The Journal of Portfolio Management

INVITED EDITORIAL COMMENT

Taking Stationarity Seriously

BRADFORD CORNELL

I wear two hats. For 40 years, I have been a professor of finance. For the last 10 years, I have run a small hedge fund. In my role as an academic, I play down the importance of stationarity to get on with research efforts. When I have to make investment decisions, it is the elephant in the room. In fact, the question of stationarity is so important that it often dominates my investment decision making and as a result renders much academic research of little practical value. The point of this commentary is to argue that finance research needs to take the question of stationarity more seriously to be more useful to investors.

Formally, a *stationary stochastic process* is a stochastic process whose joint probability distribution does not change when shifted in time. Consequently, parameters such as mean and variance, if they are relevant, also do not change over time. Nonstationarity should not be confused with unpredictability. All random processes are unpredictable. If the process is nonstationary, even the parameters of the random distribution cannot be estimated with confidence. Putting aside formal definitions, I find the example of drawing colored balls from jugs with replacement to be a great way to explain how the problem of stationarity affects investment decision making.

If there is one jug and the balls are drawn from it with replacement, the process describing the sequence of balls drawn is stationary even though the actual color of the ball to be drawn is random. If suddenly a new jug is introduced with a different mix of balls and the next series of draws is from a mixture of the two jugs, the process is nonstationary. However, this is what can be called a limited degree of nonstationarity. By simply redefining the procedure for drawing balls, a new stationary process emerges that involves two steps. In the first step, one of the two jugs is randomly selected. In the second step, a ball is drawn from the chosen jug. As long as this procedure is followed, the new process, although more complicated, is stationary. In fact, the new process can be interpreted as an example of a regime-switching model in which first the regime is chosen and then a random ball draw occurs.

The balls and jugs analogy is useful for conceptualizing differing degrees of nonstationarity. The important questions include the following: How many jugs are there? Can the number of jugs even be enumerated? What is the distribution of balls within each of the jugs? In the limit, think of the case in which there is an immense number of jugs, the contents of which are unknown, and in which the probability of selecting a given jug is also unknown and may be changing over time. I refer to this limiting case as fundamental nonstationarity. Although this may seem like an extreme case, I argue that it is a problem that investors face on a daily basis. Fundamental nonstationarity is not a rarity but rather the normal state of affairs. To explore the issue further, I consider the examples of four investment decisions.

THE SURPRISING BEHAVIOR OF THE VIX INDEX

The Volatility Index (VIX), calculated by the Chicago Board Options Exchange, measures the market's expectation of 30-day volatility. It is constructed from the implied volatilities of a wide range of S&P 500 Index options with approximately 30 days to maturity. As of October 2017, the VIX had been at near record lows for more than a year. The average was about 11% compared to a long-run historical average of 15% or more, depending on the sample period. The investment question is whether this abnormal behavior suggests taking a position in VIX derivatives.

One way to approach the question is to turn to the academic literature on fitting stochastic models to the VIX Index. It turns out that the literature is both large and highly sophisticated mathematically. A few recent examples among the many papers include Goard and Mazur's [2013] "Stochastic Volatility Models and the Pricing of VIX Options"; "Double-Jump Stochastic Volatility Model for VIX: Evidence from VVIX" by Zang et al. [2016]; and Kaeck and Alexander's [2013] "Continuous Time VIX Dynamics: On the Role of Stochastic Volatility of Volatility." In their defense, these papers, and others like them, do allow for some nonstationarity along the lines of the two-jug analogy. They do so by incorporating the possibility of random jumps or stochastic volatility. The problem I have as an investor is that I fear the process during the current quiescent period is not just a result of a random failure of jumps to materialize or a random drop in volatility in a stochastic volatility model but instead is a fundamentally different process.

Of course, if a model is fit with enough flexibility in its parameters, it will appear to account for the nonstationarity during the sample period, but in doing so, it will misstate the true nature of the process. This is critical from an investment standpoint because if the true process is fundamentally nonstationary, at some point, it will change in a manner unanticipated by investors. If the change involves drawing from an entirely new jug among a vast number of jugs, a complex process fit to historical data will simply be misleading. This is, in effect, the argument Taleb [2007] made with regard to the financial crisis. However, the observation is not limited to the dramatic, black swan events that Taleb described. If the world is fundamentally nonstationary, it is a problem that investors face continually and to varying degrees as the social, political, and economic environments evolve.

In particular, the stochastic process for the VIX will change when the social, political, and economic factors, which are yet to be delineated, that led to its historical low mean value are transformed. One such factor that could have altered market volatility was the election of Donald Trump. However, the fact that such a hypothesis is speculative is precisely the problem. As Ross [2005] observed, it is difficult even after the fact to identify events that may have altered the stochastic process of asset returns.

THE CROSS-SECTION OF EXPECTED RETURNS

Following the lead of Fama and French [2002], intense interest in factor models designed to explain the cross-section of expected returns has led to extensive research in the area. As Harvey, Liu, and Zhu [2016] documented, that research effort has produced a veritable zoo of allegedly significant factors. Based on their review of 313 articles, the authors reported the identification of 316 priced factors. This factor zoo led Harvey, Liu, and Zhu to argue for the use of higher cutoffs for statistical significance to overcome the impact of apparent data mining.

Data mining and nonstationarity are different issues, but they can have a similar impact from a practical investment standpoint. *Data mining* refers to the problems that arise when there is repeated sampling from the same historical dataset. The most common problem that results from data mining is the discovery of idiosyncratic quirks that are unique to the sample but are not actual, true relations.[1] As a result of data mining, spurious relations uncovered in the sample period will fail to hold postsample. When the data are nonstationary, a relation may be found that does, in fact, hold for the historical sample period but which is no longer true. Once again, the relation fails to hold in the postsample period, although for a different reason.

The failure of factor models estimated in one period to hold in another may be due to data mining, nonstationarity or some combination. Either way, given the vast zoo of factors that have been uncovered, we (the research professionals) are almost assured of finding a factor model that explains the cross-section of expected returns in any chosen historical sample period. However, it remains unclear what practical value this has for investors who cannot be confident that the relations will hold going forward.

INDIVIDUAL STOCKS

With regard to individual stocks, language is an impediment to appreciating the full extent of potential nonstationarity. Throughout its corporate life, Apple has always been called Apple, but the company has reinvented itself numerous times.[2] In the process, it transformed itself from a start-up maker of personal computers into a global consumer products and services powerhouse, despite having several brushes with insolvency. Of course, it is possible that the process for stock returns remained stationary while the company was continually transformed because stock returns depend on investor expectations. However, it would be foolhardy for an investor to assume that the dramatic evolution of the firm did not have a major impact on investor perceptions, including investor estimates of risk and, thereby, on stock returns.

It is worth noting that applied investment research—by that, I mean the work of security analysts—appears to take the problem of stationarity for granted. If the stochastic process generating key metrics of financial performance, such as revenue, earnings, and free cash flow, were stationary, then presumably the best way to project future financial performance would be to fit statistical models much like those used to analyze the VIX Index. This is not what analysts do. Instead, they examine the details of the company's business with the hope that the understanding they achieve will help them predict future financial performance. This can be interpreted as an effort to overcome nonstationarity by attempting to predict how future business conditions will generate revenues, earnings, and free cash flow given currently available information. In the context of the balls and jugs analogy, security analysts are using fundamental analysis to select the jug.

SMART BETA AND FACTOR PREMIUMS

As a final example, there has been an active debate recently regarding so-called smart beta and associated factor premiums. As Asness [2016] noted, smart beta and factor-based strategies have become increasingly popular in recent years. The goal of these strategies is to identify factors, of which Fama and French's small minus large factor is an early example, and then to harvest the factor premium by investing in long–short portfolios.

As Arnott et al. [2016] repeatedly stated—although they do not couch their argument explicitly in terms of stationarity—this investment strategy is based on the assumption that the stochastic process governing factor returns is sufficiently stationary that past average premiums are reasonable estimates of future expected premiums. Arnott et al. argued that the assumption is false. They claimed that research identifying historical factor premiums has failed to adequately account for the extent to which rising valuations contributed to the lofty historical returns. Based on their empirical research, Arnott et al. concluded that valuation increases have been the primary driver of smart beta returns over the short term, and even over the long term; as a result, past excess returns are not likely to be sustainable in the future. In fact, they suggested that factor portfolios that have markedly appreciated could "go horribly wrong" and potentially crash. The point here is not to evaluate whether Arnott et al. are correct—many authors, including Asness [2016], argue their conclusions are exaggerated—but to note that the entire debate is basically a dispute over stationarity.

In the context of the jugs and balls analogy, valuation increases can be thought of as drawing from a jug without replacement. Every time, say, a red ball is drawn, the probability of drawing another red ball declines. For this reason, the distribution is nonstationary. The probability of drawing a red ball can be interpreted as the probability that a factor portfolio will earn excess returns. The more the valuation increases, the more red balls are drawn, and the less likely it will be that valuations will rise in the future.

Perhaps the most controversial factor premium in this regard is momentum. Early papers, such as that by Jegadeesh and Titman [1993], found significant premiums associated with momentum. Later papers, including that by Dolvin and Foltice [2017], argued that the anomaly had disappeared. Simultaneously, Daniel and Moskowitz [2016] reported significant crash risk associated with momentum, but Barroso and Santa-Clara [2015] claimed that this risk could be ameliorated by varying the leverage of the momentum portfolio. This is just a sliver of an immense and internally contradictory literature on momentum. From the standpoint of a practical investor, the safe conclusion is that if there is a momentum effect, it is far from stationary.

The four examples offered here are by no means unique. Similar arguments apply to most every investment strategy based on estimates of statistical parameters derived from historical data. All such strategies assume, explicitly or implicitly, that the world is sufficiently stationary that such estimates are of practical value to investors.

CONCLUSIONS AND IMPLICATIONS

The basic conclusion is straightforward: Nonstationarity is not a minor statistical annoyance but a fundamental and unavoidable issue that investors face each time they make an investment decision. I argue that there is generally insufficient evidence to support the assumption that the processes underlying social institutions (including financial markets), unlike those underlying many physical systems, are stationary. Such nonstationarity includes not only the possibility of large, unexpected breaks from the past, as occurred during the financial crisis, but daily changes in the stochastic processes governing asset returns. It is not surprising, therefore, that fundamental security analysis, which takes nonstationarity for granted, remains the basis for most practitioner-based investment research.

ENDNOTES

I would like to thank Rob Arnott, Cliff Asness, John Haut, and Richard Roll for their helpful comments.

[1]My favorite example of data mining involves Richard Feynman and the expansion of pi. Feynman would reel off the first 768 digits of the expansion, the last six of which are 9-9-9-9-9-9, and then say "and so on" before breaking into laughter. The 763rd digit of pi has now become known as the Feynman point, but the six 9s have no meaning.

[2]To be fair, the original name of the company was Apple Computer, which was shortened to Apple as other devices (which are actually computers) became the predominant source of the company's revenue. However, throughout its life, the company has generally been referred to as Apple.

REFERENCES

Arnott, R., N. Beck, V. Kalesnik, and J. West. "How Can 'Smart Beta' Go Horribly Wrong." Research Affiliates, February 2016. https://www.researchaffiliates.com/en_us/publications/articles/442_how_can_smart_beta_go_horribly_wrong.html.

Asness, C.S. "My Factor Philippic." AQR Capital Management, April 2016, pp. 1-30.

Barroso, P., and P. Santa-Clara. "Momentum Has Its Moments." *Journal of Financial Economics*, Vol. 116 (2015), pp. 111-120.

Daniel, K., and T.J. Moskowitz. "Momentum Crashes." *Journal of Financial Economics*, Vol. 122 (2016), pp. 221-247.

Dolvin, S., and B. Foltice. "Where Has the Trend Gone? An Update on Momentum Returns in the U.S. Stock Market." *The Journal of Wealth Management*, Vol. 20 (2017), pp. 29-40.

Fama, E. F., and K.R. French. "The Cross-Section of Expected Stock Returns." *The Journal of Finance*, Vol. 47 (2002), pp. 427-465.

Goard, J., and M. Mazur. "Stochastic Volatility Models and the Pricing of VIX Options." *Mathematical Finance*, Vol. 23 (2013), pp. 439-458.

Harvey, C., Y. Lui, and H. Zhu. "…and the Cross-Section of Expected Returns." *Review of Financial Studies*, Vol. 29, No. 1 (2016), pp. 5-68.

Jegadeesh, N., and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance*, Vol. 48 (1993), pp. 65-91.

Kaeck, A., and C. Alexander. "Continuous Time VIX Dynamics: On the Role of Stochastic Volatility of Volatility." *International Review of Financial Analysis*, Vol. 28 (2013), pp. 46-56.

Ross, S.A. *Neoclassical Finance*. Princeton, NJ: Princeton University Press, 2005.

Taleb, N.N. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House, 2007.

Zang, X., J. Ni, J.Z. Huang, and L. Wu. "Double-Jump Stochastic Volatility Model for VIX: Evidence from VVIX." Working paper, Peking University, 2016.

Bradford Cornell *is a professor of finance at the California Institute of Technology and managing director at SMBP LLC in Pasadena, CA.*
bcornell@caltech.edu