Correlations and Technical Indicators in *QuanTek*

(Revised May 19, 2005)

by

Robert Murray, Ph.D.
Omicron Research Institute

Contents

According to the *Random Walk* model, stock price returns (the changes in price over a given time period, such as from one closing price to the next) are supposed to be independent, uncorrelated random variables. (More precisely, it is the *logarithmic* price returns that are usually considered. These are postulated to be Gaussian random variables.) Then the logarithmic prices, which are the sum of these independent price returns, follow a stochastic process called the **Random Walk**. The main consequence of the Random Walk hypothesis is that future returns are independent (and hence uncorrelated) with the past prices (or any other financial data, such as fundamental data). So, theoretically no function of past data can be used to predict future price returns. This is a statement of the **Efficient Market Hypothesis**, of which the Random Walk model is a special case.

If the market were perfectly efficient, then there would be no point to short-term trading. On the average, the expected return from short-term trading would be zero, relative to a buy-and-hold strategy. If the Random Walk process is one with drift, corresponding to the secular upward trend of the stock market, then the buy-and-hold strategy would give an overall average return over a long holding period (equal to the secular trend). This is presumably a reward for the risk inherent in stock investing, which is measured by the *variance* or *standard deviation* of the Random Walk over time. But short-term trading would only increase the risk, with no corresponding increase in expected returns over time. Thus it would be just like gambling, except that the expected return (over buy-and-hold) would be zero (rather than a loss, as with most gambling).

However, hardly anybody believes that the market is truly **efficient**. There are many people interested in short-term trading, and many others who are prudent to buy and sell securities over longer term holding periods, as the situation changes and different securities look more promising (based on past information). A simple, rough argument indicates that **the market can never be truly efficient**. If the market were perfectly efficient, then there would be no reward for short-term trading (or longer-term trading either), so people would stop trading. But it is precisely the trading activities, on all time scales, that keep the market efficient. Hence when the trading stops, inefficiencies would immediately be created, which would induce people to start trading again because they would then be able to make a profit. So the conclusion is that people trade to the extent that they can still make a profit, so the efficiency of the market is dictated by the ability of the best traders to still be able to make a profit (at the expense of the less knowledgeable traders). So we expect inefficiencies to exist at a level that the most sophisticated traders are just barely able to find and take advantage of them. At the present time, the market is *almost* efficient, but it can never be perfectly efficient. Profitable trading opportunities will always exist for the most sophisticated traders.

Measuring Correlation

In order to find **trading rules** that work, we must find certain functions of the past price data (and/or perhaps other financial data such as fundamental data) that have a non-zero **correlation** with **future returns**. (See the **Appendix** for the definition of **correlation**.) As we have stated, the *Random Walk* model states that this correlation should be zero. We can construct various functions and measure their correlation with future returns, or more precisely, we can measure the **sample correlation**. The sample correlation is an *estimate* of the actual correlation, based on a finite sample of data. The true correlation can only be determined in a hypothetical stochastic system in which there is an *infinite amount of data* available, and the stochastic process is **second-order stationary**, meaning that the correlation is constant for the whole data set. And here is a major problem regarding financial data: There is almost never a very large data set to work with, and within this data set it is almost certain that the stochastic process is **non-stationary**. So the measured correlation within one block of data will (probably) be different from that within other blocks of data in the same data set. Furthermore, within a finite data set, the *sample correlation* is itself subject to a statistical uncertainty. A totally random data set can yield a measured value of the sample correlation, which is non-zero, just because of random statistical fluctuations. The standard error for these fluctuations, for the usual Linear or Pearson's R correlation, is given by $1/\sqrt{N}$, where N is the number of data points in the set. (The standard error is slightly smaller for the *robust* correlation methods.) So, for a set of returns 100 days long, the standard error of the sample correlation for these returns is 10%, which would be a very sizable correlation if it existed. For a data set 1000 days long, which is the usual length of the data set that we work with, the standard error is 3.16%, which would be a small but non-negligible correlation if real. Furthermore, there are indications that **long-range correlations** only extend to a maximum of 1000 data days, or four years [Peters (1991, 1994)]. So, the conclusion is that any correlations that exist in the data, are likely to be "down in the statistical noise" and of the same order of magnitude as the statistical uncertainty of the sample correlations. Nevertheless, these small correlations, if real, can lead to very sizable returns from short-term trading.

As an example, suppose we find a technical indicator that has a 5% correlation with the 1-day future returns. Suppose the daily volatility is 2% (r.m.s. value of daily returns). Then, setting the daily trading position (*trading rules*) proportional to the technical indicator, the expected daily gain is the product of the correlation times the volatility, or 0.1%. Assuming 256 trading days per year, this leads to a simple annual gain from short-term trading of 25.6% and a compounded annual gain of 29.2% (over buy-and-hold), which most people would regard as excellent! However, by most standards the 5% correlation, given a standard error of 3.16%, would not even be regarded as **statistically significant**. The conclusion is that if we want to find trading rules that work, we have to search for correlations that are barely above the statistical "noise" level, and as a result we must also accept that the standard error of the gains (from short-term trading) will inevitably be of the same order of magnitude as the gains themselves. Nevertheless, if the short-term trading is done within the setting of an overall portfolio strategy, the standard error for short-term trading for the whole portfolio can be reduced while the returns remain the same. In this case the standard error of the returns will be reduced roughly by a factor $1/\sqrt{N}$, where now N is the number of securities in the portfolio. Of course, to get this $1/\sqrt{N}$ reduction in the standard error, it is necessary to do N times as much work!

Regarding the *statistical significance* of the correlation, the usual interpretation is that a correlation greater than two standard errors (from zero) is regarded as *significant*. A correlation this large, at least 6.32% in the example above, is achieved only 4.6% of the time by pure chance alone. (This corresponds to a 4.6% *significance level*.) So, we say that this correlation is **significant** at the 95.4% *confidence level*, because there is a 95.4% chance that this correlation is *not* due to chance alone. (We are calling the *confidence level* that quantity, which is 100% minus the *significance level*.) Theoretically, when estimating the "true" correlation by means of the sample correlation, the measured sample correlation will itself be a random variable with a Gaussian distribution of values. The standard error of this distribution is $1/\sqrt{N}$ as stated above, for a sample size N. Thus, if there is no actual correlation at all, then the measured values of the correlation will be distributed around zero, with a standard error $1/\sqrt{N}$. These values will lie within one standard error of zero 68.3% of the time, within two standard errors of zero 95.4% of the time, and within three standard errors of zero 99.7% of the time [Natenberg (1994)]. So, if the measured correlation is not at least two standard errors away from zero, it is usually regarded as not **statistically significant**. If there is no real correlation, then the distribution of measured correlation values will be a Gaussian distribution (approximately) as stated, distributed around zero. However, this does *not* mean that if the measured correlation is within two standard deviations of zero, then it is necessarily not a real correlation. All it means is that the measured correlation is *consistent* with zero correlation (to the 4.6% significance level). Most of the correlation we measure, at the "peaks" in the **Correlation Test** display in ***QuanTek***, are actually more than two standard deviations away from zero, so they can be regarded as **significant**. However, we prefer the following interpretation, which seems more reasonable: The measured correlation represents the *mean* or **expected value** of the actual correlation, and this value is *uncertain* by an amount given by the **standard error**, $1/\sqrt{N}$. In this way we are not forced to ignore measured correlations that are within two standard deviations of zero, and then "define" them to be zero. We regard the measured correlations to be the *most likely value* of the actual correlations, subject to a rather wide uncertainty given by $1/\sqrt{N}$. If Edgar Peters (1991, 1994) is correct and the correlations do not persist longer than 1000 days or so, then we cannot reduce this statistical uncertainty any lower than about 3% by taking a larger data set, so there is never any way to conclusively separate the correlations we are seeking from the stochastic uncertainty of the sample correlation measurement. Nevertheless, these correlations, provided they are really there (which they *seem* to be), can still be used to construct profitable (over the long term) short-term trading rules.

The ultimate point is that there is no "**Law of Large Numbers**", or mathematical limit as $N \to \infty$, that we can take in order to prove conclusively the **existence of correlation**, or measure the sample correlations to **arbitrarily high confidence levels**. This limit might be approximated by finding some trading rule, and testing it on a whole portfolio of stocks over a long period of time, say 1000 days. In this way, we may finally be able to find an unambiguous signal for a highly statistically significant correlation, and the portfolio ensemble then plays the role of the **very large statistical ensemble**. But such a calculation might take hours or days to perform, and I have not yet attempted such long calculations. In the meantime, it is still necessary to apply a certain amount of **intuition** in deciding which are meaningful correlations and which are just stochastic noise in the measured correlations (when constructing technical indicators using ***QuanTek***). So it is still a creative process, requiring some thinking and judgment, to choose an effective set of trading rules using the statistical tools in the ***QuanTek*** program. Short-term trading cannot and should not be reduced to a mere mechanical operation, at least in my view. (Having said this, I should add that the ***QuanTek*** program yields some rather clear signals for correlations between certain technical indicators and future returns, which certainly do not *look like* stochastic noise. But there is no statistical test that can prove *conclusively* that they are real correlations. Without an infinite data set, or at least a very large one, it is *impossible* to prove *anything* conclusively from Statistics.)

Technical Indicators

The usual definition of a **technical indicator** is some function of the past price data, which "signals" a buy or sell point. As a prototype, one of the most commonly used technical indicators is a combination of two (exponential, say) moving averages, one with a longer time scale than the other. When the shorter MA crosses the longer MA moving upward, this is a buy signal, and when the shorter MA crosses the longer MA moving downward, this is a sell signal. The expectation is that as long as the shorter MA is above the longer MA, the prices will be in an up-trend, and as long as the shorter MA is below the longer MA, the prices will be in a down-trend [Pring (1991)]. (Evidently there is an assumption here that the prices will be in one of two modes, either bull or bear market, and that these modes will last much longer than the time scale of the moving averages themselves.) Equivalently, we can form an *oscillator* from the two MA's, by subtracting the longer one from the shorter one (assuming logarithmic price data). This is a logarithmic version of an oscillator called the **Moving Average Convergence-Divergence (MACD)**, in which the ratio of two exponential MA's of the price data is taken [Pring (1991)]. Then the buy/sell points are marked by the points at which this MACD crosses the zero line, moving up or down respectively.

I would like to remark here that, in my opinion, most of the traditional rules of **Technical Analysis** are probably obsolete. They probably worked well in decades past, when there were far fewer players in the market and the rate of information exchange was much slower and the amount of information available much less. The markets are undoubtedly much more efficient now than they were when these traditional rules were first formulated [Edwards & Magee (1992)]. In particular, the ability to signal a long-term trend change by the crossing of two MA's of much shorter time scale seems "too good to be true", as do the other methods of signaling a trend change by means of technical patterns of short time duration. Probably in today's market the predictive power of any technical indicator formed from price data over a certain time scale is only of the order of that time scale. I will have more to say on this shortly, in connection with the *phase relationships* of the technical indicators.

I would now like to make a slight generalization of the concept of technical indicator, and regard a **technical indicator** as any function of the past prices, which is supposed to be *correlated* with *future returns*. So, for example, the implication is that the oscillator formed from two moving averages will be above zero when the (intermediate or long-term) future returns are positive, and below zero when they are negative. In other words, there is expected to be a **positive correlation** between this oscillator and the future returns over some time interval N. It is possible to form a whole variety of technical indicators of this sort, and measure their correlation with N-day future returns to determine their effectiveness. Then, either a **linear** trading rule can be used in which the position in the security is adjusted to linearly follow the value of the indicator, or a **non-linear** trading rule can be used in which the position is *long* by a fixed amount when the indicator is positive and *short* by a fixed amount when the indicator is negative. (This latter trading rule, of course, requires far fewer trades.) Likewise, the indicator itself can be a **linear** function of the past returns, such as MA's or sums and differences of MA's, or it can be a **non-linear** function of past data, such as polynomials or the hyperbolic tangent function or the error function. By using non-linear functions of the data and measuring their linear correlation with future returns, we are actually capturing some of the **higher-order statistics** of the data, which is probably important for financial data. However, for the time being we will confine the discussion to various linear combinations of various types of smoothings of the past data. However, our method can be extended to non-linear functions of the past data simply by defining and using such functions instead of linear ones. Evidently, the traditional technical indicators themselves may be regarded as very complicated non-linear functions of the past price data. Examples of this would be support/resistance levels, head and shoulders tops and bottoms, triangles, rectangles, flags, and so forth, and even trend lines for bull and bear trends. However, once again I question the validity of some of these patterns in today's market.

There appear to be two basic categories of **technical indicators**, corresponding to two basic categories of **correlation**. The most basic correlation is what is known as **return to the mean**. This implies that there is some *mean* or "correct" price, which the security returns to if the security becomes mis-priced. So, if the price is below some average level, it can be expected to move higher, and if it is above the average level, it can be expected to move lower. So the technical indicator consists of the current price relative to some longer-term average or smoothed price. The future returns are then expected to be *anti-correlated* with this indicator (or correlated with the negative of the indicator). Since the security becomes mis-priced in the first place after some up- or down- move, the presence of a return to the mean mechanism also shows up in the anti-correlation of past *returns* with future returns. There is a rather pronounced **anti-correlation** in daily returns up to about three days in the past, with the future one-day returns, and this can be explained by the return to the mean mechanism acting over these very short time intervals. It also appears to act over much longer time intervals as well. This appears to be a consequence of **fractal statistics** [Peters (1991)]. Note that this mechanism is nothing other than the famous "**Buy low – Sell high**" strategy.

A second correlation is known as **trend persistence**. This correlation corresponds to the tendency of the market to remain in either a **bull** or **bear** market. In other words, if returns are positive or negative in the past, they are the same in the future, so that there is a **positive correlation** between past and future returns. This mechanism would seem to be at variance with the return to the mean mechanism, which implies **negative correlation**. However, these two mechanisms can be reconciled by supposing that the "mean" is some smooth, *slowly* varying function of past prices and economic data. The **trend**, corresponding to a bull or bear market, is **persistent** and is related to the (usually) slowly varying rate of change of this price mean. (Or it can be thought of as the mean value of the returns.) Then, the shorter-term **fluctuations** about this mean price level are **anti-persistent**, and correspond to the **return-to-the-mean** mechanism. So, given any time scale, we may smooth the price data on this time scale, and then suppose that the smoothed long-term trend is persistent, and the short-term fluctuations about this trend are anti-persistent. Evidently in an efficient market, these two mechanisms "cancel out", leading to zero correlation and neither persistence nor anti-persistence of returns. But when inefficiencies exist, they do not cancel out, and evidently may exist simultaneously on different time scales. Evidently the true situation is much more complicated than this, and what has just been said should be regarded as merely an oversimplified "sketch" of the true picture. To our knowledge, nobody has yet formulated a complete theory of stock price correlations, although steps in this direction are outlined in The Econometrics of Financial Markets [CLM (1997)].

A possible third correlation does not really have a name, but we will call it the presence of **turning points** or **trend reversal** mechanism. According to this idea, if we can identify the turning points or changes of trend of the price data, then this will be correlated with a future positive or negative trend. In other words, if we can identify a point where the trend seems to change from negative to positive, then this should be correlated with a future positive trend, and a point where the trend seems to change from positive to negative should be correlated with a future negative trend. Examples of these change-of-trend indicators in traditional Technical Analysis are identification of top and bottom formations such as **head-and-shoulders**. However, we may also construct an oscillator-type indicator by taking the rate-of-change of the returns, which is itself a rate-of-change of the log prices themselves. In other words, the returns are the ***first derivative (velocity)*** of the log prices, while the *turning point indicator* is the ***second derivative (acceleration)*** of the log prices. The hypothesis is then that this *turning point indicator* is correlated with future returns, at some point in the future. However, this indicator may be less reliable than the first two, because it tends to emphasize the higher frequency modes, while most of the correlation seems to exist in the low frequency modes (explained in the next section).

Smoothing and Stochastic Noise

The **Random Walk** model may be thought of as a model in which each price movement is an independent *random shock*. Such a *random shock* is presumably the result of some business development or news input regarding the security. However, probably a more realistic model of stock price behavior is that it is due to **random shocks** occurring at infrequent intervals, and in between the shocks the price action is due to **investor reaction** to these shocks. This investor reaction is not instantaneous, in the real world, so the market is not perfectly efficient. The investors react to the shocks and the present state of the market with some finite time delay, which is of the order of the investment horizon of that investor. Also, many investors do not know how to properly interpret the present condition of the market, so they over-react and cause prices to swing above or below their "fair value". This combination of inefficiencies should cause some sort of dynamical behavior of asset prices in response to the *shocks* due to external influences, such as the state of the company itself or of the overall economy, or political events. So, we have a set of shocks, with large shocks occurring at infrequent intervals, and smaller shocks occurring more frequently, according to some power spectrum, say, and a dynamical reaction to the shocks, which is delayed in time according to the time horizon of each investor. So the result is a spectrum of **unpredictable random shocks**, and of **predictable dynamical responses** to those shocks. It is these dynamical responses that Technical Analysis hopes to capitalize on by means of various indicators. But the point is that, due to the finite response time of investors, the deterministic part of the price patterns are, to some extent, **smooth and slowly varying**. (For a partial theory of correlation in price tick data, see The Econometrics of Financial Markets [CLM (1997)].)

Hence we may postulate a model for stock price action. It consists of a **deterministic** part, which can be predicted (in principle, if not in practice), which is smooth and slowly varying, and hence consists of the **lower-frequency Fourier components** of the returns process. To this is added a **random** part, which may be modeled as **stochastic white noise**, with a constant spectrum. Thus most of the high-frequency variation of prices is random, stochastic noise with very little predictive power. (However, an exception to this is the apparent anti-correlation of returns over time intervals of a few days.) In order to uncover the predictable, deterministic part, it is necessary to employ **smoothing** to filter out the high-frequency components. Otherwise, the small correlations in the low-frequency deterministic part are completely drowned out in the high-frequency noise and cannot be seen. This is probably why it has been found so many times that the stock price data are statistically a Random Walk, and no clear deviations from the Random Walk can be seen by the classical statistical tests. After smoothing the data, however, we do find some clear indications of usable correlations, although it should be emphasized that these are hardly ever very far above the level of the stochastic noise.

At present there are three main types of **smoothing** used in the ***QuanTek*** program. The main type of smoothing used is called the **Savitzky-Golay smoothing filter**. This is a state-of-the-art digital smoothing filter, which has the property that it preserves the first and second moments of the price data. (In other words, if there are peaks in the data, the smoothing preserves the positions of the peaks and also their widths.) This filter uses **Fourier** methods to compute the smoothing, and it turns out that it can also be used to make an extrapolation into the future based on the past data. Essentially, the filter decomposes the data set into its Fourier components, filters out the high frequency components, and then the extrapolation consists of extending the low-frequency components that are left forward in time, preserving their phase relationships. It appears that this extrapolation itself has predictive power, indicating that these low-frequency components persist in time, at least out to one wavelength or so for each component. The Savitzky-Golay smoothing filter itself comes in two variants, the **acausal** and **causal** filters. The *acausal* filter smooths over a time window consisting of a number of days *in the past and future* around the given day, equal to the smoothing time period. (Hence, half the period of the dominant cycle is equal to two smoothing periods.) This acausal filter should have the advantage that it preserves the phase relationships of the various Fourier components. The *causal* smoothing filter smooths over a time window equal to two smoothing periods *in the past*. Hence there is an inherent time delay of (approximately) one smoothing period with the causal filter. This causal filter will probably not preserve phase relationships, which is a disadvantage. The third type of smoothing is the ordinary exponential Moving Average (MA). This type of smoothing filter is also causal, in that it does not make use of any data in the future relative to the given day. As is well known, the exponential MA also introduces a time delay of the order of one smoothing period (for a time scale of smoothing of two time periods). However, the exponential MA does not (itself) make any future extrapolation, even though technically it is a digital filter just like the Savitzky-Golay smoothing filter. This is because it does not use the Fourier transform. (We can, of course, make the future extrapolation using **Linear Prediction** or some other prediction filter.)

The **Savitzky-Golay digital filter** also has the capability of computing the smoothed *first and second (and higher) derivatives* of the price data. Thus using this filter we may directly compute the smoothed *velocity* and *acceleration* indicators mentioned above. Using the **acausal** filter, these smoothed velocities and accelerations should exhibit no time delay, and be in phase with the price data. Using the **causal** filter, on the other hand, will introduce a time delay and the various Fourier components will be out of phase to some extent, which depends on the **frequency and phase response** of the filter. This is why I prefer the acausal filtering to identify the **buy/sell points**. The time delay of the causal filter depends on the **order** of the filter which is chosen, with the lowest order causal SG filter being equivalent to an ordinary (not exponential) MA. The higher order causal SG filters still are causal and use only the past data, but they do not exhibit as much time delay, because it is compensated for in the way the data are smoothed. In other words, within a block of data $2N+1$ units long, the usual MA and lowest order SG filter fit the data to a constant average value within this block, but higher order SG filters fit the data to a straight line, parabola, and so forth. This has the effect of essentially eliminating the time delay, although this time delay is eliminated through a polynomial fit to the data. However, the phases of the various Fourier components are still evidently not preserved using the causal filter, unlike the case of acausal filtering.

Types of Technical Indicators used in *QuanTek*

There are three main types of technical indicators used by the ***QuanTek*** program. These correspond to the three main types of correlation mentioned above, namely **return to the mean** and **trend persistence**, plus **turning point** or **trend reversal**. The first type of indicator is called the ***Relative Price***. This is a difference of the (logarithmic) price levels with smoothings on two different time scales, the shorter time scale minus the longer time scale. This is similar to the oscillator consisting of the difference of two exponential MAs mentioned previously (except without the time lag). This type of indicator is a measure of the **return to the mean** mechanism, with the longer time period smoothed price level playing the role of the mean level. When the shorter time period smoothed price level is below the longer period one, the future prices are expected to *rise*, and when it is above, the future prices are expected to *fall*. There is a certain time delay here, which is of the order of the shorter smoothing time period, in which the trough or peak of this indicator *now* implies that the future returns will be positive or negative *later*, roughly by this time delay. So there is a **phase difference** between the *Relative Price* indicator and the future returns that it is supposed to predict. The negative of this indicator *leads* the expected future returns by *approximately* one time period.

The second type of indicator is called the ***Velocity***. It is the smoothed **first derivative** of the log prices, or the difference of two smoothed first derivatives with different time scales. This is the kind of indicator, which is normally called a *Momentum* indicator, but we reserve the term *Momentum indicator* for any indicator that is supposed to be positively correlated with returns, not just the *Velocity*. It is clear that if the trend is persistent, then the smoothed velocity of the log prices should be correlated with the returns. The smoothed velocity *now* should be correlated with the returns *now* and in the *immediate past and future*. So the Velocity indicator is *in phase* with the returns. The Velocity indicator may be constructed directly by using the **Savitzky-Golay smoothing filter**. It may also be constructed by taking an **exponential MA** of the log price *returns* rather than the log prices themselves (and then compensating for the time lag).

We may also construct a third type of technical indicator that is called the ***Acceleration***. It is the smoothed **second derivative** of the log prices or the difference of two smoothed second derivatives with different time scales. This indicator may be interpreted as an indicator of **turning points**, because the second derivative is positive when the prices are at a minimum (positive or upward curvature) and is negative when they are at a maximum (negative or downward curvature). Hence it can be seen that this *Acceleration* indicator will be *positive* when the *Relative Price* is *negative*, and vice-versa. Thus the Acceleration indicator is exactly *out of phase* with the Relative Price (with acausal SG smoothing). However, the Acceleration differs from the Relative Price in that, with each successive derivative, the high frequency components are emphasized more and more. Hence the Acceleration indicator contains much more high-frequency components than the Relative Price and hence is much less smooth. Evidently, after a minimum of the prices, we expect the future returns to become positive and reach a maximum after a certain time delay, which is roughly one smoothing time period. But at the minimum the smoothed returns are zero (by definition of a minimum), and the Acceleration is at a positive peak. So the Acceleration is out of phase with the returns by roughly one smoothing time period, and this indicator *leads* the returns. The Relative Price, being out of phase with the Acceleration, therefore *lags* the returns by one smoothing time period.

Normally, short-term trading presumes some kind of ***cyclic*** or ***oscillatory*** behavior of the prices. The goal is to buy at price minima, and sell on price maxima, on some arbitrary time scale. Also, the Savitzky-Golay smoothing method, and all others, is fundamentally based on the concept of Fourier analysis. The smoothing filters out the higher-frequency components of the signal, which have a period shorter than the smoothing time scale, leaving only the lower-frequency components. Then, on this smoothing time scale, the dominant frequency is the highest frequency that is left, which has a period of, let us say, four time periods (if the smoothing time scale is two time periods, corresponding to half a cycle). Then all the above indicators have an oscillatory behavior with a period of four smoothing time periods. Then we may specify the phase relationships between these indicators.

Momentum Indicator

We define a ***Momentum*** indicator as any function of the *past* prices that is supposed to show a positive correlation with *future returns*. The display of the *Momentum* indicator should be such that it is *in phase* with the ***Velocity***, and hence the returns. The value of the ***Momentum*** indicator for each day may then be interpreted as our *estimate*, based on past price data, of the *return* for that day. This then translates directly into ***Trading Rules***. The position should be varied according to the *returns* to be expected, so the position should be *positive* when the ***Momentum*** indicator is positive, and should be *negative* when the ***Momentum*** indicator is negative.

Velocity Indicator

We have just seen that the ***Velocity*** indicator is in phase with the *returns*, so it is itself a ***Momentum*** indicator. The estimate of the ***Velocity*** indicator for *future returns*, up to *N* days in the future, based on the **Price Projection**, provides the values of the ***Momentum*** indicator up to *N* days in the future. We may also subtract the ***Velocity*** two smoothing time periods in the past, since it should be out of phase with the present returns. This delayed negative *Velocity* indicator is then also a ***Momentum*** indicator. The negative *Velocity* at the *present time* should then be correlated with the *future returns* two time periods in the future.

Acceleration Indicator

The ***Acceleration*** indicator *leads* the present returns (***Velocity***) by one-quarter cycle, or 90 degrees, so it is one smoothing time period ahead of the returns. The ***Acceleration*** should therefore be *in phase* with the *future returns* one time period in the future. Hence the ***Acceleration*** one time period in the past is a ***Momentum*** indicator, since it will be *in phase* with the returns. In other words, to form a ***Momentum*** indicator from the ***Acceleration*** indicator, we take values of the ***Acceleration*** indicator *N* days in the past, where *N* is the smoothing time period.

Relative Price Indicator

As for the ***Relative Price***, we could also reverse the sign of this indicator and take its one time period (*N*-day, 90 degrees) delayed values to get a ***Momentum*** indicator, since it would then be *in phase* with the ***Acceleration***. As it stands, the ***Relative Price*** should be one time period *behind* the *present* returns, and hence the *N*-day future projection of the ***Relative Price*** should be *in phase* with the returns. Thus the one time period future projection of the ***Relative Price*** is also a ***Momentum*** indicator. We may extrapolate all of these indicators into the future (still using only *past data*) by making some sort of extrapolation of the prices. This may be done using a **Linear Prediction** filter of some sort, or it may actually be done using the **Savitzky-Golay smoothing filter** itself, just by using a linear trend-line extrapolation and then smoothing. (We can construct a technical indicator using any function of the data we want, including future extrapolations, so long as we use only *past data*. This includes future extrapolations using Linear Prediction or other types of digital filters.) Then, using this extrapolated technical indicator, we may take its extrapolated values one time period in the future, minus the value one time period in the past, for the ***Relative Price*** indicator, to be our ***Momentum*** indicator corresponding to ***Relative Price***. In reality, the phases indicated here are only approximate, and the best correlations will be obtained by adjusting these phases for each individual stock.

Buy/Sell Signals

It is implicit in the above definitions that the ***Momentum*** indicators reach a positive peak when the *expected return* is maximum. In other words, the ***Momentum*** indicators are supposed to be *surrogates* for the *expected returns* at each point in time. The optimum trading rules are to be *long* when the returns are *positive*, and *short* when the returns are *negative*. Thus, the optimum buy point is just when the ***Momentum*** indicators are crossing zero from negative to positive [Z+], and the optimum sell point is when the ***Momentum*** indicators are crossing zero from positive to negative [Z–]. More precisely, an optimal strategy would be to vary the position daily so that the position is proportional to the value of the ***Momentum*** indicator. In ***QuanTek*** a weighted sum of three ***Momentum*** indicators is taken and called the ***Trading Rules*** indicator. Then the buy/sell points are the positive/negative zero crossings [Z±] of this indicator. A continuously varying optimal trading position is also specified which is proportional to the value of the ***Trading Rules*** indicator, and ranges from somewhat greater than +100% to somewhat less than –100%. It is also easy to vary the trading strategy by simply looking at the ***Trading Rules*** and ***Momentum*** indicators, since these indicators are supposed to represent estimates of the (past and) *future returns*. You can then vary your trading strategy in accordance with these estimated future returns.

Example – MACD

One of the most basic examples of an oscillator that (reputedly) forms a ***Momentum*** indicator is the ordinary **MACD** [Pring (1991)], or difference of two MA's of (logarithmic) prices. Moving averages (let us say exponential ones) are used as technical indicators by superposing two MAs with different time scales of averaging. When the shorter time scale MA crosses the longer time scale MA moving upward, this is taken as a buy signal, and when it crosses moving downward, this is a sell signal. Hence, forming an oscillator consisting of the difference of the two MAs, the buy/sell points are marked by the zero crossings upward/downward, respectively. If there were no time lag with moving averages, then the oscillator so formed would be classified as a ***Relative Price*** indicator, because it is the difference of two smoothed prices, and the buy points would be the minima of this indicator, and the sell points the maxima. However, there is a time lag of roughly one smoothing time unit, where the smoothing time scale is two such units, and the dominant cycle hence has a period of roughly four time units. Thus there is a time lag of roughly 90 degrees, or one-quarter cycle. Since the *Relative Price* indicator already lags by one time unit the *Momentum* indicator (with acausal smoothing), this means that the MACD will lag the *Momentum* by *two* time units. Hence it should be out of phase with the *Velocity* indicator by 180 degrees. This means that, **for trading on cycles of four time units period, the MACD is exactly out of phase with the correct trading signals**. As a check, we wish to buy at price minima and sell at price maxima. The MACD corresponds to these minima and maxima, except *delayed* by one time period. Due to the delay of one quarter cycle, the *downward* zero crossing of the MACD (which is *ahead* of the price minimum) will be lined up with the actual price minimum, implying a *buy* point, and the *upward* zero crossing of the MACD (which is *ahead* of the price maximum) will be lined up with the actual price maximum, implying a *sell* point. *But this is exactly opposite to the trading signals that we are supposed to use!* Evidently, the traditional use of the MACD is confined to long-term trends that are much longer than the shorter smoothing time period, so that this indicator will be approximately in phase for these long-term cycles. But this illustrates that perhaps traditional **Technical Analysis** has been too cavalier about preserving the correct **phase relationships** between technical indicators and actual prices moves!

Trading Rules and Phase Relationships

As for ***trading rules***, these are based directly on the ***Momentum*** indicators. The goal is to buy at price minima, when the Velocity (as estimated by the Momentum indicator) is going through zero and increasing. Similarly, we want to sell at price maxima, when the Velocity (and Momentum) is going through zero and falling. So we time our optimal buy-sell points at these ***upward/downward zero crossings of the Momentum indicators***. The *amount* to buy or sell is given by the average value of the Momentum indicator between these zero crossings. This average value must be estimated from the future extrapolation of the indicator, based on **Linear Prediction** or some other type of prediction filter. Or the *amount* to buy and sell can simply be taken to be a constant, as in traditional *Technical Analysis*. Or, the amount invested can be varied on a daily basis, and made proportional to the daily value of the ***Momentum Indicator***, or of some non-linear function of it such as the hyperbolic tangent function or error function. In the ***QuanTek*** program, the trading rules are based on the daily value of the ***Trading Rules*** indicator, which is in turn formed from an adjustable sum of the three custom ***Momentum*** indicators that you can design for yourself. Then, given the daily recommended position in terms of the daily value of the ***Trading Rules*** indicator, and the optimal buy/sell points, traders can make their own decision as to actual amounts to buy and sell on any given day.

In practice, the phase relationships between these indicators and the future returns are not exactly known. It depends partly on the actual frequency and phase response of the Savitzky-Golay and other digital filters that are used to construct the indicator. Also, the above analysis was based on the premise of using the ***acausal*** Savitzky-Golay smoothing filter, which has no inherent time delay. We can also use a ***causal*** Savitzky-Golay smoothing filter, or the exponential Moving Average smoothing, which is also *causal*. In the case of a causal smoothing, there will be another time delay roughly of the order of the time period of the smoothing. So we must take into account this time delay, and take values of the indicator a number of days in the future equal to this time delay, in constructing our Momentum indicator. Thus, using the causal smoothing filter, the *Acceleration* should now be roughly in phase with the future returns, and the *Velocity* and *Relative Price* indicators must use values from their future extrapolations, by one and two time periods respectively. (Of course, to compute the *acausal* smoothing itself requires this future extrapolation, since it uses future values as well as past values.) However, we still cannot be sure of the exact phase relationships even after all this. So a provision is made in constructing the indicators to *make adjustments to the phase* of the indicators, so that the indicator is adjusted to be in phase with the “future” returns (so that the correlation is at a maximum) over the past data set. Then the hope is that the indicator will remain in phase with future returns for some time into the future. Evidently this phase relationship can persist for roughly of the order of one smoothing cycle at least, although we have not proven this.

We may summarize the ***trading rules*** for the various types of indicators in the following table. This table lists the (approximate) buy points and sell points at various points of the two types of indicators, both for *acausal SG* and *exponential MA (causal)* indicators. The buy/sell points will occur either at a *maximum (max)*, a *minimum (min)*, an *upward zero crossing (Z+)*, or a *downward zero crossing (Z–)* of the indicator. The phase lead (+) or lag (–) of the indicator relative to *returns* is in units of the smoothing time period, in which four smoothing time units make up the period of the dominant cycle. The sign in front of the indicator shows whether the positive or negative indicator is used:

Indicator	**Smoothing**	**Lead/Lag**	**Buy Point**	**Sell Point**
(+) Relative Price	Acausal SG	–1 unit	min	max
(–) Relative Price	Acausal SG	+1 unit	max	min
(+)Velocity	Acausal SG	0	Z+	Z–
(–)Velocity	Acausal SG	±2 unit	Z–	Z+
(+)Acceleration	Acausal SG	+1 unit	max	min
(–)Acceleration	Acausal SG	–1 unit	min	max
(+) Relative Price	Exp. MA	±2 unit	Z–	Z+
(–) Relative Price	Exp. MA	0	Z+	Z–
(+)Velocity	Exp. MA	–1 unit	min	max
(–)Velocity	Exp. MA	+1 unit	max	min
(+)Acceleration	Exp. MA	0	Z+	Z–
(–)Acceleration	Exp. MA	±2 unit	Z–	Z+

If the indicator *leads* the returns, then a **Momentum** indicator may be constructed by taking *past values* of the indicator, by approximately one time unit. Similarly, if the indicator *lags* the returns, then *future extrapolated values* of the indicator must be used to construct a **Momentum** indicator. (Or, you can use *past values* in such a way that the **Lead/Lag** is shifted by a multiple of –4. It appears that these past values are usually more reliable than the future projected values in most cases.)

Note that the **causal SG filter**, which is not shown, will have a time lag that is somewhere between that for the **acausal SG filter** and the **exponential MA filter**. If it were a zero-order filter, it would be equivalent to a **simple MA**, and hence would have a time lag (roughly) the same as the **exponential MA**. However, since we are actually using the *fourth-order* filter in ***QuanTek***, the time lag should be roughly the same as that for the **acausal SG filter**. Hence, when using the causal SG filter, it will probably be necessary to tweak the lead times of the indicators for best results.

The summary is that the ***Velocity*** indicator is the most direct measure of the current trend of the future returns. The ***Acceleration*** indicator is the most direct measure of the oscillatory behavior of the prices at the chosen smoothing time scale, and also depends mostly on *past* values of the indicator (as opposed to future extrapolated values). It is hence a useful timing indicator for the buy/sell points. The ***Relative Price*** indicator is the smoothest and the best indicator for the slowly varying drift of the price level. It also depends most heavily on the future price extrapolation based on the Linear Prediction, Savitzky-Golay, or other prediction filter. This is the indicator that gives the best indication of the long-term trend of price action.

Technical Indicators Dialog

The **Technical Indicators** dialog is available from the **Correlation Test** dialog, by clicking the *Indicator* button. The **Correlation Test** dialog box is available from the toolbar when you open each **Main Graph** of a stock data file, and also from the **Correlations – Stocks** dialog box. The **Technical Indicators** dialog box enables you to build a technical indicator, which is then tested for correlation with future returns in the **Correlation Test** dialog. You can also build the three **Momentum Indicators**, which are displayed in the three panes of the **Technical Indicators splitter window**. The weighted sum of these three **Momentum Indicators** then make up the **Trading Rules indicator**, which is displayed in the bottom pane of the **Trading Rules splitter window**. The **buy/sell signals** and **buy/sell points** are then derived from this **Trading Rules indicator**. They are the positive going and negative going zero crossings of this indicator, respectively. Slider bars in the **Trading and Portfolio Parameters** dialog box set the weights for the sum of the three **Momentum Indicators**.

The Technical Indicators dialog box enables you to construct a technical indicator using one of three types of smoothing filters, which are **acausal Savitzky-Golay smoothing**, **causal Savitzky-Golay smoothing**, and **Exponential Moving Average**. You can construct an indicator of one of the three types described above, which are **Relative Price**, **Velocity**, and **Acceleration**. The indicator may be constructed with a positive sign, a negative sign, or the difference of two contributions of different time scale. For each such contribution or part of the indicator, you can choose the **smoothing time scale** for the filter, and the **lead-time**, which adjusts the **phase**. So this pretty much exhausts the possibilities for oscillator-type indicators that can be constructed from past stock data.

In the **Technical Indicators** dialog box, you can also choose various combinations and types of **Linear Prediction filter** with the **Fractional Difference filter**. You can also set the **fractional difference parameter** of the **Fractional Difference filter**. It is very important that all indicators and displays use the *same* **type of filtering** with the *same* **time scale** and **fractional difference parameter**. Hence these settings apply to all the displays and indicators at once.

Correlation Test Dialog

The **Correlation Test** dialog is used to test the indicator built in the **Technical Indicators** dialog for correlation with *future returns*. The **Correlation Test** dialog can be called from the toolbar of each **Main Graph**, or from the **Correlation – Indicators** dialog, which itself is available from the **Main Window** toolbar (or from the opening dialog when you first open ***QuanTek***). From the **Correlation Test** dialog, you can call the **Technical Indicators** dialog to build a technical indicator, as described above.

The **Correlation Test** dialog has a graph, which ranges from –100 days to +100 days, with ZERO at the center of the graph. The vertical scale of the graph ranges from –100% to +100% correlation. When you construct a **technical indicator**, it is a function $f(n)$ of the past data, relative to the *present*, which corresponds to $n = 0$. Negative values of *n* correspond to the *past*, and positive values of *n* correspond to the *future*. However, **no future data** are used to construct the technical indicator. The future values are constructed from the *past* data by means of the prediction filters, consisting of the **Fractional Difference filter** and the **Linear Prediction filter**. In addition, the **Savitzky-Golay smoothing filter** itself functions as a prediction filter, out to about *N* days, where *N* is the smoothing time scale. Then, for each value of *n* ranging from –100 days to +100 days, the correlation of $f(n)$ with the *returns* over some time interval is computed. To compute the correlation, the technical indicator is computed with each day in the past as the *present* day, and the correlation with the *future* relative to that present day is computed (one term for each different "*present*" day). The result is a value of correlation between $f(n)$ and future returns for each value of *n* in the specified range. This correlation for each value of *n* is then displayed in the graph.

There is a list box for changing the **Time Horizon**, which is the number of days of daily returns over which the correlation is computed. The correlation between the future returns over this time horizon is computed with a simple average of the corresponding number of future days of the technical indicator. This then should preserve the phase of the technical indicator for varying time horizons. There is also a spin button that you can use to set the **lead-time**. This merely adjusts the **phase** of the technical indicator by shifting it left or right. The goal is to determine the amount of lead-time, or phase shift, necessary to bring the (positive) peak of the correlation graph under the ZERO line in the middle of the graph. This then corresponds to the maximum degree of correlation between the *technical indicator* and the *future returns*. Then the technical indicator fits the definition of a **Momentum indicator**, which means that the technical indicator is maximally correlated with the returns. (If the lead-time is not zero, you can go back to the Technical Indicator dialog and build this lead-time into a revised technical indicator.)

The Correlation Test dialog box also computes some numerical quantities of interest. First, the actual value of the **correlation** under the ZERO line is displayed. The **standard error** is displayed, which depends on the number of data points. The average ***N*-day volatility** for the stock data returns is displayed. From these numbers, a theoretical estimate of the annual simple and compound gain is computed, using the formula given in the Appendix below (daily correlation times daily volatility times 256 for simple returns). This gives you an idea of the theoretical gains possible with a given degree of correlation, given the volatility of the returns for that stock. There is also a set of radio buttons to change the vertical scale of the graph, but the scale is changed automatically anyway, so you usually don't need to do it yourself. Finally, the correlation can be re-computed using any of the three methods described in the Appendix, namely the **Pearson's R**, which is the normal method, and the two **robust** methods, which are **Spearman Rank-Order** and **Kendall's Tau**. (However, note that the actual gains to be obtained in trading are related to the **Pearson's R** method of correlation, not to the others.)

Correlations – Indicators dialog

This dialog box is accessed through the **Main Window** toolbar. It displays a **scatter graph** of the **correlation** between any given stock *future returns* data and a *technical indicator* constructed from that data. The stock is chosen by clicking the *Stock Data* button, which opens an **Open File** dialog. The technical indicator is chosen by clicking the *Indicator* button, which opens the **Technical Indicator** dialog box (described above). The scatter plot is then displayed, and the **correlation** and **confidence level** according to all three correlation methods are also displayed.

If you click the *Display* button, this opens the **Correlation Test** dialog box, described above. You may then view the correlation between the technical indicator and future returns for each value of lead-time from –100 days to +100 days. (The scatter plot only displays the correlation for ZERO lead-time.) The **Technical Indicator** dialog is also accessible from the **Correlation Test** dialog box. So this is a second method to access the **Correlation Test** and the **Technical Indicator** dialog boxes.

Correlations – Stocks Dialog

This dialog box is accessed through the **Main Window** toolbar. It displays a **scatter graph**, similar to the one in the **Correlations – Indicators** dialog, but this time the correlation is between the **returns** of two different stocks. You choose the two stocks using the **Data 1** and **Data 2** buttons, which both open an **Open File** dialog. The three types of **correlation** and their **confidence levels** are displayed as in the other Correlation dialog box.

If you click the **Display** button, this shows another dialog box, which contains a set of bar graphs of the correlation between the two stocks, as a function of **time lag**. The **time lag** is just the time difference, in days, between the returns that are compared in the correlation test. One graph for positive time lags and one graph for negative time lags are displayed. If you choose the same stock for both of the two stocks, then you can view the **autocorrelation** of the returns of that stock. In that case, the two bar graphs for positive and negative lags will be the same. If you change the time scale for the correlation to N days, then the correlation between N-day returns is computed. In this case, the bars of the bar graphs become N pixels wide rather than just one pixel wide.

This **Correlation – Stocks** dialog is useful for comparing the degree of **correlation** between different stocks for the purpose of **portfolio selection**. It is also useful for general studies of the **correlation structure** of the stock returns data. For example, when studying the autocorrelation of daily returns, if you look closely you can see that the autocorrelation for the first three days of lag is almost always *negative*. Often you can also spot what look like **cycles** in the correlation structure, with periods in the intermediate-term range of, say, one to several months.

Periodogram Dialog

This dialog box incorporates the standard **Periodogram** test of **Time Series Analysis**. The **Periodogram** is a method for measuring the **spectrum** of a time series, in this case stock returns data. A description of this test can be found in many standard textbooks, such as Brockwell & Davis (1991). The **Periodogram Method** is basically a **Fourier transform** of the stock returns data, and displays the amplitude of each frequency component, in steps of the lowest frequency, up to the **Nyquist frequency** (with period 2 days). For comparison, a second method of spectrum estimation, called the **Maximum Entropy Method**, is also displayed. This method relies on the **Linear Prediction filter**. The LP filter coefficients are computed from the returns data, then the ME method estimates the spectrum from these coefficients. It can be seen that the results are pretty similar in both cases.

According to the theory of the **Periodogram**, it must be **smoothed** on some time scale. If it is left unsmoothed, the *standard error* of each Fourier component is roughly 100% of the amplitude of the component. After smoothing on a time scale of *N* days, the *standard error* of the *smoothed* Fourier component is roughly $1/\sqrt{N}$. The default smoothing is set at 6 days, but you can change to a wide range of smoothing time intervals using a list box. Please consult a standard text for an explanation of the necessity for smoothing the Periodogram.

There are many peaks and valleys in the observed spectrum, but unfortunately it is not possible to show conclusively that these are any different from a random result. To demonstrate this, the **Periodogram** can be viewed using only **random Gaussian data**, generated by a random-number generator. To view the random data, click the **Random** button. Each time you click this button, a new set of random data is generated, and displayed in the two windows. The **stock returns** are displayed in **dark red**, while the **Gaussian data** are displayed in **blue**. Clicking the same button, which now reads **Restore**, returns to the same stock data. So you can repeatedly compare the stock data Periodogram with that of the random Gaussian data, with a new set of random data each time, just by repeatedly clicking this same button. It can be seen that the random Gaussian data also displays the same type of peaks and valleys, so it can be concluded that, whatever correlations are present in the stock data, this test is not very sensitive to them.

Also included in this dialog box are two standard statistical tests. The **Kolmogorov-Smirnov test** [NR (1992)] compares the spectral distribution of the Periodogram to a constant distribution. It then computes the **confidence level** that the **spectrum** is *different* from a constant distribution. This can be interpreted as the probability that the spectrum was *not* obtained from a random Gaussian distribution by random chance alone. It will be observed that, using the random Gaussian data, this confidence level ranges from 0% to 100%, and is distributed roughly equally over this range. This is what you would expect from a purely random result. The **Fisher's test** [Brockwell & Davis (1991)] computes the **confidence level** for a **periodic component** in the spectrum. This is used to determine the probability that an observed **cycle** in the data is *not* obtained from a random Gaussian distribution by random chance alone. It will likewise be seen that, using the random Gaussian data, this confidence level also ranges from 0% to 100%, and is also distributed roughly equally over this range. It would be interesting to run these two tests over a collection of stock data files, and observe whether or not the distribution of the confidence levels of the two tests is still constant from 0% to 100%.

Appendix: Definition of Correlation

The standard definition of linear correlation of two random variables, called **Pearson's R**, is given by [NR (1992)]:

$$r \equiv \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2}\sqrt{\sum_i (y_i - \bar{y})^2}}$$

Here, $\bar{x}$ and $\bar{y}$ are the mean values of the two random variables. There are two other types of correlation, which are called **robust** correlations, which are the **Spearman Rank-Order** and **Kendall's Tau**. These are called nonparametric methods of computing correlation because, unlike the linear or Pearson's R correlation, they do not presuppose a *Gaussian* distribution of the random variables. The Spearman Rank-Order correlation is the linear correlation of the *ranks*, as opposed to the linear correlation of the values of the variables as in linear correlation. To compute the ranks, the values are arranged in increasing order, and the order of each value is its rank. *Kendall's Tau* uses the correlation of the numerical *order* of the ranks (greater than, less than, or the same), as opposed to the difference in *value* of the ranks as in *Spearman Rank-Order*. These two robust methods are more reliable when the distribution of the random variables is non-Gaussian, and in particular when the distribution has "fat tails" as is the case with most financial data.

However, for our purposes a modified definition of correlation is more suitable. The problem with the above definition is that it breaks down when the buy-and-hold strategy is considered. To be specific, one of the above random variables will represent the future returns, and the other will represent the *trading rules*, or amount to be invested in a short-term trading strategy. If *s* is the number of shares, and δp is the actual returns (change in price per share), then the expected (simple) gain *g*, in dollars, is given by (summed over the trading days in a given time interval):

$$s \equiv \text{number of shares}$$

$$\delta p \equiv \text{actual returns}$$

$$g \equiv \sum_i s_i \delta p_i = \text{gain}$$

For *y* in the correlation formula we may use the logarithmic returns as a conservative estimate for the actual returns:

$$y_i \equiv \ln(p_i + \delta p_i) - \ln p_i = \ln\left(\frac{p_i + \delta p_i}{p_i}\right) = \ln\left(1 + \frac{\delta p_i}{p_i}\right) \leq \frac{\delta p_i}{p_i}$$

The amount invested, in dollars, at time i is given by:

$$d \equiv s \times p = (\#\text{shares}) \times (\text{price per share}) = (\text{dollar amount invested})$$

Thus we have:

$$\begin{aligned} g \equiv \text{gain} &= \sum_i (s_i p_i) \frac{\delta p_i}{p_i} = \sum_i d_i \frac{\delta p_i}{p_i} \\ &= (\text{dollar amount invested}) \times \\ &\quad \times (\text{fractional change in price}) \end{aligned}$$

For the *annualized simple gain* we sum over the number of trading days in a year, assuming we are dealing with daily returns, which may be taken to be 256 days.

The *trading rules* variable x is defined as the dollar amount invested at any given time, relative to the average amount of equity invested over the time period. This average equity can be either long or short, so the average equity invested is given by the *average absolute value* of the dollar amount invested over the time interval:

$$x_i \equiv \frac{d_i}{\langle |d_i| \rangle} \equiv \text{"trading rules" variable}$$

Here we define the *average absolute value* of the dollar amount invested over the time interval by:

$$\langle |d_i| \rangle \equiv \frac{1}{N} \sum_i |d_i| \equiv \text{average absolute value of equity invested}$$

The average absolute value of the equity invested, as a percentage of the total equity available to invest, is called the *average percent margin*. To compare measured correlations to measured returns from trading rules, we normalize the average percent margin to 100%. In other words, the normalized gain, denoted $\tilde{g}$, will be given by the *annualized simple gain* divided by the *average absolute value* of the dollar amount invested.

However, the correlation is expressed in terms of the root mean square of the trading rules, not the average absolute value of the trading rules (which is defined to be unity). We need to convert between one and the other. This is straightforward if we assume the random variables are distributed according to a Gaussian distribution. Denoting a Gaussian random variable by z, with standard deviation σ, it is well known that the Gaussian distribution (assuming $N \to \infty$) is normalized as follows:

$$\int_{-\infty}^{+\infty} \exp\left[-z^2 \big/ 2\sigma^2 \right] dz = \sqrt{2\pi}\sigma$$

The r.m.s. value of z is then given as the square root of the mean value of z^2, which (the latter) is defined to be the variance:

$$\langle z^2 \rangle \equiv \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} z^2 \exp\left[-z^2 \big/ 2\sigma^2 \right] dz = \sigma^2 \equiv \text{variance}$$

The average absolute value of z, on the other hand, is given as follows:

$$\langle |z| \rangle \equiv \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} |z| \exp\left[-z^2 \big/ 2\sigma^2 \right] dz = \frac{2}{\sqrt{2\pi}\sigma} \int_0^{\infty} z \exp\left[-z^2 \big/ 2\sigma^2 \right] dz$$

$$= \frac{1}{\sqrt{2\pi}\sigma} \int_0^{\infty} \exp\left[-z^2 \big/ 2\sigma^2 \right] dz^2 = \frac{2\sigma^2}{\sqrt{2\pi}\sigma} = \sqrt{\frac{2}{\pi}}\sigma$$

Thus we have the following general relationship between the average absolute value of a Gaussian variable and its standard deviation (root mean square value):

$$\sqrt{\langle z^2 \rangle} \equiv \sigma = \sqrt{\frac{\pi}{2}} \langle |z| \rangle \approx 1.2533 \langle |z| \rangle$$

Thus when any quantity is normalized to unit average absolute deviation (dividing by the average absolute deviation), it will be about 25% greater than when it is normalized to unit standard deviation (dividing by the standard deviation).

Thus the annualized gain, normalized to unit margin (average absolute amount of dollars invested) will be given by:

$$\tilde{g} \equiv \frac{g}{\langle |d_i| \rangle} \approx 1.2533 \frac{g}{\sqrt{\langle d_i^2 \rangle}}$$

This may be rewritten using the definition of the gain given above (renormalizing d_i in numerator and denominator by dividing by the average absolute deviation):

$$\tilde{g} \approx 1.2533 \sum_i \frac{d_i}{\sqrt{\langle d_i^2 \rangle}} \frac{\delta p_i}{p_i} = 1.2533 \sum_i \frac{x_i}{\sqrt{\langle x_i^2 \rangle}} \frac{\delta p_i}{p_i}$$

We may now use the inequality give above to rewrite this in terms of the logarithmic returns y_i:

$$\tilde{g} \geq 1.2533 \sum_i \frac{x_i}{\sqrt{\langle x_i^2 \rangle}} y_i = 1.2533 \sqrt{\langle y_i^2 \rangle} \sum_i \frac{x_i y_i}{\sqrt{\langle x_i^2 \rangle} \sqrt{\langle y_i^2 \rangle}}$$

Taking into account that there are 256 trading days in a year, we find:

$$\tilde{g} \geq 1.2533 \sqrt{\langle y_i^2 \rangle} \sum_i \frac{x_i y_i}{\sqrt{\frac{1}{256} \sum_i x_i^2} \sqrt{\frac{1}{256} \sum_i y_i^2}} = 1.2533 \cdot 256 \sqrt{\langle y_i^2 \rangle} \sum_i \frac{x_i y_i}{\sqrt{\sum_i x_i^2} \sqrt{\sum_i y_i^2}}$$

Let us denote the average volatility, by which we mean the r.m.s. value of the logarithmic returns, by σ:

$$\sigma \equiv \sqrt{\langle y_i^2 \rangle} \equiv \text{root mean square logarithmic volatility}$$

We may then define our modified correlation, as follows:

$$\tilde{r} \equiv \frac{\sum_i x_i y_i}{\sqrt{\sum_i x_i^2} \sqrt{\sum_i y_i^2}} \equiv \text{"modified" correlation coefficient}$$

In other words, the modified correlation is the regular correlation with the mean values of the variables not subtracted off.

The annualized gain, normalized to unit margin, is the expected dollar gain divided by the average absolute amount of dollars invested. It is thus given in terms of the quantities defined above by:

$$\tilde{g} \geq 1.2533 \cdot 256 \cdot \sigma \cdot \tilde{r}$$

$$\tilde{g} \equiv \text{annualized (simple) gain, normalized to unit margin}$$

$$\sigma \equiv \text{root mean square logarithmic volatility}$$

$$\tilde{r} \equiv \text{"modified" correlation coefficient}$$

Thus the expected annualized simple gain, normalized to *unit margin* (unit average absolute amount of equity invested) is approximately given by the modified correlation multiplied by the average (r.m.s.) daily volatility of returns, times the number of trading days in a year and a numerical factor.

Thus we see that the meaningful quantity for the estimation of trading returns is this modified correlation, computed as if the mean values of the variables were zero, rather than the standard definition of correlation. In the ideal case of daily returns that are constant, the trading rules would be simply a constant amount invested, and then the modified correlation between the trading rules and the returns would be 100%. On the other hand, according to the usual definition of correlation, the correlation would be indeterminate because the variance of both the trading rules and returns would be zero, both of these would be equal to their mean values, so there would be zero in both the numerator and denominator. If, as often happens, the trading rules are nearly constant, then there would be very small quantities in both the numerator and denominator, and the computed correlation would be dependent on minute variations in the trading rules, which has very little to do with actual investment gains or losses. The modified correlation, on the other hand, would register the gain or loss to be incurred from the nearly constant investment, so it is the appropriate measure of correlation to be employed here.

The usual routines for measuring correlation [NR (1992)] use the data with the means subtracted off, so these routines must be modified to eliminate this subtraction of the means, resulting in the formula for the modified correlation given above. The *theoretical return* is then computed as above, multiplying this *modified correlation* by the *r.m.s. (logarithmic) volatility*, times the number of trading days in a year and a numerical factor, which results in a number which is approximately the actual gain, for small values of the daily returns, and is always less than or equal to the actual gain (so it is a conservative estimate).

References

Peter J. Brockwell & Richard A. Davis, Time Series: Theory and Methods, 2nd ed. Springer-Verlag, New York (1991)

John Y. Campbell, Andrew W. Lo, & A. Craig MacKinlay (CLM), The Econometrics of Financial Markets, Princeton University Press, Princeton, NJ (1997)

Robert D. Edwards & John Magee, Technical Analysis of Stock Trends, 6th ed.
John Magee Inc., Boston, Mass. (1992)

Sheldon Natenberg, Option Volatility & Pricing,
McGraw-Hill, Inc., New York, NY (1994)

Edgar E. Peters, Chaos and Order in the Capital Markets,
John Wiley & Sons, Inc., New York, NY (1991)

Edgar E. Peters, Fractal Market Analysis,
John Wiley & Sons, Inc., New York, NY (1994)

William H. Press, Saul A. Teukolsky, William T. Vetterling, & Brian P. Flannery (NR),
Numerical Recipes in C, The Art of Scientific Computing, 2nd ed.
Cambridge University Press, Cambridge, UK (1992)

Martin J. Pring, Technical Analysis Explained, 3rd ed.,
McGraw-Hill, Inc., New York, NY (1991)

Tonis Vaga, Profiting From Chaos
McGraw-Hill, Inc., New York, NY (1994)