Erasmus University Rotterdam

Master Thesis Financial Economics

Is it different this time?

About the performance of low-volatility strategies in a changing market environment

Author	Mees Bovelander
Supervisor	P.J.P.M. Versijp
Date	xx-xx-2023

Abstract

The objective of this study is to comprehensively examine the underlying driving forces behind the low-volatility premium and test whether these forces are resilient to changes in the market structure. Utilizing and bundling a growing body of empirical research, we find that institutional investors as well as individual investors have an identifiable preference for volatile equities. This preference for volatility could explain a particularly remarkable cross-sectional pattern in asset pricing theory called the low-volatility premium. As the persistence of the low-volatility effect is explained by incentives and biases that occur in the process of active capital allocation, the shift towards passive index investing could potentially jeopardize the continuity of the low volatility premium. Consequently, this research paper has evaluated to what extent the low-volatility effect - or its decomposed variants - has been resilient to changes in the market structure. We find that, in the period 1963-2022, portfolios comprising low beta, low variance, and low idiosyncratic volatility stocks have outperformed their volatile counterparts in terms of absolute excess return, realized alpha, and the Sharpe ratio. Based on the recent performance (2005-2022) of low-volatility portfolios and the development of the alpha spread between low-volatility and high-volatility strategies, we conclude that the trend towards passive index investing and the corresponding change in the market structure has not led to a diminishing of the volatility premium.

The four most expensive words in the English language are ‘This time, it's different.’

Sir John Templeton

Introduction

In 1936, John Maynard Keynes compared investing to a beauty contest in which participants were tasked with choosing, out of a hundred faces, the six that they believed the general public would consider the most attractive. According to Keynes, *each competitor has to pick not those faces* [read: stocks] *he himself finds the prettiest, but those which he thinks likeliest to catch the fancy of other competitors*. *We thereby reach the third degree where we devote our intelligence to anticipating what the average opinion expects the average opinion to be.* The message of Keynes is simple: when investing, it is all about understanding the other market participants. Achieving superior returns is fundamentally a matter of comprehending the motivations of other market participants and possessing the foresight to anticipate.

Keynes' beauty contest metafor invites us to think about the true nature of investing. Whereas many scholars tend to treat the field as an exercise in uncovering statistical regularities that try to explain stock prices, Keynes reminds us that the market is a dynamic and complex web of human interactions, shaped by the diverse and ever-changing desires, incentives, and cognitive biases of the participants. Only by embracing this perspective can we hope to gain a true understanding of the market and its workings. (Keynes, 1936; Thaler, 2015)

The academic discipline of modern finance has long been preoccupied with the motivations and preferences of market participants. An insight gained by academics in this field is that institutional investors as well as individual investors have an identifiable preference for volatile equities. (Kahneman & Tversky, 1992; Brunnermeier et al, 2007; Barberis & Huang, 2008; Bali et al, 2011) This preference for volatility has led to a particularly remarkable cross-sectional pattern in asset pricing theory called the low volatility premium. (Black, Jensen & Scholes, 1972; Fama and Macbeth, 1973; Haugen & Heins, 1975; Clarke, de Silva, and Thorley, 2006; Ang, Hodrick, Xing & Zhang, 2006; Van Vliet and Blitz, 2007) The premium refers to the outperformance of low volatility stocks compared to high volatility stocks which goes against the fundamental principle of modern finance, which is illustrated in the Capital Asset Pricing Model, that risk is - and should be - compensated with higher expected return. (Sharpe, 1964; Lintner, 1965)

The theoretical part of this research paper aims to combine the results of various studies in the fields of behavioral economics, behavioral finance, and modern finance, in order to uncover the driving forces behind the low volatility premium. Utilizing existing literature, it shall be demonstrated that both institutional and individual investors exhibit a preference for highly volatile equities. After uncovering the driving forces behind the low volatility premium the altered market structure brought about by the emergence of passive investment options will be discussed. Using a combination of recent studies and existing literature, it will be examined whether this evolving market structure could potentially constitute a threat to the persistence of the low volatility premium. Moreover, we pose the vital question whether historical research findings on the aforementioned premium remain relevant in the face of fundamental changes to the market structure. Driven by this rationale, the empirical part of this research paper aims to evaluate, by assessing whether there has been a significant change in the low volatility premium over time, to what extent the low-volatility effect - or its decomposed variants - has been resilient to changes in the market structure. To be concise, we test whether the magnitude of the low volatility premium has significantly changed since the emergence of exchange traded funds and the corresponding shift towards passive investing.

This research paper contributes to the existing literature in several ways. Inspired by the message of Keynes, this paper will not lose sight of the human dimension of the market. It is for this reason that this paper will devote significant attention to a diverse body of literature from the fields of behavioral economics and modern finance to arrive at a plausible and comprehensive explanation for the continued existence of the low volatility premium. Furthermore, this research paper presents a unique perspective on the low volatility premium by giving due consideration to the latest developments in market structure and their potential impact on the driving forces behind the aforementioned premium. By questioning ergodicity, we pose the question whether the historical performance of low volatility stocks is an accurate indicator of future returns. After all, it may be different this time.

Part I

The Cross-Section of stock returns

Executive summary

Part I functions as the theoretical foundation for analyzing cross-sectional patterns in stock prices. According to the efficient market hypothesis, prices of all traded assets are equal to their intrinsic value. (Malkiel, 1989) Assuming the market establishes efficient prices, it is impossible to realize a risk-adjusted outperformance. Since investors dislike risk, they demand a higher compensation for more risky assets. (Falkenstein, 2009) Hence, risk exposure negatively affects the value of an asset. This perception of risk and return was first described in Markowitz' Portfolio Theory. (Markowitz, 1952) Sharpe and Lintner later used the theoretical foundation of Markowitz' Portfolio Theory to construct the Capital Asset Pricing Model which expresses a positive linear relationship between the amount of systematic risk of an asset and its expected return. (Sharpe, 1964; Lintner, 1965) The 'excess' return that remains after adjusting for systemic risk - thus unexplained by the CAPM - is called alpha. Assuming volatility is an appropriate measure of risk, the existence of a structural alpha, implies that the market is not correctly establishing prices which cannot be explained by the efficient market hypothesis. Factor premiums therefore pose a direct threat to the validity of this hypothesis or to the asset pricing model. If mispricing can exist, for which there is a large body of empirical evidence, the inefficiency of the market is rooted in a distortion of demand. (Black, Jensen & Scholes, 1972; Miller & Scholes, 1972; Black, 1972; Fama & MacBeth, 1973) The distortion of demand diverts the price of an asset from its intrinsic value, so that the compensation an investor receives becomes disproportionate to the amount of risk he bears. However, 'irrational behavior' at the demand side is not sufficient to explain structural mispricing in a stock market equilibrium. An explanation for the absence of arbitrageurs willing to exploit mispricing is equally important.

1.1 Market efficiency

According to the efficient market hypothesis, the market mechanism will ensure that prices of all traded assets are equal to their intrinsic value. Therefore, all traded securities should have a similar trade-off between risk and return. (Malkiel, 1989) Shares of different companies may yield different returns, but once controlled for the amount of risk exposure, this should be equal. The assumption is made that market participants are both rational and (completely) informed and will therefore 'agree' on stock prices that correctly reflect the intrinsic value of the corresponding company. When an inefficient price is established, which is reflected in the risk-return ratio, it should be immediately exploited and therefore 'corrected' by arbitrageurs. (Jensen, 1978) In other words, to arrive at an efficient stock market equilibrium, where price and value are aligned, the efficient market hypothesis assumes that all publicly available information is incorporated by rational agents ensuring a 'right' price for every traded asset. (Malkiel & Fama, 1970) *Even if* stock prices diverge from their intrinsic value, this will be exploited by arbitrageurs.

A related, but conceptually distinct, component of the efficient market hypothesis concerns whether it is possible to 'beat' the market. This 'no free lunch' principle builds upon the before mentioned 'right price' principle. When it is assumed that the market establishes efficient prices, it is impossible to realize a risk-adjusted outperformance on a structural basis. After all, if prices correctly reflect a trade-off between risk and return, increasing your return will only be possible by taking more risk. (Thaler, 2015) The rationale behind this so-called 'risk premium' stems from the historical idea that risk is perceived as something investors dislike. The assumption is made that investors are, in general, risk-averse. Since investors do not like risk, they demand a higher compensation for more risky assets. (Weil, 1989) Hence, the amount of risk exposure negatively affects the value of an asset. (Falkenstein, 2009) This perception of risk and return was first described in Markowitz' (1952) Portfolio Theory. According to Markowitz' theory, rational agents will invest in portfolios that deliver the highest expected returns for a specific level of risk. Within this mean-variance framework, the portfolio that offers the highest risk-return trade-off is described as efficient. (Markowitz, 1952)

…

1.2 The Capital Asset Pricing Model

Sharpe and Lintner later used the theoretical foundation of Markowitz' Portfolio Theory to construct the Capital Asset Pricing Model (CAPM). According to the standard one factor Capital Asset Pricing Model, there should be a positive relationship between the amount of risk exposure, measured by the degree of relative volatility, and the amount of return. The model expresses a positive linear relationship between the amount of systematic risk of an asset and its expected return, which can be (visually) displayed as the Security Market Line (SML). The amount of systemic risk is captured by beta, which is a regression coefficient that indicates the extent to which an asset fluctuates relative to a benchmark index that should be a value-weighted average of all available risky assets. It is therefore a measure of relative volatility. (Sharpe, 1964; Lintner, 1965)

The launch of the CRSP database in 1964 has, among other factors, enabled academics to extensively analyze the relationship between risk and return. (Thaler, 2015) Since then, numerous studies have shown that the single factor CAPM does not have sufficient explanatory power. (Black, Jensen & Scholes, 1972; Miller & Scholes, 1972; Black, 1972; Fama & MacBeth, 1973) Beta as the sole variable explaining returns on stocks was even declared 'dead' by Eugene Fama and Kenneth French as it 'completely failed to predict returns'. (Fama & French, 1992). The validity of the CAPM was further compromised by the growing body of academic literature proving the existence of cross sectional pricing anomalies such as the Value premium and the Size premium. (Banz, 1981)

1.3 Risk revisited

In the case of an anomaly, empirical evidence is found that certain types of stocks, often with a common property, offer a risk-return ratio that is not explained by the asset pricing model. This indicates that selecting a portfolio of securities based on this quantitative characteristic, which is also called a factor, yields a return that is higher (or lower) than would be expected based on the amount of risk. In other words, the return of an asset cannot be explained by just the corresponding amount of (relative) volatility. The amount of 'excess' return that remains after adjusting for market-related volatility is called *alpha*. Therefore, alpha is often described as a premium. It describes the proportion of return that is unrelated to systemic risk and is therefore interpreted as a '*free lunch'*. Investors do not have to bear any additional risk for this portion of the return. Assuming volatility is an appropriate measure of risk, the existence of a structural alpha, regardless of its sign, cannot be explained by the efficient market hypothesis. Factor premiums therefore pose a direct threat to the validity of this hypothesis. (Jensen, 1978; Thaler, 2015; Falkenstein, 2012)

But is volatility an appropriate measure of risk? According to the CAPM, stock returns are explained by their exposure to systemic risk. However, as mentioned, beta fails to predict returns as a sole explanatory variable. The observed relationship between volatility and return therefore deviates from what is expected on the basis of theory. This can be explained in two different ways. First, there is the argument that beta does not correctly reflect the total risk of an asset. If there are sources of risk that are not captured by volatility measures such as beta, it could be the case that the alpha of a factor is just a compensation for hidden risk. Explaining factor premiums as a compensation for hidden risk is a convenient way to defuse criticism on the efficient market hypothesis. If factor premiums are a compensation for risk that is not captured by volatility measures, it can be seen as a way in which the market establishes its efficient prices. Shareholders of stocks with certain factor loadings are rewarded with excess return after adjusting for market-related volatility because additional risk has been borne. In this way anomalies and the efficient market can coexist.

Since it became clear that asset prices cannot be explained just by the amount of relative volatility, the idea of multidimensional risk gained traction. In the 1992 paper *'The Cross-Section of Expected Stock Returns'* Fama and French observe that diversified portfolios of small cap stocks, which are characterized by their small market capitalization, outperform portfolios of stocks with large market capitalization. Furthermore, portfolios of Value stocks, which are characterized by a low market to book ratio, outperform portfolios of Growth stocks, which are characterized by a high market to book ratio. The authors incorporated these pricing errors (Value and Size) as additional explanatory variables in their 3-factor asset pricing model. This 3-factor model was later expanded to a 5-factor model by the addition of a Profitability factor and an Investment factor. (Fama & French, 2006) Although the authors find several factors besides systemic risk that explain stock returns, their belief that the market establishes efficient prices, which correctly reflect the trade-off between risk and return, is unprecedented. According to Fama and French, the observed factors that explain stock returns are tied to non-systemic sources of risk. Value stocks and small cap stocks are, according to the authors, inherently risky. The market is efficiently pricing these non-systemic risks by compensating investors who expose themselves to these (risk)factors. (Fama & French, 1992) The way in which Fama and French explain the existence of factor premiums symbolizes the way in which academics, since the appearance of Markowitz's

Portfolio Theory and the CAPM of Sharpe and Lintner, have approached risk. According to Falkenstein, *'The belief that risk, if properly measured, must be positively related to returns is very deep among academics'* (Falkenstein, 2009)

1.4 Joint hypothesis problem

The key takeaway of the previous passage is that Asset pricing models cannot be used to prove or disprove market efficiency. Any test of market efficiency using an asset pricing model is a joint test of market efficiency and a test of the used pricing model. In other words, when market efficiency is tested by the use of an asset pricing model, the joint hypothesis problem occurs. If an asset model suggests that the market is not incorporating information into prices, it is hard to tell whether this indicates that the market is inefficient or the model is wrong. Academics who strongly believe in the efficiency of the market therefore usually assume that the existence of factor premiums indicate that the used pricing model is inaccurate. The inaccuracy manifests itself mainly in the failure to properly incorporate risk.

1.5 Free lunch

Assuming that factor premiums are not explained by a compensation for additional risk factors, the existence of a factor premium indicates that the market mechanism is unable to establish the right price for every traded asset. The existence of alpha, both positive and negative, in that case indicates mispricing that can be exploited with long/short strategies. The question remains, of course, how structural mispricing can continue to exist in a stock market equilibrium. It is expected that mispricing, especially if it is as well documented as the factors discussed, will eventually disappear because of arbitrage opportunities. In other words, it is expected that the magnitude of alpha will decrease over time due to the existence of arbitrage opportunities.

In general, there is a shared tendency in the emergence of mispricing in the stock market. There are certain conditions that have to be met in order to end up in an equilibrium in which mispricing can exist. First, there must be a group of market participants that, by their actions, divert the price of an asset from its intrinsic value. This could, evidently, also be a lack of action in the case of changing intrinsic value. This behavior, which is generally characterized as irrational, can manifest itself in buying securities at a price above the intrinsic value or selling at a price below the intrinsic value. The causes of this behavior, which will be discussed in detail in Part II, are ranging from subject confusion - *participants not being able* *to establish correct price*s - to behavioral biases. (Hussam, Porter & Smith 2008) Therefore, *If* mispricing can exist, for which there is a large body of empirical evidence, this market inefficiency is rooted in a distortion of demand. The price of an asset cannot divert from its intrinsic value without the influence of a marginal trader. The trading behavior of this marginal trader influences the price of an asset so that the compensation another investor receives becomes disproportionate to the amount of risk he bears. In order to explain the occurrence of mispricing and the corresponding outperformance, analyzing what drives market participants to buy (sell) assets at a price that is above (below) their intrinsic value is therefore crucial. Whether this is based on behavioral fallacies or perverse incentives, there has to be a plausible motive. Once this motive is established and substantiated by empirical evidence, one can make assumptions about the development of these motives over time. *If* it is likely that certain motives will change over time, it is reasonable to assume that the degree of mispricing will change. After all, the persistence of factor premia depends on the persistence of the motives that drive mispricing in the stock market. (Thaler, 2015)

However, this is only half the explanation. It is, of course, difficult to deny that mispricing *can* occur. In fact, this does not even contradict the efficient market hypothesis. It becomes problematic, for this hypothesis, when mispricing is not exploited by arbitrageurs. Therefore, ‘irrational behavior’ at the demand side is not sufficient to explain structural mispricing in a stock market equilibrium. An explanation for the absence of arbitrageurs is equally important.

1.6 Remarks

The first passage has laid a theoretical foundation for analyzing cross-sectional patterns in stock prices. This research paper will be characterized by a similar structure. First, patterns in the cross section will be analyzed extensively. Then, it will be analyzed whether these patterns can be explained by risk factors that are not sufficiently incorporated in the asset pricing model. If this is not the case, it will be investigated what could be the cause of the demand distortion that leads to this structural mispricing. Finally, there should be an explanation for the absence of arbitrageurs willing to exploit this mispricing.

Part II

The low volatility premium

Executive summary

In Part II the empirical underpinnings for the continued existence of the low volatility premium will be discussed. According to the discussed literature, the low volatility effect is significantly present during several periods in time, the effect is significantly present in multiple geographic areas and the effect is robust to changes of the volatility measure. The results of multiple research papers are combined to arrive at a plausible explanation for the persistence of the low volatility premium. In order to maximize the outperformance for a given level of tracking error, money managers construct portfolios with a beta greater than one. Therefore, building on the decomposition of the volatility discussed by Asness, Frazzini & Pedersen (2020), the incentive structure of institutional money managers in combination with leverage constraints mainly influences the beta effect while the low idiosyncratic volatility (IVOL) is mostly caused by the behavioral bias (lottery preference) of individual investors. Therefore, the existence of the low-volatility premium is explained by the incentives and biases of active market participants.

2.1 Empirical evidence for the low volatility premium

Although there are several established anomalies in modern finance theory, a particularly remarkable one is the superior track record of low-volatility stock portfolios. The empirical evidence that low-volatility and low-beta portfolios have offered high average returns in combination with small drawdowns goes against the fundamental principle of modern finance, which is illustrated in the Capital Asset Pricing Model, that risk is - and should be - compensated with higher expected return. Numerous studies, including the papers by Eugene Fama and Kenneth French, have found evidence for the existence of a so-called *low volatility premium* during several periods in time and in multiple geographic areas (Fama & MacBeth, 1973). The empirical evidence for the existence of a low volatility premium indicates that selecting portfolios based on low volatility can, in the long run, provide equivalent returns with a lower risk profile or higher returns with an equivalent risk profile. Therefore, selecting portfolios with low volatility provides a superior risk/return ratio, which is reflected in the *alpha* coefficient (α). In short, although Markowitz's Portfolio Theory and the CAPM describe a positive relationship between the amount of systematic risk of an asset and its expected return, in practice this relationship seems rather negative.

With the research paper '*The Capital Asset Pricing Model: Some Empirical Tests'* Black, Jensen and Scholes were among the first academics to extensively analyze the predictive value of the CAPM model by Sharpe and Lintner. (Black, Jensen & Scholes, 1972). The aim of this paper was to test the relationship between systemic risk and expected returns. The authors find that the slope of the beta coefficient significantly differs from what would be expected based on the CAPM. Contrary to the CAPM, Black et al. conclude that '*expected excess return on an asset is not strictly proportional to its beta*'. These findings are supported by the later conducted tests of the CAPM by Fama and Macbeth (1973). These tests of the CAPM are the first evidence that the relationship between systematic risk and expected return may not be as straightforward as described by Sharpe and Lintner. However, the first evidence for the existence of a low volatility premium was found by Haugen and Heins (1975). The authors draw the same conclusion that there is no clear empirical evidence for the positive linear relationship between the amount of systemic risk of an asset and its returns. More importantly, Haugen and Heins find that diversified portfolios of low volatility stocks tend to outperform their high volatility counterparts. (Haugen & Heins , 1975)

Since the discovery of the low volatility premium, the superior track record of low-volatility stock portfolios has been extensively analyzed. Clarke, de Silva, and Thorley (2006) find that minimum-variance portfolios deliver returns comparable to the market portfolio while lowering the portfolio volatility by 25%, resulting in a superior Sharpe ratio. Van Vliet and Blitz (2007) find, by ranking stocks based on their past return volatility, that portfolios of stocks with the lowest volatility are '*associated with Sharpe ratio improvements*'. Ang, Hodrick, Xing & Zhang (2006) find that portfolios with low idiosyncratic volatility outperform portfolios with high idiosyncratic volatility. The results of these studies indicate that the low volatility effect is significantly present during several periods in time, the effect is significantly present in multiple geographic areas and the effect is robust to changes of the volatility measure. (Proietto, 2018)

2.2 Institutional Investors and Incentive structures

In Markowitz's portfolio theory, investors make a trade-off between risk, measured in terms of volatility and return. This assumes that volatility is an appropriate way to quantify risk. This assumption can be substantiated from the perspective of an individual investor, who trades on his own behalf. However, According to research done by Morgan Stanley, '*retail investors make up about 10% of the daily trading value of the 3,000 biggest U.S. stocks*'. (Reuters, 2021) Therefore, a large proportion of market participants trade on behalf of institutional parties such as banks, insurers and asset managers. Ultimately, these firms pursue an equal goal as retail investors; maximizing the total value of their assets (under management) while mitigating risks. However, as is often the case, the goals of these institutional parties are not necessarily aligned with the incentive structures of their employees. Due to the great influence of these firms on stock prices, it is crucial to understand what drives institutional money managers in their investment decision to explain the low volatility premium. The following passage will use the results of various research papers to explain how the incentive structure of institutional investors in combination with leverage constraints can lead to a situation where low volatility stocks are relatively risky from a managerial perspective and are therefore unattractive. The lack of attractiveness may then explain the potential mispricing of low volatility stocks.

The performance of institutional money managers is in most cases measured by the unlevered 'information ratio' relative to a fixed benchmark. The information ratio is a measure that divides the outperformance of a portfolio by the tracking error. Managers are therefore incentivised to create portfolios that maximize the ratio between returns and tracking error. In other words, the ultimate goal is to have as much correlation as possible with the benchmark while maximizing returns. (Baker, Bradley & Wurgler, 2011) Due to the way in which managers are assessed and the incentive structure that follows from this, it is questionable whether volatility is an appropriate measure of risk for these professionals. Based on the performance measure, the risk for money managers seems to manifest itself mainly in the (short term) tracking error. (De Koning & Van Vliet, 2018) In order to preserve a job, a money manager should not lag too much behind the benchmark, even if 'trailing' the benchmark benefits the long-term performance of the portfolio.

$$TE = \sqrt{\frac{\sum (F - I)^2}{N - 1}}$$

Where F = fund return,
I = index return and
N = number of periods.

In the 1992 paper '*A mean/variance analysis of Tracking Error'* Ross demonstrates the implications of a fixed benchmark mandate. According to Ross, measuring performance relative to a fixed benchmark incentivizes money managers to construct portfolios that maximize outperformance for a given level of tracking error volatility (TEV). Ross creates an agency model that alters Markowitz's portfolio theory in order to be consistent with the incentives of money managers. Ross' model shows that portfolios that are efficient from a TEV perspective generally have a beta greater than 1. Thus, this agency model shows that maximizing the outperformance relative to a benchmark for a given level tracking error volatility goes together with an increased exposure to systemic risk. (Ross, 1992) Israelsen and Cogswell (2006) come to the same conclusion in their paper 'The error of tracking error'. The authors find, by ranking mutual funds based on tracking error, that funds with the lowest tracking error '*exhibit lower alpha, higher beta, and lower average performance compared to funds with high tracking error*'. Baker, Bradley and Wurgler (2011) show, by decomposing the information ratio, that low beta stocks (even those with significant positive alpha) are more likely to be underweight because their negative influence on the tracking error exceeds their positive effect on the outperformance of the portfolio. Low beta stocks, without the use of leverage, decrease the correlation between the benchmark and the portfolio and therefore impose a managerial risk by increasing the tracking error. Although stocks with low volatility generate significant positive alpha, these stocks are actually, due to the way performance is measured, more risky from a managerial perspective.

However, this reasoning fails to mention a very important point. Money managers could, in theory, use leverage to increase the correlation between low beta stocks and the benchmark, thereby reducing the tracking error. In the absence of leverage constraints, it is difficult to argue that low beta stocks are unattractive when maximizing the outperformance for a given level tracking error volatility. The preference for high beta stocks therefore cannot be explained on the basis of the incentive structure alone. Adding potential leverage constraints makes the argument plausible. In that case, the observed preference of money managers for a portfolio beta greater than 1 (Ross, 1992; Cogswell & Israelsen, 2006) will have to be materialized by selecting stocks with a beta greater than 1. In other words, the incentive to have a portfolio beta greater than one combined with a leverage constraint ensures that only high beta stocks can be used to achieve that desired beta. Therefore, money managers shift to riskier assets when there are binding leverage constraints. (Boguth & Simutin, 2018) This link between leverage constraints and the relative attractiveness of high beta stocks has been analyzed extensively by academics. Jylhä (2018) has shown the relationship between margin requirements and the slope of the security market line; Boguth and Simutin (2018), Malkhozov et al. (2016) and Adrian et al. (2014) show respectively that leverage constraints, international illiquidity and financial intermediary leverage can be used to explain the overperformance of a betting against beta strategy.

In short, the research papers of Ross (1992), Isrealsen and Cogswell (2006) and Baker, Bradley and Wurgler (2011) show that money managers are incentivized to maximize the information ratio. The authors have found different ways to reach an equivalent conclusion; in order to maximize the outperformance for a given level of tracking error volatility, money managers construct portfolios with a beta greater than one. If this incentive to have a portfolio beta greater than one is combined with a leverage constraint, money managers shift to assets with a higher beta. (Boguth & Simutin, 2018; Black, 1972, Frazzini, Pedersen, 2014) This is possibly part of the explanation for the continued existence of the low volatility premium. If institutional investors, who are responsible for the majority of the trading volume, are attracted to high beta stocks, low beta stocks are likely to be more attractively priced. This positively affects the future returns of low beta stocks.

2.3 Individual investors

As discussed, institutional investors, who are responsible for the majority of the trading volume, are - due to their incentive structure - attracted to high beta stocks which could be an explanation for the low volatility premium. After all, this could increase demand for stocks with certain characteristics independent of their intrinsic value. However, according to the literature, individual investors have a similar preference for highly volatile stocks.

Daniel Kahneman and Amos Tversky (1992), pioneers in the field behavioral economics, were among the first academics to document that individuals are unable to grasp probabilities. The probability weighting component of their cumulative prospect theory (1992) indicates that individuals structurally underestimate large probabilities and overestimate small probabilities. For probabilities near zero, like winning the lottery, individuals overestimate the chance of it happening. The findings of these authors explain, for instance, why people, although the expected outcome is negative, buy lottery tickets. Something rational choice theory, commonly used by economists, cannot. According to Bali et al. (2011), this so-called 'lottery preference' generates demand for high volatility stocks that is '*not warranted by the stocks' fundamentals*'. Barberis and Huang (2008) have drawn the same conclusion in their paper '*Stock as lotteries: the Implications of Probability Weighting for Security Prices*'. According to the authors, the skewness preference of individual investors will lead to the overvaluation of highly volatile stocks. Brunnermeier et al. (2007) conclude that investors have a preference for lotteries and therefore seek speculative stocks. As a result, high volatility stocks become relatively expensive in comparison to low volatility stocks. The similarity between the studies is that they use biased probability weighting to explain irrational demand for stocks - demand unrelated to the fundamental value of the underlying firm - with highly volatile payoff structures. People, and thus investors, have a certain bias that occurs when interpreting probabilities. Especially in the case of extreme probabilities. As a result, small probabilities are structurally being overestimated which leads to a lottery preference. This lottery preference manifests itself in a demand for high volatility stocks.

2.4 Volatility measures

So far, volatility has been approached as a general measure to describe the return distribution of a stock. In practice, there are several methodological methods to measure the volatility of a stock or portfolio. The three most commonly used methods to measure volatility are;

exposure to systematic risk (beta), variance of the returns, and idiosyncratic volatility. Although these methods are, by construction, highly correlated (Asness et al, 2020), it appears that in practice they capture a slightly different effect. According to Asness et al, the outperformance of low beta stocks is mostly related to the mechanism discussed in 2.2 while the outperformance of stocks with low idiosyncratic volatility (IVOL) is mostly caused by behavioral biases (lottery preference) discussed in 2.3. Therefore, to prevent confusion, a distinction will be made between the low-beta effect and the low-IVOL effect. The low-beta effect refers to the risk-adjusted outperformance of low-beta stocks compared to high-beta stocks. The low-IVOL effect refers to the risk-adjusted outperformance of stocks with low idiosyncratic volatility compared to stocks with high idiosyncratic volatility.

2.5 Omitted variables

An alternative - yet challenging - explanation for the low volatility premium could be that volatility is a proxy for (known) cross sectional patterns in stock pricing. It could be possible that companies with low volatility in the distribution of their stock returns are, in practice, smaller companies with a lower market valuation relative to their book value. If this is true, the risk adjusted outperformance of low volatility stocks could be attributed to the underlying factor loading of these stocks, or *vice versa*. In the case of known factors, such as size and value, this is relatively straightforward to test by analyzing whether low volatility portfolios generate an alpha in a multifactor asset pricing model that incorporates these factors. However, in the case of hidden (risk)premia, this becomes problematic. Moreover, the process of constructing portfolios based on past return volatility can also lead to 'significant sector bets' in case of high volatility dispersion between sectors. (Edwards, Lazzara & Preston, 2018) This is not only problematic because sector-specific results influence the factor premium, but also because idiosyncratic risks become significant. *If* low-volatility portfolios carry idiosyncratic risks due to sector weighting, it makes sense from an asset pricing perspective that beta, a measure of exposure to systematic risk, cannot fully explain the returns of a low volatility strategy. Hence, the factor premium.

Although it is difficult to completely refute the impact of omitted variables on the risk-adjusted outperformance of low volatility stocks, empirical analysis has shown that the performance of low volatility stocks cannot be explained by their corresponding Size and Value factor loadings. (Falkenstein, 2013) Furthermore, although constructing portfolios based on past return volatility leads to 'significant sector bets', this large allocation tends to fluctuate over time. (Edwards, Lazzara & Preston, 2018) In other words, the allocation of a low volatility portfolio towards a certain sector changes over time which ensures an intertemporal diversification of idiosyncratic risk. Because low volatility stocks have a large allocation to different sectors at different moments in time, it is difficult to substantiate that their overperformance is attributed to sector allocation.

Finally, it is important to analyze whether the risk-adjusted outperformance of low volatility stocks is not 'just' a compensation for (hidden) risk factors. As discussed in 1.3, if there are sources of risk that are not captured by the asset pricing model, it could be the case that the alpha of a strategy is a compensation for hidden risk. In the case of low volatility stocks, it is very difficult to substantiate that the outperformance is explained by uncaptured risk factors. Blitz and Van Vliet (2007) show that stocks with low volatility generally have smaller drawdowns and smaller standard deviation which leads to narrower return distributions. Furthermore, Falkenstein (2013) shows that stocks with low volatility have, on average, healthier credit ratings and are, on average, larger firms. Although it is in principle difficult to substantiate that historical stock performance can be related to the risk profile of a company, empirical results therefore show that it is unlikely that low volatility stocks are inherently risky. In fact, the evidence points in the opposite direction.

2.6 Remarks

As discussed in part I, in order to explain the occurrence of mispricing and the corresponding outperformance, analyzing what drives market participants to buy (sell) assets at a price that is above (below) their intrinsic value is crucial. Whether this is based on behavioral fallacies or perverse incentives, there has to be a plausible motive. Part II showed that institutional investors as well as individual investors have an identifiable preference for stocks with highly volatile payoff structures. In the case of constrained institutional money managers, selecting stocks with a high beta enables them to maximize the outperformance for a given level of tracking error volatility. The low beta component of the overall low volatility effect is therefore mainly explained by institutional investors. In the case of individual investors, biased probability weighting and the corresponding lottery preference explains the demand for stocks with highly volatile payoff structures. The low IVOL component of the overall low volatility effect is therefore mainly explained by individual investors.

Part III

Passive investing and market efficiency

Executive summary

The emergence of exchange traded funds (ETF) and the corresponding accessibility of passive investing is a direct result of technological developments in the stock market. The emergence of passive investment opportunities has ensured that both individual investors and institutional investors have increased their proportion of passive investments relative to their proportion of active investments. The shift towards passive investing could improve market efficiency by reducing the number of unskilled market participants, the cost of short selling and the methodological barriers for the exploitation of factor premia.

3.1 The emergence of index funds

Due to technological advancements in recent decades, trading securities has become increasingly computer-driven. The job of market makers - those who provide liquidity in financial markets - has, as a result, been completely transferred to computer-algorithms. The emergence of exchange traded funds (ETF) and the corresponding accessibility of passive investing is a direct result of this technological development in the stock market. These developments have therefore laid the foundations for a complete restructuring of the investment industry. In Part III it will be discussed to what extent this restructuring of the investment industry - the shift from actively managed portfolios towards passively managed portfolios - influences the efficiency of the market and thus the continued existence of the low volatility premium. So far, within this study, no clear distinction has been made between active and passive management of stock portfolios. By focusing on managerial incentive structures and personal biases, the emphasis of Part II was mainly on the relation between active portfolio management and the occurrence of mispricing in the stock market. *If* all market participants actively manage their portfolios and thus make investing decisions based on the described mechanisms, this would potentially be sufficient to explain the occurrence of mispricing. However, as will be discussed, this assumption cannot be made without deviating from the market environment as we currently observe.

3.2 The shift towards passive investing

Since the development of exchange traded funds, the proportion of passively managed funds compared to actively managed funds has grown rapidly which drastically changed the ownership structure in the US stock market. Both individual investors and institutional investors have increased their proportion of passive investments relative to their proportion of active investments. (Garleanu & Pedersen, 2020; Koijen et al, 2019) As a result, the majority of the US market is currently owned by institutional investors who have, on average, decreased their allocation to active management. To be concise, the aggregate active share of institutional investors - the proportion allocated to active managed funds - has decreased from 45% in 1980 to 25% in 2019 while institutional ownership increased from 29% to 76%. (ICI investment company, 2022; Lewellen et al, 2018) In short, capital that was traditionally invested in actively managed funds has now flown, via institutional parties such as Blackrock and Vanguard, towards passive (index) funds.

It should be clear that the emergence of exchange traded funds has caused a shift from actively managed funds towards passively managed funds. The question remains to what extent this shift influences the efficiency of the market and thus the continued existence of the low volatility premium. To be clear, changing ownership structures as a result of the shift from active to passive investing are, in isolation, insufficient to explain increases or decreases in the degree of market efficiency. Stock prices are, after all, determined by supply and demand. Therefore, mispricing is *de facto* caused by the 'marginal trader'. Although an increasing share of the global market capitalization is owned by passive index funds, the trading volume of these funds is, according to the 2018 Vanguard study: '*Setting the record straight: Truths about indexing'*, limited to around 5% of the daily trading volume in US stocks. According to the 2017 Blackrock study: *'Index Investing Supports Vibrant Capital Markets'*, for every 1$ of stock traded by index funds there is 22$ worth of stock traded by 'active mandates'. This limited amount of daily trading volume of passive index funds, despite the amount of assets under management, is explained by the fact that 'creation and redemption', which requires the most trading, takes place in the secondary market. However this is beyond the scope of this paper. The important takeaway is that, although capital has shifted towards passive funds, the majority of trading is still done by actively managed capital. Therefore, according to the Vanguard study, the impact of index funds on trading activity is minimal and there is no reason to believe that indexing hinders efficient price discovery. In other words, the shift towards passive investing has had limited consequences for the role of actively managed capital in the process of price discovery. Therefore, the mechanisms - described in Part II - that lead to the undervaluation (and the corresponding long term risk adjusted outperformance) of low-volatility stock would still be in place.

However, the unchanged role of actively managed funds in the process of price discovery is still insufficient to draw conclusions about whether the degree of market efficiency has been altered as a result of the shift towards passive investing. The shift towards index funds ensures that the proportion of active management decreases. Therefore, a group of market participants that used to actively manage their funds and thereby contribute to price discovery has now ceased to do so. As described by Fama and French (2005), this has a somewhat ambiguous effect on market efficiency. In the case of an uninformed active investor, a shift towards passive investing could positively affect market efficiency. In the case of an informed active investor, a shift towards passive investing could negatively affect market efficiency. Whether the shift towards passive investing positively or negatively influences market efficiency therefore depends on the person or group that shifts their funds. To be clear, the studies by Blackrock and Vanguard show that actively managed capital (still) has an important role in the process of price discovery. However, the ‘pool’ of active participants has shrunk. This changing pool of active market participants could positively or negatively influence the efficiency of the market, depending on who is switching.

3.2 Who is left?

The degree of market efficiency is a result of market participants incorporating information into their trading behavior and therefore influencing market prices. In this process of price discovery a market needs active participants to incorporate information until the equilibrium state of perfect efficiency is reached. Without active contribution, the market is unable to establish efficient prices. A complete shift towards passive investing would therefore lead to the rise of inefficient prices. This is known as the Grossman-Stiglitz paradox (1980). In the equilibrium state, where prices correctly reflect the intrinsic value of the underlying assets, passive investing is the most efficient. The outperformance of active management, relative to passive management, is after all bounded by the degree of exploitable inefficiency. In other words, the market mechanism only needs active participants in an inefficient state but cannot obtain the efficient state without active participants. Therefore, markets cannot be perfectly efficient. (Grossman & Stiglitz, 1980; Felix, 2020)

Since the amount of assets under management for index funds are increasing, this means that capital ‘leaves’ actively managed mandates. According to Fama and French (2005), this is no direct threat for the process of price discovery. Since the amount of active participants needed to maintain efficient markets is directly related to the cost of information, a decrease in the cost of ‘uncovering and evaluating’ information - as we have seen in the recent decades as a result of technological improvements - would decrease the need for active management. A smaller pool of active participants is therefore not a direct threat for market efficiency. Pástor and Stambaugh (2012) show that the contribution of active participants to efficient market prices can be related to their ability to ‘generate alpha’. When too much capital is allocated to actively managed funds, alpha will decrease or even become negative. In this case, if alpha is negative, it would be efficient if a proportion of the actively managed capital shifts towards passive investing strategies. Especially if the active participants with the lowest alpha shift towards passive investing. Berk and Van Binsbergen (2015) show that this is the case. According to the authors, ‘investors are proficient at rewarding skilled managers with flows into their funds'. Therefore, if capital shifts from active investing towards passive investing, the skillful allocators - those who contribute the most to market efficiency - will 'survive'. In short, the market does not need an unlimited number of active participants to establish efficient prices (Fama & French, 2005). In fact, When too much capital is allocated to actively managed funds, alpha will become negative, which is inefficient. (Pástor & Stamßaugh, 2012) Since investors reward skillful managers (Berk & Van Binsbergen, 2015), the shift toward passive investing, which ensures that allocation towards actively managed funds decreases, could therefore 'eliminate' unskilled participants, making markets more efficient. The emergence of index funds therefore positively affects market efficiency.

3.3 Securities lending

The 'crowding out' of unskilled capital allocators is not the only evidence of index funds increasing market efficiency. In the paper '*Passive asset management, securities lending and asset prices*' (2019) Palia and Sokolinski show that the emergence of index funds has decreased the costs of short selling. Index funds passively hold large amounts of assets which allows them to temporarily lend these assets to short sellers. Therefore, the increase of passive index funds has increased the competition among lenders which has decreased the lending rate. As a result, the cost of short selling has decreased. This should enhance the process of price discovery - and thus market efficiency - as it decreases market distortion.

3.4 Accessibility

Finally, the emergence of passively managed funds has decreased the barriers to exploit known cross sectional patterns in asset pricing. A common point of critique, with regard to factors such as Value, Size and Momentum, is that their practical implementation is often difficult. The methodological approach used to arrive at conclusions about the profitability of these factors is usually time-intensive and costly to the point of being close to impracticable. Exploitation of anomalies, like the low volatility premium, is therefore difficult to carry out in practice. This could explain the long run persistence of pricing anomalies. However, the emergence of exchange traded funds has made it increasingly accessible for investors to exploit pricing anomalies. ETF providers such as Blackrock and Vanguard offer passive index funds that follow similar methodological approaches as used in research papers. For example, Blackrock offers the S&P 500 Low Volatility ETF which allows investors to passively invest in the least volatile S&P 500 stocks. As the methodological barriers for the exploitation of factor premia decreases, the magnitude of these premia likely decreases as well. Assuming these premia are uncorrelated to risk, a decrease in methodological barriers - especially the decrease in transaction costs - would decrease market distortion which positively affects market efficiency. This is another example of index funds contributing to the process of price discovery and thereby the efficiency of market prices.

3.5 Passive investing and the low-volatility premium

In Part II, it is discussed that the existence of the low volatility premium can be explained by the incentive structures of institutional money managers and the biases of individual investors. Therefore, the existence of the low-volatility premium was explained by the actions of active market participants. However, as discussed in Part III, the amount of active capital allocation has drastically decreased as a result of the emergence of (passive) index funds. This could, in theory, decrease the influence of active market participants on stock prices. The low-volatility premium, at least partly explained by actions of active participants, could therefore decrease as a result of the shift towards passive investing. Although recent studies by Vanguard and Blackrock indicate that the daily trading volume - most important for price discovery - is still dominated by actively managed capital, there are several reasons to believe that the emergence of index funds, and their growing assets under management, have positively influenced market efficiency.

It is worth noting that the reasoning presented is one-sided. The discussed development of the market structure may well have a beneficial impact on market efficiency, thereby reducing the low volatility premium. Consequently, it is reasonable to assess whether there has been a significant change in the premium over time. However, if it is discovered that the low volatility premium has in fact decreased, this does not necessarily confirm that the trend towards passive investment is the root cause. Conversely, if the premium has not diminished, it does not necessarily invalidate the rationale presented.

Part IV

Research Design

4.1 Research Description

The aim of this research paper is to thoroughly analyze the driving force behind the low-volatility effect and the corresponding implications for the longevity of the premium. By examining to what extent the driving force is robust to changes in the market structure, one can hypothesize about the development of the low-volatility effect - or decomposed variants of this effect - over time. In part II, empirical results from different research papers were bundled to expose what we believe to be the main drivers behind the low volatility premium. The incentive structure of institutional money managers in combination with leverage constraints mainly influences the beta effect while the behavioral biases of individual investors mainly influence the idiosyncratic volatility effect. In both components of the low-volatility effect, investors are drawn towards stocks characterized by highly volatile payoffs while stocks with low volatility payoffs are neglected. As a result low volatile stocks become attractively priced from a relative perspective which increases the future stock performance. Therefore, the existence of the low-volatility premium was primarily explained by the way in which active market participants allocate their capital. However, as discussed in Part III, the amount of active capital allocation has drastically decreased as a result of the emergence of (passive) index funds. Since we believe the low-beta effect as well as the low-IVOL effect are driven by (respectively) incentives and biases that occur in active capital allocation, this shift towards passive investing could be a potential threat to the future returns of low-volatility strategies as it positively influences market efficiency. Therefore, this research paper will examine to what extent the low-volatility effect - or decomposed variants of this effect - has been robust to changes in the market structure.

Research question: *Is the low-volatility premium significantly present since the emergence of exchange traded funds and the corresponding shift towards passive investing (2005-2022)?*

4.2 Hypotheses

According to the discussed literature, portfolios consisting of low-volatility stocks realize higher risk-adjusted returns than portfolios consisting of high-volatility stocks in the long run. Consistent with the literature, we expect to find a similar negative relationship between the past risk-adjusted returns of US stocks within the sample and their exposure to return volatility. Therefore, the following hypothesis will be the starting point of this research:

H1: *The long term alpha of low-volatility stock portfolios exceeds the long term alpha of high-volatility stock portfolios.* $[\alpha_{low} > \alpha_{high}]$

According to economic theory, as discussed in Part I, the magnitude of the low volatility premium should decrease over time due to arbitrage opportunities. After all, mispriced securities should, according to the efficient market hypothesis, be exploited by arbitrageurs who thereby converge stock prices and intrinsic value. However, as discussed in Part II, incentive structures and personal biases ensure that investors - individuals as well as institutional investors - are drawn towards stocks characterized by highly volatile payoffs while stocks with low volatile payoffs are neglected. This was used to explain the long term persistence of the low volatility premium. As discussed in Part III, the amount of active capital allocation has drastically decreased as a result of the emergence of (passive) index funds. Because the low-volatility premium was explained by incentives and biases that occur in active portfolio management, this shift towards passive investing could be a potential threat to the volatility premium. As explained, the shift towards passive investing could reduce the number of unskilled market participants, the cost of short selling and the methodological barriers for the exploitation of factor premia. Building on this reasoning, we believe the shift towards passive investing has increased market efficiency and therefore decreased the low volatility premium. Therefore the second hypothesis is as follows:

H2: *The magnitude of the low-volatility premium has decreased since the emergence of exchange traded funds and the corresponding shift towards passive investing. (2005-2022)*

4.3 Methodology

The low-volatility premium

In order to test the first hypothesis, we will compare the performance of low-volatility portfolios with the performance of high-volatility portfolios. As previously discussed, there are several methodological approaches to assessing the volatility of a given stock. To avoid biasing our results by selecting a specific measure of volatility, and to account for the fact that different measures can capture distinct aspects of volatility (see 2.4), we will evaluate volatility using three distinct methods: exposure to systematic risk (beta), variance of returns and idiosyncratic volatility. For all three measures, we want to analyze the relationship between volatility and returns. In order to expose this relationship, we will construct portfolios with different levels of volatility exposure. To construct these portfolios, all NYSE, AMEX, and NASDAQ listed stocks will be ranked based on beta, variance or idiosyncratic volatility. This depends, of course, on the applied measure of volatility. After the stocks are ranked based on volatility, they are divided into volatility quintiles. The 20% of stocks with the lowest volatility belong to the first quantile (1) and the 20% of stocks with the highest volatility belong to the fifth quantile (5). This construction process is repeated at the end of every month. In the case of beta, portfolios are constructed at the end of each month using the daily volatility of returns in the previous 60 months. In the case of return variance, portfolios are constructed at the end of each month using the total return variance of the 60 preceding days. In the case of idiosyncratic volatility, portfolios are constructed at the end of each month using the total residual return variance of the 60 preceding days. The residual return variance is calculated using the Fama & French 3-factor model. To be more specific, the idiosyncratic volatility is equal to the standard deviation of the Fama-French 3-factor residual (ε) term.

At this point, we have explained the portfolio construction process for all three measures of volatility. By tracking the returns of the constructed portfolios we can analyze the performance for all fifteen (5x3) volatility quintiles. This allows us to compare the performance of low-volatility quintiles to the performance of high-volatility quintiles. The performance of different quintiles will be assessed using the geometric average of monthly returns. Risk will be measured by the standard deviation of returns and the portfolio beta. Ultimately, (Jensen's) alpha and the Sharpe ratio will be calculated for each quantile to determine whether a significant risk adjusted outperformance was realized. The alpha and beta coefficients of a portfolio are calculated using a linear regression. The returns of a portfolio are regressed against the returns of the value weighted market portfolio. In order to test the significance of these coefficients a t-test is used. This test calculates the probability that the 'real' value of the coefficient is equal to zero. If this probability is small (<0,05) we reject this null hypothesis, meaning the coefficient is considered statistically significant.

In order to determine the significance of the disparity in returns between low-volatility and high-volatility portfolios, we have used parametric (t-test) as well as nonparametric tests (Wilcoxon Rank Sum), depending on whether the sample data is normally distributed. We have tested normality using the Shapiro–Wilk test. In some cases, we have used parametric tests even when the data is not normally distributed. This is motivated based on the findings of Lumley et al. (2002). We discuss this in more detail when applied.

The low-volatility premium over time

In order to test the second hypothesis, the performance of low-volatility portfolios will be compared to the performance of high-volatility portfolios during different periods in time. The portfolio construction process does not deviate from the aforementioned method. The first step in testing whether the low-volatility premium has declined since the shift towards passive investing is by adjusting the analyzed period. In other words, we can then test whether the difference in performance between low-volatility and high-volatility has changed in a more recent period. For this recent period, we have chosen the period 2005-2022. The reason we decided to use this - somewhat arbitrary - timeframe stems from the earlier cited Vanguard study (*Setting the record straight: Truths about indexing*). According to this study, the percentage of assets in equity funds invested in index funds reached 20% for the first time in 2005. Although the emergence of index funds began around 1993, the potential impact on the low-volatility premium was neglectable as a result of the low level of initial adaptation (<5%). Therefore, as we wanted to select a period in which the potential impact of passive investing was present, we set 20% of assets in equity funds invested in index funds as the starting point. This corresponds to the year 2005.

By testing the first hypothesis, we have gained a clear picture of the long-term performance at different levels of exposure to volatility. The performance disparity between low volatility and high volatility can therefore be considered as the long-term low volatility premium. The next step is to analyze the low-volatility premium in the 2005-2022 period. If we have established the recent low-volatility premium and compared it to the long term premium, we can analyze the development of the volatility premium in greater detail. This will be done using a rolling alpha method. The alpha of the low-volatility portfolios (low-beta, low-variance and low-IVOL) and the high-volatility portfolios (high-beta, high-variance and high-IVOL) will be calculated at the end of each month, using the past 12 months of performance data. Subsequently, we take the average alpha of the three low-volatility portfolios to calculate a general low-volatility rolling alpha. The same method is used for the high-volatility rolling alpha. Then, at the end of every month, we subtract the high-volatility rolling alpha from the low-volatility rolling alpha to calculate the *Alpha-spread*. This spread allows us to analyze the low-volatility premium over time. Since the Alpha spread is normally distributed (see 5.5) we use a two-sample t-test to test whether the alpha spread has significantly changed since the emergence of passive investing. In other words, we test whether the alpha spread sample means significantly differ during different periods of time.

Also, we will test whether the absolute performance difference between low-volatility and high-volatility has significantly changed over time. Therefore, we test whether we can find evidence for a difference in differences. Similar to the alpha spread analysis, we calculate the performance difference by taking the average performance difference of IVOL, VAR and Beta-sorted portfolios. To be concise, *Performance gap = AVG(Q1V-Q5V, Q1R-Q5R, Q1B-Q5B).* Since the performance differences are not normally distributed (see 5.6), we have utilized the Wilcoxon Rank Sum to test for significant differences in the sample mean.

Calculating Alpha

So far, we have explained that Alpha serves as the ultimate measure of risk-adjusted performance. In the process of calculating alpha, we are searching for returns that are not explained by the coefficients included within our asset pricing model (see section 2-5 of Part I). This raises the question of what coefficients we have incorporated within our asset pricing model to facilitate the calculation of Alpha. Given our examination of the established factor premia of small-cap stocks (SMB) and value stocks (HML), it would be obvious to include them within our model. However, apart from constructing portfolios based on idiosyncratic volatility, for which we have utilized the Fama-French 3-factor model, we have made the decision to not use SMB and HML as coefficients in our asset pricing model. The rationale for this decision stems from research done by Falkenstein (2012), who concludes that the factor loadings for different levels of volatility exposure do not account for the observed return disparity. Based on the SMB and HML factor loading for portfolios ranked on past volatility, the author finds that: '*If the coefficients on value, the market and size are driving expected returns* [of low volatility stocks]*, the only part of that story working is the value/low volatility connection; it does not work for size/low-volatility, nor is it symmetric with the value/high-volatility returns.'* Furthermore, we have not found compelling and consistent evidence for other (risk)factors that explain the continued existence of the low-volatility premium. Therefore, we have chosen to utilize a standard Capital Asset Pricing Model (CAPM) with systemic risk as our sole explanatory variable in determining Alpha.

4.4 Data

The sample data consists of all NYSE, AMEX, and NASDAQ listed stocks over the period 1963-2022. The daily, monthly and yearly price data, which are used to calculate periodic returns, are obtained from the CRSP database. The used volatility measures are also calculated with the daily price data obtained from the CRSP database. Furthermore, the

Kenneth French database is used to obtain the returns of the market portfolio, the risk-free rate and the returns of several cross sectional factors. This data is used to calculate the risk-adjusted outperformance of stocks and their corresponding idiosyncratic volatility. A value-weighted approach has been chosen for both the market portfolio and the constructed quintile portfolios. This mechanism ensures that the allocation of certain stock within a portfolio is positively related to the market capitalization of the firm.

Part V
Results & Conclusion

5.1 Performance of low-beta

The long term performance of low volatility portfolios compared to high volatility portfolios has been the starting point for this research paper. All NYSE, AMEX, and NASDAQ listed stocks were sorted into quintiles based on their past beta, variance and residual variance. First, in order to test hypothesis I, the long term performance of the different volatility-sorted quintiles will be analyzed. Then, in order to test the second hypothesis, it will be analyzed whether the performance of the different volatility-sorted quintiles has changed over-time.

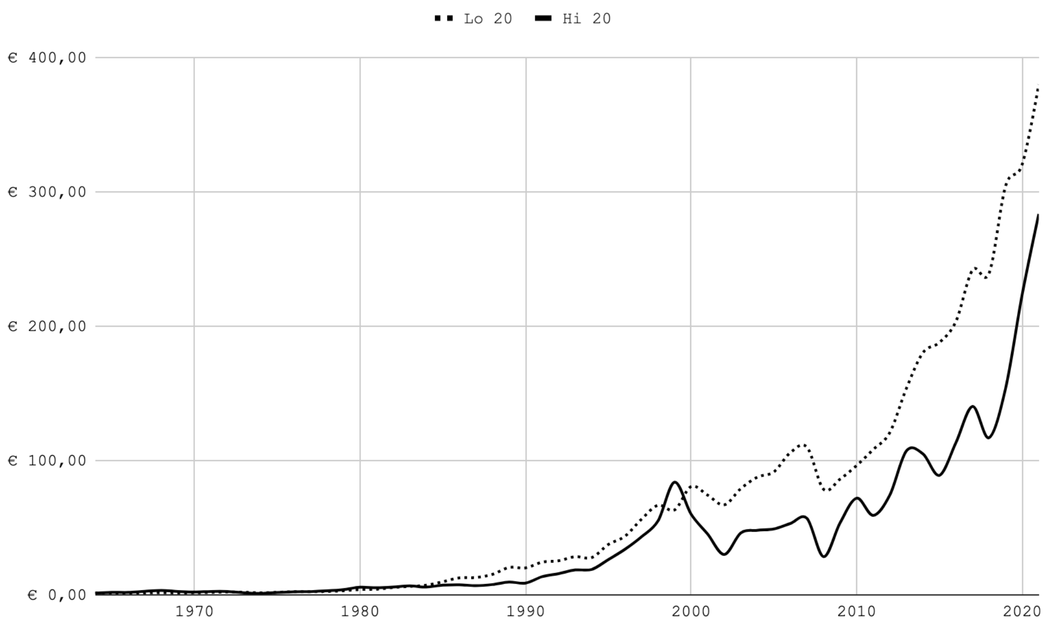

Figure 1: Performance of Lo 20 and Hi 20, beta sorted (1963-2022)
Notes: Lo 20 (Hi 20) shows the performance of a portfolio constructed by selecting the 20% stocks with the lowest (highest) beta at the start of each year. Beta's are calculated using the preceding five years of past monthly returns. Initial investment is €1. Both portfolio's are long only and unlevered.

As shown in Figure 1, the low-beta portfolio has outperformed the high-beta portfolio in the period 1963-2022. An Euro invested in the low-beta portfolio grew to € 379,68 while an Euro invested in the high-beta portfolio grew to € 283,24. The yearly geometrical return of the low-beta portfolio was 10,783% while the yearly geometrical return of the high-beta portfolio was 10,225%. Therefore, on a geometrical basis, the low-beta portfolio has outperformed the high-beta portfolio. Interestingly, on an arithmetic basis the low-beta portfolio has grown 11,769% annually while the high-beta portfolio grew 14,112%. Therefore, if an investor wants to achieve maximum short term performance, the high-beta portfolio offers a better perspective. However, in the long run, taking into account the asymmetric impact of losses on the growth rate of capital, the low-beta portfolio offers a better perspective. This tradeoff between long term and short term performance fits the agency problem discussed earlier. It could potentially explain why actively managed capital - judged / rewarded based on short term performance - is drawn towards volatile equities.

Table 1: Summary statistics of 5 portfolios sorted on past beta (1963-2022)

Portfolio	Beta	Alpha	T	P>|T|	Excess return	Std. Dev	Sharpe
Q1 - Lo	0,671	0,156	2,270	0,023*	0,472	3,51	0,134
Q2	0,908	0,120	2,500	0,013*	0,540	4,26	0,127
Q3	1,055	0,040	0,800	0,422	0,513	4,90	0,105
Q4	1,213	0,002	0,040	0,971	0,523	5,66	0,092
Q5 - Hi	1,489	-0,158	-1,480	0,138	0,416	7,23	0,058
Market	1	-	-	-	0,462	4,48	0,103

Notes: Table 1 shows the summary statistics of 5 different portfolios sorted on their past beta. To calculate the alpha and beta for a portfolio, the returns of the portfolio are regressed against the returns of the value weighted market portfolio. The T-statistic is used to determine the significance of the alpha's. The excess return is the monthly geometric portfolio return minus the risk free rate.

As displayed in Table 1, we observe a negative relationship between the degree of systemic risk exposure and alpha. The two least volatile portfolios from a beta perspective, Q1 and Q2, are the only portfolio's with a positive significant alpha. As we 'move' from the low-beta quintile towards higher volatility quintiles, alpha (although insignificant) structurally decreases. The most volatile quintile (Q5) has realized the lowest risk adjusted excess return and is the only portfolio with a negative alpha. Furthermore, looking at Sharpe ratios, the low-beta portfolio has again outperformed its volatile peers. As with the alpha, we observe a negative relationship between the degree of exposure to systematic risk and the Sharpe ratio. The negative relationship between beta and Sharpe ratio is displayed in Figure 3 (page 35).

In order to determine the significance of the disparity in returns between the low beta portfolio and the high beta portfolio, we have used a parametric (t-test) as well as a nonparametric test (Wilcoxon Rank Sum). We have tested whether the portfolio returns were normally distributed using the Shapiro–Wilk test. Based on the results of the SW test, we reject the null hypothesis that the return data of the beta sorted portfolios is normally distributed. Normally, this would imply that the use of parametric tests is not valid. Or at least, this is the common approach. However, according to Lumley et al. (2002) '*the t-test and least-squares linear regression do not require any assumption of Normal distribution in sufficiently large samples*'. The authors find that for large sample sizes of (extreme) non-normally distributed data, the t-test is still viable due to the *Central Limit Theorem*. According to this theory, the distribution of sample means converges to a normal distribution if the sample size becomes large enough. Since we do not visually observe extreme non-normality, skewness or kurtosis in the return distribution of the beta sorted portfolios, and we have a large number of observations (700+) we feel confident that the results of a parametric mean comparison test are valid. To increase the robustness of the results, we will also use a non-parametric test. This test does not assume normality of the return distribution.

Based on the Capital Asset Pricing Model, it would be expected that the mean return of the Q1 portfolio would be significantly lower than that of the Q5 portfolio, thus resulting in a positive mean difference. Therefore, to be concise, we would expect the mean difference is greater than zero (Mean: Return Q5 - Return Q1 > 0). However, upon conducting the t-test, the null hypothesis that the means of Q1 and Q5 are equivalent cannot be rejected (H0: mean(Q5-Q1) != 0, P = 0.7362). The results of the Wilcoxon test are equal. The test indicates that there is not enough statistical evidence to prove a significant difference in distributions (H0: Return(Q1B) = Return(Q5B), P = 0.3711). These findings align with our previous observation, as outlined in Table 1, that the substantial alpha of the low beta portfolio can be attributed primarily to the difference in risk exposure. The returns are not significantly different, but so is the exposure to risk. Despite this, it should be noted - as we did earlier - that the use of arithmetic means in our comparison of the sample distributions can be troublesome - particularly in situations where outliers, also known as *tail events,* heavily influence the outcome. An investor realizing a mean return of 10% could be better off than an investor realizing a mean return of 15%. Therefore, when analyzing the long-term performance of a portfolio, it is crucial to consider the difference in geometric means, as the risk-adjusted growth rate of capital is ‘what matters in the end’ (Spitznagel, 2021). Despite this limitation, the results of both tests do favor the low volatility portfolio when accounting for risk exposure. Also, building on the discussed limitation of arithmetic means in comparison to geometric means, we would expect that *if* the test results were indeed biased, it would be in favor of the high beta portfolio. After all, the low beta portfolio had lower arithmetic returns but higher geometrical returns. Taking into account this potential bias towards high beta portfolios, the results become even more compelling. The limitations of the applied methodology will be discussed in more detail in the Discussion.

To summarize, contrary to what we would expect based on the capital asset pricing model, the low-beta portfolio has, between 1963 and 2022, realized a higher growth rate of capital than the high-beta portfolio. Also, the low-beta portfolio has realized a positive significant alpha whereas the high-beta portfolio has realized a negative but insignificant alpha. These results are in line with the findings of Van Vliet and Blitz (2007) who find, by ranking stocks based on their past return volatility, that portfolios of stocks with the lowest volatility are '*associated with Sharpe ratio improvements*'. We observe a similar negative relationship between beta and the Sharpe ratio. According to the parametric T-test and the non-parametric Wilcoxon Rank Sum test, the arithmetic returns of the low beta portfolio do not significantly differ from the returns of the high beta portfolio. Taking into account the difference in risk exposure, this is more evidence in favor of the low volatility premium.

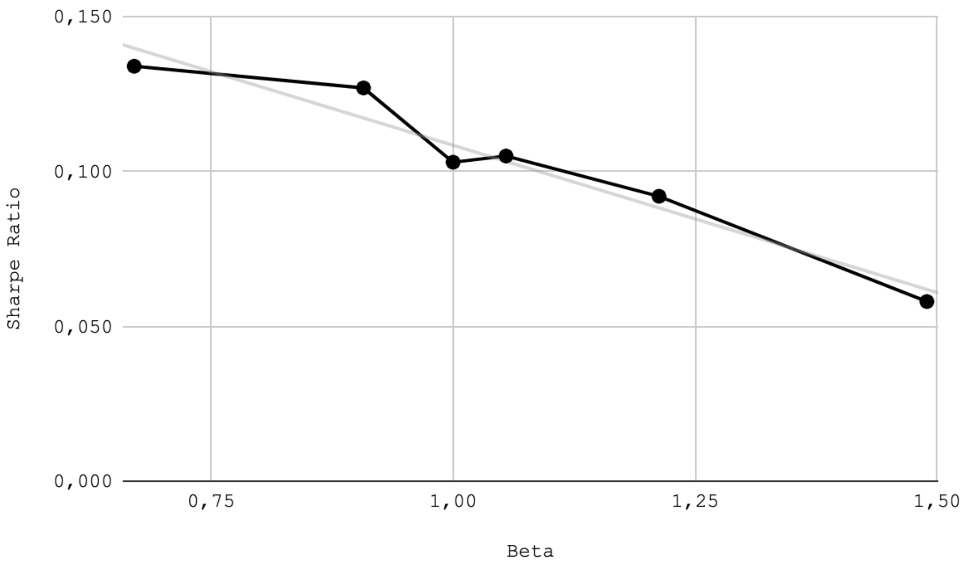

Figure 3: Relationship between portfolio beta and Sharpe ratio
Notes: Figure 3 shows the relationship between portfolio beta and Sharpe ratio. The dots represent the five different volatility quintiles and the value weighted market portfolio. The Sharpe ratio is calculated by dividing the excess returns by the standard deviation of the excess returns.

5.2 Robustness analysis

As mentioned before, there are several methodological methods to measure the performance of low volatility portfolios. To ensure that the negative relationship between systemic risk exposure and performance, as presented in 5.1, is robust to changes in the methodological approach, it will be analyzed whether this relationship is also present in variance-sorted and IVOL-sorted portfolios. The results are presented in Table 2 and Table 3 on the next page.

Table 2: Summary statistics of 5 portfolios sorted on past IVOL (1963-2022)

Portfolio	Beta	Alpha	T	P>|T|	Excess return	Std. Dev	Sharpe
Q1 - Lo	0,790	0,161	3,17	0,002**	0,534%	3,78	0,141
Q2	0,971	0,064	1,37	0,171	0,507%	4,52	0,112
Q3	1,125	0,073	1,40	0,162	0,568%	5,22	0,109
Q4	1,313	0,041	0,53	0,595	0,584%	6,21	0,094
Q5 - Hi	1,548	-0,561	-3,85	0,003**	-0,008%	7,93	-0,001
Market	1	-	-	-	0,462	4,48	0,103

Notes: Table 2 shows the summary statistics of 5 different portfolios sorted on their idiosyncratic volatility. To calculate the alpha and beta for a portfolio, the returns of the portfolio are regressed against the returns of the value weighted market portfolio. The excess return is the monthly geometric portfolio return minus the risk free rate.

Table 3: Summary statistics of 5 portfolios sorted on past Variance (1963-2022)

Portfolio	Beta	Alpha	T	P>|T|	Excess return	Std. Dev	Sharpe
Q1 - Lo	0,697	0,196	3,25	0,001***	0,526%	3,50	0,150
Q2	0,934	0,526	1,58	0,114	0,511%	4,41	0,116
Q3	1,100	0,056	1,29	0,196	0,558%	5,14	0,108
Q4	1,334	0,028	0,36	0,716	0,578%	6,30	0,092
Q5 - Hi	1,583	-0,576	-4,01	0,002**	-0,014%	8,04	-0,002
Market	1	-	-	-	0,462	4,48	0,103

Notes: Table 3 shows the summary statistics of 5 different portfolios sorted on their variance. To calculate the alpha and beta for a portfolio, the returns of the portfolio are regressed against the returns of the value weighted market portfolio. The T-statistic is used to determine the significance of the alpha's. The excess return is the monthly geometric portfolio return minus the risk free rate.

As displayed in Table 2 and Table 3, when sorting stocks on past idiosyncratic volatility and past variance, the least volatile portfolios (Q1-IVOL and Q1-VAR) have realized a significant positive alpha whereas the most volatile portfolios (Q5-IVOL and Q5-VAR) have realized a significant negative alpha. As with the beta sorted portfolios, we observe a similar negative relationship between the exposure to the volatility measures (IVOL or VAR) and the risk-adjusted excess returns of a portfolio. In the Q1-Q4 range, increasing the exposure to IVOL and variance can increase the absolute excess returns. However, increasing the exposure to IVOL or variance is in no case beneficial for the risk adjusted excess returns. After all, the Lo20 portfolio's (Beta, Variance, IVOL) have realized the highest Sharpe ratios and the highest significant alphas. Therefore, Table 2 and Table 3 show that the low-volatility effect is robust to changes in the methodological approach.

Again, in order to determine the significance of the disparity in returns between the low volatility portfolios and the high volatility portfolios, we have used a t-test as well as a Wilcoxon Rank Sum test. Based on the results of the t-test, we cannot - at the 0.05 significance level - reject the null hypothesis that the means of Q1 and Q5 are equivalent for both measures of volatility. However, at the 10% significant level, the Q5-Q1 mean is significantly below zero for both measures (Ha: Mean VAR < 0, $Pr(T > t) = 0.0998$) (Ha: Mean IVOL < 0, $Pr(T > t) = 0.0858$). The opposite of what we would expect based on the capital asset pricing model. Therefore, based on the parametric tests, we find a more significant return disparity between low volatility portfolios and high volatility portfolios when the portfolio construction is based on variance and/or residual variance. This aligns with the results presented in Table 1, Table 2 and Table 3. The alpha disparity between low volatility and high volatility portfolios is higher and more significantly present with IVOL- and VAR-sorted portfolios in comparison to the Beta sorted portfolio. In the case of VAR- and IVOL-sorted portfolios, the results of the Wilcoxon test are similar to the results discussed earlier with the Beta-sorted portfolios. The test indicates that there is not enough statistical evidence to prove a significant difference in mean between the low-volatility and the high-volatility portfolios (Q1 vs. Q5). It is important to reiterate that the finding of no significant deviation in the average returns of low-volatility and high-volatility portfolios, owing to the difference in risk exposure, in essence, serves as evidence in favor of the low volatility premium. Obtaining similar returns in combination with lower systematic risk exposure ultimately results in a positive alpha. Based on the results presented in Table 1, Table 2 and Table 3, and the results of the mean-comparison tests, we conclude that taking on more risk does not have a significant impact on the absolute return.

To summarize, in the period 1963-2022, portfolios of low volatility stocks have realized a higher growth rate of capital than high-beta portfolios while being exposed to less systemic risk. As a result, the Low-volatility portfolios; Lo20-Beta, Lo20-Var and Lo20-IVOL have each realized a positive significant alpha whereas their volatile counterparts; Hi20-Beta, Hi20-Var and Hi20-IVOL have realized a significant negative alpha. The observed alpha disparity between low-volatility portfolios and high volatility portfolios is supported by the results of the mean comparison tests (t-test and Wilcoxon Rank Sum). Based on the Wilcoxon test, we do not find a significant difference in mean between the low-volatility and the high-volatility portfolios. Based on the t-test, we also find, at the 5% significance level, no significant difference in mean between the returns of low volatility portfolios and high volatility portfolios. However, in the case of IVOL and VAR sorted portfolios, we find, at the 10% significance level, evidence for a significant negative mean (Q5-Q1). Based on these findings, we reject the notion that increasing the amount of volatility exposure has a positive significant effect on the arithmetic returns of the portfolio. Taking into account the difference in risk exposure, this explains the observed alpha disparity between low and high volatility portfolios. Based on these findings, we do not reject the first hypothesis.

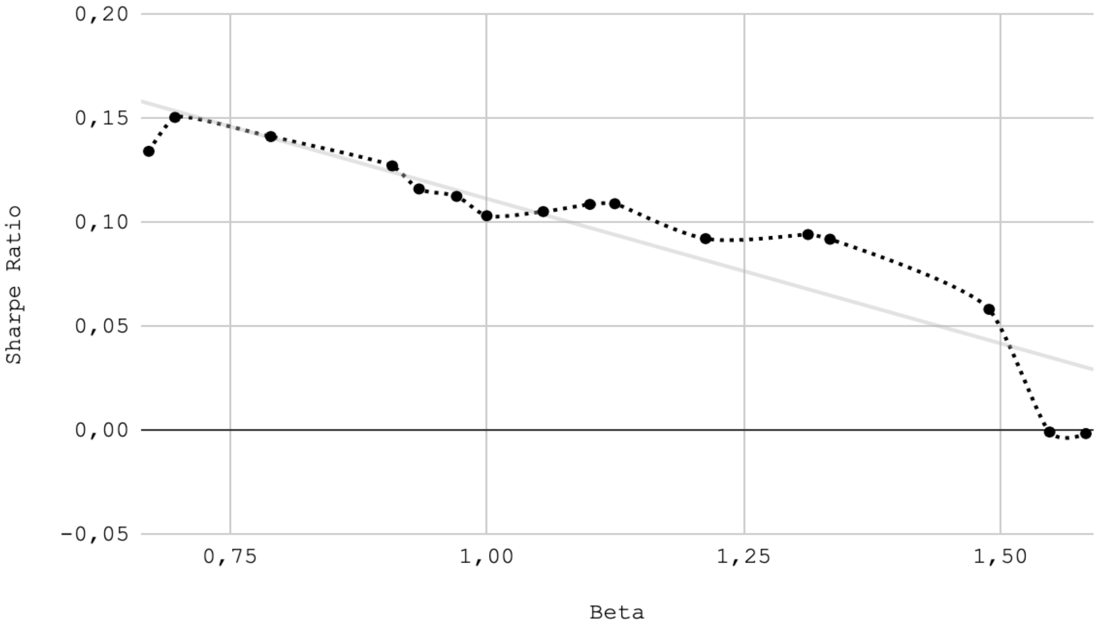

Figure 4: Relationship between portfolio beta and Sharpe ratio
Notes: Figure 4 shows the relationship between portfolio beta and Sharpe ratio. The dots represent the 15 different portfolios (3 times 5 quintiles) and the value weighted market portfolio. The Sharpe ratio is calculated by dividing the excess returns by the standard deviation of the excess returns.

5.3 Nuance

Figure 4 shows the beta and Sharpe ratio of the 15 different portfolios and the value weighted market portfolio. The figure symbolizes the global picture that emerges from back-testing volatility-sorted portfolios. Increasing the amount of exposure to volatility decreases the Sharpe ratio of a portfolio. However, this picture deserves some nuance.

The overall relationship between volatility and excess returns, and thus the low-volatility premium, is largely influenced by the diverging results of the most 'extreme' portfolios. The least volatile portfolios significantly outperform the market portfolio while the most volatile portfolios significantly underperform the market portfolio. Hence, the low-volatility premium. However, when analyzing the intermediate portfolios (Q2, Q3 and Q4), a less strong relationship emerges. Within this range, increasing the exposure to volatility does in fact increase the long run geometric return. Adjusting for risk, there remains a slightly negative relationship between volatility and excess return. Therefore, it is important to mention that for the largest part of the sample, increasing the volatility exposure will only decrease the risk-adjusted returns and not the absolute returns, as is the case when the 20% least volatile stocks are compared to the 20% most volatile stocks. Although not proportional, exposure volatility does increase returns within the second to fourth decile.

5.4 Low-volatility premium over time

In the previous passage, we have presented evidence for the existence of a long-term low-volatility premium. However, the shift from active to passive investing has, as discussed in part III, the potential to diminish or even eliminate this premium. Therefore, as most of the sample data (70%) is from before this shift took place, the observed relationship between volatility and returns could be based on the large majority of observations not impacted by this shift. In other words, *If* the shift towards passive investing has indeed resulted in a diminished premium, this effect can be 'buried' by a large amount of observations from before this period. We are, after all, describing an effect based on the average of a large number of observations. Consequently, we shall now analyze the performance of low-volatility stocks in a more recent time frame and compare this to our long term observations. The goal is to analyze discrepancies between the recent data and the observed long term trend.

As presented in Table 4, the Low-IVOL and Low-variance portfolios have, between 2005 and (August) 2022, outperformed their volatile peers (fifth quintile) on an absolute and a risk-adjusted basis. The portfolios have higher excess returns, lower standard deviation of returns and therefore a higher Sharpe ratio. The observed Sharpe ratio improvement for less volatile portfolios is also present in the beta-sorted portfolio. However, in the case of a beta-sorted portfolio, the risk adjusted outperformance is caused only by a lower standard deviation as the excess return of the low-beta portfolio is lower than the excess return of the high-beta portfolio. Therefore, the absolute low-volatility premium was only present in the IVOL-sorted and variance-sorted portfolios. The realized alphas show a similar result. Only the least volatile VAR and IVOL portfolio have realized a significant positive alpha. The most volatile IVOL and Variance portfolio have realized a significant negative alpha. The beta-sorted portfolios did not realize a significant alpha. Although it is still too early to draw conclusions, it bears mentioning that, since two out of three volatility measures indicate a significant outperformance of low volatility portfolios and all approaches show Sharpe ratio improvements, thus far, the evidence has predominantly been in favor of the persistence of the low volatility premium. In other words, the observed performance disparity between low-volatility and high-volatility portfolios, in the period 2005-2022, is very similar to the long term performance disparity (1963-2022).

Table 4: Summary statistics of 15 portfolios sorted on past volatility (2005-2022)

IVOL	Beta	Alpha	T	P>|T|	Excess return	Std. Dev	Sharpe
Q1	0,816	0,211	2,910	0,004**	0,76%	3,79	0,201
Q2	1,009	0,006	0,090	0,925	0,67%	4,61	0,146
Q3	1,162	-0,195	-2,090	0,038	0,55%	5,36	0,102
Q4	1,349	0,000	0,000	0,999	0,83%	6,42	0,129
Q5	1,570	-0,665	-2,480	0,014**	0,22%	7,99	0,027
VARIANCE	Beta	Alpha	T	P>|T|	Excess return	Std. Dev	Sharpe
Q1	0,724	0,314	3,470	0,001***	0,81%	3,48	0,231
Q2	1,021	-0,020	-0,310	0,759	0,65%	4,65	0,140
Q3	1,168	-0,103	-0,950	0,344	0,64%	5,44	0,118
Q4	1,429	-0,186	-1,180	0,241	0,68%	6,78	0,100
Q5	1,608	-0,753	-2,800	0,006**	0,14%	8,15	0,017
BETA	Beta	Alpha	T	P>|T|	Excess return	Std. Dev	Sharpe
Q1 - Lo	0,663	0,109	1,090	0,277	0,56%	3,29	0,171
Q2	0,960	0,099	1,670	0,096	0,74%	4,38	0,168
Q3	1,105	-0,016	-0,190	0,851	0,70%	5,09	0,137
Q4	1,284	0,046	0,390	0,695	0,85%	5,98	0,142
Q5 - Hi	1,500	-0,233	-1,170	0,244	0,65%	7,28	0,089
Market	-	-	-	-	0,66%	4,47	0,148

Table 4 shows the summary statistics of 15 different portfolios sorted on past volatility measures. To calculate the alpha and beta for a portfolio, the returns of the portfolio are regressed against the returns of the value weighted market portfolio. The T-statistic is used to determine the significance of the alpha's. The excess return is the monthly geometric portfolio return minus the risk free rate.

Again, in order to determine the significance of the disparity in returns between the low volatility portfolios and the high volatility portfolios, we have used a t-test as well as a Wilcoxon Rank Sum test. Based on the results of the t-test, we cannot - at the 0.05 significance level - reject the null hypothesis that the means of Q1 and Q5 are equivalent for all three measures of volatility. The results of the Wilcoxon test are equal. The tests indicate that there is not enough statistical evidence to prove a significant difference in mean returns for all three measures of volatility. Within the sample period 2005-2022 we cannot find evidence for a significant difference in arithmetic mean returns. Taking into account the difference in risk exposure, this explains the observed alpha disparity between low and high volatility portfolios.

So far, we have presented empirical evidence for the existence of the low-volatility premium in the 1963-2022 period and the more recent 2005-2022 period. Based on these results, we can already conclude that the premium has not been eliminated as of recent. However, the question remains whether the premium has significantly declined in the recent past compared to the long term trend. In other words, we still have to analyze whether the magnitude of the return disparity between low-volatility and high-volatility portfolios has changed over time.

5.5 Alpha-spread development over time

In 5.1, 5.2 and 5.4, we have compared the performance of low volatility portfolios to the performance of high volatility portfolios in a manner that can be characterized as static. For two different sample periods, the Alpha and Sharpe ratio for different levels of volatility exposure were measured in order to draw conclusions about the relationship between volatility exposure and return. While this approach is straightforward and easily interpretable, it may be inadequate for drawing conclusions about the evolution of the volatility premium over time. As mentioned, based on our presented findings we cannot prove that the magnitude of the return disparity between low-volatility and high-volatility portfolios has changed over time. We can only conclude that the premium has not been eliminated in the recent sample period. To gain a deeper - more robust - understanding of the evolution of the low-volatility premium, we introduce a measure of analysis to study the development of the premium over time. The low volatility premium, which reflects the risk-adjusted performance gap between low and high volatility portfolios, can be studied through the spread in realized alpha. This spread captures the difference in risk-adjusted performance between low and high volatility portfolios. By the use of a rolling alpha method, where the alpha of a portfolio is calculated using data from the previous twelve months, we can track the development of alpha spread - and thus the volatility premium - over time. In this process, we have decided to use all three volatility measures; beta, variance and idiosyncratic volatility, to calculate the rolling alpha and then take the average to determine a ‘general’ low volatility and high volatility alpha. By doing this over a long period of time, we can identify trends in the development of the low volatility premium. This allows us to determine whether the premium has deviated from its long-term trend since the emergence of passive investing. The long term alpha spread development is presented in Figure 5.

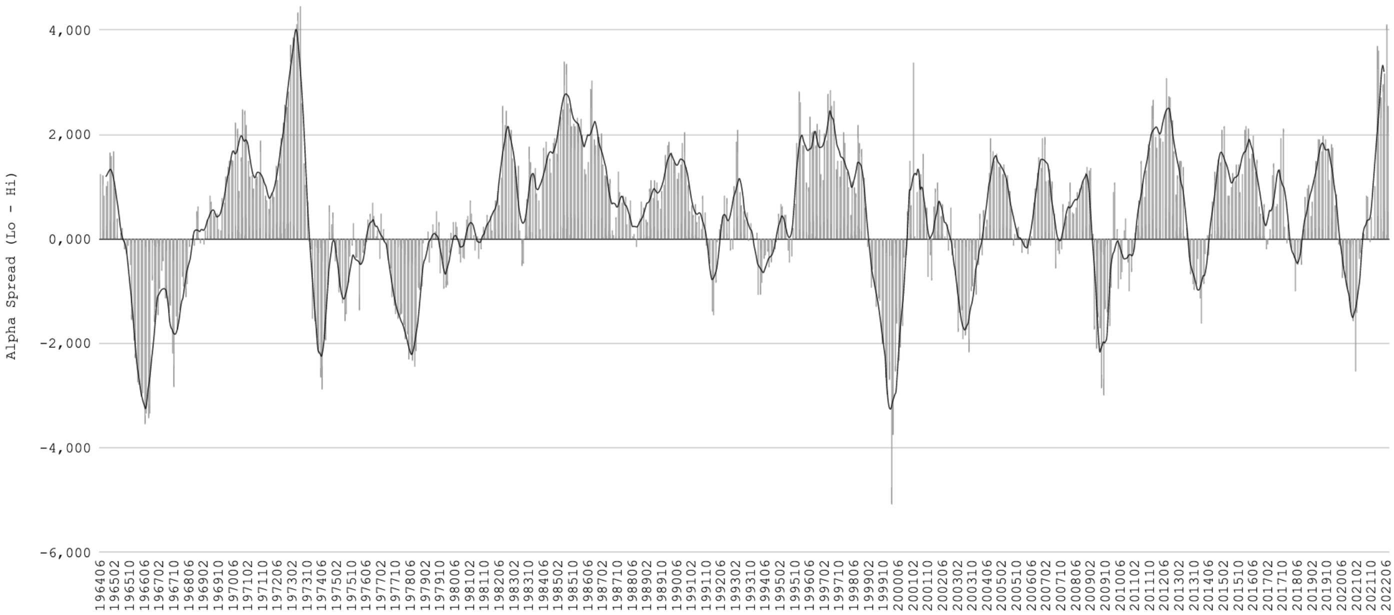

Figure 5: Rolling Alpha spread low volatility - high volatility (1963-2022)
Notes: Figure 5 shows the absolute difference between the rolling alpha of a low volatility strategy and a high volatility strategy. The rolling alpha of the low volatility strategy is calculated by taking the average of the alphas realized by the low-beta, low-variance and low-IVOL portfolio in the previous 12 months. The same method is used for the rolling alpha of the high volatility strategy. A positive spread indicates that the low volatility strategy outperformed the high volatility strategy.

As shown in Figure 5, the alpha spread between low volatility stocks and high volatility stocks fluctuates over time. Generally, periods of a positive Lo-Hi spread are followed by periods of a negative spread. This aligns with the findings of Blitz and Van Vliet (2007), who also find that there is no persistent positive volatility premium throughout history. In the period 1963-2022, we observed a (positive lo-hi) mean alpha spread of 0.448 with a 95% confidence interval of 0.347 - 0.549. This is in line with the earlier findings that indicated that, in the long run, low-volatility portfolios realize a significantly higher alpha than high-volatility portfolios.

Our first approach to determine whether the low-volatility premium has deviated from its long-term trend is by testing for a significant difference in alpha spread before and after the chosen ‘breakpoint’ (2005). We again have utilized the Shapiro-Wilk test to test whether the

Alpha spread is normally distributed. Based on the results of the SW test, we cannot reject the null hypothesis that the alpha spread is normally distributed (Prob>z: 0.412). Therefore, we have used a two-sample t-test to test whether the alpha spread has significantly changed since the emergence of passive investing. Based on the results of the two-sample t-test, we find - at the 5% significance level - evidence for an significant increase in the alpha spread in the period 2005-2022 compared to the previous period (1963-2004). Between 1963 and 2004, the mean alpha spread was 0.362 with a 95% confidence interval of 0.235 - 0.488. In the period 2005-2022, the mean alpha spread was 0.648 with a 95% confidence interval of 0.484 - 0.812. This is the opposite of what we would expect based on the theories discussed in Part II and Part III and the corresponding hypothesis.

A potential explanation for the significant increase in the alpha spread can be observed visually in Figure 5. In the period 1963-2004, we observed two distinct periods of large negative alpha spread indicating a risk-adjusted outperformance of high-volatility stocks. The timing of this negative alpha spread can be explained by an earlier observation of Falkenstein (2013) who found that high-volatility stocks perform well in a 'highly speculative market environment'. It is thus not entirely coincidental that we observe two outliers of the spread occurring precisely during the 'post-war bull market' and the 'dot-com bubble'. Since we do not observe a similar magnitude of speculative demand in the period 2005-2022, this could explain the significant difference in mean. In other words, if there is an absence of these periodical *waves* of 'speculation' in our sample period, this increases the mean alpha spread. As a result, the results of this analysis are sensitive to changes in the chosen sample period. For instance, if we compare the alpha spread in the periods 1963-1999 and 2000-2022 - so we include the dot-com bubble in our 'after' sample period - the statistical evidence for a significant difference disappears. In addition, it is possible, possibly even expected, that a similar period of 'speculation' will arise in the period after 2022, which may bring the average alpha spread down again. Therefore, we must not confuse absence of evidence for evidence of absence. This is also known as the *Inverse Turkey* problem. (Taleb, 2004) Taking into account the drawbacks of the analysis, we are cautious in interpreting the results. However, in line with the results presented in 5.4, we do not find any evidence for a decreasing volatility premium. Even *If* we change the sample period to include the latest period of large negative alpha spread, we still do not find evidence for a declining volatility premium.

5.6 Difference in difference analysis

An alternative way to analyze the development of the low-volatility premium is by testing whether the absolute performance difference between low-volatility and high-volatility has significantly changed over time. Therefore, we test whether we can find evidence for a difference in differences. Similar to the alpha spread analysis, we calculate the performance difference by taking the average performance difference of IVOL, VAR and Beta-sorted portfolios. To be concise, *Performance gap = AVG(Q1V-Q5V, Q1R-Q5R, Q1B-Q5B)*. Based on the Shapiro Wilk test, we reject the null hypothesis that the return differences between low-volatility portfolios and high-volatility portfolios is normally distributed. Therefore, we have utilized the Wilcoxon Rank Sum to test for significant differences in the sample mean. To perform this test, we created a grouping variable to identify a period *before* and *after* the chosen break point of 2005 (see Part IV for rationale). Based on the WRS test, we cannot - at the 5% significance level - reject the null hypothesis that *AVG (Period==After) = AVGD (Periode==Before).* We can not find statistical evidence for a significant difference in the - arithmetic - performance gap between low-volatility and high-volatility portfolios in the period before and after the shift towards passive investing.

A critical note that must be added to this method is that, in order to ensure its validity, it must be analyzed whether the difference in risk exposure between the low-volatility and high-volatility portfolios has changed over time. If the difference in risk exposure has changed, it logically follows that a difference in absolute performance would also be expected. At least, assuming Markowitz' portfolio theory (1952) is correct. However, we do not find evidence for an 'alarming' change. In the period 1963-2004, the average Q5-Q1 beta spread was 0,82. In the period 2005-2022, the average Q5-Q1 beta spread was 0,83. Therefore, the validity of this approach is not jeopardized by omitted beta spread.

…

6. Conclusion

Institutional investors as well as individual investors have an identifiable preference for stocks with highly volatile payoff structures. In the case of constrained institutional money managers, selecting stocks with a high beta enables them to maximize the outperformance for a given level of tracking error volatility. In the case of individual investors, biased probability weighting and the corresponding lottery preference explains the demand for stocks with highly volatile payoff structures. This preference for volatility explains the relatively low demand for low-volatility stocks and thus the low-volatility premium. However, since the magnitude of the low-volatility premium is calculated using historical performance data, the results are influenced by, among other things, the prevailing market structure. Substantial changes in the market structure, such as the shift towards passive investing - that could reduce the number of unskilled market participants, the cost of short selling and the methodological barriers for the exploitation of factor premia - could therefore influence the persistence of the low volatility premium. Also, it poses the vital question whether historical research findings on the aforementioned premium remain relevant in the face of fundamental changes to the market structure. Consequently, this research paper has evaluated to what extent the low-volatility effect has been resilient to changes in the market structure.

During the period from 1963 to 2022, portfolios of low-volatility stocks have outperformed portfolios of high-volatility stocks on a geometrical return basis. Furthermore, the low-volatility portfolios (Lo20-Beta, Lo20-Var, and Lo20-IVOL) have each realized a statistically significant positive alpha, while their volatile counterparts (Hi20-Beta, Hi20-Var, and Hi20-IVOL) have realized a significant negative alpha. Also, the arithmetic returns of the low-volatility portfolios did not significantly differ from the arithmetic returns of the high-volatility portfolios while being exposed to less systemic risk. Therefore, we have found a negative relationship between exposure to volatility and the realized Sharpe ratio.

In the period from 2005 to 2022, which we define as the recent past, the IVOL-effect and the Variance-effect were both present and statistically significant. The Beta-effect, however, was not statistically significant during this period. In both the long-term and the recent past, reducing exposure to systemic risk has been shown to increase a portfolio's Sharpe ratio. Our analysis using a rolling alpha method has revealed that the alpha spread between low-volatility and high-volatility portfolios fluctuates over time, in line with the findings of Blitz and Van Vliet (2007), who also found no persistent positive volatility premium throughout history. Furthermore, we do not find evidence for a declining trend in the alpha spread between low-volatility and high-volatility portfolios. In fact, we found evidence for a significant increase. However, as discussed, we are cautious in interpreting these results, since they are proven to be sensitive to changes in the chosen sample period. We have also analyzed the development of the absolute performance gap between low-volatility and high-volatility portfolios. We have not found evidence for a significant deviation in the performance gap in the period before and after the shift towards passive investing. Based on these findings we conclude that the low-volatility premium is significantly present over time and robust to changes in the methodological approach. Furthermore, the volatility premium has not significantly decreased in recent times.

It is not different this time.

7. Discussion

In this research paper, a lot of discussion points have been extensively covered in the text. Nevertheless, certain aspects have not been fully examined. Consequently, this section aims to address the outstanding issues that have not been adequately discussed.

We were wrong to believe that the shift towards passive investing has diminished the low-volatility premium. Although, in our view, the hypotheses were based on sound economic reasoning, we may have been naive to believe that we can understand market dynamics by (only) looking at the behavior of individual participants. According to Nassim Nicholas Taleb (2018), *'The psychological experiments on individuals showing biases do not allow us to understand aggregates or collective behavior, nor do they enlighten us about the behavior of groups.'* This is an obvious pitfall in the theoretical reasoning behind the motivation for this research paper. We have tried to explain cross sectional patterns in asset pricing based on the sum of biases and incentives of individuals. Then we rationalized, based on changes in the market structure, how this could lead to a diminished - or even eliminated - premium. However, we therefore implicitly assume that we can describe and predict market behavior based on the sum of individual behavior. This strong assumption, that we can understand the *macro* from the *micro*, may explain the difference between the hypothesized and the observed results. We may have a good understanding of the behavior of individuals, but it does not automatically mean that we can understand the market as a whole.

Another point of discussion is the methodology we applied to analyze the performance of the volatility-sorted portfolios. In the calculation of the alphas and the Sharpe ratios, which are key performance indicators, we control for the amount of risk exposure. Risk being equal to volatility. Although this is standard practice in modern finance, we do not necessarily agree with this rationale. In our view, using 'risk adjusted' returns as a performance indicator goes beyond the actual goal of investing; maximizing the long term growth rate of capital. This is, after all, what determines the actual ending value of wealth. Therefore, we are not necessarily indifferent between Portfolio A and Portfolio B if Portfolio A has twice the return and risk of portfolio B. The performance of a portfolio should, in our view, be primarily measured based on the long term growth rate of capital. This geometric average return is, by construction, adjusted for (real) risk factors. Whether we favor portfolio A or B should be based on this metric, not the arithmetic return corrected for volatility exposure.

A related, but slightly distinct point of discussion is the use of linear regression and mean comparison tests. As previously stated, we do not attach a lot of value to arithmetic returns. The growth rate of capital, and thus the end value of wealth, is primarily influenced by returns in the tail of the distribution. (Spitznagel, 2021) Extremes matter a lot. Consequently, observations around the mean become of secondary significance. Thus, also taking into account the asymmetric effect of losses compared to gains on the growth rate of capital, the utilization of arithmetic average returns becomes altogether meaningless in this class of distribution. As a result, comparing the portfolio returns based on sample mean, as we did to determine significant differences between portfolios over time, is questionable. *Do not cross a river that is 2 feet deep, on average.* (Marks, 2015) The same holds for linear regression, which we have used to determine the alpha and beta of the portfolios. The actual influence and thus the importance of outliers is not sufficiently accounted for if we calculate coefficients based on the average effect measured using arithmetic returns.

Altogether, this exposes, in our view, an important drawback to the key statistics used in modern finance. There seems to be a mismatch between the statistics used to describe the risk-return relationship and the observed (real world) distribution of the sample data. Mean-variance analysis - the bedrock of modern finance - is unsuitable in a field that is primarily non-gaussian. In other words, the assumption that returns are normally distributed, which is necessary for mean-variance, does not hold in practice. Theories and models building on this assumption, such as Markowitz' portfolio theory and the CAPM, are as a result flawed predictors for market behavior. Therefore, applying these theories in practice can, in fact, be harmful. If you use variance (or any other metric related to the second moment) as a measure of risk, while the distribution is not normal, you become fragile to events in the tail of the distribution. These events are rare by construction which can create the problem of biased sample means. This could explain how Markowitz' portfolio theory and the CAPM fail to predict the actual observed risk-return relationship.

8. Advice for investors and future research

Lastly, we shall conclude this research paper by presenting a number of tangible findings which investors may apply to enhance their performance. First of all, in the long run, portfolios of low-volatility stocks have outperformed portfolios of high-volatility stocks based on geometrical returns. Therefore, if an investor has to (or wants to) commit to an investment strategy, for example in the case of a *lump sum* investment, a portfolio of low-volatility stocks should be preferred compared to a portfolio of high-volatility stocks. Furthermore, since Asness et al (2020) find that different measures of volatility capture different effects, an investor may achieve diversification benefits by using multiple measures of volatility in the portfolio construction process. However, if an investor had to pick only one sorting method, we would advise to choose for variance of returns as the low-VAR portfolio has delivered the highest excess returns. In the case of active portfolio management, a hybrid exposure to volatility is preferred since we have found no persistent positive volatility premium throughout history. Generally, periods of positive Lo-Hi alpha spread are followed by periods of negative spread. Therefore, it could be performance enhancing to switch between low-volatility and high-volatility stocks. However, in order to perform this strategy one needs robust switching points. Although we have not identified robust switching points within this research paper, we know from Falkenstein (2013) that high-volatility stocks perform well in a 'highly speculative market environment'. In other words, alpha spread is to some extent predictable as it is correlated to the market cycle. This is our recommendation for future research. It could be particularly intriguing to discover robust switching points based on, for example, the rolling alpha spread or market cycle indicators.

Literature

Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. Journal of financial economics, 6(2-3), 95-101.

Hsu, J., & Li, F. (2013). Low-volatility investing. The Journal of Beta Investment Strategies, 4(2), 67-72.

Falkenstein, E. G. (2012). The missing risk premium: Why low volatility investing works. Eric Falkenstein.

Hussam, R. N., Porter, D., & Smith, V. L. (2008). Thar she blows: Can bubbles be rekindled with experienced subjects?. American Economic Review, 98(3), 924-37.

Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. The Journal of Finance, 19(3), pp. 425-442

Lintner, J. (1965). Security Prices, Risk, and Maximal Gains from Diversification. The Journal of Finance, 20(4), pp. 587-615.

Thaler, R. H., & Ganser, L. J. (2015). Misbehaving: The making of behavioral economics.

Fama, E. F., & French, K. R. (2006). Profitability, investment and average returns. Journal of financial economics, 82(3), 491-518.

Malkiel, B. G., & Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2), 383–417.

Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. Journal of financial economics, 6(2-3), 95-101.

Malkiel, B. G. (1989). Efficient market hypothesis. In Finance (pp. 127-134). Palgrave Macmillan, London.

Haugen, R. A., & 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.

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

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

Israelsen, C. L., & Cogswell, G. F. (2007). The error of tracking error. Journal of Asset Management, 7(6), 419-424.

Roll, R. (1992). A mean/variance analysis of tracking error. Journal of portfolio management, 18(4), 13-22.

de Koning, J., & van Vliet, P. (2018). De conservatieve belegger: Een bewezen formule voor hoog rendement met laag risico. Business Contact.

Boguth, O., & Simutin, M. (2018). Leverage constraints and asset prices: Insights from mutual fund risk taking. Journal of Financial Economics, 127(2), 325-341.

Horowitz, J. L., Loughran, T., & Savin, N. E. (2000). Three analyses of the firm size premium. Journal of Empirical Finance, 7(2), 143-153.

O'Brien, P. C., & Bhushan, R. (1990). Analyst following and institutional ownership. Journal of Accounting Research, 28, 55-76

Fichtner, J. (2020). The rise of institutional investors. In The Routledge international handbook of financialization (pp. 265-275). Routledge.

Garleanu, N., & Pedersen, L. H. (2020). Active and passive investing. SSRN Electronic Journal.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. Journal of monetary economics, 24(3), 401-421.

Lumley, T., Diehr, P., Emerson, S., & Chen, L. (2002). The importance of the normality assumption in large public health data sets. Annual review of public health, 23(1), 151-169.