The Volatility Effect Revisited

David Blitz is the head of quantitative research at Robeco in the Netherlands, d.c.blitz@robeco.com

Pim van Vliet is the head of conservative equities at Robeco, in the Netherlands, p.van.vliet@robeco.com

Guido Baltussen is the head of quant allocation at Robeco and professor of finance at Erasmus University in the Netherlands, g.baltussen@robeco.com

Abstract: High-risk stocks do not have higher returns than low-risk stocks in all major stock markets. This paper provides a comprehensive overview of this low-risk effect, from the earliest asset pricing studies in the nineteen seventies to the most recent empirical findings and interpretations since. Volatility appears to be the main driver of the anomaly, which is highly persistent over time and across markets, and which cannot be explained by other factors such as value, profitability, or exposure to interest rate changes. From a practical perspective we argue that low-risk investing requires little turnover, that volatilities are more important than correlations, that low-risk indices are suboptimal and vulnerable to overcrowding, and that other factors can be efficiently integrated into a low-risk strategy. Finally, we find little evidence that the low-risk effect is being arbitraged away, as many investors are either neutrally positioned, or even on the other side of the low-risk trade.

Keywords: low risk, low volatility, low beta, minimum variance, anomaly, factor investing, smart beta, low-volatility investing

JEL Classification: G11, G12, G14

* The authors are grateful to Eric Falkenstein, Laurens Swinkels, Winfried Hallerbach, and Bernard Breloer for valuable feedback on an earlier version of this paper.

This paper provides a comprehensive overview of the low-risk effect, i.e. the empirical finding that higher risk is not rewarded with a higher return in global stock markets, nor within other asset classes. Although the main driver of the low-risk effect appears to be volatility, which implies that it is essentially a low-volatility effect, we will use the more neutral term 'low-risk effect' throughout this paper. In the first section we review the empirical evidence for the low-risk effect. In the next section we argue that the low-risk effect cannot be explained by factors such as value, profitability, or exposure to interest rate changes. After this we make the step from theory to practice, discussing the key considerations that come into play when investing based on the low-risk effect with real money. We next argue that there is a little evidence that the low-risk effect is being arbitraged away, as many investors turn out to be either neutrally positioned, or even on the other side of the low-risk trade. Finally, we conclude.

EMPIRICAL EVIDENCE FOR THE LOW-RISK EFFECT

In this section we review the empirical evidence for the low-risk effect, which goes back all the way to the very first empirical asset pricing studies in the nineteen seventies, and which by now consists of an extensive stream of literature that also extends to asset classes other than equities. We also outline the various metrics used to measure risk and discuss commonly cited explanations for the existence of the low-risk effect.

History of the low-risk effect

The roots of the low-risk effect can be traced back to the very first empirical tests of the Capital Asset Pricing Model (CAPM). The CAPM predicts a linear relation between a security's systematic risk, measured by its beta against the market portfolio, and its return. The first empirical tests of the CAPM by Black, Jensen and Scholes [1972], Miller and Scholes [1972], and Fama and MacBeth [1973] found that although higher risk is rewarded with higher return, it is not rewarded enough. In other words, the empirical Security Market Line was observed to be flatter than expected. Haugen and Heins [1975] first recognized the existence of a low-risk anomaly, concluding that "*our empirical efforts do not support the conventional hypothesis that risk – systematic or otherwise – generates a special reward. Indeed, our results indicate that, over the long run, stock portfolios with lesser variance in monthly returns have experienced greater average returns than their 'riskier' counterparts*". Such findings were not interpreted as a major cause for concern though. The general consensus was that, although maybe not perfect, the CAPM did an adequate job at explaining stock prices. The first serious challenge to the CAPM was the finding of Banz [1981] that small stocks had higher returns than large stocks, even after correcting for the fact that the average small stock is more risky than the average large stock.

It took until the nineties until the true failure of the CAPM became clearly visible. Fama and French [1992] found that market beta is entirely unpriced in the cross-section of stock returns when size and market beta are properly disentangled from each other. Quoting from the abstract of their paper: "*when the tests allow for variation in beta that is unrelated to size, the relation between market beta and average return is flat, even when beta is the only explanatory variable*". Whereas anomalies such as size and value imply that the CAPM may need to be augmented with some additional factors, the low-risk anomaly challenges the heart of the

CAPM, i.e. the very notion that higher risk should be rewarded with higher return in the cross-section.

Unknown at the time, all these studies suffered from a delisting bias in the data that had not been identified and resolved yet. Shumway [1997] and Shumway and Warther [1999] found that stocks which are delisted, e.g. due to bankruptcy, tend to have very negative returns in the month of delisting, which were either not at all or incorrectly recorded in old CRSP tapes. This bias caused stock returns in general to be overestimated, but particularly so the returns of small and risky stocks, as the probability of delisting is higher for such stocks; see Cochrane [1999]. As a result, the relation between risk and return appeared to be more positive than it actually was, and also the magnitude of the size premium was overestimated.

The Kenneth French data library[1] nowadays contains bias-free historical return series for US stock portfolios sorted on 60-month market beta, with data going back to July 1963. In Exhibit 1 we plot the performance of ten decile portfolios sorted on beta, and in Exhibit 2 we plot the performance of the 5x5 size/beta portfolios that are also available from this source. All Exhibits are based on publicly available data only and include a second-order polynomial trendline. The graphs clearly show a flat, or even slightly negative relation between risk and return, which implies that, over a sample period which is twice as long by now, the conclusion of Fama and French [1992] that beta is not a priced factor still holds.

INSERT EXHIBIT 1 HERE

INSERT EXHIBIT 2 HERE

The rejection of the CAPM relation by Fama and French [1992] was confirmed by several studies which appeared shortly afterwards. Black [1993] finds that the raw relation between beta and return became flat in the decades following the sample period that was covered in the Black, Jensen, and Scholes [1972] study. Falkenstein [1994] even reports a negative relation between risk and return, when applying the crucial control for the size effect. Despite this compelling empirical evidence, the academic community did not abandon the CAPM relation between beta and return, and the investment community ignored the investment opportunity offered by these insights.

Fifteen years after the seminal Fama and French [1992] study, Blitz and van Vliet [2007] take a fresh look at the low-risk effect and find that, if anything, it has become even stronger over time. Over their 20-year sample period, the relation between risk and return is not merely flat, but even outright inverted, with a top-minus-bottom CAPM-alpha spread of 12% per annum. They also find that total volatility appears to be at least as effective as beta, and that the effect is not only present in the U.S. equity market, but also in international equity markets.

The Paradox Investing website[2] contains publicly available data for volatility-sorted deciles portfolios. These portfolios are constructed by sorting the 1,000 largest US stocks on their past 36-month volatility, with data starting in January 1929, as in van Vliet and de Koning [2017]. Exhibit 3 shows that the full-sample relation between risk and return is clearly flat instead of upward sloping, and even becomes inverted in the highest-risk spectrum. Exhibit

[1] http://mba.tuck.dartmouth.edu/pages/faculty/ ken.french/data_library.html
[2] https://www.paradoxinvesting.com/data

4 shows that this result is robust over time, when we separately consider the 1963-1990 sample period of Fama and French [1992], the pre-1963 "pre-sample" period which resembles the original sample period of Black, Jensen, and Scholes [1972], and the post-1990 "out-of-sample" period.

INSERT EXHIBIT 3 HERE

INSERT EXHIBIT 4 HERE

Subsequent studies by Baker, Bradley, and Wurgler [2011], Baker and Haugen [2012], and Frazzini and Pedersen [2014] all confirm and extend these empirical results, using various definitions of risk and proposing various explanations for the findings. A parallel stream of literature examined the empirical performance of the theoretical minimum-variance portfolio and also found clear evidence for the existence of a low-risk anomaly; see Haugen and Baker [1991, 2010] and Clarke, de Silva, and Thorley [2006, 2011]. Following these studies, dedicated low-risk investing has gradually involved into a widely accepted investment approach, typically labelled *low-volatility*, *managed volatility*, *minimum volatility*, *minimum variance*, *defensive*, or *conservative*. Still, theoretical thinking continues to be shaped by the CAPM. An illustration of this is that although academics disagree about the specific set of factors that should be included in asset pricing models, the classic CAPM relation between market beta and expected return remains the starting point in all such models.

The low-risk effect is universally present

The low-risk effect is remarkably robust from a geographic perspective (present in all major developed and emerging markets), from an industry perspective (present within and across industries), and from a time perspective (consistent over time). Blitz and van Vliet [2007] show that the anomaly exists in the U.S., European and Japanese equity markets, and Blitz, Pang, and van Vliet [2013] document the low-risk anomaly for emerging equity markets. The low-risk effect is further confirmed in international samples by Baker and Haugen [2012], Frazzini and Pedersen [2014], and Walkshäusl [2014]. Recently, Han, Li, and Li [2018] and Chen, Pong, and Wang [2018] find that the low-risk anomaly is also present in the local Chinese (A shares) equity market, and Joshipura and Joshipura [2016] document the anomaly for the Indian stock market.

If the relation between risk and return is flat, then a long low-risk and short high-risk hedge portfolio will show an average return of zero and a strong negative CAPM beta. Frazzini and Pedersen [2014] construct a so-called Betting-Against-Beta (BAB) factor which is designed to turn this into a positive premium with zero beta, by dynamically levering the long low-risk portfolio up, to a beta of 1, and de-levering the short high-risk portfolio down, also to a beta of 1. Exhibit 5 shows that the BAB premium is positive in all 24 countries for which data is available online.[3] The BAB study helped to popularize the low-risk effect, especially among academics. Recently, however, the methodological assumptions behind BAB have been criticized by Novy-Marx and Velikov [2018], who argue that a large part of the premium is driven by dynamic hedging and shorting highly illiquid micro caps.

INSERT EXHIBIT 5 HERE

[3] https://www.aqr.com/Insights/Datasets/Betting-Against-Beta-Equity-Factors-Monthly

The previously mentioned Paradox Investing website contains a US volatility factor portfolio which addresses this concern by staying close to the factor construction methodology of Fama and French [1993, 2015]. In order to construct this volatility factor (VOL), every month all stocks in the CRSP database are first classified as either large or small, using the NYSE median market capitalization as breakpoint, and next value-weighted low-, mid-, and high-volatility portfolios are created within both of these size groups using past 36-month volatility and the NYSE 30th and 70th percentiles as breakpoints. The long leg of the factor takes a fifty-fifty mix of the large-cap low-volatility and small-cap low-volatility portfolios, and the short leg a fifty-fifty mix of the large-cap high-volatility and small-cap high-volatility portfolios. The only deviation from the standard Fama-French factor construction methodology is that the portfolio is made beta neutral, by levering up the long leg and levering down the short leg to full-sample markets betas of 1 each. Without beta neutrality, the VOL factor would have a highly negative beta to the market factor.

Exhibit 6 shows the annualized premiums of the volatility factor and the factors in the Fama and French [2015] model by decade since 1940. The data for the profitability (RMW) and investment (CMA) factors over the 1940-1963 period is obtained from Wahal [2019].[4] The graph shows that the volatility factor is the only factor which has generated a positive premium in every decade. Unreported tests shows that the volatility factor has even been positive in every 120-month moving window period. It is also the only factor which has delivered a solid premium over the most recent decade.

INSERT EXHIBIT 6 HERE

Compared to the Fama-French factors, the volatility premium is not only stable through time, but also large in magnitude. Over the full sample period, the average premium is 5.8% per annum with a volatility of 9.0%. This translates into an annualized Sharpe ratio of 0.65 and an accompanying t-statistic of 5.7, well above all common thresholds for statistical significance. To put the statistical strength of the VOL premium into perspective Exhibit 7 shows the t-statistics of the other factor premiums. On a stand-alone basis the volatility premium is the strongest factor, while the size premium is the weakest factor.

INSERT EXHIBIT 7 HERE

Many anomalies are known to be concentrated in small-cap stocks and therefore difficult to exploit in reality, but Auer and Schuhmacher [2015] show that the low-risk anomaly is strongly present among the largest, most liquid U.S. stocks. Asness, Frazzini, and Pedersen [2014] and Baker, Bradley and Taliaferro [2014] show that the low-risk effect exists within industries and countries, and also across industries and countries. Annaert and Mensah [2014] document a clear low-risk effect for the Brussels stock exchange in the decades before World War I, when this was one of the biggest stock markets.

Further evidence for the low-risk effect is coming from studies on other asset classes. Carvalho, Dugnolle, Lu, and Moulin [2014] and Israel, Kang, and Richardson [2015] document a low-risk effect within the investment grade corporate bond market, and Houweling and Van Zundert [2017] show that the effect is not only present among investment grade corporate bonds but also among high yield corporate bonds. Falkenstein [2009] documents over twenty occurrences of the low-risk effect, including some more

[4] http://jfe.rochester.edu/data.htm

‘exotic’ ones such as the options market, movie production (De Vany and Walls [2002]), and sports books (Snowberg and Wolfers [2010]). Eraker and Ready [2015] find that very risky OTC stocks have very poor average returns, Moskowitz and Vissing-Jorgensen [2002] find a low-risk anomaly for private business returns, Adhami, Gianfrate, and Johan [2019] observe an inverse relation between risk and return in the crowdlending market, and Jordan and Riley [2015] find that the low-risk anomaly is also present in the cross-section of mutual fund returns. Altogether, there appears to be a low-risk effect within every asset class. The relation between risk and return only seems to be positive across entire asset classes, since stocks have higher returns than bonds, and corporate bond returns are higher than government bond returns, in the long run.

Low volatility or low beta?

Is the low-risk anomaly primarily a low-volatility or a low-beta anomaly? The first thing to note in this regard is that volatility and beta are closely related metrics, since the beta of a stock to the market index is equal to its volatility times its correlation with the market index, divided by the volatility of the market. In cross-sectional comparisons the latter term is a constant and hence irrelevant. As such, defining low-risk based on volatility or beta is effectively a choice on the added value of correlations. Blitz and van Vliet [2007] and Baker, Bradley, and Wurgler [2011] find slightly stronger results for volatility than for beta, although these differences are relatively small. More recently, Liu, Stambaugh, and Yuan [2018] also find weaker results for beta.

Asness, Frazzini, Gormsen, and Pedersen [2019] directly disentangle the two driving components of beta, volatility and correlation. They find that there is a clear alpha when stocks are sorted on volatility, and that next to that there is an alpha when stocks are sorted on correlation within volatility buckets. In other words, correlation matters, but only among stocks that have a similar level of volatility. These results suggest that volatility is the main driver of the low-risk effect, and that the added value of correlations is a second-order effect.

All the studies mentioned so far typically estimate risk using metrics such as beta or volatility, estimated on 1, 3 or 5 years of historical data. A closely related phenomenon is the idiosyncratic volatility (iVol) effect of Ang, Hodrick, Xing, and Zhang [2006, 2009], which they estimate using daily return data over the past 1 month. This 1-month lookback period for risk estimation, in combination with a standard 1-month holding period, results in one of the most powerful manifestations of the low-risk anomaly on paper, but one that is less suitable for practical applications because of the amount of turnover (and hence transaction costs) involved. Another practical concern is that Ang, Hodrick, Xing, and Zhang [2006] find that most of the alpha of their strategy comes from the short side, i.e. a very poor performance of the recently most risky stocks. In reality this alpha may be beyond the reach of many investors because of limits to arbitrage; see also the next section.

The previously mentioned Kenneth French data library also contains data for portfolios sorted 60-day variance and 60-day residual variance. The reason for their use of an approximately 3-month lookback period instead of the 1-month lookback period of Ang, Hodrick, Xing, and Zhang [2006] is unclear, but the results appear generally similar. Exhibit 8 shows the full-sample performance of decile portfolios sorted on the short-term risk-measures. Once again we observe that the risk-return relation goes from flat to inverted at the far end. The very low returns for the most risky stocks are particularly striking. Sorting on 60-day variance or 60-day residual variance seems to make little difference. All these results are consistent with Ang, Hodrick, Xing, and Zhang [2006], who also report dismal performance for the most risky stocks and very similar results for sorting on past 1-month total return volatility and past 1-month idiosyncratic volatility.

INSERT EXHIBIT 8 HERE

Yet another manifestation of the low-risk effect, one could argue, is the finding of Campbell, Hilscher, and Szilagi [2008] find that financially distressed stocks have delivered anomalously low returns.

Explanations

We continue with discussing the most popular explanations for the low-risk effect. For the purposes of this paper we group these explanations in five main categories: (i) constraints, (ii) relative performance objectives, (iii) agency issues, (iv) skewness preference, and (v) behavioral biases. For a more extensive overview of explanations we refer to Blitz, Falkenstein, and van Vliet [2014].

Constraints. The low-risk effect has been linked to the limits to arbitrage that arise from various practical constraints, in particular shorting and leverage constraints. Heterogeneous beliefs cause the investors with the most optimistic expectations to drive up the price of risky assets. Miller [1977] argues that in the absence of enough short-sellers this *winners' curse* flattens the risk-return relation. For many investors the possibility to sell short or use leverage is restricted by mandate or means. Theoretically, Brennan [1971] and Black [1972] showed that the security market line may become flatter than predicted by the CAPM in the presence of leverage constraints. Black [1993] argues that, if anything, leverage constraints have tightened over time. The intuition behind this explanation is as follows. In the world of the CAPM there is only one efficient portfolio, and investors simply lever or de-lever this portfolio based on their degree of risk aversion. In the presence of leverage constraints, however, investors looking to increase their return are forced to tilt their portfolios towards high-beta securities. This extra demand for high-beta securities and reduced demand for low-beta securities may explain a flattening of the security market line. In support of this explanation, Frazzini and Pedersen [2014] find that when leverage constraints are tighter, the low-risk anomaly tends to be stronger.

Relative performance objectives. A second explanation for the low-risk effect is a focus on performance relative to others instead of absolute performance. The CAPM assumes that investors only care about absolute returns, but in reality many investors focus on beating the market average. Blitz and van Vliet [2007], Falkenstein [2009], and Baker, Bradley, and Wurgler [2011] argue that if the CAPM would hold, then low-risk stocks are unattractive for benchmark-relative investors, because they involve high tracking error and lower expected return. Brennan [1993] and Brennan, Cheng, and Li [2012] assume the simultaneous presence of absolute and relative return-oriented investors and show that this implies a partial flattening of the security market line, with the degree of flattening depending on the number of relative-return investors, versus the number of absolute-return investors. Falkenstein [2009] shows that if investors only care about performance relative to others, then in equilibrium the relation between risk and return is flat. Within this relative utility framework, there are no risk premiums left at all, which is at odds with the existence of a global equity premium. Blitz [2014] reconciles the simultaneous existence of a positive equity risk premium and absence of a risk premium within asset classes by considering a two-stage investment process, where investors first make asset allocation decisions based on absolute performance criteria, but next switch to a relative performance objective when trying to identify the best managers or securities within each asset class. This explanation assumes a mental accounting bias, as in the two-layer portfolio model of Shefrin and Statman [2000], where the low-aspiration layer is designed to avoid poverty, while the high-aspiration layer aims for a shot at riches.

Agency issues. The low-risk effect has also been attributed to agency issues. Karceski [2002] argues that profit-maximizing asset managers have a strong incentive to create high-beta products due to the highly asymmetric nature of the flow-performance relationship documented by Sirri and Tufano [1998]: most flows are attracted by funds with the best performance in asset classes that also had a good performance. Falkenstein [1996] shows that mutual fund managers, who have an incentive to attract investor flows, have a preference for stocks with higher idiosyncratic volatility, and Agarwal, Jiang and Wen [2018] find that smaller, younger mutual funds with poor recent performance own more risky 'lottery-like' stocks, arguably in order to attract capital. Such tournament behavior may also be present among analyst, as Hsu, Kudoh and Yamada [2013] provide evidence that sell-side analysts prefer high-volatility stocks. Baker and Haugen [2012] generalize this agency problem and argue that all portfolio managers and their analysts implicitly or explicitly have option-like reward structures, which incentivizes them to focus on high-risk assets. A lab-in-the field experiment by Kirchler, Lindner and Weitzel [2018] with a large group of professional investors supports the notion that ranking and tournament incentives are important drivers of risk-taking.

Skewness preference. Another explanation for the low-risk effect is a preference for lottery-like payoffs or positive skewness; see Blitz and van Vliet [2007], Baker, Bradley and Wurgler [2011], Ilmanen [2012], and Hsu and Chen [2017]. Kumar [2009] shows that many investors participate in the stock market in order to gamble. Bali, Cakici and Whitelaw [2011] test the 'stocks as lotteries' hypothesis of Barberis and Huang [2008] and use this to explain the very low returns of the most risky stocks. High-risk stocks are attractive to lottery-seeking investors, because they offer limited downside risk combined with unlimited upside potential. A preference for skewness can even support the existence of an inverse risk-return relationship, i.e. that instead of requiring a compensation for taking on risk, investors may actually be willing to pay a premium for it.

Behavioral biases. A fifth explanation for the low-risk effect is behavioral biases, such as attention-grabbing bias, representativeness bias, and overconfidence, which cause investors to irrationally 'prefer' higher risk stocks over lower risk stocks; see e.g. Blitz and van Vliet [2007], and Baker, Bradley and Wurgler [2011]. For example, investors pay more attention to stocks that are very visible than to stocks which remain more under the radar. High-risk stocks are more likely to experience extreme returns that grab the attention of investors, which creates excessive buying pressure on these stocks; see Barber and Odean [2008]. With representativeness bias or overconfidence investors may become too optimistic about the future prospects of volatile stocks that might be the next Amazon or Google, causing these stocks to become overpriced and generate lower subsequent returns. Behavioral explanations can also support the existence of an inverse risk-return relationship.

The low-risk effect is not a data fluke

Harvey [2017] argues that a serious concern in finance, and science in general, is 'p-hacking'. Scientists are subject to statistical testing limitations, have several degrees of freedom (on areas like data manipulation, statistical method, and results chosen to present), publications are biased to positive results, and scientists have an incentive to publish. As a consequence, the possibility exists that evidence for the low-risk effect might actually be a false positive. As a case in a point, Harvey, Liu, and Zhu [2016] find a clear publication bias pattern in the top finance journals, and that many out of 300+ documented stock-level anomalies become questionable after analyzing these in a rigorous testing framework that allows for multiple hypotheses testing bias.

For the low-risk effect we deem the p-hacking explanation unlikely, for the following reasons. First, the effect did not originate from research that tested dozens or hundreds of different potential alpha factors, but from research that aimed to confirm the basic predictions of the CAPM. Moreover, the anomaly was ignored for many decades, until it simply could not be ignored anymore. Second, the low-risk effect is not about a slight failure of the CAPM, but about the total absence of a positive relation between risk and return. Many studies even find the relation to be clearly inverted. And, as discussed above, there is also a clear economic rationale for these findings. Third, the effect is shown to be very robust over different samples, initially observed in the U.S. over half a century ago, and since then 'out-of-sample' in many studies covering different time periods and/or markets. These include more recent decades and all main regions, including emerging markets, as outlined above. Fourth, as shown in Exhibit 6, the effect is remarkably robust over subperiods, even more so than widely accepted factors such as size and value. Fifth, the effect is also present in other asset classes where the same economic mechanisms are present, like corporate bonds.

IS LOW-RISK A DISTINCT EFFECT?

In this section we discuss the most important challenges to the low-risk effect, in particular whether it might be a manifestation of interest rate exposure, or that it can be explained by other factors such as value or profitability. We review the studies which have specifically addressed these concerns, and conclude that, at best, they are able to explain only a small part of the low-risk effect, or the performance over a very specific sub-period. None of them, however, is able to provide a comprehensive explanation for the low-risk anomaly.

Low-risk is not explained by interest rate risk

Some have suggested that the low-risk effect may be explained by an implicit exposure to interest rate risk. Standard asset pricing models only use equity-based factors, but Baker and Wurgler [2012] show that low-risk stocks have pronounced bond-like features. De Franco, Monnier, and Rulik [2017] formally examine this issue, and conclude that although low-risk stocks have a statistically significant exposure to interest rate risk, this only explains a very small part of their alpha. This is not really surprising, because the bond premium is relatively small, so the loading of low-risk stocks on bonds would have to be extremely large in order for the alpha to be fully explained. Another reason for why it is implausible that the alpha of low-risk stocks is driven by interest rate exposure is because the low-risk anomaly is not merely a phenomenon of recent decades, during which interest rates were falling, but was also present in earlier decades, during which interest rates were stable or even rising. Moreover, the interest rate risk explanation does not carry over to the low-risk anomaly in corporate bonds, as low-risk corporate bonds have short maturities and hence lower instead of higher exposure to interest rate changes. Van Vliet [2011] and Coqueret, Martellini, and Milhau [2017] argue that from an ALM perspective the interest rate exposure of low-risk stocks is actually a desirable feature, as it implies better liability-hedging properties compared to a standard equity portfolio.

Low-risk is distinct from the value effect

In the mid-2000s many investors wondered whether the low-risk effect was simply a manifestation of the well-known value effect. Like generals preparing for the next war by training to fight the previous war again, investors still had the technology bubble fresh in their minds. The technology stocks that soared during the late nineties and crashed in the early-2000s were unattractive from both a value and a low-risk perspective, while 'old economy' stocks were both cheaper and less risky. Blitz and Van Vliet [2007], Frazzini and Pederson [2014], and Walkshäusl [2014] all find that the alpha of low-risk stocks is not explained when controlling for value (and other factors), for the universes and sample periods in their studies.

That said, the alpha of a plain-vanilla US low-risk strategy seems to be explained by an implicit loading on the classic HML value factor in time-series regressions over the commonly used post-1963 period. Blitz [2016] examines this issue in detail and finds that the value effect fails to explain the performance of large-cap low-risk strategies in the US pre-1963 as well as post-1984, when the Fama–French value factor itself ceased to be effective in the large-cap segment of the market. Moreover, the value effect cannot explain the performance of small-cap low-risk strategies during any period. In other words, the value factor is only able to explain the performance of US large-cap low-volatility strategies during a period roughly in the middle of the CRSP sample, which makes up less than a quarter of the entire CRSP sample. These results reject the notion that low-risk might simply be value in disguise.

After the global financial crisis of the late 2000s the criticism that low-risk may be value waned, as low-risk stocks began trading at higher multiples than the market. In other words, investors were willing to pay up for defensiveness. This is not unprecedented, as the same happened in the wake of the great depression in the nineteen thirties. Nevertheless, it fueled a new criticism by Arnott, Beck, Kalesnik, and West [2016], who argue that low-risk (and other smart beta) indices have been designed based on relatively short-back test periods, and obtain a large part of their return from multiple expansion of their holdings that is not sustainable in the long run. Although they document that multiple expansion can indeed be a significant explanatory factor for performance in the short run, it cannot explain the long-run success of low-risk strategies. Nevertheless, investors should be aware that low-risk stocks can go through long cycles of being either value- or growth-tilted.

Low-risk is distinct from the profitability effect

A more recent challenge to the low-risk anomaly is that it is explained by profitability factors. This argument was first made by Novy-Marx [2014] using the gross profitability factor of Novy-Marx [2013], and subsequently also by Fama and French [2016] using the slightly different profitability factor of the Fama and French [2015] five-factor model. Despite the unequivocal conclusion in Fama and French [1992] that market beta is not a priced factor in the cross-section of stock returns, Fama and French [1993] chose to retain the fundamental CAPM relation between market beta and return as the basis for their three-factor model. The Fama and French [2015] five-factor model adds profitability and investment factors to the three-factor model, again without changing the CAPM basis of the model. This time, however, they specifically address whether this is justified in light of the low-risk anomaly. In Fama and French [2016] they argue that it is, because with the new factors included the five-factor model is able to explain the performance of risk-sorted portfolios in time-series regressions.

Blitz and Vidojevic [2017] acknowledge that the low-risk effect appears to be subsumed by the profitability factor in time-series regressions. They go on to argue, however, that if it were true that the CAPM relation holds when accounting for interactions with the profitability factor, then it should be possible to construct portfolios which exhibit a clear positive relation between market beta and return, provided one controls for the profitability characteristics of these portfolios. They argue that this can be tested by running Fama and MacBeth [1973] cross-sectional regressions, as the returns estimated with these regressions can be interpreted as the reward to a unit exposure to a factor, controlling for the exposures to all other factors included in the analysis. Applying this approach they find that all factors in the five-factor model are priced, except market beta. In other words, it is not possible to construct high-beta portfolios with a high return and low-beta portfolios with a low return, whether one controls for profitability or not. Based on this finding they conclude that it is premature to assume that the low-risk effect is explained by profitability or other factors.

Additional evidence is provided by Blitz, Baltussen, and van Vliet [2019], who examine how the long legs and short legs of classic factor strategies separately contribute to performance. They find that the long legs of factors tend to have a stronger risk-adjusted performance than the short legs. Moreover, while the short side of a low-risk strategy (i.e. high risk) can be explained by the short side of the new Fama-French factors (e.g. poor profitability), the long side is not, i.e. the alpha of low-risk stocks cannot explained by the alpha of high-profitability stocks. In other words, the results of Novy-Marx [2014] and Fama and French [2016] are entirely driven by the short sides of factors. In a long-only setting, which is the preferred approach of many investors in practice (see next section), low-risk clearly stands its ground as a distinct factor.

Blitz, Baltussen, and van Vliet [2019] also observe that even the long-short low-risk factor is not explained by the new Fama-French factors over the 1963-1990 subsample period, which is the period that was used by Fama and French [1992] to establish their classic size and value factors. Based on the data in Exhibit 6 the long-short low-risk factor is also not explained by the other Fama-French factors pre-1963 and post-2010. Altogether this implies that the conclusions of Novy-Marx [2014] and Fama and French [2016] are entirely driven by the twenty-year period surrounding the turn of the century.

CAPTURING THE LOW RISK EFFECT

In this section we make the step from theory to practice, discussing the key considerations that come into play when investing based on the low-risk effect with real money. We first discuss why investors who aim to profit from the low-risk effect typically use a long-only approach. We then discuss the optimal amount of turnover that is needed to capture the low-risk anomaly, the role of correlations, how to deal with currency risk, the pros and cons of passively following a generic low-risk index, and whether or not to combine low-risk with other factors.

Low-risk strategies tend to be long-only in practice

Theoretical studies typically consider long-short portfolios. For instance, the Fama-French factors for size, value, momentum, profitability and investment all assume a long portfolio of stocks with good factor characteristics and a matching short portfolio of stocks with bad factor characteristics. In practice, however, factor investing is often implemented using a long-only approach, witness for instance the popularity of 'smart beta' ETFs which track (long-only) factor indices. Not only the index-based, but also the most popular active offerings that target the low-risk effect tend to follow a long-only approach.

In principle, the low-risk effect can also be targeted with a long-short approach, in which one not only goes long the lowest-risk stocks, but simultaneously short in the highest-risk stocks. Perhaps certain individual hedge funds follow such a long-short approach, but overall it seems to be far less popular than a more straightforward long-only approach. One reason for this may be management fees: long-only funds tend to charge considerably lower fees than long-short hedge funds. Another reason may be that shorting the most risky stocks can be difficult and expensive in practice. In fact, shorting costs tend to increase (substantially) with a stock's volatility, see Drechsler and Drechsler [2016], and also the probability of recall increases with a stock's volatility, see D'Avolio [2002]. Moreover, shorting high risk stocks is relatively risky, since the price of these stocks can suddenly increase a lot over short periods of time. Beta management, i.e. making sure that the long position in low-risk stocks and the short position in high-risk stocks has zero net beta, is also challenging because volatility levels are highly time-varying with occasional spikes. Given that long-only appears to be the preferred way to harvest the low-risk effect in practice, we will take this as the starting point for the remainder of this section.

Low risk requires low turnover

The literature on low-risk investing reports annual turnover levels ranging from less than 20% to over 100%. A wide dispersion can also be observed in the turnover levels of the low-risk strategies that are offered by asset managers and index providers. This gives rise to the question how much turnover is needed in order to efficiently harvest the low-risk premium. Li, Sullivan and Gracia-Feijoo [2014] argue that high turnover and implementation costs could be a serious restriction for investors looking to profit from the low-risk anomaly. On the other hand, Novy-Marx and Velikov [2015] find that with smart trading rules most anomalies remain profitable after costs, except for the ones which involve very high amounts of turnover. Frazzini, Israel and Moskowitz [2012] also examine implementation strategies designed to reduce turnover, and argue that actual trading costs are oftentimes lower than assumed in the literature. Van Vliet [2018] conducts a meta-analysis of the low-risk literature and considers various historical portfolio simulations to conclude that an efficient low-risk strategy does not require more than 30% turnover per annum.

One way to control turnover is to estimate risk using medium or long lookback periods. For instance, the 1-month iVol effect of Ang, Hodrick, Xing, and Zhang [2006] may look great in theoretical tests with 1-month holding periods, but is not particularly suitable for a real-life investment strategy that involves transaction costs. Another crucial aspect when it comes to controlling turnover is to not create an entirely fresh portfolio every month, but to take the existing portfolio into account and only replace those holdings that are most in need of replacement. Clarke, de Silva, and Thorley [2006] find that a naïve minimum-volatility approach with monthly rebalancing results in an annual turnover of 143%, but that performance is hardly affected if turnover is reduced to 56%. The MSCI Minimum Volatility indices impose constraints that explicitly limit annual turnover to just 20%, without significantly diluting the exposure to the low-risk effect. In sum, the low-risk anomaly does not require a high turnover, but a naïve implementation can result in excessive turnover, and hence unnecessary transaction costs.

Low risk or minimum risk: mainly semantics

Although low-risk and minimum-risk approaches are closely related, there are also subtle differences. The most important difference is that a low-risk approach will only select stocks which each have a low risk individually, while a minimum-risk approach may additionally select stocks that have a medium or even high risk on their own, but which help to reduce risk at the portfolio level due to their low correlation with other stocks. Stocks of gold-mining companies are a classic example of high-risk, but low-correlation stocks, given the role of gold as a catastrophe hedge. This raises the question whether the alpha of low-risk strategies is concentrated in stocks which have a low risk on their own, or if there is also alpha in risky stocks that have low correlations with other stocks. Soe [2012] empirically compares low-volatility and minimum-volatility strategies based on the constituents of the S&P 500 index, and finds that the two approaches exhibit very similar performance characteristics. Carvalho, Lu, and Moulin [2012] investigate five popular risk-based approaches and find that they generally load on low-beta and low-volatility factors. Scherer [2011] and Chow, Hsu, Kuo and Li [2014] come to the same conclusion. In other words, differences between low-risk and minimum-risk are mostly semantics.

Currency risk

A related question is how to deal with currencies, because the volatilities and correlations of stocks can change when returns are measured in a different currency. The key question here is whether a low-risk strategy should be optimized towards the base currency of an investor, or whether it should be base-currency agnostic. To illustrate, MSCI offers, amongst others, USD-optimized, EUR-optimized, JPY-optimized, GBP-optimized, and AUD-optimized versions of their commonly used global Minimum Volatility indices. In other words, their optimal minimum-volatility portfolio is different for investors with different base-currency perspectives. Alonso and Barnes [2017] empirically investigate this issue, and find that optimizing towards a base-currency results in a very large bias towards the home market of the investor. The intuition behind this result is that investing in foreign markets is more risky than investing in the home market, because next to share price volatility in the local currency it also exposes the investor to exchange rate volatility. This currency risk is reduced by creating a home-market bias. However, this is probably not the most efficient way to deal with currency risk. Investors can also create a base-currency agnostic low-risk portfolio which only considers the risk of stocks measured in their local currency, and use derivatives to directly hedge out any undesired foreign exchange rate risk.

Low-risk indices: beware of the pitfalls

Low-risk investing is essentially active investing, because one intentionally deviates from the capitalization-weighted market portfolio. Since the riskiness of stocks can change over time, low-risk investing is also not a buy-and-hold strategy but requires (at least some) turnover. Furthermore, a wide range of active choices is required, such as the choice of risk metric, choice of lookback period, choice of weighting rules, choice of rebalancing frequency, concentration limits on e.g. countries or sectors, etc.; see also Alighanbari, Doole, and Shankar [2016]. For these reasons there is no such a thing as a passive approach towards low-risk investing. Investors can choose to passively replicate a 'smart beta' low-risk index, but this does not change the fact that the index itself is not passive in these cases. The main benefit of index-based low-risk investing is that it may be cheaper than an outright active approach, since index replication typically involves lower management fees than active management. Investors also appreciate the transparency of low-risk indices, as the index methodology is fully specified and public.

However, there is also a flip side to these advantages. Although low-risk indices can be suitable for tracking the hypothetical performance of a generic low-risk investment approach, they are less suitable for large-scale replication with real money. In particular, low-risk indices tend to use simplistic rebalancing rules that are fine for back-testing purposes, i.e. on paper, but which are clearly suboptimal when it comes to actual money management. The popular MSCI Minimum Volatility indices, for instance, conduct all their trades on just two days of the year, the last business day of May and the last business day of November. Blitz and Marchesini [2019] show that this can easily result in trades that are ten or even hundred times the average daily volume of the stocks in question, which implies high expected market impact costs, and also the risk of ending up in overcrowded positions. They argue that the larger the amount of money that is invested in a low-risk strategy, the more important it is to apply a trading strategy which makes efficient use of the liquidity that the market has to offer throughout the year. This means continuously trading in small amounts, rather than infrequently trading big chunks.

The full transparency of low-risk indices is also a double-edged sword, because it makes the investors in these strategies vulnerable to index arbitrage, which is a common strategy among hedge funds. The time gap between the announcement of the new index composition and the effectuation of these changes also offers opportunities for passive managers to game the index. By already buying stocks that enter the index between the announcement and effective dates, the more one pushes up the price of these stocks, the worse the price at which the index will 'buy' the stock at the effective date. Huij and Kyosev [2016] empirically investigate the price effects of stocks entering and exiting during MSCI Minimum Volatility index rebalances, and estimate a hidden cost to investors tracking these indices amounting to about 16 basis points per annum.

Low-risk can be combined efficiently with other factors

There is a lot of evidence for the low-risk anomaly, but also for other factor premiums such as size, value, momentum, profitability, and investment. When setting up a low-risk investment strategy, investors need to decide whether to ignore or also do something with these other factors. Haugen and Baker [1996] develop a multi-factor model which generates higher return than the market, with lower risk. Furthermore, Garcia-Feijoo, Kochard, Sullivan, and Wang [2015] find that the performance of a low-risk strategy is higher when valuation levels are lower.

A concern may be that incorporating other factors into a low-risk strategy leads to a dilution of the amount of low-risk exposure and an increase in trading costs. Van Vliet [2018] addresses these concerns and finds that with just a little bit of additional trading the exposure to other factors in a low-risk strategy can be increased sharply, especially when using negatively correlated factors such as value and momentum. Blitz and van Vliet [2018] show that a multi-factor strategy which uses three simple price-based metrics (36-month volatility, net payout yield, and 12-minus-1-month price momentum) is able to give investors simultaneous exposure to all major factor premiums. This "Conservative Formula" not only beats the market index by a wide margin, but also a generic low-volatility approach and a wide range of other single-factor indices. Thus, the low-risk anomaly lends itself well to combining with other factors.

Alighanbari, Doole, and Melas [2017] show that Environmental, Social and Governance (ESG) factors can also be efficiently integrated into a low-risk strategy. For instance, they report that improving the portfolio level ESG score by 30% comes at the cost of only a 0.5% increase in volatility.

THE LOW-RISK EFFECT IS NOT ARBITRAGED AWAY

The first dedicated low-risk products were introduced in the mid-2000s, but only after the global financial crisis of the late-2000s investors began to show serious interest in these products. Within the space of just a decade, low-risk investing has developed into a widely accepted investment style. The approach is so well-known nowadays that many investors are concerned that it might become a victim of its own success. In this section we look for tangible evidence that the low-risk effect is already in the process of being arbitraged away, in particular by investors in mutual funds, exchange-traded funds, and hedge funds, and find little justification for such concerns.

Mutual Funds – typically on the other side

Falkenstein [1996] and Sias [1996] find that mutual funds have a negative exposure towards low-risk stocks. Beveratos et al [2017] also find that mutual funds are tilted towards smaller stocks with higher volatility. Ang, Madhavan, and Sobcyzk [2017] examine the aggregate assets of all US active mutual funds since the late 1990s and find that that mutual funds have a consistently negative exposure towards low-volatility stocks. Although low-volatility investing became more popular during this period, mutual fund managers did not change their investment behavior and continued to underweight low-risk stocks. Christoffersen and Simutin [2017] argue that in an effort to beat benchmarks, fund managers tend to increase their exposure to high-risk stocks, while aiming to maintain tracking errors around the benchmark. These results are consistent with the benchmark-relative objectives explanation for the low-risk anomaly discussed before. Altogether, it seems that mutual funds are more drawn towards high-risk stocks than towards low-risk stocks.

ETFs – on aggregate neutral on low-risk

One commonly heard concern with low-risk investing is that the sizable assets in dedicated low-risk ETFs nowadays may imply that the anomaly is rapidly being arbitraged away. Ang, Madhavan, and Sobcyzk [2017] report that 24 billion USD is invested in low-volatility ETFs at the end 2016. Blitz [2017] directly addresses this concern by examining the factor exposures of all US-listed ETFs investing in US equities. The study finds that there are indeed quite a few ETFs which offer a large exposure to the low-risk factor. This includes dedicated low-risk ETFs, but also some high-dividend ETFs, and some sector ETFs on typical low-risk sectors such as utilities. Many other ETFs, however, have a clear high-risk profile. This may not be evident to investors, because, for obvious reasons, these ETFs are typically not marketed as "high-risk" funds. On aggregate, so when all ETFs are lumped together into one big total portfolio, the sizable low-risk and sizable high-risk exposures, that are so strongly present at the individual fund level, almost perfectly cancel out. In other words, ETF investors have zero net exposure to low- and high-risk stocks. The conclusion of this analysis is that although some ETF investors are targeting the low-risk anomaly, either explicitly or implicitly, there are just as many other ETF investors who do precisely the opposite, and these two groups keep each other nicely in check. Based on these findings there is little cause for concern that ETF investors are arbitraging the low-risk anomaly away.

Hedge funds – typically on the other side

One can also wonder if hedge funds are perhaps exploiting the low-risk effect, especially if leverage and benchmark constraints are at the root of the anomaly. These constraints are much less of an issue for hedge funds, because the flexibility to use leverage is one of the defining features of hedge funds, and because hedge funds typically have an absolute return objective. In other words, the limits to arbitrage that hamper many investors do not apply to hedge funds. Blitz [2018] investigates whether hedge funds are indeed exploiting the low-risk anomaly, but finds strong evidence for the exact opposite: hedge funds have higher returns when high-risk stocks outperform low-risk stocks, instead of the other way around. These results suggest that instead of arbitraging the low-risk anomaly away, hedge funds may actually be helping to create and sustain it. Put differently, hedge funds seem to be on the other side of the low-risk trade. Asness, Ilmanen, Israel, and Moskowitz [2015] also find that hedge funds are positioned against the low-risk anomaly. Perhaps hedge funds have a preference for high-risk stocks because their option-like incentive structures reward risk-seeking behavior, as suggested in Baker and Haugen [2012].

Institutional investors – discouraged by regulation

Other important market participants are institutional investors such as insurance companies, pension funds and banks. Potentially, low-risk stocks are very interesting for institutional investors, because of their downside protection and liability-hedging properties. However, regulatory frameworks, such as Solvency II for insurance companies (and similar local regulation for pension funds) and Basel III for banks, do not incentivize low-risk investing, by requiring the same capital charges for low-risk stocks as for high-risk stocks. Blitz, Hallerbach, Swinkels, and van Vliet [2018] examine this issue in detail and argue that current regulation is sustaining the low-risk effect rather than encouraging investors to benefit from it.

CONCLUSION

In this paper we have reviewed the key insights into the low-risk (or low-volatility) effect. A low-risk approach has been effective for as far as the data goes back, across all major stock markets, from developed to emerging, within and across industries, across various market regimes and using different measures of risk. The effect is also present in other asset classes, such as corporate bonds. The low-risk effect is commonly attributed to the existence of leverage constraints, a focus on benchmark-relative instead of absolute performance, agency issues, skewness preference, and various behavioral biases. As with any other anomaly, only time will tell if the low-risk effect will persist going forward. The concern that the anomaly is already being rapidly arbitraged away by for example mutual funds, ETFs or hedge funds does not appear to be justified by the empirical evidence, which more often finds such investors to be on the other side of the low-risk trade.

REFERENCES

Adhami, S., Gianfrate, G., and Johan, S. (2019). Risks and Returns in Crowdlending. SSRN working paper no. 3345874.

Agarwal, V., Jiang L., and Wen, Q. (2018). Why do mutual funds hold lottery stocks?, Georgetown McDonough School of Business research paper no. 3164692.

Alighanbari, M., Doole, S., and Melas, D. (2017). Managing Risks Beyond Volatility. *Journal of Index Investing*, 8(2), 68-76.

Alighanbari, M., Doole, S., and Shankar, D. (2016). Designing Low-Volatility Strategies. *Journal of Index Investing*, 7(3), 21-33.

Alonso, N., and Barnes, M. (2017). What Is Missing in Common Minimum Volatility Strategies? The Ignored Impact of Currency Risk. *Journal of Index Investing*, 8(2), 77-88.

Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2006). The cross-section of volatility and expected returns. *Journal of Finance*, 61(1), 259-299.

Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2009). High idiosyncratic volatility and low returns: International and further US evidence. *Journal of Financial Economics*, 91(1), 1-23.

Ang, A., Madhavan, A., and Sobczyk, A. (2017). Crowding, capacity, and valuation of minimum volatility strategies. *Journal of Index Investing*, 7(4), 41-50.

Annaert, J., and Mensah, L. (2014). Cross-sectional predictability of stock returns, evidence from the 19th century Brussels Stock Exchange (1873–1914). *Explorations in Economic History*, 52, 22-43.

Arnott, R. D., Beck, N., Kalesnik, V., and West, J. (2016). How Can 'Smart Beta' Go Horribly Wrong? SSRN working paper no. 3040949.

Asness, C.S., Frazzini, A., Gormsen, N. J., and Pedersen, L. H. (2019). Betting Against Correlation. *Journal of Financial Economics*, forthcoming.

Asness, C. S., Frazzini, A., and Pedersen, L. H. (2014). Low-risk investing without industry bets. *Financial Analysts Journal*, 70(4), 24-41.

Asness, C. S., Ilmanen, A., Israel, R., and Moskowitz, T. J. (2015). Investing with style. *Journal of Investment Management*, 13(1), 27-63.

Auer, B. R., and Schuhmacher, F. (2015). Liquid betting against beta in Dow Jones Industrial Average stocks. *Financial Analysts Journal*, 71(6), 30-43.

Baker, M., Bradley, B., and Taliaferro, R. (2014). The low-risk anomaly: A decomposition into micro and macro effects. *Financial Analysts Journal*, 70(2), 43-58.

Baker, M., Bradley, B., and Wurgler, J. (2011). Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal*, 67(1), 40-54.

Baker, N. L., and Haugen, R. A. (2012). Low risk stocks outperform within all observable markets of the world. SSRN working paper no. 2055431.

Baker, M., and Wurgler, J. (2012). Comovement and predictability relationships between bonds and the cross-section of stocks. *The Review of Asset Pricing Studies*, 2(1), 57-87.

Bali, T. G., Cakici, N., and Whitelaw, R. F. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2), 427-446.

Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1), 3-18.

Barber, B. M., and Odean, T. (2008). All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors, *Review of Financial Studies*, 21(2), 785-818.

Barberis, N., and Huang, M. (2008). Stocks as lotteries: the implications of probability weighting for security prices, *American Economic Review*, 98(5), 2066-2100.

Beveratos, A., Bouchaud, J. P., Ciliberti, S., Laloux, L., Lempérière, Y., Potters, M., and Simon, G. (2017). Deconstructing the low-vol anomaly. *Journal of Portfolio Management*, 44(1), 91.

Black, F. (1972). Capital market equilibrium with restricted borrowing. *Journal of Business*, 45(3), 444-455.

Black, F. (1993). Beta and return: Announcements of the 'Death of Beta' seem premature. *Journal of Portfolio Management*, 20(1), 11-18.

Black, F., Jensen, M. C., and Scholes, M. (1972). The capital asset pricing model: Some empirical tests. *Studies in the Theory of Capital Markets*, 81(3), 79-121.

Blitz, D. (2014). Agency-Based Asset Pricing and the Beta Anomaly. *European Financial Management*, 20(4), 770-801.

Blitz, D. (2016). The Value of Low Volatility. *Journal of Portfolio Management*, 42(3), 94-100.

Blitz, D. (2017). Are Exchange-Traded Funds Harvesting Factor Premiums?. *Journal of Investment Consulting*, 18(1), 24-30.

Blitz, D. (2018). Are Hedge Funds on the Other Side of the Low-Volatility Trade?. *Journal of Alternative Investments*, 21(1), 17-26.

Blitz, D., Baltussen, G., and van Vliet, P. (2019). The Long and Short of Factor Investing. Working paper.

Blitz, D., Falkenstein, E., and Van Vliet, P. (2014). Explanations for the Volatility Effect: An Overview Based on the CAPM Assumptions. *Journal of Portfolio Management*, 40(3), 61-76.

Blitz, D., and Marchesini, T. (2019). The Capacity of Factor Strategies, forthcoming *Journal of Portfolio Management*

Blitz, D., Pang, J., and Van Vliet, P. (2013). The volatility effect in emerging markets. *Emerging Markets Review*, 16, 31-45.

Blitz, D., and Vidojevic, M. (2017). The profitability of low-volatility. *Journal of Empirical Finance*, 43, 33-42.

Blitz, D., and van Vliet, P. (2007). The Volatility Effect. *Journal of Portfolio Management*, 34(1), 102-113.

Blitz, D., and van Vliet, P. (2018). The Conservative Formula: Quantitative Investing Made Easy. *Journal of Portfolio Management*, 44(7), 24-38.

Brennan, M. J. (1971). Capital market equilibrium with divergent borrowing and lending rates. *Journal of Financial and Quantitative Analysis*, 6(5), 1197-1205.

Brennan, M. J. (1993). Agency and asset pricing, Unpublished manuscript, UCLA and London Business School.

Brenan, M. J., Cheng, X., and Li, F. (2012). Agency and institutional investment. *European Financial Management*, 18(1), 1-27.

Campbell, J. Y., Hilscher, J., and Szilagyi, J. (2008). In search of distress risk. *Journal of Finance*, 63(6), 2899-2939.

Carvalho, R. L., Dugnolle, P., Lu, X., and Moulin, P. (2014). Low-risk anomalies in global fixed income: Evidence from major broad markets. *Journal of Fixed Income*, 23(4), 51-70.

Carvalho, R, Lu, X., and Moulin, P. (2012). Demystifying equity risk-based strategies: A simple alpha plus beta description. *Journal of Portfolio Management*, 38(3), 56-70.

Chen, A., Pong, E., and Wang, Y. (2018). Accessing the China A-Shares Market via Minimum-Variance Investing. *Journal of Portfolio Management*, 45(1), 106-117.

Chow, T. M., Hsu, J. C., Kuo, L. L., and Li, F. (2014). A study of low-volatility portfolio construction methods. *Journal of Portfolio Management*, 40(4), 89-105.

Christoffersen, S. E., and Simutin, M. (2017). On the demand for high-beta stocks: Evidence from mutual funds. *The Review of Financial Studies*, 30(8), 2596-2620.

Clarke, R., De Silva, H., and Thorley, S. (2006). Minimum-variance portfolios in the US equity market. *Journal of Portfolio Management*, 33(1), 10-24.

Clarke, R., De Silva, H., and Thorley, S. (2011). Minimum-variance portfolio composition. *Journal of Portfolio Management*, 37(2), 31-45.

Cochrane, J. H. (1999). New Facts in Finance. NBER working paper no. 7169.

Coqueret, G., Martellini, L., and Milhau, V. (2017). Equity Portfolios with Improved Liability-Hedging Benefits. *Journal of Portfolio Management*, 43(2), 37-49.

D'Avolio, G. (2002). The market for borrowing stock. *Journal of Financial Economics*, 66(2), 271-306.

De Franco, C., Monnier, B., and Rulik, K. (2017). Interest rate exposure of volatility portfolios. *Journal of Index Investing*, 8(2), 53-67.

De Vany, A. and Walls. W. D. (2002). Does Hollywood make too many R-rated movies? Risk, stochastic dominance, and the illusion of expectation. *Journal of Business*, 75(3), 425-452.

Drechsler, I. and Drechsler, Q. F. S. (2016). The Shorting Premium and Asset Pricing Anomalies, SSRN working paper no. 2387099.

Eraker, B. and Ready, M. (2015). Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks. *Journal of Financial Economics*, 115(3), 486-504.

Falkenstein, E. G. (1994). Mutual Funds. *Idiosyncratic Variance, and Asset Returns*. PhD Thesis, Northwestern University.

Falkenstein, E. G. (1996). Preferences for stock characteristics as revealed by mutual fund portfolio holdings. *Journal of Finance*, 51(1), 111-135.

Falkenstein, E. G. (2009). Risk and return in general: Theory and evidence. SSRN working paper no. 1420356.

Fama, E. F., and French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2), 427-465.

Fama, E. F., and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3-56.

Fama, E. F., and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1-22.

Fama, E. F., and French, K. R. (2016). Dissecting anomalies with a five-factor model. *The Review of Financial Studies*, 29(1), 69-103.

Fama, E. F., and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3), 607-636.

Frazzini, A., Israel, R., and Moskowitz, T. J. (2012). Trading costs of asset pricing anomalies. Fama-Miller working paper, no. 14-05.

Frazzini, A., and Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1), 1-25.

Garcia-Feijóo, L., Kochard, L., Sullivan, R. N., and Wang, P. (2015). Low-Volatility Cycles: The Influence of Valuation and Momentum on Low-Volatility Portfolios. *Financial Analysts Journal*, 71(3), 47-60.

Han, X., Li, K., and Li, Y. (2018). Investor Overconfidence and the Security Market Line: New Evidence from China. SSRN working paper no. 3284886.

Harvey, C. R. (2017). Presidential Address: The Scientific Outlook in Financial Economics, *Journal of Finance*, 72(4), 1399-1440.

Harvey, C. R., Liu, Y., and Zhu, H. (2016). … and the cross-section of expected returns. *The Review of Financial Studies*, 29(1), 5-68.

Haugen, R. A., and Heins, A. J. (1975). Risk and the rate of return on financial assets: Some old wine in new bottles. *Journal of Financial and Quantitative Analysis*, 10(5), 775-784.

Haugen, R. A., and Baker, N. L. (1991). The efficient market inefficiency of capitalization-weighted stock portfolios. *Journal of Portfolio Management*, 17(3), 35-40.

Haugen, R. A., and Baker, N. L. (1996). Commonality in the determinants of expected stock returns. *Journal of Financial Economics*, 41(3), 401-439.

Haugen, R. A., and Baker, N. L. (2010). Case closed. In Handbook of Portfolio Construction (pp. 601-619). Springer, Boston, MA.

Houweling, P., and Van Zundert, J. (2017). Factor investing in the corporate bond market. *Financial Analysts Journal*, 73(2), 100-115.

Hsu, C. C., and Chen, M. L. (2017). The timing of low-volatility strategy. *Finance Research Letters*, 23, 114-120.

Hsu, J. C., Kudoh, H., and Yamada, T. (2013). When sell-side analysts meet high-volatility stocks: an alternative explanation for the low-volatility puzzle. *Journal of Investment Management*, 11(2), 28-46.

Huij, J., and Kyosev, G. (2016). Price Response to Factor Index Additions and Deletions. SSRN working paper no. 2846982.

Ilmanen, A. (2012). Do financial markets reward buying or selling insurance and lottery tickets?. *Financial Analysts Journal*, 68(5), 26-36.

Israel, R., Palhares, D., and Richardson, S. (2018). Common Factors in Corporate Bond Returns, *Journal of Investment Management*, 16(2) 17–46.

Jordan, B. D., and Riley, T. B. (2015). Volatility and mutual fund manager skill. *Journal of Financial Economics*, 118(2), 289-298.

Joshipura, M., and Joshipura, N. (2016). The Volatility Effect: Evidence from India. *Applied Finance Letters*, 5(1), 12-27.

Karceski, J. (2002). Returns-chasing behavior, mutual funds, and beta's death. *Journal of Financial and Quantitative Analysis*, 37(4), 559-594.

Kirchler, M., Lindner, F., and Weitzel, U. (2018). Rankings and risk-taking in the finance industry. *Journal of Finance*, 73(5), 2271-2302.

Kumar, A. (2009). Who Gambles in the Stock Market? *Journal of Finance, 64(4)*, 1889-1933.

Li, X., Sullivan, R. N., and Garcia-Feijóo, L. (2014). The limits to arbitrage and the low-volatility anomaly. *Financial Analysts Journal*, 70(1), 52-63.

Liu, J., Stambaugh, R. F., and Yuan, Y. (2018). Absolving beta of volatility's effects. *Journal of Financial Economics*, 128(1), 1-15.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance*, 32(4), 1151-1168.

Miller, M. H. and Scholes, M. (1972). Rates of Return in Relation to Risk: A Reexamination of Some Recent Findings. Studies in the Theory of Capital Markets, Praeger, New York, 47-78.

Moskowitz, T. J. and Vissing-Jorgensen, A. (2002). The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review*, 92(4), 745-778.

Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1), 1-28.

Novy-Marx, R. (2014). Understanding defensive equity. NBER working paper no. 20591.

Novy-Marx, R., and Velikov, M. (2015). A taxonomy of anomalies and their trading costs. *The Review of Financial Studies*, 29(1), 104-147.

Novy-Marx, R., and Velikov, M. (2018). Betting Against Betting Against Beta. SSRN working paper no. 3300965.

Scherer, B. (2011). A note on the returns from minimum variance investing. *Journal of Empirical Finance*, 18(4), 652-660.

Shefrin, H. and Statman, M. (2000). Behavioral Portfolio Theory. *Journal of Financial and Quantitative Analysis*, 35(2), 127-151.

Shumway, T. (1997). The Delisting Bias in CRSP Data. *Journal of Finance*, 52(1), 327-340.

Shumway, T. and Warther, V. A. (1999). The Delisting Bias in CRSP's Nasdaq Data and Its Implications for the Size Effect. *Journal of Finance*, 54(6), 2361-2379.

Sias, R. W. (1996). Volatility and the institutional investor. *Financial Analysts Journal*, 52(2), 13-20.

Sirri, E. R. and Tufano. P. (1998). Costly Search and Mutual Fund Flows. *Journal of Finance*, 53(5), 1589-1622.

Snowberg, E., and Wolfers, J. (2010). Explaining the favorite–long shot bias: Is it risk-love or misperceptions?. *Journal of Political Economy*, 118(4), 723-746.

Soe, A. M. (2012). Low-volatility portfolio construction: ranking versus optimization. *Journal of Index Investing*, 3(3), 63-73.

Swinkels, L., Blitz, D., Hallerbach, W., and van Vliet, P. (2018). Equity Solvency Capital Requirements-What Institutional Regulation Can Learn from Private Investor Regulation. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 43(4), 633-652.

Van Vliet, P. (2011). Ten Things You Should Know About Low-Volatility Investing. *Journal of Investing*, 20(4), 141-143.

Van Vliet, P. (2018). Low Volatility Needs Little Trading. *Journal of Portfolio Management*, 44(3), 33-42.

Van Vliet, P., and De Koning, J. (2017). *High Returns from Low Risk: A Remarkable Stock Market Paradox*. John Wiley & Sons.

Wahal, S. (2019). The profitability and investment premium: Pre-1963 evidence. *Journal of Financial Economics*, 131(2), 362-377.

Walkshäusl, C. (2014). International low-risk investing. *Journal of Portfolio Management*, 41(1), 45-56.

Exhibit 1: Ten portfolios sorted on 60-month beta, 1963-2018

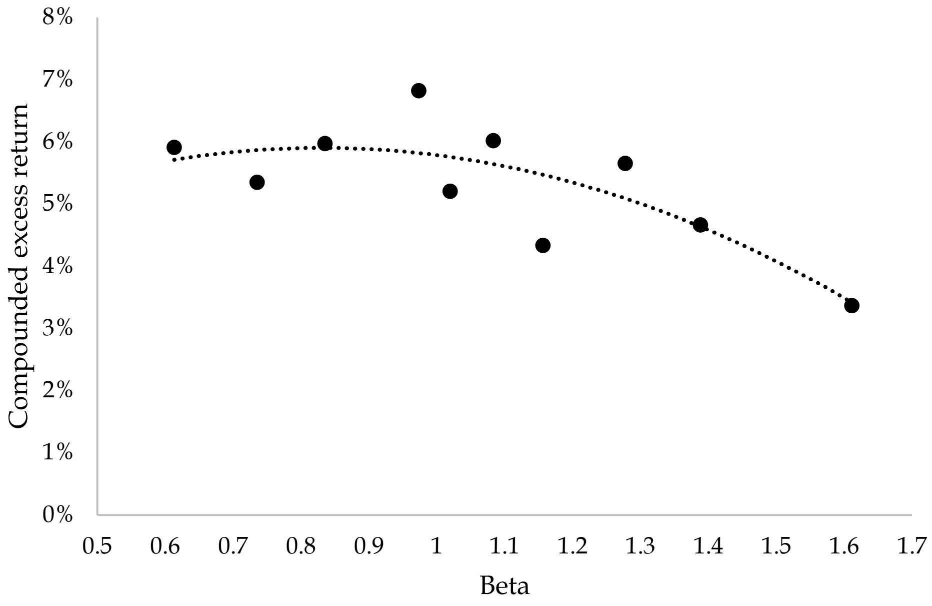

Source: Kenneth French data library

Exhibit 2: 5x5 portfolios sorted on size and 60-month beta, 1963-2018

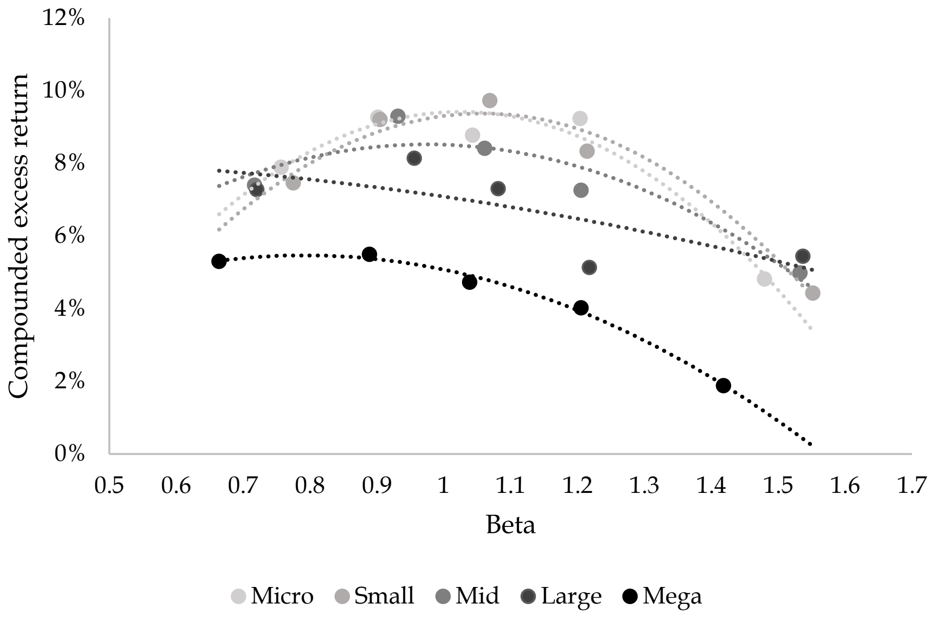

Source: Kenneth French data library

Exhibit 3: Ten portfolios sorted on 36-month volatility, 1929-2018

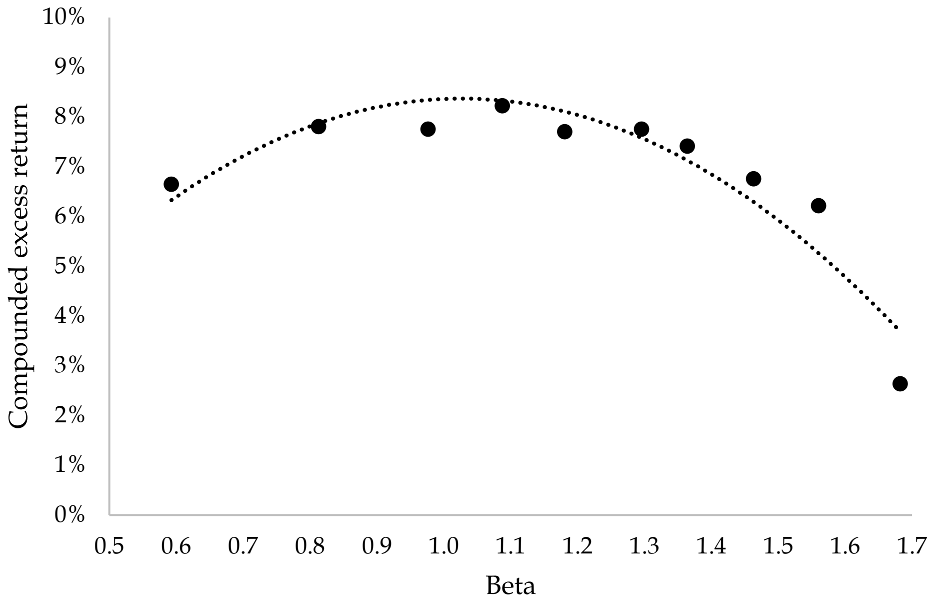

Source: paradoxinvesting.com

Exhibit 4: Ten portfolios sorted on 36-month volatility, sub-periods

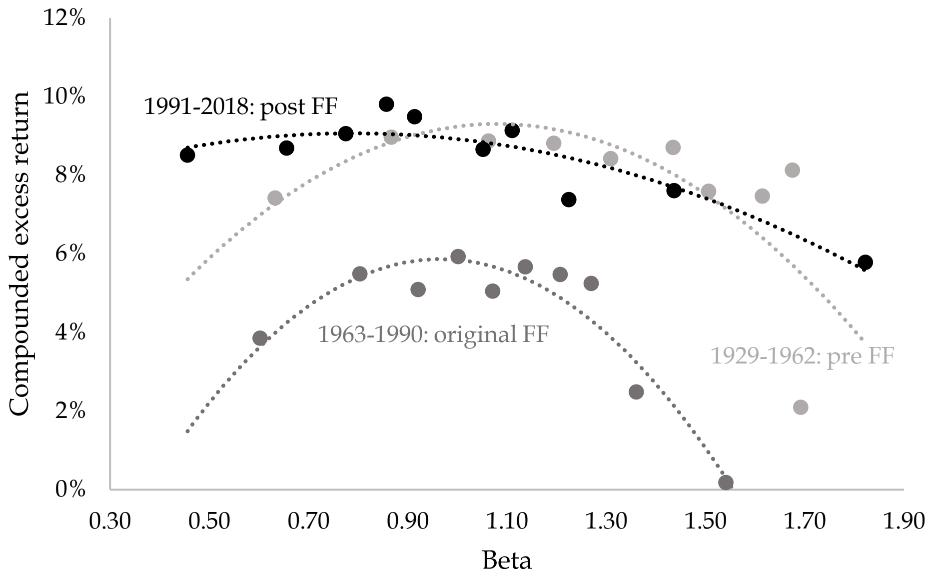

Source: paradoxinvesting.com

Exhibit 5: Betting-Against-Beta (BAB) premium around the world, various start dates until end-2018 (maximum available history for each country)

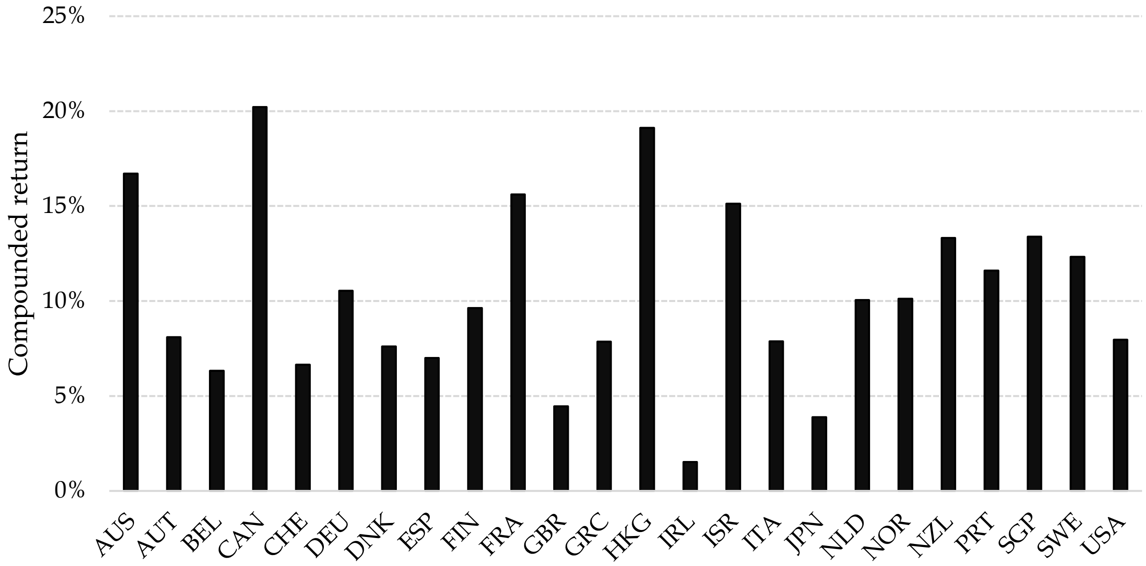

Source: AQR data library

Exhibit 6: Factor premiums by decade

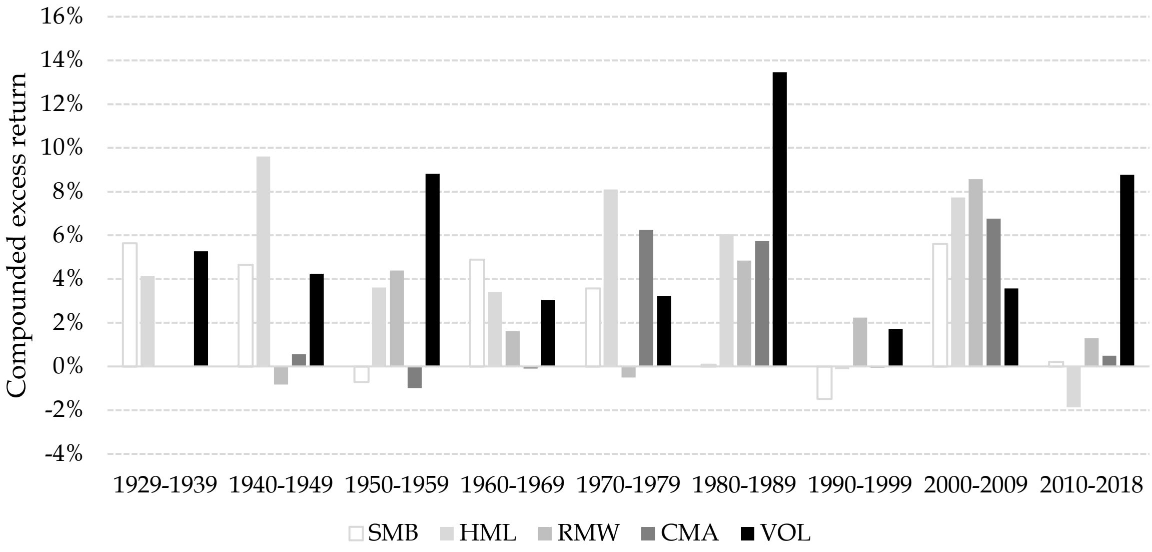

Sources: Kenneth French data library, paradoxinvesting.com, Journal of Financial Economics website

Exhibit 7: Strength of factor premiums, 1940-2018

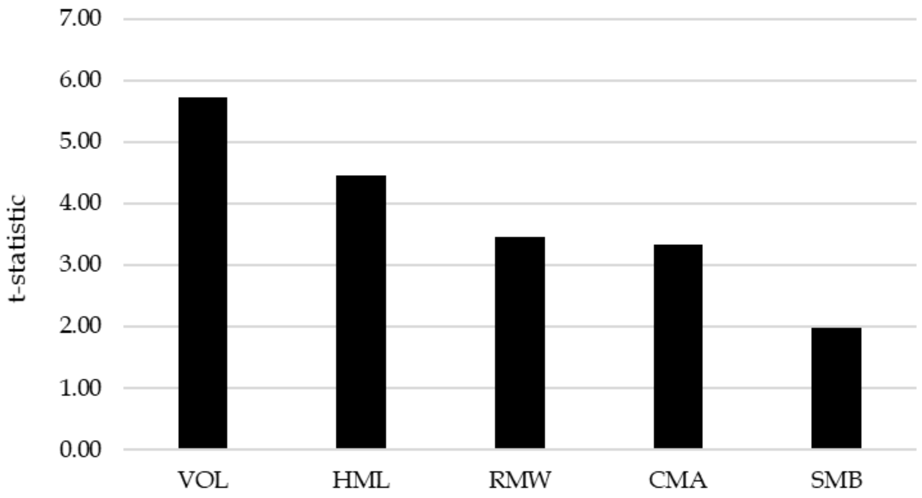

Sources: Kenneth French data library, paradoxinvesting.com, Journal of Financial Economics website

Exhibit 8: Ten portfolios sorted on 60-day variance and 60-day residual variance, 1963-2018

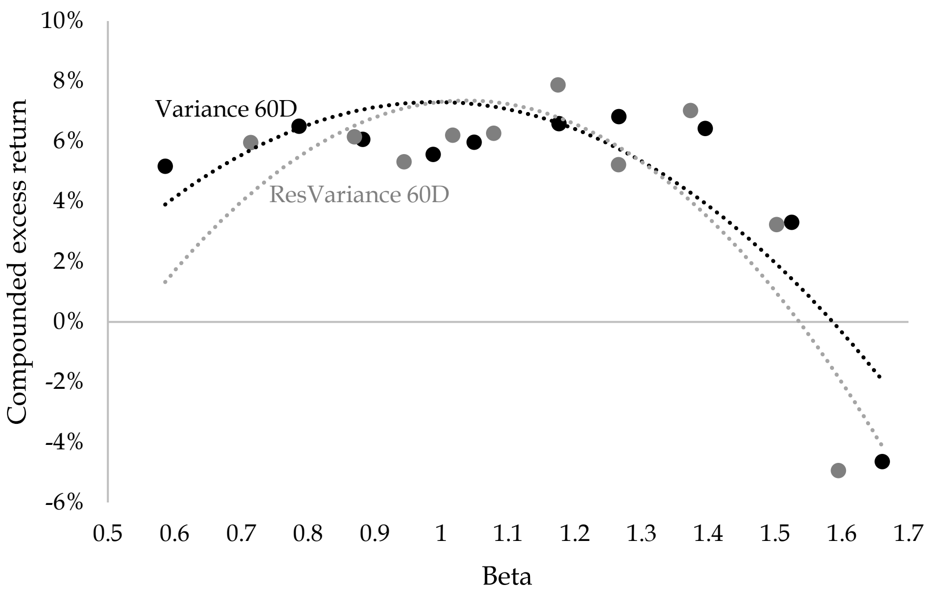

Source: Kenneth French data library