CEDAR RIVER CAPITAL LLC

Paul C Latta, CFA
October 28, 2019

Channeling Monsieur Fourier

"Mathematics is the language in which God has written the universe"

- Galileo Galilei

There is a certain fundamental purity to the study of mathematics, a discipline that stands alone in its quest for absolute truth. The quest for truth has been quite fruitful for modern society as the discovery and revelation of mathematical relations has been a key factor setting the pace for many scientific and technological advances. Most investors are keenly aware of the power of mathematical reasoning and are instinctively drawn to quantitative factors such as momentum, valuation, growth, earnings quality, mean reversion and more.

The historical chart of a stock price intuitively lends itself to mathematical modeling. The pattern-matching feature of the brain sees a price chart and automatically attempts matching it with familiar patterns. Familiar patterns could be mathematical functions such as linear growth, exponential growth or cyclical sine waves. Or the brain might match patterns with an historic library of other stocks that have had similar charts.

Naturally there are challenges to realizing profits from simple pattern recognition, not the least of which being that obvious patterns will be seen by all and will therefore tend to disappear rather quickly. Statisticians know that patterns can sometimes form in random data, leading to false positives. And psychologists describe a human weakness to see patterns in random data, a condition known as apophenia. Pattern-matching is instinctive but far from perfect.

Of course, a stock price chart is more than a simple pattern on the screen. It is also a convenient progress report for investors, a very visual record of the ups and downs on the road to realizing long-term shareholder value. At a deeper level, the price chart is a Cliff Notes summary of a more epic saga, a saga with an enormous cast of actors including company management, investors, politicians, consumers, unions, media and more. There is a lot going on under the surface!

This paper explores the various wave theories of the stock market, touching on historical visual approaches as well as modern quantitative approaches using Fourier Transforms. It also addresses real world waveform implementation issues and offers a surprisingly intuitive yet little used waveform paradigm for classifying investment managers. The paper is broken down into seven parts, in order:

- *Don't Judge a Book by its Cover:* Challenges with simple stock charts as viewed from a quantitative perspective, with an emphasis on breaking down and decomposing charts into component wave forms.
- *The Spectrum of Reality:* A brief overview of the Fourier Transform and related decomposition of common signals into constituent wave forms.
- *A Wave to the Past:* Historical context and overview of a few of the more popular wave form theories over the last 150 years.
- *Attempting Spectral Analysis of Stocks:* Challenges of applying Fourier Transforms to stocks and the stock market and selected attempts to overcome these challenges.
- *Transforming the Alpha Beta Sigma Fraternity:* Extending Fourier waveform analysis to other time series charts - namely alpha, beta and sigma - with additional and counterintuitive insights.
- *Human Factors in the Waveform:* Factors that can push investment managers outside their core competency and established comfort zones within frequency, amplitude and phasing domains.
- *FrAP Mapping the Pros:* A simple yet powerful paradigm for categorizing investment managers based on the waveform framework of Frequency, Amplitude and Phasing (FrAP).

While there are many drawbacks to using Fourier Transforms in the "real world" of the stock market, a better understanding of the frequency domain is clearly beneficial as it flows through to a better understanding of periodic market discontinuities, a deeper appreciation for the need to specialize within an average holding period, and the knowledge that a simple two-dimensional time-series chart is not the whole story.

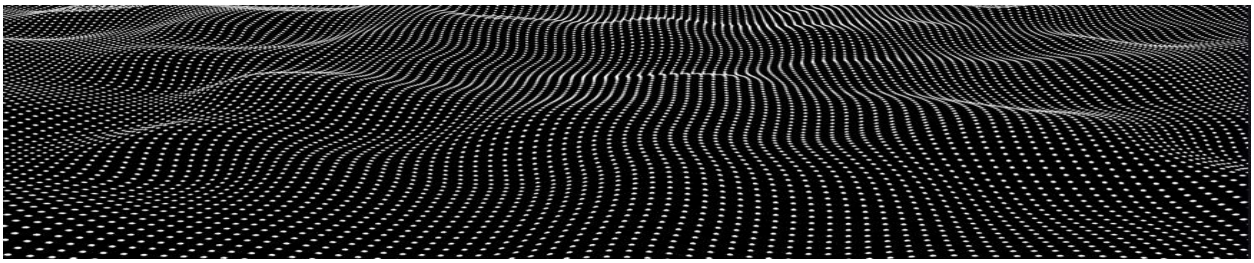

Don't Judge a Book by its Cover

It is difficult to not immediately cast judgment on a stock after seeing the stock chart. Yet as the overused legal disclosure says, "past results are no guarantee of future performance". A stock chart only graphs the market clearing price, imposing a form of survivorship bias to the historic record. Yet the market clearing price is not the whole story. There is an underlying cacophony of resonant and self-canceling investment behaviors associated with any stock and a chart of the simple market clearing price glosses over many of the true underlying signals and motivations. The daily market movement for any stock has activity from investors ranging across the spectrum from high frequency traders to passive index investors to activist hedge funds. Again, there is a lot going on under the surface.

A more complete analysis of a stock price history should look at the history of the full spectrum of transacting and non-transacting investors. While such data is generally not available, there are tools for more complete "spectral analysis" that break down variances in investment behavior by holding period, giving a better picture of why, for example, the day traders may be buying yet the weekly swing traders are selling. While these tools have limitations, they at least acknowledge that the simple price chart is an incomplete picture giving only the daily vote result without any disclosure of the closeness of the vote or which classes of votes went which way.

This is a little different than the usual criticism of chartists (a.k.a. technicians), a criticism typically leveled by fundamental analysts. That is, chart technicians tend to ignore important fundamental factors like valuation and the relative health of the actual business of a company. The traditional criticism is fair, but it ignores the fact that most technicians have short-term investment horizons. For a day trader, there is no point in understanding the details of a company business model and where earnings are headed over the next several quarters. A day trader will be long gone before the glacial moves of the fundamentals have any impact on the price of a stock. As a rule of thumb, if the average holding period is less than one calendar quarter, the analysis will likely be a chart-based technical analysis. If it is more than one quarter, the analysis will likely be fundamental.

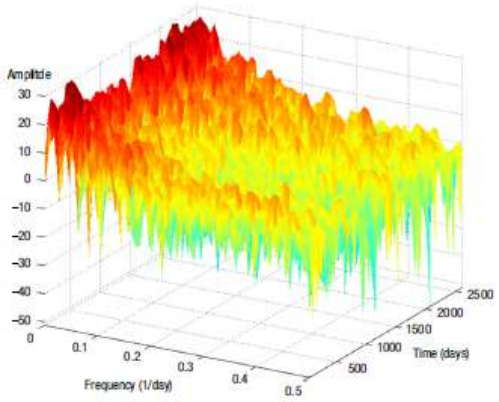

Spectrogram of Boeing stock, daily frequency

While most popular websites and books on technical analysis are aimed at visual pattern recognition, there is a more elevated level of non-fundamental analysis that is generally not accessible to the investment masses. This is the realm of quantitative analysis, an area where the physicists of Wall Street, steeped in mathematical expertise, apply their knowledge of partial differential equations and artificial intelligence to the stock market. Comparatively little is written on these strategies outside a few articles using various science fiction analogies and the occasional acknowledgement of amazing track records from the likes of Jim Simons PhD, the wealthiest of the quants with a net worth now exceeding

$20 billion. Spectral analysis is only a small part of the quantitative toolkit, but it seems to enhance appreciation for what might be hiding in plain sight.

Still, there is a siren song to the visually oriented technical analysis of stock charts. Of course, the same chart may yield different conclusions to different investors, perhaps another example that beauty is in the eye of the beholder. Yet the chart is still a touchstone that generates a quick reaction whether the investor is technical, fundamental or quantitative - a reaction that frequently occurs within the 40 milliseconds it takes for the human brain to process a visual input. The quickness of the visual technical assessment is a sharp contrast to fundamental analysis which might involve a mind-numbing dozens of hours *per company* to understand each company, the earnings model, sales segmentation, competitive environment, management competency, etc. So despite the imperfections of technical analysis, the intuitive visual appeal, ease of use and applicability across all companies means technical analysis will always be around, especially for investors with shorter term time horizons.

While fundamental analysis is clearly of limited use for short-term traders, this does not mean that technical analysis – flawed or not - is irrelevant for long-term investors. After all, stock charts are fractal in nature, meaning that a technical pattern such as a Golden Cross may occur in a stock chart looking at say the most recent 100 sample periods, but the sample periods could be as short as milliseconds or as long as years. The same pattern – a Golden Cross - can occur at any sampling frequency, short or long. This frequency independent formation of patterns means that at any given point in time, there is a noisy confusion of differing views that interact with each other in unpredictable ways. Indeed, it is possible that a simple daily chart will show a Golden Cross, but a weekly chart will show its complete opposite, a Death Cross.

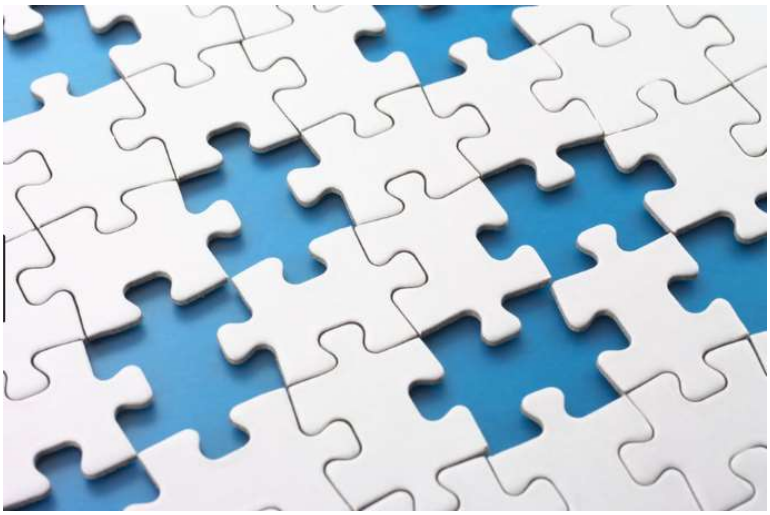

And so we are back to grasping for more in-depth spectral analysis of frequencies, knowing that a single chart of the price history is not the total history and that this whole is less than the sum of its parts.

The Spectrum of Reality

The Fourier Transform has been described as a mathematical prism that breaks up the light of a mathematical function into the rainbow of its constituent colors. There is an intuitive appeal to the Fourier Transform as an analytical tool to delve down into the details of the cycles that make up the stock market. Historically there was some degree of skepticism about the Fourier Transform. Indeed, when Fourier presented his landmark publication on heat transfer to the Paris Academy in 1807 along with the very bold claim that "any function could be represented by a summation of sines and cosines", it took another five years for the Paris Academy to award Fourier the Grand Prize. The Fourier Transform is a powerful tool.

But for investors today, it seems there is little in the way of plug and play tools to tap this power. A quick Google search will mostly show coding samples for quantitative algorithmic traders rather than graphical drag and drop tools for more casual users. The limited access for the masses partly reflects the upper-class mathematics of Fourier Transforms. But it also reflects the simple power of the tool, a tool that can be dangerous in the hands of the inexperienced.

Fourier Transform Equation

$$\dot{f}(\xi) = \int_{-\infty}^{\infty} f(x)\, e^{-2\pi i x}\ dx$$

Critics of Fourier Transforms highlight the difficulty of implementation in the noisy world of the stock market where wave patterns are subject to shifting phases, frequencies and amplitudes. The stock market is not an engineering laboratory for digital signal processing, an environment where the Fourier Transform thrives.

This is a challenge for rules-based algorithmic traders since implementation inevitably involves compromises in the form of filters, amplifiers and other processing to tease out the hidden patterns underlying the market. And if the quants are shying away, the rest of the investment world simply has no place for yet another tool, especially one that appears to be of marginal utility.

Challenges aside, seeing the constituent wave patterns of a stock or the stock market is a revelatory experience. It immediately calls into question the usefulness of even some of the simplest technical indicators such as moving averages, technical trend bands, support/resistance levels and more. It also sheds new light on other non-stock charts ranging from economic data series to portfolio alpha, beta and sigma. But investors need to delve a little deeper to see the full rainbow of colors within the Fourier paradigm and approach.

The Fourier Transform hinges on the idea that a chart or signal can be broken down into simpler constituent waves. The adjacent chart is a visual example of the combination of a low frequency sine wave with a higher frequency sine wave to arrive at a more complex wave pattern, a process known as superpositioning. Notably the Fourier Transform reverses this process by starting with the more complex combined wave pattern and decomposing it into its two simpler constituent wave patterns. That is, it gives us the "recipe" upon which any complex wave form is composed. The Fourier Transform can reveal many of the ingredients that are not apparent from a casual glance at the combined wave form. In a sense, if one were to pour a can of Coca-Cola into a Fourier Transform machine, the machine would reveal all the ingredients in Coke's 130-year-old, closely guarded secret recipe.

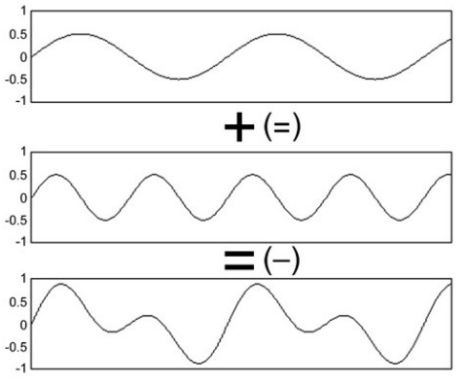

Real world signals are typically composed of far more than two simple sine waves. Most Fourier Transform software tools tend to highlight only the dominant frequencies, and still there could be a dozen or more such frequencies. Adjacent is another simple visual illustration using a hypothetical and more complex real world signal and the hypothetical output of its Fourier Transform with the dominant component wave forms. From a technical analysis perspective, this illustration is not particularly bullish or bearish when judging by the simple price action, even with the aid of various technical tools such as moving averages and momentum oscillators. And while the Fourier Transform has a rather complicated multiplicity of waveform outputs, the output still has more to go on from a forecasting perspective since the sine waves are repeatable. Further, the dominant wave frequencies hint at the types of investors that are moving the stock – day traders, long-term investors, high frequency traders and more.

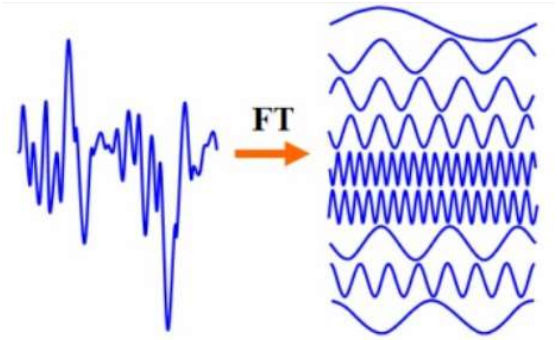

For those employed in digital signal processing (DSP) where waveform decomposition is a core skill set, there are certain common waveforms that are part of the daily DSP lexicon that have surprising relevance for equity investors. Consider the adjacent square wave reconstruction in red. Clearly the short-term traders focused on the high frequency squiggles are going to miss the giant low frequency gaps. This is the most common criticism of short-term investing, that most of the positive annual return of an investment occurs on just a handful of days of the year, days when the short-term traders are inevitably in cash and out of the market. Naturally, there is a parallel criticism of long-term investing, or for that matter, investing within any frequency band. While there are pitfalls to spectral analysis, seeing the full spectrum helps complete the puzzle of understanding.

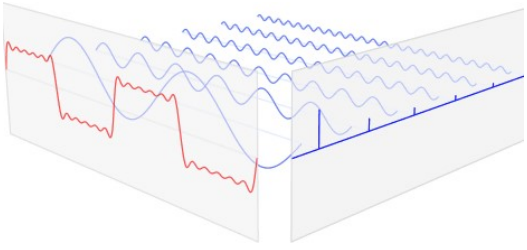

A Wave to the Past

A Wikipedia Talk website on Elliott Wave Theory includes an honest and amusing question from a contributor:" Isn't Elliott Wave Theory just a crappy Fourier Transform?" While there was no response online, the short answer is "Yes", Elliott Wave Theory is basically a restrictive and dumbed down version of Fourier Transforms. Popular wave theories have increased in sophistication over the last century beginning with rudimentary Dow Theory, followed by Elliott Wave Theory, and then Hurst Cycles. These theories aim to analyze stock price waves of varying frequencies using intuitive visual approaches and avoid the complexities of a more formal spectral analysis using Fourier Transforms.

Dow Theory is arguably the granddaddy of all wave theories, owing to its age and its provenance. While Charles Dow, the founder and first editor of the Wall Street Journal, never formally penned or promulgated his eponymous theory, Dow Theory was forensically stitched together by Dow's admirers after his death based on years of his Wall Street Journal editorials dating between 1851 and 1902.

From a wave theory perspective, the Dow Theory simply breaks down market movements into three categories: The *primary* trend, which is the long-term movement of prices typically lasting several months to several years; the *intermediate* or *secondary* trend, which is a shorter term deviation from the primary trend, deviations which are usually eliminated when a correction pushes prices back to the primary trend; the *minor* or *tertiary* trend, which is typically a daily fluctuation of lesser relevance.

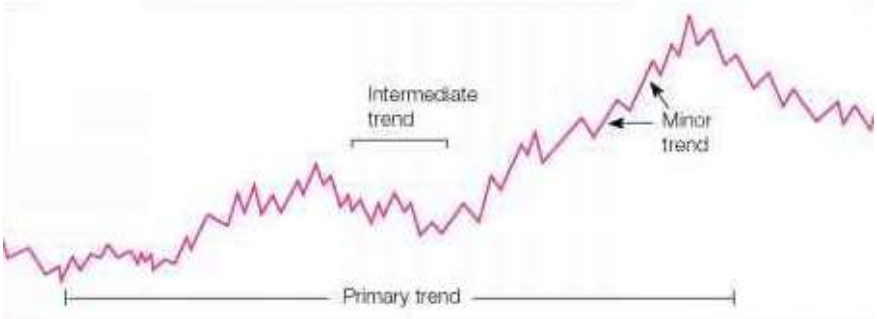

There are more aspects to Dow Theory beyond the wave forms. For example, Dow Theory does require agreement between the Dow Industrials and the Dow Transports. So a Bull Market is not a true Bull Market unless both the Industrials and the Transports are in Bull formations. Also, trends do need to be confirmed with volume. Price movements on very low volume cannot be relied upon as being representative of the whole market.

Dow Theory waves acknowledge the presence of the short-term, medium-term and long-term trends or waves, not unlike what a Fourier Transform might reveal. But Dow Theory does not acknowledge the superpositioning of waves that are revealed in a Fourier Transform, nor does it reveal wave patterns beyond the three primary, secondary and tertiary waves. Dow Theory tells us we have a can of Coke but does not reveal the secret recipe.

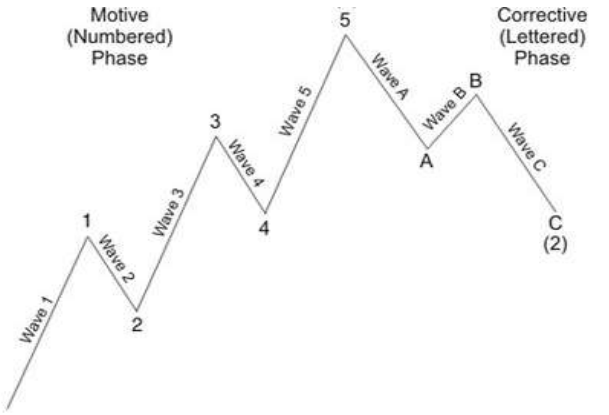

Elliott Wave Theory looks into more detail at any single trend – perhaps one of the Dow Theory primary, secondary or tertiary trends – and brings to light the persistence of a pattern of a five-wave progression (the Motive wave) being followed by a three-wave regression (the Corrective wave). The basic theory has been expanded upon since Ralph Nelson Elliott initially formulated the idea in the 1930s (e.g. Frost & Prechter, 1979).

One of the interesting features of Elliott Wave Theory is the idea that the five-wave and three-wave patterns can be found in a trend when looking at any time frame, whether looking at the last 365 days or the last 365 seconds. If you zoom in or zoom out, you will be able to spot the five-wave and three-wave patterns repeating at every level. The recurring pattern is a good visual example of a fractal, as discussed above. If you Google fractal, you will find a stunning variety of repeating artistic patterns, both in nature and computer-generated. Artistic eye candy aside, the real beauty of the fractals is the idea that stock cycles are not limited to the three Dow Theory frequencies of primary (months), secondary (weeks) and tertiary (days). In fact, there are cycles at every frequency ranging from milliseconds to decades.

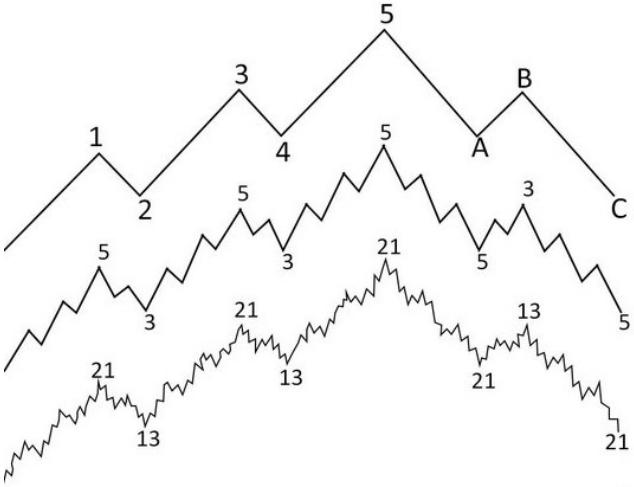

Like with Dow Theory, Elliott Wave Theory has additional rules beyond the top-level view of the five and three wave patterns, including understanding the variety of types of motive and corrective waves, projecting wave magnitude using Fibonacci ratios and more. Still, the acknowledgement of wave patterns at any frequency puts Elliott Wave Theory one step closer than Dow Theory to a Fourier Transform. That said, Elliott Wave Theory, like Dow Theory, does not acknowledge the superpositioning of wave patterns and the related interference patterns that can disrupt repetitiveness of Elliott Waves.

The Hurst Cycle Theory was arguably the first stock market wave theory to dip into the nerd world of engineering. JM Hurst was an aerospace engineer by training and published his Hurst Cycles strategy in the early 1970s. Hurst clearly acknowledged the superpositioning of wave forms and addressed items such as how short cycle waves can be swamped by movements of longer cycle waves.

At the heart of Hurst Cycle theory is the Nominal Model, a model which acknowledges the simultaneous presence of waves of varying amplitude but defined frequencies. That is, any individual stock simultaneously experiences cycles with frequencies of approximately 10 days, 20 days, 40 days, 80 days (10 weeks), 20 weeks, 40 weeks, 80 weeks, and so on. With a handful of simple visual tools, the Hurst analyst can lay out a roadmap for where the stock is headed.

Hurst identifies price targets and time frames by using a line chart of the stock, but advanced by one half of a cycle, known as the Future Line of Demarcation (FLD). When a stock price crosses up (or down) through its FLD, it implies the stock price has made it only halfway from the last trough (peak) to the next peak (trough) by both price and time. The target price and time is a seemingly elegant paean to the simple form of the sine wave, though perhaps a little less flexible than Elliott's simple Fibonacci approach.

Hurst charts can look busy. They commonly have multiple FLDs given the multiple cycles in the Nominal Model. Hurst charts also have specific Hurst-defined trend lines called Valid Trend Lines (VTLs). And there is a phasing analysis at the heart of Hurst Cycle theory. The adjacent chart is a representative Hurst analysis of the S&P 500 including approximate Nominal Phase notations, FLDs and VTLs.

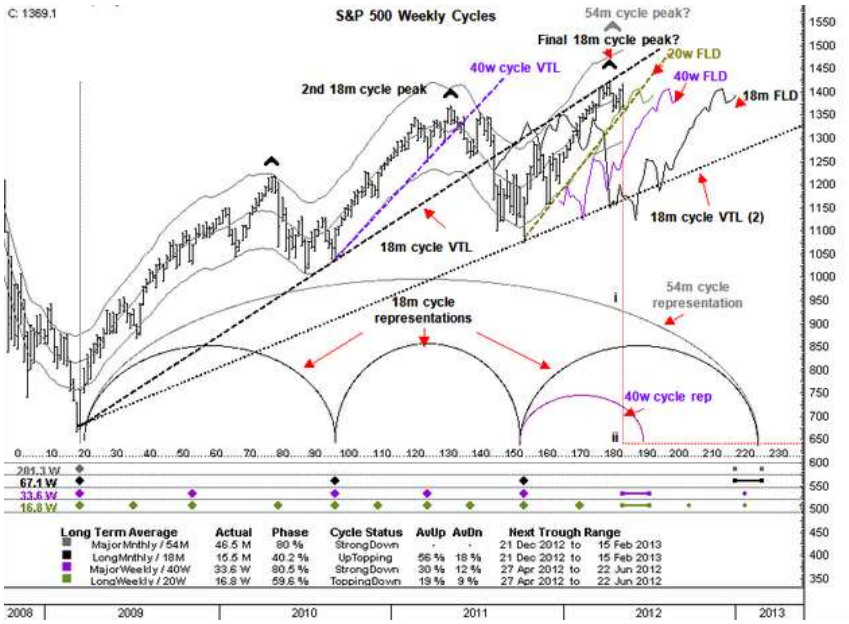

While the Hurst Cycle Theory seems to address many of the shortcomings of Elliott Wave Theory - in particular the acknowledgement of superpositioning - Hurst Theory still falls short of the mark. The visually-oriented certified technicians may be drawn to the mapping of the fluctuations of the many overlapping cycles. But the mathematical quant folks will see signs of sleight of hand. What is the rationale for the seemingly arbitrary frequency choices of the Nominal Model? Why does the phasing analysis seek out confirmation with price moves, when price moves can be disguised with superpositioning?

Hurst was aware of these and other shortcomings and indeed even suggested running a spectral analysis to either confirm the Nominal Model frequencies or shift the Nominal Model to truer and more dominant frequencies. A spectral analysis opens the door to another layer of analysis with the various Fourier Transform techniques.

Attempting Spectral Analysis for Stocks

One of the first things one notices when seeking out spectral analysis tools and techniques is the lack of simple tools for the masses. They are out there, but they are niche products. One will not find spectral analysis tools on a Bloomberg terminal, or on a Schwab Advisory platform, or in books with titles such as Spectral Analysis for Dummies or Fourier Transforms for Idiots. Instead you will find the tools where the quants gather – MATLAB, TradeStation and a variety of coding books aimed at algorithmic traders. While the math is not exactly quantum mechanics, it is advanced just enough to separate the quants from the mass market.

For do-it-yourself quant investors with time to spare, there is a free Fast Fourier Transform (FFT) tool built into Excel. It is accessed from the Data Analysis button that is under the Data tab. It is not crafted for stock market users but rather is a general Fast Fourier Transform tool used to decompose any wave form, whether the wave form arises from a stock price, electrical signals, tidal action, violin strings, cardiac rhythms, atmospheric changes and more.

There are quirks with the Excel FFT function. The input for the Excel FFT is simply the data for the stock in question, perhaps the daily closing price of IBM for the last year. Like any FFT tool, the number of input data points must be a multiple of 2^N, so a year of data would most closely approximate 256 days of data (or 2^8 days). Also, the Excel FFT output is in complex form - real plus imaginary components - so it needs a little scrubbing to be useful[1]. But the normalized output should then simply range between 0 and 1 for each frequency bar, with higher output ranges representing "dominant" frequencies. When running the Excel FFT frequency response, there should be some overlap with the Hurst Phasing analysis, especially for those Hurst charts where there is little ambiguity.

The adjacent chart on US Steel – borrowed from the sales literature of the MetaStock FFT tool – shows what ought to be a representative FFT output. While the scaling runs 0 to 1000 instead of the Excel scale of 0 to 1, the output clearly shows a dominant cycle of 205 days and two smaller sub-cycles of 39 days and 27 days.

This works well as a visual demonstration for the MetaStock salesforce since the FFT dominant cycle of 205 days happens pair up so well with a chart that basically looks like a sine wave with a period of just over 200 days.

[1] Specifically, the complex output needs to be converted to a simple vector modulus with the Excel IMABS function. And then the modulus for each frequency should be normalized by the total bar count, 256 in this case

While MetaStock may be a tad guilty of cherry picking a chart that looks an awful lot like a sine wave, the visual is still instructive by at least showing that the FFT tool sees the same sine wave that we all see. By implication, if it can see the obvious sine wave, perhaps it can also find waves hidden by superpositioning and waves of higher frequency harmonics – the things we cannot see.

The efficacy on these more difficult real-world charts is mixed at best. The May 2016 issue of the Journal of Economics and Management published an in-depth article from the University of Economics at Prague on *Fourier Analysis for Stock Price Forecasting* which bleakly concludes that the methodology is a failure. While the cycles were indeed uncovered, the ability to profit from the movement of the cycles was not realized. Similarly, the Hurst Cycles literature generally embraces Fourier spectral analysis but is quite cautious about getting lost in the waveform details due to a "triple infinity" of possible components – amplitude, frequency and phase.

The Fourier Transform seems to wrestle most with the shiftiness and discontinuities of the stock market. Fourier Transforms work best breaking down the fundamental wave form and harmonics of a single unmoving wave, say a sustained note on a violin string. They do not cope as well with a moving melody and less so with the melody and harmony from the full symphony orchestra of the overall stock market. Despite all these red flags, engineers and programmers are still drawn to the utility of Fourier Transforms, a Rosetta Stone of wave form analysis.

One of the more serious attempts to force fit Fourier spectral analysis into stock market analysis is John Ehler's book *Rocket Science for Traders (Wiley, 2001)* which details his approach, Maximum Entropy Spectral Analysis (MESA). The MESA approach uses a low latency pattern-matching filter, while restricting the output to only one dominant wave form and over a short period of time (to minimize risk of market discontinuities). The restrictions are a bit of a straightjacket, but the output does carry more weight.

The adjacent chart is a graph of Microsoft stock in the top panel accompanied by MESA's characteristic Corona chart which highlights the dominant cycle over the period. Note that in contrast to the MetaStock output, the MESA output gives only one dominant cycle and the dominant cycle is not fixed but rather is constantly moving. Because of the dynamic nature of the dominant cycle, trading signals are more conducive to walk forward optimization rather than back testing.

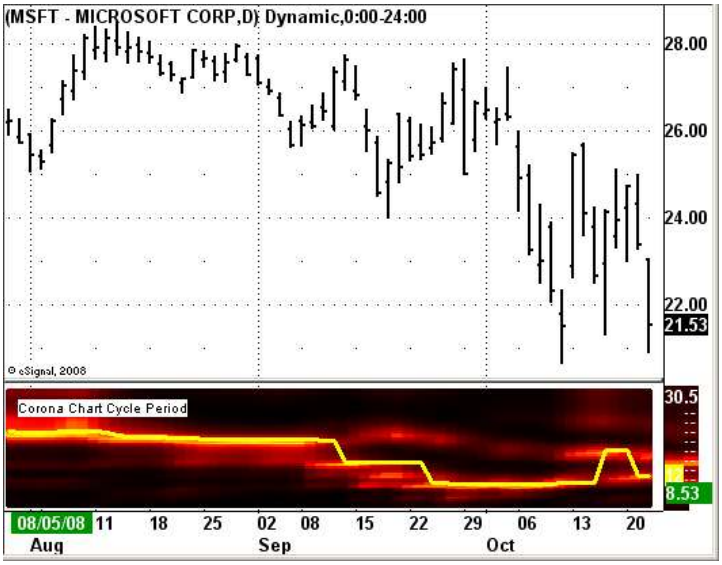

While the straightjacket restrictions on the MESA approach improve the utility of the output for traders, it does so at the expense of utility for longer term investors. The short time frame restriction was in theory motivated by fear of market discontinuity since discontinuities are akin to pushing the reset button on any waveform analysis. Second, Ehlers notes that lengthening the time frame draws attention to patterns that are more easily spotted by traders and hence patterns that simply will not last. Both of these assertions are problematic.

If the stock market is indeed a fractal, a market discontinuity to a day trader is simply white noise to long-term investors like Warren Buffett. Same story with the trader pattern recognition, patterns that Warren Buffett will neither see nor care about (and vice versa – the day trader has no use for a ten-year pattern). The tendency to use the short-term as a proving ground while either ignoring the long-term or fabricating bogus reasons to avoid the long-term is common, especially among quantitatively inclined academics. Even the likes of Eugene Fama and Andrew Lo have occasionally stepped over the line a tad here. Ironically this is the very same crowd that understands via Fourier Transform that long-term results are not simply the sum of short-term results. Perhaps the fixation of quantitative academics on short-term results reflects the fact that proving long-term inefficiencies might take decades to prove out, a time frame that is simply too long for those academics angling for the tenure track.

Transforming the Alpha Beta Sigma Fraternity

Since the Fourier Transform can be used on any mathematical function or signal, the application goes beyond the idea of examining the primary frequencies of a particular stock. Academics have advanced the idea of looking at the whole market to get a feel for whether the markets are driven by HFTs or by long-term Buffett-type investors. And why stop at stock charts? Why not look at GDP to get a sense of the short-term and long-term cycles that move GDP? Or interest rates, unemployment, consumer confidence, etc.? While there is a whole world to be transformed, equity investors may be especially interested in transforming measures of outperformance (alpha), correlation (beta) and volatility (sigma).

Investors will likely begin here by looking at realized alpha over various holding periods (a.k.a. frequency domains) to understand when managers with monthly return strategies will do well versus managers with daily return strategies. Hopefully this will lead to some understanding of how to combine disparate holding period managers without too much return correlation between the managers. In basic Spectral Portfolio Theory, it is common to use a dynamic measure of alpha that addresses forecasting power across such varying time horizons.

The dynamic alpha acknowledges that active portfolio management is predominantly about choosing portfolio weights. It is measured as the covariance of the portfolio weights with the portfolio position performances. So if the big, overweighted positions are quickly moving higher, the covariance, or dynamic alpha, will move higher as well (and vice versa).

A spectral decomposition of the dynamic alpha will reveal at which frequencies the weights and positions are in phase (positive alpha) or out of phase (negative alpha). While the tools and data feeds to decompose dynamic alpha are limited and not very user friendly, there is comparatively less ambiguity in final decompositions. Like with any report card, the decomposition is clear at explaining past alpha, a key measure for all investors, but it is less useful for projecting forward alpha.

The Fourier Transform can also be applied to betas, a fruitful exercise for anyone who has ever used an automatic beta calculator. Bloomberg was one of the first widely available services to offer a beta calculator 25-plus years ago. Just type in the tickers and the time frame and voila, there is the beta. Alas, the calculated beta never seems to square with other published sources such as S&P or Value Line. By altering the time frame or sample frequency, the beta calculator might yield widely varying results. Many casual users of such beta calculators throw in the towel here when learning it cannot dependably replicate other published sources.

The problem is not with the beta calculators. The problem is with the other published sources and even the AIMR (the predecessor to the CFA Institute) for promulgating the idea that each stock has a beta that is constant and unchanging and applicable to all investor types from day traders to buy & hold Warren Buffett wannabes. By manually sampling betas across time frames and sample frequencies, users of beta calculators are essentially reconstructing a Fourier Transform of the beta, cranking it out the hard way. It would be more elegant and straightforward to start with the Fourier Transform and let the transform lay out the resonant frequencies where betas peak and trough.

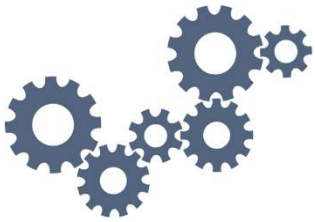

Beta itself can be a hot button topic given the ability of company-specific factors to temporarily alter correlations. With beta, past performance is no guarantee of future results. Still it is a useful concept that conveys much in the way of inference. Mention high beta and investors immediately think of recent IPOs, small caps and story stocks with high valuations. With low beta, investors think of blue chips, defensive utilities and staples. The archetypal categorization of high and low beta is useful at a theoretical level, but the real world is more subtle than simple black and white.

It is useful to know that beta varies with the holding period or frequency. As an example, during the 2000 dot com bubble bursting, two Seattle hedge funds “cleverly” sought to hide from the onslaught of selling in dot com tech names yet still stay invested as required by mandate. The idea was to seek out names that were not focused in the main areas of selling such as the tech sector and larger names that were in the indexes and hence subject to passive index selling. So the funds bought scads of microcap banks, industrials and service companies.

But as the selling intensified over the year, even the non-index microcaps eventually unwound and both funds saw losses in excess of 60%. Ouch! From the perspective of the hedge funds, the daily or weekly beta over the prior several years was close to zero, an attractive feature in a declining market. But using a quarterly or annual beta over the last twenty-plus years would have put the beta well above the market beta, a frighteningly bullish bet in a deep bear market. A Fourier Transform of the beta would have revealed the high peak at low frequencies, though that does assume the microcap data is actually available. A Bloomberg beta calculator also would have revealed the variability of the beta, if the user chose to calculate betas at all frequencies - high, medium and low.

Sigma (volatility) also varies by holding period. This can be evidenced by performing a Fourier Transform on market volatility using an indicator such as the VIX to see which frequencies (holding periods) have the dominant volatility signals. Over the last 50 years, short-term day traders and swing traders have seen about 10%-20% less volatility than investors with holding periods over one month (Chaudhuri, Lo 2016).

This observation is clearly counterintuitive and indeed flies in the face of the patient tortoise vs hare marketing message of the wealth management industry. Notably, this kinder and gentler environment for the day trading hares appears to be flagging somewhat in recent years, perhaps due to competition and technology or perhaps due to simple mean reversion. So maybe over the next 50 years, the tortoise will eventually overcome the hare.

Day traders and buy & hold investors clearly experience differing volatilities, differing betas and differing alphas for the exact same portfolio of stocks. Of course, Morningstar, eVestment and peer investment performance databases have not really blessed disclosure by frequency domain. But for the professional, it is still useful to understand if your outperformance was due to a decline in volatility within your frequency domain, a lift in the beta in your frequency domain, or true portfolio alpha.

Human Factors in the Waveform

While investment managers typically specialize in certain frequency bands as well as other waveform factors such as amplitude and phasing, there can be problematic leakages into adjacent bands**.** Such leakages outside the core competencies are usually driven by human behavioral factors.

An amusing example here of frequency band leakage is what commonly happens when short-term trades go awry. Instead of booking the loss and moving on, the investor reclassifies the short-term trade as a long-term investment. The short-term loss then comes to be regarded as mere noise on the way to longer term gains. This moving of the goalposts is a whitewashing behavior and is especially prevalent with portfolio managers that have investors/customers that can see the actual trading, a common arrangement for separate account portfolios (as opposed to pooled portfolios such as mutual funds and hedge funds). Clearly, no one wants to see a loss, or worse let the customer see a loss.

A more egregious example of frequency band leakage comes from Family Office settings and tightly controlled investment firms and can be challenging for investment team employees. Most single-family offices with more than $50 million in assets are set up from a governance perspective as something akin to a benevolent despot and not an ESG-compliant board with broad internal and external representation. Unfortunately, not all despots are benevolent and an out-of-favor investment team employee is likely to find the goalposts moved to whatever time frame is most disadvantageous. This is a good example of the Golden Rule, "He who has the gold makes the rules". This moving of the goalposts here is a blackballing behavior that seems more prevalent in family offices and tightly controlled investment firms where turnover is above average.

A more honest approach is to identify up front the average holding period of the investment style from which the alpha is to be derived and stick to it. The thinking here is that the investment paradigm that Warren Buffett uses to identify investments for his 10+ year holding period has essentially no overlap at all with the investment paradigm used by a high frequency trader with an average holding period measured in milliseconds. Investors will then build up an expertise on the factors that add the most alpha within the paradigm of their target holding period. Shifting between holding period paradigms should then be flagged as outside the core competency.

Investor/Investor Type	Average Holding Period
Warren Buffett	10+ years
Pension Funds, US Average	17 months
Mutual Funds, US Average	16 months
Investors' Business Daily IBD50	2-4 months
Swing Traders	1 week
Day Traders	1 day
Scalp Traders	<1 day
High Frequency Traders	<1 second

Building the core competency within a frequency zone takes time. It is a common observation that rookie analysts do not immediately and intuitively separate low frequency catalysts from high frequency catalysts. For example, if GM stock gaps down 10% one day, the rookie analyst will not separate lower frequency headwinds - such as market shares losses to foreign competition (an incorrect short-term explanation) – from higher frequency headwinds – such as the earnings shortfall announced just before the market open (the correct short-term explanation). Amid dozens of GM headwinds and tailwinds blowing across all frequencies, the "frequency deafness" of a new rookie analyst takes time, training and practice to overcome.

Just like with the frequency domain, professional investors also tend to specialize in a particular amplitude domain. Amplitude variances are essentially captured in the beta. And while beta can be a squirrely number varying over frequency as discussed above, it is still intuitive to aggregate highly reactive stocks (or highly unreactive stocks) within any frequency domain and build a paradigm of expertise based on the amplitude domain. High beta stocks just don't trade like low beta stocks. They are owned by different villages of investors with different value systems and it takes time to build the tool set to operate successfully within each amplitude domain. Human fear or greed usually drives the leakage into adjacent amplitude domains. This is not necessarily a bad thing if the fear/greed spider sense proves accurate about an impending recession/boom, and if the investment manager has a mandate that allows switching between offense and defense. Such a flexible mandate is common with small wealth managers but relatively uncommon with most mutual funds.

New analysts-in-training are also subject to "amplitude deafness", a condition usually characterized by overdependence on Bloomberg beta calculators rather than an informed and holistic understanding of beta and where beta is headed. Just because XYZ is a tech stock with a high Bloomberg beta does not mean it will have a high beta in the next holding period. The subtleties take time and practice to learn.

The phasing domain is also subject to human fear or greed. Phasing variances can be captured within the spectrum of contrarian versus momentum thinking. While most investors try to buy at the bottom of the sine wave and sell at the top of the sine wave, in the real world, there are clearly two camps - contrarians who believe they can call the bottom and momentum investors who wait for some sign of a bottom forming. Contrarians tend to buy a little before the bottom and momentum investors buy a little after the bottom. And, like developing expertise in a particular frequency or amplitude domain, investors will also build up expertise within a particular phasing domain, developing alpha despite the contrarian lead or the momentum lag. Phase domain leakage seems most prevalent when things are persistently working better for the "other team", an understandable and human frustration.

FrAP Mapping the Pros

By classifying investment managers into where they stand on each spectrum – Frequency, Amplitude and Phasing – a simple yet powerful visual tool, or FrAP Map, can be constructed. A FrAP map captures the essence of the waveform approach an investment manager is using to generate alpha: Frequency captures the average holding period; amplitude captures the average beta; and phasing captures the degree of contrarian/lead versus momentum/lag.

As seen in the maps below, Warren Buffett is clearly very low frequency, medium low amplitude and lead phasing. Jim Simons is very high frequency, medium high amplitude and lead phasing. Alpha Architects (a momentum-based investment firm in Chicago) is medium frequency, medium amplitude and lag phasing. Primecap (a long-term high-growth shop in Pasadena) is very low frequency, very high amplitude and phase neutral.

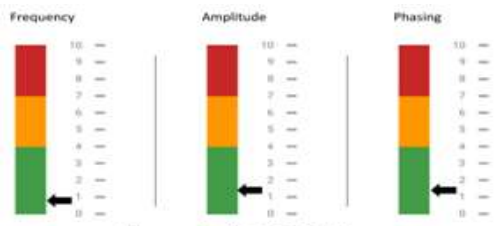

Warren Buffett FrAP Map

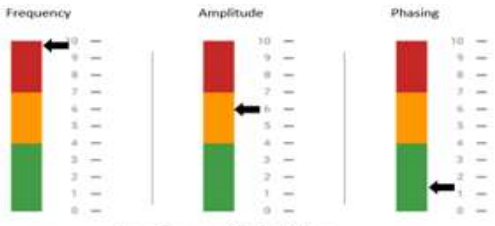

Jim Simons FrAP Map

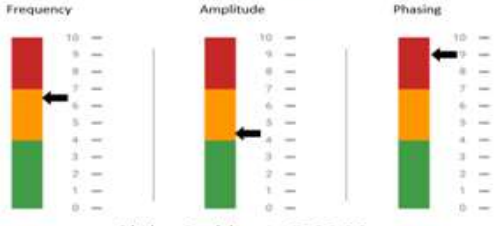

Alpha Architects FrAP Map

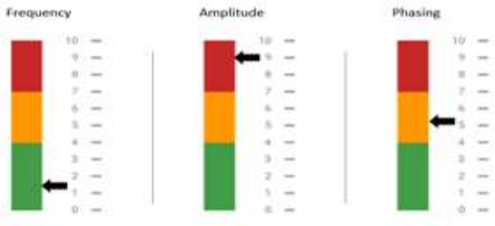

Primecap FrAP Map

Notably, there are a handful of fund choices for passive investors as well. With the rise of factor-based investing, some of the more popular ETFs are based on amplitude and phasing domains. Indeed, low amplitude has been one of the best performing factors in 2019, though it generally goes under a different name – Minimum Volatility (ticker: USMV). Its mirror image opposite, the high amplitude ETF has not fared so well this year – S&P High Beta (ticker: SPHB). Phasing ETFs also exist such as the Momentum Factor ETF (ticker: MTUM), though contrarian/reversion phasing ETFs are less prevalent.

The offering of active and passive investment managers seems to cover quite a few of the combinations of frequency, amplitude and phasing. But in reality, it barely scratches the surface. While the Minimum Volatility ETF has been quite popular this year, the fact is that the ETF industry is offering a sought-after factor (low amplitude), but in only one frequency domain and in one phasing domain. Yet the fund universe could easily be split into 10 frequency domains (decade, year, quarter, month, week, day, hour, second, millisecond, nanosecond), 6 amplitude domains based on beta ranges, and 6 phasing domains based on lead/lag spectrum. Such as breakout would total 360 investment funds to give a full menu of choices. The current passive ETF total of 3 or 4 is well short of the mark. Active managers are needed to fill the ETF shortcoming. The biggest shortcoming is in the frequency domain offerings, a niche where active manager choices are limited and passive ETF choices are non-existent.

While the FrAP Map paradigm may have been born in the arcane lingo of the digital signal processing laboratory, it clearly puts an elegantly simple and powerful label on investment styles. If there were such a thing as baseball trading cards featuring famous investors, a FrAP Map on the back of each card would be the sports equivalent of the league, the team and the position played.

Merci Beaucoup Monsieur Fourier

Catching a good wave to ride home is a common objective for both the short-term technical trader and the long-term fundamental investor. The short-term technician watches wave patterns to spot that epic wave for the best ride that day. The long-term fundamental investor watches the moon for evidence of high tide. Both approaches have merit within their respective frequency domains.

This paper covers challenges with traditional stock charts, the applicability of spectral analysis of stock charts, and the extension of Fourier techniques and frequency domain thinking within a broader investment paradigm. While there is much food for thought in the above discussion, there is a clear acknowledgement of the challenges of both technical analysis and spectral analysis:

- *Stock charts alone are not a reliable means for forecasting:* The chart is neither all-inclusive of the information for price discovery nor representative of all short-term and long-term investor motivations. Tools such as moving averages, supports & channels are of marginal utility at best.
- *Spectral analysis of individual stocks is good in theory but very difficult in practice:* A spectral analysis of stock charts using Fourier Transforms holds promise addressing some of the key shortcomings of technical analysis, but implementation is limited by shifting frequencies, amplitudes and phasing.

Besides stocks, other time series can be transformed such as economic data or interest rates. Thinking through spectral transforms on alpha, beta and sigma is especially instructive for equity investors.

- *Spectral analysis of alpha can reveal whether long holding periods surpass short holding periods:* While the data may be difficult to source and process, a dynamic alpha decomposition will yield results, giving investors a report card broken down by frequency.
- *Spectral analysis of betas can reveal beta drift and reversals for differing holding periods:* Again, the data is not easy to source nor process; but, one can intuit the transform with automatic beta calculators over varying holding periods.
- *Spectral analysis of market volatility, or sigma, shows that higher frequency trading is actually lower risk:* This counterintuitive result is based on a Fourier Transform of the VIX volatility index over the last 50 years, a result that could possibly reverse in the future.

Investors typically specialize within certain frequency domains, amplitude domains and phasing domains. But human factors can push managers outside their core competency.

- *Problematic frequency domain shifting can arise* from whitewashing bad trades, blackballing investment team members or frequency deafness of rookie investors.
- *Amplitude and Phase domain shifting can arise* from greed/fear moves which may or may not align with mandates. Amplitude deafness can arise from overdependence on beta data.

Classifying investors based on the three core waveform variables – frequency, amplitude and phasing – provides a simple yet powerful paradigm to classify and fingerprint investment managers.

- *The FrAP Map paradigm provides an easy visual mapping of the investment management world based on Frequency, Amplitude and Phasing, the three variables of waveform analysis.* While inspired by the arcane mathematics of Digital Signal Processing, the classification system has instant intuitive appeal.