear in the comparison of the conditional densities for firms that are 25 per cent of average firm size and 175 per cent of average firm size (Figure 18).

The relationship between equity betas and firm size will be partly driven by the process of extreme beta firms self-selecting themselves out of the sample by being delisted with a higher probability. Indeed, the relationship appears to be driven largely by the adjustment of newly listed firms. Figure 19 is constructed in the same way as Figure 17 using just observations occurring five or more years after the initial listing of a corporation. Therefore, it only captures the beta/firm size relationship for relatively mature firms. While there is still some tendency for the beta density to be more concentrated for larger firms, the differences are smaller than in the case where data on immature firms is included.

Figure 17: Equity-beta Densities Conditioned on Firm Size

The window width is 0.1 for beta

The window width is 0.05 for firm size

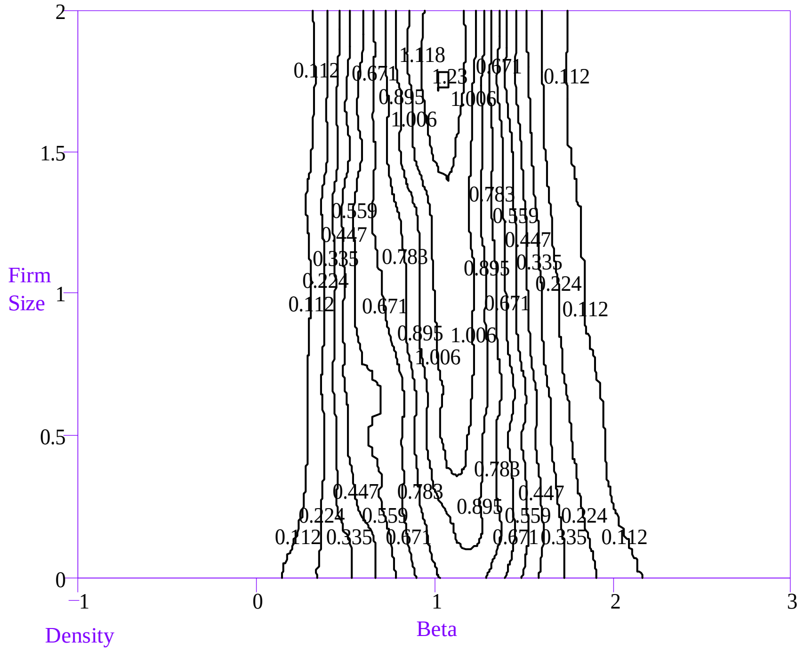

Figure 18: Equity-beta Densities Conditioned on Firm Size

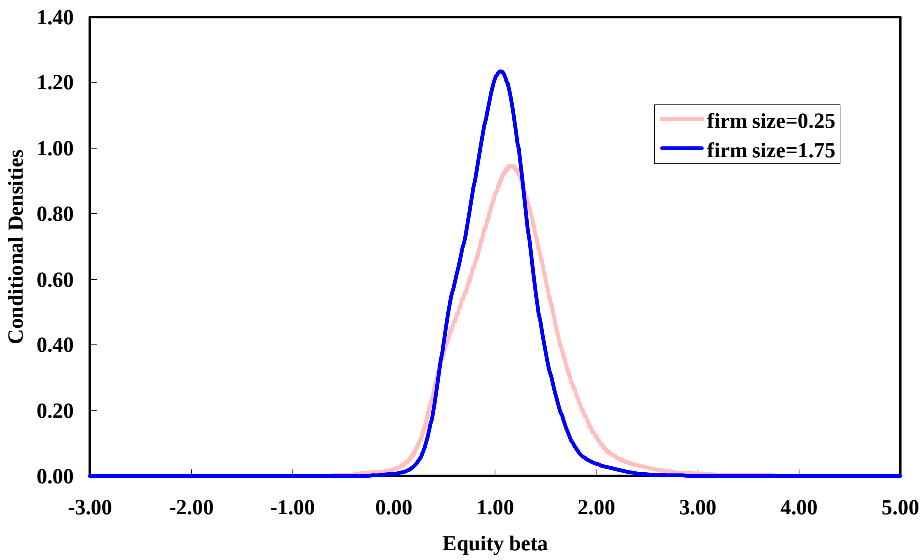

Figure 19: Equity Beta Densities Conditioned on Size (Mature Firms Sub-sample)

The window width is 0.1 for betas
The window width is 0.05 for firm size

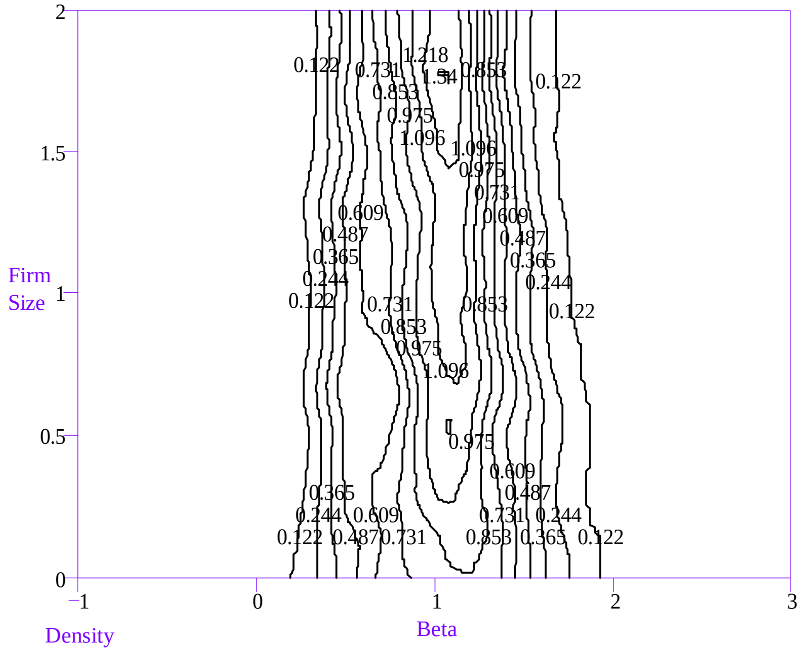

The regularity between systematic risk and firm size appears, at first glance, to be contrary to the evidence in Banz (1981), Chan and Chen (1988), Fama and French (1992), and others, all of whom find that firm size and systematic risk are strongly negatively correlated. Chan and Chen (Table I, p. 316 and Table II, p. 317) report a correlation between beta and firm size in excess of -0.9. Fama and French (Table I, Panel B, p. 435) document a similarly strong negative relationship. In Figure 17, however, the dominant regularity is that firm size is negatively related to the *dispersion* of equity betas around unity. A negative relationship seems, at first glance, quite dissimilar to a collapsing density of equity betas around unity.

How do the findings reconcile? The observed negative correlation can be explained by considering Figure 18. The cross-section equity beta density of the smaller firms has a slightly higher modal value than that for the larger firms. More importantly, however, the density for the small firms has a thick upper tail. This positive skew is substantially reduced for the larger firms. In combination, the modal shift and reduced skew explain the negative correlation. The strength of the negative correlation arises from the portfolio formation techniques used in the previous studies which disguise variation across betas. Consequently, these studies downplay the dominant feature of the data which is the reduction in the spread of the equity beta density for larger firms.

To what extent is this relationship between the concentration of the beta density and the size of firms driven by adjustment of firms as they age? The next subsection explores this issue by conditioning beta densities on the time since firms first listed on the NYSE.

4.6 Firm Age and Betas

This paper uses the number of months between the initial listing of a firm and the current return observation as a measure of age. Firms that were already listed before the first period of the sample are omitted in this section of the analysis. Having established that larger firms tend to have betas near unity, it is reassuring to observe in Figure 20 that more established firms (ones that have been listed for longer) exhibit the same patterns. This is evidenced by the higher peak (reduced dispersion) in the horizontal cross-sections for older firms. To the extent that older firms are generally also larger firms, this evidence matches that of the previous subsection.

Figure 20: Equity-beta Densities Conditioned on Firm Age

The window width is 0.1 for beta

The window width is 20 for firm age (measured in months)

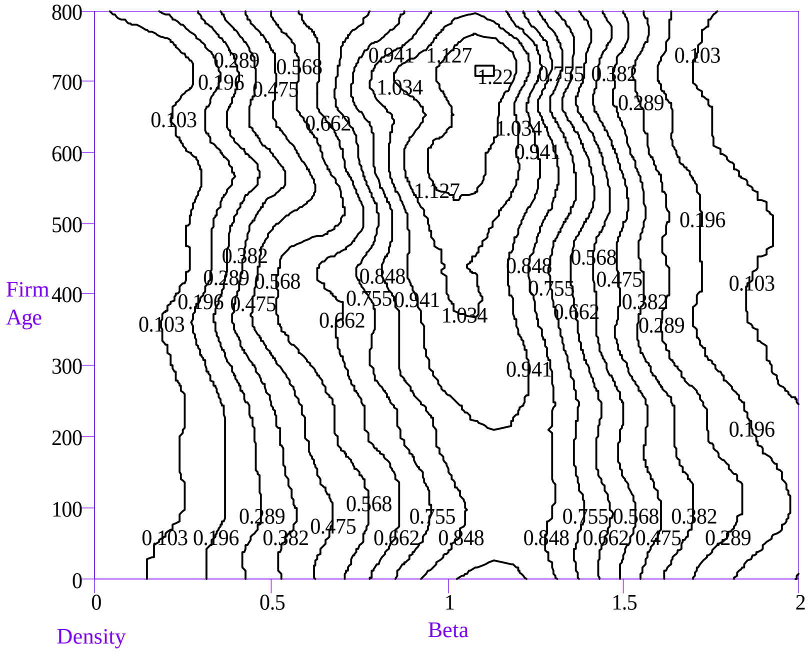

The horizontal axis, representing equity betas, has been restricted to the range 0–2 to increase the informativeness of the contours. As with the relationship between equity betas and firm size, the conditional density spreads fall as the age of firms rises. This is clearly seen in the comparison of two conditional densities in Figure 21.

Figure 21: Equity-beta Densities Conditioned on Firm Age

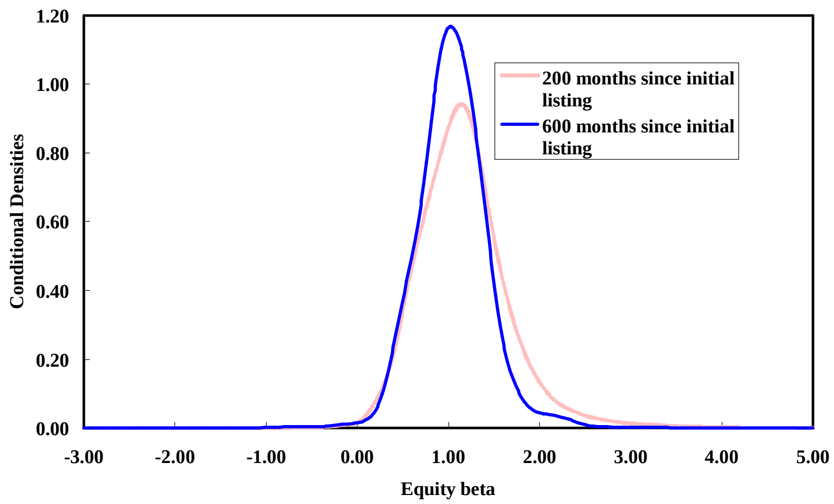

The approximation involved in the measurement of firm age since listing ignores the fact that many firms have a long and successful existence prior to listing. Without this measurement error less firms would be measured as being young when in fact they are mature. By wrongly classifying such firms, the density of equity betas for 'young' firms will be made more concentrated than it actually is. Thus, the measurement error reduces the extent to which the relationship between the dispersions of betas and the age of firms is apparent in the data.

The evidence in Figure 21 is not as convincing as that relating equity betas to firm size because a greater fraction of observations on older firms occur in the latter part of the sample. This is also the part of the sample that has the lowest measurement errors in betas, based on the standard error density functions shown in Figure 7. This, reduces the variance of cross-sectional densities of betas for older firms compared to beta densities for younger firms. However, because the differences in measurement error across time periods are quite small, it is unlikely that the fairly pronounced relationship in Figure 21 is entirely spurious.

For the youngest firms, the cross-section beta density is more concentrated than for slightly older firms. This is a reflection of the artificially tight priors imposed on betas in the estimation procedure. The fact that the data works against this artificially tight prior for young firms supports the view that the relationship between beta dispersion and firm age is not driven by measurement error issues.

5. Conclusion

The results in this paper have established the following points. There is robust and convincing evidence that equity beta convergence is a behavioural phenomenon. Evidence for this conclusion takes the form of estimated laws of motion for equity betas of individual firms. By dealing with firm level data rather than portfolio data, this study has sidestepped the complications introduced by portfolio formation while emphasising the connection between the characteristics of individual firms and their systematic risk.

The empirical analysis of the previous section strongly suggests that equity returns follow a particular pattern over the corporate life cycle of publicly listed corporations. Upon listing, a firm's beta tends to be relatively mobile and is more likely to be extreme relative to unity compared to firms that have been listed longer. However, over time, newly listed firms tend to drive their betas toward unity. Those that are not successful in forcing their beta into the range between about 0.5 and 1.5 have a substantially greater risk of being delisted than would otherwise be the case. This life-cycle view of equity betas reinforces the behavioural interpretation of equity beta convergence. The explanation of systematic risk convergence based on measurement error aggregation is strongly refuted by the direct evidence from firm level data of convergence and the robust relationship between the extent of convergence and observed characteristics of firms.

These results raise a series of questions. While many features of equity beta (and thus systematic risk) behaviour have been established, no behavioural model has been provided to explain the observations about equity returns. Why should extreme equity betas be associated with a higher risk of reduced managerial entrenchment (delisting)? Is equity beta convergence in the interests of investors? If not, why are investors unable to constrain the actions of those that do determine betas? Because of the separation between the asset-pricing literature and the corporate finance literature, these questions are difficult to frame and answer within existing theoretical frameworks.

The asset-pricing literature has generally taken the firm as an exogenous feature of the economy (with notable exceptions being Brock (1982) and Cochrane (1996) who develop asset-pricing models that explicitly incorporate profit maximising behaviour within firms). From this perspective, there has been little point asking why firms should alter their earnings characteristics in response to the equity pricing consequences of their investment profile and their financial structure.

The literature exploring the economic structure of the firm has, for the most part, developed models with asset-pricing consequences that are insufficiently rich to capture concepts like systematic risk. In most cases, these limitations arise directly from the partial equilibrium framework used to study the forces operating within firms.

Models that address systematic risk convergence need to make explicit the instruments through which firms manipulate their systematic risk. The models must also solve the optimisation problems of the agents controlling the firms to show why these instruments are used to drive the observed convergent behaviour in systematic risk. The development of such models is a fruitful direction for future research.

Data Appendix

All firm related data is drawn from the CRSP database. The market index is the value-weighted return including all distributions (labelled *VWRETD*). Individual firms' returns, *RET(t),* also include all distributions such as dividends. Prices are captured using the absolute value of the price series labelled by *PRC*. Note that zeros, when no price is observed, are replaced by the mean of the next observations on either side. Also, the absolute value of the price is used because negative values are introduced to signal whenever the closing price is the mean of the bid ask spread rather than an actual trading price. The number of shares outstanding is obtained from the series labelled by *CURSHR*. The Standard Industry Classification (SIC) uses the *SICCD* code applicable for each month.

The US three-month treasury bill rate is used as an approximation to the risk-free rate of return because it is available at a monthly frequency back to 1934. Before that date, it has had to be approximated by assuming no change in the risk-free rate until the first available observation in 1934. This assumption is fairly reasonable, in the context of this research, because interest rates on short maturity debt issues in the US between 1925 and 1934 were extremely low. For the purposes of estimating betas for excess rates of return on equity, the variations in these rates is of secondary importance compared to the variation in the raw equity returns. See Banz (1981, p. 7) for discussion of these issues.

References

Banz, R.W. (1981), ‘The Relationship between Return and Market Value of Common Stocks’, *Journal of Financial Economics*, 9(1), pp. 3–18.

Basu, S. (1977), ‘Investment Performance of Common Stocks in Relation to their Price-earnings Ratios: A Test of the Efficient Markets Hypothesis’, *Journal of Finance*, 32(3), pp. 663–682.

Black, F. (1972), ‘Capital Market Equilibrium with Restricted Borrowing’, *Journal of Business*, 45, pp. 444–454.

Black, F., M.C. Jensen and M. Scholes (1972), ‘The Capital Asset-pricing Model: Some Empirical Tests’, in M.C. Jensen (ed.), *Studies in the Theory of Capital Markets*, Preager Publishing, New York, pp. 79–124.

Blume, M.E. (1970), ‘Portfolio Theory: A Step Toward its Practical Application’, *Journal of Business*, 43(2), pp. 152–173.

Blume, M.E. (1975), ‘Betas and their Regression Tendencies’, *Journal of Finance*, 10(3), pp. 785–795.

Blume, M.E. and I. Friend (1973), ‘A New Look at the Capital Asset-pricing Model’, *Journal of Finance*, 28(1), pp. 19–33.

Blume, M.E. and F. Husic (1973), ‘Price, Beta and Exchange Listing’, *Journal of Finance*, 28(2), pp. 283–299.

Breeden, D.T. (1979), ‘An Intertemporal Asset-pricing Model with Stochastic Consumption and Investment Opportunities’, *Journal of Financial Economics*, 7(3), pp. 265–296.

Brock, W.A. (1982), ‘Asset Prices in a Production Economy’, in J. McCall (ed.), *The Economics of Uncertainty and Information*, University of Chicago Press, Chicago, pp. 1–46.

Chan, K.C. and N. Chen (1988), 'An Unconditional Asset-pricing Test and the Role of Firm Size as an Instrumental Variable for Risk', *Journal of Finance,* 43(2), pp. 309–325.

Cochrane, J.H. (1996), 'A Cross-sectional Test of an Investment-based Asset-pricing Model', *Journal of Political Economy*, 104(3), pp. 572–621.

Fama, E.F. and K.R. French (1992), 'The Cross-section of Expected Stock Returns', *Journal of Finance*, 47(2), pp. 427–465.

Fama, E.F. and J.D. MacBeth (1973), 'Risk, Return and Equilibrium: Empirical Tests', *Journal of Political Economy*, 81(3), pp. 607–636.

Friend, I. and M.E. Blume (1970), 'Measurement of Portfolio Performance Under Uncertainty', *American Economic Review*, 60(4), pp. 561–575.

Gibbons, M.R. (1982), 'Multivariate Tests of Financial Models: A New Approach', *Journal of Financial Economics*, 10(1), pp. 3–27.

Harvey, A.C. (1989), *Forecasting Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge.

Jaswinski, A.H. (1970), *Stochastic Processes and Filtering Theory*, Academic Press, New York.

Klemosky, R.C. and J.D. Martin (1975), 'The Adjustment of Beta Forecasts', *Journal of Finance*, 10(4), pp. 1123–1128.

Kolmogorov, A.N. (1933), 'Sulla Determinazione Empirica Di Una Legge Di Distribuzione', *Giornale Dell'Instituto Ital. Degli Attuari*, 4, pp. 83–91.

Lintner, J. (1969), 'The Aggregation of Investor's Diverse Judgements and Preferences in Purely Competitive Markets', *Journal of Financial and Quantitative Analysis*, 4(4), pp. 346–382.

Litzenberger, R.H. and K. Ramaswamy (1979), ‘The Effects of Personal Taxes and Dividends on Capital Asset Prices: Theory and Empirical Evidence’, *Journal of Financial Economics*, 7(2), pp. 163–195.

Markowitz, H. (1959), *Portfolio Selection*, Yale University Press, New Haven.

Miller, M. and M. Scholes (1972), ‘Rates of Return in Relation to Risk: A Re-examination of Some Recent Findings’, in M.C. Jensen (ed.), *Studies in the Theory of Capital Markets*, Preager Publishing, New York, pp. 47–78.

Mossin, J. (1969), ‘Security Pricing and Investment Criteria in Competitive Markets’, *American Economic Review*, 59(5), pp. 749–756.

Quah, D. (1996), ‘Empirics for Economic Growth and Convergence’, *European Economic Review*, 40(6), pp. 1353–1375.

Reinganum, M.R. (1981), ‘The Arbitrage Pricing Theory: Some Empirical Results’, *Journal of Finance*, 36(2), pp. 313–321.

Roll, R. (1977), ‘A Critique of the Asset-pricing Theory’s Tests: Part I, On Past and Potential Testability of the Theory’, *Journal of Financial Economics*, 4(2), pp. 129–176.

Shanken, J. (1985), ‘Multivariate Tests of the Zero–Beta CAPM’, *Journal of Financial Economics*, 14(3), pp. 327–348.

Sharpe, W.F. (1964), ‘Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance*, 19(3), pp. 425–442.

Silverman, B.W. (1986), *Density Estimation for Statistics and Data Analysis*, Monographs on Statistics and Applied Probability 26, Chapman and Hall, London.

Smirnov, N. (1939), ‘On the Estimation of the Discrepancy between Empirical Curves of Distribution for Two Independent Samples’, *Bulletin Mathèmatique de l’Université de Moscou*, 2, fasc. 2.

Stambaugh, R.F. (1982), ‘On the Exclusion of Assets from Tests of the Two-Parameter Model: A Sensitivity Analysis’, *Journal of Financial Economics*, 10(3), pp. 237–268.

Treynor, J. (1961), ‘Toward a Theory of the Market Value of Risky Assets’, unpublished manuscript.

Vasicek, O. (1973), ‘A Note on using Cross-sectional Information in Bayesian Estimation of Security Betas’, *Journal of Finance*, 8(5), pp. 1233–1239.