Smithers & Co. Ltd.

20 St. Dunstan's Hill, London EC3R 8HL

Telephone: *020 7283 3344* Facsimile: *020 7283-3345*

Web Site: www.smithers.co.uk E-mail: info@smithers.co.uk

Stock Market Volatility.

Report No. 184 **11th November, 2002.**

Derry Pickford and Andrew Smithers.

Stock Market Volatility.

Summary.

1. Value is only a weak indicator of short-term stock market direction. We are therefore looking for other indicators, which might prove more powerful than value used in isolation.

2. It is widely recognised that volatility is persistent. We show here that high volatility is also associated with poor returns. Sadly, the help provided for market timing, however, is weak.

3. The returns from the US stock market do not fit a "normal" distribution. Very high and very low returns are more frequent, as are returns close to the average. The actual distribution can be better described by assuming that the outcome is drawn from one of two normal distributions, with a random choice between the two.

4. One of these distributions, which occurs only 9% of the time, combines negative returns with high volatility. According to the Capital Asset Pricing Model, assets with high expected volatility must give high expected returns. The observed result might appear to require some explanation if the high volatility distribution is predictable.

5. The poor return/high volatility phases are not easily predictable, although the poor phases seem to have an above average chance of being followed by another poor period. There also appears to be a statistical link with changes in *q* and past returns.

6. If the relationships were straightforward and reliable this would be useful to investors, as it would enable us to predict when returns are likely to be both poor and volatile. Sadly the statistical relationships have an odd form and are not easy to explain. There is thus a risk that the apparent link may be spurious.

7. The risk of bankruptcy to leveraged investors is greater than it would be if returns were normally distributed. This makes leveraged investment more difficult. Since leveraged investment is necessary for equity risk to be efficiently distributed within an economy, barriers to leveraged investment will raise the equilibrium equity risk premium. This increases the probability that the historically observed equity risk premium is unlikely to change.

8. We also touch on the possibility that the distribution of returns provides a career incentive for fund managers to be optimists. This in turn provides a possible insight into the current problems of the European insurance industry.

1. Introduction.

The return the stock market gives from one month to the next is usually assumed to be lognormally distributed. The lognormal (rather than the normal) distribution is used because of the effect of compounding on returns[1]. It also implies that impossible returns such as –120% are given a zero probability. It has therefore been an extremely useful tool in finance theory and is the basis of the Black-Scholes pricing model. It has long been recognised that the distribution of actual log returns has a "fat tail" or excess kurtosis. [2] It has also been recognised that the returns on the index (if not individual stocks) have negative skew (log returns are not symmetrically distributed - a very adverse move is more likely than a very favourable move).

In this report we show that it is better to assume that the stock market is a combination of two lognormal distributions, rather than one. The overall picture then exhibits the observed characteristics of excess kurtosis and negative skew. Distribution 1 has a low variance (low volatility) and a high mean return. Distribution 2 has a high variance (high volatility) but a negative mean return.

It would of course be extremely helpful to investors if we were able to predict whether one of the two distributions was more or less likely than usual. (If this were possible, investors would be able to achieve well above average returns at lower volatility than is possible for "buy-and-hold" investors.) We find that the more likely that the current return was from the distribution 2 (the high volatility and negative average return distribution), the more likely that the next period's return will also be drawn from distribution 2. Unfortunately this element of persistence is weak and the predictive power of this effect is thus very small.

Our general conclusion is thus a negative one, which we naturally regret, that volatility analysis is unlikely to provide investors with any real help in making asset allocation decisions with regard to equities. Such negative information is useful, but purely pragmatic investors will, we fear, find little of interest in the remainder of this report. We hope, however, that those with a technical interest in the working of the stock market will find our detailed analysis of interest.

---

[1] If periodic returns are lognormally distributed, then the continuous return is normally distributed. Suppose the return on an asset over a year is 10%. If the interest is paid twice yearly, this would correspond to 4.88% per 6 months or an "annual rate" of 9.76%. If the interest is paid continuously, then only a rate of 9.53% is required. The continuous return is therefore 9.53%. However, if the return on an asset is –10% over a year, then the continuous rate of return is –10.54%. We hereafter refer to periodic returns being lognormally distributed or the log of 1 + the periodic return or the continuous return being normally distributed.

[2] For example see Mandelbrot (1963) and Fama (1965).

As usual, the main text of this report gives a general account of our approach and the information that it reveals. A more technical description of what we have done is set out in the Appendices.

2. The US Stock Market's Volatility.

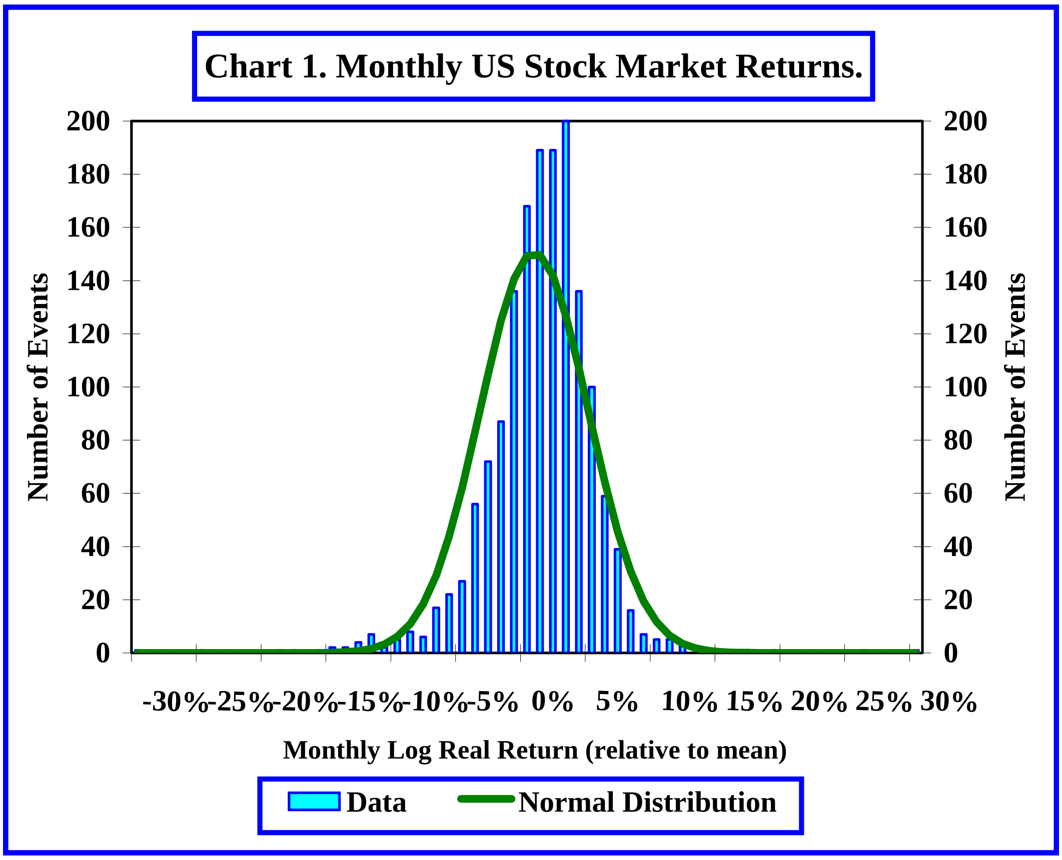

Chart 1 plots the monthly (log) real returns achieved on the US stock market, compared with the average return, relative to the normal distribution. On the basis that the past is a guide to the future, and it is certainly the only guide that we can have, this shows the probability of the market giving any particular return, above or below its average. We then compare the actual results with those that would occur if these probabilities were normally distributed.

Chart 2 focuses on the tails of the distributions in Chart 1. Very high and very low returns have been more common than would "normally" be expected. The same is true for returns that are close to the average. To make up for these differences, returns in the mid-range are less common. Because the two ends of the chart are more prominent than they would "normally" be, the actual distribution is referred to as being "fat tailed" or having "excess kurtosis".

Chart 2. Monthly US Stock Market Returns.

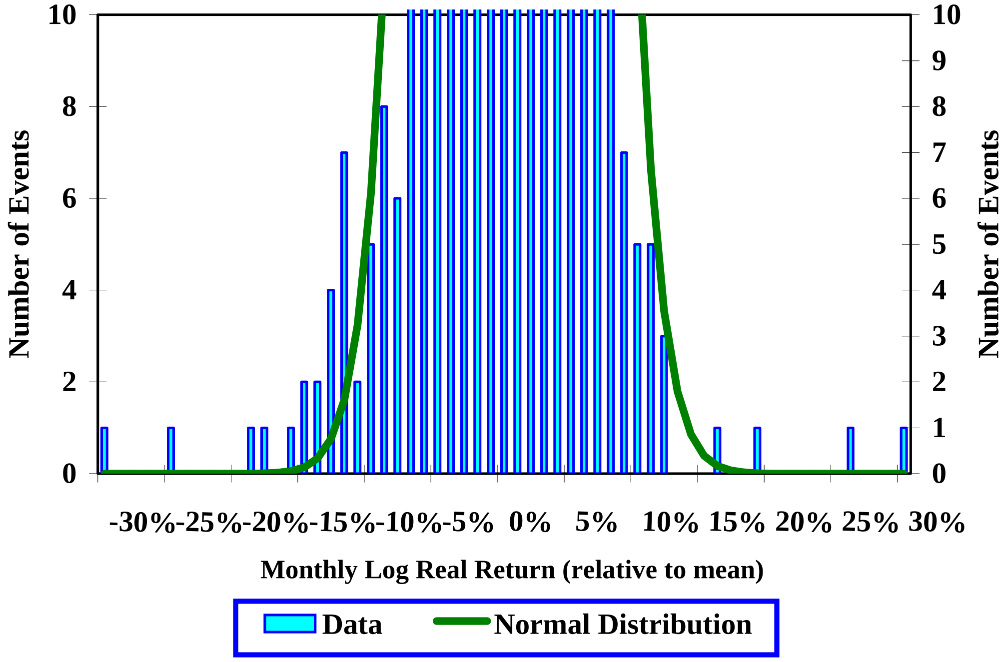

Another interesting feature of the actual distribution is that the modal value, the one most likely of all to occur, is higher than the mean value, which is the return that investors can expect over the very long-term. This is because large negative returns are more likely than large positive ones. This feature is known as “negative skew”.

We wish to find out if there are any significant market implications that can be drawn from these differences. Our first step in this enquiry is to try and find a better way of describing the observed behaviour of the market. We find that a “mixed” distribution, in which the market switches between two different “normal” ones, describes the market’s actual behaviour much better than a single normal one.

We could, in principle, “improve” the description governing the distribution of returns by allowing the return to be drawn from more than two normal distributions. If we were to use a very large number, say 1,600 distributions, we would be able to describe the actual historical distribution exactly, since there could be a distribution (with zero variance) to describe each observation! The fit using just 2 distributions we believe to be sufficiently good that, on the principle of Occam’s Razor[3], we resist the temptation to add additional distributions.

---

[3] Occam’s Razor is the principle that if two theories explain a set of observations equally well, then the simpler theory is to be preferred.

3. A "Mixed" Distribution.

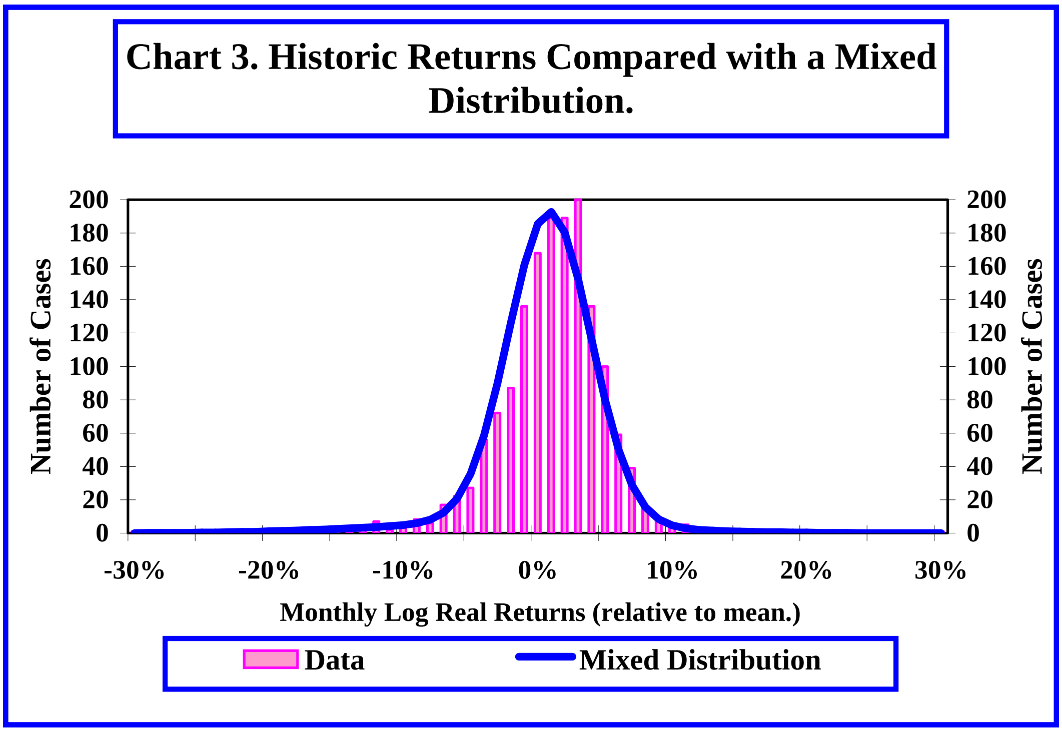

In Chart 3 we show how a mixed distribution of two different normal distributions compares with the actual distribution of log monthly returns.

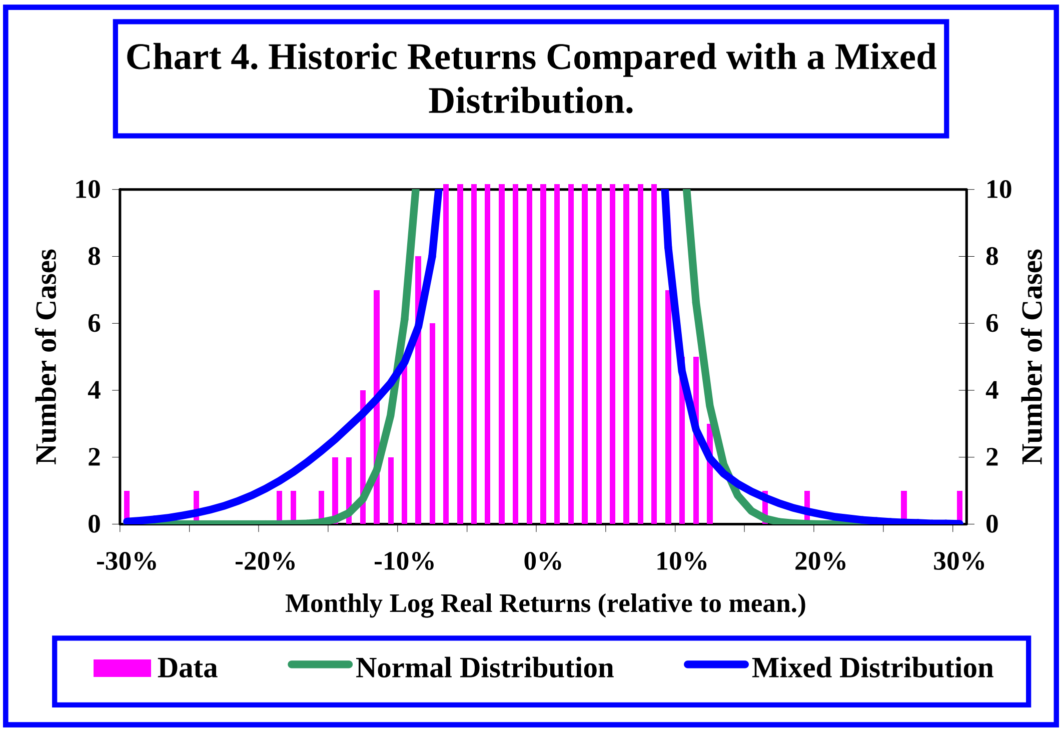

Chart 4 focuses in on the tails. This shows how the mixed distribution explains extreme values better than the normal distribution.

4. "Phase" Characteristics.

As Charts 3 and 4 make clear, a model in which the stock market has two phases, with different average returns and standard deviations, is a better description of the real world than a model that assumes that returns are normally distributed.

Table 1. Statistical Characteristics of the Distributions in Each Phase.			
	Probability %	**Mean Monthly Return %**	**Standard Deviation %**
Distribution I	90.99	0.88	3.06
Distribution II	9.01	-2.73	9.22
Whole Sample	100	0.55	4.15

Each of the two distributions is thus normal, but they have markedly different characteristics, which are set out in Table 1. Ninety percent of the time the stock market is well behaved, giving a return of nearly 0.9% per month, i.e. around 11% p.a. Nine percent of the time, however, it gives a negative return of 2.73% per month, which is equivalent to a negative annual rate of 38% p.a. Putting the two together produces a monthly return of 0.55%, which is an annual rate of 6.8%.[4] The difference is, to put it mildly, huge. If investors could predict the market's next phase and thus avoid the bad periods, their annual return would rise from 6.8% in real terms to 11%.

5. "Phase" Identification.

Our ability to improve our description of the stock market's behaviour by this mixed distribution model does not, however, enable us to predict the phase in which the stock market is currently operating or about to enter.

We cannot be certain, at any time, in which phase the market is currently operating. Once we have observed the return, we can estimate the probability of having been in distribution 2 in that period. We refer to this as the *ex-post* probability. This contrasts with the unconditional probabilities (shown in Table 1) which are the probabilities that a return will be drawn from distribution 1 or 2, when no other information is known. In Appendix 1 we discuss the estimation of *ex-ante* probabilities. The *ex-ante* probability of phase 2

[4] This is, naturally enough, the same long-term real return on equity as that which comes from most data, e.g. it is similar to the 200 year result of Jeremy Siegel. In Report No. 180 "Monthly Proxy for *q* using S&P500 Data." we showed that the stock market seems to have been rather undervalued at the beginning of the period for which we have monthly data, 1871, and a bit overvalued at the end, Q2 2002. The long-term real return on equities is thus probably below 6.8% p.a.

occurring is the probability of the next period's return coming from distribution 2, given all the information that we know today.

The chances of extreme values being generated by the first distribution are low.[5] Very high, as well as very low returns are much more likely to come from the second distribution, even though it produces poor returns on average.

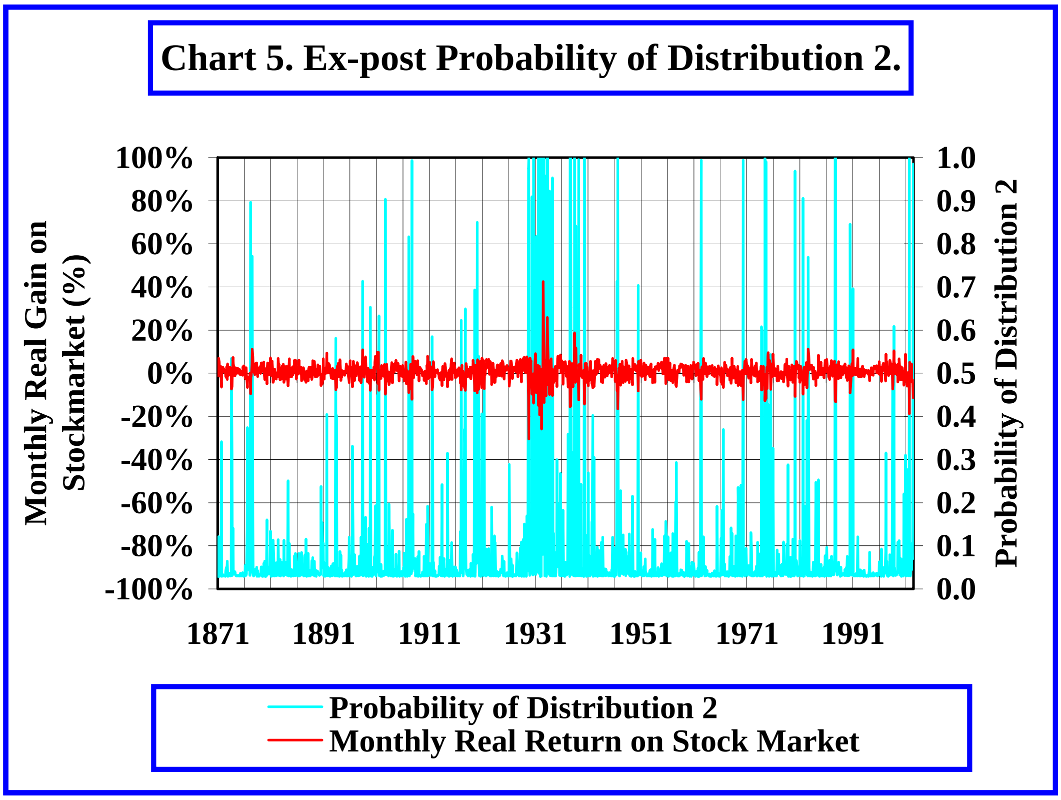

In Chart 5 we show the chances of the stock market being in the relatively rare phase that produces poor returns. It shows that such periods tend to be bunched together. They were particularly common from 1925-1950.

6. A Failure of the Capital Asset Pricing Model?

A basic principle of the capital asset pricing model (CAPM) is that investors need to be compensated for risk by higher returns. The ratio between a market's expected excess return (the expected return minus the risk free rate) relative to its riskiness is known as its Sharpe ratio. It is only relatively recently that we have had nearly risk-free assets (US TIPS, UK Index-linked bonds etc). If we are to assume that the risk-free rate of return is non-negative, then if we were able to know when the market was going to be in the negative mean/high volatility phase, this would imply that the Sharpe ratio of equities would on

[5] This is because its standard deviation is relatively low.

occasions be negative. This would have grave implications for market efficiency.[6] We therefore examine the predictability of the distribution.

Our main finding is that, if recent returns have been more likely than average to have been drawn from the high variance negative return distribution, then this period's return is also more likely to have been drawn from this distribution. In other words, the probability of distribution 2 occurring exhibits "persistence" or "positive serial correlation". This positive correlation remains if other shorter time periods are used, e.g. the post-war period.

The persistence does not, unfortunately, offer a great deal of certainty about when the return is going to be drawn from the second distribution. The "R-Squared" of the regression is statistically significant, but low.

Our second angle is to engage in an exercise in data-mining to derive a forecast of the probability of the second distribution occurring. We look at a number of other measures, including functions of our *q* measure (such as *q* vs. average logged and squared to indicate period of extreme over and under valuation), lagged returns, and various business cycle variables.

Table A4 in the Appendix lists those variables that appear to have some explanatory power. Since we are starting off by examining the data rather than formulating a coherent model of how returns should be generated, our results are likely to be spurious, especially as there is also no convincing story to explain the estimates of the coefficients that we derive. Since this again provides only limited help in forecasting when returns will be from distribution 2, we believe that these phases are essentially unpredictable.

In conclusion, our estimates of *ex-ante* probabilities when using just lagged *ex-post* probabilities are not sufficiently different from the unconditional probability to be of much use. When using more information the validity of our forecasts is questionable. Again our full results are reported in Appendix 1.

7. Some Other Consequences of the Non-Normal Distribution.

The distribution of returns may provide an insight into the current problems of the European insurance industry. One feature of the actual distribution of returns, compared with a normal one, is that 90% of the time the stock market gives a better return with less volatility than it does on average. It

[6] The high volatility low return phase may be associated with particularly buoyant periods in the real economy. This would allow investors to use the stock market as a kind of insurance policy against job losses or other economic ills, and hence be consistent with market efficiency. However we do not believe that this explanation is credible.

seems probable that this influences the behaviour of fund managers and their selection through survival.

If fund managers are rewarded for success, but liable to lose their jobs if they perform below average, they will have a 90% chance of being right if they are perennial optimists and only a 10% chance of being wrong. Furthermore, in the conditions in which they will be at risk of losing their jobs, such risks will probably be high, even for the relatively successful. Finally, while they may lose their jobs in bear markets, they will not have to repay past bonuses.

The net result is that optimism is the paying proposition. Furthermore, the longer a bull market lasts, the greater are the chances that the naturally cautious will be weeded out in favour of the successful.

If this bias towards optimism is natural among fund managers, it will need to be offset by added awareness of the risks of extreme returns at the senior management level. Because the non-normal distribution of returns adds to the risks of leverage, it can create special difficulties for UK life companies with a high proportion of "with-profits" policies.

The liabilities of these policies are not fixed, since they vary with the bonuses that the companies declare. But equally they are not "unit linked", so that the liabilities do not automatically adjust to falls in the value of the companies' assets. The net result, it seems, is as if the companies were protected against average falls in the market, but not against extreme ones. It is to be feared that UK life insurers have failed to allow adequately for the non-normal distribution of equity market returns when formulating their investment policies.

Non-normality also sheds light on the "Equity Risk Premium" puzzle. If very bad returns are more frequent than a casual examination of stock market variances would suggest, then this raises the required expected return. More importantly, for the past premium to be a puzzle, stock market risk must be able to be efficiently distributed within the economy. For this to be possible, young individuals, who will have little wealth but longer investment horizons, must be able to bear a higher proportion of this risk than individuals who are just about to retire.

Our results suggest that leveraged individuals and investment companies would go bankrupt more often than on the basis of a standard distribution. Given the high costs of bankruptcy, this will make leveraged investment more difficult, and hence will act to raise the equity risk premium.[7]

---

[7] Economists have tended to refer to this phenomenon as "Junior can't borrow".

In Appendix 2 we argue that a leveraged portfolio, with limited liability, is not very different from holding a call option. The option market provides additional evidence that constructing high beta/high expected return portfolios is difficult. This further supports the argument that equity risk is not likely to be "efficiently" distributed within the economy. This justifies a higher equity risk premium than would otherwise seem likely.

8. Conclusions.

- Equity returns are not "lognormally" distributed**.** A better description is provided by two lognormal distributions. In happy times, which rule 90% of the time, returns are well above average and volatility is low.

- It is probable that investors, including many professionals, are deceived by this into believing that stock market investment is safer than it actually is.

- Volatility is associated with poor returns and has some tendency to persist. Investors should be even more than usually cautious of investing in over-valued markets when they are volatile.

- Unfortunately the persistence is not very strong and the help given to market timing is thus weak.

- The distribution of returns creates additional constraints on leveraged investment. The inability of Junior to borrow to invest in the stock market is important in models that seek to explain the historic level of the equity risk premium.

Smithers & Co. Ltd.

Appendix 1.

Deriving Our “Mixed Distribution”.

The usual approach, when assessing a security or a market, is to measure its mean return and also to assess the variance or the standard deviation (the square root of the variance) of this return about this mean. Thus, using Professor Robert Shiller's data for the US stock market over the period from January 1871 to January 2000, we find that the mean monthly real return was 0.60% per month, while its standard deviation was 4.11% per month.

These measures are not, however, in themselves sufficient to indicate to investors the risks that they face. In order to give some indication to investors of the chances of particularly good or particularly poor outcomes, it is necessary to give a better indication of the distribution of returns. A given standard deviation can be delivered either by a distribution which is 'peaked', in which extreme values are rare, or by one which is flat, in which extreme values are common. A failure to take due account of the chances of extreme values is likely to result in difficulties, particularly for leveraged investors.

The basic tool of analysis is the probability distribution. This shows the chance of a particular outcome being observed relative to the distance of the specific outcome from the long-term average. One possibility might be that all outcomes, between a minimum and a maximum value were equally likely, in which case the distribution would be rectangular. But it is more usual to find that outcomes close to the average are more common than those a long way from the average. In such a case the distribution is bell-shaped.

The most commonly used distribution of this type is the normal distribution. Despite its complex mathematical structure, it has a number of desirable properties and represents a wide range of everyday processes reasonably well. It also forms a reference point for much of the work on probability theory. For example, basic option pricing theory (Black and Scholes, 1972) is set out on the assumption that returns to securities are normally distributed.

The basic assumption we make is that, instead of there being a single normal distribution summarising the movements of the stock market, the process is instead a mixture of two distributions with distinct means and standard deviations. At each point in time it is not known which probability distribution is going to be used to determine the outcome; one knows simply the probability of each being selected. Even after the event, one cannot say with any certainty which distribution has been used to determine the outcome. One can simply give an *ex-post* probability to the outcome being determined by one or the other distribution.

In principle, there is no limit to the number of component distributions which might be assumed to make up the mixed distributions. A degree of statistical testing is possible, but the properties of such tests are not well-known. In any case, the purpose of developing this analysis is to provide a useful summary of the properties of market returns, and there is no point in taking the modelling exercise to the point at which the results are no longer easily comprehensible. We therefore limit the bulk of our discussion to the possibility that there are two underlying normal distributions that together produce the dispersion of market returns observed in Shiller's data.

With two underlying distributions there are five parameters of interest. These are the means and standard deviations of the two normal distributions concerned, and the unconditional probability that the outcome is chosen from the first distribution. The chance of it being selected from the second is one minus this probability.

There are a number of possible ways of determining these parameters, and the most satisfactory is generally regarded to be the method of maximum-likelihood. The most likely values of the parameters are chosen on the assumption that the basic model is the correct one. The resulting values have the property that the sum of the means of the two distributions, weighted by the probabilities of each mean being selected, is the mean of the overall distribution. The standard deviations are similarly coherent with the standard deviation of the overall sample.

Applying this to Shiller's data, we find that, taking the period Jan 1871-June 2002 as a whole, the two distributions have the characteristics shown in Table A1, which is also shown as Table 1 in the text.

Table A1. Statistical Characteristics of the Distributions.

	Probability	**Mean**	**Standard Deviation**
Distribution I	0.9099	0.0088	0.0306
Distribution II	0.0901	-0.0273	0.0922
Whole Sample	1	0.0055	0.0415

Thus our analysis suggests that for just over 90% of the time the stock market appears well-behaved with a real growth rate of 0.9% per month and a standard deviation of this growth rate of just over 3 percentage points. However, on nine occasions out of one hundred, the situation is much gloomier, with an average decline of 2.75% in the month and a much higher level of uncertainty - a standard deviation of over nine percentage points. Thus, far from investors being rewarded by higher returns in times of greater uncertainty, they are penalised by much poorer returns.

While we cannot be certain whether particular realisations are delivered by Distribution 1 or Distribution 2, we can estimate the probability that each belongs to the first or the second distribution. Unconditional probabilities are given by the numbers in Table 1 and are, of course, uniform. But *ex-post* we can calculate the probability that particular observations are generated by one or the other.

The chances of extreme values being generated by the first distribution are low, because its standard deviation is relatively low. Thus both high and low extreme values are much more likely to come from the second distribution, with its lower mean making relatively little difference to this.

In Chart 5 we showed the probabilities that each observation is taken from the second rather than the first distribution. Perhaps the most interesting point is that the phase 2 periods appear to be bunched together rather than evenly dispersed. They appear to be much more frequent in the period 1925-1950 than either before or afterwards. This is confirmed by the period averages shown in Table A2.

Table A2. Probability Distributions of Monthly Stock Market Returns.

1901-1950	**Probability**	**Mean**	**Standard Deviation**
Distribution I	0.8864	0.0093	0.0347
Distribution II	0.1136	-0.0340	0.1120
1951-2002			
Distribution I	0.8505	0.0098	0.0270
Distribution II	0.1415	-0.0179	0.0608

For the two halves of the twentieth century the probabilities of the two distributions are similar. But in the first half the bad periods were much worse, with lower means and larger standard deviations. The good periods were very similar in both halves.

Thus, the basic point survives. A mixed distribution represents the data better than a single normal distribution. But the fact that the means of the less frequent distributions differ markedly in the two sub-periods might suggest that we should try more and more complicated mixtures. However, fitting them starts to become complicated; the distributions are meant to be approximations summarising the process, and a the mixture of two gives a much better indication of the nature of the risks investors face than does a single one. The gain from complicating things further is likely to be more limited.

We therefore investigate whether the probability that the market return is taken from distribution 2 rather than distribution 1 is predictable. Firstly, we explore whether it is a function of past probabilities. [8]

Table A3. Forecasting the probability that the return will be from the second distribution.			
1573 monthly observations used for estimation from June 1871 to June 2002.			
Dependent variable: **PROB (probability that the return will be from the second distribution.)**			
Regressors	**Coefficient**	**Standard Error**	**T-Ratio [Prob]**
CONSTANT	0.05	.0056	8.72 [.000]
PROB lagged 1 month	0.19	.0254	7.33 [.000]
PROB lagged 2 months	0.08	.0257	3.04 [.002]
PROB lagged 3 months	0.01	.0258	0.42 [.673]
PROB lagged 4 months	0.10	.0257	3.96 [.000]
PROB lagged 5 months	0.08	.0254	3.18 [.002]
R-Squared 7.9658%		**R-Bar-Squared**	7.6722%

Secondly, we engage in an exercise in data-mining and explore whether we can predict the distribution with a range of different variables. The regression coefficients over both the full sample period and the post-war period are reported in table A4 (overleaf). There appears to be no consistency or logic to the estimation of the coefficients and we feel that they're unlikely to have any future predictive power. Although the R-Squared for this regression might at first suggest that the results are statistically significant, the fact that our explanatory variables have been chosen to fit the data means that the R-squared are almost certainly insignificant.

These exercises should first of all caution investors against making the assumption that stock market behaviour follows a normal distribution. The reality is more complicated. There appear to be periods of volatility in which investors can expect a poor average return. These appear to come between one month in ten and one month in seven, depending on whether one studies the stock market since 1871, or only the period since 1951, with some evidence

[8] Probabilities have to lie between 0 and 1. Regression equations, however, usually make it possible for variables to take any values. Thus the use of a regression gives rise, at least notionally, to the possibility that it may result in a predicted probability below zero or above one. The usual solution, for any probability, P, is to consider the natural logarithm of P/(1-P). This can indeed take any real value; no problem arises from its use in a regression equation. We have tried both approaches and find little difference between the two.

that periods of volatility are clustered together. The exercise does, however, invite speculation beyond the capacity of statistical analysis. The second quarter of the 20th century was one of much greater volatility than the rest of the period considered. Since this has happened once it is impossible to rule it out in the future. But market history offers no real guide as to the chances of such a period of disturbances re-emerging.

Table A4. Forecasting the probability that the return will be from the second distribution.			
Dependent variable: PROB (probability that the return will be from the second distribution.)			
1573 monthly observations used for estimation from June 1871 to Sept. 2002			
Regressors	**Coefficient**	**Standard Error**	**T-Ratio [Prob]**
CONSTANT	. 061947	.017464	3.5471[.000]
PROB lagged 1 month	.19182	.027088	7.0813[.000]
PROB lagged 5 months	-.0098322	.046367	-.21205[.832]
RETURN lagged 1 month	-.095506	.11162	-.85561[.392]
RETURN lagged 3 months	-221.9198	40.6050	-5.4653[.000]
RETURN lagged 4 months	433.3883	73.7166	5.8791[.000]
RETURN lagged 5 months	-212.6605	43.7230	-4.8638[.000]
VOLATILITY lagged 5 months	.64768	.24906	2.6005[.009]
ChangeQ lagged 3 months	221.0558	40.4273	5.4680[.000]
ChangeQ lagged 4 months	-432.9161	73.5466	-5.8863[.000]
ChangeQ lagged 5 months	212.5155	43.7266	4.8601[.000]
R-Squared	9.6945%	R-Bar-Squared	9.1164%
624 monthly observations used for estimation from Aug. 1950 to Sept. 2002.			
CONSTANT	.061110	.024245	2.5206[.012]
PROB lagged 1 month	.076281	.043873	1.7387[.083]
PROB lagged 5 months	-.14823	.077016	-1.9247[.055]
RETURN lagged 1 month	-.51378	.18036	-2.8487[.005]
RETURN lagged 3 months	-340.2468	58.7290	-5.7935[.000]
RETURN lagged 4 months	561.8869	124.8162	4.5017[.000]
RETURN lagged 5 months	-223.9291	86.0918	-2.6010[.010]
VOLATILITY lagged 5 months	1.3595	.42939	3.1662[.002]
ChangeQ lagged 3 months	339.0384	58.5614	5.7894[.000]
ChangeQ lagged 4 months	-561.2402	124.5900	-4.5047[.000]
ChangeQ lagged 5 months	223.2484	86.1048	2.5928[.010]
R-Squared	13.220%	**R-Bar-Squared**	11.804%

Appendix 2.

Leveraged Equity Portfolios.

Efficient distribution of equity risk is an essential assumption in models that suggest that there may be an equity risk premium “puzzle”. Individuals with a long time horizon would wish to hold a higher beta portfolio than individuals with a short time horizon. In this appendix we wish to show that in a world with limited liability, leveraged portfolios are similar to holding call options. Furthermore, call options do not appear to be an easy route to generating high beta, high expected return portfolios. This will make the efficient distribution of equity risk difficult.

By definition, a leveraged portfolio is long equities and short a loan. There is a possibility that the value of the equities will be less than the value of the loan, so the bank will require a rate of interest above the risk free rate. This spread can be considered as the premium the bank is receiving for effectively selling a put option on the market[9]. The leveraged investor is therefore long equities, long a put option, and short a risk free loan. According to the principle of put-call parity[10], this is equivalent to being long a call option.

The expected return on a call option should be approximately equal to the expected return on the stock multiplied by the value of the stock, then multiplied by the delta of the option divided by the value of the option.[11] All options will therefore lie on the CAPM if line they’re “fairly” priced. Options could therefore represent a mechanism by which high beta portfolio could be created.

A puzzle in finance has been that options appear "over-priced" on the basis of the Black-Scholes model. The implied volatility (the level of volatility then would need to be assumed to generate the market price) for options is generally greater than the actual volatility. This means that although the expected return on at-the-money call options is greater than on the underlying asset, the expected return is not as great as one would anticipate from

[9] The put option can be considered as a plain vanilla option if the bank was to take the unusual step of lending the money for a fixed period of time with no requirement to maintain a minimum level of cover. If the bank, for example, required that the fund had cover of at least 20%, and this was reassessed daily, the bank would be effectively selling a series of 1 day put options with the strike 20% below the previous day’s close. These options would last until one is exercised.

[10] Strictly speaking put-call parity only holds for European options.

[11] If a stock has a price of $100 and it is expected to increase in value by 5%, a call option that has a delta of 0.5 will be expected to rise in price by $2.50. If it has a price of $25, then the option will be expected to increase in value by approximately 10%.

examining the beta of the option.[12] The work in our report helps partially explain one feature of option pricing, known as the volatility "smile". This is used to describe the fact that the implied volatility of "non at-the-money options" is greater than for "at the money options" and more so for "out-of-the-money" put options. The smile is probably the result of the fact that very low returns occur more frequently than the Black-Scholes model would suggest.

In order to fully explain option pricing, a more complex general equilibrium approach is required which recognizes that stochastic volatility of asset returns is an important factor in pricing assets. Attempts at shorting an option and then dynamically replicating that option will at times require a substantial injection of funds. One possibility is that such states of the world may also be times when other income, and hence consumption, are depressed. Writers of options therefore are likely to require an additional risk premium over and above that suggested by the beta of the option.

The high price of call options implies that they will not have an expected return that is implied by their beta. We believe therefore that neither purchases of call options nor leveraged portfolios can be used to create very high expected return portfolios.

www.smithers.co.uk

[12] See, for example, "Expected Option Returns" JD Coval and T Shumway, July 9th, 2000, University of Michigan Business School. http://www-personal.umich.edu/~shumway/papers.dir/optret.html