‘The Greatest Anomaly in Finance:’ Understanding and Exploiting the Outperformance of Low-Beta Stocks

By Geoff Considine
February 14, 2012

If I told you that there is an easy-to-exploit market anomaly that has enabled investors to consistently and substantially outperform the market with less risk for more than four decades, your first instinct might be to roll your eyes. After all, the unending quest to improve returns while lowering risk has yielded countless methods with initial promise that subsequently collapse under further scrutiny.

Not so fast. What if I could show that this market anomaly is well-documented in the academic literature – that it is not just some esoteric theory? And that now some newly created ETFs provide a convenient way for advisors to access this strategy?

You might listen a little closer.

As you might have guessed, this is not actually a hypothetical – I’m referring to the demonstrated outperformance of low-beta and low-volatility stocks. Since financial theory dictates that lower risk should imply lower return, this is an anomaly.

In this article, I explore the historical evidence for this powerful anomaly, explain what market mechanisms permit it to exist and persist, and examine the ETFs that exploit it to determine whether it is likely that these funds will provide future outperformance.

”The greatest anomaly in finance”

Support for the effect I describe comes from a recent research paper by Malcolm Baker of Harvard Business School and his co-authors, who called it the “greatest anomaly in finance.” While financial theory dictates that the only way to achieve higher returns is taking on more risk, this has not been historically true for equities. In their 2010 paper titled *Benchmarks as Limits to Arbitrage: Understanding the Low Volatility Anomaly*, Baker and co-authors Brendan Bradley of Acadian Asset Management and Jeffrey Wurgler of NYU Stern School of Business found that selectively investing in portfolios of either low-beta or low-volatility stocks over the 41-year period spanning 1968 through 2008 would have resulted in annualized alphas of 2.6% and 2.1%, respectively. The swings these portfolios experienced were also far less extreme than those of the broader market. If the universe of stocks under consideration is limited to the 1,000 stocks with the largest market capitalizations, the low-beta and low-volatility portfolios generated 3.49% and 2.1% in annualized alpha, respectively.

As this anomaly has become more well-known, ETFs have been specifically designed to allow investors to exploit this source of alpha. The natural question that emerges, of course, is whether this source of outperformance will disappear as investors, recognizing this arbitrage, will exploit this anomaly until it no longer exists. The intriguing conclusion of Baker et al. was that this effect is not disappearing. They proposed a theory that explains their observations and made a compelling case for the future attractiveness of low-volatility and low-beta portfolios.

Baker et al. analyzed historical market data from CRSP, a survivorship-corrected historical research database of equities, and found that grouping stocks into either low-volatility or low-beta cohorts resulted in substantial future outperformance. They found this effect both when they looked at a universe of all publicly traded stocks and when they limited the universe of potential investments to the 1,000 stocks with the highest market capitalizations.

Why it exists, and why it persists

Two puzzling facts must be explained in order to justify Baker et al.'s counterintuitive findings. First, why would some fraction of investors might behave so irrationally as to continue to select a class of equities that have historically underperformed (e.g., high-beta equities)? Second, why don't sophisticated investors arbitrage this imbalance away by buying up the underpriced low-beta stocks?

The answer to the first question is simple. Behavioral finance has demonstrated overwhelmingly that individual investors have a substantial preference for 'lottery ticket' types of returns, in which they have the potential, albeit at very low probability, of receiving a massive windfall. This is a central finding of behavioral finance, and it has been demonstrated time and again.

The second piece of the puzzle is more challenging, and its solution was the key innovation in Baker et al. Many institutional investors, such as pension fund or mutual fund managers, have a mandate to beat a specific benchmark portfolio (typically a market-cap-weighted index) and to minimize tracking error relative to that benchmark. These benchmarks give money managers a substantial disincentive to invest in low-beta or low-volatility stocks, even if they believe that these stocks have substantially positive alphas. The reason is that low-beta stocks, even when they have high expected alpha, will increase tracking error (since, by definition, they are uncorrelated to the index). As a result, as Baker et al. explain, institutional managers with target benchmarks actually have an incentive to exacerbate the low-beta anomaly.

Baker et al. were not the first to propose this mechanism, but they were the first researchers to quantify it, in research they published in the peer-reviewed *Financial Analysts Journal*. In representative cases, they found that an equity fund manager will not start to overweight a stock with a moderately low beta of 0.75 until the expected alpha

exceeds 2.5% per year. Further, the researchers concluded that a rational fund manager would actually underweight a stock if its alpha is on the order of 2% per year.

This is a stunning conclusion – it shows that those managers systematically reject attractive investments.

An important secondary conclusion in Baker et al.'s paper is that low beta is actually the fundamental driver of outperformance, not low volatility. A screen for low beta selects many of the same stocks as a screen for low volatility, but the reason why institutional investors will not arbitrage outperformance away is specifically the low-beta character of these stocks. Recall that the 41-year historical alpha is substantially greater for the portfolios selected on the basis of low beta than for those selected on the basis of low volatility, an empirical observation that buttresses Baker et al.'s explanation of their results.

Baker et al. calculated the low-beta and low-volatility effects in various ways. The simplest and most striking result arose when they created a portfolio from the 200 lowest-beta stocks among the largest 1,000 stocks by market capitalization. This portfolio was formed using trailing five-year monthly return data and updated monthly. This low-beta portfolio generated average annual returns that were 2% per year greater than the average for a market index of large-cap stocks, with less risk than the market as a whole. The alpha that the authors calculated for this portfolio was 3.49%.

The future of low-beta and low-volatility ETFs

Let's now turn to the challenge of diagnosing the potential *future* outperformance of low-beta stocks. None among the new crop of ETFs that focus on low-volatility and low-beta stocks has enough of a historical track record to draw any meaningful conclusions from their past performance. We can, however, analyze these ETFs using a number of standard statistical methods, especially by performing forward-looking Monte Carlo simulations using my Quantext Portfolio Planner (hereafter abbreviated as QPP).

Here are the most popular funds currently available for investing on the basis of low-beta or low-volatility:

Fund	Ticker	Benchmark (Representative Ticker)
S&P500 Low Volatility	SPLV	S&P500 (SPY)
Russell 1000 Low Volatility	LVOL	Russell 1000 (IWB)
Russell 1000 Low Beta	LBTA	Russell 1000 (IWB)
Russell 2000 Low Volatility	SLVY	Russell 2000 (IWM)
Russell 2000 Low Beta	SLBT	Russell 2000 (IWM)

The S&P 500 Low Volatility Index ETF, offered by PowerShares, is designed to follow the S&P Low Volatility Index. The index holds the 100 stocks among the S&P 500 with the lowest volatility. The weight of each stock within the index is determined by the inverse of its volatility. The index uses trailing one-year volatility, calculated from daily returns, and it is recalculated and rebalanced quarterly.

The Russell Low Volatility and Low Beta indices are calculated for both the Russell 1000 (large-cap stocks) and Russell 2000 (small-cap stocks), drawing the lowest-volatility or lowest-beta stocks from each index. The Russell low-volatility and low-beta ETFs in the table are designed to track these indices.

Two unique features distinguish the Russell methodologies from their S&P counterpart. First, their selection of low-beta stocks is derived from projected (rather than historical) beta, which is a more complex but also in some ways more useful approach. Second, their low-volatility selection criterion is trailing two-month volatility for each stock in the respective indexes.

These indices are rebalanced monthly.

I obtained the stock-by-stock holdings for each of the five low-beta and low-volatility ETFs mentioned above. Because a 40-holding portfolio is the largest my analytical tools can process, the portfolios I used to stand in for these ETFs comprised the largest 40 holdings (by market value) in each of these ETFs, maintained at the same relative weights as they appear in the indexes. The results of my projections are below.

Trailing three-year risk, return, and correlations of portfolios comprised of the top-40 holdings of these ETFs

		Three-Years of Historical Data for Current Holdings			
Name	Ticker	Trailing Return	Trailing Volatility	Return / Risk	Average Holding Correlation
S&P Low Volatility	SPLV	15.1%	10.6%	142%	42%
S&P500	SPY	18.7%	18.3%	102%	44%
Russell 1000 Low Volatility	LVOL	20.9%	15.1%	139%	39%
Russell 1000 Low Beta	LBTA	20.1%	13.5%	149%	34%
Russell 1000	IWB	18.6%	18.1%	103%	43%
Russell 2000 Low Volatility	SLVY	29.1%	24.3%	120%	40%
Russell 2000 Low Beta	SLBT	28.9%	16.4%	176%	28%
Russell 2000	IWM	40.1%	27.7%	145%	38%

The upshot of the simulation is unambiguous – QPP expects the outperformance of low-beta and low-volatility portfolios to persist. This is crucial support for the findings of Baker et al, and strong evidence that the continued outperformance of these stocks is well-grounded in basic market theory.

One might miss the significance of these results if they were to look only at the first result, the return of each portfolio. While the trailing average three-year returns of the low-volatility and low-beta portfolios are less than those of their respective parent indices, we must look beyond that simple measure.

Doing so, we see that their return-*to-risk* ratio is markedly higher. The Russell 1000 Low Beta's projected 149% return-to-risk ratio, for instance, is striking for being higher by almost half than the broader Russell 1000's 103% projected ratio.

The Russell 2000 Low Volatility selections have the lowest trailing return-to-risk ratio – lower, in fact, than that of the full Russell 2000 – and they also exhibit the least difference in volatility from their parent index. I believe that this is in part due to that index's use of trailing two-month volatility, which is quite noisy, in selecting and weighting stocks.

Further cause for the decreased benefit seen in low-volatility small-cap stocks, however, is the fact that volatility increases as market cap decreases – a fact that Baker et al. note. Thus the stocks with lowest volatility will also be among the larger-cap stocks in the Russell 2000. This reality means seeking out low-volatility Russell 2000 stocks costs us some of the small-cap boost to performance. For this reason, a small-cap low-volatility index would be the least attractive of these strategies, generally speaking, as the results above suggest.

Digging deeper into the results

Further calculations can better isolate the positive effects of low-beta and low-volatility and help investors choose among the various ETF alternatives.

Because the low-volatility and low-beta equity portfolios have lower risk than the parent indices from which they are built, their performance benefit is not entirely apparent. What we really want to see is the degree of outperformance between the low-beta or low-volatility strategies and the baseline indices when the risks are the same. To make this comparison, we can use hypothetical leverage to assign each of the strategies the same risk level.

If we were to use leverage to bring the risk of the portfolio of the top-40 holdings of SPLV to a level equal to that of the S&P 500, for example, the additional expected return of the leveraged portfolio is projected to be 2.1% per year. This calculation shows up as *Risk-Equivalent Excess Return* in the table below.

Projected risk and return of top 40 fund holdings

Name	Ticker	Expected Return	Expected Volatility	Beta vs. S&P500
S&P Low Volatility	SPLV	7.2%	9.5%	40%
Risk-Equivalent Additional Return vs. S&P500		*2.1%*		
Russell 1000 Low Volatility	LVOL	10.2%	14.8%	75%
Risk-Equivalent Excess Return vs. Russell 1000		*1.0%*		
Russell 1000 Low Beta	LBTA	9.8%	13.0%	63%
Risk-Equivalent Excess Return vs. Russell 1000		*2.1%*		
Russell 2000 Low Volatility	SLVY	15.1%	23.1%	119%
Risk-Equivalent Excess Return vs. Russell 2000		*-0.3%*		
Russell 2000 Low Beta	SLBT	13.8%	16.6%	79%
Risk-Equivalent Excess Return vs. Russell 2000		*4.4%*		
Average Risk-Equivalent Excess Return		**1.9%**		

Presenting the data this way aims to make it even easier to see that the QPP results are consistent with what Baker et al. found: There is a persistent performance advantage from low-beta and low-volatility portfolios, and this advantage is stronger for low-beta portfolios than for low-volatility portfolios. While Risk-Equivalent Excess Return provides a consistent way to see these advantages of the low-beta and low-volatility effects, the calculation using leverage is *not* to suggest that advisors will, or should, use leverage. Rather, it is to show the expected outperformance of low-beta and low-volatility strategies on a consistent basis.

These data also answer an important question most investors will have: How to choose among the various ETF options? Specifically, they suggest that the best way to exploit the low-volatility and low-beta anomalies is to use SPLV. The S&P low volatility strategy has a number of benefits, including simplicity in construction and expected lower turnover. The portfolio of top holdings also has the lowest overall beta with respect to the S&P 500. Another attractive ETF is the low-beta Russell 2000 fund (SLBT). While the more complex methodology and higher potential turnover in this fund are a concern, the high risk-equivalent return benefit for this fund and its low beta are encouraging signs that this fund will be able to reliably exploit the low-beta anomaly going forward.

The real deal

Not only do my Monte Carlo simulations support the findings of Baker et al., they actually take them a step further – the data suggest that one need not even assume low-beta and low-volatility stocks will be underpriced on an individual basis going forward. The advantage the QPP simulation discerns results directly from the portfolio effects of diversification.

This anomaly is real, and there's real reason to believe it's not going away. Baker et al. and other researchers have increasingly mounted a persuasive case that the

outperformance of low-beta stocks is both a real and persistent phenomenon, and their proposed mechanism comprehensively explains why it is likely to remain viable in the future.

The ETFs that now exploit these anomalies are the real deal as well. The low-beta stocks, and to a lesser-extent their low-volatility counterparts, that make up these ETFs should outperform, on a risk-adjusted basis, the S&P 500, Russell 1000 and Russell 2000 indices.

Investors and their advisors are typically not constrained by needing to track a specific equity benchmark, and there's no reason they shouldn't exploit that powerful built-in advantage. The new ETFs in this space provide the simplest, lowest-cost way to do so. But some attention must be paid to aggregate risk targets. The S&P Low Volatility Index, after all, still has lower expected return than the S&P 500.

Those who substitute the S&P Low Volatility ETF for an allocation to the S&P500 can, however, adjust the total portfolio risk levels to maintain risk targets by adding higher allocations to other risky asset classes, such as emerging markets, REITs, technology stocks, or commodities. An even easier alternative: If an equity allocation's risk is reduced (by substituting low-beta or low-volatility ETFs in place of market-cap weighted index funds), compensate by simply increasing the allocation to equities as a percentage of the total portfolio.

In the end, there's little excuse not to embrace this approach. Both historical analysis and forward-looking simulations support the viability of this powerful effect, and the underlying mechanisms proposed by Baker et al. provide further confidence that low-beta and low-volatility equity portfolios will continue to provide substantial benefits. Faced with this preponderance of evidence, even the most hardened skeptic should be asking – isn't it time I exploited "the greatest anomaly in finance"?