

Spectral Portfolio Theory^{*}

Shomesh E. Chaudhuri[†] and Andrew W. Lo[‡]

This Draft: June 2, 2016

Abstract

Economic shocks can have diverse effects on financial market dynamics at different time horizons, yet traditional portfolio management tools do not distinguish between short- and long-term components in alpha, beta, and covariance estimators. In this paper, we apply spectral analysis techniques to quantify stock-return dynamics across multiple time horizons. Using the Fourier transform, we decompose asset-return variances, correlations, alphas, and betas into distinct frequency components. These decompositions allow us to identify the relative importance of specific time horizons in determining each of these quantities, as well as to construct mean-variance-frequency optimal portfolios. Our approach can be applied to any portfolio, and is particularly useful for comparing the forecast power of multiple investment strategies. We provide several numerical and empirical examples to illustrate the practical relevance of these techniques.

Keywords: Portfolio Theory; Portfolio Optimization; Spectral Analysis; Cycles; Active/Passive Decomposition.

JEL Classification: G11, G12, C32, E32

^{*}We thank Tom Brennan, Hui Chen, Paul Mende, and participants at the 2015 IEEE Signal Processing and Signal Processing Education Workshop, the MIT Sloan Finance Lunch, and the London Quantitative Finance Seminar at Imperial College for helpful comments and discussion. The views and opinions expressed in this article are those of the authors only, and do not necessarily represent the views and opinions of any institution or agency, any of their affiliates or employees, or any of the individuals acknowledged above. Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged.

[†]Department of Electrical Engineering and Computer Science, and Laboratory for Financial Engineering, MIT

[‡]Charles E. and Susan T. Harris Professor, MIT Sloan School of Management; director, MIT Laboratory for Financial Engineering; Principal Investigator, MIT Computer Science and Artificial Intelligence Laboratory. Corresponding author: Andrew W. Lo, MIT Sloan School of Management, 100 Main Street, E62-618, Cambridge, MA 02142-1347, a1o-admin@mit.edu (email).

Contents

1	Introduction	1
2	Literature Review	4
3	Spectral Analysis	6
3.1	A Brief History	7
3.2	The Fourier Transform	8
3.3	The Power Spectrum	11
3.4	The Business Cycle	14
4	Spectral Portfolio Theory	16
4.1	Volatility	16
4.2	Correlation and Beta	18
4.3	Alpha, Tracking Error, and Information Ratios	21
4.4	Numerical Examples	24
4.5	Implications for Portfolio Optimization	29
5	An Empirical Example	32
5.1	Mean Reversion	32
5.2	Momentum	34
6	Conclusion	35
A	Appendix	36
A.1	General Moment Properties of the Power Spectrum	36
A.2	Confidence Intervals for Dynamic Correlations	37
A.3	Standard Errors and F -Tests for Dynamic Betas	37

1 Introduction

Although portfolio optimization models have explicitly incorporated a time dimension ever since the stochastic dynamic programming approach of Samuelson (1969) and Merton (1969, 1971, 1973), the decision-making horizon of investors has rarely been the main focus of attention. Portfolio weights are assumed either to be rebalanced continuously over time or at arbitrary but fixed discrete intervals. In both cases, the process by which portfolio decisions are rendered is determined by dynamic optimization, yielding optimal portfolio weights that are functions of state variables evolving through time according to their laws of motion. The generality of this approach can obscure important features of the underlying process by which information is reflected in investment decisions. For example, although high-frequency trading and long-term investing can both be profitable—and both can be modeled as a dynamic optimization problem—they operate at very different frequencies using very different methods.

In this paper, we propose a new approach to analyzing and constructing portfolios in which the frequency component is explicitly captured. Using the tools of spectral analysis—the decomposition of time series into a sum of periodic functions like sines and cosines—we show that investment strategies can differ significantly in the frequencies with which their expected returns and volatility are generated. Slower-moving strategies will exhibit more “power” at the lower frequencies while faster-moving strategies will exhibit more power at the higher frequencies. By identifying the particular frequencies that are responsible for a given strategy’s expected returns and volatility, an investor will have an additional dimension with which to manage the risk/reward profile of his portfolio.

In fact, because time-domain statistics such as means, standard deviations, correlations, and beta coefficients all have frequency-domain counterparts, it is possible to apply spectral analysis to virtually all aspects of portfolio theory, linear factor models, performance and risk attribution, capital budgeting, and risk management. In each of these areas, we can decompose traditional time-series measures into the sum of frequency-specific subcomponents. For example, for a specific set of historical asset returns, we can decompose the correlation matrix into the sum of high, medium, and low frequency correlations so that an investor can determine the best and worst sources of diversification and change his portfolio accordingly.

Specifically, if the high-frequency correlations are very large, this might suggest changing the composition of the portfolio to include slower-moving longer-term assets or strategies.

To motivate the practical relevance of frequency in the portfolio context, consider the simple market-neutral mean-reversion strategy of Lo and MacKinlay (1990). This strategy holds long positions in stocks that underperformed the average stock q days ago and holds short positions in stocks that outperformed the average q days ago, i.e., $w_{it}(q) = -(r_{it-q} - r_{mt-q})/N$ where $r_{mt-q} = \sum_i r_{it-q}/N$ is the average stock return on date $t-q$ and N is the total number of stocks. Each lag q defines a different strategy, one intended to exploit mean reversion over a q -day horizon. Common intuition might suggest that the returns of strategies with similar horizons would be highly correlated. Therefore, it may be surprising to learn that the correlation between the returns of the $q=1$ and $q=2$ strategies is -1.1% over the period from 16 July 1962 to 31 December 2015.¹ However, during the month of August 2007, these strategies all suffered significant losses as part of the “Quant Meltdown” (Khandani and Lo, 2007). During that month, the correlation between the $q=1$ and $q=2$ strategies spiked to 65.8% . Figure 1 provides a more dynamic view of this correlation, computed over 125-day rolling windows from 3 January 2006 through 31 December 2008. The correlation began increasing in July 2007 but the spike occurred in August 2007 and declined steadily until the correlation turned negative in the first half of 2008, only to reverse itself during the second half as the Financial Crisis unfolded.

These strange dynamics illustrate the relevance of frequency effects in financial asset returns, and spectral analysis is the most natural tool for quantifying these effects.

We begin in Section 2 with a literature review and then provide a brief introduction to spectral analysis for non-specialists in Section 3. Our main results are contained in Section 4 where we provide spectral decompositions for portfolio expected returns, volatility, correlation, and beta coefficients, and describe how to use them to construct frequency-optimal portfolios. We provide an empirical illustration of these techniques in Section 5 and conclude in Section 6.

¹Specifically, the strategies are implemented using data from the University of Chicago’s Center for Research in Securities Prices (CRSP). Only U.S. common stocks (CRSP share code 10 and 11) are included, which eliminates REIT’s, ADR’s, and other types of securities, and we drop stocks with share prices below \$5 and above \$2,000.




Figure 1: 125-day rolling-window correlation between daily mean-reversion strategies $\{w_{it}(q)\}$ with $q=1$ and $q=2$ where $w_{it}(q) = -(r_{it-q} - r_{mt-q})/N$ and $r_{mt-q} = \sum_i r_{it-q}/N$ is the average stock return on date $t-q$. The gray lines delineate 2-standard-deviation bands around the correlations under the null hypothesis of zero correlation.

2 Literature Review

The frequency domain has long been part of economics (Granger and Hatanaka, 1964; Engle, 1974; Granger and Engle, 1983; Hasbrouck and Sofianos, 1993), and spectral theory has also been used in finance to derive theoretical pricing models for derivative securities (Linetsky, 2002; Linetsky, 2004a; Linetsky, 2004b; Linetsky, 2008). However, econometric and empirical applications of spectral analysis have been less popular in economics and finance, in part because economic time series are rarely considered stationary. However, there has been a recent rebirth of interest in economic applications in response to modern advances in non-stationary signal analysis (Baxter and King, 1999; Carr and Madan, 1999; Croux, Forni, and Reichlin, 2001; Ramsey, 2002; Crowley, 2007; Huang, Wu, Qu, Long, Shen, and Zhang, 2003; Breitung and Candelon, 2006; Rua, 2010; Rua, 2012). This rebirth motivates our interest in the spectral properties of financial asset returns.

Spectral and co-spectral power, often calculated using either the Fourier or wavelet transform, provide a natural way to study the cyclical components of variance and covariance, two important measures of risk in the financial domain. Specifically, spectral power decomposes the variability of a time series resulting from fluctuations at a specific frequency, while co-spectral power decomposes the covariance between two real-valued time series, and measures the tendency for them to move together over specific time horizons. When the signals are in phase at a given frequency (i.e., their peaks and valleys coincide), the co-spectral power is positive at that frequency, and when they are out of phase, it is negative.

In a recent empirical study, Chaudhuri and Lo (2015) perform a spectral decomposition of the U.S. stock market and individual common stock returns over time. They noticed that measures related to risk and co-movement varied not only across time, but also across frequencies over time. Such changes were especially apparent throughout the 1990s during the advent and proliferation of electronic trading. Studying this connection between technology and market dynamics has become especially important as recent events, including the Flash Crash of 2010, have led many to question the negative impact electronic trading could have on markets. Only by understanding the sources of feedback among these automated trading programs will we be able to construct robust portfolios and implement well-designed policies and algorithms to manage risk. Moreover, identifying asset-return harmonics may have

important implications for measuring and managing systematic risk.

Our framework also suggests that investors may benefit by diversifying not only across assets, but also across strategies and securities with different trading harmonics. Along these lines, Chaudhuri and Lo (2015) develop a band-limited mean-variance optimization, which becomes particularly useful when portfolio goals differ across time horizons, and when investors wish to target specific horizons because of their preferences and life cycle. Their framework utilizes frequency specific measures of correlation and beta, introduced to the economic literature by Croux, Forni, and Reichlin (2001) and Engle (1974), respectively. In this article, we show how these statistics can be calculated using the DFT, and demonstrate their usefulness in financial applications. Specifically, as the frequency band-limited counterparts to correlation and beta, they can be applied to almost any theory of risk, reward, and portfolio construction.

In addition to improving passive investment strategies, spectral analysis can also be used to characterize and refine active strategies. The standard tools used for performance attribution originate from the Capital Asset Pricing Model of Sharpe (1964) and Lintner (1965). The difference between an investment’s expected return and the risk-adjusted value predicted by the CAPM is referred to as “alpha”, and Treynor (1965), Sharpe (1966), and Jensen (1968, 1969) applied this measure to quantify the value-added of mutual-fund managers. Since then a number of related measures have been developed including the Sharpe, Treynor, and information ratios. However, none of these measures explicitly depend on the relative timing of portfolio weights and returns in gauging investment skill.

In contrast, Lo (2008) proposed a novel measure of active management—the active/passive (AP) decomposition—that quantified the predictive power of an investment process by decomposing the expected portfolio return into the covariance between the underlying security weights and returns (the active component) and the product of the average weights and average returns (the passive component). In this context a successful portfolio manager is one whose decisions induce a positive correlation between portfolio weights and returns. Since portfolio weights are a function of a manager’s decision process and proprietary information, positive correlation is a direct indication of forecast power and, consequently, investment skill.

In this article, as an extension of this AP decomposition, we introduce the frequency

(F) decomposition, which uses spectral analysis to measure the forecast power of a portfolio manager across multiple time horizons. An investment process is said to be profitable at a given frequency if there is positive correlation between portfolio weights and returns at that frequency. When aggregated across frequencies, the F decomposition is equivalent to the AP decomposition, and therefore provides a clear indication of a manager’s forecast power across time horizons. This connects spectral analysis to the standard tools of modern portfolio theory, and allows us to study the time horizon properties of performance attribution.

To address the non-stationarity of financial time series, our analysis relies on the short-time Fourier transform, which applies the discrete Fourier transform (DFT) to windowed sub-samples of the entire sample (Oppenheim and Schaffer, 2009). Recently, wavelets (Ramsey, 2002; Crowley, 2007; Rua, 2010; Rua, 2012) and other transforms (Huang, Wu, Qu, Long, Shen, and Zhang, 2003) have also been used to study financial data in the time-frequency domain, and depending on the specific context, these alternative techniques can provide substantial benefits in terms of implementation. For example, the sinusoids used in the short-time Fourier transform do not efficiently characterize discontinuous processes, whereas the flexibility of wavelets can be used to overcome this difficulty. Moreover, the wavelet transform provides better time resolution at high frequencies, and better frequency resolution at low frequencies, although similar results can be obtained by varying the window length used with the short-time Fourier transform. However, in this article, we refrain from using the wavelet transform for two reasons: the Fourier transform is more intuitive and expositionally simpler, and all our results for the Fourier transform carry over directly to the wavelet transform (albeit with greater mathematical complexity).

3 Spectral Analysis

Although spectral methods are not new to finance, as our literature review shows, current applications are sufficiently rare that a brief overview of spectral analysis may be appropriate before we turn to our own application. We begin in Section 3.1 with some historical context, and provide the formulation of the DFT in Section 3.2. We then present the main mathematical results on the co-spectrum in Section 3.3 that will be the basis of our applications to portfolio theory, and include an example using business cycle data in Section 3.4. Readers

familiar with spectral methods may prefer to skip this section and proceed to Section 4.

3.1 A Brief History

Over the past 200 years, Fourier analysis has made fundamental contributions to fields ranging from signal processing, communications, and neuroscience, to partial differential equations, astronomy, and geology. However, in contrast to its modern ubiquity, its origins stem from a very specific problem—modeling the orbits of celestial bodies.

The seeds of the theory were sown in the mid-18th century when the mathematicians Leonhard Euler, Joseph-Louis Lagrange, and Alexis Clairaut observed that orbits could be approximated as linear combinations of trigonometric functions, i.e., sines and cosines. In fact, to estimate the coefficients from the data, Clairaut published the first explicit formulation of the DFT in 1754, 14 years before Jean-Baptiste Joseph Fourier was born. While studying the orbit of the asteroid Pallas in 1805, Carl Friedrich Gauss discovered a computational shortcut. His calculation, which appeared posthumously as an unpublished paper in 1866, was the first clear use of the Fast Fourier Transform (FFT)—an efficient way to compute the DFT. Gauss' algorithm was largely forgotten until it was independently rediscovered in a more general form almost a century later by James Cooley and John Tukey in 1965.

Nourished by this half century of progress, the theory blossomed when Fourier presented his seminal paper on heat conduction to the Paris Academy in 1807. In his treatise, Fourier claimed that *any* arbitrary function could be represented by the superposition of trigonometric functions. This broader claim was initially received with much skepticism, and it would take another 5 years before the Paris Academy awarded his paper the grand prize in 1812. Despite the award, the Academy's panel of judges, which included Lagrange, Laplace and Legendre, still held reservations about the rigor of his analysis, especially in relation to the challenging question posed by convergence. Further advances by Dirichlet, Poisson, and Riemann addressed these subtle issues, and provided the foundation for today's Fourier transform, upon which many modern mathematical applications are based (Briggs and Hen-son, 1995).

3.2 The Fourier Transform

In particular, one of the most structurally revealing analyses that can be performed on a time series is to express its values as a linear combination of trigonometric functions. This procedure relies on the Discrete-Time Fourier Transform (DTFT), and allows the data to be transformed to the frequency domain. Specifically, given a finite-energy time series x_t , the DTFT is given by,

$$X(\omega) = \sum_{t=-\infty}^{\infty} x_t e^{-j\omega t}, \quad \omega \in [0, 2\pi) \quad (1)$$

where the frequency ω has units of radians per sample and j denotes the imaginary unit $\sqrt{-1}$. When x_t is real-valued, the inverse DTFT can be written in rectangular form as,

$$x_t = \frac{1}{2\pi} \int_0^{2\pi} \left[\Re[X(\omega)] \cos(\omega t) - \Im[X(\omega)] \sin(\omega t) \right] d\omega, \quad t \in (-\infty, \infty) \quad (2)$$

or in polar form as,

$$x_t = \frac{1}{2\pi} \int_0^{2\pi} |X(\omega)| \cos(\omega t + \angle X(\omega)) d\omega, \quad t \in (-\infty, \infty) \quad (3)$$

where $\Re[X(\omega)]$ and $\Im[X(\omega)]$ are the real and imaginary components of $X(\omega)$, and $|X(\omega)|$ and $\angle X(\omega)$ are its magnitude and phase, respectively.

If only a finite sample of x_t is available, or only a local portion of x_t needs to be analyzed, the DTFT reduces to the DFT. Specifically, given a sample of x_t from times $t = 0, \dots, T-1$, the T -point DFT is given by:²

$$X_k = \sum_{t=0}^{T-1} x_t e^{-j\omega_k t}, \quad k \in [0, T-1] \quad (4)$$

where $\omega_k = 2\pi k/T$. Again, when x_t is real-valued, the inverse DFT can be written in

²In general, for finite T , $X(\omega_k) \neq X_k$ as multiplying x_t by a rectangular window results in the convolution of $X(\omega)$ with the window's DTFT in the frequency domain.

rectangular form as,

$$x_t = \frac{1}{T} \sum_{k=0}^{T-1} \left[\Re[X_k] \cos(\omega_k t) - \Im[X_k] \sin(\omega_k t) \right], \quad t \in [0, T-1] \quad (5)$$

or in polar form as,

$$x_t = \frac{1}{T} \sum_{k=0}^{T-1} |X_k| \cos(\omega_k t + \angle X_k), \quad t \in [0, T-1]. \quad (6)$$

In this real-valued case, $X_k = X_{T-k}^*$, and so $|X_k| \cos(\omega_k t + \angle X_k) = |X_{T-k}| \cos(\omega_{T-k} t + \angle X_{T-k})$. Therefore, the lowest non-zero frequency occurs at $k=1$, and the highest frequency occurs at $k = \lfloor T/2 \rfloor$. The relation $h = TT_s/k$, where T_s is the time between samples and $0 \leq k \leq T/2$, can be used to convert the k th harmonic frequency to its corresponding time horizon.

As a concrete example, consider the rectangular pulse x_t for $t = 0$ to $t = 9$ shown in panel A of Figure 2. The real and imaginary components of the 10-point DFT are plotted in panels B and C, and their magnitude and phase in panels D and E. Panel F shows the reconstruction of x_t using only the constant X_0 term in (5), which is equivalent to the average value of the time series x_t . Panel G shows the reconstruction of x_t using both the constant term and the first non-zero low-frequency terms. These low-frequency terms are dominated by the sine term in (5) which has an amplitude proportional to the magnitude of the imaginary coefficients at $k=1$ and $k=9$ in panel C. More precisely, the amplitude of the low-frequency sinusoid can be seen in panel D, and its phase in terms of a shifted cosine in panel E. As more frequencies are included (see panel H), the output of the reconstruction begins to converge to the original time series. Ultimately, the reconstruction exactly matches the original rectangular pulse in panel A when all frequency terms are included.

Since the Fourier transform is a unitary operator that changes the basis function representation of a time series from impulses to sinusoids, Parseval's theorem states that, when represented as a vector, the Euclidean length of the time series is preserved under the transformation (with proper normalization). This observation forms the foundation of spectral decomposition, and provides a method to visualize the data in the frequency domain. This representation, known as the power spectrum, characterizes how much of the variability in

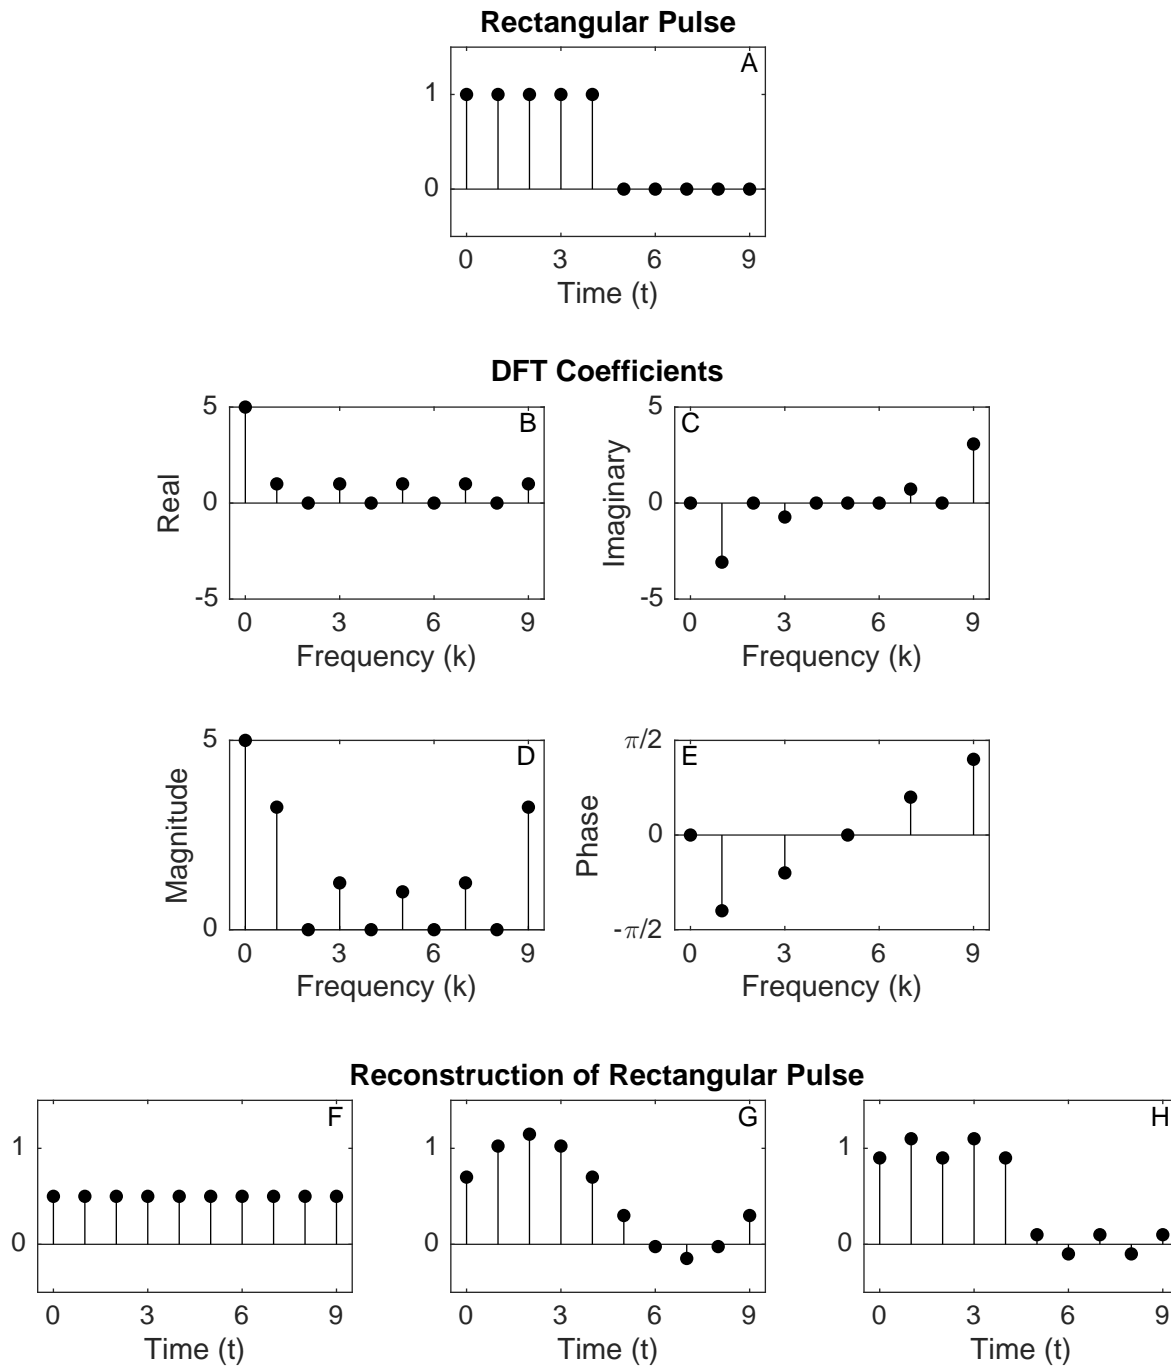


Figure 2: Panels B and C plot the real and imaginary components of the 10-point DFT coefficients of the rectangular pulse x_t shown in panel A. Panels D and E show the magnitude and phase of the DFT coefficients. Panels F through H show reconstructions of x_t using progressively more frequencies. The reconstruction matches the original pulse exactly when all frequencies are included.

the data comes from low- versus high-frequency fluctuations.

3.3 The Power Spectrum

In many situations, a time series can be modeled as the realization of a stochastic process, which can often be characterized by its first and second moments. The DTFT of the auto- and cross-covariance functions can then be interpreted as the frequency distribution of the power contained within the variance and covariance of these time series, respectively. Similarly, the inverse DTFT can be used to find the lagged second moments as functions of the auto- and cross-power spectra.

Let $\{x_t\}$ and $\{y_t\}$ form real-valued discrete-time wide-sense stationary stochastic processes with means m_x and m_y , and cross covariance function $\gamma_{xy}[m] = \mathbb{E}[(x_{t+m} - m_x)(y_t - m_y)]$.³ Assuming the cross-covariance function has finite energy, let $P_{xy}(\omega)$ be its DTFT,

$$P_{xy}(\omega) = \sum_{m=-\infty}^{\infty} \gamma_{xy}[m] e^{-j\omega m}. \quad (7)$$

The function $P_{xy}(\omega)$ is known as the cross-spectrum. Its real component, known as the co-spectrum, can be interpreted as the frequency decomposition of the covariance between x_t and y_t . Specifically, the covariance between $\{x_t\}$ and $\{y_t\}$ can be calculated using the inverse DTFT of $P_{xy}(\omega)$,

$$\text{cov}(x_t, y_t) \equiv \gamma_{xy}[0] = \frac{1}{2\pi} \int_0^{2\pi} \Re[P_{xy}(\omega)] d\omega. \quad (8)$$

We denote the co-spectrum⁴ as $L_{xy}(\omega) \equiv \Re[P_{xy}(\omega)]$.

This calculation of the power and cross-power spectra from the auto- and cross-covariance functions assumes the first and second moments of the stochastic process are known and do not change with time; however, for practical applications, especially those in finance, the underlying distributions are often unknown and nonstationary. To address this issue, we

³Specifically, the stochastic processes $\{x_t\}$ and $\{y_t\}$ are said to be wide-sense stationary if and only if $\mathbb{E}[x_t]$ and $\mathbb{E}[y_t]$ are constants independent of t , and $\mathbb{E}[x_{t_1}x_{t_2}]$, $\mathbb{E}[y_{t_1}y_{t_2}]$ and $\mathbb{E}[x_{t_1}y_{t_2}]$ depend only on the time difference $(t_1 - t_2)$.

⁴The co-spectrum, $L_{xy}(\omega)$, is the real part of the cross-spectrum, $P_{xy}(\omega)$. The imaginary part, $Q_{xy}(\omega)$, is called the quadrature spectrum.

compute the short-time Fourier transform to decompose rolling-window covariances into their frequency components. This approach uses the DFT to express windowed subsamples of x_t and y_t in the frequency domain, and then analyzes their magnitude and phase. When the time series are in phase at a given frequency, the contribution that frequency makes to the sample covariance is positive; when they are out of phase, that particular frequency's contribution will be negative. Longer windows will provide better frequency resolution, but will conflict with our ability to resolve changes in the statistical properties of signals over time.

Specifically, consider a real-valued subsample of x_t and y_t from times $t = 0, \dots, T-1$. The sample covariance over this interval can be calculated as:

$$\text{cov}\langle x_t, y_t \rangle = \frac{1}{T} \sum_{t=0}^{T-1} (x_t - \bar{x})(y_t - \bar{y}), \quad (9)$$

where \bar{x} and \bar{y} are the sample means of x_t and y_t over the same subperiod. This calculation is exactly equivalent to the one formed using the T -point DFT:

$$\text{cov}\langle x_t, y_t \rangle = \frac{1}{T} \sum_{k=1}^{T-1} \hat{L}_{xy}[k] \quad , \quad \hat{L}_{xy}[k] \equiv \frac{1}{T} \Re[X_k^* Y_k] \quad (10)$$

where X_k and Y_k are the T -point DFT coefficients of the subsample of x_t and y_t . Thus, the sum over $\hat{L}_{xy}[k]$ is proportional to the sample covariance of x_t and y_t . Moreover, the sum of $\hat{L}_{xy}[k]$ over a band of frequencies, $\text{cov}_K\langle x_t, y_t \rangle \mid K \subseteq \{1, \dots, T-1\}$, is proportional to that band's contribution to the sample covariance. For this reason the function $\hat{L}_{xy}[k]$, called the cross-periodogram, is an estimate of the co-spectrum at the harmonic frequency ω_k , and can be interpreted as the frequency distribution of the power contained in the sample covariance. It can be shown that these estimators are asymptotically unbiased, but inconsistent. Practical implementation details, including the standard errors of these estimators, are discussed in the Appendix. Further references on the statistical properties of spectrum estimates can be found in, for example, Jenkins and Watts (1968), Hannan (1970), Anderson (1971), Priestly (1981), Brockwell and Davis (1991), Brillinger (2001), Velasco and Robinson (2001), Phillips, Sun, and Jin (2006, 2007), Shao and Wu (2007), Oppenheim and

Schafer (2009), and Wu and Zaffaroni (2016).

Note that $k=0$, the zero frequency, is not involved in (10) since adding or subtracting a constant to either time series does not change the sample covariance. In addition, as mentioned in Section 3.2, values of k that are symmetric about $T/2$ (e.g., $k=1$ and $k=T-1$) have the same frequency and their contributions to the sample covariance are equivalent. Therefore, pairs of elements that correspond to the same frequency should be included together in the frequency band K to form the one-sided spectrum.⁵

As an illustrative example, suppose that,

$$x_t = \alpha_x + \beta_x F_t + u_t, \quad (11)$$

$$y_t = \alpha_y + \beta_y F_{t-1} + v_t, \quad (12)$$

where α_x , α_y , β_x and β_y are constants, and F_t , u_t and v_t are white-noise random variables that are uncorrelated at all leads and lags.




Figure 3: The spectral decomposition of (A) the variance of x_t , (B) the variance of y_t , and (C) the covariance of x_t and y_t .

Panels A and B display the one-sided co-spectrums, $L_{xx}(\omega)$ and $L_{yy}(\omega)$. Since x_t and y_t are serially uncorrelated at all leads and lags, their power spectrums are flat, and each frequency contributes equally to the variance. In this example, the lagged dependence of y_t on F_t relative to x_t suggests that x_t and y_t will be in phase over longer time horizons, and out of phase over shorter time horizons. As shown in panel C, this leads to a positive contribution to

⁵For example, see the one-sided and two-sided power spectrums in Figure 4. For real-valued time series, the cross-spectrum is conjugate symmetric causing the quadrature spectrum components to cancel. For this reason, we focus on the co-spectrum.

the covariance at low frequencies, and a negative contribution at high frequencies. Moreover, since x_t and y_t are uncorrelated, we find that $L_{xy}(\omega)$ integrates to 0.

3.4 The Business Cycle

One of the most natural applications of spectral analysis is to measure the business cycle, which many studies have done (Granger and Hatanaka, 1964; King and Watson, 1996; Baxter and King, 1999). Consider U.S. real GDP from the onset of the Great Moderation in the mid-1980s to 2015. The annualized quarterly percentage change in seasonally adjusted real GDP is plotted in panel A of Figure 4. Notice the data exhibit longer-scale cyclical patterns in accordance with recessions and expansions, as well as high-frequency oscillations related to more transitory dynamics.

As a first step, in panel B of Figure 4 we subtract the mean, to view fluctuations about the long-term growth rate. We then apply the DFT to decompose this adjusted time series into its frequency components, and plot the estimated two-sided power spectrum in panel C of Figure 4. The horizontal axis of this graph is now frequency instead of time, and the spectrum is symmetric about the center frequency. Therefore, it is common to aggregate coupled frequencies into the one-sided power spectrum as shown in panel D of Figure 4. In this form, it is clear that a substantial portion of the signal's power resides in low frequencies less than 1 cycle every 5 years. These periods correspond to economic expansions and recessions, i.e., the business cycle. A reconstruction of the original time series using only these low frequencies is shown in panel E of Figure 4. Notice the more transitory components have been removed, and what remains features the recession of the early 1990s, the internet bubble, the Financial Crisis, and the subsequent recovery. A second reconstruction using frequencies less than or equal to 1 cycle per year, is shown in panel F of Figure 4, and provides a more realistic reconstruction of the original data with more transitory effects. This example demonstrates how spectral analysis can often reveal structure in a time series not immediately evident in the raw data.

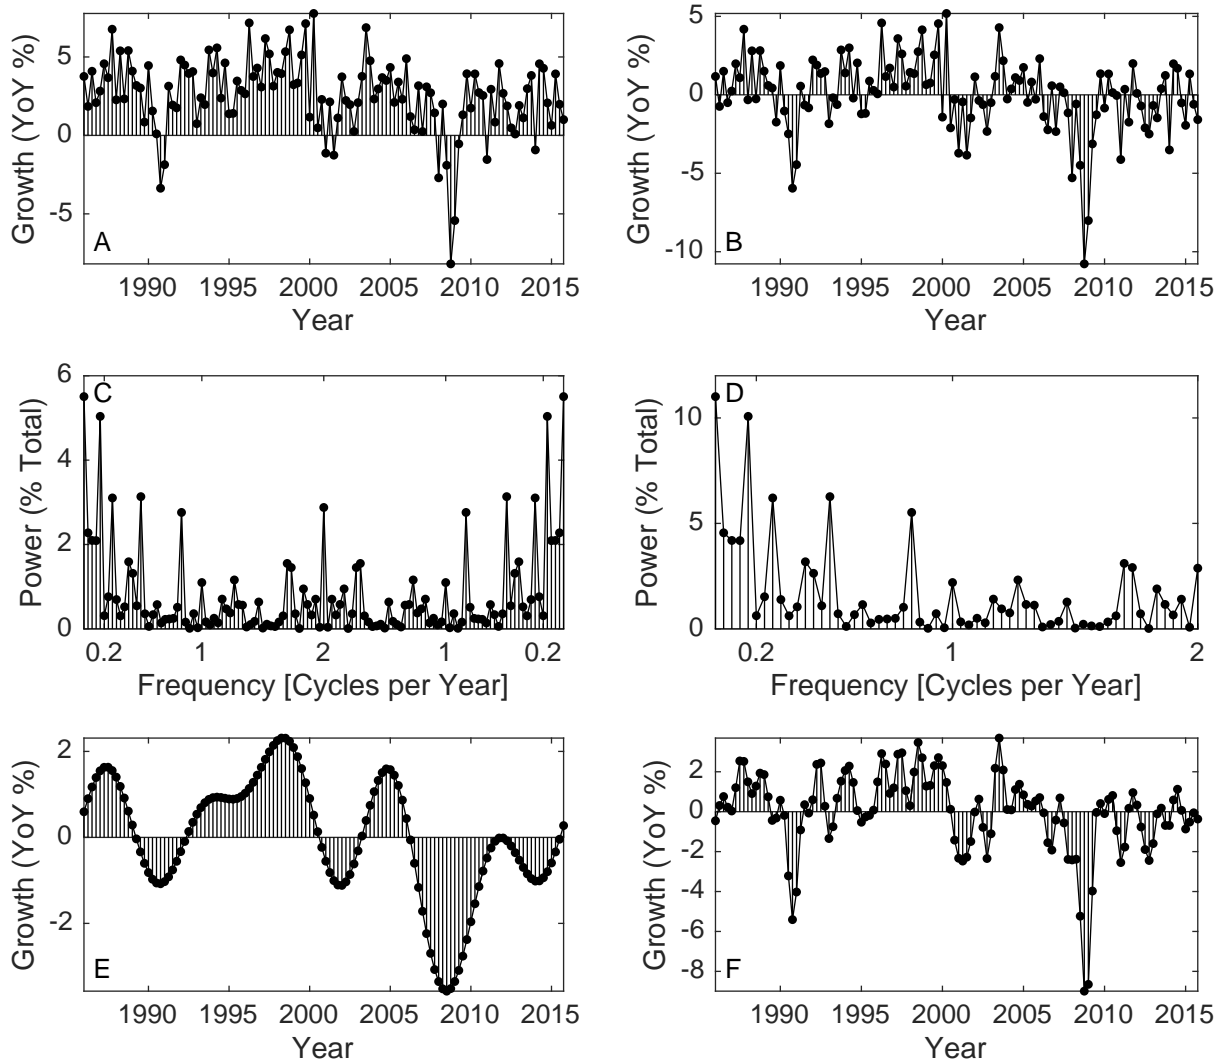


Figure 4: Illustration of how the DFT can be used to implement the spectral decomposition of a time series. The annualized quarterly percentage change in seasonally adjusted US real GDP from 1986 to 2015 is plotted in panel A. The same time series minus its mean is shown in panel B. Panel C shows its two-sided power spectrum after applying the DFT. Note that the horizontal axis represents frequency, and the vertical axis represents the relative contribution of each frequency to the overall variability of the time series. Panel D aggregates pairs of equivalent frequencies into the one-sided power spectrum. Panels E and F plot reconstructions of the time series using frequencies less than 1 cycle per 5 years, and less than or equal to 1 cycle per year, respectively.

4 Spectral Portfolio Theory

It has been observed that the properties of financial securities are not constant, but change over time. Given that economic shocks produce distinct effects on financial assets over different time horizons, these dynamics are likely to have important implications for any theory of risk, reward, and portfolio choice. However, the traditional inputs into these analytics—means, variance, covariances, alphas, and beta—are static, and do not distinguish between the short- and long-term components of these dynamics. The fact that the standard estimators of these statistics are invariant to how the data are ordered suggests that traditional portfolio analytics are incapable of capturing the dynamic properties of asset returns.

In this section, we apply spectral analysis to develop dynamic, frequency-specific analogs for each of these portfolio analytics. In Section 4.1 we provide a spectral decomposition of volatility, and do the same for correlation and beta in Section 4.2. We then turn to alpha, tracking error, and the information ratio in Section 4.3. In Section 4.5, we show how to incorporate these concepts into the traditional mean-variance portfolio optimization framework. Our exposition focuses on sample statistics and DFT-based estimates of the underlying power spectrums, but we note that each equation has a population statistic analog based on the DTFT.

4.1 Volatility

Estimating volatility is central to mean-variance portfolio management, performance attribution, and risk management. A spectral decomposition of returns allows us to measure the fraction of variability that can be attributed to fluctuations at different time scales.

Let x_t be the one-period return of a security between dates $t-1$ and t . The sample variance of returns over an interval from $t = 0, \dots, T-1$ can be decomposed into its frequency components using (10):

$$\text{var}\langle x_t \rangle = \frac{1}{T} \sum_{k=1}^{T-1} \hat{L}_{xx}[k] \quad , \quad \hat{L}_{xx}[k] \equiv \frac{1}{T} |X_k|^2 \quad (13)$$

where X_k are the T -point DFT coefficients of the subsample of x_t . As an illustrative example, Figure 5 decomposes the 10-year rolling sample variance of the daily returns of the CRSP

value-weighted and equal-weighted market indices from 1926 to 2015 into its low (less than 1 cycle per month), medium (between 1 cycle per month to 1 cycle per week), and high (more than 1 cycle per week) frequency components. This spectral decomposition is compared to a white noise null hypothesis where the windowed returns were rendered serially uncorrelated by generating random permutations of their order. This exercise was repeated 10,000 times from which 95% confidence intervals were formed around the flat-band, white-noise null hypothesis.




Figure 5: Spectral decomposition of the 10-year rolling sample variance of the daily returns of CRSP value-weighted market index (panel A), and the CRSP equal-weighted market index (panel B) from 1926 to 2015. Frequency components are grouped into 3 categories: high frequencies (more than 1 cycle per week), mid frequencies (between 1 cycle per week and 1 cycle per month), and low frequencies (less than 1 cycle per month).

Our analysis shows that, from the mid-1960s to late-1990s, the variance of both value-weighted and equal-weighted market returns exhibited smaller fluctuations at short time scales (between 2 and 5 days), and greater fluctuations at longer time scales (greater than 1 month), than would be expected if returns were serially uncorrelated. In fact, this effect is more pronounced, and continues into the late 2000s, for the equal-weighted market returns. This low-frequency power is in agreement with the large positive serial correlation in weekly returns described in Lo and MacKinlay (1990), which would tend to shift power from high frequencies to lower frequencies. However, this spectral decomposition also shows these dynamics weakening over the subsequent decades, most likely in response to increased

competitive forces and technological advances such improved telecommunications, standardized electronic information exchange protocols, and automated trading. This simple example demonstrates the usefulness of the frequency domain in visualizing complex dynamics that may exist over a wide range of leads and lags in the time domain. For example, in addition to tests in the time domain that detect local correlation between neighboring samples, the power spectrum allows us to detect departures from white noise caused by periodic effects such as seasonal variation.

4.2 Correlation and Beta

The dynamic correlation, a spectral based measure of correlation introduced to the economics literature by Croux et al. (2001), gauges the degree of synchronization in the fluctuation of two time series at different frequencies. It is derived by normalizing the band-limited sample covariance by the square root of the band-limited sample variances. Specifically, for a frequency band K , the dynamic correlation is given by,

$$\rho_K \langle x_t, y_t \rangle = \frac{\sum_{k \in K} \hat{L}_{xy}[k]}{\sqrt{\sum_{k \in K} \hat{L}_{xx}[k]} \sqrt{\sum_{k \in K} \hat{L}_{yy}[k]}}. \quad (14)$$

Since returns are real-valued, $-1 \leq \rho_K \langle x_t, y_t \rangle \leq 1$, and is computationally equivalent to calculating the sample correlation of the inverse DFT reconstructions of the time series, restricted to the frequencies specified by K . Confidence intervals for this estimator are provided in the Appendix.

Similarly, linear factor models are often used in financial applications, including market model regressions, the CAPM, the APT, and the Fama-French 5-factor model. As we noted above, the estimated beta coefficients in these models are static measures that are incapable of capturing dynamic relationships among the variables. Band spectrum regression, proposed by Engle (1974), captures the sensitivity of the dependent variable to the fluctuations in the independent variables over different time horizons. More precisely, for a frequency band

$K \subseteq \{1, \dots, T-1\}$, the dynamic beta coefficients for an M -factor model are given by,

$$\beta_K \langle y_t; x_{1,t}, \dots, x_{M,t} \rangle = \left[\sum_{k \in K} \hat{\mathbf{L}}_{\mathbf{xx}} \right]^{-1} \left[\sum_{k \in K} \hat{\mathbf{L}}_{\mathbf{xy}} \right], \quad (15)$$

where,

$$\sum_{k \in K} \hat{\mathbf{L}}_{\mathbf{xx}} \equiv \begin{bmatrix} \sum_{k \in K} \hat{L}_{x_1, x_1}[k] & \cdots & \sum_{k \in K} \hat{L}_{x_1, x_M}[k] \\ \vdots & \ddots & \vdots \\ \sum_{k \in K} \hat{L}_{x_M, x_1}[k] & \cdots & \sum_{k \in K} \hat{L}_{x_M, x_M}[k] \end{bmatrix}, \quad \sum_{k \in K} \hat{\mathbf{L}}_{\mathbf{xy}} \equiv \begin{bmatrix} \sum_{k \in K} \hat{L}_{x_1, y}[k] \\ \vdots \\ \sum_{k \in K} \hat{L}_{x_M, y}[k] \end{bmatrix}. \quad (16)$$

When only one factor is present, (15) reduces to the familiar expression,

$$\beta_K \langle y_t; x_t \rangle = \frac{\text{cov}_K \langle x_t, y_t \rangle}{\text{var}_K \langle x_t \rangle} = \rho_K \langle x_t, y_t \rangle \frac{\sqrt{\text{var}_K \langle y_t \rangle}}{\sqrt{\text{var}_K \langle x_t \rangle}}. \quad (17)$$

Intuitively, these calculations are computationally equivalent to estimating the beta coefficients by regressing the inverse DFT reconstruction of the time series, restricted to the frequencies specified by K . Standard errors and the F statistic for this band-spectrum regression are provided in the Appendix.

As an illustrative example, Table 1 tests the hypothesis that long- and short-term components of several hedge-fund style-category returns are equally sensitive to market index returns across all frequencies. Specifically, we analyzed the monthly returns of the HFRI ED: Distressed/Restructuring, HFRI FOF: Market Defensive, and HFRI EH: Quantitative Directional indices relative to the monthly returns of the CRSP value-weighted market index using the dynamic correlation and beta measures described above. The short-term component was assumed to include frequencies higher than 1 cycle per year, and the original series were de-meanned.

The F statistic value of 43.116 ($p < 0.001$) for the Distressed/Restructuring index rejects the hypothesis that the sensitivity of this strategy's returns to market movements over this period did not differ between short- and long-term components. The F statistics for the Market Defensive and Quantitative Directional indices have p -values of 0.056 and 0.009, respectively. Clearly, the null hypothesis that the sensitivity of these indices to market

	$\hat{\rho}$ (95% CI)	$\hat{\rho}_{LF}$ (95% CI)	$\hat{\rho}_{HF}$ (95% CI)	$\hat{\beta}$ (SE)	$\hat{\beta}_{LF}$ (SE)	$\hat{\beta}_{HF}$ (SE)	$F_{1,309}$
Distressed/Restructuring	0.581 (0.502,0.650)	0.726 (0.564,0.833)	0.558 (0.468,0.637)	0.250 (0.020)	0.500 (0.067)	0.190 (0.018)	43.116***
Market Defensive	0.087 (-0.025,0.196)	0.107 (-0.171,0.369)	0.081 (-0.041,0.201)	0.033 (0.022)	0.043 (0.056)	0.030 (0.023)	0.056
Quantitative Directional	0.858 (0.826,0.885)	0.842 (0.739,0.907)	0.862 (0.827,0.890)	0.704 (0.024)	0.709 (0.064)	0.703 (0.026)	0.009

*** $p < 0.001$

Table 1: All-, low-, and high-frequency estimates of the correlation and beta of hedge fund index monthly returns from 1990 to 2015 with the CRSP value-weighted market index returns. Frequencies are grouped into two categories: low frequencies (less than or equal to 1 cycle per year), and high frequencies (more than 1 cycle per year). F statistics are formed to compare the restricted and unrestricted regression models.

movements is the same across short and long horizons cannot be rejected at the standard significance levels.




Figure 6: Cumulative returns of hedge fund indices alongside the CRSP value-weighted market index from 1990 to 2015.

Figure 6 plots the cumulative returns of these hedge fund indices alongside the cumulative return of the market index. The figure illustrates that the Market Defensive index had low sensitivity to fluctuations in market returns across all frequencies ($\hat{\beta}_{LF} = 0.043$; $\hat{\beta}_{HF} = 0.030$), while the Quantitative Directional index responded strongly to all market return fluctuations ($\hat{\beta}_{LF} = 0.709$; $\hat{\beta}_{HF} = 0.703$). This observation corresponds to the flat correlation and beta values across low to high frequencies listed for these strategies in Table 1. However,

the Distressed/Restructuring index appeared sensitive primarily to long-term fluctuations, leading to significantly larger correlation and beta values at low frequencies relative to high frequencies ($\hat{\beta}_{LF} = 0.500$; $\hat{\beta}_{HF} = 0.190$). This suggests that while this strategy had a lower overall beta ($\hat{\beta} = 0.250$), it was asymmetrically exposed to low-frequency, systematic risk.

4.3 Alpha, Tracking Error, and Information Ratios

From microseconds to years, the shortest decision interval of today’s investment strategies span a wide range of time horizons. While legendary value investor Warren Buffett changes his portfolio weights rather slowly, the same cannot be said for famed day-trader Steven Cohen of SAC Capital, yet both manage to generate value through active investment. While alpha, tracking error, and information ratios are the standard tools for gauging the value-added of a portfolio manager, they are unable to capture the essence of active management: time-series predictability. Specifically, none of these standard performance metrics directly measure the dynamic relationship between weights and returns which is often the central focus of active investment strategies.

In this section, we propose an explicit measure of the value of active management—dynamic alpha—that takes into account forecast power across multiple time horizons. Expanding on the framework of the active/passive or “AP” decomposition developed by Lo (2008), we use a frequency or “F” decomposition to separate the expected return of a portfolio into distinct components that depend on the correlation between portfolio weights and returns at different frequencies. The result is a static component that measures the portion of a portfolio’s expected return due to passive investments, and multiple active components that capture the manager’s timing ability across a range of time horizons. Our method closely parallels Hasbrouck and Sofianos (1993), however we make a novel modification to their analysis to make it applicable to the expected returns of portfolios.

Our approach uses the DFT to express the portfolio’s underlying security weights and returns in the frequency domain and then analyzes their phase. When the weights and returns are in phase at a given frequency, the contribution that frequency makes to the portfolio’s expected return is positive. When they are out of phase, then that particular frequency’s contribution will be negative.

If we consider a portfolio with N securities, then for $t=0, \dots, T-1$, the average one-period

portfolio return can be calculated as:

$$\bar{r}_p = \frac{1}{T} \sum_{i=1}^N \sum_{t=0}^{T-1} w_{i,t} r_{i,t}, \quad (18)$$

where $w_{i,t}$ and $r_{i,t}$ are the realized weight and return of the i th stock at time t , respectively. Using the definition of covariance, Lo (2008) showed the average portfolio return can be decomposed into an active component (α_p) and a passive component (ν_p) as follows,

$$\bar{r}_p = \alpha_p + \nu_p, \quad (19)$$

$$\alpha_p = \sum_{i=1}^N \text{cov}(w_{i,t}, r_{i,t}), \quad \nu_p = \sum_{i=1}^N \overline{w_{i,t}} \cdot \overline{r_{i,t}}. \quad (20)$$

The value of the passive component arises from the manager's average position in a security, and can be thought of as the portion of the portfolio's return that results from collecting risk premiums. The value of the active component estimates the profitability of the portfolio manager's conscious decision to buy, sell, or avoid a security by aggregating the sample covariances between the portfolio weights, $w_{i,t}$, and security returns, $r_{i,t}$. In particular, if a manager has positive weights when security returns are positive and negative weights when returns are negative, this implies positive covariances between portfolio weights and returns, and will have a positive impact on the portfolio's average return. In effect, the covariance term captures the manager's timing ability, asset by asset.

Using (10) allows us to decompose this covariance term further, capturing the manager's timing ability over multiple time horizons:

$$\alpha_p = \sum_{k=1}^{T-1} \alpha_{p,k} \quad , \quad \alpha_{p,k} = \frac{1}{T} \sum_{i=1}^N \hat{L}_{w_i r_i}[k] \quad , \quad (21)$$

where $\hat{L}_{w_i r_i}[k]$ is the co-spectrum estimate between the weights and returns for stock i . This spectral decomposition first deconstructs the weights and returns into their various frequency components. At each frequency, if the weights and returns are in phase, then that time horizon's contribution to the average portfolio return will be positive. If the two signals are out of phase, then that particular frequency's contribution will be negative. Note that

in this form, $\alpha_{p,0} = \nu_p$, and often it is convenient to include $\alpha_{p,0}$ when computing the F decomposition.

In addition to quantifying the value added from active management across time horizons, we can also gauge the consistency of a portfolio manager's timing ability. Historically, the consistency of investment skill has been characterized by the volatility of the tracking error, which is a measure of the variability of the difference between the portfolio return and some benchmark return. Low tracking error volatility and a positive excess return (i.e., alpha) indicates that the manager is reliably adding value through active management. Dividing alpha by the tracking error volatility measures the efficiency with which a manager generates excess returns and is called the information ratio. The higher the information ratio, the better the manager.

These measures can be incorporated into our framework by defining the active risk, σ_A , as the variability of the difference between the portfolio return, $r_{p,t}$, and the passive component, $\nu_{p,t} = \sum_{i=1}^N \overline{w_{i,t}} \cdot r_{i,t}$. Specifically,

$$\sigma_A = \sqrt{\text{var}\langle r_{p,t} - \nu_{p,t} \rangle}, \quad (22)$$

where σ_A is a measure of the risk taken by the portfolio manager in an attempt to generate higher returns by engaging in timing decisions. The information ratio, I , can then be defined as,

$$I = \frac{\alpha_p}{\sigma_A}, \quad (23)$$

and is a risk-adjusted measure of the active component. These performance metrics can be calculated for a specific range of time horizons by aggregating the frequency components of α_p and σ_A over the band of interest. This provides us with a risk-adjusted measure of the manager's timing ability for a specific frequency band. Intuitively, it quantifies the manager's predictive power across a range of time horizons, but also attempts to identify the consistency of this power.

4.4 Numerical Examples

To develop further intuition for our spectral decomposition, consider the following simple numerical example of a portfolio of two assets, one that yields a monthly return that alternates between 1% and 2% (Asset 1) and the other that yields a fixed monthly return of 0.15% (Asset 2). Let the weights of this portfolio, called A1, be given by 75% in Asset 1 and 25% in Asset 2. Table 2 illustrates the dynamics of this portfolio over a 12-month period, where the average return of the portfolio is 1.1625% per month, all of which is due to the passive component. In this case, because the weights are constant, the active risk measure will also be 0%.

Month	w_1	r_1	w_2	r_2	r_p	
Strategy A1						
1	75%	1.00%	25%	0.15%	0.7875%	
2	75%	2.00%	25%	0.15%	1.5375%	
3	75%	1.00%	25%	0.15%	0.7875%	
4	75%	2.00%	25%	0.15%	1.5375%	
5	75%	1.00%	25%	0.15%	0.7875%	
6	75%	2.00%	25%	0.15%	1.5375%	
7	75%	1.00%	25%	0.15%	0.7875%	
8	75%	2.00%	25%	0.15%	1.5375%	
9	75%	1.00%	25%	0.15%	0.7875%	
10	75%	2.00%	25%	0.15%	1.5375%	
11	75%	1.00%	25%	0.15%	0.7875%	
12	75%	2.00%	25%	0.15%	1.5375%	
Mean:	75%	1.50%	25%	0.15%	1.1625%	
F decomposition of \bar{r}_p						
ν_p	$2\alpha_{p,1}$	$2\alpha_{p,2}$	$2\alpha_{p,3}$	$2\alpha_{p,4}$	$2\alpha_{p,5}$	$\alpha_{p,6}$
1.1625%	0%	0%	0%	0%	0%	0%

Table 2: The expected return of a constant portfolio depends only on the passive component.

Now consider portfolio A2, which differs from A1 only in that the portfolio weight for Asset 1 alternates between 50% and 100% in phase with Asset 1's returns which alternates between 1% and 2% (see Table 3). In this case, the total expected return is 1.2875% per month, of which 0.1250% is due to the positive correlation between the portfolio weight for Asset 1 and its return at the shortest-time horizon (i.e., highest frequency). In addition, the active risk for this portfolio is 0.3375%, and the information ratio is about 0.37.

Month	w_1	r_1	w_2	r_2	r_p
Strategy A2					
1	50%	1.00%	50%	0.15%	0.5750%
2	100%	2.00%	0%	0.15%	2.0000%
3	50%	1.00%	50%	0.15%	0.5750%
4	100%	2.00%	0%	0.15%	2.0000%
5	50%	1.00%	50%	0.15%	0.5750%
6	100%	2.00%	0%	0.15%	2.0000%
7	50%	1.00%	50%	0.15%	0.5750%
8	100%	2.00%	0%	0.15%	2.0000%
9	50%	1.00%	50%	0.15%	0.5750%
10	100%	2.00%	0%	0.15%	2.0000%
11	50%	1.00%	50%	0.15%	0.5750%
12	100%	2.00%	0%	0.15%	2.0000%
Mean:	75%	1.50%	25%	0.15%	1.2875%

F decomposition of $\overline{r_p}$						
ν_p	$2\alpha_{p,1}$	$2\alpha_{p,2}$	$2\alpha_{p,3}$	$2\alpha_{p,4}$	$2\alpha_{p,5}$	$\alpha_{p,6}$
1.1625%	0%	0%	0%	0%	0%	0.1250%

Table 3: The dynamics of the portfolio weights are positively correlated with returns at the shortest time horizon, which adds value to the portfolio and yields a positive contribution from the highest frequency ($\alpha_{p,6}$).

Finally, consider a third portfolio A3 which also has alternating weights for Asset 1, but exactly out of phase with Asset 1's returns. When the return is 1%, the portfolio weight is 100%, and when the return is 2%, the portfolio weight is 50%. Table 4 confirms that this is counterproductive as Portfolio A3 loses 0.1250% per month from its highest frequency component, and its total expected return is only 1.0375%. In this case, the active risk is 0.3375%, and the information ratio is -0.37.

Month	w_1	r_1	w_2	r_2	r_p
Strategy A3					
1	100%	1.00%	0%	0.15%	1.0000%
2	50%	2.00%	50%	0.15%	1.0750%
3	100%	1.00%	0%	0.15%	1.0000%
4	50%	2.00%	50%	0.15%	1.0750%
5	100%	1.00%	0%	0.15%	1.0000%
6	50%	2.00%	50%	0.15%	1.0750%
7	100%	1.00%	0%	0.15%	1.0000%
8	50%	2.00%	50%	0.15%	1.0750%
9	100%	1.00%	0%	0.15%	1.0000%
10	50%	2.00%	50%	0.15%	1.0750%
11	100%	1.00%	0%	0.15%	1.0000%
12	50%	2.00%	50%	0.15%	1.0750%
Mean:	75%	1.50%	25%	0.15%	1.0375%

F decomposition of $\overline{r_p}$						
ν_p	$2\alpha_{p,1}$	$2\alpha_{p,2}$	$2\alpha_{p,3}$	$2\alpha_{p,4}$	$2\alpha_{p,5}$	$\alpha_{p,6}$
1.1625%	0%	0%	0%	0%	0%	-0.1250%

Table 4: The dynamics of the portfolio weights are negatively correlated with returns at the shortest time horizon, which subtracts value from the portfolio and yields a negative contribution from the highest frequency ($\alpha_{p,6}$).

Note that in all three cases, the lowest frequency components are identical at 1.1625% per month because the average weight for each asset is the same across all three portfolios. The only differences among A1, A2, and A3 are the dynamics of the portfolio weights at the shortest time horizon, and these differences give rise to different values for the highest frequency component. As shown in (21), contributions from higher frequencies ($k > 0$) sum to Lo's active component. These higher frequency contributions can then be interpreted as the portion of the active component that arises from a given time horizon.

For a more realistic example, consider the long/short equity market-neutral strategy of

Lo and MacKinlay (1990) described in the introduction:

$$w_{i,t}(q) = -\frac{1}{N}(r_{i,t-q} - r_{m,t-q}), \quad (24)$$

$$r_{m,t-q} = \frac{1}{N} \sum_{i=1}^N r_{i,t-q} \quad (25)$$

for some $q > 0$.

By buying the date- $t - q$ losers and selling the date- $t - q$ winners at the onset of each date t , this strategy actively bets on mean reversion across all N stocks and profits from reversals that occur within the subsequent interval. For this reason, Lo and MacKinlay (1990) termed this strategy “contrarian” as it benefits from market overreaction and mean reversion (i.e., when underperformance is followed by positive returns and outperformance is followed by negative returns). By construction, the weights sum to zero and therefore the strategy is also considered a “dollar-neutral” or “arbitrage” portfolio. This implies that much of the portfolio’s return should be due to active management and value will be added near frequencies inversely related to q .

Now suppose that stock returns satisfy the following simple MA(1) model,

$$r_{i,t} = \varepsilon_{i,t} + \lambda \varepsilon_{i,t-1}, \quad (26)$$

where the $\varepsilon_{i,t}$ are serially and cross-sectionally uncorrelated white-noise random variables with variance σ^2 . In this case, the expected one-period portfolio return can be calculated as,

$$E[r_p] = \begin{cases} -\lambda \sigma^2 (1 - \frac{1}{N}) & \text{if } q = 1 \\ 0 & \text{if } q > 1. \end{cases} \quad (27)$$

We see that the expected return is proportional to the mean reversion factor λ and the volatility factor σ^2 when $q = 1$. When $q > 1$, the expected return yields 0 since there is no correlation in the returns between times $t - q$ and t . Applying the F decomposition, we find

that,

$$\alpha_p(\omega) = -\sigma^2 \left(1 - \frac{1}{N}\right) \left[\lambda \cos(\omega(q+1)) + (1 + \lambda^2) \cos(\omega q) + \lambda \cos(\omega(q-1)) \right], \quad (28)$$

$$\omega \in [0, 2\pi).$$

Panel A of Figure 7 plots the dynamic alpha for the case of no serial correlation ($\lambda=0$) when $q=1$. The dynamic alpha at high frequencies is positive indicating that the weights and returns are in phase over these short time horizons. However, this value added gets cancelled out since the weights and returns are out of phase at longer time horizons, resulting in zero net alpha.

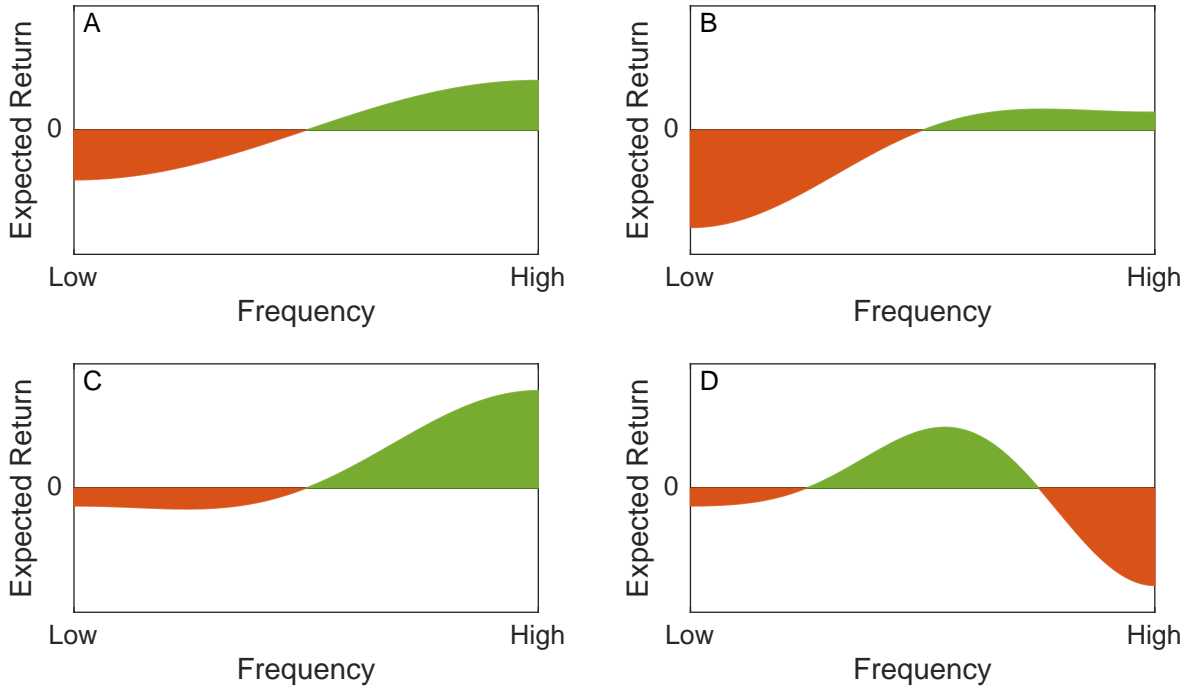


Figure 7: F decomposition of the contrarian trading strategy with $q=1$ applied to the serially uncorrelated (Panel A), momentum (Panel B), and mean reversion (Panel C) implementations of (26). Panel D shows the case of the mean reversion realization of (26) but with q increased from 1 to 2.

Panels B and C of Figure 7 show the dynamic alpha for the cases of momentum ($\lambda > 0$) and mean reversion ($\lambda < 0$) in the first lag of returns, respectively. For the mean reversion case we notice that both the lowest and highest frequencies are more profitable relative to the serially uncorrelated case. This result is intuitive since both weights and returns

now have more variability in these higher frequency fluctuations. These high-frequency components will be in phase leading to a large positive contribution and an overall positive alpha. The momentum case is the opposite. Relative to the serially uncorrelated case, both the lowest and highest frequencies are less profitable, and the net contribution over all frequency components is negative.

Panel D of Figure 7 shows the dynamic alpha for the case of one-period mean reversion ($\lambda < 0$), but when we increase q from 1 to 2. We notice that when $q = 2$, the portfolio loses most of its profits at the highest frequencies. This occurs because the returns receive much of their variability from the highest frequencies, however they will tend to be out of phase with the weights at these frequencies. By reducing q from 2 to 1, we improve the strategy's timing at the shortest time horizon and convert these losses into gains. This example provides one simple illustration of how the F decomposition can identify expected-return "leakages" in an investment process that can be exploited to improve overall performance.

Finally, suppose that in each of the above cases, the volatility factor doubles. In this case, the contribution to the average portfolio return from each frequency quadruples which is a characteristic that we will encounter when we apply the F decomposition to empirical returns in Section 5.

4.5 Implications for Portfolio Optimization

We have shown that volatilities and correlations change not only across time, but also across frequencies. One implication of this finding is that frequency can now be included as a parameter in portfolio design. This additional parameter is particularly useful when portfolio goals differ across time horizons, and when investors wish to target specific horizons because of their preferences and life cycle. For example, short-run volatility, even if correlated with an investor's portfolio, may not affect his investment goals if his time horizon is much longer, e.g., Warren Buffett and Berkshire Hathaway. Similarly, long-term fluctuations may be unimportant to a high-frequency trader who does not operate at the same timescale. The frequency domain provides a systematic framework for incorporating these considerations into a portfolio.

In mean-variance portfolio theory, given a target value, $\tilde{\mu}$, for the expected portfolio return, the efficient portfolio weights, $\tilde{\mathbf{w}}$, are those that minimize the portfolio variance for

all portfolios with expected return $\tilde{\mu}$. Mathematically, the optimization problem can be written as,

$$\tilde{\mathbf{w}} = \arg \min_{\mathbf{w}} \mathbf{w}^T \Sigma \mathbf{w} \quad (29)$$

subject to the constraints

$$\mathbf{w}^T \boldsymbol{\mu} = \tilde{\mu} \quad \text{and} \quad \mathbf{w}^T \mathbf{1} = 1 \quad (30)$$

where w_i is the portfolio weight on the i th security, $\mu_i = E[r_i]$ and $\Sigma_{i,j} = \text{cov}(r_i, r_j)$.

The inputs for this optimization problem are the expected returns and covariance matrix of these securities. Similarly, a time-horizon-specific mean-variance optimization, restricted to the frequency band K , can be developed by simply replacing the covariance matrix estimates with those based on the co-spectrum. As a first-order approximation, sample estimates based on (10) can be used for these values. This band-limited framework has the attractive feature that optimization techniques developed to solve for the efficient frontier are still valid as the formulation of the problem has not been not affected.

Consider the monthly returns of the Distressed/Restructuring and Market Defensive indices in section 4.2. Figure 8 shows the cumulative percentage of the total variance of these returns from January 1990 to December 2002, as a function of increasing frequency. We see that the Distressed/Restructuring index has substantial low-frequency variability, consistent with the low-frequency systematic risk described previously. An investor with a long time horizon may therefore consider weighting this risk more heavily when considering their asset allocation.

As a simple example, suppose at the end of 2002 an investor considers forming a portfolio that consists of these two indices. Assuming a risk-free rate of 0%, standard mean-variance optimization suggests the portfolio that maximizes the Sharpe ratio allocates 58% of capital in the Distressed/Restructuring strategy and 42% in the Market Defensive strategy. If only frequencies less than 1 cycle per 9 months are considered when estimating the covariance matrix, then the optimization suggests an allocation of only 39% in the Distressed/Restructuring strategy. In the subsequent period, from January 2003 to December 2015, the annualized

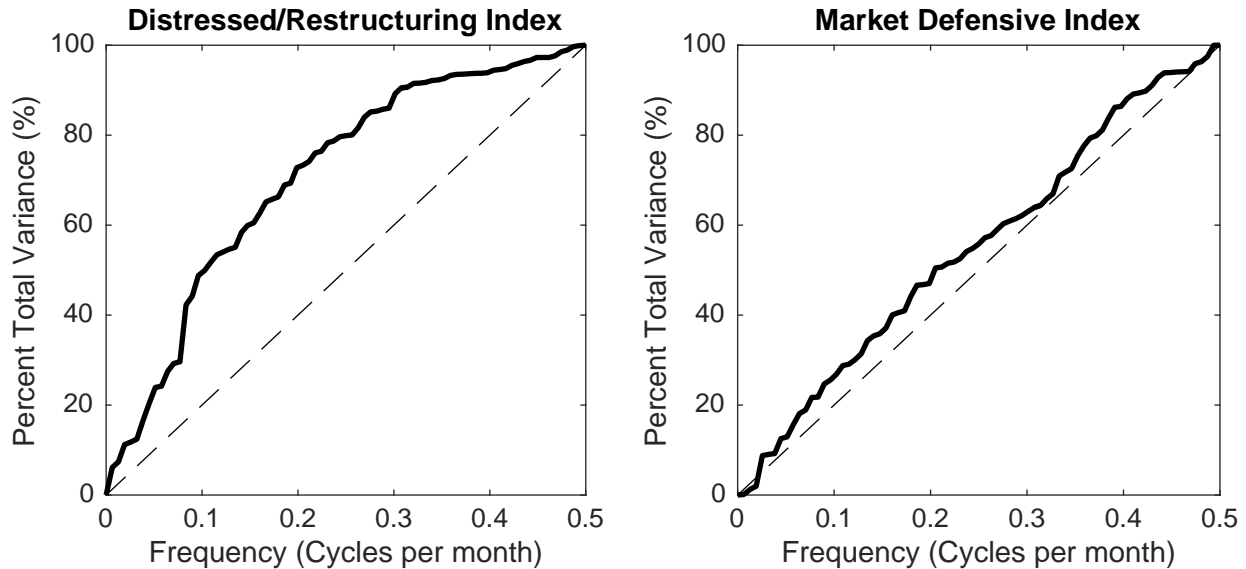


Figure 8: Cumulative percentage of the total variance in the monthly returns of the Distressed/Restructuring index and Market Defensive index from January 1990 to December 2002 as a function of increasing frequency. Notice that a substantial percentage of the Distressed/Restructuring index’s variance can be attributed to low frequencies.

Sharpe ratio of the monthly portfolio returns will be 1.18 and 1.11 for the standard and band-limited optimizations, respectively. However, if the Sharpe ratios are calculated using annual returns instead, their respective performance in the latter period changes to 0.69 and 0.78. Thus, if an investor has a longer time horizon and considers performance at yearly rather than monthly intervals, then the low-frequency, band-limited mean variance optimization provides better performance.

In the previous example, a similar result would have been achieved had we based the portfolio optimization on longer holding-period returns. However, this band-limited mean-variance optimization can be generalized such that the optimization attempts to shape the power spectrum of the portfolio’s returns into an any functional form. Such an objective would be needed when an investor faces risk constraints imposed at different frequencies. For example, an investor who wants to minimize both long and short term fluctuations may try to diversify his risk across frequencies. Moreover, a portfolio manager may have to satisfy one set of investors focused on short-term fluctuations, and another focused on longer horizons. One method to accomplish this goal would be to add a regularization term to the objective

function,

$$\tilde{\mathbf{w}} = \arg \min_{\mathbf{w}} \mathbf{w}^T \Sigma \mathbf{w} + \lambda H(\mathbf{w}). \quad (31)$$

This term would penalize the cost function at a rate λ if the resulting portfolio’s power spectrum was concentrated too highly in a particular frequency band. Concentration measures based on information-theoretic entropy or the Herfindahl-Hirschman index would be suitable candidates for H if the objective were to spread risk across frequencies. Moreover, functional distance measures such as KL divergence or total variation distance could be used if one wanted to approximate a more general form for the portfolio return power spectrum. An interesting area for future research is to investigate the practical advantages of such a framework in a broader variety of portfolios and strategies.

5 An Empirical Example

To develop a better understanding of the characteristics of the F decomposition, we apply our framework to two market-neutral equity trading strategies that, by construction, are particularly dynamic: Lo and MacKinlay’s (1990) contrarian (mean reversion) trading strategy, and a simplified version of Jegadeesh and Titman’s (1993) momentum strategy.⁶ We apply this analysis to weekly and monthly returns on all S&P 500 stocks from January 1, 1964 to December 31, 2015.

5.1 Mean Reversion

Panels A and B of Figure 9 plot the 1-year rolling average of the mean-reversion trading strategy’s portfolio return for $q = 1$ week and $q = 2$ weeks, respectively. Panels D and E

⁶Note that since the weights of these strategies sum to zero, their return for a given interval can be calculated as the profit-and-loss of the strategy’s positions over that interval, divided by the capital required to support those positions. In the following analysis, we assume that Regulation T applies, and so the minimum amount of capital required is one-half the total capital invested (often stated as 2:1 leverage, or a 50% margin requirement). The unleveraged (Reg T) portfolio return, $r_{p,t}$ is given by:

$$r_{p,t} = \frac{\sum_{i=1}^N w_{i,t} r_{i,t}}{I_t} \quad , \quad I_t = \frac{1}{2} \sum_{i=1}^N |w_{i,t}| .$$

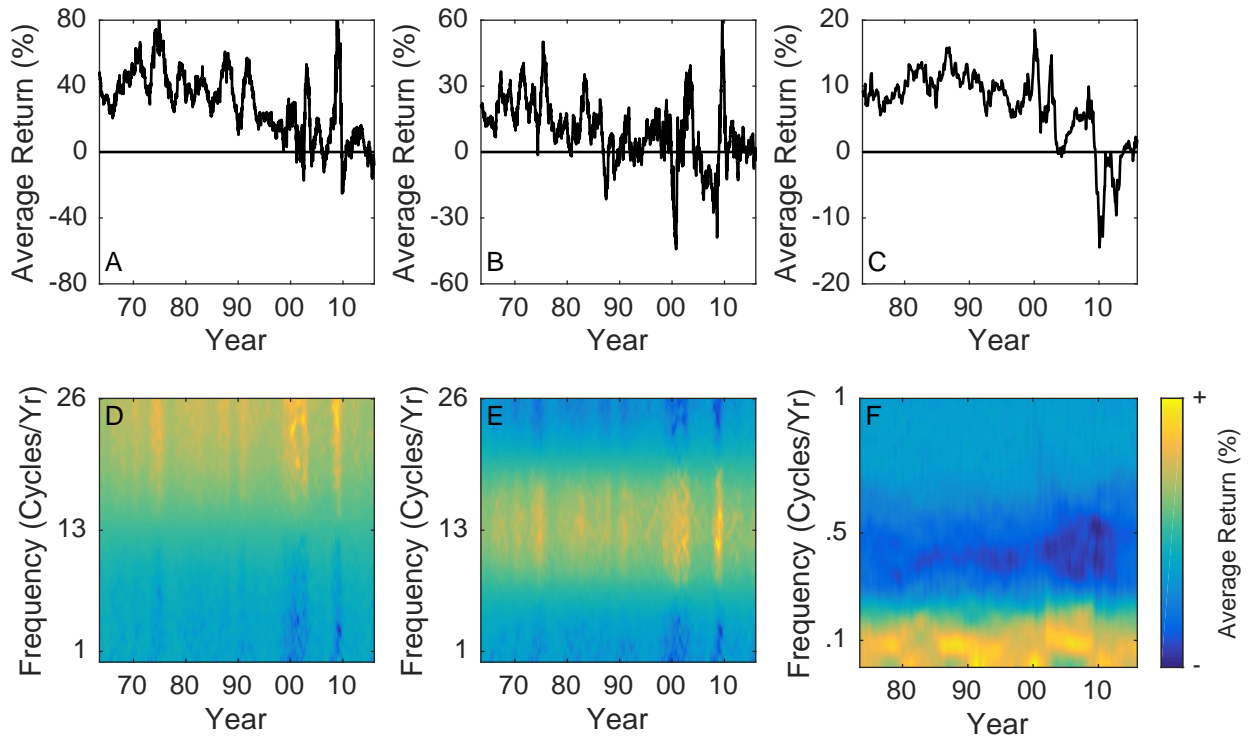


Figure 9: The 1-year rolling average of the mean reversion trading strategy applied to all S&P 500 stocks from 1964 to 2015 with $q = 1$ week and $q = 2$ weeks are shown in panels A and B, respectively. Their corresponding F decompositions are displayed in panels D and E. The 10-year rolling average of the calendar year momentum strategy applied to the same dataset is plotted in panel C. Its corresponding F decomposition is shown in Panel F.

apply the F decomposition to decompose these average returns into their frequency components. As expected, we see that the value-added for both these strategies occurs from active management at the targeted time horizons. Conversely, a negative risk premium subtracts value from the average portfolio returns at the passive and low frequencies. Moreover, their non-overlapping profitability bands subject them to diverse market dynamics, resulting in a correlation between their annual returns over the sample period of only 0.46. Notice that a large component of this correlation results from their decreasing profitability over time, a trend driven by increased competition and greater market efficiency. In addition, as described in Section 4.1, we see that periods of increased volatility, such as the early and late 2000s, amplify the contribution to the average portfolio return at each frequency.

5.2 Momentum

Panel C analyzes the 10-year rolling average of the monthly returns of our momentum trading strategy, which consists of buying the winners and selling the losers from the previous calendar year. Specifically the securities in the top decile of returns from the previous year are bought, and the securities in the bottom decile are sold. These equally weighted positions are held for one year and rebalanced each month such that the portfolio has no net position, and the Reg T requirements are satisfied. Notice that this strategy's profitability also decreased over time, and that it suffered heavy losses during the Financial Crisis.

Panel F decomposes these average returns into their frequency components. In general, this strategy earns profits at very low frequencies (less than 1 cycle per 2 years), yet performs poorly in response to oscillations on the order of 1 cycle per 2 years, which tend to move opposite to the strategy's weights. We see that, during the Financial Crisis and subsequent recovery, reversals on the order of 2 years would have caused momentum strategies operating at these frequencies to suffer severe losses. On the other hand, the one- and two-week mean reversion trading strategies were robust to these dynamics, yet were sensitive to changes in other market fluctuations.

This analysis suggests that, for strategies with investment power at specific timescales, we may consider diversifying not only across assets, but also across the frequency components of trading strategies. As was demonstrated for the momentum trading strategy, market dislocations can be isolated to certain frequency bands, and therefore, a portfolio with its returns

spread over multiple frequencies may diversify both idiosyncratic and systemic sources of risk. Since these time-horizon-specific strategies can be implemented contemporaneously, they can be viewed as separate assets with varying risk-reward characteristics and correlations. Similar to the concepts of mean-variance optimization and risk parity, one could then consider allocating risk and capital across different frequency bands. In fact, the frequency band-limited counterparts to alpha, beta, volatility, and correlation described herein can be applied to almost any theory of risk, reward, and portfolio construction.

6 Conclusion

In this article, we have applied spectral analysis to develop dynamic measures of volatility, correlation, beta, and alpha. These frequency-specific measures allow us to distinguish between short- and long-term components risk and covariances, providing additional insights into portfolio and risk management above and beyond their static counterparts. These considerations are particularly useful when portfolio goals differ across time horizons, and when investors wish to target specific horizons because of their preferences and life cycle.

We have also developed a technique—the F decomposition—that allows us to determine whether portfolio managers are capturing alpha and over what time horizons their investment processes have forecast power. In this context, an investment process is said to be profitable at a given frequency if there is positive correlation between portfolio weights and returns at that frequency. When aggregated across frequencies, the F decomposition is equivalent to the AP decomposition, and provides a clear indication of a manager’s forecast power and, consequently, active investment skill. Moreover, the F decomposition can identify alpha leakages in an investment process and suggest possible methods for improving performance.

Frequency-dependent alphas, betas, variances and auto- and cross-covariances can be used to incorporate dynamics into many other financial applications. For example, dynamic versions of performance attribution, linear factor models, asset-allocation models, risk management, and measures of systemic risk can all be constructed using spectral analysis. Finally, our DFT-based framework can be extended to other time-frequency decompositions including the wavelet transform so as to address the impact of time-varying frequencies and other nonstationarities. We hope to explore these extensions in future research.

A Appendix

In this Appendix, we derive statistical properties of the main estimators in the paper that are required for conducting standard inferences such as hypothesis tests and significance-level calculations.

A.1 General Moment Properties of the Power Spectrum

Consider the real-valued discrete-time wide-sense stationary stochastic processes $\{x_t\}$ and $\{y_t\}$ with means m_x and m_y , and cross-covariance function $\gamma_{xy}[m] = \text{E}[(x_{t+m} - m_x)(y_t - m_y)]$. Assuming the cross-covariance function has finite energy, let $P_{xy}(\omega)$ be its Discrete-Time Fourier Transform (DTFT) such that,

$$P_{xy}(\omega) = \sum_{m=-\infty}^{\infty} \gamma_{xy}[m]e^{-j\omega m}, \quad (\text{A.1})$$

$$\gamma_{xy}[m] = \frac{1}{2\pi} \int_0^{2\pi} P_{xy}(\omega)e^{j\omega m} d\omega. \quad (\text{A.2})$$

The function $P_{xy}(\omega)$ is known as the cross spectrum, and can be interpreted as the frequency distribution of the power contained in the covariance between x_t and y_t . A rectangular window of length T can be used to select a finite-length subsample of x_t and y_t . Forming the cross spectrum estimate from the DFT of this finite subsample we find that $\text{E}[C_{xy}[k]]$ is not generally equal to $P_{xy}(\omega_k)$, where $\omega_k = 2\pi k/T$, and is therefore a biased estimator. The bias results from the convolution of the true power spectrum, $P_{xy}(\omega)$, with the DTFT of the aperiodic autocorrelation function of the window. As the window length increases, its DTFT approaches a Dirac delta function, and so the bias approaches 0. Thus, $\text{E}[C_{xy}[k]]$ is an asymptotically unbiased estimator of $P_{xy}(\omega_k)$ (Oppenheim and Schaffer, 2009). Moreover, over a wide range of conditions, it can be shown that,

$$\text{var}[L_{xy}[k]] \approx \frac{1}{2}(P_{xx}(\omega_k)P_{yy}(\omega_k) + \Lambda_{xy}^2(\omega_k) - \Psi_{xy}^2(\omega_k)), \quad (\text{A.3})$$

where $\Lambda_{xy}(\omega)$ and $\Psi_{xy}(\omega)$ are the theoretical co-spectrum and quadrature spectrum between x_t and y_t , respectively. At the harmonic frequencies, which are separated in frequency by $1/T$, these frequency-specific estimators of the co-spectrum are approximately uncorrelated (Jenkins and Watts, 1968). This property can be used to estimate the variance of the sum of co-spectrum estimators, $L_{xy}[k]$.

A few important implementation details still remain. Notice that the variance of the co-spectrum estimates are not consistent as they do not asymptotically approach 0 as T increases. Averaging the co-spectrum estimates calculated over overlapping time intervals can reduce the variance of the spectral estimates at the expense of introducing bias. In addition, windowing procedures (e.g., multiplication by a Hamming window) can be applied to the data before calculating the DFT. This procedure will generally decrease spectral leakage at the expense of reducing spectral resolution. An estimate of the co-spectrum can

also be calculated from the Fourier transform of the estimated cross-covariance function. Finally, if x_t and y_t are sampled at a low frequency relative to the rate at which their properties change, then the decomposition will be biased due to a phenomenon known as aliasing. See Oppenheim and Schaffer (2009) and Jenkins and Watts (1968) for a more detailed discussion of these advanced implementation techniques.

A.2 Confidence Intervals for Dynamic Correlations

Fisher's z -transformation is often required to produce a sample correlation coefficient estimator which is approximately normal, and whose shape is independent of the correlation coefficient parameter. For the estimated correlation coefficient $\hat{\rho}$, based on T independent sample pairs, the estimator,

$$z(\hat{\rho}) \equiv \frac{1}{2} \ln \left(\frac{1 + \hat{\rho}}{1 - \hat{\rho}} \right) = \operatorname{arctanh}(\hat{\rho}), \quad (\text{A.4})$$

is approximately normal with mean $\frac{1}{2} \ln \left(\frac{1 + \rho}{1 - \rho} \right)$, and standard error $\frac{1}{\sqrt{T-3}}$. An approximate $100(1 - \alpha)\%$ confidence interval for ρ_K based on the DFT is therefore,

$$\left[\tanh \left(z(\hat{\rho}_K) - \frac{\Phi^{-1}(1 - \frac{\alpha}{2})}{\sqrt{T_K - 3}} \right), \tanh \left(z(\hat{\rho}_K) + \frac{\Phi^{-1}(1 - \frac{\alpha}{2})}{\sqrt{T_K - 3}} \right) \right], \quad (\text{A.5})$$

where T_K is the number of DFT coefficients associated with the frequency band K . The number of DFT coefficients at harmonic frequencies in the band of interest is used because, under the assumptions of Fisher's z -transformation, the denominator should consist of the number of independent samples used in constructing the correlation coefficient (Whitcher, Guttorp, and Percival, 2000). As described in the above section, the assumption of approximately uncorrelated DFT coefficients is justified.

A.3 Standard Errors and F -Tests for Dynamic Betas

This section summarizes the relevant band-spectrum regression properties from Engle (1974). If we define \tilde{y} to be a $T \times 1$ column vector, and \tilde{x} to be a $T \times M$ matrix, where the k th rows are the DFT coefficients of Y_k and $[X_{1,k}, \dots, X_{M,k}]$, respectively, then an M -factor band-spectrum regression can be specified as:

$$A\tilde{y} = A\tilde{x}\beta + A\tilde{\varepsilon}, \quad (\text{A.6})$$

where A is a $T \times T$ matrix with ones on the diagonals corresponding to included frequencies and zeros elsewhere. Using this form, (15) can be rewritten as:

$$\beta_K = ((A\tilde{x})^\dagger(A\tilde{x}))^{-1}((A\tilde{x})^\dagger(A\tilde{y})) \quad (\text{A.7})$$

$$\text{var}(\beta_K) = ((A\tilde{x})^\dagger(A\tilde{x}))^{-1}\sigma^2, \quad (\text{A.8})$$

where “ \dagger ” denotes the conjugate transpose, and an estimator of σ^2 is given by:

$$A\tilde{u} = A\tilde{y} - A\tilde{x}\beta_K, \quad (\text{A.9})$$

$$\hat{\sigma}^2 = \frac{(A\tilde{u})^\dagger(A\tilde{u})}{T_K - M}. \quad (\text{A.10})$$

From these equations, the standard errors of β_K can be estimated. If a regression model that forces the β 's to fit all frequencies (the restricted model) holds, then a regression model that allows the β 's to differ across frequency bands (the unrestricted model) will be relatively inefficient. In this case, Fisher's F -test can be used to determine if the unrestricted model yields a significantly better fit. Assuming the same T' frequencies are used for both models, and letting \tilde{u} and \tilde{v} be the unrestricted model's and restricted model's residuals, respectively, the F -statistic is given by,

$$F = \frac{\left(\frac{v^\dagger v - u^\dagger u}{n_u - n_v}\right)}{\left(\frac{u^\dagger u}{T' - n_u}\right)}, \quad (\text{A.11})$$

where the unrestricted model has n_u parameters, and the restricted model has n_v parameters. Under the null hypothesis that the unrestricted model does not provide a significantly better fit than the restricted model, F will have an F distribution with $(n_u - n_v, T' - n_u)$ degrees of freedom.

References

- Anderson, T., 1971, *The Statistical Analysis of Time Series*. John Wiley & Sons, Inc., New York, NY.
- Bauwens, L., S. Laurent, and J. Rombouts, 2006, “Multivariate GARCH Models: a Survey,” *Journal of Applied Econometrics*, 21(1), 79–109.
- Baxter, M., and R. King, 1999, “Measuring business cycles: Approximate band-pass filters for economic time series,” *Review of Economics and Statistics*, 81(4), 575–593.
- Breitung, J., and B. Candelon, 2006, “Testing for short- and long-run causality: A frequency-domain approach,” *Journal of Econometrics*, 132(2), 363–378.
- Briggs, W., and V. Henson, 1995, *The DFT: An Owners’ Manual for the Discrete Fourier Transform*. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Brillinger, D., 2001, *Time Series: Data Analysis and Theory*. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Brockwell, P., and R. Davis, 1991, *Time Series: Theory and Methods*. Springer-Verlag, New York, 2nd edn.
- Broto, C., and E. Ruiz, 2004, “Estimation Methods for Stochastic Volatility Models: A Survey,” *Journal of Economic Surveys*, 18(5), 613–649.
- Campbell, J., A. Lo, and A. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.
- Carr, P., and D. Madan, 1999, “Option valuation using the fast Fourier transform,” *Journal of Computational Finance*, 2(4), 61–73.
- Chaudhuri, S., 2014, “A time and frequency domain analysis of contrarian trading strategies,” Master’s thesis, Massachusetts Institute of Technology.
- Chaudhuri, S., and A. Lo, 2015, “Spectral analysis of stock-return volatility, correlation, and beta,” *IEEE Signal Processing and Signal Processing Education Workshop (SP/SPE)*, pp. 232–236.
- Croux, C., M. Forni, and L. Reichlin, 2001, “A measure of comovement for economic variables: Theory and empirics,” *Review of Economics and Statistics*, 83(2), 232–241.
- Crowley, P., 2007, “A guide to wavelets for economists,” *Journal of Economic Surveys*, 21(2), 207–267.
- Engle, R., 1974, “Band Spectrum Regression,” *International Economic Review*, 15(1), 1–11.
- , 1982, “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 50(4), 987–1007.

- Granger, C., and R. Engle, 1983, “Applications of Spectral Analysis in Econometrics,” *Handbook of Statistics*, 3, 93–109.
- Granger, C., and M. Hatanaka, 1964, *Spectral Analysis of Economic Time Series*. Princeton University Press, Princeton, NJ.
- Hannan, E., 1970, *Multiple Time Series*. John Wiley & Sons, Inc., New York, NY.
- Hasbrouck, J., and G. Sofianos, 1993, “The trades of market makers: An empirical analysis of NYSE specialists,” *The Journal of Finance*, 48(5), 1565–1593.
- Huang, N., M. Wu, W. Qu, S. Long, S. Shen, and J. Zhang, 2003, “Applications of Hilbert-Huang transform to non-stationary financial time series analysis,” *Applied Stochastic Models in Business and Industry*, 19(3), 245–268.
- Jegadeesh, N., and S. Titman, 1993, “Returns to buying winners and selling losers: Implications for stock market efficiency,” *The Journal of Finance*, 48(1), 65–91.
- Jenkins, G. M., and D. G. Watts, 1968, *Spectral Analysis and Its Applications*. Holden-Day, San Francisco, CA.
- Jensen, M., 1968, “The performance of mutual funds in the period 1945–1964,” *The Journal of Finance*, 23(2), 389–416.
- , 1969, “Risk, the pricing of capital assets, and the evaluation of investment portfolios,” *The Journal of Business*, 42(2), 167–247.
- Khandani, A., and A. Lo, 2007, “What happened to the quants in August 2007?,” *Journal of Investment Management*, 5(4), 5–54.
- King, R., and M. Watson, 1996, “Money, prices, interest rates and the business cycle,” *Review of Economics and Statistics*, 78, 35–53.
- Linetsky, V., 2002, “Exotic Spectra,” *RISK*, pp. 85–89.
- , 2004a, “The Spectral Decomposition of the Option Value,” *International Journal of Theoretical and Applied Finance*, 7, 337–384.
- , 2004b, “Spectral Expansions for Asian (Average Price) Options,” *Operations Research*, 52, 856–867.
- , 2008, “Spectral Methods in Derivatives Pricing,” in *Handbooks in Operations Research and Management Science*, ed. by J. R. Birge, and V. Linetsky. Elsevier, Amsterdam, vol. 15, pp. 223–299.
- Lintner, J., 1965, “The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets,” *The Review of Economics and Statistics*, 47(1), 13–37.
- Lo, A., 2008, “Where do alphas come from?: A new measure of the value of active investment management,” *Journal of Investment Management*, 6(2), 1–29.

- Lo, A., and A. MacKinlay, 1990, "When are contrarian profits due to stock market overreaction?," *Review of Financial Studies*, 3(2), 175–205.
- Markowitz, H., 1952, "Portfolio Selection," *The Journal of Finance*, 7(1), 77–91.
- Merton, R., 1969, "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics*, 51, 247–257.
- , 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.
- , 1973, "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.
- Officer, R., 1973, "The Variability of the Market Factor of the New York Stock Exchange," *Journal of Business*, 46(3), 434–453.
- Oppenheim, A., and R. Schaffer, 2009, *Discrete-Time Signal Processing*. Prentice Hall, Upper Saddle River, NJ, 3rd edn.
- Phillips, P., Y. Sun, and S. Jin, 2006, "Spectral Density Estimation and Robust Hypothesis Testing Using Steep Origin Kernels without Truncation," *International Economic Review*, 47(3), 837–894.
- , 2007, "Long run variance estimation and robust regression testing using sharp origin kernels with no truncation," *Journal of Statistical Planning and Inference*, 137(3), 985–1023.
- Priestly, M., 1981, *Spectral Analysis and Time Series*. Academic Press, London-New York.
- Ramsey, J., 2002, "Wavelets in economics and finance: Past and future," *Studies in Nonlinear Dynamics & Econometrics*, 6(3), 1–27.
- Roll, R., 1984, "A simple implicit measure of the effective bid-ask spread in an efficient market," *The Journal of Finance*, 39(4), 1127–1139.
- Rua, A., 2010, "Measuring comovement in the time-frequency space," *Journal of Macroeconomics*, 32(2), 685–691.
- , 2012, "Wavelets in economics," *Economic Bulletin and Financial Stability Report Articles*, pp. 71–79.
- Samuelson, P. A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 239–246.
- Shao, X., and W. Wu, 2007, "Asymptotic spectral theory for nonlinear time series," *The Annals of Statistics*, 35(4), 1773–1801.
- Sharpe, W., 1964, "Capital asset prices: A theory of market equilibrium under conditions of risk," *The Journal of Finance*, 19(3), 425–442.
- , 1966, "Mutual fund performance," *The Journal of Business*, 39(1), 119–138.

- Treynor, J., 1965, “How to rate management of investment funds,” *Harvard Business Review*, 43(1), 63–75.
- Velasco, C., and P. Robinson, 2001, “Edgeworth Expansions for Spectral Density Estimates and Studentized Sample Mean,” *Econometric Theory*, 17(3), 497–539.
- Whitcher, B., P. Guttorp, and D. B. Percival, 2000, “Wavelet analysis of covariance with application to atmospheric time series,” *Journal of Geophysical Research*, 105, 941–962.
- Wu, W., and P. Zaffaroni, 2016, “Asymptotic Theory for Nonparametric Multivariate Spectral Density Estimates,” *Working Draft*.