The Persistence of Smart Beta

Hamish Preston
Researcher
Department of Economics
Birmingham University

Tim Edwards
Head
European Index Investment Strategy
S&P Dow Jones Indices

Craig Lazzara
Global Head
Index Investment Strategy
S&P Dow Jones Indices

"Knowledge of the fact differs from knowledge of the reason for the fact." – Aristotle

The notion that patterns in securities prices can be predicted and exploited has given rise to at least two industries: quantitative fund management and, more recently, the index-based alternative operating under the ambitious moniker "smart beta." The performance of such systematic strategies poses a challenge to the "efficient" markets of classical theory, and has therefore produced a third cottage industry for academics—alternatively quantifying, explaining, or refuting the strategies' supposed outperformance. As funds or indices gain in popularity and usage, or as academic papers exploring their themes are celebrated, there is frequently a resultant change in performance. This creates a particular challenge for investors interested in extrapolating the past into the future.

At a general level, there are two (not mutually exclusive) reasons that explain why a particular strategy might outperform, above and beyond sheer luck.[1] The first reason is that the outperformance might simply be compensation for increased risk. For example, Fama and French[2] famously documented that cheap stocks outperform more expensive stocks over time. Perhaps this effect arises because cheap stocks are more volatile than expensive ones—in which case one might argue that the effect is simply a reward for bearing the incremental risk of cheapness. On the other hand, a strategy's incremental performance might not be a compensation for risk, but might represent a true anomaly.[3] In our example, this would imply that the incremental outperformance of cheap stocks *more than* compensates for their putative higher risk.

The question of whether the incremental returns attributed to a given factor (e.g., the outperformance of stocks with high momentum or low volatility) will persist is impossible to answer definitively. Yet investment vehicles tracking non-standard indices have become increasingly popular.[4] The vast majority posit both the *existence* and *persistence* of an anomaly in the market (the undervaluation of value stocks, for example) and systematically exploit them. When evaluating such investments, investors ranging from the individual to the largest institution must ask themselves not only if a particular vehicle is well-designed to exploit the anomaly but, first, if the anomaly is expected to persist?

We argue that the third industry—academic research—can have a material impact on factor persistence.[5] We illustrate this by identifying four distinct types of anomalies, only two of which show any degree of persistence.

- As the name suggests, **disappearing anomalies** don't last. The disappearance category includes strategies whose returns are arbitraged away after discovery, indicating that the returns themselves are neither a compensation for risk nor difficult to replicate. In such cases, once the average investor becomes aware of the anomaly, its benefits are completely eroded.
- Worse yet are **statistical anomalies**. Here we illustrate the pitfalls of investing based on spurious relationships that appear to exist due to chance. In these circumstances, expecting a predictable pattern of returns to emerge is naïve; we caution against the high false-positive rate to be expected with modern computing power.[6]
- Moving to the positive side of the ledger, we consider **attenuated anomalies**, the risk-adjusted returns of which diminish as they become more widely known. Attenuation shows the importance of assessing returns on a risk-adjusted basis; seemingly persistent returns may simply be a compensation for bearing additional downside risk.
- Finally, there are **persistent anomalies**. This final type shows that persistent returns can exist, even after adjustment for risk—and reminds us of the importance of conducting risk analysis to distinguish the character of anomalies.

This is not a purely academic exercise, as these four categories provide investors with a toolkit to use when assessing the anomalous returns on various strategies. In particular, we hope to provide a deeper insight into what may happen to anomalous returns—and "smart beta" indices—in the future.

Disappearance

"Tell me why? I don't like Mondays." – Bob Geldof, The Boomtown Rats

In 1973, Frank Cross' paper was the first published research to document the difference in returns between Fridays and Mondays. His research showed that the distribution of positive (negative) returns on Mondays preceded by positive (negative) returns on Fridays differed significantly from the corresponding daily differences in returns for the rest of the week. Cross also provided evidence that the difference in the probability of positive returns on Fridays (62%) and Mondays (39.5%) was statistically significant.[7] Taken together, these results highlighted an example of non-random movement in stock prices, therefore raising questions about the validity of the Efficient Market Hypothesis (EMH). Given the prominence of EMH at this time, the weekend effect became one of the hallmark anomalies of the period.

Whilst 1973 is viewed as the birth of literature on what is now called the "Weekend Effect," it was Kenneth French who coined the term in his 1980 paper supporting Cross' findings. In many cases, the unexpected returns were explained with recourse to a behavioral observation: companies tended to release bad news after the market's close on Fridays, and market participants did not fully account for this phenomenon in their day-to-day trading. However, following a period when many further, supportive papers were published, there began a growing movement against the initial literature.

Connolly (1989) argued that the whole effect disappeared after the 1970s, while Rogalski (1984) asserted that the anomaly could be entirely attributed to the period between Friday's close and Monday's open, and that Monday's returns from open to close did not differ significantly from those on Friday. More recently, Brusa, Liu, and Schulman (2000) showed the existence of a *reverse* weekend effect, whereas Sullivan, Timmerman, and White (2001) are skeptical that the historical results are not examples of data mining. The latest development appears to draw upon the short-selling theory to explain this violation of the EMH.[8]

To determine the impact of all this research, it is convenient to examine investment strategies based on their results. Exploiting the Weekend Effect is simple: buy stocks at the market close on Monday, and sell them at the close on the subsequent Friday. The cumulative returns attributed to this strategy as hypothetically applied to the S&P 500®, compared to the S&P 500 itself, are shown in Exhibit 1.

The log scale of Exhibit 1 allows us to observe the growth rate of cumulative returns. Until the early 1970s, the strategy's returns increased at a fairly constant rate, which appears to be reduced after this period; there appears to have been a change in the pattern of excess returns.[9]

This change is better illustrated when looking at the difference in average daily returns between the strategy and the market, i.e. the difference between the average return of the S&P 500 on Tuesdays, Wednesdays, Thursdays, and Fridays, and the average return on all five days of the trading week including Monday. As Exhibit 2 shows, a downward trend clearly started in the early 1970s, with the exception of the late 1980s, and a reverse in the downward trend emerges around 2000.

If the research confirming the anomaly's existence was convincing enough at the time, we might suppose late 1970s investors frequently sold stocks late on Fridays and bought them back on Mondays to capture the *ex-ante* returns. The expected consequence is that the more investors exploit the Weekend Effect, the worse the performance on Fridays would be, the better the performance on Mondays would be, and the lower the returns would be for such investors going forward.

This is exactly what we see in Exhibit 2; the downward trend starting in 1974 came one year after Cross' paper. The sharp increase in the difference just after 1984 coincides with Rogalski's paper questioning the Weekend Effect—and if Rogalski's paper dissuaded investors from avoiding Mondays, it takes little imagination to suppose that the "Black Monday" of October 1987 provided grounds to reconsider. A positive trend emerged

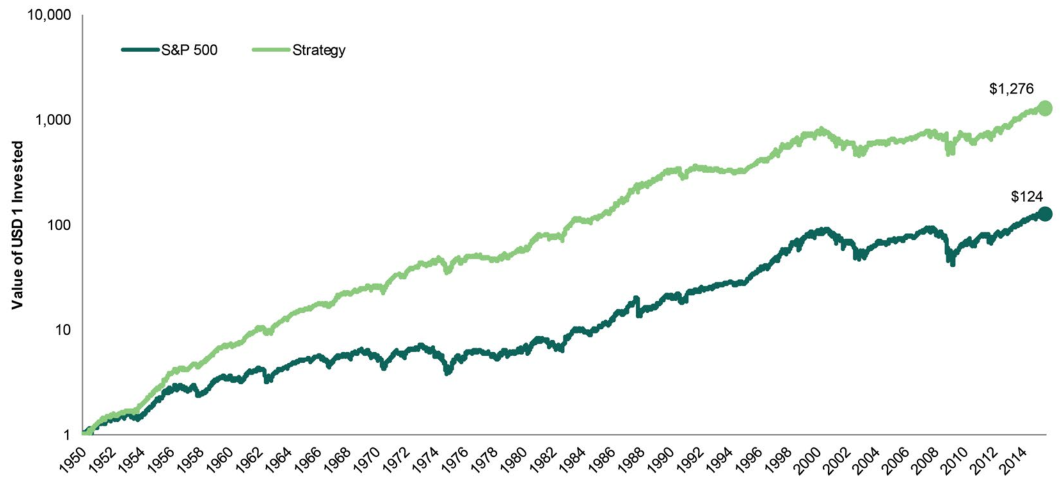

Exhibit 1: Exploiting the Weekend Effect in U.S. Equities
Source: S&P Dow Jones Indices LLC. Data from December 1949 to June 2015.
Past performance is no guarantee of future results. Chart is provided for illustrative purposes and reflects hypothetical historical performance. Please see the Performance Disclosures at the end of this document for more information regarding the inherent limitations associated with back-tested performance.

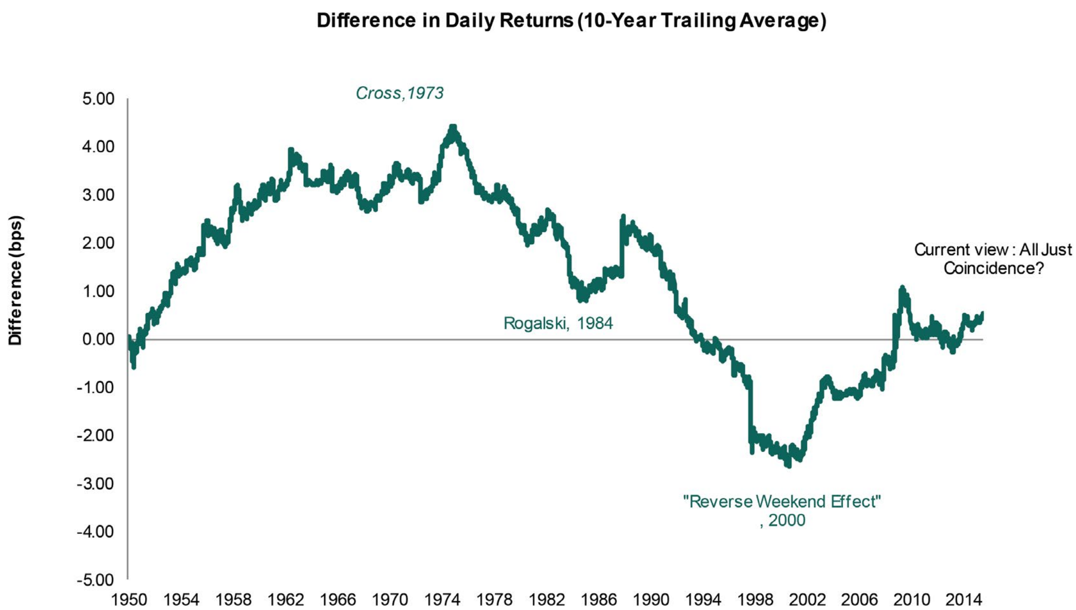

Exhibit 2: What a Difference a Day Makes
Source: S&P Dow Jones Indices LLC.
Data from December 1949 to June 2015. Line represents difference in performance between the average return of the S&P 500 on Tuesdays, Wednesday, Thursdays, and Fridays and the average return on all five days of the trading week including Monday. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

around 2000, during which there was growing skepticism about the statistical techniques used in previous research.[10] Brusa, Liu, and Schulman (2000) also published evidence in favor of a reverse Weekend Effect. Hence, the inflection points and overriding trend in the data appear to be explained by the stance of prominent research papers of the time.

As a result, the Weekend Effect exemplifies the disappearing anomaly; the pattern of returns is impacted as expected, and the returns themselves are arbitraged away as investors become aware of the anomaly's existence. The strategy itself is also easy to understand and act upon without suffering undue trading costs (using futures, for example); a characteristic that most certainly accelerated its disappearance.

Statistical Anomalies

"Get your facts first, then you can distort them as you please."
– Mark Twain

We have assumed so far that anomalies, and their disappearance, can be explained by some coherent economic or behavioral argument. In the case of the Weekend Effect, a behavioral argument involving the timing of bad news created the anomaly, and arbitrageurs' responses diminished it. But is this always a valid assumption?

The quantity of information at our fingertips today is without historical precedent. Coupled with advances in computer processing power, these data enable investors to fit many relationships within financial markets that, they believe, will provide some competitive edge. Unsurprisingly, a large number of relationships have been identified and many strategies continue to be proposed in order to obtain anomalous returns. It is possible, however, that the people proposing these investment ideas are, knowingly or otherwise, distorting the facts. In particular, what if there is no explainable pattern in returns because the returns only ever existed due to chance?

Competing with the Dutch tulip market for historical infamy, the stock market crash of the 1720s has become known as the "South Sea Bubble." After the British South Sea Company made extravagant claims about the potential value of trade deals with the New World, investors readily bought stock. But after the company's share price increased tenfold during 1720, many began selling. This downward pressure caused prices to fall, which created a liquidity crisis as leveraged investors faced margin calls. Individuals were left bewildered by the stock's wild gyrations; one of the numerous people to be left out of pocket, Isaac Newton, commented after the crash, "I can calculate the motion of heavenly bodies, but not the madness of people."

In 1992, David Dolos began to use the daily price records of South Sea Company stock to generate extraordinary profits trading the Dow Jones Industrial Average. His trading rule was simple: starting in December 1992 (for the Dow®) and starting with the South Sea Company's stock price as of August 11, 1719, if the South Sea Company's daily price increased (decreased), Dolos bought (sold short) the Dow. The next month, his position in the Dow was determined by the next day's return from the South Sea Company. Exhibit 3 shows the cumulative returns from this strategy through the end of March 2008.[11]

The strategy performed admirably, delivering triple the Dow's increase over the period. Since Dolos' discovery was not widely publicized, it is unsurprising that the anomaly persisted; if arbitrageurs were unaware of the relationship then their behavior could not have diminished it. Consequently, using such a strong predicative indicator should have made Dolos a rich man, especially during 2008-2009, when relatively few investors were able to avoid the effects of the global financial crisis. As Exhibit 4 shows, however, Dolos had no such luck.

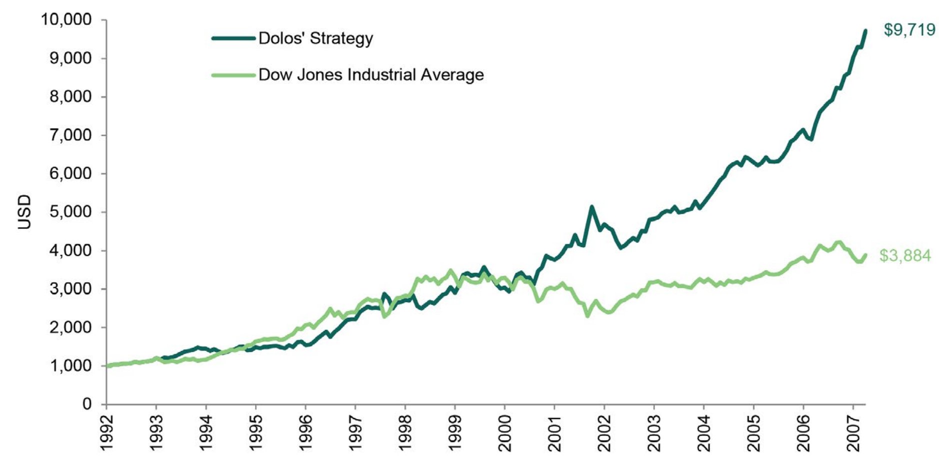

Exhibit 3: Dolos' South Sea Strategy
Source: S&P Dow Jones Indices LLC.
Data from December 1992 to March 2008. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

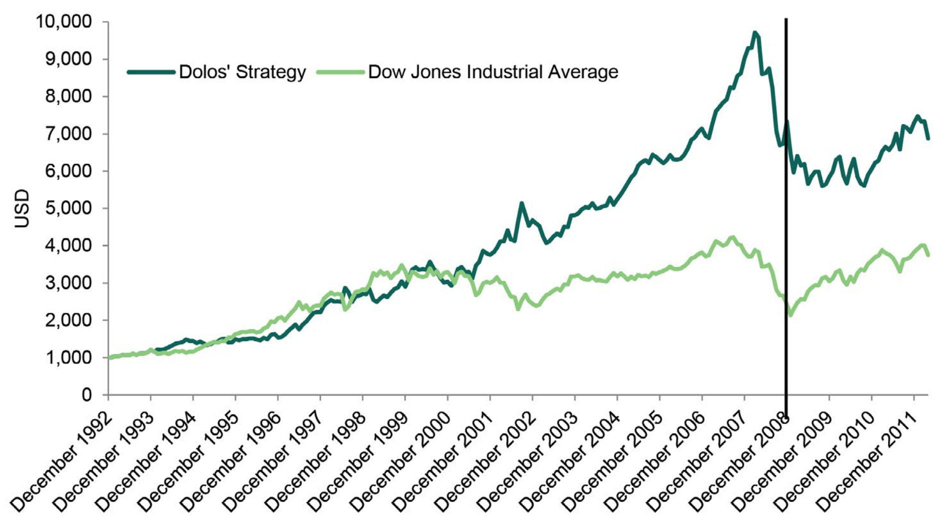

Exhibit 4: Dolos' South Sea Strategy Unravels
Source: S&P Dow Jones Indices LLC.
Data from December 1992 to December 2011. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

The strategy's cumulative returns fell dramatically after 2007, reflecting a breakdown in the predictive relationship. So what changed to influence this trend? The answer is: nothing!

David Dolos never discovered, traded, or wrote about this strategy; in fact, David Dolos never existed at all. (Scholars of Greek mythology may recall that Dolos is the spirit of trickery and guile.) The purpose of this trickery was to show how easy it can be to "mine" data using large datasets; by assigning 1s and 0s to prices that went up or down, respectively, it is straightforward to find a match using the power of computer processing. The relationship broke down because there was no more reason for its existence in the first place than coincidence—some string of 1s and 0s will yield the longest match, and it just so happens that this match has been shown on the graph between December 1992 and March 2008.

Another way to view the chance aspect of this type of anomaly is through statistics. As John Allen Paulos pointed out, "uncertainty is the only certainty there is."[12] Relatedly, the discovery of an anomaly via the use of statistical techniques is accompanied by a confidence level. This confidence level provides an indication of how likely it is that the relationship found may have arisen by chance, simply through random variations in the data.

Confidence intervals are powerful tools for isolated tests, but they are increasingly meaningless as the search broadens, a fact that means that the risk of statistical anomalies is frequently underestimated. For example, suppose an investment is proposed exploiting the predictive power of an accounting statistic—revenue per salesperson, for example. The proposer states that he has identified a profitable relationship with share prices and tells you, with a 95% degree of confidence, that the relationship has not arisen through chance alone. Dangerously, the proposer also looked at 100 different accounting statistics before finding one that worked. However, if the 95% confidence interval is correct, then by chance alone one might expect to find relationships for 5 of the 100 accounting statistics with similarly strong—yet entirely misplaced—confidence. In such circumstances, the high confidence interval provides scant comfort; if there were only one relationship found at that level of confidence, it would seem much more likely to be casual than causal. Combined with the real-world truth that researchers have tested the predictive power of thousands of statistics in manifold combinations, we should be exceedingly cautious of those few showing sufficiently convincing performance to merit inclusion in a sales pitch.

The statistical anomaly category acts as a note of caution to investors. Worse, its appearance is not limited to pure coincidence; how do you distinguish between a strong relationship and weak relationship when the weak relationship benefits from recent good fortune? There is no silver bullet to distinguish meaningful from meaningless coincidences, but there is an armory of more prosaic weapons.[13] Two types of analysis are particularly useful; the first is to extend samples beyond the time frame (or assets) in which the relationship was found. Second, and arguably more important, is a robust and critical examination of the economic reasoning behind relationships. If possible, the reasoning should be tested in other ways; for example if for U.S. stocks a high revenue per salesperson in one quarter predicts an increase in share prices the next, does the same hold in each sector? Does it work for smaller stocks and larger stocks? Does it work for Canadian companies? What happens during and after mergers of companies with differing statistics?

Nonetheless, it remains difficult to distinguish the merit of newly found strategies with sparse history, or when the proposed explanations are conceptually challenging.

Attenuation

"Every side of a coin has another side." – Myron Scholes

Risk and return in financial markets are two sides of the same coin—investors should be extremely wary of considering one without the other. Our analysis thus far has focused only on the return side of the coin, since the disappearance of arbitrageable or chance returns does not warrant an analysis of risk. Some observed effects, however, are *attenuated* by greater awareness. Our attenuation category includes anomalies which can, in principle, be impacted by increasing awareness, but where the impact is to increase the associated risk (or otherwise to adjust the balance of risk and reward). If the returns are simply a reward for risk, this is obviously grounds to expect their persistence, an explanation for why they are unlikely to be arbitraged away, and a reason for caution in investment.

In order to provide an example of an attenuated anomaly, we turn to momentum. There is a stark simplicity to the concept of trend-following and—as an informal heuristic to capital allocation—it is probably as old as commerce itself. Momentum was first formalized into a systematic investment strategy no later than the late 19th century, as a part of Dow Theory. At least as early as the 1930s, the question of its effectiveness was the subject of celebrated academic pursuits.[14] The history of momentum is rich in controversy and characters, with the post-war development of both modern financial theory and computing power, a stream of papers debated its existence and potential genesis.[15] However, the field was stacked with oddballs and fans of esoteric technical analysis; it took a different approach to bring momentum to wider prominence.

The most influential paper in the field is arguably Mark Carhart's 1997 study, which showed that adding a momentum factor to the Fama-French three-factor model considerably increased the model's explanatory power.[16] With momentum understood as a key factor in describing *cross-sectional* returns, the returns to that factor began to be broadly incorporated into risk management and active management processes; a multitude of investors took notice of its performance. Momentum has a complicated interaction with its own popularity. In the case of the Weekend Effect, its systematic exploitation acted to diminish returns, but in the case of momentum, greater awareness is initially self-reinforcing: the greater the demand for winners, the more they should continue winning. We argue that this feedback loop may give rise to a systematic instability, with continued outperformance leading to a risk of increasingly material drawdowns.

To examine the performance of momentum, the natural starting place is the so-called 12-month-1-month momentum strategy (12M-1M). It forms the basis of Carhart's extension of the Fama-French three-factor model and has since become the default expression of momentum's performance in the investment community more generally. It is also a simple strategy: as first documented in Jegadeesh and Titman's 1993 paper, the 12M-1M momentum of a security is simply its 11-month return up to one month ago. Practically, it can be viewed as an 11-month momentum strategy executed with a one-month delay.

Another justification for using 12M-1M momentum is that its prominence has resulted in the wide availability of long-term data for analysis. Exhibit 5 shows one such example, the hypothetical performance of a momentum strategy based on U.S. equities going back to 1947.[17] The performance shown in Exhibit 5 is constructed as follows: calculated monthly, the return of the momentum strategy is the difference in performance between two hypothetical portfolios, each constructed from a broad universe of listed U.S. stocks. The first portfolio comprises stocks with momentum in the top tertile among all stocks, the second portfolio comprises stocks in the bottom tertile, and the weight of each stock in each portfolio is calibrated so that neither company size nor book-to-market value differs significantly between the two hypothetical portfolios.[18] Thus, the performance of the strategy approximates those returns to momentum that are not generated by an unintended bias for cheap or smaller stocks.

As Exhibit 5 shows, between 1944 and 2015, there was a definite upward trend in the cumulative returns attributed to the momentum factor. The near straight-line performance of the strategy from 1943 to the end of the century implies a consistent growth rate more or less unvaried over decades. There appears to be some change in the pattern of returns beginning in the late 1990s, which coincides (among other things) with Carhart's influential 1997 paper, but the upward trend remains. Indeed, if we discount the performance during the 2008-2009 financial crisis, an outlier event, the returns attributed to momentum are more or less persistent. In summary, advertisement of the strong performance of the 12M-1M strategy seems to have had little impact on its returns.

But the pattern of returns *did* change. The graph in Exhibit 5 clearly becomes more volatile after the late 1990s; successes come at an increased cost. As noted in the start of this section, momentum strategies can be initially self-reinforcing. Stocks with strong price performance are bought by momentum followers, which drives up prices further and subsequently provides momentum with an even more compelling track record and more followers. As long as this continues without correction, bubbles

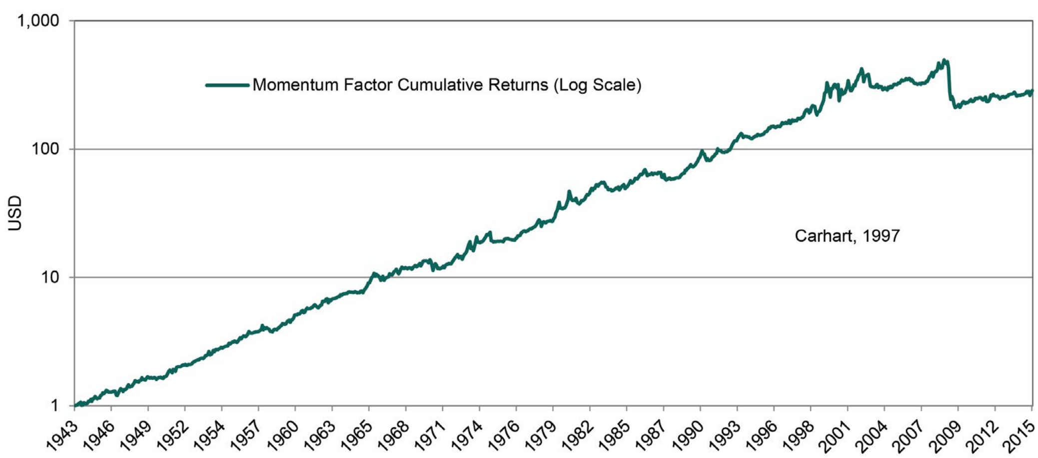

Exhibit 5: The Momentum of Momentum
Source: S&P Dow Jones Indices LLC. Data from December 1943 to June 2015.
Line shows cumulative hypothetical return of difference between high and low momentum portfolios. Past performance is no guarantee of future results. Chart is provided for illustrative purposes and reflects hypothetical (back-tested) historical performance. Back-tested data is subject to inherent limitations because it reflects application of a methodology in hindsight.

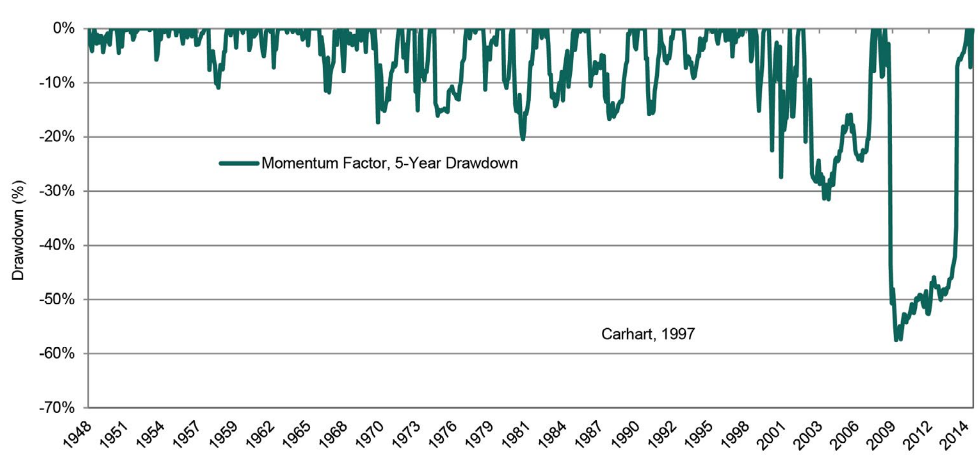

Exhibit 6: Increasing Drawdowns Over Time in Momentum
Source: S&P Dow Jones Indices LLC. Data from 1948 to 2014.
Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

in the valuations of single equities are likely to form and become exaggerated. But even the most committed follower of momentum has a modicum of historical awareness, and experience tells us that *at some point*, stock valuations become so excessive that reality bites. Previous winners will become viewed as the most overpriced; a downturn hurts those stocks with positive momentum harder. As winners become losers, momentum chasers rush to sell. Those investors who wait a month to reassess their positions are hit harder still. Experience therefore suggests that as momentum strategies become increasingly popular, their propensity to generate losses during market corrections should increase.

Exhibit 6 demonstrates the increasing drawdown risks faced by the 12M-1M strategy. Specifically, the exhibit compares the cumulative return of the strategy at any point to its highest level over the previous five years, a measure of the hypothetical losses faced at the time by an investor who entered at the recent "top".

Exhibit 6 shows that while the 12M-1M momentum strategy may have continued to add returns, its downside risk has increased, especially since 1997. Carhart's paper seems relevant because such a widely read piece of research is likely to have increased the awareness and popularity of momentum strategies; certainly its publication marks a period of dramatically increased drawdowns. On a longer time scale it would appear that in fact the downside risk in momentum has been increasing since the end of WWII.

In conclusion, 12M-1M momentum epitomizes the existence of strategies for which research and popularity have not—as yet—triggered a disappearance of returns. On the surface, such persistence would appear attractive. However, the returns have come at an increasing risk, with the current risk profile appearing more elevated than ever. It may well be that the risk attributable to momentum strategies normalizes in the future, with the additional return attributable to momentum varying commensurately with the (informed) risk preferences of market participants. Or, the risk may continue to increase until its realization convinces a wide audience (including academics) to demote 12M-1M momentum from its current position as a celebrated anomaly. In either case, this risk-based attenuation of anomalous returns is conceptually possible for a majority of popular strategies, and analyzing the risk-adjusted returns attributable to strategies becomes a vital component of their assessment.

Persistence

"No matter how beautiful the theory, one irritating fact can dismiss the entire formulism, so it has to be proven" – Michio Kaku.

Some of the most elegant financial theories are also those with results that can be digested easily and have significant ramifications for investors' behavior. In our attempts to identify anomalies that can, in principle, be affected by popularity but which show return persistence without an increase in downside risk, it seems reasonable to consider an anomaly with a fairly stable risk profile.[19]

The idea that investments should offer returns commensurate to their risk, as put forward by the CAPM, is one of the cornerstones of financial theory. However, the irritating fact that contradicts this theory is the low-volatility anomaly. It was first discovered by Haugen and Heins in 1975, when they found that stocks with lower volatility in monthly returns experienced greater average returns than for the high-volatility stocks.

Rather than this discovery standing alone against a bank of literature questioning Haugen and Heins, many other papers have supported the initial findings. Similar to Haugen and Baker's (1991) work, Jagannathan and Ma (2003) showed that investing in a minimum variance portfolio delivered higher returns and lower risk in the U.S. than for the cap-weighted benchmark. In global markets, Carvalho, Xiao, and Moulon (2012) found the highest Sharpe ratio of many investment strategies was a minimum variance portfolio, while Blitz and van Vliet (2007) found a 12% spread between low- and high-volatility decile portfolios, even after accounting for value and momentum effects. More recently, various authors have shown that such anomalous effects appear to be present in most equity markets, globally.[20]

With broad evidence of a low-volatility anomaly in different markets and timeframes, and cogent behavioral and economic arguments available in support, it seems there is more than a spurious relationship at work. However, there has been growing demand for low-volatility strategies after the financial crisis of 2008, while easily accessible vehicles such as ETFs have removed barriers to constructing portfolios exploiting the anomaly and popularized the concept. The increasing awareness and popularity of low-volatility strategies leads us to wonder if the return patterns for strategies based on this anomaly have been affected—by either increased risk or diminished return.

However, if we look at the cumulative returns to the S&P 500 Low Volatility Index—either since its launch in 2011 or to the full extent of its back-tested performance since 1990, this is not the case.[21] Exhibits 7 and 8 demonstrate this persistence—first by a direct comparison of total return and, second, by comparing the risk-adjusted excess return of the S&P 500 Low Volatility Index to that of the benchmark S&P 500.

Exhibit 7 demonstrates the persistence of an excess return, but it requires us to check that such persistence has not come at the expense of increased risk. It's appropriate to evaluate the strategy's risk on a relative basis (i.e., in comparison to a market benchmark) and over a suitably long period to capture longer-term trends.[22] The risk-adjusted relative return shown in Exhibit 8 is calculated as follows: at each point in time, the previous six-year daily volatility of returns for both the S&P 500 Low Volatility Index and the S&P 500 are calculated, and the six-year total return of the S&P 500 is multiplied by the ratio of the two volatilities to derive a "risk-adjusted benchmark return." The risk-adjusted benchmark return is thus the return of the S&P 500, but scaled to the volatility of the low-volatility strategy. The risk-adjusted *relative* return is the six-year return of the S&P 500 Low Volatility Index, minus the risk-adjusted benchmark return. Thus, the risk-adjusted relative return is the excess (or deficit) return in the strategy compared to the volatility-scaled benchmark's return. If the risk-adjusted relative return is greater than zero, we appear to be earning a greater return than might be expected given the strategy's risk, and vice-versa. The results are shown in Exhibit 8.

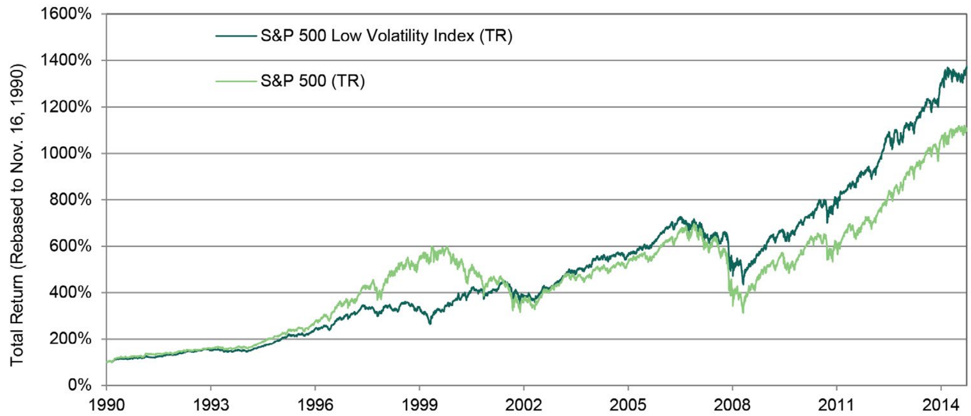

Exhibit 7: S&P 500 Low Volatility Index Outperformance
Source: S&P Dow Jones Indices LLC.
Data from November 1990 to August 2015. Past performance is no guarantee of future results. Chart is provided for illustrative purposes. Some data for the S&P 500 Low Volatility Index reflect hypothetical historical performance. Please see the Performance Disclosures at the end of this document for more information regarding the inherent limitations associated with back-tested performance.

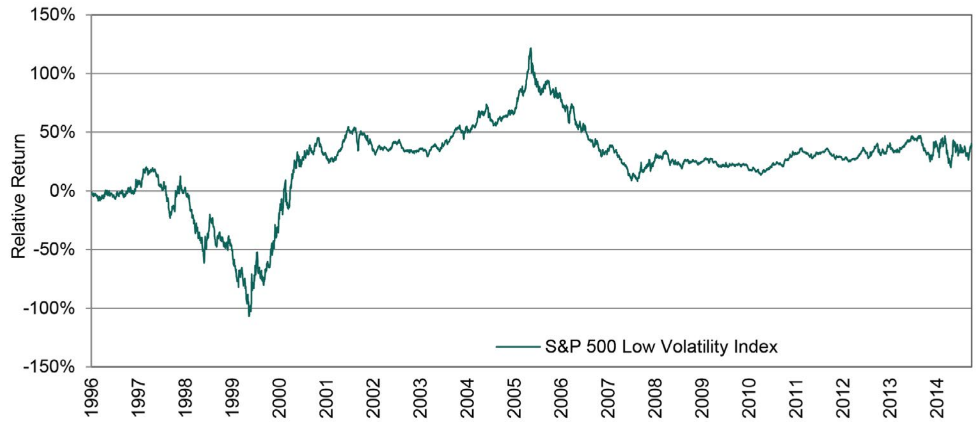

Exhibit 8: S&P 500 Low Volatility Index Six-Year, Risk-Adjusted Relative Return
Source: S&P Dow Jones Indices LLC. Data from 1990 to 2014.
Past performance is no guarantee of future results. Chart is provided for illustrative purposes and reflects hypothetical historical performance. Please see the Performance Disclosures at the end of this document for more information regarding the inherent limitations associated with back-tested performance.

Aside from two periods around 2000 and 2008, the pattern of risk-adjusted annual returns remains relatively flat; the oscillations persist around a stable, positive mean. If anything, notwithstanding those two major events, the level of the long-term, risk-adjusted relative returns would appear to be *increasing* over time. In particular, the current reading (covering the years since the market for U.S. equities began its remarkable bull run) is as good as, if not better than, what might be expected from history and current circumstances.

The S&P 500 Low Volatility Index provides a particularly resonant example of persistent anomalous returns that are not easily dismissed as a compensation for risk. However, a note of caution is still needed. All that Exhibits 7 and 8 demonstrate conclusively is that, *so far,* the investment and attention directed toward low-volatility strategies has not been sufficient to temper their returns or attenuate their risk/return profile. This can be taken as an indication that, whatever investment flows or perspectives give rise to the anomaly, they exceed those set to exploit it—by several orders of magnitude. As such, this analysis may provide a degree of comfort to investors considering such strategies.

Conclusion

"In theory, there is no difference between theory and practice. In practice, there is." – Yogi Berra

Some might see our attempts to categorize anomalies as a fact-finding mission that has little practical benefit or a zoo-like menagerie of some things that have happened to some anomalies and may happen to others, but this would miss the bigger point. In particular, we stress that investors should be wary of analyzing returns in isolation without any consideration for the associated risk, and that seemingly persistent returns may actually be a reward for thus far *unappreciated* risks**.**

More important, arguably, is an awareness of the chance relationships in large datasets; the power of computers means that an increasing number of these relationships can be found at an exponentially increasing risk of confusing the spurious with the causal. Moreover, the sophisticated explanations proposed for some statistical anomalies can make this effect fiendishly difficult to identify and avoid. To reduce the possible impact of unanticipated changes in the returns' patterns, solutions such as extending samples and thinking about the economic reasoning are on offer.

It would be naïve to expect persistent performance from anomalies that rely on investors behaving insensibly, are easy to trade, and that are not a reward for risk—unless evidence suggests that the bank of investors offering to be exploited is deep pocketed and broadly populated. Examining the performance of strategies as they are popularized by broadly cited academic papers and offered in products made widely available allows us to glean information about what is driving their unexpected returns, and the potential for those returns either to continue or to come at the price of increased risk. This provides a toolkit to use when assessing the success of many strategies.

References

Asness, Cliff, "How Can a Strategy Still Work If Everyone Knows About It?" August 31, 2015.

Black, Fischer, Michael C. Jensen & Myron S. Scholes, "The Capital Asset Pricing Model: Some Empirical Tests", *Praeger Publishers Inc.,* 1972.

Blitz, D and P. Van Vliet, "The volatility effect: Lower risk without lower return", *Journal of Portfolio Management,* July 2007.

Brusa, Jorge, Pu Liu and Craig Schulman, "The Weekend Effect, 'Reverse' Weekend Effect, and Firm Size", *Journal of Business Finance and Accounting,* June 2000.

Carhart, Mark M., "On Persistence in Mutual Fund Performance" *The Journal of Finance,* March 1997.

Chan, Fei Mei and Craig J. Lazzara, "Is the Low Volatility Anomaly Universal?" April 2015.

Chan, Louis K., Narasimhan Jegadeesh and Josef Lakonishok, "Momentum Strategies", *The Journal of Finance,* December 1996.

Chen, Honghui and Vijay Singal, "Role of Speculative Short Sales in Price Formation: The Case of the Weekend Effect", *The Journal of Finance,* March 2003.

Connolly, Robert A. An Examination of the Robustness of the Weekend Effect", *Journal of Financial and Quantitative Analysis,* June 1989.

Cowles III, Alfred, "Can Stock Market Forecasters Forecast", *Econometrica*, Jul 1933

Cross, Frank, "The Behavior of Stock Prices on Fridays and Mondays", *Financial Analysts Journal*, November/December 1973.

Daniel, Kent, David Hirshleifer and Avanidhar Subrahmanyam, "Investor Psychology and Security Market Under- and Overreactions", *The Journal of Finance,* December 1998.

Fama, Eugene F., "Market Efficiency, long-term returns, and behavioral finance", *The Journal of Financial Economics,* September 1998.

Fama, Eugene F. and Kenneth R. French, "The Cross-Section of Expected Stock Returns," *The Journal of Finance*, June 1992.

French, Kenneth, "Stock Returns and the Weekend Effect", *Journal of Financial Economics,* February 1980.

Harvey et al "… and the Cross Section of Expected Returns", *National Bureau of Economic Research,* February 2015.

Haugen, Robert A. and A. James Heines, "Risk and the Rate of Return on Financial Assets: Some Old Wine in New Bottles", *Journal of Financial and Quantitative Analysis,* December 1975.

Haugen, Robert A. and Nardin L. Baker, "The efficient market inefficiency of capitalization -weighted stock portfolios", *The Journal of Portfolio Management,* Spring 1991.

Jagannathan, Ravi and Tongshu Ma, "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps", *The Journal of Finance,* August 2003.

Jegadeesh, Narasimhan and Sheridan Titman, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency", *The Journal of Finance,* March 1993.

Lazzara, Craig J., "The Limits of History," January 2013.

Leote de Carvalho, R., X. Lu, P. Moulin, "Demystifying Equity Risk-Based Strategies: A Simple Alpha plus Beta Description", *The Journal of Portfolio Management,* Spring 2012.

Levy, Robert, "Relative strength as a criterion for investment selection", *Journal of Finance,* December 1967.

McLean & Pontiff, "Does Academic Research Destroy Stock Return Predictability", Forthcoming.

Keim, Donald B., "The Cost of Trend Chasing and the Illusion of Momentum Profits", *The Rodney L. White Center for Financial Research,* May 2003.

Paulos, John Allen, *A Mathematician Plays the Stock Market*, *Basic Books,* 2003.

Rogalski, Richard J., "A Further Investigation of the Weekend Effect in Stock Returns: Discussion", *The Journal of Finance,* July 1984.

Rouwenhorst, K. Geert, "International Momentum Strategies", *The Journal of Finance,* February 1998.

Sullivan, Ryan, Allan Timmermann and Halbert White, "Dangers of data mining: The case of calendar effects in stock returns", *Journal of Econometrics,* November 2001.

Endnotes

1. Asness, Cliff, "How Can a Strategy Still Work If Everyone Knows About It?" Aug. 31, 2015.
2. Fama, Eugene F. and Kenneth R. French, "The Cross-Section of Expected Stock Returns," The Journal of Finance, June 1992.
3. Asness (op. cit.) argues that the anomalies come about "because investors make errors."
4. See BlackRock Global ETP Landscape, December 2014, p. 4. "Organic growth for smart beta is 18%, twice that of market-cap weighted equity ETPs."
5. The authors acknowledge their debt in particular to two papers that inspired their approach, namely Harvey et al "… and the Cross Section of Expected Returns" (2015) and McLean & Pontiff, "Does Academic Research Destroy Stock Return Predictability?" (forthcoming).
6. To clarify, disappearing anomalies really do exist, for a while, until they become widely appreciated, at which point they vanish. Statistical anomalies, in the sense used here, are mirages—there's really nothing there, and never was—although with enough data mining, an effect may appear to be real.
7. The results given in Cross' paper are for the S&P Composite 1500® between Jan. 2, 1953, and Dec. 21, 1970. Similar results were found for the Dow Jones Industrial Average® and the New York Stock Exchange Composite Index, but these were not included in the paper.
8. See Chen and Singal (2003).
9. The fact that the October 1987 crash occurred on a Monday might cause concern over the dominance of extreme events in such results. In fact, once removing extremes from the data, both the original Weekend Effect and its disappearance during the 1980s remain evident.
10. See Sullivan, Timmerman, and White (2001) for a more detailed discussion on the critiques of statistical techniques used to derive evidence in favor of the Weekend Effect.
11. South Sea daily returns are those between Aug. 11, 1719, and June 29, 1720 (source: International Center for Finance at Yale). The monthly returns on the DJIA are those between Dec. 31, 1992, and March 31, 2008.
12. See Paulos, John Allen, A Mathematician Plays the Stock Market, 2003.
13. See Lazzara, Craig J., "The Limits of History," January 2013.

14. See Cowles (1933)
15. See Swinkels (2003) for an overview.
16. Carhart, Mark M., "On Persistence in Mutual Fund Performance." The paper has 8,985 citations on Google Scholar, as of Aug. 18, 2015, which ranks highest for all the research papers on momentum we analyzed. See also Fama and French, op. cit.
17. In fact, performance is available going back to 1924; we exclude the pre-war period in part acknowledgement of the very different market environment of the time, but the reader may be interested to know that the market crash of 1929 represented a reversal in momentum's performance far greater than any seen since.
18. Full details on the construction of the momentum factor, as well as a downloadable return series, are available in the French Factor Library at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library.html.
19. The result was anticipated by the observation that market beta appeared to be negatively correlated to returns, found in Black, Jensen, and Scholes' earlier 1972 paper; "The Capital Asset Pricing Model: Some Empirical Tests."
20. This spread was found using data between 1986 and 2006 and the paper provides potential explanations for the existence of the anomaly: leverage-confined investors being unable to arbitrage away the returns; inefficient decentralized investment approaches; and behavioral biases among private investors. See also Chan, Fei Mei and Craig J. Lazzara, "Is the Low Volatility Anomaly Universal?" April 2015.
21. The S&P 500 Low Volatility Index comprises 100 stocks that are members of the S&P 500 and have the lowest levels of realized volatility over the previous 12 months. Rebalancing occurs quarterly, with the index weights of each component set at each rebalance in inverse proportion to realized volatility.
22. We chose six years so that the most recent values capture the strong bull market in equities that began in March 2009 and encompass the period over which low-volatility may be said to have gained its current popularity, but the results are not particularly sensitive to the length of period chosen.

About S&P Dow Jones Indices

S&P Dow Jones Indices LLC, a part of McGraw Hill Financial, Inc., is the world's largest, global resource for index-based concepts, data and research. Home to iconic financial market indicators, such as the S&P 500® and the Dow Jones Industrial Average™, S&P Dow Jones Indices LLC has over 115 years of experience constructing innovative and transparent solutions that fulfill the needs of institutional and retail investors. More assets are invested in products based upon our indices than any other provider in the world. With over 1,000,000 indices covering a wide range of assets classes across the globe, S&P Dow Jones Indices LLC defines the way investors measure and trade the markets. To learn more about our company, please visit www.spdji.com.

Authors' Bios

Hamish Preston
Researcher
Department of Economics
Birmingham University

Hamish Preston has a BSc in Economics from the London School of Economics and Political Science and is currently engaged in postgraduate research at the Department of Economics at Birmingham University. At the time this article was written, Hamish was an analyst at S&P Dow Jones Indices.

Tim Edwards
Head
European Index Investment Strategy
S&P Dow Jones Indices

Tim Edwards is head of European Index Investment Strategy for S&P Dow Jones Indices. The group provides research and commentary on the entire S&P Dow Jones Indices product set, including U.S. and global equities, commodities, fixed income, and economic indices.

Prior to joining S&P Dow Jones Indices, Tim worked at Barclays Capital, where he had global responsibility for product development of exchange-traded notes across all asset classes, covering commodities, volatility, foreign exchange, fixed income and emerging markets.

Tim holds a PhD in mathematics from University College London.

Craig Lazzara
Global Head
Index Investment Strategy
S&P Dow Jones Indices

Craig Lazzara is global head of index investment strategy for S&P Dow Jones Indices. The index investment strategy team provides research and commentary on the entire S&P Dow Jones Indices' product set, including U.S. and global equities, commodities, fixed income, and economic indices.

Craig previously served as product manager for S&P Indices' U.S. equity and real estate indices. These include the S&P 500® and the S&P/Case-Shiller Home Price Indices, two of the most widely tracked benchmarks in the world.

Prior to joining S&P Indices in 2009, Craig was a managing director of Abacus Analytics, a quantitative consulting firm serving the brokerage and investment management communities. He previously directed marketing and client service for ETF Advisors and Salomon Smith Barney's Global Equity Index Group, as well as for the Equity Portfolio Analysis group at Salomon Brothers. Earlier, Craig served as chief investment officer of Centurion Capital Management and Vantage Global Advisors, as a managing director of TSA Capital Management, and as a vice president and portfolio manager for Mellon Bank and T. Rowe Price Associates.

A Chartered Financial Analyst, Craig is a graduate of Princeton University and Harvard Business School.