Econometric Analysis in *QuanTek*

(Revised October 7, 2017)

Robert Murray, Ph.D.
Omicron Research Institute

2.0 Contents

2.1 Financial Econometrics and QuanTek

The ***QuanTek*** Econometric program is designed to utilize the most state-of-the-art ideas from the rapidly evolving field of **Financial Econometrics** to develop profitable trading rules that have the greatest chance of providing ***maximum returns*** with ***minimum risk***. The central feature is the use of a **Wavelet Adaptive Filter** to provide an improved estimate of future returns, and also an estimate of future volatility that is interpreted as risk. The estimated future **returns** from the **Wavelet Adaptive Filter** are utilized in a short, medium, or long-term trading strategy, but the trading **risk** is controlled by a trading strategy incorporated within a diversified **Optimal Portfolio**.

The prediction of future returns of financial assets is tricky business, of course, and according to the **Efficient Market Hypothesis**, it cannot be done at all. But some estimate of future returns must be made, otherwise how does an investor know what investments to make? An excellent book on the **Kalman** filter, with applications to Finance, is Forecasting, structural time series models, and the Kalman filter (1989), by Andrew C. Harvey [H]. (The Adaptive filter incorporated in ***QuanTek*** is similar to the **Kalman** filter.) Here is a quote from the preface to this book (p.xi), which illustrates the situation perfectly:

> "The inclusion of 'forecasting' in the title of the book is perhaps a little rash. It is always very difficult to predict the future on the basis of the past. Indeed it has been likened to driving a car blindfolded while following directions given by a person looking out of the back window. Nevertheless, if this is the best we can do, it is important that it should be done properly, with an appreciation of the potential errors involved. In this way it should at least be possible to negotiate straight stretches of road without a major disaster. Too many forecasting procedures seem to attribute the person in the back seat with supernatural powers when in fact his behaviour is more consistent with that of someone who is mildly inebriated."

With this in mind, it becomes very important to be able to discern the limits of our ability to forecast expected returns, and to develop methods for separating the small 'signal' present in the data from the much larger stochastic noise. In particular, it is very easy to confuse the stochastic noise with a 'signal', and end up developing trading rules that merely follow the random patterns of the stochastic noise. The methods of ***Technical Analysis*** are particularly vulnerable to this malady, and these 'intuitive' methods end up in the similar situation to someone looking up at the sky and perceiving faces and other patterns in the clouds that are not, in fact, there at all. This fallacy is due to the nature of the human mind to perceive patterns in random data, and one must therefore be careful to try to use rigorous methods of Statistics to discern the patterns that are really there from the apparent patterns that are due merely to stochastic noise.

Signal Processing & Linear Prediction

What is the problem that the ***QuanTek*** program attempts to solve? There are many different trading strategies and methods of estimating future price behavior of stocks and other securities. Foremost among these is the study of ***Technical Analysis***. Our interpretation of Technical Analysis, from a Signal Processing point of view, is that the financial data are mostly random noise, with a small "signal" component buried in the noise. Then Technical Analysis may be viewed as an attempt to intuitively "estimate" this small signal buried in the noise, and use this to devise profitable trading rules. One of the most basic technical indicators is the prevailing **trend**, over some time period. If **trend persistence** holds, then over some time scale you might expect the most recent trend over this time scale to persist into the future over a comparable time scale. The problem is then how to separate the data into time scales over which trend persistence might be expected to hold.

One way of viewing the Wavelet technique is that it separates the financial data into octaves corresponding to different time scales, and isolates the time variation of the returns on each time scale. Then on each time scale the current (averaged) returns could be expected to persist into the future for a time approximately the same as that time scale. So, one way to view the **Wavelet Adaptive Filter** is that it predicts the trend on each time scale to persist a corresponding length of time into the future. This is in contrast to traditional Fourier techniques, which separate the financial data into pure sine waves of different frequencies. Then each sine wave is expected to persist indefinitely into the future. This is what you would expect if the time series were **stationary** (or better yet, **deterministic**), as for example a periodic signal in an electronic circuit. But this is not really the appropriate picture for financial time series, and the Wavelet approach is much more realistic. In the Wavelet approach, the high-frequency part of the data that lies in the distant past is eliminated from the signal estimation, and only the most recent data on each time scale is used for the future estimation of the signal. This then eliminates about half the degrees of freedom of the financial data as irrelevant noise, thereby improving the chance of finding a true signal that may be present, buried deep in the stochastic noise.

On the other hand, the traditional **Efficient Market Hypothesis** and **Random Walk** model [PB] stipulate that there is no "signal" in the financial data, and future price action cannot be predicted based on past price action. We take the point of view that the markets are, indeed, very efficient, but not 100% efficient. There are small inefficiencies that can be uncovered and exploited to make a profit, because traders and investors are not completely rational, not completely well informed, and do not react instantaneously as the Efficient Market Hypothesis supposes. So it *should* be possible for a knowledgeable trader or investor with a sound trading strategy to "beat the market". On the other hand, it is well known by now that the Random Walk model is only a crude approximation to market behavior. It supposes that price returns are independent random variables that satisfy a Gaussian distribution. But it is clear that the actual distribution is not Gaussian and exhibits "fat tails", such as the stable Paretian or Levy distributions [M, PB]. These fat-tailed distributions give rise to **risk** that is far greater than the Gaussian distribution would predict, especially when it comes to market "bubbles" and the subsequent "crashes" which inevitably result. So the upside and downside risk needs to be accurately estimated, taking into account the fact that the distribution of returns is non-Gaussian. This then involves the **GARCH** models of time-varying volatility [G]. The ***QuanTek*** program will implement (in a future version) this strategy using Wavelet techniques to estimate both the future time-varying returns and volatility, based on the fact that these quantities are not exactly independent random variables, but may exhibit correlation with the past.

Wavelet Analysis of Time Series

One of the main features of the ***QuanTek*** program is the use of Wavelet analysis [PW] of the financial time series data. This is an alternative to analyzing the time series directly in the time domain as in a conventional regression approach, on the one hand, and using the Fourier analysis or frequency domain, on the other hand. Here, the time domain means dividing the whole time series up into its components with respect to time, whereas the frequency domain means taking linear sums of the time series and decomposing it into signals with a definite frequency.

In the conventional **Fourier** analysis, the time series data are separated into pure sine waves of a single frequency, extending over all time. Usually the **Fast Fourier Transform** (FFT) is used, which requires that the data satisfy circular boundary conditions. So the data must be a fixed length, in our case 2048 units (8 years) long, and the data are arranged to form a circular data set (with zero-padding). Then in the Fourier analysis the circular data are decomposed into component frequency signals that extend over the whole circle. This decomposition is ideal for periodic signals such as waveforms in an electronic circuit, but not so good for financial data. For one thing, the Fourier transform assumes the time series to obey **stationary statistics**, whereas we expect financial time series to obey **non-stationary statistics**. The non-stationarity means that the statistical properties such as the mean and covariance matrix vary with time. So due to the "infinite" time extent of the component waves in the Fourier analysis, it is not well adapted to non-stationary time series such as financial time series.

It turns out that the **Wavelet** analysis is much better suited for many types of realistic signals than the **Fourier** analysis. In the Wavelet analysis the data are also circular with a fixed length of 2048 units, like the FFT. But Wavelet analysis is halfway between the time domain of the signal as a function of time, and the frequency (Fourier) domain of the signal as a function of frequency. The basis functions of the Wavelet analysis are waves that extend over a finite range of both time and frequency. By comparison, the basis functions of the frequency domain (Fourier) are pure sine waves of a single frequency which extend over all time, while the basis functions of the time domain are the signals at any specific time, but which extend over all frequencies as a result. The Wavelet basis functions are divided into octaves of frequency, and the high frequency octave functions extend over short time intervals, while the low frequency octave functions extend over long time intervals. This corresponds to what we expect in financial data – that the low frequency fluctuations have an influence over long time periods, while the high frequency fluctuations have an influence only over short time periods. Thus using the Wavelet analysis enables an immediate improvement in the signal to noise ratio, in that the short-term fluctuations occurring in the distant past can be eliminated right away, since they are unlikely to have any influence on the present or future. Working in the Wavelet basis enables half of the stochastic noise to be eliminated and a drastic reduction in the number of degrees of freedom that have to be fitted in the regression, reducing (to some extent) the problem of "overfitting".

In the Wavelet regression model, therefore, each wavelet level, corresponding to an octave of frequencies, is approximately independent and is fitted independently. Then the low-frequency octaves make a long-term prediction that extends for many time units into the future, while the high-frequency octaves make only a short-term prediction that extends only a few time units into the future. Thus, in effect, for the returns the long-term trends are extended into the future a number of days corresponding to the length of the trend, while the shorter-term trends are extended into the future a correspondingly shorter number of days. The extension of the trends on each wavelet level is also dependent on the time-dependent correlation in the data, which in turn weights the trends on each wavelet level differently. In fact, if the trend on each wavelet level is extended with equal weight, the future projected return adds up to zero, corresponding to no correlation. The correlation enters due to the difference of the weights of the projected return on the different wavelet levels. These different weights are determined by the regression procedure, which is greatly simplified because the regression is performed only on the different wavelet levels, not on each past data day separately as in the usual time-domain approach. So, this separation of the data into wavelets results in a vast simplification of the regression problem, and helps separate the important degrees of freedom in the financial time series from the rest, which make up the "stochastic noise".

Predictors of Future Returns

In the usual approach in **Econometrics**, a stochastic model is postulated and then this model is "estimated" or fit to the past time series data, and a statistical analysis of the fit is performed. However, in the ***QuanTek*** program we use a more pragmatic method, which can also be computed within a reasonable time. This method is to search for a set of **predictors**, which are functions of the past data with various filterings applied. Each predictor is tested for **correlation with future returns**. If a meaningful correlation is found, then this predictor is used as part of a linear regression in the **Wavelet Adaptive Filter**. Using a small number of predictors, which have been tested and verified to have meaningful correlation with future returns, greatly simplifies the regression problem as opposed to regressing on *all* the past data.

By default, a set of three predictors is used as a starting point, which corresponds to different types of smoothings or filterings of the past data of the same security. These we call the **Relative Price**, **Velocity**, and **Acceleration**. The smoothing has the effect of eliminating the high-frequency stochastic noise so the correlation of the lower-frequency components is revealed more clearly. Actually, in the Wavelet approach, on each wavelet level these three predictors are similar, only shifted from each other by a time lag corresponding to a quarter-cycle. So, on each wavelet level when the three predictors are fitted to future returns, the phase lag can be fitted for optimum correlation. This is done using a recursion algorithm on the past data to "train" the **Wavelet Adaptive Filter** for an optimum estimation of future returns. Once this fit is made, the three predictors together can be used to define and display the optimal **buy/sell points**, in which the three predictors correspond to the three basic **trading rules**: **Buy Low – Sell High (Relative Price)**, **Trend Persistence (Velocity)**, and **Turning Points (Acceleration)**.

The **Relative Price** is just a smoothing of the logarithmic price itself, relative to a long-term smoothing. For example, we can take the *M*-day Wavelet smoothing of the log prices minus, say, the 128-day smoothing as a reference. We expect an anti-correlation between the Relative Price (at some point in the past) and the future returns due to the **buy low – sell high** mechanism. In other words, if the *M*-day Wavelet smoothing yields frequencies, or "cycles" with period of *M* days, then we expect a low point a quarter cycle in the past to be correlated with positive returns at the present, and a high point a quarter cycle in the past to be correlated with negative returns at the present. This correlation between the *smoothed* Relative Prices and *N*-day future returns is revealed in the ***QuanTek*** **Correlation Test**. We expect the Relative Price predictor to be the best one because the price is the sum over returns, leading to noise reduction, and on top of that the Relative Price is smoothed.

The **Velocity** indicator is just the smoothed returns themselves. We expect a correlation between past returns and future returns at the immediate present on a smoothing scale of *M* days due to the mechanism of **trend persistence** or **buy high – sell higher**. In this case it helps in defining trial functions and measuring correlation to project the predictor into the future using, perhaps, Standard Linear Prediction (SLP). This is perfectly acceptable, with regard to measuring correlation between past and future returns, so long as only *past* data are used in the projection, so that the future projection is actually a function of *past* data. This predictor also shows respectable correlation in the Correlation Test, depending on the method used to make the future projection.

Finally, we have the **Acceleration** predictor, which is the smoothed first-difference of the returns. This predictor should have the opposite correlation from the Relative Price, and may be thought of as marking the **turning points**, or minima-maxima of the prices. This indicator is less effective than the other two, since it emphasizes the higher frequencies more than the lower frequencies, but the Correlation Test sometimes shows a correlation here as well.

In a more sophisticated approach, using nonlinear predictors, we might for example suppose that the past price action has a *memory*. This is the same idea used in Technical Analysis to determine, for example, support and resistance levels. A *potential* function could be defined to incorporate this memory, with a damping factor to represent the fading of the memory. The potential function could measure the number of trades, for example, at different price levels, and define a "potential barrier" at the levels where many trades took place in the past, to represent the **support** and **resistance** levels. This is an idea for the future, but it represents the method by which predictors can be defined to represent the nonlinear behavior of the price action. Another nonlinear predictor that could be defined is a potential centered on a definite level of *return*, which becomes deeper the longer that return is maintained. This would then represent what is known in Technical Analysis as an **uptrend** or **downtrend**. But the simple *linear* predictors described above should give a reasonable account of market dynamics, to first approximation, ignoring the nonlinear effects.

In a future version of ***QuanTek***, I plan to generalize the predictor approach to a regression of each security over predictors defined in terms of the whole portfolio of securities and in addition various market indexes and fundamental or economic data (such as interest rates). Generally, a simplified approach will be adequate, such as an approach similar to the **CAPM** (Capital Asset Pricing Model) [SAB] where only the portfolio as a whole is used as a predictor or **factor**, and its correlation with each individual security measured and used as input to the price projection. Of course a market index should also be included in a CAPM-like model. Or a **Factor Model** can be used where several such factors are identified and used as predictors. This strategy can also be implemented in the form of a **Kalman** filter. The only limitation to how far we can go with this is, of course, the computational power of the computer, and also the availability of the data.

Optimal Trading Strategies

The question arises: "What is the optimal trading strategy?" The thing many people may not realize is that any trading strategy relies on the presumption of some kind of **correlation between past price action and future returns**. On the one hand, short-term traders may think that the shorter the time scale that they can trade, the more of the up-down swings or volatility they can utilize to make a profit. However, this is only true if there is a **correlation** between the technical indicators or trading rules they are using, and future price swings. In fact, if the **Random Walk** model is true, then there is no such correlation, and the **expected return** is always zero no matter what kind of trading is done. However, the more rapid the trading is, the more the **risk** is increased, resulting in a wider **variance** in the eventual outcome. In addition, the transaction costs are much greater. At the other extreme, a buy-and-hold investor ignores the short-term swings in the market, and only invests to take advantage of a long-term up-trend in prices that they believe to exist. This corresponds to an assumption that a **Random Walk model with drift** applies, so the returns have a non-zero **mean** value, although the de-trended returns will have zero correlation. Thus the buy-and-hold investor is assuming a particular **stochastic model** for the price action, which it turns out, is a gross oversimplification of the true situation and is not really correct. This assumption of the existence of a long-term trend seems dangerous, because of the clear existence of market crashes and long periods of flat markets. Also it implicitly relies on the presumption of a trending Random Walk, which is a **stationary** time series model, but it is clear that the actual financial time series are much better explained in terms of a **non-stationary** stochastic model.

Our point of view, then, is that the optimal trading strategy must start with a determination of the potential **correlation** in the returns data. This appears to vary from one security to the next, so the correlation must be measured between a set of chosen **predictors** and the future returns, for each security, using a **Wavelet Adaptive Filter**. The **Wavelet Adaptive Filter** then predicts the future returns and volatility for each time scale, and trades are made based on this predicted price action. The trader may choose the time scale on which to measure the correlation and set up trading rules. But this must take place within the context of an **optimal portfolio** to **control risk**. In standard portfolio construction the portfolio is *rebalanced* at periodic intervals, say every month, to maintain the computed optimal mixture of securities as the prices change [FFK]. But instead of rebalancing at fixed time intervals, our **active trading** method can be thought of as a rebalancing for each security as predicted trading opportunities arise in that security, for any chosen time scale N. In this way it is hoped that the **maximum returns** may be realized by short-term trading, while at the same time incurring **minimum risk** by doing it within an **optimal portfolio**. Also if the market appears overextended and vulnerable to a crash, our **active trading strategy** gives the best hope of getting out of the market with minimum loss. This is achieved by gradually decreasing positions in those securities that become “overextended”. Clearly this is a complicated problem, which is why a sophisticated approach such as utilized by the ***QuanTek*** program is the only adequate way to address the problem.

2.2 Efficient Market Hypothesis and Trading Strategies

According to the **Efficient Market Hypothesis** [FFK], the market is 100% efficient and it is impossible to predict future price action on the basis of any past information. Most people agree that this hypothesis is too extreme, and that small inefficiencies do exist which can enable the knowledgeable trader or investor to make a profit. Indeed, this is the case with all economic activity – it is frequently true that assets or property are mispriced, and those individuals who have a deep understanding of the true value of these assets or properties can make a profit by buying or selling them. On the other hand, with financial securities the price series are well approximated by the **Random Walk** model, and it is very difficult to forecast the future price action based on the past price series alone. At least there are no clearly measurable correlations in the price data that can be used for a forecast – the returns series appears to be random white noise in almost all cases. However, upon closer inspection it does appear that small correlations in the data do exist, although these are buried in the "white noise" and are also no doubt time dependent. The question, then, is how to utilize these small, time-dependent correlations, assuming they exist, to formulate a profitable set of trading rules.

I like to think of stock trading as similar to gambling. If you play a game of Craps in a casino with fair dice, then over the long run (Law of Large Numbers) you will slowly lose, because the odds of winning or losing are even except for a small "percentage" that goes to the house. However, suppose the dice are slightly "loaded". Then, even though the advantage might not be noticeable over short time periods, over long time periods you can win big money because the odds are now slightly in your favor. The trading of financial securities is the same way. If you can find a set of trading rules that puts the odds of winning slightly in your favor, then over the long run you can make large profits, at least potentially. This is the goal of **Technical Analysis** – to find the profitable trading rules, buried deep in the stochastic noise, which can lead to consistent trading profits when the stochastic noise is averaged out over the long run. This may also be viewed from a **Signal Processing** point of view. There is a small "signal" with significant correlation with future returns, buried in the stochastic noise, and the problem is how to isolate or "estimate" this small signal buried in the noise, and use it to formulate a profitable trading strategy. Then Technical Analysis may be viewed as a "heuristic" attempt to estimate this small correlated signal buried in the stochastic noise of the Random Walk model. But perhaps now that we have powerful desktop computers at our disposal, a more scientific approach to financial **Signal Processing** is called for.

Autocorrelation & Power Spectrum

Then, given that the financial returns closely follow a Random Walk, except perhaps for a small "signal" buried in the noise, how is a trading or investment portfolio to be constructed? The choice of portfolio requires an estimate of both **future returns** and **future risk**. How are these to be estimated? One idea is to estimate the future returns or trend as equal to some long-term average of past returns, utilizing the concept of **trend persistence**. This implies some kind of correlation between past returns and future returns on some time scale. Now here is a mathematical fact from the theory of *stationary* time series: The autocovariance series is the Fourier transform of the power spectrum. This is called the Wiener-Khinchin theorem [NR, p.498]. So if the Fourier transform of the time series is taken and the Fourier components (corresponding to each "frequency") are squared to give the power spectrum, this is the Fourier transform of the autocovariance series. (The autocovariance is the autocorrelation of the returns times the variance.) So, the result is that if the power spectrum is "flat", corresponding to a "white noise" spectrum, then there are no correlations in the returns series, and the Random Walk model applies. Conversely, for there to be correlations in the returns series, there must be a non-constant power spectrum, so some of the "frequencies" have more "power" than others. So the search for correlations and profitable trading rules is equivalent to the search for those time scales or "frequencies" that have more "power" than the average. This can be extended to the *non-stationary* case if we suppose that over short time periods the stationary case holds approximately, and over longer time periods the power spectrum and autocovariance change with time. Then the **trend persistence** results from identifying those frequencies or time scales with a higher power spectrum than average, at any given time, which then results in a positive correlation between the past trend and future trend *on those time scales at any given time*.

There is a caveat to this. If we measure the power spectrum directly, in the form of the **Periodogram Spectrum** for a stationary time series or its **Wavelet Spectrum** counterpart for a non-stationary time series, then this power spectrum itself is dominated by stochastic noise. Smoothing techniques can reduce this noise, but in general a direct measurement of the power spectrum from the returns data is not an effective method of detecting the underlying correlation.

So instead of this, the parameters corresponding to the power spectrum are fit directly to the future returns data in a recursive algorithm which uses the past data in a **training** period, which adapts the **Wavelet Adaptive Filter** to make the optimal prediction of future returns. This then may be interpreted as a method of determining the "true" wavelet variance or power spectrum, and hence the correlation, more accurately than a direct measurement would yield.

Trading Time Scales

The Wiener-Khinchin theorem may be visualized in terms of the **Efficient Market Hypothesis** by making the idealization that different traders or investors trade on *cycles* of varying time scales. The trading activity on various time scales is measured by the power spectrum of the returns time series on those time scales. The market is efficient when the reaction to new information is instantaneous and future returns are uncorrelated with past returns. This corresponds to equal trading activity on all time scales and a constant or "white noise" power spectrum. Because, giving an arbitrage argument, if any cycles were to predominate over the others, traders would take advantage of these apparent cycles to make a profit, and then the trading opportunity would vanish. So if the market is perfectly efficient, the trading activity on all time scales has equal intensity, the power spectrum is perfectly flat, and the returns consist of random "white noise".

However, in reality this condition may be violated, especially if the activity is measured over localized time intervals rather than taking an average. In the past, it used to be said that there were strong correlations in the returns over time intervals of tick data shorter than about 15 minutes. This was no doubt due to most traders being without real-time quotes, so the trading activity on these very short time scales was below the average. But in recent years the availability of real-time data has been more prevalent, so these very short-term correlations are no doubt largely eliminated by now. But some traders still tend to be "behind the curve" with regard to new information, so they do not react immediately, and this should still result in a deficit of trading activity on the shorter time scales. This will result in positive correlations of returns or **trend persistence**. However, it may also happen that traders over-react and trade too strongly on shorter time scales in response to news, and in particular in response to the perception of a falling market. This will lead to a surplus of trading activity on the shorter time scales, which results in negative correlations of returns or **trend anti-persistence**. The price moves tend to be too large in this case and then must correct back to their "true" values, as traders correct for their initial over-reaction. This correction mechanism for over-reaction to past prices and anti-persistence then leads to the **return-to-the-mean** mechanism in the price series. (There can be more than one return-to-the-mean mechanism at work at the same time, on different time scales. Then, for example, over a certain time scale the prices may be reverting to a strongly up-trending mean price, then suddenly switch and revert to a longer-term trend, resulting in a sudden crash back to this longer-term mean price. This, however, undoubtedly involves *non-linear or higher-order correlations* in the data, such as perhaps a simultaneous correlation with the entire market. Another name for this is **critical phenomenon**, which is an active field of research in **EconoPhysics** [PB] – a market crash can be thought of as a "phase transition" in the market dynamics.) In any case, the response of the market to changing conditions has not been completely **efficient**, correlations are introduced between past and future returns, and this results in further trading opportunities by the arbitrage argument. This correction mechanism also implies that any short-term trading imbalances are then corrected by the market over the longer term, so although the power spectrum imbalances and the resulting correlations may exist over the short-term, taking a long-term average of the power spectrum leads to almost no visible correlation in the returns data. In particular, in the *stationary* approximation it is difficult to see correlation in the returns series – these short-term trading imbalances are *non-stationary* phenomena.

So it is due to this property of the autocovariance sequence and the *non-stationarity* of the time series that it doesn't make much sense to presuppose the existence of *permanent* correlations over particular time scales, such as very long time scales. The very long-term trend is just one particular range of frequencies or time scales, and there is no more reason to assume the existence of a positive trend on these time scales than on any other. However, it is believed that in the generic case for financial time series (as well as other time series occurring in Nature) long-range positive autocorrelation does exist in the data, due to the occurrence of **fractal statistics**, so the series is of the type known as a **Long Memory** or **Fractionally Differenced (FD)** process, as advocated by Mandelbrot [M]. The Wavelet techniques are particularly well adapted to the analysis of such processes [PW]. Such an FD stationary process may indeed be a good approximation for the actual non-stationary process of financial time series.

Law of Large Numbers

Rather than *assume* the existence of a certain type of correlation in the price data, a more plausible approach would be to *measure* correlations of various types, on any chosen time scale, and then base trading rules on these measured correlations, with the understanding that they are constantly changing with time as the market dynamics changes. This makes the statistical problem more difficult, however, as there is now no longer a "Law of Large Numbers" as in the case of stationary time series.

In fact, it should be noted that, when modeling financial time series as stochastic time series, there is the problem that there is only *one instance of each time series*. If the stochastic time evolution of the price of a security could be "played" over and over, and the results averaged, then the Law of Large Numbers could be applied to uncover the true "signal" in the stochastic noise, and the true statistical properties of the time series could be measured. As it happens, there is only one instance of each time series to work with, and since the statistics are non-stationary, there is no Law of Large Numbers for any individual time series.

Thus in the case of financial time series, the *ergodic hypothesis* [Hay] must be used. This means that, instead of applying the Law of Large Numbers to a large ensemble of "instances" of the financial time series, all with identical statistical properties (which of course does not exist), the averaging is instead applied to the single time series over the time index of the time series. This necessarily means that the statistics must be in some sense *stationary*. However, we do not have to assume that the statistics are stationary in the narrow sense of constant mean and covariance matrix over the whole time series. We may instead look for more complicated, nonlinear or time dependent, functions of the past data, or *predictors*, and test these for correlation with future returns. However, in order to perform the test, the assumption must be made that these nonlinear correlations are stationary, at least over the duration of the correlation test, which will usually take up most or all of the given time series. So the generalized notion of stationarity is that the *formula* for computing these complex *predictors* is constant, at least over the duration of the given time series. (The formula itself could conceivably change from one block of data to the next, for example if it uses some input parameter such as interest rates that change slowly with time.)

Actually the complex formula that remains constant with time can be identified with the trading strategy itself or components of it. Then this trading strategy will necessarily involve the whole portfolio of securities, not just a single security. The only way to apply the Law of Large Numbers and try to judge the effectiveness of a trading strategy is to do it within a diversified **Optimal Portfolio**. Then, with a large enough portfolio, the stochastic noise can be averaged out and the growth (or otherwise) of the whole portfolio can be measured and trading strategies compared. This is why it is crucial to devise the trading rules within the context of the whole portfolio, and not just individual securities – otherwise there is no real way to measure how well it works, and also no way to control **risk**.

Overall Investment Strategy

Of course, along with **Technical Analysis**, there is also **Fundamental Analysis**. In order to achieve superior returns and "beat the market", it is necessary to make use of *all* available information. This includes studying each company or security that you want to invest in, to determine whether it will be a favorable investment over the long-term. For stocks the **earnings** are very important, as well as the price in relation to the earnings or P/E ratio. A good understanding of the macroeconomic factors and how they affect each market sector is also crucial. This information is separate from the information about the price action. In fact the price action can be viewed as a *response* to these fundamental and macroeconomic factors, with, however, a time delay and possible over-response due to the fact that traders and investors are not completely well-informed and not completely rational in their decisions. In the absence of this information, the approach of **Technical Analysis** is to try to take this information into account by means of its effect on the price action itself. Then you can try to spot the initial effect of the new information on the price action, and then try to respond quickly before others are able to respond. However, this always puts you "behind the curve" to some extent and it is much better to be able to take into account both the **Fundamental** and the **Technical** information together. One way to do this would be to construct a long-term investment portfolio based on the known Fundamental and macroeconomic data. Then doing short-term trading in the portfolio based on Technical signals and data can enhance the investment returns. Another way of looking at this is that any long-term investment portfolio must periodically be *rebalanced* [FFK], and the short-term trading based on Technical signals can be viewed as a way of rebalancing the portfolio. Another consideration is that if something goes wrong and the individual security or the whole market has a sudden decline or "crash", then it is important to act quickly and get out of the position. If this can be done then it can result in avoiding potentially large losses in such a situation, which would result if the "Buy and Hold" strategy were followed.

2.3 Financial Returns as a Stochastic Process

According to the **Random Walk Model**, as first proposed by Bachelier in his Ph.D. dissertation in 1900 [B], financial returns (price differences) may be described by a *stochastic process* which is a simple Random Walk. This stochastic process is also known as **Brownian motion**. More precisely, this has been refined to the statement that the *logarithmic* returns, or the differences of the *logarithms* of the prices (logarithm of the price *ratios*), follow a Random Walk process that is called **Geometric Brownian motion**. Unless the price ratios are large, these two processes are nearly the same. This process may be further generalized by considering the **Random Walk with drift**. This is the model many long-term buy-and-hold investors have in mind, when they consider averaging out the short-term random price fluctuations in favor of reaping a return on the long-term trend (deterministic drift). (Unfortunately, it does not seem to be that simple in reality.)

At any rate, if the returns follow a Random Walk, then they are serially uncorrelated. In fact, the probability distribution of returns is also usually assumed to be Gaussian, in which case the (daily, let's say) returns are *independent* random variables. Thus they appear as random white noise, with a flat power spectrum. In fact, this is how they actually appear (in most cases) when viewed using the **Periodogram** display or the graph of the raw returns in ***QuanTek***. Thus the (naïve) search for correlation between the N-day future returns and daily past returns is likely to be fruitless. However, in **Technical Analysis** it is believed that there do exist past price patterns that are correlated with future returns. But these correlations are buried in the stochastic noise of the daily returns, so the problem becomes one of *signal extraction* or *noise reduction* to ferret out the underlying "signal" buried in the noise.

Noise Reduction and Smoothing

In order to extract the *signal* buried in the stochastic noise, the main technique that is used in ***QuanTek*** is **smoothing** of the past data. This has the effect of eliminating the high-frequency stochastic noise and exposing the correlation of the lower-frequency *signal* with the future returns. To illustrate this, suppose that a signal consists of a small correlation between the daily past returns and daily future returns, which persists over a time scale of M days of past returns and N days of future returns. This small correlation has the value ε, where ε is assumed to be very much smaller than unity:

$$\operatorname{cor}\left(\xi_{-m}, \xi_{+n}\right) \equiv \left\langle \xi_{-m}, \xi_{+n} \right\rangle \Big/ \left\langle \xi_{-m}^2 \right\rangle^{1/2} \left\langle \xi_{+n}^2 \right\rangle^{1/2} \simeq \varepsilon \ll 1$$
$$\text{where} \quad \xi_{-m} \equiv \text{past return}, \quad \xi_{+n} \equiv \text{future return}$$

The correlation between the two variables is, by definition, the covariance between them divided by the standard deviation (square root of the variance) of each variable. So this small correlation ε between each past daily return and future daily return is too small to measure compared to the variance of the daily returns. Now, to illustrate the principle of *noise reduction* via **smoothing**, let us consider a simple situation where the correlation is constant over a range of M values of the past returns and N values of the future returns, and the variance of all the daily returns is assumed to be constant, given by $\left\langle \xi^2 \right\rangle$. Now, we can define *smoothed* past and future returns, over a time scale M and N, respectively, by:

$$\bar{\xi}_- \equiv \sum_{m=1}^{+M} \xi_{-m+1} \qquad \bar{\xi}_+ \equiv \sum_{n=1}^{+N} \xi_{+n}$$

Now, it is a property of independent random variables that the variance of a sum of such variables is given by the sum of the variances. Hence we can write (ignoring the small correlation) for the variance of each of the above smoothed variables:

$$\operatorname{var}\left(\bar{\xi}_-\right) = \operatorname{var}\left(\sum_{m=1}^{+M} \xi_{-m+1}\right) = M \operatorname{var}\left(\xi\right) \equiv M \left\langle \xi^2 \right\rangle$$
$$\operatorname{var}\left(\bar{\xi}_+\right) = \operatorname{var}\left(\sum_{n=1}^{+N} \xi_{+n}\right) = N \operatorname{var}\left(\xi\right) \equiv N \left\langle \xi^2 \right\rangle$$

Now the correlation of the smoothed variables is (approximately) given by:

$$\begin{aligned}
\operatorname{cor}\left(\bar{\xi}_-, \bar{\xi}_+\right) &\equiv \left\langle \sum_{m=1}^{+M} \xi_{-m+1}, \sum_{n=1}^{+N} \xi_{+n} \right\rangle \Big/ \left(M \left\langle \xi^2 \right\rangle\right)^{1/2} \left(N \left\langle \xi^2 \right\rangle\right)^{1/2} \\
&= \sum_{m=1}^{+M} \sum_{n=1}^{+N} \left\langle \xi_{-m+1}, \xi_{+n} \right\rangle \Big/ M^{1/2} N^{1/2} \left\langle \xi^2 \right\rangle \\
&= M^{-1/2} N^{-1/2} \sum_{m=1}^{+M} \sum_{n=1}^{+N} \left[\left\langle \xi_{-m+1}, \xi_{+n} \right\rangle \Big/ \left\langle \xi^2 \right\rangle \right] \\
&\equiv M^{-1/2} N^{-1/2} \sum_{m=1}^{+M} \sum_{n=1}^{+N} \left[\varepsilon\right] = M^{-1/2} N^{-1/2} \left(MN\right) \varepsilon \\
&= M^{+1/2} N^{+1/2} \varepsilon
\end{aligned}$$

Thus it can be seen that the correlation between the smoothed *M*-day past returns and the *N*-day future return has been increased by a factor of $M^{+\frac{1}{2}}N^{+\frac{1}{2}}$ compared to the correlation between the individual past and future returns. So the smoothing has *amplified* the correlation and reduced the level of stochastic noise. We call the sum of the past returns a **predictor** because it is a function of the past returns data. Note that the correlation cannot be greater than unity, so the product $M^{+\frac{1}{2}}N^{+\frac{1}{2}}\varepsilon$ must be less than unity. Otherwise, it would be impossible for the small correlation ε to persist over the interval of *M* past returns and *N* future returns, as was assumed in the above simple calculation. So the smoothing is a means of separating out the *low-frequency correlation* from the *high-frequency noise*.

In general, a **predictor** can be *any* function of the past price data or any other economic data. The only condition is that all the data that goes into the predictor be in the *past* of the *future* returns that we are trying to predict. Then the projection of the future returns can be computed by doing a **linear regression** on all the predictors that are chosen. This regression has a much greater chance of success than a simple regression over the raw daily returns, because as shown above the correlation between each predictor and the future returns is *amplified* by means of the smoothing. Even though a linear regression is used, the model can actually be highly nonlinear because any *nonlinear* function of the past prices or other data can be used for the predictor. So, for example, we can look for a correlation between the past volatility (smoothed variance) and the future returns. In fact it has been shown that the variance is much more highly predictable than the returns themselves, and this fact has been utilized in recent years in the **GARCH** (Generalized AutoRegressive Conditional Heteroskedasticity) models. (Heteroskedasticity is another word for time-dependent variance – the future variance is conditional on the past variance and the dependence is modeled as an autoregression.)

Financial Risk & GARCH

To devise a profitable trading strategy and at the same time control risk is a difficult mathematical problem and is still not completely solved within the field of Econometrics. An excellent recent text explaining the latest research in this area is Financial Modeling of the Equity Market: From CAPM to Cointegration (2006), by Frank J. Fabozzi, Sergio M. Focardi, & Petter N. Kolm [FFK]. In particular, the nature of financial risk is still not completely understood, and according to Benoit Mandelbrot [M], the financial markets are far riskier than is taken into account in the standard Modern Portfolio Theory of the past decades. The ***QuanTek*** program uses Wavelet techniques [PW] to try to estimate both the optimal trading strategy and the future downside risk (in a future version), both of which are time varying, so the Wavelet approach takes into account the **non-stationary** nature of financial time series. This is also related to the **GARCH** models of time-varying volatility [G], which are currently undergoing rapid development in the Econometric literature.

Stochastic vs. Deterministic Process

It is widely believed that financial returns follow a stochastic process – indeed, this was Bachelier's hypothesis in his Ph.D. dissertation of 1900 [B] and was followed up with the development of **Modern Portfolio Theory** in the 1950's. But what does this really mean? It is important to clarify the concept of Stochastic (or Random) process vs. a Deterministic process. A more precise statement would be that financial returns *resemble* a stochastic process. Because, it could be argued that *ultimately*, there are no stochastic processes in the real world, pure randomness is an abstract mathematical concept, and everything in the real world is deterministic on the "exact" level. The ultimate laws that govern the Universe are believed to be deterministic at the most fundamental level, and hence if it were possible to comprehend the Universe in the minutest detail, it would be seen that all human activity is deterministic and stock prices in particular are "predetermined". Of course, this is not a very good way of looking at it from a practical point of view. The practical point of view is that human activity, such as making investment decisions and trading in the stock market, is so fantastically complex that it can never be comprehended on the "fundamental" level. All anybody can hope to do with regard to financial markets is to comprehend a tiny subset of all the possible aspects and degrees of freedoms. Then the definition of randomness from the practical point of view, as for example in the tossing of a coin, is that the real process is far to complex to comprehend, our knowledge is incomplete, and therefore the outcome is random just because of our lack of knowledge of all the details of the system.

Another example of this is the random number generator in your computer. If you start the random number generator, it generates a series of numbers that appear random, according to most statistical tests. They are distributed randomly within some interval, such as the interval from 0 to 1. However, the random number generator is actually a *deterministic* algorithm in your computer. If you start the random number generator exactly the same way each time, it generates exactly the same sequence of “random” numbers over and over. And if you knew the exact algorithm by which it generates these numbers, you could *predict* the exact sequence of numbers that it generates each time. It is only the lack of knowledge of this algorithm that makes the sequence of “random” numbers *appear* random.

So a more accurate statement is that the financial markets *resemble* a stochastic system and financial returns *resemble* a stochastic process. Then the question arises, “What kind of stochastic process best describes the returns series of a financial asset?” Bachelier believed that it was a simple Gaussian Random Walk, but we have come to understand [M, PB] that it is far more complex than this. But this is the crucial question when it comes to designing the optimal trading strategy. For example, if the markets are described by the simple Random Walk with drift, then short-term trading is fruitless, the expected return from trading is zero, and the only sensible strategy is Buy-and-Hold to take advantage of the long-term drift. However, all evidence suggest that the financial markets are far more unstable and unpredictable than the simple Random Walk model would suggest, so to simply trust in the long-term upward drift in the markets might be dangerous. An active trading strategy should be adopted, and securities bought or sold on a relatively long time scale in order to adjust the portfolio to changes in the market, the economy, and the situation with respect to each individual security such as changes in projected earnings or the market position of the company.

That having been said, it should also be pointed out that short-term trading greatly increases the variance or **risk** of the investment strategy, and it also greatly increases the transaction cost, so unless the trader has a very good understanding of the underlying correlation in the market (that a given investment is undervalued or overvalued), the best strategy is a long-term one. But my own view is that the middle ground is best, and depends on the frequency of the data that is available. From the above analysis of the amplification of correlation by smoothing, it would appear that to ferret out shorter-term correlation requires shorter-term data. As a rule of thumb, I would suggest trading no more frequently than 10 times the smallest time unit of the data that is available. So for daily data, the shortest trading time scale should be 10 trading days or two weeks. For real-time tick data, on the other hand, it might be possible to take advantage of much shorter-term correlations that might be present, but the trading time scale should still be long relative to the frequency of the incoming data.

Here is another subtle point with regard to randomness. Mathematically, it is a theorem by Kolmogorov [PW] that (considering an infinite, stationary time series) if the power spectrum is positive everywhere (zero only on a set of "measure zero") then the time series is *stochastic*. But if the power spectrum is discrete or has any finite gaps where it is zero, then the time series must be *deterministic*. Now consider the Fourier transform or Wavelet transform for a series. Both of these utilize data that are "cyclic", which form a closed loop, in our case of 2048 data points. These transforms are actually designed for *periodic* signals that repeat after every such cycle – if the data does not repeat then it is being approximated as cyclic. This cyclic data has a *discrete* Fourier or Wavelet power spectrum. But this cyclic, periodic data, like a periodic signal in Electronics, is actually a *deterministic* signal. Rigorously, it must be deterministic if the power spectrum consists of discrete frequencies. When all these discrete frequencies are mixed together, and the (discrete) power spectrum is flat, then the result is a signal that resembles stochastic white noise, even though in reality it is *deterministic*. In such a signal each individual frequency is independent and may be continued indefinitely into the future, resulting in complete predictability if the amplitudes and phases of each frequency component are known. But these amplitudes and phases are just what the Fourier or Wavelet transforms of the (periodic) signal give.

So here is an amusing thought: What if the returns series, which appears to be stochastic white noise, is in reality at least partly an admixture of deterministic cycles that exist independently of one another? Then if a trader picks out any one of these deterministic cycles and trades on that cycle, the other cycles will average out and the trader will make a profit from short-term trading, because that cycle will be *predictable* since it is deterministic. This would be the hallmark of a *linear* system, but in reality we expect the financial market to be an extremely complex *non-linear* system. Then all the individual cycles will *mix* together and predictability will be lost. If the system is completely random, with complete mixing of the component cycles, then there can be no predictability. However, in reality we might hope that the markets are not exactly efficient, the returns series is not exactly random, and due to the slight inefficiency there could be a small deterministic component to the signal. Then, using the Wavelet decomposition, we could hope that on each wavelet level the signal is partially predictable out to the time scale of that wavelet level. This could be possible even though the data appear to be random white noise, if there is a small deterministic component due to slight inefficiency in the market. Indeed, this is also the fundamental assumption underlying **Technical Analysis**.

Predictability of Market Processes

It must be concluded that, given the available statistical tests on a finite data set, it cannot be determined conclusively whether the data are random or deterministic, or some combination of the two. If the time series were infinitely long and stationary, and somehow an infinitely fine Fourier analysis, with a continuous Fourier spectrum, could be computed, and from this a continuous Periodogram (display of the Fourier power spectrum, the square of the Fourier coefficients), then one could determine from this Periodogram whether the data are random white noise, a correlated stochastic process, or a deterministic process. If the (continuous) Periodogram were non-zero everywhere and had a constant power spectrum, then the (stationary) stochastic process would be completely random white noise. If the Periodogram were non-zero everywhere but non-constant, then it would be random but correlated, and hence partially predictable. But by Kolmogorov's theorem, if the Periodogram had any finite gaps in which it were zero, and in particular if it consisted only of discrete frequencies, then it would be completely deterministic and perfectly predictable.

However, in real situations we have a finite data set, such as 2048 daily data units corresponding to 8 years of daily data. Then we may use the Fast Fourier Transform on the data of 2048 units, arranged as a circular data set. The Periodogram is then also discrete with 2048 values. Then it is difficult to tell absolutely whether the data are random or deterministic just by looking at the Periodogram. The (smoothed) Periodogram may appear to be that of random white noise, with a constant average value, hence with no predictability. On the other hand, if the Periodogram really consisted of discrete values, then it would be completely deterministic, and perfectly predictable. This would correspond to the case of something like a periodic electronic signal, with a fundamental period of 2048 time units, which repeats exactly every period. Then the signal is perfectly predictable – it simply repeats every fundamental cycle. So the actual financial data could be somewhere in between random white noise and a perfectly deterministic signal, and it would be impossible to tell from standard statistical tests alone which is the case.

In reality, the situation is more complex than the cases of stationary random or deterministic behavior, because we also suppose that the financial time series are **non-stationary**, with statistical properties that depend on time. This concept can be defined precisely only if we are given an infinite ensemble of infinitely long time series, all with the same (non-stationary) statistical properties, for then we can measure the statistical properties by averaging over the infinite ensemble of "instances" of the series. For the case of real data, there is only one instance of each financial time series, and we must rely on sample averages over finite sections of the data, or **smoothings** of the data, to approximately measure the statistical properties. Ultimately, we must define and measure the predictability of the time series by computing various smoothings and indicators based on the past data, and measuring the average **correlation** of these indicators with future returns. When these correlations are averaged over many different time series and many different securities in a portfolio, they become a measure of the validity of the indicators and the trading rules that are based upon them.

Correlations in Financial Time Series

According to the standard **Modern Portfolio Theory** of the '70s, a stochastic process consisting of just Gaussian white noise described financial returns. A famous hedge fund, Long-Term Capital Management [FFK, p.264][M, pp.105-107] based its trading strategy on this assumption, which also underlies the Black-Sholes options pricing formula, with disastrous results (in September 1998), and the rest is history. Finally it was realized that the statistical distribution of returns is not Gaussian, but in fact has "fat tails" which means that the likelihood of extreme events such as the market crash of 1987 is much more likely than Gaussian statistics would predict. Also it is known now that the volatility is not constant but is time-varying. Mandelbrot [M] has explained this behavior as being due to **fractal statistics**. The time series is described as a **fractionally differenced** process with an **anomalous fractal dimension**. These types of process were originally discovered by studying the statistical distribution of the Nile river levels. What this means for us is that there should exist **long-range correlations** in the financial returns, but whereas in ordinary correlated processes these correlations decay away exponentially with the time interval, for the case of fractal statistics the correlations decay more slowly, as a power law. This then gives rise to the observed erratic behavior of financial markets, with a much greater likelihood of extreme "outlier" events such as market crashes than would be expected on the basis of Gaussian statistics.

Thus, fundamentally, due to this long-range correlation of returns due to fractal statistics, we can expect to find underlying **trend persistence** in financial data. This trend persistence is also a cornerstone of **Technical Analysis** [Pr]. When viewed in terms of the Periodogram, this fractal trend persistence would take the form of a sharp spike at the very low-frequency end of the power spectrum. Unfortunately, this is very difficult to see in the actual computed Periodogram because there are very few degrees of freedom at these low frequencies, and the spike is obscured by stochastic noise in the Periodogram. The Wavelet analysis is better suited to this type of spectrum due to long-range fractal statistics.

However, we also know from experience that there is a definite **return to the mean** mechanism in financial markets. This means that when prices are low, they tend to rise and when prices are high, they tend to fall. So this correlation is between past *price levels* and future *returns*. However, this type of behavior could also be explained as persistence of the returns on individual Wavelet levels. Each wavelet level represents an octave of Fourier frequencies, and so both the (relative) prices and the returns have an approximately cyclical behavior corresponding to the average frequency of the wavelet level. In other words, on each wavelet level the **Relative Price**, **Velocity**, and **Acceleration** indicators are each shifted relative to each other by one-quarter cycle. So on an individual wavelet level, the **trend persistence** correlation of past and future returns, and the **return to the mean** (anti-)correlation of past prices a quarter cycle in the past and future returns, are really two aspects of the same thing. So the conclusion is, instead of just postulating a long-range persistence of returns due to fractal statistics, we could instead generalize this and postulate persistence on each wavelet level of the returns, with a time scale roughly the same as that of the wavelet level. The persistence due to fractal statistics would then appear as just the persistence on the highest or "scaling" level of the wavelet analysis.

Non-linear Behavior and Market Crashes

One particular situation deserves further comment. This is the situation such as occurred in 1929 and 1987, where the market enters a phase with a strong uptrend, where the **Velocity** predictor and **Relative Price** predictor are both strongly positive. Then due to the nonlinear effect, the longer the uptrend stays in place, the more positively correlated the future returns are to the **Velocity** predictor, but at the same time the **Relative Price** gets higher and higher resulting in it being more negatively correlated with future returns. Then the market enters a delicate balance, where the slightest fluctuation downward can break the uptrend and result in prices crashing due to the highly inflated **Relative Price**. But the dangerous thing about this situation is that the exact break point is in principle unknown and unpredictable. This phenomenon has been compared to avalanches, earthquakes, and cascading piles of sand. In each case the probability of the sudden occurrence can be estimated, but the exact point in time when it occurs is unpredictable. Using the linear **Relative Price** and **Velocity** predictors, their correlation with future returns would cancel, and the projected return would be somewhere around the neutral point. However, the time-dependent volatility could also be estimated (in principle) by the **Wavelet Adaptive Filter** technique, and it could yield a strong probability for a large downward move in prices, due to the large (positive) absolute values of the two predictors.

An interesting point to note about market crashes, for example the one in 1987, is that in many cases they seem to illustrate a **return to the mean** mechanism. If you look at the long-term (say, 10 years) logarithmic price graph (say, the S&P 500) of the 1987 crash, you will note that the market entered a strong uptrend a couple of years before the crash, from a more moderate longer-term uptrend. Then when the market crashed, the price levels fell very rapidly from the level of the shorter-term strong uptrend back to the lower level longer-term uptrend. It appears to be a jump from the shorter-term trending mean to a longer-term trending mean. This seems to imply that if we can identify the trends on various time scales, they might have predictive power because we can suppose that prices will revert from the shorter-term trends back to the levels of the longer-term trends on time scales comparable to those of the trends. The sudden jumps are no doubt nonlinear phenomena, but a linear model in which the trends on various time scales can be identified can approximate this behavior. But this is just what the **Wavelet** analysis is designed to accomplish.

2.4 Portfolio Optimization and Rebalancing

The objective of the ***QuanTek*** program is to provide a trading strategy that will yield ***maximum returns with minimum risk***. However, the primary problem with active trading is that it increases risk without any corresponding increase in returns. This is why in many cases, the buy-and-hold strategy is the one that yields the maximum return with the minimum risk. If there is no correlation present and the trading rules do not work, then active trading increases the variance of the returns (risk) with no corresponding increase in the mean return. Actually, a difficult milestone to surpass is the return of the overall market. To simply buy and hold an index fund is a prudent strategy for many investors. More generally, it is a general principle that greater investment returns are accompanied by greater risk, as exemplified by high-grade bonds vs. junk bonds. So if you decide to actively trade, the important question is how to increase the returns above the market return, while keeping risk under control.

The way to accomplish the risk control is to trade within an overall **optimal portfolio** strategy rather than in each security independently. Typically one computes the optimal portfolio that yields the overall maximum return with minimum risk, and then one must do **portfolio rebalancing** [FFK] periodically in order to maintain the optimum weights of all securities in the portfolio. This optimal weight is determined, typically, as in the classical **Markowitz** method [SAB], using the **expected return** of each security along with the **covariance matrix** of the returns of all the securities in the portfolio. As a simplification of this, a **Factor Model** [SAB], utilizing only a few combinations of the securities (factors) instead of all the securities, or the **Capital Asset Pricing Model (CAPM)** [SAB], utilizing only the *market portfolio*, may be used. (In fact, according to the CAPM, the **market portfolio** as a whole is the "ultimate" **optimal portfolio**.) Then, as the better performing securities rise in price and become over-weighted, they are sold at periodic intervals, and the under-weighted securities are bought. This mechanism automatically incorporates the **buy low – sell high** strategy, thereby taking advantage of this important correlation with future returns. Thus the trading strategy for each security is linked to its performance and optimal weight within the whole portfolio.

This begs the question of how the securities are to be chosen in the first place. Ultimately, it is best to take fundamentals into account when selecting the securities for the portfolio. Also to achieve diversification, securities from different sectors should be chosen, since the securities within each sector tend to move in tandem and be correlated with each other. More generally, if the covariance matrix is used, securities are chosen to be *anti-correlated* in order to achieve the lowest possible portfolio risk for a given return. There is a problem with this, however, since it has been shown that the covariance matrix of securities is in reality mostly stochastic noise [LCBP]. Only a few combinations of the securities correspond to statistically significant eigenvalues of the covariance matrix, and one of these is simply the sum of the securities with a certain set of weights corresponding to the **market portfolio**. So one easy rule of thumb, rather than having to diagonalize the covariance matrix to compute the portfolio and arrive at a solution which is an awkward mixture of the securites [FFK], is to just maintain equal weights (price per share times number of shares) of each security in the portfolio. However, better results might be obtained by utilizing more sophisticated methods such as Value-At-Risk (VAR) [BP].

Our approach is a modification of the **Markowitz** method [SAB], which involves maximizing a **Q-function** (or "Quality" function) which uses a portfolio constraint *quadratic* in the portfolio weights. The Q-function contains the **expected returns**, **covariance matrix**, and is *quadratic* in the **portfolio weights**. The portfolio constraint is to ensure that the portfolio value remains constant. Normally, in the Markowitz method, the portfolio constraint is linear in the portfolio weights, but this becomes a difficult problem in **constrained quadratic optimization**. It is called constrained because the portfolio weights are constrained to be positive. With the quadratic portfolio constraint the portfolio weights are allowed to be negative, corresponding to the possibility of short-selling, so the portfolio weights themselves are unconstrained. I therefore call this method **unconstrained quadratic optimization**. This is a practical way to calculate the optimal portfolio which is computationally tractable, yet corresponds to the usual situation of the individual investor. But it should be clear that the solution that is obtained depends on the choice of the optimization procedure that is used.

Optimal Buy/Sell Points

The remaining question is when to buy and sell. In general, it is difficult to predict optimal buy/sell points – the apparent buy/sell points are generally just due to stochastic fluctuations of price. In the standard portfolio-rebalancing scheme, the portfolio is rebalanced back to its optimal weighting at fixed time intervals, say every month. However, the ***QuanTek*** program does attempt to identify optimal **buy/sell points**, in the hope that if the portfolio is rebalanced at these points, then there is a chance to obtain a little better price. First, a time scale N for trading is identified, which is the time scale for holding at which the return and risk is to be optimized. Then the three predictors **Relative Price**, **Velocity**, and **Acceleration** are filtered on these time scales. These three filtered predictors are used to identify optimal **buy/sell points** by the following rule: For a **buy signal**, the **Relative Price** should be minimum, the **Velocity** should be going through zero in the positive direction, and the **Acceleration** should be a maximum (positive turning point). For a **sell signal**, the **Relative Price** should be maximum, the **Velocity** should be going through zero in the negative direction, and the **Acceleration** should be a minimum (negative turning point). More generally, a range of permissible buy/sell points are identified using these predictors: A **buy signal** corresponds to the quarter cycle with negative **Relative Price** and positive **Velocity**, and a **sell signal** corresponds to the quarter cycle with positive **Relative Price** and negative **Velocity**. If the individual security is bought or sold only at these points, then this increases the chance of a favorable outcome over the *N*-day holding period, and decreases the chance of buying too high or selling too low.

The caveat here is that the values of the smoothed predictors at the present time depend on the **Price Projection**, so the predictors must incorporate this future projection. Then, the **optimal portfolio** usually also takes into account the **Price Projection** in the form of **expected returns**. So, consistent with the **buy/sell points**, a security with a positive expected return should be bought or held, and a security with a negative expected return should be sold or held, depending on the price level. But the **expected return** derived from the **Price Projection** depends on time, which in turn leads to the **active trading** strategy, continually rebalancing the portfolio in response to the ever-changing expected return and the specified optimal weighting.

It should be pointed out that the **expected return** is an important element of the estimation of the optimal weighting of the portfolio. Even though returns are hard to predict, they must be predicted in order to arrive at some estimate for the expected return for the **optimal portfolio** calculation. So this is the difficult challenge, making full use of the latest results from the science of **Financial Econometrics** (and **EconoPhysics**) that the ***QuanTek*** program attempts to meet.

2.5 References

[B] Louis Bachelier, Théorie de la Spéculation (doctorial dissertation),
Annales Scientifiques de l'Éciole Normale Supérieure (iii),
Vol.17, pp.21-86 (1900) Translation: Cootner (1964)

[BP] Jean-Philippe Bouchaud & Marc Potters,
Theory of Financial Risks: From Statistical Physics to Risk Management, 2nd ed.,
Cambridge University Press, Cambridge, UK (2000)

[FFK] Frank J. Fabozzi, Sergio M. Focardi, & Petter N. Kolm,

Financial Modeling of the Equity Market: From CAPM to Cointegration,
John Wiley & Sons, Hoboken, New Jersey (2006)

[G] Christian Gourieroux, ARCH Models and Financial Applications,
Springer-Verlag, New York (1997)

[H] Andrew C. Harvey,
Forecasting, structural time series models and the Kalman filter,
Cambridge University Press (1989)

[Hay] Simon Haykin, Adaptive Filter Theory, 4th ed.,
Prentice Hall, Upper Saddle River, NJ (2002)

[LCBP] Laurent Laloux, Pierre Cizeau, Jean-Philippe Bouchaud, & Marc Potters,
"Noise Dressing of Financial Correlation Matrices",
Physical Review Letters, Vol.83, No.7 (16 August 1999)

[M] Benoit Mandelbrot, The (Mis)behavior of Markets,
Basic Books, New York (2004)

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

[Pr] Martin J. Pring, Technical Analysis Explained, 3rd ed.,
McGraw-Hill, New York (1991)

[PB] Wolfgang Paul & Jörg Baschnagel,
Stochastic Processes: From Physics to Finance,
Springer-Verlag, New York (1999)

[PW] Donald B. Percival & Andrew T. Walden,
Wavelet Methods for Time Series Analysis,
Cambridge University Press, Cambridge, UK (2000)

[SAB] W.F. Sharpe, G.J. Alexander, & J.V. Bailey, Investments, 5th ed.,
Prentice Hall, Englewood Cliffs, NJ (1995)