The Predictive Performance of Implied and Historical Volatility

Evidence from Swedish Equity Options and Portfolio Risk Minimization

Erit Aliu
Elias Selin

Department of Economics
Lund University School of Economics and Management
Bachelor's Thesis – Finance
Supervisor: Andreas Johansson
NEKH02
January 2026

Abstract

This thesis examines whether implied volatility derived from Swedish equity options provides superior predictive information relative to historical volatility for future realized volatility and stock returns, and whether this information improves portfolio risk outcomes. Using monthly option and stock data for OMXS30 firms, the study applies predictive regressions and constructs minimum variance portfolios based on implied volatility and historical volatility. The results show that implied volatility consistently outperforms historical volatility in forecasting realized volatility and exhibits stronger return predictability, while historical volatility shows no return predictive power. These forecasting gains translate into economically meaningful portfolio improvements, as an implied volatility-based minimum variance portfolio achieves lower realized variance and downside risk than standard benchmarks. The findings highlight the practical relevance of implied volatility for risk management in the Swedish equity market, extending prior evidence beyond U.S. equity markets.

Acknowledgement

This bachelor's thesis was written as part of the Bachelor of Science in Business and Economics, major in Economics, at the Lund University School of Economics and Management.

We would like to express our sincere gratitude to our supervisor, Andreas Johansson, for his valuable guidance, constructive feedback, and support throughout the research process. His insights and suggestions were greatly appreciated and contributed significantly to the development of this thesis.

We would also like to extend our heartfelt thanks to Rita Aliu and Theresia Selin for their constant encouragement, patience, and support throughout the writing process. Their motivation and belief in our work were invaluable.

Lund, January 2026

Erit Aliu

Elias Selin

Contents

1 INTRODUCTION **6**

2 THEORY **9**

2.1 Volatility 9
2.2 Option Theory 11
2.2.1 Black-Scholes-Merton Model 12
2.3 Risk and Return Framework 18
2.4 Portfolio Theory 20
2.4.1 Portfolio Construction Strategies 20
2.4.2 Portfolio Performance Measure 23

3 EMPIRICAL DATA **27**

4 METHODOLOGY **29**

4.1 Research Design and Empirical Strategy 29
4.2 Ordinary Least Square Assumptions 30
4.3 Portfolio Construction: Minimum Variance Portfolio 31
4.4 Portfolio construction: IV Minimum Variance Portfolio 36
4.5 Portfolio construction: Naive Diversification Portfolio 39
4.6 Statistical Inference in Predictive Regressions 41
4.7 Statistical Inference in Portfolio Risk 42

5 EMPIRICAL RESULTS **43**

5.1 Volatility Forecasting Results 43
5.2 Return Predictability Results 45
5.3 Portfolio Results 47
5.4 Discussion 51

6 CONCLUSION **54**

REFERENCES **55**

APPENDIX **58**

STATEMENT ON THE USE OF GENERATIVE AI **62**

List of Abbreviations

ATM: At-the-money

BSM: Black-Scholes-Merton

CAGR: Compound Annual Growth Rate

CAPM: Capital Asset Pricing Model

EHM: Efficient Market Hypothesis

FE: Fixed Effects

GBM: Geometric Brownian Motion

HV: Historical Volatility

IV: Implied Volatility

IV-MVP: Implied Volatility-Based Minimum Variance Portfolio

ITM: In-the-money

MDD: Maximum Drawdown

MVP: Minimum Variance Portfolio

OLS: Ordinary Least Squares

OOS: Out-of-Sample

OTM: Out-of-the-money

RE: Random Effects

RV: Realized Volatility

STD: Standard Deviation

VRP: Variance Risk Premia

1 Introduction

Forecasting of volatility and returns is central to asset pricing, portfolio allocation, and risk management (Bodie, Kane, & Marcus, 2019). In both academic research and investment practice, volatility forecasts are used to assess future risk, evaluate expected returns, and construct portfolios that balance risk and reward. Traditionally, volatility has been estimated using historical return data, resulting in backward-looking measures that implicitly assume stability in the risk environment (Hull, 2022). However, option-implied volatility (IV) provides a forward-looking alternative by extracting market expectations of future uncertainty from option prices (Christensen & Prabhala, 1998). Whether IV offers superior predictive information relative to historical volatility (HV) for future realized volatility (RV) and stock returns remains a fundamental question in financial economics (Andersen, Bollerslev, Diebold, & Labys, 2003). Moreover, if IV improves volatility forecasts, an important question is whether this informational advantage can be translated into economically meaningful outcomes through improved portfolio construction, particularly within a minimum variance portfolio (MVP) framework (Markowitz, 1952).

Early studies show that HV is highly persistent and contains substantial information about future volatility (Figlewski, 1994; Andersen et al., 2003). A major contribution is Christensen and Prabhala (1998), who demonstrate that IV from index options outperforms HV in forecasting future RV. Furthermore, Blair, Poon and Taylor (2001) show that IV is particularly useful at short horizons and conclude that IV is a better predictor of future RV than HV. Recent work demonstrates that IV has economic value in portfolio allocation, improving risk and covariance estimates (Guidolin & Wang, 2023). Overall, the literature suggests that IV captures forward-looking market expectations about future RV that are not fully reflected in past returns.

Evidence on whether volatility predicts future returns is mixed. Some studies find a positive relationship between IV and future returns (Banerjee, Doran & Peterson, 2007; Ammann, Verhofen, and Süss, 2009), while others show that IV alone has little predictive power (Bali & Hovakimian, 2009). Bollerslev, Tauchen, and Zhou (2009) argue that return predictability mainly arises from the variance risk premia (VRP), approximated as the difference between IV and HV, a result further supported by Drechsler and Yaron (2011). In contrast, the evidence on HV as a predictor of future returns is generally weak and inconsistent. Studies show that HV is often statistically insignificant and unstable, largely due to the backward-looking nature of HV (Campbell, French, Schwert & Stambaugh, 1988; Ghysels, Santa-Clara & Valkanov, 2005; Corrado & Miller, 2006). Taken together, the literature highlights the importance of IV, particularly in forecasting future RV and in risk management. Importantly, the studies are based on the U.S. equity market, which raises the question of whether these findings also hold in the Swedish equity market.

Therefore, this study aims to examine whether IV derived from Swedish equity options provides superior predictive power relative to HV in forecasting future RV and stock returns. In addition, the study evaluates whether incorporating IV into an MVP framework leads to improved realized risk outcomes when compared to commonly used benchmark portfolio strategies. Finally, the study investigates whether the predictive performance of IV in the Swedish equity market resembles from that documented in prior studies predominantly based on U.S. equity markets.

In the first stage, monthly panel regressions are used to evaluate the ability of lagged IV and HV to predict future RV. The study then examines whether lagged IV, HV, and VRP predict future excess stock returns. In addition to the panel analysis, firm-level time-series regressions are conducted to evaluate whether the RV forecasting relationships hold at the individual stock level. All regressions are specified using lagged explanatory variables to ensure a clear information structure and avoid look-ahead bias.

To answer the second research question, the statistical forecasting results are translated into an economic evaluation through portfolio construction. Forecasts of future RV based on IV and HV are used as inputs in an MVP framework. To strengthen the link between the regression evidence and the portfolio analysis, the implied volatility-based minimum variance portfolio (IV-MVP) is constructed using only those stocks for which IV is found to be a statistically significant predictor of future RV at the firm level. Portfolio performance is then evaluated out-of-sample (OOS) and compared against benchmark portfolios to assess whether superior forecasting performance leads to improved risk outcomes.

By combining predictive regressions with portfolio-based performance analysis, the methodology allows the study to assess whether this information can be economically exploited in portfolio allocation. This integrated approach ensures that the findings are relevant for both academic research and practical investment decision-making. The results are subsequently compared to evidence documented in prior studies based on U.S. equity markets, allowing an assessment of whether the predictive performance of IV differs in the Swedish context.

The findings of the study show that IV derived from Swedish equity options provides stronger predictive information for future RV than HV and exhibits superior return predictability. Importantly, the statistical evidence translates into economic value, as an IV-MVP achieves lower realized risk and improved risk-adjusted performance relative to standard benchmark portfolios. While these results are broadly consistent with prior evidence from U.S. equity markets, the study extends the existing literature by demonstrating that the predictive and economic relevance of IV also holds in the Swedish equity market. By jointly examining volatility forecasting, return predictability, and portfolio risk outcomes within a unified framework, the study contributes to a broader understanding of the practical applicability of IV as a forward-looking risk measure in a smaller and less studied equity market.

2 Theory

2.1 Volatility

Volatility is a central concept in financial theory and is commonly used as a quantitative measure of investment risk. It captures the variability of asset returns over time and reflects the risk investors face regarding future price movements. As emphasized by Bodie et al. (2019), asset prices fluctuate in response to firm-specific information and macroeconomic news, implying that there is no theoretical benchmark for a "natural" level of risk in financial markets. Moreover, neither risk nor expected returns are directly observable and must therefore be inferred from historical return data, which introduces estimation error and uncertainty. Consequently, commonly used volatility measures are largely backward-looking and rely on the assumption that historical return behavior provides useful information about future risk, making the choice of volatility measure crucial for risk assessment and interpretation in investment analysis (Bodie et al. 2019).

Historical and Realized Volatility

HV measures the degree of variation in an asset's past returns and is therefore a backward-looking indicator of risk. It is typically calculated as the sample standard deviation (STD) of historical returns over a specified period and reflects the variation that has been realized in the asset's return process. As discussed by Bodie et al. (2019), historical return volatility is commonly used as a proxy for risk in portfolio analysis and empirical asset pricing due to its simplicity and intuitive interpretation.

Annualized daily volatility:

$$\sigma = \sqrt{\frac{1}{n-1}\sum_{t=1}^{n}(r_t - \bar{r})^2} \times \sqrt{252} \tag{2.1}$$

$$\bar{r} = \frac{1}{n}\sum_{t=1}^{n} r_t \tag{2.2}$$

$$r_t = ln(\frac{S_t}{S_{t-1}}) \tag{2.3}$$

The annualized HV is calculated as the sample STD of continuously compounded returns over a given sample period and adjusted by the factor $\sqrt{252}$ to annualize daily volatility. The term n represents the number of observations, r_t denotes the continuously compounded return at time t, and $\bar{r}$ is the sample mean return as defined in Equation (2.2). The continuously compounded return r_t is calculated according to Equation (2.3) as the natural logarithm of the ratio of the asset price at time t to its price at time t - 1.

The factor 252 corresponds to the approximate number of trading days in a calendar year, and its square root is used to convert daily volatility into annualized volatility, consistent with standard practice in financial markets (Hull, 2022).

As noted by Hull (2022), HV may fail to incorporate new information in a timely manner, which can render it misleading when future market conditions differ from those observed in the past. Consequently, while HV provides a useful and intuitive measure of risk, it may be insufficient in environments characterized by rapid market changes.

2.2 Option Theory

Options are financial derivatives that grant the holder the right to buy or sell an underlying asset at a predetermined price at a predetermined time. The predetermined price is called the strike price, and the underlying asset can include instruments such as stocks, equity indices, commodities, currencies, interest rate contracts and other financial instruments. Options are traded both on exchanges and in the over-the-counter market. The most common form of options are vanilla options, which have standardized payoff structures and include call and put options. Vanilla options can further be classified as American or European options, depending on when they may be exercised. American options can be exercised at any time up to and including the expiration date, whereas European options can only be exercised at maturity (Hull, 2022).

A call option grants the holder the right to buy the underlying asset at the strike price, whereas the put grants the right to sell it. Option holders are not obligated to exercise their right, thus they provide flexibility and are commonly used for hedging, speculation and portfolio management. However, the option writer (seller) is obligated to fulfill the terms of the contract if the option holder chooses to exercise it. This asymmetric relationship between the buyer and the writer implies that the buyer pays the value of the option, known as the premium, to compensate the writer for taking on the obligation. The value of the option contract depends on several factors such as current asset price, strike price, time to maturity, the risk-free rate and expected volatility of the underlying asset (IV) (Hull, 2022).

$$Payoff_{Call} = \max(S_t - K, 0) \tag{2.4}$$

$$Payoff_{Put} = \max(K - S_t, 0) \tag{2.5}$$

Where S_t denotes the stock price at time t, and K denotes the strike price.

The relationship between the current price of the underlying asset and the strike price is referred to as an option's *moneyness*, and it determines whether an option is classified as in-the-money (ITM), at-the-money (ATM), or out-of-the-money (OTM). For a call option, the option is ITM when the underlying asset price exceeds the strike price, ATM when the two prices are equal, and OTM when the underlying price is below the strike price. For a put option, the definitions are reversed: it is ITM when the underlying price is below the strike price, ATM when they are equal, and OTM when the underlying price is above the strike price. Moneyness is closely linked to the option's intrinsic value, which represents the immediate exercise value of the contract, and its time value, which reflects the value of the remaining time to maturity. ITM options have positive intrinsic value, ATM options have no intrinsic value but typically the highest sensitivity to price changes in the underlying asset, and OTM options have zero intrinsic value but may still carry time value due to the possibility of becoming profitable before expiration (Hull, 2022).

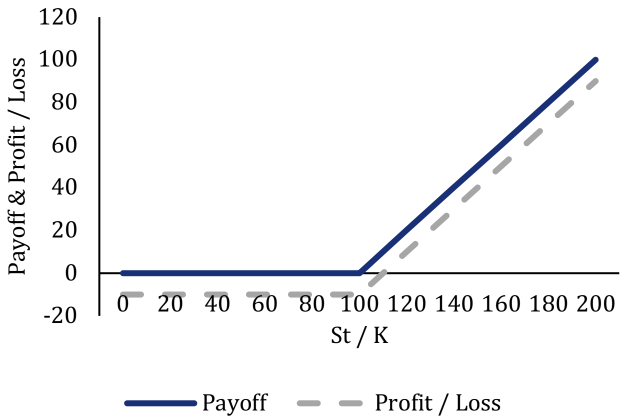

Figure 2.1. Example of payoff and profit/loss of a European call option illustrating option moneyness

Notes: K = 100, and c = 10.

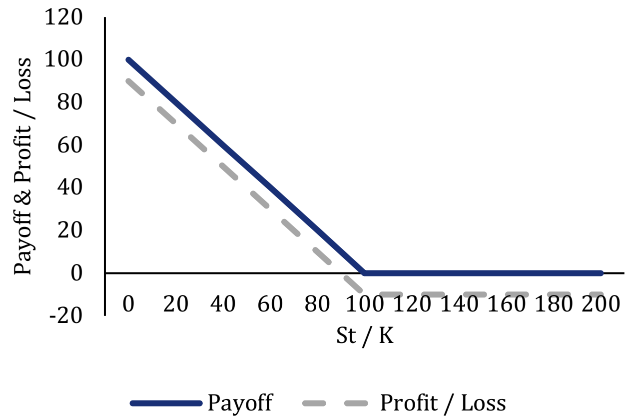

Figure 2.2. Example of payoff and profit/loss of a European put option illustrating option moneyness

Notes: K = 100, and p = 10.

2.2.1 Black-Scholes-Merton Model

Stochastic Modeling of Stock Prices

Before introducing the Black-Scholes-Merton (BSM) model, it is useful to outline the stochastic framework used to describe the evolution of stock prices. Modern option pricing theory assumes that asset prices evolve randomly over time and can therefore be modeled as continuous-time stochastic processes. This approach captures the inherent uncertainty in financial markets and provides a mathematically consistent foundation for derivative pricing (Hull, 2022).

Stock prices are assumed to follow a geometric Brownian motion (GBM). This specification implies that price changes are proportional to the current price level and consist of two components: a deterministic component reflecting the expected rate of return, and a stochastic component capturing random market fluctuations. An important motivation for GBM is that it ensures stock prices remain strictly positive and leads to lognormally distributed prices, a property that is consistent with empirical observations of financial data (Hull, 2022).

The stochastic behavior of asset prices in continuous time is driven by a Wiener process, which provides a mathematical representation of random and unpredictable price movements. Together with Ito's lemma, which governs how functions of stochastic processes evolve, this framework allows option prices to be analyzed. These concepts form the theoretical foundation upon which the BSM model is derived and justify its underlying assumptions regarding price dynamics and probability distributions (Hull, 2022).

Wiener Process

A Wiener process W_t is defined by the following properties:

- $W_0 = 0$
- The increments of W_t over disjoint time intervals are independent
- Over a short time interval dt, the increment is normally distributed with mean zero and variance dt: $dW_t \sim N(0, dt)$

The Wiener process captures the continuous and unpredictable nature of asset price movements and serves as the fundamental source of randomness in continuous time financial models (Hull, 2022).

Geometric Brownian Motion

Under the assumption of GBM, the dynamics of a stock price S_t are given by the following stochastic differential equation:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{2.6}$$

Where μ denotes the expected rate of return on the stock and σ denotes the volatility of the stock, measuring the magnitude of random fluctuations.

Solving this stochastic differential equation using Ito's Lemma yields a closed-form expression for the stock price process:

$$S_t = S_0 exp((\mu - \frac{1}{2}\sigma^2)t + \sigma W_t) \tag{2.7}$$

This solution implies that stock prices are lognormally distributed, ensuring that prices always remain strictly positive. The lognormal property constitutes a key assumption underlying the BSM model and plays a central role in the derivation of option pricing formulas (Hull, 2022).

Risk-Neutral Representation

For option pricing purposes, the stock price process is expressed under a risk-neutral probability measure. In a risk-neutral economy, investors are indifferent to risk, and the expected return on all assets equals the risk-free interest rate. Importantly, the volatility parameter remains unchanged under this transformation. Therefore, in Equations (2.6) and (2.7), μ is replaced with the continuously compounded risk-free interest rate *r*.

This transformation allows derivative prices to be expressed as discounted expected payoffs under the risk-neutral measure and provides the final building block required for the derivation of the BSM model.

Black-Scholes-Merton Model and Implied Volatility

The BSM model represents a major milestone in financial economics and derivative pricing. Developed in the early 1970s by Fischer Black, Myron Scholes, and Robert Merton, the model provided the first closed-form solution for pricing European stock options and fundamentally changed the way derivatives are valued and hedged in practice. Its importance is reflected not only in its widespread use by market practitioners, but also in its lasting influence on modern asset pricing theory. In recognition of its profound impact, Myron Scholes and Robert Merton were awarded the Nobel Prize in Economic Sciences in 1997 for their contributions to the development of the model (Hull, 2022).

Under the assumptions of the BSM framework and using risk-neutral valuation, closed-form pricing formulas for European call and put options can be derived. These formulas express the option price as a function of the current stock price, the strike price, the time to maturity, the risk-free interest rate, and the volatility of the underlying stock. Importantly, volatility is the only parameter in the BSM model that cannot be directly observed in the market, which later motivates the concept of IV.

The pricing relationship in BSM model for a European call and put option are given by formulas below:

$$c = S_0 \cdot N(d_1) - Ke^{-rT} \cdot N(d_2) \tag{2.8}$$

$$p = Ke^{-rT} \cdot N(-d_2) - S_0 \cdot N(-d_1) \tag{2.9}$$

$$d_1 = \frac{ln(\frac{S_0}{K}) + (r + \frac{\sigma_{IV}^2}{2}) \cdot T}{\sigma_{IV} \cdot \sqrt{T}}, d_2 = d_1 - \sigma_{IV} \cdot \sqrt{T} \tag{2.10}$$

Where c and p denote the prices of a European call and put option, S_0 is current asset price, K the strike price, T the time to maturity, r the risk-free rate, σ and $N(\cdot)$ the cumulative distribution of the standard normal distribution.

Hull (2022) explains that $N(d_2)$ can be interpreted as the risk-neutral probability that the option will expire ITM, that is, the probability that the stock price at maturity exceeds the strike price. Consequently, the term $Ke^{-rT}N(d_2)$ represents the present value of the expected strike payment, conditional on the option being exercised. The term $N(d_1)$, while not a probability in a strict sense, is closely related and reflects the option's sensitivity to changes in the underlying stock price, commonly referred to as the delta.

The value of a European call option is increasing in the current stock price and volatility and decreasing in the strike price. An increase in volatility raises option values because it increases the likelihood of extreme outcomes, which benefits option holders due to the asymmetric payoff structure. Similarly, longer time to maturity generally increases option value, as it allows more time for favorable price movements to occur. While volatility and time to maturity affect call and put option values symmetrically, the effects of the underlying stock price and the strike price are reversed for put options. Together, the interpretations of $N(d_1)$ and $N(d_2)$ and the monotonicity properties of the pricing formulas provide important intuition for how the BSM model links option prices to underlying economic forces, and why volatility plays such a central role in option valuation (Hull, 2022).

Given Equations (2.8)-(2.10), the only parameter not directly observed in the market is the volatility of the returns of the underlying price process. Hence, IV can be obtained by numerically solving the BSM formula for volatility, such that the theoretical option price equals the observed market price.

IV therefore represents a forward-looking, market-based measure of expected future volatility embedded in option prices. Unlike HV, which is constructed from past returns, IV reflects investors' expectations about future risk and is widely used for forecasting volatility as well as for option pricing and hedging. Because future volatility is not directly observable, IV provides a practical and informative proxy for market expectations of risk (Hull, 2022).

Empirically, IV is not constant across strike prices, as assumed by the BSM model. Instead, equity options typically exhibit a volatility smirk, characterized by higher IV for lower strike prices and a downward sloping IV curve as the strike price increases. This pattern implies that OTM put options command relatively higher IV compared to ATM and ITM options.

According to Hull (2022), the volatility smirk reflects the shape of the risk-neutral return distribution implied by option prices, which displays negative skewness and a heavier left tail relative to the lognormal distribution assumed under the BSM framework. This asymmetry indicates that market participants assign a higher probability, or greater severity, to large negative price movements than to extreme positive outcomes. As a result, options that provide protection against downside risk are priced at higher IVs.

The presence of a volatility smirk therefore represents a systematic deviation from the constant-volatility assumption underlying the BSM model and highlights how option prices embed market perceptions of asymmetric downside risk.

Hence, the IV is identical for the call option and the put option. Consequently, European call and put options on equities with the same strike price and maturity exhibit an identical volatility smirk.

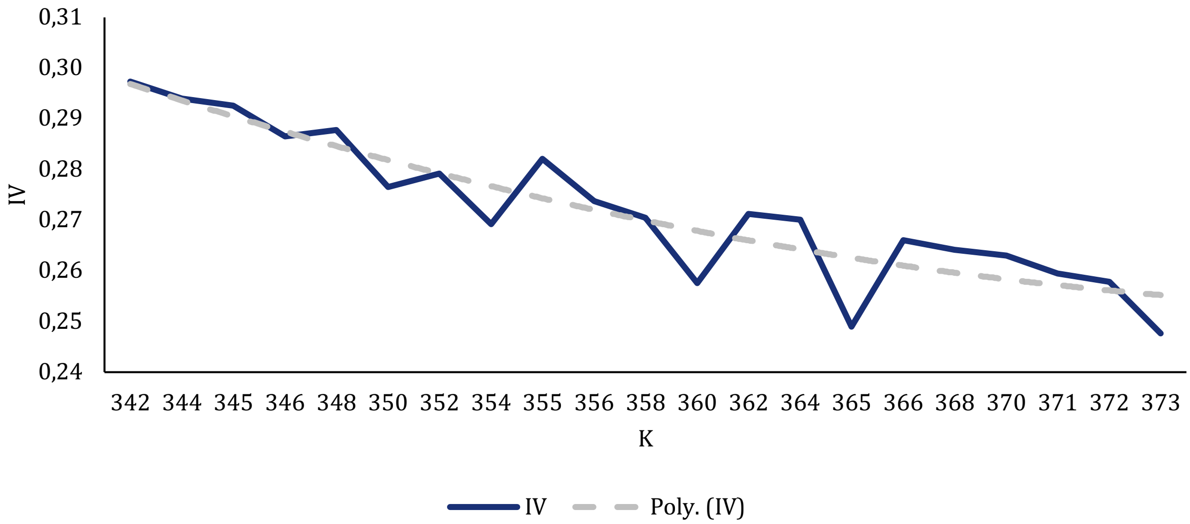

Figure 2.3. Example of volatility smirk for European call and put options with identical maturity

Notes: The figure is an example that illustrates implied volatility as a function of the strike price. S_0 = 358.56, r = 0.01, q = 0.02, and T = 28/365. Where q denotes the continuous dividend yield, and T is the time to maturity measured in years.

Variance Risk Premia

The VRP reflects the compensation investors require for bearing volatility risk. Prior research shows that the VRP captures time variation in risk premia rather than volatility levels per se. Bollerslev et al., (2009) demonstrate that fluctuations in the VRP contains significant information about expected stock returns, particularly during periods of heightened market uncertainty.

VRP is defined as the difference between expected future variance under risk-neutral probability measures and the expected future variance under physical probability measures. In empirical applications, the VRP is often proxied using volatility-based measures (Bollerslev et al., 2009).

Following Bollerslev et al., (2009), the VRP is approximated as the difference between IV and HV:

$$VRP_t = IV_t - HV_t \tag{2.11}$$

A positive risk premium indicates that IV exceeds HV, reflecting investors' aversion to volatility risk and willingness to pay a premium for insurance against it. Conversely, a negative VRP suggests that HV surpasses market expectation. This measure is widely used in empirical finance to capture compensation for bearing risk (Bollerslev et al., 2009).

2.3 Risk and Return Framework

Capital Asset Pricing Model

The capital asset pricing model (CAPM) emerged during the 1960s through Sharpe, Lintner and Mossin (Sharpe, 1964; Lintner, 1965; Mossin, 1966) and has become one of the most influential frameworks in asset pricing and financial theory. The model explains how investors should be compensated for bearing systematic risk (Bodie et al. 2019). In CAPM, the only source of priced risk is assets' sensitivity to fluctuations in the market portfolio captured by the beta coefficient.

Systematic risk refers to the portion of an asset’s total risk that arises from market-wide factors that cannot be diversified away. In CAPM, this component is priced because it reflects the extent to which the asset co-moves with the market. In contrast, firm-specific risk can be diversified away and therefore investors do not require a risk premium. CAPM thus implies that the expected returns only compensate investors for exposure to systematic risk.

According to the model, the expected excess return on an asset is proportional to its beta and market risk premium. This relationship is expressed as:

$$E(r_i) = r_f + \beta_i \cdot \left[E(r_m) - r_f\right] \tag{2.12}$$

Where β_i measures a security’s systematic risk by indicating how sensitive its returns are to movements in the overall market. A beta greater than one implies that the security exhibits stronger movements than the market, while a beta below one indicates lower exposure to systematic risk.

The model relies on several assumptions where investors are risk-averse, form similar expectations about the market that are frictionless and can borrow or lend at the risk-free rate (Bodie et al. 2019). These assumptions rarely hold in practice, but the model remains central because it established a clear link between risk and expected return. This relationship is relevant for our thesis, even though CAPM frames beta rather than volatility. CAPM provides theoretical support for the idea that greater expected risk should correspond with higher expected returns.

Efficient Market Hypothesis

The efficient market hypothesis (EMH) was formally articulated by Fama (1970), who synthesized existing theoretical and empirical research on market efficiency. The hypothesis posits that financial asset prices adjust to incorporate all available information, implying that securities are generally priced at their fundamental values. As a result, consistently achieving abnormal returns through trading strategies based on available information is not feasible in an efficient market.

Within the EMH framework, excess returns can only be obtained by assuming higher levels of systematic risk. Market efficiency is defined by the extent to which prices fully and rapidly reflect information. Fama (1970) categorizes market efficiency into three distinct forms based on the type of information reflected in prices.

- **Weak form efficiency:** Asset prices fully reflect all historical price information, implying that trading strategies based on past prices or return patterns cannot consistently generate abnormal returns.
- **Semi-strong form efficiency:** Asset prices incorporate both historical information and all publicly available information, such as financial statements, news, and macroeconomic data. As a result, neither technical nor fundamental analysis should enable investors to achieve persistent excess returns.
- **Strong form efficiency:** Asset prices reflect all available information, including both public and private (insider) information. Consequently, even investors with access to privileged information are unable to systematically outperform the market.

2.4 Portfolio Theory

2.4.1 Portfolio Construction Strategies

Minimum Variance Portfolio

MVP theory is a foundational framework in finance introduced by Markowitz (1952). The theory establishes that investors should evaluate portfolios based on expected returns and risk. By combining assets with imperfectly correlated returns, investors can reduce overall risk through diversification even if assets are risky on their own.

A central concept in this framework is the efficient frontier, which presents the set of portfolios that deliver the highest expected return for each level of risk. Among these portfolios, the MVP is unique in that it achieves the lowest achievable level of risk among all possible combinations of assets. As a result, the MVP is often used as a benchmark for understanding the risk structure of financial markets and for evaluating the performance of alternative asset allocations. In the Markowitz framework, the MVP represents a special case of the general mean-variance optimization problem and is frequently used as a benchmark portfolio (Bodie et al., 2019).

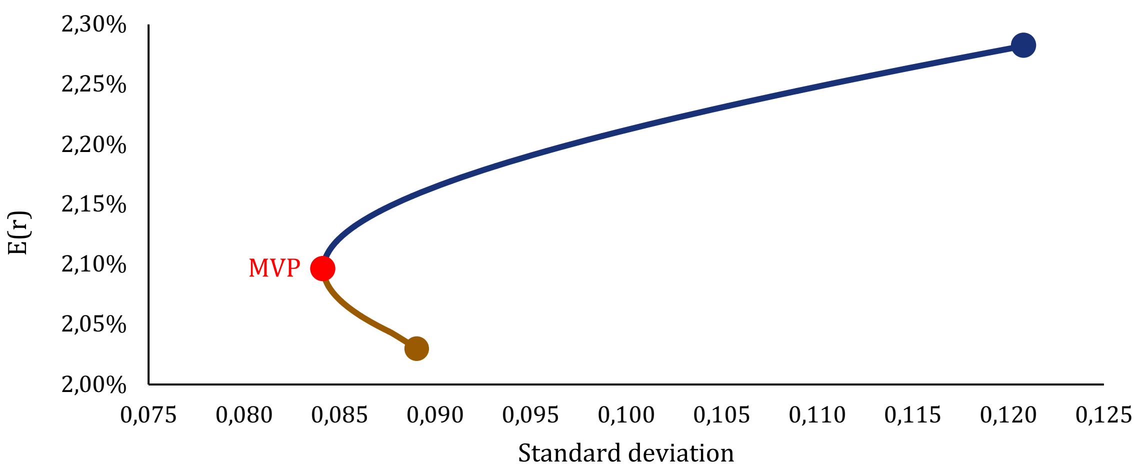

Figure 2.4. Portfolio frontier generated from two risk assets and the minimum variance portfolio

Figure 2.4 illustrates an example of a portfolio frontier generated from a set of risky assets. The blue curve represents the efficient frontier showing the portfolio that provides the highest expected return for a given level of risk. The red point marks the MVP, which corresponds to the lowest achievable level of risk among all feasible combinations of assets.

According to theory, the MVP is the portfolio with the lowest achievable risk. This is attained by combining assets with correlations less than one. The MVP represents the left-most point on the portfolio frontier in Figure 2.4 and serves as a natural benchmark for portfolio risk, as the portfolio diversifies away all diversifiable risk. Unlike general mean-variance efficient portfolios, the MVP depends only on the covariance matrix of asset returns and is independent of expected returns. Because the MVP relies solely on the covariance matrix, it is less sensitive to estimation errors than portfolios that depend on expected return estimates. The MVP is in mathematical terms a minimizing problem, where the portfolio's variance is minimized. The portfolio can be constructed from at least two assets and an unlimited number of assets.

The optimization setup is as follows:

$$\begin{gathered} \arg min_w (w^T \Sigma w) \\ s.t.\ 1^T w = 1 \end{gathered} \tag{2.13}$$

The weights formula for more than three assets:

$$w^{MVP} = \frac{\Sigma^{-1} 1}{1^T \Sigma^{-1} 1} \tag{2.14}$$

Where $w = (w_1, w_2, \dots, w_N)^T$ denotes the vector of portfolio weights, with w_i representing the proportion of total portfolio value invested in asset i, Σ is the covariance matrix of asset returns, and 1 is a vector of ones (Guidolin & Wang, 2023, Eq. 5).

Implied Volatility-Based Minimum Variance Portfolio

The IV-MVP is constructed within the same Markowitz mean-variance framework as the traditional MVP. The underlying optimization problem remains unchanged and aims to minimize portfolio variance subject to a full-investment constraint. The key difference lies in the estimation of the covariance matrix.

Instead of relying on historical return data, the IV-MVP incorporates forward-looking information from IV when constructing the covariance matrix. As argued by Guidolin and Wang (2023), option markets reflect market participants' expectations about future volatility and may therefore provide more timely information about changes in risk during turbulent market conditions.

Following Guidolin and Wang (2023), expected returns are set equal across assets to reduce the impact of estimation error, implying that portfolio selection focuses exclusively on variance minimization. The resulting portfolio weights are given by:

$$w^{IV-MVP} = \frac{\widehat{\Sigma_{IV}^{-1}}\ 1}{1^T\ \widehat{\Sigma_{IV}^{-1}}\ 1} \tag{2.15}$$

Where $\widehat{\Sigma_{IV}}$ denotes the IV covariance matrix and *1* is a vector of ones (Guidolin & Wang, 2023, Eq. 5).

In this framework, IVs are used to estimate the variances of individual assets, while historical correlations are employed to construct the off-diagonal elements of the covariance matrix. This approach results in a quasi-implied covariance matrix, which combines forward-looking volatility information from option prices with historically estimated correlations.

Naive Diversification

The naive diversification portfolio, commonly referred to as naive diversification, assigns equal weights to all assets in the investment universe and serves as a widely used benchmark in empirical portfolio studies. Despite its simplicity, DeMiguel, Garlappi, and Uppal (2009) show that the naive strategy often outperforms optimized portfolios OOS, including mean-variance portfolios constructed using historical estimates of expected returns and covariances. The key reason is that optimized portfolios are highly sensitive to estimation error, particularly in expected returns, which can severely degrade OOS performance. By avoiding parameter estimation altogether, the naive portfolio provides a robust allocation that is less affected by model misspecification and sampling noise. Consequently, the naive strategy constitutes a natural benchmark against which the economic value of more sophisticated allocation methods, such as MVPs based on historical estimates or IV, can be evaluated (DeMiguel et al., 2009).

Weight distribution in the naive portfolio:

$$w_i = \frac{1}{N} \tag{2.16}$$

Where w_i denotes the weight in asset *i* and *N* denotes the total number of assets.

2.4.2 Portfolio Performance Measure

The selected measures are consistent with previous research in the field of portfolio evaluation (Maillard, Roncalli, & Teiletche, 2010) and allow for a comprehensive comparison of the risk, return, and structural characteristics of the considered strategies. Together, these measures provide an in-depth assessment of the differences among portfolio construction approaches and facilitate an evaluation of their relative economic relevance.

Annual Return

The annual return of each portfolio is calculated as the Compound Annual Growth Rate (CAGR), which represents the annualized geometric average return over the evaluation period.

$$CAGR = \left(\left(\frac{Final\ value}{Initial\ value}\right)^{\frac{1}{years}}\right) - 1 \tag{2.17}$$

Cumulative Return

Cumulative return measures the total return achieved over the evaluation period. It can be computed either by compounding periodic returns or directly as the relative change between the final and initial portfolio value.

$$CR = \frac{Final\ value\ -\ Initial\ value}{Initial\ value} \tag{2.18}$$

Sharpe Ratio

The Sharpe ratio, originally introduced by Nobel laureate Sharpe (1966), is a widely used measure of risk-adjusted performance. It evaluates the return of an investment or portfolio relative to the risk undertaken, thereby allowing comparisons across portfolios with different volatility characteristics. The Sharpe ratio is defined as:

$$S_p = \frac{E\left[r_p - r_f\right]}{\sigma_p} \tag{2.19}$$

Where r_P denotes the portfolio return and r_f represents the risk-free rate of return. The denominator σ_P corresponds to the STD of the portfolio's excess return, capturing total risk. Consequently, the Sharpe ratio quantifies the excess return earned per unit of total risk, with higher values indicating superior risk-adjusted performance (Sharpe, 1966).

Sortino Ratio

The Sortino ratio is a risk-adjusted performance measure that evaluates portfolio returns relative to downside risk rather than total volatility. Unlike the Sharpe ratio, which penalizes both positive and negative return fluctuations, the Sortino ratio focuses exclusively on negative deviations below a target or minimum acceptable return. This makes the measure particularly relevant when investors are primarily concerned with adverse return outcomes.

In this study, the Sortino ratio for each portfolio is calculated as:

$$Sortino\ Ratio = \frac{E\left[r_p - r_f\right]}{LPSD} \tag{2.20}$$

Where LPSD represents the lower partial standard deviation, computed using only returns below the target return. A higher Sortino ratio indicates superior performance per unit of downside risk. The theoretical motivation for this measure follows standard portfolio theory as discussed in Bodie et al. (2019).

Maximum Drawdown

Maximum drawdown (MDD) measures the largest decline in portfolio value from a historical peak to a subsequent trough over a given time horizon. It captures the worst cumulative loss experienced by the portfolio before a new peak is attained. The MDD for each portfolio is computed as:

$$MDD(T) = max\{0, max\ P(t)\ - P(T)\} \tag{2.21}$$

Where T denotes the final time of the evaluation period ($t \in 0, t$) and *P(t)* represents the portfolio value at time t. The MDD thus reflects the most severe peak-to-trough loss observed during the sample period.

Calmar Ratio

The Calmar ratio is a performance metric designed to evaluate return relative to downside risk. Like other risk-adjusted measures, it relates portfolio performance to a measure of risk, but instead of volatility it focuses on the MDD. The Calmar ratio for each portfolio is calculated as:

$$CalR = \frac{E\left[r_p - r_f\right]}{MDD} \tag{2.22}$$

Where MDD represents the MDD over the evaluation period. A higher Calmar ratio indicates a more favorable balance between return and drawdown risk.

Portfolio Turnover

Portfolio turnover measures the extent of trading activity within a portfolio over a given period and captures how frequently portfolio holdings are adjusted. A higher turnover indicates more frequent rebalancing and trading, which is typically associated with higher transaction costs and lower net returns. By relating the minimum of purchases and sales to the average portfolio value, this measure reflects the proportion of the portfolio that is effectively replaced during period t (Dow, 2007).

$$Turnover_t = \frac{\min\left(Purchases_t, Sales_t\right)}{Average\ Portfolio\ Value_t} \tag{2.23}$$

3 Empirical Data

This study employs daily stock return data for equities included in the OMXS30 index, obtained from Nasdaq. Daily closing prices are used to compute HV at a monthly frequency, ensuring that each monthly volatility estimate reflects realized return variability over the corresponding period. IV is obtained from Bloomberg and is observed on the last trading day of each month, thereby reflecting market expectations of volatility over the subsequent month. While this approach ensures consistency across assets and over time, the exact computational methodology used by Bloomberg is proprietary and cannot be independently verified. Consequently, the IV estimates should be interpreted as market-based volatility measures rather than exact model-implied parameters.

The sample period spans from October 2020 to October 2025, allowing the analysis to capture recent market conditions across different market regimes. Due to data availability constraints and variation in option market liquidity across individual stocks, the final sample includes 28 of the 30 OMXS30 constituents, as IV data are not available for Addtech B and Lifco B over the full sample period. Although the OMXS30 consists of the most liquid equities in the Swedish market, option liquidity varies across stocks, and lower trading activity in some equity options may introduce noise into the IV estimates.

The use of monthly observations is motivated by several considerations. First, aggregating daily data to a monthly frequency reduces the influence of short-term price fluctuations and market microstructure noise that may distort volatility estimates. Second, the monthly frequency ensures consistency between HV and IV measures, facilitating a meaningful comparison between backward-looking and forward-looking volatility measures. Third, the use of a monthly frequency is consistent with the portfolio rebalancing strategy employed in the study, as higher-frequency rebalancing (e.g., daily or weekly) would likely result in excessive transaction costs and unrealistically high portfolio turnover.

When using the OMXS30 as a benchmark in the portfolio performance analysis, index-level return data are employed. Since the historical individual constituent weights of the index are not publicly available, it is not possible to compute portfolio turnover or concentration measures for the OMXS30 benchmark, unlike for the constructed portfolios.

Table 3.1 provides a summary of all variables used in the empirical analysis, including their construction, frequency, and data sources.

Table 3.1 Summary of variables

Variable	**Construction**	**Frequency**	**Source**
Returns	Equation (2.3)	Monthly	Nasdaq
Excess Returns	$R_{i,t}^{excess} = r_{i,t} - r_t^f$	Monthly	Nasdaq; Sveriges Riksbank
Risk-Free Rate	$r_t^f = (1 + r_{t,annual}^f)^{1/12} - 1$	Monthly	Sveriges Riksbank
HV	Equation (2.1)	Monthly	Nasdaq
IV	Bloomberg one-month ATM IV from European options (100% moneyness), observed at the end of the previous month.	Monthly	Bloomberg
RV	Equation (2.1)	Monthly	Nasdaq
VRP	Equation (2.11)	Monthly	Bloomberg

4 Methodology

4.1 Research Design and Empirical Strategy

This study adopts a two-step empirical strategy to examine whether IV derived from Swedish equity options provides superior predictive power relative to HV in forecasting future RV and asset returns.

In the first step, predictive regressions are estimated using both panel data and firm-level time-series methods in R. Panel-based ordinary least squares (OLS) regressions are employed to assess the average predictive relationship between lagged volatility measures (IV and HV) and future RV. In addition, panel-based OLS regressions are used to examine the predictive power of IV, HV, and the VRP for future stock returns across OMXS30 constituent firms. Finally, separate time-series regressions are estimated for each individual firm to evaluate whether IV and HV possess predictive power for future RV at the stock-specific level.

The volatility predictive regressions follow:

$$RV_t = \alpha + \beta_1 \cdot IV_{t-1} + \varepsilon_t \tag{4.1}$$

$$RV_t = \alpha + \beta_1 \cdot HV_{t-1} + \varepsilon_t \tag{4.2}$$

The return predictive regressions follow:

$$R_t = \alpha + \beta_1 \cdot IV_{t-1} + \varepsilon_t \tag{4.3}$$

$$R_t = \alpha + \beta_1 \cdot HV_{t-1} + \varepsilon_t \tag{4.4}$$

$$R_t = \alpha + \beta_1 \cdot VRP_{t-1} + \varepsilon_t \tag{4.5}$$

In the second step, the regression results are used to inform portfolio construction within an MVP framework. Forecasts based on IV and HV are translated into variance estimates that serve as inputs in the portfolio optimization. Importantly, the IV-MVP is constructed only using stocks for which lagged IV derived from equity options is found to be a statistically significant predictor of future RV at the firm level. This screening procedure ensures that the portfolio construction is directly linked to the stock-level forecasting evidence and reduces noise from assets for which IV does not contain predictive information.

By comparing the performance of portfolios based on IV and HV, the study evaluates whether superior statistical forecasting performance translates into economically meaningful portfolio outcomes.

In the portfolio construction, transaction costs are not explicitly considered. Although transaction costs are relevant in real-world portfolio management, they are excluded to simplify the analysis. Instead, portfolio turnover is incorporated to reflect trading activity, allowing for an indirect assessment of the potential impact of transaction costs across portfolios. Annual portfolio turnover is computed as the average of period turnovers over the 36 months.

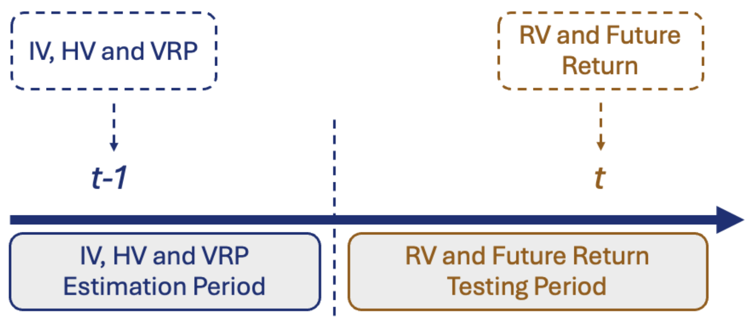

Figure 4.1 illustrates the timing consistency between the volatility measures and the forecast horizon, ensuring that all regressions are based solely on information available at time t-1

4.2 Ordinary Least Square Assumptions

The empirical analysis is conducted using panel data regression models estimated within an OLS framework, including both fixed-effects and random-effects specifications, which rely on several standard assumptions. First, the regression models are assumed to be linear in parameters, and the explanatory variables are treated as predetermined using lagged volatility measures.

Although linear predictive regressions are standard in the literature, the relationship between volatility measures and future outcomes may be nonlinear or state dependent, which cannot be fully captured within the OLS framework.

In panel data settings, key econometric concerns in volatility regressions include the presence of firm-specific effects, as well as serial correlation and heteroskedasticity in the error terms. Volatility measures are known to exhibit persistence over time, which may induce serial correlation within firms. In addition, periods of market turbulence may lead to time-varying error variances and heterogeneity across firms.

Therefore, statistical inference in the panel regressions is based on heteroskedasticity-robust standard errors with clustering that reflects the underlying dependence structure of the residuals. Specifically, standard errors are clustered at the firm level in specifications with firm fixed effects (FE) or random effects (RE), at the time level in specifications with time FE, and two-way clustered by firm and time in two-way fixed-effects models.

Random-effects specifications are estimated as robustness checks, and the Hausman test is employed to assess the validity of the random-effects assumption. In line with common practice in empirical finance, firm FE are adopted as the baseline specification. All estimation choices are made ex ante and applied consistently across specifications to ensure comparability of results.

For FE models, explanatory power is assessed using within R^2, which measures the share of within-firm variation explained by the regressors. For the RE specification, overall R^2 is reported. As these measures capture different sources of variation, R^2 values are not directly comparable across specifications.

4.3 Portfolio Construction: Minimum Variance Portfolio

To assess whether the statistically identified predictability of IV has economic relevance, this thesis translates the regression evidence into a portfolio allocation exercise. In this study construction of an MVP based on historical (backward-looking) risk estimates is made, which serves as a benchmark allocation. Its OOS performance is subsequently compared to that of an IV based MVP in Section 4.4. The portfolios are constructed in October 2022 and rebalanced at the end of every month, until September 2025.

The MVP framework is particularly well suited for this purpose because portfolio weights depend exclusively on the covariance structure of asset returns, rather than on expected return estimates. This is an important advantage in empirical applications, as expected returns are notoriously difficult to estimate and highly sensitive to sampling error. By focusing solely on risk minimization, the MVP allows for a clean economic evaluation of whether improved volatility forecasts translate into lower realized portfolio variance.

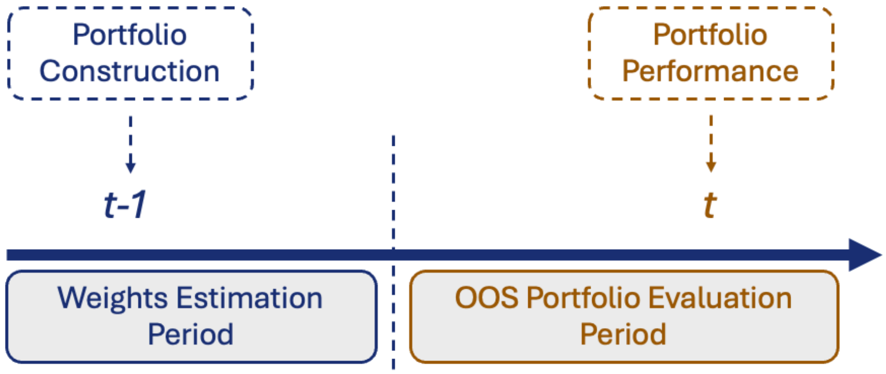

Figure 4.2. Portfolio construction at time t-1 and out-of-sample performance evaluation at time t

Portfolio Universe and Asset Screening

The investable universe consists of OMXS30 constituent stocks for which consistent return and IV data are available over the full sample period. To ensure that the portfolio construction is tightly linked to the empirical results, an econometric screening step is applied prior to portfolio formation.

Specifically, the volatility forecasting regressions described in Section 4.1 are estimated, where lagged implied volatility IV_{t-1} is used to explain future realized volatility RV_t. Stocks are retained for portfolio formation only if the estimated coefficient on lagged IV is statistically significant at the chosen significance level.

This screening procedure is economically motivated. If IV contains statistically significant information about future RV, it should also improve forecasts of asset-specific variances. Since variance forecasts are the key inputs in minimum variance optimization, restricting the portfolio universe to assets with demonstrable volatility predictability ensures that the portfolio exercise directly reflects the forecasting evidence. By excluding assets for which IV does not exhibit predictive power for future risk, the analysis reduces noise and strengthens the link between statistical predictability and economic outcomes.

Estimation Window and Rolling Covariance Matrix

For the selected set of stocks, the benchmark MVP requires an estimate of the return covariance matrix. In line with the traditional Markowitz framework, this covariance matrix is constructed using historical realized returns.

The estimation procedure follows a rolling-window design:

- ***Return series:*** Monthly logarithmic returns are computed for each of the stocks, as defined in the Data section.
- ***Rolling estimation window:*** At each portfolio formation date, the covariance matrix is estimated using a 12-month (one-year) rolling window of daily returns.
- ***Rolling update scheme:*** The window advances by one month at a time. When moving from month *t* to *t+1*, the oldest return observation is dropped and the most recent month is added, ensuring that the covariance matrix reflects only information available at the time of portfolio formation.
- ***Implementation:*** All pairwise covariances are computed in Microsoft Excel, resulting in a n×n covariance matrix for each rolling window. This implementation ensures transparency and full reproducibility of the inputs used in the portfolio optimization.

This rolling approach captures time variation in return comovements while preserving a strictly ex-ante information structure.

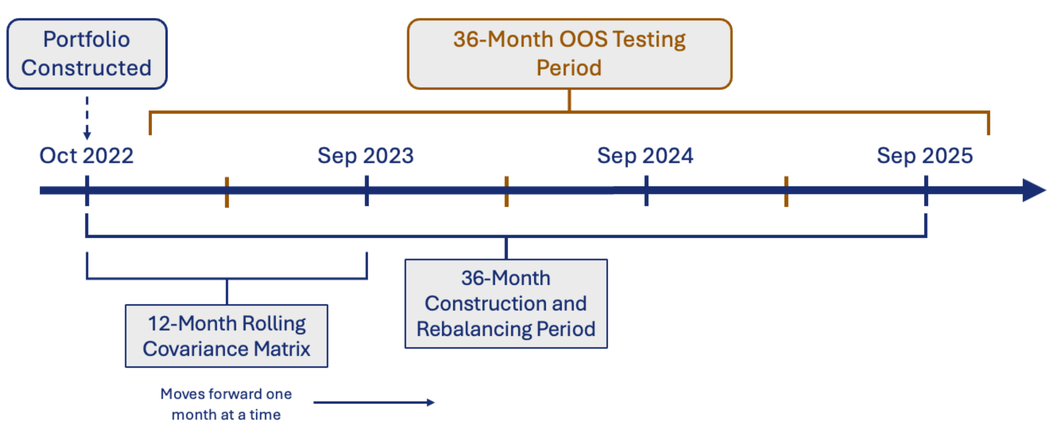

Figure 4.3. Rolling portfolio construction using a 12-month estimation window and 36-month OOS testing period

Weight Computation and Rebalancing Rule

Given the estimated covariance matrix at each month-end, portfolio weights are computed using the closed-form MVP solution presented in Section 2.4.1. To avoid unnecessary repetition, the optimization problem and its analytical solution are not restated here; instead, the theoretical expression is applied mechanically using the estimated rolling covariance matrix.

The operational procedure is as follows:

- ***Information set at month-end t:*** The covariance matrix is estimated using returns from the preceding 12 months.
- ***Weight computation:*** MVP weights are calculated based on this covariance matrix using the theoretical expression introduced in Section 2.4.1 and Equation (2.14).
- ***Rebalancing:*** The portfolio is rebalanced at a monthly frequency, with weights updated at each month-end.
- ***Constraints:*** The portfolio satisfies the standard full-investment constraint that weights sum to one. No short-sale constraints are imposed, implying that portfolio weights are allowed to take negative values. This unconstrained setup follows standard practice in the academic minimum variance literature and ensures that the analytical MVP solution is not distorted by binding inequality constraints.

This process yields a time series of monthly portfolio weights and an associated realized return series.

Out-of-Sample Performance Design

To avoid in-sample overfitting and to reflect a feasible investor decision process, portfolio performance is evaluated strictly OOS.

The OOS design is implemented as follows:

- ***Weight formation:*** At the end of month *t*, portfolio weights are computed using only information available up to that point.
- ***Holding period:*** These weights are held fixed throughout the subsequent month *t+1*.
- ***Realized return:*** The portfolio return for month *t+1* is calculated using the fixed weights formed at t and the realized asset returns observed during *t+1*.

For example, weights computed using data available at the end of September 2025 are evaluated on realized returns during October 2025. Repeating this procedure over the sample yields an OOS back test covering 36 monthly observations (three years), with portfolio weights updated each month based on a one-year rolling covariance estimate.

Role of The Benchmark MVP In the Thesis Design

The historically estimated MVP serves as a foundational benchmark in the empirical design of this thesis. Its primary role is to represent the standard, backward-looking approach to risk minimization that relies exclusively on realized return data and does not incorporate any forward-looking information from option markets.

In addition to this historical MVP, the analysis constructs two alternative portfolio allocations: an IV-MVP and a naive portfolio. The inclusion of these portfolios allows for a comprehensive economic evaluation of whether information contained in IV forecasts can be translated into superior portfolio outcomes.

The comparison across portfolios is deliberately designed to be clean and internally consistent. The optimization framework, asset universe, rebalancing frequency, rolling estimation window, and OOS evaluation procedure are held constant across the historical MVP and the IV-MVP. The naive portfolio is included as a model-free benchmark that avoids estimation altogether and is known to perform competitively OOS in many empirical settings.

This structure ensures that any observed differences in realized portfolio variance or risk-adjusted performance can be attributed solely to differences in the underlying risk estimates. If the IV-MVP exhibits lower realized variance than both the historical MVP and the naive portfolio, this provides direct evidence that the predictive relationship between lagged IV and future RV, documented in the regression analysis, has tangible economic value.

In this sense, the portfolio exercise serves as an economic validation of the statistical results. If IV at time t-1 contains information about RV at time t, then using this information to forecast asset variances should improve the construction of risk-minimizing portfolios. Superior OOS performance of the IV-MVP therefore illustrates how forward-looking volatility expectations embedded in option prices can be exploited to reduce portfolio risk relative to traditional backward-looking and naive diversification strategies.

4.4 Portfolio construction: IV Minimum Variance Portfolio

To evaluate whether the predictive content of IV documented in the regression analysis translates into economically meaningful portfolio improvements, this section constructs an IV-MVP. The IV-MVP is deliberately designed to differ from the benchmark portfolio introduced in Section 4.3 only in the estimation of the variance-covariance matrix of asset returns. All other elements of the portfolio construction and evaluation framework are held fixed.

The construction closely follows established practice in the literature on option-implied risk estimation, specifically the approach proposed by Guidolin and Wang (2023), in which IV is used to forecast asset-specific volatilities while the cross-asset dependence structure is obtained from historical return data. This design isolates the economic contribution of forward-looking volatility information embedded in option prices.

This design choice ensures that any differences in OOS portfolio performance can be attributed exclusively to the incorporation of IV-based risk estimates, rather than to changes in the optimization framework, asset universe, or rebalancing scheme.

Rolling Correlation Matrix

As in the benchmark MVP, the dependence structure across assets is captured through a rolling correlation matrix estimated from realized returns. Specifically, at each portfolio formation date *t*, a N×N correlation matrix $\widehat{R_t}$ is computed using a 12-month rolling window of monthly logarithmic returns for the N assets in the portfolio universe.

The rolling-window scheme advances one month at a time, ensuring that only information available at the time of portfolio formation is used. By construction, the correlation matrices employed in the benchmark MVP and the IV-MVP are identical at each point in time. This ensures that any differences in portfolio outcomes are not driven by changes in the estimated correlation structure.

Importantly, the rolling correlation matrices used in the IV-MVP are identical to those employed in the benchmark MVP at each point in time. As a result, the dependence structure across assets is held constant across portfolios, and any differences in portfolio outcomes cannot be attributed to changes in estimated correlations.

IV-Based Volatility Matrix

Forward-looking volatility forecasts are constructed using IV observed at the end of month *t*. For each asset *i*, IV is interpreted as a forecast of the STD of returns over month *t+1*, consistent with the volatility forecasting evidence presented in Section 4.1.

These IV-based STD forecasts are assembled into a diagonal volatility matrix $\widehat{D_t}$, defined as:

$$\widehat{D}_t = \begin{pmatrix} \widehat{\sigma_{1,t}^{IV}} & \cdots & 0 \\ \vdots & \widehat{\sigma_{2,t}^{IV}} & \vdots \\ 0 & \cdots & \widehat{\sigma_{N,t}^{IV}} \end{pmatrix} \tag{4.6}$$

Where $\widehat{\sigma_{i,t}^{IV}}$ denotes the IV-based forecast of asset *i*'s STD and all off-diagonal elements are set to zero.

This representation makes explicit that IV enters the portfolio construction exclusively through asset specific volatility forecasts and does not affect the correlation structure of returns (Guidolin & Wang, 2023).

IV-based Covariance Matrix Construction

The IV-based covariance matrix is obtained by combining the rolling correlation matrix with the IV-based volatility matrix. Specifically, the covariance matrix at time t is obtained mechanically as:

$$\widehat{\Sigma_t^{IV}} = \widehat{D_t} \times \widehat{R_t} \times \widehat{D_t} \quad (4.7)$$

This matrix multiplication yields individual covariance elements of the form:

$$\widehat{\Sigma_{ij,t}^{IV}} = \widehat{\rho_{ij,t}} \times \widehat{\sigma_{i,t}^{IV}} \times \widehat{\sigma_{j,t}^{IV}} \quad (4.8)$$

Where $\widehat{\rho_{ij,t}}$ is the historical correlation between asset *i* and *j*, while $\widehat{\sigma_{i,t}^{IV}}$ and $\widehat{\sigma_{j,t}^{IV}}$ are IV-based volatility forecasts.

By construction, this hybrid approach preserves the empirically observed co-movement structure of returns while replacing backward-looking variance estimates with forward-looking volatility forecasts derived from IV. The resulting covariance matrix therefore differs from that used in the benchmark MVP only through its diagonal elements.

Portfolio Optimization and Rebalancing

Given the IV-based covariance matrix $\widehat{\Sigma_t^{IV}}$, portfolio weights are computed using the same closed-form minimum variance solution employed for the benchmark MVP. The portfolio is subject to the full-investment constraint that weights sum to one and imposes no short-sale restrictions.

The portfolio is rebalanced at a monthly frequency. At the end of month *t*, weights are computed using $\widehat{\Sigma_t^{IV}}$ and held fixed throughout month *t+1*. Realized portfolio returns are then computed using observed asset returns over the holding period.

The rebalancing frequency, asset universe, rolling estimation window, and OOS evaluation design are identical to those used for the benchmark MVP.

Interpretation

The IV-MVP represents a forward-looking approach to risk minimization in which market expectations of future RV, as reflected in option prices, are directly incorporated into portfolio construction. If IV contains incremental information about future RV beyond that contained in historical returns, the IV-MVP should achieve lower realized portfolio variance than both the historical MVP and the naive portfolio.

Accordingly, superior OOS performance of the IV-MVP would provide direct economic validation of the regression evidence, demonstrating that the predictive relationship between IV and future RV can be successfully exploited in a risk-minimizing portfolio framework. Therefore, this approach would be a useful technique in risk and asset management.

4.5 Portfolio construction: Naive Diversification Portfolio

To provide a simple, model-free benchmark for portfolio performance, this section constructs an equally weighted portfolio, commonly referred to as the naive portfolio. The naive portfolio serves as a naive diversification strategy that does not rely on any estimation of expected returns, variances, or covariances. As such, it provides a useful reference point against which the performance of the optimized portfolios can be evaluated.

The inclusion of the naive portfolio allows the analysis to assess whether the gains from sophisticated risk estimation and optimization outweigh the benefits of simple equal-weight diversification, which has been shown to perform competitively in many OOS settings.

Portfolio Weights

The naive portfolio assigns equal weights to all assets in the investable universe. At each portfolio formation date, each of the N assets is assigned a portfolio weight of:

$$w_{i,t} = \frac{1}{N}, \text{for all } i = 1, \ldots, N. \tag{4.9}$$

These weights are fixed across assets and do not depend on any estimates of volatility, correlation, or expected returns. Consequently, the naive portfolio avoids all parameter estimation errors.

Rebalancing and OOS Evaluation

The naive portfolio is rebalanced at the same monthly frequency as the benchmark MVP and the IV-MVP. At the end of each month t, equal weights are assigned using the same asset universe as in the optimized portfolios and held fixed throughout the subsequent month t+1.

Portfolio returns for month t+1 is then computed using the realized asset returns over the holding period. This procedure ensures that the naive portfolio is evaluated strictly OOS and is fully comparable to the optimized portfolios.

Role of the Naive Portfolio in the Empirical Design

The naive portfolio serves as a model-free benchmark that provides a lower bound on the complexity required to achieve satisfactory diversification. Unlike the historical MVP and the IV-MVP, the naive portfolio does not exploit any information about return variances, correlations, or volatility predictability.

By comparing the OOS performance of the optimized portfolios to that of the naive portfolio, the analysis evaluates whether the additional complexity introduced by volatility forecasting and covariance estimation delivers economically meaningful improvements over naive diversification.

If the IV-MVP achieves lower realized variance than the naive portfolio, this provides strong evidence that the predictive information contained in IV can be translated into tangible portfolio risk reductions beyond what can be achieved through equal weighting alone.

4.6 Statistical Inference in Predictive Regressions

Statistical inference in the predictive regressions is based on standard t-tests of the estimated slope coefficients. For each regression specification, the null hypothesis states that the predictor does not contain predictive information for the dependent variable.

Formally, inference is conducted by testing:

$$H_0: \beta_1 = 0$$
$$H_1: \beta_1 \neq 0$$

Where β_1 denotes the coefficient on the lagged volatility measure (IV, HV, or the VRP).

A rejection of the null hypothesis indicates that the predictor contains statistically significant predictive content. All tests are conducted using robust standard errors as specified in Section 4.2.

4.7 Statistical Inference in Portfolio Risk

To assess whether observed differences in portfolio variance are statistically significant, formal hypothesis testing is employed.

Specifically, a one-sided F-test for equality of variances is used to test whether the variance of the IV-MVP portfolio is lower than that of the benchmark portfolios. The hypotheses are specified as:

$$H_0: \sigma^2_{IV-MVP} = \sigma^2_{benchmark}$$

$$H_1: \sigma^2_{IV-MVP} < \sigma^2_{benchmark}$$

Where σ_i^2 denotes the variance of portfolio returns for portfolio *i*.

The test statistic is defined as:

$$F = \frac{s^2_{IV-MVP}}{s^2_{benchmark}} \tag{4.10}$$

Where s_i^2 represents the sample variance of portfolio returns for portfolio *i*.

Rejection of the null hypothesis indicates that the IV-MVP achieves statistically significant variance reduction relative to the benchmark portfolio.

Given that the IV-MVP is designed to minimize risk rather than maximize returns, differences in returns and risk-adjusted performance measures are discussed descriptively rather than subjected to formal statistical testing.

5 Empirical Results

5.1 Volatility Forecasting Results

Implied Volatility and Realized Volatility

Table 5.1. Panel Regression Results: Implied Volatility predicting Realized Volatility

$RV_t = \alpha + \beta_1 \cdot IV_{t-1} + \varepsilon_t$				
Variable	**(1)**	**(2)**	**(3)**	**(4)**
β_{IV}	0.875***	0.732***	0.899***	0.634***
	(0.044)	(0.060)	(0.034)	(0.113)
Firm FE	No	Yes	No	Yes
Time FE	No	No	Yes	Yes
Observations	1,708	1,708	1,708	1,708
R^2	0.383	0.197	0.411	0.115

Notes: Robust standard errors clustered by firm (Firm FE and Two-Way FE) and by time (Time FE). *, **, *** denote significance at the 10%, 5%, and 1% levels. Reported R^2 values correspond to within R^2 for fixed-effects models and overall R^2 for the RE model.

The results in Table 5.1 show that IV is a strong and statistically significant predictor of future RV across all model specifications. The estimated coefficients are consistently positive, indicating that higher IV is associated with higher subsequent RV.

Column (2), which reports the baseline specification with firm FE, shows that IV remains highly statistically significant with an estimated coefficient of 0.732. This indicates that, within firms, increases in IV are followed by economically meaningful increases in RV. Columns (3) and (4) additionally control for common time effects and for both firm and time FE, respectively. While the magnitude of the coefficient is somewhat reduced when time effects are included, IV remains positive and statistically significant, suggesting that the predictive relationship is not solely driven by common market-wide volatility components.

These findings are in line with Christensen and Prabhala (1998), who document that IV significantly predicts RV in the U.S. equity market. The results presented here extend this evidence to the Swedish equity market, suggesting that the forward-looking informational content of IV is not limited to large and highly liquid markets.

Historical Volatility and Realized Volatility

Table 5.2. Panel Regression Results: Historical Volatility Predicting Realized Volatility

$RV_t = \alpha + \beta_1 \cdot HV_{t-1} + \varepsilon_t$				
Variable	**(1)**	**(2)**	**(3)**	**(4)**
β_{HV}	0.424*** (0.062)	0.210*** (0.053)	0.432*** (0.037)	0.105*** (0.064)
Firm FE	No	Yes	No	Yes
Time FE	No	No	Yes	Yes
Observations	1,708	1,708	1,708	1,708
R^2	0.180	0.045	0.186	0.011

Notes: Robust standard errors clustered at the firm level (Firm FE and RE), at the time level (Time FE), and two-way clustered by firm and time (Two-Way FE). *, **, *** denote significance at the 10%, 5%, and 1% levels. Reported R^2 values correspond to within R^2 for fixed-effects models and overall R^2 for the RE model.

The results in Table 5.2 show that HV is a statistically significant predictor of future RV across all specifications. The estimated coefficients are consistently positive, indicating that higher HV is associated with higher subsequent RV. This finding is consistent with Blair et al., (2001), who document that HV often predicts future RV.

In the baseline specification with firm FE, HV remains statistically significant, indicating that within-firm changes in HV are associated with future RV. When additional FE are included, either controlling for common time effects or for both firm and time effects, the estimated relationship remains positive but is reduced in the most restrictive specification. This pattern suggests that HV primarily captures persistent risk differences across firms and common market-wide volatility conditions, rather than new firm-specific information about future risk.

Implied Volatility and Realized Volatility Firm Level

Table A.4 reports firm-level time-series regression results assessing whether IV at the individual firm level predicts future RV for firms included in the OMXS30 index. The table presents estimates for 28 out of the 30 constituent firms, with 61 observations per firm. The estimated coefficient on IV is positive and statistically significant for 27 of the 28 firms, with Svenska Handelsbanken being the only firm for which the coefficient is not statistically significant.

5.2 Return Predictability Results

Implied Volatility and Future Returns

Table 5.3. Panel Regression Results: Implied Volatility Predicting Excess Returns

$\boldsymbol{R_t = \alpha + \beta_1 \cdot IV_{t-1} + \varepsilon_t}$				
Variable	**(1)**	**(2)**	**(3)**	**(4)**
β_{IV}	0.047* (0.028)	0.094** (0.045)	0.019 (0.037)	0.061 (0.066)
Firm FE	No	Yes	No	Yes
Time FE	No	No	Yes	Yes
Observations	1,708	1,708	1,708	1,708
R^2	0.002	0.005	0.000	0.001

Notes: Robust standard errors clustered at the firm level (Firm FE and RE), at the time level (Time FE), and two-way clustered by firm and time (Two-Way FE). *, **, *** denote significance at the 10%, 5%, and 1% levels. Reported R^2 values correspond to within R^2 for fixed-effects models and overall R^2 for the RE model.

Focusing on the baseline specification with firm FE reported in Table 5.3, IV is positive and statistically significant, although the estimated effect is economically small. When time FE are included, either alone or together with firm FE, the estimated relationship becomes statistically insignificant, indicating that the return predictability associated with IV is sensitive to model specification and not robust to controlling for common time effects.

These findings are broadly consistent with Ammann et al., (2009) and Banerjee et al., (2007), who document statistically significant return predictability associated with IV. At the same time, both studies emphasize that explanatory power is low and that the relationship is unstable, a pattern that is also reflected in the present results.

Historical Volatility and Future Returns

Table 5.4. Panel Regression Results: Historical Volatility Predicting Excess Returns

$R_t = \alpha + \beta_1 \cdot HV_{t-1} + \varepsilon_t$				
Variable	**(1)**	**(2)**	**(3)**	**(4)**
β_{HV}	-0.013 (0.020)	-0.016 (0.021)	-0.011 (0.026)	-0.014 (0.021)
Firm FE	No	Yes	No	Yes
Time FE	No	No	Yes	Yes
Observations	1,708	1,708	1,708	1,708
R^2	0.000	0.000	0.000	0.000

Notes: Robust standard errors clustered at the firm level (Firm FE and RE), at the time level (Time FE), and two-way clustered by firm and time (Two-Way FE). *, **, *** denote significance at the 10%, 5%, and 1% levels. Reported R^2 values correspond to within R^2 for fixed-effects models and overall R^2 for the RE model.

The results in Table 5.4 show no evidence that HV predicts future excess returns. All specification estimated coefficients are consistently negative, economically negligible and statistically insignificant across all specifications with zero explanatory power. This indicates that backward-looking volatility measures do not contain useful information about returns. The findings are not surprising considering Campbell et al. (1988), Ghysels et al., (2005), and Corrado and Miller (2006) all show HV is not a strong predictor of future returns.

Variance Risk Premia and Future Returns

Table 5.5. Panel Regression Results: Variance risk premia Predicting Excess Returns

$R_t = \alpha + \beta_1 \cdot VRP_{t-1} + \varepsilon_t$				
Variable	**(1)**	**(2)**	**(3)**	**(4)**
β_{VRP}	0.060** (0.027)	0.060** (0.026)	0.034 (0.027)	0.034* (0.018)
Firm FE	No	Yes	No	Yes
Time FE	No	No	Yes	Yes
Observations	1,708	1,708	1,708	1,708
R^2	0.004	0.005	0.001	0.001

Notes: Robust standard errors clustered at the firm level (Firm FE and RE), at the time level (Time FE), and two-way clustered by firm and time (Two-Way FE). *, **, *** denote significance at the 10%, 5%, and 1% levels. Reported R^2 values correspond to within R^2 for fixed-effects models and overall R^2 for the RE model.

The results in Table 5.5 show that the VRP is positively related to future excess returns. Across all specifications, the estimated coefficient on the VRP is positive. Focusing on the baseline specification with firm FE, the VRP is statistically significant, indicating that higher VRP are associated with higher future excess returns. The estimated relationship remains positive when time FE are included, although statistical significance is somewhat sensitive to model specification.

Although the explanatory power of the regressions is low, this pattern is consistent with previous literature. The findings are in line with Bollerslev et al., (2009) and Drechsler and Yaron (2011), who show that the VRP contains information about time-varying risk premia and predicts future stock market returns.

5.3 Portfolio Results

This section presents the empirical performance results for the four portfolio strategies over the full OOS period from November 2022 to October 2025. The evaluated portfolios are the IV-MVP, the MVP, the naive portfolio, and the OMXS30 index. Performance is reported using return-based, risk-based, and risk-adjusted measures. Since IV is found to be a statistically significant predictor of future RV for 27 out of the 28 firms at the firm level according to Table A.4, the portfolio constructions are based exclusively on these 27 significant stocks.

Return and Volatility

Table 5.6. Portfolio return and volatility performance, during the whole testing period

Return and Volatility				
Portfolio	**Annual Return (%)**	**Cumulative return (%)**	**Annual Volatility (%)**	**Maximum Drawdown (%)**
IV-MVP	13.00	44.27	*11.86*	**-15.67**
MVP	12.29	41.60	12.44	-17.17
Naive Portfolio	*11.22*	*37.59*	14.90	-21.06
OMXS30	**13.26**	**45.29**	**15.81**	*-21.68*

In Table 5.6, return and volatility measures for all four portfolios are presented, while Figure 5.1 shows the portfolio price development. The OMXS30 index and the IV-MVP record the highest returns, while the naive portfolio has the lowest average annual return. The OMXS30 index and the naive portfolio display the highest volatility, whereas the IV-MVP records the lowest volatility. The OMXS30 index records the highest MDD in absolute terms, while the IV-MVP records the lowest in absolute terms.

Guidolin and Wang (2023) report OOS results for the IV-MVP, historical MVP, and naive portfolio under weekly rebalancing using an unconstrained portfolio setting. In their baseline specification, the IV-MVP achieves lower RV than both the historical MVP and the naive portfolio but also delivers the lowest average return among the compared strategies.

The results in Table 5.6 display both similarities and differences relative to these findings. Consistent with Guidolin and Wang (2023), the IV-MVP in the present study achieves the lowest RV among all portfolios. In contrast to their results, the IV-MVP also delivers strong return performance, ranking second only to the OMXS30 index over the full OOS period. The naive portfolio displays higher volatility and lower returns, while the historical MVP lies between the IV-MVP and the naive portfolio in terms of risk. With respect to MDD, Guidolin and Wang (2023) report that the historical MVP records the lowest drawdown, while the IV-MVP reduces drawdown relative to the naive portfolio but does not dominate the historical MVP. In Table 5.6, the IV-MVP records the lowest MDD among all portfolios, although the absolute drawdown levels differ across the two studies.

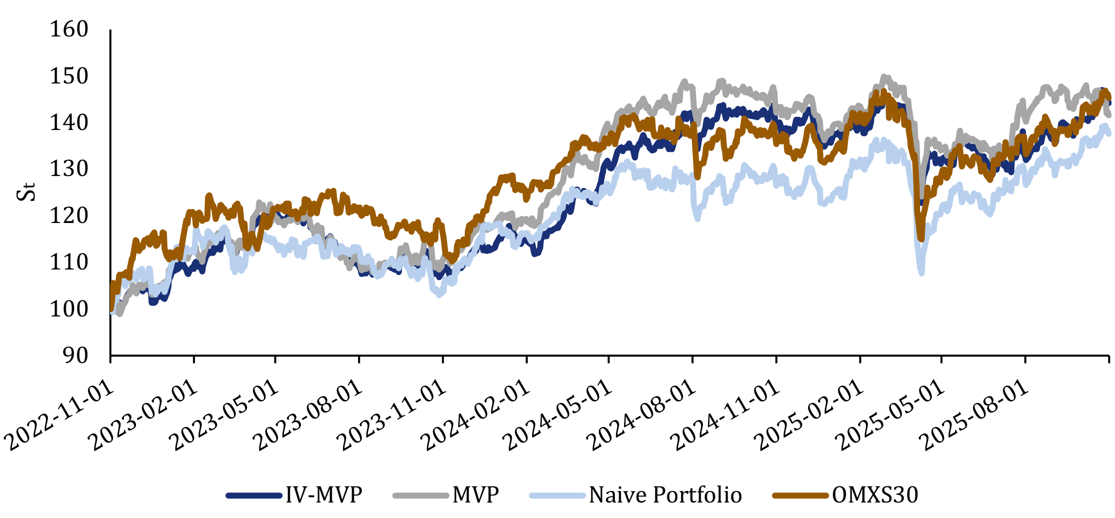

Figure 5.1. Cumulative price performance over time of IV-MVP, MVP, Naive portfolio and OMXS30

Risk-Adjusted Performance

Table 5.7. Portfolio Sharpe Ratio, Sortino Ratio and Calmer Ratio, for the whole testing period

Risk-Adjusted Performance			
Portfolio	**Sharpe Ratio**	**Sortino Ratio**	**Calmar Ratio**
IV-MVP	**0.8451**	**1.1961**	**0.8295**
MVP	0.7493	1.0598	0.7157
Naive Portfolio	*0.5538*	*0.7789*	*0.5329*
OMXS30	0.6508	0.9020	0.6118

In Table 5.7, risk-adjusted performance measures for all four portfolios are presented. The IV-MVP achieves the highest Sharpe ratio, while the naive portfolio records the lowest. Once again, the IV-MVP outperforms the other portfolios by achieving the highest Sortino ratio, while the naive portfolio records the lowest value. The IV-MVP achieves the highest Calmar ratio, whereas the naive portfolio again records the lowest value.

Guidolin and Wang (2023) report the Sharpe and Sortino ratios for the portfolios. In their baseline specification, the MVP achieves the highest Sharpe and Sortino ratios, while the IV-MVP ranks below the MVP but above the naive portfolio.

In contrast to Guidolin and Wang (2023), the IV-MVP in the present study attains the highest Sharpe and Sortino ratios, while the MVP ranks second. The naive portfolio records the lowest risk-adjusted performance in both studies.

Variance Hypothesis Testing

Table 5.8. Hypothesis testing for IV-MVP variance difference against benchmarks, during the whole period, using one-side F-tests

One-Sided F-Tests of Portfolio Variance		
Test	**F-Value**	**P-Value**
IV-MVP vs. MVP	0.909	0.096*
IV-MVP vs. Naive	0.634	0.000***
IV-MVP vs. OMXS30	0.563	0.000***

This section examines whether the IV-MVP exhibits a statistically significantly lower volatility than the benchmark portfolios. The test statistics and corresponding p-values for the variance differences between the IV-MVP and the benchmark portfolios are reported in Table 5.8. The results indicate statistically significant differences in variance between the IV-MVP and both the naive portfolio and the OMXS30 index at conventional significance levels.

The variance difference between the IV-MVP and the MVP is marginally significant at the 10% level, while the variance differences relative to the naive portfolio and the OMXS30 index are significant at the 1% level.

Portfolio Turnover

Table 5.9. Portfolio turnover over the whole testing period

Portfolio	Average Annual Turnover
IV-MVP	0.676
MVP	0.203
Naive Portfolio	0.000

Table 5.9 presents the average annual turnover for the portfolio strategies over the full evaluation period. The IV-MVP exhibits the highest turnover, with an average value of 0.676, indicating relatively frequent rebalancing. In contrast, the MVP shows a substantially lower turnover of 0.203. The naive portfolio displays zero turnover, as it is constructed at the beginning of the sample period and held without rebalancing throughout the entire horizon. As noted in the methodology section, turnover for the OMXS30 index cannot be computed due to data limitations.

Guidolin and Wang (2023) report that IV-MVP exhibit substantially higher turnover than both historical MVPs and naive portfolios.

The results in Table 5.9 are consistent with these findings. The IV-MVP displays the highest average annual turnover over the full evaluation period, while the MVP exhibits a markedly lower turnover. As expected, the naive portfolio records zero turnover.

5.4 Discussion

Volatility Forecasting: Implied Volatility vs. Historical Volatility

The volatility regressions provide the clearest empirical support for our study. Focusing on the firm FE specification, IV emerges as a strong and robust predictor of future RV, while HV displays substantially weaker power. This difference is central to our findings and highlights an important asymmetry in the informational content of option-implied and backward-looking volatility measures (Tables 5.1 and 5.2).

The strong performance of IV, when comparing it to HV is consistent with Christensen and Prabhala (1998) and Blair et al., (2001). A natural explanation for this asymmetry is that IV is forward-looking by construction, as it is extracted from options prices that reflect market participants expectations about future volatility. Christensen and Prabhala (1998) emphasize that option prices aggregate information about anticipated macroeconomic conditions, firm-specific news and changes in risk perceptions. By contrast, HV is backward looking and only resects to shocks after they have been realized. As a result, it primarily captures persistent component of volatility rather than changes in expected future risk.

Blair et al., (2001) show that IV contains most of the information that can be found in HV. This supports the view that option markets quickly incorporate expectations about future risk. Our results suggest that this is also true for the Swedish equity option market, even though it is smaller and less liquid.

Return Predictability

The return predictability results provide more nuanced evidence than the volatility forecasting regressions. IV is positively related to future excess returns, although the estimated effect is economically small and characterized by low explanatory power. This suggests that IV captures aspects of time-varying risk premia rather than providing a strong or stable return forecasting signal. These findings are consistent with Ammann et al., (2009) and Banerjee et al., (2007), who document weak and sensitive return predictability associated with IV.

In contrast, HV shows no predictive power for future excess returns. The consistently insignificant coefficients indicate that backward-looking volatility measures do not contain relevant information about expected returns, in line with earlier evidence from Campbell et al. (1988), Ghysels et al., (2005), and Corrado and Miller (2006).

Evidence of return predictability is also observed for the VRP. The positive relationship between the VRP and future excess returns is consistent with the interpretation that investors require compensation for bearing volatility risk, as documented by Bollerslev et al., (2009) and Drechsler and Yaron (2011). However, the explanatory power is low and comparable to that obtained for IV alone, suggesting that both measures capture related but economically modest components of time-varying risk compensation.

From the perspective of asset pricing theory, these results do not contradict market efficiency or the intuition of the CAPM. This is because the CAPM links expected returns to systematic market risk exposure, not to the level of volatility itself, which is therefore not a directly priced factor in the model. The absence of return predictability from HV is consistent with weak-form efficiency, as past price information should not generate abnormal returns. While IV and the VRP display weak statistical return predictability, the estimated effects are economically small, unstable across specifications, and exhibit very low explanatory power. As a result, this predictability is unlikely to support a trading strategy capable of generating abnormal returns and is therefore compatible with semi-strong efficiency rather than a violation of the EMH. Prior evidence suggests that the VRP is related to expected returns (Bollerslev & Zhou, 2007), but the weaker coefficients and low explanatory power observed in the Swedish equity market indicate that the pricing of volatility-related risk is limited in magnitude and difficult to exploit empirically.

Portfolio Implications

The empirical results show that the IV-MVP achieves the lowest realized variance among all considered portfolios, outperforming the MVP, the naive portfolio, and the OMXS30 benchmark. This confirms the theoretical prediction of the minimum variance framework that portfolio risk can be minimized when the covariance matrix is accurately estimated. Importantly, the results demonstrate that the effectiveness of the MVP strategy is enhanced when IV, rather than HV, is used as the volatility input.

This outcome is consistent with classical portfolio theory introduced by Markowitz (1952), which emphasizes that the MVP depends exclusively on second moments. By replacing backward-looking HV with forward-looking IV, the IV-MVP benefits from more timely and informative estimates of future risk. This is further supported by option pricing theory, as IV is extracted from market prices within the Black–Scholes–Merton framework and reflects investors' expectations about future risk rather than past realizations (Hull, 2022).

The finding that the IV-MVP exhibits significantly lower variance than the naive portfolio and the OMXS30, and weakly significantly lower variance than the MVP, aligns closely with previous empirical evidence. Guidolin and Wang (2023) show that incorporating IV improves portfolio risk forecasts relative to historical-based strategies, and the outperformance relative to the naive diversification portfolio is particularly notable given its documented robustness to estimation error (DeMiguel et al., 2009). While prior studies primarily rely on descriptive comparisons of realized variance, the present study formally tests whether differences in realized variance are statistically significant. The present study therefore extends these results to the Swedish equity market.

In addition to lower variance, the IV-MVP also displays a substantially lower MDD, indicating improved protection against severe downside risk. Although the IV-MVP exhibits higher portfolio turnover, implying higher potential transaction costs, the overall results suggest that the economic benefits from improved risk estimation remain substantial. Differences relative to prior studies may partly reflect the lower rebalancing frequency employed in this study, as IV retains its forward-looking informational content even when portfolios are rebalanced less frequently.

Overall, the findings demonstrate that using IV as a forward-looking risk measure enables the construction of a portfolio that outperforms traditional benchmarks in terms of risk minimization. These results extend the validity of the MVP framework beyond the U.S. equity market and highlight the relevance of IV-based portfolio strategies for risk-averse investors such as pension funds and institutional asset managers.

6 Conclusion

This study examined whether IV derived from Swedish equity options provides superior information relative to HV in forecasting future RV and stock returns, and whether this information can improve portfolio risk outcomes. Motivated by the U.S.-based literature, the objective was to assess whether option-implied information retains its relevance in a smaller and less liquid equity market.

The empirical evidence delivers a clear message. IV outperforms HV in predicting both future RV and future stock returns. While the predictive power for returns is economically small, IV consistently dominates HV, which exhibits no predictive ability for returns. These results confirm that forward-looking information embedded in option prices is more informative than backward-looking measures based on past returns and that IV primarily reflects expectations about future risk rather than serving as a strong return forecasting signal.

The main contribution emerges in the portfolio analysis. An IV-MVP achieves the lowest realized variance and lower downside risk than standard benchmarks, including the traditional MVP, the naive portfolio, and the OMXS30 index. Importantly, these improvements are obtained using a quasi-IV-based construction, where IV is used to estimate variances while historical correlations are retained. Even with this conservative implementation, the performance gains are economically meaningful, suggesting that the benefits of IV may be understated. Overall, the results for the Swedish equity market are closely aligned with those documented in prior studies based on U.S. equity markets.

Taken together, the findings demonstrate that IV enhances the effectiveness of the classical minimum variance framework proposed by Markowitz (1952) and that this insight extends beyond U.S. equity markets to the Swedish equity market. More broadly, the results highlight the practical relevance of IV for portfolio risk management. The findings also point to promising avenues for future research, including the construction of a fully IV-MVP that relies exclusively on option-implied measures. The evaluation of such strategies over longer sample periods, alternative rebalancing frequencies, and the assessment of their performance across varying market conditions characterized by high and low volatility.

References

Ammann, M., Verhofen, M., & Süss, S. (2009). Do Implied Volatilities Predict Stock Returns, *Journal of Asset Management*, vol. 10, no. 4, pp. 222–234, https://dx.doi.org/10.2139/ssrn.1670909

Andersen, T.G., Bollerslev, T., Diebold, F.X., & Labys, P. (2003). Modeling and Forecasting Realized Volatility, *Econometrica*, vol. 71, no. 2, pp. 579–625, https://doi.org/10.1111/1468-0262.00418

Bali, T.G., & Hovakimian, A. (2009). Volatility Spreads and Expected Stock Returns, *Journal of Financial Economics*, vol. 92, no. 3, pp. 384–411, https://doi.org/10.1287/mnsc.1090.1063

Banerjee, P.J., Doran, J.S., & Peterson, D.R. (2007). Implied Volatility and Future Portfolio Returns, *Journal of Banking & Finance*, vol. 31, no. 10, pp. 3183–3199. https://doi.org/10.1016/j.jbankfin.2006.12.007

Blair, B., Poon, S.H., & Taylor, S.J. (2001). Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High-Frequency Index Returns, *Journal of Econometrics*, vol. 105, no. 1, pp. 5–26, https://doi.org/10.1016/S0304-4076(01)00068-9

Bodie, Z., Kane, A., & Marcus, A.J. (2019). *Investments*, 11th edn, New York: McGraw-Hill Education.

Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected Stock Returns and Variance Risk Premia, *Review of Financial Studies*, vol. 22, no. 11, pp. 4463–4492, https://doi.org/10.1093/rfs/hhp008

Campbell, J.Y., French, K.R., Schwert, G.W., & Stambaugh, R.F. (1988). The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies*, vol. 1, no. 3, pp. 195–228, https://doi.org/10.1093/rfs/1.3.195

Christensen, B.J., & Prabhala, N.R. (1998). The Relation Between Implied and Realized Volatility, *Journal of Financial Economics*, vol. 50, no. 2, pp. 125–150, https://doi.org/10.1016/S0304-405X(98)00034-8

Corrado, C.J., & Miller, T.W. (2006). Estimating Expected Excess Returns Using Historical And Option-Implied Volatility, https://doi.org/10.1111/j.1475-6803.2006.00168.x

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal Versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?, *Review of Financial Studies*, vol. 22, no. 5, pp. 1915–1953, https://doi.org/10.1093/rfs/hhm075

Doran, J.S., & Peterson, D.R. (2007). Implied Volatility and Expected Returns, *Journal of Banking & Finance*, vol. 31, no. 5, pp. 1357–1376, https://doi.org/10.1002/fut.20279

Dow, C.G. (2007). *Portfolio Turnover and Common Stock Holding Periods*, Dow Publishing Company, Inc.

Drechsler, I., & Yaron, A. (2011). What's Vol Got to Do with It?, *Review of Financial Studies*, vol. 24, no. 1, pp. 1–45, https://doi.org/10.1093/rfs/hhq085

Fama, E.F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance*, vol. 25, pp. 383–417, https://doi.org/10.2307/2325486

Ghysels, E., Santa-Clara, P., & Valkanov, R. (2005). There Is a Risk–Return Trade-off After All, *Journal of Financial Economics*, vol. 76, no. 3, pp. 509–548, https://doi.org/10.1016/j.jfineco.2004.03.008

Guidolin, M., & Wang, H. (2023). Implied Volatility, Portfolio Choice, and Economic Value, *Journal of Financial Economics*, vol. 147, no. 1, pp. 1–29, https://doi.org/10.1016/j.physa.2023.128496

Hull, J.C. (2022). *Options, Futures, and Other Derivatives*, 11th edn, Harlow: Pearson Education.

Lintner, J. (1965). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, *Review of Economics and Statistics*, vol. 47, no. 1, pp. 13–37, https://doi.org/10.2307/1924119

Maillard, S., Roncalli, T., & Teiletche, J. (2010). The Properties of Equally Weighted Risk Contribution Portfolios, *The Journal of Portfolio Management*, vol. 36, no. 4, pp. 60–70, https://doi.org/10.3905/jpm.2010.36.4.060

Markowitz, H. (1952). Portfolio Selection, *Journal of Finance*, vol. 7, no. 1, pp. 77–91, https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Mossin, J. (1966). Equilibrium in a Capital Asset Market, *Econometrica*, vol. 34, no. 4, pp. 768–783, https://doi.org/10.2307/1910098

Sharpe, W.F. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance*, vol. 19, no. 3, pp. 425–442, https://doi.org/10.2307/2977928

Sharpe, W. F. (1966). ”Mutual Fund Performance”. The Journal of business, vol.39 (no.1), pp.119–138, https://www.jstor.org/stable/2351741

Sveriges Riksbank. (2025). Benchmark government bond yields, https://www.riksbank.se

Appendix

A.1 Numerical Solution for Implied Volatility Derived from Put-Call Parity

IV is obtained by numerically inverting the BSM pricing formula such that the theoretical option price equals the observed market price. Since the BSM option pricing formula cannot be analytically solved for volatility, IV is defined as the solution to the nonlinear pricing equation:

$$p_M - p_{BSM}(\sigma) = 0 \tag{A.1}$$

or, equivalently, as the value of volatility that minimizes the squared pricing error:

$$\sigma_{imp} = arg\ min_{\sigma > 0}\ (p_M - p_{BSM}(\sigma))^2 \tag{A.2}$$

Under the absence of arbitrage, put–call parity must hold for both theoretical BSM prices and observed market prices. For BSM prices, put–call parity is given by:

$$c_{BSM} - p_{BSM} = Se^{-qT} - Ke^{-rT} \tag{A.3}$$

PCP also holds for observed market prices:

$$c_M - p_M = Se^{-qT} - Ke^{-rT} \tag{A.4}$$

Since the right-hand sides are identical, it follows that:

$$c_M - p_M = c_{BSM} - p_{BSM} \Leftrightarrow c_M - c_{BSM} = p_M - p_{BSM} \tag{A.5}$$

Let σ_{imp} denote the IV obtained from the put option and insert $\sigma = \sigma_{IV}$ into the BSM pricing formulas for both the call and the put:

$$c_M - c_{BSM}(\sigma_{IV}) = p_M - p_{BSM}(\sigma_{IV}) \tag{A.6}$$

By definition of IV:

$$p_M - p_{BSM}(\sigma_{IV}) = 0 \tag{A.7}$$

It therefore follows that:

$$c_M - c_{BSM}(\sigma_{IV}) = 0 \tag{A.8}$$

Hence, the IV is identical for the call option and the put option. Consequently, European call and put options with the same strike price and maturity exhibit an identical volatility smile or volatility smirk. While equity options typically exhibit a volatility smirk, other asset classes, such as foreign exchange options, more often display a symmetric volatility smile (Hull, 2022).

A.2 Hausman Specification Tests

Table A.1. Hausman Specification Test statistics for panel data regressions.

Regression	Dependent Variable	Explanatory Variable	P-Value	Preferred Model
(1)	RV	IV	0.000***	FE
(2)	RV	HV	0.000***	FE
(3)	Excess Returns	IV	0.039**	FE
(4)	Excess Returns	HV	0.767	RE
(5)	Excess Returns	VRP	0.804	RE

A.3 Portfolio Correlation

Table A.2. Portfolio correlation over the whole testing period

Portfolio Correlation						
	Annual Return	Annual Volatility	Correlations			
Portfolio	(%)	(%)	IV-MVP	MVP	Naive	OMXS30
IV-MVP	13.00%	11.86%	1	0.9068	0.6426	0.3383
MVP	12.29%	12.44%		1	0.6618	0.3355
Naive	11.22%	14.90%			1	0.4432
OMXS30	13.26%	15.81%				1

A.4 List of Stocks

Table A.3. Complete list of stocks, data acquired from Yahoo Finance

Company	Ticker	Sector
ABB	ABB	Industrials
Alfa Laval	ALFA	Industrials
Assa Abloy AB	ASSA B	Industrials
Astra Zeneca	AZN	Healthcare
Atlas Copco	ATCO A	Industrials
Boliden	BOL	Basic Materials
Epiroc	EPI A	Industrials
EQT	EQT	Financial Services
Ericsson	ERIC B	Technology
Essity	ESSITY B	Consumer Defensive
Evolution	EVO	Consumer Cyclical
Hexagon	HEXA B	Technology
HM	HM B	Consumer Cyclical
Industrivärden	INDU C	Financial Services
Investor	INVE B	Financial Services
NIBE	NIBE B	Industrials
Nordea	NDA SE	Financial Services
Sandvik	SAND	Industrials
SCA	SCA B	Basic Materials
SEB	SEB A	Financial Services
Skanska	SKA B	Industrials
SKF	SKF B	Industrials
Svenska Handelsbanken	SHB A	Financial Services
SAAB	SAAB B	Industrials
Swedbank	SWED A	Financial Services
Tele2	TEL2 B	Communication Services
Telia	TELIA	Communication Services
Volvo	VOLV B	Industrials

A.5 Implied Volatility and Realized Volatility Regression Firm Level

Table A.4. Implied Volatility and Realized Volatility Firm Level Regression

Company	Constant	β (IV)	R^2	n
ABB	0.032*	0.636***	0.194	61
Alfa Laval	0.007	0.991***	0.421	61
Assa	0.042	0.770***	0.364	61
Astra Zeneca	0.030	0.852***	0.178	61
Atlas Copco	0.032	0.881***	0.284	61
Boliden	0.108*	0.798***	0.238	61
Epiroc	0.044	0.882***	0.335	61
EQT	0.058	0.845***	0.360	61
Ericsson	0.01	0.914***	0.328	61
Essity	0.045	0.889***	0.249	61
Evolution	0.097	0.767***	0.290	61
Hexagon	0.082	0.690***	0.135	61
HM	0.014	0.969***	0.305	61
Industrivärden	0.032	0.887***	0.376	61
Investor	0.052	0.733***	0.265	61
NIBE	0.141**	0.689***	0.218	61
Nordea	0.090**	0.554***	0.140	61
Sandvik	0.055	0.805***	0.390	61
SCA	0.080**	0.640***	0.200	61
SEB	0.036	0.865***	0.226	61
Skanska	0.028	1.125***	0.371	61
SKF	0.062	0.810***	0.273	61
Sv. Handelsbanken	0.215**	0.066	0.001	61
SAAB	0.028	0.980***	0.286	61
Swedbank	0.083	0.624***	0.112	61
Tele2	0.061	0.817**	0.071	61
Telia	0.013	1.033***	0.266	61
Volvo	0.041	0.883***	0.211	61

Notes: Inference in the firm-level regressions is based on heteroskedasticity- and autocorrelation-consistent (HAC) standard errors using the Newey–West estimator. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels.

Statement on the Use of Generative AI

This bachelor's thesis is an independent academic work conducted at the undergraduate level, with the purpose of demonstrating the authors' ability to independently carry out a study in economics. The authors bear full responsibility for the content of the thesis and have produced all text themselves.

With ongoing technological development, generative artificial intelligence has become increasingly prevalent, including in academic writing. The growing use of such tools raises important questions regarding transparency and academic integrity.

In this thesis, generative AI has been used in accordance with the guidelines of the Department of Economics. Specifically, the following tool has been employed:

- ChatGPT 5.2

The model has been used as a language-support tool to suggest grammatical improvements, refine formulations, and provide feedback on language, clarity, and academic style. All substantive content and arguments are authored by the students. The language model was only used to revise and improve text that had already been written and served as a complement to standard word-processing software.

ChatGPT has also been used as technical support for programming tasks in R, including assistance with code generation and clarification of practical and technical programming-related questions.

All use of generative AI has been carried out with caution and without compromising academic independence or integrity. The authors assess that the use of generative AI has had a limited influence on the outcome and that the thesis constitutes an independent academic work.