Energy Economics
Volume 48, March 2015, Pages 288-294
https://doi.org/10.1016/j.eneco.2014.12.021

A Comparison of Implied and Realized Volatility in the Nordic Power Forward Market

(forthcoming in Energy Economics)

Ole Henrik Birkelund, Erik Haugom, Peter Molnár, Martin Opdal and Sjur Westgaard*

February 7, 2015

Abstract

In this paper we study implied and realized volatility for the Nordic power forward market. We create an implied volatility index with a fixed time to maturity. This index is compared to a realized volatility time series calculated from high-frequency data. The results show that the implied volatility has a positive bias against the realized volatility measure indicating that there is a risk premium imposed by option traders. The results are consistent with previous research in other markets.

1 Introduction

Understanding and managing risk is crucial for all participants involved in financial transactions. In order to price assets, hedge production, or hedge financial positions, the risk characteristics need to be understood. Electricity is different from other commodities in that there is yet to exist a technology that lets us economically store electricity. Therefore, mismatches in electricity demand and generation must be covered immediately, resulting in short spikes or troughs in prices and transient periods of high volatility. The non-storability of electricity makes understanding risks more important, but also increases complexity. Financially settled forwards and options on these forwards help participants manage risks. These contracts also create the opportunity of designing models that describe and predict the market's expectation of volatility.

Volatility, as implied from option prices, is a commonly used measure of the market's expectation of future risk and it has been extensively studied, particularly for equities. Previous research shows that implied volatility (IV)-indices provide better forecasts for volatility than traditional time series methods such as GARCH (Martens and Zein, 2004). Christensen and Prabhala (1998) show this for the VIX index for S&P500 volatility and Haugom et al. (2014a)

*All the authors are associated with the Norwegian University of Science and Technology, Trondheim, Norway. Erik Haugom is associated also with Lillehamer University College, Lillehamer, Norway. Correspondence should be sent to Peter Molnár, Norwegian University of Science and Technology, Trondheim, Norway; email: peto.molnar@gmail.com, tel: +4773593578.

for WTI futures[1]. However, Agnolucci (2009) found that a Component-GARCH model performs slightly better than IV in forecasting the volatility of crude oil futures.

The main contribution of this paper is to introduce an implied volatility index for the Nordic power forward market. This is, to the best of our knowledge, the first electricity implied volatility index in the world. We also test how well it predicts realized volatility calculated from intraday data. We use methods commonly used in finance. It would be possible to use more advanced approaches, but we wanted to make the results as comparable as possible to results from other markets.

Volatility in the Nordic power forward market has been studied extensively in the academic literature. Haugom et al. (2011) was the first paper to utilize high-frequency data to analyze the Nord Pool forward market and to apply known market measures to forecast the future volatility. Haugom et al. (2010) compared forecasts of day-ahead volatility obtained from GARCH models with forecasts obtained with various auto-regressive models of realized volatility. They find that latter approach outperforms the GARCH framework.

However, these papers are concerned with realized volatility only, and do not make any use of the implied volatility. Implied volatilities from options on power forwards in the Nordic Market have not been studied yet. In this paper a unique dataset on bid and ask prices of options from market makers allows us to create an implied volatility index. The implied volatility index is then compared to realized volatility calculated from intra-daily returns. We find that implied volatility on average is greater than realized volatility, and hence that there is evidence of a volatility risk premium in the Nordic power market.

The rest of this paper is organized as follows. Section 2 describes derivatives trading in the Nordic power market, Section 3 the methodology, Section 4 the data, and Section 5 the results. Finally, Section 6 concludes.

2 Nordic Power Market

Many countries have liberalized their power markets in the past 30 years and the Nordic countries are no exception. The deregulation of the power market resulted in competitive markets and sometimes large movements in spot prices. With prices fluctuating, a healthy and increasingly liquid derivatives market sprung from the need to control risk. Consequently, academic research studying the pricing of electricity derivatives also emerged, see Vehvilainen (2002), Benth et al. (2007), Weron (2008) and Kiesel et al. (2009).

Norway and Sweden established the Nord Pool electricity and power market exchange in 1996, as the world's first multinational exchange for trading and clearing financial power contracts. Clearing of standardized financially settled contracts was introduced in 1997, and standardized options on forward and futures contracts were introduced in 1999. Nord Pool Clearing was in 2008 acquired by Nasdaq OMX, and the exchange changed name to Nasdaq OMX

[1]VIX is a trademark ticker symbol for the Chicago Board Options Exchange (CBOE) Market Volatility Index. WTI is short for 'West Texas Intermediate' and is a light sweet crude oil product that is the underlying commodity of the New York Mercantile Exchange's oil futures contracts.

Commodities Europe in 2010. However, Nord Pool remains in existence today as an independent exchange for spot electricity.

The Nasdaq OMX Commodities Europe exchange is open for trading on power derivatives between 08:00 and 15:30 (CET) and both the underlying forward and the option contracts are cleared within this time span. Options on forwards are mostly traded over-the-counter (OTC) at various brokerage firms and trades are cleared the same day as long as they are submitted before the deadline of 15:30. Trading after the deadline is cleared the next day. Closing prices are fixed at a random time between 15:25 and 15:30 and the contracts are settled financially. Forwards are available for daily, weekly, monthly, quarterly and yearly contracts.

A forward at the exchange is an obligation to buy or sell a predetermined amount of power at a given price with delivery each hour for the time covered by the forward. The minimum size of the contract is 1 MW and the minimum ticker is 0.01 EUR. The contract is settled financially.

An option on a forward is the right to buy/sell a forward contract for a given price K at time T in the future. The maturity of the option is 10 working days before the maturity of the underlying forward and the payoff is a function of the forward price only (Vehvilainen, 2002). Forwards and options on those forwards are standardized agreements, making comparisons possible without introducing unnecessary variables.

3 Methodology

3.1 Implied Volatility

Soon after the option pricing model of Black and Scholes (1973) and Merton (1973)[2] it was observed that the function could be reversed to calculate implied volatility (Latane and Rendleman, 1976). All of the input variables in the BSM model, except for the volatility, are observable in the market. This makes it possible to calculate the volatility based on the current option price, current forward price, F_0, the strike price K, the time to maturity T and the risk free interest rate, r_f. Such a method was used to create several implied volatility indices, most notably the VIX from the Chicago Board Options Exchange (CBOE) in 1993.

In 2003 the CBOE decided to change the method for calculating the VIX. The previous model, the BSM-IV was replaced by the framework developed by Britten-Jones and Neuberger (2000), the model-free implied volatility, for the VIX on the S&P500 stock index (CBOE, 2003).[3] This method makes the implied volatility independent of any option price model and calculates implied volatility from the full set of available strikes for European puts and calls (Andersen and Bondarenko, 2007). Jiang and Tian (2005) generalized the method to include jumps and showed that model-free implied volatility subsumes all information contained in BSM implied volatilities and that it gives a more efficient forecast of future realized volatilities.

[2] Henceforth called the BSM model

[3] The old BSM implied volatility index method is still in use, but with the ticker VXO for the S&P100 index

Our main objective is to use a measure that is readily understood by practitioners which relies on standard models in the financial literature. Volatility traders usually quote option prices not in dollars, but in implied volatilities. These implied volatilities are calculated from the BSM. In the Nordic financial markets brokerages also report implied volatilities based on the BSM. Even though the BSM model has its limitations, it is still the most commonly used approach for calculating implied volatility in practice. This approach is also frequently applied in empirical studies examining implied volatility for other markets (see Dufour et al. (2012) for a recent example). As financial electricity prices are found to exhibit similar properties as more traditional financial assets (Haugom, 2011), we believe the BSM is an appropriate approach for the construction of a meaningful volatility index also for the Nordic power forward market.

3.1.1 Creating an IV-Index

IVs calculated from options with different moneyness,[4] but with the same maturity will give different implied volatilities. Without further discussion of the sample period, Figure 1 shows this for different dates in our data. Two factors should be highlighted; the concavity and the variability over time.

First, we observe that the IV is convex across moneyness. This convexity is referred to as the smile effect (Taylor, 2005; Alexander, 2008b). The reasons for this effect varies, but some suggest transaction costs or traders including a risk premium for out/in-of-the-money options as possible explanations (Peña et al., 1999; Taylor, 2005). At-the-money options are most suitable for creation of a volatility index, particularly in illiquid markets.

Many methods exist to find or use ATM option prices when an exact ATM option is not available and Taylor (2005) suggests that the most liquid option nearest ATM is a natural choice. An alternative is to use weighting as suggested by Siriopoulos and Fassas (2009) and Ederington and Guan (2002). Our IV index is therefore computed from bid prices[5] of the two nearest out-of-the-money put options and the two nearest out-of-the-money call options. Ederington and Guan (2002) showed and our preliminary results (not included in the paper) confirm that it is not necessary to use more options in the weighting. Hence, equation 1 describes our model for calculating the implied volatility for a given date.

$$IV = \frac{1}{M}(m_a IV_a + m_b IV_b + m_c IV_c + m_d IV_d), \qquad (1)$$

where m_i is the moneyness in case of call options and inverse of moneyness in case of put options, with volatility IV_i and M is the sum of these weights from all four options. This weighting scheme gives larger weights to near at-the money options. However, all the weights are close to one and this weighting scheme does not differ much from simple arithmetic average. All IV's are calculated on the same day with the same time to maturity, H.

[4] By moneyness we mean the ratio $= F_0/K$

[5] Using bid prices might introduce a downward bias in the volatility estimates. However, as reader will later see, we find that implied volatility is on average higher than realized volatility. Therefore, this possible bias makes our results even stronger.

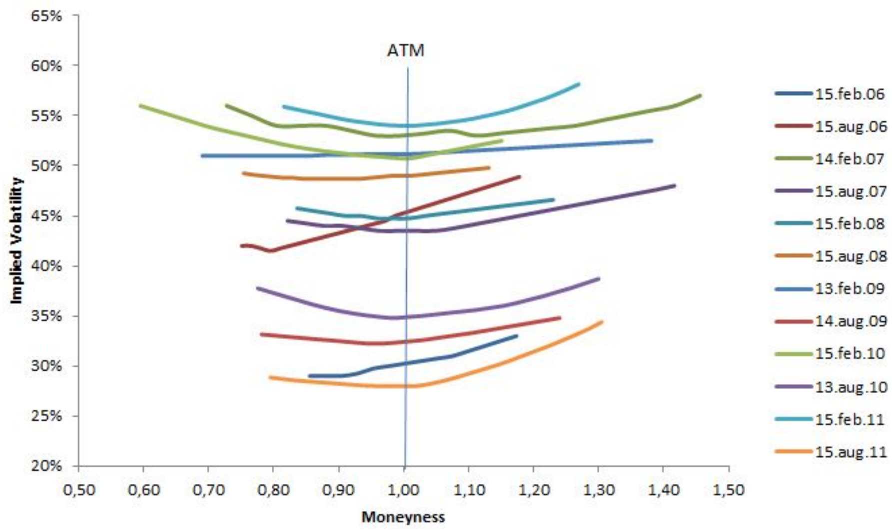

Figure 1: Illustration of volatility smile and its development throughout our sample

The second factor to notice in Figure 1 is how the smile varies with time. This is because the volatility changes over time. Our index should capture these changes, but to create a coherent measure of volatility time to maturity must be constant. We follow Martens and Zein (2004) and adjust IV to the desired time horizon by using linear interpolation between options of different maturities. For two options with maturities at T_1 and T_2 on day t, and with $T_1 < H < T_2$ we create an IV measure with time horizon H from equation 2.

$$IV_{t,H} = IV_{t,T_1} + \frac{H - T_1}{T_2 - T_1}(IV_{t,T_2} - IV_{t,T_1}) \tag{2}$$

This method for creating an IV-index is analogous to the previously mentioned method for VIX on the S&P Index (Martens and Zein, 2004).

3.2 Realized volatility

The theory of quadratic variation suggests that the following holds if the discrete sampled returns exhibit no serial correlation and the sample path for the volatility, σ_t, is continuous (Karatzas and Shreve, 1991):

$$\operatorname*{plim}_{N\to\infty} \left(\int_0^1 \sigma^2_{t+\tau} d\tau - \sum_{j=1}^{N} r^2_{t,j} \right) \to 0, \tag{3}$$

where N is the sampling frequency, $\sigma^2_{t+\tau}$ is the integrated (unobservable) variance, and $r_{t,j}$ is the intraday return for time j of a given day, t. Volatility is unobservable in the market, but the last term in equation 3, which is known as the realized variance, measures this theoretical integrated variance almost perfectly when the sampling frequency, N, is sufficiently high. Formally, denote total number of days as T. The realized variance for day t is given by:

$$RVar_t = \sum_{j=1}^{N} r_{t,j}^2, \quad t = 1, \ldots, T. \tag{4}$$

Anderesen and Bollerslev (1998) showed that a well behaving estimator of volatility is realized volatility as measured by the square root of realized variance. However, in order to obtain an efficient estimate of the true volatility over the whole day, we include the overnight return as described in Hansen and Lunde (2005). This method is found to be optimal from a theoretical perspective and performs well in empirical applications (see Hansen and Lunde (2005) and Haugom et al. (2014b)). The procedure is also simple to implement.

Using this methodology we calculate the realized volatility as a weighted sum of the intradaily and overnight return by finding weights that minimize the squared error between the realized volatility and the true volatility. Let or_t denote the overnight return between the last price on day $t-1$ and first price on day t, and let $RVar_t$ denote the intradaily squared returns on day t. We then define the following measures:

$$\hat{\mu}_0 = \frac{1}{n} \sum_{t=1}^{n} (or_t^2 + RVar_t) \tag{5}$$

$$\hat{\mu}_1 = \frac{1}{n} \sum_{t=1}^{n} or_t^2 \tag{6}$$

$$\hat{\mu}_2 = \frac{1}{n} \sum_{t=1}^{n} RVar_t \tag{7}$$

$$\hat{\eta}_1^2 = Var(or_t^2) \tag{8}$$

$$\hat{\eta}_2^2 = Var(RVar_t) \tag{9}$$

$$\hat{\eta}_{12}^2 = Cov(or_t^2, RVar_t) \tag{10}$$

where n is the total number of observations. The relative importance factor is calculated in the following way:

$$\hat{\varphi} = \frac{\hat{\mu}_2^2 \hat{\eta}_1^2 - \hat{\mu}_1 \hat{\mu}_2 \hat{\eta}_{12}}{\hat{\mu}_2^2 \hat{\eta}_1^2 + \hat{\mu}_1^2 \hat{\eta}_2^2 - \hat{\mu}_1 \hat{\mu}_2 \hat{\eta}_{12}} \tag{11}$$

The optimal weights are:

$$\hat{\omega}_1^* = (1 - \hat{\varphi}) \frac{\hat{\mu}_0}{\hat{\mu}_1} \quad \text{ } and \quad \hat{\omega}_2^* = \hat{\varphi} \frac{\hat{\mu}_0}{\hat{\mu}_2} \tag{12}$$

In our model, 24-hour realized volatility for day t is then calculated by the following formula:

$$RV_t = \sqrt{\hat{\omega}_1^* or_t^2 + \hat{\omega}_2^* RVar_t} \tag{13}$$

as suggested by Hansen and Lunde (2005). In order to find the realized volatility that matches the time horizon of the IV-index we average the annualized realized volatilities over the desired time horizon, H.

$$RV_{t,H} = \frac{1}{H} \sum_{i=t}^{t+H} RV_i. \tag{14}$$

Choosing the right sampling method and frequency is important for the validity of our estimates. We sample by extracting ticker prices prior to every minute as suggested by Wasserfallen and Zimmermann (1985) and discussed by Hansen and Lunde (2006). From equation 3 we see that to fully capture the information content in high frequency data, the sampling frequency should be as high as possible. However, higher sampling frequencies give biased measures of RV due to microstructure effects such as bid-ask bounce (Alexander, 2008a; Taylor, 2005). Andersen et al. (2001) suggests using 5 minute intervals, a practice that has been followed by many researchers in other markets (Haugom et al., 2014a; Patton, 2011; Martens, 2002). To resolve the tradeoff between statistically high information content and microstructure problems we use a volatility signature plot as suggested by Andersen et al. (2000) and used by Bollerslev et al. (2008) and Haugom et al. (2014b). It shows the average realized volatility such that

$$\overline{RV}^{(N)}_{t_0,T} = \frac{1}{T} \sum_{t=t_0}^{t_0+T} RV_t, \tag{15}$$

where N is the number of samples per day from equation 4, and T is the number of observations. The plot is obtained by varying N. The highest number of N where the plot is flat, is the point where the RV measure is approximately free of microstructure bias (Andersen et al., 2000).

3.3 Regression

Bias of implied volatility in comparison to subsequent realized volatility is usually studied using the following regression (see Ederington and Guan (2002) and Martens and Zein (2004)):

$$RV_{t,H} = \alpha + \beta IV_{t,H} + \varepsilon_t. \tag{16}$$

These tests usually find that $\alpha > 0$ and $\beta < 1$. When we run this regression, we find the same result, but α is not statistically significant even on 5% level, whereas β is significant even at 0.01% level. Result that $\alpha > 0$ and $\beta < 1$ is usually interpreted as finding that IV is upward biased. However, such result does not really imply that IV is upward biased, it actually implies that IV is upward biased for large values, but downward biased for low values. We therefore run the regression without constant:

$$RV_{t,H} = \beta IV_{t,H} + \varepsilon_t \tag{17}$$

and test for possible bias in IV by testing the hypothesis $H_0 : \beta \geq 1$ against the alternative hypothesis $H_1 : \beta < 1$.

Since the RV observations in this regression are overlapping by construction (two consecutive observations share $H-1$ of daily observations), the assumption of no autocorrelation in the error term is violated. One possible solution would be to use non-overlapping observations only. However, this would disregard a lot of available information and make the test less powerful. Instead we use all the observations and report Newey-West standard errors (with 65 lags).

4 Data

Data for implied volatility is collected from ICAP Energy, a commodity brokerage firm that provides OTC brokering and advisory services, and is presented to us in a refined form by Fred Espen Benth of the Center of Mathematics for Applications, University of Oslo. The raw data contains daily prices of options on forwards with different strikes for quarterly and yearly contracts. Previous studies find that longer time horizons provide the best predictive powers for IV against RV (Taylor, 2005), we therefore use quarterly contracts.[6]

The liquidity of options on forwards is low and it is possible to argue that there are too few trades in the market to successfully create an IV-index. However, ICAP provides the bid and ask prices from the market makers and market participants trust these prices not to be affected by orders below 10 MW. Implied volatilities are calculated by ICAP from the BSM model and we use the implied volatilities from the closing bid prices in our analysis. We create an index with a constant time horizon from the front forward, called 1pos, and the forward with between one and two quarters to maturity, the 2pos contract. These contracts have time to maturity T_1 and T_2 respectively, and the choice of H is then naturally the average number of working days in one quarter, 66.

The raw data containing continuously recorded ticker prices of forwards for trades performed in the opening hours of Nasdaq OMX was obtained from Montel. Figure 2 shows the development of liquidity in the market as measured by ticks per trading day. We observe that, with a few exceptions, liquidity is stable over time. Days where trades are recorded outside of the opening hours are few and the return is treated as overnight return.[7] Our estimate of realized volatility is based on prices from the nearest quarter, the 1pos contract. This deviates slightly from the practice of Martens and Zein (2004) as they use 2pos contract prices when the liquidity of the 2pos contract surpasses that of the 1pos contract. This usually happens as the contract is close to, but before, maturity. However, since Nordic electricity forwards are settled financially, there is no drop in liquidity before maturity and therefore no need to roll over contracts before this point in time.

On average there are 204 trades per day, a trade every 2 min and 11 seconds. With this level of liquidity the impact of microstructure noise is low. This can be observed in the volatility signature plot in Figure 4 on page 10. The expected parabolic decrease in volatility for longer tick intervals is not observed and the plot fails to give a clear indication of the best sampling frequency. Lien et al. (2012) chose a 30 minute sampling interval when studying the electricity forward market. We see from the volatility signature plot that the volatility in our data

[6]We consider the yearly forward contracts to have too low liquidity to be efficiently studied using realized variance.

[7]Special circumstances and late clearing explains these examples. The dates include: October 12, 2005, March 24, 2006, April 25, 2006, May 2, 2006, May 29, 2006, June 23, 2006, July 27, 2006, August 23, 2006, October 30, 2006, December 4, 2006, January 2, 2007, June 29, 2007, August 28, 2007, November 22, 2007, February 28, 2008, March 6, 2008, June 10, 2008, August 28, 2008, December 3, 2008, December 5, 2008, December 10, 2008, March 6, 2009, May 12, 2009, May 14, 2009, July 1, 2009, and December 28, 2009.

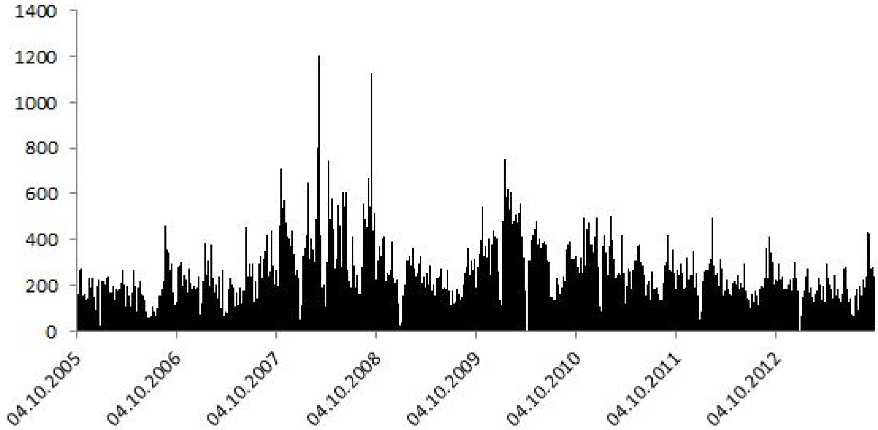

Figure 2: Development of ticks per day throughout the sample

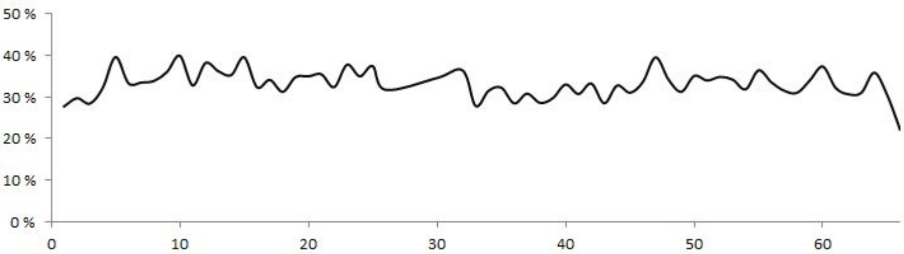

Figure 3: Average realized volatility as a function of time to maturity for 1pos contracts in the Nordic electricity forward market

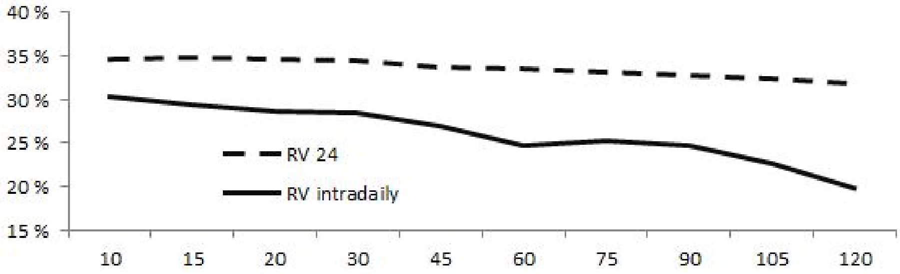

Figure 4: Volatility signature plot

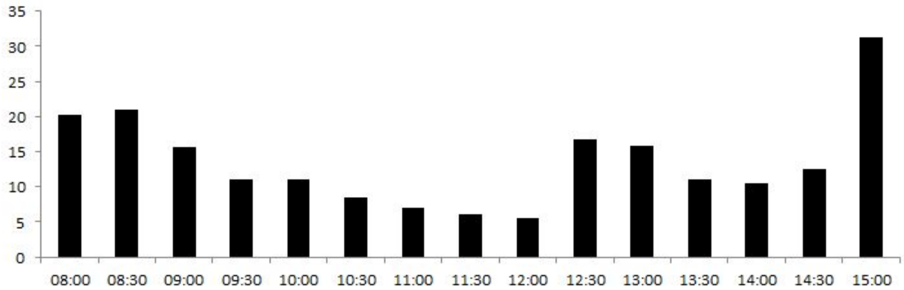

Figure 5: Average number of trades per half hour

is relatively stable at this frequency. We therefore use a sampling frequency of 30 minutes in this study.

The choice of this interval in realized volatility calculations has ramifications for what days that should be removed from our sample. With 30 minute intervals, we must sample 15 ticks per day. Figure 5 shows that the trades are not evenly distributed during the day. Hence, to avoid imposing a negative bias in our measurements, we remove days with fewer than 20 trades. Following these adjustments, our sample starts October 10, 2005 and ends September 14, 2011, in total 1368 daily observations. 14 outliers or days with low liquidity are removed[8] which leaves us with 1354 days for the empirical study.

5 Results

In Figure 6 we present the 24 hour realized volatility measure, RV_t. It is clear from the figure that volatility varies significantly. Days with volatility of up to 200% (annualized) illustrate significant uncertainty in the market. Such extreme events are to be expected in power markets where disruption at one large facility will induce large uncertainties about future prices.

Table 1 shows the descriptive statistics for IV and RV with a time horizon of one quarter. On average, IV is higher than RV with IV having a sample mean of 42% as opposed to 36% for RV. This suggests that there is a risk premium

[8]The dates include: December 23, 2005, December 28-30, 2005, June 23, 2006, March 26, 2007, January 2, 2008, June 19, 2008 March 19, 2009, June 18-19, 2009, July 23, 2009, September 17, 2009 and June 17, 2009

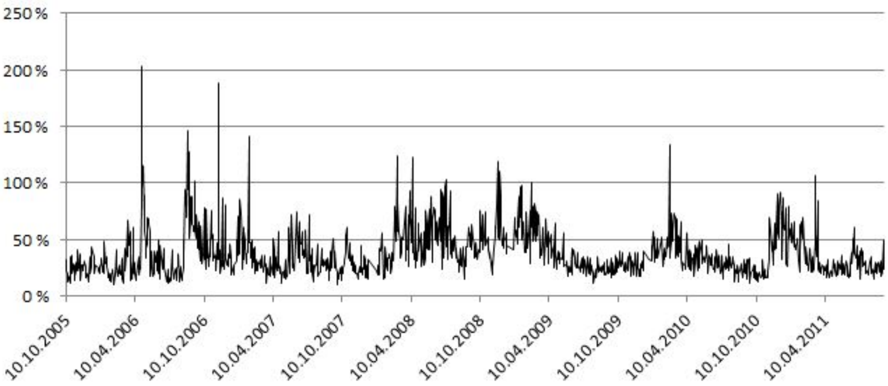

Figure 6: Annualized 24 hour daily realized volatility with 30 minute tick intervals

Table 1: Summary statistics. Volatilities are annualized.

	IV	RV
Mean	42%	36%
Median	41%	33%
Maximum	66%	59%
Minimum	28%	21%
Std. Dev.	7.6%	10%
Skewness	0.43	0.52
Kurtosis	2.60	1.97
Observations	1354	1354

in the option contract prices and the result is similar to finding of Ederington and Guan (2002) for S&P 500 futures.

Figure 7 on page 12 reports the IV-index compared to the realized volatility over the full sample period. Note that the realized volatility is much smoother than the IV-index. This is because realized volatility is measured as the average volatility over the next 66 working days whereas the IV-index shows data from individual days. From the same figure we also observe that the IV measure lag behind the RV measure. This is also caused by the fact that the realized volatility is measured as the average volatility over the next 66 working days whereas the IV-index shows data from individual days. Therefore, RV contains future information (it is average over next 66 days) whereas IV contains only

Table 2: Regression results. Standard errors are Newey-West standard errors with 65 lags. P-value is p-value for the test $H_0 : \beta \geq 1$ against $H_1 : \beta < 1$.

	Coefficient	Standard error	p-value
β	0.866	0.036	1.8%
R^2	0.21		
Adjusted R^2	0.21		

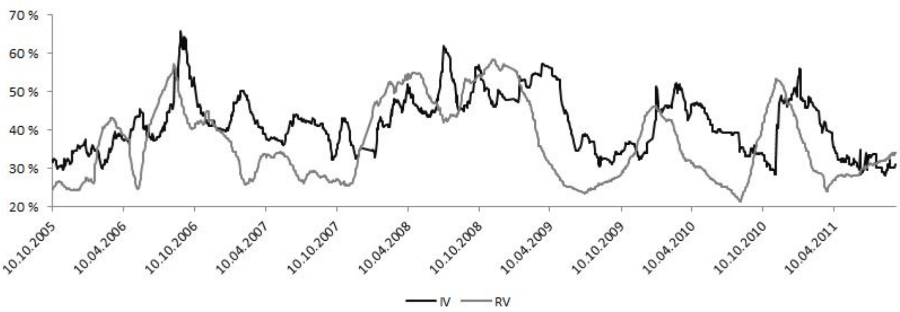

Figure 7: Development of implied and realized volatility over the sample.

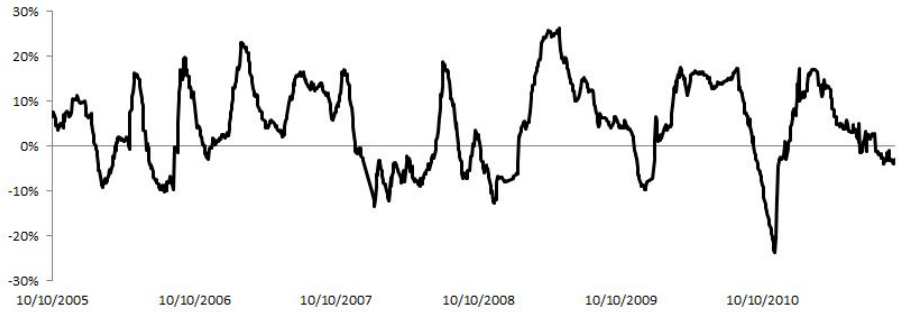

Figure 8: The volatility risk premium defined as the difference between implied and realized volatility over the sample period.

present information. As a result, RV leads IV. However, this is not a mistake, our goal is to compare present IV with the subsequent (future) RV.

Figure 8 shows the evolution of the volatility risk premium (defined as the difference between the option implied volatility and the realized volatility, see e.g. Eraker (2009), Dufour et al. (2012), Ribeiro et al. (2012)) over time. Note that this is ex-post volatility risk premium and therefore it is sometimes negative. However, the difference is on average positive.

To formally test the relation between IV and RV we perform the regression analysis described in Section 3.3. We find that $\hat{\beta}$ is 0.87 and we can reject the hypothesis that $\beta \geq 1$ in favor of the alternative that $\beta < 1$. We therefore conclude that there is a positive volatility risk premium in the Nordic electricity market. This result is in line with findings for other markets, see Ederington and Guan (2002).

6 Conclusion

This is the first paper to calculate an implied volatility index for the Nordic power forwards market. The creation of the index was made possible due to a unique dataset on option prices provided by ICAP. We compare the newly developed index with the observed realized volatility as calculated from intra-daily data.

Our results suggest that there exists a positive volatility risk premium in

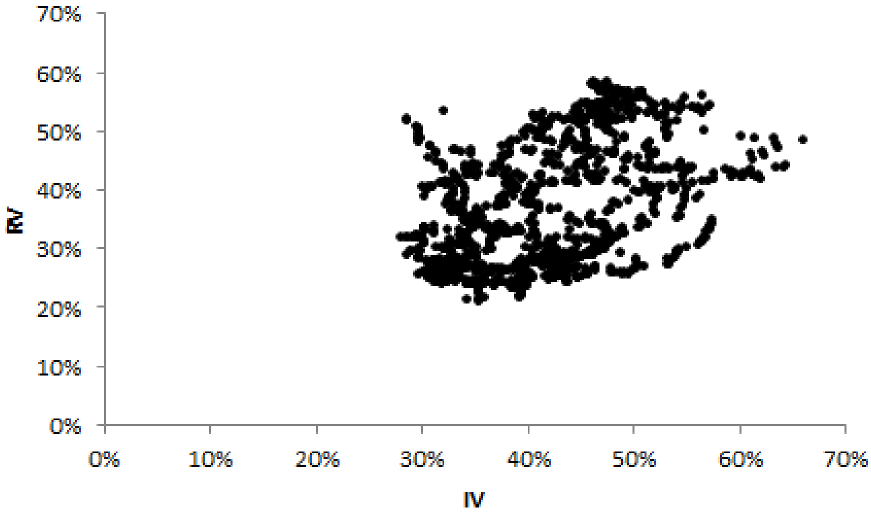

Figure 9: Scatter plot of RV and IV.

options prices for power forwards contract at Nasdaq OMX commodities. On average, the implied volatility index is 42% and the realized volatility is 36% over the whole sample period. This is also supported when we formally test the relation between these two volatility measures. The volatility index is a biased predictor of the observed realized volatility. These findings for the Nordic power market are similar to findings for other, more traditional, financial markets.

Acknowledgements

We would like to thank to Fred Espen Benth of the University of Oslo and Knut Strømnes of Montel for providing the data.

References

Agnolucci, P. (2009). Volatility in crude oil futures: A comparison of the predictive ability of GARCH and implied volatility models. *Energy Economics*, 31(2):316 – 321.

Alexander, C. (2008a). *Market Risk Analysis, Volume 2: Practical Financial Econometrics*. Wiley.

Alexander, C. (2008b). *Market Risk Analysis, Volume 3: Pricing, Hedging and Trading Financial Instruments*. Wiley.

Anderesen, T. G. and Bollerslev, T. (1998). Aneswering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4):885–905.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1):43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2000). Great realisations. *Risk*, pages 105 – 108.

Andersen, T. G. and Bondarenko, O. (2007). Construction and interpretation of model-free implied volatility. Technical report, National Bureau of Economic Research.

Benth, F. E., Kallsen, J., and Meyer-Brandis, T. (2007). A non-gaussian ornstein–uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, 14(2):153–169.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3):637 – 654.

Bollerslev, T., Law, T. H., and Tauchen, G. (2008). Risk, jumps, and diversification. *Journal of Econometrics*, 144(1):234–256.

Britten-Jones, M. and Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, 55(2):839–866.

CBOE (2003). Vix white paper. *URL: http://www. cboe. com/micro/vix/vixwhite. pdf.*

Christensen, B. J. and Prabhala, N. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50:125 – 150.

Dufour, J.-M., Garcia, R., and Taamouti, A. (2012). Measuring high-frequency causality between returns, realized volatility, and implied volatility. *Journal of Financial Econometrics*, 10(1):124–163.

Ederington, L. H. and Guan, W. (2002). Measuring implied volatility: Is an average better? which average? *Journal of Futures Markets*, 22(9):811–837.

Eraker, B. (2009). The volatility premium. Technical report, Working paper.

Hansen, P. R. and Lunde, A. (2005). A realized variance for the whole day based on intermittent high-frequency data. *Journal of Financial Econometrics*, 3:525–554.

Hansen, P. R. and Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business and Economic Statistics*, 24(2):127–161.

Haugom, E. (2011). Some stylized facts about high-frequency nord pool forward electricity prices. *The Journal of Energy Markets*, 4(1):21 – 49.

Haugom, E., Langeland, H., Molnár, P., and Westgaard, S. (2014a). Forecasting volatility of the u.s. oil market. *Journal of Banking & Finance*, 47:1 – 14.

Haugom, E., Veka, S., Lien, G., and Westgaard, S. (2014b). Estimating and evaluating value-at-risk forecasts based on realized variance: Empirical evidence from ice brent crude oil futures. *Forthcoming in the OPEC Energy Review.*

Haugom, E., Westgaard, S., Solibakke, B. B., and Lien, G. (2011). Realized volatility and the influence of market measures on predictability: Analysis of nord pool forward electricity data. *Energy Economics*, 33:1206 – 1215.

Haugom, E., Westgaard, S., Solibakke, P., and Lien, G. (2010). Modelling day ahead nord pool forward price volatility: Realized volatility versus GARCH models. *The European Energy Market (EEM), 2010, 7th International Conference on the European Energy Markets*, pages 1–9.

Jiang, G. J. and Tian, Y. S. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4):1305–1342.

Karatzas, I. A. and Shreve, S. E. (1991). *Brownian motion and stochastic calculus*, volume 113. Springer.

Kiesel, R., Schindlmayr, G., and Borger, R. (2009). A two-factor model for the electricity forward market. *Quantitative Finance*, 9(3):279 – 278.

Latane, H. A. and Rendleman, R. J. (1976). Standard deviations of stock price ratios implied in option prices. *The Journal of Finance*, 31(2):369–381.

Lien, G., Haugom, E., Westgaard, S., and Solibakke, P. B. (2012). Covariance estimation using high-frequency data: An analysis of nord pool electricity forward data. *Journal of Energy and Power Engineering*, 6(4).

Martens, M. (2002). Measuring and forecasting s&p 500 index-futures volatility using high-frequency data. *The Journal of Futures Markets*, 22(6):491 – 518.

Martens, M. and Zein, J. (2004). Predicting financial volatility: High frequency time-series forecasts vis-à-vis implied volatility. *The Journal of Futures Markets*, 24(11):1005 – 1028.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):pp. 141–183.

Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246 – 256.

Peña, I., Rubio, G., and Serna, G. (1999). Why do we smile? on the determinants of the implied volatility function. *Journal of Banking and Finance*, 23(8):1151 – 1179.

Ribeiro, R., Kolanovic, M., Silvestrini, D., and Lee, T. (2012). Risk premia in volatility markets: Exploiting volatility spillover and clustering. Technical report, J.P. Morgan report available at https://mm.jpmorgan.com/EmailPubServlet?h=-2s2dbt6&doc=GPS-999930-0.pdf.

Siriopoulos, C. and Fassas, A. (2009). *Implied Volatility Indices - A Review.* http://ssrn.com/abstract=1421202.

Taylor, S. J. (2005). *Asset Price Dynamics, Volatility, and Prediction.* Princeton University Press.

Vehvilainen, I. (2002). Basics of electricity derivative pricing in competitive markets. *Applied Mathematical Finance*, 9(1):45–60.

Wasserfallen, W. and Zimmermann, H. (1985). The behavior of intra-daily exchange rates. *Journal of Banking and Finance*, 9(1):55 – 72.

Weron, R. (2008). Market price of risk implied by asian-style electricity options and futures. *Energy Economics*, 30:1098 – 1115.