Bachelor Thesis [programme International Bachelor Economics and Business Economics]

Earnings uncertainty: Investor's ability to anticipate post-announcement realized volatility

On the informational content of implied volatility

Cedric S. Haberland

Date final version: August 6th, 2020

Student Number: 476680
Erasmus University Rotterdam: Erasmus School of Economics
Major: Financial Economics
Seminar: Stock Pricing and Investment Strategies
Supervisor: Esad Smajlbegovic
Second Assessor: Dr. Tim Eisert
The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics or Erasmus University Rotterdam

Abstract

The implied volatility of options gives the market forecast for the future volatility of the underlying security, over the remaining life of the option. Using standardized option prices, this paper assesses the accuracy of this volatility forecast for the 30 days following an earnings announcement for all historical components of the S&P 500 index over the time frame 1996 to 2019 and compares this to forecasts obtained using historical volatility. I find that implied volatility is a biased estimator of future volatility that produces forecasts superior to those obtained using historical volatility, while also containing all of the latter's informational content. The accuracy of forecasts obtained using implied volatility is furthermore shown to first increase with ex-ante uncertainty, before it begins to deteriorate again. The favored reasons for the bias in implied volatility are related to deficiencies in the Black-Scholes option pricing model that lead the calculated implied volatility to differ from the markets real volatility forecast.

Contents

1 Introduction

The implied volatility (IV) of options has seen a lot of interest from the academic field as a measure of the market's forecast of the underlying asset's volatility over the remaining life of the option. By reverting a formula such as that of Black & Scholes (1973), one can use the formula's other inputs, like the time to expiration, strike price, and interest rate, to obtain the volatility that sets the formula equal to the currently observed market price. Much existing literature (Canina & Figlewski, 1993; Christensen & Prahabla, 1998) focuses on the IV's predictive ability, especially in comparison with more traditional volatility measures such as the historical volatility (HV) of the underlying asset.

In the case of stocks as the underlying security, often implied and historical volatilities peak around a firms corporate events. According to Dubinsky, Johannes, et al. (2006), volatility in stock returns is indeed mostly concentrated in the days around earnings announcements. These announcements constitute one of the most important reoccurring corporate events, as they involve the disclosure of large amounts of material information about the company's ongoing and future performance, and therefore often have profound effects on a company's stock price. Earnings announcements usually occur each quarter that a firm is in operation after the firm becomes publicly traded. Since their occurrence is known, investors anticipate these events and form an expectation regarding the likely prospects of the firm. Similarly, investors anticipate the likelihood and potential magnitude of swings in the stock price as a result of the information release. Given that all other inputs are known for the relatively short time frame surrounding an earnings announcement, investor volatility expectations are reflected in the implied volatility of options in the days leading up to the event and can be calculated.

After the event has occurred, expectations will be adjusted. Since the new information has been released, implied volatility usually drops, a phenomenon day traders refer to as "IV-Crush", as the decrease in implied volatility leads to a drop in option prices. However, an unpublished working paper by Subramanyam, Marquardt, & Zhang (2005) suggests that this need not be the case. According to them, the earnings surprise affects the expectations of future volatility such that large absolute surprises can lead to an increase in post-announcement uncertainty, rather than a decrease.

Although existing research (such as Diavatopoulos, Doran, Fodor, & Peterson (2012)) outlines the informational value of IV regarding the direction of stock price movements following earnings announcements, little attention has been paid to whether investors correctly anticipate the magnitude of the price swings that follow. The goal of this research is therefore to investigate with what accuracy investor expectations of volatility induced by an earnings announcement, captured by the implied volatility of options prior to the announcement date, correspond to the volatility that is actually realized in response to the announcement.

By comparing the explanatory value of IV with that of the historical volatility of firms stock prices following earnings announcements, this paper furthermore aims to contribute a new facet to the existing literature discussing whether IV or HV is more valuable in making inferences regarding future volatility. Finally, the analysis may give further intuition on the idea that investors trading in option markets tend to have superior information than their stock-trading counterparts. If investors systematically over- or underestimate the volatility of a stock following earnings announcements, one should furthermore be able to construct trading strategies that can yield abnormal returns by employing non-directional option positions that focus on the underlying's volatility. This research may thus yield valuable insights regarding investor's ability to accurately forecast and anticipate risk, and whether potential flaws in investor behavior regarding IV occur systematically and can thus be exploited.

The following section gives an overview of the existing literature regarding the implied volatility of options and the applications it has found. In different subsections, I first discuss existing evidence on the explanatory power of IV and HV for the future volatility of a security. Afterwards, I pay attention to the different models one can use to measure IV, and how to arrive at one value for the IV, as its values usually differ over options with different strike prices on the same underlying security. Finally, the paper takes a brief look at literature concerning the behavior of IV around different corporate events, and whether investors in option markets indeed have information advantages. Section 3 discusses the data and sampling procedure. The methodology applied in analyzing the data to arrive at the results is outlined in section 4. In section 5 I present the results of the analysis, and section 6 concludes.

2 Theoretical Framework

Mayhew (1995) compiled an extensive review about the early literature regarding the implied volatility of options. The main research areas under this topic are the usefulness of implied volatility in volatility forecasting, the implied volatility vs. historical volatility debate, the adequacy of option pricing models like the Black-Scholes, as well as the information contained in differing IV estimates across options with different strike prices on the same underlying and expiration date, the so-called volatility smile. The term emerges from the shape obtained when plotting IV against strike prices, as options that are further in-the-money (ITM) or out-of-the-money (OTM) tend to have higher IV than at-the-money (ATM) options as described in Dumas, Fleming, & Whaley (1998).

2.1 Implied or historical volatility?

Early research indicated that IV gives a more accurate forecast of future realized volatility than the historical volatility of the underlying asset (Beckers, 1981; Mayhew, 1995). Beckers (1981) notes that most research over this period used the closing prices of options to calculate IV, leaving open the question of whether results may be influenced by whether the contract closed at the bid or ask price. In his own research, he found that adding HV as an explanatory variable added explanatory power to his model regressing realized volatility on the IV of ATM options, suggesting that options markets were not fully efficient at that time. Early results also indicated that which volatility measure is better suited for a given forecast may depend on the time horizon (Mayhew, 1995).

Canina & Figlewski (1993) reject the notion of IV being a better predictor of future volatility than the underlying's historical volatility. By investigating the forecasting power of the IV calculated from options on the underlying S&P 100 (OEX) index, the most commonly traded options at the time of their research, they found that IV lacks any meaningful correlation with future volatility. Furthermore, they compared their results to forecasts made using historical volatility and found the latter to be a far better predictor than IV. They initially state that their results do not necessarily mean that IV is a bad predictor, as they may be driven by forecast errors. However, since HV leads to much better predictions than IV, they eventually conclude that it is more likely that IV is simply a bad predictor of future volatility.

Christensen & Prabhala (1998) refute the findings of Canina & Figlewski (1993). According to them, the method of Canina & Figlewski contains a major flaw, as the use of option pairs with overlap in their remaining time to expiration leads the forecast errors for the implied volatilities calculated from these options to be correlated. This problem was acknowledged by Canina & Figlewski (1993), who tried to correct for time dependence in order to handle the serial correlation, as creating nonoverlapping observations would require aggregation and exclusion of some data, which would have reduced the power of their statistical tests. Christensen & Prahabla (1998) replicate the research of Canina & Figlewski (1993) using a longer time period and nonoverlapping data, where they find IV to be a good predictor for future volatility and to clearly outperform historical volatility in forecasting, in accordance with research such as that by Day & Lewis (1988), Harvey & Whaley (1992), and Sheikh (1989). They partly attribute the difference between their findings and those of Canina & Figlewski (1993) to a "regime shift around the October 1987 crash [that] explains why implied volatility is more biased in previous work" (Christensen & Prahabla, 1998), as the sample used by Canina & Figlewski (1993) ended in 1987.

Dumas et al. (1998) also reject HV as a predictor of future volatility since the measure is by definition backwards-looking. The conflicting results of Canina & Figlewski (1993) therefore indeed seem to have been caused by forecast errors, resulting in further support for IV's advantage over HV in forecasting applications. The two authors furthermore add that how well IV can predict the future volatility in part depends on its accuracy, which is generally higher the easier it is to conduct arbitrage on the underlying asset, so that option prices and prices of the underlying are in line. Since their research focuses on an index consisting of 100 stocks, they attribute their results in part to the excessive costs and complications involved in trying to conduct arbitrage on the underlying asset in their study. Consequently, IV should be a far better predictor of future volatility when it comes to analyzing individual securities, rather than indices. Furthermore, Christensen & Prahabla (1998) extended their research regarding potential bias and inefficiency of IV by applying an instrumental variable approach to circumvent potential measurement error caused by non-adjustment for dividends and their use of American, rather than European options for the OEX. Their findings provide evidence that IV as a predictor of future volatility is both unbiased and efficient.

2.2 Option pricing models

How well IV serves in forecasting future volatility may also depend on how it is being measured. Many researchers have touched upon the irony of using a model such as the Black-Scholes to derive the IV for options with different strike prices, as the model itself assumes volatility to be constant across strike prices. This assumption, however, does not hold empirically, leading to different IV estimates for different strike prices (Beckers, 1980; Canina & Figlewski, 1993; Dumas et al., 1998; Mayhew, 1995). While the Black-Scholes formula is accurate for relatively short-expiration options with strike prices near the current price of the underlying, for other options there are large and systematic differences between observed market and Black-Scholes prices (Mayhew, 1995). For put options, this effect is at least in part caused by risk aversion of investors. In order to protect themselves against large losses, they are willing to pay a premium when purchasing out-of-the-money (OTM) put options in order to hedge their positions.

Those apparent deficiencies in the Black-Scholes model prompted other researchers to look for better option pricing models. Dumas et al. (1998) compare a valuation model based on a deterministic volatility function (DVF) used by Derman & Kani (1994), Dupire et al. (1994), and Rubinstein (1994) to a procedure that simply smooths the different Black-Scholes implied volatilities across exercise prices and times to expiration. They describe the DVF as assuming that "the local volatility rate is a flexible but deterministic function of asset price and time" (Dumas et al., 1998) and find that the DVF model is not more valid regarding the underlying volatility function than the Black-Scholes procedure. According to Christensen & Prahabla (1998), for at-the-money (ATM) options there is no large difference between Black-Scholes IV and the expected future return volatility, even when returns follow a stochastic volatility model. Furthermore, stochastic volatility models do not always work well in practice, as their use requires a parameter for the market price of risk, which is hard to estimate accurately. Thus, despite its shortcomings, the Black-Scholes model remains the most commonly used in empirical research.

2.3 Implied volatility measures

Due to the different IVs obtained for different strike prices, applying the Black-Scholes formula leads to the question of which measurement to apply as the IV of the underlying asset. According to Dumas et al. (1998), options with Black-Scholes IV higher than average will be valued too low, and vice versa. Early literature suggests using an average. The simplest schemes use an equally-weighted average; however, this method does generally not result in a precise IV measurement. As the Black-Scholes model prices some options more accurately than others, the IV implied by those options prices should receive larger weights. Options that are near-the-money are most accurately priced by the Black-Scholes formula, contain the most information about future volatility, and thus generally receive higher weights.

Furthermore, in the older literature, it was found that options that have a high vega, which measures the option prices sensitivity to changes in the implied volatility of the underlying security, produce some of the best volatility forecasts (Mayhew, 1995). The CBOE volatility index (VIX) is similarly calculated using the weighted average IV of 4 calls and 4 put options, using the options with strike prices nearest the money. According to Xing, Zhang, & Zhao (2010) ATM calls are most often used as a benchmark for IV as they tend to be the most liquid option contracts, thus giving the most complete reflection of investor consensus about future volatility.

2.4 Corporate Events

IV has also been shown to be an important indicator when it comes to corporate events. Levy & Yoder (1993) investigate IV around M&A announcements, and research by Barone-Adesi, Brown, Harlow, et al. (1994) indicates that the IV of options with the stock of the target firm as the underlying may indicate the probability of a successful takeover. Xing et al. (2010) as well as Diavatopoulos, Doran, Fodor, & Peterson (2012) investigate the predictive power of information implied in the IV-smile calculated from options on a firm's stock in predicting stock returns, with Diavatopoulos et al. (2012) paying special attention to earnings announcements. They find that the implied skewness and kurtosis of IV contain information about stock returns around announcement dates, and that changes in implied skewness indicate whether the stock price is more likely to jump or dip following the announcement. Thus, if the prices of OTM call options increase relative to the prices of ATM or in-the-money (ITM) call options on the same underlying, this points to investor optimism regarding the announcement, whereas the opposite holds if the prices of OTM put options increase relative to their ATM and ITM counterparts.

In some situations, investors may be relatively sure that material information with large price implications will be released, but unsure whether the news will be positive or negative. According to Diavatopoulos et al. (2012), in such cases prices of both OTM calls and puts will increase relative to the prices of the same options with strike prices at-the-money, since the increased uncertainty leads to thicker tails (kurtosis) in the implied price distribution of the underlying asset.

Other studies that investigate the behavior of IV around earnings announcements include that of Subramanyam, Marquardt, & Zhang (2005). In this paper the authors relate analysts earnings forecast errors to changes in the level of implied volatilities around earnings announcements and find that the relationship between the earnings surprise and uncertainty regarding the value of the firm, as measured by IV, follows a V-shape. Small absolute surprises thus result in decreases in IV. This leads them to give earnings announcements with small absolute surprises a sort of confirmational role, where despite no information being released uncertainty decreases, possibly because investors were unsure about whether material information would be released or not. Their findings are interesting since most existing literature assumes that the release of new information, as for example per earnings announcements, decreases investor uncertainty (Isakov & Perignon, 2001; Patell & Wolfson, 1979; Truong, Corrado, & Chen, 2012). Subramanyam et al. (2005), however, find that large absolute earnings surprises may increase uncertainty about the firms value, leading to higher IVs, contrary to the commonly expected 'IV Crush'.

2.5 Information advantage of option traders

Diavatopoulos et al. (2012), and Xing et al. (2010) also strongly support the notion that there is more informed trading going on in option markets than in conventional stock markets, as they assume option traders to have informational advantages. For earnings announcements, this point of view is supported through findings by Amin & Lee (1997), that indicate large differences in volume changes between options and stock markets in the 4 days leading up to an announcement. During this time, the volume for options that have the firms stock as the underlying security increases by 10%, whereas stock volume only picks up by 5%. By investigating skewness and kurtosis of the implied volatility smile, Diavatopoulos et al. (2012) provide empirical evidence that option prices incorporate some information contained in earnings announcements before the announcement has actually happened, which may be an indication for information leaks or insider trading. They state that “at day (-5), option prices have virtually fully adjusted to the forthcoming earnings announcements” (Diavatopoulos et al., 2012).

Easley, O’hara, & Srinivas (1998) investigate the informational role of transaction volumes in options markets by sorting options into positive- and negative news trades. The former is measured as buyer-induced and the latter as seller-induced trades. They find that the volumes of those categories contain information about future stock prices, adding that if informed traders prefer the options market, information may be incorporated into options prices before even having affected the stock price. Indeed, Xing et al. (2010) find that implied volatility smirks contain information that can be used to predict firm performance for an astonishing time frame of up to 6 months due to sluggishness of the equity market in incorporating this information into prices, and that this information is related to firm fundamentals. They note that if the information advantage is substantial enough, the information is incorporated into prices as per the equilibrium model of Garleanu, Pedersen, & Poteshman (2008). Here, the additional demand by informed traders affects the price of options, as market makers may not always be able to perfectly hedge their positions and thus demand higher premiums for certain contracts.

There are many reasons why more informed traders may prefer option or other derivative markets. First, these markets are usually less regulated (regulatory arbitrage) and traders often incur lower transactions costs than when trading equities. Second, derivatives provide cheaper ways of betting against a firm, i.e. lower short-selling costs, and higher leverage is available to investors (Black, 1975). Indeed, the findings by Easley et al. (1998) are more pronounced for negative news trades, which points to options markets being more attractive especially for traders with negative information regarding the future prospects of a firm.

2.6 Hypotheses

Reviewing the existing literature provides the motivation behind the different hypotheses of this research. Given that traders in option markets seem to have an informational advantage, even though they may at times over- or underestimate the volatility following the earnings announcement, these misestimations should even out so that on average, investors get the volatility right. Thereby follows *H1: Over the whole sample, IV is an unbiased predictor of the post-earnings realized volatility.*

The amount of time that has passed between this research and that of Canina & Figlewski (1993) and Christensen & Prahabla (1998) implies that investors, if HV even had any incremental value over IV in explaining future volatility, should by now have learned to implement this into their own volatility forecasts that are reflected in today's IV. This leads to the second hypothesis behind this research: *H2: HV does not offer any explanatory power regarding realized volatility if IV is included in the model.*

High values of pre-earnings IV are caused by large amounts of uncertainty regarding the contents of the earnings announcements and, if investors can estimate post-announcement realized volatility with reasonable accuracy, thus tend to precede large absolute earnings surprises (Amin & Lee, 1997; Patell & Wolfson, 1979; Subramanyam et al., 2005). As outlined in Subramanyam et al. (2005), such surprises may even lead to increased uncertainty following the announcement. While a higher level of pre-earnings IV is more likely to lead to larger absolute differences between IV and the post-announcement realized volatility, there is no reason to believe that the market forecast should become less accurate when evaluating realized volatility relative to pre-earnings IV. Therefore, the accuracy of IV as a volatility forecast should be independent of the level of IV itself. Thus follows *H3: The accuracy of the forecasts of post-earnings realized volatility, produced by IV, does not depend on the value of pre-earnings IV.*

Finally, Diavatopoulos et al. (2012) touch on the idea that changes in uncertainty preceding the announcement may be more informative than levels of IV. On one hand, non-public information may only become available to some individuals relatively short-term before an earnings announcement is made. On the other, individuals that possess such information would want to enter positions as soon as possible to avoid missing out on profitable trades in case the information somehow becomes publicly available before the announcement. However, since the information may simply not have been created or sufficiently verified until very shortly before it is announced, it seems like the former effect should outweigh the latter. This leads to the final hypothesis of this research. *H4: Changes in the pre-earnings IV have explanatory power regarding post-earnings realized volatility.*

3 Data and sample selection

In order to analyze the above-mentioned hypotheses, this research compiles data from different sources. I obtain implied and historical, as well as realized volatilities from *Option Metrics*. The *I/B/E/S* database is consulted to obtain the earnings announcement dates for all current and historical component companies of the S&P 500 index, providing a broad sample representative of the American market. The sampling period ranges from January 1996 to December 2019. As this research mainly examines the explanatory power of implied volatilities, the sample is naturally tilted towards larger companies, as in order to be included, a company must have stock options being traded on their equity. This results in a data panel containing observations on 95 quarters and over 740 firms.

The paper employs data on additional explanatory variables such as firm betas and financial ratios to construct control variables. This data is obtained from the *CRSP* database, as well as the *Financial Ratios Suite* provided by the *Wharton Research Data Service (WRDS)*. All databases used are widely regarded as sources of high-quality, extensive, and recent data and have been extensively used in prior research, such as that of Xing et al. (2010) and Diavatopoulos et al. (2012).

In order to avoid possible maturity-mismatch problems caused by using IV calculated from options with differing amounts of days left until expiration on the day before earnings are announced, this research utilizes standardized implied volatilities provided by *Option Metrics*. By calculating the forward price of the underlying and then using the volatility surface to linearly interpolate to the forward price and target expiration, one arrives at an at-the-money-forward implied volatility. Therefore, every implied volatility in the sample measures investor expectations regarding the volatility of the underlying security over the following 30 days. I chose this relatively short time frame in order to minimize influences that other factors than the upcoming announcement may have on the IV. An even more exact measure would be to try to confine the research to only the announcement day or just a few days after. However, such research would have to be confined to firms with earnings announcements such that they nearly coincide with option expiries, leading to issues regarding external validity, or find a way to proxy the implied volatility for only the announcement day. At-the-money call options were chosen as per the reasons outlined in Section 2.3.

Daily historical volatilities are calculated based on a 30-day trailing window using the following formula

$$vol = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(R_i - \overline{R})} \tag{1}$$

Where R_i refers to the realized return between last days close and todays close, $\overline{R}$ is the average daily return over the last 30 trading days, and vol is the realized volatility on the day in question. In order to match the time left to maturity used to calculate the implied volatilities, I calculate the post-announcement realized volatility as the average of the daily realized volatilities over the 30 days following the announcement. As the true volatility is not observable, this paper approximates it by calculating the stock price volatility using daily closing prices. This procedure can lead to measurement error, as the calculated volatility is not necessarily equal to the true volatility. Some studies use intraday returns in order to get an estimate closer to the true volatility. Blair et al. (2010) conduct research similar to that of Canina & Figlewski (1993) and Christensen & Prahabla (1998) and compare results when using daily returns and 5-minute returns to calculate historical volatility. Their results indicate that intraday returns do not add significant incremental forecasting information, suggesting that using daily returns, the calculated historical volatility reasonably approximates the true volatility. Nonetheless, the possibility of measurement error remains.

To make sure that the time periods indeed match the time at which the information contained in the announcement first becomes integrated into the stock price, I measure IV on the day prior to the announcement if the announcement took place before market open, and on the day of the announcement in case the information was released after market close. The

HV measure contains the average historical volatility of the underlying security on the 30 days following the companies last ten earnings announcements.

Betas on the market, small-minus-big, and high-minus-low portfolios were calculated based on a 1-year timeframe. Many of the control variables used are accounting variables such as P/E ratios and profit margins. Due to the way these variables are calculated, one can end up with extremely large outliers that may be legitimate values, but are so far from all other, more ‘normal’ values that they have large distortionary effects on means and standard deviations, and thus the results of linear regression methods as used in this research. In order to retain these values, as they still contain important information being at the extremes of their respective distributions, the paper winsorizes these variables at the 1st and 99th percentile. Table 1 below shows summary statistics.

Table 1: Descriptives

Variable	**Mean**	**St. dev.**	**Min.**	**Max.**	**Observations**
IV(-1)	0.3705	0.1872	0.0122	2.2786	41541
IV(-10)	0.3573	0.1797	0.0288	2.7851	41479
IVchange	0.0129	0.0757	-1.5239	1.5239	41479
HVavg	0.3534	0.2251	0.0121	4.2745	41541
HVavgp10	0.3602	0.1718	0.0979	2.2178	40797
after close	0.3729	0.4836	0	1	41541
market beta	1.0463	0.3847	-0.7099	3.6148	41338
smb beta	0.1815	0.5313	-2.8011	4.7665	41338
hml beta	0.1588	0.8216	-6.6565	7.5138	41338
B/M	0.4859	0.3839	0.0174	2.4455	37647
P/B	4.1149	4.8837	0.3583	38.3378	37647
P/S	2.5705	3.2185	0.0987	27.0468	38640
Dividend Payout Ratio	0.3704	0.7010	0	6.9694	35323
Dividend Yield	0.0140	0.0151	0	0.0696	41541
Net Profit Margin	0.0707	0.1903	-1.7435	0.4767	38640
Debt/Assets	0.6199.	0.2184	0.1010	1.3582	38624
Debt/Equity	2.6955	7.0129	-47.3914	49.8087	38622

As can be taken from table 1, IV tends to increase slightly between the 10th day and the day prior to an earnings announcement, with an average change of 0.0129 (1.29%). IV(-10) also has a wider distribution with more extreme values than IV(-1), although the standard deviation is smaller. IV(-1) has a mean of 0.3705, which is slightly higher than the mean of the average post-announcement realized volatility, which is 0.3534. This is also slightly lower than the average post-announcement realized volatility following a firm's last ten earnings announcements, which has a mean of 0.3602. IV(-1) and IV(-10) both have positively skewed distributions, meaning that their means are larger than their medians, which are 0.3243 and 0.3123, respectively. The same holds for the variables containing historical volatilities, although the skew is less pronounced here. Furthermore, the distributions of all variables containing either implied or historical volatility have large excess kurtosis, meaning they have more observations with extreme values than assumed under a normal distribution, a phenomenon also referred to as 'fat tails'. After using a logarithmic transformation the distributions become more close to normal ones, however, some excess kurtosis still remains. A comparison of the results to hypothesis one using the nontransformed variables and their logarithmic counterparts reveals almost no discernible differences in outcomes. Consequently, I use the nontransformed values throughout the rest of this research.

4 Methodology

In order to determine whether investors correctly estimate the volatility following and earnings announcement, the paper applies the same procedure as in Canina & Figlewski (1993) and Christensen & Prahabla (1998). In this procedure, a linear regression of the form

$$RV = \alpha + \beta IV_{-1} + \epsilon \tag{2}$$

is estimated. RV refers to the post-announcement realized volatility, the average daily volatility over the 30 days following the announcement and IV refers to the implied volatility on the day before the announcement (t=-1). Given that investors perfectly anticipate the realized volatility, without any bias, and that this is completely reflected in option prices, the IV coefficient ß should be equal to 1, and the constant zero. If results significantly differ from these values, it follows that investors cannot or do not perfectly predict the volatility resulting as a response to the information released through the event. As long as the coefficient is nonzero, pre-earnings IV contains at least some information regarding the post-earnings realized volatility.

Since earnings announcements occur once every three months, and I have calculated IV on options with a time left to maturity of 30 days, the sample is nonoverlapping and thus evades the problem faced by Canina & Figlewski (1993), where IVs computed from option pairs with overlapping remaining lifetimes led to correlated forecast errors.

In order to test whether HV offers any incremental explanatory power over IV, I add HV as an explanatory variable in the regression procedure outlined in Equation (2) above. The resulting regression has the form

$$RV = \alpha + \beta_1 IV_{-1} + \beta_2 HV + \epsilon \tag{3}$$

Hypotheses 2 is therefore confirmed if the regression coefficient for HV is 0. If this is not the case and HV does indeed offer incremental explanatory power, the conclusion would be that option markets are not efficient, for investors would not be integrating all available information in their forecasts of future volatility as measured by IV. Afterwards, the paper adds several variables regarding firm characteristics as additional control variables similar to HV in Equation (3).

In order to test hypothesis 3, whether forecast accuracy is independent of the level of pre-earnings IV, quantile regressions are used. This method is similar to OLS but more useful to examine the distribution of data, as instead of estimating the mean of some distribution, one can estimate different quantiles. This makes it possible to examine whether the informational value of IV remains constant or differs across different parts of the sample. The method is also more robust to outliers than OLS, which is however unlikely to make a difference for this research due to many of the variables that do show extreme outliers having been winsorized. Quantile regressions make use of the whole sample, but assign different weights to observations, so that those observations closest to the respective quantile receive the highest weights. If the hypothesis holds, IV coefficients as well as their standard errors should remain relatively constant across the different quantiles.

As outlined earlier in this paper, prior research by Diavatopoulos et al. (2012) has indicated that changes in implied skewness and kurtosis preceding earnings announcements are more informative than levels in forecasting the subsequent stock price move. Thus, after assessing the predictive content of

IV levels, I pay further attention to changes in IV preceding the announcement. For this, the change in IV occurring over the ten days prior to the announcement is calculated, and a regression of the form

$$RV = \alpha + \beta IV_{change} + \epsilon \tag{4}$$

is run. By examining the coefficient of the change in IV, one can make an inference regarding the predictive value of the change in IV for the post-announcement realized volatility. The paper also considers whether the change in IV adds incremental information to IV's level by including it in a model with IV, again similar to HV in Equation (3).

There are a few issues with the data as taken from *Option Metrics*. Since this research utilizes IVs calculated from European at-the-money call options, the non-adjustment for dividends could be a problem as it leads to a downward bias in implied volatilities (Subramanyam et al., 2005). In order to test whether this issue affects the results of this research, I compare the full-sample results for hypothesis one as tested by Equation (2) to the results obtained using only the observations with a dividend yield of zero, which make up about one-third of the sample. This issue would result in an upwards bias in the constant, leaving the regression coefficients unaffected.

Another possible problem is measurement error, which could have causes like the exchanges on which options are traded having different closing times than the ones where their underlying is traded, leading to nonsynchronous closing prices, or the deficiencies of the Black-Scholes model leading to IV estimates that differ from the market's true volatility forecast. Furthermore, the use of closing prices to calculate IV may results in measurement error as one does not know whether the options contract closed at its current bid or ask price (Beckers, 1981). If the values of the explanatory values used throughout the research contain errors, there could also be attenuation bias, resulting in the regression coefficient estimates being biased towards 0.

Since stock prices exhibit widely differing average levels of volatility both across firms and across different time frames, it is unreasonable to assume that the model parameters are truly random. The same holds for implied volatilities and some of the control variables used, such as the betas, which tend to be persistently higher for some firms than others. To account for these issues when estimating the predictive value of IV for post-announcement realized volatility, I estimate the regression models with 2-dimensional fixed effects accounting for both the firm and the quarter that an observation belongs to. This deviation from the methodology of Canina & Figlewski (1993) and Christensen & Prahabla (1998) is explained by their use of options on the S&P 100 index. While this research employs panel data, their research simply used a time-series on one entity, meaning that they did not need to control for time invariant effects. The Hausman test confirms that the fixed-effects model specification is indeed appropriate. Similarly, for hypotheses 1,2, and 4, this paper clusters standard errors along firm and quarter dimensions in order to avoid their misspecification and make valid inferences regarding the coefficients resulting from the model. The quantile regressions used in order to examine hypothesis 3 are estimated using fixed effects, but the lack of an option to cluster standard errors by firm means that coefficients need to be interpreted with some caution. Additionally, fixed-effects models traditionally do not report a constant, as a different constant is computed for each firm in the sample. This research is however interested in the value of the constant, or more specifically, whether it is equal to zero. The constant that is reported in the models pertaining to hypotheses 1, 2, and 4 is thus an average of the individual firm constants.

Finally, after each regression I test for whether the IV coefficients and the constant are different from one and zero, respectively. Thus, the number of asterisks following coefficients in the results sections indicates whether the coefficient is different from one at significance levels of 5%, 1%, and 0.1% (*, **, ***, respectively) for all coefficient estimates regarding variables containing the IV as measured one or ten days prior to the announcement. For coefficient estimates regarding any other variable used and the constant, the number of asterisks indicates whether the coefficient is significantly different from zero instead. Additionally, this paper carries out the Wooldridge test for autocorrelation in panel data to test whether there is remaining first-order autocorrelation in the model residuals, which would point to the model results not providing an efficient estimate of average post-announcement realized volatility.

5 Results

5.1 Hypothesis 1

Table 2: Results regarding Hypothesis 1.

Variables	Model 1	Model 2	Model 3 (no dividends)	Model 4 (no dividends)
IV(-1)	0.8583***		0.8180***	
St. Error	(0.0437)		(0.0436)	
95% Confidence Interval	[0.7715, 0.9451]		[0.7313, 0.9047]	
IV(-10)		0.8496**		0.8043***
St. Error		(0.0451)		(0.0519)
95% Confidence Interval		[0.7601, 0.9390]		[0.7012, 0.9075]
Constant	0.0354*	0.0495**	0.0608**	0.0792***
St. Error	(0.0163)	(0.0162)	(0.0202)	(0.0231)
95% Confidence Interval	[0.0031, 0.0677]	[0.0174, 0.0816]	[0.0208, 0.1009]	[0.0333, 0.1251]
R^2	0.715	0.691	0.683	0.659
Observations	41532	41470	13841	13809
Wooldridge p-value	0.031	0.064	0.003	0.344

Table 2 shows the results of the various regression Models pertaining to Hypothesis one. In Model 1, the IV as measured one day before the announcements has a coefficient of 0.8583, which is close, but still significantly different from one at the 1% significance level. Thus, it appears that IV on average underestimates post-earnings volatility by about 14%. The constant is also very low at 0.0354, but is still significantly different from zero at 5%. As outlined in the methodology section, the constant here represents the average of the firm- and quarter-level fixed effects. Although I do not report the individual fixed effects here due to the sample containing more than 90 quarters and over 700 firms, there is little variation in their values across firms and quarters. Furthermore the R^2 of 0.715 is lower than what would be expected if investors accurately predicted volatility, as almost 30% of the variation in the volatility following the announcements remains unexplained by IV. The null-hypothesis of the Wooldridge test, that the error term shows no first-order autocorrelation, must be rejected at 5% significance.

Results for Model 3 are calculated using the same regression equation as Model 1, but only employing observations on firms that do not pay dividends. The IV coefficient drops to 0.818, while its standard error virtually remains the same as in Model 1. Since the constant is higher than in Model 1 at 0.0608, with significance increasing to 1%, it appears that the nontreatment of dividends when calculating IVs has not led to an upward bias in the constant in Model 1, meaning that results are not affected by this issue.

Models 2 and 4 show the same results using the IV measured ten days prior to an earnings announcement. Again, the constant increases from Model 2 to Model 4. Furthermore the IV coefficients and R^2 are slightly lower, indicating that there is still informed trading going on in the ten days prior to an announcement, as the estimates of post-announcement volatility become more accurate over this time frame.

As neither the IV coefficients nor any constant in Table 2 are equal to one and zero, respectively, investors do not on average produce an accurate prediction of post-earnings volatility as measured by the IV of at-the-money call options on the stock of the firm in question. In fact, investors seem to systematically overestimate volatility following an earnings announcement. According to the value of the IV coefficient in Model 1, on average the actual volatility that is realized over the 30 days following the announcement is only 85.83% of that predicted by the IV. Therefore, I reject the first hypothesis.

5.2 Hypothesis 2

Table 3: Regression results regarding Hypothesis 2.

Variables	**Model 5**	**Model 6**	**Model 7**
HVp10	0.5142***	0.0548	0.0788*
St.Error	(0.0375)	(0.0284)	(0.0309)
95% Confidence Interval	[0.4398, 0.5886]	[-0.0017, 0.1112]	[0.0176, 0.1401]
IV(-1)		0.8403**	
St. Error		(0.0543)	
95% Confidence Interval		[0.7325, 0.9481]	
IV(-10)			0.8208**
St. Error			(0.0583)
95% Confidence Interval			[0.7050, 0.9366]
Constant	0.1675***	0.0222	0.0312*
St. Error	(0.0135)	(0.0135)	(0.0134)
95% Confidence Interval	[0.1407, 0.1944]	[-0.0047, 0.0490]	[0.0045, 0.0579]
R^2	0.579	0.718	0.694
Observations	40789	40789	40776
Wooldridge p-value	0.000	0.022	0.022

In Table 3 I report the results pertaining to hypothesis 2, which addresses the debate of whether HV offers any incremental value over IV in predicting post-announcement realized volatility. On its own, the average HV following the last ten earnings announcements of a company contains some information regarding the volatility that will be realized following such an event with a highly significant coefficient of 0.5142. Thus, its predictive value is relatively similar to when it is measured over longer periods of time instead of for specific events, as is done in Canina & Figlewski (1993). The R^2 of 0.579 implies that the historical volatility can model more than half of the variation in post-announcement realized volatility, but this can be explained by the fact that companies with more volatile stocks on average tend to realize larger volatility following an earnings announcement than other firms, regardless of its content. Furthermore, the Wooldridge p-value signals that there still remains first-order autocorrelation in the model residuals.

Though significant by itself, the coefficient decreases to 0.0548 when the IV measured one day before the announcement is included in the model, and no longer differs significantly from zero at 5% significance. Furthermore, the constant increases to 0.0222, and is similarly no longer significant at 5%. The coefficient for IV decreases slightly from when HV is not included, to 0.8403, and remains significantly different from one at 1% significance. By including IV, the R^2 increases from 0.579 to 0.718. Unfortunately, the Wooldridge null-hypothesis of no first-order autocorrelation again has to be rejected at 5% significance. Similar findings result when switching out the IV measured one day before the announcement to that measured ten days prior. The IV coefficient estimate again decreases, whereas the constant and HV coefficient increase. Given these results, this paper does not reject hypothesis 2.

5.3 Control variables

Table 4: IV results using controls.

Control	IV Coefficient	Control Coefficient	Constant	R^2	Observation
after close	0.8586**	0.0029	0.0343*	0.715	41532
	(0.0437)	(0.0026)	(0.0164)		
mkt beta	0.8366***	0.0356***	0.0063	0.718	41328
	(0.0459)	(0.0054)	(0.0165)		
smb beta	0.8528**	0.0091	0.0359*	0.717	41328
	(0.0484)	(0.0063)	(0.0172)		
hml beta	0.8591**	0.0030	0.0347*	0.716	41328
	(0.0439)	(0.0030)	(0.0163)		
Book/Market	0.8537**	0.0267**	0.0243	0.712	37643
	(0.0457)	(0.0089)	(0.0188)		
Price/Book	0.8659**	0.0000	0.0177	0.711	37643
	(0.0475)	(0.0002)	(0.0177)		
Price/Sales	0.8534**	0.0003	0.0365*	0.709	38635
	(0.0452)	(0.0009)	(0.0166)		
Dividend Payout Ratio	0.8399**	0.0031	0.0393*	0.697	35317
	(0.0482)	(0.0016)	(0.0172)		
Dividend Yield	0.8575**	-0.1391	0.0376*	0.715	41532
	(0.0439)	(0.1472)	(0.0169)		
Net Profit Margin	0.8469**	-0.0221	0.0413	0.709	38635
	(0.0459)	(0.0137)	(0.0175)		
Debt/Assets	0.8571**	0.0132	0.0278	0.710	38619
	(0.0449)	(0.0134)	(0.0170)		
Debt/Equity	0.8575**	0.0002*	0.0352*	0.710	38618
	(0.0447)	(0.0001)	(0.0168)		

Table 5: IV results using controls cont'd.

		Model 8	$R^2 = 0.718$		
IV(-1)	**market beta**	**smb beta**	**hml beta**	**Constant**	**Observations**
0.8328**	0.0349***	0.0057	-0.0015	0.0076	41328
(0.0496)	(0.0061)	(0.0066)	(0.0035)	(0.0179)	
		Model 9	$R^2 = 0.711$		
IV(-1)	**Book/Market**	**Price/Sales**	**Price/Book**	**Constant**	**Observations**
0.8466**	0.0315***	0.0012	0.0004	0.0197	37617
(0.0461)	(0.0096)	(0.0009)	(0.0009)	(0.0194)	
		Model 10	$R^2 = 0.697$		
IV(-1)	**Dividend Payout Ratio**	**Dividend Yield**		**Constant**	**Observations**
0.8394**	0.0034*	-0.0955		0.0409*	35317
(0.0488)	0.0016)	(0.1684)		(0.0187)	
		Model 11	$R^2 = 0.709$		
IV(-1)	**Net Profit Margin**	**Debt/Assets**	**Debt/Equity**	**Constant**	**Observations**
0.8470**	-0.0215	0.0095	0.0002	0.0347*	38611
(0.0459)	(0.0138)	(0.0129)	(0.0001)	(0.0174)	

In Tables 4 and 5 I show the results of various regressions of RV on IV including individual or combinations of different control variables. The IV(-1) coefficient barely varies across models, staying between 0.83 and 0.86, but most often remains slightly below its value of 0.8583 in Model 1. Its standard error similarly stays virtually unchanged in between 0.04 and 0.05, keeping it significantly different from one at at least 1% significance. Out of all control variables reported, only the market beta, book/market ratio and debt/equity ratio return significant coefficients, although the latter is much too small in magnitude at 0.0002 to have any practical impact. I have estimated various further models containing control variables falling into similar categories as those reported here, such as Shiller's cyclically adjusted P/E ratio and profitability measures such as return on assets or equity. However, since IV coefficient estimates in these further models were, like for many of the variables reported here, similar, and control variable coefficients insignificant, they are omitted. Furthermore, I have omitted the Wooldridge p-value in tables 4 and 5, as the null hypothesis of no serial correlation of the first order remaining had to be rejected for all models at 5% significance.

The coefficients of the market beta and book/market ratio are 0.0356 and 0.0267 and significant at 0.1% and 1%, respectively. Thus, firms with a higher market beta generally have larger post-announcement realized volatility. While this seems logical given that these firms are exposed to more systematic risk, which induces higher volatility, this is a well known fact that investors should on average be able to anticipate and thus should be reflected in the IV. Nonetheless, given that market betas are usually relatively small in value, the effect is not too large. Ceteris paribus, the average volatility over the 30 days following an earnings announcement of a stock that has a market beta that is larger than that of another stock by exactly one is 3.56% higher.

It also seems that so-called value stocks with a higher book/market ratio tend to be more volatile following earnings announcements than their glamour stock counterparts. Here, if the book/market ratio increases by one, average volatility increases by 2.67%. The explanation seems relatively straightforward as value stocks tend to be more volatile, which is however another well known phenomenon amongst investors, and countless articles have been written on the difference between firms with high and low book/market ratios. It therefore surprising to find that IV, the market's forecast of volatility, does not fully account for this relationship.

While the IV coefficient remains significantly different from 1, the inclusion of controls leads to what are mostly decreases in the constant term. For the models containing the controls with statistically significant coefficient estimates, the constant is not significantly different from zero at 5% significance.

The adjusted R^2 values are similar to those obtained in Model 1 or slightly lower due to the penalty for including insignificant explanatory variables, showing that even the few significant controls are not very helpful in explaining additional variation in the realized volatility. The number of observations differs across models due to the differing availability of the control variables, but results should remain generalizable as even the model with the lowest amount of observations still includes over 35 000 earnings announcements, and data availability is unlikely to be correlated with any of the variables used.

In Table 5 I report the results of models containing different combinations of control variables that fall into the same or similar categories (i.e. beta measures, valuation ratios, dividend measures, and measures of profitability/financial health). IV(-1) coefficients again decrease slightly but stay close to 0.86, and remain different from one at 1% significance. The market beta coefficient stays virtually unchanged upon the inclusion of the small-minus-big and high-minus-low betas, whereas that on the book/market ratio increases to 0.0315 when the price/sales and price/book ratios are included. Both coefficients remain significant even at 0.1%, whereas the constant is not significant in either of the two models. Interestingly, when including IV and the dividend yield, the dividend payout ratio becomes significant at the 5% level. However, the relationship is very weak, with a coefficient of 0.0034. Since dividend payout ratios are usually very low, this effect is not really meaningful, as the average 30-day post-announcement realized volatility of a company that pays out all of its net income each year in dividends would only be 0.034% larger than that of a company that pays no dividends at all.

Estimating a further regression including both the market beta and book/market ratio as explanatory variables next to the implied volatility does not yield any different results, which I thus have omitted here. The coefficients of each variable remain stable when the other is included in the model. The IV coefficient similarly barely deviates from its initial level. All coefficients are highly significant, with the exception being the constant. These results are as expected, as the correlation between the market beta and book/market ratio is very low at 0.0863.

5.4 Hypothesis 3

Table 6: Regression results regarding Hypothesis 3

Quantile	IV(-1)	Observations
0.20	0.6291***	41541
	(0.0476)	
0.40	0.7570***	41541
	(0.0359)	
0.60	0.8870***	41541
	(0.0300)	
0.80	1.0684**	41541
	(0.0370)	

Results of the different quantile regressions regarding hypothesis 3 are shown in Table 6. Interestingly, the IV(-1) coefficient estimates start out very far from one, but increase with percentiles. This means that investors actually tend to produce more accurate forecasts of post-announcement realized volatility as measured by the IV when there is larger uncertainty around the contents of the announcement, i.e. when pre-earnings IV is higher. At the 20th percentile, IV completely overstates the volatility that will actually be realized. A coefficient estimate of 0.6291 implies that at this percentile, the average realized volatility over the 30 days following the event will only be around 63% of what is predicted by IV. The coefficient estimates for the 40th and 60th percentile are much larger at 0.757 and 0.887 respectively, but still remain relatively far from one. At the 80th percentile, investors actually tend to underestimate post-announcement realized volatility. Here, the coefficient is 1.0684, implying that the realized volatility is on average 6.84% higher than what is predicted by IV. Furthermore, IV coefficients are different from one at 1% significance or more. Between the 20th, 40th, and 60th percentiles of pre-earnings IV standard errors decrease, but increase again between the 60th and 80th percentile. Based on these results, I reject hypothesis 3.

The IVs at the 20th, 40th, 60th, and 80th percentile are 0.2316, 0.2914, 0.3625, and 0.4762 respectively. Using the coefficient that is calculated for a specific IV percentile and the corresponding value of IV, these results suggest that investors could systematically buy or sell options based on whether they belong to a percentile at which IV tends to under- or overestimate post-event realized volatility, and earn an abnormal profit.

5.5 Hypothesis 4

Table 7: Regression results regarding Hypothesis 4

Variables	**Model 12**	**Model 13**
IVchange	0.2628***	-0.3494***
St.Error	(0.0450)	(0.0424)
95% Confidence Interval	[0.1734, 0.3523]	[-0.4335, -0.2653]
IV(-1)		0.9474
St.Error		(0.0468)
95% Confidence Interval		[0.8546, 1.0402]
Constant	0.3497***	0.0068
St.Error	(0.0006)	(0.0170)
95% Confidence Interval	[0.3485, 0.3508]	[-0.0270, 0.0406]
R^2	0.534	0.726
Observations	41470	41470
Wooldridge p-value	0.000	0.081

In Table 7 I present the results to Equation (4), as well as Equation (2) while employing the change in IV between ten days and one day prior to the announcement, instead of HV. When considering the change in IV on its own, I find a positive coefficient of 0.2628, suggesting that when the expected post-event volatility by investors increases during this time frame, the average realized volatility over the 30 days following the announcement also increases. The coefficient is highly significant even at the 0.1% level, and so is the constant with a value of 0.3498. Nonetheless, the change in IV remains a poor predictor of post-event realized volatility, as it can only explain slightly more than 50% of its variation, and there is remaining first-order autocorrelation in the model residuals.

However, the picture changes drastically when IV(-1) is included in the model. The coefficient for the change in IV flips sign and now has a value of -0.3494, remaining significant at 0.1%. Furthermore, the IV coefficient increases in value from the prior models. It now lies at 0.9474 and is no longer significantly different from one at 5%. Similarly, the constant decreases to 0.0068 and is no longer significantly different from zero at 5% significance. While the change in IV is highly significant, the models explanatory power only increases minimally, with an R^2 of 0.726. However, the forecast produced now seems to be efficient, as the Wooldridge test's null hypothesis of there not being any autocorrelation of the first order remaining in the residuals cannot be rejected at 5% significance. Given these results, I do not reject hypothesis 4.

6 Discussion & Conclusion

This research paper set out to investigate the accuracy with which investors can estimate the average realized volatility of a stock over the 30 days following the firms earnings announcement, as reflected in the implied volatility of at-the-money call options expiring 30 days after the announcement, measured on the day before the announcement takes place.

The findings indicate that by itself, IV is neither an unbiased nor an efficient estimator of the post-announcement realized volatility. In fact, investors tend to overestimate this volatility as, on average, it only comes out at about 86% of the pre-earnings implied volatility. By how much the effect is overestimated furthermore seems to depend on the pre-announcement uncertainty. Thus, the lower the IV is one day before the earnings announcement, the larger on average the overestimation. This goes completely against the expectation formulated in hypothesis 3 and implies that instead of the accuracy of investor predictions of the average post-announcement volatility staying relatively constant across the distribution of pre-earnings IV, investors tend to overestimate the volatility following the announcement by more when there is relatively little uncertainty to begin with. As uncertainty increases, their estimations become more accurate, until they begin to underestimate the average post-earnings volatility around the 80th IV percentile. Furthermore, I find that, although neither unbiased nor efficient, IV has far more explanatory value regarding the average post-announcement realized volatility than the historical volatility, putting the results more in line with those of Christensen & Prahabla (1998), rather than Canina & Figlewski (1993). Using various control variables, the results furthermore show that the size of the IV coefficient is relatively robust to these different model specifications. Moreover, stocks with a larger exposure to systematic risk as measured by their beta to the market portfolio, as well as so-called value stocks (stocks of firms with large book/market ratios) tend to have larger average post-announcement volatilities, of which the effect is not fully internalized in the implied volatility, meaning that although empirical knowledge has long proven that these kinds of stocks tend to be more volatile than others, investors do not sufficiently take this into account in their volatility estimates as measured by IV. Finally, the findings give further confirmation that Diavatopoulos et al. (2012) were right in their assessment that changes in IV leading up to the announcement may be more informative than just levels. Indeed, the change in IV between the 10th and last day prior to the announcement can help in estimating the average post-announcement volatility. Although the R^2 only increases minimally, still leaving about 27.4% of the variation in post-announcement volatility unexplained, when I include the change in IV in a model with the level of IV on the day prior to the announcement, the model gives an efficient estimate as its residuals no longer contain first-order autocorrelation.

Since investors seem to on average overestimate the average post- announcement realized volatility, this implies that option prices as measured by IV on the day before earnings are too expensive, and that investors may in theory achieve abnormal returns by exploiting the on average lower than expected volatility through non directional option strategies, such as for example selling straddles. A straddle consists of one put and one call option with the same underlying stock and exercise price. Thus, purchasing a straddle allows you to profit from large stock price movements independent of direction, whereas selling straddles can allow you to profit from smaller than expected movements. As option prices are positively related to IV, and the results indicate that IV tends to overestimate post-announcement realized volatility, selling one put and one call option with the same strike price, on the day before the firm corresponding to the underlying stock reports their earnings, and that both expire 30 days after, should on average turn a profit. However, great care should be taken when trying to implement such a strategy. First of all, selling options exposes the investor to extreme potential losses. Therefore, even if a strategy works on average, investors could always get extremely unlucky and lose most or even all of their money before beginning to profit. Second, about 30% of the variation in the average post-announcement realized volatility remains unexplained by IV, meaning that there may well be models that are better at predicting this, especially considering the remaining residual autocorrelation when using only IV. Finally, the 95% confidence interval for the IV coefficient ranges from 0.7715 to 0.9451, meaning that even at 5% significance the average level of overestimation may be much lower than expected, leading to lower possible profits. It is thus not entirely clear whether employing such a strategy is worth the risk even if results are entirely accurate. To circumvent this, traders could focus on options with an IV of around 0.2316, which in the sample used in this paper marks the 20th percentile. Around these values, the average overestimation is much larger, with the average post-announcement volatility only coming in at about 63% of IV. In either case caution should be used, as the results only hold for stocks that are, or used to be, components of the S&P 500 index and have not been confirmed on an out-of-sample basis.

Even given these limitations to implementing a possible trading strategy, the results are unexpected in the sense that it seems rather illogical that investors continuously and systematically underestimate the average post-announcement realized volatility. One possible explanation may be that IV is affected by other things than only investor volatility expectations. For example, the IV derived from option prices may be too high because risk-averse investors use options to hedge against unexpected outcomes from the earnings announcement, and are thus willing to pay a premium for these options, pushing up the price and therefore the implied volatility that is derived using the Black-Scholes formula. The explanation is not completely satisfactory, however, as especially for hedging purposes when trying to protect against low probability, large magnitude movements, investors tend to prefer options that are relatively far out-of-the-money. Furthermore, it is very interesting that the effect is quite homogeneous across firms. While investors in many cases systematically over- or underestimate future returns based on firm characteristics such as size, in terms of volatility following earnings announcements the misconception, if the deviation in IV from the volatility that is actually realized really stems from false expectations, seems to be a more general one. Therefore, if there is an explanation for this phenomenon, it is more likely to be found in reasons concerning investor behaviour and preferences when it comes to option trading. Canina & Figlewski (1993) touch upon a possible reason for why findings may differ from what is expected when using IV to try to explain future volatility. When using a wrong model, "the computed IV will differ from the market's real volatility forecast" (Canina & Figlewski, 1993). Indeed, the IV used in this research may be affected by possible problems with the Black-Scholes model as outlined in section 2.2 and is furthermore taken from standardized options with interpolated prices as described in section 3, meaning that these prices have not been directly observed on the market. Trusting the correctness and accuracy of the interpolated option prices as calculated by *Option Metrics*, as well as the reasons for using the IV of ATM call options as outlined in section 2.3 and Xing et al. (2010), this paper would have to conclude that investors on average indeed systematically underestimate the volatility that is realized over the 30 days following an earnings announcement. However, I favour an explanation that takes into account the deficiencies of the Black-Scholes model. The model assumes that returns are normally distributed, which is well known not to be the case, as return distributions tend to be skewed and have large kurtosis, the same 'fat tails' that this research has observed for the volatility variables, as described in section 3. Sinclair (2013) gives a description as to why implied volatility tends to be higher than realized volatility. According to him, the reasons are twofold. First, in options trading most of the risk lies on the seller's side, who cannot always create a perfect hedge and therefore demands an insurance premium. Furthermore, the fat tails of the return distribution mean that the Black-Scholes model understates the risk of large movements which benefit the buyer of the option, as by holding an option contract they are technically 'long volatility'. Prices need to be adjusted for this, which is why Black-Scholes implied volatilities tend to be larger than realized volatilites. Other authors similarly used "the fact that along with investor's volatility forecasts an option's market price also impounds the net effect of the many factors that influence option supply and demand but are not in the option model" (Canina & Figlewski, 1993) to explain why IV is not a more accurate forecast of future volatility.

According to Canina & Figlewski (1993) IV measurements also tend to be more accurate (closer to the actual volatility forecast of investors) when arbitrage is easy. Given that this research used options on individual stocks rather than an index, it follows that IV should be more accurate, and thus yield better results, which is indeed the case. Nonetheless, there still remains some bias in the volatility forecast produced using IV. It thus seems as if sufficient arbitrage alone is not enough to bring IV down to more closely match realized volatility, which makes sense given the explanation by Sinclair (2013) above. Another explanation favouring investor rationality may be that demand & supply are even more affected by factors outside the option pricing model in periods closely preceding earnings announcements. Comparing the statement of Diavatopoulos et al. (2012) that option prices have adjusted to the earnings 5 days before they take place to the findings of Amin & Lee (1997), that option volume picks up by 10% in the 4 days prior to the announcement (compared to an increase in stock volume of only 5%), it seems likely that this increase in volume is mostly caused by uninformed traders hoping to make a quick profit, and not necessarily due to informed trading as hypothesized by Diavatopoulos et al. (2012). By driving up the demand for options, these investors may thus artificially prop up option prices and thus IVs even further, which may explain why IV on average overstates the actual post-announcement realized volatility by such a large extent. However, in that case I would expect to find more accurate forecasts using for instance the IV measured ten days before the announcement, which is also not the case for this sample. This further reiterates the notion of the bias in IV stemming from the wrong assumptions the Black-Scholes model makes about return distributions.

It is not possible to entirely reconcile the results that I found in this paper with those from Canina & Figlewski (1993) and Christensen & Prahabla (1998) due to the large differences between this research, and the research that they have conducted. While the two papers mentioned focus on OEX options, which have the S&P 100 index as their underlying security, this paper focused on options that have different individual firms stock as their underlying securities, namely the current and historical components of the S&P 500 index. Furthermore, their research was about the more general use of IV in forecasting the underlyings future volatility, while this paper examined the same relationship but focusing on its properties around earnings announcements. The obtained results very clearly suggest siding with Christensen & Prahabla (1998) on the debate of whether IV or HV is more useful in forecasting future volatility, as IV clearly outperforms HV both in terms of how close their coefficients come to one, and on how much of the variation in post-event realized volatility could be explained. IV is more favoured than HV as a forecast of future volatility not only by Christensen & Prahabla (1998), but also by other researchers (Day & Lewis, 1998; Harvey & Whaley, 1992; Sheikh, 1989) as well as contemporary academics for example due to the fact that HV is "inherently backwards-looking" (Dumas et al., 1998). However, it should be kept in mind that the HV measurement I use in this paper is relatively crude, as it simply consists of the average volatility that was realized over the 30 days following each of the last ten (or if there had not been ten prior, as many as are available) earnings announcements by the firm in question. Using other measures of HV based on longer or shorter time spans may therefore yield different results. Mayhew (1995) furthermore comments that which of IV and HV produces better volatility estimates may depend on the forecasting horizon. If that is true, it is no surprise that the results in this paper side with those of Christensen Prahabla (1998) as their research, like this one, employed a forecasting horizon of 30 days. Further research may thus be conducted on whether coefficient values and significance change when forecasting post-announcement realized volatility over a longer time span, such as the whole 3-month gap to the next earnings announcement of the firm in question.

While Beckers (1981) found that options markets were not fully efficient at the time of his research, as HV provided incremental explanatory power when included in a model forecasting future volatility with IV, this does not seem to have be the case between 1996 and 2019. Including HV in equation (3) does not yield a statistically significant coefficient, meaning that investors nowadays properly take the information that is contained in a firm's past volatility into account when estimating future volatility. Nonetheless, given that this paper finds IV to be a biased indicator that leaves large parts of the variation in future volatility unexplained, as well as the effects of larger market betas and book/market ratios on post-announcement volatility that aren't contained in IV, I cannot conclude that options markets today are fully efficient either. Although no indication for it can be found in the results of this paper, there may be more ex-ante known factors that influence the volatility after earnings and are not completely integrated into investors volatility forecasts. Sinclair (2013) states that IV tends to overstate actual future volatility especially for firms that are large, have large book/market ratios, and are very profitable. Given that large stocks also tend to have higher market betas, and that positive coefficients were found for the market beta as well as the book/market ratio when included in the model predicting post-announcement realized volatility besides IV, the results of this paper reiterate that claim. However, I do no find such evidence for profitability using both the gross and net profit margin.

If the inaccuracy of IV is indeed caused by factors outside of the model that is used to calculate IV influencing option prices, the interpretation of the results this paper finds regarding hypothesis 3 is relatively straightforward. Since such factors are based on investor behaviour and preferences, they are likely to be more or less constant, and thus not directly linked to the level of IV, such that the percentage deviation of IV from the true market volatility forecast is larger when IV is low than when it is large. Thus, the overestimation of future volatility will be more extreme when IV is relatively small in the first place, as its deviation from the true volatility forecast will also be larger when compared to the true level of IV. In the light of this explanation, it is thus not surprising to see that IV becomes more accurate as a predictor of post-announcement volatility as its value increases. The question that remains and withstands this explanation is why IV suddenly begins to underestimate the post-announcement realized volatility around the 80th percentile.

Finally, the results in this paper agree with the claim of Diavatopoulos et al. (2012) that changes in IV tend to be more informative than levels. Indeed, including the change in implied volatility between ten days prior and one day prior to the announcement in a model regressing the post-announcement realized volatility on IV(-1) increases the models explanatory power, although only by a very slight margin. The fact that the change in IV has a negative coefficient furthermore supports the idea that option prices are propped up in the days immediately before the announcement, likely by excess demand from less-informed traders. I therefore conclude that although existing literature states that investors trading in option markets tend to have superior information to those exclusively trading in stocks, there still seem to be large amounts of less informed, speculative trading going on, especially in the days preceding earnings announcements.

References

Amin, K. I., & Lee, C. M. (1997). Option trading, price discovery, and earnings news dissemination. *Contemporary Accounting Research*, *14*(2), 153–192.

Barone-Adesi, G., Brown, K., Harlow, W. V., et al. (1994). On the use of implied stock volatilities in the prediction of successful corporate takeovers. *Advances in Futures and Options Research*, *7*(1), 147–65.

Beckers, S. (1981). Standard deviations implied in option prices as predictors of future stock price variability. *Journal of Banking & Finance*, *5*(3), 363–381.

Black, F. (1975). Fact and fantasy in the use of options. *Financial Analysts Journal*, *31*(4), 36–41.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of political economy*, *81*(3), 637–654.

Blair, B. J., Poon, S.-H., & Taylor, S. J. (2010). Forecasting s&p 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. In *Handbook of quantitative finance and risk management* (pp. 1333–1344). Springer.

Canina, L., & Figlewski, S. (1993). The informational content of implied volatility. *The Review of Financial Studies*, *6*(3), 659–681.

Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of financial economics*, *50*(2), 125–150.

Day, T. E., & Lewis, C. M. (1988). The behavior of the volatility implicit in the prices of stock index options. *Journal of Financial Economics*, *22*(1), 103–122.

Derman, E., & Kani, I. (1994). The volatility smile and its implied tree. *Goldman Sachs Quantitative Strategies Research Notes*, *2*, 45–60.

Diavatopoulos, D., Doran, J. S., Fodor, A., & Peterson, D. R. (2012). The information content of implied skewness and kurtosis changes prior to earnings announcements for stock and option returns. *Journal of Banking & Finance*, *36*(3), 786–802.

Dubinsky, A., Johannes, M., et al. (2006). Fundamental uncertainty, earning announcements and equity options. In *Working paper.*

Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. *The Journal of Finance*, *53*(6), 2059–2106.

Dupire, B., et al. (1994). Pricing with a smile. *Risk*, *7*(1), 18–20.

Easley, D., O'hara, M., & Srinivas, P. S. (1998). Option volume and stock prices: Evidence on where informed traders trade. *The Journal of Finance*, *53*(2), 431–465.

Garleanu, N., Pedersen, L. H., & Poteshman, A. M. (2008). Demand-based option pricing. *The Review of Financial Studies*, *22*(10), 4259–4299.

Harvey, C. R., & Whaley, R. E. (1992). Market volatility prediction and the efficiency of the s & p 100 index option market. *Journal of Financial Economics*, *31*(1), 43–73.

Isakov, D., & Perignon, C. (2001). Evolution of market uncertainty around earnings announcements. *Journal of Banking & Finance*, *25*(9), 1769–1788.

Levy, H., & Yoder, J. A. (1993). The behavior of option implied standard deviations around merger and acquisition announcements. *Financial Review*, *28*(2), 261–272.

Mayhew, S. (1995). Implied volatility. *Financial Analysts Journal*, *51*(4), 8–20.

Patell, J. M., & Wolfson, M. A. (1979). Anticipated information releases reflected in call option prices. *Journal of accounting and economics*, *1*(2), 117–140.

Rubinstein, M. (1994). Implied binomial trees. *The journal of finance*, *49*(3), 771–818.

Sheikh, A. M. (1989). Stock splits, volatility increases, and implied volatilities. *The Journal of Finance*, *44*(5), 1361–1372.

Sinclair, E. (2013). *Volatility trading.* John Wiley & Sons.

Subramanyam, K., Marquardt, C., & Zhang, C. (2005). *Earnings surprises and uncertainty: Theory and evidence from option implied volatility* (Tech. Rep.). unpublished working paper.

Truong, C., Corrado, C., & Chen, Y. (2012). The options market response to accounting earnings announcements. *Journal of International Financial Markets, Institutions and Money*, *22*(3), 423–450.

Xing, Y., Zhang, X., & Zhao, R. (2010). What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis*, 641–662.