Norwegian University
of Life Sciences

Master's Thesis 2024 30 ECTS
School of Economics and Business

Risk-return relationship and information content of implied and realized volatility

Vemund Almvik Tronsmoen
Master of Science in Business and Administration, Finance

Acknowledgements

This thesis is written as part of my Master of Science degree in Business and Administration with a profile in Finance at the Norwegian University of Life Sciences (NMBU). This thesis marks the end of my total five-year education in Business and Administration comprised of two years education at NMBU following my BSc degree in Business Law.

Writing this thesis has been both demanding and exciting, and I have learnt a lot on the way.

Finally, I would like to thank my supervisor, Associate Professor Daumantas Bloznelis for his guidance in writing this paper. I would especially highlight his patience, availability, and responsiveness, and I deeply appreciate him challenging my ideas throughout my work and being able to discuss back and forth on topics. I would also like to thank my family and girlfriend for support during the entirety of my education.

Norwegian University of Life Sciences

Ås, May 2024

Vemund Almvik Tronsmoen

Abstract

The capital asset pricing model (CAPM), Arbitrage Pricing Theory (APT), and Fama & French's three-factor model (FF3F) all assume a positive linear risk-return relationship. Findings of the contrary are deemed anomalous. In this thesis, I use portfolio sorts of short-term, long-term, and trade-weighted implied volatilities and their corresponding realized volatility as the main methodological framework. I investigate: i) the presence of a low-volatility anomaly, and whether it can be exploited. ii) Difference in risk-return relationship between realized and implied volatility, and iii) the predictive power of the call-put implied volatility spread and volatility risk premium on future stock returns. To my knowledge trade-weighted implied volatility has not been used before in such analysis, and my analyzed sample-period of 2009-2023 provides an update of previous findings on these topics. I find that i) the low-volatility anomaly is present but cannot be exploited. ii) Trade-weighted implied volatility gives the most theoretically consistent risk-return relationship out of all measures analyzed, and iii) the call-put implied volatility spread of short-term options predicts future stock returns.

Table of contents

Table of figures

Table of tables

Table of equations

1. Introduction

1.1 Motivation and background

Financial theory is essentially built upon the assumption of a positive risk-return trade-off typically based on some variation of a utility function arguing that an investor wishes to maximize their utility dependent on the factors expected return and expected risk (Markowitz, 1952; Tobin, 1958; Sharpe, 1964). This notion of a positive risk-return trade-off is the background of linear factor models such as CAPM by Sharpe (1964) and Lintner (1965), Ross' (1976) Arbitrage Pricing Theory, and Fama & French's (1993) three-factor model. These models not only assume a positive risk-return trade-off but also a linear relationship of risk and return. Findings of the contrary is deemed anomalous, such as the "low-volatility anomaly".

There has been research done on the low-volatility anomaly and its variations. Most prominently Haugen and Heins (1972) discovered that low-beta portfolios outperformed high-beta portfolios. In 2017 Pim van Vliet and Jan de Koning published a book called "High returns from low risk: a remarkable stock market paradox" which further perpetuates the notion that low risk portfolios outperforms high risk portfolios, or that "the tortoise beats the hare". In the book van Vliet and de Koning (2017) claims that the anomaly still persists when using realized volatility instead of beta as risk metric. Additionally, Ang et al. (2006) finds that portfolios containing low idiosyncratic volatility stocks outperform portfolios made up of high idiosyncratic stocks.

These three findings of different versions of the low-volatility anomaly all use ex-post risk measures. A market participant could perceive the stocks as low risk based on their previous volatility history, however market participants would often heuristically assign the level of volatility to a stock based on many factors outside of historic figures. News, general economic situation, and reports are examples of factors that shape a market participant's perception of risk. If a market participant suddenly feels that the price of an asset is about to change in either direction, it would probably not stem from historical data alone. Therefore, when testing the low-volatility anomaly it arguably makes sense to use ex-ante risk measures, as that is the one supposedly containing market expectation based on current information. Option-implied volatility is an example of forward-looking volatility due to the contracts giving the right to buy or sell an asset in the future, meaning that the implied volatility derived from option prices contains beliefs about the future underlying price.

Additionally, using option-implied metrics instead of historical metrics comes with additional information. For example, the volatility risk premium is often proxied by the spread between the implied and realized volatility (Bali & Hovakimian, 2009; Cremers & Weinbaum, 2010; Bollerslev et al., 2009). Moreover, the price of a call and an otherwise equal put contract does not have to be equal in the market. The spread between the call and put implied volatilities of an otherwise equal contract may very well have information on the anticipated upside and downside volatility.

The recurring way of proving or disproving the low-volatility anomaly is by using portfolio sorts, where portfolios are ranked based on the volatility metric, then the returns of each portfolio are compared to each other. A simple enough exercise is by doing the exact same portfolio sort but instead of ranking on realized volatility, the ranking is done on implied volatility.

So, does the tortoise beat the hare? Can the excess information be used to further predict stock market returns?

1.2 Research question and objectives

The objective of this thesis is threefold:

i) Investigate whether a low-volatility anomaly is present in the American equity market, and whether it can be exploited.

ii) Assessing whether the risk-return relationship of option-implied volatility is more theoretically consistent than of realized volatility.

iii) Exploring whether the excess information contents of option-implied volatility over realized volatility have predictive power of future stock returns.

The cumulative results will be used to answer the following research question:

Is option-implied volatility theoretically superior to realized volatility in terms of linear risk-return relationship, and does its excess information content predict future stock returns?

1.3 Structure

The thesis is divided into seven sections. Section 2 consists of the theoretical background and is split in two parts; the first part being a historical and theoretical recollection of risk-return relationship, or “ex-post” part. The second part is a theoretical explanation of implied volatility and its information content, or “ex-ante” part. Section 3 outlines the data gathered, while section 4 covers the methodological framework as well as preparation and use of data in this thesis. Section 5 is a presentation of the results obtained corresponding to each objective in chronological order, which are subsequently discussed in section 6. Lastly, section 7 concludes the findings of the thesis and answering the research question.

2. Theoretical background

2.1 Part 1 (ex-post)

2.1.1 Risk-return trade-off

Markowitz (1952) introduces the “E-V rule”, a rule of investment behavior when selecting a portfolio, that the investor should consider “expected return a desirable thing and variance of return an undesirable thing” (Markowitz, 1952, p.77). In other words, the investor would choose a portfolio that maximizes its expected return and minimizes the variance of return. The expected returns of portfolios will differ, as will the variance of return. One can choose a portfolio with a higher expected return, at the cost of undertaking higher variance and oppositely minimize variance by trading off expected return (Markowitz, 1952). This implies a theoretical risk-return trade-off. The portfolios with the best possible expected relationship between variance and return can be plotted, forming the “efficient frontier”.

2.1.2 Capital asset pricing model (CAPM)

Sharpe (1964) carries the theory of Markowitz further, extending it to capital asset pricing. Sharpe (1964) introduces a “capital market line” which is the efficient boundary of risk and return. Assets with an optimal risk-return relationship will be invested in, while assets not on the optimal line will not be invested in. This uneven distribution of investments will cause changes in prices of those assets, so that the price of “optimal” assets increases, while the price of sub-optimal assets decreases. This shift in price makes other combinations of assets optimal, which in turn leads to changes in prices. Therefore, under the assumption of rational investing, the price of assets is related to its risk-return relationship (Sharpe, 1964). The slope of the capital market line is equal to the Sharpe ratio of the optimal portfolio, and if more than one portfolio lies on the capital market line, then they all must have identical Sharpe ratios (Sharpe, 1965).

2.1.3 Critique of CAPM

While the notion that undertaking more risk will give higher expected return, i.e the existence of a risk-premium is intuitive, there are findings which do not support this theory. Haugen and Heins (1972) critiques the CAPM and its assumption of a risk-return trade-off in the following ways: firstly, the impossible problem of testing the CAPM empirically is that ex-post values are used to determine ex-ante expectations. Secondly, the timing of the sample could skew the results depending on the market situation of the time sampled. Thirdly, the sample could be biased towards survivors, meaning that assets used in the sample period have survived for the whole period, while excluding "terminated operations" (Haugen & Heins, 1972). Haugen and Heins (1972) try to solve the latter two problems by selecting stocks listed on NYSE in 1926, and if a stock is de-listed, then another stock will be introduced to the sample. The time period of the sample used is 46 years, divided into 5-year periods which may or may not be biased towards bullish or bearish markets. They find that over the long run, the portfolios with the lowest beta (risk) have had greater return than the portfolios with higher beta (risk) (Haugen & Heins, 1972).

2.1.4 Low-volatility anomaly

Haugen and Heins' (1972) findings were in regard to the risk-return relationship postulated by the CAPM, which uses the beta as risk measure. Since then, Frazzini and Pedersen (2014) created a betting against beta model based on the assumption that when leveraged equally, low beta assets outperforms high beta assets. An anomalous risk-return relationship has not only been found in regard to beta. Findings from Blitz and van Vliet (2007) suggest the existence of another low volatility anomaly, this time based on the realized volatility rather than beta. Using a trailing 3-year trailing standard deviation of weekly returns, they find the existence of a low volatility anomaly globally, and for the US, European, and Japanese markets separately. Baker and Haugen (2012) find a similar result using 2-year trailing standard deviation of monthly returns in 33 markets. De Carvalho et al. (2015) also find evidence of the low volatility anomaly for several markets, even within sectors using 2-year trailing standard deviation of monthly returns. Additionally, Ang et al. (2006) find a low idiosyncratic volatility anomaly in the US market using one-month trailing idiosyncratic volatility of daily returns.

The presence of low volatility anomaly has been found for beta, total volatility, and idiosyncratic volatility using different calculations, over different time periods, for different markets, and different sectors. This has created a rise in popularity for low volatility investing, causing MSCI and S&P to construct their first minimum volatility indices in 2008 and 2011 respectively, with at least 70 more indices globally that are tracked by exchange-traded funds (ETF) (Baronayan & Rothbarth, 2019).

2.1.5 Realized volatility

The aforementioned different risk measures have been used to prove the existence of different low volatility anomalies. Although the three ex-post volatility measures themselves are just that, measures of volatility, they convey different information. The beta is the systematic risk of an asset, which is the volatility of the asset in relation to a benchmark most often the market. This risk is non-diversifiable (Sharpe, 1964). The idiosyncratic volatility can be thought of as the complement to the beta, as it is the volatility that is not correlated to the market. This volatility measure can be diversified away, as each asset's unsystematic component of returns will even out through optimal portfolio management (Sharpe, 1964).

The realized volatility, often used interchangeably with historical volatility, is an ex-post estimate of the volatility. The realized volatility can be thought of as an estimate of the total volatility, containing both the systematic and the idiosyncratic volatility. A common way of calculating realized volatility is to take the standard deviation of however many trailing periods, then multiply it by the square root of the number of time periods. For example, if one has daily returns and wants to compute the annualized volatility, the standard deviations of the 1-day returns are multiplied by the square root of the number of trading days in a year, most commonly: $\sqrt{252}$ (Diebold et al., 1998). Diebold et al. (1998) laments such use of scaling for high-frequency asset returns. The reason is that for the square root of time rule to hold, the data has to be independent and identically distributed, which high-frequency returns are not. The use of square root of time rule to scale volatility for non-i.i.d data will either under- or overstate the actual volatility, while temporal aggregation should in actuality dampen the fluctuations of volatility over time (Diebold et al., 1998). These potential measurement errors can introduce complications when analyzing the realized volatility's relationship with returns.

2.2 Part 2 (ex-ante)

2.2.1 Implied volatility

Haugen and Heins' (1972) "impossible problem" of using ex-post values to determine ex-ante expectations is indeed impossible to circumvent as we cannot know what the investors' expectations are, and instead assume this to regress to the mean. However, by using option implied metrics instead of historical metrics we have a proxy for the ex-ante, or forward-looking measure of future volatility. This is because the price of an option should be determined by the options contract's attributes such as volatility, underlying price, time to maturity et cetera (Mayhew, 1995). As the price of an option depends on a volatility parameter among others, and is forward looking in nature, a pricing model for options can be used to reveal the market's expected future volatility of the option's underlying. This means that for options with one month to maturity, the implied volatility during the time period *t* will represent the average volatility of the options remaining time to maturity (Mayhew, 1995). Whether or not implied volatility is a good predictor of future volatility is heavily researched and contended, of which the answer is not directly concerning the objectives of the thesis.

Implications from options contracts have been used before to find a risk-return relationship that is close to linear, and therefore congruent to asset pricing models such as the CAPM. Buss and Vilkov (2012) use the thought-process that a forward-looking measure of volatility is desired to test the risk-return relationship the CAPM posits, in an attempt to resolve Haugen and Heins' (1972) "impossible problem". What they found is that using their calculated implied betas, a more linear risk-return relationship emerges in line with Sharpe's (1964) CAPM, compared to using standard rolling-window betas (Buss & Vilkov, 2012). However, the problem is that implied volatility is not a forward-looking beta. Buss and Vilkov (2012) have to construct implied betas by making modeling choices that may or may not be a correct representation of the market's anticipation of beta. When testing the existence of a low realized/total volatility anomaly and not the CAPM, the need for modeling and intermediate calculation is drastically lessened, though not fully removed.

2.2.2 Options

Black and Scholes (1973) define an option contract as “a security giving the right to buy or sell an asset, subject to certain conditions, within a specified period of time” (Black & Scholes, 1973, p. 637). To further explain what an option contract is, I will provide explanations using definitions from The Options Industry Council (OIC, n.d.) of the most relevant terminology and attributes.

An option contract specifies the asset which can be bought or sold if exercised. This asset is called the “underlying”. The contract also specifies at what price the underlying will be bought or sold for if the option is exercised, i.e the “strike”. If the contract gives the owner the right to buy the underlying, it is a “call” option. Conversely, if the contract gives the owner the right to sell the underlying, it is a “put” option. The contract’s lifespan is denoted by its “expiration date” or “maturity”, which is the last day at which the contract can be exercised. The contracts themselves can be bought or sold; the buyer is called a “holder” while the seller is called a “writer”.

Options can be described in different ways depending on their intrinsic or extrinsic value. The intrinsic value is the value of an option if exercised. Considering a call option, if the strike price is below the current underlying price, then one could buy the underlying for cheaper than the asset’s current price by exercising the option. One could then sell the asset at the current price, meaning one would make a profit. The intrinsic value is therefore the difference in strike price and underlying price. If an options contract has intrinsic value, then the options is in-the-money (ITM). The extrinsic value or “time value” of a call option is whatever the contract is worth apart from its intrinsic value. A call option with a strike price above the current price has no intrinsic value and is therefore out-of-the-money (OTM). However, the contract is not worthless, as the underlying price can change in the remaining time to maturity. The price of the option is then made up purely of time value. If an option has a strike price equal to the underlying price, then it is at-the-money (ATM) (OIC, n.d).

Rewriting Black and Scholes' (1973) definition above, "one can hold or write an option that gives the right to buy (call) or sell (put) the underlying for the strike price within the contract's expiration date".

There are several styles of option contracts, however the most common are "American" and "European". The difference between the two styles is that the "American" style contract can be exercised whenever in the lifespan of the contract. "European" contracts can only be exercised at maturity, when the lifespan of the contract runs out (Black & Scholes, 1973). The implied volatilities analyzed in this paper are from American style options only.

2.2.3 Cox-Ross-Rubinstein binomial model (BOPM)

The binomial option pricing model (BOPM) was developed by Cox, Ross and Rubinstein (1979) as an alternative of pricing options that is less mathematically advanced than the famous Black-Scholes model (Cox et al., 1979). Additionally, the model allows for more correct pricing of American options, meaning the model accounts for dividend yield and early expiration, which the Black-Scholes model cannot.

Cox et al. (1979) describes the main idea behind this pricing model, which is to create a binomial tree in which the stock price has probabilities of either moving up or down. By assuming riskless arbitrage is impossible, the price of the option on the underlying stock can be calculated. Considering a stock with the price S, the price can move up (S_u) or down (S_d) with probabilities q and $1-q$ respectively. Similarly, to value a call (C) or put (P) option, the option price can move to C_u and $C_d$with probabilities q and $1-q$ respectively (Cox et al., 1979). If a call option is in-the-money, the difference between the strike price (K) and the stock price has to be positive. If not, then the contract is out-of-the-money and is worthless if exercised.

The value of a call option is therefore:

$$C = \max(0, S - K) \qquad \{1\}$$

Inversely, the value of a put option is:

$$P = \max(0, K - S) \qquad \{2\}$$

Furthermore, the model takes the risk-free interest rate into account as it assumes no riskless arbitrage. Additionally, as the previously mentioned probabilities q and $1 - q$ are subjective perceived probabilities, these are instead replaced by risk-neutral probabilities denoted by p.

To calculate p, one would need the upward movement u and downward movement d, as well as the one-plus risk-free interest rate r:

$$p = \frac{r - d}{u - d} \quad \text{and} \quad 1 - p = \frac{u - r}{u - d} \qquad \{3\}$$

This gives the following equation for valuing a call option one period before expiration:

$$C = [pC_u + (1 - p)C_d]/r \qquad \{4\}$$

So far, the risk-free interest rate, as well as the upward movement and downward movement are fixed for one unit of calendar time. To allow for valuation of options over smaller time intervals, these need to be modified to properly scale with the risk-neutral probabilities. Otherwise, the probability of up or down movement in prices in the span of some minutes can end up being the same as the span of a couple months. Considering a length of h in which prices change, t being fixed time intervals, and n being number of periods of length h, then $h = t/n$. To make the risk-free interest rate dependent on time intervals, the risk-free interest rate is expressed as $\hat{r} = r^{\frac{t}{n}}$. Additionally, to make the price movements dependent on number of periods n; $u = e^{\sigma\sqrt{t/n}}$, and $d = e^{-\sigma\sqrt{t/n}}$. Since d is the same as u, however in the negative power, then $d = 1/u$.

Accounting for dividend yield (δ), the stock owner will receive dividend of δS_u or δS_d. This means that the stock price ex-dividend will be $u(1-\delta)^v S$ or $d(1-\delta)^v S$, where $v = 1$ if the date is ex-dividend and $v = 0$ if else.

Altogether, the value of a call option for i periods before maturity is (Cox et al., 1979):

$$C(n,i,j) = \max\left[u^j d^{n-i-j}(1-\delta)^{\bar{v}(n,i)}S - K, \frac{[pC(n,i-1,j+1) + (1-p)C(n,i-1,j)]}{\hat{r}}\right] \quad \{5\}$$

$$\text{for } j = 0,1,2,\dots,n-i$$

To value a put option, one simply has to reverse the difference between the strike price and stock price at every n in the binomial framework (Cox et al., 1979).

2.2.4 Deriving Implied volatility from the BOPM

The binomial option pricing model (BOPM) calculates option prices using the following inputs:

$S =$ Stock price
$K =$ Strike price
$p =$ Risk-neutral probability calculated using u, d, and $\hat{r}$
$u =$ Upside movement with respect to volatility (σ), time (t) and number of periods (n)
$d =$ Downside movement with respect to volatility (σ), time (t) and number of periods (n)
$\hat{r} =$ Risk-free interest rate with respect to time (t) and number of periods (n)
$\delta =$ Dividend yield if applicable

All variables in the formula are observable, however the volatility (σ) is not. This means that accurately pricing options would mean knowing the implied volatility, or ex-ante knowledge of the volatility. The best one could do when pricing options with this model is therefore merely an estimate of the implied volatility. This also means that by using this model to observe the implied volatility, one would need the market price of the option.

There are multiple ways of deriving implied volatility numerically from option prices, but an efficient approach is the Newton-Raphson method (van der Hoek & Elliott, 2006). The general idea behind the method is to compare the option prices calculated through the binomial option pricing model to the market prices of the option. If there is a difference in prices, then the volatility variable is changed in the BOPM until equal prices are obtained. The volatility that makes the two prices equal is the estimated implied volatility.

2.2.5 What drives implied volatility?

When pricing options with a model such as the BOPM, the theoretical option prices calculated will have all but one variable efficiently priced in, as only the implied volatility is unobservable. Additionally, option prices and implied volatility have a monotonic relationship (Rodriguez et al., 2015). This means that option prices should equal its difference from the theoretical price multiplied by the implied volatility's effect on the price, Vega. Because deriving the implied volatility numerically is estimated by comparing the option's theoretical price to the market price, the implied volatility can be seen as a "fudge factor". One could then argue that the observed implied volatility means nothing at all and is simply a measure of the BOPM's mispricing. Or is it the market options that are mispriced?

Continuing this train of thought, at least a part of the implied volatility of an option can be seen as a reflection of supply and demand. Haug and Taleb (2011) argues that option traders rely on "sophisticated heuristics" instead of option pricing models when pricing options. Additionally, Gârleanu et al. (2009) finds a relationship between an option's end-user demand and its "expensiveness", or implied volatility, indicating that implied volatility is a reflection of supply and demand. If the supply and demand of an option changes, the bid-ask spread of the option change as well, since the writers can take a higher price or be willing to take a lower price depending on the change. This shift in price will perturb the estimated implied volatility. However, the changes of an option's supply and demand must obviously stem from investors' anticipation of movement in the underlying price, hence implied volatility being a parameter of the market's anticipated volatility.

By this rationale, different implied volatilities of options can say something about market sentiment. This is especially relevant when considering a put and a call option of equal maturity, underlying and strike. The put-call parity is a no-arbitrage relationship which holds for European options, but not for American options. In relation to implied volatility, the put-call parity states that a call and a put option otherwise equal will have the same implied volatility (Cremers & Weinbaum, 2010). Because the equity options in the data used in this paper are American, the put-call parity itself will not be discussed extensively, as deviations from this occurs frequently in practice (Cremers & Weinbaum, 2010). However, these deviations are an effect of pricing, so the implied volatilities do not only contain information about the anticipated movement of the underlying price, but perhaps also in which direction when accounting for put-call differences in implied volatility.

The observed implied volatility is a function of the option price as shown earlier. There is a sort of a “chicken or the egg” conundrum when it comes to implied volatility. Is the market price affected by the market’s anticipation of future volatility, or is the implied volatility affected by the market price of the option in question? A well-known use of options is for hedging. As a basic example, one can buy a stock then take the opposite position in the form of a put option at a certain strike price to reduce the potential losses should the stock price fall. When people buy this put option, the demand for this option increases, causing the price of these options to increase. This results in an increased observed implied volatility. The problem with this effect on implied volatility is that hedging one-self this way is not reflecting the hedger’s anticipation of future volatility, but rather as an insurance policy to reduce the downside risk. These are acts driven by the aversion to risk, rather than the anticipation of risk. The CBOE volatility index (VIX) uses implied volatility to convey the overall anticipated market risk, but it is colloquially called the “fear index” perhaps appropriately due to the role of options in hedging strategies.

2.2.6 Call-put implied volatility spread

The call-put implied volatility spread (CPIV) is the spread between a call option’s implied volatility and a put option’s implied volatility with otherwise equal specifications such as strike, maturity, and underlying. As previously mentioned, the put-call parity states that the implied volatility of such contracts should be equal for European options. This parity does not hold for American options due to possibility of earlier exercise, as well as taxes, transaction costs, borrowing rate inequal to lending rates, and margin requirements (Cremers & Weinbaum, 2010). Because of all these factors, deviations from the put-call parity exist and contain information on the underlying asset. The formula for the call-put implied volatility spread for asset i at time t is:

$$CPIV_{i,t} = IV_{i,t}^{Call} - IV_{i,t}^{Put} \qquad \{6\}$$

Considering the observed implied volatility is derived from market prices of options, a call and put option with the otherwise same specifications but different implied volatilities is a result of one being more expensive than the other. Applying the logic of supply and demand and assuming implied volatility is the market's anticipation of future volatility, if the put option in question has a higher price and subsequently a higher implied volatility, then the market anticipates a shift in the underlying price accordingly. The implied volatility supposedly conveys the probability of future price movement, and I hypothesize that the call-put implied volatility spread provides further information about the direction of the movement as well.

Cremers and Weinbaum (2010) find that the call-put implied volatility spread contain information about future stock returns in their sample period of 1996-2005. They also find that a positive spread gives economically significant positive returns, while negative spread gives negative returns. This relationship seems to decline over time as these findings are less significant in their second subperiod, crediting this to reduced mispricing over time, decrease in trading costs, and hedge fund capital growth (Cremers & Weinbaum, 2010). Similar findings are also found by Bali and Hovakimian (2009).

2.2.7 Implied-realized volatility spread/Volatility risk premium

The implied-realized volatility spread is the difference between the implied volatility and the realized volatility. As the implied volatility reflects expectations of the future volatility, while the realized volatility is calculated using past values, the realized volatility during time period t should be comparable to the implied volatility during the period of time $t-1$. Thus, the formula for the implied-realized volatility spread for asset i during the time period t is:

$$IVRV_{i,t} = IV_{i,t-1} - RV_{i,t} \quad \{7\}$$

The implied-realized volatility spread can be considered a proxy for the volatility risk premium (Bali & Hovakimian, 2009; Bollerslev et al., 2009). Considering options serving as insurance, and observed implied volatility being derived from market prices of options, then a positive implied-realized volatility spread reflects the premium the options writers demand for undertaking the risk. Eraker (2021) finds that on average across the market, there is a substantial annualized volatility risk premium of 3,3%. The information content in the volatility risk premium is in what degree the options writers are anticipating a price shift in either direction. The writers of the options profit off of little to no movement in price, meaning they will charge a higher price to hedgers if they believe there to be a higher probability of a significant price movement before expiration.

Bollerslev et al. (2009) find that stock returns are predictable by the volatility risk premium, however using different methods of estimating implied and realized volatility. They also find that this predictability is significant for horizons of up to six months (Bollerslev et al., 2009). Bali and Hovakimian (2009) find that portfolios containing assets with low (high) “inverse” volatility risk premium give high (low) returns. They do however subtract the implied volatility from realized volatility, hence being an “inverse” volatility risk premium.

3. Data

All the data used in this thesis come from Refinitiv Eikon's Datastream. The universe consists of 513 stocks listed on NYSE and NASDAQ in the time period of June 1st, 2008 until and including January 1st, 2024.

The reason for this specific start date is due to unavailability of options data before June 2008. The availability of the options data is therefore deciding what stock data is used. For many stock options, there are lacking implied volatilities at some periods due to the time series often being scrapped if the company is acquired or changes name. So even though the total return index and price index may have available data for the entire period, the options data of the same company might be shorter. One could find inactive, or "dead" options data chains for the company up until the beginning of the now active, or "new", options data chains and splice them together. However, that would be too time-consuming and the risk of matching some faulty data is high, considering much of the "dead" data are probably inactive for a reason, such as a change of company ownership or mergers. As a result of this, the majority of implied volatility data are on "popular" or successful companies. In this universe, 471 out of the 513 companies are included in the S&P 500 as of January 2024. This might be a problem of survivorship bias, as some of the current S&P 500 companies were growth stocks in 2008, and in hindsight are sure-to-succeed companies. The remaining 42 companies are not included in the current S&P 500, and some are even delisted companies.

Additionally, the one-month US treasury bill rate for the same period is gathered from Refinitiv Eikon's Datastream to proxy for the risk-free rate.

For calculating returns and realized volatility, the total return index of each stock was obtained. All the different implied volatilities of the 513 stock options were already available and are calculated in-house by Datastream using the BOPM, as all the equity options are American. The different implied volatilities (IV) obtained are:

1M = Constant one-month maturity of at-the-money strike (30 days)

1Y = Constant one-year maturity of at-the-money strike (360 days)
TW = Trade-weighted IV across all strikes by trading volume, with a maturity of the nearest contract.

For all the different implied volatility metrics, the average of the call and put implied volatilities will be calculated to determine the single “stock IV”.

4. Methodologies

4.1 Introduction

The first objective of the thesis requires establishing whether a low-volatility anomaly is present in this data universe. This will be done following the method of previous literature.

Subsequently, the same analysis will be done by using different volatility metrics implied by the options to assess the second objective of the thesis. To assess the theoretical consistency of the different volatility metrics in terms of risk-return relationship, a test for monotonic relationship will be conducted as well as analysis of the portfolios' Sharpe ratios.

Furthermore, to explore the information content of options, portfolios will be conditionally double sorted based on the implied volatility and their corresponding call-put implied volatility spread as well as their implied-realized volatility spread. Additionally, long-short portfolios will be created by sorting on the call-put implied volatility spreads, as well as weighted by the same spreads to test the hypothesis of the spread having directional predictability of returns. That is, whether or not going long a positive call-put implied volatility spread gives positive returns and going short a negative spread gives positive returns.

All the analysis will be performed on the full sample period of July 2009 throughout December 2023, as well as two sub-periods of July 2009 throughout September 2016 and October 2016 throughout December 2023 for robustness.

4.2 Preparation of data

4.2.1 Realized volatility

There are differences in the data used to calculate the realized volatility. In the book "High returns from low risk" de Koning and van Vliet (2017) uses a trailing three-year standard deviation of monthly returns to establish the stock volatilities. Ang et al. (2006) as well as Bali and Hovakimian (2009) use one month of daily returns. Blitz and van Vliet (2007) use three years of weekly returns, while Baker and Haugen (2012) use two years of monthly returns to compute the realized volatility.

Considering this paper aims to look at the differences between using option-implied volatility metric instead of realized volatility to potentially exploit a possible low-volatility anomaly, the availability of the implied volatilities dictates the use of data for calculating the realized volatility. The chosen implied volatilities in this paper are therefore the constant one-month maturity, the constant one-year maturity and lastly the trade-weighted by volume with a maturity equal the nearest contract.

The constant one-month maturity reconciles both the period used in some previous literature and the availability obstacle. The one-year constant maturity is chosen as comparison to a longer horizon of realized volatility, although the most prominent articles regarding low-volatility anomaly seem to prefer a longer horizon than one-year trailing realized volatility, the longest horizon available for implied volatilities is one year. Lastly, the trade-weighted implied volatility is measured using the contracts nearest maturity. The equity options of different underlying usually expire on the same date, which is the third Friday of the month, with some exceptions (CBOE, n.d.). This means that the expiration day of the month is different for each month. This then also means that the trade-weighted implied volatilities have different information content for each period, however, they should be comparable across equities in the same periods. The reason for including the trade-weighted implied volatility is due to the desire of a correct stock implied volatility. As the implied volatility reflects the market's ex-ante perception of volatility, and if the most popular strike is not at-the-money, then the implied volatility of the most traded strikes reflects more of the market sentiment and should therefore be weighted more. As all the implied volatility data are captured on the first day of every month, the nearest month option maturity is approximately three weeks ahead. To more correctly compare the trade-weighted implied volatility with its realized counterpart, a trailing three-week realized volatility is constructed.

The realized volatilities used in this paper are therefore trailing one-month of daily returns, trailing one-year of monthly returns, and trailing three-weeks of daily returns.

These realized volatilities are annualized as the implied volatility is by default expressed annually. The scaling will be done using the square-root of time rule. Although there might be unwanted implications explained by Diebold et al. (1998), it seems to be the general standard practice.

4.2.2 Implied volatility

As previously mentioned, all the implied volatilities gathered are all computed in-house by Refinitiv Eikon's Datastream using the BOPM. The implied volatilities obtained are the call and put implied volatilities for each underlying at either one-month constant maturity, one-year constant maturity, and trade-weighted by volume. Considering the computation of implied volatilities for call and put options, ceteris paribus, the implied volatilities will be the same if the two contracts also have the same price. For some options and periods in this data universe, the price difference is negligible, however for others there are distinct differences. This causes differences in the call and put implied volatilities. A higher put implied volatility might indicate that the market anticipates a downward price movement, relative to the belief of an upward price movement. To reconcile the two beliefs in the direction of price movement, an average of the call and put implied volatilities is calculated to determine the individual stock implied volatility. This follows Bali and Hovakimian's (2009) methodology and partly the calculation of the CBOE volatility index (VIX) (Mayhew, 1995).

4.2.3 Risk-free rate

The one-month US treasury bill is used to proxy the risk-free rate. The risk-free rate will only be used when calculating the Sharpe ratio of the portfolios. The one-month US treasury bill rate is expressed annually, so to calculate the portfolio return in excess of the risk-free rate, the rate is converted into a monthly interest rate.

4.2.4 Equalizing data

To further ensure that the different volatility metrics are comparable, the datasets will be equalized. There are lacking data at different periods and for different stocks for the six volatility metrics analyzed. The empty data in each dataset will decide the empty data in the other datasets. If there is a stock at period t that contains no data in one dataset, then the data in the other datasets are removed as well. Conclusively, the six different datasets will have the exact same amount of data. For example, the constant one-year maturity implied volatility starts in July 2009, therefore all other datasets used will start in July 2009 as well for more correct comparison between the different volatility metrics.

4.2.5 Call-put implied volatility spread (CPIV)

The implied volatilities of all options are for both the call and put. To calculate the call-put implied volatility spread, the implied volatility of the put option is subtracted from the implied volatility of the corresponding call option.

4.3 Part 1: Differences in realized volatility and implied volatility

The difference in periods between implied and realized volatility is illustrated as follows:

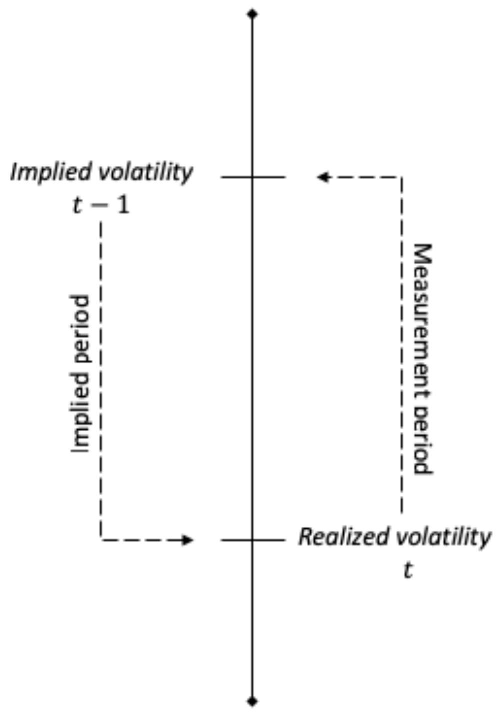

Figure 1: Illustration of corresponding periods of implied and realized volatility

4.3.1 Implied-realized volatility spread (IVRV)

When using different implied and realized volatility metrics to assess the risk-return relationship, the results will obviously be the same if there are no differences between the volatility metrics. Although differences are expected, they have to be quite different to give a completely contrasting result. A way to evaluate the potential result differences beforehand is to calculate the difference between the implied volatility and the realized volatilities for each stock at each time. Because the implied volatility is forward-looking, while the realized volatility is backward-looking, the implied-realized volatility spread at time t will be calculated by subtracting the realized volatility during the time period t from the implied volatility during the time period $t-1$.

4.3.2 Correlation

Additionally, the correlation of each stock's implied volatility and their realized volatility will be calculated. The implied-realized volatility spread alone cannot say anything about potential result differences as long as they are both highly positively correlated. Hypothetically, if one finds the presence of a low-volatility anomaly using realized volatility, then to expect the complete opposite result when using implied volatility would suggest having negatively correlated realized and implied volatilities. Therefore, the correlations of all implied volatilities to their corresponding realized volatilities are calculated. Once again, the implied volatility is ex-ante, while realized volatility is ex-post, meaning the correlations will be calculated accordingly.

4.3.3 Portfolio sorts

After getting an indication of the potential result differences, I will employ the standard procedure of checking for presence of a low-volatility anomaly by forming decile portfolios sorted on the volatility metric for the period of July 2009 throughout December 2023. The portfolios are ranked from low to high volatility, meaning portfolio 1 contains the least volatile stocks, and portfolio 10 contains the most volatile stocks. The portfolio returns are calculated using equal weights, allowing no shorting of stocks. The amount of stocks sorted into portfolios differ from the data universe due to differences in available data. Each option-implied volatility metric and corresponding realized volatility is equalized in the way that empty data in either dataset are the same. This results in fewer stocks than initially gathered are being sorted throughout the whole period.

On average, the amount of stocks sorted from July 2009 to December 2023 is 384, going from 268 at the minimum, to 486 at maximum.

To begin with, I will perform the sorting of the stocks based on the realized volatility to assess whether or not the low-volatility anomaly is present in this data universe. At the beginning of each month, stocks are sorted based on their trailing one-month, or one-year, or three-week standard deviation. The portfolios will be rebalanced each month, and the captured returns on each portfolio are the returns measured at the same time. This is a practically infeasible sorting, as the realized volatility is measured using ex-post values and cannot be replicated as a strategy as the future realized volatility is unknown. A potential presence of the low-volatility anomaly has to be interpreted as whether or not the low-volatility stocks have had greater returns than high-volatility stocks looking backwards. I will not do the same procedure for the implied volatilities, as these do not have a measurement period of ex-post values unlike the realized volatilities. Therefore, one cannot employ the same logic of the stocks that have had the lowest or highest implied volatility have simultaneously earned the most returns, because the implied volatility has not been measured alongside the return.

Instead, I perform the same sorting of portfolios for both the option-implied side and the realized side with the notion of whether or not the potential low-volatility anomaly can be exploited by either volatility metric. At the beginning of each month, stocks are sorted based on their current implied and realized volatilities, and the captured returns on each portfolio are instead the returns measured over the following month. Subsequently I will compare the risk-return relationship of the six different volatility metrics.

4.3.4 Sharpe ratio

To evaluate the risk-return relationship for the different portfolio sorts, the returns excess of the risk-free rate of each period will be calculated along with the portfolio's total excess return and the standard deviation of the excess returns. The Sharpe ratio of each decile portfolio gives an indicator of their performance accounting for both return and risk. For example, a portfolio can have a higher return, yet unproportionate to the amount of risk needed to attain this return. A simple linear regression model of Sharpe ratios on the portfolio index is run. This effectively creates just a trendline of the Sharpe ratios across portfolios, and the slope of this line as well as their residual standard error will be evaluated to determine which volatility metric gives the most linear and therefore "theoretically" consistent risk-return relationship. Consider the efficient frontier and the capital market line (CML); the slope of this line is the same as the Sharpe ratio. The optimal portfolio in terms of risk-adjusted return theoretically lies on this line. If all portfolios lie on this line, then the volatility metric used to sort these portfolios is the most "theoretically" consistent in terms of risk-return relationship, as one is sufficiently compensated per unit of risk is undertaken. Consequently, a slope coefficient of zero with minimal residual standard error is desired.

4.3.5 Monotonicity

In addition to the Sharpe ratio, each portfolio sort will be evaluated on their monotonic relationship, as Buss and Vilkov (2012) did when investigating their implied betas. The test of monotonic relationship will be Patton and Timmermann's (2010) nonparametric test. The test will be done on quintile portfolio sorts rather than decile portfolio sorts as advocated by Romano and Wolf (2013). Running the test on quintile portfolios should ensure sufficient power of the test (Romano & Wolf, 2013). To evaluate the different portfolio sorts, the test will be subsequently run on both sides, that is assuming an increasing and a decreasing monotonic relationship. To tell which of the volatility metrics is most concurrent with a linear risk-return relationship, a rejection of decreasing monotonic relationship and simultaneously failing to reject an increasing monotonic relationship is desired.

4.4 Part 2: Excess information content of implied over realized volatility

4.4.1 Double sorts

To explore the excess information content of option-implied volatility over realized volatility, the portfolios will be sorted as before based on their implied volatility. Additionally, these portfolios will be sorted conditionally (hierarchically) based on either a positive or negative call-put implied volatility spread. The results will show in which way the spread contributes to the volatility portfolios' returns. The excess returns that the positive spread gathers over the negative spread will be subjected to a t-test with a confidence level of 0,95 to assess whether or not these average excess returns are statistically different from zero.

The same procedure will be performed using the implied-realized volatility spread, to assess its contributions to the volatility portfolios' returns. The returns that the positive spread earns in excess of the negative spread will also be subjected to a t-test with a confidence level of 0,95.

4.4.2 Long-short portfolio sorts

Lastly, long-short portfolios will be created by ranking on the spread, while simultaneously weighting based on a positive or negative spread. This is done to test the hypothesis of the call-put implied volatility spread being a predictor of increase or decrease in the underlying price. The portfolios will have equal weights, however some negative and some positive.

5. Results

The results presented in this section is for the full sample period of July 2009 throughout December 2023 only. The results for the two sub-periods of July 2009 throughout September 2016 and October 2016 throughout December 2023 will be referenced throughout this section but will be found in the appendix section to avoid clutter.

5.1 Correlations and volatility risk premium

The table below shows median and mean correlation coefficients and implied-realized volatility spread between the implied and realized volatility metrics. As mentioned earlier, the calculations are done using implied volatility during the time period $t - 1$ and realized volatility during the time period t. The results show a positive correlation between the implied and realized volatility, yet not a very high correlation. The least amount of correlation is found between the trade-weighted implied and three-week realized volatility, while the highest correlation is found between the one-year implied and realized volatility. Additionally, the implied-realized volatility spreads are all positive as expected, and in line with Eraker's (2021) findings of a positive average volatility risk premium. These findings indicate that there will not be substantial differences when comparing the risk-return relationship of realized and implied volatility later on.

Table 1: Correlations and IVRV spread

Vol metric	*Median*		*Mean*	
	Correlation	*IVRV spread*	*Correlation*	*IVRV spread*
1-month	*0,54*	*2,33*	*0,50*	*4,03*
1-year	*0,62*	*1,20*	*0,51*	*1,65*
3-week	*0,38*	*6,82*	*0,34*	*9,68*

5.2 Portfolio sorts

5.2.1 Realized volatility and returns over the time period t

When looking at the decile portfolio's performance from 2009-2023, both sorting on realized volatility and capturing returns during the time period t, there is a clear pattern of low-volatility anomaly. The portfolios containing the least volatile stocks have been outperforming the most volatile portfolios. This pattern of a low volatility anomaly is not as apparent in the realized volatility using one-year trailing standard deviation, where it looks quite flat. This is however looking back on both the measurement period of volatility and returns, which has to be interpreted as the least volatile stocks have simultaneously had greater returns than the more volatile stocks. If one were to exploit this phenomenon, one would need to know the realized volatility beforehand. Similar pattern of a low-volatility anomaly can be found for both the two subperiods; see Figures 8 and 9 in the appendix.

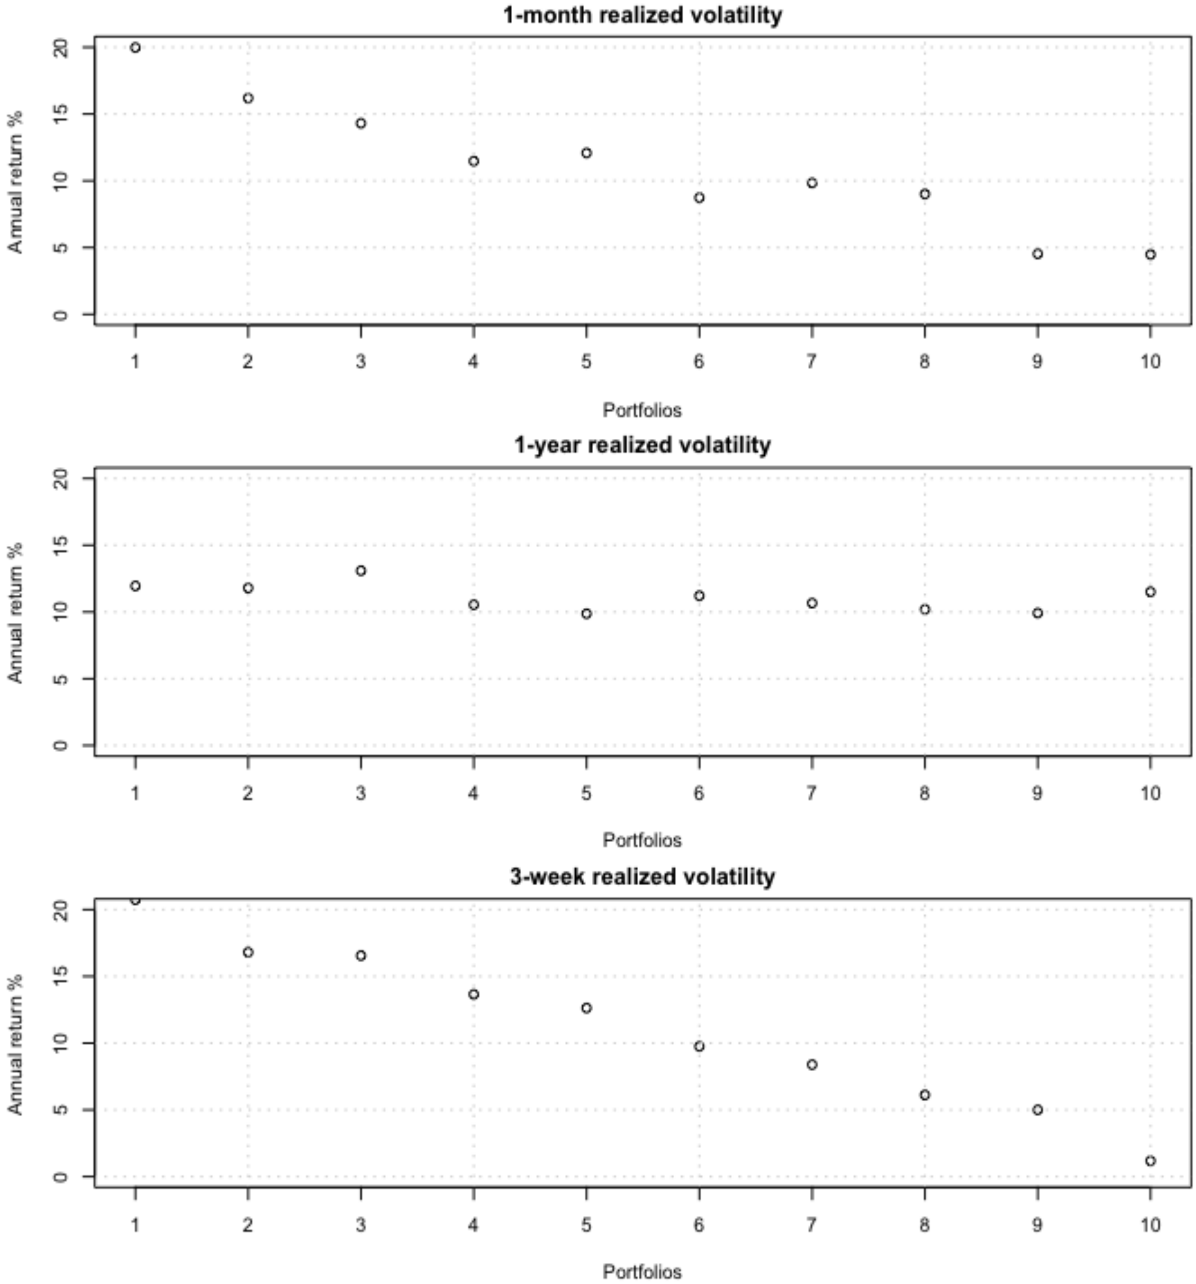

Figure 2: Portfolio sorts of realized volatilities 2009-2023, capturing returns at same month

5.2.2 Realized and implied volatility, returns over the following period

The implied volatility is supposedly the anticipated future volatility. Therefore, portfolios are now sorted and rebalanced using implied volatilities during the time period $t-1$, while still capturing returns over time t. This is done to assess whether one could exploit the low volatility anomaly by using the implied volatility of equity options as a predictor of the next months realized volatility. As seen in the figure below, the pattern is now completely different. For all three implied volatilities, there is a clear trend in high volatility gathering higher returns in the period of 2009-2023. Additionally, the portfolios are sorted and rebalanced using realized volatilities during the time period t, capturing returns over $t + 1$. The possible low-volatility anomaly observed previously cannot be exploited using either volatility metric in this data universe. The same trend can be found in the two sub-periods; see Figures 10 and 11 in the appendix.

Looking closely at the returns gathered for the decile portfolios, it is clear the portfolios sorted on different implied volatility metrics have the largest gap between the highest and lowest volatile portfolios. The least volatile portfolio for all three implied volatility metrics has lower returns than their realized counterpart. Moreover, the most volatile portfolios sorted on implied volatilities has the highest returns.

Even though portfolios sorted on implied volatility give higher returns for higher volatility, the risk-return relationship as portrayed by the Sharpe ratios in the dashed lines below shows that the excess volatility undertaken in the portfolios is not always compensated sufficiently by greater returns. For portfolios sorted on both one-month realized or implied volatility, the best performing portfolios in terms of risk-return relationship is portfolio 2 and 3 respectively. Both being relatively low volatility portfolios. For portfolios sorted on one-year realized volatility, the best performing portfolio is portfolio 6, which is slightly volatile. For the one-year implied volatility sort, portfolio 3 has the best risk-return relationship. The largest difference between the realized and implied volatility metric in terms of Sharpe ratio comes in the three-week realized, and the trade-weighted implied volatility. For the three-week realized volatility, the Sharpe ratio is generally decreasing after its peak at portfolio 3, yet for the corresponding trade-weighted implied volatility the Sharpe ratio is generally increasing at least up until its peak at portfolio 7.

The peak Sharpe ratios are different in the two sub-periods, where the period of 2009-2016 have their peak Sharpe ratio in low volatility portfolios except for the trade-weighted implied volatility. In the period of 2016-2023 all volatility metrics have a peak Sharpe ratio in a high volatility portfolio; see Figures 10 and 11 in the appendix.

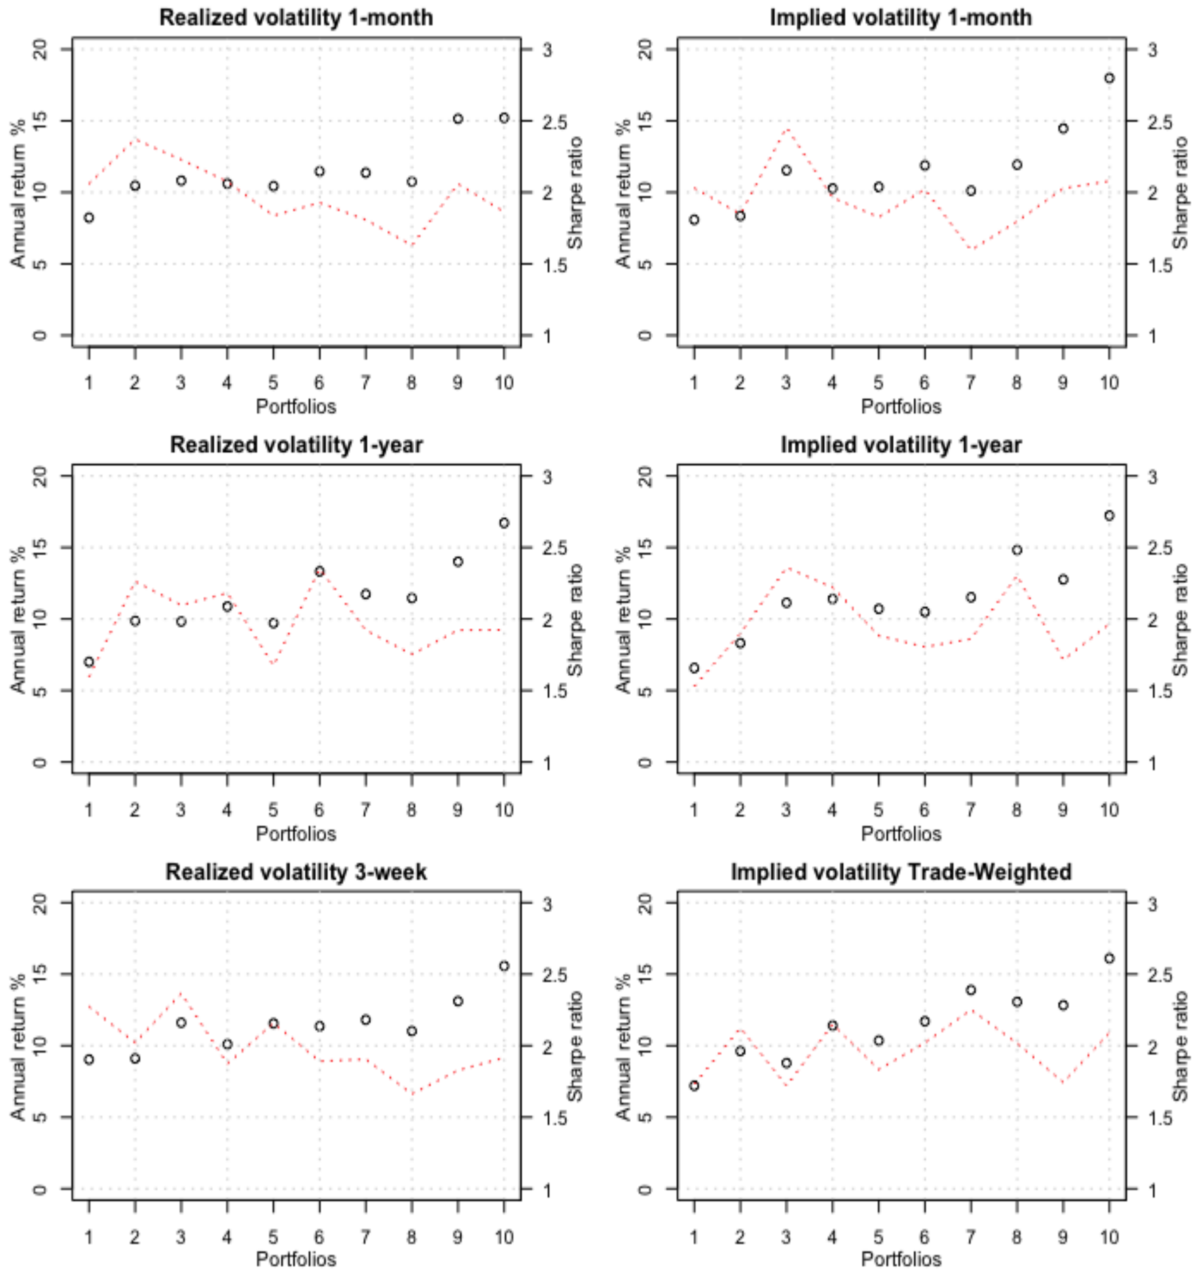

Figure 3: Portfolio sorts of realized and implied volatilities 2009-2023, capturing returns at following month

Looking more closely at the trendline for the portfolio Sharpe ratios of each volatility metric, we see that the implied volatilities all have slightly higher slope coefficients than their realized counterpart, which are all negative. The slope of one-month and three-week realized volatility is statistically different from zero (marked with an asterisk) with a confidence level of 0,95. However, the residual standard errors are also larger for all implied volatilities than their realized counterpart.

The statistical significance of the slope coefficients being different from zero is even stronger and applies for one-year realized volatility as well in the first sub-period. In the second sub-period, only the slope coefficient of one-year realized volatility is statistically different from zero; see Tables 8 and 9 in the appendix.

Table 2: Linear regression of Sharpe ratios (trendline) 2009-2023

Vol metric	*Intercept*	*Slope coefficient*	*Residual standard error*
1-month RV	*2,24*	*-0,05**	*0,18*
1-month IV	*2,05*	*-0,02*	*0,24*
1-year RV	*2,00*	*-0,01*	*0,27*
1-year IV	*1,91*	*0,01*	*0,28*
3-week RV	*2,26*	*-0,05**	*0,16*
Trade-weighted IV	*1,88*	*0,02*	*0,20*

* *statistically distinguishable from zero*

As expected from looking at the previous plots, we can never reject the null hypothesis of an increasing monotonic relationship for portfolio returns sorted on any volatility metric. On the other hand, we can reject the null hypothesis of a decreasing monotonic relationship for portfolio returns sorted on trade-weighted implied volatility at a 1% significance level. Other than for the three-week horizon, there are seemingly not much difference in monotonic relationship of portfolio returns between the realized and implied volatility metrics. Similar results are found in the first sub-period; see Table 10 in the appendix.

In the second sub-period, we reject only the p-values of a decreasing monotonic relationship for portfolio returns sorted on one-year and three-week realized volatility; see Table 11 in the appendix.

Table 3: Monotonicity test: p-values 2009-2023

Vol metric	H_0: *increasing*	H_0: *decreasing*
1-month RV	*0,98*	*0,08*
1-month IV	*0,99*	*0,10*
1-year RV	*0,95*	*0,08*
1-year IV	*0,65*	*0,23*
3-week RV	*0,82*	*0,07*
Trade-weighted IV	*0,45*	*0,01*

5.3 Excess information content in implied volatility

5.3.1 Double sort of IV and CPIV

When performing a conditional double sort of implied volatilities and their call-put implied volatility spread, the results show differences between the horizons. For one-month implied volatility a positive call-put implied volatility spread gives higher returns than a negative spread. This applies across all volatilities, and the largest return difference between the spreads is in the most volatile portfolio. However, for the one-year horizon, there really is not much of a pattern. In the first, fourth, and last portfolio, the positive spread gives higher returns. For portfolio 2 and 3, a positive spread performs worse than a negative one. Lastly, for the trade-weighted implied volatility, a negative spread performs the best across all volatilities. This effect does however seem to dampen as the volatility is increased, since the difference in returns between the spread seemingly decreases as the volatility increases.

The excess returns gathered by the positive spread are subjected to a t-test with a confidence level of 0,95 to assess whether the results are statistically different from zero.

As seen in Table 4 below, a positive call-put implied volatility spread gathers on average positive excess annual return of the negative spread for the one-month and one-year implied volatility. The average excess returns gathered by the one-month implied volatility is statistically distinguishable from zero. This does not apply for the one-year implied volatility.

However, a negative excess return is found for the trade-weighted implied volatility. This negative excess return is also statistically distinguishable from zero. These excess returns are obtained through effectively going long the stocks with a positive call-put implied volatility spread, and short the stocks with negative spread.

Table 4: Excess returns gathered by positive CPIV 2009-2023

Vol metric	*Mean excess returns*	*H_0: Indistinguishable from zero (p-value)*
One-month IV	*2,32%*	*0,04*
One-year IV	*0,99%*	*0,32*
Trade-weighted IV	*-1,47%*	*0,03*

In the first sub-period, there are no clear pattern of positive spread outperforming negative spread for any implied volatility metrics, as well as no significant mean excess returns; see Figure 12 and Table 13 in the appendix.

In the second sub-period, there is a clear pattern of positive spread outperforming negative spread for one-month and one-year implied volatilities, both their mean excess returns are statistically different from zero; see Figure 13 and Table 14 in the appendix.

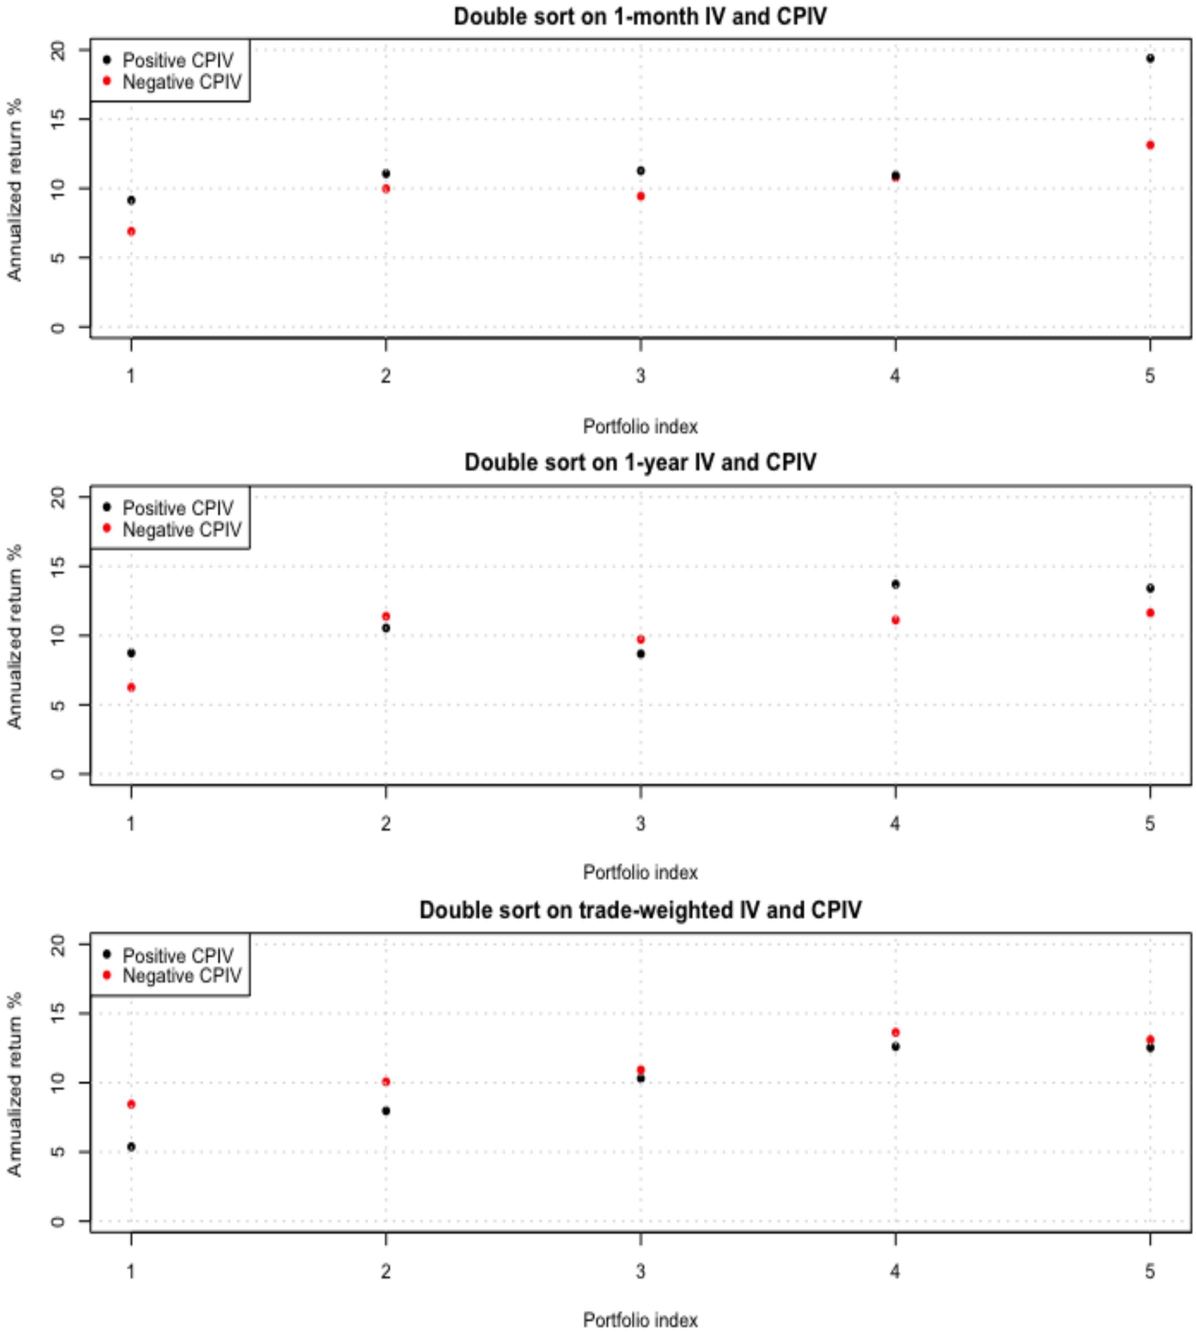

Figure 4: Double sorts of implied volatilities and call-put implied volatility spread 2009-2023

5.3.2 Long-short portfolios

Furthermore, to investigate whether or not the call-put implied volatility is a good proxy for anticipated increase or decrease in underlying price, long-short portfolios sorted on the spread as well as weighted by the spread are constructed. The higher (lower) call-put implied volatility spread generally earns higher (lower) returns. This applies to the call-put spread of all three implied volatility metrics. Similar results are found for the two sub-periods; see Figures 14 and 15 in the appendix. This result indicates predictive power of a positive spread signaling increase in underlying price, however not a negative spread signaling decrease in underlying price. If the call-put implied volatility spread was a good proxy of anticipated movement of both directions in underlying price, then the results should have shown a smile pattern.

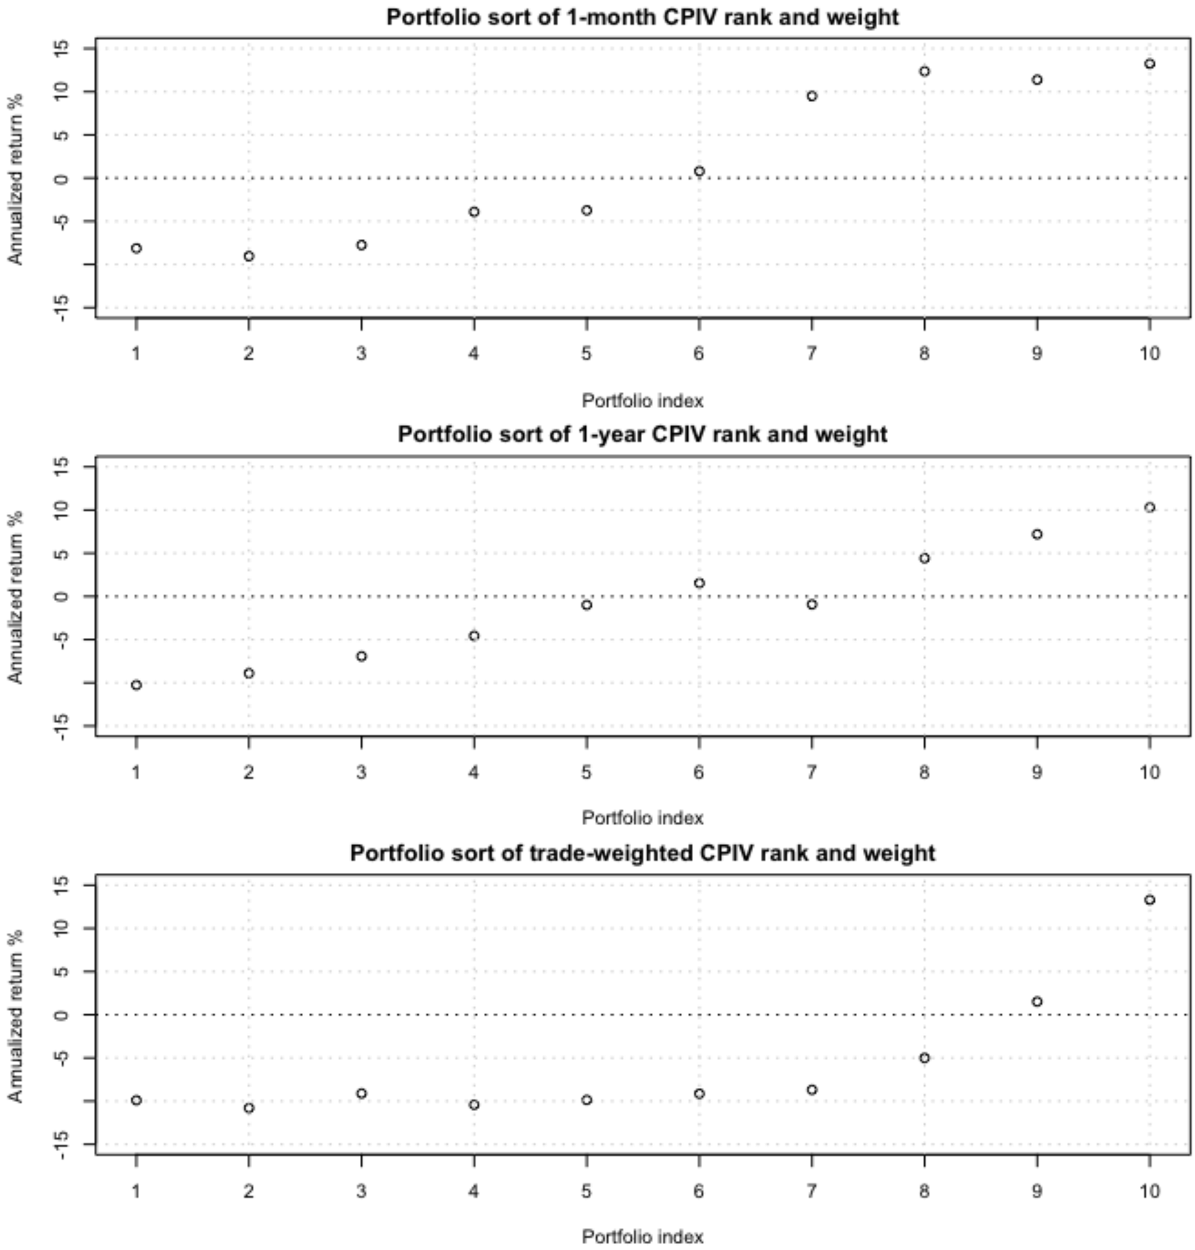

Figure 5: Portfolios sorted and weighted by call-put implied volatility spread 2009-2023

5.3.3 Double sort of IV and volatility risk premium

Double sorting on implied volatility and their implied-realized volatility spread reveals no clear pattern of predictability of returns for all three horizons across volatilities. Effectively going long the stocks with a positive spread and short the stocks with a negative spread gives an average annual return of 0,43% for one-month implied volatility, 0,07% for one-year implied volatility, and lastly -0,73% for trade-weighted implied volatility. These returns are inconsistent across volatilities, and are all not statistically different from zero, with p-values of 0,71, 0,94, and 0,55 respectively. Similar results are found for the two sub-periods; see Figures 16 and 17, and Tables 15 and 16 in the appendix.

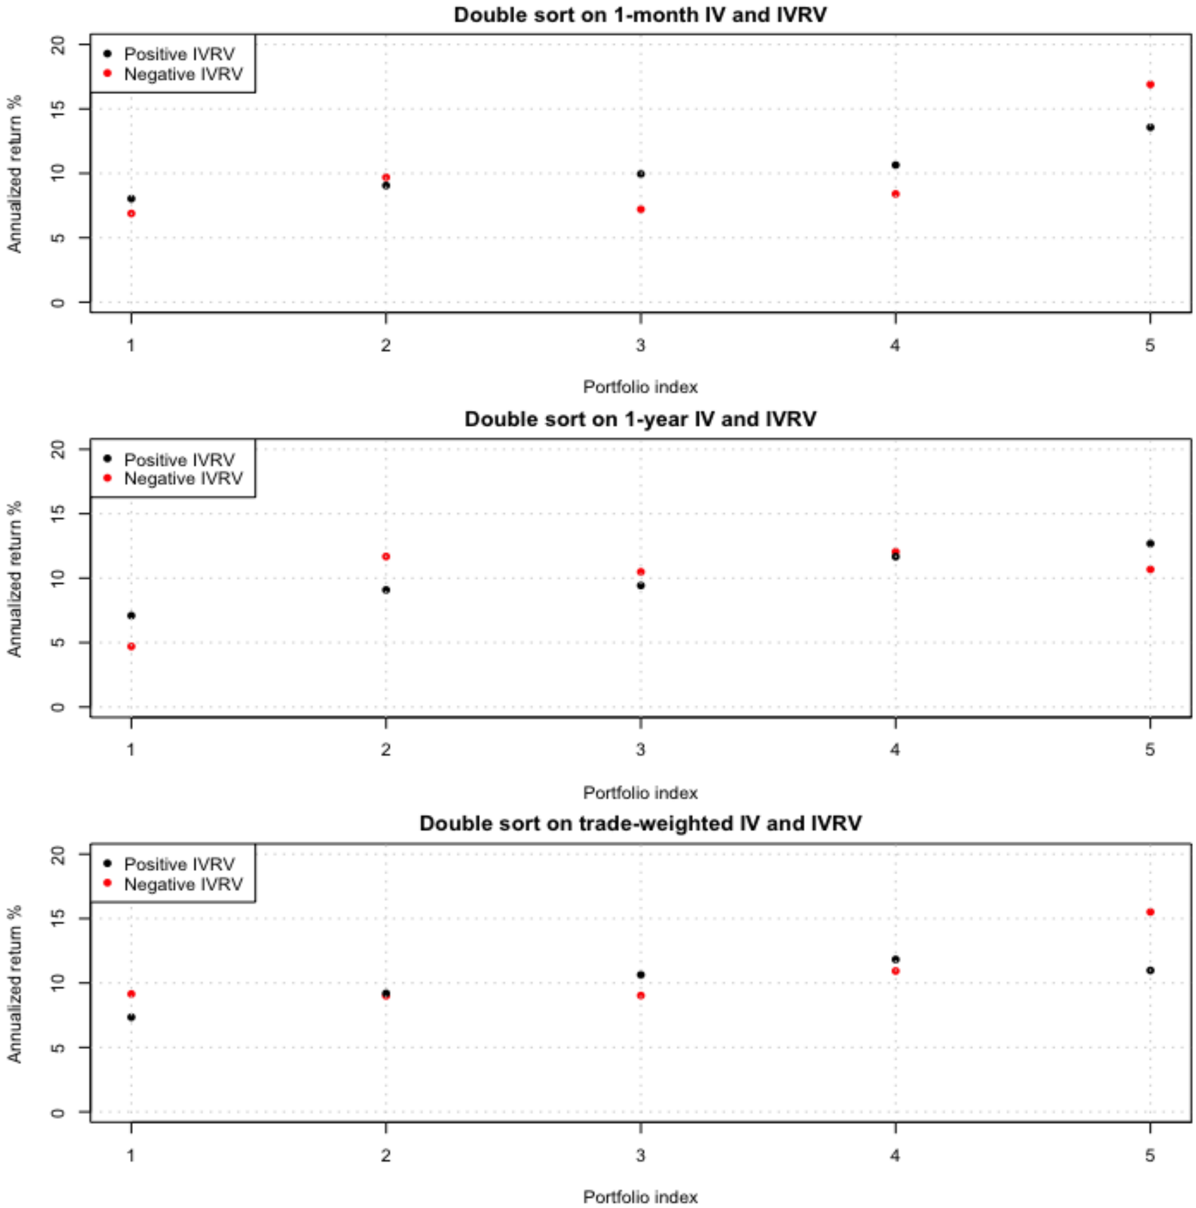

Figure 6: Double sort on implied volatilities and volatility risk premium 2009-2023

6. Discussion

6.1 Part 1: Data

Haugen and Heins (1972) had three main critiques to testing of the CAPM and how the models assumes a positive risk-return relationship. Although the CAPM is not of interest in this paper, a positive risk-return relationship is still a widely accepted theory, hence a negative risk-return relationship being called anomalous. The first of their critiques were the "impossible problem" of using ex-post values to determine expected values. The purpose of using implied volatility instead of realized volatility is indeed to circumvent this "impossible problem". The implied volatility is supposedly an ex-ante value of future volatility, but the question of whether or not it does a good job of that remains. When running cross-correlation between implied and realized volatility for all stocks in the dataset, the result is that an overwhelming majority of the highest correlation is at lag 0, and not at lag -1 as it theoretically should be considering implied volatility anticipates future volatility. It seems that implied volatility still is heavily dependent on lagged realized values. Apparently, the implied volatility is not leading the realized volatility. On the contrary, the realized volatility is leading the implied volatility.

Table 5: Cross-correlations 2009-2023

Volatility metric	*Amount of highest correlations*	
	Lag 0	*Lag -1*
1-month IV - RV	*361*	*88*
1-year IV - RV	*199*	*138*
Trade-weighted IV - 3-week RV	*361*	*40*

Regarding Haugen and Heins' second critique, the sample timing might be skewed due to the market situation of that time period. My dataset contains values from 2009 throughout 2023. This includes the aftermath of the financial crisis of 2008, which was generally bullish up until the COVID-19 pandemic of 2020, and the stock market decline of 2022. I also have divided the sample period into two sub-periods of July 2009 to October 2016 and October 2016 throughout 2023 for robustness. These two periods are almost equal in terms of percentage growth for the S&P 500 index;

however, the second sub-period includes two major stock market declines of 2020 and 2022 (SPX, 2024). Although both periods have been generally bullish in total, the second sub-period has been more volatile, which may be the reason for differences in the results for the two periods.

Lastly, the third critique of Haugen and Heins (1972) is that the data has bias towards survivorship. Considering my dataset contains mostly S&P 500 stocks, it obviously suffers from this bias. Many of the stocks that are now included in the S&P 500 were not in the S&P 500 in 2009. Therefore, these stocks in question were destined to succeed. There are however some "dead" stocks in my data universe, and the option data suffer more than stock data when it comes to missing data. Because I equalize all the data, there certainly should be some of the S&P 500 constituents that are effectively "dead" at different points during the period of 2009 and 2023. All in all, the sample does probably suffer of survivorship bias which might skew the results. However, had the data deliberately consisted of more "dead" companies and less popular stocks, the implied volatilities might have been less correct in the context of an anticipated future volatility. This is because a high traded volume is important in order for the data to contain the most information as possible. Consider a hypothetical where ten people bid on an options contract, and the market price is settled giving an observed implied volatility. This implied volatility would theoretically only be an anticipated future volatility by those ten people. If this anticipation is irrational or incongruent with the market, this would of course be arbitraged away in the end reflecting the market's anticipation. However, how much input is needed to sufficiently proxy the market's anticipation? Extrapolating Roll's (1977) critique of asset pricing theory tests, one would need to know the true market's anticipation, that is effectively the aggregate anticipation of every single person. Either way, traded volume is in this case effectively a measurement of "shared opinions" of future anticipation.

6.2 Part 2: Low-volatility anomaly and risk-return relationship

The first result from sorting realized volatility during the time period t, and capturing returns during the same time period found a presence of a low-volatility anomaly in this data universe. Conversely, sorting realized volatility during the time period of t and capturing returns over the time period $t + 1$ found the opposite. This shows that while the least volatile stocks have simultaneously earned higher returns, the same stocks will not earn higher returns over the following month. This begs the question of whether the stocks with the higher returns will have lower volatility, or whether it is the least volatile stocks that gather the highest returns. Either way, trying to exploit the low-volatility anomaly by betting on the currently observed least volatile stocks, which is measured by the past period, does work. Neither does betting on the current implied volatility, as this supposed forward-looking metric seems to be biased towards lagged measurement of realized volatility.

As sorting on both implied and realized volatility and capturing returns over the following month gave similar return structure across portfolios, the question of which volatility metric is more “theoretically” consistent in terms of risk-return relationship remains. As the portfolios are using total volatility as “risk” metrics, the risk-return relationship is measured by the Sharpe ratio (unlike Jensen’s alpha, it does not depend upon an asset pricing model such as the CAPM). For both the one-month volatilities the peak Sharpe ratio is in a low-volatility portfolio, suggesting that a low-volatility anomaly can technically be attributed to this sample, as these portfolios had greater risk-adjusted returns. This cannot be said for the one-year volatility metrics, as the Sharpe ratio peaks at different degrees of volatile portfolios. For the three-week realized volatility, the peak Sharpe ratio is in a low volatility portfolio and declines as the volatility increases. Conversely, for the trade-weighted implied volatility, the Sharpe ratio kept increasing up until the semi-volatile portfolio 7. However, the peak Sharpe ratio does not accurately explain which of the volatility metrics provide the most “theoretically” consistent risk-return relationship. The Sharpe ratios of portfolios sorted on one-month and trade-weighted implied volatility metrics had a higher and more neutral slope coefficient than their realized equivalent, which on top of that were statistically non-zero. The slope coefficient of the one-year implied and realized volatility were both statistically indifferent from zero and had similar residual standard error.

The monotonicity test of portfolio returns fails to reject the null hypothesis of an increasing monotonic relationship for all volatility metrics. The same applies for the null hypothesis of a decreasing monotonic relationship assuming a confidence level of 0,95, with the exception of trade-weighted implied volatility. Although the p-values of monotonic relationship for the other implied volatility metrics are generally "less statistically significant" than their realized equivalent, it is incorrect to declare them as such. The test is concerning a hypothesis, and whether to reject or fail to reject the null. Whether one is "more" statistically significant than the other is unknown.

Combining the results from the Sharpe ratios and the test of monotonic relationship, there are no clear differences in the risk-return relationship between realized and implied volatility metrics across the board, as has been found in previous literature. Buss and Vilkov (2012) concluded that their constructed option-implied betas give a positive and monotone risk-return relationship, that is significantly better than the historical betas in this regard. Additionally, Mateus and Kongsilp (2014) find that implied idiosyncratic volatility is the best predictor of stock returns among the volatility metrics they tested. Based on this we may expect the total implied volatility, which should be made up of the hypothetical implied beta and implied idiosyncratic volatility, to give similar results. Yet, my results suggest the implied volatility metrics are not altogether "better" than realized volatility in terms of risk-return relationship. However, because the implied beta and implied idiosyncratic volatility are both modelled and not directly observed, then perhaps these two modelled metrics combined still leave out an unexplained piece of information.

Looking at each volatility metric separately instead of categories of implied and realized volatility, there are some statistically significant differences. The implied volatility is more consistent to a traditional risk-return relationship at least for the shorter measurement and implied periods as shown by the slope coefficients of the Sharpe ratio regression. As argued in section 4.2.1, the trade-weighted implied volatility was assumed to be the most precise stock implied volatility, as it is interpolated between different strikes, rather than observed at-the-money. This is the only metric that rejects a decreasing monotonic relationship, results in a peak Sharpe ratio in a volatile portfolio, and in addition gives a preferred risk-return relationship in regard to the regression line of Sharpe ratios.

6.3 Part 3: Excess information of implied over realized volatility

6.3.1 Call-put implied volatility spread

The individual stock implied volatilities were calculated as the average of the call and put implied volatilities. These are not always identical, which by averaging them out may leave out important information about future stock returns. This begs the question of whether simply averaging the call and put implied volatility is the best way of determining individual stock implied volatility. Consider a case where a call and put for one stock option both have 30% implied volatility, while another stock option has a call and put implied volatility of 35% and 25% respectively. The average implied volatility for both stock options would be 30%, however the call-put implied volatility spread would be 0% for the first case, and 10% for the latter case. Perhaps a low average implied volatility with a near-zero spread between the call and put should be classified as low volatility, while the same average implied volatility however with a large spread should be classified as more volatile. The general standard practice seems to just be averaging the call and put implied volatility, however the information lost in the process has been shown to contain additional predictions of stock returns.

Cremers and Weinbaum (2010) found that the degree of predictability by the call-put implied volatility spread has decreased over time due to reduced mispricing, trading costs, and increased hedge fund capital. They used a sample period of 1996-2000 and 2001-2005, while my sample period goes from 2009-2023. Extrapolating from Cremers and Weinbaum (2010), the degree of predictability should be insignificant for my sample period.

What I find in the double sort is that for one-month implied volatility, a positive call-put implied volatility spread outperforms the negative spread across all volatilities. The return difference between the spreads are statistically significant. The degree of predictability is even stronger for the second sub-period, suggesting that the degree of predictability may be stronger in times of market decline. However, the same cannot be said for the one-year implied volatility, as this return difference is not statistically significant for the full period, nor is it consistent across volatilities. The call-put implied volatility spread was hypothesized as a proxy for anticipated directional movement in price. Presumably, the spread has less predictive power for the one-year implied volatility, because anticipating a full year of stock price movement is significantly more uncertain than for one month only. Following this logic, one can assume that people that do anticipate stock price movement, do so for as short of a period as possible. This logic is in accordance with findings from Eaton et al. (2022), that retail option traders prefer buying short-dated and out-of-the-money calls and writing long-dated puts.

The trade-weighted implied volatility is implied of a shorter period than a month, so why does a positive spread not predict an increase in stock price? Why does the negative spread predict the increase in stock price better than the positive spread? The probable answer is because the implied volatilities are trade-weighted and interpolated between strikes. This means that the call implied volatility of a contract can effectively be for a completely different strike than the put implied volatility. The put-call parity does not hold for contracts of different strikes, which means that the call-put implied volatility spread for the trade-weighted metric is not deviation of the put-call parity, it is instead expected. Using this spread when researching anticipated increase or decrease in stock price is therefore futile.

6.3.2 Long-short portfolios

Looking at the long-short portfolio sort, where the stocks are ranked as well as weighted based on their call-put implied volatility spread, there is no indication of a smile pattern of the portfolio's returns. Considering the results of the double sort, accompanied by the results of the long-short portfolio, it seems that the positive spread predicts future stock returns. However, this does not go both ways. The negative spread should have predicted negative returns, meaning that shorting these stocks would yield positive returns. Yet my results suggest that using a negative spread to predict negative returns does not work.

6.3.3 Volatility risk premium

Using the square-root of time rule to annualize the standard deviation, or realized volatility, amplifies the fluctuation of realized volatility (Diebold et al., 1998). This suggests that the raw difference of implied and realized volatility is not what it “should” be. For example, a stock can be considered to have a negative (positive) spread, however the “real” spread may be positive (negative). Therefore, the raw difference of implied and realized volatility may not be a good proxy for the volatility risk premium. Bollerslev et al. (2004) briefly mentions that this raw difference is very noisy, but I have not accounted for this noise in my analysis as it would entail using “model-free” volatility metrics. This is outside the scope of this thesis.

Both Bali and Hovakimian (2009) as well as Bollerslev et al. (2009) found that the implied-realized volatility spread, or the volatility risk premium predicts future stock returns. Despite these previous results, the implied-realized volatility spread did not seem to predict stock returns in my data universe.

Bali and Hovakimian (2009) did use a different calculation of the spread than what I presented. First off, they subtracted the implied volatility from the realized volatility, which just inverts the results. Secondly, they calculated the volatility risk premium using realized and implied volatility both during the time period t. I have argued in section 2.2.7 that the realized volatility during the time period t is comparable to the implied volatility during the time period $t-1$. This is due to the realized volatility being measured over the previous period, while the implied volatility is implied over the following period. The volatility risk premium is supposed to reflect the premium option writers demand to offset the potential risk of stock price movement. If the option writers and buyers knew the realized volatility in advance, then there would be no premium. No one would pay for insurance that they knew would never be claimed.

Bollerslev et al. (2009) recognizes this, and instead forecasts the future realized volatility using a heterogeneous autoregressive model. Because of this forecast, they are able to capture returns at time t, leading to a result of higher (lower) volatility risk premium earning higher (lower) returns.

I instead use the observed realized volatility at time t and the implied volatility at time $t-1$, meaning that to use the volatility risk premium as a strategy, the returns have to be captured at $t+1$. This is two periods ahead of the implied volatility, and presumably whatever information the volatility risk premium had over the period of $t-1$ to t should already be priced in before the period of t to $t + 1$. Not surprisingly, the results show no pattern with respect to returns.

Had I instead captured returns at time t, which would correspond to an infeasible trading strategy, then the results would have been similar to what is found by Bollerslev et al. (2009):

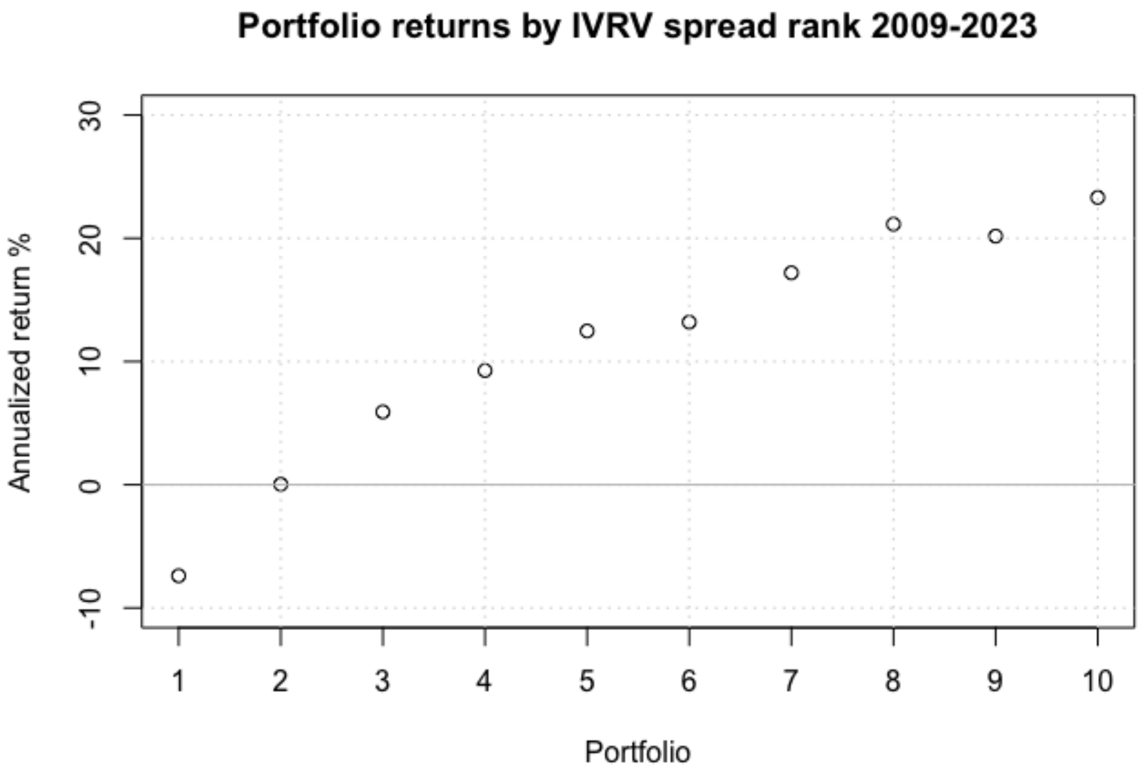

Figure 7: Portfolio sort of one-month volatility risk premium 2009-2023, capturing returns at time t

7. Conclusion

This thesis aims to answer a main research question, containing three objectives. This thesis has investigated whether a low-volatility anomaly is present in the American equity market, and whether it can be exploited by either implied or realized volatility metrics. Additionally, the risk-return relationship of implied and realized volatility divided into a total of six different volatility metrics has been studied and compared. Lastly, the thesis researches the possible predictability of the excess information content of implied volatility over realized volatility.

To answer these questions the implied and realized volatility as well as the returns of 513 stocks in the American equity market during the period of 2009-2023 has been gathered and analyzed as described in section 4.

To begin with, the thesis concludes that portfolios sorted on stocks with the lowest volatility has simultaneously earned the most returns. Yet, the anomaly cannot be exploited using either implied or realized volatility as forecasts of future volatility.

Secondly, the thesis concludes that sorting portfolios on implied volatility does not give a more theoretically consistent risk-return than portfolios sorted on realized volatility overall. However, out of the six total volatility metrics, sorting portfolios on trade-weighted implied volatility gives the most theoretically consistent risk-return relationship. This is concluded due to its Sharpe ratio regression's slope coefficient is consistently indistinguishable from zero, and it is the only metric which rejects a decreasing monotonic relationship of risk and return for both the full sample period and sub-sample period of 2009-2016. However, the three-week realized volatility seems to be better in this regard when major stock market declines constitutes a larger portion of the time period.

Lastly, the thesis concludes that the call-put implied volatility spread of one-month options contracts predicts future stock returns in the American market. The predictability is also stronger when major stock market declines constitutes a larger portion of the time period. However, positive spread does predict positive future stock returns, while negative spread does not predict negative future stock returns. The thesis found no predictability of the volatility risk premium proxied by the implied-realized volatility spread for any period or implied volatility metric.

Finally, to answer the main research question:

Is option-implied volatility theoretically superior to realized volatility in terms of linear risk-return relationship, and does its excess information content predict future stock returns?

The trade-weighted implied volatility generally gives the most theoretically consistent risk-return relationship out of the six metrics analyzed, and the call-put implied volatility spread of one-month options contracts does predict future stock returns. The proxy used for volatility risk premium does however not predict future stock returns.

References

Ang, A., Hodrick, R. J., Xing, Y. & Zhang, X. (2006). The Cross-Section of Volatility and Expected Returns. *The Journal of Finance*, *61*(1), 259–299 https://doi.org/10.1111/j.1540-6261.2006.00836.x

Baker, N. L. & Haugen, R. A. (2012). Low Risk Stocks Outperform within All Observable Markets of the World. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.2055431

Bali, T. G. & Hovakimian, A. (2009). Volatility Spreads and Expected Stock Returns. *Management Science*, *55*(11), 1797–1812. http://www.jstor.org/stable/40539244

Baronayan, S. & Rothbarth, C. (2019). *The global evolution of low-volatility investment in asset management.* EPFR. https://epfr.com/insights/papers/global-evolution-low-volatility-investment-in-asset-management/

Black, F. & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, *81*(3), 637–654. http://www.jstor.org/stable/1831029

Blitz, D. C. & van Vliet, P. (2007). The Volatility Effect: Lower Risk Without Lower Return. *Journal of Portfolio Management, 34(1),* 102-113. DOI: 10.3905/jpm.2007.698039

Bollerslev, T., Gibson, M. & Zhou, H. (2004). Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics, 160(1)*, 235-245. https://doi.org/10.1016/j.jeconom.2010.03.033

Bollerslev, T., Tauchen, G. & Zhou, H. (2009). Expected Stock Returns and Variance Risk Premia. *The Review of Financial Studies, 22(11)*, 4463–4492. http://www.jstor.org/stable/40468365

Buss, A., & Vilkov, G. (2012). Measuring Equity Risk with Option-Implied Correlations. *Review of Financial Studies, 25(10)*, 3113-3140. https://doi.org/10.1093/rfs/hhs087

Chicago Board Options Exchange. (n.d). *Equity Options Product Specifications.* Cboe. https://www.cboe.com/exchange_traded_stock/equity_options_spec/

Cox, J.C., Ross, S.A. & Rubinstein, M. (1979) Option Pricing: A Simplified Approach. *Journal of Financial Economics, 7(3)*, 229-263. https://doi.org/10.1016/0304-405X(79)90015-1

Cremers, M. & Weinbaum, D. (2010). Deviations from Put-Call Parity and Stock Return Predictability. *The Journal of Financial and Quantitative Analysis*, *45*(2), 335–367. http://www.jstor.org/stable/27801488

de Carvalho, R. L., Zakaria, M., Xiao, L. & Moulin, P. (2015). Low Risk Anomaly Everywhere - Evidence from Equity Sectors. In E. Jurczenko (Ed.), *Risk-Based and Factor Investing* (pp. 265-289). ISTE Press – Elsevier. https://doi.org/10.1016/B978-1-78548-008-9.50011-3

Diebold, F.X., Hickman, A., Inoue, A. & Schuermann, T. (1998). *Converting 1-Day Volatility to h-Day Volatility: Scaling by Root-h is Worse than You Think*. Wharton Financial Institutions Center. Working Paper 97-34. Retrieved from: https://www.sas.upenn.edu/~fdiebold/ResearchPapersChronological.htm

Eaton, G. W., Green, T. C., Roseman, B. S. & Wu, Y. (2022). Retail Option Traders and the Implied Volatility Surface. Working paper. https://dx.doi.org/10.2139/ssrn.4104788

Eraker, B. (2021). The volatility premium. *The Quarterly Journal of Finance*, *11(3)*, 1-35. https://doi.org/10.1142/S2010139221500142

Fama, E. F. & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics, 33(1)*, 3-56. https://doi.org/10.1016/0304-405X(93)90023-5

Frazzini, A. & Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics, 2014, 111(1)*, 1-25. https://doi.org/10.1016/j.jfineco.2013.10.005

Gârleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-Based Option Pricing. *The Review of Financial Studies, 22(10),* 4259–4299. http://www.jstor.org/stable/40468358

Haug, E. G. & Taleb, N. N. (2011). Option traders use (very) sophisticated heuristics, never the Black–Scholes–Merton formula. *Journal of Economic Behavior & Organization, 77(2)*, 97-106. https://doi.org/10.1016/j.jebo.2010.09.013

Haugen, R. A. & Heins, A. J. (1972). On the Evidence Supporting the Existence of Risk Premiums in the Capital. *ERN: Econometric Modeling in Financial Economics.* http://dx.doi.org/10.2139/ssrn.1783797

Lintner, J. (1965). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *The Review of Economics and Statistics, 47(1)*, 13–37. https://doi.org/10.2307/1924119

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, *7*(1), 77–91. https://doi.org/10.2307/2975974

Mateus, C. & Kongsilp, W. (2014). Implied Idiosyncratic Volatility and Stock Return Predictability. *Journal of Mathematical Finance, 4*, 338-352. DOI: 10.4236/jmf.2014.45032.

Mayhew, S. (1995). Implied Volatility. *Financial Analysts Journal*, *51*(4), 8–20. http://www.jstor.org/stable/4479853

Options Industry Council. (n.d). *Options Glossary.* Options Education. https://www.optionseducation.org/referencelibrary/optionsglossary?filter=S

Patton, A. J. & Timmermann, A. (2010). Monotonicity in asset returns: New tests with applications to the term structure, the CAPM, and portfolio sorts. *Journal of Financial Economics, 98(3)*, 605-625. https://doi.org/10.1016/j.jfineco.2010.06.006

Rodriguez, J. L., Fang, X. & Li, Q. (2015). *Implied Volatility for Options on Futures Using the Cox-Ross-Rubinstein (CRR) Model* [Project]. Loyola University Chicago. DOI:10.13140/RG.2.1.5052.0168

Roll, R. (1977). A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. *Journal of Financial Economics, 4(2)*, 129-176. https://doi.org/10.1016/0304-405X(77)90009-5

Romano, J. P. & Wolf, M. (2013). Testing for Monotonicity in Expected Asset Returns. *Journal of Empirical Finance, 23*, 93-116. https://doi.org/10.1016/j.jempfin.2013.05.001

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory, 13(3)*, 341-360. https://doi.org/10.1016/0022-0531(76)90046-6

Sharpe, W.F. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *Journal of Finance, 19(3)*, 425-442. https://doi.org/10.1111/j.1540-6261.1964.tb02865.x

Sharpe, W.F. (1965). Mutual Fund Performance. *The Journal of Business, 39(1)*, 119-138. http://www.jstor.org/stable/2351741

SPX. (2024). *S&P 500 INDEX (^SPX)*. Yahoo! Finance. Retrieved from https://finance.yahoo.com/quote/%5ESPX/

Tobin, J. (1958). Liquidity Preference as Behavior Towards Risk. *The Review of Economic Studies, 25(2)*, 65–86. https://doi.org/10.2307/2296205

van der Hoek, J. & Elliott, J. R. (2006). *Binomial models in finance*. Springer.

van Vliet, P. & de Koning, J. (2017). *High returns from low risk: a remarkable stock market paradox.* Wiley.

Appendix

Table 6: Portfolio annualized returns % 2009-2023

Vol	Portfolios										
metric	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*10-1*
RV1M	*8,23*	*10,46*	*10,82*	*10,61*	*10,43*	*11,47*	*11,37*	*10,75*	*15,14*	*15,21*	*6,98*
IV1M	*8,09*	*8,35*	*11,54*	*10,27*	*10,37*	*11,88*	*10,12*	*11,93*	*14,47*	*17,99*	*9,90*
RV1Y	*7,00*	*9,86*	*9,82*	*10,88*	*9,71*	*13,33*	*11,74*	*11,47*	*14,01*	*16,72*	*9,72*
IV1Y	*6,57*	*8,31*	*11,13*	*11,40*	*10,71*	*10,50*	*11,52*	*14,83*	*12,76*	*17,23*	*10,66*
RV3W	*9,03*	*9,11*	*11,60*	*10,10*	*11,56*	*11,36*	*11,81*	*11,02*	*13,11*	*15,57*	*6,54*
IVTW	*7,20*	*9,62*	*8,78*	*11,41*	*10,36*	*11,70*	*13,89*	*13,06*	*12,83*	*16,10*	*8,90*

Table 7: Standard deviation of portfolios 2009-2023

Vol	Portfolios									
metric	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
RV1M	*3,56*	*4,03*	*4,44*	*4,67*	*5,17*	*5,46*	*5,77*	*6,02*	*6,88*	*7,64*
IV1M	*3,53*	*4,01*	*4,31*	*4,75*	*5,16*	*5,41*	*5,75*	*6,11*	*6,65*	*8,17*
RV1Y	*3,82*	*3,94*	*4,23*	*4,55*	*5,24*	*5,25*	*5,61*	*6,00*	*6,78*	*8,17*
IV1Y	*3,70*	*3,89*	*4,31*	*4,70*	*5,18*	*5,29*	*5,68*	*6,01*	*6,88*	*8,25*
RV3W	*3,56*	*4,04*	*4,50*	*4,89*	*4,91*	*5,50*	*5,70*	*6,06*	*6,65*	*7,61*
IVTW	*3,63*	*4,09*	*4,57*	*4,86*	*5,13*	*5,32*	*5,74*	*6,00*	*6,80*	*7,22*

Table 8: Portfolio Sharpe ratios 2009-2023

Vol	Portfolios									
metric	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
RV1M	*2,05*	*2,36*	*2,22*	*2,07*	*1,83*	*1,93*	*1,81*	*1,63*	*2,06*	*1,86*
IV1M	*2,03*	*1,85*	*2,45*	*1,96*	*1,83*	*2,02*	*1,60*	*1,80*	*2,03*	*2,08*
RV1Y	*1,59*	*2,26*	*2,10*	*2,18*	*1,67*	*2,35*	*1,92*	*1,75*	*1,92*	*1,92*
IV1Y	*1,53*	*1,90*	*2,36*	*2,22*	*1,88*	*1,80*	*1,86*	*2,30*	*1,71*	*1,97*
RV3W	*2,27*	*2,02*	*2,37*	*1,87*	*2,16*	*1,89*	*1,91*	*1,66*	*1,83*	*1,92*
IVTW	*1,73*	*2,12*	*1,72*	*2,15*	*1,83*	*2,02*	*2,25*	*2,02*	*1,74*	*2,09*

Realized volatility and returns over the time period t

Sub-period 1: 2009-2016

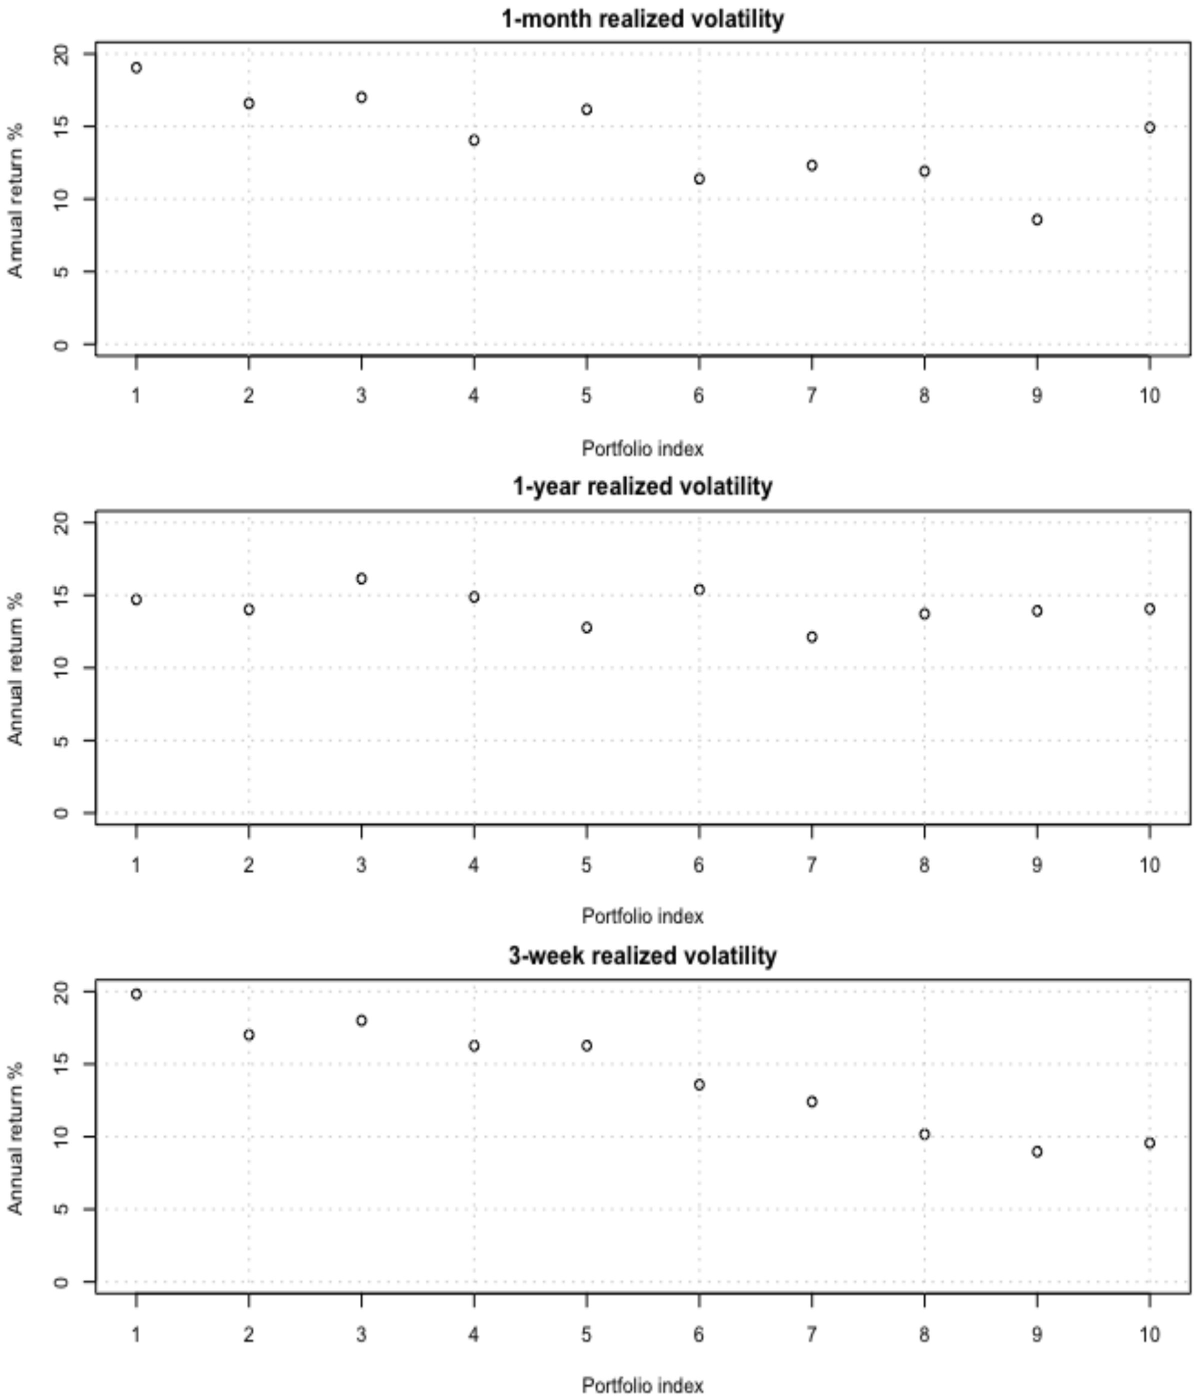

Figure 8: Portfolio sorts of realized volatilities 2009-2016, capturing returns at same month

Realized volatility and returns over the time period t

Sub-period 2: 2016-2023

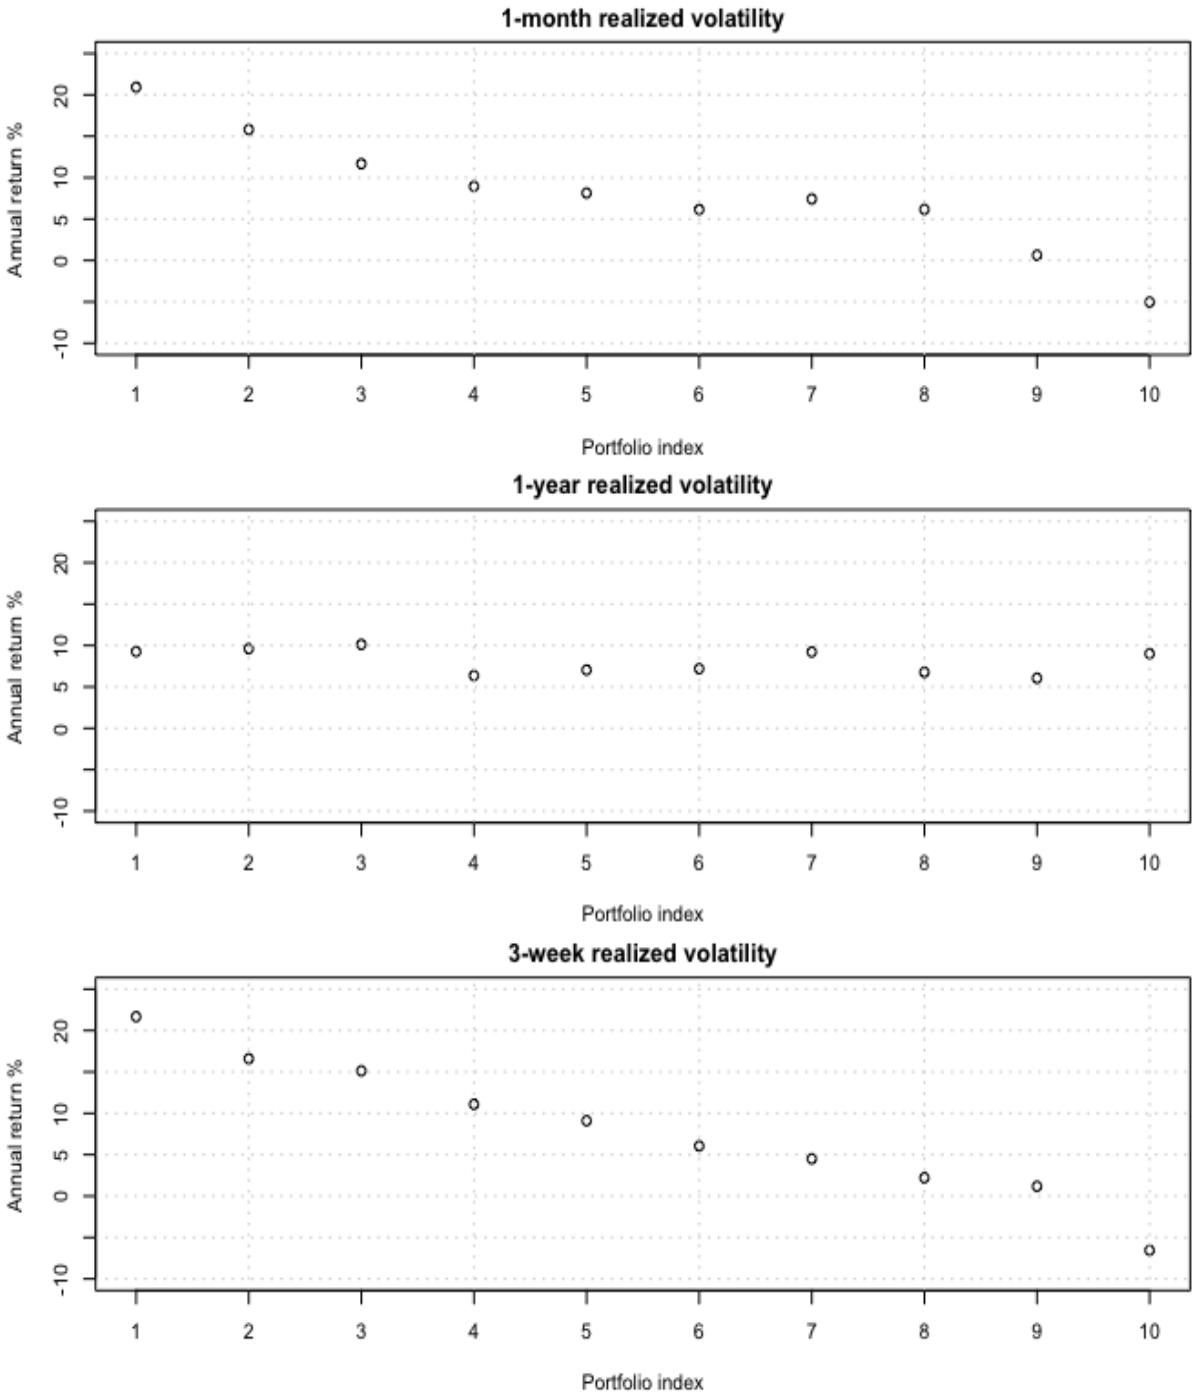

Figure 9: Portfolio sorts of realized volatility 2016-2023, capturing returns at same month

Realized and implied volatility, returns over the following period

Sub-period 1: 2009-2016

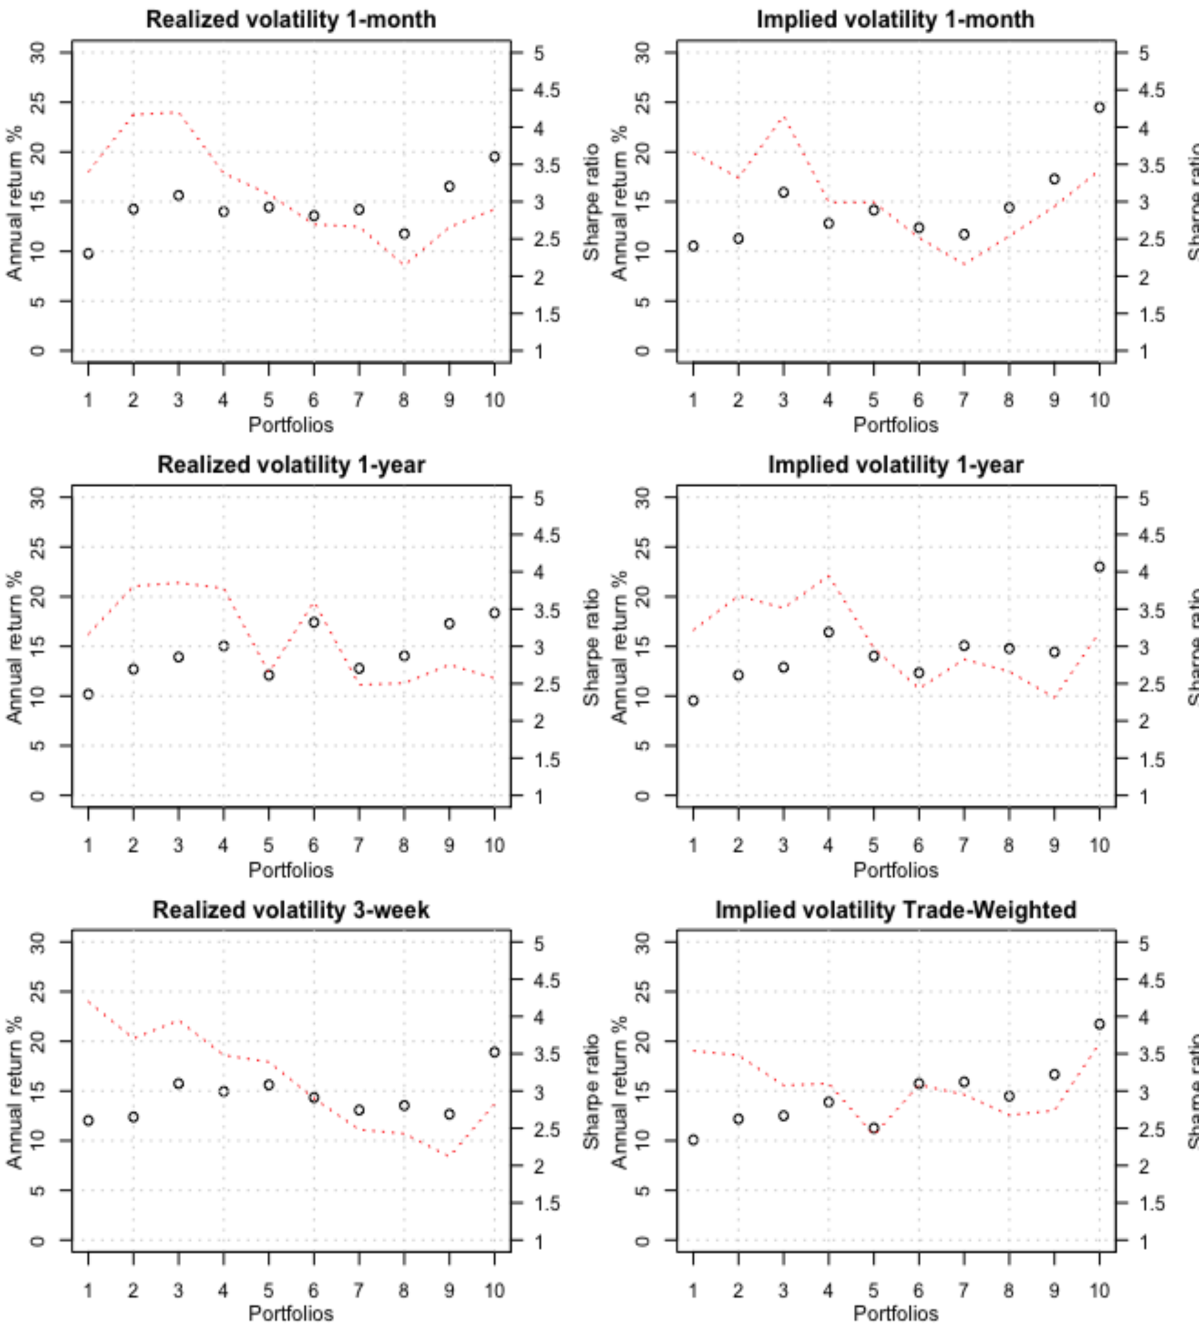

Figure 10: Portfolio sorts of realized and implied volatility 2009-2016, capturing returns at following month

Realized and implied volatility, returns over the following period

Sub-period 2: 2016-2023

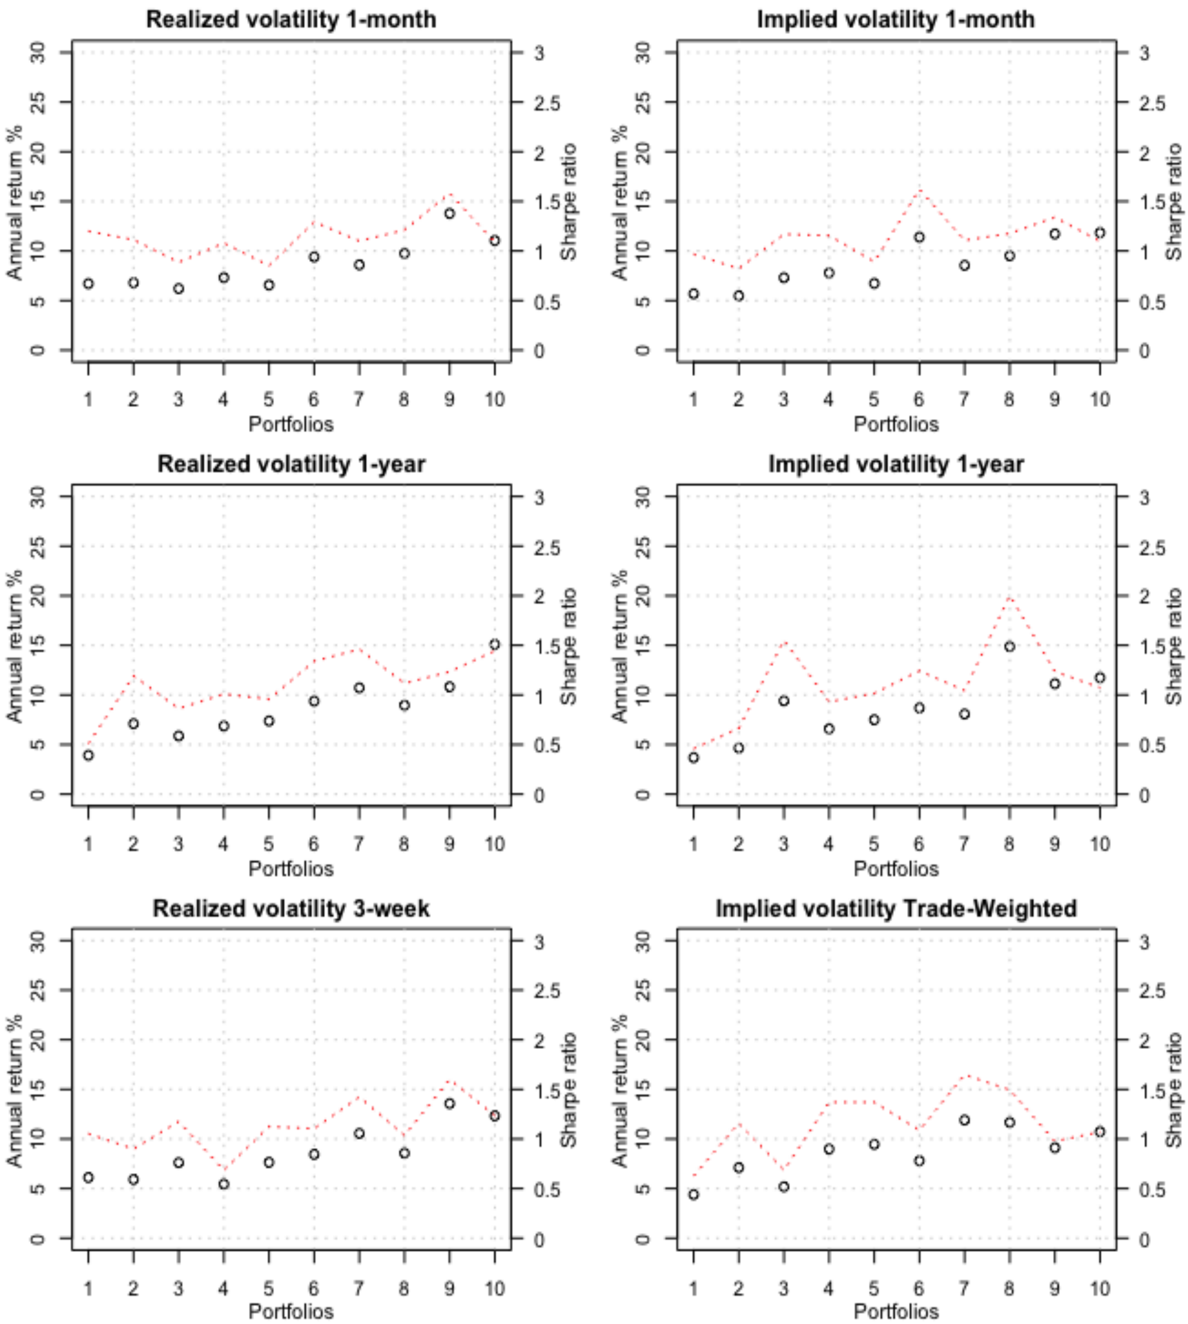

Figure 11: Portfolio sorts of realized and implied volatility 2016-2023, capturing returns at following month

Sharpe regression

Sub-period 1: 2009-2016

Table 9: Linear regression of Sharpe ratios (trendline) 2009-2016

Vol metric	*Intercept*	*Slope coefficient*	*Residual standard error*
1-month RV	*4,06*	*-0,17**	*0,45*
1-month IV	*3,60*	*-0,10*	*0,54*
1-year RV	*3,86*	*-0,14**	*0,45*
1-year IV	*3,67*	*-0,11*	*0,45*
3-week RV	*4,30*	*-0,21**	*0,32*
Trade-weighted IV	*3,27*	*-0,04*	*0,41*

** statistically distinguishable from zero*

Sub-period 2: 2016-2023

Table 10: Linear regression of Sharpe ratios (trendline) 2016-2023

Vol metric	*Intercept*	*Slope coefficient*	*Residual standard error*
1-month RV	*0,99*	*0,03*	*0,20*
1-month IV	*0,95*	*0,03*	*0,22*
1-year RV	*0,72*	*0,07**	*0,21*
1-year IV	*0,71*	*0,07*	*0,39*
3-week RV	*0,87*	*0,05*	*0,22*
Trade-weighted IV	*0,91*	*0,04*	*0,32*

** statistically distinguishable from zero*

Monotonicity

Sub-period 1: 2009-2016

Table 11: Monotonicity test: p-values 2009-2016

Vol metric	H_0: *increasing*	H_0: *decreasing*
1-month RV	*0,91*	*0,30*
1-month IV	*0,96*	*0,31*
1-year RV	*0,97*	*0,38*
1-year IV	*0,70*	*0,33*
3-week RV	*0,54*	*0,38*
Trade-weighted IV	*0,72*	*0,05*

Sub-period 2: 2016-2023

Table 12: Monotonicity test: p-values 2016-2023

Vol metric	H_0: *increasing*	H_0: *decreasing*
One-month RV	*0,91*	*0,13*
One-month IV	*0,77*	*0,11*
One-year RV	*0,77*	*0,04*
One-year IV	*0,27*	*0,12*
Three-week RV	*0,76*	*0,03*
Trade-weighted IV	*0,46*	*0,20*

Double sort of IV and CPIV

Sub-period 1: 2009-2016

Table 13: Excess returns gathered by positive CPIV 2009-2016

Vol metric	*Mean excess returns*	H_0*: Indistinguishable from zero (p-value)*
One-month IV	*1,73%*	*0,31*
One-year IV	*-0,77%*	*0,58*
Trade-weighted IV	*-0,97%*	*0,52*

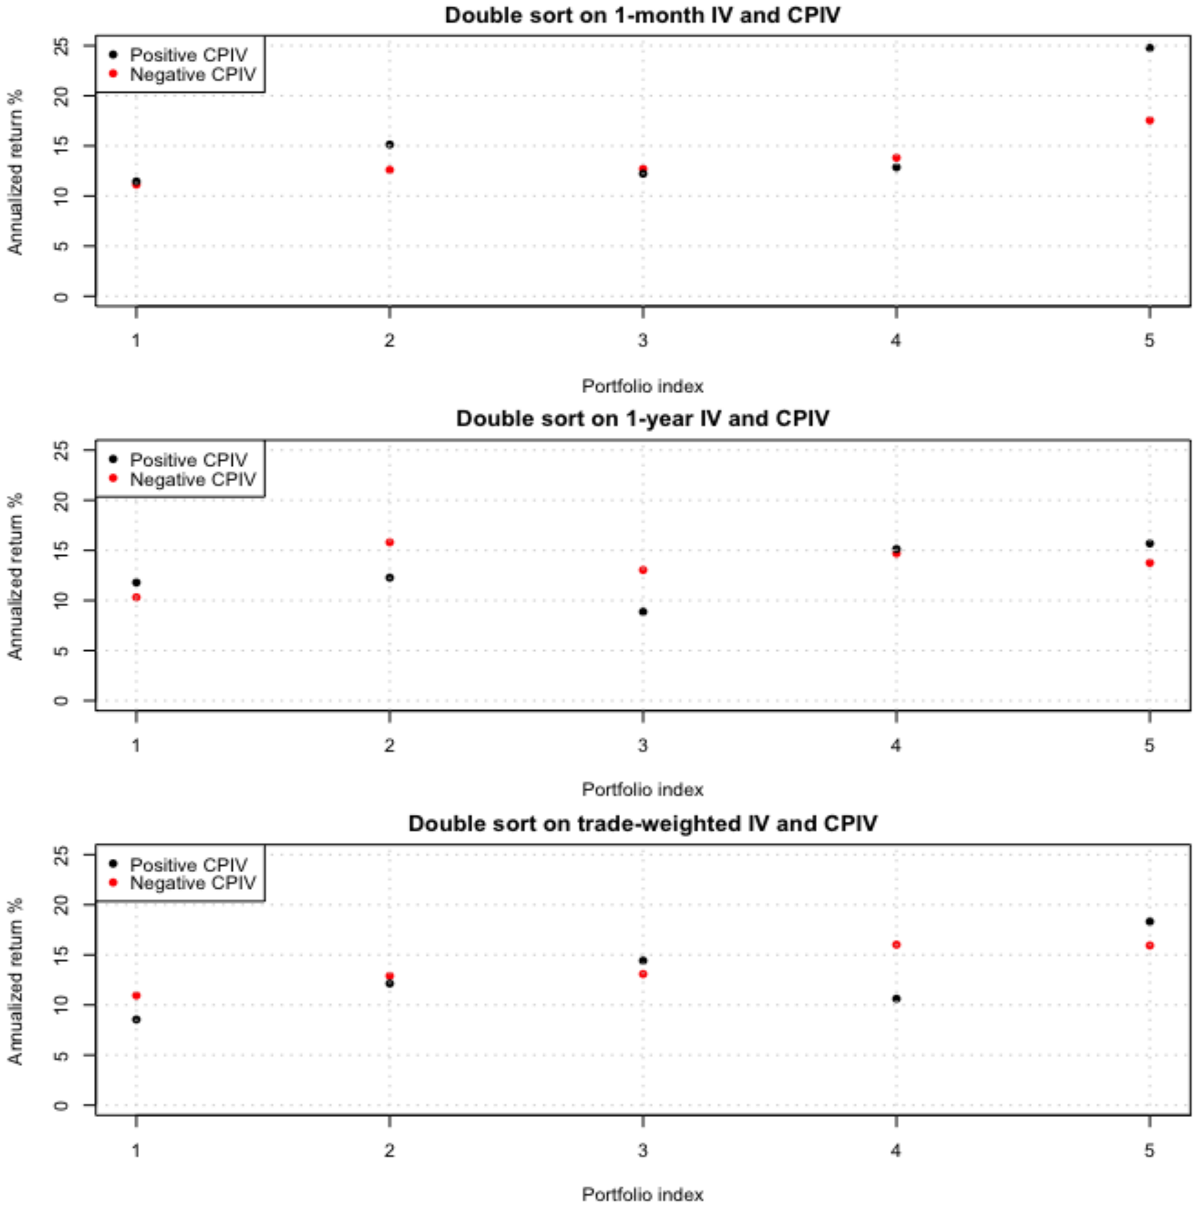

Figure 12: Double sorts of implied volatilities and call-put implied volatility spread 2009-2016

Double sort of IV and CPIV

Sub-period 2: 2016-2023

Table 14: Excess returns gathered by positive CPIV 2016-2023

Vol metric	*Mean excess returns*	H_0*: Indistinguishable from zero (p-value)*
One-month IV	*2,94%*	*0,05*
One-year IV	*2,67%*	*0,01*
Trade-weighted IV	*-1,90%*	*0,23*

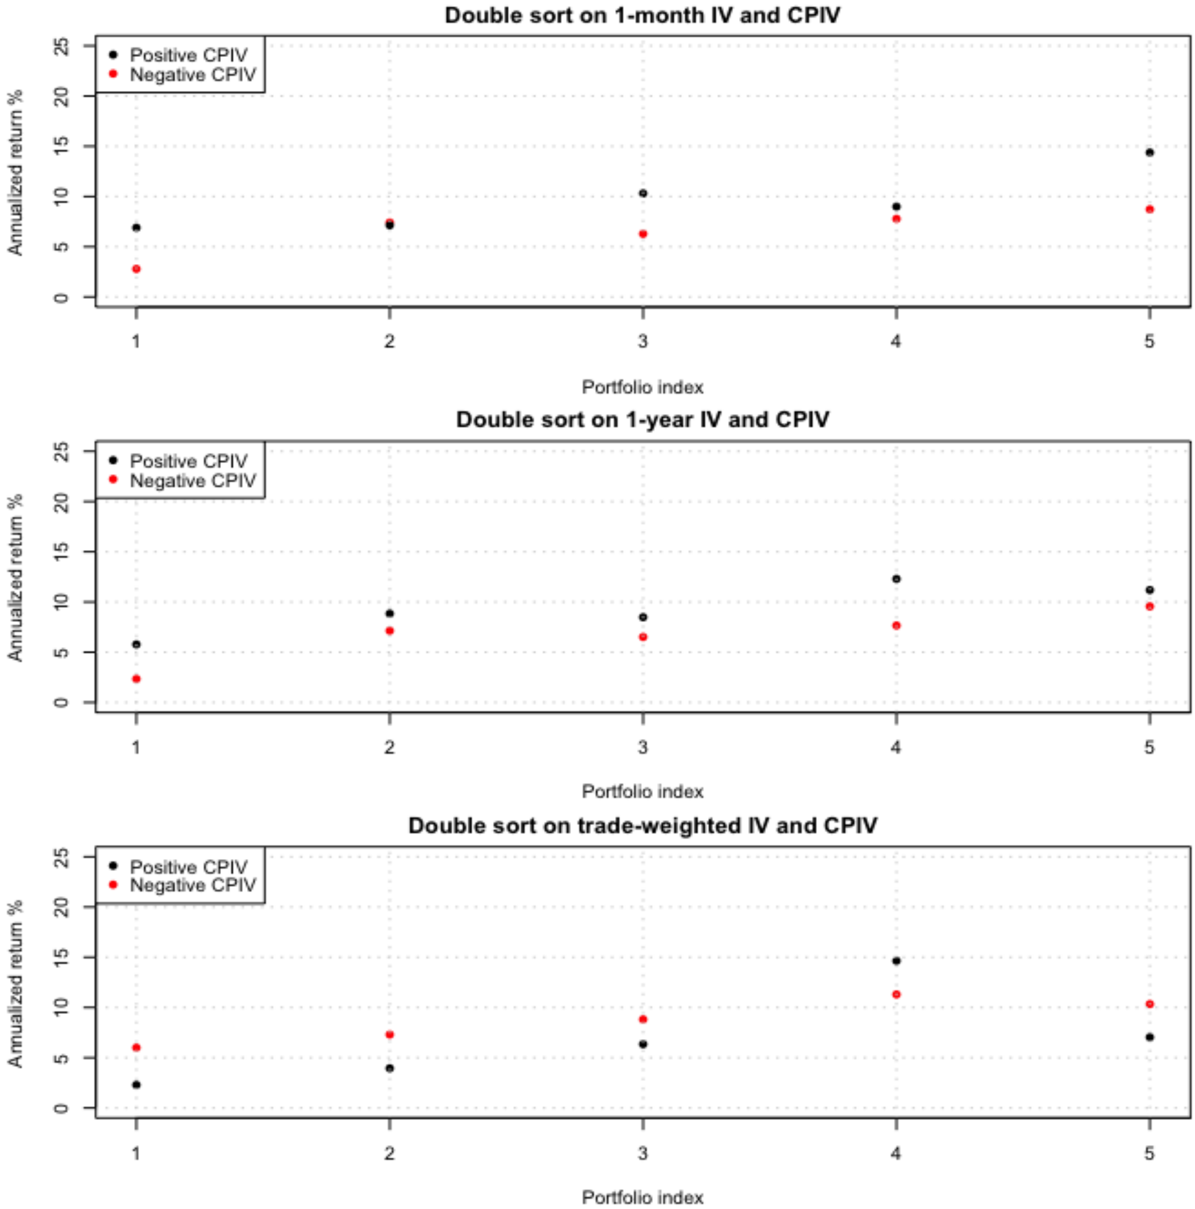

Figure 13: Double sorts of implied volatilities and call-put implied volatility spread 2016-2023

Long-short portfolios

Sub-period 1: 2009-2016

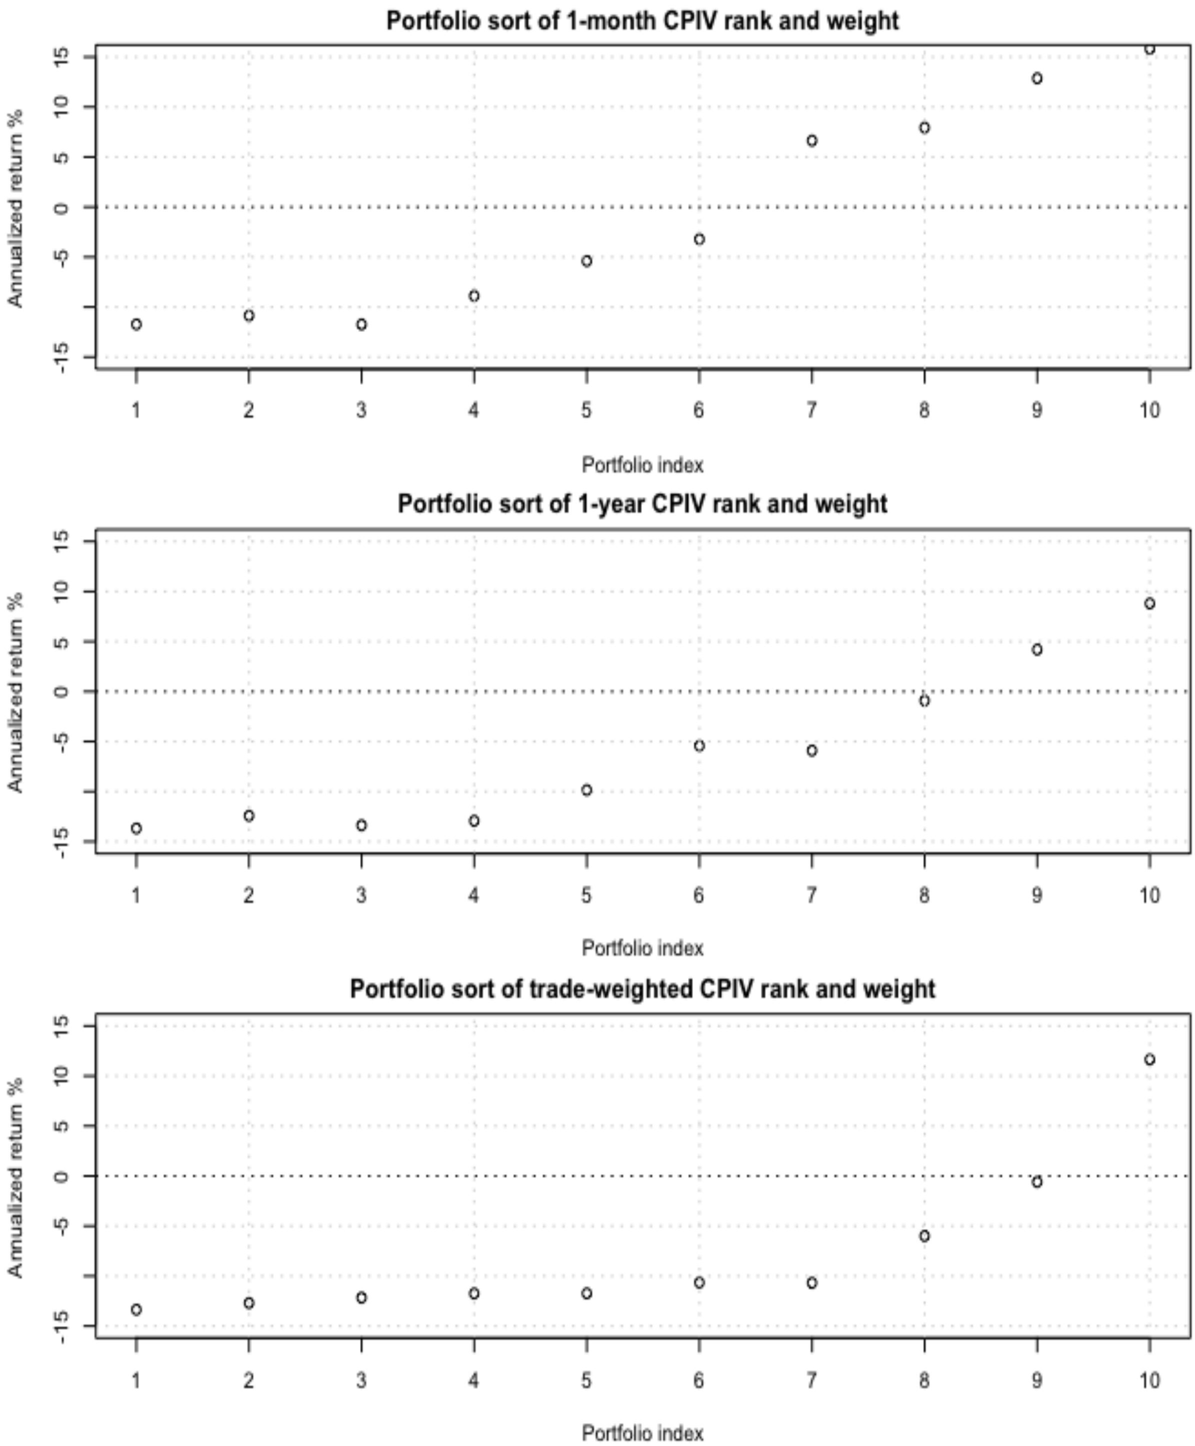

Figure 14: Portfolios sorted and weighted by call-put implied volatility spread 2009-2016

Long-short portfolios

Sub-period 2: 2016-2023

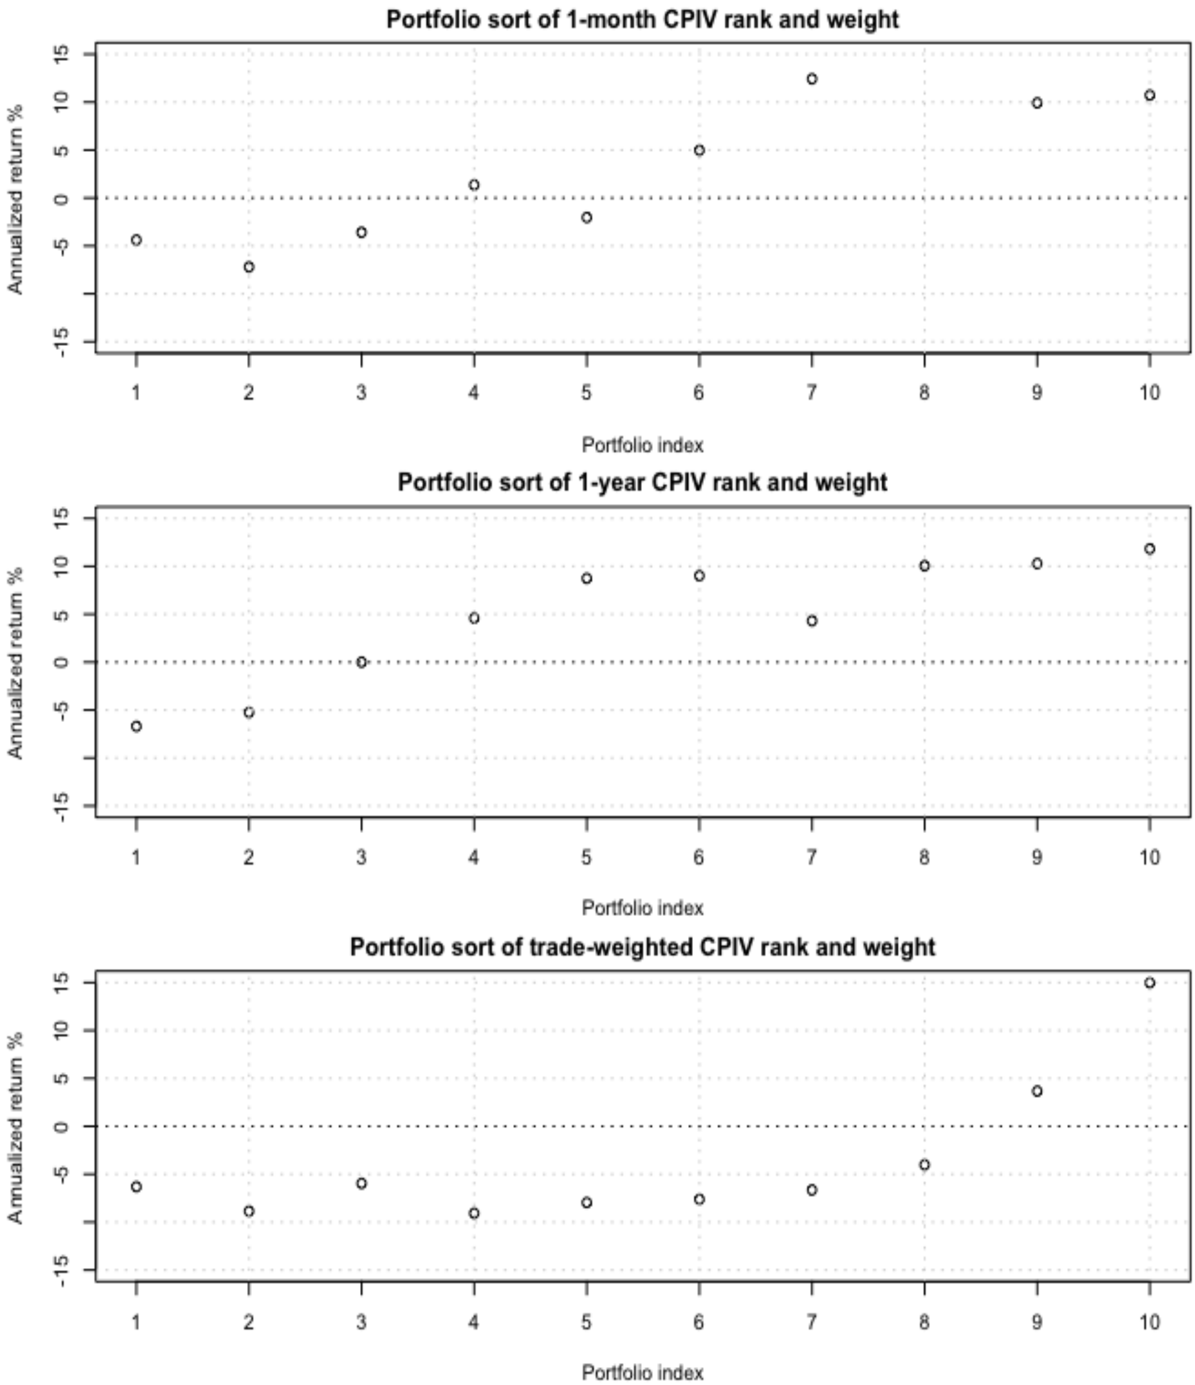

Figure 15: Portfolios sorted and weighted by call-put implied volatility spread 2016-2023

Double sort of IV and volatility risk premium

Sub-period 1: 2009-2016

Table 15: Excess returns gathered by positive IVRV 2009-2016

Vol metric	*Mean excess returns*	*H_0: Indistinguishable from zero (p-value)*
One-month IV	*1,50%*	*0,49*
One-year IV	*1,29%*	*0,68*
Trade-weighted IV	*2,43%*	*0,14*

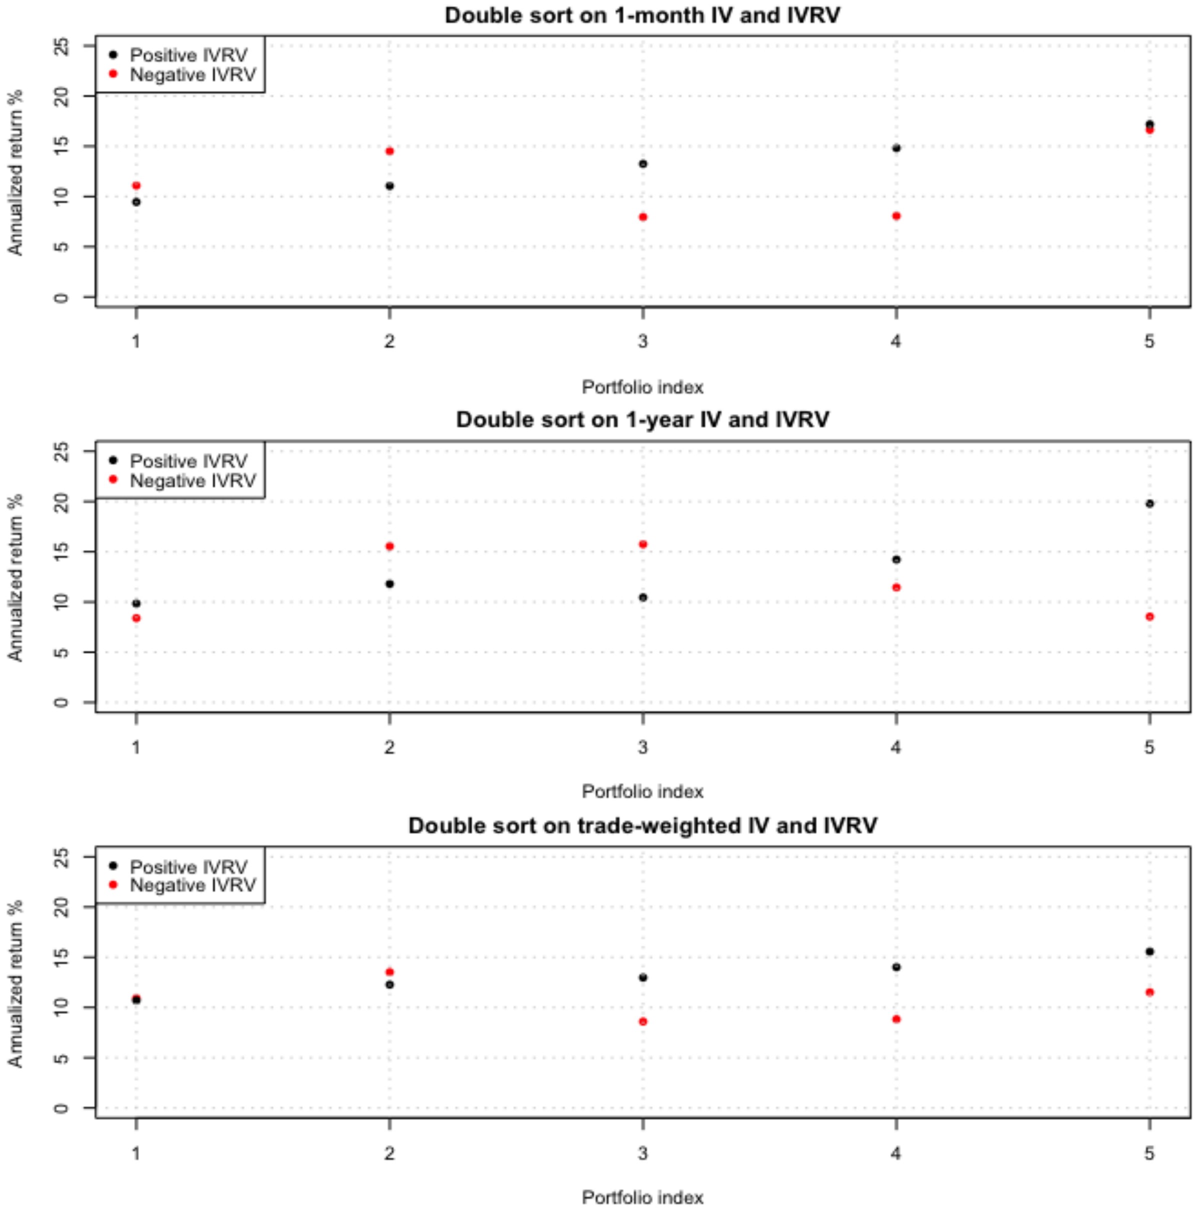

Figure 16: Double sort on implied volatilities and volatility risk premium 2009-2016

Double sort of IV and volatility risk premium

Sub-period 2: 2016-2023

Table 16: Excess returns gathered by positive IVRV 2016-2023

Vol metric	*Mean excess returns*	H_0*: Indistinguishable from zero (p-value)*
One-month IV	*-0,85%*	*0,66*
One-year IV	*-1,22%*	*0,58*
Trade-weighted IV	*-3,99%*	*0,18*

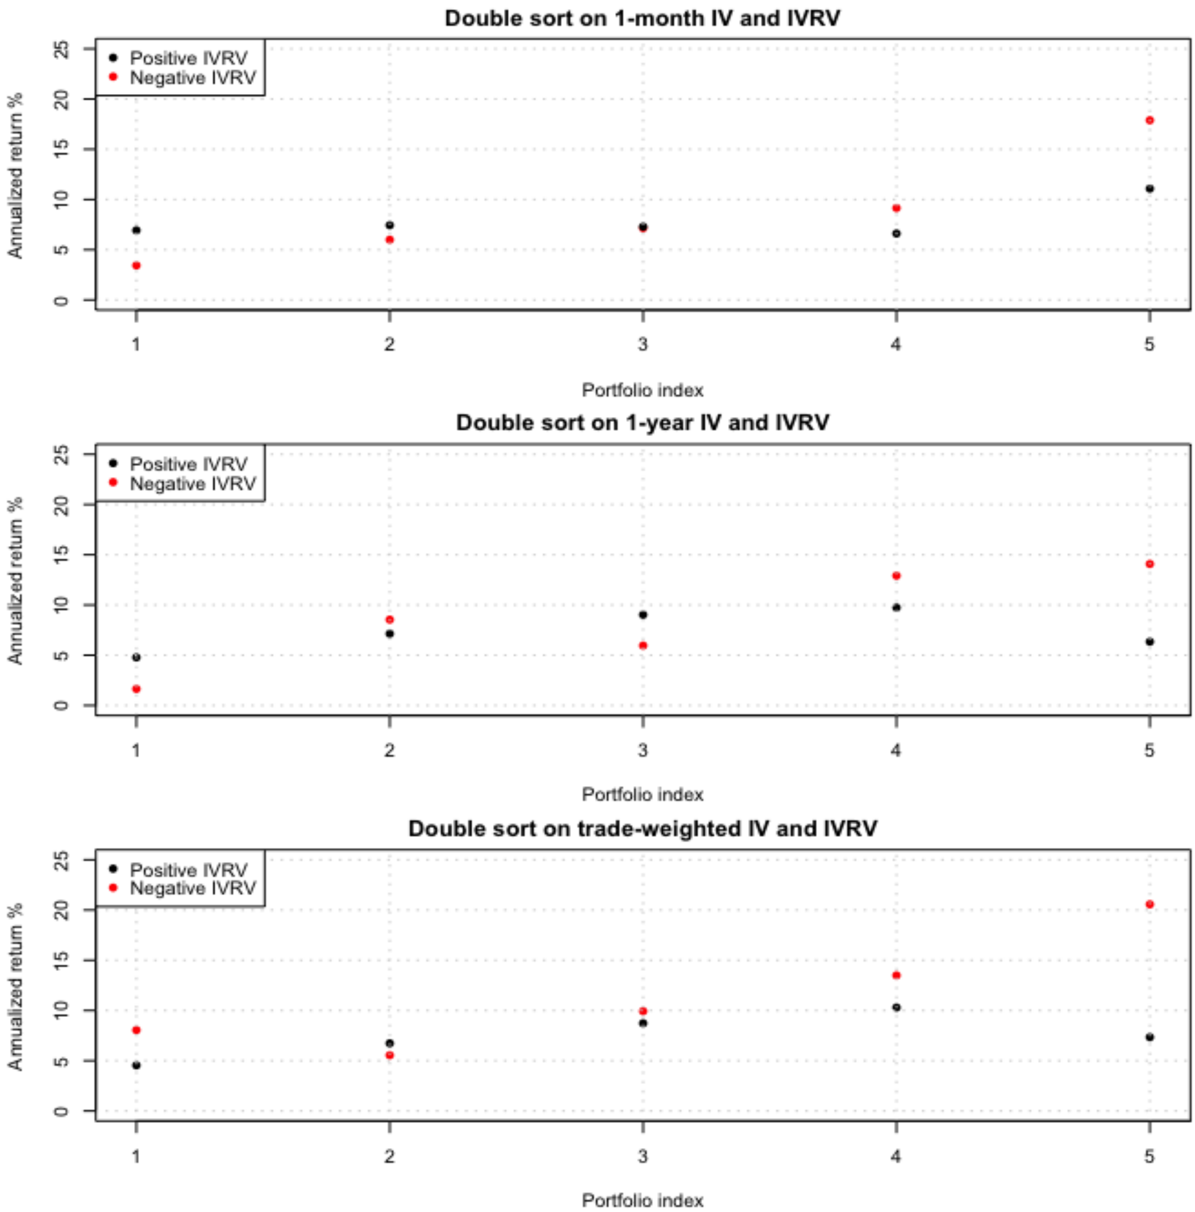

Figure 17: Double sort on implied volatilities and volatility risk premium 2016-2023

Norges miljø- og biovitenskapelige universitet
Noregs miljø- og biovitskapelege universitet
Norwegian University of Life Sciences

Postboks 5003
NO-1432 Ås
Norway