



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

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

Shomesh E. Chaudhuri, Andrew W. Lo

To cite this article:

Shomesh E. Chaudhuri, Andrew W. Lo (2018) Dynamic Alpha: A Spectral Decomposition of Investment Performance Across Time Horizons. Management Science

Published online in Articles in Advance 18 Oct 2018

. <https://doi.org/10.1287/mnsc.2018.3102>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, INFORMS

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

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

Shomesh E. Chaudhuri,^a Andrew W. Lo^{b,c}

^a MIT Department of Electrical Engineering and Computer Science, MIT Laboratory for Financial Engineering, Cambridge, Massachusetts 02139-4307; ^b MIT Sloan School of Management, MIT Laboratory for Financial Engineering, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, Massachusetts 02142; ^c Santa Fe Institute, Santa Fe, New Mexico 87501

Contact: shomesh@mit.edu,  <http://orcid.org/0000-0002-4141-570X> (SEC); alo-admin@mit.edu,

 <http://orcid.org/0000-0003-2944-7773> (AWL)

Received: June 2, 2016

Revised: June 28, 2017; November 27, 2017;
March 12, 2018

Accepted: March 26, 2018

Published Online in *Articles in Advance*:
October 18, 2018, and updated October 24, 2018

<https://doi.org/10.1287/mnsc.2018.3102>

Copyright: © 2018 INFORMS

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 resulting from 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.

History: Accepted by Gustavo Manso, finance.

Funding: Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged.

Supplemental Material: The online supplement is available at <https://doi.org/10.1287/mnsc.2018.3102>.

Keywords: investments • performance attribution • portfolio management • alpha • spectral analysis

1. Introduction

The shortest decision interval of a modern investment strategy may range from microseconds to years—a wide span of time horizons. Although 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-add 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 measures 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 such as 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

exhibit more “power” at the lower frequencies, and faster-moving strategies exhibit more power at the higher frequencies. By identifying the particular frequencies that are responsible for a given strategy's expected returns, an investor has an additional dimension with which to manage the risk/reward profile of the investor's 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, in which 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 nonstationary signal analysis (Baxter and King 1999; Croux et al. 2001; Ramsey 2002; Huang et al. 2003; Breitung and Candelon 2006; Crowley 2007; Rua 2010, 2012; Dew-Becker and Giglio 2016; Bandi et al. 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-add 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 depends 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. Because 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 nonstationarity of financial time series, our analysis relies on the short-time Fourier transform, which applies the discrete Fourier transform (DFT) to windowed subsamples of the entire sample

(Oppenheim and Schaffer 2009). Recently, wavelets (Ramsey 2002; Crowley 2007; Rua 2010, 2012) and other transforms (Huang et al. 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 resulting from 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 $t = 0, \dots, T - 1$, the average one-period portfolio return can be calculated as

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

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 (δ_p) and a static component (v_p) as follows:

$$\bar{r}_p = \delta_p + v_p, \quad (2)$$

$$\delta_p = \sum_{i=1}^N \text{Cov}\langle w_{i,t}, r_{i,t} \rangle, \quad v_p = \sum_{i=1}^N \bar{w}_{i,t} \cdot \bar{r}_{i,t}. \quad (3)$$

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 has 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,

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

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 in the online supplement) 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 $k \in \{0, \dots, T-1\}$, 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 $0 \leq k \leq T/2$, can be used to convert the k th harmonic frequency to its corresponding time horizon, h . We also note that $\delta_{p,0} = v_p$, 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 nonzero 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 nonzero 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) indicate 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*, σ_δ , as the variability of the difference between the portfolio return, $r_{p,t}$, and the static component, $v_{p,t} = \sum_{i=1}^N \bar{w}_{i,t} \cdot r_{i,t}$. Specifically,

$$\sigma_\delta = \sqrt{\text{Var}\langle r_{p,t} - v_{p,t} \rangle}, \quad (5)$$

where σ_δ 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_δ , can then be defined as

$$I_\delta = \frac{\delta_p}{\sigma_\delta}, \quad (6)$$

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 δ_p and σ_δ 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 and 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,

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

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 $E[\varepsilon_{i,t} | F_{1,t}, \dots, F_{M,t}] = 0$. This specification is consistent with Merton's (1973) intertemporal capital asset pricing model and Ross's (1976) arbitrage pricing theory. Because our expected-return decomposition is considerably more general than any particular asset-pricing model or linear-factor structure, we allow for an intercept, α_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

$$\begin{aligned} \bar{r}_p = & \underbrace{\text{Risk-Free Rate} + \text{Risk Premia} + \text{Security Selection}}_{v_p \equiv \text{Static Component}} \\ & + \underbrace{\text{Factor Timing} + \text{Security Timing}}_{\delta_p \equiv \text{Dynamic Component}}, \end{aligned} \quad (8)$$

where

$$\text{Risk-Free Rate} \equiv \bar{r}_{f,t} \quad (9)$$

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

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

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

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

Because of the structure of the linear multifactor 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 α_i 's, a factor-timing component that depends on the covariance between the portfolio's factor

exposures and the underlying factors, and a security-timing component that depends on the covariance between weights and the idiosyncratic component of security returns. Note that the factor- and 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 α_i 's, which are reflected in the static component, v_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 α_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, during which 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%.

Table 1. The Expected Return of a Constant Portfolio Depends Only on the Static Component

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%	
Spectral decomposition of \bar{r}_p						
v_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%

Note. Mean portfolio return is highlighted in bold.

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	
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%	
Spectral decomposition of \bar{r}_p						
v_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%

Note. Mean portfolio return is highlighted in bold.

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 alternate 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 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 .

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.

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 ($\delta_{p,6}$)

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%	
Spectral decomposition of \bar{r}_p						
v_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%

Note. Mean portfolio return is highlighted in bold.

For a more realistic example, consider the long/short equity market-neutral strategy of Lo and MacKinlay (1990)

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

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

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:

$$r_{i,t} = \varepsilon_{i,t} + \lambda \varepsilon_{i,t-1}, \quad (16)$$

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

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

We see that the expected return is proportional to the mean reversion factor, λ , and the volatility factor, σ^2 . Applying our spectral decomposition (see Section A.2 in the online supplement), we find that

$$\delta_{p,\omega} = -\sigma^2\left(1 - \frac{1}{N}\right)(\lambda \cos(2\omega) + (1 + \lambda^2) \cos(\omega) + \lambda), \tag{18}$$

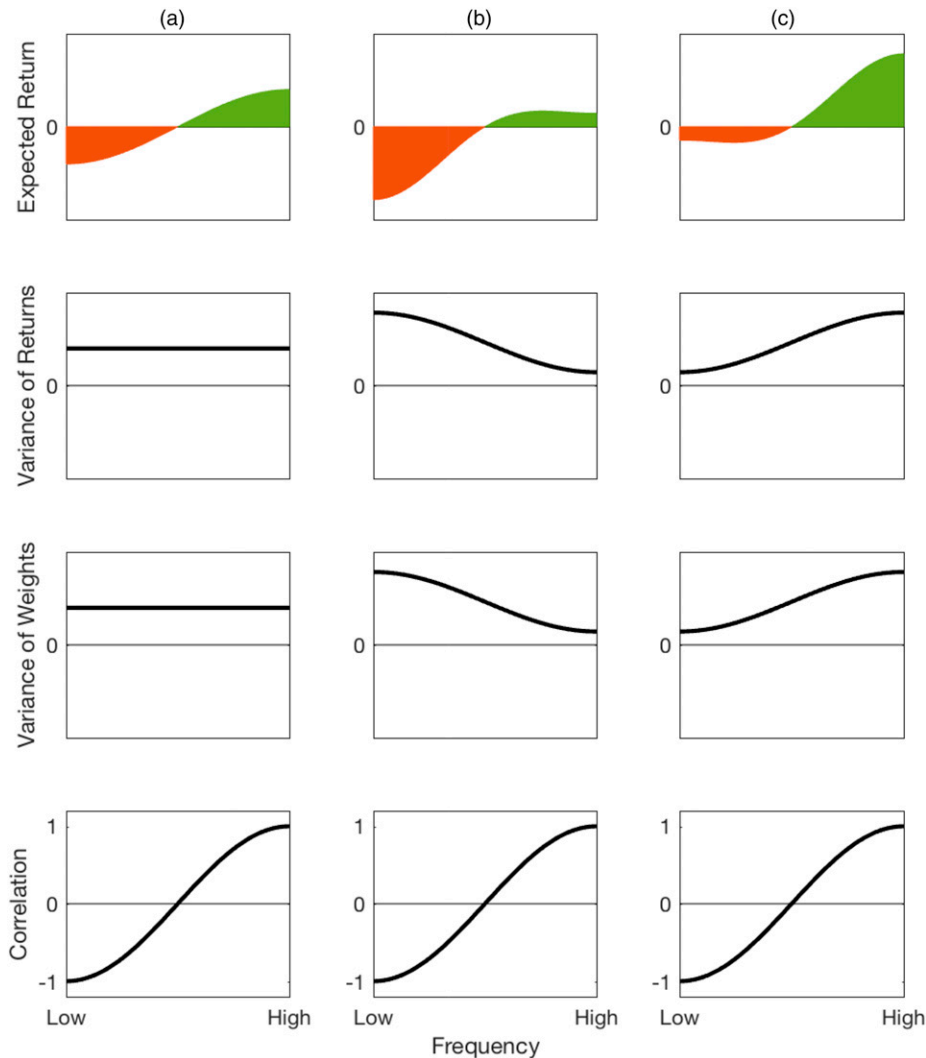
$\omega \in [0, 2\pi)$.

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 ω to its corresponding time horizon, h .

Figure 1(a) 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 because the weights and returns are out of phase at longer time horizons, resulting in zero net alpha.

Figure 1, (b) and (c), shows 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 because 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

Figure 1. (Color online) Dynamic Alpha of the Contrarian Trading Strategy Applied to the (a) Serially Uncorrelated, (b) Momentum, and (c) Mean Reversion Implementations of (16)



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, because 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 of 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

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

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 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 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 one-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 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 on the low-mean assets and short on the high-mean assets. Therefore, the static component, that is, 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

Figure 2. (Color online) 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




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

Statistic	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	One-day	Two-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
Minimum	-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%
Maximum	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

Note. The Sharpe ratio (SR) is calculated relative to a 0% risk-free rate.

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 low-frequency “noise.” Because 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.

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 nonoverlapping 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.

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

Statistic	One-day	Two-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 ($h \geq 5d$)	-44.7%	-19.1%
Medium frequency ($3d \leq h < 5d$)	6.3%	33.7%
High frequency ($h < 3d$)	70.6%	-0.4%

Note. 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).

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 the 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 et al. (2013) using total assets to equity.

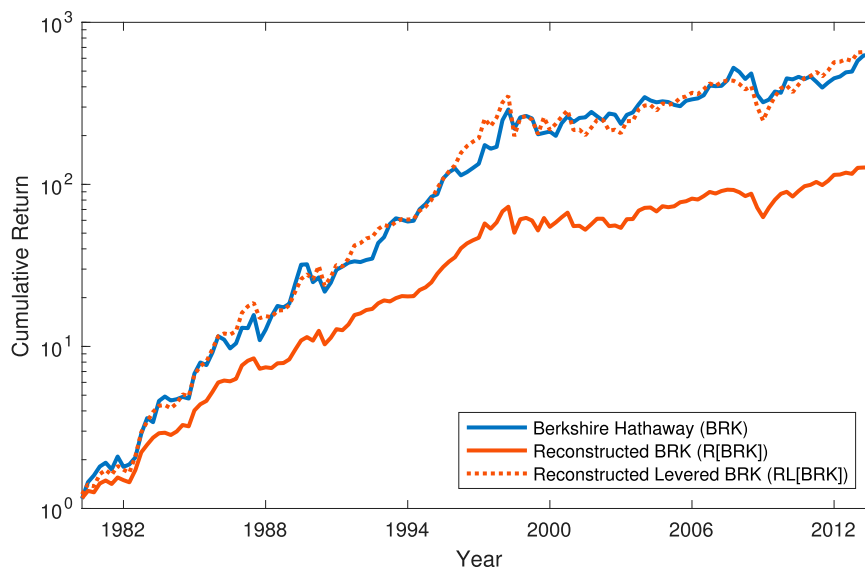
Table 6 contains summary statistics for the monthly returns of each time series. With an average annualized return of 22.9% over more than 30 years, Berkshire clearly has positive alpha when compared with traditional risk factors. Frazzini et al. (2013) find that Buffett’s returns are due more to security selection than his effect on management, which suggests that a large component of his returns must be static alpha, that is, high average weights on securities with large α_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 because of 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.

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 five 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

Figure 3. (Color online) 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



Note. 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]).

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

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
Minimum	0.0%	-23.7%	-30.1%	-30.9%	-43.6%
Median	1.2%	3.9%	4.4%	4.2%	6.0%
Maximum	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

Note. 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]).

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 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

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

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 ($h \geq 5y$)	4.3%
Med frequency ($1.5y \leq h < 5y$)	1.1%
High frequency ($h < 1.5y$)	-1.2%

Notes. 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 because of rounding.

of systemic risk can all be constructed using spectral analysis (see Chaudhuri and Lo 2016 for details). 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 nonstationarities.

Acknowledgments

The authors thank Tom Brennan, Hui Chen, Ian Dew-Becker, and Paul Mende, as well as 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 herein.

References

- Bandi FM, Perron B, Tamoni A, Tebaldi C (2017) The scale of predictability. Accessed April 8, 2018, <https://ssrn.com/abstract=2184260>.
- Baxter M, King RG (1999) Measuring business cycles: Approximate band-pass filters for economic time series. *Rev. Econom. Statist.* 81(4):575–593.
- Breitung J, Candelon B (2006) Testing for short- and long-run causality: A frequency-domain approach. *J. Econometrics* 132(2):363–378.
- Carr P, Madan D (1999) Option valuation using the fast Fourier transform. *J. Comput. Finance* 2(4):61–73.
- Chaudhuri S, Lo A (2016) Spectral portfolio theory. Accessed April 8, 2018, <https://ssrn.com/abstract=2788999>.
- Croux C, Forni M, Reichlin M (2001) A measure of comovement for economic variables: Theory and empirics. *Rev. Econom. Statist.* 83(2):232–241.
- Crowley PM (2007) A guide to wavelets for economists. *J. Econom. Surveys* 21(2):207–267.
- Dew-Becker I, Giglio S (2016) Asset pricing in the frequency domain: Theory and empirics. *Rev. Financial Stud.* 29(8):2029–2068.

- Engle RF (1974) Band spectrum regression. *Internat. Econom. Rev.* 15(1):1–11.
- Frazzini A, Kabiller D, Pedersen LH (2013) Buffett's alpha. NBER Working Paper w19681. Accessed April 8, 2018, <http://ssrn.com/abstract=2360949>.
- Granger CWJ, Engle RF (1983) Applications of spectral analysis in econometrics. Brillinger DR, Krishnaiah PR, eds. *Handbook Statist.* (Elsevier Science Publishers, New York) 3:93–109.
- Granger CWJ, Hatanaka M (1964) *Spectral Analysis of Economic Time Series* (Princeton University Press, Princeton, NJ).
- Hasbrouck J, Sofianos G (1993) The trades of market makers: An empirical analysis of NYSE specialists. *J. Finance* 48(5):1565–1593.
- Huang NE, Wu ML, Qu W, Long SR, Shen SSP, Zhang JE (2003) Applications of Hilbert-Huang transform to non-stationary financial time series analysis. *Appl. Stochastic Model. Bus. Indust.* 19(3):245–268.
- Jensen MC (1968) The performance of mutual funds in the period 1945–1964. *J. Finance* 23(2):389–416.
- Jensen MC (1969) Risk, the pricing of capital assets, and the evaluation of investment portfolios. *J. Bus.* 42(2):167–247.
- Lintner J (1965) The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Rev. Econom. Statist.* 47(1):13–37.
- Lo AW (2008) Where do alphas come from? A new measure of the value of active investment management. *J. Invest. Management* 6(2):1–29.
- Lo AW, MacKinlay AC (1990) When are contrarian profits due to stock market overreaction? *Rev. Financial Stud.* 3(2):175–205.
- Merton R (1973) An intertemporal capital asset pricing model. *Econometrica* 41(5):867–887.
- Oppenheim AV, Schaffer RW (2009) *Discrete-Time Signal Processing*, 3rd ed. (Prentice Hall, Upper Saddle River, NJ).
- Ramsey JB (2002) Wavelets in economics and finance: Past and future. *Stud. Nonlinear Dynam. Econometrics* 6(3):1–27.
- Ross SA (1976) The arbitrage theory of capital asset pricing. *J. Econom. Theory* 13(3):341–360.
- Rua A (2010) Measuring comovement in the time-frequency space. *J. Macroeconom.* 32(2):685–691.
- Rua A (2012) Wavelets in economics. Bank of Portugal Economic Bulletin. Accessed April 8, 2018, https://www.bportugal.pt/sites/default/files/anexos/papers/ab201208_e.pdf.
- Sharpe WF (1964) Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance* 19(3):425–442.
- Sharpe WF (1966) Mutual fund performance. *J. Bus.* 39(1):119–138.
- Treynor JL (1965) How to rate management of investment funds. *Harvard Bus. Rev.* 43(1):63–75.

CORRECTION

In this article, “Dynamic Alpha: A Spectral Decomposition of Investment Performance Across Time Horizons” by Shomesh E. Chaudhuri and Andrew W. Lo (first published in *Articles in Advance*, October 18, 2018, *Management Science*, DOI:10.1287/mnsc.2018.3102), Table 7 has been corrected to include six additional rows of data: “Risk-free rate,” “Risk premium,” and “Static alpha” have been added under “Static component $\times 4$ ” and “Low frequency ($h \geq 5y$),” “Medium frequency ($1.5y \leq h < 5y$),” and “High frequency ($h < 1.5y$)” have been added under “Dynamic component $\times 4$ ” along with their corresponding leveled returns.