Final Project Report

On

Financial Time Series Signal Processing

By

Harshavardhan Bapat (2015B3A30580H)

Sajal Jain (2015B3A80597H)

Kallol Bairagi (2015B3A30615H)

Under the supervision and guidance of

Dr Manish Narwaria

SUBMITTED IN FULLFILLMENT OF THE REQUIREMENTS OF

EEE/INSTR F376: DESIGN PROJECT

BIRLA INSTITUTE OF TECHNOLOGY AND SCIENCE PILANI (RAJASTAN)

HYDERABAD CAMPUS

(April 29, 2019)

Acknowledgements

I would like to express my gratitude and sincere thanks to Dr Manish Narwaria, Assistant Professor, Department of Electrical and Electronics, BITS Pilani Hyderabad Campus, for providing me the opportunity to pursue this project and the related research on Financial Signal Processing techniques. The guidance and inputs provided by him have been invaluable to this project.

I would also like to thank Dr Sanket Goel, Head of Department of Electrical and Electronics, BITS Pilani Hyderabad Campus for permitting me to do this project. This project would not be possible without the resources provided by the department.

Lastly, I would like to thank Prof. G Sundar, Director, BITS Pilani Hyderabad Campus for enabling me to pursue this project by providing the necessary infrastructure and permitting access to the same without which this thesis would have been impossible to pursue.

Birla Institute of Technology and Science-Pilani,

Hyderabad Campus

Certificate

This is to certify that the project report entitled "Financial Signal Processing" submitted by Mr. Harshavardhan Bapat (ID No 2015B3A30580H), Ms. Sajal Jain (ID No 2015B3A80597H), and Mr. Kallol Bairagi (ID No 2015B3A30615H) in partial fulfilment of the requirements of the course EEE/INSTR F376, Design Oriented Project Course, embodies the work done by him under my supervision and guidance.

Date: (Dr Manish Narwaria)

BITS- Pilani, Hyderabad Campus

Abstract

Important financial data such as stock prices, index prices, exchange rates change on a daily basis. These changes are brought about by effects from various macroeconomic factors, trading patterns and general public sentiment. The ability to forecast in such cases is highly desirable as it allows potential investors to identify profitable investments, and inform investors when it is ideal to close out positions to lock in profits. As this data continuously changes over time, it can be represented as a time-variant signal. Processing techniques applicable to such time variant systems can provide important information relevant to forecasting. It is possible that signal decomposition techniques such as Wavelet Transforms may provide superior capacity to forecast time series signals such as financial data. The same is to be verified in this study. Before signal decomposition techniques are analysed, traditional forecasting techniques were studied and implemented, to note the shortcomings of such models. Statistical approaches to treating financial data have proven useful and deliver statistically significant results. However, there is reason to believe a technical approach using analytical tools having properties which present a significant advantage in treatment of financial data may provide useful information regarding the data, thereby procuring valuable information that is required as inputs for artificial neural networks, which help accurately predict future behaviour of the observed securities. This report comprises of a summary of such relevant tools and techniques to better analyse financial data.

Table of Contents

Summary of Statistical Approach

Traditional econometric models were analysed and employed to fit time series data such as USDINR exchange rates. The models were then used to generate forecasts for the time series by using univariate forecasting methods. Autocorrelation was observed among residuals of values forecasted by Holt's Exponential forecasting, which signified a scope for improvement in the model.

ARIMA model forecasting is considered superior to Exponential forecasting for the reason that it employs no assumption regarding the residuals of errors as assumed by the exponential forecasting methods. On fitting the data in the appropriate ARIMA model, it was deduced that the data closely replicated an ARMA (2, 2) model, where the current value of the series depended on the previous two lagged values of itself, as well as the residuals of the lagged values. Forecasts of this model were also analysed, and some autocorrelation among residuals was observed. A slightly skewed normal distribution of residuals of forecasts once again suggested that a better model may be fitted.

Traditional econometric models of forecasting are limited in their scope to perfectly map and replicate possible trends in data. The authors believe some decomposition techniques employed on time-variant systems may help deduce further information regarding financial signals, which may provide more accurate forecasts.

Technical Approach

Fourier Transforms

Fourier transform theory states a given time series can equivalently be characterized either in time domain or in frequency domain. In general, transformation of the signal representation between the time domain and the frequency domain (also known as the spectral domain) is achieved by computing the Fourier transform (FT) and the inverse Fourier transform (IFT) as given in equations (1.a) and (1.b), respectively

$$X(f) = \int_{-\infty}^{\infty} x(t)\, e^{-j2\pi ft}\, dt \qquad \textbf{(1.a)}$$

$$x(t) = \int_{-\infty}^{\infty} X(f)\, e^{j2\pi ft}\, df \qquad \textbf{(1.b)}$$

where $x(t)$ is the time function evaluated at time t, $X(f)$ is the Fourier transform evaluated at frequency f, and j =-1 is the unit imaginary number.

For discrete-time problems, the discrete Fourier transform (DFT) and the inverse discrete Fourier transform (IDFT) need to be used instead of the continuous time FT and IFT pair given in Equation 1. The Fast Fourier Transform algorithm (FFT) and its inverse (IFFT) are computationally optimized signal processing tools that can be used to compute Fourier transform pairs for discrete signals. While the characterization of discrete data in time domain requires the techniques for the analysis of time series, characterization in frequency domain calls for the techniques of discrete spectral analysis, both producing essentially the same results.

Alternatives to Fourier Transforms

The spectral content of the series varies as time progresses, rendering the conventional Fourier theory inadequate to fully describe the cyclical characteristics of the series. The joint Time-Frequency Representation (TFR) techniques overcome this problem as they are capable of analyzing a given function of time (continuous or discrete) in time domain and in frequency domain simultaneously.

Both continuous and discrete Fourier transforms have proved indispensable as data analysis tools for stationary signals. Yet, if the statistical properties of a time signal are time-variant and hence, its spectral content varies as time progresses, the conventional Fourier theory becomes inadequate to fully describe the signal characteristics. The TFR techniques overcome this problem as they are capable of analyzing a given function of time (continuous or discrete) in time domain and in frequency domain simultaneously. In other words, TFRs can characterize a given time signal in the two-dimensional joint time-frequency domain enabling localization both in time and frequency within the resolution limits allowed by the uncertainty principle.

Time series may be analyzed either in the time domain or in the spectral (or frequency) domain, both producing essentially the same results. Transformation of the signal representation from one domain to another is achieved by computing the Fourier transform (FT) and the inverse Fourier transform (IFT). If the spectral content of a given signal varies as time progresses, however, the conventional Fourier theory fails to fully describe the contribution of arbitrarily chosen spectral components over certain time bands. The time frequency representation (TFR) techniques have emerged as viable solutions to this challenging problem as they analyze a given function of time (continuous or discrete) in time domain and in frequency domain simultaneously. Gabor transform (GT), short time Fourier transform (STFT) and Wavelet transform (WT) are linear time-frequency representations whereas Wigner distribution (WD), Spectrogram (magnitude square of STFT), Scalogram (magnitude square of WT), Choi-Williams distribution (CWD) and Page distribution (PD) are some of the well-known quadratic time-frequency representations.

Wavelet Transforms

Wavelets have had popular usage in natural sciences, especially in earth sciences, and the past couple decades have seen usage of wavelet methods in engineering as well. However, the first applications of wavelets in economics and finance emerged only recently, despite its evident suitability for this discipline. One of the fundamental advantages of wavelet analysis over Fourier analysis is the capability to decompose time series into different components by virtue of time and frequency localization properties it holds. An observed time series may contain several structures, each occurring on a different time scale. Wavelet techniques possess an inherent ability to decompose this kind of time series into several sub-series which may be associated with a particular time scale. The problem of Fourier analysis is that the time information is lost completely. The assumption of "natural" periods and stationarity that are inherent in the Fourier methods are also problematic.

A wavelet is a rapidly decaying wave like oscillation that has zero mean.
Unlike sinusoids which extend to infinity a wavelet exists for a finite duration.
Wavelets come in wide range of shapes and sizes. The availability of a wide range of wavelets is a key strength of wavelet analysis.

There are two important wavelet concepts – scaling and shifting.

Scaling refers to the process of stretching or shrinking this signal in time with the following equation. –

$$\psi(t\ /\ s)s > 0$$

S is the scaling factor and corresponds to how much the wavelet is scaled in time. The scale factor is inversely proportional to frequency. For wavelets, there is a reciprocal relationship between the scale and the frequency with a constant of proportionality. This constant of proportionality is called the center frequency of the wavelet. This is because, unlike the sine wave, the wavelet has a band pass characteristic in the frequency domain. Mathematically, the equivalent frequency is defined using this equation –

$$F_{eq} = \frac{C_f}{s\delta t}$$

Therefore, when you scale a wavelet by a factor of 2, it results in reducing the equivalent frequency by an octave. A larger scale factor results in a stretched wavelet, which corresponds to a lower frequency, and vice versa. A stretched wavelet helps in capturing the slowly varying changes in a signal, while a compressed wavelet helps in capturing the abrupt changes.

Shifting a wavelet simply means delaying or advancing the onset of the wavelet along the length of the signal. A shifted wavelet represented using this notation means the wave is shifted and centered at k. We need to shift the wavelet to align with the feature we are looking for in the signal. The two major transforms in wavelet analysis are continuous and discrete wavelet transforms. Continuous wavelet analysis is useful in time frequency analysis and filtering of time localized frequency components.

One can use this transform to obtain a simultaneous time frequency analysis of a signal. Analytical wavelets are best suited for time frequency analysis as these wavelets do not have negative frequency components. Examples of analytical wavelets that are suitable for continuous wavelet analysis are Morse wavelets, bump wavelets and analytic Morlet wavelet. The output of cwt are coefficients which are functions of scale or frequency and time. When you scale a wavelet by a factor of 2, it results in reducing the equivalent frequency by an octave. In cwt you have the added flexibility to analyze the signal at intermediary scales within each octave. This allows for fine scale analysis. This parameter is referred to as the number of scales per octave. The higher the number of scales per octave, the finer the scale discretization. Typical values for this parameter are 10, 12, 16, and 32. The scales are multiplied by the sampling interval to obtain a physical significance. Each scaled wavelet is shifted in time along the entire length of the signal and is compared with the original signal. This process can be repeated for all the scales resulting in coefficients that are a function of wavelet scale and shift parameters. For example, a signal with 1000 samples analyzed with 20 scales results in 20000 coefficients. In this way you can better characterize oscillatory behavior in signals with the continuous wavelet transform.

Mathematical Definitions

The mathematical definitions of GT, STFT, WD and PD are as follows:

Gabor Transform

The Gabor expansion coefficients *Gx* (*n*, *k*) of a given time signal *x*(*t*) are implicitly defined by

$$x(t) = \sum_{n} \sum_{k} G_x(n,k)\, g_{nk}(t) \tag{2.a}$$

with

$$g_{nk}(t) = g(t - nT)\, e^{j2\pi(kF)t} \tag{2.b}$$

being the basis functions of the expansion, which were originally taken to be time frequency shifted Gaussian functions by Gabor (1946), as the Gaussian functions are well concentrated both in time and frequency domain. Then, the expansion coefficient *Gx* (*n*, *k*) is expected to indicate the signal's time and frequency content around the point (nT, kF) in the joint time-frequency domain.

Short Time Fourier Transform

The STFT of a given time signal *x*(*t*) is computed by

$$STFT_x^{(\gamma)}(t,f) = \int_{t'} x(t')\, \gamma^*(t'-t)\, e^{-j2\pi f t'}\, dt' \tag{3}$$

where *g*(*t*¢ - *t*) is the chosen *window of analysis* which is centered at *t* ¢ = *t* and the superscript * denotes complex conjugation. As implied by this definition, the STFT of a signal may be interpreted as the *local Fourier transform* of the signal around the analysis time t.

Wigner Distribution

The auto-Wigner distribution of a given time signal $x(t)$ is given by

$$W_x(t,f) = \int_\tau x(t+\frac{\tau}{2})\, x^*(t-\frac{\tau}{2})\, e^{-j2\pi f\tau}\, d\tau \quad \textbf{(4)}$$

The WD is a real-valued quadratic TFR preserving time shifts and frequency shifts of the signal. The frequency (time) integral of the WD corresponds to the signal's instantaneous power (spectral energy density) as the WD satisfies the so-called *marginals*. As a matter of fact, the WD is the only quadratic TFR satisfying all of the desired properties of the *energetic* time-frequency representations.

Page Distribution

The Page distribution of a given time signal $x(t)$ is defined as

$$PD_x(t,f) = \frac{d}{dt}\left|\int_{-\infty}^{t} x(t')\, e^{-j2\pi f t'}\, dt'\right|^2 \quad \textbf{(5)}$$

The PD is also an energetic, shift-invariant, quadratic TFR like the WD. Most of the desirable properties satisfied by the WD are also satisfied by the PD except for a few of them such as the property of having a finite frequency support.

Wavelets in Finance

The following presents a literature review of wavelet applications in finance. It predominantly focuses on decomposition applications of wavelet methods, while also shedding some light on applications to interdependence studies with wavelets.

Capobianco (2004) applies wavelet methods to the multiresolution analysis of high frequency Nikkei stock index data, by applying the matching pursuit algorithm suggested by Mallat and Zhang (1993). He argues that the algorithm suits financial data, and shows how the wavelet matching pursuit algorithm can be used to uncover hidden periodic components.

Gencay et al (2001a) used wavelet methods to investigate the scaling properties of foreign exchange rates. They found that foreign exchange rate volatilities are described by different scaling laws on different horizons. They too used the maximal overlap discrete wavelet transform estimator of the wavelet variance to decompose variance of the process to come across this observation. Gencay et al. (2003) decompose a given time series on a scale-by-scale basis. The wavelet variance of the market return and the wavelet covariance between the market return and a portfolio are calculate on each scale to obtain an estimate of a portfolio's beta. It was concluded that the estimations of the CAPM are more relevant in the medium and long run than on short time horizons. Gencay et al. (2004) present a powerful method to analyze the relationship between stock market returns and volatility on multiple time scales using wavelet decomposition. The leverage effect was found to be weak at high frequencies but becomes prominent at low frequencies. It was also found that positive correlation between the current volatility and future returns becomes dominant on the timescales of one day and higher, lending evidence that risk and return are positively correlated.

The frequency components of European business cycles were analyzed by Crowley and Lee (2005) using wavelet multiresolution analysis. They use a real GDP as a proxy for the business activity of European countries. The analysis is performed using the maximal overlap discrete wavelet transform, and significant differences between the countries is found, where the degree of integration varies significantly. They also found out that most of the energy in these economic time series can be found in longer term fluctuations. Also, indications were found that recessions are a result of a simultaneous dip in growth cycles at all frequencies.
Vuorenmaa (2005, 2006) analyses stock market volatility using the maximal overlap discrete wavelet transform. He observes that the global scaling laws and long memory of stock's volatility may not be time-invariant.

Wavelets have also been widely used to study interdependence of economic and financial time series. The studies presented in the following have also decomposition aspects but their main aspect is in the interdependence of processes.

In & Kim (2006c, 2007), In et al. (2008) and Kim & In (2005, 2006, 2007) have conducted many studies in finance using the wavelet variance, wavelet correlation and cross-correlation. Kim & In (2005) study the relationship between stock markets and inflation using the MODWT estimator of the wavelet correlation. They conclude that there is a positive relationship between stock returns and inflation on a scale of one month and on a scale of 128 months, and a negative relationship between these scales. Furthermore, they stress how the wavelet-based scale analysis is of utmost importance in the economics studies since their results solve many puzzles around the Fisher hypothesis previously noted in literature. In et al. (2008) study the performance of US mutual funds using wavelet multi-scaling methods and the Jensen's alpha. The results reveal that none of the funds are dominant over all timescales. In & Kim (2006c) study the relationship between stock and futures markets with the MODWT based estimator of wavelet cross-correlation. There is a feedback relationship between the stock and the futures markets on every scale. The results also reveal that correlation increases as time scale increases. In & Kim (2007) examine how well the Fama- French factor model works on different time scales. They conclude that the SMB (small capital business minus big capital business) and the HML (high book-to market minus low book-to-market) share much of the information with alternative investment opportunities in the long run but not in the short run. Therefore, the importance of scale dimension is verified again. Kim & In (2006) find that correlation between industry returns and inflation does not vary along with the scale. Furthermore, they find indications that industry returns can be used as a hedge against inflation, depending on the particular industry.

Hilbert Transforms using Market Modes

Market Modes

Much research has been done to prove that the market is indeed efficient. However, the fact that there exists a number of traders who are continuously successful is adequate proof that markets are not necessarily completely efficient. The failure of the efficiency hypothesis in several cases is sufficient evidence to invalidate the hypothesis itself.

Classical efficient market models are often concerned with the adjustment of security prices to three information subsets. Weak form tests comprise the first subset, in which we are simply given the historical prices. The second subset is semi-strong form tests that concern themselves with whether prices efficiently adjust to other publicly available information. Strong form tests, the third subset, are concerned with whether investors have monopolistic access to any information relevant to price formation. The general conclusion, particularly for the weak form tests, is that the markets can be only marginally profitable to a trader. In fact, only the strong form tests are viewed as benchmarks against deviations from market efficiency. These strong form tests point to activities such as insider trading and the market-making function of specialists.

The efficient-markets-model statement that the price fully reflects available information implies that successive price changes are independent of one another. In addition, it has usually been assumed that successive changes are identically distributed. Together, these two hypotheses constitute the Random Walk Model, which says that the conditional and marginal probability distributions of an independent random variable are identical. In addition, it says that the probability density function must be the same for all time. The simplest modification of the Random Walk is to allow the coin toss to determine the persistence of motion. This is called the Drunkard's Walk solution. In other words, with probability p the drunkard makes his next step in the same direction as the last one, and with probability l-p he makes a move in the opposite direction. The Drunkard's Walk solution can describe two market conditions. In the first condition, the probability is evenly divided between stepping to the right or the left, resulting in the Trend Mode. The second condition, the probability of motion direction is skewed, results in the Cycle Mode.

The market only has a single dominant cycle most of the time. When multiple cycles are simultaneously present, they are generally harmonically related. This is not to say that nonharmonic simultaneous cycles cannot exist-just that they are rare enough to be discounted in simplified models of market action. The general observation of a single dominant cycle tends to support the notion that the natural response to a disturbance is monotonic harmonic motion. A more complete model of the market can be achieved by dividing the market action into a Cycle Mode and a Trend Mode. By having only two modes in our market model, we can switch trading strategies back and forth between them, using the more appropriate tool according to our situation.

Hilbert Transform

The Hilbert Transform is a procedure to create complex signals from the simple chart data familiar to all traders. Once we have the complex signals, we can compute indicators that are more accurate and responsive than those computed using conventional techniques.

When data are sampled at a sampling frequency, that sampling frequency acts like a radio carrier signal. That is, the real data being sampled are heterodyned into upper and lower sidebands of the sampling frequency. Mathematically, heterodyning is multiplying two frequencies (and then filtering to select the desired output). So, if we have a baseband data frequency of fb, the heterodyning can be described as the product of two signals. By a trigonometric identity, this product results in the sum and difference frequencies as the lower sideband can be considered as a negative frequency relative to the sampling frequency, and the upper sideband can be considered as a positive frequency relative to the sampling frequency. Furthermore, every harmonic of the sampling frequency exists. Each harmonic also has an upper and lower sideband containing the baseband signals. Since the lower sideband of the sampling frequency exists, it could extend down into the baseband range of frequencies. This is called the Nyquist sampling criteria. In trading, this means the absolute shortest period we can use is a 2-bar cycle, or a frequency of 0.5 cycles per bar. The sampling frequency can be weekly, daily, hourly, and so on, but the shortest period we can consider in any time frame is a 2-bar cycle. We can synthesize the analytic signal by summing the two complex signals. When the real component is summed with the imaginary component in the equations, the two complex signals can be called the Inphase (i.e., the Cosine) component and the Quadrature (i.e., the Sine) component. Quadrature means being rotated by 90 degrees.

We can approximate the Hilbert Transformer by truncating the extent. For example, we could truncate the filter at n = 7. In this case, where the detrended Price is represented by P, the Quadrature component (Q) of the Hilbert Transform can be written as

$$Q = (P/7 + P[2]/5 + P[4]/3 + P[6] - P[8] - P[10]/3 - P[12]/5 - P[14]/7)/(1 + 1/3 + 1/5 + 1/7)$$

This short Hilbert Transformer has a lag of only 3 bars.

We attempt to construct a Hilbert Transformer in Python with assistance from John Ehler's Rocket Science for Traders.

Conclusions

The shortcomings of the Fourier Transform in analysing time variant financial data are overcome by various other transformation tools such as Wavelets and Hilbert Transforms. These transforms allow localisation of frequency in time and space, thereby providing an advantage over Fourier transforms. This information is important and relevant in the case of financial analysis, as it allows the analyst to separate a time series of financial data into multiple individual series each with its individual characteristics. This multiresolution allows an analyst to localise significant events which may allow her to extract information regarding the recurrence of such events. Wavelet and Hilbert transform resolution also allow analysts to identify and separate various trends and cycles in financial data. As information is crucial in terms of performing consistently profitable trades, knowledge about repetitions in cycles and trends is extremely useful to an analyst or a trader. Derivation of frequencies of such trends and cycles can provide the analyst with the necessary inputs one may need to build an artificial neural network to predict the movement of the stock, index or security under observation.

References

1. Box, G.E.P., Jenkins, G.M., & Reinsel, G.C. 1994. Time-Series Analysis, Forecasting and Control. Holden-Day, San Francisco.
2. Dahlhaus, R. 1997. Fitting time series models to nonstationary processes. Annals of Statistics.
3. Drakakis, Konstantinos. Application of Signal Processing to the Analysis of Financial Data
4. Drakakis, Konstantinos. 2008. Analysis of Financial Data through Signal Processing Techniques
5. Ehler, John. Rocket Science for Traders
6. Grandell, Jan. Time Series Analysis
7. Jothimani, D., Shankar, R., Yadav, S.S., 2015. Discrete Wavelet Transform-Based Prediction of Stock Index
8. Ranta, Mikko. 2010. Wavelet Multiresolution Analysis of Financial Time Series
9. Tsay, R.S. 2002. Analysis of Financial Time Series. Wiley, New York.
10. https://www.statmethods.net/advstats/timeseries.html
11. https://datascienceplus.com/time-series-analysis-using-arima-model-in-r/
12. https://www.thoughtco.com/the-augmented-dickey-fuller-test-1145985
13. https://www.analyticsvidhya.com/blog/2018/09/non-stationary-time-series-python