

Mathematics of the Random Walk Model

(Revised April 12, 2004)

Robert Murray, Ph.D.
Omicron Research Institute

(Copyright © 2004 *Omicron Research Institute*. All rights reserved.)

It is common wisdom that the stock market is *efficient*, meaning that all available information is instantly taken into account by the market, and stocks are fairly priced at all times. This implies that past information cannot be used to predict future stock prices and make a consistent profit. The mathematical expression of this idea is (presumably) that stock price fluctuations (returns) are random Gaussian noise, and hence the price movements of stocks correspond to Brownian motion. This is the content of the *Random Walk* theory of the stock market. At best, according to this theory, there is a secular trend, which reflects the volatility (risk) of the stock investment.

However, it is clear to most traders and investors that, while the financial markets are indeed very efficient, they are not 100% efficient. The remaining inefficiencies may be utilized to try to make a profit and “beat the market”. The financial markets cannot be completely efficient, because investors and traders are not perfectly well informed, perfectly rational, and they do not react instantaneously to new market information. Instead, some players in the market act on the basis of incomplete information, behave irrationally and display a “herd mentality”, and they have a finite “time horizon”. These investor behaviors introduce a small degree of correlation between different market data, which may be exploited by the astute trader or investor. It is our purpose here to try to model these small correlations and develop trading rules to best exploit them. If these trading rules were to become widely known among traders and investors, then they would become ineffective because the market would become more efficient. If the market were totally efficient, then all financial assets would be fairly priced, volatility would be minimized, and dangerous market manias and crashes would be eliminated. Ideally, we would like to see the market become as efficient as possible, so that all assets are fairly priced. This will be a happy day when it finally arrives, but in the meantime, astute traders and investors can make use of the inefficiencies of the market to make a profit and “beat the market”, at the expense of the less astute traders and investors.

When the market is efficient, and the price of stocks is a fair and true reflection of their economic value, then this is the desired situation. Then the most investment capital will be allocated to those companies with the most potential to grow and produce the greatest wealth, and all investors will gain by this. However, if the market is inefficient, then traders who make money on these short-term inefficiencies can only do so when other traders sustain a loss, so it is a *zero-sum game*. If somebody or other has to lose because the market is inefficient, it is better for it to be the other guy. However, it would be better for everyone if the market could be made as efficient as possible, so that investment gains are a reflection of an increase in true economic value. This will happen if traders can

react to existing market inefficiencies in such a way as to reduce or eliminate them, making the market more stable (and less volatile), rather than causing the market to become even more inefficient and unstable (and more volatile). This should happen, in general, if traders can use good technical indicators to judge the state of the market, and take appropriate action to anticipate impending moves in the market.

The central issue in any active trading strategy is the expected *return* for a given level of *risk*. The return is the most likely gain, in percentage terms, per year, and by risk we mean the variance, or average spread of the returns (squared), to be expected per year. In a (stationary Gaussian) stochastic process, both the expected return and the variance (square of the expected “spread” of the returns, or square of the standard deviation) increase proportionally to time. Hence we study the ratio of expected return to variance as a measure of the ratio of return/risk. We begin by studying some of the consequences of the Random Walk model and Efficient Market Hypothesis, and study the expected return and risk associated with an active trading strategy as compared with the “Buy & Hold” strategy.

Efficient Market Hypothesis

An *efficient market* is defined as follows [Sharpe, et.al. (1995), p.105]:

A (perfectly) efficient market is one in which every security’s price equals its investment value at all times.

This implies that any attempt to identify mispriced securities is futile. An equivalent definition is as follows [Sharpe, et.al. (1995), p.106]:

A market is efficient with respect to a particular set of information if it is impossible to make abnormal profits by using this set of information to formulate buying and selling decisions.

This implies that in an efficient market, investors should expect to make only “normal” profits and earn a “normal” rate of return on their investments. There are actually three forms of market “efficiency”:

- 1) **Weak-form efficiency** means it is impossible to make abnormal profits by using *past prices* to make decisions about when to buy and sell securities.
- 2) **Semi-strong-form efficiency** means it is impossible to make abnormal profits by using *all publicly available information* to make decisions about when to buy and sell securities.
- 3) **Strong-form efficiency** means it is impossible to make abnormal profits by using *all information, both public and private* to make decisions about when to buy and sell securities.

The evidence seems to suggest that the market conforms to *weak-form efficiency*, but it does not conform so well to *semi-strong-form efficiency*, and even less so to *strong-form efficiency* [Sharpe, et.al. (1995)]. However, we will present other evidence to suggest

that even the weak-form efficiency is only approximately valid, and it is indeed possible to use past prices alone to formulate profitable short-term trading rules.

But what is the precise meaning of the Efficient Market Hypothesis? Can this hypothesis be formulated in a precise mathematical manner? How can one devise a statistical test to determine the validity of the Efficient Market Hypothesis?

It is already clear that the Efficient Market Hypothesis is rather vague in its formulation. What is the definition of an “abnormal profit”? Presumably this refers to a specific theory of asset pricing such as the Capital Asset Pricing Model (CAPM), in which the expected rate of return is the risk-free rate (T-bill or discount rate) plus a risk premium that is proportional to the volatility of the investment. This model regards the volatility as due to Gaussian stochastic noise superimposed on a constant expected rate of return, which depends on the volatility. So perhaps the Efficient Market Hypothesis should be re-formulated in mathematical terms as follows:

A market is efficient with respect to a particular set of information if the de-trended price returns (over a given time interval) of securities are statistically independent of this set of information.

This would then imply that the prices of securities obey the Random Walk model (with drift), in which the drift velocity is a simple function of the (constant) volatility, and the price fluctuations are due to Gaussian white noise. So we may test this hypothesis by testing for the statistical independence of price returns with the information in question, in particular with the past price returns.

In The Econometrics of Financial Markets (1997), p.20, CLM offers another definition of an efficient market, which in turn is quoted from Malkiel (1992)¹:

A capital market is said to be efficient if it fully and correctly reflects all relevant information in determining security prices. Formally, the market is said to be efficient with respect to some information set...if security prices would be unaffected by revealing that information to all participants. Moreover, efficiency with respect to an information set...implies that it is impossible to make economic profits by trading on the basis of [that information set].

These authors make the point that “perfect efficiency is an unrealistic benchmark that is unlikely to hold in practice”. But they consider *relative efficiency*, or “the efficiency of one market measured against another”, to be a more useful concept. They go on to make the following statement [CLM (1997), p.24]:

Few engineers would ever consider performing a statistical test to determine whether or not a given engine is perfectly efficient—such an

¹ Malkiel, B., 1992, “Efficient Market Hypothesis”, in Newman, P., M. Milgate, & J. Eatwell (eds.), New Palgrave Dictionary of Money and Finance, Macmillan, London.

engine exists only in the idealized frictionless world of the imagination. But measuring relative efficiency—relative to the frictionless ideal—is commonplace. Indeed, we have come to expect such measurements for many household products: air conditioners, hot water heaters, refrigerators, etc. Similarly, market efficiency is an idealization that is economically unrealizable, but that serves as a useful benchmark for measuring relative efficiency.

So the conclusion is that the Efficient Market Hypothesis is a useful idealization and is approximately true in the real world, but that market inefficiencies can and do exist. It is then up to the astute trader to find and make use of these inefficiencies in order to make profitable trades and “beat the market”.

Sources of Inefficiency

If investors had perfect access to all economic and market data, had perfect judgment and foresight and knew the correct way to react to all situations, and could react instantaneously, then the financial markets would be perfectly efficient. However, this is not the case. Some investors are more astute than others in these areas, and these astute investors will be able to win and “beat the market”. Some of the main sources of inefficiency in the stock market are the following (arranged in approximate order of importance):

- 1. Herd mentality.** Investors get caught up in the emotion of *greed* and buy as stock prices steadily become more and more overvalued, until prices are driven up way beyond the true value of the stock, resulting in a *mania*. Then investors are caught up in the emotion of *fear* and sell all at once, as stock prices fall suddenly, resulting in a *crash*. Stock market crashes are probably the most important form of inefficiency in the market. Manias and crashes can probably only be explained by a *nonlinear* theory or model of price dynamics.
- 2. Inability to judge true economic value.** This is mainly a question of fundamental valuation, so it will depend on fundamental and economic data. The ability to judge the true economic potential of various businesses and invest accordingly is necessary to fulfill the true purpose of the stock market. When this decision is made correctly and the best companies are given the most investment capital, then all investors are rewarded as those companies create economic wealth.
- 3. Short time horizon.** Investors and traders take a short-term perspective and react to short-term economic developments and price fluctuations, ignoring the overall long-term picture. This may result in price fluctuations that reflect only the dynamics of traders reacting to the actions of other traders, rather than to the true economic value of the company. Traders who take a long-term perspective can take advantage of these price fluctuations, and profit over the long-term. Traders, who over-react to short-term developments and ignore the long-term picture, will lose over the long-term.
- 4. Time delay in reacting to situations.** The situation in this regard has improved dramatically in recent years, over the situation in decades past. A time lag in reacting to news or changing situations can be taken advantage of by traders who

are able to react faster. For example, news of a decrease in earnings will result in a response by traders over a range of time delays. Those who are able to respond the fastest will be in the best position to profit from the changing earnings situation by being the first to sell. A time lag in traders reacting to price fluctuations generated by other traders will result in various technical price patterns that can be used for profit, by those traders who can *anticipate* the resulting price moves and react fast enough. On the very short time scales, those traders who have access to real-time quotes have an advantage over those whose quotes are 15-minute delayed. This probably accounts for the observation that short-term “ticks” are highly correlated over time scales less than about 20 minutes, but essentially uncorrelated over longer time scales.

Some of the ways in which technical (and/or fundamental) indicators can be used to profit in the market are the following (arranged in approximate order of importance):

- 1. Avoiding stock market crashes.** If investors can tell that the market is in a *mania*, leading ultimately to an impending *crash*, then they can take steps to get out of the market before the crash occurs, thereby protecting themselves from serious losses. Technical indicators that can distinguish overall overvalued and undervalued conditions, and those that indicate the overall trend of prices, are of value in this regard. Also indicators that can show an impending reversal of trend would be of great value, but this may be difficult because trend reversals may be sudden and of a random nature (like avalanches or earthquakes). The best solution is to take the appropriate action *before* the trend reversal (and resulting *crash*) occurs.
- 2. Judging true fundamental value.** Using all available economic indicators, as well as information about the state of a particular business sector and the state of the individual company (fundamental data), enable an investor to make the best possible judgment about the prospects for long-term success of the company. Also, a deep understanding of business and economics, and an intuitive understanding of how well the company is positioned in its economic niche, will also be of great importance in determining the ultimate success or failure of a particular company. This is the very essence of successful long-term investing (as exemplified by Warren Buffett).
- 3. Judging the overall state of the economy and stock market.** Indicators that show the state of the overall economy and the stock market as a whole are of the most importance in making long-term investing decisions. Short-term trading decisions should be made with the long-term picture in mind. Using indicators that show whether a stock is overvalued or undervalued in the short-term, relative to the long-term, enable a trader to make a profit from short-term trades. These trades should be superimposed on an investment strategy based on long-term indicators of the overall state of the economy and market, as well as the individual stock. The short-term situation should always be viewed within the context of the long-term outlook.
- 4. Profiting from short-term fluctuations.** This is the traditional *technical analysis* method of trading. Various technical patterns can be recognized in the price data, due to correlations induced by the actions and reactions of other traders to price moves, with a corresponding range of time delays. Then a profit can be

made by correctly anticipating the price moves before the majority of other traders have a chance. Also, it may be possible to notice a change in situation due to forthcoming news, *before* the news actually becomes public, due to the reactions of “insiders” to the news and the resulting price moves. Making use of these price changes as indicators of forthcoming news, it is possible to react quickly and take advantage of the changing situation to make a profit. This would involve using a technical indicator that can detect an unanticipated change in price and anticipate the resulting price action, before it actually unfolds.

From these lists of market inefficiencies and ways to take advantage of them, it should be clear that it is of paramount importance to maintain an accurate picture of the *long-term* outlook for the economy, the market, and the individual company. Trading strategies based only on short-term indicators, that ignore the long-term picture, will lose, because there is no way to judge short-term value apart from the long-term context. But when viewed within the context of the long-term picture, technical indicators can be used for short-term trading to increase average *returns* (with a corresponding increase in *risk*), while maintaining an appropriate long-term investment strategy. Thus technical indicators should be used within the context of overall portfolio selection and strategy, to maximize *returns* while minimizing *risk*.

It should be added that, logically, the market must be inefficient to some extent. The market is kept efficient by the actions of short-term traders and arbitrageurs, but these types of traders are motivated by the profit motive, based on their perception that the market has inefficiencies that can be exploited. If the market were perfectly efficient, then there would be no incentive to do any short-term trading, and then the market would become inefficient because there would be no mechanism to keep it efficient. The conclusion is that the market is maintained at a level of inefficiency that is just sufficient to motivate short-term traders to keep it at that level of efficiency. Thus, if a trader is able to exploit the residual inefficiency to a greater degree than other traders, then he/she should be able to profit from short-term trading. Another way to look at it is that the market should be considered *efficient for the average investor*. But the average investor limits the complete efficiency of the market as mentioned above. More astute traders and investors will be able to spot inefficiencies and take advantage of them to make a profit. At the same time, less astute investors will tend to lose money in the market, by succumbing to the various pitfalls mentioned above. Trading and investing are very much a competition with definite winners and losers, not merely a game of chance in which winning is a matter of pure dumb luck. It is more like a poker game, in which a skilled player can win by taking the best advantage of lucky situations as they arise, and avoid serious losses in unlucky situations.

Mathematical Description of the EMH

It is not exactly clear how the Efficient Market Hypothesis implies a particular mathematical description of the stochastic process of price returns. First of all, it is not clear what is meant, exactly, by “abnormal profits”. Evidently some people do not admit the existence of a *secular trend* in securities prices, although market averages (in the United States) have been growing roughly exponentially for the past 200 years. However, if the distribution of stock price returns has zero mean for the log prices, then

there will be a natural upward bias for the actual price returns due to the volatility, as will be shown below. So we could postulate that the distribution of *logarithmic* price returns has zero mean, and the secular trend in actual prices arises solely due to the volatility. This is a small effect that is always smaller than the volatility of the prices themselves. But it appears that the secular trend is actually much larger than this effect, so we must include it explicitly in our model. However, from an intuitive point of view, we would expect the exponential secular rise in prices to be due to the corresponding exponential expansion of the overall economy. This being the case, it makes sense to invest for the long term to take advantage of this secular trend. Another point of view, derived from the CAPM, is that the secular trend is a “normal profit”, which compensates investors for the *risk* that they must bear due to the volatility. However, it is not clear *a-priori* what value this *risk premium* should have.

The next question that is unanswered by the EMH is the precise statistical nature of the stochastic process of (log) price returns. Perhaps the most extreme point of view is that the log price returns constitute a *strictly stationary time series of Gaussian random variables with constant mean and constant variance* σ^2 , $N[0, \sigma^2]$. If this is the case, then it is indeed impossible to predict any short-term price fluctuations, and any short-term trading strategy is fruitless. The only trading strategy that makes any sense is the buy-and-hold strategy, making use of the long-term secular trend due to the (constant) mean of the log price returns. This specification of the stochastic process as a Gaussian stationary process implies that the mean, covariance, and all higher-order moments of the probability distribution functions are constant and are given by the Gaussian distribution. In particular, the correlation of the next day’s log price return with any function of the past returns, whether linear or nonlinear, will be zero, and no technical indicator will have any predictive power. This is a very strong assumption, and is probably unrealistic. A weaker specification of the stochastic process of log price returns is that it is a *second-order stationary time series of independent, identically distributed (IID) random variables with constant mean and constant variance* σ^2 , $\text{IID}[0, \sigma^2]$. This implies that the price returns are not necessarily Gaussian, and although the random variables are uncorrelated, there may be higher-order correlations that are nonzero. In fact, the *Higher-Order Statistics (HOS)* [Nikias & Petropulu (1993)] are associated with non-Gaussian distributions. This would imply that no technical indicator that is a linear function of the past prices, including any “block” linear filter, would have predictive power. But a technical indicator, which is a nonlinear function of past prices, such as an “adaptive” linear filter in which the filter coefficients are actually functions of the past prices, might well have predictive power. So the EMH could either imply a strictly stationary process or a second-order stationary process for the log price returns. If we assume that the EMH is to be interpreted in the sense of a strictly stationary Gaussian process, we will see that inefficiencies do indeed exist, and the EMH is only approximately correct.

Nature of Correlation in Price Data

If the past price returns and other data are uncorrelated with future returns, then the EMH is valid and no technical indicator can have predictive power. On the other hand, if some correlation does exist between past prices or other data and future returns, then this

correlation may be used to make a partial prediction of the future returns. This can lead to a set of profitable trading rules. The central problem is then to identify the existence of this *correlation* and then formulate an optimal set of trading rules to take advantage of it. We may *define* a technical indicator to be some function of the past stock prices that is correlated with future returns. What is the nature of this correlation? In an “ordinary” time series, such as an ARMA (Auto-Regressive Moving Average) series, the correlation is greatest for short time “lags”, and decreases exponentially for longer lags. However, there is evidence that in the case of financial data, a “fractional” process, which has the property that the correlation decreases like a power law rather than exponentially, may more accurately describe the correlation. This makes sense, if investors are basing their investing decisions on the long-term price history of the stock. So it is important to look for correlation going far back into the past, not just in the recent price returns.

On the other hand, it has been demonstrated that there is rather strong correlation in the price “ticks” of stock prices, for time scales less than about 20 minutes. This inefficiency evidently results from the fact that most traders do not have access to real-time quotes, and use 15-minute delayed quotes instead. However, this type of real-time trading is most appropriate for floor traders on the exchange floor, or other financial professionals. The ordinary individual investor or trader would have difficulty doing this kind of real-time trading. This topic is properly the subject of an entirely separate book, so we will not consider it further here.

Of greater interest are the correlations obtained by acting on the basis of technical patterns in the stock prices (as well as earnings and other fundamental and economic data). Let us now examine the nature of this correlation. From Technical Analysis, there would appear to be two basic kinds of correlation. These are *price trends* and *reversion to the mean*. The *reversion to the mean* mechanism must (theoretically) exist to some extent in stock prices, because it is a fundamental mechanism in the establishment of an equilibrium price for the security. If the price fluctuates below this equilibrium price then buying pressure increases, driving the price back up to its equilibrium value. If the price fluctuates above the equilibrium price, then selling pressure drives the price back down to equilibrium. So this phenomenon is fundamental to the pricing mechanism of the market. If the market were perfectly efficient, of course, then the price would never fluctuate from its “true” equilibrium price. But investors and traders sometimes over-react to news stories or other developments, causing a price fluctuation, which then corrects back to equilibrium when the market regains its long-term perspective. This is just the “Buy Low – Sell High” mechanism of the market. So there should be a (negative) correlation between the variation of the price from some equilibrium value, and future returns. The other main type of correlation is *price trends*. If the price is trending higher, then it should continue to do so. This is the most basic kind of correlation to measure in the financial price return data, as it takes the form of *autocorrelation* of the price returns series. Unfortunately, it appears that when the long-term correlation of price returns is measured, there seems to be virtually no correlation present. However, this seems to run contrary to real-life experience, in which trending behavior definitely seems to be present. What appears to be going on is that *short-term correlation* seems to exist in the data, so that a trend will develop, run for a while, and

then dissipate again. Evidently what is happening is that the correlation of the returns is *time-dependent*. Over various intervals of time the correlation of returns is either positive or negative, and it all averages out to zero over a long time interval. So to measure this correlation, it is necessary to measure it over short time intervals, and consider a time-dependent autocorrelation function. Likewise, the *reversion to the mean* correlation is also probably time dependent. Another way of saying this is that as a stochastic process, the price returns process is *nonstationary*.

In any event, the correlations in question must inevitably be very small. If there is a correlation of, say, 1% between some technical indicator and the daily price returns, and variance in the price returns is also 1%, then this is an expected gain per day of 0.01%. This sounds small, but when compounded over 256 trading days in a year, it yields a compounded annual gain of 2.6%, which is a worthwhile (extra) gain. A 10% correlation between the indicator and the daily returns leads to an annual gain of 29%, which is excellent! On the other hand, a correlation of 3% between a technical indicator and the daily price returns leads to an annual gain of 8%, which is very respectable. But if the autocorrelation function of daily returns is measured as the sample autocorrelation, using, say, 1000 days of past price data (about four years), then the measurement error of this autocorrelation has a standard deviation of $1/\sqrt{1000} = 3.2\%$. At the 95% confidence level, the bounds of the measurement error are $\pm 1.96/\sqrt{1000} = \pm 6.2\%$. But a correlation at this level leads to an annual gain of 17.2%! Thus we see that practically all the usable correlation in the price data is going to be smaller than the bounds of the measurement error, whether at one standard deviation or at the 95% confidence level. This means that the variance in the returns, or *risk*, is always of the same order of magnitude as the *expected returns*. To detect the presence of these small correlations below the level of the measurement error, therefore, it will be necessary to develop some alternative statistical tests, and to be content with a rather large variance in the returns. Nevertheless, even though the value of each correlation coefficient, for a given time lag, is very uncertain, we can verify the existence of this correlation by measuring the *square* of the correlation coefficient, averaged over a range of time lags. Statistical measures for doing this will be described shortly. However, it should be pointed out that part of the problem arises when we attempt to measure the correlation as a long-term sample average, based on the premise that the statistics are *stationary*. In fact, as mentioned above, the statistics are almost certainly *nonstationary*, so some measure of correlation over short time intervals, as a function of time, would be more appropriate. In this regard, the wavelet methods appear very promising for financial data, as opposed to working with the time series directly (time domain) or its Fourier transform (frequency domain), both of which rely on the assumption of stationarity [GSW (2002)].

Validity of Statistical Inference in Finance

In statistics, various inferences and deductions are made on the assumption that some sort of “Law of Large Numbers” applies, so that taking some kind of limit as $N \rightarrow \infty$ leads to certain conclusions. For example, the theoretical definition of correlation and expectation depends on having a very large “statistical ensemble” consisting of N similar systems, which vary from each other only in the values of their random variables, and

then inferences are drawn based on taking the limit as $N \rightarrow \infty$. However, in finance there are not available a large number N of “instances” of IBM stock, for example, over a given time interval. In the case of a stationary stochastic time series, we may invoke the postulate that the system is *ergodic*, which means that time averages may be substituted for ensemble averages. Thus, for example, to calculate mean values or correlations, we take a sum over a range of data extending over a large time interval, rather than a sum over the ensemble (which does not exist). But if the time series is nonstationary, then the assumption of ergodicity cannot be valid either. If the financial time series is nonstationary, as it almost certainly is, then it cannot be a stationary series of Independent, Identically Distributed Normal random variables, as in the Random Walk model. On the other hand, the nonstationarity will make it difficult to actually identify any correlation in the data, using the usual simple tests.

What can we do about this? One solution is to employ more sophisticated tests of correlation in the data, rather than basic sample averages over the time series. For example, instead of a simple measure of correlation as a sample average, we can employ a Linear Prediction Filter to find the best linear regression of the future returns on the past data, over a given time interval. This calculation can be repeated every day, thus leading to a new fit to the data every day, so that the LP coefficients now become time-dependent. Then we can define a new measure of correlation that measures the correlation between the predicted and actual future price returns, calculated in this manner, for price returns over a given time period. (Thus we are using a measure of correlation that depends on the degree of fit of the linear regression.) Depending on the degree of nonstationarity, the time interval used for the calculation of the LP coefficients can be longer or shorter, to capture slowly or rapidly varying correlation in the data. Likewise, other technical indicators can be defined, which utilize only the most recent data, and the degree of correlation between these indicators and future returns can be measured. The other solution to this problem is to replace the *ensemble* averages with *portfolio* averages. Thus, considering only a single stock, the advantage due to using a technical indicator may be smaller than the variance of the returns, but if the same trading strategy is used for a large diversified portfolio, then the advantage will be averaged, but the variance will be reduced, and the trading advantage will end up significantly larger than the variance. Also, the correlation in the portfolio data shows up more clearly than for individual stocks, due to the fact that the effects of “exogenous influences” on the individual stock prices cancel out in the average over the portfolio. So it is important to consider the trading strategy in the context of the portfolio as a whole, rather than each stock separately. Thus, in summary, the ensemble averages in a statistical analysis are replaced in the case of financial data by the use of time-dependent technical indicators whose correlation with the future returns is measured by means of a *sample* average, and by means of a *portfolio* average.

Measurement and Verification of Correlation

In statistical time series, the statistical properties may be specified (in part) by means of *expectations*. The expectation is an average over a statistical ensemble. Suppose we had a (hypothetical) statistical ensemble of “systems” consisting of N “instances” of the same

stock, for example. Then the expected value of some function of the (present and past) data at time n is given by:

$$\langle f(n) \rangle \equiv E(f(n)) \equiv \frac{1}{N} \sum_{\text{ensemble}} [f(n)]$$

This expected value of the function of the past data $f(n)$, at time n , is also called the *mean* of the function.

The ability to predict future returns from past data depends on the degree of *correlation* between the past data and the future returns. Suppose we had some function of the past data $f(n)$, at time n , and the future returns $r(n+1)$. Then we could calculate the *covariance* between the past data and the future return by taking the *expected value* of their product:

$$\text{cov}[f(n), r(n+1)] \equiv \langle f(n)r(n+1) \rangle \equiv E(f(n)r(n+1))$$

The *variance* is similarly defined as the covariance of a quantity with itself:

$$\text{var}[r(n+1)] \equiv \langle r^2(n+1) \rangle \equiv E[r^2(n+1)]$$

$$\text{var}[f(n)] \equiv \langle f^2(n) \rangle \equiv E[f^2(n)]$$

Then the *correlation* is defined in terms of the covariance and variance as follows:

$$\rho_{fr}(n) \equiv \frac{\langle f(n)r(n+1) \rangle}{\sqrt{\langle f^2(n) \rangle} \sqrt{\langle r^2(n+1) \rangle}}$$

The correlation coefficient $\rho(n)$ varies between the values of -1 and $+1$. If the two quantities are uncorrelated, then the correlation coefficient is 0, but if they are (positively or negatively) correlated, then the value of the correlation coefficient is different from zero. If a non-zero correlation coefficient exists between the function of the past data $f(n)$, at time n , and the future returns $r(n+1)$, then the past data can be used to make a (partial) forecast of the future returns.

In the real world of finance, there are no statistical ensembles to work with. Therefore, we have to try to approximate the ensemble averages by some other methods. The main method that is used instead of an ensemble average is a *sample average*, invoking the assumption that the statistics are *ergodic*. Then the mean value over the ensemble is replaced by the sample mean:

$$\langle f(n) \rangle \equiv E(f(n)) \equiv \frac{1}{n} \sum_{i=1}^n [f(i)]$$

If the statistics are *stationary*, then this sample mean should not change (on the average) with time. But the more usual case is that the statistics are *non-stationary*, so this sample mean will be a function of the time n . In this case, we can take the sum only over m of the most recent values of the data, where m is the typical time scale over which the statistics vary. Similarly, in the covariance and variance the ensemble averages are replaced by time averages, invoking the assumption of ergodicity. When the correlation is calculated as above, using sample averages, then the correlation thus calculated is also known as *Pearson's r*. Perhaps a more sensitive method of detecting correlation is to use the nonparametric or rank correlation, such as the *Spearman rank-order correlation* [Numerical Recipes (1992)]. This type of correlation uses only the relative ordering of

the different values of the quantities involved, and is considered more *robust* (insensitive to the requirement of Normal or Gaussian statistics) than the usual sample correlation.

The Linear Prediction method amounts to a *linear regression* of the future returns on the past price and other data (to be described below). It turns out that this linear regression involves only the variances and covariances as described above. If we regress the future returns on all the N past returns, or on N functions of the past returns or other data, then the result is the *linear* function of the N functions which is the best fit to the future returns. We could also consider a non-linear function of the past prices and data, and perform a *non-linear regression*. This nonlinear regression would then involve expectations of products of the future returns with quadratic or higher functions of the past data. These higher-order expected values constitute what are known as *higher-order statistics (HOS)* [Nikias & Petropulu (1993)]. Using nonlinear regression methods, and example of which is the *neural network*, it is possible to model phenomena that cannot be modeled by linear methods. However, it should be noted that a *nonlinear* model refers to a model that is nonlinear in its *parameters*. We may still perform a linear regression on a set of technical indicators that are themselves highly nonlinear functions of the prices. In this case we are simply solving for the best *linear* combination of the nonlinear technical indicators, to fit the data. In most cases of interest to us (in finance), this is more than adequate, and it is not really necessary to use models that are nonlinear in their parameters. This is what might be called a *parametric* approach. In a *nonparametric* approach, such as a neural network, we apply the past prices and other data to the input, and let the network itself find the optimal nonlinear function of this data to approximate the future returns. In the parametric approach, we postulate the nonlinear technical indicators a-priori, and simply find the linear combination of these indicators that produces the best fit to the future returns.

Testing for the Presence of Correlation

For a stationary stochastic process, we would expect any departure from the Random Walk model to show up in the autocorrelation coefficient at lag k , $\rho(k)$. This is defined over the entire data set, consisting of T data points. However, we expect that in reality financial returns data is highly non-stationary. Hence, there may be temporary correlations present, which form and dissipate due to changing conditions. Since the Random Walk appears (at first glance) to be very closely approximated by the financial time series, we might conjecture that over a long time average, short-term correlations average out. But if these short-term correlations are present, they should show up in the *square* of the correlation coefficient. Thus we can use the Box-Pierce Q -statistic as a measure of the presence of correlation in the price series returns [CLM (1997)]:

$$Q_m \equiv T \sum_{k=1}^m \rho^2(k)$$

For large sample sizes T , it is shown that the sample correlation coefficients are asymptotically independent and normally distributed [CLM (1997)]:

$$\sqrt{T} \hat{\rho}(k) \overset{a}{\sim} N(0,1)$$

Thus, if the price returns series is Gaussian distributed, then the Q -statistic is distributed like the sum of squares of m Gaussian random variables. So this statistic is asymptotically distributed as the chi-square distribution with m degrees of freedom, χ_m^2 .

Another test of correlation is the Variance Ratio test [CLM (1997)]. This test is based on the fact that for a Gaussian Random Walk, the variance of the returns after q time steps should be q times the variance of each time step. This is a consequence of the fact that the variances add over any time interval. So we can detect the presence of correlations by measuring the variance after q time steps and comparing this with q times the variance for a single time step. If the q -time step return (at time t) is denoted by

$$r_t(q) \equiv r_t + r_{t-1} + \dots + r_{t-q+1}$$

then a short calculation shows that the Variance Ratio may be expressed in terms of the correlation coefficients at lags from 1 to $q-1$ by [CLM (1997), p.49]:

$$VR(q) \equiv \frac{\text{var}[r_t(q)]}{q \cdot \text{var}[r_t]} = 1 + 2 \sum_{k=1}^{q-1} \left(1 - \frac{k}{q}\right) \rho(k)$$

In contrast to the Box-Pierce Q -statistic above, which is a sum of squares of the correlation coefficients, the Variance Ratio statistic is a sum of the correlation coefficients themselves (at different time lags). Hence, if our hypothesis is correct that the correlation coefficients of the non-stationary process tend to average out over time, then the Box-Pierce Q -statistic, that measures the sum of squares of the correlations, will be a more sensitive indicator than the Variance Ratio statistic. However, Campbell, Lo, & MacKinlay (1997) show that in addition to the Box-Pierce Q -statistic, the Variance Ratio statistic also indicates significant departures from the null hypothesis of a Random Walk in various stock market data.

One could also construct a variant of the Q -Statistic by using the Rank-Order correlation in place of the ordinary correlation. Another variant could be constructed by first calculating the Linear Prediction coefficients for the regression of the future returns on the past returns, then expressing the correlation coefficient in terms of these LP coefficients. This method should be equivalent, since the LP coefficients are calculated in terms of the covariances between the future returns and past returns, as explained later in connection with Linear Regression. A nonzero result in the calculation of the LP coefficients, at a given percentile level, implies the existence of correlation in the returns at this significance level.

Let us quote here some of the numbers reported by Campbell, Lo, & MacKinlay (1997), in The Econometrics of Financial Markets, Ch. 2. Not all of the results will be quoted here; for the rest please see the book. But first, it should be pointed out that much greater correlation was detected for the *market averages* than for *individual stocks*. Furthermore, CLM report “significant positive autocorrelation for the equal-weighted portfolio, and autocorrelation close to zero for the value-weighted portfolio” [CLM (1997), p.72]. Then they go on to say:

“That the returns of individual securities have statistically insignificant autocorrelation is not surprising. Individual returns contain much

company-specific or *idiosyncratic* noise that makes it difficult to detect the presence of predictable components. Since the idiosyncratic noise is largely attenuated by forming portfolios, we would expect to uncover the predictable *systematic* component more readily when securities are combined. Nevertheless, the weak negative autocorrelations of the individual securities are in interesting contrast to the stronger positive autocorrelation of the portfolio returns.”

Then CLM go on to explain why the positive autocorrelations exist in the portfolio returns[CLM (1997), p.74]:

“Despite the fact that individual security returns are weakly negatively autocorrelated, portfolio returns—which are essentially averages of individual security returns—are strongly positively autocorrelated. This somewhat paradoxical result can mean only one thing: large positive cross-autocorrelations across individual securities across time.”

Thus we see that in order to utilize whatever correlations are present, it must be done within a trading strategy that applies to the whole portfolio, not just to individual stocks. We also see that the construction of this portfolio must take into account the cross-correlation between the different stocks. This may be done using the Markowitz Model portfolio or the Capital Asset Pricing Model (CAPM).

We list here the measured correlation in the *daily* returns for the CRSP Value-Weighted Index and the CRSP Equal-Weighted Index, over a period from 1962 to 1994 [CLM (1997), p.67]. We also note that the *weekly* returns for the CRSP Equal-Weighted Index also showed significant correlation (not shown). The correlation coefficient for the first four days is shown, as well as the 5-day and 10-day day Box-Pierce *Q*-Statistic. For comparison with the latter, we note that the chi-square value corresponding to the 99.5-percentile is 16.7 for the 5-day statistic (and slightly greater for the 10-day). Thus, a value of the *Q*-Statistic greater than this value indicates that there is less than a 0.5% probability that this is due to random chance, and greater than a 99.5% chance that it is significant. As can be seen, the *Q*(5) and *Q*(10) statistics listed below are *much* greater than this. Also, the correlation coefficient is statistically significant if it is greater than $3/\sqrt{T}$ (three standard deviations or 99.7-percentile), which equals 3.30% for $T = 8179$, and 4.68% for $T = 4090$. Most of the correlation coefficients $\rho(1), \rho(1), \rho(1), \rho(1)$, are also much greater than this:

CRSP Value-Weighted Index (Daily Returns)

Period	Size (T)	$\rho(1)$	$\rho(2)$	$\rho(3)$	$\rho(4)$	Q(5)	Q(10)
‘62-‘94	8,179	+17.6%	-0.7%	+0.1%	-0.8%	263.3	269.5
‘62-‘78	4,090	+27.8%	+1.2%	+4.6%	+3.3%	329.4	343.5
‘78-‘94	4,089	+10.8%	-2.2%	-2.9%	-3.5%	69.5	72.1

CRSP Equal-Weighted Index (Daily Returns)

Period	Size (T)	$\rho(1)$	$\rho(2)$	$\rho(3)$	$\rho(4)$	Q(5)	Q(10)
'62-'94	8,179	+35.0%	+9.3%	+8.5%	+9.9%	1,301.9	1,369.5
'62-'78	4,090	+43.1%	+13.0%	+15.3%	+15.2%	1,062.2	1,110.2
'78-'94	4,089	+26.2%	+4.6%	+2.0%	+4.9%	348.9	379.5

Thus we see that there are very significant correlations for the daily data for both Value-Weighted and Equal-Weighted indexes, but especially the latter. For the weekly and monthly data there are still some significant correlations, but less than for the daily data. It should also be observed that the correlations are much greater for the first period than for the second, indicating that the market has become more *efficient* over time.

Of course, these correlations were measured only between the one-day returns and the returns at lags of 1,2,3,4 days (as well as the 5 and 10 day *Q*-Statistics). There are many other possible functions of the past data that could be used for technical indicators, and which could display significant correlations. If the returns time series is modeled by a *fractional* process, then this indicates that long-term *persistent* correlations are present, which decay over time like a power law rather than exponentially (as with an ordinary time series). If this is the case, then a technical indicator consisting of a long-term average of the past data could be constructed, and the correlation of this indicator then compared to the future returns. Various nonlinear functions of the data can be used as technical indicators, and their correlation with the future returns measured. Once we have decided on a set of technical indicators to use, we may use *linear regression* to find the best fit of the future returns to some linear combination of these indicators. This fit could be a simple *Linear Prediction* filter, which is the optimal fit of the future returns to the most general linear function of the past returns themselves (going back *N* days), or it could be a regression on a set of *N* functions of the past data (linear or nonlinear), considered as technical indicators. An example of the latter could be, calculating a set of *N+1* moving averages of the past *price* data (not *returns*), and then taking the *N* differences between adjacent moving averages as the *N* independent technical indicators. Of course, many other technical indicators are possible.

Linear Regression

If the Random Walk model were to hold (exactly), then each daily log price return would be a Gaussian random variable, uncorrelated and statistically independent from all past log price returns. (That the random variables are distributed Normally, according to the Gaussian distribution, implies that if they are uncorrelated then they are statistically independent. Thus there are no higher-order correlations either.) On the other hand, in real financial data we expect the market to be almost, but not quite, efficient, and hence that there should be weak correlations present in the data. These correlations may then be found by computing the correlation coefficients, or more generally the covariances, between various functions of the data and the future returns. These functions are, in general, what we mean by *technical indicators*, and they may be linear or nonlinear functions of the past price returns, prices, fundamental data, or other economic data. If the correlations exist, we may then use these to calculate the expected value of the future returns. We do this by means of a general **linear regression** of the future returns on the

set of technical indicators. If the regression of the future returns is done directly on the past returns, then this is called an **autoregression**. The linear regression is the linear combination of the set of technical indicators that provides the best, or *minimum mean square error*, fit to the future returns. The linear combination of the technical indicators consists of a set of linear regression coefficients, each one multiplying each technical indicator in the set. In the more general **nonlinear regression** problem, we could have a situation where a technical indicator depends nonlinearly on a parameter, and we need to find the best value of this nonlinear parameter in a fit to the data. However, generally, since the correlations are so weak and the fit is not very precise anyway, we see no need to employ nonlinear regression. If a special situation arises where a technical indicator depends nonlinearly on a parameter that must be fit to the returns data, then this can be dealt with by means of one of the standard methods of **parametric** nonlinear regression. But this will generally not be needed. A more general technique, an example of which is the **Neural Network**, is to try to make a **nonparametric** nonlinear model of the data. This technique tries to make as few assumptions about the form of the nonlinear fit to the data, so it tries to find a “general” nonlinear model of the data. However, we prefer the method of choosing a plausible set of *technical indicators* to try to fit to the data using linear, parametric methods. This allows us to make maximum use of any “prior” knowledge we might have regarding the possible dynamics of the financial markets and sources of inefficiency, based upon our understanding of investor behavior. We may, of course, incorporate **nonlinearity** into the linear regression by choosing technical indicators that are themselves nonlinear functions of the past price and other data.

We must also make a sharp distinction between the case of a *stationary* time series and *non-stationary* time series. In the *stationary* case, the correlation structure, and indeed the joint probability distribution in general, is constant in time. Hence the technical indicator $z(n)$ is the *same* function of the past prices and economic data, for all times. In other words, the technical indicator $z(n)$ has no *explicit* time dependence. In the *non-stationary* case, on the other hand, the technical indicator $z(n)$ as a function of the past prices and economic data changes explicitly with time. This is evidently the usual case in financial time series. Thus we must search for correlations that are time dependent and continually changing with time. We must find an empirical way to measure the correlation over the recent past, for a given range of the past data over that time interval, and use this to estimate the correlation extrapolated a given time interval into the future. This empirical correlation estimate, however, must not be a fixed function of the past price and economic data. Rather, it can be a fixed *nonlinear* function (as for example, the definition of the correlation coefficient itself), but the correlation coefficients themselves are time dependent rather than being fixed constants. In other words, we measure the (time averaged) correlation coefficient over a short time interval in the recent past, rather than taking a long-term time average over all the data. This short-term correlation coefficient is then itself a function of time. We may also consider various nonlinear functions of the recent short-term data, and measure the correlation coefficient between these and the future returns. This correlation coefficient will again be time dependent.

We need a notation for the sample covariance of two random variables. Let us define the sample covariance as follows:

$$\langle \xi(n)\eta(n-h) \rangle_m \equiv \sum_{k=0}^{m-1} \xi(n-k)\eta(n-k-h)$$

$$\langle \xi(n-h)\eta(n) \rangle_m \equiv \sum_{k=0}^{m-1} \xi(n-k-h)\eta(n-k)$$

This notation takes into account the sum over only a finite data set of length m , with the present time indicated by n , and the time lag between the two variables denoted by h . Thus this covariance function is an explicit function of time, taking into account non-stationarity. It also takes into account the fact that if the variables ξ and η are different, then the covariance is different depending on which variable lags the other one by h time units.

Suppose the one-step future price return (or whatever variable we are estimating) is given as a function of N other variables, which could be the N past returns, some nonlinear functions of these returns, economic data, and so forth. Thus we wish to estimate the following equation:

$$u(n+1) = v(n+1) + \sum_{k=0}^{N-1} a_k(n|m)\eta(n;k) \quad [\eta(n;k) = u(n-k) \text{ or other variables.}]$$

We now calculate the covariance with each of the N variables on the r.h.s. The condition that the regression minimizes the sum of squares of the error terms $v(n+1)$ is precisely the condition that these error terms are “white noise”, uncorrelated with the regressor variables. (This is the part of the future return that is “unexplained” by the regressor variables.) The covariance of the regressors with the noise term is then taken to be zero as the condition for the optimal least squares fit. This then yields:

$$\langle \eta(n;l)u(n+1) \rangle_m = \sum_{k=0}^{N-1} a_k(n|m) \langle \eta(n;l)\eta(n;k) \rangle_m \quad (l = 0 \cdots N-1)$$

We define the covariance matrix for this set of variables at time n as follows:

$$\Gamma_{lk}(n|m) \equiv \langle \eta(n;l)\eta(n;k) \rangle_m \quad (l, k = 0 \cdots N-1)$$

We also define a covariance vector between the “desired response” $u(n+1)$ and the variable at time n :

$$\Delta_l(n|m) \equiv \langle \eta(n;l)u(n+1) \rangle_m \quad (l = 0 \cdots N-1)$$

Hence the linear regression equation becomes:

$$\Delta_l(n|m) = \sum_{k=0}^{N-1} \Gamma_{lk}(n|m)a_k(n|m) \quad (l = 0 \cdots N-1)$$

Assuming the random variables are linearly independent (so that the covariance matrix is not singular), we may invert the covariance matrix and solve for the LP filter coefficients:

$$a_k(n|m) = \sum_{l=0}^{N-1} \Gamma_{kl}^{-1}(n|m)\Delta_l(n|m) \quad (k = 0 \cdots N-1)$$

Then this equation yields the optimal coefficients for the best linear prediction of the one-step ahead price return $u(n+1)$ in terms of the N variables. Note that these filter coefficients are, in general, functions both of the current time n and of the time interval m over which the covariances are calculated.

If we can approximate the random variables as uncorrelated, as for example in the case where the random variables are the past price returns, then we have approximately:

$$\begin{aligned}\Gamma_{lk}(n|m) &\approx \Gamma(n|m)\delta_{lk} & \Delta_k(n|m) &\equiv \Delta(n|m)\rho_k(n|m) \\ a_k(n|m) &\approx \frac{\Delta(n|m)}{\Gamma(n|m)}\rho_k(n|m) & (k=0\cdots N-1)\end{aligned}$$

In other words, apart from a normalization factor (involving the possibly time-dependent variances), the filter coefficients are proportional to the correlation coefficients between the random variable η and the desired future price return u . Hence we can contemplate defining a more general Portmanteau statistic as the sum of squares of the LP filter coefficients:

$$\tilde{Q}_N(n|m) \equiv T \sum_{k=0}^{N-1} \tilde{\rho}_k^2(n|m) \quad \tilde{\rho}_k(n|m) \equiv \frac{\Gamma(n|m)}{\Delta(n|m)} a_k(n|m)$$

This statistic seems to combine the usual Q -Statistic with the Variance Ratio test. Thus a direct measure of the correlation in the data and hence the predictive power, is to compare this statistic with the chi-square distribution χ_N^2 .

This leads to the interesting idea of computing the filter coefficients for a variety of N values. Then, for each N , the Portmanteau statistic can be compared with the corresponding chi-square distribution, and the value of N for which the percentile is maximum can be chosen. This is a way to choose the optimum value of N for the greatest predictive power with the least variance.

We may choose any set of N technical indicators that we want for our linear regression model. The future returns are regressed on this set of N technical indicators, to yield the best one-step prediction of the future returns based on the present values of the indicators. In general, these technical indicators could be linear or nonlinear functions of the past log returns, log price levels, or other fundamental or economic data. However, it should be noted that an autoregression on the past log price returns is the most general linear function of the log returns. Any set of technical indicators that can be written as a linear function of past log returns, should be encompassed within the linear autoregression on past log returns. For example, if we consider a set of moving averages of the log prices, and take as our technical indicators the differences between the moving averages corresponding to different time scales, or if we consider log price levels relative to various averages on different time scales, then these should all be encompassed within the general linear autoregression (LP filter). The only advantage to using specific linear functions as technical indicators would be that the fit could be accomplished with far fewer fitting parameters, leading to a reduction in the “noise” or variance of the fit, or in “over fitting” with too many parameters. This is a very worthwhile consideration, so the approach of using specific technical indicators rather than a general autoregression on the past time series, is definitely worthwhile. We can also consider various nonlinear functions of the past data, such as the hyperbolic tangent of log price deviations from some long-term average, thereby possibly leading to new nonlinear effects in the model (such as the prediction of manias and crashes).

Time-Dependent Correlation

Let us now consider in more detail the case of “time-dependent correlation”. We consider a simple case of the one-step autocorrelation of a variable, which is time-dependent. We postulate the following equation (for $N=1$), which we would like to estimate:

$$u(n+1) = v(n+1) + \sum_{k=0}^{N-1} a_k(n|m)\eta(n;k) \equiv v(n+1) + a(n|m)u(n)$$

We calculate the (time average) covariance, over m time steps, of both sides with the regressor $u(n)$. The coefficients $a(n|m)$ are optimized over these m time steps, in the least-squares sense, and the condition that they are optimized is equivalent to the condition that the innovations $v(n+1)$ are uncorrelated “white noise”. Hence we arrive at the following optimization conditions, for each time n :

$$\langle u(n)u(n+1) \rangle_m = a(n|m)\langle u(n)u(n) \rangle_m \Rightarrow a(n|m) = \frac{\langle u(n)u(n+1) \rangle_m}{\langle u(n)u(n) \rangle_m}$$

At each point in time n , this is supposed to be the optimum linear regression over the previous m time steps. Hence the product $a(n|m)u(n)$, which we call the “regression function”, itself constitutes a technical indicator function which is an independent function of n (since the regression coefficients are not constant, but depend on n). We can thus write:

$$\eta(n|m) \equiv a(n|m)u(n) \quad u(n+1) = v(n+1) + \eta(n|m)$$

This is the best linear regression over the past m time units, for each time n . We can clearly generalize this procedure taking into account a set of N technical indicators, and calculating their best linear regression over the past m -day interval. This will then lead to a set of N regression coefficients, which are themselves time-dependent and complicated nonlinear functions of the past data.

If we calculate the regression over a different value of m time units, we arrive at different time-dependent regression coefficients, and a different technical indicator. Presumably, for some value of m there will be an optimal regression, in the sense that the mean-square error of the long-time regression of the future data on the indicator is minimized. This value of m , then, is what we mean by the “time scale” of the correlation. This optimal time scale will be different depending on how many time units ahead we are predicting. So to calculate a prediction at one-day intervals over the next 100 days, for example, we can choose 100 different values of m for the regression, one value for each number of days of prediction. The least squares fit for each m , for each number of days ahead predicted, can be made and the optimal value of m in each case can be determined.

Note that, since the coefficients $a(n|m)$ are no longer constants, but are ratios of two quadratic functions of the past data, the regression function $a(n|m)u(n)$ becomes a highly nonlinear function of the variables $u(n)$. The upshot of this is that a *time-dependent correlation* is equivalent to *higher-order statistics*. Unless we know of some explicit time dependence to take into account, such as the business cycle for example, the only way to estimate the time-dependent correlation is through nonlinear functions of the past data, embodying higher-order statistics. In the example used here, the nonlinear

function in question was the covariance of the past data over the m -day interval, which was then used to estimate the *present* value of the (two-point) covariance (correlation) of the data. However, other nonlinear functions could be used to estimate higher-order correlations, which may themselves be time dependent. In fact, any statistical indicator we can think of, when calculated over the past m -day interval, can be used as a time-dependent, nonlinear, technical indicator taking into account time-dependent higher-order statistics in the data. In particular, arbitrary products of these statistical indicators with themselves, each other, and with the covariance functions between themselves and the future data (and each other), may be used as new technical indicators. So we have at our disposal an infinite hierarchy of ever more highly nonlinear functions for use as possible technical indicators.

Return/Risk in the Random Walk Model

First let us consider the implications of the straight Random Walk model, according to the Efficient Market Hypothesis, so that there are no technical indicators that are effective. Then the daily (log) returns are given simply by a Gaussian random variable, which is uncorrelated with other daily returns, both past and future, and with constant variance (volatility squared), in addition to the constant secular trend. In the pure Random Walk case, the returns are given simply as follows:

$$\begin{aligned} u(n) &\equiv (\log) \text{ price return at time } n \text{ days} \\ v(n) &\equiv \text{Gaussian random variable (innovation)} \\ \delta &\equiv \text{secular trend (drift velocity per day)} \end{aligned}$$

$$u(n) = v(n) + \delta \quad (\text{Random Walk hypothesis})$$

The expected value of the return for day n and the variance of (logarithmic) returns are thus:

$$E[u(n)] = \delta \quad V[u(n)] = \sigma_v^2$$

Considering this process from day 0 to day n , the change in the log prices is given in terms of the returns as follows:

$$\log \frac{P_n}{P_0} = \sum_{k=1}^n u(k) = \sum_{k=1}^n \{v(k) + \delta\}$$

Since the Gaussian random variables are independent, the expected values and variances in the sum of the log price changes add, and the sum is given by:

$$E \left[\log \frac{P_n}{P_0} \right] = n\delta \quad V \left[\log \frac{P_n}{P_0} \right] = n\sigma_v^2$$

The Return/Risk ratio R/R for the log prices is defined to be the expected n -day return divided by the standard deviation, which is the square root of the variance. It is thus given by:

$$R/R \left[\log \frac{P_n}{P_0} \right] = \sqrt{n} \frac{\delta}{\sigma_v}$$

Thus for the log prices, in the Random Walk (Gaussian) case (with drift), the Return/Risk ratio increases in proportion to the square root of the time interval. Eventually, for a long

enough holding period n , the cumulative return will be much larger than the uncertainty (standard deviation) of the return. This is why, for the Random Walk model, the Buy & Hold strategy is the most sensible one. For a Buy & Hold strategy, the logarithmic gains are additive, since the amount invested each day is proportional to the price on that day, for a constant number of shares. In other words, this corresponds to a continuously compounded gain.

Suppose we try a different trading strategy, corresponding to day trading in the Random Walk model. Each day we invest a constant amount, hold it for one day, and then take our profit or loss. Since we are talking about the Random Walk model and there are no technical indicators, there is nothing on which to base our decision as to how much to invest each day. We begin by considering what happens when we invest a constant amount each day, and then consider investing a different amount each day and observe what effect this has on the variance of the returns. (It should not have an effect on the expected returns, provided that the average amount invested is the same, because the expected returns are due solely to the secular trend.)

In the day-trading strategy, we need to consider the actual price returns, not just the logarithmic price returns. In this case, the mathematics becomes a little more complicated, because in terms of actual price returns the distribution is not Gaussian. What process do the price returns themselves follow, if the log price returns follow a Random Walk? The log price return is given in terms of the prices by:

$$u(n) \equiv \log p(n) - \log p(n-1) = \log \frac{p(n)}{p(n-1)}$$

The fractional dollar return or gain per day is thus given by:

$$G(n) \equiv \frac{p(n) - p(n-1)}{p(n-1)} = \exp u(n) - 1 = e^\delta \exp v(n) - 1$$

The cumulative (dollar) gain over time, assuming that a constant amount is invested each day, would be obtained by adding together all of the daily fractional dollar returns, and then multiplying by the (constant) amount invested. In the more general case, the daily fractional dollar returns would be multiplied by the (non-constant) amount invested each day, then the results added up to produce the cumulative dollar gain. If the non-constant amount invested each day is assumed to be given by a constant dollar amount times a log-Normal probability distribution (since there are no technical indicators and it is random), then this simply corresponds to adding another log-Normal random variable to the daily (log) price returns. This should then have the effect of adding the variances of the two random variables, increasing the variance but not the average returns. We will investigate this question shortly.

To calculate expectation values of the fractional daily dollar return, we must now calculate the expectation of the exponential of white noise, $E[\exp v(n)]$. To do this, we must make use of the various moments of the Gaussian distribution. These are given by the moment generating function of a normal random variable, which is given by means of a Gaussian integral as [Ghahramani (1996)]. (This is discussed in the Appendix below):

$$M_Z(t) \equiv E[e^{tZ}] = \exp\left(\frac{t^2}{2}\right) \quad (Z \text{ a standard normal variable})$$

We define the Gaussian variable $v(n)$ in terms of the standard normal variable Z with variance unity as follows:

$$v(n) \equiv \sigma_v Z(n)$$

Hence we have:

$$E[\exp v(n)] = E[\exp \sigma_v Z(n)] = \exp\left(\frac{\sigma_v^2}{2}\right)$$

Thus the expected return on the investment in the Random Walk model is given by:

$$\langle G(n) \rangle \equiv E[G(n)] = E[e^\delta \exp v(n) - 1] = e^\delta \exp\left(\frac{\sigma_v^2}{2}\right) - 1$$

Let us note here that even when the expectation of the logarithmic price returns is zero, the expectation of the price returns themselves is (slightly) positive, due to the variance of the (log) returns and the fact that the distribution of the actual prices is not Gaussian.

Likewise, the variance of the fractional one-day dollar return is given by:

$$\begin{aligned} V[G(n)] &= E[G(n) - \langle G(n) \rangle]^2 \\ &= E\left[e^{2\delta} \left\{ \exp v(n) - \exp\left(\frac{\sigma_v^2}{2}\right) \right\}^2\right] \\ &= e^{2\delta} E\left[\exp 2v(n) - 2 \exp v(n) \exp\left(\frac{\sigma_v^2}{2}\right) + \exp(\sigma_v^2) \right] \\ &= e^{2\delta} \left[\exp\left(\frac{(2\sigma_v)^2}{2}\right) - 2 \exp\left(\frac{\sigma_v^2}{2}\right) \exp\left(\frac{\sigma_v^2}{2}\right) + \exp(\sigma_v^2) \right] \\ &= e^{2\delta} \left[\exp(2\sigma_v^2) - \exp(\sigma_v^2) \right] = e^{2\delta} \exp(\sigma_v^2) \left[\exp(\sigma_v^2) - 1 \right] \end{aligned}$$

Once again, this variance vanishes if the variance of the log price returns series itself vanishes.

Similarly, the (*compounded*) fractional price gain over n days in the Random Walk model is given by:

$$\frac{p(n) - p(0)}{p(0)} = \exp\{u(n) + u(n-1) + \dots + u(1)\} - 1 = \prod_{k=1}^n \exp u(k) - 1$$

Due to the statistical independence of the Gaussian random variables, we have:

$$E\left[\prod_{k=1}^n \exp u(k) - 1\right] = \prod_{k=1}^n E[G(k) + 1] - 1 = \prod_{k=1}^n e^\delta \exp\left(\frac{\sigma_v^2}{2}\right) - 1 = e^{n\delta} \exp\left(n \frac{\sigma_v^2}{2}\right) - 1$$

Thus the expected (compounded) fractional price gain over n days is given by:

$$f(n) = e^{n\delta} \exp\left(n \frac{\sigma_v^2}{2}\right) - 1$$

Notice that the variance increases the expected return, over what it would be with zero variance ($e^{n\delta} - 1$). Using the same method, we may calculate the variance of the n day (*compounded*) fractional price returns:

$$\begin{aligned}
V\left[\prod_{k=1}^n \exp u(k) - 1\right] &= E\left[\prod_{k=1}^n (G(k) + 1) - \prod_{k=1}^n \langle G(k) + 1 \rangle\right]^2 \\
&= E\left[\prod_{k=1}^n e^\delta \exp v(k) - \prod_{k=1}^n e^\delta \exp\left(\frac{\sigma_v^2}{2}\right)\right]^2 \\
&= e^{2n\delta} E\left[\prod_{k=1}^n \exp v(k) - \exp\left(n \frac{\sigma_v^2}{2}\right)\right]^2 \\
&= e^{2n\delta} E\left[\prod_{k=1}^n \exp 2v(k) - 2\left\{\prod_{k=1}^n \exp v(k)\right\} \exp\left(n \frac{\sigma_v^2}{2}\right) + \exp(n\sigma_v^2)\right] \\
&= e^{2n\delta} \left[\exp\left(n \frac{(2\sigma_v)^2}{2}\right) - 2 \exp\left(n \frac{\sigma_v^2}{2}\right) \exp\left(n \frac{\sigma_v^2}{2}\right) + \exp(n\sigma_v^2)\right] \\
&= e^{2n\delta} \left[\exp(2n\sigma_v^2) - \exp(n\sigma_v^2)\right] = e^{2n\delta} \exp(n\sigma_v^2) [\exp(n\sigma_v^2) - 1]
\end{aligned}$$

Thus the compounded variance is given by:

$$v(n) = e^{2n\delta} \exp(n\sigma_v^2) \{\exp(n\sigma_v^2) - 1\}$$

Both of these results follow simply from the fact that the variance of the log returns increases linearly with time.

Now when we calculate the Return/Risk ratio for the dollar fractional price returns, we find:

$$R/R \left[\frac{p(n) - p(0)}{p(0)} \right] = \frac{f(n)}{\sqrt{v(n)}} = \frac{1 - e^{-n\delta} \exp\left(-n \frac{\sigma_v^2}{2}\right)}{\sqrt{\exp(n\sigma_v^2) - 1}}$$

Note that for small δ and σ_v , this result is approximately the same as the result for the log prices, as it must be:

$$R/R \left[\frac{p(n) - p(0)}{p(0)} \right] \approx \frac{1 - \left(1 - n\delta - n \frac{\sigma_v^2}{2}\right)}{\sqrt{n\sigma_v^2}} \approx \sqrt{n} \frac{\delta}{\sigma_v} \quad (n\delta, \sqrt{n}\sigma_v \ll 1)$$

However, for large times the behavior is different. We find:

$$R/R \left[\frac{p(n) - p(0)}{p(0)} \right] \approx \frac{1}{\sqrt{\exp(n\sigma_v^2)}} = \exp\left(-n \frac{\sigma_v^2}{2}\right) \rightarrow 0 \quad (n\delta, \sqrt{n}\sigma_v \gg 1)$$

This is disturbing news. It indicates that, for the fractional prices, the Return/Risk ratio for a Random Walk process must ultimately go to zero for long times. Ultimately, the standard deviation of the process becomes larger than the expected return, no matter how

large the secular trend. This is because the standard deviation in the *actual* prices itself increases in proportion to the factor $e^{n\delta}$, along with the expected price return, in addition to other factors that depend on the variance. Thus the variance always wins in the end, no matter what the secular trend, for a Random Walk process. So, although the expected *return* increases exponentially with time, the *risk* increases at an even faster rate, so that the *uncertainty* of the returns in the Random Walk model increases exponentially faster than the returns themselves. This revelation provides a great deal of incentive to try to control this risk, somehow, through active trading (if possible). (However, it is likely that active trading will actually *increase* the risk, and *decrease* the Return/Risk ratio, even though the expected return will *increase*.)

Actually, this result might not be as bleak as it sounds, because most of the uncertainty is going to be on the upside, not the downside. It is clear that, since the price can never fall below zero, if we were to calculate the r.m.s. risk on the downside only, the R/R ratio for this downside risk alone must be a number greater than unity. (Hence it cannot go to zero as in the calculation above.) It is only on the upside that the “risk” of a large price move can be much greater than the expected price. For this reason, in the case of non-Gaussian distributions such as the distribution of actual price returns, it makes sense to consider other measures of risk, such as the *Value at Risk* (VaR) [Bouchard & Potters (2000)]. However, in any event, we must still realize that in terms of the actual prices, a “law of large numbers” does not appear to operate, in the sense that the width of the distribution in actual dollars, as a fraction of the expected return, does not go to zero for large times.

The *simple* fractional price gain over n days in the Random Walk model is given by the *sum* of the fractional price gains over the time interval:

$$\sum_{k=1}^n \left[\frac{p(k) - p(k-1)}{p(k-1)} \right] = \sum_{k=1}^n [\exp u(k) - 1] = \sum_{k=1}^n [e^\delta \exp v(k) - 1]$$

The expected simple price gain is thus given by:

$$E \left[\sum_{k=1}^n G(k) \right] = \sum_{k=1}^n \langle G(k) \rangle = \sum_{k=1}^n \left[e^\delta \exp \left(\frac{\sigma_v^2}{2} \right) - 1 \right] = n \left[e^\delta \exp \left(\frac{\sigma_v^2}{2} \right) - 1 \right]$$

Thus the expected (simple) fractional price gain over n days is given by:

$$f_{\text{simp}}(n) = n \left[e^\delta \exp \left(\frac{\sigma_v^2}{2} \right) - 1 \right]$$

As expected, the expected gain increases linearly with time. Since the random variables are uncorrelated, the variance is given by the sum of the variances (as can be verified explicitly):

$$\begin{aligned} V \left[\sum_{k=1}^n G(k) \right] &= \sum_{k=1}^n V [G(k)] \\ &= \sum_{k=1}^n \left\{ e^{2\delta} \exp(\sigma_v^2) \left[\exp(\sigma_v^2) - 1 \right] \right\} \\ &= n \left\{ e^{2\delta} \exp(\sigma_v^2) \left[\exp(\sigma_v^2) - 1 \right] \right\} \end{aligned}$$

Thus

$$v_{\text{simp}}(n) = ne^{2\delta} \exp(\sigma_v^2) [\exp(\sigma_v^2) - 1]$$

Thus, in this case, the variance is also the sum of the one-day variances. This makes sense, since the amount invested each day is constant, rather than changing from day to day with each profit or loss.

Now when we calculate the Return/Risk ratio for the simple dollar fractional price returns, we find:

$$R/R \left[\sum_{k=1}^n \left[\frac{p(k) - p(k-1)}{p(k-1)} \right] \right] = \frac{f_{\text{simp}}(n)}{\sqrt{v_{\text{simp}}(n)}} = \sqrt{n} \frac{1 - e^{-\delta} \exp\left(-\frac{\sigma_v^2}{2}\right)}{\sqrt{\exp(\sigma_v^2) - 1}}$$

We see that in this case, the R/R ratio depends, for large times n , on the square root of n , just as in the case of the log price changes. This is because the random variables are again additive, rather than multiplicative. Thus the risk is greatly reduced, by adding to or subtracting from, the position in order to maintain a constant dollar amount all the time. Only in this way can the risk grow at a smaller rate than the return over time, so that the secular trend wins out over the variance. So in this way, an active trading strategy is actually much safer than the Buy & Hold trading strategy. At the same time, though, since the risk in the Buy & Hold strategy is actually mostly on the *upside*, rather than the *downside*, the strategy of adjusting the amount invested to a constant amount can result in the loss of substantial gains, which would otherwise occur due to the exponential growth of the investment. So for the Random Walk model, the Buy & Hold strategy is the optimal one. (Note that the constant-investment strategy described here is actually very similar to what is called the “Dollar Cost Averaging” strategy.)

Effects of Portfolio Diversification

A portfolio consists of a collection of stocks or other assets of N different companies, each with its own price per share. If the logarithmic prices were additive, then computing the portfolio statistics would be simple because the statistics would be Gaussian (assuming the Random Walk model). However, additive log prices imply that the actual (percentage) price returns are multiplied together, which is the case for a single stock in the Buy & Hold strategy but not for different stocks. In the portfolio as a whole it is the actual price returns that must be added together, of course. Hence we will consider the situation of simple returns in a portfolio in which the amount invested in each security is held constant.

We wish to consider the effect of adding together the returns of N securities, with statistically independent returns, on the overall return and risk of the portfolio. In general, even though the distribution is not Gaussian for the actual price returns, due to the statistical independence of the returns it is still true that the expected returns and variances of the individual securities are additive. Thus we may simply add the expressions derived above for the mean values and variances of the returns of the securities, weighted by the market value of each security. If the market values were equal, along with the return and variance of each security, this implies that a factor of N corresponding to N securities multiplies the individual returns, and N also multiplies the

variances. Hence the standard deviation is multiplied by the square root of N , so the square root of N also multiplies the ratio of expected return to risk. Hence the R/R ratio is improved by a factor of the square root of N , for a diversified portfolio consisting of N securities.

We consider the N securities to each have their own secular trend and daily innovation, with the (constant) dollar amount invested in security s being denoted by C_s , so the total simple return for the portfolio is given by:

$$\sum_{s=1}^N C_s \sum_{k=1}^n \left[\frac{p_s(k) - p_s(k-1)}{p_s(k-1)} \right] = \sum_{s=1}^N C_s \sum_{k=1}^n [G_s(k)] = \sum_{s=1}^N C_s \sum_{k=1}^n [e^{\delta_s} \exp v_s(k) - 1]$$

The expected simple price gain is thus given by:

$$E \left[\sum_{s=1}^N C_s \sum_{k=1}^n [G_s(k)] \right] = \sum_{s=1}^N \sum_{k=1}^n E[C_s G_s(k)] = n \sum_{s=1}^N C_s \left[e^{\delta_s} \exp \left(\frac{\sigma_{v_s}^2}{2} \right) - 1 \right]$$

Thus the expected (simple) fractional price gain over n days is given by:

$$f_{\text{simp}}(n; N) = n \sum_{s=1}^N C_s \left[e^{\delta_s} \exp \left(\frac{\sigma_{v_s}^2}{2} \right) - 1 \right]$$

Similarly, since the random variables are uncorrelated, the variance is given by the sum of the variances of the individual securities:

$$\begin{aligned} V \left[\sum_{s=1}^N C_s \sum_{k=1}^n G_s(k) \right] &= \sum_{s=1}^N \sum_{k=1}^n V[C_s G_s(k)] = \sum_{s=1}^N \sum_{k=1}^n C_s^2 V[G_s(k)] \\ &= n \sum_{s=1}^N C_s^2 \left\{ e^{2\delta_s} \exp(\sigma_{v_s}^2) [\exp(\sigma_{v_s}^2) - 1] \right\} \end{aligned}$$

Thus

$$v_{\text{simp}}(n; N) = n \sum_{s=1}^N C_s^2 \left\{ e^{2\delta_s} \exp(\sigma_{v_s}^2) [\exp(\sigma_{v_s}^2) - 1] \right\}$$

Now when we calculate the Return/Risk ratio for the simple dollar fractional price returns for N securities, we find:

$$R/R \left[\sum_{s=1}^N C_s \sum_{k=1}^n \left[\frac{p_s(k) - p_s(k-1)}{p_s(k-1)} \right] \right] = \frac{f_{\text{simp}}(n; N)}{\sqrt{v_{\text{simp}}(n; N)}} = \sqrt{n} \frac{\sum_{s=1}^N C_s \left[e^{\delta_s} \exp \left(\frac{\sigma_{v_s}^2}{2} \right) - 1 \right]}{\sqrt{\sum_{s=1}^N C_s^2 \left\{ e^{2\delta_s} \exp(\sigma_{v_s}^2) [\exp(\sigma_{v_s}^2) - 1] \right\}}}$$

For small secular trend and variance, compared to unity, we find the following approximation:

$$R/R \left[\sum_{s=1}^N C_s \sum_{k=1}^n \left[\frac{p_s(k) - p_s(k-1)}{p_s(k-1)} \right] \right] \approx \sqrt{n} \frac{\sum_{s=1}^N C_s \left(\delta_s + \frac{\sigma_{v_s}^2}{2} \right)}{\sqrt{\sum_{s=1}^N C_s^2 \sigma_{v_s}^2}} \quad (\delta_s, \sigma_{v_s}^2 \ll 1)$$

We see that, once again, the Return/Risk ratio improves with time n , proportionally to \sqrt{n} . However, we now see that, if the secular trend and variance (and amount invested) for each individual security were the same, then the above expression would reduce to the corresponding expression for each individual stock, multiplied by a factor \sqrt{N} . In other words, the Return/Risk ratio (for simple returns) also improves with the number of securities N , roughly in proportion to \sqrt{N} .

In the Random Walk model, the different securities are uncorrelated. Then it can be seen from the above expression that, given the secular trend and variance of each individual security in the portfolio, the optimal R/R ratio is given by adjusting the amounts invested in each security so that the above expression is maximized. If the returns of different securities are correlated, then we can utilize the Markowitz model, which optimizes the return on the portfolio for a given level of risk given the correlation matrix of the overall portfolio. In either case, the greater the “risk aversion”, the lower the risk, but the lower the expected returns as well, from the portfolio. This will be discussed below.

Return/Risk for an Active Trading Strategy

Now we need to investigate how the return and risk are modified by an active trading strategy. We expect the returns to be increased by such a strategy, but the risk is also increased. Given the presence of autocorrelation in the price data, or more generally correlation with past price returns and other fundamental and economic data, we would like to determine to what extent the return and risk are increased by an active trading strategy.

In order to take into account correlation in the log price data, we now admit the possibility that the future returns are a function of past prices and other economic data, whose influence is contained in a small correction to the Random Walk model. We would thus like to express the next day’s (log) return as the sum of a Gaussian random variable, the secular trend, and a small function of all the known past returns and other data:

$$u(n+1) = v(n+1) + \delta + z(n; X[n])$$

The function $v(n+1) \sim N(0, \sigma_v^2)$ is a Gaussian white noise variable, which we call the *innovation*, with mean 0 and variance σ_v^2 . The constant δ represents the *secular trend*. We can call the function $z(n; X(n))$ the *deterministic factor*; it is a function of the past log price returns (and perhaps other data) which partially determines the future returns as a function of the past data. The Random Walk Model specifies that the deterministic factor $z(n; X(n))$ is zero. If the market is not perfectly inefficient, then this factor will be small but nonzero. However, the catch is that this deterministic factor is in general *unknown*. We can attempt to estimate it in terms of various functions of the past price returns data and other data; these functions are known as *technical indicators*. We denote the *estimate* of the *deterministic factor* $z(n; X(n))$ as the *technical indicator* $\tilde{z}(n; X(n))$,

which is distinguished by a tilde. Explicitly, the technical indicator $\tilde{z}(n; X(n))$ is given as follows:

$$\tilde{z}(n; X(n)) \equiv \tilde{z}[u(n), u(n-1), u(n-2) \cdots; \text{other economic data}]$$

We can designate the technical indicator at time n , defined over a short-term time interval of length m units into the past, by:

$$\tilde{z}(n; X(n) | m) \equiv \tilde{z}[u(n), u(n-1), \dots, u(n-m); X(n), X(n-1), \dots, X(n-m); \cdots]$$

In the limit as $m \rightarrow \infty$, we have the case of a stationary time series.

The most basic example of such a technical indicator is the case where it is a linear function of the past log price returns data. This corresponds to an *autoregressive* model, in which the future returns are regressed on the past returns series. In the stationary case, we can define the sample autocorrelation coefficient as follows:

$$\rho(h) \equiv \frac{1}{\|u\|^2} \sum_{n=0}^{\infty} u(n)u(n+h) \quad \|u\|^2 \equiv \sum_{n=0}^{\infty} u^2(n)$$

Then the technical indicator is defined by:

$$\tilde{z}(n) \equiv \sum_{h=1}^{\infty} a(h)u(n-h) \quad a(h) \equiv \sum_{k=1}^{\infty} [\rho^{-1}](h-k)\rho(k)$$

Note that the Linear Prediction coefficients $a(h)$ are a function of the inverse correlation matrix and of the one-step-ahead correlation coefficients. Each of these, in turn, are ratios of quadratic functions of the price series. So this technical indicator is, in reality, a nonlinear function of the price data. However, if we suppose that the time series is stationary, then these LP coefficients are constant, so the technical indicator is a linear function of the past data. However, in the non-stationary case, it becomes a complicated non-linear function of the past data. We may also calculate the (time dependent) correlation of the future returns with higher order functions of the past data, thereby leading to technical indicators that are even more highly nonlinear.

Let us assume that the technical indicator $\tilde{z}(n; X(n))$ is a Gaussian random variable $\tilde{z}(n) \sim N(0, \sigma_{\tilde{z}}^2)$, of zero mean and variance $\sigma_{\tilde{z}}^2$, with a certain correlation coefficient with the log price returns. Since the technical indicator is assumed to be uncorrelated with the innovations $v(n+1)$ (as is everything else), the degree of correlation of the technical indicator with the future (log) price returns is the same as the correlation between the *technical indicator* and the *deterministic factor*:

$$\rho(z(n), \tilde{z}(n)) \equiv \frac{\langle z(n; X(n)) \tilde{z}(n; X(n)) \rangle}{\sqrt{\langle z^2 \rangle \langle \tilde{z}^2 \rangle}}$$

In the nonstationary case, we postulate that this correlation is actually time-dependent. Since the purpose of the technical indicator is to simulate the deterministic factor as closely as possible, the better the technical indicator, the closer this correlation is to perfect correlation. (Generally, however, the correlation will be very small.)

Since we measure the correlation between the technical indicator (assumed to have mean zero) and the actual returns, not just their deterministic part, a more useful expression is the correlation between the technical indicator and the actual returns:

$$\rho(u(n), \tilde{z}(n)) \equiv \frac{\langle u(n; X(n)) \tilde{z}(n; X(n)) \rangle}{\sqrt{\langle u^2 \rangle \langle \tilde{z}^2 \rangle}}$$

However, the variance of the returns is equal to the sum of the variance of the deterministic part and the innovations, since they are uncorrelated:

$$\langle u, \tilde{z} \rangle = \langle (v + z), \tilde{z} \rangle = \langle z, \tilde{z} \rangle \quad \langle u^2 \rangle = \langle v^2 \rangle + \langle z^2 \rangle \quad (\langle v, z \rangle = \langle v, \tilde{z} \rangle = 0)$$

Thus we have the following relationships:

$$\rho_{u\tilde{z}}\sigma_u = \rho(u(n), \tilde{z}(n))\sqrt{\langle u^2 \rangle} = \frac{\langle z(n; X(n)) \tilde{z}(n; X(n)) \rangle}{\sqrt{\langle \tilde{z}^2 \rangle}} = \rho(z(n), \tilde{z}(n))\sqrt{\langle z^2 \rangle} = \rho_{z\tilde{z}}\sigma_z$$

Using this, we would like to express all quantities in terms of the correlations and variances with respect to u rather than z , since u is observable and z is not.

In general, a *trading rule* is connected with any given *technical indicator*, in that the amount invested is equal to a constant investment plus an amount that is an increasing function of the technical indicator. Exactly how the investment should increase as the technical indicator increases is an arbitrary decision. For example, we could simply make an investment that is proportional to the technical indicator itself, $\tilde{z}(n; X(n))$.

However, this technical indicator is supposed to be a model of the log price returns, not the actual fractional price returns. On the other hand, to calculate the actual return from short-term trading, we need to multiply the amount invested by the actual fractional price returns. Thus, in order to make the math simpler and more straightforward, we invest an amount that is proportional to our expectation of the actual (daily) fractional price returns (plus a constant, in general). Thus the amount invested, as a function of time n , is given by:

$$\$I_s(n) \equiv \$C_s(\delta_s) + \$\kappa_s [\exp \tilde{z}_s(n; X(n)) - 1]$$

The constants of proportionality could in general be different for each security s , possibly depending on the secular trend, as well as the volatility, of each security. We take this as our standard trading rule, given a technical indicator that is supposed to model the log price returns. This form also has the advantage that the maximum amount invested on the short side is limited to a 100% short position (100% of $\$ \kappa$), although the margin on the long side is unlimited. So this trading rule has an intrinsic safety feature built in. We assume that this technical indicator is itself a Gaussian random variable with zero mean, except that since it uses only past data, its value is always known. But as a function of all the past log price innovations, which are themselves (hypothetically) Gaussian random variables, the technical indicator is itself a Gaussian random variable if it is a linear function of those variables. If the technical indicator is a nonlinear function of the past data, then it will not in general be Gaussian, but we take it to be approximately Gaussian anyway, for the purpose of calculational expediency.

The actual expected dollar gain per day is given by the expectation of the percentage gain per day multiplied by the amount invested each day. The percentage gain per day is given by:

$$\frac{\Delta P(n)}{P(n)} = \exp u(n+1) - 1$$

Let us define a new variable $\tilde{u}(n)$ that has mean zero, so that we have:

$$\begin{aligned} u(n+1) &\equiv v(n+1) + \delta + z(n; X[n]) \\ &\equiv \delta + \tilde{u}(n+1) \end{aligned}$$

The expected gain per day is thus given by:

$$\begin{aligned} \langle \Delta G(n) \rangle &\equiv \left\langle \$I(n) \frac{\Delta P(n)}{P(n)} \right\rangle \\ &= \left\langle \left\{ \$C + \$\kappa (e^{\tilde{z}} - 1) \right\} (e^{(\delta + \tilde{u})} - 1) \right\rangle \\ &= \$C \langle e^{(\delta + \tilde{u})} - 1 \rangle + \$\kappa_s \langle (e^{\tilde{z}} - 1) (e^{(\delta + \tilde{u})} - 1) \rangle \end{aligned}$$

Using the results from Appendix B (see in particular Case 2), we have for the two expectation values:

$$\begin{aligned} \langle e^{(\delta + \tilde{u})} - 1 \rangle &= \left[\exp \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) - 1 \right] \\ \langle (e^{\tilde{z}} - 1) (e^{(\delta + \tilde{u})} - 1) \rangle &= \exp \left(\frac{\sigma_{\tilde{z}}^2}{2} \right) \left[\exp \left\{ \delta + \frac{\sigma_{\tilde{u}}^2}{2} + \rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} \right\} - 1 \right] - \left[\exp \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) - 1 \right] \end{aligned}$$

Putting all this together, we have for the expected return:

$$\begin{aligned} \langle \Delta G(n) \rangle &= \$C \left[\exp \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) - 1 \right] + \\ &+ \$\kappa \left\{ \exp \left(\frac{\sigma_{\tilde{z}}^2}{2} \right) \left[\exp \left\{ \delta + \frac{\sigma_{\tilde{u}}^2}{2} + \rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} \right\} - 1 \right] - \left[\exp \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) - 1 \right] \right\} \end{aligned}$$

If we make an expansion to first order in the secular trend and variance of the returns, but to all orders in the variance of the trading rules, we find:

$$\langle \Delta G(n) \rangle \approx \left\{ \$C + \$\kappa \left[\exp \left(\frac{\sigma_{\tilde{z}}^2}{2} \right) - 1 \right] \right\} \left[\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right] + \$\kappa \exp \left(\frac{\sigma_{\tilde{z}}^2}{2} \right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + \frac{(\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2}{2} \right]$$

From this form, it can be seen that even if the technical indicator is completely uncorrelated with the price returns, there is still a small gain from short-term trading due to the increased volatility. This is contained in the first term, which is proportional to the volatility of the price returns and the secular trend. However, it can also be seen that this additional gain can equally well be achieved simply by investing an additional constant amount. The second term contains the dependence on the correlation between the

technical indicator and the daily price returns. The leading term shows that the gain from short term trading is simply proportional to the product of the standard deviations of the log price returns and log trading rules and their correlation, as is to be expected. However, if the variance of the trading rules is large, there are higher order terms that enter as well.

To calculate the variance due to short-term trading, we compute:

$$V[\Delta G(n)] = \left\langle \left\{ \$C + \$\kappa (e^{\tilde{z}} - 1) \right\}^2 \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle - \left\langle \left\{ \$C + \$\kappa (e^{\tilde{z}} - 1) \right\} \left(e^{(\delta + \tilde{u})} - 1 \right) \right\rangle^2$$

We may break the first bracket down term-by-term as follows:

$$\begin{aligned} & \left\langle \left\{ \$C + \$\kappa (e^{\tilde{z}} - 1) \right\}^2 \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle = \\ & = \$C^2 \left\langle \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle + 2\$C\$ \kappa \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle + \$\kappa^2 \left\langle \left(e^{\tilde{z}} - 1 \right)^2 \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle \end{aligned}$$

The second bracket may be broken down term-by-term as follows:

$$\begin{aligned} & \left\langle \left\{ \$C + \$\kappa (e^{\tilde{z}} - 1) \right\} \left(e^{(\delta + \tilde{u})} - 1 \right) \right\rangle^2 = \\ & = \$C^2 \left\langle e^{(\delta + \tilde{u})} - 1 \right\rangle^2 + 2\$C\$ \kappa \left\langle e^{(\delta + \tilde{u})} - 1 \right\rangle \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta + \tilde{u})} - 1 \right) \right\rangle + \$\kappa^2 \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta + \tilde{u})} - 1 \right) \right\rangle^2 \end{aligned}$$

In general, even for this simple model of a lognormal distribution, the variance becomes very complicated to calculate. The exact expression is not particularly illuminating, and is very tedious to compute. However, we can compute an exact expression for the first bracket above without too much trouble, corresponding to the buy-and-hold part of the trading strategy. Using the formulas in Appendix B, we find for the first term in the first bracket:

$$\begin{aligned} \left\langle \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle & = e^{2\delta} \left\langle e^{2\tilde{u}} \right\rangle - 2e^\delta \left\langle e^{\tilde{u}} \right\rangle + 1 \\ & = e^{2\delta} \exp\left(2\sigma_{\tilde{u}}^2\right) - 2e^\delta \exp\left(\frac{\sigma_{\tilde{u}}^2}{2}\right) + 1 \end{aligned}$$

We previously arrived at the result:

$$\left\langle e^{(\delta + \tilde{u})} - 1 \right\rangle = \left[\exp\left(\delta + \frac{\sigma_{\tilde{u}}^2}{2}\right) - 1 \right]$$

The first bracket may then be re-written as:

$$\begin{aligned} \left\langle \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle & = e^{2\delta} \left\{ \exp\left(2\sigma_{\tilde{u}}^2\right) - \exp\left(\sigma_{\tilde{u}}^2\right) \right\} + \left\{ e^{2\delta} \exp\left(\sigma_{\tilde{u}}^2\right) - 2e^\delta \exp\left(\frac{\sigma_{\tilde{u}}^2}{2}\right) + 1 \right\} \\ & = e^{2\delta} \exp\left(\sigma_{\tilde{u}}^2\right) \left\{ \exp\left(\sigma_{\tilde{u}}^2\right) - 1 \right\} + \left\{ e^\delta \exp\left(\frac{\sigma_{\tilde{u}}^2}{2}\right) - 1 \right\}^2 \\ & = e^{2\delta} \exp\left(\sigma_{\tilde{u}}^2\right) \left\{ \exp\left(\sigma_{\tilde{u}}^2\right) - 1 \right\} + \left\langle e^{(\delta + \tilde{u})} - 1 \right\rangle^2 \end{aligned}$$

Thus the exact form for the variance due to the first bracket is:

$$\boxed{\left\langle \left(e^{(\delta+\tilde{u})} - 1 \right)^2 \right\rangle - \left\langle e^{(\delta+\tilde{u})} - 1 \right\rangle^2 = e^{2\delta} \exp(\sigma_{\tilde{u}}^2) \left\{ \exp(\sigma_{\tilde{u}}^2) - 1 \right\} \approx \sigma_{\tilde{u}}^2}$$

As expected, the variance for the buy-and-hold part of the trading strategy is approximately given by the variance of the log price returns themselves. However, if this variance were large, there would be higher-order corrections due to the fact that the distribution of the actual prices is lognormal, not normal.

The second term in the first bracket in the expression for the variance is given by:

$$\begin{aligned} \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right)^2 \right\rangle &= \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{2(\delta+\tilde{u})} - 2e^{(\delta+\tilde{u})} + 1 \right) \right\rangle \\ &= \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{2(\delta+\tilde{u})} - 1 \right) \right\rangle - 2 \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right) \right\rangle \end{aligned}$$

From Case 1 in Appendix B, the first term may be approximated by:

$$\left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{2(\delta+\tilde{u})} - 1 \right) \right\rangle \approx 2 \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2 \right] + 2 \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \left(\delta + \sigma_{\tilde{u}}^2 \right)$$

From Case 2 in Appendix B, the second term may be approximated by:

$$\left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right) \right\rangle \approx \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + \frac{(\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2}{2} \right] + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right)$$

Thus we find:

$$\begin{aligned} \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right)^2 \right\rangle &= \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{2(\delta+\tilde{u})} - 1 \right) \right\rangle - 2 \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right) \right\rangle \\ &\approx \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \right\} \sigma_{\tilde{u}}^2 \end{aligned}$$

Using results arrived at previously, the second term in the second bracket is given by:

$$\begin{aligned} \left\langle e^{(\delta+\tilde{u})} - 1 \right\rangle \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right) \right\rangle &\approx \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + \frac{(\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2}{2} \right] + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) \right\} \\ &\approx 0 \end{aligned}$$

To first order in the secular trend and variance of the log price returns, this term is zero.

Thus the second term of the variance is given by:

$$\boxed{\begin{aligned} \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right)^2 \right\rangle - \left\langle e^{(\delta+\tilde{u})} - 1 \right\rangle \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta+\tilde{u})} - 1 \right) \right\rangle &\approx \\ \approx \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \right\} \sigma_{\tilde{u}}^2 \end{aligned}}$$

To lowest order in the variance of the trading rules, this term is proportional to the product of the variance of the trading rules and the variance of the log price returns. Note that it is nonzero even if the correlation between the trading rules and the log price returns is zero. This is again due to the fact that the distribution of the actual prices is lognormal, not normal.

The third term of the first bracket is given by:

$$\begin{aligned} \left\langle (e^{\bar{z}} - 1)^2 (e^{(\delta+\bar{u})} - 1)^2 \right\rangle &= \left\langle (e^{2\bar{z}} - 2e^{\bar{z}} + 1)(e^{2(\delta+\bar{u})} - 2e^{(\delta+\bar{u})} + 1) \right\rangle \\ &= \left\langle (e^{2\bar{z}} - 1)(e^{2(\delta+\bar{u})} - 1) \right\rangle - 2 \left\langle (e^{2\bar{z}} - 1)(e^{(\delta+\bar{u})} - 1) \right\rangle - \\ &\quad - 2 \left\langle (e^{\bar{z}} - 1)(e^{2(\delta+\bar{u})} - 1) \right\rangle + 4 \left\langle (e^{\bar{z}} - 1)(e^{(\delta+\bar{u})} - 1) \right\rangle \end{aligned}$$

From Case 3 in Appendix B, the first term may be approximated by:

$$\left\langle (e^{2\bar{z}} - 1)(e^{2(\delta+\bar{u})} - 1) \right\rangle \approx 4 \exp(2\sigma_{\bar{z}}^2) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + 2(\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2 \right] + 2 \left[\exp(2\sigma_{\bar{z}}^2) - 1 \right] (\delta + \sigma_{\bar{u}}^2)$$

From Case 4 in Appendix B, the second term may be approximated by:

$$\left\langle (e^{2\bar{z}} - 1)(e^{(\delta+\bar{u})} - 1) \right\rangle \approx 2 \exp(2\sigma_{\bar{z}}^2) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + (\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2 \right] + \left[\exp(2\sigma_{\bar{z}}^2) - 1 \right] \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \right)$$

The sum of the first two terms is therefore given by:

$$\begin{aligned} \left\langle (e^{2\bar{z}} - 1)(e^{2(\delta+\bar{u})} - 1) \right\rangle - 2 \left\langle (e^{2\bar{z}} - 1)(e^{(\delta+\bar{u})} - 1) \right\rangle &\approx \\ &\approx \left\{ 4 \exp(2\sigma_{\bar{z}}^2) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + 2(\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2 \right] + 2 \left[\exp(2\sigma_{\bar{z}}^2) - 1 \right] (\delta + \sigma_{\bar{u}}^2) \right\} \\ &\quad - 2 \left\{ 2 \exp(2\sigma_{\bar{z}}^2) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + (\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2 \right] + \left[\exp(2\sigma_{\bar{z}}^2) - 1 \right] \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \right) \right\} \\ &= \left\{ 4 \exp(2\sigma_{\bar{z}}^2) (\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}})^2 + \left[\exp(2\sigma_{\bar{z}}^2) - 1 \right] \right\} \sigma_{\bar{u}}^2 \end{aligned}$$

From Case 1 in Appendix B, the third term may be approximated by:

$$\left\langle (e^{\bar{z}} - 1)(e^{2(\delta+\bar{u})} - 1) \right\rangle \approx 2 \exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + (\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2 \right] + 2 \left[\exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) - 1 \right] (\delta + \sigma_{\bar{u}}^2)$$

From Case 2 in Appendix B, the fourth term may be approximated by:

$$\left\langle (e^{\bar{z}} - 1)(e^{(\delta+\bar{u})} - 1) \right\rangle \approx \exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + \frac{(\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2}{2} \right] + \left[\exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) - 1 \right] \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \right)$$

The sum of the last two terms is therefore given by:

$$\begin{aligned} -2 \left\langle (e^{\bar{z}} - 1)(e^{2(\delta+\bar{u})} - 1) \right\rangle + 4 \left\langle (e^{\bar{z}} - 1)(e^{(\delta+\bar{u})} - 1) \right\rangle &\approx \\ &\approx -2 \left\{ 2 \exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + (\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2 \right] + 2 \left[\exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) - 1 \right] (\delta + \sigma_{\bar{u}}^2) \right\} \\ &\quad + 4 \left\{ \exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) \left[\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}} + \frac{(\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}} \sigma_{\bar{u}})^2}{2} \right] + \left[\exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) - 1 \right] \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \right) \right\} \\ &= -2 \left\{ \exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) (\rho_{\bar{z}\bar{u}} \sigma_{\bar{z}})^2 + \left[\exp\left(\frac{\sigma_{\bar{z}}^2}{2}\right) - 1 \right] \right\} \sigma_{\bar{u}}^2 \end{aligned}$$

Using results arrived at previously, the third term in the second bracket is given by:

$$\begin{aligned} \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta + \tilde{u})} - 1 \right) \right\rangle^2 &\approx \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + \frac{(\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2}{2} \right] + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) \right\}^2 \\ &\approx \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} \right\}^2 = \exp(\sigma_{\tilde{z}}^2) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2 \end{aligned}$$

Thus we have for the third bracket of the variance:

$$\begin{aligned} \left\langle \left(e^{\tilde{z}} - 1 \right)^2 \left(e^{(\delta + \tilde{u})} - 1 \right)^2 \right\rangle - \left\langle \left(e^{\tilde{z}} - 1 \right) \left(e^{(\delta + \tilde{u})} - 1 \right) \right\rangle^2 &\approx \\ &= \left\{ 4 \exp(2\sigma_{\tilde{z}}^2) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp(2\sigma_{\tilde{z}}^2) - 1 \right] \right\} \sigma_{\tilde{u}}^2 \\ &\quad - 2 \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \right\} \sigma_{\tilde{u}}^2 \\ &\quad - \exp(\sigma_{\tilde{z}}^2) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 \sigma_{\tilde{u}}^2 \end{aligned}$$

Note that this result is again proportional to the variance of the log price returns, to this order of approximation.

We can now write the final result for the variance, valid to first order in the secular trend and variance of the log price returns, as follows:

$$\begin{aligned} V[\Delta G(n)] &\approx \$C^2 \sigma_{\tilde{u}}^2 + \\ &\quad + 2\$C\$K \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \right\} \sigma_{\tilde{u}}^2 + \\ &\quad + \$K^2 \left\{ 4 \exp(2\sigma_{\tilde{z}}^2) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp(2\sigma_{\tilde{z}}^2) - 1 \right] \right\} \sigma_{\tilde{u}}^2 \\ &\quad - 2\$K^2 \left\{ \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] \right\} \sigma_{\tilde{u}}^2 \\ &\quad - \$K^2 \exp(\sigma_{\tilde{z}}^2) (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}})^2 \sigma_{\tilde{u}}^2 \end{aligned}$$

In this expression, the variance is a complicated function of the variance of the technical indicator, multiplied by the variance of the log price returns. Hence for a given trading strategy, the risk is directly proportional to the log price variance. For a given log price variance, the risk is a complicated function of the variance of the technical indicator, but it generally increases exponentially with the variance of the technical indicator.

To get an estimate of the ratio of return to risk for short-term trading, in general we see that this is going to be very complicated. However, we can get a substantial simplification if we consider the situation where the variance due to the short-term trading rules is much greater than that due to the log price returns themselves. Then we can consider only the dominant terms in the variance of the technical indicator, and drop all the other terms. In the expression for the expected return, if we consider the

approximation in which the secular trend and variance of the log price returns is negligible compared to the variance of the trading rules, then we obtain:

$$E[\Delta G(n)] \approx \$\kappa \exp\left(\frac{\sigma_z^2}{2}\right) [\exp(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0\right)$$

In this approximation, we obtain for Case 1:

$$\langle (e^{\bar{z}} - 1)(e^{2(\delta + \bar{u})} - 1) \rangle \approx \exp\left(\frac{\sigma_z^2}{2}\right) [\exp 2(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0\right)$$

In this approximation, we obtain for Case 2:

$$\langle (e^{\bar{z}} - 1)(e^{(\delta + \bar{u})} - 1) \rangle \approx \exp\left(\frac{\sigma_z^2}{2}\right) [\exp(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0\right)$$

In this approximation, we obtain for Case 3:

$$\langle (e^{2\bar{z}} - 1)(e^{2(\delta + \bar{u})} - 1) \rangle \approx \exp(2\sigma_z^2) [\exp 4(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0\right)$$

In this approximation, we obtain for Case 4:

$$\langle (e^{2\bar{z}} - 1)(e^{(\delta + \bar{u})} - 1) \rangle \approx \exp(2\sigma_z^2) [\exp 2(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0\right)$$

In this approximation, we keep only the largest terms in the variance of the technical indicator, so we drop the first and second brackets in the expression for the variance, and keep only the two largest terms in the third bracket. We thus have:

$$\begin{aligned} \langle (e^{2\bar{z}} - 1)(e^{2(\delta + \bar{u})} - 1) \rangle - 2 \langle (e^{2\bar{z}} - 1)(e^{(\delta + \bar{u})} - 1) \rangle &\approx \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0, \quad \sigma_z^2 \gg 0\right) \\ &\approx \exp(2\sigma_z^2) \{ [\exp 4(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] - 2 [\exp 2(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1] \} \\ &= \exp(2\sigma_z^2) [\exp 2(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1]^2 \end{aligned}$$

Thus the variance is given in this approximation by:

$$V[\Delta G(n)] \approx \$\kappa^2 \exp(2\sigma_z^2) [\exp 2(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) - 1]^2 \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0, \quad \sigma_z^2 \gg 0\right)$$

This gives us an estimate for the Reward/Risk ratio, in this approximation:

$$R/R \equiv \frac{E[\Delta G(n)]}{\sqrt{V[\Delta G(n)]}} \approx \left\{ \exp\left(\frac{\sigma_z^2}{2}\right) [\exp(\rho_{\bar{u}\bar{z}} \sigma_{\bar{u}} \sigma_{\bar{z}}) + 1] \right\}^{-1} \quad \left(\delta + \frac{\sigma_{\bar{u}}^2}{2} \approx 0, \quad \sigma_z^2 \gg 0\right)$$

Note that this formula may not be accurate when the correlation is near zero, since both the expected return and the variance approach zero. In this case, the lower order terms will evidently dominate. In fact, from the above expression for the variance, keeping only the highest-order terms (in the variance of the technical indicator), we find the

following approximate expression for the variance of the returns in the case that the correlation is zero (and the technical indicators do not work):

$$V[\Delta G(n)] \approx \left\{ \$C^2 + \$\kappa^2 \left[\exp(2\sigma_z^2) - 1 \right] \right\} \sigma_u^2 \quad (\rho_{z\bar{u}} = 0, \quad \sigma_z^2 \gg 0)$$

This formula shows that the risk due to short-term trading, when the correlation is zero, increases exponentially with the variance of the technical indicator, and is proportional to the variance of the log price returns (in the approximation that this is small). However, when the correlation is substantial, most of the “risk” may be on the upside, not the downside. In fact, it can be seen from the above formula that the “risk” actually increases as the correlation increases.

Thus we see that, for active trading, even though the expected returns increase exponentially with the variance of the trading rules, the variance of the returns increases so much faster that the ratio of return to risk decreases exponentially with the variance of the trading rules. This is true no matter how well correlated the technical indicators are with the log price returns. Thus, to control risk, it is necessary to do short-term trading using only a fraction of the total portfolio equity. To trade with the entire portfolio equity is risky in the extreme.

Hence we see that the variance of the outcome depends in a very complicated way on the variance of the trading rules. However, generally speaking, the variance of the trading rules plays a similar role to the log variance of the price returns. But since the former is likely to be much greater than the latter, it will usually be the determining factor in the risk.

(Consider also small variance of technical indicators, corresponding to a “smooth” indicator.)

Appendix A: Moment Generating Functions

When considering lognormal random variables, the expectation values of interest to us take the form of moment-generating functions, or are directly related to them. For a random variable Z , Ghahramani (1996) defines the moment-generating function:

$$M_Z(t) = E[\exp(tZ)]$$

We are especially interested in Gaussian random variables. For a *standard Normal* variable (a Gaussian random variable with mean zero and unit variance), the probability distribution function is:

$$f(z)dz \equiv \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) dz$$

Hence the moment generating function is given by a Gaussian integral:

$$M_Z(t) = \int \exp(tz) f(z) dz \equiv \frac{1}{\sqrt{2\pi}} \int \exp(tz) \exp\left(-\frac{z^2}{2}\right) dz$$

We then solve this Gaussian integral as follows [Ghahramani (1996), p.346]. We first complete the square in the exponent:

$$-\frac{z^2}{2} + tz = -\frac{1}{2}(z^2 - 2tz + t^2 - t^2) = -\frac{1}{2}([z-t]^2 - t^2)$$

This yields for the moment-generating function:

$$M_z(t) = \frac{1}{\sqrt{2\pi}} \int \exp\left\{-\frac{1}{2}[z-t]^2 + \frac{1}{2}t^2\right\} dz$$

We now define a new variable as follows:

$$u \equiv z - t \quad du = dz$$

Then the integral becomes:

$$M_z(t) = \exp\left(\frac{1}{2}t^2\right) \frac{1}{\sqrt{2\pi}} \int \exp\left\{-\frac{1}{2}u^2\right\} du$$

We note the following value of the Gaussian integral:

$$\frac{1}{\sqrt{2\pi}} \int \exp\left\{-\frac{1}{2}u^2\right\} du = 1$$

Thus we see that the moment-generating function for a single standard Normal variable has the value:

$$\boxed{M_z(t) = E[\exp(tZ)] = \exp\left(\frac{t^2}{2}\right)}$$

Thus, if Z is a standard Normal random variable, representing a lognormal distribution of returns, then the moment-generating function becomes a convenient method of finding the expectation values of the actual price returns.

We now repeat the above calculation for the case of a sum of two standard Normal random variables Z_1 and Z_2 , which are correlated with correlation coefficient ρ . In this case we define the moment-generating function as follows:

$$\boxed{M_{z_1, z_2}(t_1, t_2) = E[\exp(t_1 Z_1 + t_2 Z_2)]}$$

It is a theorem [Ghahramani (1996), p.414] that the joint probability distribution for two such standard Normal random variables has the general form:

$$Q(z_1, z_2) \equiv z_1^2 - 2\rho z_1 z_2 + z_2^2$$

$$f(z_1, z_2) dz_1 dz_2 \equiv \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left\{-\frac{1}{2} \frac{Q(z_1, z_2)}{(1-\rho^2)}\right\} dz_1 dz_2$$

In order to reduce this to Gaussian form, we must first complete the square in the quadratic form Q . To do this, let us write it in the following form:

$$a^2(z_1 + z_2)^2 + b^2(z_1 - z_2)^2 = (a^2 + b^2)(z_1^2 + z_2^2) + 2(a^2 - b^2)z_1 z_2 \equiv \frac{Q(z_1, z_2)}{(1-\rho^2)}$$

To equate this to the quadratic form Q , we make the following assignments:

$$(a^2 + b^2) = \frac{1}{(1-\rho^2)} \quad (a^2 - b^2) = \frac{-\rho}{(1-\rho^2)}$$

Now we may solve for the coefficients:

$$a^2 = \frac{1}{2} \frac{1-\rho}{(1-\rho^2)} = \frac{1}{2(1+\rho)} \quad b^2 = \frac{1}{2} \frac{1+\rho}{(1-\rho^2)} = \frac{1}{2(1-\rho)}$$

We now define two new variables as follows:

$$u_+ \equiv a(z_1 + z_2) = \frac{(z_1 + z_2)}{\sqrt{2(1+\rho)}} \quad u_- \equiv b(z_1 - z_2) = \frac{(z_1 - z_2)}{\sqrt{2(1-\rho)}}$$

We may write this transformation as follows:

$$\begin{bmatrix} u_+ \\ u_- \end{bmatrix} \equiv \begin{bmatrix} +a & +a \\ +b & -b \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \quad |\det J| \equiv \left| \frac{\partial u}{\partial z} \right| = 2ab = \frac{1}{\sqrt{1-\rho^2}}$$

Thus we have in the Gaussian integral:

$$u_+^2 + u_-^2 = \frac{Q(z_1, z_2)}{(1-\rho^2)} \quad du_+ du_- = \left| \frac{\partial u}{\partial z} \right| dz_1 dz_2 = \frac{1}{\sqrt{1-\rho^2}} dz_1 dz_2$$

Finally the joint probability distribution takes the form:

$$f(z_1, z_2) dz_1 dz_2 = \frac{1}{2\pi} \exp\left\{-\frac{1}{2}(u_+^2 + u_-^2)\right\} du_+ du_-$$

Thus we have arrived at a diagonal form for the joint distribution function of the two correlated standard Normal random variables.

To proceed, we now need to express the old variables in terms of the new ones. We have:

$$z_1 = \frac{1}{2} \left(\frac{u_+}{a} + \frac{u_-}{b} \right) \quad z_2 = \frac{1}{2} \left(\frac{u_+}{a} - \frac{u_-}{b} \right)$$

Thus:

$$t_1 z_1 + t_2 z_2 = \frac{t_1}{2} \left(\frac{u_+}{a} + \frac{u_-}{b} \right) + \frac{t_2}{2} \left(\frac{u_+}{a} - \frac{u_-}{b} \right) = \frac{t_1 + t_2}{2a} u_+ + \frac{t_1 - t_2}{2b} u_- \equiv t_+ u_+ + t_- u_-$$

Now we may write the moment-generating function as:

$$\begin{aligned} M_{z_1, z_2}(t_1, t_2) &= \iint \exp(t_1 z_1 + t_2 z_2) f(z_1, z_2) dz_1 dz_2 \\ &= \frac{1}{2\pi} \iint \exp(t_+ u_+ + t_- u_-) \exp\left\{-\frac{1}{2}(u_+^2 + u_-^2)\right\} du_+ du_- \end{aligned}$$

Now it can be seen that this moment-generating function has been reduced to the product of two independent Gaussian integrals of a single variable. In fact, in terms of the one-variable moment-generating function derived above, we have:

$$M_{z_1, z_2}(t_1, t_2) = M_{U_+}(t_+) M_{U_-}(t_-) = \exp\left(\frac{t_+^2 + t_-^2}{2}\right)$$

Now it is just a matter of substituting back the original t variables:

$$t_+^2 = \left(\frac{t_1 + t_2}{2a} \right)^2 = (1+\rho) \frac{(t_1 + t_2)^2}{2} \quad t_-^2 = \left(\frac{t_1 - t_2}{2b} \right)^2 = (1-\rho) \frac{(t_1 - t_2)^2}{2}$$

Thus we find for the final result:

$$M_{z_1, z_2}(t_1, t_2) = \exp \left\{ (1 + \rho) \left(\frac{t_1 + t_2}{2} \right)^2 + (1 - \rho) \left(\frac{t_1 - t_2}{2} \right)^2 \right\}$$

This result may be re-arranged in a form that is a little easier to work with. The sum of the two terms involving the t 's is given by:

$$\left(\frac{t_1 + t_2}{2} \right)^2 + \left(\frac{t_1 - t_2}{2} \right)^2 = \left(\frac{t_1^2 + 2t_1t_2 + t_2^2}{4} \right) + \left(\frac{t_1^2 - 2t_1t_2 + t_2^2}{4} \right) = \left(\frac{t_1^2 + t_2^2}{2} \right)$$

The difference of the two terms yields:

$$\left(\frac{t_1 + t_2}{2} \right)^2 - \left(\frac{t_1 - t_2}{2} \right)^2 = \left(\frac{t_1^2 + 2t_1t_2 + t_2^2}{4} \right) - \left(\frac{t_1^2 - 2t_1t_2 + t_2^2}{4} \right) = t_1t_2$$

Thus the above expression may be written in the nicer form:

$$M_{z_1, z_2}(t_1, t_2) = \exp \left\{ \frac{t_1^2 + t_2^2}{2} + \rho t_1 t_2 \right\}$$

In case the two standard Normal random variables are uncorrelated, it can be seen that this expression reduces to the product of the moment-generating functions for the individual variables, as it should. If they are perfectly correlated, on the other hand, then the two Z variables are equivalent, and the sum of the two t 's plays the role of a single t , and once again the result coincides with the single-variable result.

Appendix B: Mean Values and Variances

In this appendix, we calculate some mean values and variances that are needed in the main body of the text. Suppose u is a Normal variable with variance σ_u^2 . Then we can write u in terms of the standard Normal variable Z (with unit variance) as $u \equiv \sigma_u Z$. Then from the result in Appendix A, the expectation value of the exponential of this variable is given by:

$$M_u(t) \equiv E[\exp(tu)] \equiv E[\exp(t\sigma_u Z)] = \exp \left(\frac{(t\sigma_u)^2}{2} \right)$$

We thus have our first result:

$$E[\exp(u) - 1] \equiv \langle \exp(u) - 1 \rangle = \exp \left(\frac{\sigma_u^2}{2} \right) - 1$$

When the variance is small, this expectation value is approximately given by:

$$E[\exp(u) - 1] \approx \frac{\sigma_u^2}{2}$$

This is due to the fact that the distribution is lognormal, and hence the distribution for the actual prices has a "fatter tail" in the direction of higher prices.

Let us now calculate the variance of this quantity. We have:

$$\begin{aligned}
\langle [\exp(u) - 1]^2 \rangle &= \langle \exp(2u) - 2\exp(u) + 1 \rangle \\
&= \exp\left(4\frac{\sigma_u^2}{2}\right) - 2\exp\left(\frac{\sigma_u^2}{2}\right) + 1 \\
&= \left\{ \exp(2\sigma_u^2) - \exp(\sigma_u^2) \right\} + \left\{ \exp(\sigma_u^2) - 2\exp\left(\frac{\sigma_u^2}{2}\right) + 1 \right\} \\
&= \exp(\sigma_u^2) [\exp(\sigma_u^2) - 1] + \langle \exp(u) - 1 \rangle^2
\end{aligned}$$

Thus the variance is given by:

$$V[\exp(u) - 1] \equiv \langle [\exp(u) - 1]^2 \rangle - \langle \exp(u) - 1 \rangle^2 = \exp(\sigma_u^2) [\exp(\sigma_u^2) - 1]$$

It can be seen that when the variance of u is small compared to unity, this expression reduces to the variance of u :

$$V[\exp(u) - 1] \approx \sigma_u^2$$

This makes sense, since it is approximately the variance of the logarithmic returns.

For our next expectation value, consider two variables u_1 and u_2 , with correlation coefficient ρ_{12} . The expectation of the exponentials of these variables was derived in Appendix A and is given by:

$$M_{u_1, u_2}(t_1, t_2) \equiv E[\exp(t_1 u_1 + t_2 u_2)] = \exp\left\{ \frac{(t_1 \sigma_{u_1})^2 + (t_2 \sigma_{u_2})^2}{2} + \rho_{12} (t_1 \sigma_{u_1})(t_2 \sigma_{u_2}) \right\}$$

Now we have:

$$\begin{aligned}
\langle [\exp(t_1 u_1) - 1][\exp(t_2 u_2) - 1] \rangle &= \langle \exp(t_1 u_1 + t_2 u_2) \rangle - \langle \exp(t_1 u_1) \rangle - \langle \exp(t_2 u_2) \rangle + 1 \\
&= \langle \exp(t_1 u_1 + t_2 u_2) - 1 \rangle - \langle \exp(t_1 u_1) - 1 \rangle - \langle \exp(t_2 u_2) - 1 \rangle
\end{aligned}$$

Using the previous results, we find:

$$\begin{aligned}
\langle [\exp(t_1 u_1) - 1][\exp(t_2 u_2) - 1] \rangle &= \left[\exp\left\{ \frac{(t_1 \sigma_{u_1})^2 + (t_2 \sigma_{u_2})^2}{2} + \rho_{12} (t_1 \sigma_{u_1})(t_2 \sigma_{u_2}) \right\} - 1 \right] \\
&\quad - \left[\exp\left(\frac{(t_1 \sigma_{u_1})^2}{2} \right) - 1 \right] - \left[\exp\left(\frac{(t_2 \sigma_{u_2})^2}{2} \right) - 1 \right]
\end{aligned}$$

It is clear by inspection that when the variances are small compared to unity, this expression reduces to the following simple form:

$$\langle [\exp(t_1 u_1) - 1][\exp(t_2 u_2) - 1] \rangle \approx \rho_{12} (t_1 \sigma_{u_1})(t_2 \sigma_{u_2})$$

This makes sense, since it is approximately the covariance between the two logarithmic returns $t_1 u_1$ and $t_2 u_2$.

Finally, the results are modified if the mean value of the log variables is not zero. In this case we add a constant to the variables to represent the mean value. Thus we have:

$$\begin{aligned}\langle \rangle &\equiv \langle [\exp(\delta_1 + t_1 u_1) - 1][\exp(\delta_2 + t_2 u_2) - 1] \rangle \\ &= \langle e^{\delta_1 + \delta_2} \exp(t_1 u_1 + t_2 u_2) \rangle - \langle e^{\delta_1} \exp(t_1 u_1) \rangle - \langle e^{\delta_2} \exp(t_2 u_2) \rangle + 1 \\ &= \langle e^{\delta_1 + \delta_2} \exp(t_1 u_1 + t_2 u_2) - 1 \rangle - \langle e^{\delta_1} \exp(t_1 u_1) - 1 \rangle - \langle e^{\delta_2} \exp(t_2 u_2) - 1 \rangle\end{aligned}$$

This leads to the result:

$$\begin{aligned}\langle \rangle &\equiv \langle [\exp(\delta_1 + t_1 u_1) - 1][\exp(\delta_2 + t_2 u_2) - 1] \rangle \\ &= \left[\exp \left\{ \delta_1 + \delta_2 + \frac{(t_1 \sigma_{u_1})^2 + (t_2 \sigma_{u_2})^2}{2} + \rho_{12} (t_1 \sigma_{u_1})(t_2 \sigma_{u_2}) \right\} - 1 \right] \\ &\quad - \left[\exp \left(\delta_1 + \frac{(t_1 \sigma_{u_1})^2}{2} \right) - 1 \right] - \left[\exp \left(\delta_2 + \frac{(t_2 \sigma_{u_2})^2}{2} \right) - 1 \right]\end{aligned}$$

When we keep terms of second order in the mean values, and first order in the variances and covariances, we arrive at the following approximate result:

$$\langle [\exp(\delta_1 + t_1 u_1) - 1][\exp(\delta_2 + t_2 u_2) - 1] \rangle \approx \delta_1 \delta_2 + \rho_{12} (t_1 \sigma_{u_1})(t_2 \sigma_{u_2})$$

These expectation values are all that we need to calculate any mean values or variances involving two independent variables, such as the daily price returns and the daily values of some technical indicator.

We are going to need four special cases of the above formula for the calculation of variance in the main text. We expect the secular trend δ and the variance of the log price returns σ_u^2 to be small relative to unity, so we keep only the first order terms in these quantities. However, the variance of the technical indicator (trading rules) σ_z^2 could be arbitrarily large, so we calculate the exact result in terms of this quantity.

Case 1: We wish to calculate $\langle (e^z - 1)(e^{2(\delta + \bar{u})} - 1) \rangle$:

$$\begin{aligned}\langle (e^z - 1)(e^{2(\delta + \bar{u})} - 1) \rangle &= \left[\exp \left\{ 2\delta + \frac{(\sigma_z)^2 + (2\sigma_{\bar{u}})^2}{2} + \rho_{z\bar{u}} (\sigma_z)(2\sigma_{\bar{u}}) \right\} - 1 \right] - \\ &\quad - \left[\exp \left(\frac{(\sigma_z)^2}{2} \right) - 1 \right] - \left[\exp \left(2\delta + \frac{(2\sigma_{\bar{u}})^2}{2} \right) - 1 \right]\end{aligned}$$

This may be rewritten in the following form:

$$\langle (e^z - 1)(e^{2(\delta + \bar{u})} - 1) \rangle = \exp \left(\frac{\sigma_z^2}{2} \right) \left[\exp 2 \{ \delta + \sigma_{\bar{u}}^2 + \rho_{z\bar{u}} \sigma_z \sigma_{\bar{u}} \} - 1 \right] - \left[\exp 2(\delta + \sigma_{\bar{u}}^2) - 1 \right]$$

When we keep terms of first order in the mean value and variance of u , and to all orders in the variance of z , we arrive at the following approximate result:

$$\langle (e^{\tilde{z}} - 1)(e^{2(\delta + \tilde{u})} - 1) \rangle \approx 2 \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + (\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2 \right] + 2 \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] (\delta + \sigma_{\tilde{u}}^2)$$

Note that only the first term depends on the correlation between z and u .

Case 2: We wish to calculate $\langle (e^{\tilde{z}} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle$:

$$\begin{aligned} \langle (e^{\tilde{z}} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle &= \left[\exp\left\{ \delta + \frac{(\sigma_{\tilde{z}})^2 + (\sigma_{\tilde{u}})^2}{2} + \rho_{\tilde{z}\tilde{u}} (\sigma_{\tilde{z}})(\sigma_{\tilde{u}}) \right\} - 1 \right] - \\ &\quad - \left[\exp\left(\frac{(\sigma_{\tilde{z}})^2}{2}\right) - 1 \right] - \left[\exp\left(\delta + \frac{(\sigma_{\tilde{u}})^2}{2}\right) - 1 \right] \end{aligned}$$

This may be rewritten in the following form:

$$\langle (e^{\tilde{z}} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle = \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\exp\left\{ \delta + \frac{\sigma_{\tilde{u}}^2}{2} + \rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} \right\} - 1 \right] - \left[\exp\left(\delta + \frac{\sigma_{\tilde{u}}^2}{2}\right) - 1 \right]$$

When we keep terms of first order in the mean value and variance of u , and to all orders in the variance of z , we arrive at the following approximate result:

$$\langle (e^{\tilde{z}} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle \approx \exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + \frac{(\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2}{2} \right] + \left[\exp\left(\frac{\sigma_{\tilde{z}}^2}{2}\right) - 1 \right] (\delta + \frac{\sigma_{\tilde{u}}^2}{2})$$

Note that only the first term depends on the correlation between z and u .

Case 3: We wish to calculate $\langle (e^{2\tilde{z}} - 1)(e^{2(\delta + \tilde{u})} - 1) \rangle$:

$$\begin{aligned} \langle (e^{2\tilde{z}} - 1)(e^{2(\delta + \tilde{u})} - 1) \rangle &= \left[\exp\left\{ 2\delta + \frac{(2\sigma_{\tilde{z}})^2 + (2\sigma_{\tilde{u}})^2}{2} + \rho_{\tilde{z}\tilde{u}} (2\sigma_{\tilde{z}})(2\sigma_{\tilde{u}}) \right\} - 1 \right] - \\ &\quad - \left[\exp\left(\frac{(2\sigma_{\tilde{z}})^2}{2}\right) - 1 \right] - \left[\exp\left(2\delta + \frac{(2\sigma_{\tilde{u}})^2}{2}\right) - 1 \right] \end{aligned}$$

This may be rewritten in the following form:

$$\langle (e^{2\tilde{z}} - 1)(e^{2(\delta + \tilde{u})} - 1) \rangle = \exp(2\sigma_{\tilde{z}}^2) \left[\exp 2\left\{ \delta + \sigma_{\tilde{u}}^2 + 2\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} \right\} - 1 \right] - \left[\exp 2(\delta + \sigma_{\tilde{u}}^2) - 1 \right]$$

When we keep terms of first order in the mean value and variance of u , and to all orders in the variance of z , we arrive at the following approximate result:

$$\langle (e^{2\tilde{z}} - 1)(e^{2(\delta + \tilde{u})} - 1) \rangle \approx 4 \exp(2\sigma_{\tilde{z}}^2) \left[\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}} + 2(\rho_{\tilde{z}\tilde{u}} \sigma_{\tilde{z}} \sigma_{\tilde{u}})^2 \right] + 2 \left[\exp(2\sigma_{\tilde{z}}^2) - 1 \right] (\delta + \sigma_{\tilde{u}}^2)$$

Note that only the first term depends on the correlation between z and u .

Case 4: We wish to calculate $\langle (e^{2\tilde{z}} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle$:

$$\begin{aligned} \langle (e^{2z} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle &= \left[\exp \left\{ \delta + \frac{(2\sigma_z)^2 + (\sigma_{\tilde{u}})^2}{2} + \rho_{z\tilde{u}}(2\sigma_z)(\sigma_{\tilde{u}}) \right\} - 1 \right] - \\ &\quad - \left[\exp \left(\frac{(2\sigma_z)^2}{2} \right) - 1 \right] - \left[\exp \left(\delta + \frac{(\sigma_{\tilde{u}})^2}{2} \right) - 1 \right] \end{aligned}$$

This may be rewritten in the following form:

$$\langle (e^{2z} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle = \exp(2\sigma_z^2) \left[\exp \left\{ \delta + \frac{\sigma_{\tilde{u}}^2}{2} + 2\rho_{z\tilde{u}}\sigma_z\sigma_{\tilde{u}} \right\} - 1 \right] - \left[\exp \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right) - 1 \right]$$

When we keep terms of first order in the mean value and variance of u , and to all orders in the variance of z , we arrive at the following approximate result:

$$\langle (e^{2z} - 1)(e^{(\delta + \tilde{u})} - 1) \rangle \approx 2\exp(2\sigma_z^2) \left[\rho_{z\tilde{u}}\sigma_z\sigma_{\tilde{u}} + (\rho_{z\tilde{u}}\sigma_z\sigma_{\tilde{u}})^2 \right] + \left[\exp(2\sigma_z^2) - 1 \right] \left(\delta + \frac{\sigma_{\tilde{u}}^2}{2} \right)$$

Note that only the first term depends on the correlation between z and u .

References

- Jean-Philippe Bouchard & Marc Potters, Theory of Financial Risks, 2nd ed.
Cambridge University Press, Cambridge, UK (2000)
- 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)
- Wayne A. Fuller, Introduction to Statistical Time Series
John Wiley & Sons, New York (1996)
- Ramazan Gençay, Faruk Selçuk, & Brandon Whitcher (GSW),
An Introduction to Wavelets and Other Filtering Methods in Finance and Economics,
Academic Press, San Diego, CA (2002)
- Saeed Ghahramani, Fundamentals of Probability
Prentice Hall, Upper Saddle River, NJ (1996)
- William H. Greene, Econometric Analysis, 5th ed.
Prentice Hall, Upper Saddle River, N.J., (2003)
- Rosario N. Mantegna & H. Eugene Stanley,
An Introduction to Econophysics, Correlations and Complexity in Finance
Cambridge University Press, Cambridge, UK (2000)
- Chrysostomos L. Nikias & Athina P. Petropulu,

Higher-Order Spectra Analysis, A Nonlinear Signal Processing Framework
PTR Prentice Hall, Upper Saddle River, NJ (1993)

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

William F. Sharpe, Gordon J. Alexander, & Jeffery V. Bailey, Investments, 5th ed.
Prentice Hall, Englewood Cliffs, NJ (1995)