

On the realized volatility of the ECX emissions 2008 futures contract: distribution, dynamics and forecasting

Julien Chevallier (Imperial College London)
Benoît Sévi (Université d'Angers)


International Energy Workshop – Venice – June 17-19, 2009

The European carbon market: banking, pricing and risk-hedging strategies

- ▶ One EU allowance (EUA) is equal to one ton of CO₂ emitted in the atmosphere.
- ▶ Covers up to 46% of CO₂ emissions from European energy-intensive industries (combustion, iron and steel, pulp and paper, refineries, cement, etc.).
- ▶ A new commodity market attracting increasing academic interest, in a moving institutional context. A literature still in its “infancy”, with influential early work by D. Ellerman and co-authors.


Design issues: transactions

- ▶ The volume of transaction has been increasing from 262 million tons in 2005 to 1,443 million tons in 2007.



Design issues: allocation

- ▶ 2.2 billions of allowances per year were allocated to 10,600 installations ($>20\text{MWTh}$) across 27 EU Member States in 2005-07.
- ▶ 2.08 billions of quotas per year will be distributed during 2008-2012.




Design issues: calendar


- ▶ The delivery of allowances is made on a yearly basis:
- ▶ On February 28 of year N , European operators receive their allocation for the commitment year N ;
- ▶ March 31 of year N is the deadline for the submission of the verified emissions report during year $N - 1$, from each installation to the European Commission;
- ▶ April 30 of year N is the deadline for the restitution of quotas utilized by operators during year $N - 1$;
- ▶ May 15 of year N corresponds to the deadline of the official publication by the the European Commission of verified emissions for all installations covered by the EU ETS during year $N - 1$.

Design issues: price development

- ▶ High spot price volatility during 2005-07; current medium-term price signal around €20 per ton.
- ▶ Various European-based market places for spot, futures and derivatives prices.



The divorce between Phases I and II prices



Our motivation

- ▶ A new option market in October, 2006: how to price derivatives on a new asset?
- ▶ Is the standard Brownian motion plus a drift adapted for this market in light of the behaviour of the 2008 futures contract?

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) \quad \text{with} \quad 0 \leq t \leq T$$

- ▶ Using intraday data and recent econometric techniques, we show that our series has not standard distributional properties advocated by Clark (1973).

Our motivation (con't)

- ▶ Our main issue: distributional properties. **Why is the distribution of interest?**
- ▶ Models such as in Hull and White (1987) assume stochastic volatility
- ▶ If volatility is not constant (true!) then we need as much information as possible to model volatility: distribution is one (the most important) of them.
- ▶ In addition, are increments normally distributed?
- ▶ And finally, are returns and volatility independent?

Outline of the presentation

Introduction

Distributional properties

- Intraday data

- Realized measures

- Unconditional distribution of futures returns and realized volatility

The dynamics of realized volatility

- A taste of long memory

- The HAR model

A forecasting exercise

- Method

- Results

Some concluding remarks

Intraday data

- ▶ Tick-by-tick transactions for the ECX futures contract of maturity December 2008, going from January 2 to December 15, 2008.
- ▶ This is equivalent to 240 days of trading after cleaning the data for outliers, and until the expiration of the contract.
- ▶ The total amount of intraday observations in our sample is equal to 167,004.
- ▶ The average number of transactions for the ECX carbon futures tick-data is equal to 700 trades per day (50 seconds between each transaction) vs 163 for the Eurodollar futures, 3,366 for the S&P500 futures, and 1,710 for T-bonds.

Realized measures of volatility

- ▶ $p(t)$ denote a logarithmic asset price at time t .
- ▶ With no jump, the continuous-time diffusion process generally employed in asset and derivatives pricing may be expressed by a stochastic differential equation as:

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) \quad \text{with} \quad 0 \leq t \leq T$$

- ▶ $\mu(t)$ a continuous and locally bounded variation process, $\sigma(t)$ a strictly positive càdlàg (right continuous with left limits) stochastic volatility process, and $W(t)$ a standard Brownian motion.

Realized measures of volatility (con't)

- ▶ Next, let us consider the quadratic variation (QV) for the cumulative return process $r(t) \equiv p(t) - p(0)$:

$$[r, r]_t = \int_0^t \sigma^2(s) ds$$

- ▶ The QV simply equals the integrated volatility of the process described in the previous equation.

Realized measures of volatility (con't)

- ▶ The *realized volatility* (RV) is defined as the sum of returns at a frequency $1/\Delta$, or:

$$RV_{t+1}(\Delta) \equiv \sum_{j=1}^{1/\Delta} r_{t+j.\Delta,\Delta}^2$$


- ▶ When $\Delta \rightarrow 0$, using theory of quadratic variations, it can be shown (cf. Andersen *et al.* (2001), Barndorff-Nielsen and Shephard (2002a, b)) that:

$$RV_{t+1}(\Delta) \rightarrow \int_0^t \sigma^2(s) ds$$

Realized measures of volatility (con't)

- ▶ Theory thus suggests that optimal sampling corresponds to sampling at the highest possible frequency.
- ▶ This is not true in light of the microstructure effects (bid-ask spread, rounding, non-synchronicity, etc.) which introduce noise in the price process.
- ▶ The observed price is not the “efficient price”.

The volatility signature plot for the realized variance (full sample)




- As frequency becomes higher and higher, the realized variance includes an increasing noise component.

Realized measures of volatility (con't)

- ▶ To mitigate the impact of microstructure noise, a series of methodologies have been employed in the empirical financial literature:
 1. Subsampling schemes as in Zhang *et al.* (2005);
 2. Pre-filtering methods as in Andreou and Ghysels (2002), Podolskij and Vetter (2006) and Jacod *et al.* (2007);
 3. Kernel-based methodologies as in Zhou (1996) or Hansen and Lunde (2006), among many others;
 4. Newey-West correction of the estimator as in Fleming and Paye (2006)
 5. Modeling the noise component as in Aït-Sahalia, Mykland and Zhang (2005)
 6. or methods based on the inspection of the volatility signature plot, what we do

Volatility estimates (realized variance, volatility and log volatility)



- ▶ We sample using a 15-min interval between observations and the last tick method (superior to interpolation).
- ▶ Subsampling and kernel-based estimators are similar in nature.


Descriptive statistics

	Mean	SD	Skewness	Kurtosis	L-B (20)
Naive estimator					
RV_t	0.0130	0.0184	3.2097	16.6500	7.1942
$RV_t^{1/2}$	0.0948	0.0636	1.2998	5.1144	82.886
$\log(RV_t^{1/2})$	-2.5652	0.7369	-0.3279	2.9927	420.63
Zhang <i>et al.</i> (2005) subsampling estimator					
RV_t	0.0085	0.0104	2.9814	15.7588	3.3318
$RV_t^{1/2}$	0.0798	0.0467	1.1008	4.7061	74.181
$\log(RV_t^{1/2})$	-2.6966	0.6551	-0.4657	3.4233	376.51
Bartlett kernel-based estimator					
RV_t	0.0065	0.0079	3.0012	15.4313	2.2043
$RV_t^{1/2}$	0.0702	0.0403	1.1365	4.8489	59.803
$\log(RV_t^{1/2})$	-2.8119	0.6264	-0.3229	3.3850	334.17

The optimal sampling frequency and the maturity effect

- ▶ Testable using the recent test by Awartani *et al.* (2009)
- ▶ For the full sample, it appears that the choice of 15 minutes returns should allow to minimize the impact of the microstructure noise, while ensuring for each day a sufficient number of observations.
- ▶ Somewhat different patterns between the full sample and end-of-year sub-samples but a 15-min interval choice does not appear unreasonable.
- ▶ In addition, results from a volatility signature plot based on two months of data have to be taken with care.


The volatility signature plot for the November-December period



Unconditional distribution of futures returns and realized volatility




- ▶ Following Clark's (1973) seminal contribution on the cotton futures returns, the standardized returns should follow a normal distribution if the process governing the realized volatility is log-normal and the process governing returns is normal.
- ▶ In Clark's vocabulary, the volatility process is the “*directing process*” and the distribution of standardized returns is said to be “*subordinated*” to the distribution of returns.
- ▶ The resulting process is thus a lognormal-normal mixture. It has been extended by Tauchen and Pitts (1983) to give rise to the so-called “*mixture-of-distribution hypothesis*” (MDH)
- ▶ Other transformations may be used [Gonçalves and Meddahi (2008)]

Unconditional distributions of realized variance, volatility and log volatility




- ▶ Realized variance is right-skewed, realized volatility is also right-skewed, log realized volatility seems to approach normality despite standard statistical tests reject this hypothesis.
- ▶ We may accept the Gaussian hypothesis in light of our small sample.

Corresponding theoretical quantile-quantile plots



Unconditional distributions of daily (open-to-close) returns, RV-standardized returns and GARCH-standardized returns



	Mean	SD	Skewness	Kurtosis	Jarque-Bera	$Q(20)$	$Q^2(20)$
Daily returns R_t	0.0000337	0.0296	-0.0472	3.2425	0.6919	75.609	51.66
RV-standardized daily returns	0.0019	0.4984	0.8936	8.8460	381.4887	66.923	152.95
GARCH-standardized daily returns	0.3078	46.31	0.1034	3.4476	2.4622	72.154	19.500

Unconditional distribution of futures returns and realized volatility

- ▶ The distribution of daily returns is quite normal, the distribution of RV-standardized returns is not.
- ▶ The distribution of daily returns standardized with a GARCH estimate is more normally distributed than using a realized volatility.
- ▶ From the ECX emissions futures data, it is clear that the standardized returns are not normally distributed (no need to use advanced tests) both for the realized and GARCH standardisations.


Unconditional distribution of futures returns and realized volatility

- ▶ As in Areal and Taylor (2002), the rejection of the MDH may be due to:
 1. the imperfect estimation of the logarithmic volatility through the realized estimator (BPV would be an alternative)
 2. the extreme outlier occurring the 13th of October, which strongly deforms our distribution.
 3. the overestimated microstructure noise for large returns thus implying a bias in kurtosis (cf. Fleming and Paye, 2005 and Andersen et al., 2007).
- ▶ The rejection of Clark's conjecture for the 2008 ECX futures provide evidence that a latent variable governing the information flow (through volatility) for the European emissions market does not seem to exist.

Conditional distribution of realized volatility

- ▶ We are now interested in the modelling of the dynamic of realized measures.
- ▶ Major practical application is volatility forecasting which is of interest for option pricing, portfolio rebalancing and risk management activities.
- ▶ The main characteristic of volatility is its persistence.

Estimated autocorrelation functions for realized variance, volatility and log volatility



- ▶ The “hyperbolic” decay in the autocorrelation of the log series may characterize long memory.
- ▶ The unit-root hypothesis is rejected.

A taste of long memory

- ▶ We first estimate the fractional integration coefficient
- ▶ First, let S_T be the variance of the sum of T consecutive observations of, say, logarithm of the realized measure $\log(RV_t^{1/2})$. For long memory processes, the variances S_T follow a scaling law such that


$$T^{-(2d+1)} S_T \rightarrow C$$

- ▶ as $T \rightarrow \infty$ with $d > 0$ and C a constant.¹ The regression coefficient corresponds to $2d + 1$, and thus leads to an implicit value of the fractional integration coefficient.

¹In comparison, setting the value $d = 0$ is a feature of short memory.

A taste of long memory (con't)

- ▶ Figure below plots the sample variances S_T of the partial sums of the realized logarithmic standard deviations versus the logarithm of the aggregation level for T



A taste of long memory (con't)

	ADF test	$d(GPH)$	$\hat{d}$ from regression
Naive estimator			
RV_t	-13.9122	0.4376	—
$RV_t^{1/2}$	-11.1715	0.3318	—
$\log(RV_t^{1/2})$	-4.2934	0.6849	0.4634
Zhang <i>et al.</i> (2005) subsampling estimator			
RV_t	-14.6932	0.4399	—
$RV_t^{1/2}$	-11.3561	0.3247	—
$\log(RV_t^{1/2})$	-4.4725	0.6964	0.4588
Bartlett kernel-based estimator			
RV_t	-15.0757	0.4306	—
$RV_t^{1/2}$	-11.8635	0.3066	—
$\log(RV_t^{1/2})$	-3.7696	0.6520	0.4711

The HAR-RV model

- ▶ Our forecasting exercise is based on the Corsi's (2009) HAR model and Mincer-Zarnowitz technique.
- ▶ The HAR model has been used in Liu and Maheu (2009), Andersen *et al.* (2007), and Corsi *et al.* (2008), among others.
- ▶ The economic motivation for this model is that different groups of investors have different investment horizons and consequently behave differently (see Müller *et al.* (1997) for the presentation of the HARCH original model relying on the Heterogeneous Hypothesis).
- ▶ Note that ARFIMA estimation does not appear suitable alternatives for the one-year ECX emissions futures with tick-by-tick data, since the estimation of formal long memory models would require several years of data.

The HAR-RV model (con't)

- ▶ The original HAR-RV model by Corsi (2009) is formally a constrained AR(22). For the realized volatility, it has the form:

$$\sqrt{RV_t} = \alpha_0 + \alpha_d \sqrt{RV_{t-1}} + \alpha_w (\sqrt{RV})_{t-5:t-1} + \alpha_m (\sqrt{RV})_{t-22:t-1} + u_t$$

- ▶ While not formally a long memory model, it seems to be able to reproduce some long memory properties of realized variance time series.

The HAR-RV model for realized variance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
α_0	0.0130 (0.0028)	0.0137 (0.0022)	0.0130 (0.0028)	0.0090 (0.0019)	0.0093 (0.0015)	0.0090 (0.0019)	0.0074 (0.0016)	0.0075 (0.0013)	0.0074 (0.0016)
α_1	0.0810 (0.0746)	0.1013 (0.0645)		0.0323 (0.0741)	0.04683 (0.0648)		0.0139 (0.0739)	0.0211 (0.0649)	
α_2	0.0762 (0.1505)		0.1556 (0.1315)	0.0594 (0.1580)		0.0916 (0.1395)	0.0283 (0.1619)		0.0424 (0.1435)
R^2	0.0109	0.0102	0.0059	0.0026	0.0021	0.0018	0.0005	0.0004	0.0003
Adj. R^2	0.0024	0.0061	0.0017	-	-	-	-	-	-
Log-lik.	484.29	494.19	483.69	564.60	575.98	564.50	595.53	607.52	595.51
AIC	-	-	-	-	-	-	-	-	-
	4.0960	4.1187	4.0995	4.7796	4.8031	4.7873	5.0428	5.0671	5.0511
SC	-	-	-	-	-	-	-	-	-
	4.0519	4.0896	4.0701	4.7354	4.7740	4.7578	4.9986	5.0380	5.0217

The HAR-RV model for realized volatility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
α_0	0.0490 (0.0107)	0.0680 (0.0076)	0.0497 (0.0108)	0.0388 (0.0088)	0.0577 (0.0061)	0.0390 (0.0089)	0.0366 (0.0081)	0.0538 (0.0054)	0.0367 (0.0081)
α_1	0.1904 (0.0770)	0.3118 (0.0615)		0.1522 (0.0773)	0.2960 (0.0619)		0.1181 (0.0775)	0.2554 (0.0627)	
α_2	0.3174 (0.1234)		0.5013 (0.0996)	0.3782 (0.1255)		0.5281 (0.1004)	0.3787 (0.1298)		0.4964 (0.1046)
R^2	0.1211	0.0975	0.0980	0.1207	0.0877	0.1060	0.0971	0.0653	0.0881
Adj. R^2	0.1136	0.0937	0.0941	0.1131	0.0839	0.1022	0.0894	0.0613	0.0842
Log-lik.	290.55	293.69	287.49	359.07	362.22	357.12	382.36	386.22	381.19
AIC	-	-	-	-	-	-	-	-	-
	2.4472	2.4409	2.4297	3.0304	3.0144	3.0223	3.2286	3.2152	3.2271
SC	-	-	-	-	-	-	-	-	-
	2.4030	2.4118	2.4003	2.9862	2.9853	2.9929	3.1844	3.1861	3.1977

The HAR-RV model for the log transform

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
α_0	-	-	-	-	-	-	-	-	-
	0.6329 (0.1850)	1.2134 (0.1479)	0.6452 (0.1884)	0.6663 (0.1893)	1.1683 (0.1493)	0.6866 (0.1960)	0.7345 (0.2039)	1.2884 (0.1581)	0.7512 (0.2101)
α_d	0.2480 (0.0789)	0.5275 (0.0549)		0.3299 (0.0776)	0.5678 (0.0534)		0.3058 (0.0781)	0.542864 (0.0545)	
α_w	0.5041 (0.1043)		0.7473 (0.0713)	0.4226 (0.1022)		0.7449 (0.0711)	0.4325 (0.1045)		0.7322 (0.0733)
R^2	0.3477	0.2798	0.3200	0.3691	0.3226	0.3200	0.3429	0.2946	0.2995
Adj. R^2	0.3421	0.2768	0.3171	0.3637	0.3197	0.3171	0.3373	0.2916	0.2965
Log-lik.	-	-	-	-	-	-	-	-	-
	220.98	228.18	225.88	190.05	200.23	198.87	184.02	194.03	191.54
AIC	1.9062	1.9851	1.9394	1.6430	1.6923	1.7095	1.5917	1.6404	1.6471
SC	1.9504	2.014230	1.9688	1.6872	1.7214	1.7389	1.6358	1.6695	1.6766

The HAR-RV model

- ▶ The log-likelihood for the log transformation is the best for the model with the daily and the week components.
- ▶ In view of the R^2 , the explanatory power has no common measure between realized variance (around 1%), realized volatility (around 10%) and the log series (around 35%).
- ▶ This is similar to the findings of ABDL (2003) for FX rates which use a VAR(5) and finally an AR(5) for the modeling of dynamic.

A modest forecasting exercise


- ▶ Our sample is very short (as in Taylor and Xu (1997) for instance) and precludes from any long memory (ARFIMA) tentative
- ▶ Forecasting accuracy is evaluated using one-day-ahead out-of-sample (100 observations are used for estimate each model).
- ▶ Mincer-Zarnowitz (encompassing) regressions for our purpose is as follows:

$$(v_{t+1})^{1/2} = b_0 + b_1(v_{t+1|t,HAR-RV})^{1/2} + b_2(v_{t+1|t,GARCH})^{1/2} + u_{t+1}$$

Forecasting: main result

	b_0	b_1	b_2	$R - squ.$
Daily realized variance (RV_t)				
HAR-RV	0.006327 (0.02075)	0.5678 (1.3777)		0.0011
GARCH daily	0.01301 (0.00434)		1970.17 (3120.91)	0.0028
HAR-RV + GARCH daily	0.00699 (0.0208)	0.4156 (1.4074)	1788.50 (3190.7)	0.0033
Daily realized volatility in standard deviation form ($RV_t^{1/2}$)				
HAR-RV	-0.00654 (0.0240)	1.0408*** (0.2419)		0.1139
GARCH daily	0.05527 (0.0130)		45.8000*** (13.879)	0.0703
HAR-RV + GARCH daily	-0.0069 (0.0238)	0.8526*** (0.2735)	23.403 (15.322)	0.1281
Daily realized volatility in logarithmic form ($\log(RV_t^{1/2})$)				
HAR-RV	0.1479 (0.2517)	1.0656*** (0.0942)		0.4704
GARCH daily	2.3640*** (0.8599)		0.6945*** (0.1188)	0.1917
HAR-RV + GARCH daily	1.2419* (0.7032)	0.9724*** (0.1090)	0.1854* (0.1113)	0.4800

Graphical results: actual vs. one step ahead forecast for the naive estimator



- ▶ The model used to make prediction is the HAR with a daily and a weekly component.
- ▶ Forecast accuracy is improved for the logarithmic transformation.

Forecasting: summary of results

- ▶ The forecast power of the HAR-RV model is about 50% for the log-series.
- ▶ Such a high level of explanatory power is representative of the strong persistence of volatility highlighted in the last section for our series and in the literature for all other financial assets.

Some concluding remarks: our main results

- i)* Realized volatility is quite lognormal ($\log R_v$ is quite normal) $\rightarrow$ implications for stochastic volatility modelling and then for derivatives pricing
- ii)* Standardized returns are not normally distributed $\rightarrow$ rejection of Clark's conjecture
- iii)* A simple HAR-RV model does a better job in volatility modelling than any GARCH model
- iv)* This is confirmed by forecasting accuracy

Some possible extensions

- ▶ More robustness (distribution, dynamics) will be achieved using more and more data
- ▶ Our paper does not consider the issue of “jumps” in the time series; jumps may explain some anomalies about the distribution.
- ▶ Our results suggest to resort to more complex model than the standard diffusive process for option pricing.
- ▶ What about the maturity effect in futures markets? Not explored while of importance (jump, noise).