

Preliminary draft. Please do not quote.

Volatility and Portfolio Protection Over 107 Years*

Elroy Dimson, Paul Marsh and Mike Staunton

edimson@london.edu pmarsh@london.edu mstaunton@london.edu

London Business School, Regents Park, London NW1 4SA, UK. Tel: +44 20 7000 7000

First draft: 4 May 2007

Abstract: We create and analyse a long-term history for various types of equity investment that are protected against downside risk. To do this, we estimate the performance achievable by employing modern investment technologies and markets, such as purchasing or dynamically creating portfolio insurance. We report the returns and risks that would have been experienced by a tax-exempt investor following these strategies, before deduction of management fees. We conclude that protected portfolios are relatively unattractive for the long-term investor, who would typically be better served by multi-asset diversification than by purchasing short-term portfolio insurance.

JEL classification: G11, G13, G18, G24, G28, N20, N22, N24

Keywords: Portfolio protection, financial history, downside risk, volatility, long-term returns

* Presented at the 3rd Beta-Workshop in Historical Economics: Financial History, Strasbourg, 4-5 May 2007

Volatility and Portfolio Protection Over 107 Years

What should investors expect if they seek the upside while avoiding the downside of stock market investing? How would guaranteed products have performed in the past? How would their risk have compared to unprotected investment in equities? By taking a long-term historical perspective, we reveal the potential attractions and the costs of structured products for investors with a long investment horizon.

We investigate strategies that eliminate downside scenarios that, at the beginning of the investment period, appear unlikely but potentially serious. An example would be the purchase of portfolio protection, perhaps by buying a put option at a strike that is, say, 10% or 15% below today's market level. We also look at other approaches to insuring a portfolio against adverse outcomes. We find that risk-controlled products, which curtail downside exposure, offer limited benefits to long-term investors. They would often be better served by allocating a larger proportion of a portfolio to cash, and a smaller proportion to risky securities. In addition, whereas diversified passive strategies can nowadays replicate index performance almost costlessly, the same cannot be said of protected investment strategies: after fees, the long-term risk-return profile of protected products is less attractive than low-cost indexation. For long-term investors, the evidence favours simple approaches to strategic asset allocation.

The structure of our paper is as follows. In Section 1, we describe a variety of approaches to protecting portfolios against adverse investment performance. Our evidence is based on a database of daily stock market returns for the US since 1900 and for the UK since 1930. This enables us, in Section 2, to consider traditional stop-loss rules for portfolio management; and in Section 3, to examine the impact of stock market volatility over the long term. Section 4 outlines more contemporary portfolio hedging strategies, and in Section 5, we study their performance over short and longer horizons. Section 6 examines alternative methods for risk control, and Section 7 concludes. The Appendix describes the volatility estimates that underpin our analysis.

1. Protected portfolios

Many European countries recently introduced pension reforms. Each country is different, but Germany's Retirement Savings Act of 2001 is of particular interest. This reform triggered extensive changes in the market for savings and investment products. While the Act introduced tax-preferred treatment of contributions to private pension plans, at the same time it required a nominal capital guarantee for pension contributions; see Maurer and Schlag (2003) and Seier (2006). The rationales were, first, to protect uninformed individuals through a statutory minimum return; and second, since the state pension provides a floor to retirement income, to protect the welfare system against the moral hazard of excessive risk taking by pension funds.

Protected portfolios are marketed to pension fund, mutual fund, bank and insurance company investors; see Bethel and Ferrel (2007). They are promoted to less sophisticated savers among the public, such as those making counter purchases in a UK Post Office, without a financial advisor's guidance. These products aim to guarantee some minimum final value, or to control downside exposure on the path from initial investment to final maturity. Many approaches can be used in structuring portfolio returns, such as buy-and-hold, stop-loss, constant-mix, constant proportion portfolio insurance, option based portfolio insurance, and ratcheted lock-ins. In this section, we briefly review these alternatives.

The buy-and-hold strategy is a 'do-nothing' solution to asset allocation, in which an initial mix is bought and then held—for example, an initial 60% in shares plus 40% in cash, which is never rebalanced. No matter what happens, the portfolio's nominal value cannot fall below the value of the initial cash investment, which can be set at a level that meets a maturity guarantee.

Stop-loss strategies enable the investor to start with a higher commitment to equities. A minimum level of performance may be assured through setting a downside threshold, at which risky securities are exchanged for the risk-free asset. Apart from any jump in security prices from the time of the sell signal to completion of the stop-loss trade, this ensures that there is a floor to performance.

The constant-mix strategy requires a set proportion of the portfolio to be invested in equities. When asset values change, the investor buys and sells to revert to the target mix. In general, rebalancing to a constant-mix requires the investor to buy shares as they fall in value and sell as they rise in value. Rebalancing could be as frequent as daily, or could be triggered when the asset holding deviates from the long-term policy mix by more than a specified margin, say 2%. As long as the stock market trends upwards, the constant-mix strategy underperforms a buy-and-hold benchmark, since it is constantly selling equities into a rising market. If the stock market trends downwards, the constant-mix strategy also underperforms a buy-and-hold benchmark, since it is constantly buying equities from a declining market. However, in a flat but oscillating market, the constant-mix strategy beats a buy-and-hold benchmark, since it involves buying equities low and selling when they are high.

When the manager guarantees a full return of the initial investment, after a number of years, the simplest products start with an asset allocation of just enough bonds so that the par value at expiry, together with re-invested coupons, equates to 100% of the initial investment; the rest of the initial investment is then allocated to equity. If the equities rise, the buffer increases in value and the manager can increase the equity proportion. This is a stylised version of portfolio insurance, which can be implemented more formally as constant-proportion portfolio insurance (CPPI).

CPPI involves buying equities as they rise and selling as they fall. The investor selects a floor below which the portfolio value is not allowed to fall. The floor rises in value at the rate of return on cash. Defining the difference between the assets and floor as a 'cushion,' CPPI keeps equity exposure at a constant multiple of the cushion. Consider a portfolio worth \$100 and a floor of \$75, giving an initial cushion of \$25, and assume the multiple is 2. The initial equity holding is therefore \$50 (2 x the initial cushion), with the other \$50 in cash. Suppose the equity market falls by 10%, so the portfolio is now worth \$95 and the cushion is now only \$20. Equity exposure should now be \$40 (2 x the new cushion), which requires the sale of \$5-worth of equities and a \$5 increase in cash. If equities fall further, more will be sold; if they rise, more will be bought.

Option-based portfolio insurance (OBPI) is an alternative to CPPI. In an OBPI structure, a zero-coupon note is purchased at the outset to cover the guarantee of the product. The rest of the capital is used to purchase options on the underlying investments, which are often indices. With this structure, if the options fall in value there is no need for assets to be reallocated as the zero has covered the guarantee, and the investor is not directly reliant on the market remaining continuously liquid.

Buy-and-hold is a special case of CPPI, where the floor is the value of the cash holding and the multiplier is equal to 1. The constant-mix strategy is another special case of CPPI, where the floor is zero and the multiplier is less than 1. Which strategy gives the highest returns? In a flat oscillating market, CPPI does relatively poorly, as the investor buys on strength only to see the market weaken, and sells on weakness only to see the market rebound. If the stock market trends upwards, without too many reversals, the CPPI strategy outperforms a buy-and-hold benchmark, since it is buying more equities as they rise. In a downward-trending market, CPPI puts more and more into cash as the equities decline, eliminating equity exposure as the portfolio approaches the floor. Provided a collapse does not impede rebalancing, the CPPI portfolio must give a return that at least exceeds the floor.

In most cases, products that are described as guaranteed or protected provide assurance to the investor that he will receive at least a specified return after a fixed term, with the potential for more substantial performance depending on various factors. The guaranteed return is typically reimbursement of the original investment—a money-back guarantee. However, in other products there is a ratcheted capital guarantee, which offers the potential for an additional return that is often linked to an index but may be subject to caps or floors. The ratchet enables gains to be locked in periodically, for example every month or quarter. One firm explains its system like this: “Our portfolio managers use a ‘ladder’ mechanism to lock in performance throughout the life of the mandate: as soon as there is a significant increase in the portfolio’s value, the protection target is raised incrementally.” With a ratchet, if the fund has grown in value beyond a pre-specified step-size, this will be largely preserved through a new level of protection. Ratcheted protection has become popular as a means of reducing downside exposure, over single and multiple time periods.

2. Traditional protection strategies

The traditional approach to risk management—now replaced by more recent approaches such as CPPI—is to provide capital protection through a stop-loss arrangement. To implement this, the investor selects an investment horizon over which cumulative performance must not violate a specified floor. Risky assets are then sold when the floor is breached, and the proceeds put into cash.

For example, if the horizon is one year, the portfolio may start at New Year with a full allocation to equities. With a floor of -25%, an investment of \$100 would be kept in equities unless the value declines to \$75 or less. As the portfolio reaches \$75, the portfolio is liquidated and is reinvested in Treasury bills. There it remains until the investment horizon ends. The simple stop-loss approach can serve as a benchmark, against which to compare the effectiveness of the solutions outlined earlier.

Stop-loss strategies

Table 1 reports the results of managing the equity index portfolio, with a stop-loss at various levels of year-to-date (YTD) underperformance. A stop-loss of -100% corresponds to a buy-and-hold strategy. Beyond this, the trigger points are -35%, -30%, and so on to -10%. We follow this approach from the start of 1900 for the US, and from 1930 for the UK. Each year, we calculate the portfolio's return, which is the maximum of the index return or the return achieved after breach of the stop-loss.

Table 1 around here

The upper panel starts with summary statistics for the sequence of buy-and-hold stock market returns for the US over 1900–2006. The first row shows that the geometric mean return on US equities was 9.77%; the arithmetic mean of the 107 annual returns was 11.65% and their standard deviation was 19.79%; and the Sharpe Ratio of the US market was 0.38. Subsequent rows report the performance of US stop-loss strategies that have successively greater downside protection. For example, with a floor of 15%, the annualised return shrinks to 8.99%; and the arithmetic mean and standard deviation fall to 10.57% and 18.61% respectively. This decline in risk is because the

stop-loss shifts the portfolio into cash for at least part of the time, which limits volatility; and in addition, the portfolio is more likely to exit from equities in years that have experienced a severe downturn and are relatively volatile. The decline in return is because the portfolio quits equities after a market fall, yet, as Figure 1 illustrates, there is no reason to expect the fall to continue. Too often, the portfolio misses a recovery and, by being in cash some of the time, the stop-loss portfolio on average underperforms the market.

Figure 1 around here

Moving down the rows of the upper panel of Table 1, the portfolio exits equities with increasing probability. Consequently, its beta is lower; the number of switches is higher; and the proportion of the time spent in equities falls. As the trigger gets closer to zero, stop-loss portfolio returns decline faster than the standard deviation. Accordingly, even when transaction costs are ignored, the Sharpe Ratio is impaired by the stop-loss. The lower panel of Table 1 reports the corresponding results for UK equities over 1930–2006, and the pattern is similar. The stop-loss attenuates return as well as risk, and does not enhance performance, as compared to the buy-and-hold alternative. Even selecting the stop-loss with the best Sharpe Ratio—chosen with hindsight—we are unable to beat the Sharpe Ratio of the overall equity market.

Profit lock-ins

As with stop-loss rules, it is also possible to reduce risk by locking in upside gains. The process is one in which a portfolio, upon reaching a target level of appreciation, liquidates its position in risky securities until the end of the investment horizon—here taken to be a year. The reduction in risk involves exchanging an uncertain return, over the final part of the year, for a known return. This eliminates the possibility of slippage in which initial gains are ceded as equities suffer a late fall in value.

The upper panel of Table 2 reports the results of a lock-in strategy for US equities. The top line displays the summary statistics for a buy-and-hold equity strategy (represented here by a lock-in of +999%). Over the long term, lock-ins attenuate the annualised return from 9.77% to 5.86%, which is only a small margin above the return

on US bonds and bills of 4.94% and 4.00%. The lower panel presents the outcome from a lock-in strategy for the UK. Lock-ins attenuate the annualised equity return from 11.48% to 6.78%, which is only a small margin above the return on UK bonds and bills of 6.82% and 5.56%. Compared to the buy-and-hold policy, the Sharpe Ratios from a lock-in are never more than marginally higher (before costs) and can be considerably lower (after costs).

Table 2 around here

The costs of implementing a stop-loss or lock-in include commissions, stamp duty, bid-ask spreads, organization of short-notice program trades, and the price pressure when selling into a declining market. We have also assumed that trading prices can be backdated, so that trades can be executed at the day's closing index level, and that there are no adverse tax consequences from total switches into or out of the market. Additionally, we have ignored management fees. Consequently, our assumed bid-ask spread for equities (of zero, one or two percent) probably present an impression of performance that could have been achieved in practice. Finally, we have not presented results with an investment horizon that spans multiple years, since a long investment horizon can take a portfolio out of equities and into cash for such extended periods that the stop-loss strategy can cease to provide 'equity-like' stock market exposure.

We have seen how, in principle, a floor can underpin performance through a mechanism as simple as a stop-loss strategy. There is considerable inefficiency in such an approach, however. As the portfolio's value heads towards the floor, there is no anticipatory rebalancing of the portfolio, and the asset switch is 'saved up' for one all-embracing transaction that might or might not happen. If the value passes through the floor, there is a total sell-out from the risky asset, and therefore no participation in a potential subsequent recovery.

The alternative is to enter into a transaction before financial disaster strikes, in order that the floor is not violated. This involves insuring against adverse events, rather than trading down once the worst has happened. To put this alternative in context, it is helpful if we review the volatility of the stock market over the long term. At this point, we should insert CPPI results (based on alternative multipliers).

3. Stock market volatility

To understand the context for protecting portfolios, it is important to consider stock market volatility. In Dimson, Marsh and Staunton (2002, 2007), we present long-term return indices for multiple countries. For two countries, the US and UK, we have especially high-quality index series that measure monthly total returns. We fit daily data within each month using the broadest available daily index for each market, with the result that we have accurate measures of cumulative total return over the very long term, accompanied by good estimates of intra-month volatility.

To estimate the intra-month variability of stock returns, we use the daily Wilshire 5000 from December 1979. Before that, the indices are the S&P 500 from inception on 4 March 1957, preceded by the S&P 90; before 1928, we use the Dow Jones Industrial Average. For the UK, we use the daily ABN AMRO/LBS UK All Equities index from end-1999, which follows principles articulated in Dimson and Marsh (2001). Before that, the indices are the FTSE All Share from inception on 10 April 1962, preceded by the FT 30 Share from inception on 1 July 1935; and before then, the Financial News 30 Share, which was published daily from the start of 1930.

Realised volatilities

Using this daily data, we compute the standard deviation of index returns for every month since 1900 (1930 for the UK). These volatility estimates, computed as the root mean squared daily returns over (typically) 21 trading days, are plotted in Figure 2. Expressed as an annualised standard deviation, the volatility of daily US equity index returns has varied from a low of 2.9% (in February 1964) to a high of 83% (in October 1987). In the UK, the annualised standard deviation of daily equity index returns has varied from a low of 1.2% (in March 1944) to highs of 59% (in October 1931) and 60% (in October 1987).

Figure 2 around here

The chart confirms that volatility is mean reverting, with long intervals of subdued stock fluctuations and occasional spikes of extreme variability. Apart from volatility

clustering and mean reversion, there are other stylised features of financial market volatility. Engle and Patton (2001) highlight asymmetric responses of volatility to stock price changes, fat-tailed return distributions, and the impact of exogenous variables. Such attributes are well documented in the derivatives and volatility forecasting literature, and similar patterns are reported for markets in other countries (Panetta et al, 2006) and periods (Gerlach, Ramaswamy and Scatigna, 2006).

Implied volatilities

To take a closer look at the market's view of volatility, we report in Figure 3 on the VIX index, which measures the annualised standard deviation implied by S&P 500 option prices. The VIX (and pre-1990, the VXO) has varied from a low of 9.3% to a high of 150.2%, with an average of 19.9%. As Carr and Wu (2006) note, VIX implied volatilities exceed realised volatilities for the S&P 500 by an average of around five percentage points. On 21 November 2006, VIX implied volatility was down to 9.9%, and on 15 December, the VIX reached an intra-day low of 9.4%.

Figure 3 around here

While many broad stock market indices in the US and UK breached all-time highs at the turn of 2007, the volatility indicated by the VIX index fell from 45% (on 5 August 2002) to around 10% in early 2007. The volatility of four-and-a-half years ago was associated with the collapse of stock prices after the tech boom. The relative calm of contemporary markets reflects the steadiness of the recovery since then. The low VIX suggests low investor concern about future volatility. This may be because the consensus is that volatility is expected to be low or because of an increased appetite for risk. For those who are concerned about potential stock price declines, today's levels for the VIX offer downside protection at historically low premia.

Volatility forecasts

For forecasting future volatility, there is a choice between four types of predictor. These are the standard deviation implied by option prices (the ISD); a variety of measures of historical volatility (referred to generically as HISVOL); several

autoregressive conditional heteroskedasticity models (referred to generically as GARCH); and stochastic volatility models (SV). There is a consensus that the ISD is superior to other predictors. Poon and Granger (2003) review 34 studies that compare ISD-based volatility forecasts with HISVOL-based forecasts, reporting that in 76% of studies, ISD-based predictions are more accurate. They also review 19 studies that compare ISD-based forecasts with GARCH/SV-based forecasts, reporting that in 95% of studies, ISD-based predictions are more accurate. Apart from a few exceptions, such as Aboura and Villa (2003), research on ISDs is US based—the recency with which volatility started to be traded in other markets, coupled with illiquidity problems, have dissuaded researchers from studying non-US markets.

Almost tautologically, the cost of portfolio protection is determined more by the ISD than by forecasts of realised volatility, so the measure of choice is the ISD. In the Appendix, we present an empirical model that estimates ISDs from HISVOLs. The model is accurate, at least on an in-sample basis. In the absence of direct evidence from the financial markets, we use it to estimate ISDs from historical volatilities for all years when daily returns are available but the VIX is unavailable. Chain-linking this series to the VIX, we compute a 107-year history of estimated ISDs for the US. Applying the same model parameters to the UK, we also compute a history of estimated ISDs for the UK, starting in 1930 when daily index data began to be published. The estimated ISDs are presented in Figure 4.

Figure 4 around here

4. Hedging against decline

We estimate the returns, over the long haul, from running a portfolio that is protected ex ante against an unacceptable loss. The task is not to compute what portfolio performance would have been for a person who lived a century ago and applied investment principles that were generally accepted in his or her lifetime. Rather, we apply modern investment approaches to our long run of historical data.

In principle, the hedge can be implemented by buying put options, and we can simply accumulate portfolio returns net of the cost of the puts. The results will indicate the

impact on performance, net of the cost of put options valued using our choice of option-pricing model. The key input to the latter is, of course, the estimated volatility of the market. To achieve robustness in our analysis, we also select a hedging strategy that is insensitive to the estimated volatility. This involves purchasing portfolio protection through a collar strategy. In a collar, option purchase is financed by taking a short position in another option. Though the individual option values are sensitive to the volatility estimate, the combined position is relatively insensitive to measurement error in volatility. Consequently, we can be more confident about using the estimated ISDs from Figure 4 to assess the terms for insuring a portfolio's value against decline.

Option-based insurance

The principle of buying portfolio protection is to eliminate adverse returns. In Figure 5, we show the distribution of annual returns from the UK equity market. The histogram highlights the dispersion of historical stock market returns, and the precise years in which returns of various magnitudes occurred. Some years were wonderful, others terrible. What investors want is a means of enjoying good equity returns while avoiding the downside.

Figure 5 around here

Figure 6 segments the tails of the return distribution. In the top-left, we show the distribution of UK equity returns, truncated so that no year was worse than -10% . This return distribution avoids much of the impact of extreme financial disasters. The net effect is that those years which would have been particularly poor (returns below -10%) are stacked up with an outcome of exactly -10% . They are identified in the histogram by grey shading. This kind of skew to returns can be constructed through dynamic management of the portfolio or by buying a suitable put option. As noted above, this is the first option-based strategy that we examine.

Figure 6 around here

On the top-right of the chart, we show the distribution of annual returns truncated on the right, so that in no year was the return better than $+30\%$. This return distribution

experiences financial disasters, but curtails the benefit from the best years. In the years that are shaded in grey, returns are limited to exactly +30%. In effect, the investor has sold covered call options. In the bottom-left panel, we conjoin the diagrams in the upper panels, truncating stock market returns on both the downside and the upside. The original returns survive in the middle of the distribution, but there is an increased likelihood of a return equal to the floor or cap. This distribution of returns is safer, in the sense that very poor returns are eliminated (albeit at the expense of very favourable returns).

A costless collar

What will it cost to make stock market performance safer in this way? The bottom-right of Figure 6 shows the truncated tails of the return distribution. Buying a put (avoiding the adverse performance on the left of this distribution) has a cost, but selling a call (foregoing favourable performance on the right) generates an income. In the collar strategy, the investor sells a claim on the portfolio's performance in excess of a stated cap. The cap is selected so that the proceeds from sale of the call option are sufficient to buy a protective put that will insure the portfolio against a loss of selected magnitude. The investor participates in her portfolio's return only up to the call strike price. This strategy is referred to as a 'costless collar'.

With the costless collar, the investor has eliminated price risk below the put strike price, and cannot lose more than a specified proportion of the value of her portfolio. The payoffs to which she is exposed look like the solid line in Figure 7: whatever happens to the market, there is a floor to portfolio performance, accompanied by a cap on the upside. For example, if an investment of \$100 were protected with a floor of \$85, the hedge would protect against a 15% decline over the maturity of the collar. Ignoring fees and taxes, there is a zero net premium for this method of insuring against adverse outcomes. Because the premium is zero, the sloping part of the solid line (the protected portfolio) coincides with, and does not fall beneath, the corresponding part of the dashed line (which is essentially an unprotected index-tracking portfolio).

Figure 7 around here

To convey the broad nature of the trade and the general magnitude of costs, assume that the index's at-the-money (ATM) volatility is 15%, and the interest rate and dividend yield are 4% and 2% respectively. To provide a -15% floor over the next year, the put option is priced using a volatility smile with a skew of 6%. This choice of skew parameter is based on Rubinstein (1994) and Allen, Einchcomb and Granger (2006), and we also refer to Demeterfi, Derman, Kamal and Zou (1999). The skew raises the volatility for the put by 6% for each 10% that the strike is out of the money, namely by 9% ($6\% \times 15\% / 10\%$). The put is therefore priced on a volatility of 24% (the ATM volatility of 15% plus an additional 9%), and it therefore has a value of \$2.75. A call option with a strike of \$111 would also sell for \$2.75. Putting these two elements together, buying the put and writing the call would require a net premium of zero.

In practice, of course, a dealer or market maker would not do the trade on these terms: there would be commissions and spreads to consider. We apply option and equity spreads of 0.02 and 0.01 respectively. Our model thus includes incremental dealing costs, but no management fees. Following this procedure, we estimate the long-term performance of the protected portfolio strategy over a sample period of up to 107 years.

5. Performance of protected portfolios

Ratcheted protection

We first examine the long-term performance of investing in equities, with downside risk hedged by means of a protective put or costless collar, and with the stated level of protection reinstated monthly. Table 3 presents the results. The upper panel reports on the US equity market from the start of 1900 to date, while the lower panel reports on the UK equity market from the start of 1930 to date. The Buy & Hold column reports the cumulative return on the market index when there is no downside protection, as reported in Dimson, Marsh and Staunton (2007). The remaining two blocks report long-term performance if the downside is protected by purchase of a protective put or by implementing a costless collar, in each case with a monthly floor of -10%, -8%, -6%, -4%, or -2%.

Table 3 around here

With a base value in 1900 (for the US) or 1930 (UK), the start-2007 value of the equity index is large. As protection becomes more stringent, the underlying exposure of a protected portfolio is attenuated, and the start-2007 index value is reduced. This pattern is also apparent in the reduction in geometric mean returns and arithmetic mean returns as protection progresses from none (buy-and-hold) to very high protection (a floor of -2%). The decline in risk is equally apparent. For the US over the last 107 years, the annualised standard deviation of portfolio returns declines from a value of 19.79% (for an unprotected buy-and-hold portfolio) to 3.90% (to a portfolio tightly protected portfolio by a collar strategy). For the UK over the last 77 years, the annualised standard deviation declines from a value of 24.32% to 3.73% . Similarly, the beta—measuring the portfolio’s sensitivity to fluctuations in the national equity market—declines from 1.00 to 0.13 (in the US) or 0.14 (in the UK).

The bottom rows of each panel record the extent to which return or risk declines faster. The Sharpe Ratio measure the mean return in excess of the risk-free interest rate, expressed as a proportion of the standard deviation. The Sharpe Ratio is in general reduced by portfolio protection. That is, the standard deviation of a monthly-protected strategy is generally higher than that of a simple equities/cash mix that has the same mean return.

In Figure 8, we plot the combinations of arithmetic mean return and standard deviation that were attainable through various levels of portfolio protection. The points on the chart are taken from the arithmetic means and standard deviations in Table 3, and from additional entries that were omitted from the table for reasons of space. Also shown is a capital market line that connects the risk/return profiles of a risk-free investment (taken here to be Treasury bills) and an unprotected passive investment in the equity market. Any point on the line can be attained by selecting the appropriate combination of Treasury bills and equities. The slope of this line measures the reward per unit of risk in the stock market—in other words, the Sharpe Ratio of the overall stock market.

Figure 8 around here

In the chart, the top-right point represents unprotected investment in the stock market—essentially, an equity index fund. As one moves from that point towards the left, risk and return are both reduced. The US protected portfolios plot entirely beneath the capital market line, while the UK portfolios plot above and then beneath the capital market line. In the US, for a given expected return, standard deviations are lower if the investor were simply to split the portfolio between cash and equities. Similarly, for a given standard deviation, returns are higher if the investor were to split the portfolio between cash and equities. In the UK, the record is mixed, but is essentially like the US. In both countries, the impact of a few extreme months has a marked impact, because protection against one or two outliers may provide a benefit that contributes negatively or positively to the Sharpe Ratio. All portfolios protected through a collar strategy have a lower mean return and lower standard deviation than an unprotected index fund. Add a second line

In reality, it is relatively expensive to engineer a protected portfolio and relatively inexpensive to buy and hold a passive portfolio of equities plus Treasury bills. The fees, embedded by the intermediary creating the structured product, are larger than the fees normally charged for managing an index fund. There is no indication that monthly protection has a significant risk-reducing impact, beyond what could easily be achieved from an unprotected portfolio split between equities and short-term fixed income deposits.

Annual protection

Downside protection is typically over an investment horizon that exceeds a month, and is usually not ratcheted. Here we consider an investment horizon of a year (or more). Unfortunately, data series are rarely available to estimate the properties and cost of such strategies. As Zumbach (2007) points out, “even a time series of 150 years will not resolve the statistical problem for risk evaluation at the scale of one year.” Our approach is to follow a process of monthly protection that is not ratcheted. Protection over intervals longer than a month—here, a year—then follows monthly rollovers of either the protective put or the costless collar. We use the same monthly predictions of standard deviation as before.

We vary the floor set at the start of each month, so that, from February onward, we continue to protect the dollar amount established at the start of the year. For example, if the start-year floor is -10% , and the market index rises from 100 to 120 over the course of January, then the floor for the month of February is set to -25% . This insures against the index falling below 90. Similarly, with the same start-year floor, if the market index falls from 100 to 90 over the course of January, then the floor for the month of February is set to -0% . This forces the portfolio into cash, thereby ruling out any upside for February or subsequent months until the year-end. The resulting annual returns are aggregated to produce a cumulative return series for each protected strategy.

Table 4 reports the results, with the US in the upper panel and the UK in the lower panel. As in the preceding table, based on monthly protection, the first column reports the cumulative return when there is no downside protection. The remaining columns of the table indicate the long-term risk and return attributes if capital losses are constrained to a floor of -20% , -15% , -10% , -5% , or -2% in a single year, using a protective put or costless collar to implement the strategy.

Table 4 around here

In Figure 9, we plot the combinations of arithmetic mean return and standard deviation that were attainable through protection over an interval of a full year. Similar to the earlier diagram, the points on this chart are taken from the lower panel of Table 4 and from additional, omitted entries. The curves in Figure 9 are less smooth than in Figure 8 and are no longer monotonic. Whereas the returns in Figure 8 were based on up to 1284 months, in Figure 9 we use only 107 or 77 annual observations. We note in Dimson, Marsh and Staunton (2007a) that there have been some extreme UK market returns, such as 1974 (-48.8%) or 1975 ($+145.6\%$), which have a marked influence on the performance of an investment strategy and give rise to greater imprecision in Figure 9. Add a second line to the plot for protective put.

Figure 9 around here

The pattern documented earlier for monthly protection nevertheless carries over to the

strategies built around an annual investment horizon. Although a longer horizon may control turnover, it introduces complexities like time-varying moneyishness for the protective option. At the start of the year, the put will be out of the money. As the year progresses, the market may rise and the cost of a mid-year put may become low. On the other hand, the market may fall, and the put moves towards being at the money. Once the portfolio is thereby forced into cash, it cannot share in a stock market recovery. This dilemma obviously extends to protection that spans intervals longer than a year, and reflects the increasing cost of portfolio insurance as the investment horizon grows longer—a factor pointed out by Bodie (1995).

6. Risk and diversification

The results above describe the disappointing risk/reward attributes over the long term of protective put and costless collar strategies. The performance generated by the protective put depends crucially on the estimated ISD, but the collar strategy is insensitive to the ISD because estimation errors for the cost of the protective put are offset by estimation errors in the premium from sale of the call. In Table 5 we examine the sensitivity of our findings to key assumptions. The table demonstrates the impact of volatility estimation when the put is acquired without a collar. Apart from that, the Sharpe Ratios, and our general conclusions, are broadly insensitive to the parameter values used in the collar strategy.

Table 5 around here

It may therefore be unwarranted to place heavy reliance on protective portfolio strategies as the primary source of risk control for a portfolio. What is the alternative? In Table 6, we report on the simplest form of risk management—diversifying a portfolio between domestic equities and government bonds. The upper panel of the table analyses 107 years of monthly total returns on these two asset categories in the US; the lower panel analyses 107 years of equivalent data for the UK. The data are from the Dimson, Marsh and Staunton (2007) database.

Table 6 around here

The first block of numbers summarises asset class attributes from Dimson, Marsh and Staunton (2007). Equities have had an identical nominal return in the US and UK, averaging an annual 9.8% (geometric mean) or 11.6% (arithmetic mean) in both countries. Government bonds had a somewhat higher realised return in the UK, as compared to the US. The standard deviation of annual equity and bond returns has been higher in the UK, so the UK has a slightly lower Sharpe Ratio for both assets.

The eight columns on the right of Table 6 show that risk can be reduced effectively by spreading assets across equities and bonds. In the US panel, the historical Sharpe Ratio is slightly enhanced by diversifying, in appropriate proportions, across both assets. In the UK, the proportion of bonds that gives rise to an enhanced Sharpe Ratio is under 10%. This, however, is a 107-year historical figure, and the actual benefit from diversification depends on the extent to which the two asset categories have returns that are uncorrelated or, ideally, negatively correlated. Fortunately, as Figure 10 reveals, the correlation between equities and bonds, which was markedly positive through most of the 1980s and 1990s, has moved into negative territory. There is today renewed scope to benefit from diversifying across asset classes. International diversification offers further opportunities for risk reduction.

Figure 10 around here

7. Conclusion

We have constructed a return series for the US and UK equity markets, with performance constrained not to breach a prespecified limit. Buying protection through the derivatives-based strategy, described here, reduces not only risk but also return. The net effect is that these protected strategies offer limited benefits, but also some drawbacks, as compared to reducing equity exposure through conventional asset allocation. Changing the strategic asset mix can give rise to Sharpe Ratios that are similar, and quite often superior, to an option-based solution.

We have examined 107 years of returns for the US and 77 for the UK, with similar results, and we have tested their sensitivity. Why might the long-run performance of a protected strategy be somewhat disappointing, compared with a diversified strategy

without protection? An important explanation is that many investors would like to be protected against severe financial setbacks. This makes it relatively costly to buy protection against large downward market moves, and the volatility implied by derivative prices has tended to be larger than realised volatilities. This has been the case especially in relation to options with an exercise price well below the level of the index—the very instruments that best protect an investor against financial disaster.

Our findings may be contrasted with the apparently attractive returns from writing covered calls. Hill, Balasubramanian, Gregory and Tierens (2006) demonstrate the profitability of buy-write strategies in which the investor buys stocks and writes call options against the stock position. Writing covered calls generates income that more than offsets limitations in upside potential. The CBOE's buy-write index (ticker BXM) implements this investment process, and measures the hypothetical performance from buying a portfolio of S&P 500 stocks and simultaneously writing covered call options on the S&P 500 Index. Callan Associates (2006) confirm favourable performance from the BXM product, and the CBOE reports that, in just two years, more than 40 new buy-write investment products raised over \$20 billion in assets.

Covered calls are equivalent to taking a short position in a put option. If shorting a put is return-enhancing, then buying a put is return-attenuating. It is likely that the historical record of underperformance from protected investment strategies is a permanent feature of securities markets. Clearly, downside protection has a particular appeal for some investors—those with a short investment horizon and whose performance is measured over a relatively brief interval. Long-term investors, who are not penalised for shortfalls over short horizons, should rarely buy put options. They should favour selling insurance to other investors. A popular way of doing this is by setting a long-term asset mix and sticking to it, rebalancing as necessary. This is a simple form of tactical asset allocation in which more is invested in equities after a price fall—quite the opposite of buying portfolio protection.

Appendix: Estimating implied volatility

The cost of portfolio protection is largely determined by the market's consensus view of volatility at the date protection is initiated. We therefore base our analysis on the market's ISD, where the latter is available. For the US, we use the VIX index, which measures implied volatility over 30 calendar days into the future, as our ISD series. We chain-link the VIX, which began in 1990, to the previous VXO volatility index. This provides a series of daily ISDs from 1986–2006. For periods before this, which pre-date explicit trading in volatility indices, we use an estimate of the ISD, derived from historical volatility. We discuss historical volatility estimates and ISDs in turn.

EWMA forecasts

In Dimson and Marsh (1990), we evaluated a volatility forecasting method based on an exponentially weighted moving average (EWMA), and demonstrated that it was superior to a variety of alternative predictors. In this study, we follow the same approach to predict month-ahead standard deviations, using the model below:

$$\sigma_t = \lambda \sigma_{t-1} + (1-\lambda) s_{t-1} \quad [1]$$

where σ_t denotes the prediction, available at the start of month t , of the standard deviation for month t , s_t is the realised standard deviation in month t , and λ is a decay parameter. The standard deviation is based on the dispersion of returns around an expected return of zero, and is therefore equal to the mean of the month's squared daily returns. The daily returns are continuously compounded and are measured without deduction of the sample mean. Realised volatility is then computed on an annualised basis following Carr and Wu (2006).

To interpret equation [1], if the decay parameter, λ , is equal to 1.0, then the prediction for next month's volatility is the equally-weighted average of all prior historical volatilities. If $\lambda=0$ then volatility is implicitly assumed to follow a random walk, whereby the prediction for next month's volatility is for it to be at the same level as the prior month's realised volatility. If $0<\lambda<1$ then the prediction for next month's volatility is an exponentially weighted average of all prior realised volatilities, with older estimates receiving progressively lower weights.

Using monthly observations of intra-month daily standard deviations, we search for the optimal decay parameter, namely the value for λ that minimises the mean squared error (MSE) when predicting the standard deviation of index returns. Over the entire interval from 1986–2006, the optimal value for λ is 0.69. This provides a basis for volatility prediction using a moving window of historical returns. The weighting coefficient, λ , is important for predicting volatilities, since it generates errors (MSEs) that are much below the benchmarks of the long-term mean ($\lambda=0$) or short-term random walk ($\lambda=1$). Note, however, that forecasting accuracy is relatively insensitive to the precise value assumed for λ , and while $\lambda=0.69$ is the optimal value, the MSE rises by only 0.2% if $\lambda=0.60$, and by just 0.7% for $\lambda=0.80$.

The RiskMetricsTM volatility predictions resemble model [1] except that they project variance, σ_t^2 , from the prior period's EWMA forecast, σ_{t-1}^2 , and the realised variance, s_{t-1}^2 . The predicted standard deviation is then taken as the square root of the predicted variance. In the 'classic' RiskMetricsTM approach, Zangari (1996) models variance as:

$$\sigma_t^2 = \lambda \sigma_{t-1}^2 + (1-\lambda) s_{t-1}^2 \quad [2]$$

and identifies a decay parameter of $\lambda=0.94$ as the most suitable value, an estimate that has since been revised upward to $\lambda=0.97$ (Zumbach, 2007). This high value for λ means that the EWMA prediction is relatively insensitive to an extreme realisation of volatility, which circumvents the disproportionate influence of squared returns.

In their extensive survey, Poon and Granger (2003) favour the standard deviation [1] over the variance [2] as a measure of volatility. This is because the MSE (mean squared error) criterion for evaluating variance forecasts is the 4th power of the return, and hence has a wide confidence interval that precludes choosing between forecasting methods. In our analysis of standard deviation forecasts using model [1], the decay parameter is smaller than when forecasting variances. This indicates that the most recent realised volatility is more informative about future volatility when volatility is measured as a standard deviation.

Implied standard deviations

While the EWMA estimation process [1] may be among the better approaches to volatility forecasting, implied standard deviations (ISDs) are the measure that is appropriate for estimating the economic cost of hedging stock market exposure. We therefore determine the market value of portfolio protection from actual ISDs when they are available; and when they are not, we use an estimate of the implied standard deviation. The estimated ISD is based on the following regression:

$$\text{ISD}_t = a + b \sigma_t + \varepsilon_t \quad [3]$$

where ISD_t denotes the implied standard deviation as at the start of month t , and σ_t denotes the prediction, available at the start of month t , of the standard deviation that will be realised over the course of month t . The prediction σ_t is based on model [1], with a decay parameter, λ , which is set at a level that maximises forecast accuracy.

Equation [3] is a more accurate representation of ISDs than can be achieved by splicing data from VIX with historical volatilities, as in Bloom's (2007) recent paper on stock market volatility. There are a variety of adjustments that can take account of volatility skew and other features of option pricing (some of which underpin the model used in this paper). However, Li and Pearson (2006) demonstrate that the most important ingredient for accurate option pricing is a good estimate of the at-the-money ISD.

To interpret equation [3], if the historical volatility predictions σ_t have no forecasting power, for example because $\lambda=1$, then $b=0$ and the regression intercept a will be equal to the mean level of the ISD. If the historical volatility predictions σ_t are perfect forecasts of ISDs, then $b=1$, the regression intercept $a=0$, and $\text{Adj } R^2=100$ percent. If ISDs deviate from historical volatilities, then it is likely that a good forecast of ISD will give some weight to the historical volatility predictions σ_t and some to the long-term mean ISD, implying $0 < b < 1$.

The results of regression [3] are reported in Table 7. The decay parameter that minimises the MSE is $\lambda=0.44$, which is lower than the optimal value when

forecasting realised standard deviations. As compared to predicting realised standard deviations, when predicting ISDs more weight is given to recent observations and less weight is given to observations from long ago (i.e. λ is lower). However, the influence of recent fluctuations is nevertheless curtailed by the regression intercept a in equation [3], which pulls in the forecast ISD towards the long-term mean ISD.

Table 7 around here

The value of $\lambda=0.44$, which produces the best ISD forecasts [3], has an Adj R^2 of 75.2 percent and generates an estimated ISD of $5.59+0.81\sigma_t$. The fit of models, with substantially differing weights λ for forming the EWMA historical volatility predictions, are shown in the table. Using markedly different values for λ still gives rise to very similar levels of explanatory power, as indicated by the Adj R^2 estimates; and the estimates vary over a relatively narrow range.

Figure 11 plots actual month-end ISDs (i.e., VIX data) for the US equity market. Superimposed on the actual ISDs are estimates based on regression model [3]. During periods when volatility indices are unavailable, model [3] is used to estimate ISDs. It is also used, with the same parameters, to estimate ISDs for the UK.

Figure 11 around here

References

1. Sofiane Aboura and Christophe Villa, 2003, International market volatility indices: a study on VX1, VDAX and VIX, Paris Dauphine University working paper, available at: <http://ssrn.com/abstract=394420>.
2. Peter Allen, Stephen Einchcomb, and Nicolas Granger, 2006, Variance Swaps, London: JP Morgan (November).
3. Jennifer Bethel and Allen Ferrel, 2007, Policy issues raised by structured products, in Yasuki Fuchita and Robert Litan (Eds), Brookings-Nomura Papers on Financial Services, DC: Brookings Institution Press.
4. Nick Bloom, 2007, The impact of uncertainty shocks: a firm-level estimation and a 9/11 simulation, CA: Stanford University Economics Department working paper, available at: <http://www.stanford.edu/~nbloom/>.
5. Zvi Bodie, 1995, On the risk of stocks in the long run, *Financial Analysts Journal*, 51(3): 18-22.
6. Tim Bollerslev and Hao Zhou, 2006, Expected stock returns and variance risk premia, Duke University working paper, available at <http://ssrn.com/abstract=948309>.
7. Callan Associates, 2006, An historical evaluation of the CBOE S&P 500 BuyWrite index strategy, CA: Callan Associates (October).
8. Peter Carr and Liuren Wu, 2006, A tale of two indices, *Journal of Derivatives*, 13(3): 13-29.
9. Kresimir Demeterfi, Emanuel Derman, Michael Kamal and Joseph Zou, 1999, More than you ever wanted to know about volatility swaps, *Journal of Derivatives*, 6(4): 9-32.
10. Elroy Dimson and Paul Marsh, 1990, Volatility forecasting without data-snooping, *Journal of Banking and Finance* 14(2): 399-421.
11. Elroy Dimson and Paul Marsh, 2001, UK financial market returns 1955–2000, *Journal of Business* 74(1) 1-31.
12. Elroy Dimson, Paul Marsh and Mike Staunton, 2002, *Triumph of the Optimists: 101 Years of Global Investment Returns*, Princeton, NJ: Princeton University Press.
13. Elroy Dimson, Paul Marsh and Mike Staunton, 2007, *Global Investment Returns Yearbook 2007*, London: ABN AMRO.
14. Elroy Dimson, Paul Marsh and Mike Staunton, 2007a, The worldwide equity premium: a smaller puzzle, in Rajnish Mehra (Ed.) *Handbook of Investments: Equity Risk Premium*, Amsterdam: Elsevier.

15. Robert Engle and Andrew Patton, 2001, What good is a volatility model? *Quantitative Finance*, 1(2): 237-245.
16. Eugene Fama and Ken French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics*, 33(1): 3-56.
17. Lars Forsberg and Eric Ghysels, 2007, Why do absolute returns predict volatility so well?, *Journal of Financial Econometrics*, 5(1): 31-67.
18. Stephen Gerlach, Srichander Ramaswamy and Michela Scatigna, 2006, 150 years of financial market volatility, *BIS Quarterly Review*, September: 77-91.
19. Joanne Hill, Venkatesh Balasubramanian, Krag Gregory and Ingrid Tierens, 2006, Finding alpha via covered index writing, *Financial Analysts Journal*, 62(5): 29-46.
20. Minqiang Li and Neil Pearson, 2006, A horse race among competing option pricing models using S&P 500 index options, GA: Georgia Tech working paper, available at <http://ssrn.com/abstract=952770>.
21. Raimond Maurer and Christian Schlag, 2003, Money-Back Guarantees in Individual Pension Accounts: Evidence from the German Pension, chapter 6 of Olivia Mitchell and Kent Smetters (Eds) *The Pension Challenge: Risk Transfers and Retirement Income Security*, Oxford Scholarship Online.
22. Fabio Panetta (Chair), Paolo Angelini, Giuseppe Grande, Aviram Levy, Roberto Perli, Pinar Yesin, Stefan Gerlach, Srichander Ramaswamy, and Michela Scatigna, 2006, The recent behaviour of financial market volatility, *Bank for International Settlements: BIS Occasional Paper no 29*.
23. Ser-Huang Poon and Clive Granger, 2003, Forecasting volatility in financial markets: a review, *Journal of Economic Literature*, 41(2): 478-539.
24. Mark Rubinstein, 1994, Implied binomial trees, *Journal of Finance*, 49(3): 771-818.
25. Andrea Seier, 2006, *A Micro-Perspective on Funding in the German Pension Reform of 2001*, Lohmar, Germany: Josef Eul Verlag.
26. Peter Zangari, 1996, An improved methodology for measuring VaR, *RiskMetrics Monitor*, 2: 7-25.
27. Gilles Zumbach, 2007, Backtesting risk methodologies from one day to one year, *RiskMetrics Journal*, 7(1): 17-59.

Table 1: Long-term performance of annual stop-loss strategies in US and UK

Market	Floor	Geometric mean %	Arithmetic mean %	Standard deviation	Beta v stock market	Number of switches	% of time in equities	Sharpe Ratio for equity spreads of		
								0.00	0.01	0.02
US	-100%	9.77	11.65	19.79	1.00	0	100.0	0.38		
1900–	-35%	9.40	11.39	20.23	1.01	5	98.8	0.36		
2006	-30%	9.55	11.47	20.03	1.00	7	98.3	0.37		
	-25%	9.49	11.39	20.00	1.00	11	97.2	0.37		
	-20%	9.78	11.54	19.52	0.96	16	94.7	0.38		
	-15%	8.99	10.57	18.61	0.81	24	89.0	0.35		
	-10%	8.83	10.17	17.40	0.70	38	78.3	0.35		
UK	-100%	11.48	13.80	24.32	1.00	0	100.0	0.34		
1930–	-35%	11.33	13.63	24.41	0.99	2	98.8	0.33		
2006	-30%	11.27	13.59	24.46	0.99	4	98.4	0.32		
	-25%	11.36	13.62	24.34	0.99	6	97.1	0.33		
	-20%	11.55	13.71	24.02	0.97	8	95.8	0.34		
	-15%	11.53	13.63	23.89	0.95	14	92.4	0.33		
	-10%	11.70	13.67	23.39	0.91	23	81.7	0.34		

Table 2: Long-term performance of annual lock-in strategies in US and UK

Market	Lock-in	Geometric mean %	Arithmetic mean %	Standard deviation	Beta v stock market	Number of switches	% of time in equities	Sharpe Ratio for equity spreads of		
								0.00	0.01	0.02
US	+999%	9.77	11.65	19.79	1.00	0	100.0	0.38		
1900–	+35%	9.66	11.43	19.00	0.93	13	97.1	0.39		
2006	+30%	9.23	10.90	18.32	0.90	23	94.6	0.37		
	+25%	9.00	10.52	17.31	0.81	35	89.9	0.37		
	+20%	8.37	9.73	16.17	0.73	49	82.5	0.35		
	+15%	8.06	8.96	13.16	0.50	63	70.3	0.37		
	+10%	5.86	6.57	11.53	0.43	74	57.8	0.22		
	UK	+999%	11.48	13.80	24.32	1.00	0	100.0	0.34	
1930–	+35%	10.73	12.43	18.77	0.69	12	94.3	0.36		
2006	+30%	10.54	12.16	18.12	0.64	14	92.1	0.36		
	+25%	10.35	11.89	17.55	0.61	21	88.4	0.36		
	+20%	10.07	11.45	16.31	0.52	36	80.9	0.36		
	+15%	8.72	9.93	15.13	0.47	45	69.9	0.28		
	+10%	6.78	7.78	13.47	0.39	53	58.2	0.16		

Table 3: Long-term performance with monthly protection of the equity portfolio

Summary statistics	Buy+hold	Protective put with various floors					Collar strategy with various floors				
		-10%	-8%	-6%	-4%	-2%	-10%	-8%	-6%	-4%	-2%
United States											
Index (1/1900=100)	2,153,579						1,365,532	1,133,279	519,401	67,330	2,552
Geom mean return %	9.77						9.31	9.12	8.32	6.28	3.07
Arith mean return %	11.65						10.93	10.57	9.46	6.98	3.32
Standard deviation %	19.79						18.44	17.48	15.48	12.10	7.10
Beta rel. to market	1.00						0.78	0.72	0.61	0.46	0.25
Sharpe Ratio	0.385						0.374	0.374	0.350	0.243	-0.101
United Kingdom											
Index (1/1930=100)	430,991						181,862	137,792	83,567	22,552	1,956
Geom mean return %	11.48						10.24	9.84	9.13	7.29	3.94
Arith mean return %	13.80						11.76	11.15	10.14	7.91	4.16
Standard deviation %	24.32						17.89	16.63	14.64	11.43	6.80
Beta rel. to market	1.00						0.79	0.72	0.62	0.46	0.26
Sharpe Ratio	0.335						0.341	0.331	0.307	0.198	-0.218

Table 4: Long-term performance with annual protection of the equity portfolio

Summary statistics	Buy+hold	Protective put with various floors					Collar strategy with various floors				
		-20%	-15%	-10%	-5%	-2%	-20%	-15%	-10%	-5%	-2%
United States											
Index (1/1900=100)	2,153,579						340,206	458,390	123,674	63,510	1,516
Geom mean return %	9.77						7.90	8.20	6.88	6.22	2.57
Arith mean return %	11.65						9.86	9.87	8.46	7.25	3.10
Standard deviation %	19.79						20.16	18.98	18.55	15.13	10.80
Beta rel. to market	1.00						0.78	0.68	0.58	0.41	0.23
Sharpe Ratio	0.385						0.289	0.307	0.238	0.212	-0.087
United Kingdom											
Index (1/1930=100)	430,991						206,100	228,023	108,954	27,421	1,332
Geom mean return %	11.48						10.42	10.56	9.51	7.56	3.42
Arith mean return %	13.80						12.37	12.28	11.10	8.58	3.99
Standard deviation %	24.32						21.17	20.08	19.20	15.35	11.48
Beta rel. to market	1.00						0.84	0.76	0.67	0.50	0.30
Sharpe Ratio	0.335						0.318	0.330	0.284	0.191	-0.144

Table 5: Robustness of results for annual protection

Scenario	Mkt	Summary statistics	Buy+hold	Floor for protective put option			Floor for costless collar		
				-20%	-10%	-5%	-20%	-10%	-5%
Base case: Skew = .06 2% option spread 1% equity spread	US	Geom mean return %	9.77				7.90	6.88	6.22
		Arith mean return %	11.65				9.86	8.46	7.25
		Standard deviation %	19.79				20.16	18.55	15.13
		Beta relative to market	1.00				0.78	0.58	0.41
		Sharpe Ratio	0.385				0.289	0.238	0.212
	UK	Geom mean return %	11.48				10.42	9.51	7.56
		Arith mean return %	13.80				12.37	11.10	8.58
		Standard deviation %	24.32				21.17	19.20	15.35
		Beta relative to market	1.00				0.84	0.67	0.50
		Sharpe Ratio	0.335				0.318	0.284	0.191
Increase every ISD by 10%	US	Geom mean return %	9.77						
		Arith mean return %	11.65						
		Standard deviation %	19.79						
		Beta relative to market	1.00						
		Sharpe Ratio	0.385						
	UK	Geom mean return %	11.48						
		Arith mean return %	13.80						
		Standard deviation %	24.32						
		Beta relative to market	1.00						
		Sharpe Ratio	0.335						
Halve the trading spreads	US	Geom mean return %	9.77				8.11	7.55	6.81
		Arith mean return %	11.65				10.04	9.09	7.85
		Standard deviation %	19.79				20.04	18.36	15.21
		Beta relative to market	1.00				0.78	0.59	0.41
		Sharpe Ratio	0.385				0.299	0.275	0.250
	UK	Geom mean return %	11.48				10.50	9.90	8.42
		Arith mean return %	13.80				12.46	11.47	9.43
		Standard deviation %	24.32				21.19	19.11	15.34
		Beta relative to market	1.00				0.84	0.68	0.52
		Sharpe Ratio	0.335				0.321	0.305	0.247
Halve the volatility skew	US	Geom mean return %	9.77				8.11	7.39	6.53
		Arith mean return %	11.65				10.04	8.93	7.62
		Standard deviation %	19.79				20.04	18.37	15.65
		Beta relative to market	1.00				0.79	0.59	0.42
		Sharpe Ratio	0.385				0.300	0.266	0.229
	UK	Geom mean return %	11.48				10.69	9.98	8.56
		Arith mean return %	13.80				12.65	11.62	9.66
		Standard deviation %	24.32				21.39	19.59	16.01
		Beta relative to market	1.00				0.86	0.69	0.53
		Sharpe Ratio	0.335				0.328	0.305	0.250
Replace ISDs by perfect-foresight volatilities	US	Geom mean return %	9.77				7.88	6.83	5.99
		Arith mean return %	11.65				9.85	8.41	7.01
		Standard deviation %	19.79				20.18	18.55	15.04
		Beta relative to market	1.00				0.78	0.58	0.40
		Sharpe Ratio	0.385				0.288	0.236	0.197
	UK	Geom mean return %	11.48				10.52	9.72	7.21
		Arith mean return %	13.80				12.54	11.32	8.22
		Standard deviation %	24.32				21.68	19.36	15.29
		Beta relative to market	1.00				0.86	0.67	0.50
		Sharpe Ratio	0.335				0.318	0.293	0.168

Table 6: Risk and return from portfolios comprising equities plus bonds

Summary statistics	Asset class attributes			Proportion in equities (remaining proportion in government bonds)							
	Equities	Bonds	Bills	90%	80%	70%	60%	50%	40%	30%	20%
US 1900–2006											
Geom mean return %	9.77	4.94	4.00	9.48	9.14	8.76	8.34	7.87	7.37	6.82	6.24
Arith mean return %	11.65	5.24	4.04	11.01	10.37	9.73	9.09	8.44	7.80	7.16	6.52
Standard deviation	19.79	8.15	2.79	17.91	16.07	14.30	12.62	11.07	9.71	8.63	7.95
Beta rel. to market	1.00	.04	-.01	.90	.81	.71	.62	.52	.42	.33	.23
Sharpe Ratio	.385	.147	.000	.390	.394	.398	.400	.398	.388	.362	.312
UK 1900–2006											
Geom mean return %	9.76	5.35	5.04	9.42	9.07	8.69	8.29	7.86	7.41	6.93	6.43
Arith mean return %	11.60	5.97	5.11	11.04	10.48	9.92	9.35	8.79	8.23	7.66	7.10
Standard deviation	21.64	12.01	3.79	20.14	18.70	17.34	16.07	14.91	13.90	13.08	12.46
Beta rel. to market	1.00	.30	.03	.93	.86	.79	.72	.65	.58	.51	.44
Sharpe Ratio	.300	.072	.000	.295	.287	.277	.264	.247	.224	.195	.160

Table 7: Regression [3] of monthly ISDs on historical volatilities, 1986-2006

Coefficient	$\lambda=0$	$\lambda=.2$	$\lambda=.3$	$\lambda=.4$	$\lambda=.44$	$\lambda=.5$	$\lambda=.6$	$\lambda=1$
a	9.17	7.45	6.65	5.89	5.59	5.17	4.51	43.89
(t-value)	(16.6)	(13.9)	(12.4)	(10.7)	(10.1)	(9.0)	(7.2)	(0.8)
b	0.60	0.70	0.75	0.79	0.81	0.83	0.87	-1.43
(t-value)	(21.7)	(25.4)	(26.8)	(27.6)	(27.7)	(27.5)	(26.3)	(-0.5)
Adj R ² %	65.1	71.8	74.0	75.1	75.2	74.9	73.2	0.00

Figure 1: How the stop-loss method worked in the US during 1929 and in the UK during 2001




Figure 2: Annualised standard deviation of daily equity returns for the US and UK, monthly 1900-2006

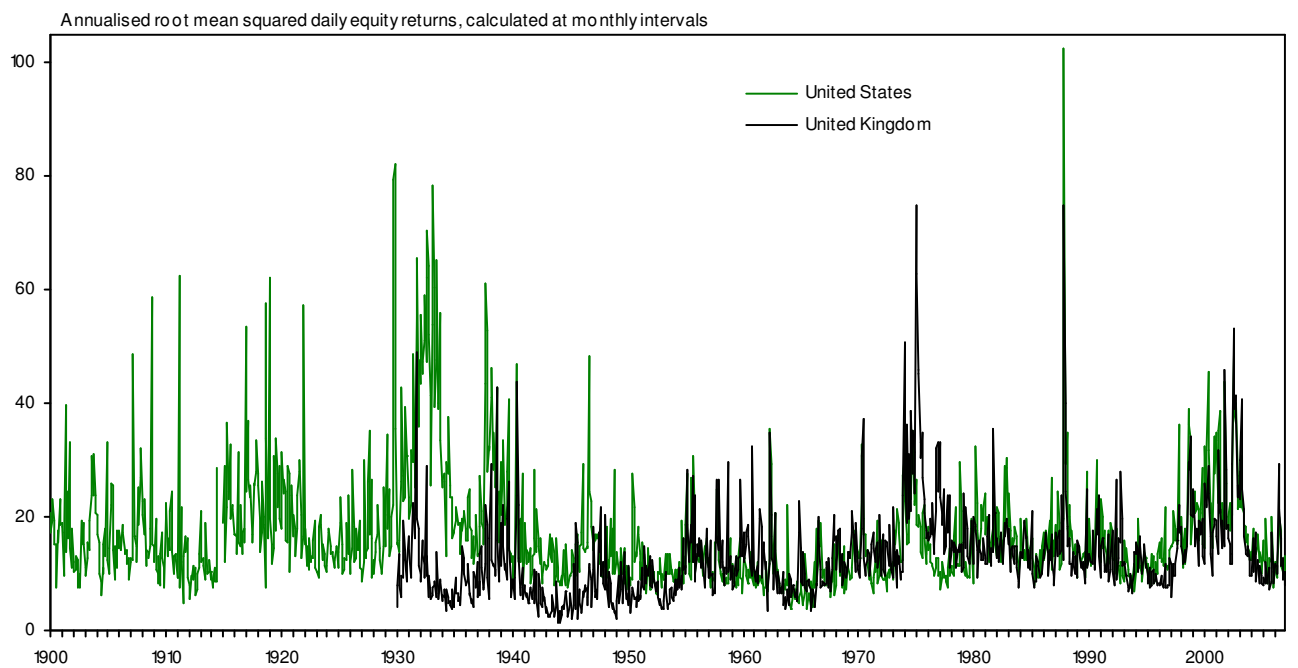


Figure 3: The VIX index of implied volatilities on the US equity market, 1986-2006


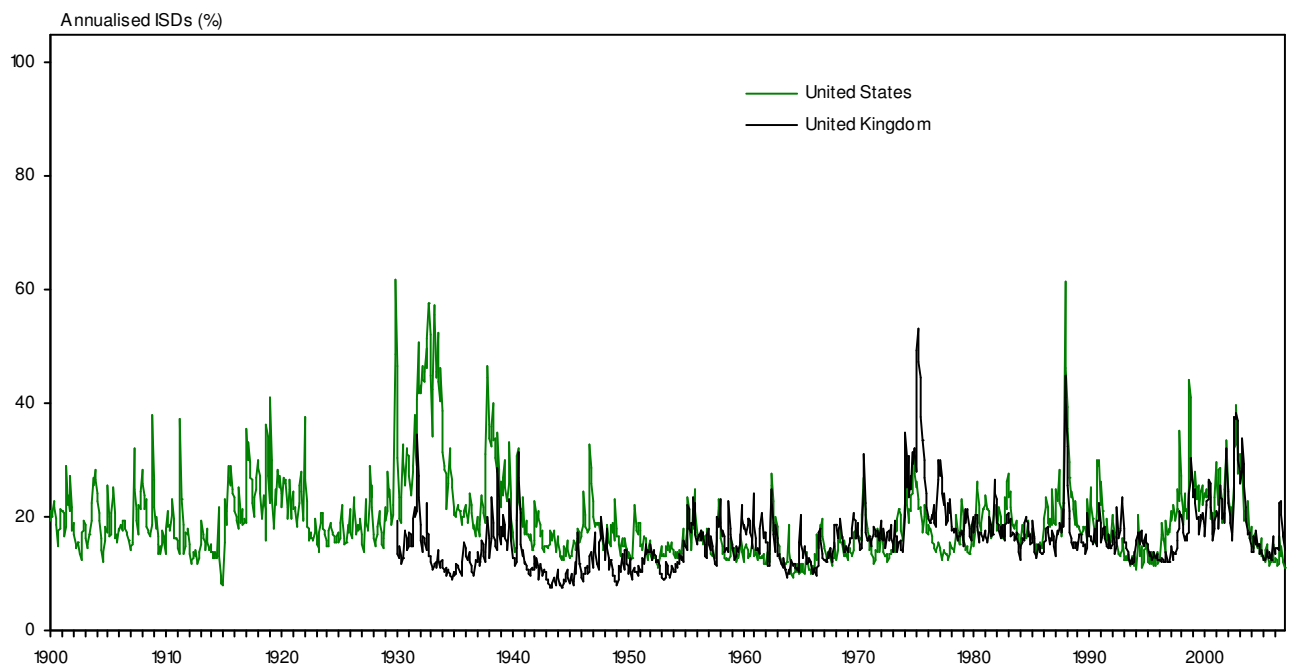


Figure 4: Estimated ISDs for the US and UK equity markets, monthly 1900-2006



Note: ISDs are estimated as described in the text and as detailed in the Appendix

Figure 5: Distribution of equity returns in the UK, annually 1900–2006




Figure 6: Distribution of truncated returns for UK, annually 1900–2006



Figure 7: Payoff diagram for a ‘costless collar’




Figure 8: Return vs standard deviation with equity portfolio protection over monthly horizons




Figure 9: Return vs standard deviation with equity portfolio protection over annual horizons




Figure 10: Rolling 60-month equity-bond correlation, US from 1900 and UK from 1930

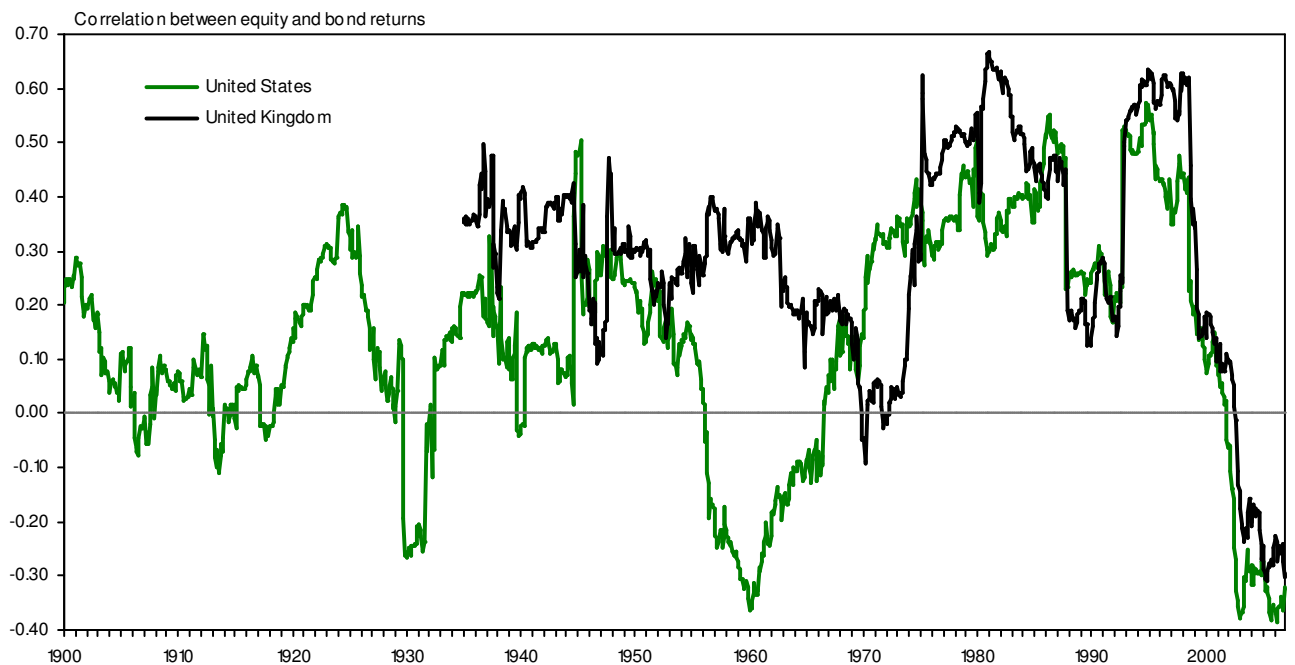


Figure 11: Time-plot of estimated ISDs based on regression [3], monthly 1986-2006

