Alpha Generation and Risk Smoothing using Managed Volatility

Tony Cooper

Quantitative Analyst

Double-Digit Numerics*

August 6, 2010

Abstract

It is difficult to predict stock market returns but relatively easy to predict market volatility. But volatility predictions don't easily translate into return predictions since the two are largely uncorrelated. We put forward a framework that produces a formula in which returns become a function of volatility and therefore become somewhat more predictable. We show that this strategy produces excess returns giving us the upside of leverage without the downside.

As a side-effect the strategy also smoothes out volatility variation over time, reduces the kurtosis of daily returns, reduces maximum drawdown, and gives us a dynamic timing signal for tilting asset allocations between conservative and aggressive assets.

It has been said that diversification is the only free lunch in investing. It appears that once you have diversified away some risk you can get a further free lunch by smoothing what risk remains.

Keywords and phrases: volatility timing, volatility of volatility, extreme volatility, volatility drag, managed volatility, continuously dynamic leverage, leveraged exchange traded funds, alpha generation, kurtosis, maximum drawdown

JEL Classification: C22, C5, G1

---

*`www.ddnum.com`. This paper differs from the NAAIM Wagner Award (`naaim.org/wagneraward.aspx`) winning paper in having a title change, minor text improvements, and a new section on maximum drawdown and kurtosis. Thanks to Aaron Brown and Florin Spinu for their related comments.

1 Introduction

Whether or not you believe in the efficient market hypothesis it is difficult to predict market returns (Ferson, Simin, and Sarkissian 2003) whereas market volatility is clearly forecastable (Poon and Granger 2003). Fischer Black (Miller 1999) reasoned that estimating variances is orders of magnitude easier than estimating expected returns. This conclusion does not violate market efficiency since accurate volatility forecasting is not in conflict with underlying asset and option prices being correct.

The purpose of this paper is to use the predictability of volatility to generate excess returns (returns over and above the market when adjusted for risk). For such predictions to be useful volatility itself must be volatile.

The remainder of this article is structured as follows. The next section presents the mathematics of compounding daily leveraged returns which gives rise to the formula that motivates the investment strategies. The next two sections discuss the concept of the volatility of volatility and introduce the strategies that exploit the concept. The following section forms the main body of the article where the econometric results are discussed. Finally, the last section summarizes the empirical findings, draws some conclusions, and points out avenues for future research.

All index data comes from Yahoo.com except for the US stock market data from 1885 to 1962 which comes from (Schwert 1990) (the capital index is used). We augment this dataset using the S&P 500 index to produce a price series up to 2009.

2 Simplified Mathematics of leverage

To introduce our key investment concept of volatility drag and to see how leverage works we start with an example involving leveraged exchange traded funds (ETFs).

2.1 Definition of Leveraged Exchange Traded Funds

Leveraged ETFs are funds that aim to magnify the daily moves of an index. For example in a double-leveraged fund (a 2x fund), if the index goes up, then the fund goes up twice that amount. The magnification multiple is specified in the prospectus for the ETF and is a fixed constant. Typical values are 1, 1.5, 2 and 3. Negative multiples such as -3, -2, and -1 are used as well. For example if an ETF promises a return of 2 times the S&P 500 index then if the S&P 500 index goes up 1.2% in one day the ETF will go up 2 x 1.2% = 2.4%.

The salient point about this definition is that the return is marked to the benchmark *daily*, not annually. This causes confusion for investors because the daily leverage multiple does not translate to the same multiple when applied to annual returns of the ETFs or to compound returns. The dynamics of leveraged ETFs are more complicated than just applying a multiple to a return and the correct formula provides the key to our volatility strategies.

2.2 Myth about Long Term Holding of Leveraged ETFs

Leveraged ETFs confuse many investors because of the difference between arithmetic and compound returns and because of the effects of volatility drag (explained below).

One of the common myths is:

> Leveraged ETFs are not suitable for long term buy and hold

This myth is expressed in various ways. Some quotes from the internet about leveraged ETFs:

> "unsuitable for buy-and-hold investing," "leveraged ETFs are bound to deteriorate," "over time the compounding will kill," "leveraged ETFs verge on insanity," "levered ETFs are toxic," "levered ETFs [are] a horrible idea," "...practically guarantees losses," "in the long run [investors] are almost sure to lose money," "anyone holding these funds for the long term is an uneducated lame-brain." "Leveraged ETFs are leaky," "Warning: Leveraged and Inverse ETFs Kill Portfolios."

There is even an article comparing these ETFs to swine flu.

The explanation popularly given for this myth is that volatility eats away at long term returns. If this were true then non-leveraged funds would also be unsuitable for buy and hold because they too suffer from volatility. We need to more closely examine the effect of volatility.

2.3 Volatility Drag

Daily volatility hurts the returns of leveraged ETFs (including those with leverage 1x). This is due to the equality

$$(1-x)(1+x) = 1 - x^2$$

Suppose the market goes down by x and then the next day it goes up by x. For example the market goes down by 5% then up by 5%. Then the net result is that the market has gone to (1-0.05) times (1+0.05) = 0.9975 which is a drop of 0.0025 or 0.25%.

That doesn't seem fair. The market has gone down by 5% then up by 5% but our ETF that has a leverage of 1x has gone *down* by 0.25%.

This drop always occurs because x^2 is always positive and the sign in front is negative. So whenever the market has volatility we lose money. We call this loss *volatility drag*.

The larger x is, the larger x^2 is, so the larger the volatility drag. For a leveraged ETF the leverage multiplies x and so multiplies the volatility drag. Even an ETF with a leverage of 1x has volatility drag.

The myth has resulted from the belief that volatility drag will drag any leveraged ETF down to zero given enough time. But we know that leverage of 1x (i.e. no leverage) is safe to hold forever even though leverage 1x still has volatility drag. If 1 times leverage is safe then is 1.01 times leverage safe? Is 1.1 times safe? What's so special about 2 times? Where are you going to draw the line between safe and unsafe?

Note that even if stocks fluctuate logarithmically rather than arithmetically (so that changes of $(1-x)$ and $(1+x)$ are not equally likely) we still get volatility drag because it results from volatility *per se* rather than the nature of the volatility.

2.4 Sample Market Data

Figure 1 shows 135 years worth of daily US index prices going back to 16 February 1885. The construction of this index is described in Schwert (1990) and the index used is the capital index (no dividends reinvested). S&P 500 data from Yahoo.com has been used to augment the data up to 2009.

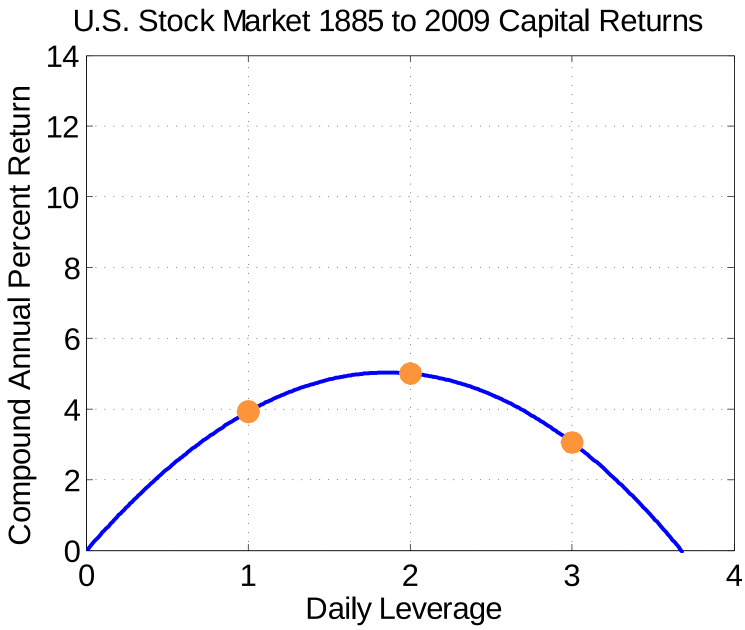

Figure 1: US Stock Market Returns for a Range of Leverages

The orange circles show popular leverage rates 1, 2, 3, and, where visible, 4. It can be seen that increasing leverage from zero to 1 increases the annualised return as would be expected. But then, contrary to what the myth propagators say, increasing the leverage even further still keeps increasing the returns.

There is nothing magic about the leverage value 1. There is no mathematical reason for returns to suddenly level off at that leverage. We can see that returns do drop off once leverage reaches about 2. That is the effect of volatility drag.

What the myth propagators have forgotten is that there are two factors that decide leveraged ETF returns: benchmark returns and benchmark volatility. If the benchmark has a positive return then leveraged exposure to it is good and compensates for volatility drag. Since the return is a multiple of leverage and the drag a multiple of the leverage squared then eventually the drag overwhelms the extra return obtained through leverage. So there is a limit to the amount of leverage that can be used.

The tradeoff between return acceleration and volatility drag and the management of the tradeoff is the basis for our volatility strategies.

2.5 The Formula for the ETF Long Term Return

The formula for the long term compound annual growth rate of a leveraged ETF cannot be written in terms of just the benchmark return and volatility. It also involves terms containing the skewness and kurtosis of the benchmark. It is derived using a Taylor series expansion (details available from the author). It does not assume that benchmark returns are Gaussian or that returns are continuous as do formulae derived using Ito's lemma.

But it turns out that for the world's stock markets and for low levels of leverage (up to about 3) the formula can be approximated by this formula:

$$R = k\mu - 0.5k^2\sigma^2/(1 + k\mu) \tag{1}$$

where R is the compound daily growth rate of the ETF, k is the ETF leverage (not necessarily an integer or positive), μ is the mean daily return of the benchmark, and σ is the daily volatility (i.e. standard deviation) of the daily return of the benchmark.

R is the quantity you use to calculate the long term buy-and-hold return of the ETF. You can see from the formula that if the volatility is zero then $R = k\mu$ so that the return of the ETF is k times the return of the benchmark. The $0.5k^2\sigma^2/(1 + k\mu)$ term is the volatility drag. Since $k^2\mu^2$ is always positive and $(1 + k\mu)$ is always close to 1 then the volatility drag is always positive.

R is a quadratic function of k with a negative coefficient for the square term. That means we will always get the parabola shape shown in Figure 1 and we will always have a maximum for some value of k. Some algebra shows that the maximum is approximately (for small $k\mu$) at

$$k = \mu/\sigma^2 \tag{2}$$

This clearly shows the return/volatility trade-off that determines the optimal leverage.

This formula occurs in an appropriate form in the Kelly Criterion (Thorp 2006) and Merton's Portfolio Problem (Merton 1969). Its appearance here as the result of an optimisation is no surprise.

2.6 A Look at Some Stock Markets

Let's look at some other markets and time frames in Figure 2.

The pattern is quite clear. Over various markets over various time periods (mostly the last two or so decades) except for the Nikkei 225 the optimal leverage is about 2.

2.7 Leveraged ETFs Fees

Unfortunately there may be a reason not to hold leveraged ETFs for the long term but it has nothing to do with volatility drag. It is because of fees.

Most leveraged ETFs where the leverage is greater than one charge an annual fee of about one percent. This imposes a "fee drag" on the ETF.

Also leveraged ETFs suffer from tracking error. They do not exactly hit their target return for the day every day. From studying a few leveraged ETFs it has been observed that the tracking error is a complicated error and beyond the scope of this document. But tracking error is small and is ignored for this article.

An annual fee of 0.95% (a typical value for recent leveraged ETFs) applied daily subtracts 0.95% from the annual return of the ETF. For the 1885 to 2009 data this fee reduces the return of a 2x leveraged ETF to that of a 1x fee-free ETF. So all the benefits of leveraging have been lost.

For other markets over shorter time frames the fees aren't as destructive. The optimal leverage is still about 2 and even after fees 2x ETFs outperform the benchmark over several decades.

2.8 Leveraged ETFs Risks

Leveraged ETFs can be held long term provided the market has enough return to overcome volatility drag. It usually does. For most markets in recent times the optimal leverage is about 2. But some markets and time frames will reward a leverage of up to 3. No markets will reward a leverage of 4. So it would seem that using leveraged ETFs is a good idea for enhancing returns. But this is not true for the following reasons:

- Leveraged ETFs do not generate alpha. Any leverage that multiplies return also multiplies

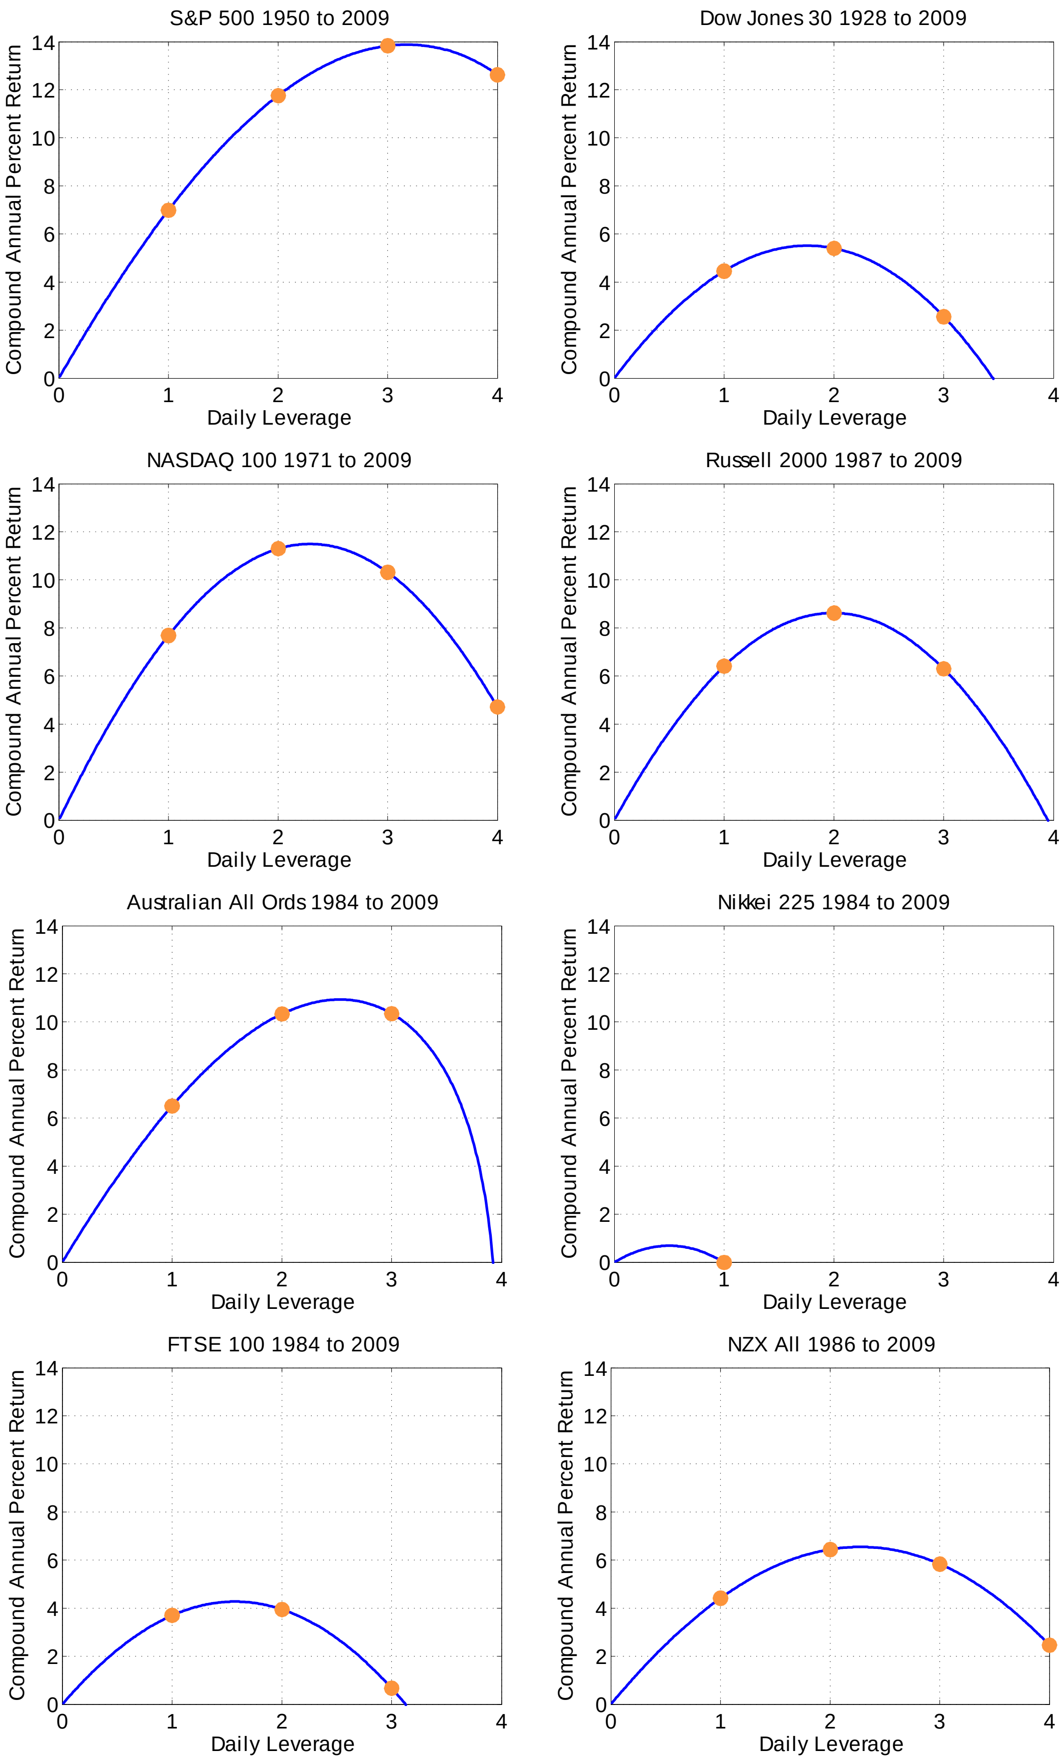

Figure 2: Returns vs Leverage for a Rangle of Markets and Time Periods

volatility by the same multiple. So risk-adjusted returns are not enhanced.

- Nearing the peak of the parabolas from the left we are adding leverage and thus volatility and getting not much extra compounded return for it. And going past the peak we are adding volatility and losing return. This is dangerous and since we do not know in future markets where the peak is exactly we run the risk of negative alpha.
- Leveraged ETFs run the risk of ruin—losing all money invested.

So to use leveraged ETFs effectively we want to invest as close to the left as possible and to closely watch volatility to avoid straying over to the right. This is where the volatility of volatility comes into play.

3 Volatility and Volatility of Volatility

The optimal leverage $k = \mu/\sigma^2$ gives the optimal value of R. For the charts of the markets in Figure 2 μ and σ are the mean values over the lifetime of the ETF. But in reality μ and σ are not regarded as constant but reasonably can be seen to vary over time. So it seems likely that the optimal value of k is time varying also. Can we predict μ and σ well enough to estimate the optimal value of k for the immediate future?

We cannot easily predict μ and it is outside the scope of this volatility paper anyway.

This is a dynamic asset allocation problem where we want to maximise some measure of terminal wealth with decisions made at discrete time intervals such as days or months. This is a dynamic stochastic programming problem (Samuelson 1969) but it is outside the scope of this paper to provide an optimal solution. Instead we provide an indication of the possible gains from solving the problem.

3.1 EGARCH(1, 1) Volatility Forecasting

σ, the daily volatility, is a random variable which can vary from day to day. It is defined to be the instantaneous standard deviation of price fluctuations. Measuring volatility is not as straightforward as measuring price. It requires measuring price changes over time but during the time interval measured the volatility itself may change. For example, if only daily closing prices are observable then to get a reasonable measure of standard deviation 10 prices over 10 days need to be observed. But over those 10 days the volatility itself is sure to have changed. So instantaneous volatility is unobservable. Instantaneous volatility can only be estimated in the context of a model.

We shall use the EGARCH(1, 1) model for volatility forecasting because it is a relatively standard choice and Poon and Granger (2003) report studies where it outperforms other methods, particularly when used on market indexes. It is outside the scope of this expository paper to find the best method for our purposes so an indicative method will suffice. The EGARCH(1, 1) model is

$$\log(\sigma_t^2) = m(1-\alpha-\beta) + \beta\log(\sigma_{t-1}^2) + \alpha\left|\frac{\varepsilon_{t-1}}{\sigma_{t-1}}\right| + \gamma\frac{\varepsilon_{t-1}}{\sigma_{t-1}} \tag{3}$$

This models the log variance on day t, $\log(\sigma_t^2)$, as a weighted average of the long run log variance m, the previous day's log variance $\log(\sigma_{t-1}^2)$, and the absolute value of the previous day's return ε_{t-1} (scaled for consistency by the previous day's volatility σ_{t-1}). The γ term is a correction[1] term that allows for the fact that negative daily returns tend to add positive amounts to the daily volatility. After estimating m, α, β,and γ we can estimate volatility on day t from quantities calculated on day $t-1$.

Even though volatility is not a directly observable quantity this method gives us estimates of volatility for days $1, \ldots, t$. But, further, we can plug in values calculated on day t to forecast the volatility for day $t+1$.

Figure 3 shows the EGARCH estimates for the S&P 500 daily volatility from 1950 to 2009. The left hand charts show the volatility estimates, the right the estimates on a log scale. It is apparent that logs of the volatilities have more symmetry than the raw volatilities; they also are less "spiky." This gives us justification for the use of the log estimator in Equation (3) over an estimator that estimates σ_t^2 directly.

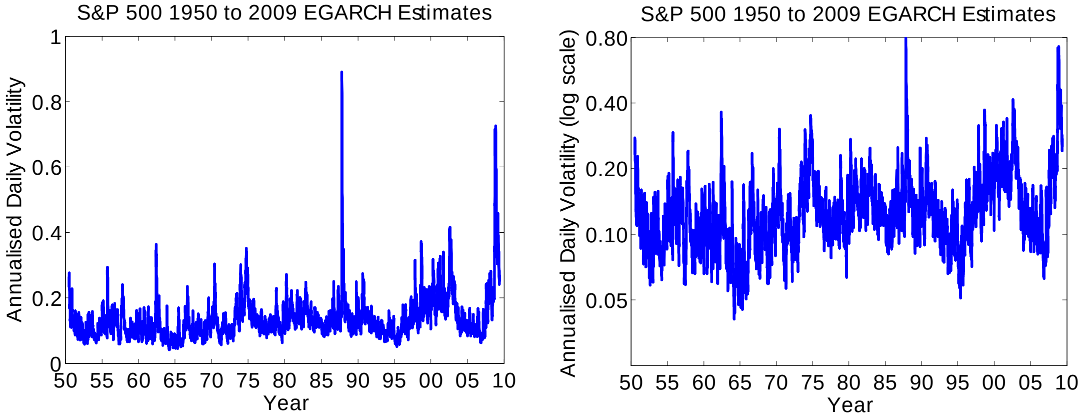

Figure 3: EGARCH Estimates of Daily Volatility

3.2 Volatility of Volatility

A feature of the charts in Figure 3 is that the estimates of volatility vary over time. The most likely reason is that the actual volatility being estimated varies over time. Let us define the volatility of volatility (vovo for short) as the standard deviation of the daily volatilities σ_t for $t = 1, \ldots, n$ where n represents the number of days of observation.

Then for the purposes of this paper we shall estimate vovo by the intuitively reasonable standard deviation of the EGARCH(1, 1) estimates of the daily volatility. It is beyond the scope of this exploratory paper to examine the statistical properties of this estimator. For our purposes it suffices to know that volatility does vary and that our EGARCH model gives us estimates of the variable volatility.

We don't take the standard deviation of the daily volatilities, however. We deal with log volatility where possible and calculate vovo on the log scale.

3.3 Costs of Volatility of Volatility

In traditional asset allocation vovo has a cost. As an example, consider a balanced-fund manager who uses, say, bonds to reduce the volatility of a fund which contains equities. Or a financial planner who assesses the risk tolerance of a client investor and proposes, say, a 60-40 equities-bonds mix. The traditional way of managing these investments is to look at the long term historical volatility of the component asset classes to decide the proportions to invest in. Usually no consideration is given to the vovo of the asset classes and no constant volatility target is set.

[1]The γ term is frequently called the "leverage" term but we won't use that volatility term here because it could be confusing.

So the investor or manager with a static or reasonably constant 60-40 mix has to suffer the varying volatility of the markets—sometimes sleeping well at night, occasionally not. In order to ensure that the worst volatility is minimised the asset allocation will have erred on the side of conservative. This will have cost returns.

The larger the vovo the more conservative the asset allocation will have to be and therefore the more the cost to returns.

In addition to the loss of returns there is a cost to the fund manager due to the extra expected drawdown of the fund arising from vovo. Drawdowns cause worried investors to redeem funds so funds under management will drop. If we have two fund managers with the same annual returns and same annual volatility the manager with the higher vovo will lose the most funds in a market drawdown.

3.4 Volatility of Volatility Strategies

We now create three portfolio strategies that use volatility prediction and the μ/σ^2 formula in a short-term look-ahead manner. The strategies are:

- The Constant Volatility Strategy (CVS)
- The Optimal Volatility Strategy (OVS)
- Mean Only Strategy (MOS)
- The Optimal Volatility Plus Mean Strategy (OVPMS)

The MOS is not a strategy that we consider for investment because it is included only as a benchmark for the the OVPMS strategy.

We also construct a data-snooping version of OVPMS called OVPMS+. The purpose of this version is to give an indication of how well a calibrated version of OVPMS would have done.

Every day just before the close of trading we take the closing price for the day and put it into our EGARCH estimator to estimate our volatility for the next day. Call our estimate of the ith day's volatility s_i. Then we calculate our leverage k_i for the ith day using our leverage calculation formula. Then just on the closing bell we execute trades to bring our portfolio to leverage k_i so that it begins the next day at this level of leverage. The fact that the EGARCH estimation and ensuing trading have to be done in the instant just before closing and that there are transaction costs are implementation inconveniences that we overlook for the purposes of this expository study. Also, in practice, we would likely be using intra-day prices for estimates of the volatility rather than end of day prices since they can provide better predictions (Poon and Granger 2003).

These strategies are not pattern-detection market timing strategies or technical analysis. The strategies are not obtained by backtesting. They are purely mathematical day-ahead optimisation strategies that use past data only (OVPMS+ the exception) for volatility prediction.

3.5 Leveraged ETFs and the Probability of Ruin

The strategies can be implemented as overlays on existing portfolios or as investments in their own right. They can be implemented by using combinations of leveraged ETFs provided that suitable ETFs exist. For example suppose that leverage of 2.4 is required. This can be achieved by investing 0.6 of the asset into a 2x leveraged ETF and 0.4 into a 3x leveraged ETF.

A typical range of leverages such as -3, -2, -1, 1, 2, and 3 may exist for any one index. That gives a range of possibilities for mixing them. For example 2.4 = 0.6x2 + 0.3x3 but also 2.4 = 0.3x1 + 0.7x3 and also 2.4 = 0.15x(-1) + 0.85x3. More than two ETFs may be used. Which is the best to choose? One possible answer considers the possibility of the ETF being ruined.

An ETF with leverage k will drop to zero if the market drops by $1/k$. So in the 2.4 = 0.6x2 + 0.3x3 case the strategy will be ruined if the index drops by 50% or more in one day. But the 2.4 = 0.15x(-1) + 0.85x3 case will still have 0.15x(0.5) left.

The probability of ruin when invested in an index is virtually zero. But a leveraged ETF where the absolute value of the leverage exceeds one will always have a finite chance of ruin. The 1987 crash would have ruined any 5x leveraged ETF. Fortunately all the strategies reduce leverage when volatility is high so this makes the chances of ruin smaller. We examine the behaviour of the strategies at the 1987 crash below. Non-market catastrophic events that hit the market without warning would be the most deadly risk.

4 Volatility of Volatility Strategies

Our three strategies differ in the formula used to derive the daily values of the k_i.

4.1 The Constant Volatility Strategy (CVS)

This strategy sets the leverage to be $k_i = c/s_i$ for some constant c. This strategy is not optimal in that it uses s_i in the denominator rather than s_i^2. But it has a simplicity and appeal that we find irresistible. If tomorrow is predicted to be a high volatility day then it lowers the leverage for that day in proportion to the predicted volatility. This means that it is aiming for a constant volatility of c every day.

This has four useful implications for financial planning and risk management:

- it provides a target risk level that can be sold as a number (and put into the prospectus)
- it reduces the vovo experienced by the strategy.
- it increases the portfolio return by reducing the costs of vovo as described in subsection 3.3
- it generates alpha.

4.2 The Optimal Volatility Strategy (OVS)

This strategy sets the leverage to be $k_i = c/s_i^2$ for some constant c. The problem with this strategy is that c cannot be easily set in advance to target a specific volatility as it can with CVS. For a given prediction s_i and leverage c/s_i^2 the predicted strategy volatility of the strategy for day i is c/s_i rather than the constant c for CVS. So OVS is not targeting constant volatility but targeting extra low volatility the higher predicted volatility is.

OVS is not trying to reduce vovo so even though it is optimal it will lose some of its desirableness in a management setting when the cost of vovo is taken into account. And a further implementation problem arises in that in targeting a mean level of volatility the value of c required is not known in advance. It can be calibrated by examining historical data but this introduces estimation error and the assumption that the future will behave as the past.

A further problem with OVS that we will see below is that vovo introduces extra volatility into the k_i values which sometimes get very high. Higher k_i values have higher probabilities of portfolio ruin.

4.3 The Mean Only Strategy (MOS)

This strategy sets the leverage to be $k_i = cm(s_i)$ for some constant c and some function $m(s)$ which is described more fully in the next strategy. We are not interested in the MOS strategy in its own right—just as a means of looking at the effect of the μ term contribution in the expression μ/σ^2. That is, this strategy acts as a benchmark.

4.4 The Optimal Volatility Plus Mean Strategy (OVPMS, OVPMS+)

This strategy sets the leverage to be $k_i = cm(s_i)/s_i^2$ for some constant c and some function $m(s)$. Here we have some estimate of the mean return *as a function of volatility only.* Any other prediction is outside the scope of this volatility paper.

The difference between OVPMS and OVPMS+ is that the former only uses past data to estimate the function $m(s)$ and every day when the EGARCH forecast is calculated the estimate of $m(s)$ is updated as well. OVPMS+ uses only one estimate of $m(s)$ and that is calculated from the entire date range. Therefore, for example, it uses in 1951 an estimate of $m(s)$ that was calculated from data up to 2009. The reason for doing this is to give us an idea of the performance of this strategy from 2009 onwards by seeing how well the 2009 parameters did when applied from 1950 onwards. This is a valid use of data-snooping in a calibration context but the usual caveats of the future resembling the past apply.

4.5 Estimating Return as a Function of Volatility

For OVPMS we require an estimate of market return as a function of volatility $m(s)$. More specifically, we want a prediction of market return as a function of our *prediction* of market volatility.

The left hand chart of Figure 4 shows for the S&P 500 index from 1950 to 2009 a plot of the market return versus our predicted EGARCH(1, 1) values for 14823 days. The daily returns d_i are annualised by scaling to $(1 + d_i)^{252}$ and the daily predicted volatility is annualised by

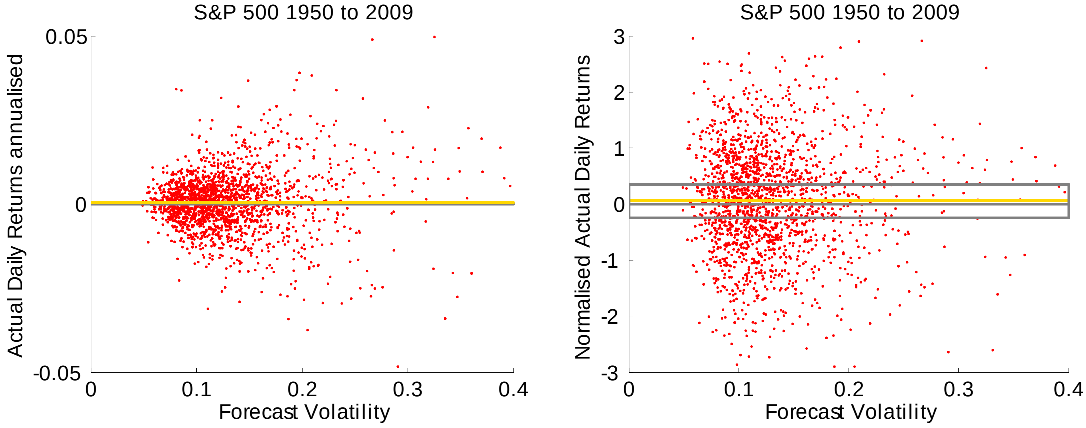

Figure 4: Actual Return versus Forecast Volatility

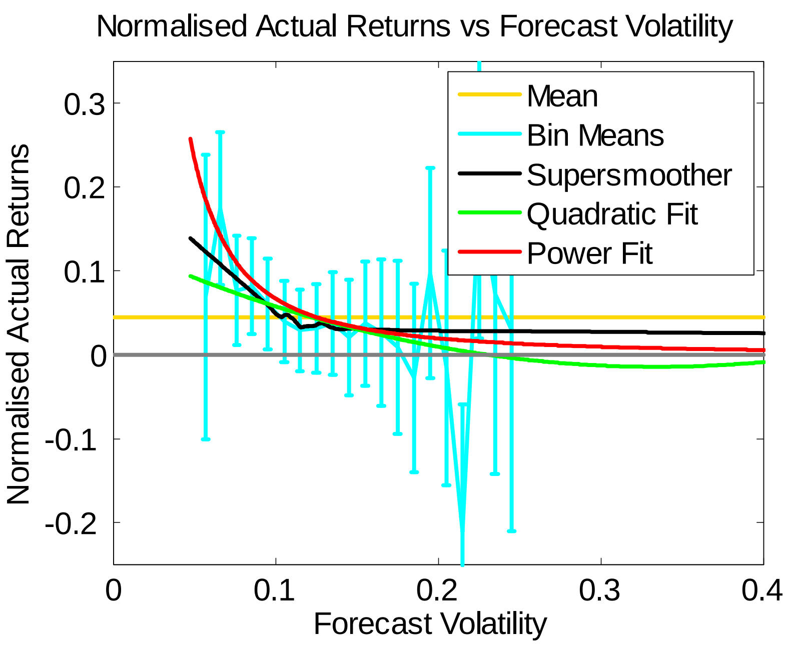

Figure 5: Normalised Actual Returns versus Forecast Volatility

multiplying by $\sqrt{252}$. Some extreme points such as the 1987 crash of over 20% have been omitted to improve readability.

At first glance, no obvious relationship between return and volatility can be seen. The correlation between the two is 0.005557 (not significantly different from zero, P = 0.50 significance level). Let's make some changes to the figure to get the right hand chart. Noticing that the spread of points about the horizontal line is expanding as we move left to right we aim for a uniform spread by normalising the daily returns by dividing them by the predicted volatility. This produces a uniform spread which allows us to fit curves using unweighted least squares regression. Now we zoom in on the chart into the rectangle shown in the chart. We get Figure 5. We have divided the interval of forecast volatility from 0.05 to 0.25 (which is where 96% of the points lie) into 20 equally spaced bins and calculate the mean normalised return for each bin. These means are plotted in cyan along with two standard deviation error bars. We see convincing evidence that returns decrease with increasing volatility.

In an effort to get an idea of the relationship between returns and volatility amongst what is almost all noise we apply some smoothing techniques to all the points, not just the ones shown in the figure. The gold line is the overall mean. The black line is an application of Friedman and Stuetzle's (1982) *supersmoother* which uses local linear fits. This smooths the data without imposing any constraints onto it other than smoothness so is the most truthful smooth. But it does not give us a clean functional form. So we try some parametric functions.

The green line is a quadratic fit. This imposes a turning point onto the relationship and allows the scaled return to go below zero. The red line a robust fit of a power function $f(s) = as^b$. This imposes positivity onto the relationship and has no requirement for a turning point. The parametric fits are sensitive to the constraints we impose and we have no justification for imposing them.

But we have to start somewhere and the final portfolio may not be sensitive to the constraints anyway. So we focus on the power function. This function seems to capture the relationship quite well and it multiplies well with Formula 2 so we use this function for all further work. Table 1 shows the estimates for b and the 95% confidence intervals for these estimates.

The power function fit to the scaled returns is $m/s = as^b$ or $m = as^{b+1}$ so our OVPMS strategy becomes $k_i = cm(s_i)/s_i^2 = cas_i^{b-1}$ and we drop the a by absorbing it into the constant c. Thus

$$k_i = cs_i^{b-1}$$

where for OVPMS we update our estimate of b daily and for OVPMS+ we use a single estimate for the whole period.

4.6 Another look at Return as a Function of Volatility

Return as a smooth function of volatility may not be the way the market works. It could be that there is a threshold of volatility where the returns jump to zero or negative above that threshold level. There is evidence that such a threshold exists for all the markets studied and that it has a similar value in all cases. But estimating that value is a statistical exercise

Index	Time Period	Annual Volatility	b	95% Confidence Interval for b
US Stocks	1885–2009	0.1672	-1.69	(-2.28, -1.09)
S&P 500	1950–2009	0.1531	-1.76	(-2.63, -0.90)
Dow Jones Industrial 30	1928–2009	0.1846	-1.52	(-2.27, -0.76)
NASDAQ 100	1971–2009	0.2020	-1.71	(-2.24, -1.19)
Russell 2000	1987–2009	0.2024	-1.82	(-2.60, -1.03)
Australian All Ords	1984–2009	0.1584	-1.87	(-2.79, -0.94)
Nikkei 225	1984–2009	0.2349	-2.49	(-4.05, -0.92)
FTSE 100	1984–2009	0.1833	-3.10	(-4.50, -1.71)
NZX All	1986–2009	0.1529	-3.18	(-4.43, -1.93)
NZX 10	1988–2009	0.1783	-2.80	(-4.93, -0.66)

Table 1: Indexes, Volatilities, and Power Function Exponents

(estimating peaks in noisy time-series signals) that is outside the scope of this paper.

4.7 Backtesting Strategies

Backtesting refers to the use of past data to:

- test a hypothesis
- estimate how a strategy worked in the past *as a guide to how it will work in the future.*

In this study we use backtesting for both purposes. Our main hypothesis that we are testing is that returns can be increased and risk reduced by using our leverage formulae. This is a valid and reasonably safe use of backtesting.

The problems come when you use the data to test hypotheses suggested by the data itself. This dual use of the data produces biases which are explained below.

The problems with the second use of backtesting are more straightforward. To use the past as a guide to the future assumes that the future will behave as the past. There is a risk that it does not. We cannot quantify this risk or correct for it and simply take it as given that any predictions that our models make in this paper are subject to this risk.

For the purposes of this paper we classify the backtesting biases into three categories based on how dangerous they are when used to predict the future:

- Data snooping bias – this occurs when we use data in the past that was obtained from future data. This bias can be small when we use future data to calibrate a model and then use calibrated parameters in the past. An example is when we test an asset allocation strategy that requires a measure of the volatility of the market. We can estimate the volatility from 50 years worth of data then go back in time and use the strategy as though we knew that estimate right at the beginning. Provided that the estimate is reasonably close to the actual future volatility this bias may be small. For example, our estimate of the market volatility may turn out to be too high or too low with roughly equal probability and this may make our estimate of future returns too high or too low with equal probability. The equality of the probabilities means that the bias is low. And in any case it does give us an idea of the potential returns of an accurately calibrated model.

- Data mining bias – this occurs when we use data in the past that we obtained from future data where we expect some *regression to the mean* in future excess returns. This bias usually occurs when we optimise some parameter in our strategy. For example, we might have a strategy based on a moving average of period λ. We use an optimising algorithm to find the value of λ that gives the best excess returns. In that case we expect considerable bias in backtesting because the true value of λ in the future will be different than our estimated value and our actual excess returns will be less than our estimated value no matter whether our estimate was too high or too low. Our actual excess returns will regress towards zero. Hopefully not all the way to zero—our optimal value of λ will hopefully be in the neighbourhood of the true value—and so data mining can be useful even though biased.

  Another situation that creates data mining bias is testing a hypothesis where the data itself suggested the hypothesis. The correlation between the test and the suggestion means that the significance levels of the hypothesis test are really not as good as the test statistics suggest. This bias arises frequently when the data is "mined" for ideas which are then tested using the same data. We avoid this bias in this study by formulating our equations and strategies before we look at the data and only use the data for calibration not idea generation.

- Data torturing bias – this occurs when we "torture" the data into revealing information it does not actually contain. In this case we expect excess returns to regress all the way to zero. An example is when we notice that Friday the 13ths have a larger mean excess return than any other day so our strategy is to only invest on Friday the 13. Since the outperformance of that day is (presumably) just due to chance then the strategy will produce zero excess returns in future.

The expression comes from the apothegm amongst statisticians that "if you torture a dataset long enough it will confess to anything."

5 Volatility of Volatility Strategies Performance

All the vovo strategies are evaluated in a return-risk framework. We use three different risk formulations for this:

- The traditional risk – the standard deviation of daily returns (SDDR)
- The extreme risk – vovo-oriented risk – the standard deviation of daily returns plus two times the vovo (SDDR+2)
- The maximum drawdown – the maximum peak to trough loss experienced by the strategy equity curve

The first risk measure is traditionally the one used in portfolio risk-adjusted measures of performance such as the Sharpe Ratio (Sharpe 1966). But in the presence of vovo the standard deviation of daily returns gives a conservative measure of risk in that it averages volatility over the whole period. For example if volatility moves between 0.10 and 0.20 then SDDR will give a value around 0.15. In contrast SDDR+2 will give a value closer to 0.20 thus reflecting the worst volatility during the investment period.

We call the SDDR+2 value the *extreme volatility* value. It is roughly the volatility value that is exceeded only 5% of the time[2].

We tested all the strategies on the full list of indexes in Table 1 but for brevity we only show here the results for the S&P 500 for 1950–2009. The patterns in the results are similar and consistent across all markets and all indexes except for the magnitudes of the returns. In particular the strategies show the same relative rankings.

[2] It deals with the volatility of volatility and should not be confused with the concept of value-at-risk which deals with the volatility of returns

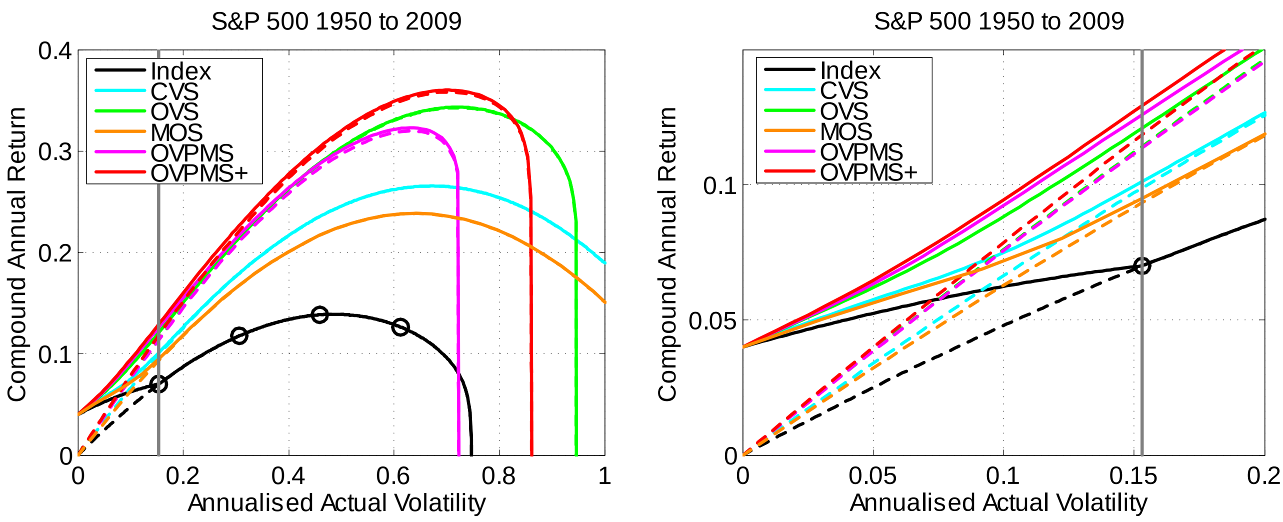

Figure 6: Return versus Volatility for the Strategies

5.1 Traditional Risk Performance

Figure 6 (right hand chart is a zoomed-in version) shows the returns of the strategies versus the annualised daily standard deviation (SDDR) for the S&P 500 from 1950 to 2009. The curves are traced out as a function of increasing volatility by increasing the constant scaling factor c of Section 4. The vertical grey line is the mean volatility of the index over the time period. For each strategy two versions are shown: the solid line is that strategy that invests in the risk free return for the balance of any assets not invested in the market when the leverage is less than one. The dotted line assumes that when leverage is less than one any uninvested funds return zero. (Some dotted lines are hard to see but they all converge onto the same point at zero). The actual value of the risk free rate of return (in this case 4%) can be seen as the intercept of the solid curves when the leverage is zero. We have assumed for simplicity that the risk free rate has been constant over the period. The effect of this investment is to raise the returns a little when leverage is low—the curve is twisted up near zero.

The black curve is a leveraged investment in the index itself as would be produced by a leveraged ETF without fees. The black circles show integer values for the leverage. So the black circle that intersects the grey vertical line is a 1x leveraged ETF. The next black circle is 2x etc. Up to 4x leverage is shown.

The effect of volatility drag is evident on all the strategies including the leveraged ETFs. All the curves have a peak and then drop down to (not shown on these charts) a return of -1. The sharp dropoff down to -1 is a result of the 1987 crash where the S&P 500 had a return of

-20%. Any strategy having a leverage of 5 or more lost *all* their assets on that day. Due to the one-off nature of the crash and to the high volatility of the strategies on the right hand side of the chart we shall confine our interest to the zoomed-in area of the chart near the vertical grey line which is a realistic level of volatility for a fund manager to consider.

We choose the value of c that gives the strategy the same volatility as the index volatility. We call this the 1x instance of the strategy. Here 1x refers to a multiple of the index volatility rather than the index return[3]. The strategy returns can be read off the chart and are given in table 3. All the strategies beat the index but, more importantly, met our expectations. CVS did well but not as well as OVS and this was expected because OVS used the exponent 2 in the optimum Formula (2) whereas CVS only used an exponent 1. OVPMS did better than OVS because it used the μ term in the formula whereas OVS set $\mu = 1$. And OVPMS+ did better than OVPMS because it used data-snooping to optimise the formula $m(s_i)$ used to estimate μ as a function of σ. OVPMS+ gives us an idea of how well OVPMS will perform in the future with a calibrated $m(s_i)$ function. It only provides a small improvement over OVPMS which suggests that the daily dynamic updating of the $m(s_i)$ worked quite well as an adaptive method of estimating the optimal formula.

MOS uses just the μ part of Formula (2) and did less well than any of the strategies that used σ. This indicates that our excess returns over the index did not come just from the fact that returns are high when volatility is low. Further returns came from the denominator of the formula. For all our strategies and for all the markets studied the correlation between daily returns and leverages is about 0.005 to 0.02. This is very small—too small to be noticed if plotted on a chart and it is not statistically different from zero (P = 0.50 to 0.10 typically). But it is large enough to create a significant timing boost to returns.

We have charts for all the other markets and for brevity they are omitted here. The results are summarised in Table 2 and are consistent over all the markets and time periods studied. Some special bear, sideways, and bull market time periods are considered below.

Figure 7 shows for the strategies CVS 1x, OVS 1x, MOS 1x, and OVPMS+ 1x the daily leverages used. We see that only CVS keeps maximum leverage down below 3 most of the time. The other strategies concern us a little because we start getting nervous when volatility

[3]For a leveraged ETF the designations 1x, 2x, ... with equal validity could refer to a multiple of the return *or* the volatility

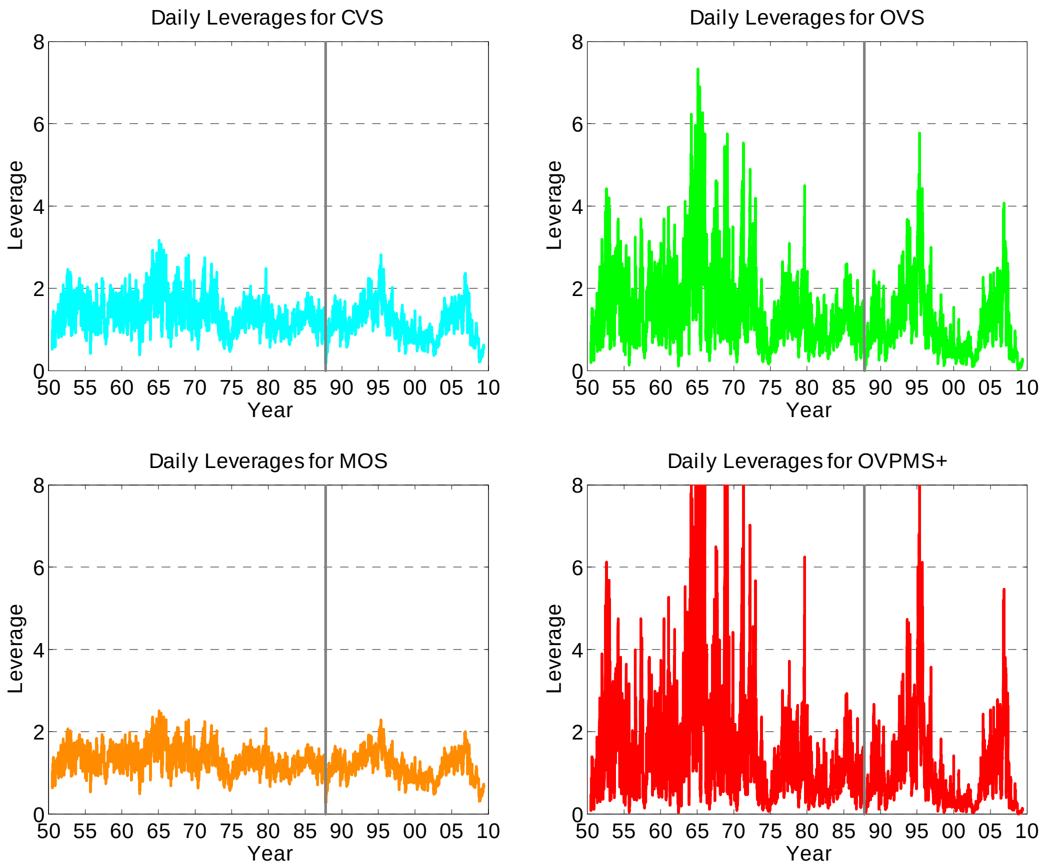

Figure 7: Daily Leverages used for the 1x versions of the Strategies

exceeds that level because of the encroaching risk of ruin. The OVPMS+ strategy actually gets leverage as high as 12.5 which must surely qualify as dangerous because an 8% drop in the market would ruin the investment. The grey vertical line shows the 1987 crash and we can easily see that the strategies anticipated the event by reducing leverage. This is discussed further below.

Figure 8 shows, on a log scale, the equity curve for the strategies and the equity curves for the S&P 500 at 1x, 2x, 3x, 4x. The strategies do not do as well at times as the 3x and 4x leveraged ETFs but they all have volatility 1x whereas the ETFs have volatility 3x and 4x. The ETFs also suffer severely at the 1987 crash.

Table 2 shows the performance of the strategies for the same range of markets and time frames as Table 1. As a rough rule of thumb the CVS strategy seems to return about twice as much as the market and OVPMS+ three times. All the strategies are the instances which produce the same volatility as the market (1x instances). This table provides compelling evidence that the strategies have worked very well in the past. Whether they will continue to do so in the future is an open question.

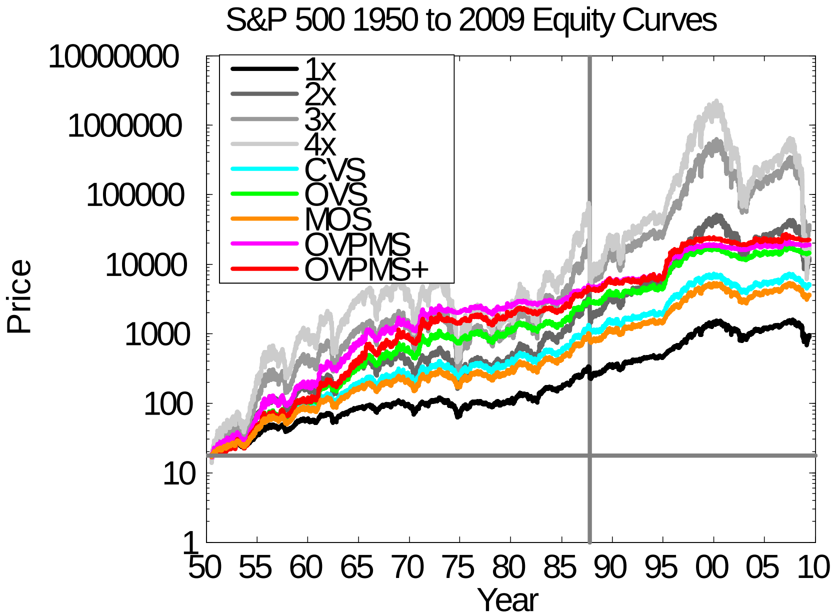

Figure 8: Equity Curves for Leveraged ETFs and the Strategies

Index	**Annual Return**	**CVS Return**	**OVS Return**	**MOS Return**	**OVPMS Return**	**OVPMS+ Return**
US Stocks	0.0393	0.0647	0.0837	0.0572	0.0953	0.0917
S&P 500	0.0699	0.1012	0.1209	0.0950	0.1260	0.1292
Dow Jones Industrial 30	0.0446	0.0855	0.1046	0.0678	0.0904	0.1071
NASDAQ 100	0.0769	0.1883	0.2540	0.1708	0.2606	0.2681
Russell 2000	0.0640	0.1393	0.1920	0.1296	0.1857	0.2062
Australian All Ords	0.0650	0.1095	0.1371	0.1050	0.1413	0.1516
Nikkei 225	-0.0000	0.0317	0.0920	0.0605	0.1228	0.1552
FTSE 100	0.0370	0.0578	0.0710	0.0720	0.0767	0.0747
NZX All	0.0441	0.1020	0.1565	0.1650	0.1952	0.2060
NZX 10	0.0504	0.0744	0.1030	0.0973	0.1116	0.1394

Table 2: Indexes and Returns for 1x Instances of the Strategies

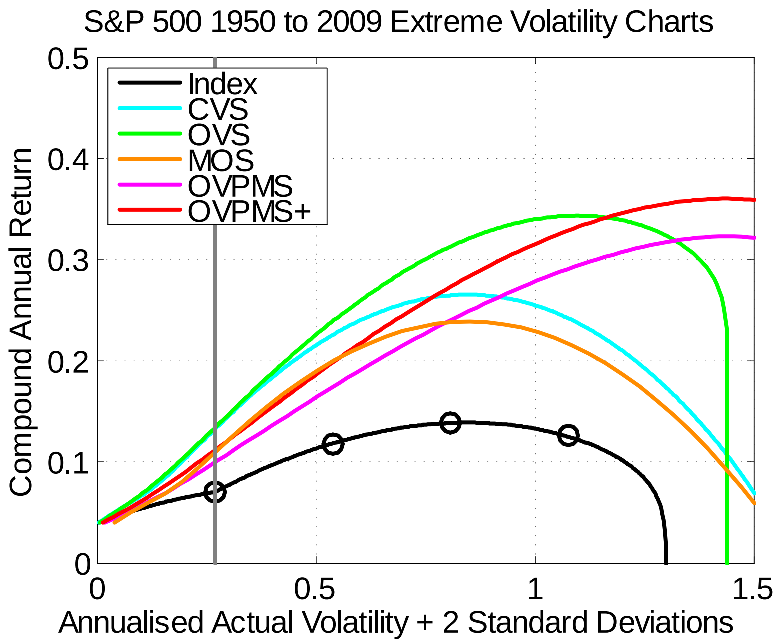

Figure 9: Extreme Volatility Chart for the Strategies

5.2 Extreme Volatility Risk Performance

Figure 9 shows the returns of the strategies versus the extreme volatility (SDDR+2). This produces dramatic changes in the rankings of the strategies.

CVS with its almost constant volatility and smaller vovo produces returns as good as OVS and both strategies beat the strategies that use the $m(s_i)$. It seems that any method that incorporates $m(s_i)$ gets an extra boost in vovo which reduces its effectiveness in this risk framework.

How good is the constant volatility of CVS? Since many managers estimate volatility using historical volatility and since investor perception of volatility comes from experience of recent volatility let us estimate the volatility of CVS using 63-day (corresponding to a three monthly reporting period) historical volatility. This is plotted in Figure 10 with the S&P 500 index 63-day volatility in the background. A SDDR+2 horizontal line is drawn for each series. Both series have the same mean level of volatility (shown as a grey horizontal line) but CVS is less extreme both upwards and downwards. CVS is still not as smooth as we might like and this is an area for further research and optimisation.

We have looked at extremes in the volatility - what happens to extremes in the returns? One measure of extreme returns is the kurtosis of the returns which measure how 'fat' the tails are. It is known (Poon and Granger 2003) that vovo increases kurtosis in the returns

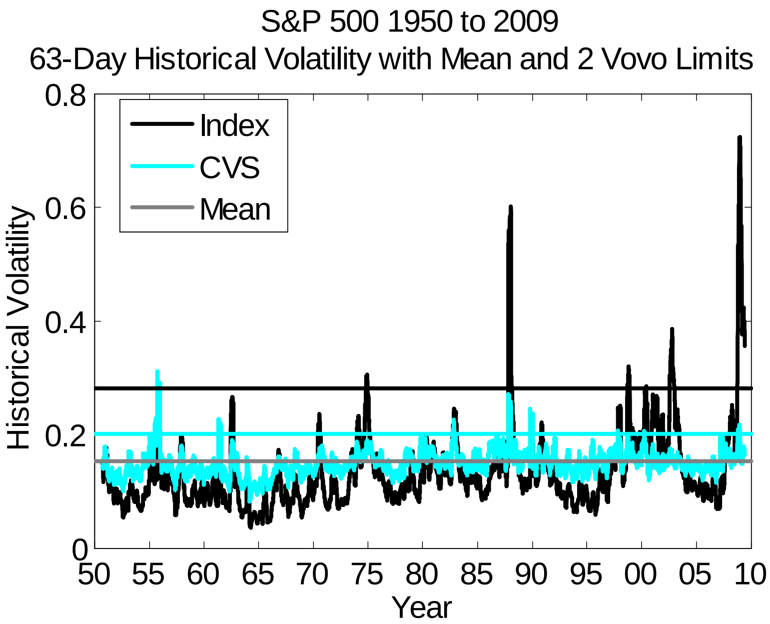

Figure 10: Historical Volatility for the Index and CVS

so a reasonable question to ask is whether smoothing the volatility reduces vovo. Figure 11 shows the kurtosis of the returns for the strategies for various indexes (the lines connecting the points have no meaning and are there to guide the eye). We see that the CVS, OVS, and MOS strategies produce less kurtosis but the other strategies increase the kurtosis to exceed that for the index.

5.3 Maximum Drawdown

Kurtosis measures the tail risk of returns and SDDR+2 measure the tail risk of volatility. But investors are not concerned with these risks *per se* because volatility brings both rewards along with losses. What really keeps investors awake at night is a sustained sequence of losses that produces drawdown in the investment. So we look at maximum drawdown which is defined as the greatest peak-to-trough loss in the investment in any given time period. We use a time period of one year for our investigation so that we get more observations although in practice longer periods such as a decade might be more useful.

A sustained sequences of losses can occur with both high and low accompanying volatility so it is not obvious that reducing volatility reduces drawdowns. Magdon-Ismail, Atiya, Pratap, and Abu-Mostafa (2004) find that in a simplified random walk model expected maximum drawdown approximately increases with σ^2 and decreases with μ where these quantities refer to the

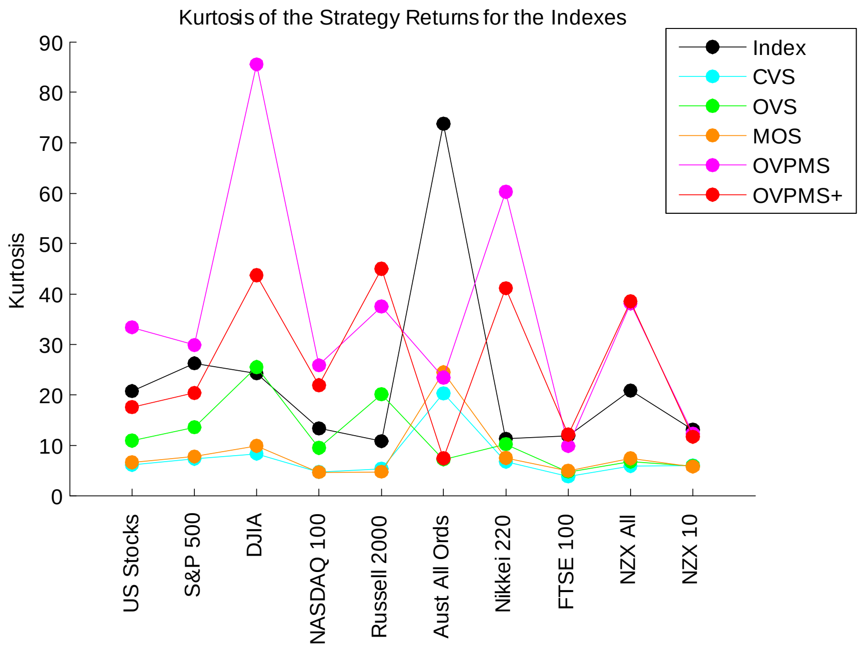

Figure 11: Kurtosis of the Strategy Returns

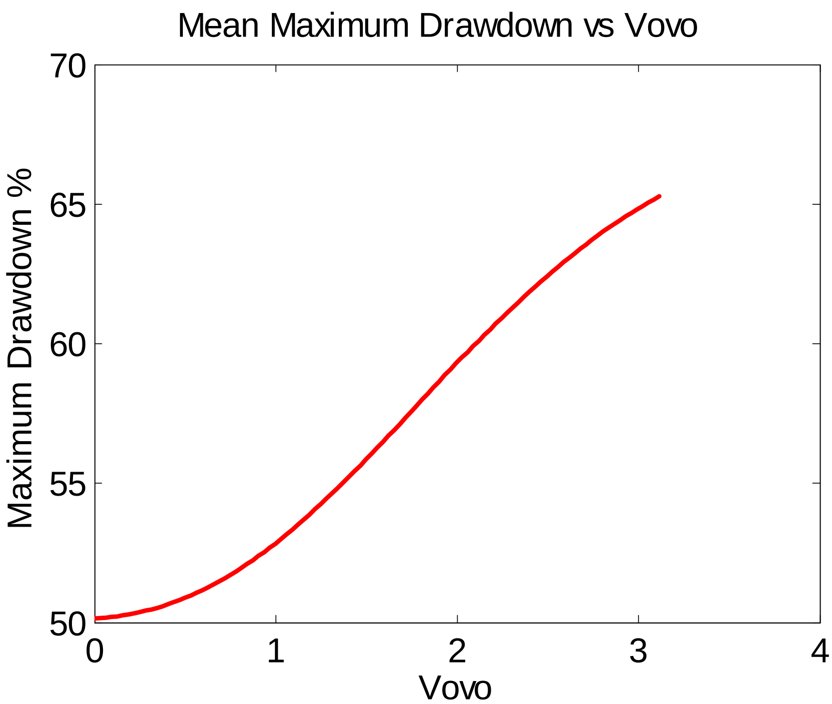

Figure 12: Maximum Drawdown versus Vovo Simulation Results

daily returns[4]. Even though stockmarkets follow paths more complicated than simple random walks—particularly in regards to the interaction between returns and changes in volatility—it seems reasonable to expect drawdowns to reduce with reduced volatility.

It is even less clear what the effect of vovo is on maximum drawdown. Vovo both reduces and increases volatility so may reduce or increase drawdown which may reduce or increase expected maximum drawdown. In the absence of a formula for expected maximum drawdown for processes with vovo we did some simulations of EGARCH models. We defined vovo to be the standard deviation of the log daily EGARCH(1, 1) volatilities and allowed this to range from 0 to 3. We fixed the mean daily volatility and the mean daily return to be the same as those for the S&P 500. Then we ran enough simulations to get a smooth curve of mean maximum drawdown versus vovo and plot the curve in Figure 12. We see that mean maximum drawdown does increase with increasing vovo. This is a useful result because it suggests that even if risk smoothing does not produce benefits in the tradition risk framework (i.e. produces alpha) it can still reduce maximum drawdowns. Let's look at some real data.

In Figure 13 the period from 1950 to the start of 2010 has been divided up into individual years and the maximum drawdown has been calculated for each year for the 1x version of each strategy. Since the strategies apply leverage inversely proportional to volatility then they increase maximum drawdowns in years with low volatility and decrease maximum drawdowns in years with high volatility. This makes it difficult to get an overall impression of which strategy works best. But it is evident that the market (black) has the highest maximum drawdowns so the strategies improve on the market. Comparing the CVS (cyan) and OVS (green) we see that OVS almost always has less maximum drawdown than CVS even though CVS has less vovo. This is because OVS has a higher mean return than CVS and return is more important than vovo in reducing drawdown.

[4]In fact they show that expected maximum drawdown over an extended period of time is approximately proportional to σ^2/μ which, interestingly, is inversely proportional to our optimum leverage. It appears that management of the optimum leverage may be equivalent to management of expected maximum drawdown.

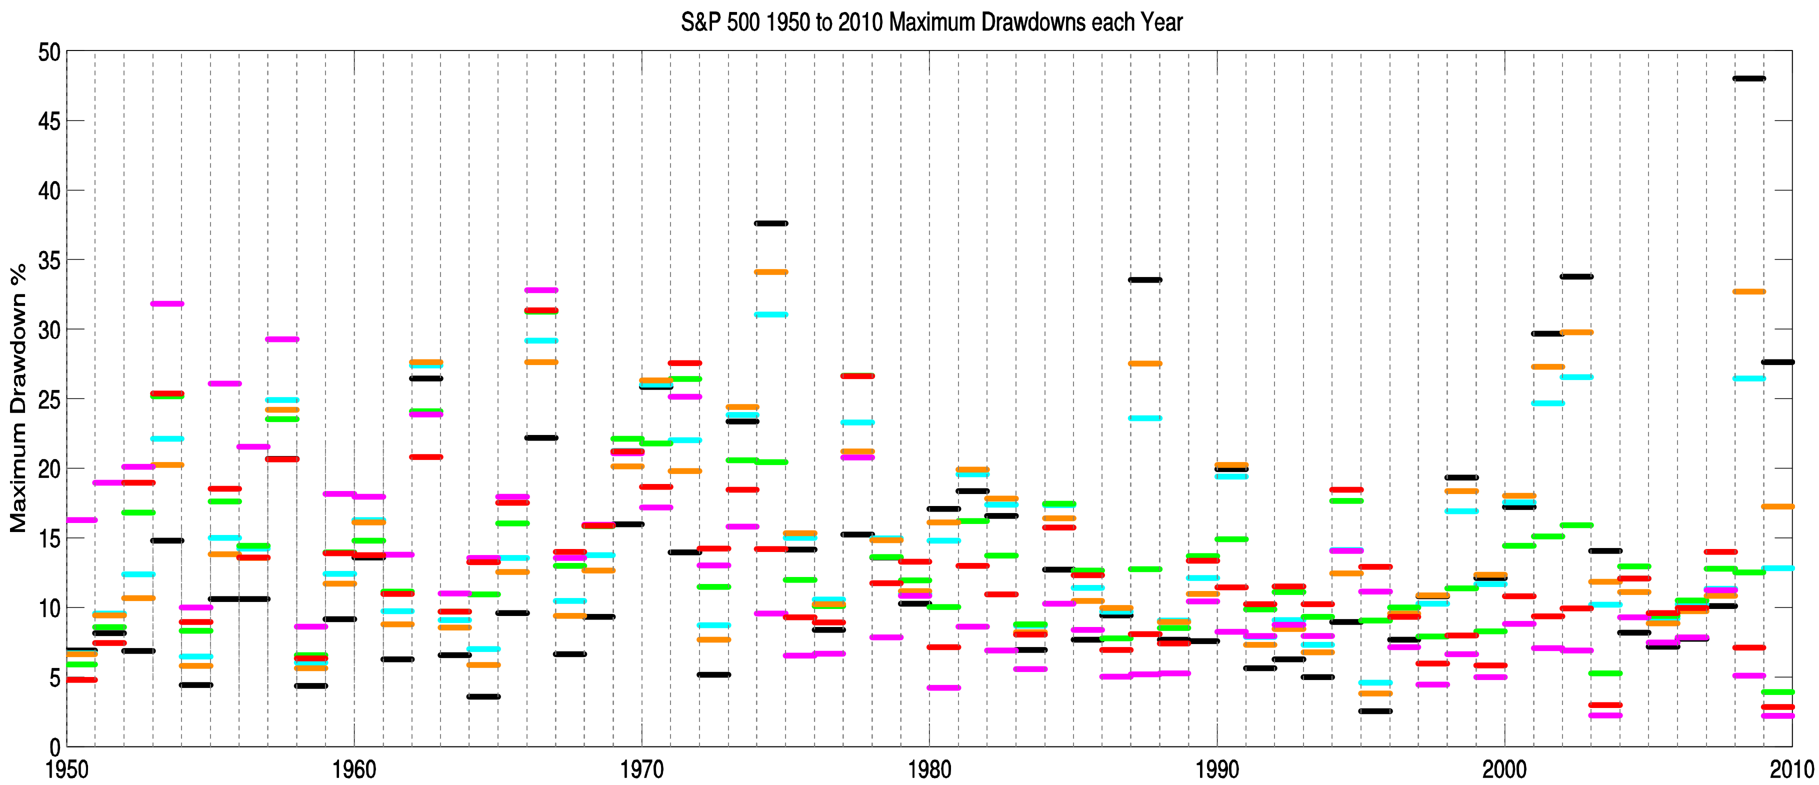

Figure 13: Maximum Drawdowns

Strategy	Annualised Volatility	Annualised Return	SDDR+2	Mean Leverage	Daily Alpha	Beta
S&P 500	0.153	0.070	0.283	1.000	0.000000	1.000
CVS 1x	0.153	0.101	0.191	1.303	0.000147	0.896
OVS 1x	0.153	0.121	0.222	1.380	0.000270	0.731
MOS 1x	0.153	0.095	0.205	1.257	0.000113	0.932
OVPMS 1x	0.153	0.126	0.291	1.310	0.000330	0.594
OVPMS+ 1x	0.153	0.129	0.264	1.352	0.000336	0.612

Table 3: Statistics of the Strategies

5.4 Alpha Beta charts

Now we get to a paradox about leverage and beta. Traditionally they are thought of as measures of the same thing. But here we find it is not so.

Figure 14 shows plots of the strategy daily returns versus the S&P 500 index daily returns for the instances of the strategies that have the same volatility as the S&P 500 index (the 1x versions, OVPMS 1x omitted to save space since it is similar to OVPMS+ 1x). Alpha and Beta are the intercept and slope of the regression lines—the least squares fit to the equation $StrategyReturn = \alpha + \beta IndexReturn$. β measures the amount the strategy returns for each unit of return in the index. It is thus a measure of leverage. Table 3 shows that all the strategies have a β less than one. On average, the strategies move less than the index moves.

The table also shows the mean daily leverage that the strategies applied to the S&P 500 index over the duration of the study. The figure shows the mean leverages as a blue line. All the strategies have a mean leverage greater than one. This explains why the strategies beat the index—the index has a long term trend that is upwards and the strategies have a leveraged exposure to this trend. The vovo strategies in a sense apportion out this leverage so that the overall volatility is not increased yet the overall leverage is increased.

So why is β telling us the exposure to the index decreased overall? The reason for the paradox is that in the regression of daily returns the points with the most leverage have the lowest volatility and so tend to be near the origin of the charts and have less influence on the regression line[5].

[5]Statisticians would say that those points have less "leverage" on the regression line but we won't use that statistical term here because it could be confusing.

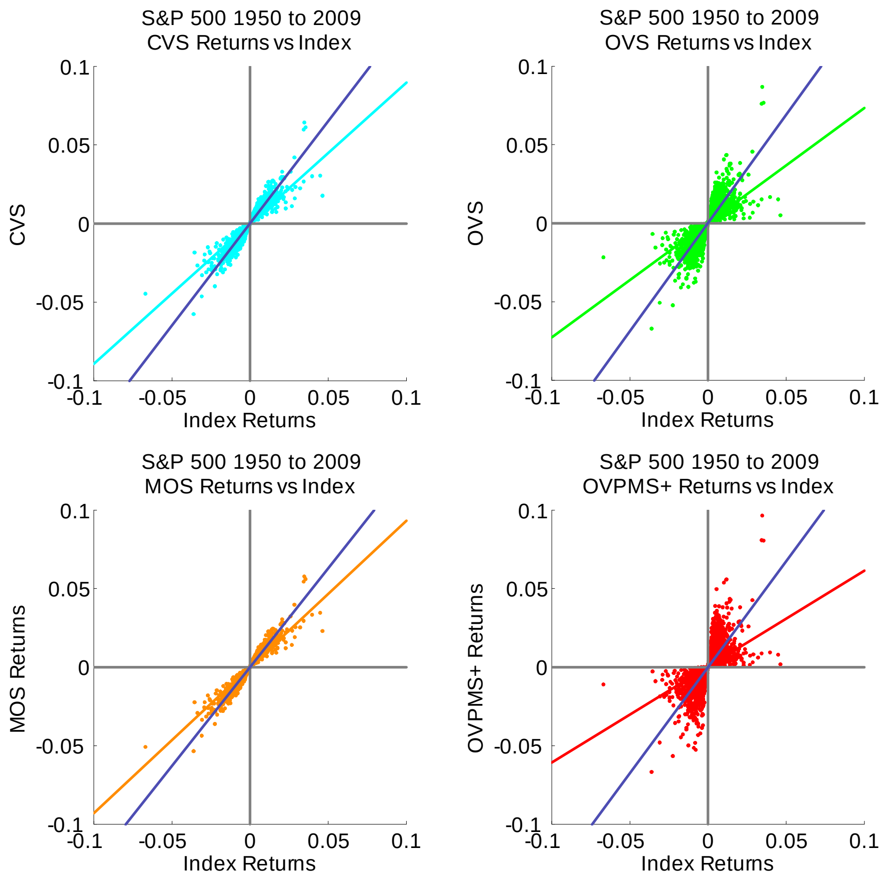

Figure 14: Strategy Returns versus Index Returns

5.5 Sideways and Bear Market Performance

The Nikkei 225 index from 1984 to 2009 is an interesting case of a sideways market. This index was 10149.00 on 14 June 1984 and 10135.82 on 12 June 2009. So the annualised rate of return was -0.0053% over this period which is practically zero. This lets us see how the strategies performed in a market that would not benefit from leverage under a buy-and-hold strategy.

The strategies that use mean prediction (Table 2 OVPMS and OVPMS+) did very well in this market. It appears that using volatility to predict returns is an effective strategy for the Nikkei 225—we got 12.5% or 16% annual returns out of a market that went nowhere.

The next index we look at (Figure 15) is the S&P 500 from 24 March 2000 to 9 March 2009 when the index went from 1527.46 to 676.53—a sustained bear market where the annualised return was -8.72%. All the strategies beat or equalled the index but they still lost money over the period. So it appears that in a bear market when leverage is expected to amplify losses the vovo strategies do no worse than the index. But they do not turn a bear market into a bull.

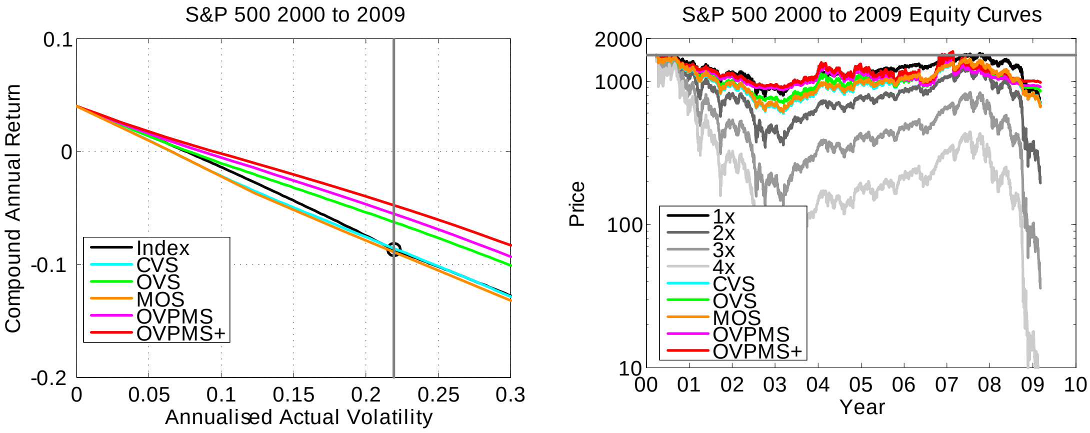

Figure 15: Strategy Performance for a Bear Market

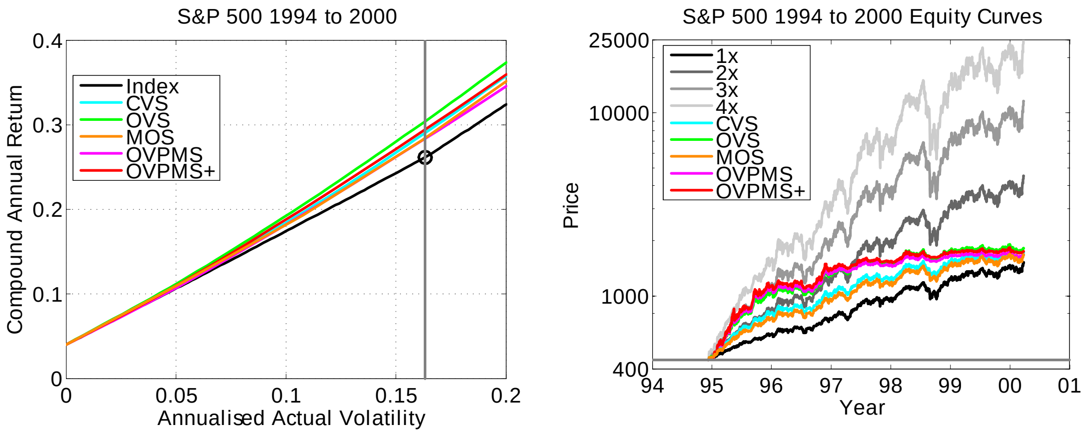

Figure 16: Strategy Performance for a Bull Market

5.6 Bull Market Performance

9 Dec 1994 to 24 March 2000 was a bull market where the S&P 500 index went from 446.96 to 1527.46—an annualised rate of 26.11%. Figure 16 shows that the vovo strategies outperformed the index but not by much. The reason is possibly because this was a period of low volatility and low vovo. The vovo strategies did not have much vovo to work on.

5.7 1987 Crash Protection

We are interested in two aspects of the 1987 crash — did the strategies protect the portfolio from the downside of the crash and did the crash come close to ruining the strategies?

Table 4 shows the leverage that the strategies determined the day before the crash could have been used, successfully, the day of the crash.

It can be seen that the strategies greatly reduced their leverage for the day after the crash so to a certain extent they reacted as panicked investors and missed some of the next day's

Strategy	Leverage 19 October 1987	Leverage 20 October 1987
CVS	0.455	0.160
OVS	0.152	0.019
MOS	0.568	0.256
OVPMS	0.034	0.002
OVPMS+	0.058	0.003

Table 4: Special Leverage Values of the Strategies

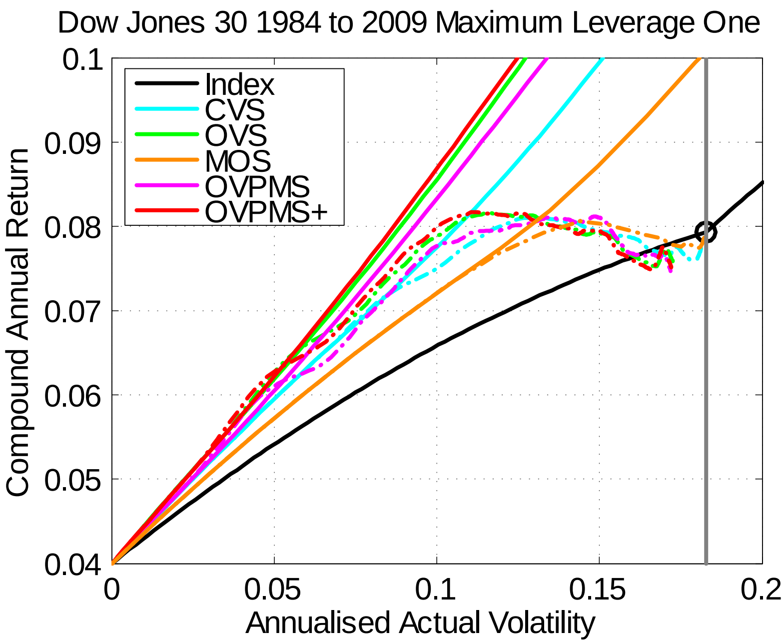

Figure 17: Return versus Volatility when Maximum Leverage is One

bounce. But the important point here is that all the strategies saw the crash coming and had a greatly reduced exposure to it thereby gaining overall (relative to the market) from the crash. Any conclusions we make here constitute data torturing because the crash was a one-time-only event and is (hopefully) unlikely to ever be repeated due to measures put into place since then (such as closing the markets when they move too far).

5.8 The Long-Only Maximum Leverage 1 Portfolio

Our vovo strategies had no limit to the amount of leverage applied but we need to look at how well they perform when they are restricted to a maximum leverage of one. Figure 17 shows the results for the Dow Jones Industrial Average for 1984 to 2009. Setting the maximum leverage to one means that the volatility of the strategy cannot exceed the volatility of the index. So we only show the region of the chart to the left of the index volatility.

The most interesting result is that our returns are reduced to about the index return. This is perhaps to be expected since our excess returns come from having mean leverage greater than one in favourable markets so the maximum leverage of one will reduce the mean leverage to below one. But this is not a law, however. By reducing leverage when the index is doing poorly means that the mean return *can* exceed the return of the index. For some indexes such as the NASDAQ 100 (omitted for brevity) the returns do exceed the index returns. But not by much. A limit to the actual and mean leverage does impact strategy returns.

The chart shows that we can approximately halve the annualised volatility down to about 0.1 while slightly increasing returns. So the effect of the strategies on the traditional portfolio is likely to be a reduction in volatility rather than an increase in returns.

6 Conclusion and Directions for Further Research

Returns are not easily predictable but volatility is. By allowing for leveraged investing we have introduced a formula for compound returns that depends on volatility thereby introducing an element of predictability into returns. We have identified three sources of alpha from the volatility predictability:

- The extra overall leverage allowed by the apportionment of leverage without increasing overall portfolio volatility
- The relationship between returns and volatility and the use of extra leverage when returns are higher
- The deleveraging prior to the 1987 crash

We have introduced three simple intuitive investment strategies that use these sources to generate alpha in the various markets that we have tested it in.

The three strategies are: the Constant Volatility Strategy (CVS) which aims to achieve a constant daily volatility by predicting the volatility and setting a leverage value to aim at the target; the Optimal Volatility Strategy (OVS) which sets the daily leverage to be inversely proportional to the predicted volatility squared; and the Optimal Volatility Plus Mean Strategy (OVPMS) which estimates the expected return as a power function of volatility and sets the daily leverage to be the expected return divided by the predicted volatility squared.

All three strategies manage to get the upsides of leverage without the downsides—they generate excess returns in bull and sideways markets and do no worse than the market in bear markets. If they are restricted to leverage no more than one then they manage to greatly reduce the volatility of a portfolio without significantly reducing returns. And all three strategies reduce the volatility of volatility of the portfolio which allows the portfolio manager to deliberately seek more volatility than normal.

The order of performance of the strategies in increasing order when measured using risk-adjusted returns is CVS, OVS, OVPMS. But of the three strategies the one that we prefer is the worst performing one CVS. We prefer this because:

- It allows us to seek a specified constant target volatility that can be put into a prospectus and in practice the realised volatility is much less variable than the market and other strategies
- It usually has leverage less than 3 so has the least maximum risk of all the strategies (and the least probability of ruin) and can be implemented using existing 3x leveraged ETFs
- It has the best return profile when measured against the maximum lifetime volatility of the portfolio (as measured by mean volatility plus two times the standard deviation of volatility)
- It produces daily returns with the lowest kurtosis of all the strategies

It has been said that diversification is the only free lunch in investing. It appears that once you have diversified away some risk you can get a further free lunch by smoothing what risk remains.

The following subsections discuss a few directions for further research currently under investigation.

6.1 Multiple Assets Model

We only considered the volatility of stock market indexes in any detail which is a univariate situation. When we consider multiple assets we get not only volatilities but correlations between returns. Then Formula (2) becomes $\underline{k} = \Sigma^{-1}\underline{\mu}$ where Σ is a variance-covariance matrix and $\underline{k}$ and $\underline{\mu}$ are vectors. We might consider a multivariate EGARCH model for estimating Σ.

6.2 Interaction of Volatility Changes and Return

We have ignored the relationship between volatility changes and return. We have only estimated return in as much as it relates to absolute *levels* of volatility rather than to *changes* in volatility. But there is evidence that returns and volatility changes are highly negatively correlated and in more complicated ways than given by γ in Equation (3). For example Bouchaud, Matacz, and Potters (2001) find in their "Retarded Volatility Model" that the correlation effect for individual stocks is moderate and decays over 50 days, while for stock indices it is much stronger but decays faster.

6.3 Implications for Technical Trading Rules

The above analysis investigated dynamic leverage as an overlay for buy-and-hold strategies. But some investment strategies (such as those using technical trading rules) that use leverage of either zero or one cannot use a dynamic leverage overlay. An example is a strategy that says to be in the market (leverage one) when the market is above its 200 day moving average else it is out of the market (leverage zero). For those strategies a volatility overlay has to either be converted into modifications of the existing timing signals or modifications of the zero and one leverage values (position sizing).

To achieve this the technical rules have to be represented in a framework where the predictability of volatility can be used. This is an interesting area for research.

References

Bouchaud, Jean-Philippe, Andrew Matacz, and Marc Potters, 2001, Leverage Effect in Financial Markets: The Retarded Volatility Model, *Phys. Rev. Lett.* 87, 228701.

Ferson, Wayne, Timothy Simin, and Sergei Sarkissian, 2003, Is Stock Return Predictability Spurious?, *Journal of Economic Literature* 1, 10–19.

Friedman, J.H., and W. Stuetzle, 1982, Smoothing of scatterplots, *Stanford Linear Accelerator Center Technical Report PUB-3013 ORION* 003.

Magdon-Ismail, Malik, Amir F. Atiya, Amrit Pratap, and Yaser S. Abu-Mostafa, 2004, On the maximum drawdown of a Brownian motion, *Journal of Applied Probability* 41, 147–161.

Merton, Robert C., 1969, Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case, *The Review of Economics and Statistics* 51, 247–257.

Miller, M.H., 1999, The history of finance, *Journal of Portfolio Management*.

Poon, S.H., and C.W.J. Granger, 2003, Forecasting volatility in financial markets: A review, *Journal of Economic Literature* 41, 478–539.

Samuelson, Paul A., 1969, Lifetime portfolio selection by dynamic stochastic programming, *The Review of Economics and Statistics* 51, 239–246.

Schwert, G.W., 1990, Indexes of U.S. stock prices from 1802 to 1987, *Journal of Business* 63, 399–426.

Sharpe, William F., 1966, Mutual Fund Performance, *The Journal of Business* 39, 119.

Thorp, E.O., 2006, The Kelly criterion in blackjack, sports betting and the stock market, *Handbook of Asset and Liability Management, Volume A: Theory and Methodology* A.