

# MIT Open Access Articles

# *Dynamic Alpha: A Spectral Decomposition of Investment Performance Across Time Horizons*

The MIT Faculty has made this article openly available. *Please share* how this access benefits you. Your story matters.

**As Published:** 10.1287/MNSC.2018.3102

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: https://hdl.handle.net/1721.1/135075

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



# Dynamic Alpha: A Spectral Decomposition of Investment Performance Across Time Horizons<sup>\*</sup>

Shomesh E. Chaudhuri<sup> $\dagger$ </sup> and Andrew W. Lo<sup> $\ddagger$ </sup>

This Draft: May 22, 2018

#### Abstract

The value added by an active investor is traditionally measured using alpha, tracking error, and the information ratio. However, these measures do not characterize the dynamic component of investor activity, nor do they consider the time horizons over which weights are changed. In this paper, we propose a technique to measure the value of active investment that captures both the static and dynamic contributions of an investment process. This dynamic alpha is based on the decomposition of a portfolio's expected return into its frequency components using spectral analysis. The result is a static component that measures the portion of a portfolio's expected return due to passive investments and security selection, and a dynamic component that captures the manager's timing ability across a range of time horizons. Our framework can be universally applied to any portfolio, and is a useful method for comparing the forecast power of different investment processes. Several analytical and empirical examples are provided to illustrate the practical relevance of this decomposition.

**Keywords**: Investments; Performance Attribution; Portfolio Management; Alpha; Spectral Analysis.

<sup>\*</sup>We thank Tom Brennan, Hui Chen, Ian Dew-Becker, Paul Mende, and participants at the 2015 IEEE Signal Processing and Signal Processing Education Workshop, the MIT Sloan Finance Lunch, the London Quantitative Finance Seminar at Imperial College, and the 2017 Western Finance Association Conference 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.

<sup>&</sup>lt;sup>†</sup>Department of Electrical Engineering and Computer Science, and Laboratory for Financial Engineering, MIT

<sup>&</sup>lt;sup>‡</sup>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, alo-admin@mit.edu (email).

## Contents

Introduction	1
Literature Review	1
Dynamic Alpha	3
Alpha vs. Beta	6
Numerical Examples	8
An Empirical Example	12
Warren Buffett's Alpha	17
Conclusion	20
Spectral Analysis A.1 The Fourier Transform	<b>22</b> 22
A.2 The Power Spectrum	$\frac{23}{25}$
	Introduction         Literature Review         Dynamic Alpha         Alpha vs. Beta         Numerical Examples         An Empirical Example         Warren Buffett's Alpha         Conclusion         Spectral Analysis         A.1 The Fourier Transform         A.2 The Power Spectrum         A.3 General Moment Properties of the Power Spectrum

## 1 Introduction

The shortest decision interval of a modern investment strategy may range from microseconds to years, a wide span of time horizons. While the legendary value investor Warren Buffett tends to change his portfolio weights rather slowly, the same cannot be said for famed day trader Steven Cohen of SAC Capital, yet both manage to generate enormous value through active investment. Although alpha, tracking error, and the information ratio are the standard tools for gauging the value-added of a portfolio manager, they can obscure important features of the underlying process by which information is reflected in investment decisions. Specifically, none of these standard performance metrics directly measure the dynamic relationship between weights and returns, which is the central focus of active investment strategies.

In this paper, we propose a new approach to analyzing investment strategies 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 the sine and cosine functions—we show that investment strategies can differ significantly in the frequencies with which their expected returns 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, an investor will have an additional dimension with which to manage the risk/reward profile of his portfolio.

We begin in Section 2 with a brief review of the financial spectral analysis literature. Our main results are contained in Sections 3 and 4, where we provide spectral decompositions for an investment strategy's forecast power. We provide numerical and empirical illustrations of these techniques in Sections 5–7, and conclude in Section 8.

## 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 the Fourier transform has been used in finance to efficiently evaluate theoretical pricing models for derivative securities (Carr and Madan, 1999). However, econometric and empirical applications of

spectral analysis in economics and finance are less common, in part because economic time series are rarely considered stationary. Recently, there has been a rebirth of interest in their application to economics in response to modern advances in non-stationary signal analysis (Baxter and King, 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; Dew-Becker and Giglio, 2016; Bandi, Perron, Tamoni, and Tebaldi, 2017). This rebirth motivates our interest in the spectral properties of financial asset returns.

In this article, we show that spectral analysis can be used to characterize and refine active investment strategies. The standard tools used for performance attribution originated from the Capital Asset Pricing Model (CAPM) 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 that quantifies the predictive power of an investment process by decomposing the expected portfolio return into the covariance between the underlying security weights and returns and the product of the average weights and average returns. 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.

As an extension of this decomposition, we introduce the concept of *dynamic alpha*, 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 a positive correlation between portfolio weights and returns at that frequency. When aggregated across frequencies, dynamic alpha is equivalent to Lo's (2008) active component, 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, allowing us to study the time-horizon properties of investment performance.

To address the non-stationarity of financial time series, our analysis relies on the shorttime Fourier transform, which applies the discrete Fourier transform (DFT) to windowed subsamples of the entire sample (Oppenheim and Schafer, 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. These techniques can provide substantial benefits in practice. For example, the sinusoids used in the short-time Fourier transform do not efficiently characterize discontinuous processes, but 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 simpler in exposition, and all our results for the Fourier transform carry over directly to the wavelet transform (albeit with greater mathematical and expositional complexity).

## 3 Dynamic Alpha

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 decomposition developed by Lo (2008), we use the DFT 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 one component that measures the portion of a portfolio's expected return due to passive investments and active security selection, and multiple dynamic 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, ..., T-1, the average one-period portfolio return can be calculated as,

$$\overline{r_p} = \frac{1}{T} \sum_{i=1}^{N} \sum_{t=0}^{T-1} w_{i,t} r_{i,t},$$
(1)

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, the average portfolio return can be decomposed into a dynamic alpha component  $(\delta_p)$  and a static component  $(\nu_p)$  as follows,

$$\overline{r_p} = \delta_p + \nu_p,\tag{2}$$

$$\delta_p = \sum_{i=1}^{N} \operatorname{Cov} \langle w_{i,t}, r_{i,t} \rangle , \quad \nu_p = \sum_{i=1}^{N} \overline{w_{i,t}} \cdot \overline{r_{i,t}} .$$
(3)

The value of the static 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, as well as the ability to select securities with favorable long-term prospects. This distinction contrasts with Lo's (2008) use of the term "passive" for the static component—in our setting, we wish to acknowledge the possibility that active management is responsible for long-term bets on specific securities, in which case a portion of a portfolio's static component may, in fact, be alpha rather than risk premia.

The value of the dynamic alpha component consists of 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.

Spectral analysis allows us to decompose this covariance term further, capturing the

manager's timing ability over multiple time horizons,

$$\delta_p = \sum_{k=1}^{T-1} \delta_{p,k} \quad , \quad \delta_{p,k} = \frac{1}{T^2} \sum_{i=1}^N \Re[W_{i,k}^* R_{i,k}] \quad , \tag{4}$$

where  $\Re[z]$  and  $z^*$  denote the real part and complex conjugate of a complex number z, respectively, and  $W_{i,k}$  and  $R_{i,k}$  are the *T*-point DFT coefficients (see Section A.1) of the weights and returns for stock *i*. In this form, the contribution to the average portfolio return by the *k*th harmonic frequency, where  $k \in \{0, ..., T - 1\}$ , is clearly visible. The lowest frequency occurs at k = 0, and the highest frequency occurs at the value of *k* closest to T/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 average portfolio return are equivalent. The relation  $h = TT_s/k$ , where  $T_s$  is the time between samples and  $0 \le k \le T/2$ , can be used to convert the *k*th harmonic frequency to its corresponding time horizon, *h*. We also note that  $\delta_{p,0} = \nu_p$ , and it is often convenient to include  $\delta_{p,0}$  when computing the DFT.

Simply put, 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. For this reason, a value-weighted portfolio of all securities, which is traditionally considered passive, will contain no dynamic alpha across all frequencies as long as the individual security returns are serially uncorrelated (i.e., the Random Walk Hypothesis holds for all securities). On the other hand, if returns are serially correlated, then it is possible for a buy-and-hold portfolio to yield a non-zero dynamic alpha because changes in its weights will contain information related to future returns. To distinguish between dynamically managed alpha and passive portfolios that unintentionally contain non-zero dynamic alpha, we must therefore rely on the manager's stated intentions.

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. The ratio of alpha to 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 *dynamic risk*,  $\sigma_{\delta}$ , as the variability of the difference between the portfolio return,  $r_{p,t}$ , and the static component,  $\nu_{p,t} = \sum_{i=1}^{N} \overline{w_{i,t}} \cdot r_{i,t}$ . Specifically,

$$\sigma_{\delta} = \sqrt{\operatorname{Var}\langle r_{p,t} - \nu_{p,t} \rangle},\tag{5}$$

where  $\sigma_{\delta}$  is a measure of the risk taken by the portfolio manager in an attempt to generate higher returns by engaging in timing decisions. The dynamic information ratio,  $I_{\delta}$ , can then be defined as,

$$I_{\delta} = \frac{\delta_p}{\sigma_{\delta}},\tag{6}$$

and is a risk-adjusted measure of the dynamic alpha component. These performance metrics can be calculated for a specific range of time horizons by aggregating the frequency components of  $\delta_p$  and  $\sigma_{\delta}$  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 Alpha vs. Beta

To distinguish explicitly between manager outperformance and portfolio exposure to factor risk, we have to impose additional structure on the returns of the individual assets. Specifically, we consider a linear M-factor model,

$$x_{i,t} = \alpha_i + \beta_{i,1}F_{1,t} + \dots + \beta_{i,M}F_{M,t} + \varepsilon_{i,t} , \qquad (7)$$

where  $x_{i,t}$  is defined to be the excess return of asset *i*, in excess of the risk-free rate of return,  $r_{f,t}$ ,  $F_{m,t}$  are excess factor returns, and  $E[\varepsilon_{i,t} | F_{1,t}, \ldots, F_{M,t}] = 0$ . This specification is consistent with Merton's (1973) Intertemporal Capital Asset Pricing Model and Ross's (1976) Arbitrage Pricing Theory. Since our expected-return decomposition is considerably more general than any particular asset-pricing model or linear-factor structure, we allow for an intercept,  $\alpha_i$ , in our framework.

Under these assumptions, the portfolio's exposure to factor m is  $\beta_{p,m,t} = \sum_{i=1}^{N} w_{i,t}\beta_{i,m}$ . The average return of a portfolio of assets (2) can then be rewritten as,

$$\overline{r_p} = \underbrace{\text{Risk-Free Rate} + \text{Risk Premia} + \text{Security Selection}}_{\nu_p \equiv \text{Static Component}} + \underbrace{\text{Factor Timing} + \text{Security Timing}}_{\delta_p \equiv \text{Dynamic Component}}$$

where,

Risk-Free Rate 
$$\equiv \overline{r_{f,t}}$$
 (9)

Risk Premia 
$$\equiv \sum_{m=1}^{M} \overline{\beta_{p,m,t}} \cdot \overline{F_{m,t}}$$
 (10)

Security Selection 
$$\equiv \sum_{i=1}^{N} \overline{w_{i,t}} \cdot \alpha_i$$
 (11)

Factor Timing 
$$\equiv \sum_{m=1}^{M} \text{Cov}\langle \beta_{p,m,t}, F_{m,t} \rangle$$
 (12)

Security Timing 
$$\equiv \sum_{i=1}^{N} \operatorname{Cov} \langle w_{i,t}, \varepsilon_{i,t} \rangle$$
 (13)

Due to the structure of the linear multi-factor model, (8) is a more refined decomposition than (2). The average portfolio returns are now the sum of five components: a risk-free rate component, a risk-premia component that represents the return from the passive exposures to factor risk, a security selection component that depends on the  $\alpha_i$ 's, a factor-timing component that depends on the covariance between the portfolio's factor exposures and the underlying factors, and finally, a security-timing component that depends on the covariance between weights and the idiosyncratic component of security returns. Note that the factorand security-timing terms can be decomposed further into their frequency components.

This factor-based decomposition demonstrates that investment expertise can manifest itself in two distinct ways: identifying cheap sources of expected return (i.e., the  $\alpha_i$ 's, which are reflected in the static component,  $\nu_p$ ), and creating additional expected return through factor- and security-specific timing across different time horizons (i.e., the covariance terms, which are reflected in the spectral decomposition of the dynamic component,  $\delta_{p,k}$ ). Thus, even if all  $\alpha_i$ 's are zero, as most asset-pricing models claim, there can still be substantial value-added from active management if the investment process has the ability to time price movements over certain time horizons.

#### 5 Numerical Examples

To develop additional intuition for our spectral decomposition, we extend the following simple numerical example provided by Lo (2008). Consider 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 1 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 static component. In this case, because the weights are constant, the dynamic risk measure will also be 0%.

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 2). 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., the highest frequency). In addition, the dynamic risk for this portfolio is 0.3375%, and the dynamic information ratio is about 0.37.

Finally, consider a third portfolio A3, which also has alternating weights for Asset 1, but is 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 3 confirms that this is counterproductive, as Portfolio A3 loses 0.1250% per month from its highest frequency

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%	<b>1.1625</b> %		

Table 1: The expected return of a constant portfolio depends only on the static component.

	-	$r_1$	$w_2$	$r_2$	$r_p$			
	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%			

	$ u_p$	$2\delta_{p,1}$	$2\delta_{p,2}$	$2\delta_{p,3}$	$2\delta_{p,4}$	$2\delta_{p,5}$	$\delta_{p,6}$
-	1.1625%	0%	0%	0%	0%	0%	0.1250%

Table 2: 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  $(\delta_{p,6})$ .

Month	$w_1$	$r_1$	$w_2$	$r_2$	$r_p$
		Strate	egy 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%	<b>1.0375</b> %

component, and its total expected return is only 1.0375%. In this case, the dynamic risk is 0.3375%, and the dynamic information ratio is -0.37.

Spectral decomposition of $\overline{r_p}$							
$ u_p$	$2\delta_{p,1}$	$2\delta_{p,2}$	$2\delta_{p,3}$	$2\delta_{p,4}$	$2\delta_{p,5}$	$\delta_{p,6}$	
1.1625%	0%	0%	0%	0%	0%	-0.1250%	

Table 3: 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  $(\delta_{p,6})$ .

Note that in all three cases, the static 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. These differences give rise to different values for the highest frequency component. As shown in (4), contributions from higher frequencies (k > 0) sum to the overall dynamic component. These higher frequency contributions can be interpreted as the portion of the dynamic 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):

$$w_{i,t} = -\frac{1}{N} (r_{i,t-1} - r_{m,t-1}), \tag{14}$$

$$r_{m,t-1} = \frac{1}{N} \sum_{i=1}^{N} r_{i,t-1} \quad .$$
(15)

By buying the losers and selling the winners from date t-1 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, that is, 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 that value will be added near frequencies inversely related to the mean reversion period.

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

$$r_{i,t} = \varepsilon_{i,t} + \lambda \varepsilon_{i,t-1}, \tag{16}$$

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

$$\mathbf{E}[r_p] = -\lambda \sigma^2 (1 - \frac{1}{N}) \quad . \tag{17}$$

We see that the expected return is proportional to the mean reversion factor,  $\lambda$ , and the volatility factor,  $\sigma^2$ . Applying our spectral decomposition (see Section A.2), we find that,

$$\delta_{p,\omega} = -\sigma^2 \left( 1 - \frac{1}{N} \right) \left( \lambda \cos(2\omega) + (1 + \lambda^2) \cos(\omega) + \lambda \right), \quad \omega \in [0, 2\pi).$$
(18)

The relation  $h = 2\pi T_s/\omega$ , where  $T_s$  is the time between samples and  $\omega \in [0, \pi]$ , can be used to convert frequency  $\omega$  to its corresponding time horizon, h.

Column A of Figure 1 plots the dynamic alpha for the case of no serial correlation ( $\lambda = 0$ ).

The dynamic alpha is positive at high frequencies, indicating that the weights and returns are in phase over these short time horizons. However, this added value is cancelled out since the weights and returns are out of phase at longer time horizons, resulting in zero net alpha.

Columns B and C of Figure 1 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 is an intuitive result 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 opposite in effect. 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.

#### 6 An Empirical Example

To develop a better understanding of the characteristics of dynamic alpha, we apply our framework to Lo and MacKinlay's (1990) contrarian (mean reversion) trading strategy using historical stock market data. The fact that the weights given by (14) sum to zero at each date t implies very little market-beta exposure. Also, since the weights are so dynamic, much of this portfolio's return should be due to active management near frequencies inversely related to the decision period. The 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; therefore, the amount of capital required is one-half the total capital invested (often stated as a 2:1 leverage, or a 50% margin requirement). The unleveraged portfolio return,  $r_{p,t}$  is given by:

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

We apply (14) to the one-day and two-day returns of the five smallest size-decile portfolios of all NASDAQ stocks, as constructed by the University of Chicago's Center for Research in



Figure 1: Dynamic alpha of the contrarian trading strategy applied to the serially uncorrelated (Column A), momentum (Column B), and mean reversion (Column C) implementations of (16).

Security Prices (CRSP), from January 2, 1990 to December 29, 1995. We selected this time period purposely because of the emergence of day trading in the early 1990s, an important source of profitability for statistical arbitrage strategies. Of course, trading NASDAQ size deciles is obviously unrealistic in practice, but our purpose is to illustrate the empirical relevance of our framework, not to derive an implementable trading strategy.



Figure 2: Cumulative return of the mean reversion strategy of Lo and MacKinlay (1990) over one-day and two-day returns applied to the five smallest CRSP-NASDAQ size deciles from January 2, 1990 to December 29, 1995.

Figure 2 illustrates the performance of the contrarian strategy for one-day and two-day mean reversion over the 1990–1995 sample period, and Table 4 contains summary statistics for the daily returns of the two trading strategies. For 1-day mean reversion, with an annualized average return of 31.6% and standard deviation of 7.9%, the strategy's performance is considerably better than that of a passive buy-and-hold strategy, which is one indication that active management is playing a significant role in this case.

This intuition is confirmed by the decomposition of the strategy's expected return into its dynamic alpha components in Table 5. On an annualized basis, the dynamic component yields 32.2%, which exceeds the strategy's total expected return of 31.6%, implying a slightly

Statistic	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	1-day	2-day
Mean $\times 250$	27.5%	17.4%	13.9%	13.7%	12.8%	31.6%	13.3%
$SD \times \sqrt{250}$	12.2%	9.8%	8.9%	9.1%	9.5%	7.9%	7.8%
$SR \times \sqrt{250}$	2.25	1.77	1.56	1.51	1.35	3.98	1.69
Min	-2.9%	-2.7%	-2.7%	-3.3%	-3.5%	-2.2%	-5.2%
Median	0.1%	0.1%	0.1%	0.1%	0.1%	0.1%	0.0%
Max	6.7%	3.6%	2.0%	2.1%	2.3%	2.4%	1.7%
Skew	0.6	0.1	-0.5	-0.7	-0.9	-0.1	-0.8
XSKurt	5.1	2.4	2.0	3.1	3.9	1.7	8.9

Table 4: Summary statistics of the daily returns of the one-day and two-day mean reversion strategies of Lo and MacKinlay (1990) applied to the daily returns of the five smallest CRSP-NASDAQ size deciles, from January 2, 1990 to December 29, 1995. The Sharpe ratio (SR) is calculated relative to a 0% risk-free rate.

negative static component. In this case, more than all of the strategy's expected return is coming from active management over a daily time horizon, and the low-frequency components are subtracting value.

The explanation for this rather unusual phenomenon was provided by Lo and MacKinlay (1990), who observed that because the contrarian strategy is, on average, long losers and short winners, it will typically be long the low-mean assets and short the high-mean assets. Therefore, the static component, i.e., the sum of average portfolio weights multiplied by average returns, will consist of positive average weights for low-mean stocks and negative average weights for high-mean stocks for this strategy—a losing proposition in the absence of mean reversion. Fortunately, the positive correlation between weights and returns at high frequencies is more than sufficient to compensate for this long-term negative component.

To mitigate the loss caused by the static component, we can filter out the trend component of each size-decile portfolio before calculating the mean-reversion weights. Intuitively, the mean-reversion trading strategy will no longer place a negative bias on the weights of the smallest deciles simply because they achieve relatively large average returns. Similarly, if we perfectly filter out the low-frequency dynamics of the portfolio returns, then we can extract the profitability in the high-frequency component of returns, while not suffering the substantial losses of the low-frequency component. In other words, the mean-reversion trading strategy will be trading on the relevant high-frequency signal, and not the lowfrequency "noise." Since a perfect high-pass filter cannot be implemented in practice, these low-frequency components would have to be forecasted. Therefore, rather counterintuitively, our spectral framework reveals that forecast power at low frequencies can be used to improve the overall performance of a high-frequency trading strategy.

Statistic	1-day	2-day
Portfolio Mean $\times 250$	31.6%	13.3%
Static Component $\times 250$	-0.6%	-1.0%
Dynamic Component $\times 250$	32.2%	14.2%
Low Frequency $(h \ge 5d)$	-44.7%	-19.1%
Med Frequency $(3d \le h < 5d)$	6.3%	33.7%
High Frequency $(h < 3d)$	70.6%	-0.4%

Table 5: Estimates of the dynamic alpha of the daily returns of the one-day and two-day mean reversion strategies of Lo and MacKinlay (1990) applied to the five smallest CRSP-NASDAQ size-decile returns, from January 2, 1990 to December 29, 1995. Frequency components are grouped into three categories: high frequencies (more than one cycle per three days), medium frequencies (between one cycle per three days and one cycle per week), and low frequencies (less than one cycle per week).

For mean reversion over two days, with an annualized average return of 13.3% and a standard deviation of 7.8%, the strategy's performance is considerably worse than that of the one-day mean reversion strategy. Active management is playing a significant but less productive role. Here, the positive correlation between weights and returns at medium frequencies remains sufficient to compensate for the negative correlation between weights and returns at the low and high frequencies.

The correlation of these two strategies' returns is only 0.26. This low correlation can be attributed to the fact that their performance is determined by market dynamics occurring in distinct and non-overlapping frequency bands. Moreover, these frequency-specific strategies can be implemented simultaneously, and can therefore be viewed as separate assets. These assets can then be combined in a portfolio to achieve diversification across multiple frequencies. In our sample period, these diversification benefits result in the Sharpe ratio being maximized when 84.6% of our capital is used to implement the one-day mean-reversion trading strategy, and the remaining capital is used to trade with mean reversion over two days. However, Table 5 makes it clear that both assets in our portfolio would be negatively affected by a low-frequency market shock.

## 7 Warren Buffett's Alpha

For a more realistic application of our dynamic alpha framework, we examine the returns of Warren Buffett's multinational conglomerate holding company, Berkshire Hathaway Inc., which is known for its long-term investments in public and private companies. We obtain quarterly holdings data for Berkshire Hathaway from Thomson Reuters Institutional (13F) Holdings database (based on Berkshire's SEC filings) from 1980 to 2013, and stock return data from the CRSP Monthly Stock database.

One consequence of Warren Buffett's longer decision interval is that we are less likely to be affected by aliasing when applying our decomposition to quarterly weights and returns of his portfolio; the same cannot be said for higher-frequency trading strategies. Figure 3 displays the cumulative returns for Berkshire Hathaway (BRK) and a simulated reconstruction (R[BRK]) of these returns using the holdings data based on SEC filings. The correlation between these return series is 0.7, and their Sharpe ratios are 0.69 and 0.66, respectively. The high correlation and similar Sharpe ratios indicate that the reconstructed returns capture a significant fraction of Berkshire Hathaway's price dynamics. Equating the mean of the reconstructed returns with the realized returns, we use a leverage ratio of 1.41 to reconstruct Warren Buffett's levered returns (RL[BRK]). This is similar to the average leverage ratio of 1.4 estimated by Frazzini, Kabiller, and Pedersen (2013) using total assets to equity.

Table 6 contains summary statistics for the monthly returns of each time series. With an average annualized return of 22.9% over more than 30 years, Berkshire clearly has positive alpha when compared to traditional risk factors. Frazzini, Kabiller, and Pedersen (2013) find that Buffett's returns are due more to security selection than his effect on management, which suggests that a large component of his returns must be static alpha, i.e., high average weights on securities with large  $\alpha_i$ 's. In other words, Buffett is able to select securities that provide high average returns above and beyond the expected return resulting from passive exposures to factor risk. Moreover, if Warren Buffett has a positive long-term effect on returns due to his managerial and advisory competence, then we would also expect to find a substantial component of his returns derived from lower frequencies. Finally, Buffett is a practitioner of value investing, and so we should not expect to find a significant correlation between his portfolio weights and returns at high frequencies.



Figure 3: Cumulative realized returns of Berkshire Hathaway (BRK) and a simulated reconstruction (R[BRK]) using holdings data for Berkshire Hathaway from Thomson Reuters Institutional (13F) Holdings database (based on Berkshire's SEC filings) from 1980 to 2013. Equating the mean of the reconstructed returns with the realized returns, we use a leverage ratio of 1.41 to reconstruct the levered returns (RL[BRK]).

Statistic	Risk-Free	Market	BRK	R[BRK]	RL[BRK]
Mean $\times 4$	4.7%	12.8%	22.9%	16.3%	22.9%
$SD \times \sqrt{4}$	1.7%	17.4%	26.2%	17.5%	24.7%
$SR \times \sqrt{4}$	0	0.47	0.69	0.66	0.74
Min	0.0%	-23.7%	-30.1%	-30.9%	-43.6%
Median	1.2%	3.9%	4.4%	4.2%	6.0%
Max	3.8%	21.3%	46.1%	28.8%	40.7%
Skew	0.6	-0.6	0.3	-0.5	-0.5
XSKurt	0.3	0.5	0.9	1.8	1.8

Table 6: Summary statistics of the quarterly returns of the one-month Treasury Bill (Risk-Free) rate, the value-weighted CRSP market index (Market), Berkshire Hathaway (BRK), and a simulated reconstruction (R[BRK]) using holdings data for Berkshire Hathaway from Thomson Financial Institutional (13F) Holdings Database (based on Berkshire's SEC filings) from 1980 to 2013. Equating the means of the reconstructed returns with the realized returns we use a leverage ratio of 1.41 to reconstruct the levered returns (RL[BRK]).

The decomposition of Berkshire Hathaway's reconstructed average portfolio return into its dynamic alpha components in Table 7 confirms this intuition. The static component yields an annualized return of 18.9%. In comparison, the value-weighted CRSP market index yielded an average annualized return of 12.8% over the same interval, and the annualized risk-free interest rate (one-month Treasury Bill rate) was 4.7%. The static component of the portfolio's realized market beta over this interval using quarterly returns was 0.84, which implies a risk premium component of 6.8% and a static alpha component of 7.3%. This demonstrates that a substantial component of Berkshire Hathaway's returns results from Buffett's ability to select securities with favorable long-term prospects. The dynamic alpha component contributes an additional annualized return of 4.1% to the portfolio, most of which can be attributed to dynamics occurring at time horizons greater than 5 years. The annualized dynamic risk is 10.3%, which yields a dynamic information ratio of 0.40. This result can be attributed to Buffett's ability as a manager to improve firm performance over the long run while Berkshire maintains a position in the company, and also to his ability to time transactions based on fundamental valuations.

In contrast, the dynamics at the shortest time horizons—less than 18 months—subtract 1.2% annually from the average portfolio return. Here, the negative correlation between weights and returns can be attributed in part to transaction costs and market impact. However, the quarterly sampling frequency of the holdings data restricts our ability to study these higher frequency dynamics. By observing only quarter-end weights and cumulative returns, we have no way of inferring the profitability of dynamics occurring at these higher frequencies.

A spectral decomposition of Berkshire Hathaway's returns demonstrates conclusively that Buffett is not only a consummate long-term investor, but that the horizon of his timing ability stretches far beyond the reaches of most other portfolio managers.

#### 8 Conclusion

In this article, we have applied spectral analysis to develop a dynamic measure of alpha that allows us to determine whether portfolio managers are generating alpha and over what time horizons their investment processes have forecast power. In this context, an investment

Statistic	RL[BRK]
Portfolio Mean $\times 4$	22.9%
Static Component $\times 4$	18.9%
Risk-Free Rate	4.7%
Risk Premium	6.8%
Static Alpha	7.3%
Dynamic Component $\times 4$	4.1%
Low Frequency $(h \ge 5y)$	4.3%
Med Frequency $(1.5y \le h < 5y)$	1.1%
High Frequency $(h < 1.5y)$	-1.2%

Table 7: Estimates of the static and dynamic alpha of the simulated quarterly returns of Berkshire Hathaway using holdings data for Berkshire Hathaway from Thomson Financial Institutional (13F) Holdings Database (based on Berkshire's SEC filings) from 1980 to 2013. Frequency components are grouped into three categories: high frequencies (more than one cycle per 1.5 years), medium frequencies (between one cycle per 1.5 years and one cycle per five years), and low frequencies (less than one cycle per five years). Note that table entries may not sum due to rounding.

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, dynamic alpha is equivalent to Lo's (2008) active component, and provides a clear indication of a manager's forecast power and, consequently, active investment skill. By separating the dynamic and static components of a portfolio, it should be possible to study and improve the performance of both.

Frequency-domain representations of auto- and cross-covariances can be applied to many other financial statistics in addition to alpha. 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. Our framework can also be extended to other time-frequency decompositions, including the wavelet transform, to address the impact of time-varying relationships and other non-stationarities.

#### A 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. We provide the formulation of the DFT in Section A.1. We then present the main mathematical results on the co-spectrum in Section A.2, and 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 in Section A.3.

#### A.1 The Fourier Transform

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,

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

where the frequency  $\omega$  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,

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

or in polar form as,

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

where  $\Re[X(\omega)]$  and  $\Im[X(\omega)]$  are the real and imaginary parts 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  $t = 0, \ldots, T-1$ , the *T*-point DFT is given by:<sup>1</sup>

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

<sup>&</sup>lt;sup>1</sup>In 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.

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

$$x_{t} = \frac{1}{T} \sum_{k=0}^{T-1} \left[ \Re[X_{k}] \cos(\omega_{k}t) - \Im[X_{k}] \sin(\omega_{k}t) \right], \quad t \in [0, T-1]$$
(A.5)

or in polar form as,

$$x_t = \frac{1}{T} \sum_{k=0}^{T-1} |X_k| \cos(\omega_k t + \angle X_k), \quad t \in [0, T-1].$$
 (A.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  $0 \le k \le T/2$ , can be used to convert the kth harmonic frequency to its corresponding time horizon.

Since the Fourier transform 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 the data comes from low- versus high-frequency fluctuations.

#### A.2 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  $\gamma_{xy}[m] = \mathbb{E}[(x_{t+m} - m_x)(y_t - m_y)]^2$ . Assuming the cross-covariance function has finite energy, let  $P_{xy}(\omega)$  be its DTFT,

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

<sup>&</sup>lt;sup>2</sup>Specifically, the stochastic processes  $\{x_t\}$  and  $\{y_t\}$  are said to be wide-sense stationary if and only if  $E[x_t]$  and  $E[y_t]$  are constants independent of t, and  $E[x_{t_1}x_{t_2}]$ ,  $E[y_{t_1}y_{t_2}]$  and  $E[x_{t_1}y_{t_2}]$  depend only on the time difference  $(t_1 - t_2)$ .

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  $P_{xy}(\omega)$ ,

$$\operatorname{Cov}(x_t, y_t) \equiv \gamma_{xy}[0] = \frac{1}{2\pi} \int_0^{2\pi} \Re[P_{xy}(\omega)].$$
(A.8)

We denote the co-spectrum<sup>3</sup> 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 non-stationary. To address this issue, we 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, ..., T-1. The sample covariance over this interval can be calculated as:

$$\operatorname{Cov}\langle x_t, y_t \rangle = \frac{1}{T} \sum_{t=0}^{T-1} (x_t - \overline{x})(y_t - \overline{y}), \qquad (A.9)$$

where  $\overline{x}$  and  $\overline{y}$  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:

$$\operatorname{Cov}\langle x_t, y_t \rangle = \frac{1}{T} \sum_{k=1}^{T-1} \hat{L}_{xy}[k] , \quad \hat{L}_{xy}[k] \equiv \frac{1}{T} \Re[X_k^* Y_k]$$
 (A.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,  $\operatorname{Cov}_K \langle x_t, y_t \rangle$  where  $K \subseteq \{1, \ldots, T-1\}$ , 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  $\omega_k$ , 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

<sup>&</sup>lt;sup>3</sup>The 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.

not consistent. Practical implementation details, including the standard errors of these estimators, are discussed in Section A.3. 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), Phillips, Sun, and Jin (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 (A.10), since adding or subtracting a constant to either time series does not change the sample covariance. In addition, as mentioned in Section A.1, 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. 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.

#### A.3 General Moment Properties of the Power Spectrum

In this section, 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.

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  $\gamma_{xy}[m] = \mathbb{E}[(x_{t+m} - m_x)(y_t - m_y)]$ . Assuming the cross-covariance function has finite energy, let  $P_{xy}(\omega)$  be its Discrete-Time Fourier Transform (DTFT) such that,

$$P_{xy}(\omega) = \sum_{m=-\infty}^{\infty} \gamma_{xy}[m] e^{-j\omega m}, \qquad (A.11)$$

$$\gamma_{xy}[m] = \frac{1}{2\pi} \int_0^{2\pi} P_{xy}(\omega) e^{j\omega m} \mathrm{d}\omega.$$
 (A.12)

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  $E[C_{xy}[k]]$ 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,  $E[C_{xy}[k]]$  is an asymptotically unbiased estimator of  $P_{xy}(\omega_k)$  (Oppenheim and Schafer, 2009). Moreover, over a wide range of conditions, it can be shown that,

$$\operatorname{Var}[L_{xy}[k]] \approx \frac{1}{2} (P_{xx}(\omega_k) P_{yy}(\omega_k) + \Lambda^2_{xy}(\omega_k) - \Psi^2_{xy}(\omega_k)), \qquad (A.13)$$

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 is 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 Schafer (2009) and Jenkins and Watts (1968) for a more detailed discussion of these advanced implementation techniques.

## References

- Anderson, T., 1971, The Statistical Analysis of Time Series. John Wiley & Sons, Inc., New York, NY.
- Bandi, F., B. Perron, A. Tamoni, and C. Tebaldi, 2017, "The scale of predictability," SSRN.
- 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 frequencydomain approach," *Journal of Econometrics*, 132(2), 363–378.
- 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.
- Carr, P., and D. Madan, 1999, "Option valuation using the fast Fourier transform," *Journal* of Computational Finance, 2(4), 61–73.
- 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.
- Dew-Becker, I., and S. Giglio, 2016, "Asset pricing in the frequency domain: theory and empirics," *The Review of Financial Studies*, 29(8), 2029–2068.
- Engle, R., 1974, "Band Spectrum Regression," International Economic Review, 15(1), 1–11.
- Frazzini, A., D. Kabiller, and L. Pedersen, 2013, "Buffett's alpha," National Bureau of Economic Research.
- 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 Mod*els in Business and Industry, 19(3), 245–268.

26

- 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.
- 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.
- Merton, R., 1973, "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.
- Oppenheim, A., and R. Schafer, 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.
- Ross, S., 1976, "The arbitrage theory of capital asset pricing," *Journal of economic theory*, 13(3), 341–360.
- 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.
- 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.

27

- 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.
- Wu, W., and P. Zaffaroni, 2016, "Asymptotic Theory for Nonparametric Multivariate Spectral Density Estimates," *Working Draft.*