A Method For Creating a Return Enhancing Overlay with Near-Zero Correlation to the Portfolio

January 1, 1998

Robert J. Bannon
Managing Director
Analytic•TSA Global Asset Management
700 S. Flower Street
Los Angeles, CA 90017
(213) 787-9727

The Separation of Portfolio Management from Enhancement

The existence of an active derivatives market, both exchange-traded and over-the-counter, has allowed for the physical separation of portfolio management into two distinct tasks: (1) minimizing the tracking error of generating a benchmark index's return; (2) maximizing the value-added from a return-enhancing overlay. Previously an investment manager's actions could be theoretically broken into index achievement vs. value-added only in an *ex post* sense, by subtracting the benchmark return from the actively managed portfolio's return. Manager prowess at such activities as asset allocation and security selection could only be judged *en toto*, by relative benchmark performance. Despite being judged relative to a benchmark, the traditional manager was unable to separate his actual trading activity into index management and return enhancement.

Derivatives allow for separate targetting of index returns from value added returns due to their ability to be used in an overlay capacity. Derivatives can provide a vehicle for return enhancement to a portfolio without significantly altering the manager's ability to direct the underlying portfolio independently. Typically these return-enhancing overlays require between 2% and 15% of the portfolio to be available as cash or short-term fixed income instruments to serve as collateral for the derivative instruments. Such amounts are often easily met through the portfolio's existing allocation to cash.

Generating Low Correlation Overlay Programs

Overlay programs can derive their returns from such diverse avenues as currency management, yield curve arbitrage, call overwriting, volatility forecasting, long/short equities, and many other sources. The Holy Grail of overlay managers (aside from just generating positive returns) is to achieve a low correlation of overlay returns to underlying portfolio returns. While low correlation does nothing to make the absolute return of an overlay larger (that value is simply an accounting value), it certainly makes the risk-adjusted relative return of the overlay appear more attractive.[i]

Designing and implementing a near-zero correlation overlay program is easier said than done. Even those overlay managers that claim to run a near-zero correlation overlay program rely more on their empirical track record (which may have achieved low correlation more through luck than skill) than on any *a priori* economic argument as to why their program <u>should</u> have a near-zero correlation to the underlying portfolio's benchmark. Indeed, most overlay programs are run in such a way that a moderate to strong degree of correlation should not be unexpected: call overwriting programs favor a bearish equity market over a bullish one; most currency overlay programs use trend-following technical programs that effectively boil down to a direct exposure to the direction of the US dollar, and programs that buy or sell volatility often experience the well-known empirical relationship that implied volatility rises in market downturns and falls in market upturns (which in turn effectively imposes directional market exposure despite any intention to avoid it).

An important question that might be asked regarding overlay programs then is whether there is any approach which might help to insure that the performance of an overlay program achieves a near-zero correlation to an underlying index. The goal of this article is to establish, via simulation, a set of sufficient conditions which if met can allow an overlay manager to deliver a near-zero correlation to the underlying manager's performance.

Three Conditions for a Near-Zero Correlation Overlay

The first condition for the construction of a near-zero correlation overlay is the existence of a symmetric, zero mean investment process. By "symmetric zero mean investment process" we mean a market-derived set of returns that averages to zero in reasonably large samples and whose distribution is roughly evenly distributed between negative values and positive values. Note that this is a weaker condition than what statisticians call "stationarity," since symmetry in the distribution does not necessarily require the process to have constant volatility over time.

Do processes exist in the financial markets that meet this first condition? Undeniably there are. One simple case is that of interest rate changes. Chart 1 presents a histogram of month-to-month changes in the forward-adjusted[ii] yield of the 10-year US Treasury Constant Maturity series. The average monthly yield change over the 26 year period is only -1 bp (almost identical to the median change and certainly qualifies for the "zero mean" aspect of the first condition). The number of yield increases roughly equals the number of yield decreases (qualifiying this time series for the second part of the first condition, that the time series be symmetric although not necessarily constant variance).

Chart 2 offers another possibility, that of the forward-adjusted monthly change in the US dollar/Deutschemark exchange rate over a 21 year period. While not perfectly mean zero (+0.11% per month average gain), the average gain is still small relative to the average absolute gain (2.58%), and the process does appear to be roughly symmetrical. Chart 3 shows the monthly returns to the S&P 500 risk premium over an 18 year period. As can clearly be seen, the S&P 500 monthly returns are <u>not</u> mean zero <u>nor</u> symmetric around zero, although they are roughly symmetric around their mean[iii]. Chart 4 shows one of the best candidates for a zero mean symmetric process: monthly changes in forward-adjusted US Treasury yield spreads. Both the symmetry and the zero mean are clear from the graph and data. Chart 5 shows an interesting and unusual example: assuming that a product existed that could turn changes in actual volatility into an income stream, the monthly change in the intramonth actual volatility of the S&P 500 index (measured in non-overlapping monthly increments) has the same basic properties of the previously mentioned candidate processes.[iv]

Once a candidate investment process has been identified the second condition for achieving a near-zero correlation overlay program is to develop a randomly accurate binary choice model for predicting whether the process will provide positive or negative returns in a given time period. By "randomly accurate" is meant that the signal has a degree of accuracy better than 50%, but that the occurrence of its accuracy is unpredictable. In particular, in order to insure near-zero correlation, the binary choice model cannot favor, as measured by its *ex post* success rate, the prediction of positive vs. negative returns.[v]

Models exist that can channel independent explanatory power into forecasting the probability of an event rather than into a forecast of the size of an event. Such models as Logit, Probit, Tobit and Markov chains fall into this category. Models such as discriminant analysis and CART also offer a similar approach to event analysis, although the signal must be interpreted differently since it is not a direct probability estimate. The key is to be able to forecast, with better than 50% accuracy, whether the upcoming return from a zero mean symmetric investment process will be positive or negative, and to forecast that outcome without a bias that favors either positive or negative outcomes. The specification and estimation of any particular probability forecasting model (i.e., choice of variables, lag structure, estimation period, criterion for judging goodness of fit) is an econometric issue that must be tackled on a case-by-case basis. An example will be presented below of one such set of models that have been built for yield curve spread changes.

The final condition for a near-zero correlation overlay program is that there exist an effective method for trading the aforementioned zero mean symmetric investment process efficiently in both directions. That is, the portfolio manager must have available a method for capturing the beneficial outcome (or suffering the negative consequences) of his binary choice, regardless of whether the market goes up or down. In the five markets presented in Charts 1-5 a method exists for flexibly entering the market either long or short. US Treasury bonds can be financed both long and short, or they can be bought and sold in a relatively active forwards market. These same methods extend to trading yield spreads - additionally, there are both over-the-counter and futures based alternatives for trading yield spreads on a conveniently duration and convexity neutral basis. Foreign exchange can be traded on a forward as well as spot basis. Numerous alternatives can be employed for trading the S&P on either a long or short basis, including the use of futures, options, depository receipts, or outright baskets. Finally, the options markets offer avenues for trading actual volatility (straddles, strangles), while the over-the-counter markets offer the aforementioned volatility swaps. Other markets may not be as flexible as these for going both long and short (i.e., emerging market debt), but clearly there are a large number of candidate markets that can be considered.

Benefits of Adhering to the Three Conditions for a Near-Zero Overlay

Figure 1 consolidates the three conditions described above, and lists their benefits to the portfolio manager. The first benefit, as mentioned previously, is that a zero correlation overlay program can be constructed. A mathematical proof is offered in Appendix A demonstrating that the expected correlation between the return performance of such an overlay program and the return performance of the underlying market is in fact zero under ideal circumstances. Simulation results presented in the next section suggest that even if the ideal conditions do not hold (as Chart 3 suggested might be the case for the S&P) the conclusions may not be dramatically altered.

The second benefit of the program is that it defines its value-added directly in terms of the degree of accuracy of the binary choice model. The simulation process described in the following section offers an approach that allows the portfolio manager to define in advance the expected return to an overlay as a function of the degree of accuracy of the binary choice model. This focuses the investment effort on developing an accurate binary choice model, while the simulation process allows the manager to gauge whether the highest degree of accuracy he believes can be achieved from the model will generate sufficient profit to make the program worth the attempt. The simulation process will also aid in assessing the risk to the expected profit.

The third benefit is simply that under many circumstances this program can be implemented as a near-zero cost overlay, assuming the portfolio manager has the internal authority or the external credit-standing to engage in the necessary instruments. Either exchange-traded derivatives can be utilized, or the forwards markets, or offsetting financing arrangements (repos and reverse repos). Sufficiently credit-worthy organizations can enter into over-the-counter swap contracts or other structures that directly deliver a payout targetted to the zero mean symmetric process (such as the volatility swaps, which pay or receive a fixed amount for each percentage point that actual volatility rises or falls). The total initial cost for the overlay can range anywhere from zero (for a pure OTC swap by a large, creditworthy institutional investor), to moderately low (for initial margins on exchange traded derivatives or net repo costs on cash traded bonds), to moderately high (for domestically regulated stock managers who are constrained by 50% margin rules for the trading of stock baskets).

A Simulation Study

The proof offered in Appendix A validates the proposition that, under ideal conditions, an overlay investment process generated by employing a randomly accurate binary choice model applied to an underlying zero mean symmetric investment process will yield performance returns that are uncorrelated with the returns from the underlying investment process. Answering the question as to how robust this conclusion would be under less than ideal conditions is a more difficult endeavour. For true validation one would need to identify a tradable zero mean symmetric investment process, build a predictive binary choice model, backtest it, and trade it for a number of years. Even if such a test-case overlay program was successful (in that it had positive returns and was uncorrelated with the underlying market), there is no guarantee that the results occurred due to use of the binary choice model. Any number of factors could obscure the conclusion, including a breakdown of the model's out-of-sample predictive power, a new-found non-symmetry to the underlying process, or even just trading implementation errors.

A properly designed simulation study can minimize some of these obscuring factors. In particular, the need to rely on whether or not the binary choice model retains its in-sample accuracy through the out-of-sample period, as well as the need to rely on a trader's ability to implement the program (not to mention the time spent waiting for a track record to be developed), can all be sidestepped by the use of simulation. Only the reliance on whether the historical market data are "symmetric enough" and "close enough to mean zero" remains as a question mark in the simulation study.[vi]

The methodology of the simulation needed to validate the near-zero correlation of a randomly accurate overlay program is relatively simple, and follows the approach used in the proof found in Appendix A. Given an actual time series that is believed to be zero mean and symmetric, a Bernouilli random variable is generated for each time period (observation) within the time series. The Bernoulli random variable takes the value of +1 with probability p and -1 with probability (1-p). If the Bernoulli variable is +1, the trader is presumed to have been correct about that period's market move, and to have earned the absolute value of the investment process that period (even if the investment process is negative during that period, the trader is presumed to have taken a short position in advance of it and thus profited). If the Bernoulli variable is -1, the trader is presumed to have lost the absolute value of the investment process that period. For a given underlying investment process with T time periods, a performance track record for the overlay investment process can be calculated, also with T time periods. This overlay track record can be cumulated, measured for risk, and correlated to the underlying process.

The choice of p in the generation of the Bernoulli random variable is exactly analogous to the presumption that the binary choice model will have p% accuracy in its out-of-sample use. Unlike building an econometric model, however, the simulation can insure, within the limits of sampling accuracy, that a given level of binary choice accuracy is achieved. The full simulation study for a given underlying investment process will thus consist of N simulations per T time periods for K different levels of accuracy in the binary choice. The K levels of accuracy are the 21 evenly spaced 5% intervals from 0% to 100%. The number of simulations (N) will be set to 1000 to insure large sample precision.

Charts 6, 7 and 8 present key summary information for overlay program simulations performed on three different markets. The left vertical axis for each graph measures expected monthly performance, in percentage points. The right vertical axis measures correlation, also in percentage points. The horizontal axis measures the degree of random accuracy within the simulation that has been structured to occur for the binary choice model (from 0% to 100% at 5% intervals). Performance of the overlay for 1000 separate simulations at each discrete accuracy interval is plotted (the solid line), along with a confidence band around this performance (the dotted lines). The confidence band represents a two standard deviation range across each set of 1000 simulations. Also plotted on the graphs (as triangles) is the correlation between program performance and a consistently long position in the underlying index, at each level of binary choice accuracy.

Chart 6 examines the performance and correlation that might be achieved if a randomly accurate overlay program were applied to the monthly forward-adjusted foreign exchange market returns originally portrayed in Chart 2. As can be seen, there is a nearly linear relationship between the degree of forecasting accuracy and the level of expected performance. The confidence intervals around the performance are of a relatively fixed distance, except near the extremes where the bands collapse to the expected value. Correlations of the simulated overlay to the underlying investment process are bounded between +17% and -17%, with a slight positive relationship to the degree of forecast accuracy (although in a reasonable range of forecasting accuracy, say 35% to 65%, the correlations are virtually zero and insensitive to forecasting accuracy).

As an example of the data plotted in Chart 6, consider the 55% degree of forecasting accuracy. Expected returns would be roughly 250 basis points (per month), with a standard deviation around that expectation of about 250 basis points. The average correlation observed across the 1000 simulations at 55% accuracy was less than 1%, with a standard deviation (which does not appear on the graph) of 11.4%. Clearly the average correlation is not statistically significant from zero, and should be less than 25% in absolute value more than 95% of the time. Given the large swings in the original foreign exchange returns, it is not surprising that the level of return can get quite large the greater the deviation from a 50/50 guess. At an 80% degree of forecasting accuracy, the foreign exchange overlay program would expect to deliver almost 19.5% (again on just a monthly basis), still with a standard deviation of 240 basis points, and with a correlation to the underlying dollar/DM return of only 7.8%. Of course similar size losses would accrue if the binary choice model's forecasting accuracy was 80% wrong rather than 80% right.

Chart 7 repeats the same analysis for the US equity market. Using the data from Chart 3 on monthly total returns to the S&P 500 (net of the cost of financing), 1000 simulations were again run for all 21 accuracy intervals from 0% to 100%. Oddly, despite its non-zero mean in Chart 3, the overlay return performance was stronger for the equity market than for the foreign exchange market, and with a much tighter range on the levels of correlation to the underlying market (within ±1% across the entire range of accuracy). The range of possible returns was much wider, however. As an example of the data, consider the 55% degree of forecasting accuracy. After 1000 simulations, the expected monthly return was 280 basis points, with a 370 basis point standard deviation. The average correlation across the simulations at 55% accuracy was 1%, with a 16.7% standard deviation.

The final evidence offered is for the forward-adjusted returns to trading the 2yr/10yr Treasury yield spread on a cash return basis. Chart 8 presents the monthly performance and correlation data. Yield spread changes have been scaled to a cash return basis that assumes the 2yr/10yr trade covered a 5 year note's notional size.

Conclusion

Still in process

Appendix A

Proof That a

Randomly Filtered Zero Mean Symmetric Process

has Zero Correlation to the Unfiltered Process

Let X_t be a random time series measured at discrete equal time intervals. X_t has distribution function F() and density function f(). X_t is assumed to have mean zero and finite variance, although variance is not necessarily constant over time. The variance, if it is time-varying, is assumed to vary randomly around a long run mean. X_t is assumed to be independent of all past and future values of X_t.

Let I_t be a Bernoulli random variable with a probability distribution defined as:

$$\mathbf{I_t} = \begin{cases} +1 & \textbf{prob} = \mathbf{p} \\ -1 & \textbf{prob} = (1 - \mathbf{p}) \end{cases}$$

I_t is assumed to be drawn independently at each time interval, and to be independent of the sign or size of X_t.

Finally, define Y_t as the outcome of a "binary filtering process" applied to X_t. A binary filtering process is defined as the Bernoulli random variable times the absolute value of the zero mean symmetric random variable:

$$\mathbf{Y_t} = \mathbf{I_t} \cdot |\mathbf{X_t}| = \begin{cases} +|\mathbf{X_t}| & \textbf{prob} = \mathbf{p} \\ -|\mathbf{X_t}| & \textbf{prob} = (1 - \mathbf{p}) \end{cases}$$

In terms of the financial markets, X_t can be viewed as the market return each period t that the overlay manager believes he can receive (or lose) by trading either long or short, such as the return stream from buying or selling forward foreign exchange and closing out the trade on the forward date. I_t can be viewed as the outcome of a binary choice model that is attempting to forecast $Sgn(X_t)$ each period t. The probability p associated with I_t can be viewed as the degree of forecasting accuracy assumed to be achievable by the binary choice model. Y_t will be the *ex post* performance in period t of the manager that attempts to trade the market using the signal from the binary choice model.

The goal of this proof is to show that $Corr(X_t,Y_t) = 0$. This can be interpreted as saying that an overlay manager achieves a zero correlation in his overlay performance to an underlying index manager that would be running an equal-sized long-only index portfolio in the underlying market. For example, assume Manager A ran a $100 million index portfolio composed entirely of 10 year Treasury notes. Manager A does not attempt to add value and does not alter the size of the portfolio's exposure to 10 year notes over time. Manager B runs a $100 million overlay to this bond index portfolio, taking long or short exposure to the same 10 year note return in each period according to the outcome of his binary choice model. The following proof suggests that the returns earned by the overlay manager, trading the same market as the underlying index manager, will have zero correlation to the returns of the index portfolio. Note that the zero correlation does not depend on whether the overlay manager is able to cumulatively build the sum of the Y_t's to a positive number, nor does it depend on whether the underlying index manager can cumulatively build the sum of the X_t's to be a positive number.

Consider the definition of correlation as it would apply to X_t and Y_t:

$$\mathrm{Corr}(X_t, Y_t) = \frac{E(X_t - E(X_t))(Y_t - E(Y_t))}{\sqrt{E(X_t - E(X_t))^2 E(Y_t - E(Y_t))^2}} = \frac{E(X_t Y_t)}{\sqrt{\sigma_t^2 E(Y_t - E(Y_t))^2}}$$

where the simplification in the numerator makes repeated use of the fact that X_t is a zero mean process. It suffices from this point to establish that the numerator is zero. We examine the numerator in further detail:

$$E(X_t Y_t) = E(X_t I_t |X_t|) = E(I_t) E(X_t |X_t|)$$

It is a trivial exercise at this point to show that the expectation of I_t is (2p-1), which guarantees that the correlation between X_t and Y_t is zero for p = 50%. The broader interest, though, is to show that the correlation is zero for values of p other than 50%. If we continue to expand the remaining term in the numerator:

$$\begin{aligned} E(X_t |X_t|) &= E(X_t \cdot \mathrm{Sgn}(X_t) \cdot X_t) = E(X_t^2 \cdot \mathrm{Sgn}(X_t)) \\ &= E(X_t^2) \cdot E(\mathrm{Sgn}(X_t)) = E(X_t^2) \cdot 0 = 0 \end{aligned}$$

The first equality in the second line follows from the fact that the square of a variable (which is always positive) should be independent of the sign of the unsquared variable. The second equality in the second line follows from the fact that X_t is symmetric around zero: given that, there should be 50% probability that the sign of X_t will be +1 and 50% probability that the sign of X_t will be -1, thus the expectation of the sign is zero. If this expression is always zero then the numerator of the correlation function will always be zero, and therefore the correlation is zero. QED.

[i] While it may seem irrational on the face, managers often prefer a near zero correlated overlay to a negatively correlated overlay. If an overlay has a positive return and a near zero correlation to the underlying portfolio's benchmark, this implies that the overlay must deliver its positive return a relatively high percentage of the time. While statistically speaking a negatively correlated overlay program may offer better risk-adjusted returns, many managers are reluctant to have the overlay cause them underperformance in a strong market.

[ii] All the candidate processes examined have been adjusted for the cost of implementing a trade in that market. For the interest rate and foreign exchange data this amounts to using the forward prices at the beginning of the month and the spot prices at the end of the month. For the S&P 500, this amounts to subtracting out the opportunity cost of the funds committed from the S&P's total return, i.e., focusing only the equity market risk premium over the month.

[iii] A zero mean symmetric process is not a necessary condition for low correlation under the program described herein. In particular, a positive mean process can also be accommodated such that it will have low correlation, but the variability of the overlay returns will be higher.

[iv] Such an investment opportunity is not that far-fetched: there is a relatively active over-the-counter market for "volatility swaps," in which dealers are willing to pay or receive a fixed sum in exchange for receiving or paying an amount based on the revealed level of actual volatility in an instrument or index over an upcoming period.

[v] It should be intuitively obvious that if a process was, say, 75% accurate, but that that accuracy resulted from 100% accuracy for all positive returns and 50% accuracy for all periods of negative return, the performance of such a strategy would have a strong degree of correlation with the underlying direction of the market. This does not preclude, however, a model whose binary choice signal predicts not only the success rate of a given move but also its size. For example, if signals of 75% or better (for both positive and negative returns) are associated with stronger market moves than signals of between 50% and 75%, this will not inject an unwarranted correlation in the overlay program's correlation to the underlying process (so long as the association of signal strength with payoff size does not favor negative vs. positive market returns).

[vi] Simulation studies have the added risk that the random numbers generated may not be independent, either within or across samples.

Chart 1: Monthly Yield Changes
10 Year US Treasury

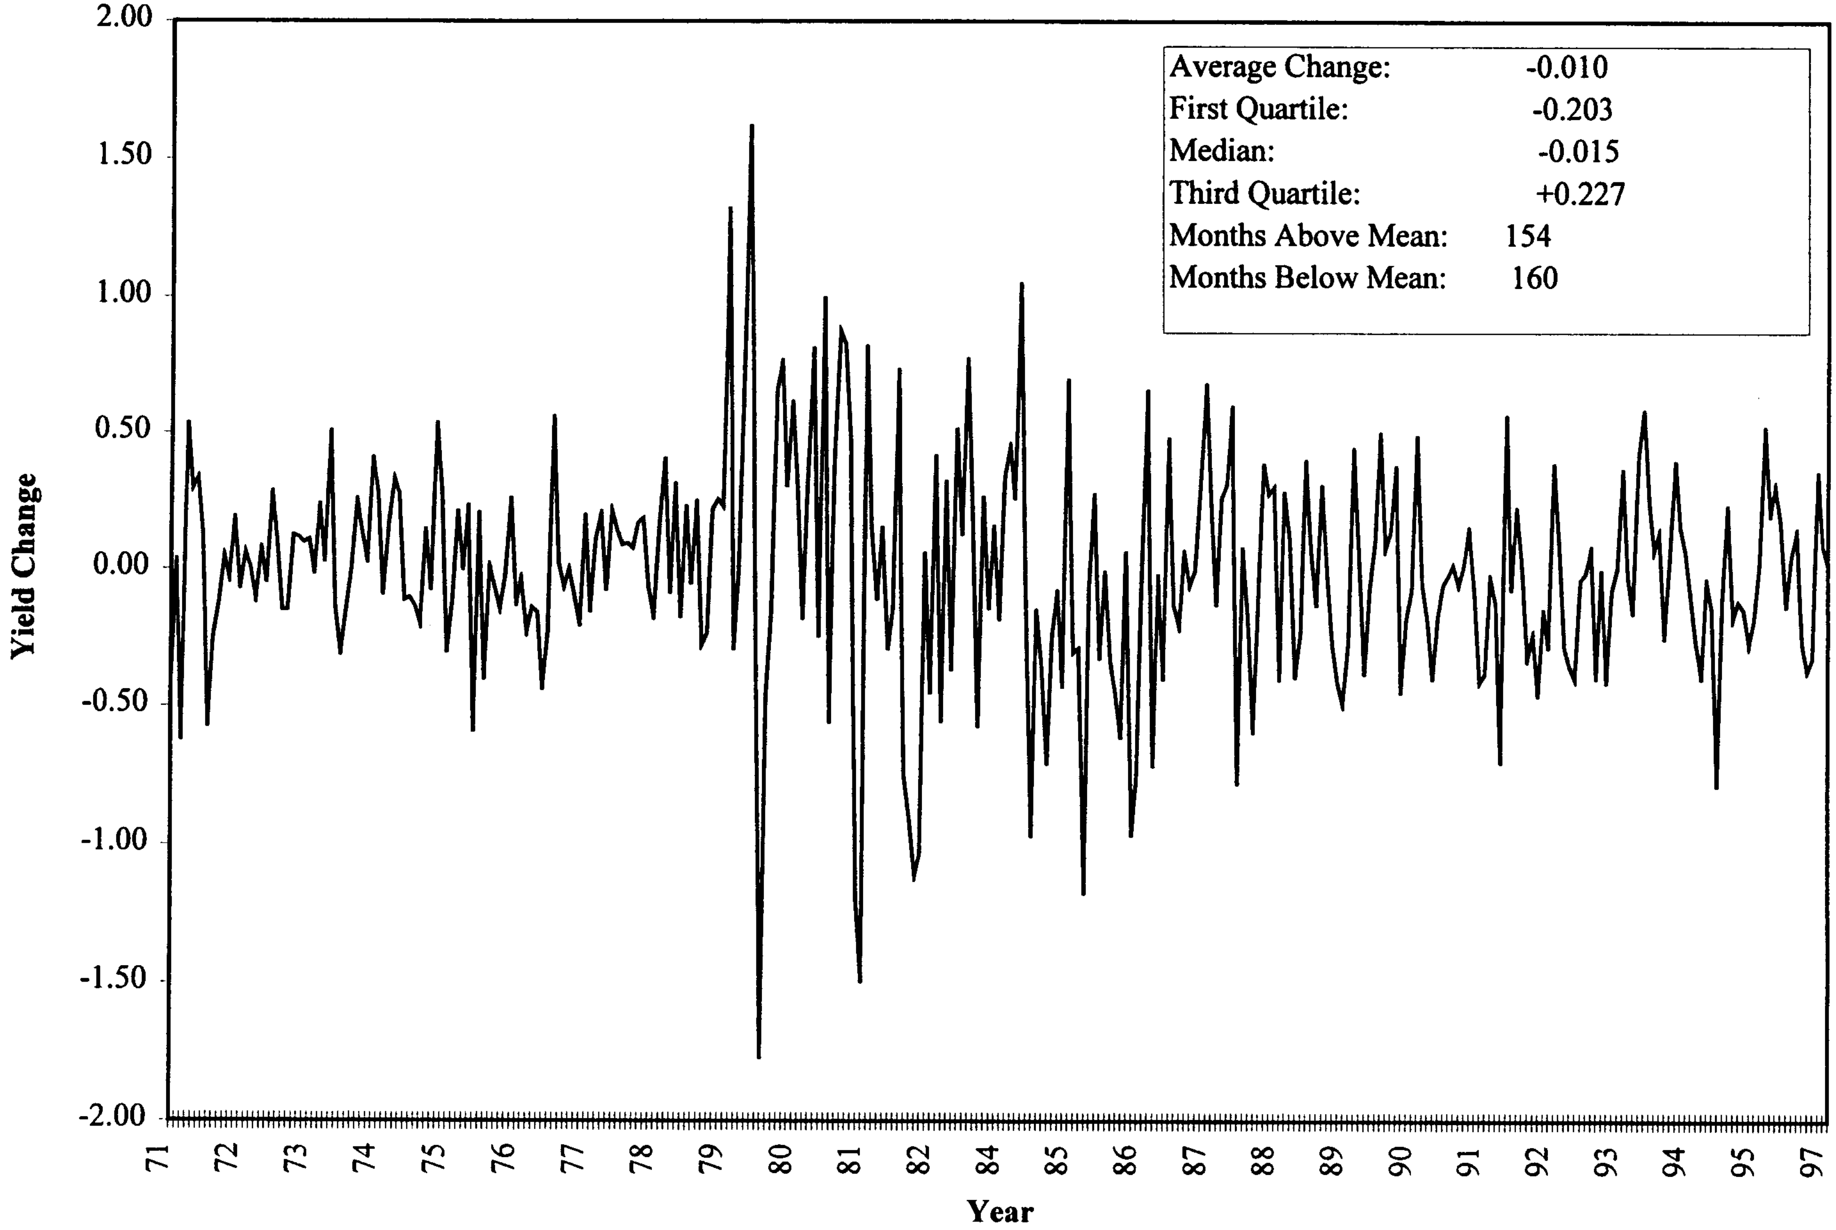

Chart 2: Monthly Forward to Spot Return
Dollar/DM Exchange Rate

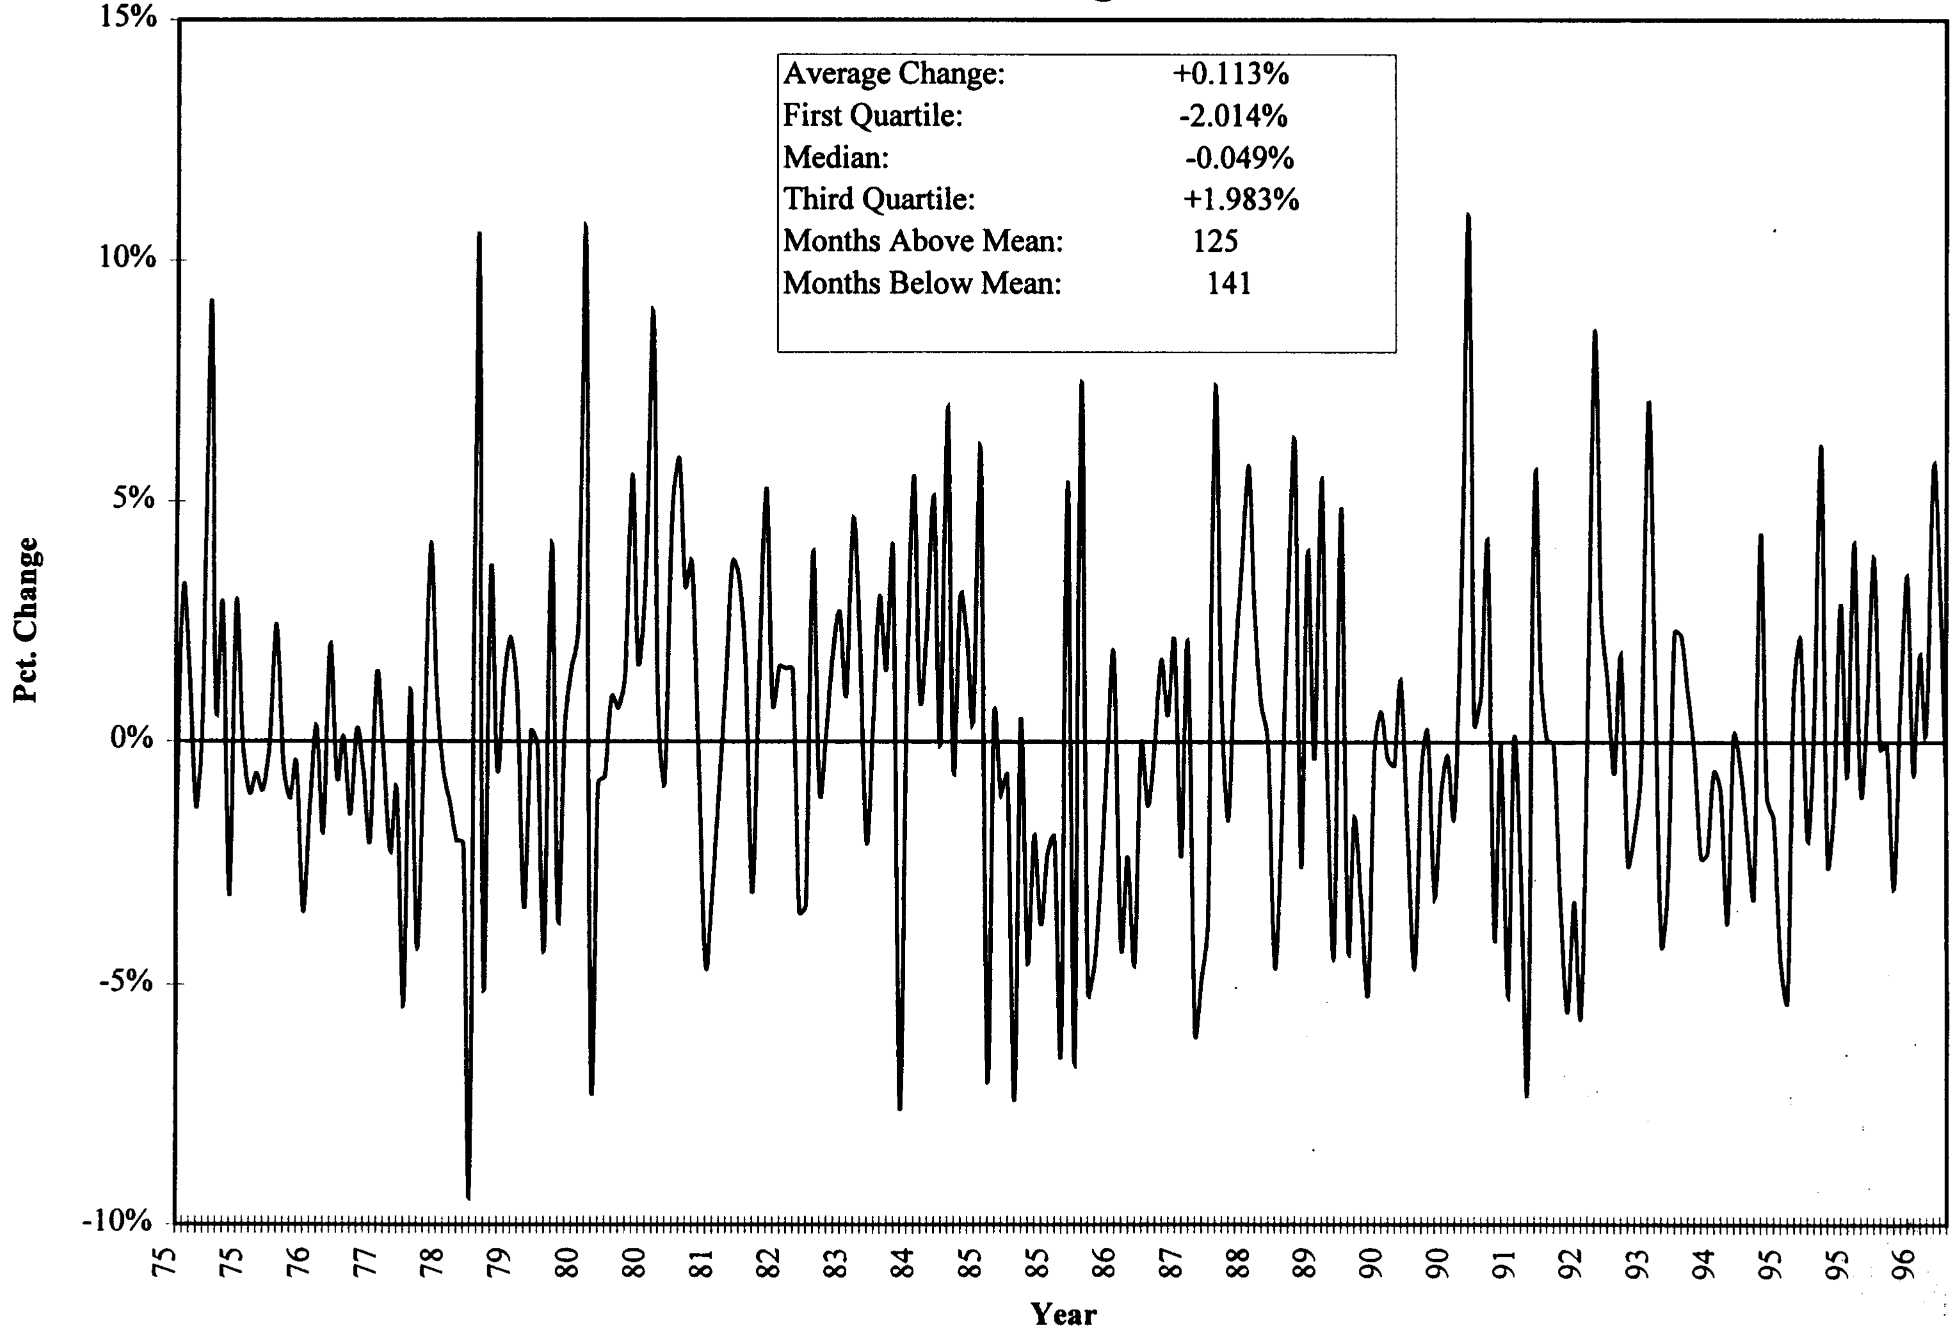

Chart 3: S&P 500 Monthly Total Return Less the 1 Month Euro Rate

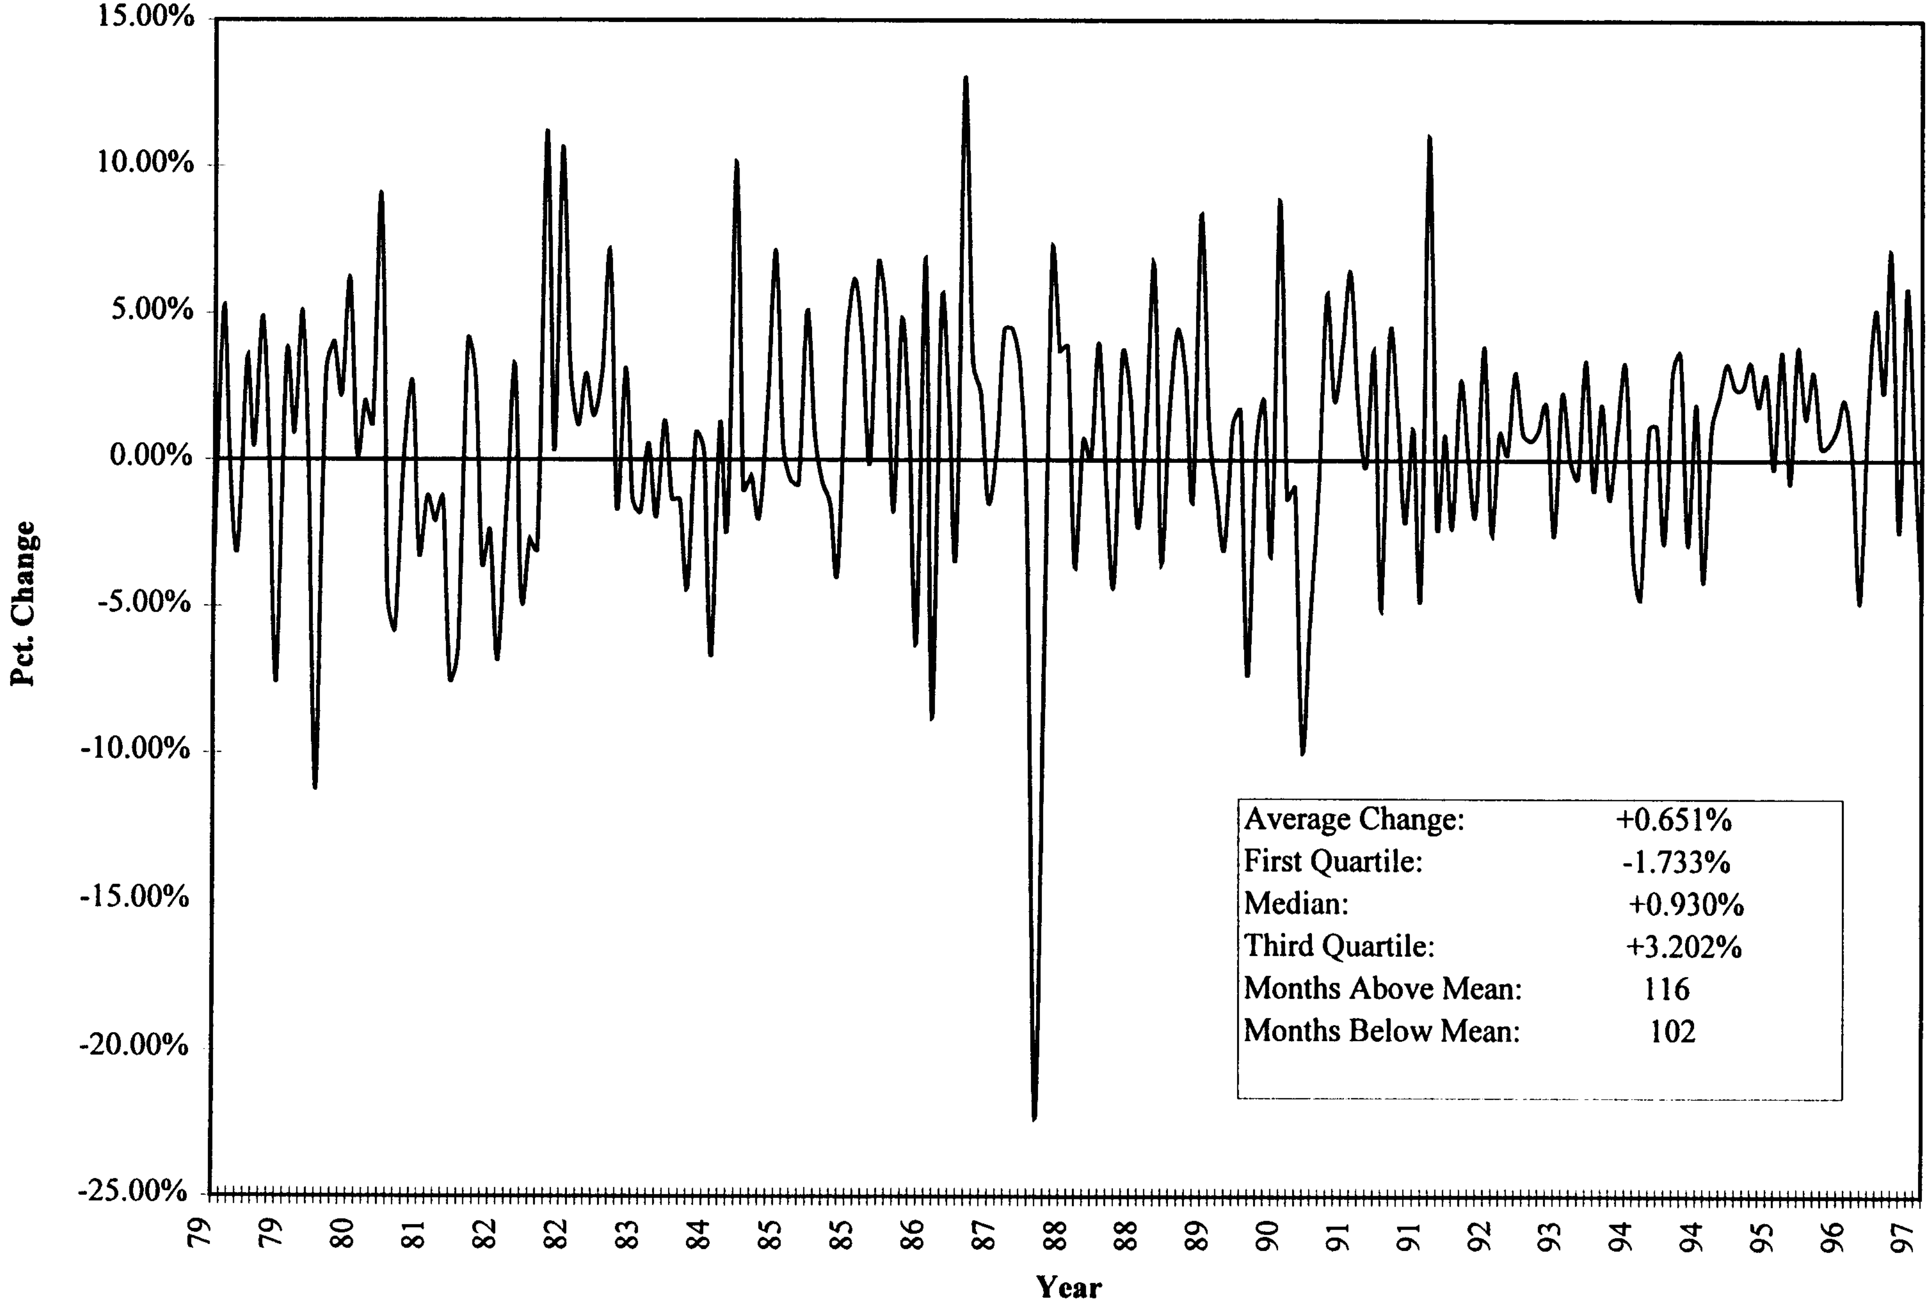

Chart 4: 2 Year-10 Year Monthly Spread Change
Forward Spread Change to Spot Spread

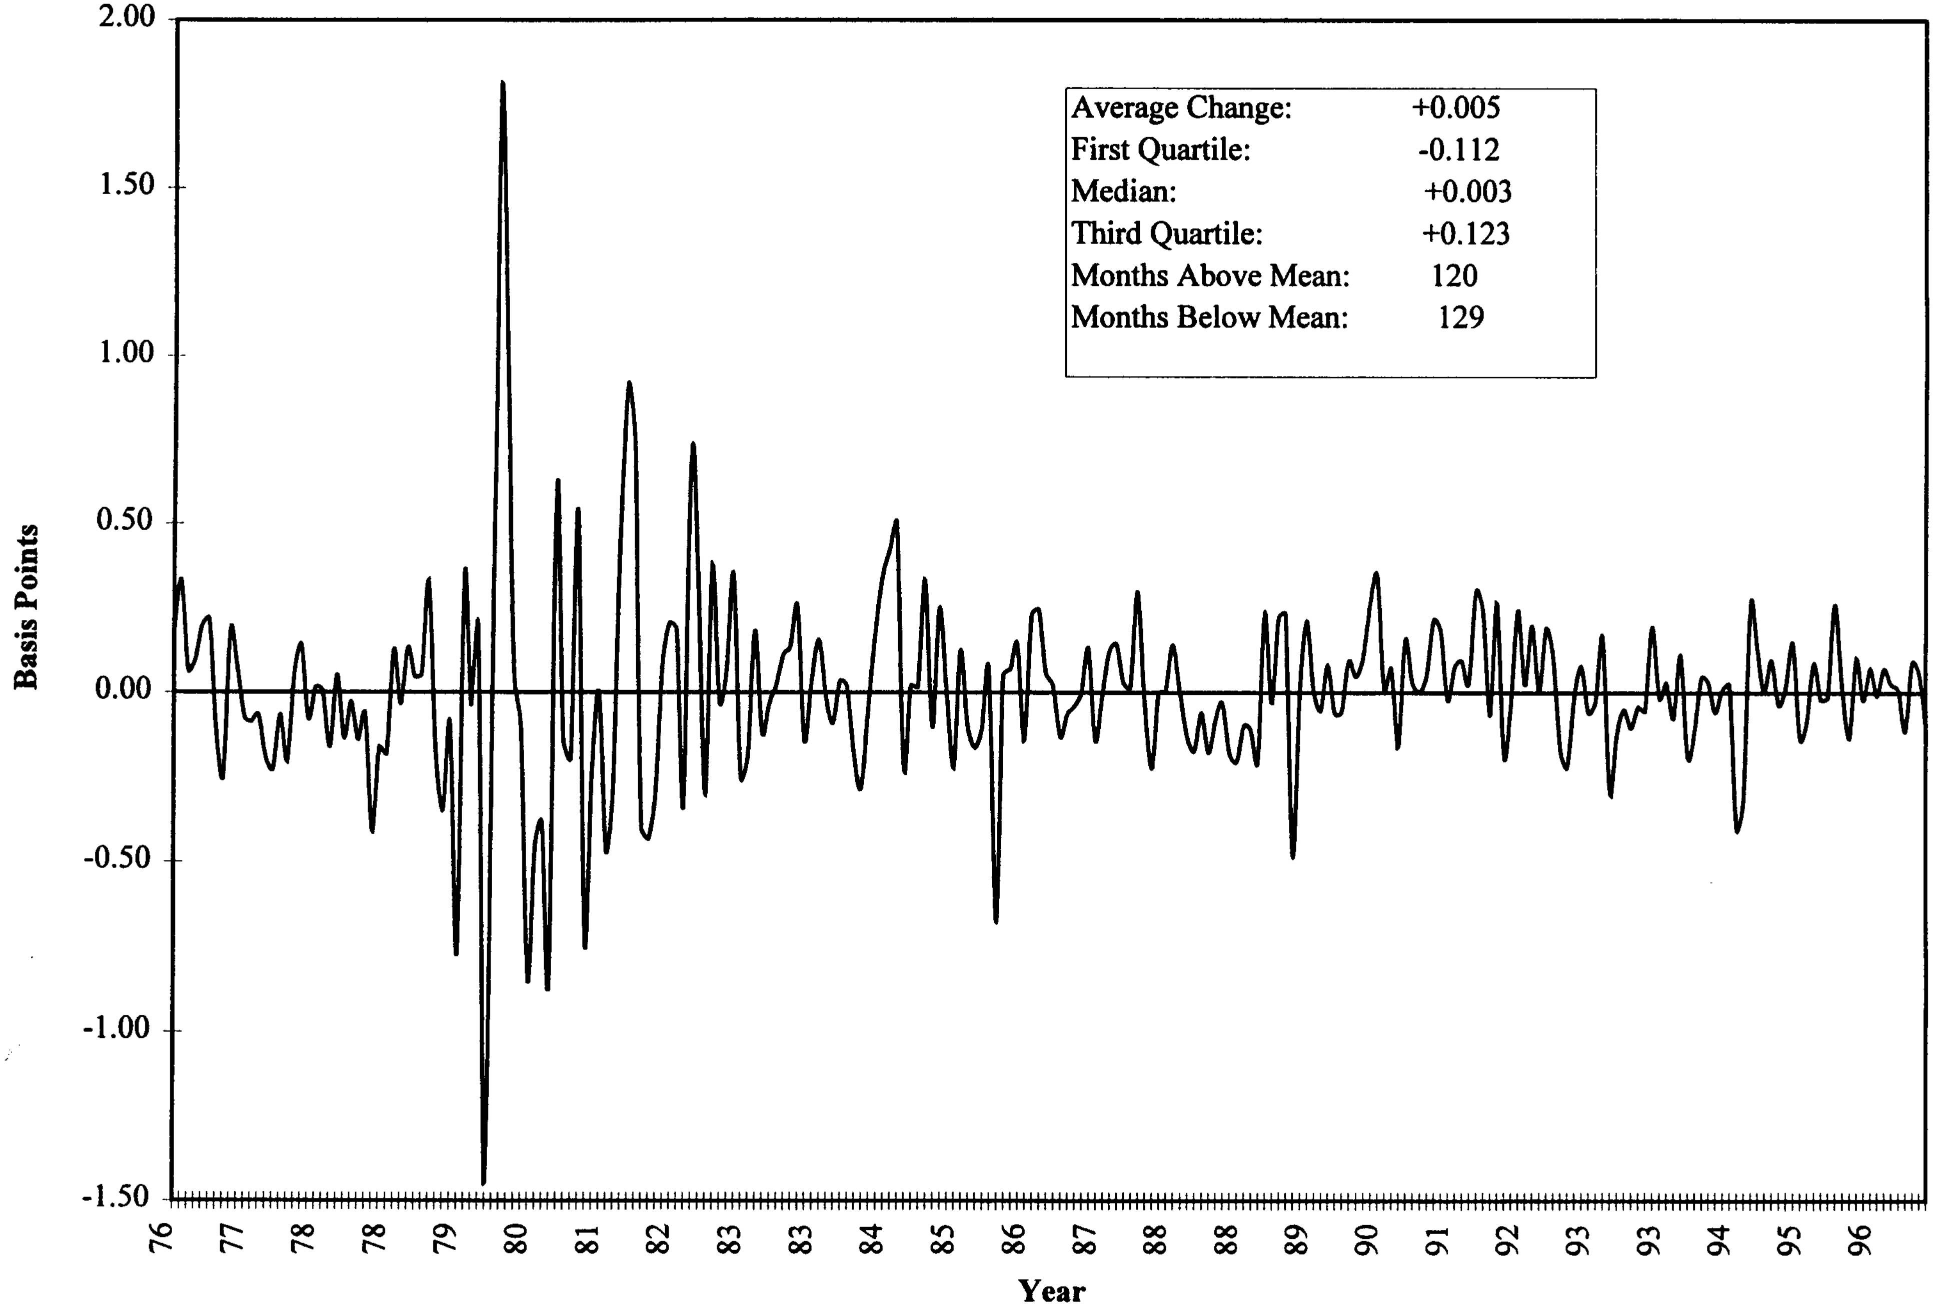

Chart 5: Monthly Changes in Actual Volatility
S&P 500 Intramonth Std. Deviations

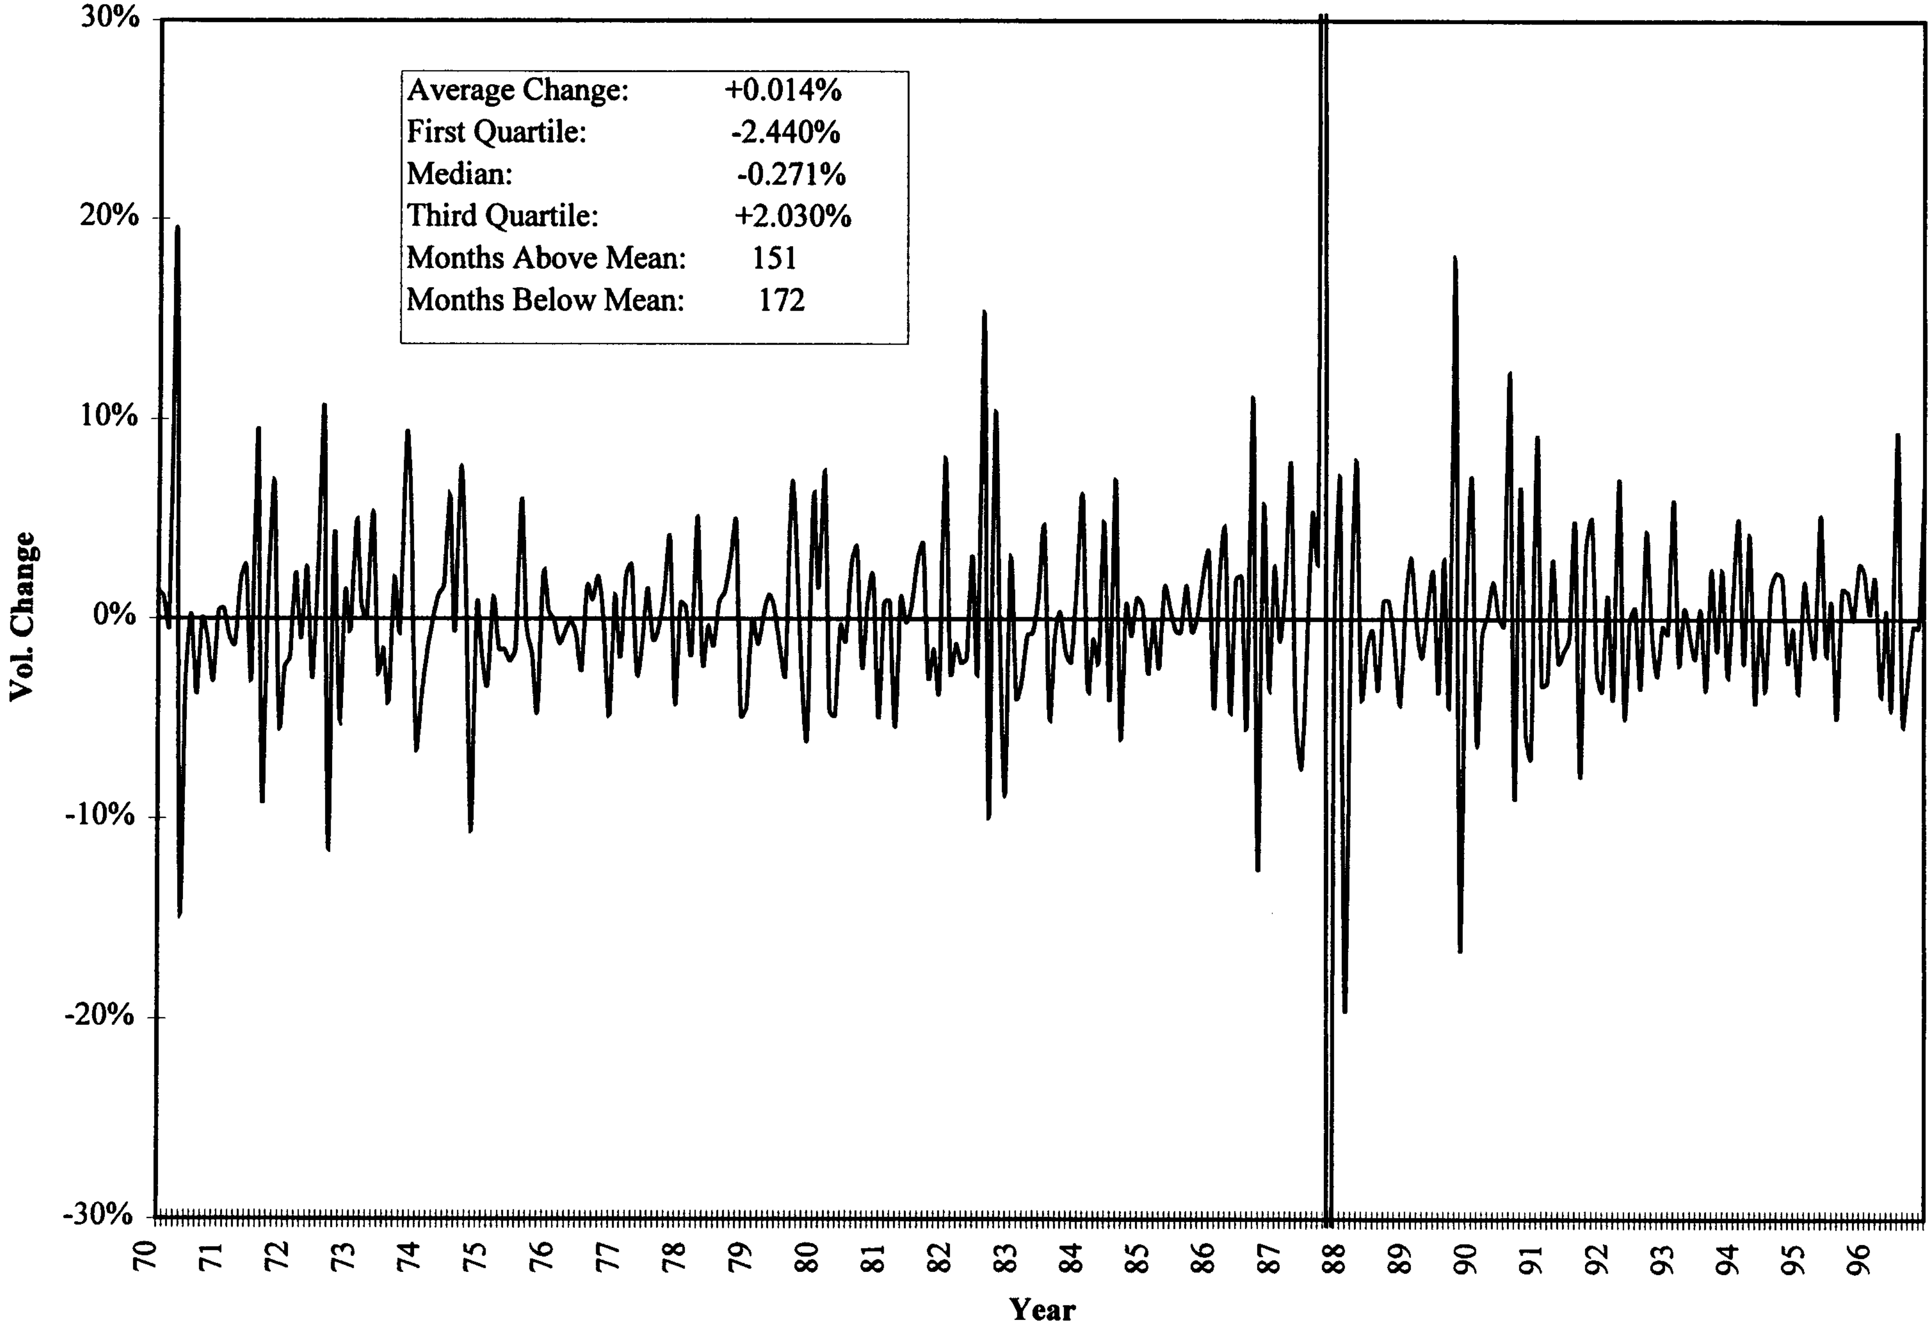

Chart 6: Overlay Performance and Correlations
Forward-Adjusted Dollar/DM Rate

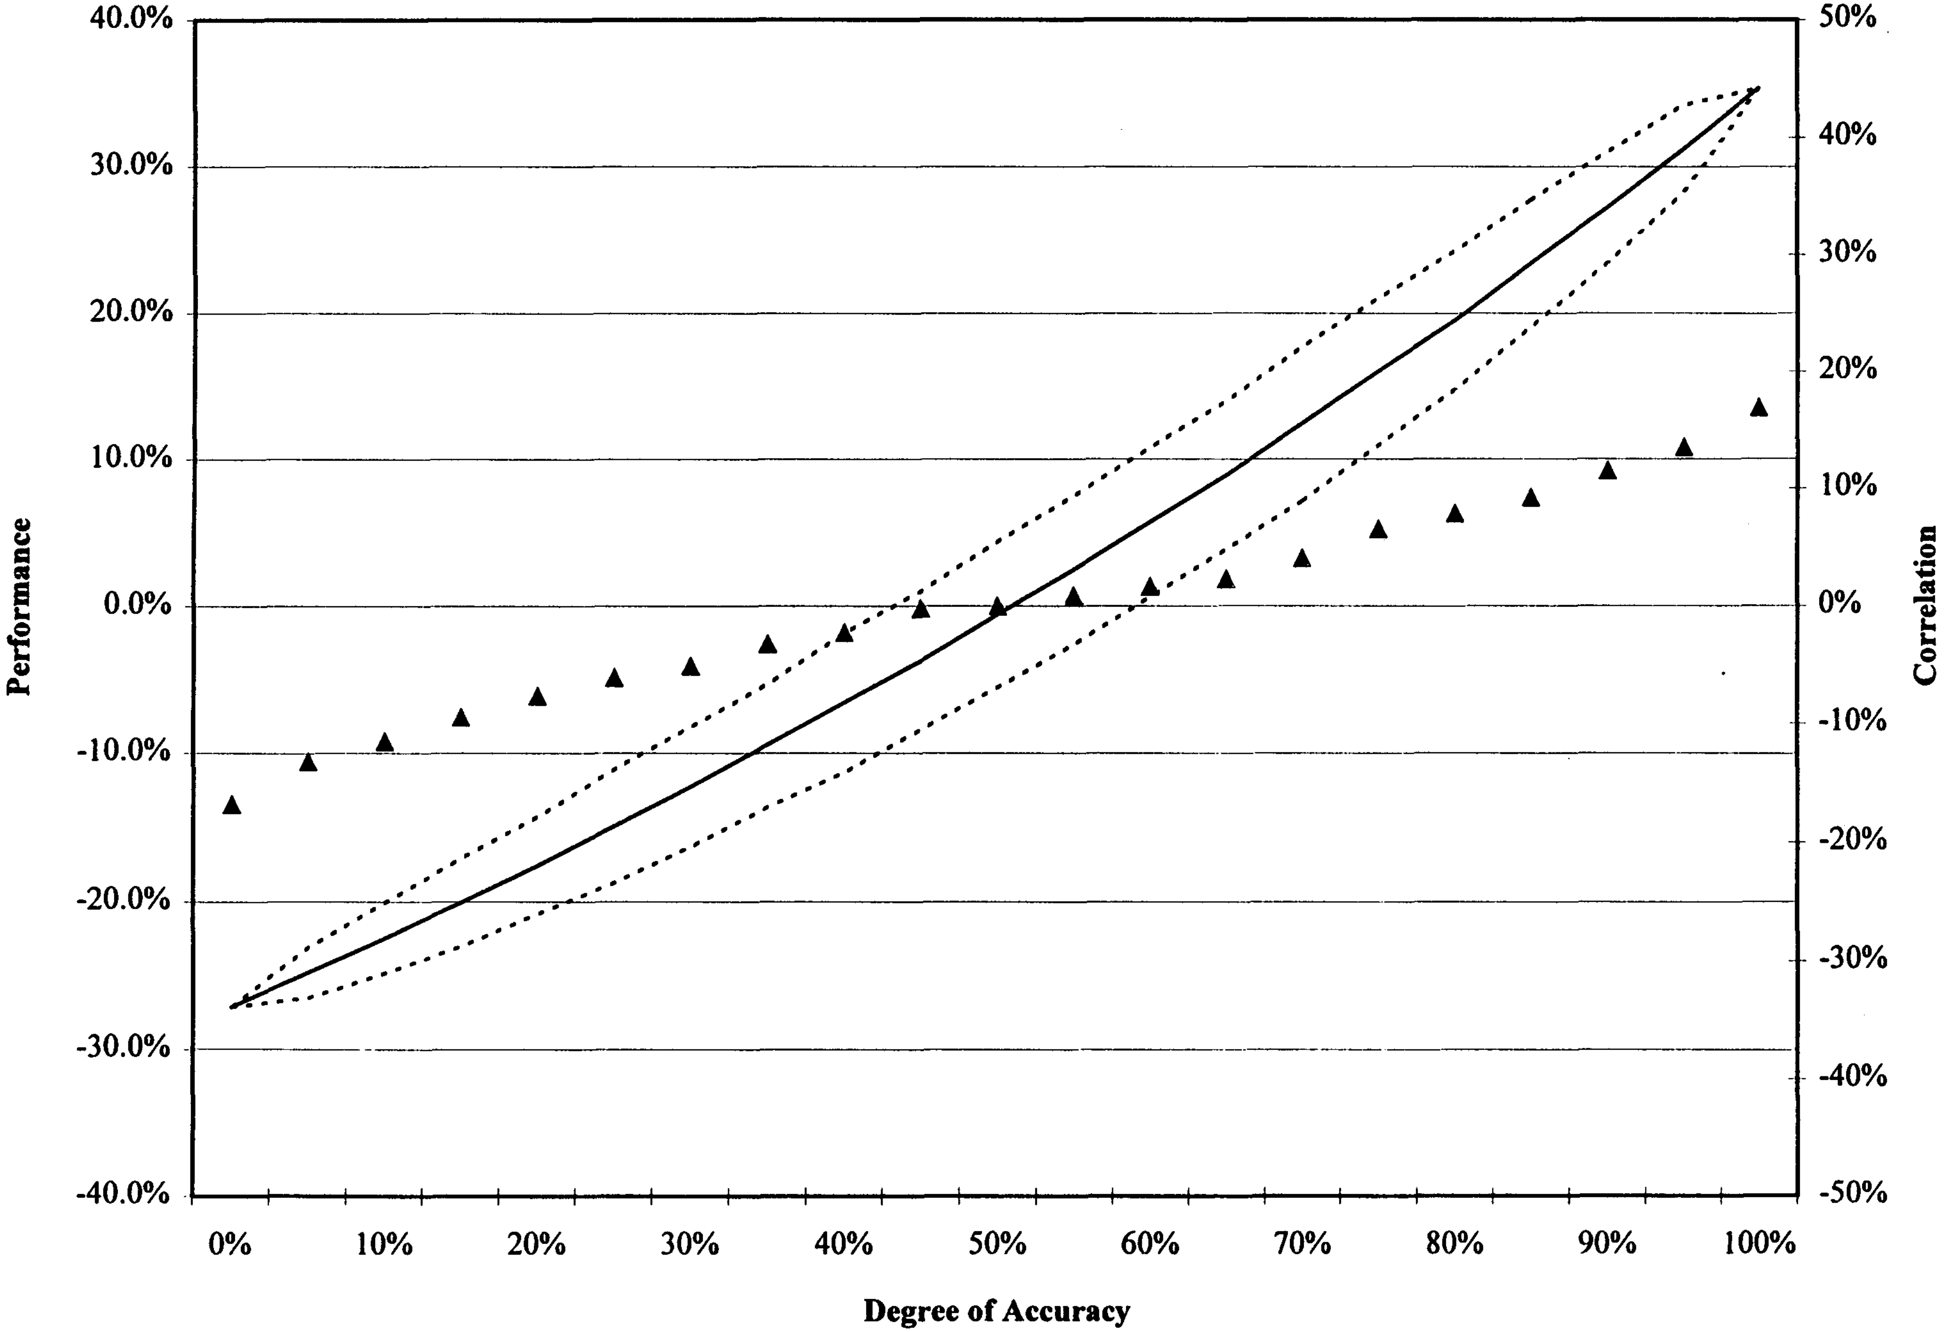

Chart 7: Overlay Performance and Correlations
S&P 500 Total Return less Riskless Rate

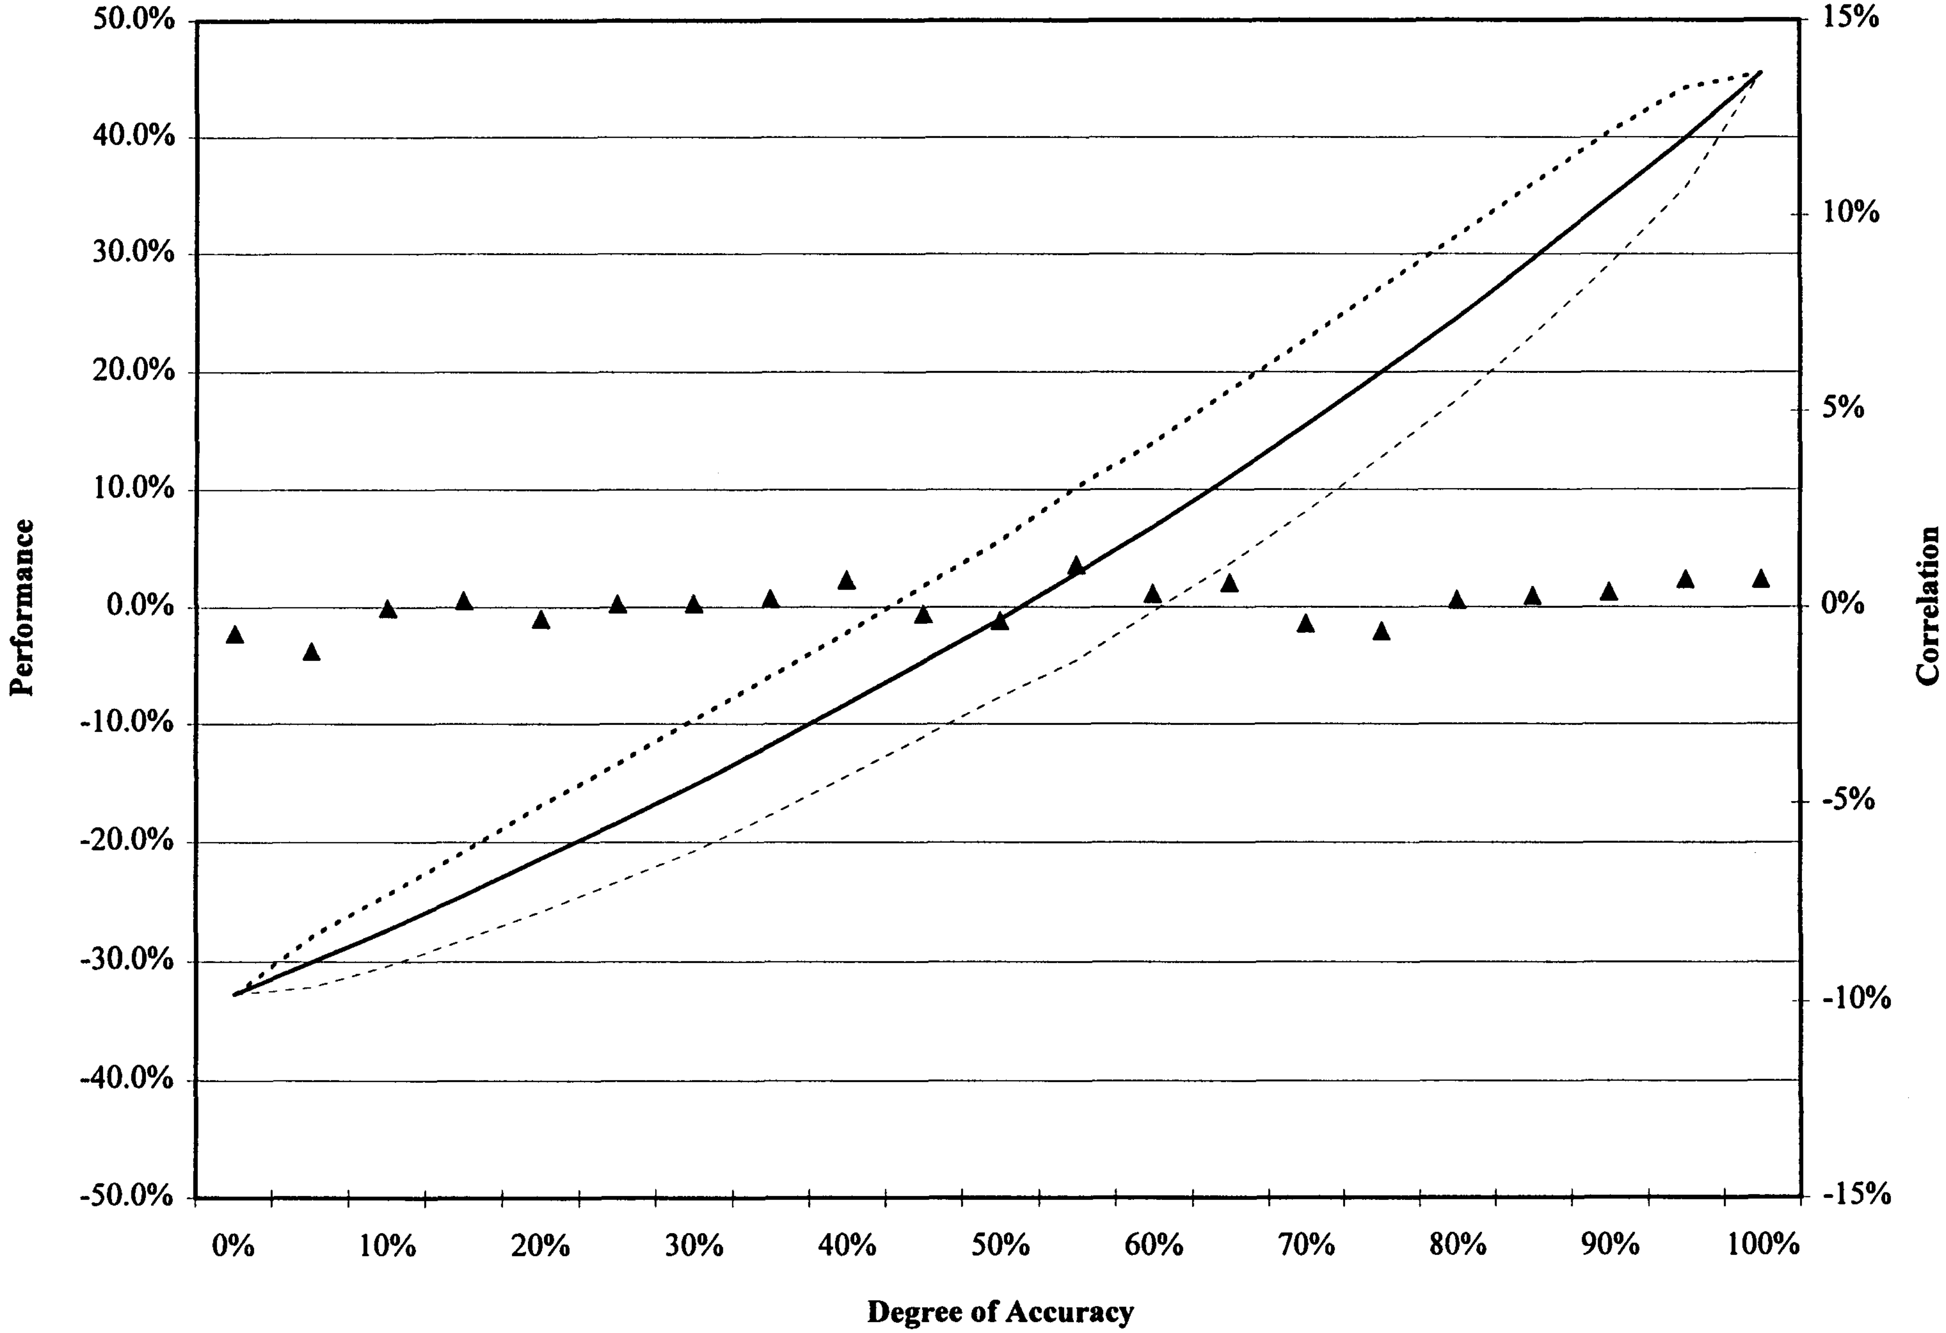

Chart 8: Overlay Performance and Correlations
Forward Adjusted 2/10 Yield Spread