

SFB 649 Discussion Paper 2017-025

Realized volatility of CO₂ futures

Thijs Benschop *
Brenda López Cabrera *



* Humboldt-Universität zu Berlin, Germany

This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".

<http://sfb649.wiwi.hu-berlin.de>
ISSN 1860-5664

SFB 649, Humboldt-Universität zu Berlin
Spandauer Straße 1, D-10178 Berlin



SFB 649 ECONOMIC RISK BERLIN

Realized volatility of CO₂ futures*

Thijs Benschop[†], Brenda López Cabrera[‡]

November 2, 2017

Abstract

The EU Emission Trading System (EU ETS) was created to reduce the CO₂ and other greenhouse gas emissions at the lowest economic cost. In reality market participants are faced with considerable uncertainty due to price changes and require price and volatility estimates and forecasts for appropriate risk management, asset allocation and volatility trading. Although the simplest approach to estimate volatility is to use the historical standard deviation, realized volatility is a more accurate measure for volatility, since it is based on intraday data. Besides the stylized facts commonly observed in financial time series, we observe long-memory properties in the realized volatility series, which motivates the use of Heterogeneous Autoregressive (HAR) class models. Therefore, we propose to model and forecast the realized volatility of the EU ETS futures with HAR class models. The HAR models outperform benchmark models such as the standard long-memory ARFIMA model in terms of model fit, in-sample and out-of-sample forecasting. The analysis is based on intraday data (May 2007-April 2012) for futures on CO₂ certificates for the second EU-ETS trading period (expiry December 2008-2012). The estimation results of the models allow to explain the volatility drivers in the market and volatility structure, according to the Heterogeneous Market Hypothesis as well as the observed asymmetries. We see that both speculators with short investment horizons as well as traders taking long-term hedging positions are active in the market. In a simulation study we test the suitability of the HAR model for option pricing and conclude that the HAR model is capable of mimicking the long-term volatility structure in the futures market and can be used for short-term and long-term option pricing.

Keywords: EU ETS, Realized Volatility, HAR, Volatility Forecasting, Intraday Data, CO₂ Emission Allowances, Emissions Markets, Asymmetry, SHAR, HARQ, MC Simulation

*The financial support from the Deutsche Forschungsgemeinschaft via SFB 649 “Ökonomisches Risiko” and IRTG 1792 “High Dimensional Non Stationary Time Series”, Humboldt-Universität zu Berlin is gratefully acknowledged.

[†]Ladislav von Bortkiewicz Chair of Statistics, Humboldt-Universität zu Berlin, Spandauer Straße 1, 10178 Berlin, Germany. Email: thijs.benschop@hu-berlin.de

[‡]Ladislav von Bortkiewicz Chair of Statistics, Humboldt-Universität zu Berlin, Spandauer Straße 1, 10178 Berlin, Germany. Email: lopezcab@wiwi.hu-berlin.de

1 Introduction

The price dynamics and determinants of European Union Allowances (EUAs) in the EU Emission Trading System (EU ETS) are of great importance for participating industries, and for sound risk management and hedging strategies of financial intermediaries as well as for policy makers who use them to evaluate the performance of the EU ETS. Moreover, the market for EUAs is constantly growing, which makes it important for market participants to have a valid pricing model. Volatility forecasting is now becoming increasingly important as CO₂ prices are very volatile. Volatility has a direct impact on futures prices and is important for pricing options. This need for volatility models is testified by investments in quantitative market analysis by major market participants.

The returns of CO₂ futures data exhibit the same stylized facts as other commodity returns, which are amongst others volatility clustering and fat tails. However, besides these stylized facts there are indications of long memory and an asymmetric response of volatility to positive and negative returns in the data, which are not captured by standard GARCH models. Therefore, we investigate the volatility process with another class of models using realized volatility (RV). RV is a measure for the latent volatility, which relies on the squared returns in high-frequency intra-day data and is an accurate measure of the volatility. Corsi (2009) introduced the Heterogeneous Auto-Regressive (HAR) model for RV, which models the long-memory behavior in the data. However, the HAR model cannot capture the observed asymmetry and several extensions have been proposed to model this asymmetry. Our work focusses on this issue by using the Semi-variance HAR (SHAR) model (see Patton and Sheppard, 2015) and the HAR Quarticity (HARQ) model (see Bollerslev et al., 2016). The HAR class models are easier to estimate and have superior modeling performance and a more compelling interpretation than standard GARCH, ARFIMA and stochastic volatility models. For the volatility in the European CO₂ markets the HAR class models yield both better model fits and forecasting results. They also allow to shed new light on the volatility structure and its drivers.

Since the introduction of the EU ETS in 2005 a new field of applied econometrics, which investigates the behavior of prices and volatility of EUAs and their derivatives, has emerged. A number of studies have focused on the price determinants of EUA daily spot prices (e.g., Mansanet-Bataller et al., 2007; Alberola et al., 2007, 2008a,b; Chevallier, 2009; Hintermann, 2010; Hammoudeh et al., 2014a,b) and found long term relationships between EUA prices and energy prices, temperature, extreme weather events and economic activity. Other studies investigated the stochastic behavior of daily prices and provide an econometric analysis, such as Paoletta and Taschini (2008), Seifert et al. (2008), Daskalakis et al. (2009), Benz and Trück (2009), Conrad et al. (2012), Benschop and Lopez Cabrera (2014) and Gil-Alana et al. (2016). Gil-Alana et al. (2016) examine the presence of long-memory and structural breaks in the spot market, concluding that spot prices show both structural breaks and long memory. These findings support our proposal of using long memory models, such as HAR class models, for CO₂ future dynamics. Segnon et al. (2017) review the state-of-the-art models for price volatility and compare these by using several forecasting evaluation criteria and a superior predictive ability test for different forecasting horizons. They conclude that Markov-switching multifractal models perform at least as good as GARCH-type models.

Estimates of volatility using intraday data are more accurate. However, only few studies have used intraday data of CO₂ emissions futures contracts. Chevallier and Benoît (2011) is one of them, studying the conditional and unconditional distributions of the RV for the 2008 futures contract in the European Climate Exchange (ECX) based on intraday data. The author uses a simplified version of the HAR model for the long-memory effect and find that this model outperforms a simple GARCH model. Due to the shorter dataset, he does not use the ARFIMA model as a benchmark model, which is a natural choice for modeling long memory. Rotfuß (2009) analyzes intraday volatility patterns and observes a deviation from the typical U-shape with higher intra-day volatility at the beginning and end of the trading day. Hitzemann et al. (2015) use intraday data to examine the impact of the yearly announcements of realized emissions on the price, trading volume and volatility and find abnormal returns, increased trading volumes and volatility on the announcement day.

The contribution of this analysis is fourfold. First, we introduce HAR class models to model on an intradaily basis the volatility in the CO₂ markets, which capture the long memory and asymmetry in the data. Second, we use the estimates of the models to gain a better understanding of the volatility drivers in the market. Third, we use these models for volatility forecasting, which is important for risk modeling, such tail risk measures via value-at-risk and option pricing. Fourth, we use the calibrated HAR models to simulate option prices and evaluate structure of the implied volatility. To the best of our knowledge this is the first study to use HAR class models to model, forecast and explain the volatility process on a long time series of EUA futures prices. Previously, the available time series were too short to estimate such models. However, for the above-mentioned reasons, the HAR class models are an appealing alternative.

Our focus is on the forecasting performance of the models, as this is the most important aspect for risk management, option pricing and value-at-risk calculations, but we also seek to understand and explain the volatility process. The models allow us to explain the volatility structure in the market by applying the Heterogeneous Markets Hypothesis. The proposed methods are evaluated based on in-sample fit and their predictive accuracy. We find that HAR class models provide a better fit and forecasting performance, which indicates that long memory is present in the data and identify drivers of the volatility process by interpreting the models' estimated parameters.

The remainder of this chapter is organized as follows. Section 2 gives a brief overview of the EU ETS, the EU carbon market and the characteristics of EUAs. The models for modeling the RV of EUA futures prices, the HAR, SHAR and HARQ models, are presented in Section 3. Section 4 presents the empirical results of our analysis by starting with a description of the data used. Subsequently, we interpret the estimated parameters and evaluate the model fit and forecasting performance. We also describe the forecasting methodology and model comparison criteria. Section 5 presents an empirical application of the calibrated HAR model to option pricing in a simulation study. Section 6 concludes and makes suggestions for further research. All computations for this study were carried out in R.

2 EU ETS and Data

The EU ETS is the key tool of the European Commission to reduce the emissions of greenhouse gases (GHG). The EU ETS covers large facilities from GHG intensive industries in the EU. One EUA gives the right to emit one tonne of CO₂ or its equivalent of another GHG into the atmosphere. The system entered into force in 2005 through EU Directive 2003/87/EC. Since then there have been three trading periods.

The first trading period, Phase I, lasted from 2005 until 2007 and served as a pilot period to test the market infrastructure. In Phase I the EUAs were freely distributed to the emitting installations. However, the liquidity in the market was low and due to oversupply and the fact that the allowances lost their value at the end of the trading period (no bankability to second trading period), prices collapsed towards the end of the trading period. Phase II, which lasted from 2008 until 2012, was the first Kyoto commitment period. Since Phase II banking of allowances between years and trading periods is allowed, which reduces the risk of prices to collapse towards the end of the trading period (European Commission, 2012). Both in Phases I and II the allowances were distributed by the principle of grandfathering, i.e., the number of allowances a firm received were relative to the historical emission levels of its installations. In the current Phase III, which runs from 2013 until 2020, free allocation of EUAs is gradually being replaced by auctioning. Yet, there is a vivid secondary market for EUAs.

The EU ETS is a cap-and-trade system, which means that the regulator, the European Commission, fixes the total amount of emissions and allowances issued in a period. If a firm's emissions exceed the allocated volume of allowances, they can either buy allowances on the market or take abatement measures. Similarly, surplus allowances can be sold. In this way, the right to emit CO₂ becomes a tradable asset. The advantage of cap-and-trade system is that the marginal abatement costs are equalized among the firms, independent of the initial allocation of allowances (Hintermann, 2010). Each year on April 30 firms have to surrender the number of allowances corresponding to the emissions of the previous year. If they fail to do so, the firms have to pay a penalty, 40 EUR and 100 EUR per ton CO₂ emitted in Phases I and II respectively, and have to surrender the lacking allowances next year.

Besides the EAU spot market, there is also a large market for derivatives on EUAs, such as options and futures, which are traded on several exchanges in Europe. In fact, the futures market for contracts near to the expiry date is more liquid than the spot market and has larger trade volumes and is therefore the focus of our study. Future contracts with expiry dates in March, June, September and December are traded.

3 Methodology

In this section we introduce the definition of RV for different time horizons and we present the HAR model as proposed by Corsi (2009). As the simple HAR model cannot model asymmetries in the volatility process which we observed, we also introduce two extensions to the HAR model that allow to model asymmetries.

3.1 Realized volatility

As shown in the seminal papers on integrated volatility (IV) and realized volatility (RV), the integrated variance can be consistently estimated by using the sum of intraday squared returns, which is the RV (Andersen et al., 2001, 2003). We refer to Andersen et al. (2001, 2003) for the derivation of the RV definition. RV converges uniformly in probability to IV. The common definition for daily RV is

$$RV_t^{(d)} = \sqrt{\sum_{j=0}^{M-1} r_{t-j\cdot\Delta}^2}, \quad (1)$$

where the interval length $\Delta = 1d/M$, M the number of time intervals per day, $r_{t-j\cdot\Delta} = p(t - j \cdot \Delta) - p(t - (j + 1) \cdot \Delta)$ defines continuously compounded returns, and t refers to the day. We use the annualized volatility, meaning that we multiply $RV_t^{(d)}$ by $\sqrt{252}$, the square root of the number of business days per year. The RV is, under certain conditions, consistent for the true latent volatility (Andersen et al., 2003). Dictated by the availability of the data, we use a 15-minute interval ($\Delta = 15$) which yields 40 intervals per day ($M = 40$). This is in line with the findings of Chevallier and Benoît (2011), who analyze the optimal sampling frequency for computing the RV of futures on EUAs in 2008. They come to the conclusion that 15 minutes is the optimal sampling frequency to minimize the effect of the microstructure noise. Yet, this sampling frequency yields a sufficient number of daily observations to compute the RV. For our empirical analysis we do not use the returns across days, therefore we lose one observation. The RV for a different time horizon can be computed using the average over the daily returns over the past period, i.e.

$$RV_t^{(pd)} = \frac{1}{p} \left(RV_t^{(d)} + RV_{t-1d}^{(d)} + \dots + RV_{t-(p-1)d}^{(d)} \right) \quad (2)$$

where p is 5 for weekly RV (5 days) and 22 for monthly RV (22 days). The series of weekly and monthly RV are smoothed curves of the daily RV.

3.2 HAR model

The HAR model, as introduced by Corsi (2009), is motivated by the Heterogenous Market Hypothesis, which acknowledges the differences between market participants. In case of financial markets, such as the EU ETS market, this heterogeneity emanates, amongst others, from differences in participants' information and investment strategies and horizons. These differences lead to different trading frequencies, where we distinguish between short-term (daily or more frequent), medium-term (weekly) and long-term (monthly or lower frequency) traders. These different actors cause different types of interrelated volatility patterns. The advantages of the model are its computational simplicity, the compelling interpretation and the good out-of-sample forecasting performance for data exhibiting long memory properties.

The HAR model is an additive cascade of partial volatilities. The model considers volatility components with three time horizons: 1 day (d), 1 week (5 days) (w) and 1 month

(22 days) (m) and is therefore coined HAR-RV(1,5,22). We refer to the HAR-RV(1,5,22) model as HAR model for conciseness, acknowledging that other lags can be used. The model is defined as follows:

$$\tilde{\sigma}_{t+1m}^{(m)} = \alpha^{(m)} + \rho^{(m)} \text{RV}_t^{(m)} + \tilde{\omega}_{t+1m}^{(m)} \quad (3)$$

$$\tilde{\sigma}_{t+1w}^{(w)} = \alpha^{(w)} + \rho^{(w)} \text{RV}_t^{(w)} + \gamma^{(m)} \text{E}_t[\sigma_{t+1m}^{(m)}] + \tilde{\omega}_{t+1w}^{(w)} \quad (4)$$

$$\tilde{\sigma}_{t+1d}^{(d)} = \alpha^{(d)} + \rho^{(d)} \text{RV}_t^{(d)} + \gamma^{(w)} \text{E}_t[\sigma_{t+1w}^{(w)}] + \tilde{\omega}_{t+1d}^{(m)} \quad (5)$$

where $\tilde{\sigma}_t^{(\cdot)}$ is the volatility generated by a certain market component, $\tilde{\omega}_{t+1m}^{(m)}$, $\tilde{\omega}_{t+1m}^{(w)}$ and $\tilde{\omega}_{t+1d}^{(d)}$ are contemporaneously and serially independent zero-mean innovations. The true latent partial volatility is modeled as a function of the observed backward-looking RV at the same time horizon and the expectation of the true latent partial volatility at the longer time horizon. The short-term volatility is influenced by the longer-term volatility, but not vice versa. Recursively substituting Equations (2) and (3) into Equation (4) yields:

$$\text{RV}_{t+1d}^{(d)} = c + \beta^{(d)} \text{RV}_t^{(d)} + \beta^{(w)} \text{RV}_t^{(w)} + \beta^{(m)} \text{RV}_t^{(m)} + \omega_{t+1d} \quad (6)$$

where $\omega_{t+1d} = \tilde{\omega}_{t+1d}^{(d)} + \omega_{t+1d}^{(d)}$. Corsi (2009) has shown in simulations that this simple model can reproduce the dynamics in the empirical data. Note that the HAR model can be reformulated as a restricted AR(p) model, where $p = 22$, the number of business days per month:

$$\text{RV}_{t+1d}^{(d)} = \phi_0 + \sum_{i=1}^{22} \phi_i \text{RV}_{t-(i-1)d}^{(d)} + \omega_{t+1d} \quad (7)$$

with restrictions $\phi_1 = \beta^{(d)} + \frac{1}{5}\beta^{(w)} + \frac{1}{22}\beta^{(m)}$, $\phi_2 = \dots = \phi_5 = \frac{1}{5}\beta^{(w)} + \frac{1}{22}\beta^{(m)}$, and $\phi_6 = \dots = \phi_{22} = \frac{1}{22}\beta^{(m)}$. These restrictions can be F -tested for statistical significance.

3.3 Extensions to the HAR model

In high-frequency financial data an asymmetric response of volatility to positive and negative returns is often observed (Bollerslev et al., 2006). The basic HAR model, as proposed by Corsi (2009), cannot capture this asymmetry. Recently, several extensions to the HAR model have been proposed to model this asymmetry. In this analysis we investigate the use of the SHAR model introduced by Patton and Sheppard (2015), who observe that the impact of positive and negative volatility on future volatility is asymmetric. The SHAR model is based on the concept of semi-variance as introduced by Barndorff-Nielsen et al. (2010). Semi-variance decomposes the RV using signed returns into positive semi-variance

(RV_t^+) and negative semi-variance (RV_t^-), which are defined as

$$RV_t^+ = \sum_{i=1}^M r_{t,i}^2 \mathbf{I}_{\{r_{t,i}>0\}} \text{ and} \quad (8)$$

$$RV_t^- = \sum_{i=1}^M r_{t,i}^2 \mathbf{I}_{\{r_{t,i}<0\}}, \quad (9)$$

where \mathbf{I} is the indicator function. Note that $RV^{(d)} = RV_t^+ + RV_t^-$ and hence the semi-variance is a complete decomposition of RV. The SHAR model is defined as

$$RV_{t+1d} = c + \beta_+^{(d)} RV_t^+ + \beta_-^{(d)} RV_t^- + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + u_{t+1d}. \quad (10)$$

The SHAR model was shown to outperform other extensions to the HAR model capturing these asymmetries in empirical applications for financial time series. In our empirical analysis we show that this implication holds for CO₂ as well.

RV converges uniformly in probability to IV, but in finite samples there is always an estimation error. Recently, Bollerslev et al. (2016) have proposed a new family of adaptive HAR models that allow the parameters of the model to vary with the degree of estimation error of RV measures. They observe that the size of the measurement error changes over time and varies with the value of RV. Therefore, they propose to capture the changes over time in the measurement error by using time-varying auto-regressive parameters. The model is called HARQ and is defined as

$$RV_{t+1d} = c + \left(\beta^{(d)} + \beta_Q^{(d)} RQ_t^{1/2} \right) RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + u_{t+1d}, \quad (11)$$

where $RQ_t = \frac{M}{3} \sum_{i=1}^M r_{t,i}^4$ is the so-called Realized Quarticity (RQ) and is a consistent estimator of the Integrated Quarticity (IQ). A logic extension to the HARQ model is the HARQF model (Bollerslev et al., 2016), which has time-varying coefficients for all lags. However, this model is too large to estimate with the data at hand.

Both the SHAR and HARQ models incorporate different asymmetries into the simple HAR model and have been shown to yield better model fits and forecasting performance.

4 Empirical analysis

For our empirical analysis we use 15 minute intraday futures price data from the European Climate Exchange (ECX), the world's largest carbon exchange, with tickers CFI2YZ8, CFI2Z8, CFI2Z9, CFI2Z0, CFI2Z1, CFI2Z2. 'CFI' stands for ECX Carbon Financial Instrument, '2' for the second trading period of the EU ETS, 'Z' for the month of expiry (December) and the last digit for the year of expiry. We use data from the period May 21, 2007 until April 13, 2012, which covers with 49,279 observations 1,250 trading days. The data include futures contracts on six different underlying assets, which are all expiring in

the second trading period (December 2008 - December 2012). For a description of CO₂ futures we refer to Benz and Klar (2008). For computational purposes we perform our analysis on log returns of the futures prices for 15-minute time intervals, which are defined as

$$r_{t,i} = \log \left(\frac{F_{t,i}}{F_{t,i-1}} \right) \quad (12)$$

where t indicates the day and $F_{t,i}$ is the futures contract last trading price in the i^{th} 15 minutes interval on day t on the futures market. Trading takes place from 7:00 AM until 5:00 PM GMT, which yields 40 intervals of 15 minutes per day and $i \in \{1, \dots, 40\}$. By not computing the log returns across days, we lose one observation per day.

For our analysis we use the data for the futures December contract closest to expiry. We use the December contracts, since these are the most liquid contracts compared with futures contract with other expiry months. Due to the erratic behavior of the futures prices close to expiry, we roll-over to the next contract on December 1 of each year, whereas the expiry date is December 15. Due to the non-bankability of EUAs from Phase I to Phase II, the December 2007 futures lost most of their value towards the end of 2007. Since our focus is on Phase II, we use the December 2008 futures rather than the December 2007 futures in 2007. Therefore, there is a roll-over in September 2008 instead of in December 2007, where we roll over from the contract with ticker CFI2YZ8 to a contract with ticker CFI2Z8, both expiring in December 2008. This yields one log-return series for the period of analysis.

Figure 1 presents a plot of the EUA futures prices at 15-minute intervals for all trading days. Figure 2 shows the log returns and clearly shows volatility clustering and heteroskedasticity. The vertical red bars in both figures indicate the roll-overs to the next futures contract. The future prices in Phase II are, contrary to the prices in Phase I, always positive in the period under consideration. The plot of the prices shows well a decrease of the prices in 2009 and 2011, which corresponds to the effect of the economic crises in both periods. Table 1 presents the descriptive statistics of the futures prices and log returns, which indicate fat tails in the log returns.

	Futures prices	Log returns
Minimum	6.05	-1.92
Maximum	29.57	1.30
Mean	16.10	-0.00
Standard deviation	5.00	0.07
Skewness	0.41	-0.40
Kurtosis	2.51	38.89

Table 1: Descriptive statistics of futures prices and log returns

While exhibiting the stylized facts of financial time series, such as heterogeneity, leptokurticity and skewness, the data also shows long memory properties. Figure 3 shows the slowly decaying ACF of the log returns series, which is an indication for a long memory process. The ACF lies outside the confidence band well past lag 20.




Figure 1: Future prices series from May 21, 2007 until April 13, 2012 (red vertical lines indicate roll-overs to next maturity of futures contract)




Figure 2: Log return series from May 21, 2007 until April 13, 2012 (red vertical lines indicate roll-overs to next maturity of futures contract)




Figure 3: ACF of log return series, blue dashed lines indicate confidence interval

Figure 4 plots the daily, weekly and monthly RV as defined in Equations (1) and (2) on a log scale. The series of weekly and monthly RV are smoothed curves of the daily RV. We can clearly observe long-memory properties, volatility clustering and peaks, especially in periods when the price decreases volatility seems to be higher. This can be explained by the fact that the supply of EUAs is inelastic. When the demand decreases due to an external shock, there might rise doubt about the overall shortage of certificates on the EUA market, which Benschop and Lopez Cabrera (2014) show by interpreting different regimes.

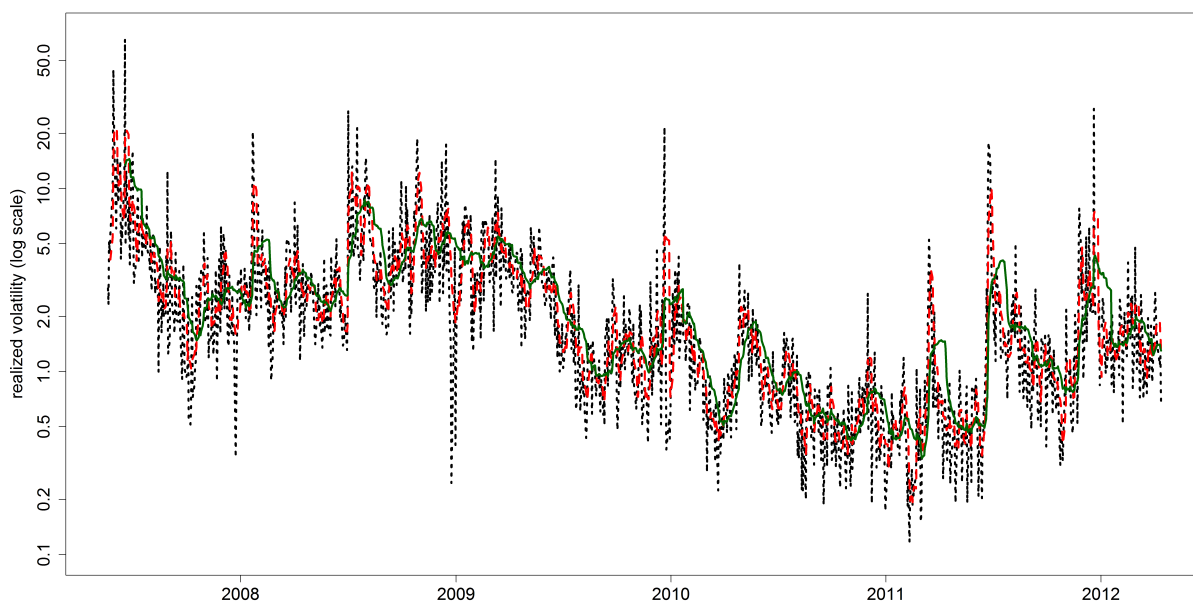


Figure 4: Daily (dotted black), weekly (dashed red) and monthly (solid green) RV series from May 21, 2007 until April 13, 2012 on a log scale

4.1 Estimation results

In this section we present the results of estimating the HAR models on the log returns in the complete sample (June 19, 2007 until April 11, 2012) and interpret the estimated parameters of the HAR class models. First, we perform the augmented Dickey-Fuller test, which indicates the stationarity of the data.

Table 2 shows the OLS coefficients when estimating the HAR model in Equation (6) on the complete sample. All coefficients are highly significant. We apply an F -test to test the restrictions in the HAR model w.r.t. the AR(22) model and find that the restrictions hold. One interpretation of the coefficients is the share of impact on the volatility by different market participants. We note that the weekly component has the largest impact. This indicates that market participants with longer horizons have more impact on the volatility, since fewer speculators were active in the market.

	HAR	SHAR	HARQ
c	0.533 (0.000)	0.546 (0.000)	0.300 (0.002)
$\beta^{(d)}$	0.183 (0.000)		0.473 (0.000)
$\beta_+^{(d)}$		-0.707 (0.361)	
$\beta_-^{(d)}$		5.147 (0.000)	
$\beta_Q^{(d)}$			-0.071 (0.000)
$\beta^{(w)}$	0.398 (0.000)	0.325 (0.000)	0.177 (0.000)
$\beta^{(m)}$	0.171 (0.027)	0.279 (0.000)	0.244 (0.000)

Table 2: Estimated coefficients for HAR, SHAR and HARQ models, p values in brackets

 RegimeSwitching

Similar to the results of Patton and Sheppard (2015) estimating the SHAR model for S&P 500 data, we find that also for the EU ETS futures data the negative semi-variance has a much larger (and positive) effect on the future volatility than its positive counterpart does. Table 2 shows the coefficients yielded from estimating the SHAR model as in Equation 9 with OLS. $\beta_+^{(d)}$ is small, negative and not significant, whereas $\beta_-^{(d)}$ is highly significant and large as compared to the combined $\beta^{(d)}$ in the HAR model in Equation 6. This shows a clear distinction between the effects of volatility attributable to negative jumps and positive jumps with negative returns causing much higher volatility. The signed jump variation can be explained by the setup of the EU ETS.

All coefficients in the HARQ model are significant as the p values of the t -test in brackets

in Table 2 indicate. We observe that more weight moves to the daily lag as compared to the HAR model. This is opposite to the findings of Bollerslev et al. (2016), who find that more weight moves to the 'cleaner' and older lags. This results indicates that the measurement error of the RV is not as large.

4.2 Model fit evaluation

To compare the performance of the HAR models we also estimate several benchmark models: autoregressive (AR) models without long memory properties (the HAR model is a restricted AR(22)) as well as a standard long-memory model, the ARFIMA model. The disadvantage of using fractional integrations is the lack of an economic interpretation. Furthermore, models with fractional integration require long time series to estimate the model. To allow for comparability we omit the first 21 days from our analysis for all models, as this is required to estimate the autoregressive models with the highest lag order. In order to evaluate the performance of the different models, we use the Akaike information criterion (AIC) as well as the adjusted R^2 . The AIC (Akaike, 1973) is defined as follows

$$\text{AIC} = -2\ell + 2k \tag{13}$$

where ℓ is the value of the estimated log likelihood function and k the number of parameters in the model.

Table 3 reports the goodness-of-fit measures for the HAR class models together with the ARFIMA and AR benchmark models. We estimate the fractional integration parameter d in the ARFIMA model on the complete sample and use subsequently an AR(5) model. The information criteria and adjusted R^2 show the best in-sample fit for the HARQ model. The HAR class models clearly outperform the standard short-memory models (AR) as well as the ARFIMA model. Besides the better model fit, the HAR class models also have a much more compelling interpretation than the fractional integration parameter of the ARFIMA model.

Model	AIC	adjusted R^2
HAR-RV	5396.892	0.363
HARQ	5274.074	0.424
SHAR	5267.936	0.379
AR1	6048.736	0.254
AR3	5967.159	0.310
AR22	5934.512	0.361
ARFIMA	5309.148	0.257

Table 3: Comparison of model fit

4.3 Forecast evaluation

Forecasting of prices and volatility is important for risk management and option pricing. Therefore we evaluate the forecasting performance of the HAR model and the extensions thereof with several benchmark models. We use the complete sample with , 1229 observations (June 19, 2007 until April 11, 2012) for the in-sample forecasts. We perform one-day ahead in-sample forecasts by estimating the model on the full sample and make forecasts for each day in the full sample. In order to test out-of-sample performance we forecast the log-returns for the out-of-sample period (June 21, 2011 until April 11, 2012).

We omit the first 21 observations, since these cannot be forecasted with the HAR models and AR model with lag length 22. For the out-of-sample forecast we use a rolling window approach with window length 200 and re-estimate the models daily. The comparison is based on the forecasts for 200 days (June 21, 2011 until April 11, 2012). We use several benchmark models to evaluate the forecasting performance. A logic benchmark model for the HAR models is the ARFIMA model, which is used in the literature as standard long-memory model. Furthermore, since the standard HAR model is a restricted version of an AR(22) model, we also include AR models with different lag lengths for comparison.

In order to evaluate the forecasting performance we compare the root mean squared error (RMSE), mean absolute error (MAE) as well as the R^2 of Mincer-Zarnowicz regressions. The focus of our analysis is on forecasting $\text{RV}_t^{(d)}$. The RME and MAE compare the true value and the forecasted value and are respectively defined as

$$\text{RMSE} = \sqrt{\frac{1}{h} \sum_{t=T+1}^{T+h} (\widehat{\text{RV}}_t^{(d)} - \text{RV}_t^{(d)})^2} \quad (14)$$

$$\text{MAE} = \frac{1}{h} \sum_{t=T+1}^{T+h} |\widehat{\text{RV}}_t^{(d)} - \text{RV}_t^{(d)}| \quad (15)$$

where $\widehat{\text{RV}}_t^{(d)}$ is the point forecast for time t , $\text{RV}_t^{(d)}$ is the true observed value and h is the forecasting horizon.

The Mincer-Zarnowicz regressions are the regressions of the RV on the model forecast and a constant, i.e.,

$$\text{RV}_t^{(d)} = b_0 + b_1 E_{t-1} \left[\widehat{\text{RV}}_t^{(d)} \right] + \varepsilon_t \quad (16)$$

If $\widehat{\text{RV}}_t^{(d)}$ is an accurate forecast of $\text{RV}_t^{(d)}$, we expect that b_0 is equal to 0 and b_1 equal to 1. We compare the R^2 of this regression, as this is an indicator of how well the forecasts predict the true values.

Patton (2011) observes that the comparison of forecasting performance with these general loss functions causes distortions if used for comparing volatility proxies, such as RV. The reason is that RV is an approximation of the integrated volatility, but the error term

remains in finite samples (see also section on SHAR). Therefore, we also use a robust loss function proposed by Patton (2011). The QLIKE loss function is defined as

$$\text{QLIKE} = \frac{1}{h} \sum_{t=T+1}^{T+h} \left(\log \text{RV}_t^{(d)} + \frac{\widehat{\text{RV}}_t^{(d)}}{\text{RV}_t^{(d)}} \right), \quad (17)$$

where $\text{RV}_t^{(d)}$ is the true conditional RV, $\widehat{\text{RV}}_t^{(d)}$ the forecasted conditional RV and h is the forecasting horizon.

4.3.1 In-sample forecasts

Table 4 reports the SSE, RMSE, MAE, R^2 of the Mincer-Zarnowicz regression and the robust QLIKE measure for the in-sample forecasting results of the HAR class models as well as the benchmark models. We observe that the HAR class models outperform the short-memory models and perform similarly to the ARFIMA model according to the SSE and RMSE. Looking at the robust QLIKE loss functions, these differences are less pronounced. Values of the QLIKE measure close to 1 indicate a better performance. Among the HAR class models, the SHAR and HARQ models perform better, which indicates that the asymmetry in the volatility process as discussed previously is present in the data.

Model	SSE	RMSE	MAE	R^2	QLIKE
HAR	5164.473	2.075	1.021	0.366	1.041
SHAR	5088.939	2.059	1.018	0.381	1.044
HARQ	5146.610	2.071	1.000	0.403	1.041
AR(1)	6762.333	2.374	1.192	0.245	1.103
AR(3)	5718.152	2.183	1.002	0.365	1.038
AR(22)	5158.471	2.073	0.986	0.392	1.052
ARFIMA	5083.560	2.058	1.027	0.260	1.055

Table 4: Comparison of in-sample forecasting

 RegimeSwitching

4.3.2 Out-of-sample forecasts

Table 5 presents the RMSE, MAE and R^2 of the Mincer-Zarnowicz regression one-day ahead out-of-sample forecasting results. Based on these measures, we conclude that the HAR class models perform much better than the benchmark models, but also than the ARFIMA model.

Model	RMSE	MAE	R ²	QLIKE
HAR-RV	2.238	0.939	0.283	0.974
HARQ	2.187	0.973	0.310	1.020
SHAR	2.100	0.916	0.354	0.989
AR(1)	2.486	1.106	0.141	1.031
AR(3)	2.881	1.149	0.050	1.006
AR(22)	6.371	2.828	0.000	1.218
ARFIMA	3.007	1.224	0.116	1.026

Table 5: Comparison of daily RV out-of-sample forecasting for the period June 21, 2011 until April 11, 2012

 RegimeSwitching

Table 6 reports the p-value of the Diebold-Mariano (DM) test (Diebold and Mariano, 1995) to compare the forecasts of the different models. In the DM test, we conduct a pairwise test on the equality of the mean squared forecast errors by analyzing the difference between the squared forecast errors of two models. The null hypothesis of equal performance is that $H_0 : \mu = 0$ in the regression $e_{t,1}^2 - e_{t,2}^2 = \mu + \epsilon_t$, where $e_{t,1}^2$ and $e_{t,2}^2$ denote the forecast errors in the two models. We focus on the t -statistics of parameter μ , denoted as DM t -stat, which supports the model in the column if it is significantly negative and the model in the row if it is significantly positive (significance level marked by asterisks). We find that the HAR class models are superior to the ARFIMA model in out-of-sample forecast. We also see that the SHAR model outperforms the HAR and HARQ models.

Model	HAR	HARQ	SHAR	ARFIMA
HAR		1.030	1.693**	-1.902**
HARQ	-1.030		1.625*	-1.839**
SHAR	-1.693**	-1.625*		-1.910**
ARFIMA	1.902**	1.839**	1.910**	

Table 6: Test statistic of DM-test to test for superiority of the model in the rows over the model in the columns, i.e., $H_0 : \mu = 0$ in the regression $e_{t,1}^2 - e_{t,2}^2 = \mu + \epsilon_t$, where $e_{t,1}^2$ and $e_{t,2}^2$ denote the forecast errors in respectively the model in the column and the model in the row. The test is modified with robust Newey-West variances for heteroscedasticity and autocorrelation with the lags equal to the forecast horizon. * denotes a significance level of 10%, ** denotes a significance level of 5%.

 RegimeSwitching

To analyze the out-of-sample forecasting performance for longer horizons (weekly and monthly), we use direct projection, i.e. we reestimate the models for RV_w and RV_m

and forecast these for the next time period. The SHAR and HARQ models do not allow for forecasting more than one day ahead as we do not forecast RQ , RV^+ or RV^- respectively for the models. For the HARQ model we use a slightly adapted model with the time-varying parameter at the lag of interest. The models for the weekly and monthly forecasting horizons are coined HARQ-h by Bollerslev et al. (2016) and defined as

$$RV_{t+1d}^w = c + \beta^{(d)} RV_t^{(d)} + \left(\beta^{(w)} + \beta_Q^{(w)} RQ_t^{(w)1/2} \right) RV_t^{(d)} + \beta^{(m)} RV_t^{(m)} + u_{t+1d}, \quad (18)$$

where $RQ_t^{(w)} = \frac{1}{5} \sum_{i=1}^5 RQ_{t-i+1}$ is the weekly RQ.

$$RV_{t+1d}^m = c + \beta^{(d)} RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \left(\beta^{(m)} + \beta_Q^{(m)} RQ_t^{(m)1/2} \right) RV_t^{(m)} + u_{t+1d}, \quad (19)$$

where $RQ_t^{(m)} = \frac{1}{22} \sum_{i=1}^{22} RQ_{t-i+1}$ is the monthly RQ. The results of the weekly and monthly out-of-sample forecast of RV are presented in Table 7. Here we also see that the HAR class models perform better than the benchmark models, including the ARFIMA model.

Model	Weekly RV				Montly RV			
	RMSE	MAE	R ²	QLIKE	RMSE	MAE	R ²	QLIKE
HAR-RV	1.432	0.842	0.878	0.983	0.164	0.075	0.970	1.017
HARQ-h	0.574	0.274	0.878	0.984	0.163	0.074	0.971	1.015
SHAR	0.707	0.522	0.886	1.053	0.576	0.561	0.972	1.112
AR(1)	0.898	0.622	0.812	1.038	0.479	0.437	0.962	1.063
AR(3)	1.848	0.735	0.332	1.024	2.336	1.007	0.115	1.086
AR(22)	1.089	0.493	0.623	1.016	0.992	0.317	0.538	1.037
ARFIMA	0.744	0.346	0.821	0.995	0.183	0.088	0.964	1.016

Table 7: Comparison of weekly and monthly RV out-of-sample forecasting for the period June 21, 2011 until April 11, 2012

 RegimeSwitching

5 Application to option pricing

Accurate volatility forecasts can be used for pricing of options as volatility is the main driver of option prices. Therefore, to test the suitability of the HAR models for option pricing, we use the calibrated HAR model to simulate options prices on futures. We show that this model allows to reproduce a realized term structure of the implied volatility that mimics those of an efficient market. Hitzemann and Uhrig-Homburg (2013) show for the years 2010 to 2012 that the option market on CO₂ futures exhibits a downward-sloping

smile. They interpret this as an indication for an efficient market that reflects well the distributional properties of emission permit prices based on the system design.

To simulate the option prices using the HAR model for RV, we use an approach based on the RV option pricing model as proposed by Corsi et al. (2013), which consists of several steps. First, we use the coefficients of the HAR model calibrated on the data of the in-sample period to simulate one year of daily volatilities. Second, using these volatilities, we simulate Monte Carlo price paths for the futures contracts based on a simple pricing process using the simulated stochastic volatility series and calibrated model parameters under the historical measure. With these simulated Monte Carlo price paths, we compute the call option price c at time $t = 0$ as

$$c = \frac{1}{M} \sum_{m=1}^M \max_{t \in \{1, \dots, T\}} (F_{t,m} - K, 0), \quad (20)$$

where $F_{t,m}$ is the simulated futures prices at time t in simulation path m , K is the strike price, M the number of Monte Carlo price paths and T the number of days until maturity. In our analysis, we set the number of simulated Monte Carlo price paths $M = 10,000$ and $T = 252$, which corresponds to one year. We repeat this analysis for a set of strike prices. Subsequently, we derive the implied volatility using the Black (1976) formula for pricing call options on futures. We solve the Black formula for the implied volatility σ .

$$c = e^{-rT} [F\mathcal{N}(d_1) - K\mathcal{N}(d_2)], \quad (21)$$

where $d_1 = \frac{\ln(F/K) + (\sigma^2/2)T}{\sigma\sqrt{T}}$ and $d_2 = d_1 - \sigma\sqrt{T}$ and r the risk-free rate.

Since the results depend on the start date of the simulation, we apply this approach for 20 different start dates. The first start date is August 1, 2007 and the subsequent start dates are the first trading date of the month with two month intervals. This means that the next start dates are October 1, 2007, December 3, 2007, etc. until April 2, 2012. In this way we obtain 20 different implied volatility curves for an option of one year to maturity. Figure 5 shows the average of these curves in thick black as well as the individual curves. The individual curves are ordered by their shade of grey: the lightest grey represents the curve obtained by using as starting day August 1, 2007 and the darkest grey corresponds to the starting date April 2, 2012. We observe that the the curves corresponding to earlier dates are flatter and above the average line. This indicates that there is a learning effect in the market and market participants have learned about the implications of the design of the system for option prices.

The average realized term structure shows a clear downward-sloping smile, which is in line with the findings of Hitzemann and Uhrig-Homburg (2013). This indicates that the HAR model can well mimic the long-term volatility structure in the futures market and is suitable for short-term and long-term option pricing. The volatility smile has its lowest point at a moneyness of approximately 1.12, which means that options with lower strike prices (deeper in-the-money) have a higher demand than option with higher strike prices (out-of-the-money). From a risk management perspective, this means that market participants want to speculate on falling prices.




Figure 5: Volatility smile based on simulation of options prices with HAR model one year ahead for different start dates (grey) and the average curve (black)

 RegimeSwitching

6 Conclusion

Modeling the volatility is important for traders and risk management in CO₂ permit markets under the EU-ETS. Especially forecasts of volatility are important for option pricing. We use intra-day data for over five years of futures prices on contracts with maturity in the second phase of the EU-ETS to compute an RV measure. In this study we show that the HAR class models have very good modeling and forecasting properties for the RV of CO₂ futures contracts and present an alternative to the ARFIMA model for long-memory processes. The performance for in-sample fit is comparable to the performance of the standard long memory model, ARFIMA, which we use as benchmark model and outperforms this for out-of-sample forecasts. Acknowledging the shortcomings of the simple HAR model and observing the asymmetries in the volatility process, we continue by using the SHAR and HARQ models, which are extensions of the standard model taking into account the asymmetric effect of positive and negative returns on the volatility (SHAR) and the effect of the size of the measurement error of RV on the volatility forecast (HARQ). The HARQ and SHAR models perform better than the simple HAR model, indicating that the characteristics described above are present in the data.

The HAR class models clearly outperform simple short memory models (AR). This is an indication that long-memory is present in the data. Yet, the HAR class models are parsimonious and have a compelling interpretation, which the ARFIMA model lacks, and are also applicable to shorter time-series. Looking at the estimated parameters of the HAR class models, we observe that the all three components (day, week and month) are important for explaining the volatility structure of CO₂ futures. This can point at different

volatility drivers. The estimated parameters for the realized semi-variance in the SHAR class model show that negative returns are coupled with higher volatility than positive returns, which is in line with the design of the EU ETS. In our forecasting exercise, we show that the SHAR and HARQ model outperform the ARFIMA and standard HAR model at the daily, weekly and monthly horizons.

By simulating option prices on futures using the HAR model, we show that this model can well mimic the volatility structure of an efficient market, which is also in line with the observed term structure in the CO₂ option market. This indicates that the models are well suited for short-term and long-term option pricing. With HAR class models, which are in fact simple auto-regressive models and can all be estimated with least-squares, we obtain better forecasting results and a better understanding of the volatility dynamics of EU-ETS futures contracts.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In Petrov, B. and Csaki, F., editors, *Second international symposium on information theory*, pages 267–281. Budapest: Akademiai Kiado.
- Alberola, E., Chevallier, J., and Chèze, B. (2007). European carbon prices fundamentals in 2005-2007: the effects of energy markets, temperatures and sectorial production. *EconomiX Working Papers 2007-33*, University of Paris West - Nanterre la Défense, EconomiX.
- Alberola, E., Chevallier, J., and Chèze, B. (2008a). The EU emissions trading scheme : The effects of industrial production and CO₂ emissions on carbon prices. Open Access publications from Université Paris-Dauphine urn:hdl:123456789/4223, Université Paris-Dauphine.
- Alberola, E., Chevallier, J., and Chèze, B. (2008b). Price drivers and structural breaks in European carbon prices 2005-2007. *Energy Policy*, 36(2):787–797.
- Andersen, T., Bollerslev, T., Diebold, F., and Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1):43–76.
- Andersen, T., Bollerslev, T., Diebold, F., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625.
- Barndorff-Nielsen, O. E., Kinnebrock, S., and Shephard, N. (2010). Measuring downside risk - realised semivariance. In Bollerslev, T., Russell, J., and Watson, M., editors, *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, pages 117–136. Oxford University Press.
- Benschop, T. and Lopez Cabrera, B. (2014). Volatility modeling of CO₂ emission allowance spot prices with regime-switching GARCH models. Technical report, SFB 649 Discussion Paper 2014-050, Humboldt-Universität zu Berlin.
- Benz, E. and Klar, J. (2008). Liquidity and price discovery in the European CO₂ futures market: an intraday analysis. Technical report, Bonn Graduate School of Business.
- Benz, E. and Trück, S. (2009). Modeling the price dynamics of CO₂ emission allowances. *Energy Economics*, 31(1):4–15.
- Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3:167 – 179.
- Bollerslev, T., Litvinova, J., and Tauchen, G. (2006). Leverage and volatility feedback effects in high-frequency data. *Journal of Financial Econometrics*, 4(3):353–384.
- Bollerslev, T., Patton, A. J., and Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18.
- Chevallier, J. (2009). Carbon futures and macroeconomic risk factors: A view from the EU ETS. *Energy Economics*, 31(4):614–625.

- Chevallier, J. and Benoît, S. (2011). On the realized volatility of the ECX CO₂ emissions 2008 futures contract: distribution, dynamics and forecasting. *Annals of Finance*, 7:1–29.
- Conrad, C., Rittler, D., and Rotfuß, W. (2012). Modeling and explaining the dynamics of European Union Allowance prices at high-frequency. *Energy Economics*, 34(1):316–326.
- Corsi, F. (2009). A simple approximate long-memory model for realized volatility. *Journal of Financial Econometrics*, 7:174–196.
- Corsi, F., Fusari, N., and Vecchia, D. L. (2013). Realizing smiles: Options pricing with realized volatility. *Journal of Financial Economics*, 107(2):284 – 304.
- Daskalakis, G., Psychoyios, D., and Markellos, R. N. (2009). Modeling CO₂ emission allowance prices and derivatives: Evidence from the European trading scheme. *Journal of Banking & Finance*, 33(7):1230–1241.
- Diebold, F. and Mariano, R. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3):253–63.
- European Commission (2012). The EU emissions trading system (EU ETS).
- Gil-Alana, L. A., Gupta, R., and de Gracia, F. P. (2016). Modeling persistence of carbon emission allowance prices. *Renewable and Sustainable Energy Reviews*, 55:221–226.
- Hammoudeh, S., Nguyen, D. K., and Sousa, R. (2014a). Energy prices and CO₂ emission allowance prices: A quantile regression approach. *Energy Policy*, 70(C):201–206.
- Hammoudeh, S., Nguyen, D. K., and Sousa, R. (2014b). What explains the short-term dynamics of the prices of CO₂ emissions? *Energy Economics*, 46(C):122–135.
- Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental Economics and Management*, 59(1):43–56.
- Hitzemann, S. and Uhrig-Homburg, M. (2013). Empirical performance of reduced-form models for emission permit prices. Discussion papers, Karlsruhe Institute of Technology (KIT) - Financial Engineering and Derivatives Department.
- Hitzemann, S., Uhrig-Homburg, M., and Ehrhart, K.-M. (2015). Emission permits and the announcement of realized emissions: Price impact, trading volume, and volatilities. *Energy Economics*, 51:560–569.
- Mansanet-Bataller, M., Pardo, A., and Valor, E. (2007). CO₂ prices, energy and weather. *The Energy Journal*, 0(3):73–92.
- Paoletta, M. S. and Taschini, L. (2008). An econometric analysis of emission allowance prices. *Journal of Banking and Finance*, 32(10):2022–2032.
- Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160:246–256.

- Patton, A. J. and Sheppard, K. (2015). Good Volatility, Bad Volatility: Signed Jumps and The Persistence of Volatility. *The Review of Economics and Statistics*, 97(3):683–697.
- Rotfuß, W. (2009). Intraday price formation and volatility in the European Union Emissions Trading Scheme: An introductory analysis. ZEW Discussion Papers 09-018, ZEW - Zentrum für Europäische Wirtschaftsforschung / Center for European Economic Research.
- Segnon, M., Lux, T., and Gupta, R. (2017). Modeling and forecasting the volatility of carbon dioxide emissions allowance prices: A review and comparison of modern volatility models. *Renewable and Sustainable Energy Reviews*, 69:692–704.
- Seifert, J., Uhrig-Homburg, M., and Wagner, M. (2008). Dynamic behavior of CO₂ spot prices. *Journal of Environmental Economics and Management*, 56(2):180–194.


SFB 649 Discussion Paper Series 2017

For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 001 "Fake Alpha" by Marcel Müller, Tobias Rosenberger and Marliese Uhrig-Homburg, January 2017.
- 002 "Estimating location values of agricultural land" by Georg Helbing, Zhiwei Shen, Martin Odening and Matthias Ritter, January 2017.
- 003 "FRM: a Financial Risk Meter based on penalizing tail events occurrence" by Lining Yu, Wolfgang Karl Härdle, Lukas Borke and Thijs Benschop, January 2017.
- 004 "Tail event driven networks of SIFIs" by Cathy Yi-Hsuan Chen, Wolfgang Karl Härdle and Yarema Okhrin, January 2017.
- 005 "Dynamic Valuation of Weather Derivatives under Default Risk" by Wolfgang Karl Härdle and Maria Osipenko, February 2017.
- 006 "RiskAnalytics: an R package for real time processing of Nasdaq and Yahoo finance data and parallelized quantile lasso regression methods" by Lukas Borke, February 2017.
- 007 "Testing Missing at Random using Instrumental Variables" by Christoph Breunig, February 2017.
- 008 "GitHub API based QuantNet Mining infrastructure in R" by Lukas Borke and Wolfgang K. Härdle, February 2017.
- 009 "The Economics of German Unification after Twenty-five Years: Lessons for Korea" by Michael C. Burda and Mark Weder, April 2017.
- 010 "Data Science & Digital Society" by Cathy Yi-Hsuan Chen and Wolfgang Karl Härdle, May 2017.
- 011 "The impact of news on US household inflation expectations" by Shih-Kang Chao, Wolfgang Karl Härdle, Jeffrey Sheen, Stefan Trück and Ben Zhe Wang, May 2017.
- 012 "Industry Interdependency Dynamics in a Network Context" by Ya Qian, Wolfgang Karl Härdle and Cathy Yi-Hsuan Chen, May 2017.
- 013 "Adaptive weights clustering of research papers" by Larisa Adamyan, Kirill Efimov, Cathy Yi-Hsuan Chen, Wolfgang K. Härdle, July 2017.
- 014 "Investing with cryptocurrencies - A liquidity constrained investment approach" by Simon Trimborn, Mingyang Li and Wolfgang Karl Härdle, July 2017.
- 015 "(Un)expected Monetary Policy Shocks and Term Premia" by Martin Kliem and Alexander Meyer-Gohde, July 2017.
- 016 "Conditional moment restrictions and the role of density information in estimated structural models" by Andreas Tryphonides, July 2017.
- 017 "Generalized Entropy and Model Uncertainty" by Alexander Meyer-Gohde, August 2017.
- 018 "Social Security Contributions and the Business Cycle" by Anna Almosova, Michael C. Burda and Simon Voigts, August 2017.
- 019 "Racial/Ethnic Differences In Non-Work At Work" by Daniel S. Hamermesh, Katie R. Genadek and Michael C. Burda, August 2017.
- 020 "Pricing Green Financial Products" by Awdesch Melzer, Wolfgang K. Härdle and Brenda López Cabrera, August 2017.
- 021 "The systemic risk of central SIFIs" by Cathy Yi-Hsuan Chen and Sergey Nasekin, August 2017.
- 022 "Das deutsche Arbeitsmarktwunder: Eine Bilanz" by Michael C. Burda and Stefanie Seele, August 2017.

SFB 649, Spandauer Straße 1, D-10178 Berlin
<http://sfb649.wiwi.hu-berlin.de>

This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".



SFB 649 Discussion Paper Series 2017

For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 023 "Penalized Adaptive Method in Forecasting with Large Information Set and Structure Change" by Xinjue Li, Lenka Zbonakova and Wolfgang Karl Härdle, September 2017.
- 024 "Smooth Principal Component Analysis for High Dimensional Data" by Yingxing Li, Wolfgang K. Härdle and Chen Huang, September 2017.
- 025 "Realized volatility of CO2 futures" by Thijs Benschop and Brenda López Cabrera, September 2017.

SFB 649, Spandauer Straße 1, D-10178 Berlin
<http://sfb649.wiwi.hu-berlin.de>

This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".

