

The fine structure of volatility feedback II: overnight and intra-day effects

Pierre Blanc^{1,2*} Rémy Chicheportiche^{1,3}
Jean-Philippe Bouchaud¹

¹ Capital Fund Management, 75 007 Paris, France

² Université Paris-Est, CERMICS, Project team MathRISK ENPC-INRIA-UMLV
École des Ponts, 77 455 Marne-la-vallée, France

³ Chaire de finance quantitative, École Centrale Paris, 92 295 Châtenay-Malabry, France

September 24, 2013

Abstract

We decompose, within an ARCH framework, the daily volatility of stocks into overnight and intraday contributions. We find, as perhaps expected, that the overnight and intraday returns behave completely differently. For example, while past intraday returns affect equally the future intraday and overnight volatilities, past overnight returns have a weak effect on future intraday volatilities (except for the very next one) but impact substantially future overnight volatilities. The exogenous component of overnight volatilities is found to be close to zero, which means that the lion's share of overnight volatility comes from feedback effects. The residual kurtosis of returns is small for intraday returns but infinite for overnight returns. We provide a plausible interpretation for these findings, and show that our IntraDay/Overnight model significantly outperforms the standard ARCH framework based on daily returns for Out-of-Sample predictions.

*Corresponding author: blancp@cermics.enpc.fr. P. B. realized most of this work during an internship at Capital Fund Management in 2012. He is currently a PhD student at CERMICS (Université Paris-Est) and benefits from a half-scholarship from Fondation Natixis.

1 Introduction

The ARCH (auto-regressive conditional heteroskedastic) framework was introduced in [8] to account for volatility clustering in financial markets and other economic time series. It posits that the current relative price change r_t can be written as the product of a “volatility” component σ_t and a certain random variable ξ_t , of zero mean and unit variance, and that the dynamics of the volatility is self-referential in the sense that it depends on the past returns themselves as:

$$\sigma_t^2 = s^2 + \sum_{\tau=1}^q K(\tau) r_{t-\tau}^2 \equiv s^2 + \sum_{\tau=1}^q K(\tau) \sigma_{t-\tau}^2 \xi_{t-\tau}^2, \quad (1)$$

where s^2 is the “baseline” volatility level, that would obtain in the absence of any feedback from the past, and $K(\tau)$ is a kernel that encodes the strength of the influence of past returns. The model is well defined and leads to a stationary time series whenever the feedback is not too strong, i.e. when $\sum_{\tau=1}^q K(\tau) < 1$. A very popular choice, still very much used both in the academic and professional literature, is the so called “GARCH” (Generalized ARCH), that corresponds to an exponential kernel, $K(\tau) = g e^{-\tau/\tau_p}$, with $q \rightarrow \infty$. However, the long-range memory nature of the volatility correlations in financial markets suggests that a power-law kernel is more plausible — a model called “FIGARCH”, see [4, 2] and below.

Now, the ARCH framework implicitly singles out a *time scale*, namely the time interval over which the returns r_t are defined. For financial applications, this time scale is often chosen to be one day, i.e. the ARCH model is a model for daily returns, defined for example as the relative variation of price between two successive closing prices. However, this choice of a day as the unit of time is often a default imposed by the data itself. A natural question is to know whether other time scales could also play a role in the volatility feedback mechanism. In our companion paper [6], we have studied this question in detail, focussing on time scales *larger than* (or equal to) the day. We have in fact calibrated the most general model, called “QARCH” [16], that expresses the squared volatility as a quadratic form of past returns. This encompasses all models where returns defined over arbitrary time intervals could play a role, as well as (realized) correlations between those — see [6] for more precise statements, and [14, 19, 5] for earlier contributions along these lines.

The main conclusion of our companion paper [6] is that while other time scales play a statistically significant role in the feedback process, the *dominant effect* for daily returns is indeed associated with past daily returns. In a first approximation, a FIGARCH model based on daily returns, with an exponentially truncated power-law kernel $K(\tau) = g \tau^{-\alpha} e^{-\tau/\tau_p}$, provides a good model for stock returns, with $\alpha \approx 1.1$ and $\tau_p \approx 50$ days. This immediately begs the question: if returns on time scales larger than a day appear to be of lesser importance,¹ what about returns on time scales smaller than a day? For one thing, a trading day is naturally decomposed into trading hours, that define an ‘Open to Close’ (or ‘Intraday’) return, and hours where the market is closed but news accumulates and impacts the price at the opening auction, contributing to the ‘Close to Open’ (or ‘Overnight’) return. One may expect that the price dynamics is very different in the two cases, for several reasons. One is that many company announcements are made overnight, that can significantly impact the price. The profile of market participants is also quite different in the two cases: while low-frequency participants might choose to execute large volumes during the auction, higher-frequency participants and market-makers are mostly active intraday. In any case, it seems reasonable to distinguish two volatility contributions, one coming from Intraday trading, the second one from Overnight activity. Similarly, the feedback of past returns should also be disentangled into an Intraday contribution and an Overnight contribution. The calibration of an ARCH-like model that distinguishes between Intraday and Overnight returns is the aim of the present paper, and is the content of Section 2. We have in fact investigated the role of higher frequency returns as well. For the sake of clarity we will not present this study here, but rather summarize briefly our findings on this point in the conclusion.

The salient conclusions of the present paper are that the Intraday and Overnight dynamics indeed behave completely differently — for example, while the Intraday (Open-to-Close) returns

¹Note a possible source of confusion here since a FIGARCH model obviously involves many time scales. We need to clearly distinguish time *lags*, as they appear in the kernel $K(\tau)$, from time scales, that enter in the definition of the returns themselves.

impact both the future Intraday and Overnight volatilities in a slowly decaying manner, Overnight (Close-to-Open) returns essentially impact the next Intraday but very little the following ones. However, Overnight returns have themselves a slowly decaying impact on future Overnights. Another notable difference is the statistics of the residual factor ξ_t , which is *nearly Gaussian* for Intraday returns, but has an *infinite kurtosis* for Overnight returns. We discuss further the scope of our results in the conclusion Section 4, and relegate several more technical details to appendices.

2 The dynamics of Close-to-Open and Open-to-Close stock returns and volatilities

Although the decomposition of the daily (Close-to-Close) returns into their Intraday and Overnight components seems obvious and intuitive, very few attempts have actually been made to model them jointly (see [10, 17]). In fact, some studies even discard overnight returns altogether. In the present section, we define and calibrate an ARCH model that explicitly treats these two contributions separately. We however first need to introduce some precise definitions of the objects that we want to model.

2.1 Definitions, time-line and basic statistics

We consider equidistant time stamps t with $\Delta t = 1$ day. Every day, the prices of traded stocks are quoted from the opening to the closing hour, but we only keep track of the first and last traded prices. For every stock name a , O_t^a is the open price and C_t^a the close price at date t . (In the following, we drop the index a when it is not explicitly needed). We introduce the following definitions of the geometric returns, volatilities, and residuals:

$$\text{Intraday return: } r_t^D = \ln(C_t/O_t) \quad \equiv \sigma_t^D \xi_t^D \quad (2a)$$

$$\text{Overnight return: } r_t^N = \ln(O_t/C_{t-1}) \quad \equiv \sigma_t^N \xi_t^N \quad (2b)$$

$$\text{Daily return: } r_t = \ln(C_t/C_{t-1}) = r_t^D + r_t^N \quad \equiv \sigma_t \xi_t. \quad (2c)$$

The following time-line illustrates the definition of the three types of return:

$$\dots \longrightarrow C_{t-1} \quad \left| \quad \begin{array}{c} \xrightarrow{\text{Night } t} \quad O_t \quad \xrightarrow{\text{Day } t} \quad C_t \\ \underbrace{\hspace{10em}}_{r_t} \\ \underbrace{\hspace{4em}}_{r_t^N} \quad \underbrace{\hspace{4em}}_{r_t^D} \end{array} \quad \left| \quad \begin{array}{c} \xrightarrow{\text{Night } t+1} \quad O_{t+1} \longrightarrow \dots \\ \underbrace{\hspace{4em}}_{r_{t+1}^N} \end{array} \quad (3)$$

To facilitate the reading of our tables and figures, intraday returns are associated with the green color (or light gray) and overnight returns with blue (or dark gray).

Before introducing any model, we discuss the qualitative statistical differences in the series of Open-Close returns r_t^D and Close-Open returns r_t^N . First, one can look at Fig. 1 for a visual impression of the difference: while the intraday volatility is higher than the overnight volatility, the relative importance of “surprises” (i.e. large positive or negative jumps) is larger for overnight returns. This is confirmed by the numerical values provided in Tab. 1 for the volatility, skewness and kurtosis of the two time series r_t^D and r_t^N .

J	$[\langle r^J \rangle]$	$\sigma^J = \sqrt{[\langle r^{J2} \rangle]}$	$[\langle r^{J3} \rangle / \langle r^{J2} \rangle^{\frac{3}{2}}]$	$[\langle r^{J4} \rangle / \langle r^{J2} \rangle^2]$
D	$1.2 \cdot 10^{-4}$	0.022	-0.12	12.9
N	$-1.0 \cdot 10^{-4}$	0.013	-1.5	62.6

Table 1: Distributional properties of intraDay and overNight returns (first four empirical moments). $\langle \cdot \rangle$ means average over all dates, and $[\cdot]$ average over all stocks.




Figure 1: Example of a historical time serie of stock day returns (top) and overnight returns (bottom).

It is also visible on Fig. 1 that periods of high volatilities are common to the two series: two minor ones can be observed in the middle of year 2000 and at the beginning of year 2009, and an important one in the middle of year 2002.

An important quantity is the correlation between intraday and overnight returns, which can be measured either as $\langle r_t^N r_t^D \rangle / \sigma^N \sigma^D$ (overnight leading intraday) or as $\langle r_t^D r_{t+1}^N \rangle / \sigma^N \sigma^D$ (intraday leading overnight). The statistical reversion revealed by the measured values of the above correlation coefficients (-0.021 and -0.009 , respectively) is slight enough (compared to the amplitude and reach of the feedback effect) to justify the assumption of i.i.d. residuals. If there were no linear correlations between Intraday and Overnight returns, the squared volatilities would be exactly additive, i.e. $\sigma_t^2 \equiv \sigma_t^{D^2} + \sigma_t^{N^2}$. Deviations from this simple addition of variance rule are below 2%.

2.2 The model

The standard ARCH model recalled in the introduction, Eq. (1), can be rewritten identically as:

$$\sigma_t^2 \equiv s^2 + \sum_{\tau=1}^q K(\tau) r_{t-\tau}^{N^2} + \sum_{\tau=1}^q K(\tau) r_{t-\tau}^{D^2} + 2 \sum_{\tau=1}^q K(\tau) r_{t-\tau}^N r_{t-\tau}^D, \quad (4)$$

meaning that there is a unique kernel $K(\tau)$ describing the feedback of past intraday and overnight returns on the current volatility level.

If however one believes that these returns are of fundamentally different nature, one should expand the model in two directions: first, the two volatilities σ^{D^2} and σ^{N^2} should have separate dynamics. Second, the kernels describing the feedback of past intraday and overnight returns should *a priori* be different. This suggests to write the following generalized model for the intraDay volatility:

$$\begin{aligned} \sigma_t^{D^2} = & s^{D^2} + \sum_{\tau=1}^{\infty} L^{D \rightarrow D}(\tau) r_{t-\tau}^{D^2} + \sum_{\tau=1}^{\infty} K^{DD \rightarrow D}(\tau) r_{t-\tau}^{D^2} + 2 \sum_{\tau=1}^{\infty} K^{ND \rightarrow D}(\tau) r_{t-\tau}^D r_{t-\tau}^N \\ & + \sum_{\tau=0}^{\infty} L^{N \rightarrow D}(\tau+1) r_{t-\tau}^N + \sum_{\tau=0}^{\infty} K^{NN \rightarrow D}(\tau+1) r_{t-\tau}^{N^2} + 2 \sum_{\tau=0}^{\infty} K^{DN \rightarrow D}(\tau+1) r_{t-\tau-1}^D r_{t-\tau}^N, \end{aligned} \quad (5)$$

where we have added the possibility of a “leverage effect”, i.e. terms linear in past returns that can describe an asymmetry in the impact of negative and positive returns on the volatility. The notation used is, we hope, explicit: for example $K^{\text{DD} \rightarrow \text{D}}$ describes the influence of squared intraDay past returns on the current intraDay volatility. Note that the mixed effect of intraDay and overNight returns requires two distinct kernels, $K^{\text{DN} \rightarrow \text{D}}$ and $K^{\text{ND} \rightarrow \text{D}}$, depending on which comes first in time. Finally, the time-line shown above explains why the τ index starts at $\tau = 1$ for past intraDay returns, but at $\tau = 0$ for past overNight returns. We posit a similar expression for the overNight volatility:

$$\begin{aligned} \sigma_t^{\text{N}^2} = s^{\text{N}^2} &+ \sum_{\tau=1}^{\infty} L^{\text{N} \rightarrow \text{N}}(\tau) r_{t-\tau}^{\text{N}} + \sum_{\tau=1}^{\infty} K^{\text{NN} \rightarrow \text{N}}(\tau) r_{t-\tau}^{\text{N}}{}^2 + 2 \sum_{\tau=1}^{\infty} K^{\text{ND} \rightarrow \text{N}}(\tau) r_{t-\tau}^{\text{D}} r_{t-\tau}^{\text{N}} \\ &+ \sum_{\tau=1}^{\infty} L^{\text{D} \rightarrow \text{N}}(\tau) r_{t-\tau}^{\text{D}} + \sum_{\tau=1}^{\infty} K^{\text{DD} \rightarrow \text{N}}(\tau) r_{t-\tau}^{\text{D}}{}^2 + 2 \sum_{\tau=1}^{\infty} K^{\text{DN} \rightarrow \text{N}}(\tau) r_{t-\tau-1}^{\text{D}} r_{t-\tau}^{\text{N}}. \end{aligned} \quad (6)$$

The model is therefore fully characterized by two base-line volatilities $s^{\text{D}}, s^{\text{N}}$, four leverage (linear) kernels $L^{\text{J} \rightarrow \text{J}}$, eight quadratic kernels $K^{\text{J}^{\text{J}} \rightarrow \text{J}}$, and the statistics of the two residual noises $\xi^{\text{D}}, \xi^{\text{N}}$ needed to define the returns, as $r^{\text{J}} = \sigma^{\text{J}} \xi^{\text{J}}$. We derive in Appendix A conditions on the coefficients of the model under which the two volatility processes remain positive at all times. The model as it stands has a large number of parameters; in order to ease the calibration process and gain in stability, we in fact choose to parameterize the τ dependence of the different kernels with some simple functions, namely an exponentially truncated power-law for $K^{\text{J}^{\text{J}} \rightarrow \text{J}}$ and a simple exponential for $L^{\text{J} \rightarrow \text{J}}$:

$$K(\tau) = g_{\text{p}} \tau^{-\alpha} \exp(-\omega_{\text{p}} \tau); \quad L(\tau) = g_{\text{e}} \exp(-\omega_{\text{e}} \tau). \quad (7)$$

The choice of these functions is not arbitrary, but is suggested by a preliminary calibration of the model using a generalized method of moments (GMM), as explained in the companion paper, see Appendix C.2 in Ref. [6].

As far as the residuals $\xi_t^{\text{D}}, \xi_t^{\text{N}}$ are concerned, we assume them to be i.i.d. centered Student variables of unit variance with respectively $\nu^{\text{D}} > 2$ and $\nu^{\text{N}} > 2$ degrees of freedom. Contrarily to many previous studies, we prefer to be agnostic about the kurtosis of the residuals rather than imposing a priori Gaussian residuals. It has been shown that while the ARCH feedback effect accounts for volatility clustering and for some positive kurtosis in the returns, this effect alone is not sufficient to explain the observed heavy tails in the return distribution (see for example [6]). These tails come from true ‘surprises’ (often called jumps), that cannot be anticipated by the predictable part of the volatility, and that can indeed be described by a Student (power-law) distribution of the residuals.

2.3 Dataset

The dataset used to calibrate the model is exactly the same as in our companion paper [6]. It is composed of US stock prices (four points every day: Open, Close, High and Low) for $N = 280$ stocks present permanently in the S&P-500 index from Jan. 2000 to Dec. 2009 ($T = 2515$ days). For every stock a , the daily returns (r_t^a), day returns ($r_t^{\text{D}^a}$) and overnight returns ($r_t^{\text{N}^a}$) defined in Eq. (2) are computed using only Open and Close prices at every date t . In order to improve the statistical significance of our results, we consider the pool of stocks as a statistical ensemble over which we can average. This means that we assume a universal dynamics for the stocks, a reasonable assumption as we discuss in Appendix B.

Bare stock returns are “polluted” by several obvious and predictable events associated with the life of the company, such as quarterly announcements. They also reflect low-frequency human activity, such as a weekly cyclical pattern of the volatility, which is interesting in itself (see Fig. 2). These are of course real effects, but the ARCH family of models we investigate here rather focuses on the endogenous self-dynamics *on top of* such seasonal patterns. For example, the quarterly announcement dates are responsible for returns of typically much larger magnitude (approximately three times larger on average for daily returns) that have a very limited feedback in future volatility.

We therefore want to remove all obvious seasonal effects from the dataset, and go through the following additional steps of data treatment before estimating the model. For every stock a , the

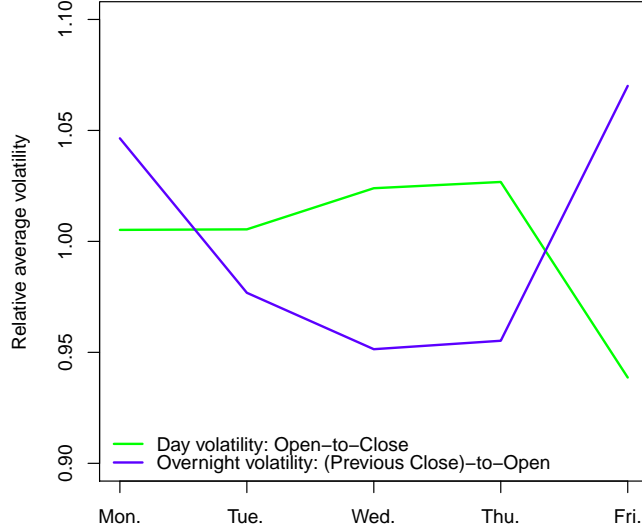


Figure 2: Normalized weekly seasonality of the volatility. The overnight volatility is that of the previous night (i.e. the volatility of the weekend for Monday and that of Thursday night for Friday).

average over time is denoted $\langle r^a \rangle := \frac{1}{T} \sum_t r_t^a$, and for each date t , the cross-sectional average over stocks is $[r_t^a] := \frac{1}{N} \sum_a r_t^a$. All the following normalizations apply both (and separately) for day returns and overnight returns.

- The returns series are first centered around their temporal average: $r_t^a \leftarrow r_t^a - \langle r^a \rangle$. In fact, the returns are already nearly empirically centered, since the temporal average is small, see Tab. 1 above.
- We then divide the returns by the cross-sectional dispersion:

$$r_t^a \leftarrow r_t^a / \sqrt{[r_t^{\neq a^2}]}$$

This normalization² removes the historical low-frequency patterns of the volatility, for example the weekly pattern discussed above (Fig. 2). In order to predict the “real” volatility with the model, one must however re-integrate these patterns back into the σ^j ’s.

- Finally, we normalize stock by stock all the returns by their historical standard deviation: for all stock a , for all t ,

$$r_t^a \leftarrow r_t^a / \sqrt{\langle r^{\cdot a^2} \rangle},$$

imposing $\langle r^{\cdot a^2} \rangle = \langle r^{\cdot N a^2} \rangle \equiv 1$.

This data treatment allows to consider that the residual volatility of the returns series is independent of the effects we do not aim at modelling, and that the series of all stocks can reasonably be assumed to be homogeneous (i.e. identically distributed), both across stocks and across time. This is necessary in order to calibrate a model that is translational-invariant in time (i.e. only the time lag τ enters Eqs. (5,6)), and also to enlarge the calibration dataset by averaging the results over all the stocks in the pool — see the discussion in Appendix B.

² If the element a is not excluded in the average, the tails of the returns are artificially cut-off: when $|r_t^a| \rightarrow \infty$, $|r_t^a| / \sqrt{[r_t^{\neq a^2}]}$ is capped at $\sqrt{N} < \infty$.

2.4 Model estimation

Assuming that the distribution of the residuals is a Student law, the log-likelihood per point of the model can be written as ($J = D, N$):

$$\mathcal{L}(\Theta^J, \nu^J | \{r^D, r^N\}) = \frac{\nu^J}{2} \ln |(\nu^J - 2) \sigma_t^2(\Theta^J)| - \frac{\nu^J + 1}{2} \ln |(\nu^J - 2) (\sigma_t^2(\Theta^J) + r_t^{J2})|, \quad (8)$$

where $\sigma_t^{J2} = \sigma_t^2(\Theta^J | \{r^D, r^N\})$ are defined in Eqs. (5,6), ν^J are the degrees of freedom of the Student residuals, and Θ^J denote generically the sets of volatility feedback parameters.

Conditionally on the dataset, we maximize numerically the likelihood of the model, averaged over all dates and all stocks.

Calibration methodology: As mentioned above, we in fact choose to parameterize the feedback kernels as suggested by the results of the method of moments, i.e. exponentially truncated power-laws for K 's and simple exponentials for L 's. Imposing these simple functional forms allows us to gain stability and readability of the results. However, the functional dependence of the likelihood on the kernel parameters is not guaranteed to be globally concave, as is the case when it is maximized "freely", i.e. with respect to all individual kernel coefficients $K(\tau)$ and $L(\tau)$, with $1 \leq \tau \leq q_{\text{free}}$. This is why we use a three-step approach:

1. A first set of kernel estimates is obtained by the Generalized Method of Moments (GMM), see [6], and serves as a starting point for the optimization algorithm.
2. We then run a Maximum Likelihood Estimation (MLE) of the model-free kernels based on Eq. (8), over $6 \times q_{\text{free}}$ parameters for both Θ^D and Θ^N , with a moderate value of maximum lag $q_{\text{free}} = 63 \simeq$ three months. Taking as a starting point the coefficients of step 1 and maximizing with a gradient descent, we obtain a second set of (short-range) kernels.
3. Finally, we perform a MLE estimation of the parametrically constrained kernels with the functional forms (7) for K 's and L 's, which only involves $4 \times 3 + 2 \times 2$ parameters in every set Θ^D and Θ^N , with now a large value of the maximum lag $q_{\text{constr}} = 512 \simeq$ two years. Taking as a starting point the functional fits of the kernels obtained at step 2, and maximizing with a gradient descent, we obtain our final set of model coefficients, shown in Tabs. 2,3.

Thanks to step 2, the starting point of step 3 is close enough to the global maximum for the likelihood to be locally concave, and the gradient descent algorithm converges in a few steps. The Hessian matrix of the likelihood is evaluated at the maximum to check that the dependency on all coefficients is indeed concave.

The numerical maximization of the likelihood is thus made on 2 or 3 parameters per kernel, independently of the chosen maximum lag q_{constr} , that can thus be arbitrarily large with little additional computation cost.

Finally, the degrees of freedom ν of the Student residuals are determined using two separate one-dimensional likelihood maximizations (one for $q = q_{\text{free}}$ and one for $q = q_{\text{constr}}$) and then included as an additional parameter in the MLE of the third step of the calibration. Note that ν does not vary significantly at the third step, which means that the estimation of the volatility parameters at the second step can indeed be done independently from that of ν .

This somewhat sophisticated calibration method was tested on simulated data, obtaining very good results.

The special case $s^2 = 0$: We ran the above calibration protocole on intra-day and overnight volatilities separately.

For the overnight model, this led to a slightly negative value of the baseline volatility s^{N2} (statistically compatible with zero). But of course negative values of s^2 are excluded for a stable and positive volatility process. For overnight volatility only, we thus add a step to the calibration protocole, which includes the constraint $s^{N2} = 0$ in the estimation of $K^{DD \rightarrow N}$ and $K^{NN \rightarrow N}$ (which are the two main contributors to the value of the baseline volatility). For simplicity, we consider

here that $\langle \sigma^{N^2} \rangle = \langle r^{D^2} \rangle = \langle r^{N^2} \rangle = 1$. We take the results of the preceding calibration as a starting point and freeze all the kernels but $K^{DD \rightarrow N}$ and $K^{NN \rightarrow N}$, expressed (in this section only) as follows:

$$K^{DD \rightarrow N}(\tau) = g \tau^{-\alpha_1} \exp(-\omega_1 \tau), \quad K^{NN \rightarrow N}(\tau) = \gamma g \tau^{-\alpha_2} \exp(-\omega_2 \tau), \quad (9)$$

where $g = g(\gamma, \alpha_1, \omega_1, \alpha_2, \omega_2)$ is fixed by the constraint $s^{N^2} = 0$:

$$g(\gamma, \alpha_1, \omega_1, \alpha_2, \omega_2) = \frac{1 - c}{h(\alpha_1, \omega_1) + \gamma h(\alpha_2, \omega_2)}; \quad h(\alpha, \omega) = \sum_{\tau=1}^q \tau^{-\alpha} \exp(-\omega \tau), \quad (10)$$

with $\gamma > 0$ the ratio of the two initial amplitudes, and c the (low) contribution of the fixed ‘cross’ kernels $K^{ND \rightarrow N}$ and $K^{DN \rightarrow N}$ to s^{N^2} . We then maximize the likelihood of the model with respect to the five parameters $\gamma, \alpha_1, \omega_1, \alpha_2$ and ω_2 , for which a gradient vector and a Hessian matrix of dimension 5 can be deduced from equations (9) and (10). The coefficients and confidence intervals of the kernels $K^{DD \rightarrow N}$ and $K^{NN \rightarrow N}$ are replaced in Sect. 3.1 by the results of this final step, along with the corresponding value of the overnight baseline volatility, $s^{N^2} = 0$ in Sect. 3.3.

For intra-day volatility instead, the results are given just below, in Sect. 3.1.

3 Intraday vs. overnight: results and discussions

The calibration of our generalized ARCH framework should determine three families of parameters: the feedback kernels K and L , the statistics of the residuals ξ and finally the ‘‘baseline volatilities’’ s^2 . We discuss these three families in turn in the following sections.

3.1 The feedback kernels


In this section, we give the results of the ML estimation of the regression kernels for a maximum lag $q = 512$: estimates of the parameters are reported in Tabs. 2,3, and the resulting kernels are shown in Fig. 3.

Kernels	$K(\tau) = g_p \tau^{-\alpha} e^{-\omega_p \tau}$			$L(\tau) = g_e e^{-\omega_e \tau}$	
	$g_p \times 10^2$	α	$\omega_p \times 10^2$	$g_e \times 10^2$	$\omega_e \times 10^2$
$K^{DD \rightarrow D}$	7.99 ± 0.06	0.71 ± 0.003	0.64 ± 0.02	---	---
$K^{NN \rightarrow D}$	6.53 ± 0.22	2.30 ± 0.07	0.04 ± 0.97	---	---
$K^{ND \rightarrow D}$	1.52 ± 0.17	1.03 ± 0.11	1.3 ± 1.2	---	---
$K^{DN \rightarrow D}$	1.35 ± 0.22	1.03 ± 0.17	3.0 ± 4.6	---	---
$L^{D \rightarrow D}$	---	---	---	-4.97 ± 0.25	18.3 ± 1.3
$L^{N \rightarrow D}$	---	---	---	-2.83 ± 0.30	22.3 ± 2.5


Table 2: Day volatility: estimated kernel parameters for K ’s and L ’s, with their asymptotic confidence intervals of level 95%, as computed using the Fisher Information matrix with the Gaussian quantile (1.98).

Kernels	$K(\tau) = g_p \tau^{-\alpha} e^{-\omega_p \tau}$			$L(\tau) = g_e e^{-\omega_e \tau}$	
	$g_p \times 10^2$	α	$\omega_p \times 10^2$	$g_e \times 10^2$	$\omega_e \times 10^2$
$K^{DD \rightarrow N}$	6.59 ± 0.33	0.80 ± 0.02	1.4 ± 0.4	---	---
$K^{NN \rightarrow N}$	3.64 ± 0.17	0.58 ± 0.01	0.58 ± 0.04	---	---
$K^{ND \rightarrow N}$	1.39 ± 0.11	0.74 ± 0.03	0.42 ± 0.12	---	---
$K^{DN \rightarrow N}$	-1.00 ± 0.29	4.22 ± 2.44	0.02 ± 23	---	---
$L^{D \rightarrow N}$	---	---	---	-2.09 ± 0.05	5.5 ± 0.2
$L^{N \rightarrow N}$	---	---	---	-2.03 ± 0.20	13.1 ± 2.2

Table 3: Overnight volatility: estimated kernel parameters for K ’s and L ’s, with their asymptotic confidence intervals of level 95%, as computed using the Fisher Information matrix with the Gaussian quantile (1.98).



(a) Quadratic $K(\tau)$ kernels in log-log scale. For 'DN \rightarrow N', the absolute value of the (negative) kernel is plotted.



(b) Linear $L(\tau)$ kernels in lin-log scale. All leverage kernels are negative, so $-L(\tau)$ are plotted.

Figure 3: Estimated kernels, impacting IntraDays in (light) green, OverNights in (dark) blue.

We define the exponential characteristic times $\tau_p := 1/\omega_p$ and $\tau_e := 1/\omega_e$, for which qualitative interpretation is easier than for ω_p and ω_e . In the case of the quadratic kernels (of type K), τ_p represents the lag where the exponential cut-off appears, after which the kernel decays to zero quickly. One should note that in three cases, we have $\omega_p - \Delta\omega_p < 0$. These correspond to kernels with $\alpha > 1$, which means that the power-law decays quickly by itself. In these cases the identifiability of ω_p is more difficult and cut-off times are ill-determined, since the value of ω_p only matters in a region where the kernels are already small. The exponential term $\exp(-\omega_p \tau)$ could be removed from the functional form of equation (7), for these kernels only (the calibration would then modify very slightly the value of the power-law exponent α).

Intraday volatility: From Tab. 2 and Fig. 3(a), we see that all intraday quadratic kernels are positive. However, a clear distinction is observed between intraday feedback and overnight feedback: while the former is strong and decays slowly ($\alpha = 0.71$ and $\tau_p = 157$ days), the latter decays extremely steeply ($\alpha = 2.30$) and is quickly negligible, except for the intraday immediately following the overnight, where the effect is as strong as that of the previous intraday. The cross kernels (${}_{ND} \rightarrow {}_D$ or ${}_{DN} \rightarrow {}_D$) are both statistically significant, but are clearly smaller, and decay faster, than the ${}_{DD} \rightarrow {}_D$ effect.

As far as the leverage effect is concerned, both L 's are found to be negative, as expected, and of similar decay time: $\tau_e \approx 5$ days (one week). However, the amplitude for their immediate impact is two times smaller for past overnight returns: $L^{N \rightarrow D} = -0.0283$ versus $L^{D \rightarrow D} = -0.0497$.

In summary, the most important part of the feedback effect on the intraday component of the volatility comes from the past intradays themselves, except for the very last overnight, which also has a strong impact — as intuitively expected, a large return overnight leads to a large immediate reaction of the market as trading resumes. However, this influence is seen to decay very quickly with time. Since a large fraction of company specific news release happen overnight, it is tempting to think that large overnight returns are mostly due to news. Our present finding would then be in line with the general result reported in [12]: volatility tends to relax much faster after a news-induced jump than after endogenous jumps.

Overnight volatility: In the case of overnight volatility, Tab. 3 and Fig. 3(a) illustrate that the influence of past intradays and past overnights is similar: $K^{DD \rightarrow N}(\tau) \approx K^{NN \rightarrow N}(\tau)$, in particular when both are large. The cross kernels now behave quite differently: whereas the behavior of $K^{ND \rightarrow N}$ is not very different from that of $K^{DD \rightarrow N}$ or $K^{NN \rightarrow N}$ (although its initial amplitude is four times smaller), $K^{DN \rightarrow N}(\tau)$ is negative and small, but is hardly significant for $\tau \geq 2$. Interestingly, as pointed out above, the equality $K^{ND \rightarrow N} = K^{DD \rightarrow N} = K^{NN \rightarrow N}$ means that it is the full Close-to-Close return that is involved in the feedback mechanism on the next overnight. What we find here is that this equality very roughly holds, suggesting that, as postulated in standard ARCH approaches, the daily close to close return is the fundamental object that feedbacks on future volatilities. However, this is only approximately true for the *overnight volatility*, while the intraday volatility behaves very differently (as already said, for intraday returns, the largest part of the feedback mechanism comes from past intradays only, and the very last overnight).

Finally, the leverage kernels behave very much like for the intraday volatility. In fact, the $N \rightarrow N$ leverage kernel is very similar to its $N \rightarrow D$ counterpart, whereas the decay of the $D \rightarrow N$ kernel is slower ($\tau_e \approx 18$ days, nearly one month).

Stability and positivity: We checked that these empirically-determined kernels are compatible with a stable and positive volatility process. The first obvious condition is that the system is stable with positive baseline volatilities s^{j^2} . The self-consistent equations for the average volatilities read: (neglecting small cross correlations):

$$\langle \sigma^{D^2} \rangle = s^{D^2} + \langle \sigma^{D^2} \rangle \sum_{\tau=1}^{\infty} K^{DD \rightarrow D}(\tau) + \langle \sigma^{N^2} \rangle \sum_{\tau=1}^{\infty} K^{NN \rightarrow D}(\tau), \quad (11)$$

$$\langle \sigma^{N^2} \rangle = s^{N^2} + \langle \sigma^{D^2} \rangle \sum_{\tau=1}^{\infty} K^{DD \rightarrow N}(\tau) + \langle \sigma^{N^2} \rangle \sum_{\tau=1}^{\infty} K^{NN \rightarrow N}(\tau). \quad (12)$$




Figure 4: CDF $\mathbb{P}(|\xi| > y)$ of the residuals ξ^D and ξ^N , in log-log scale.

This requires that the two eigenvalues of the 2×2 matrix of the corresponding linear system are less than unity, i.e.

$$\frac{1}{2} \left| \hat{K}^{DD \rightarrow D} + \hat{K}^{NN \rightarrow N} \pm \sqrt{(\hat{K}^{DD \rightarrow D} - \hat{K}^{NN \rightarrow N})^2 + 4\hat{K}^{NN \rightarrow D}\hat{K}^{DD \rightarrow N}} \right| < 1, \quad (13)$$

where the hats denote the integrated kernels, schematically $\hat{K} = \sum_{\tau=1}^{\infty} K(\tau)$. This is indeed verified empirically, the eigenvalues being $\lambda_1 \simeq 0.94$, $\lambda_2 \simeq 0.48$.

Moreover, for intra-day and overnight volatilities separately, we checked that our calibrated kernels are compatible with the two positivity conditions derived in Appendices A.2,A.3 : the first one referring to the cross kernels $K^{ND \rightarrow}$ and $K^{DN \rightarrow}$, and the second one to the leverage kernels $L^{D \rightarrow}$ and $L^{N \rightarrow}$. For $q = 512$, the criteria fail for two spurious reasons. Firstly, for lags greater than their exponential cut-offs, the quadratic kernels vanish quicker than the ‘cross’ kernels, which makes the “ τ by τ ” criterium fail. Secondly, the criterium $L^\dagger K^{-1} L \leq 4s^2$ cannot be verified for overnight volatility with $s^{N^2} = 0$ (for lower values of q , using the same functional forms and coefficients for the kernels, s^{N^2} rises to a few percents). These two effects can be considered spurious because the long-range contributions have a weak impact on the volatilities and cannot in deed generate negative values, as we checked by simulating the volatility processes with $q = 512$. We thus restricted the range to $q = 126$ (= six months) in order to test our results with the two positive volatility criteria (again, see Appendix A). For the ill-determined exponential decay rates ω_p , the upper bounds of the confidence intervals are used. The two criteria are then indeed verified for both intra-day and overnight volatilities.

3.2 Distribution of the residuals

As mentioned above, we assume that the residuals ξ (i.e. the returns divided by the volatility predicted by the model) are Student-distributed. This is now common in ARCH/GARCH literature and was again found to be satisfactory in our companion paper [6]. The fact that the ξ are not Gaussian means that there is a residual surprise element in large stock returns, that must be interpreted as true ‘jumps’ in the price series.

The tail cumulative distribution function (CDF) of the residuals is shown in Fig. 4 for both intraday and overnight returns, together with Student best fits, obtained with long feedback kernels

q	21	42	512
$\nu^D(\pm 0.3)$	10.7	12.6	13.5
$\nu^N(\pm 0.02)$	3.49	3.58	3.61
ρ^D	18.1%	12.8%	8.5%
ρ^N	14.0%	7.3%	0.0%

Table 4: Baseline volatility and tail index of the Student residuals for several maximum lags q .

($q = 512$). This reveals a clear difference in the statistical properties of ξ^D and ξ^N . First, the Student fit is better for overnight residuals than for intraday residuals, in particular far in the tails. More importantly, the number of degrees of freedom ν is markedly different for the two types of residuals: indeed, our MLE estimation yields $\nu^D = 13.5$ and $\nu^N = 3.61$ as reported in Tab. 4, resulting in values of the residual kurtosis $\kappa^D = 3 + 6/(\nu^D - 4) = 3.6$ and $\kappa^N = \infty$. This result must be compared with the empirical kurtosis of the returns that was measured directly in Sect. 2. The intuitive conclusion is that the large (infinite?) kurtosis of the overnight returns cannot be attributed to fluctuations in the volatility, but rather, as mentioned above, to large jumps related to overnight news. This clear qualitative difference between intraday and overnight returns is a strong argument justifying the need to treat these effects separately, as proposed in this paper.

We have also studied the evolution of ν^D and ν^N as a function of the length q of the memory kernel, see Tab. 4. If longer memory kernels allow to account for more of the dynamics of the volatility, less kurtic residuals should be found for larger q 's. This is indeed what we find, in particular for intraday returns, for which ν^D increases from 10.7 for $q = 21$ to 13.5 for $q = 512$. The increase is however much more modest for overnight returns. We propose below an interpretation of this fact.

3.3 Baseline volatility

Finally, we want to study the ratio $\rho^j = s^{j^2}/\langle\sigma^{j^2}\rangle$, which can be seen as a measure of the relative importance of the baseline component of the volatility, with respect to the endogenous feedback component.³ The complement $1 - \rho^j$ gives the relative contribution of the feedback component, given by $\hat{K}^{DD \rightarrow j} + \hat{K}^{NN \rightarrow j}$ in the present context.⁴

The results for ρ^j are given in Tab. 4 for $q = 21, 42$ and 512 : as mentioned above, a larger q explains more of the volatility, therefore reducing the value of both ρ^D and ρ^N . While ρ^D is small (~ 0.1) and comparable to the value found for the daily ARCH model studied in the companion paper [6], the baseline contribution is *nearly zero* for the overnight volatility. We find this result truly remarkable, and counter-intuitive at first sight. Indeed, the baseline component of the volatility is usually associated to exogenous factors, which, as we argued above, should be dominant for σ^N since many unexpected pieces of news occur overnight!

Our interpretation of this apparent paradox relies on the highly kurtic nature of the overnight residual, with a small value $\nu^N \approx 3.6$ as reported above. The picture is thus as follows: most overnights are news-less, in which case the overnight volatility is completely fixed by feedback effects, set by the influence of past returns themselves. The overnights in which important news is released, on the other hand, contribute to the tails of the residual ξ^N , because the large amplitude of these returns could hardly be guessed from the previous amplitude of the returns. Furthermore, the fact that $K^{NN \rightarrow D}$ decays very quickly is in agreement with the idea, expressed in [12], that the impact of news (chiefly concentrated overnight) on volatility is short-lived.⁵

In conclusion, we find that most of the predictable part of the overnight volatility is of endogenous origin, but that the contribution of unexpected jumps reveals itself in the highly non-Gaussian statistics of the residuals. The intraday volatility, on the other hand, has nearly Gaussian residuals but still a very large component of endogenous volatility ($1 - \rho^D \approx 90\%$).

³In fact, the stability criterion for our model reads $\rho^j > 0$, which is found to be satisfied by our calibration, albeit marginally for the overnight volatility.

⁴The cross terms $K^{j'j'' \rightarrow j}$ are not exactly zero since day/night and night/day next correlations are small but finite, but their contribution is less than one order of magnitude lower than the $K^{j'j' \rightarrow j}$.

⁵This effect was confirmed recently in [7] using a direct method: the relaxation of volatility after large overnight jumps of either sign is very fast, much faster than the relaxation following large intraday jumps.

3.4 In-Sample and Out-of-Sample tests

In order to compare our bivariate IntraDay/Overnight volatility prediction model with the standard ARCH model for daily (close-to-close) volatility, we ran In-Sample (IS) and Out-of-Sample (OS) likelihood computations with both models. Of course, in order to compare models, the same quantities must be predicted. A daily ARCH model that predicts the daily volatility σ_t^2 at date t can predict intraday and overnight volatilities as follows:

$$\widehat{\sigma}_t^{\text{D}}{}^2 = \frac{[\langle r^{\text{D}2} \rangle]}{[\langle r^2 \rangle]} \sigma_t^2, \quad \widehat{\sigma}_t^{\text{N}}{}^2 = \frac{[\langle r^{\text{N}2} \rangle]}{[\langle r^2 \rangle]} \sigma_t^2, \quad (14)$$

where $[\langle \cdot \rangle]$ is the average over all dates and all stocks, and as in Sect. 2, $r_t = r_t^{\text{D}} + r_t^{\text{N}}$ is the daily (close-to-close) return of date t . Similarly, our bivariate intraday/overnight model that provides predictions for intraday and overnight volatilities separately can give the following estimation of the daily volatility:

$$\widehat{\sigma}_t^2 = \sigma_t^{\text{D}2} + \sigma_t^{\text{N}2} + 2[\langle r^{\text{D}}r^{\text{N}} \rangle]. \quad (15)$$

For each of the 6 predictions (of the intraday, overnight and daily volatilities by the two models separately, bivariate intraday/overnight and standard ARCH), we use the following methodology:

- The pool of stock names is split in two halves, and the model parameters are estimated separately on each half.
- The “per point” log-likelihood (8) is computed for both sets of parameters, once with *the same* half dataset as used for the calibration (In-Sample), and once with *the other* half dataset (Out-of-Sample, or “control”). We compute the log-likelihoods for intraday and overnight volatilities ($J \in \{\text{D}, \text{N}\}$), and for daily volatility:

$$\begin{aligned} \text{Bivariate models : } & \mathcal{L}^{\text{biv}}(\sigma^J, \nu^J, r^J) \quad \text{and} \quad \mathcal{L}^{\text{biv}}(\widehat{\sigma}, \nu^{\text{daily}}, r^{\text{daily}}); \\ \text{Standard ARCH models : } & \mathcal{L}^{\text{std}}(\widehat{\sigma}^J, \nu^J, r^J) \quad \text{and} \quad \mathcal{L}^{\text{std}}(\sigma, \nu^{\text{daily}}, r^{\text{daily}}); \end{aligned}$$

where \mathcal{L}^{biv} is computed in the bivariate model (i.e. with six regressors: four quadratic and two linear) and \mathcal{L}^{std} is computed in the standard ARCH model (i.e. with two regressors: one quadratic and one linear), and $\widehat{\sigma}^J$ and $\widehat{\sigma}$ are as given by equations (14) and (15).

- The IS log-likelihood \mathcal{L}_{IS} of the model is computed as the average of the two In-Sample results, and similarly for the OS log-likelihood \mathcal{L}_{OS} . We call $l_{\text{IS}} = \exp(\mathcal{L}_{\text{IS}})$ and $l_{\text{OS}} = \exp(\mathcal{L}_{\text{OS}})$ the average likelihood per point (ALpp) of the model IS and OS, expressed as percentages, that are two proxies of the “probability that the sample data were generated by the model”.

We then use the values of l_{IS} and l_{OS} to compare models. For a “good” model, these values must be as high as possible, but they must also be close to each other. As a matter of fact, if l_{IS} is significantly greater than l_{OS} , the model may be over-fitting the data. On the contrary, if l_{OS} is greater than l_{IS} , which seems counter-intuitive, the model may be badly calibrated. The results of this model comparison are given in Tab. 5: the bivariate intraday/overnight model has a higher likelihood than the standard daily ARCH model, both In-Sample (this was to be expected even from pure over-fitting due to additional parameters) and Out-of-Sample (thus outperforming in predicting the “typical” random realization of the returns).

The likelihoods of the predictions obtained with equations (14) and (15) are marked with the exponent † in Tab. 5. For these likelihoods, “In-Sample” simply means that the same half of the stock pool was used for the calibration of the model and for the estimation of the likelihood, although different types of returns are considered. Similarly, “Out-of-Sample” likelihoods are estimated on the other half of the stock pool. These values serve as comparison benchmarks between the two models.

We see that in all cases, the bivariate IntraDay/Overnight significantly outperforms the standard daily ARCH framework, in particular concerning the prediction of the total (Close-Close) volatility.

Prediction	σ^{D^2}		σ^{N^2}		σ^2 (daily)	
	l_{IS}	l_{OS}	l_{IS}	l_{OS}	l_{IS}	l_{OS}
Biv. IntraDay/Overnight	44.423	44.418	50.839	50.826	45.512 [†]	45.509 [†]
Standard ARCH	44.227 [†]	44.225 [†]	50.598 [†]	50.595 [†]	44.931	44.928

Table 5: In-Sample and Out-of-Sample average per point likelihoods. Figures with [†] pertain to reconstructed volatilities Eqs. (14) or (15).

4 Conclusion and extension

The main message of this study is quite simple, and in fact to some extent expected: overnight and intraday returns are completely different statistical objects. The ARCH formalism, that allows one to decompose the volatility into an exogenous component and a feedback component, emphasizes this difference. The salient features are:

- While past intraday returns affect equally both the future intraday and overnight volatilities, past overnight returns have a weak effect on future intraday volatilities (except for the very next one) but impact substantially future overnight volatilities.
- The exogenous component of overnight volatilities is found to be close to zero, which means that the lion’s share of overnight volatility comes from feedback effects.
- The residual kurtosis of returns (once the ARCH effects have been factored out) is small for intraday returns but infinite for overnight returns.
- The bivariate IntraDay/Overnight model significantly outperforms the standard ARCH framework based on daily returns for Out-of-Sample predictions.

Intuitively, a plausible picture for overnight returns is as follows: most overnights are news-less, in which case the overnight volatility is completely fixed by feedback effects, set by the influence of past returns themselves. Some (rare) overnights witness unexpected news releases, which lead to huge jumps, the amplitude of which could hardly have been guessed from the previous amplitude of the returns. This explains why these exogenous events contribute to residuals with such fat tails that the kurtosis diverge, and *not* to the baseline volatility that concerns the majority of news-less overnights.

Having decomposed the Close-Close return into an overnight and an intraday component, the next obvious step is to decompose the intraday return into higher frequency bins — say five minutes. We have investigated this problem as well, the results are reported in [3]. In a nutshell, we find that once the ARCH prediction of the intraday average volatility is factored out, we still identify a causal feedback from past five minute returns onto the volatility of the current bin. This feedback has again a leverage component and a quadratic (ARCH) component. The intraday leverage kernel is close to an exponential with a decay time of ≈ 1 hour. The intraday ARCH kernel, on the other hand, *is still a power law*, with an exponent that we find to be close to unity, in agreement with several studies in the literature concerning the intraday temporal correlations of volatility/activity — see e.g. [15, 13, 18], and, in the context of Hawkes processes, [1, 11]. It would be very interesting to repeat the analysis of the companion paper [6] on five minute returns and check whether there is also a dominance of the diagonal terms of the QARCH kernels over the off-diagonal ones, as we found for daily returns. This would suggest a microscopic interpretation of the ARCH feedback mechanism in terms of a Hawkes process for the trading activity.

Acknowledgements We thank S. Ciliberti, N. Cosson, L. Dao, C. Emako-Kazianou, S. Hardiman, F. Lillo and M. Potters for useful discussions. P. B. acknowledges financial support from Fondation Natixis (Paris).

References

- [1] E. Bacry, K. Dayri, and J.-F. Muzy. Non-parametric kernel estimation for symmetric Hawkes processes. Application to high frequency financial data. *The European Physical Journal B*, 85:1–12, 2012.
- [2] R. T. Baillie, T. Bollerslev, and H. O. Mikkelsen. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74(1):3–30, 1996.
- [3] P. Blanc. Modélisation de la volatilité des marchés financiers par une structure ARCH multi-fréquence. Master’s thesis, Université de Paris VI Pierre et Marie Curie, Sep 2012. Available upon request.
- [4] T. Bollerslev, R. F. Engle, and D. B. Nelson. *ARCH models*, pages 2959–3038. Volume 4 of Engle and McFadden [9], 1994.
- [5] L. Borland and J.-P. Bouchaud. On a multi-timescale statistical feedback model for volatility fluctuations. *The Journal of Investment Strategies*, 1(1):65–104, December 2011.
- [6] R. Chicheportiche and J.-P. Bouchaud. The fine-structure of volatility feedback. *arXiv preprint q-fin.ST/1206.2153*, 2013. This issue (reference to be updated).
- [7] N. Cosson. Analysis of realized and implied volatility after stock price jumps. Master’s thesis, Université de Paris VI Pierre et Marie Curie, Jun 2013. Available upon request.
- [8] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, pages 987–1007, 1982.
- [9] R. F. Engle and D. L. McFadden, editors. *Handbook of Econometrics*, volume 4. Elsevier/North-Holland, Amsterdam, 1994.
- [10] G. M. Gallo. Modelling the impact of overnight surprises on intra-daily volatility. *Australian Economic Papers*, 40(4):567–580, 2001.
- [11] S. J. Hardiman, N. Bercot, and J.-P. Bouchaud. Critical reflexivity in financial markets: a Hawkes process analysis. *arXiv preprint q-fin.ST/1302.1405*, 2013.
- [12] A. Joulin, A. Lefevre, D. Grunberg, and J.-P. Bouchaud. Stock price jumps: News and volume play a minor role. *Wilmott Magazine*, pages 1–7, September/October 2008.
- [13] Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C.-K. Peng, and H. E. Stanley. Statistical properties of the volatility of price fluctuations. *Physical Review E*, 60:1390–1400, Aug 1999.
- [14] U. A. Müller, M. M. Dacorogna, R. D. Davé, R. B. Olsen, O. V. Pictet, and J. E. von Weizsäcker. Volatilities of different time resolutions — analyzing the dynamics of market components. *Journal of Empirical Finance*, 4(2):213–239, 1997.
- [15] M. Potters, R. Cont, and J.-P. Bouchaud. Financial markets as adaptive systems. *EPL (Europhysics Letters)*, 41(3):239, 1998.
- [16] E. Sentana. Quadratic ARCH models. *The Review of Economic Studies*, 62(4):639, 1995.
- [17] I. Tsiakas. Overnight information and stochastic volatility: A study of European and US stock exchanges. *Journal of Banking & Finance*, 32(2):251–268, 2008.
- [18] F. Wang, K. Yamasaki, S. Havlin, and H. E. Stanley. Scaling and memory of intraday volatility return intervals in stock markets. *Physical Review E*, 73(2):026117, 2006.
- [19] G. Zumbach and P. Lynch. Heterogeneous volatility cascade in financial markets. *Physica A: Statistical Mechanics and its Applications*, 298(3-4):521–529, 2001.

A Non-negative volatility conditions

In this appendix, we study the mathematical validity of our volatility regression model. The first obvious condition is that the model is stable, which leads to the condition (13) in the text above. This criterion is indeed obeyed by the kernels calibrated on empirical results. Secondly, the volatility must remain positive, which is not a priori guaranteed with multiple kernels associated to signed regressors. We now establish a set of sufficient conditions on the feedback kernels to obtain non-negative volatility processes.

A.1 One correlation feedback kernel, no leverage coefficients

We consider first the simpler model for daily volatility, without linear regression coefficients:

$$\sigma_t^2 = s^2 + \sum_{\tau=1}^q K^{\text{DD}\rightarrow}(\tau) r_{t-\tau}^{\text{D}}{}^2 + \sum_{\tau=1}^q K^{\text{NN}\rightarrow}(\tau) r_{t-\tau}^{\text{N}}{}^2 + 2 \sum_{\tau=1}^q K^{\text{ND}\rightarrow}(\tau) r_{t-\tau}^{\text{D}} r_{t-\tau}^{\text{N}}.$$

This modification of the standard ARCH model can lead to negative volatilities if (at least) one term in the last sum takes large negative values. This issue can be studied more precisely with the matrix form of the model:

$$\sigma_t^2 = s^2 + R_t^\dagger K R_t,$$

with

$$K = \begin{pmatrix} K^{\text{DD}\rightarrow}(1) & & & & K^{\text{ND}\rightarrow}(1) & & & & \\ & \ddots & & & & \ddots & & & \\ & & K^{\text{DD}\rightarrow}(q) & & & & K^{\text{ND}\rightarrow}(q) & & \\ K^{\text{ND}\rightarrow}(1) & & & & K^{\text{NN}\rightarrow}(1) & & & & \\ & \ddots & & & & \ddots & & & \\ & & & & & & & & K^{\text{NN}\rightarrow}(q) \\ & & & & & & & & K^{\text{ND}\rightarrow}(q) \end{pmatrix}, \quad R_t = \begin{pmatrix} r_{t-1}^{\text{D}} \\ \vdots \\ r_{t-q}^{\text{D}} \\ r_{t-1}^{\text{N}} \\ \vdots \\ r_{t-q}^{\text{N}} \end{pmatrix},$$

and where $K^{\text{DD}\rightarrow}$ and $K^{\text{NN}\rightarrow}$ coefficients are assumed to be all positive (which is the case empirically). This formula highlights the fact that the volatility remains positive as soon as the symmetric matrix K is positive semidefinite. We now determine a sufficient and necessary condition under which K has negative eigenvalues. The characteristic polynomial of K is

$$\chi_K(X) = \prod_{\tau=1}^q \left[(K^{\text{DD}\rightarrow}(\tau) - X)(K^{\text{NN}\rightarrow}(\tau) - X) - K^{\text{ND}\rightarrow}(\tau)^2 \right],$$

and the eigenvalues of K are the zeroes of $\chi_K(X)$, solutions $\chi_K(\lambda) = 0$, i.e. such that

$$\lambda^2 - (K^{\text{DD}\rightarrow}(\tau) + K^{\text{NN}\rightarrow}(\tau))\lambda + K^{\text{DD}\rightarrow}(\tau)K^{\text{NN}\rightarrow}(\tau) - K^{\text{ND}\rightarrow}(\tau)^2 = 0.$$

Hence, K has at least one negative eigenvalue iff $\exists \tau \in \{1, \dots, q\}$ s.t.

$$K^{\text{DD}\rightarrow}(\tau) + K^{\text{NN}\rightarrow}(\tau) - \sqrt{(K^{\text{DD}\rightarrow}(\tau) - K^{\text{NN}\rightarrow}(\tau))^2 + 4K^{\text{ND}\rightarrow}(\tau)^2} < 0,$$

and finally,

$$K \text{ is positive semidefinite} \Leftrightarrow \forall \tau \in \{1, \dots, q\}, \quad K^{\text{ND}\rightarrow}(\tau)^2 \leq K^{\text{DD}\rightarrow}(\tau)K^{\text{NN}\rightarrow}(\tau). \quad (16)$$

When the quadratic kernel K is positive semidefinite, σ_t^2 remains positive for all t . For example, in the standard ARCH model, the inequality is saturated for all τ by construction, and the condition (16) is satisfied.

The positive semi-definiteness of K , equivalent to

$$\forall v = (v_1^{\text{D}}, \dots, v_q^{\text{D}}, v_1^{\text{N}}, \dots, v_q^{\text{N}})^\dagger \in \mathbb{R}^{2q}, \quad v^\dagger K v \geq 0,$$

and clearly, a sufficient condition for the sum to be non-negative is that each term is non-negative:

$$\begin{aligned}
\forall t, \sigma_t^2 \geq 0 &\Leftrightarrow K \text{ is positive semidefinite} \\
&\Leftrightarrow \exists \beta \in]0, 1[, \forall v = (v_1^D, \dots, v_q^D, v_0^N, v_1^N, \dots, v_q^N)^\dagger \in \mathbb{R}^{(2q+1)}, \forall \tau \in \{1, \dots, q\}, \\
&\quad 0 \leq \beta K^{\text{NN}\rightarrow}(\tau) v_\tau^{\text{N}^2} + (1 - \beta) K^{\text{NN}\rightarrow}(\tau - 1) v_{\tau-1}^{\text{N}^2} \\
&\quad + K^{\text{DD}\rightarrow}(\tau) v_\tau^{\text{D}^2} + 2K^{\text{ND}\rightarrow}(\tau) v_\tau^{\text{D}} v_\tau^{\text{N}} + 2K^{\text{DN}\rightarrow}(\tau) v_\tau^{\text{D}} v_{\tau-1}^{\text{N}} \\
&\Leftrightarrow \exists \beta, \forall v, \forall \tau \in \{1, \dots, q\}, \quad \exists \alpha_\tau \in [0, 1], \\
&\quad \bullet |K^{\text{ND}\rightarrow}(\tau) v_\tau^{\text{D}} v_\tau^{\text{N}}| \leq \frac{1}{2} \left(\alpha_\tau K^{\text{DD}\rightarrow}(\tau) v_\tau^{\text{D}^2} + \beta K^{\text{NN}\rightarrow}(\tau) v_\tau^{\text{N}^2} \right) \\
&\quad \bullet |K^{\text{DN}\rightarrow}(\tau) v_\tau^{\text{D}} v_{\tau-1}^{\text{N}}| \leq \frac{1}{2} \left((1 - \alpha_\tau) K^{\text{DD}\rightarrow}(\tau) v_\tau^{\text{D}^2} + (1 - \beta) K^{\text{NN}\rightarrow}(\tau - 1) v_{\tau-1}^{\text{N}^2} \right) \\
&\Leftrightarrow \exists \beta, \forall \tau \in \{1, \dots, q\}, \quad \exists \alpha_\tau \in [0, 1], \\
&\quad \bullet K^{\text{ND}\rightarrow}(\tau)^2 \leq \beta \alpha_\tau K^{\text{DD}\rightarrow}(\tau) K^{\text{NN}\rightarrow}(\tau) \\
&\quad \bullet K^{\text{DN}\rightarrow}(\tau)^2 \leq (1 - \beta)(1 - \alpha_\tau) K^{\text{DD}\rightarrow}(\tau) K^{\text{NN}\rightarrow}(\tau - 1).
\end{aligned}$$

The last condition is equivalent to a simpler one, with α_τ saturating one of the two inequalities: denoting $\delta^{(\text{nn})}(\tau) = K^{\text{NN}\rightarrow}(\tau)/K^{\text{NN}\rightarrow}(\tau - 1)$ and $\delta^{(\text{nd})}(\tau) = K^{\text{DN}\rightarrow}(\tau)/K^{\text{ND}\rightarrow}(\tau)$, K is positive semidefinite if (but not only if), $\exists \beta \in]0, 1[, \forall \tau \in \{1, \dots, q\}$,

$$\begin{aligned}
\mathcal{M}(\beta, \tau) \equiv \max &\left\{ \frac{\beta \delta^{(\text{nn})}(\tau)}{1 - \beta} \left(\frac{(1 - \beta) K^{\text{DD}\rightarrow}(\tau) K^{\text{NN}\rightarrow}(\tau - 1)}{K^{\text{ND}\rightarrow}(\tau)^2} - \delta^{(\text{nd})}(\tau)^2 \right), \right. \\
&\left. \frac{1 - \beta}{\beta \delta^{(\text{nn})}(\tau)} \left(\frac{\beta K^{\text{DD}\rightarrow}(\tau) K^{\text{NN}\rightarrow}(\tau)}{K^{\text{DN}\rightarrow}(\tau)^2} - \frac{1}{\delta^{(\text{nd})}(\tau)^2} \right) \right\}
\end{aligned}$$

is larger than one, yielding the following a.s. positive volatility criterion:

$$\forall t, \sigma_t^2 \geq 0 \Leftrightarrow 1 \leq \sup_{\beta \in]0, 1[} \min_{1 \leq \tau \leq q} \mathcal{M}(\beta, \tau).$$

A.3 With leverage coefficients

We now add leverage terms to the volatility equation:

$$\sigma_t^2 = s^2 + R_t^\dagger K R_t + R_t^\dagger L = \tilde{R}_t^\dagger M \tilde{R}_t,$$

with

$$M = \begin{pmatrix} K & \frac{1}{2}L \\ \frac{1}{2}L^\dagger & s^2 \end{pmatrix},$$

and appropriate vectors R_t, L_t, \tilde{R}_t . It is easy to show that, assuming a positive definite K ,

$$M \text{ is positive semidefinite} \Leftrightarrow L^\dagger K^{-1} L \leq 4s^2. \tag{17}$$

B Universality assumption

To obtain a better convergence of the parameters of the model, the estimates are averaged over a pool of 280 US stocks. The validity of this method lies on the assumption that the model is approximately universal, i.e. that the values of its coefficients do not vary significantly among stocks.

We check that this assumption is relevant by splitting the stock pool in two halves and running the estimation of the model on the two halves independently. We obtain a set $\Theta_1 \in \mathbb{R}^{17}$ of coefficients calibrated on the first half and a set $\Theta_2 \in \mathbb{R}^{17}$ on the second half (each set contains 17 parameters, three for each of the four K kernels, two for each of the two L kernels, plus ν).

If the (normalized) returns series for each stock were independent realizations of the same process, the differences between the coefficients of the two half stockpools would be explained by statistical noise only. To quantify how close the observed differences are to statistical noise, we run a series of Wald tests and study the obtained p-values. We test $H_0 = \{f(\Theta) = 0\}$ against $H_1 = \{f(\Theta) \neq 0\}$, where $\Theta = (\Theta_1, \Theta_2) \in \mathbb{R}^{34}$, $f(\Theta) = \Theta_1 - \Theta_2$, $f : \mathbb{R}^{34} \mapsto \mathbb{R}^{17}$, by comparing the statistic

$$\Xi_n = n f(\Theta)^\dagger \Sigma(\Theta)^{-1} f(\Theta), \quad \text{with} \quad \Sigma(\Theta) = \frac{\partial f}{\partial \Theta}(\Theta) I(\Theta)^{-1} \left(\frac{\partial f}{\partial \Theta}(\Theta) \right)^\dagger, \quad (18)$$

to the quantiles of a χ^2 variable, where $n = \frac{1}{2} \times 280 \times 2515$ is the sample size for each half stockpool, $I(\Theta)$ is the Fisher Information matrix of the model and $\frac{\partial f}{\partial \Theta}$ is the Jacobian matrix of $f(\Theta)$.

For intraday volatility, the p-value is close to zero if all the 17 coefficients are included, but becomes very high (p-val = 12.3%) if we exclude $\alpha^{\text{NN} \rightarrow \text{D}}$ from the test. One can conclude that all the parameters but $\alpha^{\text{NN} \rightarrow \text{D}}$ can be considered universal, with a high significance level of 10%. It is not surprising that at least one coefficient varies slightly among stocks (it would indeed be a huge discovery to find that 280 US stocks can be considered as identically distributed!).

In the case of overnight volatility, we first test the universality of the parameters in $K^{\text{DD} \rightarrow \text{N}}$ and $K^{\text{NN} \rightarrow \text{N}}$, for which the constraint $s^{\text{N}^2} = 0$ is included in the estimation. We then test the other 11 parameters for universality. Once again, a few of them ($g_{\text{p}}^{\text{ND} \rightarrow \text{N}}$, $\alpha^{\text{NN} \rightarrow \text{N}}$ and $\omega_{\text{p}}^{\text{NN} \rightarrow \text{N}}$) must be excluded from the tests to obtain acceptable p-values. We then obtain p-val = 1% for the first test and p-val = 7.5% for the second.

It is then natural to wonder whether the four coefficients that cannot be considered as statistically universal differ significantly in relative values. That is why we compute a second comparison criterion: for a pair $(c^{(1)}, c^{(2)})$ of coefficients estimated on the first and second half stockpools respectively, we compute the relative difference, defined as:

$$\frac{|c^{(1)} - c^{(2)}|}{\max\{|c^{(1)}|, |c^{(2)}|\}}.$$

The values of this criterion are summarized in Tabs. 6,7. The first observation is that no relative difference exceeds 50% (except for three of the ill-determined ω_{p}) which indicates that the signs and orders of magnitude of the coefficients of the model are invariant among stocks. The coefficients for which the relative difference is high but the statistical one is low do not contradict the universality assumption: the ML estimation would need a larger dataset to determine them with precision, and the difference between the two halves can be interpreted as statistical noise.

Three of the four “non-universal” coefficients, $\alpha^{\text{NN} \rightarrow \text{D}}$, $g_{\text{p}}^{\text{ND} \rightarrow \text{N}}$ and $\omega_{\text{p}}^{\text{NN} \rightarrow \text{N}}$ also show a significant relative difference between the two stockpools (above 10%). These are thus the only coefficients of the model for which averaging over all stocks is in principle not well-justified, and for which the confidence intervals given in Sect. 3.1 should be larger. However, these variations do not impact the global shapes of the corresponding kernels in a major way, and our qualitative comments on the feedback of past returns on future intraday and overnight volatilities are still valid.

The results of this section indicate that most coefficients of the model are compatible with the assumption of universality. Although some coefficients do show slight variations, our stock aggregation method (with proper normalization, as presented in Sect. 2.3) is reasonable.

Kernels	$K(\tau) = g_p \tau^{-\alpha} e^{-\omega_p \tau}$			Kernels	$L(\tau) = g_e e^{-\omega_e \tau}$	
	g_p	α	ω_p		g_e	ω_e
$K^{DD \rightarrow D}$	9.4%	7.3%	13.1%	$L^{D \rightarrow D}$	10.4%	1.8%
$K^{NN \rightarrow D}$	2.6%	17.2%	34.4 %	$L^{N \rightarrow D}$	5.8%	2.3%
$K^{ND \rightarrow D}$	13.3%	9.6%	77.6 %			
$K^{DN \rightarrow D}$	2.0%	9.9%	94.8 %			

Table 6: Intra-day volatility: relative differences between the two half stockpools ($q = 512$). For ν^D , the value is 7.1%. Bold figures are above 20%.

Kernels	$K(\tau) = g_p \tau^{-\alpha} e^{-\omega_p \tau}$			Kernels	$L(\tau) = g_e e^{-\omega_e \tau}$	
	g_p	α	ω_p		g_e	ω_e
$K^{DD \rightarrow N}$	3.1%	9.7%	17.0%	$L^{D \rightarrow N}$	18.3%	32.0 %
$K^{NN \rightarrow N}$	8.3%	7.6%	33.0 %	$L^{N \rightarrow N}$	19.4%	6.5%
$K^{ND \rightarrow N}$	30.1 %	1.7%	80.0 %			
$K^{DN \rightarrow N}$	35.5 %	21.0 %	0.2%			

Table 7: Overnight volatility: relative differences between the two half stockpools ($q = 512$). For ν^N the value is 0.03%. Bold figures are above 20%.