oxford
CAPITAL STRATEGIES

Deconstructing Noise

A closer look at the notion of noise, its definition, and potential impact in today's market environment

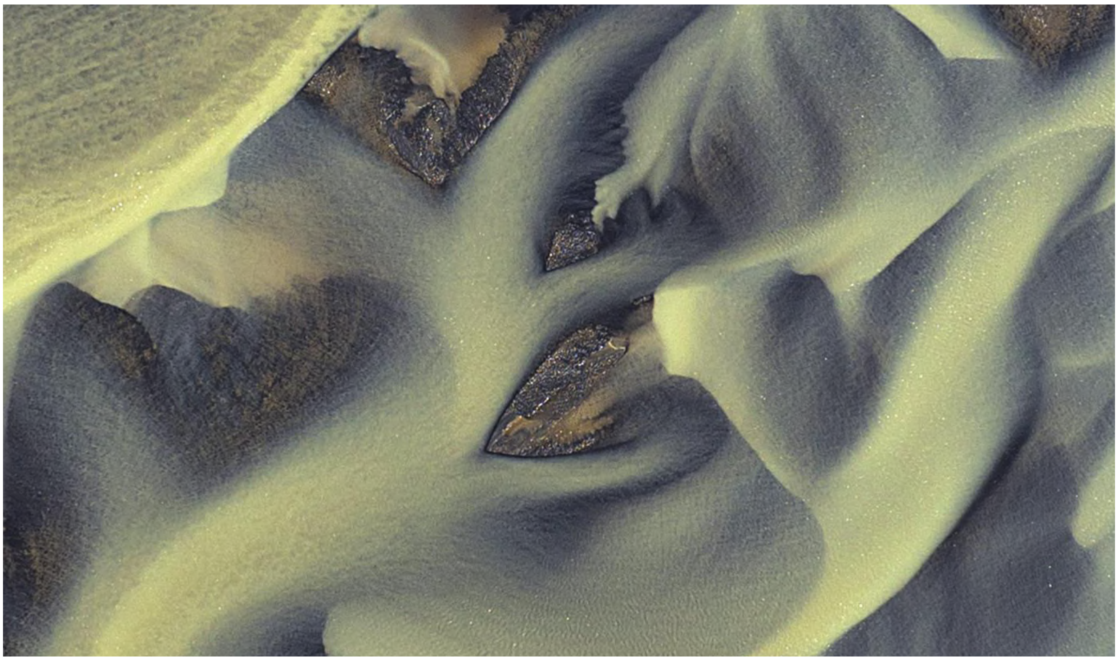

Introduction

Noise tends to mean different things to different people. It is a subject that raises a number of technical challenges, as well as deep philosophical questions. The term is used in a wide range of scientific disciplines, including physics, engineering, acoustics, biology and natural sciences. In some way, noise is so deeply embedded in nature and in the physical world that it is impossible to drive out. The notion of noise is also intrinsic to the statistical analysis of the variability of data in almost every domain of empirical enquiry.

Signal-to-noise ratios and other equivalent measures have been used extensively to relate the power or intensity of a desired “signal” (meaningful information) to the level of background “noise”. In this context, noise is viewed as an undesirable feature of data sets – a challenge to overcome in the quest for an authentic signal. In finance too, most of the empirical research and academic literature has focused on the opposition of noise to information. Yet we would argue the distinction between noise and information remains somehow elusive, making it virtually impossible to draw a clear line between “random” and “non-random” market movements.

As a matter of fact, there is still no universally accepted definition of noise in finance, nor is the notion always fully and consistently understood in financial circles. For the purpose of this paper, we will broadly refer to noise as the erratic movement that surrounds the underlying direction of prices across different timeframes. We will also refrain from characterizing noise as either positive or negative, and explore its behaviour and characteristics taking a neutral stance.

Noise vs Volatility

Unlike noise, volatility is a more established concept, and a term that finds its way into almost all conversations about the financial markets. Market volatility is generally understood as a measure of the uncertainty associated with large changes in the value of a particular asset or financial instrument. Perhaps unsurprisingly, noise is often confused with volatility, even if the two are distinct phenomena that account for different characteristics of financial time series.

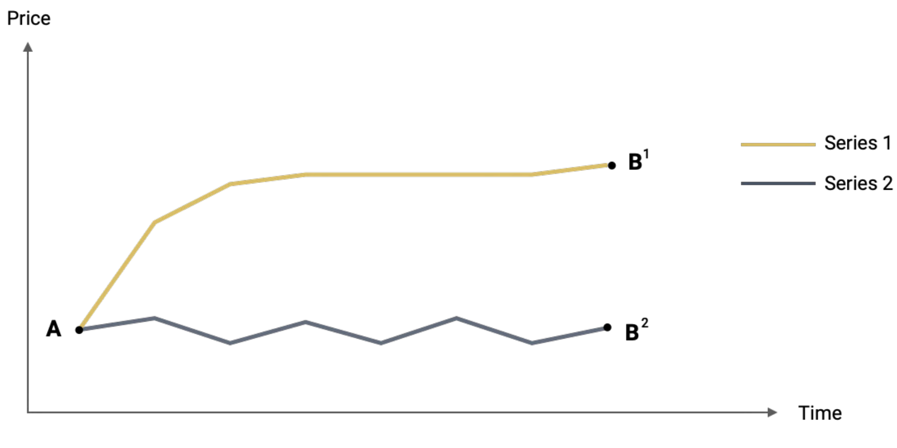

Figure 1: Two time series illustrating the difference between noise and volatility.

Figure 1 shows two time series with different levels of noise and volatility. The first time series comprises a string of smaller price changes with a large spike bringing the price significantly higher. The second series, on the other hand, is very choppy and gains little ground over the same period.

At first glance, the second series may appear to be more volatile because it is extremely noisy, yet the opposite is true. The standard deviation of returns is 1.49 for Series 1 and 1.10 for Series 2. The dichotomy of these two series exemplifies the difference between volatility and noise: Series 1 represents a high-volatility but low-noise price path, while Series 2 illustrates low volatility with high noise. Simply put, noise can be thought of as a measure of the "choppiness" or "roughness" of a market's price path, while volatility typically measures the average "magnitude" of the market's price changes.

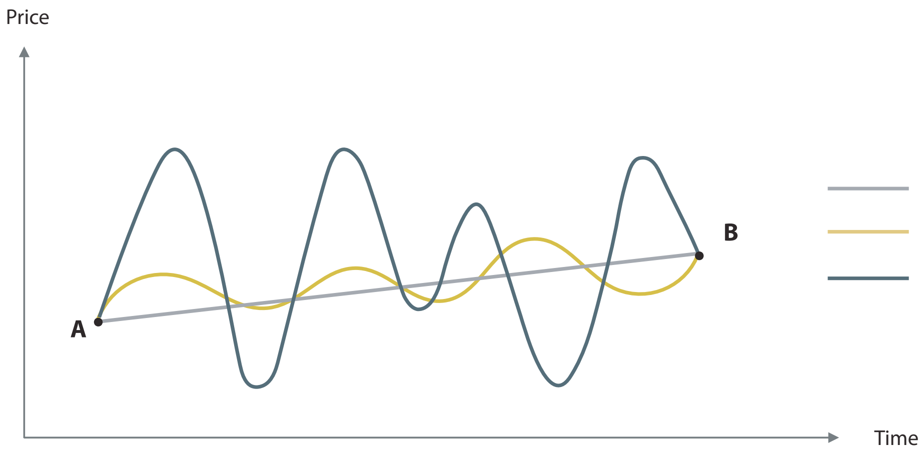

Figure 2: Different noise levels underlying the same net price change from point A to point B.

Figure 2 illustrates three different noise levels surrounding the same net price change from point A to point B. The straight line represents a hypothetical market movement with no noise, the line with smaller price changes above and below the straight line illustrates a moderate level of noise, and the line with larger swings represents a highly noisy time series.

As is clear from the foregoing, noise is also a measure of the uncertainty in market direction. It signals the lack of a clearly definable trend. As we will see, its impact turns out to be more subtle and significant than is generally understood. Unfortunately, noise often resembles volatility, because in a high-noise environment the market often changes direction quite dramatically, in an erratic fashion. With that being said, the relationship between noise and volatility remains a complex one, with periods of low volatility and high noise (or high volatility and low noise) being less frequent, but not uncommon.

Measuring Noise

Market noise can be assessed in a number of different ways, the most common assessment measures being price density, the Efficiency Ratio (ER), wavelet transforms, and Fractal Dimension. For the purpose of this paper, we will focus our discussion on Fractal Dimension, which we believe lends itself to a more scientific and granular analysis of market noise. Fractal Dimension requires a relatively complex calculation, and can only be estimated over a certain period of time. On an intuitive level, it can be thought of as a measure of "roughness". It is a ratio that compares how detail in a pattern (strictly speaking, a fractal pattern) changes with the scale at which it is measured[1]. Fractal Dimension has also been characterized as a measure of the space-filling capacity of a pattern that tells us how a fractal scales differently from the space in which it is embedded.

The most widely accepted measure of Fractal Dimension is the Minkowski–Bouligand dimension, also known as the Minkowski dimension or box-counting dimension (we will refer to it simply as box-counting dimension). The reason for its dominance lies in its relative ease of use, and the fact that it lends itself to the measurement of shapes and patterns in higher dimensional spaces (using three- or n-dimensional boxes). To calculate this dimension for a fractal S, we can imagine this fractal as lying on an evenly spaced grid, and count how many boxes are required to cover the entire set. The box-counting dimension is calculated by observing how this number changes as we make the grid finer by applying a box-counting algorithm (a classic example is illustrated in Fig. 3, below).

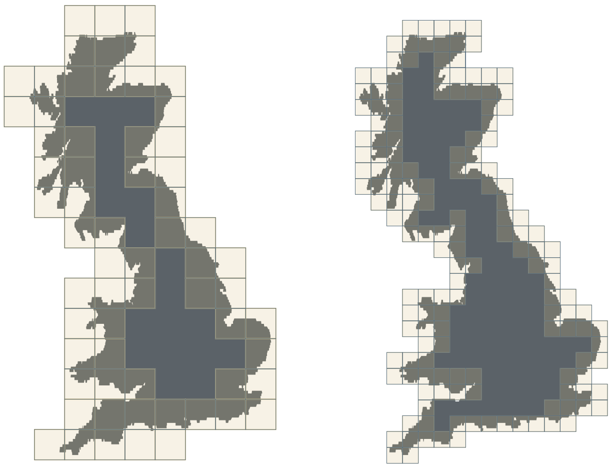

Figure 3: Estimating the box-counting dimension of the coast of Great Britain.

From a mathematical perspective, if we assume that N(ε) is the number of boxes required to cover the set and ε is the reduction factor, then the box-counting dimension may be simply defined as:

$$\dim_{\text{box}}(S) = \lim_{\varepsilon \to 0} \frac{\log N(\varepsilon)}{\log\left(1/\varepsilon\right)}$$

Now, how does Fractal Dimension translate into financial market noise? The charts in Figure 4 provide two extreme examples of market environments with a very low and very high level of noise respectively. The first graph shows the evolution of the price of Copper futures from June 2005 to May 2006, with a Fractal Dimension of 1.29.

[1] The essential idea of "fractured" dimensions has a long history in mathematics, but the term itself was brought to the fore by Benoit Mandelbrot with his famous paper on self-similarity: "How Long Is the Coast of Britain? Statistical Self-Similarity and Fractional Dimension", Science, Vol. 156, No. 3775, pp. 636-638 (May 5, 1967).

The second graph illustrates the price path of Copper futures from July 2013 to June 2014, where the Fractal Dimension is now 1.49[2]. As Figure 4 clearly illustrate, the Fractal Dimension (and hence the level of noise) that is temporarily observed in a given market is far from being static, and constantly evolves over time. Nevertheless, it provides vital information about the underlying market environment, and can be exploited in a number of different ways.

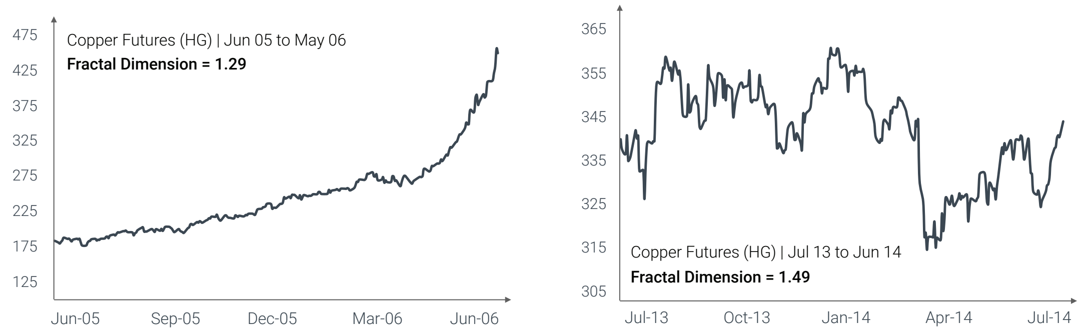

Figure 4: The Fractal Dimension of Copper futures from Jun 2005 to May 2006 (1,29) and Jul 2013 to Jun 2014 (1.49).

Drivers of Noise

While identifying the underlying drivers of market noise is beyond the scope of this paper, on an intuitive level, noise has been explained as the result of numerous market participants buying and selling for different reasons across different timeframes. Day-to-day market movements, from that perspective, are affected by various actors buying and selling to meet different operational needs - such as rebalancing investment portfolios, processing fund redemptions and subscriptions, hedging specific risks, or adhering to different mandates. Market news are also seen as a rich source of noise, and can drive price away from "fair value" over the short term.

Is Noise Increasing?

The next question we may want to ask ourselves is whether noise has increased, or whether there has been any significant shift in noise levels over the last decades. To answer this question, let us calculate the rolling average of Fractal Dimension using a diversified portfolio of 48 global futures markets, including currencies, equity indices, interest rates and commodities. The results are presented in Figures 5 and 6. Market noise, calculated as the average Fractal Dimension of all instruments using a rolling window of 252 days, has gradually increased from 1.29 to 1.45 over a 38-year period (from 1980 to 2018). The findings are fairly stable across markets, asset classes and geographies, and seem to indicate quite unequivocally that market noise has increased over the last 40 years (Fig. 6).

2. The charts were built using the back-adjusted continuous contract for CME Copper futures (HG), rolled based on Volume and Open Interest readings.

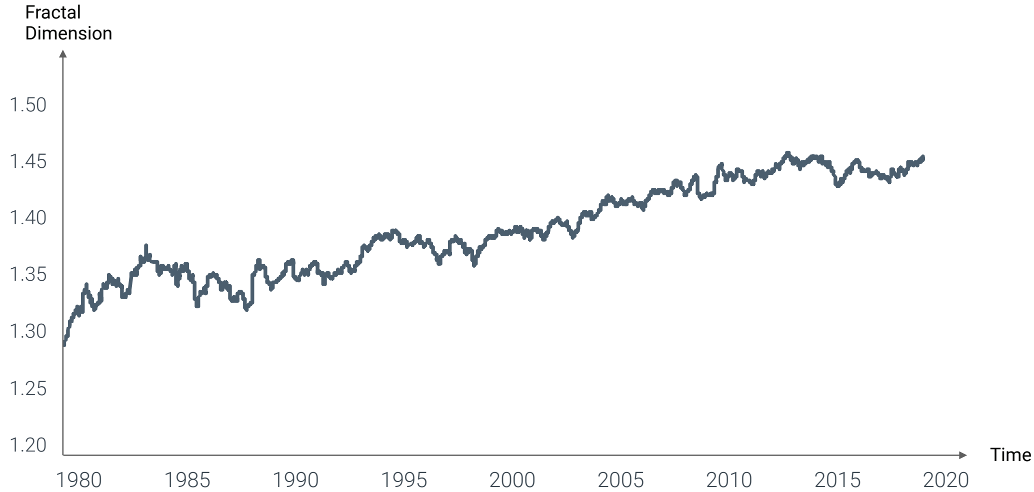

Figure 5: Average Fractal Dimension of a diversified portfolio of 48 markets from 1980 to 2018.

This tendency, which some analysts have occasionally referred to as a possible explanation of the deteriorating returns of trend-following CTAs and long-term momentum strategies, seems to have further intensified over the last decade. With that being said, it is very difficult to say whether this situation will persist or reverse in the future.

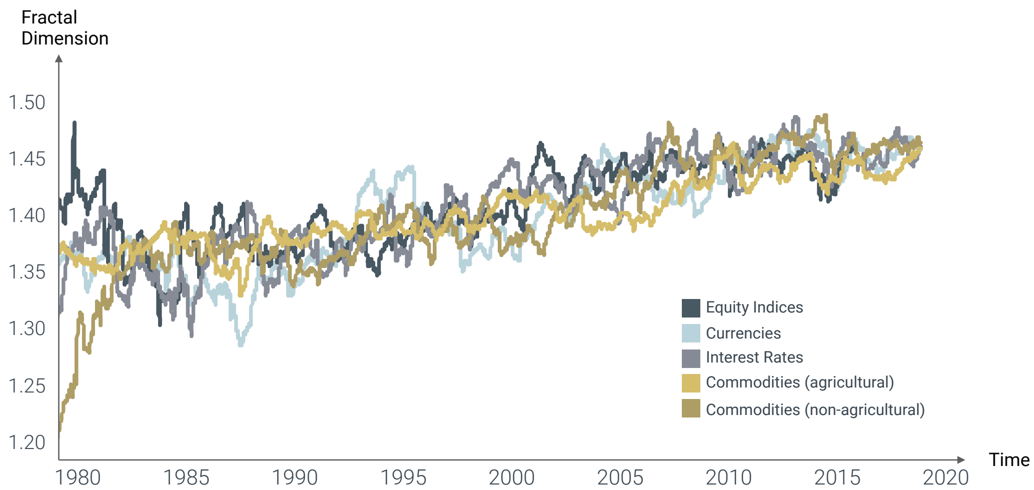

Figure 6: Average Fractal Dimension of different asset classes from 1980 to 2018.

Some empirical studies have suggested that a direct relationship may exist between the activity level of market participants and the level of noise that is observed in a given market. In this context, high noise has been characterized as a byproduct of increased trading activity, which has risen in recent years in response to increases in liquidity, lower transaction costs, and more efficient information flows. Our opinion is that these findings are still preliminary, and remain largely inconclusive. However, if we accept the above explanation for increased market noise, then we should expect noise levels to remain reasonably high in the foreseeable future, driven by the increasing flow of financial news, and growing levels of trading volumes and liquidity across the world.

The Impact of Noise

What is perhaps more important, is trying to understand the potential impact of increased noise on market participants. What are the practical implications? And can anything be done to adapt to increasing noise levels? For illustrative purposes, we will look at the potential impact of noise on a simple trend-following strategy. By extension, similar considerations also apply to other directional strategies that rely on persistent and extended price moves to generate profits.

By definition, a trend-following model will reap the greatest and cleanest profits when a market breaks out in one direction and never looks back. As price noise increases, entry signals begin to lag, and it takes longer to sort out a "real" trend from the underlying noise. In other words, a bigger price move is required for momentum confirmation. The mirror logic obviously applies to the identification of the "end" of a market trend, and the timing of exits. In an extreme scenario, market noise could become so significant that by the time entry/exit signals are identified, most trades generate net losses for an investor, and the profitability of the strategy is compromised (Fig. 7).

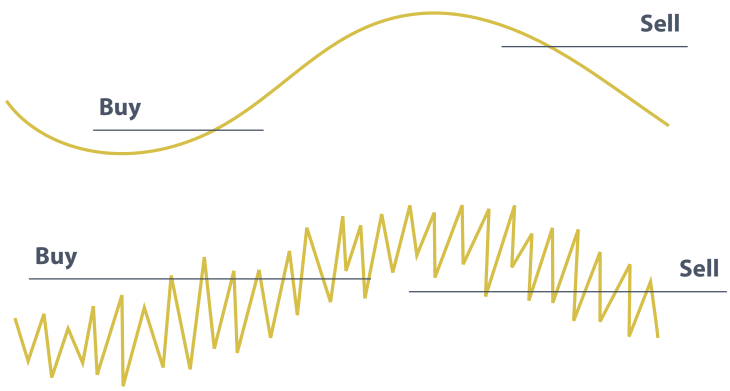

Figure 7: Hypothetical impact of increased market noise on a basic trend-following strategy.

It would be tempting to take this observation further, and argue that a higher-noise environment must be favorable for mean-reverting strategies. Caution is warranted here, as evidence suggests a more nuanced reality. Simple agent-based models (ABMs) have been used to test similar hypotheses, but they are often loaded with assumptions and require high levels of parameterization, making it difficult to reach definite conclusions.

Overall, we believe that the observation of signal lagging as a result of increased market noise is plausible and holds its ground, and that it could have far-reaching implications on the profitability of simple directional strategies. It has become a tangible phenomenon in a growing number of markets, and is increasingly manifesting itself through large market swings, sharp reversals, and recoveries across asset classes and timeframes.

Conclusion

In markets, as in life, nobody can stay in the same place without drifting backwards. Markets change, and models need constant adaptation to survive. Based on empirical evidence and research findings, market noise has been gradually increasing over the last decades (Fig. 5). While the underlying causes of this tendency remain an issue of debate, noise increase may be due to a combination of higher market activity and more efficient information flows. This would suggest that noise could possibly persist as an increasing number of markets evolve and become more mature. Although the impact of heightened noise on market participants and trading strategies is not straightforward, it can be intuitively understood by looking at the mechanics and behaviour of a simple directional strategy such as the one briefly described in this paper.

The most common responses to market noise have been to try to filter, suppress or eliminate it from financial time series, or to identify separate market environments (e.g. trending vs mean-reverting) that would be more conducive to specific sets of strategies. We would argue that a better way to address some of the challenges presented in this paper, is to embrace noise rather than discard it. Managers would be better served by treating noise as an intrinsic feature of the financial markets, and exploring new ways to leverage it and incorporate it into their trading approaches.

About the Authors

Stefan Martinek

Mr. Stefan Martinek is the Founder and Managing Director of Oxford Capital Strategies. He brings over 15 years of experience working as an independent research consultant for institutional clients, including some of the world's largest CTAs and quantitative futures managers. His research covers a variety of well-documented factors as well as alternative concepts blending fundamental, technical and behavioural finance elements. Over the course of his career, Mr. Martinek has also worked as a trader and a private equity manager in Europe and the US. He studied Nuclear Engineering at the Moscow Institute of Energetics, and holds an MBA in Finance from Oxford University.

Michele Cocchiglia

Mr. Michele Cocchiglia is the Head of Strategic Partnerships and Business Development at Oxford Capital Strategies. Prior to that, he was the Founder and Managing Partner of Kelonia Capital Management, a diversified quantitative futures manager registered in Switzerland and Luxembourg. He studied economics, statistics and financial mathematics at Bocconi University, and holds a master's degree in information systems from the London School of Economics (LSE). Earlier in his career, Mr. Cocchiglia worked as a risk modelling expert with the United Nations, where he served as the Regional Coordinator of a disaster risk assessment initiative in Asia and the Pacific.

Bibliography

Black, F. (1986), "Noise", The Journal of Finance, Vol. XLI No. 3.

Bloomfield, R., O'Hara, M., Saar, G. (2009), "How Noise Trading Affects Markets: An Experimental Analysis", Review of Financial Studies, 22(6), June 2009: 2275-2302.

Cont, R. (2005), "Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models", in Teyssière, G. & Kirman, Alan P. (Eds.) Long Memory in Economics.

Falconer, K. (1990), "Fractal Geometry, Mathematical Foundations and Applications" (Wiley).

Hu, G. X., Wang, J., Pan, J. (2013), "Noise as Information for Illiquidity", Journal of Finance, 68: 2223-2772.

Kaufman, P. (2013), Trading Systems and Methods, 5th edn (Wiley Trading).

Khandani, A., Lo, A. (2011), "What Happened to the Quants in August 2007? Evidence from Factors and Transactions", Journal of Financial Markets, 14: 1-46.

Lo, A. (2004), "The Adaptive Markets Hypothesis: Market Efficiency from and Evolutionary Perspective", The Journal of Portfolio Management, 30(5): 15-29.

Malaspina, C. (2018), An Epistemology of Noise (Bloomsbury Academic).

Mandelbrot, B. (1997), Fractals and Scaling in Finance: Discontinuity, Concentration, Risk (Springer).

Mandelbrot, B. (1999), Multifractals and $1/f$ Noise - Wild Self-Affinity in Physics (1963–1976) (Springer).

Mandelbrot, B., Hudson, R. (2007), The Misbehavior of Markets: A Fractal View of Financial Turbulence (Basic Books).

Oppenheim, A. V., Schafer, R. W. (1999), Discrete-Time Signal Processing (Prentice Hall).

Peitgen, H. O., Jürgens, H., Saupe, D. (2004), Chaos and Fractals: New Frontiers of Science (Springer).

Percival, D. B., Walden, A. T. (2006), Wavelet Methods for Time Series Analysis (Cambridge University Press).

Poggio, T., Lo, A., LeBaron, B., Chan, N. (2001), "Agent-Based Models of Financial Markets: A Comparison with Experimental Markets", MIT Sloan Working Paper No. 4195-01.

Shiller, R. J. (1981), "The Use of Volatility Measures in Assessing Market Efficiency", The Journal of Finance, 36(2): 291-304.

Silver, N. (2012), The Signal and the Noise: Why So Many Predictions Fail - but Some Don't (Penguin Press).

Taleb, N. N. (2012), Antifragile: Things That Gain From Disorder (Random House).

oxford
CAPITAL STRATEGIES

Oxford Capital Strategies Ltd.
30 Bankside Court Stationfields
Kidlington, Oxford - OX5 1JE, United Kingdom