This is the accepted and peer reviewed manuscript of: Christian Francq & Genaro Sucarrat (2023) *Volatility Estimation When the Zero-Process is Nonstationary*, Journal of Business & Economic Statistics, 41:1, 53-66, DOI: 10.1080/07350015.2021.1999821

Volatility Estimation when the Zero-Process is Nonstationary

Christian Francq[1] and Genaro Sucarrat[2]

This version: 11 October 2021

Abstract

Financial returns are frequently nonstationary due to the nonstationary distribution of zeros. In daily stock returns, for example, the nonstationarity can be due to an upwards trend in liquidity over time, which may lead to a downwards trend in the zero-probability. In intraday returns, the zero-probability may be periodic: It is lower in periods where the opening hours of the main financial centres overlap, and higher otherwise. A nonstationary zero-process invalidates standard estimators of volatility models, since they rely on the assumption that returns are strictly stationary. We propose a GARCH model that accommodates a nonstationary zero-process, derive a 0-adjusted QMLE for the parameters of the model, and prove its consistency and asymptotic normality under mild assumptions. The volatility specification in our model can contain higher order ARCH and GARCH terms, and past zero-indicators as covariates. Simulations verify the asymptotic properties in finite samples, and show that the standard estimator is biased. An empirical study of daily and intradaily returns illustrate our results. They show how a nonstationary zero-process induces time-varying parameters in the conditional variance representation, and that the distribution of zero returns can have a strong impact on volatility predictions.

JEL Classification: C01, C22, C51, C58
Keywords: Financial return, ARCH models, volatility, zero-inflated return, nonstationary return

.

[1]CREST and Université de Lille

[2]Department of Economics, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. Email genaro.sucarrat at bi.no

1 Introduction

Financial returns are frequently zero. This can be due to liquidity issues (e.g. low trading volume), price discreteness or rounding error, data issues (e.g. imputation due to missing values), market closures, and other market-specific characteristics and developments.

A number of approaches accommodate the occurrence of zeros. In continuous time approaches, for example, zeros occur when the assumed underlying price process is not observed. Lesmond et al. (1999) used this idea to construct a popular measure of liquidity based on observed zeros. More recently, the role of zeros has been repositioned in a continuous time framework by, amongst others, Bandi et al. (2017), and Bandi et al. (2020). Building on these developments, Buccheri et al. (2020) derive a bias-correction for Realised Volatility (RV), and Buccheri et al. (2020) propose a way to improve portfolio management in the presence of zeros. Supplemental appendix D outlines the connection between continuous time approaches to zeros and the model proposed here. In a second body of literature, zeros naturally occur due to the discreteness of price changes. Hausman et al. (1992) proposed an ordered probit model for discrete price changes. Russell and Engle (2005) proposed an Autoregressive Conditional Multinomial (ACM) model in combination with their Autoregressive Conditional Duration (ACD) model from Engle and Russell (1998). Liesenfeld et al. (2006) critised this approach, and proposed instead a dynamic integer count model. This was extended to the multivariate case in Bien et al. (2011). Rydberg and Shephard (2003) propose a model where the price increment is decomposed multiplicatively into three components: Activity, direction and integer magnitude. Catania et al. (2020) propose a discrete mixture approach to discrete price changes. In a third body of literature, price changes are continuous except at zero. Hautsch et al. (2013) propose a zero-inflated model for volume. Kümm and Küsters (2015) propose a zero-inflated model, where zeros occur either because there is no information available or because of rounding. In Harvey and Ito (2020), zeros occur due to censoring of an underlying continuous variable. Finally, the Generalised Autoregressive Conditional Heteroscedasticity (GARCH) class of models provides a fourth body of literature, since it accommodates zero-returns as long as the innovation can be zero, see the discussion in Sucarrat and Grønneberg (2020). In particular, if the standardised innovation is stationary, the parameters of a GARCH specification can be consistently estimated by the Standard Quasi Maximum Likelihood Estimator (QMLE) even when the conditional zero-probability is time-varying, see e.g. Escanciano (2009).

While the aforementioned contributions accommodate zeros in one or another way, very few of them pay attention to the fact that the zero-process can be nonstationary. This is striking, since the zero-process is frequently nonstationary. In daily stock returns, for example, a downwards (upwards) trend in the zero-probability can be due to an upwards (downwards) trend in liquidity over time, or an upwards (downwards) trend in the price level of the stock. Sucarrat and Grønneberg (2020) found widespread evidence of a trend in the zero-probability of daily stock returns at the New York Stock Exchange (NYSE). (We revisit a selection of their stocks in Section 5.1.) In intraday returns, the zero-probability is often nonstationary periodic: It is lower in periods with low liquidity (e.g. when the opening hours of the main financial centres do not overlap), and higher in periods with high liquidity (e.g. in hours where the main financial centres are open at the same time). An example is Kolokolov et al. (2020), who find clear evidence of a periodic zero-probability in intraday stock returns.

Here, in this paper, we propose volatility models that accommodate nonstationary zeros, where the zero-probability can be trend-like or periodic in nature, or both. To this end, volatility is specified as a generic scale (i.e. the conditional variance is a special case). We derive a modified QMLE, which we label the 0-adjusted QMLE, and prove its consistency and asymptotic normality. We start with the standard

GARCH(1,1) model for which the regularity conditions are more explicit, then we extend the results to more general models which allow for higher order lags, asymmetries and also indicators of lagged zero returns. In the stationary case, our regularity conditions coincide with the sharpest assumptions given in the literature for CAN of the QMLE. Our asymptotic results mainly rely on the ergodic theorem for nonstationary processes introduced in Francq and Gautier (2004). Variations of it have also been used in Azrak and Mélard (2006), Phillips and Xu (2006) and Regnard and Zakoian (2010). Section 2 is devoted to the simple GARCH(1,1) model. In Section 3 we extend our results to more general specifications. In particular, we consider a model where lags of zero-indicators are added as covariates. This specification is of special interest, since empirical evidence suggests jumps may follow zeros, see Kolokolov and Reno (2019). Supplemental appendix A collects the proofs of our theorems, propositions and lemmas. Section 4 contains finite sample simulations of our estimator. They show that the Standard QMLE is biased in our experiments, and verify our asymptotic results. In particular, the empirical standard errors correspond well to the asymptotic ones in finite samples. Section 5 contains an empirical application of our results. They show how a nonstationary zero-process induces time-varying parameters in the conditional variance representation. Accordingly, the distribution of zero returns can have a strong impact on volatility predictions. Finally, Section 6 concludes and suggests lines for further research.

2 Structure and estimation of the GARCH(1,1) specification

Let (I_t) a *bitstream* sequence, *i.e.* a sequence valued in $\{0, 1\}$. This bitstream sequence is said to be well fed in zeros (resp. ones) if, for all t, there exists $u \leq t$ such that $I_u = 0$ (resp. $I_u = 1$). The value $I_t = 0$ indicates a zero return and $I_t = 1$ indicates a nonzero return at time t. Conditionally on (I_t), we will consider time series (ϵ_t) such that $\epsilon_t = 0$ if $I_t = 0$ and ϵ_t follows a non degenerated GARCH-type model when $I_t = 1$. First consider a simple zero-inflated GARCH(1,1) model of the form

$$\epsilon_t = \sigma_t \eta_t I_t, \quad \sigma_t^2 = \omega_0 + \alpha_0 \epsilon_{t-1}^2 + \beta_0 \sigma_{t-1}^2, \tag{2.1}$$

with a sequence (η_t) of non-degenerated real random variables, and nonnegative parameters ω_0, α_0 and β_0. Note that, for the moment, we do not make any precise assumption on the model. In particular, the sequence (I_t) can be the realization of a nonstationary sequence. Therefore, the model (2.1) can be considered as being semi-parametric. Moreover, if a solution of (2.1) exists, in general it is nonstationary. The following proposition gives a condition for the existence of such a solution.

Proposition 2.1. *Given sequences (η_t) and (I_t), and parameters $\omega_0 > 0$, $\alpha_0 \geq 0$ and $\beta_0 \geq 0$, there exists a (unique) (non anticipative) finite solution to (2.1) if*

$$\gamma_t := \limsup_{k\to\infty} \frac{1}{k} \sum_{i=1}^{k} \log(\alpha_0 \eta_{t-i}^2 I_{t-i} + \beta_0) < 0 \qquad a.s. \quad \forall t. \tag{2.2}$$

This condition is satisfied if for all t there exists $s > 0$ such that $\limsup_{k\to\infty} \frac{1}{k} \sum_{i=1}^{k} \left|\alpha_0 \eta_{t-i}^2 I_{t-i} + \beta_0\right|^s < 1$ a.s. There exists no finite solution if $\gamma_t > 0$ for some t.

In the previous proposition, a non anticipative solution means that σ_t is measurable with respect to the sigma-field $\mathcal{F}_{t-1}$ generated by $\{\eta_u, I_u; u < t\}$, and a finite solution means that $\sigma_t < \infty$ a.s. for all t, and (σ_t) is bounded in probability, in the sense that $\forall \varepsilon > 0$, $\exists M > 0$ and $n > 0$ such that $P(\sigma_t > M) < \varepsilon \;\; \forall |t| > n$.

Recall that the necessary and sufficient strict stationarity condition of the standard GARCH(1,1) model is

A1 $\gamma := E \log(\alpha_0 \eta_t^2 + \beta_0) < 0$.

Note that, when (η_t) is supposed to be a stationary and ergodic sequence, **A1** implies (2.2) (simply because $\log(\alpha_0\eta_t^2 I_t + \beta_0) \leq \log(\alpha_0\eta_t^2 + \beta_0)$). The condition is not necessary, however, because when $\beta_0 = 0$ and (I_t) is well fed in zeros, it is easy to see that (2.2) is satisfied without any restriction on α_0, in particular even when $\gamma = E\log(\alpha_0\eta_t^2) > 0$. Note also that we cannot conclude when $\gamma_t = 0$ because we do not make assumptions of the distributions of the zeros in (I_t).

For stationary GARCH models with iid innovations, it is known that the strict stationarity condition $\gamma < 0$ entails the existence of a marginal moment (see Lemma 2.3 in Berkes, Horváth, and Kokoszka, 2003b). The following proposition is a direct extension of that result.

Proposition 2.2. *If (η_t) is iid, $E|\eta_t|^r < \infty$ for $r > 0$ and* **A1** *holds, then the finite solution to* (2.1) *is such that $\sup_t E\sigma_t^{2s} < \infty$ and $\sup_t E|\epsilon_t|^{2s} < \infty$ for some $s > 0$.*

Assume that

A2 $P(\eta_t = 0) = 0$.

Under this assumption, $I_t = 0$ if and only if $\epsilon_t = 0$, and the sequence (I_t) is then observable whenever (ϵ_t) is observed. Given observations $\epsilon_1, \dots, \epsilon_n$, it is then possible to estimate the parameter $\theta_0 = (\omega_0, \alpha_0, \beta_0) \in \Theta \subset (0,\infty)^2 \times [0,\infty)$ by

$$\widehat{\theta}_n = \arg\min_{\theta\in\Theta} \widetilde{\boldsymbol{l}}_n(\theta), \quad \widetilde{\boldsymbol{l}}_n(\theta) = \frac{1}{n}\sum_{t=r_0+1}^{n} \widetilde{\ell}_t(\theta), \quad \widetilde{\ell}_t(\theta) = I_t\left\{\frac{\epsilon_t^2}{\widetilde{\sigma}_t^2(\theta)} + \log\widetilde{\sigma}_t^2(\theta)\right\}, \quad (2.3)$$

where $r_0 \geq 1$ is a fixed integer and $\widetilde{\sigma}_t^2(\theta) = \omega + \alpha\epsilon_{t-1}^2 + \beta\widetilde{\sigma}_{t-1}^2(\theta)$, with a fixed initial value $\widetilde{\sigma}_1^2(\theta)$. To show the consistency of this modified version of the QMLE, we need additional assumptions. We would like to deal with situations where the occurrence of the zeros may be random or/and periodic of period $T \in \mathbb{N}^*$ ($T = 1$ meaning no periodicity). To this aim, we assume that I_t is determined by a realization of a T-dimensional stationary process, at least for large t. Each date $t = (N-1)T + \nu$ corresponds to a cycle $N = N_t \in \mathbb{Z}$ and a season $\nu = \nu_t \in \{1, \dots, T\}$. More precisely, we have $N = \lceil t/T \rceil$, where $\lceil \cdot \rceil$ denotes the ceiling function.

A3 Let $(\eta_t)_{t\in\mathbb{Z}}$ and $(\boldsymbol{S}_N)_{N\in\mathbb{Z}}$ be two independent stationary and ergodic processes defined on some probability space $(\Omega, \mathcal{A}, \mathbb{P})$, respectively valued in $\mathbb{R}$ and $\mathcal{S} := \{0,1\}^T$. Let $\boldsymbol{S}_N = (S_{(N-1)T+1}, \dots, S_{(N-1)T+T})'$. Assume that there exists an almost surely finite random time t_0 such that, with probability one, $I_t = S_t$ for all $t \geq t_0$.

For daily returns there is usually no seasonality in the zero-process, and Figure 1 shows that the frequency of zeros stabilizes after a certain point for the stocks studied in Section 5.1. For these series, it therefore seems reasonable to assume **A3** with $T = 1$ and $t_0(\omega)$ corresponding to a certain date.

It is important to emphasize that in model (2.1), the sequence I_t is given. Therefore, even when I_t is the realization of a stationary process, *i.e.* in **A3** $T = 1$ and $I_t = S_t(\omega)$ for all t, conditionally on (I_t), the sequence (ϵ_t) is not stationary. Indeed, it is clear that ϵ_t and ϵ_{t+1} can not have the same distribution when $I_t \neq I_{t+1}$. We will work with random variables of the form $f(I_t, I_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots)$ which, conditionally on (I_t), are not stationary. The following lemma shows that a kind of law of large numbers can however be applied to such nonstationary sequences under **A3**. Similar results appear in Azrak and Mélard (2006), Francq and Gautier (2004), Phillips and Xu (2006) and Regnard and Zakoian (2010).

Lemma 2.1. *Let $f(\cdot;\cdot) : \{0,1\}^{\mathbb{N}} \times \mathbb{R}^{\mathbb{N}} \to \mathbb{R}$ be a measurable function. Assume that for $t = 1, \dots, T$ we have $Ef^+(S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots) < \infty$. Then, given any sequence (I_t) satisfying* **A3**, *we have*

$$\frac{1}{n}\sum_{t=1}^{n} f(I_t, I_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots) \to \frac{1}{T}\sum_{t=1}^{T} Ef(S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots) \in [-\infty, \infty)$$

almost surely. If the condition $Ef^+(S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots) < \infty$ is replaced by $Ef^-(S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots) < \infty$ then the limit belongs to $(-\infty, \infty]$.

Example 2.1 (Trivial application of Lemma 2.1)**.** *Let (η_t) be an independent sequence* *of* $\mathcal{N}(0,1)$*-distributed random variables,* $0 < \sigma(0) < \sigma(1)$*,* $T = 2$*, and* $\pi_1 = P(S_1 = 1) = 1 - P(S_1 = 0)$ *and* $\pi_2 = P(S_2 = 1) = 1 - P(S_2 = 0)$*. Given a sequence* (I_t) *satisfying* **A3***, defined the process* $X_t = \sigma(I_t)\eta_t$*. Reasoning conditionally to* (I_t)*, the sequence* (X_t) *is not stationary because the distribution of* X_t *is either* $\mathcal{N}(0, \sigma^2(0))$ *or* $\mathcal{N}(0, \sigma^2(1))$*. We have however the almost sure convergence*

$$\frac{1}{n}\sum_{t=1}^{n} X_t^2 \to \frac{1}{2}\left\{\pi_1\sigma^2(1) + (1-\pi_1)\sigma^2(0)\right\} + \frac{1}{2}\left\{\pi_2\sigma^2(1) + (1-\pi_2)\sigma^2(0)\right\} \text{ as } n \to \infty.$$

For proper definition of $\widehat{\theta}_n$ and identifiability of the GARCH parameters, assume the following assumptions, which are also required for consistency of the QMLE of stationary GARCH models.

A4 $\theta_0 = (\omega_0, \alpha_0, \beta_0) \in \Theta \subset (0,\infty)^2 \times [0,1)$ and Θ is compact.

A5 $\beta < 1$ for all $\theta \in \Theta$.

A6 Conditionally on (I_t), the sequence (η_t) is iid, $E\eta_t^2 = 1$ and $P(\eta_t^2 = 1) \neq 1$.

Remark 2.1 (Interpretation of σ_t)**.** To facilitate interpretation, suppose that $E\eta_t = 0$, as is generally the case for GARCH processes. Under **A6** and (2.2), we have $\sigma_t^2 = \mathrm{Var}(\epsilon_t \mid \mathcal{F}_{t-1}, I_t = 1)$. Thus σ_t corresponds to the volatility of ϵ_t when this return is non-zero. When $I_t = 0$, the variable σ_t does not have such an interpretation. Given the observations $\epsilon_1, \dots, \epsilon_n$, one can thus interpret σ_{n+1} as the volatility of the future return ϵ_{n+1} under the scenario that the latter is nonzero. Since I_t is taken as exogenous variable, our model is not sufficient to predict ϵ_t^2. Indeed we have $E\left(\epsilon_t^2 \mid \mathcal{F}_{t-1}\right) = \sigma_t^2 P(I_t = 1 \mid \mathcal{F}_{t-1})$, see Section 5.

Obviously, to be able to estimate the parameter of the volatility process (σ_t), it is also necessary to assume that (I_t) is well fed in 1s. We even have to assume that if $I_t = 1$ then I_{t-1} is not always equal to zero, otherwise, in the ARCH(1) case, $I_t\widetilde{\sigma}_t^2(\theta) = I_t(\omega + \alpha\sigma_{t-1}^2\eta_{t-1}^2 I_{t-1})$ would not depend on α. In **A3**, we thus assume that

A7 for some $j_0 \in \{1, \dots, T\}$, $P(S_{j_0} = 0) \neq 1$ and $P(S_{j_0-1} = 0) \neq 1$.

In general S_t is not independent of $\mathcal{F}_{t-1}^{S,\eta}$, where $\mathcal{F}_t^{S,\eta}$ denotes the sigma-field generated by $\{S_u, \eta_u; u \le t\}$. We however assume that the conditional distribution of S_t given $\mathcal{F}_{t-1}^{S,\eta}$ is not degenerated in the following sense.

A8 for j_0 defined in **A7** and a constant $\tau > 0$, we have $E(S_{j_0} \mid \mathcal{F}_{j_0-1}^{S,\eta}) \ge \tau$ a.s.

Theorem 2.1. *Let (I_t) be a given bitstream sequence, and (ϵ_t) a zero-inflated GARCH(1,1) model satisfying (2.1). Under* **A1**-**A8**, *the estimator defined by (2.3) satisfies $\widehat{\theta}_n \to \theta_0$ almost surely.*

For the asymptotic normality of the QMLE, it is necessary to assume the following.

A9 $\theta_0 \in \overset{\circ}{\Theta}$, where $\overset{\circ}{\Theta}$ denotes the interior of Θ.

A10 $\kappa := E\eta_t^4 < \infty$.

Assumptions **A9** and **A10** are also required to show the asymptotic normality of the QMLE of standard stationary GARCH models (see *e.g.* Theorem 7.2 in Francq and Zakoïan, 2019). Under **A5**, let $\sigma_t^2(\theta) = \sum_{i=0}^{\infty} \beta^i(\omega + \alpha\epsilon_{t-i-1}^2)$. Let $\ell_t(\theta)$ and $\boldsymbol{l}_n(\theta)$ be defined by substituting $\sigma_t(\theta)$ for $\widetilde{\sigma}_t(\theta)$ in $\widetilde{\ell}_t(\theta)$ and $\widetilde{\boldsymbol{l}}_n(\theta)$. Note that

$$
\begin{aligned}
\frac{\partial \ell_t(\theta)}{\partial \theta} &= \dot{\ell}(\theta; I_t, I_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots), \\
\frac{\partial^2 \ell_t(\theta)}{\partial \theta \partial \theta'} &= \ddot{\ell}(\theta; I_t, I_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots)
\end{aligned}
$$

for some measurable functions $\dot{\ell} : \Theta \times \{0,1\}^{\mathbb{N}} \times \mathbb{R}^{\mathbb{N}} \to \mathbb{R}^3$ and $\ddot{\ell} : \Theta \times \{0,1\}^{\mathbb{N}} \times \mathbb{R}^{\mathbb{N}} \to \mathbb{R}^3 \times \mathbb{R}^3$. Let

$$
\begin{aligned}
I &= \frac{1}{T} \sum_{t=1}^{T} E\dot{\ell}(\theta_0; S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots)\dot{\ell}'(\theta_0; S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots), \\
J &= \frac{1}{T} \sum_{t=1}^{T} E\ddot{\ell}(\theta_0; S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots).
\end{aligned}
$$

Theorem 2.2. *Let (I_t) be a given bitstream sequence, and (ϵ_t) a zero-inflated GARCH(1,1) model satisfying (2.1). Under* **A1**-**A10**, *the estimator defined by (2.3)* *satisfies*

$$\sqrt{n}\left(\widehat{\theta}_n-\theta_0\right)=J^{-1}\frac{1}{\sqrt{n}}\sum_{t=1}^{n}I_t\left(\eta_t^2-1\right)\frac{1}{\sigma_t^2}\frac{\partial\sigma_t^2(\theta_0)}{\partial\theta}+o_P(1)$$
$$\xrightarrow{\mathcal{L}}\mathcal{N}\left(0,\Sigma\right),\qquad \Sigma=(\kappa-1)J^{-1}. \tag{2.4}$$

The matrix Σ *can be consistently estimated by* $\widehat{\Sigma}=(\widehat{\kappa}-1)\widehat{J}^{-1}$*, where*

$$\widehat{\kappa}=\frac{\sum_{t=r_0+1}^{n}I_t\widehat{\eta}_t^4}{\sum_{t=r_0+1}^{n}I_t},\qquad \widehat{J}=\frac{1}{n}\sum_{t=r_0+1}^{n}I_t\frac{1}{\widetilde{\sigma}_t^4(\widehat{\theta}_n)}\frac{\partial\widetilde{\sigma}_t^2(\widehat{\theta}_n)}{\partial\theta}\frac{\partial\widetilde{\sigma}_t^2(\widehat{\theta}_n)}{\partial\theta'},\qquad \widehat{\eta}_t=\frac{\epsilon_t}{\widetilde{\sigma}_t(\widehat{\theta}_n)}.$$

The derivatives involved in $\widehat{J}$ can be computed recursively by

$$\frac{\partial\widetilde{\sigma}_t^2(\widehat{\theta}_n)}{\partial\theta}=\begin{pmatrix}1\\ \epsilon_{t-1}^2\\ \widetilde{\sigma}_{t-1}^2(\widehat{\theta}_n)\end{pmatrix}+\beta\frac{\partial\widetilde{\sigma}_{t-1}^2(\widehat{\theta}_n)}{\partial\theta}\text{ for }t=2,\dots,n,\quad\text{with }\frac{\partial\widetilde{\sigma}_1^2(\widehat{\theta}_n)}{\partial\theta}=\begin{pmatrix}1\\0\\0\end{pmatrix}$$

in the case $\widetilde{\sigma}_1^2(\theta)=\omega$.

3 Extension to general volatility models and model checking tests

We now extend Model (2.1) by considering the general zero-inflated volatility model

$$\epsilon_t=\sigma_t\eta_tI_t,\quad \sigma_t=\sigma_t(\theta_0)=\sigma(\epsilon_{t-1},\epsilon_{t-2},\dots;\theta_0), \tag{3.1}$$

where $\theta_0\in\Theta\subset\mathbb{R}^{d_0}$ and $\sigma:\mathbb{R}^\infty\times\Theta\to(0,\infty)$. Note that this general formulation includes all GARCH(p,q) models, as well as numerous volatility models with asymmetries, such as the Asymmetric Power ARCH model (APARCH) of Ding et al. (1993). Given observations $\epsilon_1,\dots,\epsilon_n$ and arbitrary initial values $\widetilde{\epsilon}_t$ for $t\leq 0$, for

$\theta \in \Theta$ let

$$\widetilde{\sigma}_t(\theta) = \sigma(\epsilon_{t-1}, \epsilon_{t-2}, \dots \epsilon_1, \widetilde{\epsilon}_0, \widetilde{\epsilon}_{-1}, \dots; \theta).$$

Define $\widehat{\theta}_n$ by (2.3) and assume the following.

B1: There exists a finite solution to Model (3.1), which is of the form $\epsilon_t = \mathfrak{e}(\theta_0; I_t, I_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots)$ for some function $\mathfrak{e} : \mathbb{R}^\infty \to \mathbb{R}$. Moreover $\sup_t E\sigma_t^{2s} < \infty$ and $\sup_t E\varsigma_t^{2s}(\theta) < \infty$ for some $s > 0$ and all $\theta \in \Theta$, where $\varsigma_t(\theta) = \sigma(\mathfrak{e}_{t-1}, \mathfrak{e}_{t-2}, \dots; \theta)$ with $\mathfrak{e}_t = \mathfrak{e}(\theta_0; S_t, S_{t-1}, \dots; \eta_t, \eta_{t-1}, \dots)$

B2: For any real sequence (x_i), the function $\theta \mapsto \sigma(x_1, x_2, \dots; \theta)$ is continuous on Θ and belongs to $(\underline{\omega}, \infty]$ for all $\theta \in \Theta$ and for some $\underline{\omega} > 0$. Moreover $S_t(\varsigma_t^2(\theta) - \varsigma_t^2(\theta_0)) = 0$ for $t = 1, \dots, T$ iff $\theta = \theta_0$.

B3: There exist a random variable K measurable with respect to $\{\epsilon_u, u \le 0\}$ and a constant $\rho \in (0,1)$ such that $\sup_{\theta \in \Theta} |\sigma_t(\theta) - \tilde{\sigma}_t(\theta)| \le K\rho^t$.

B4: There exist no non-zero $\lambda \in \mathbb{R}^{d_0}$ such that $S_t \lambda' \frac{\partial \varsigma_t(\theta_0)}{\partial \theta} = 0$ a.s. for $t = 1, \dots, T$.

B5: The function $\theta \mapsto \sigma(x_1, x_2, \dots; \theta)$ has continuous second-order derivatives, and

$$\sup_{\theta \in V(\theta_0)} \left\| \frac{\partial \sigma_t(\theta)}{\partial \theta} - \frac{\partial \tilde{\sigma}_t(\theta)}{\partial \theta} \right\| \le K\rho^t$$

where K and ρ are as in **B3** and $V(\theta_0)$ is some neighbourhood of θ_0.

B6: There exists a neighborhood $V(\theta_0)$ of θ_0 such that, for $t = 1, \dots, T$, the following variables have finite expectation:

$$\sup_{\theta \in V(\theta_0)} \left\| \frac{1}{\varsigma_t(\theta)} \frac{\partial \varsigma_t(\theta)}{\partial \theta} \right\|^4, \quad \sup_{\theta \in V(\theta_0)} \left\| \frac{1}{\varsigma_t(\theta)} \frac{\partial^2 \varsigma_t(\theta)}{\partial \theta \partial \theta'} \right\|^2, \quad \sup_{\theta \in V(\theta_0)} \left| \frac{\varsigma_t(\theta_0)}{\varsigma_t(\theta)} \right|^4.$$

For Model (2.1), we have seen that **B1** is satisfied under **A1**, and the first part of **B2** is satisfied under **A4** and **A6**. The identifiability condition in **B2** and **B4** are entailed by **A7**-**A8**. Assumption **A5** entails **B3** and **B5**, as well as the existence of $\sigma_t(\theta)$ and its derivatives for all $\theta \in \Theta$. Relations (A.8) and (A.10) of the proof of Theorem 2.2 show that **B6** also holds true under the assumptions of Theorem 2.2.

Theorem 3.1. *Let* (I_t) *be a given bitstream sequence, and* (ϵ_t) *a zero-inflated volatility model satisfying* (3.1)*. Under* **B1**-**B3**, **A2**, **A3** *and* **A6**, *we have* $\widehat{\theta}_n \to \theta_0$ *a.s. Assume in addition* **B4**-**B6** *and* **A9**-**A10**, *then the convergence in distribution* (2.4) *holds true.*

A model of the form (3.1) of particular interest is

$$\epsilon_t = \sigma_t \eta_t I_t, \quad \sigma_t^2 = \omega_0 + \sum_{i=1}^{r} \tau_{i0} \mathbb{1}_{\epsilon_{t-i}=0} + \alpha_0 \epsilon_{t-1}^2 + \beta_0 \sigma_{t-1}^2, \tag{3.2}$$

with $\sum_{i=1}^r \tau_{i0} > -\omega_0$ and the same constraints and notations as for (2.1). If $\tau_{i0} > 0$ then zero returns tend to increase the volatility, as could be expected when zero returns reflect liquidity issues, but we do not impose this sign constraint a priori. It is clear that, for identifiability of the τ_{i0} coefficients, it is necessary to assume that (I_t) is well fed in zeros and ones. There exist less trivial reasons for non-identifiability of the parameters. For example, if $P(\epsilon_t = 1 \mid \epsilon_{t-1} = 0) = 1$ and $P(\epsilon_t = 0 \mid \epsilon_{t-1} = 1) = 1$ then all the pairs (τ_{01}, τ_{02}) such that $\tau_{01} + \tau_{02}$ is fixed are equivalent. We thus reinforce **A7** and **A8** by assuming that

A8* for $j_0 \in \{1, \dots, T\}$ and $\tau > 0$, we have $E(S_{j_0-i} \mid \mathcal{F}_{j_0-i-1}^{S,\eta}) \in [\tau, 1-\tau]$ a.s. for $i = 0, 1, \dots, r \vee 1$.

Note that, by convention, Model (3.2) with $r = 0$ corresponds to (2.1). In this case, **A8*** reduces to the conditions $E(S_{j_0} \mid \mathcal{F}_{j_0-1}^{S,\eta}) \in (0,1)$ and $E(S_{j_0-1} \mid \mathcal{F}_{j_0-2}^{S,\eta}) \in (0,1)$ a.s., which is an alternative to **A7**-**A8**.

Corollary 3.1. *Let* $\widehat{\theta}_n$ *be the 0-adjusted QMLE of the parameter* $\theta_0 = (\omega_0, \alpha_0, \beta_0, \tau_{10}, \dots, \tau_{r,0})'$ *of Model* (3.2)*. Under* **A1**-**A6**, **A8***, *with obvious changes in* **A4**, *in particular assuming*

$$\Theta \supset \{(\omega, \alpha, \beta, \tau_1, \dots, \tau_r) \in \mathbb{R}^{3+r} : \omega > 0, \alpha > 0, \beta \in [0,1), \textstyle\sum_{i=1}^r \tau_i + \omega > 0\},$$

we have $\widehat{\theta}_n \to \theta_0$ a.s. Assume in addition **A9**-**A10**, *then the convergence in distribution* (2.4) *holds true.*

It is common to assess the adequacy of a time series model by testing the whiteness of the residuals, plotting their empirical (partial) autocorrelations of using formal portmanteau tests, see the monograph by Li (2004). To test the goodness-of-fit of volatility models, Li and Mak (1994) proposed portmanteau tests based on the autocovariances of the squares of the residuals. The asymptotic distribution of these tests has been studied in particular by Berkes et al. (2003a) for the standard GARCH models, Carbon and Francq (2011) for APARCH models, Francq et al. (2018) for Log-GARCH and EGARCH models.

First note that $\widehat{\eta}_t = 0$ when $I_t = 0$, so that $\widehat{\eta}_t$ should only be a good proxy of η_t when $I_t = 1$. Let $n_1 = \sum_{t=r_0+1}^{n} I_t$ and $t_1, \dots, t_{n_1}$ the increasing subsequence of the times $t \in \{r_0+1, \dots, n\}$ such that $I_t = 1$. For fixed integers $h < n_1$ and $m < n_1$, let

$$\widehat{r}_h = \frac{1}{n_1} \sum_{i=h+1}^{n_1} \widehat{s}_{t_i} \widehat{s}_{t_{i-h}}, \qquad \widehat{s}_{t_i} = \widehat{\eta}_{t_i}^2 - 1, \qquad \widehat{\boldsymbol{r}}_m = (\widehat{r}_1, \dots, \widehat{r}_m)'.$$

We will determine the asymptotic distribution of the vector $\widehat{\boldsymbol{r}}_m$ of autocovariances of the squares residuals under the null hypothesis

$$H_0 : \text{the process } (\epsilon_t) \text{ satisfies (3.1)}.$$

Define the $m \times d_0$ matrix whose h-th row is

$$\widehat{\mathbb{K}}_m(h, \cdot) = \frac{1}{n_1} \sum_{i=h+1}^{n_1} \widehat{s}_{t_{i-h}} \frac{\partial \log \widetilde{\sigma}_{t_i}^2(\widehat{\theta}_n)}{\partial \theta'}. \tag{3.3}$$

A random variable of the form $\lambda \eta_{t-i}^2 + \mu' \log \partial \sigma_t^2(\theta)/\partial \theta$ is $\mathcal{F}_{t-1}$-measurable, but in general it is not $\mathcal{F}_{t-i-1}$-measurable. In particular, it is shown in appendix that the following assumption is satisfied under the assumptions of Corollary 3.1.

B7: If $\lambda\eta_{t-i}^2 + \mu'\partial\log\sigma_t^2(\theta_0)/\partial\theta = 0$ a.s. for $i \geq 1$ and $t \geq 1$ then $\lambda = 0$.

Let $\mathbb{I}_m$ the identity matrix of size m and $p_1 = T^{-1}\sum_{t=1}^T P(S_t = 1)$ the asymptotic proportion of 1's in the bitstream sequence, which can be estimated by $\widehat{p}_1 = n_1/n$.

Theorem 3.2. *Under H_0, the assumptions of Theorem 3.1 and* **B7** *we have*

$$T_n := n_1 \widehat{\boldsymbol{r}}'_m \widehat{D}^{-1} \widehat{\boldsymbol{r}}_m \xrightarrow{d} \chi_m^2$$

where $\widehat{D} = (\widehat{\kappa}-1)^2\mathbb{I}_m - 2(\widehat{\kappa}-1)\widehat{p}_1\widehat{\mathbb{K}}_m\widehat{J}^{-1}\widehat{\mathbb{K}}'_m + (\widehat{\kappa}-1)\widehat{p}_1\widehat{\mathbb{K}}_m\widehat{J}^{-1}\widehat{\mathbb{K}}'_m$.

It can be seen that an alternative consistent estimator of D is the empirical variance of $\Upsilon_{t_1}, \dots, \Upsilon_{t_{n_1}}$, where

$$\Upsilon_{t_i} = \widehat{s}_{t_i}\left\{\widehat{\boldsymbol{s}}_{t_{i-1}} - \frac{n_1}{n}\widehat{\mathbb{K}}_m\widehat{J}^{-1}\frac{\partial\log\widetilde{\sigma}_{t_i}^2(\widehat{\theta}_n)}{\partial\theta'}\right\}, \qquad \widehat{\boldsymbol{s}}_{t_{i-1}} = (\widehat{s}_{t_{i-1}}, \dots, s_{t_{i-m}})'.$$

The portmanteau test of Li and Mak (1994) consists in rejecting H_0 at the asymptotic level $\underline{\alpha} \in (0,1)$ if $\{T_n > \chi_m^2(1-\underline{\alpha})\}$, where $\chi_m^2(\underline{\alpha})$ is the $\underline{\alpha}$-quantile of the χ_m^2 distribution.

4 Simulations

To study the finite sample properties of the 0-adjusted QMLE, we undertake a set of Monte Carlo simulations. In the simulations the GARCH specifications are nested in

$$\begin{aligned} \epsilon_t &= \sigma_t\eta_t I_t, \qquad \eta_t \sim iid(0,1), \qquad t = 1,2,\dots,n, &(4.1)\\ \sigma_t^2 &= \omega_0 + \alpha_0\epsilon_{t-1}^2 + \beta_0\sigma_{t-1}^2 + \tau_0 1_{\{\epsilon_{t-1}=0\}}, &(4.2)\\ (\omega_0,\alpha_0,\beta_0,\tau_0) &= (0.2, 0.1, 0.8, 1.0), &(4.3) \end{aligned}$$

where equation (4.2) is a particular case of model (3.2) for which the 0-adjusted QMLE is studied in Corollary 3.1. The parameter values correspond (approximately) to the median values of the estimates in Table 3. The zero-probability $\pi_{0t} = Pr(I_t = 0)$ is governed by one of the following DGPs:

$$\begin{aligned}
&\text{DGP 1:} \quad \pi_{0t} = 0 \text{ for all } t. \\
&\text{DGP 2:} \quad \pi_{0t} = \begin{cases} 0.5 - (t-1)\cdot 0.49/(n\cdot 0.7) & \text{if } t \leq n\cdot 0.7 \\ 0.05 & \text{if } t > n\cdot 0.7 \end{cases} \\
&\text{DGP 3:} \quad \pi_{0t} = 0.1 \text{ if } t \text{ is odd and } \pi_{0t} = 0.4 \text{ if } t \text{ is pair.}
\end{aligned}$$

This means $\{I_t\}$ is stationary in DGP 1, but not in DGPs 2 and 3. In DGP 2 the zero-probability π_{0t} is downwards trending in a way that is characteristic among the daily returns of Section 5.1, see Figure 1. For $t = 1$ the probability is $\pi_{0t} = 0.5$, and then it declines until $t = n \cdot 0.7$, i.e. at 70% of the sample, where $\pi_{0t} = 0.05$. Thereafter, π_{0t} remains constant and equal to 0.05. This is in line with **A3**. In DGP 3 the zero-probability is periodic – as is common in intraday financial data, and varies between $\pi_{0t} = 0.1$ and $\pi_{0t} = 0.4$ as in our illustration in Section 5.2.

The results for the GARCH(1,1) model are contained in the upper part of Table 1. For comparison, we include the results of the Standard QMLE in addition to the 0-adjusted QMLE. Note that in DGP 1 the two QMLEs – and therefore also their results – are identical. When $n = 10000$, the average finite sample error is 0.004 or less in absolute value for the 0-adjusted QMLE. For the Standard QMLE, by contrast, the finite sample error ranges from 0.02 to 0.13 (in absolute value) in DGP 2, and from 0.01 to 0.056 (in absolute value) in DGP 3. This can be substantial in empirical applications. The asymptotic standard errors of the 0-adjusted QMLE are contained in the columns labelled $ase(.)$, see the supplemental appendix for their computation. The values correspond well to their empirical counterparts – contained in the columns labelled $se(.)$, since they differ a maximum of 0.001 (in absolute value) across the DGPs. When $n = 3000$, the 0-adjusted QMLE also produces substantially less biased estimates than the ordinary QMLE, and the empirical standard errors correspond reasonably well to their asymptotic counterparts. The only exception is β in DGP 3, where the Standard QMLE is slightly less biased.

The results for the GARCH(1,1) model with the lagged zero-indicator as covariate are contained in the lower part of Table 1. Note that simulations under DGP 1 is not possible due to exact colinearity. Qualitatively, the simulation results are similar to those of the plain GARCH(1,1). When $n = 10000$, the average finite sample bias is low in absolute value for the 0-adjusted QMLE (0.006 or less), whereas it is high for the Standard QMLE (0.010 to about 0.504 in absolute value). The largest bias is for τ_0 in DGP 2. The empirical standard errors of the 0-adjusted QMLE again correspond quite well to the asymptotic ones, since the bias is always 0.002 or less in absolute value. When $n = 3000$, the average finite sample bias is 0.016 or lower in absolute value for the 0-adjusted QMLE, and the associated discrepancy between the empirical standard errors and the asymptotic ones are never larger than 0.009 in absolute value. In other words, in these experiments the finite sample properties of the 0-adjusted QMLE are also quite good. Similarly, the biases of the Standard QMLE are again quite large, since they range from 0.011 to about 0.495 in absolute value. Also here is the largest bias for τ_0.

5 Empirical illustrations

Standard estimators of volatility, e.g. the Standard QMLE, provide estimates of the conditional variance. The volatility σ_t^2 in our model, by contrast, is not at the same scale-level. To facilitate comparison, the conditional variance representation of our model is therefore obtained as

$$E(\epsilon_t^2|\mathcal{F}_{t-1}) = \sigma_t^2 E(I_t \eta_t^2|\mathcal{F}_{t-1}) = \sigma_t^2 \pi_{1t}, \tag{5.1}$$

where $\pi_{1t} = Pr(I_t = 1|\mathcal{F}_{t-1})$, recall Remark 2.1. In other words, the conditional variance representation can be written as a GARCH with time-varying parameters. In particular, in the case of a GARCH(1,1) with the lagged 0-indicator as covariate, the conditional variance representation is

$$\sigma^2_{t,0adj} = \omega_{0t} + \alpha_{0t}\epsilon^2_{t-1} + \beta_{0t}\sigma^2_{t-1,0adj} + \tau_{0t}1_{\{\epsilon_{t-1}=0\}}, \tag{5.2}$$

where

$$\sigma^2_{t,0adj} = \pi_{1t}\sigma^2_t, \quad \omega_{0t} = \pi_{1t}\omega_0, \quad \alpha_{0t} = \pi_{1t}\alpha_0, \quad \beta_{0t} = \frac{\pi_{1t}}{\pi_{1,t-1}}\beta_0, \quad \tau_{0t} = \pi_{1t}\tau_0. \tag{5.3}$$

A higher value on the zero-probability $\pi_{0t} = 1 - \pi_{1t}$ thus implies a lower "volatility-level" ω_{0t}, a lower "sensitivity" α_{0t} to non-zero price increments in the previous period, a lower impact β_{0t} from the conditional variance (i.e. $\sigma^2_{t-1,0adj}$) in the previous period, and a lower impact τ_{0t} from a zero-return in the previous period. Note also that, when the change in π_{1t} from $t-1$ to t is sufficiently small, then $\beta_{0t} \approx \beta_0$.

5.1 Daily returns at the NYSE

We revisit a subset of the NYSE stocks studied in Sucarrat and Grønneberg (2020). The subset of stocks, 24 in total, together with descriptive statistics of their daily returns, are contained in Table 2. The daily returns are computed as $\epsilon_t = 100 \cdot (\ln S_t - \ln S_{t-1})$, where S_t is the closing price of the stock in question at day t. The datasource is Bloomberg. To be included in the subset, the NYSE stock must satisfy four criteria. First, at least $n = 1000$ daily price observations must be available over the period 3 January 2007 – 4 February 2019. Second, the proportion of zero returns must be greater than 10% over the available sample. Third, a moving average ($n = 500$) estimate of the zero-probability should clearly indicate that the zero-process is nonstationary. Graphs of the moving averages are contained in Figure 1. One of our anonymous reviewers suggested that a trend-like evolution in the zero-probability may be due to a corresponding trend-like evolution in the price level: The higher (lower) the nominal price, the lower (higher) the zero-probability due to discrete price changes. Plots of the prices (see the supplemental appendix) suggest such an effect may indeed be present in several of the stocks. Finally, the fourth criterion is that the graphs suggest assumption **A3** holds.

GARCH estimates of the daily returns obtained with the 0-adjusted QMLE are contained in Table 3. As noted above, the estimates are not directly comparable to standard GARCH estimates – recall (5.1) and (5.2), and must therefore be adjusted before comparison. As an example, suppose the estimate on the ARCH coefficient α is 0.375 (as for the CPS stock) and that the zero-probability at t is $\pi_{0t} = 0.3$. Then the estimate of the time-varying ARCH-coefficient α_t in the conditional variance representation is obtained as $\widehat{\alpha}_t = \pi_{1t}\widehat{\alpha} = (1 - \pi_{0t})\widehat{\alpha} = 0.263$. In periods where the zero-probability is 0, the estimates can be interpreted as those of the conditional variance representation. Figure 2 contains estimates of coefficients in the conditional variance representation for different values on the zero-probability π_{0t}. The vertical lines in the plots are 95% Confidence Intervals (CIs) of the estimates. When $\pi_{0t} = 0$, nine of the estimates of α_t lie in the 0.1 to 0.4 range. This is markedly higher than the typical estimate of a stationary and liquid index or stock, whose estimate is typically below 0.1. This suggests the volatility of this type of stocks can be much more sensitive to price changes at $t-1$ (when π_{0t} is zero or close to zero). But more studies are needed before firm conclusions of general nature can be made. As π_{0t} increases to 0.6, almost all estimates go below 0.1. The estimates of β_t are obtained under the assumption that $\pi_{0t} = \pi_{0,t-1}$. This is why they do not change with π_{0t} in the plots. Four estimates are lower than 0.7. For liquid indices and stocks, they are typically above 0.8. All-in-all, therefore, the plots do not suggest the estimates of β_t tend to be very different from those of liquid indices and stocks. The estimates of τ_t provide an indication of whether a zero in the previous period tends to increase ($\tau > 0$) or decrease ($\tau < 0$) volatility in the next period. In 9 out of 24 stocks the 95% CIs do not contain the value 0 (see also Table 3), so the hypothesis of an effect is supported in these cases (at 5%). For one of these stocks the effect is estimated to be negative, whereas for the other 8 it is estimated to be positive. Finally, the portmanteau test in the final column suggest there is room for improvement (at the 10% significance level) in two of the stocks.

Let $\widehat{\sigma}_{t,0adj}$ denote the estimated conditional standard deviation of our 0-adjusted QMLE, and let $\widehat{\sigma}_t$ denote the estimated conditional standard deviation of the Standard QMLE. To investigate the properties of their discrepancy, we study the distance $x_t = \widehat{\sigma}_{t,0adj} - \widehat{\sigma}_t$. To obtain an estimate of $\sigma_{t,0adj}$, an estimate of π_{1t} is needed. To this end, we device a nonstationary smoothing filter based on the first order Autoregressive Conditional Logit (ACL) of Russell and Engle (2005). Specifically, the smoothing filter is specified as

$$\pi_{1t} = \frac{1}{1+\exp(-h_t)}, \quad h_t = \phi w_{t-1} + h_{t-1}, \quad w_t = \frac{I_t - \pi_{1t}}{\sqrt{\pi_{1t}(1-\pi_{1t})}}, \quad \phi = 0.01, \quad (5.4)$$

where $\phi > 0$ controls the smoothness: The closer to 0, the smoother. Instead of fixing ϕ to 0.01 (as we do), one could instead consider estimating it by, say, maximum likelihood. However, this leads to considerably more erratic paths of π_{1t} for our stocks. Plots of (5.4) against the moving averages from Figure 1, and GARCH estimates of the Standard QMLE, are both contained in the supplemental appendix. Table 4 reports the properties of x_t. The first column contains the results of a test of whether $E(x_t) \neq 0$. The test is implemented via the regression $x_t = \mu + u_t$ with Newey and West (1987) standard errors, $H_0 : \mu = 0$ and $H_A : \mu \neq 0$. In all but four cases is the null rejected at 5%. So the results provide comprehensive support in favour of the alternative hypothesis that the volatility paths differ significantly. In all of the significant cases, the average of x_t is negative. So the Standard QMLE tends to provide volatility estimates that are too high, on average, for the stocks we consider. The next two columns contain the maximum and minimum values of x_t, respectively. These provide an indication of how the conditional volatilities differ on a day-to-day basis. As is clear, they show that the discrepancy can be huge, since they range from -20.9 (the lowest minimum) to 4.3 (the highest maximum). This can have important implications for risk and hedging purposes.

5.2 Intraday 5-minute USD/EUR returns

Intraday financial returns are frequently characterised by a periodic nonstationary zero-process, see e.g. Kolokolov et al. (2020). An example is the intraday 5-minute USD/EUR exchange rate return. Let S_t denote the exchange rate at the end of a 5-minute interval, and let r_t denote the log-return in basis points from the end of one interval to the end of the next: $\epsilon_t = 100^2 \cdot (\ln S_t - \ln S_{t-1})$. The left graph in Figure 3 contains the returns from 2 January 2017 to 31 December 2018, a total of $n = 147\,347$ returns. The source of the data is Forexite. Only trading days are included in the sample, and a typical trading day contains $24 \times 12 = 288$ returns. The first return of a trading day covers the interval from 00:00 CET to 00:05 CET, whereas the last covers 23:55 CET to 00:00 CET. The upper part of Table 5 contains descriptive statistics of the returns. As usual, the returns are characterised by excess kurtosis relative to the normal distribution, and 1st. order autocorrelation in ϵ_t^2. The proportion of zero-returns over the sample is 20.3%, and the right graph of Figure 3 depicts how the zero-proportion varies intradaily across the 24-hour trading day. In the beginning of the day, only the Asian markets are active, so the zero-probability is higher. As European markets open, activity increases and so the zero-probability falls. The zero-probability remains low until the close of the European markets, and then gradually increases again as only the American markets remain active. The zero-probability reaches its peak at the close of the American markets.

The middle part of Table 5 contains the GARCH estimates. In both the Standard and 0-adjusted cases, τ is estimated to be negative, and the 95% CIs for τ do not contain the value 0. In other words, the results suggest a zero-return in the previous period tends to reduce volatility in the next period at the 5-minute frequency for this exchange rate during the sample period of the data at the trading platform in question. To obtain estimates of π_{1t} and π_{0t}, we use a centred moving average of length 12 – i.e. one hour of trading – made up of the intradaily zero-proportions of the 5-minute intervals. The zero-proportions over the trading day, together with the estimate $\widehat{\pi}_{0t}$, are both depicted in the right graph of Figure 3. Note that the periodic cycle is 288. Figure 4 contains the estimates of the time-varying parameters implied by (5.2) together with the estimates of the Standard QMLE. As is clear, the Standard estimate of ω is biased downwards throughout the day, and it is also outside the 95% CI throughout the day. The Standard estimate of α is biased upwards throughout the day, and most of the time outside the 95% CI. The intraday evolution of the 0-adjusted estimate $\widehat{\alpha}_t$ is similar to that of $\widehat{\omega}_t$: It is at its highest in the middle of the day when trading is at its highest, and at its lowest in the beginning and end of the day when trading is at its thinnest. The 0-adjusted estimate of β_t oscillates about the Standard estimate of 0.857, and only in a couple of instances is the Standard estimate outside the 95% CI. The estimates of τ are both negative. The Standard estimate is biased upwards, but it is always within the 95% CI of the zero-adjusted estimate. So they are not significantly different from each other at 5%.

One of our anonymous reviewers asked us to compare the estimates of the 0-adjusted GARCH, which is of observed return, with those of a GARCH model of the efficient return process as defined in Bandi et al. (2020). There, zeros occur when the efficient return process is unobserved. To this end, we derive a modified version of the moment-based estimator of Kristensen and Linton (2006), see supplemental appendix D for the details. The estimates are also contained in the middle part of Table 5.

Note that an estimate of τ is not available for this estimator. Compared with the 0-adjusted estimates depicted in Figure 4, the ω and α estimates are lower, whereas the estimate of β is higher. The α and β estimates of 0.019 and 0.958, respectively, are particularly different, since they are always substantially outside the 95% CIs of the 0-adjusted estimates.

To investigate to what extent the Standard and 0-adjusted QMLEs produce different volatility estimates, we study the distance $x_t = \widehat{\sigma}_{t,0adj} - \widehat{\sigma}_t$, just as in Section 5.1 above. The lower part of Table 5 reports the properties of x_t. Again the test of whether $E(x_t) = 0$ or not is implemented via the regression $x_t = \mu + u_t$ with Newey and West (1987) standard errors. The average of x_t is -0.037, and a two-sided test with 0 as null is rejected at all the usual significance levels. Accordingly, the results suggests the Standard QMLE produces conditional volatilities that are too high, on average. Unconditionally, the value of -0.037 is not large. Conditionally, the range between the maximum and minimum values of x_t suggests the discrepancy can be large on a day-to-day basis. Figure 5 contains the graph of x_t. Most of the time x_t lies between 0.3 and -1.0. Recalling that the 5-minute returns are expressed in basis points, these differences do not appear to be large in economic terms.

6 Conclusions

Financial time series are frequently nonstationary due to a nonstationary zero-process. In these situations, standard estimators are not consistent. We propose a GARCH model that accommodates a nonstationary zero-process, and derive a 0-adjusted QMLE. The nonstationary zero-process can either be trend-like in nature, as is common in daily data, or periodic, as is common in intraday data, or both. The volatility specification in our model can contain higher order ARCH and GARCH terms, asymmetry terms ("leverage") and past zero-indicators as covariates. The latter is of special interest in the current context, since it enables us to study the effect of a zero return on volatility in the subsequent period. Consistency and asymptotic normality of the 0-adjusted QMLE is proved under mild assumptions. Moreover, under stationarity of the zero-process the estimator will still be CAN, so there is no harm in applying our estimator under stationarity. Finite sample simulations verify that the estimator has good finite sample properties, and confirm that the Standard QMLE is biased when the zero-process is nonstationary. Two empirical studies illustrate our results. One is on 24 daily stock returns at NYSE, and one is on intraday 5-minute USD/EUR exchange rate returns. In both studies we find that the time-varying zero-probability affects the dynamics in substantial ways, that the fitted volatilities can differ significantly, and that a zero-return in the previous day can have a substantial effect on volatility in the subsequent day. Interestingly, however, we do not always find that the effect is positive.

While a nonstationary zero-process is frequent in financial time-series, only recently have researchers directed their attention towards this characteristic. Several lines of future research suggest themselves. First, the extension to more general volatility models outlined in Section 3 accommodates models with asymmetry ("leverage"). An interesting line of further research is to study how the evolution of the zero-probability impacts on the effect of asymmetry. Second, it is well-known that financial time series – both daily and intradaily – can be nonstationary due to changes in the level of the unconditional volatility. How frequent are such changes due to a nonstationary zero-process? To the best of our knowledge, this has not been investigated before. Third, to obtain the conditional variance representation of our model, estimates of the time-varying probabilities of a nonstationary zero-process is required. This is challenging. More research is needed to ascertain what the most suitable approach is, and under which assumptions. Fourth, as noted by by one of our anonymous reviewers, the zero-process may not be the only source of nonstationarity. In addition, the volatility intercept (ω), and the ARCH and GARCH parameters, may also be time-varying. To the best of our knowledge, nobody has developed methods for situations where both types of nonstationarities are present. Finally, knowledge about the relation between observed zeros and the underlying efficient return process is limited, so more research on this is needed.

Acknowledgements

We are grateful to the Editor, the Associate Editor and two anonymous reviewers for their helpful comments and suggestions. Christian Francq is grateful to the Agence Nationale de la Recherche (ANR), which supported this work via the Project MLE-forRisk (ANR-21-CE26-0007-01) and also thanks the labex ECODEC.

References

Azrak, R. and G. Mélard (2006). Asymptotic Properties of Quasi-Maximum Likelihood Estimators for ARMA Models with Time-Dependent Coefficients. *Statistical Inference for Stochastic Processes 9*, 279–330.

Bandi, F. M., A. Kolokolov, D. Pirino, and R. Reno (2020). Zeros. *Management Science*. Forthcoming. DOI: https://doi.org/10.1287/mnsc.2019.3527.

Bandi, F. M., D. Pirino, and R. Reno (2017). Excess idle time. *Econometrica 85*, 1793–1846.

Berkes, I., L. Horváth, and P. Kokoszka (2003a). Asymptotics for GARCH squared residual correlations. *Econometric Theory 19*, 515–540.

Berkes, I., L. Horváth, and P. Kokoszka (2003b). GARCH processes: structure and estimation. *Bernoulli 9*, 201–227.

Bien, K., I. Nolte, and W. Pohlmeier (2011). An inflated multivariate integer count hurdle model: an application to bid and ask quote dynamics. *Journal of Applied Econometrics 26*, 669–707.

Buccheri, G., G. Livieri, D. Pirino, and A. Pollastri (2020). A closed-form formula characterization of the Epps effect. *Quantitative Finance 20*, 243–254.

Buccheri, G., D. Pirino, and L. Trapi (2020). Managing liquidity with portfolio staleness. *Decisions in Economics and Finance*. DOI: https://doi.org/10.1007/s10203-020-00300-z.

Carbon, M. and C. Francq (2011). Portmanteau goodness-of-fit test for asymmetric power GARCH models. *Austrian Journal of Statistics 40*, 55–64.

Catania, L., R. Di Maria, and P. Santucci de Magistris (2020). Dynamic Discrete Mixtures for High-Frequency Prices. *Journal of Business and Economic Statistics*. DOI: https://doi.org/10.1080/07350015.2020.1840994.

Ding, Z., C. Granger, and R. Engle (1993). A long memory property of stock market returns and a new model. *Journal of Empirical Finance 1*, 83–106.

Engle, R. F. and J. R. Russell (1998). Autoregressive Conditional Duration: A New Model of Irregularly Spaced Transaction Data. *Econometrica 66*, 1127–1162.

Escanciano, J. C. (2009). Quasi-maximum likelihood estimation of semi-strong GARCH models. *Econometric Theory 25*, 561–570.

Francq, C. and A. Gautier (2004). Large sample properties of parameter least squares estimates for time-varying arma-models. *Journal of Time Series Analysis 25*, 765–783.

Francq, C., O. Wintenberger, and J.-M. Zakoïan (2018). Goodness-of-fit tests for log and exponential GARCH models. *TEST 27*, 27–51. http://dx.doi.org/10.1007/s11749-016-0506-2.

Francq, C. and J.-M. Zakoïan (2019). *GARCH Models*. New York: Wiley. 2nd. Edition.

Harvey, A. C. and R. Ito (2020). Modeling time series when some of the obervations are zero. *Journal of Econometrics 214*, 33–45.

Hausman, J., A. Lo, and A. MacKinlay (1992). An ordered probit analysis of transaction stock prices. *Journal of financial economics 31*, 319–379.

Hautsch, N., P. Malec, and M. Schienle (2013). Capturing the zero: a new class of zero-augmented distributions and multiplicative error processes. *Journal of Financial Econometrics 12*, 89–121.

Kolokolov, A., G. Livieri, and D. Pirino (2020). Statistical inferences for price staleness. *Journal of Econometrics 218*, 32–81.

Kolokolov, A. and R. Reno (2019). Jumps or staleness. Working paper available at SSRN: http://dx.doi.org/10.2139/ssrn.3208204.

Kristensen, D. and O. Linton (2006). A Closed-Form Estimator for the GARCH(1,1) Model. *Econometric Theory 22*, 323–337.

Kümm, H. and U. Küsters (2015). Forecasting zero-inflated price changes with a markov switching mixture model for autoregressive and heteroscedastic time series. *International Journal of Forecasting 31*, 598–608.

Lesmond, D. A., J. P. Ogden, and C. A. Trzcinka (1999). A New Estimate of Transaction Costs. *The Review of Financial Studies 12*, 1113–1141.

Li, W. (2004). *Diagnostic checks in time series*. Boca Raton, Florida: Chapman and Hall.

Li, W. and T. Mak (1994). On the squared residual autocorrelations in non-linear time series with conditional heteroscedasticity. *Journal of Time Series Analysis 15*, 627–636.

Liesenfeld, R., I. Nolte, and W. Pohlmeier (2006). Modelling Financial Transaction Price Movements: A Dynamic Integer Count Data Model. *Empirical Economics 30*, 795–825.

Ljung, G. and G. Box (1979). On a Measure of Lack of Fit in Time Series Models. *Biometrika 66*, 265–270.

Newey, W. and K. West (1987). A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*, 703–708.

Phillips, P. C. and K. Xu (2006). Inference in autoregression under heteroskedasticity. *Journal of Time Series Analysis 27*, 289–308.

Regnard, N. and J.-M. Zakoian (2010). Structure and estimation of a class of nonstationary yet nonexplosive GARCH models. *Journal of Time Series Analysis 31*, 348–364.

Russell, J. R. and R. F. Engle (2005). A Discrete-State Continuous-Time Model of Financial Transaction Prices and Times: The Autoregressive Conditional Multinomial-Autoregressive Conditional Duration Model. *Journal of Business and Economic Statistics 23*, 166–180.

Rydberg, T. H. and N. Shephard (2003). Dynamics of Trade-by-Trade Price Movements: Decomposition and Models. *Journal of Financial Econometrics 1*, 2–25.

Sucarrat, G. and S. Grønneberg (2020). Risk Estimation with a Time Varying Probability of Zero Returns. *Journal of Financial Econometrics*. Forthcoming. DOI: https://doi.org/10.1093/jjfinec/nbaa014.

Table 1: Simulations of the Standard and 0-adjusted QMLEs (Section 4)

n	QMLE	DGP	$m(\widehat{\omega})$	$se(\widehat{\omega})$	$ase(\widehat{\omega})$	$m(\widehat{\alpha})$	$se(\widehat{\alpha})$	$ase(\widehat{\alpha})$	$m(\widehat{\beta})$	$se(\widehat{\beta})$	$ase(\widehat{\beta})$	$m(\widehat{\tau})$	$se(\widehat{\tau})$	$ase(\widehat{\tau})$
GARCH(1,1):														
10000	Stand.	1	0.204	0.027	0.027	0.100	0.009	0.009	0.798	0.019	0.019			
10000	Stand.	2	0.070	0.019	−	0.080	0.012	−	0.870	0.024	−			
10000	Stand.	3	0.144	0.024	−	0.070	0.009	−	0.810	0.026	−			
10000	0-adj.	1	0.204	0.027	0.027	0.100	0.009	0.009	0.798	0.019	0.019			
10000	0-adj.	2	0.201	0.028	0.029	0.100	0.010	0.011	0.799	0.022	0.023			
10000	0-adj.	3	0.203	0.035	0.034	0.100	0.012	0.012	0.798	0.027	0.027			
3000	Stand.	1	0.209	0.052	0.049	0.100	0.016	0.016	0.795	0.036	0.035			
3000	Stand.	2	0.074	0.038	−	0.080	0.021	−	0.868	0.045	−			
3000	Stand.	3	0.151	0.051	−	0.069	0.016	−	0.805	0.053	−			
3000	0-adj.	1	0.209	0.052	0.049	0.100	0.016	0.016	0.795	0.036	0.035			
3000	0-adj.	2	0.214	0.060	0.053	0.101	0.019	0.020	0.791	0.046	0.042			
3000	0-adj.	3	0.218	0.077	0.061	0.102	0.022	0.022	0.787	0.059	0.049			
GARCH(1,1) w/covariate:														
10000	Stand.	2	0.296	0.046	−	0.090	0.009	−	0.751	0.027	−	0.496	0.081	−
10000	Stand.	3	0.163	0.046	−	0.077	0.009	−	0.783	0.026	−	0.868	0.086	−
10000	0-adj.	2	0.205	0.029	0.030	0.100	0.010	0.010	0.798	0.018	0.018	1.006	0.097	0.098
10000	0-adj.	3	0.205	0.053	0.051	0.100	0.011	0.011	0.799	0.022	0.022	1.004	0.099	0.097
3000	Stand.	2	0.308	0.087	−	0.089	0.017	−	0.747	0.051	−	0.505	0.150	−
3000	Stand.	3	0.182	0.111	−	0.077	0.016	−	0.774	0.058	−	0.879	0.163	−
3000	0-adj.	2	0.209	0.056	0.054	0.099	0.018	0.018	0.798	0.033	0.033	1.005	0.178	0.179
3000	0-adj.	3	0.216	0.101	0.092	0.100	0.020	0.020	0.795	0.041	0.039	1.010	0.186	0.178

$m(.)$, average of estimate. $se(.)$, empirical standard error (computed as the sample standard deviation of estimates). $ase(.)$, asymptotic standard error. 1000 replications with sample size n in each replication.

Table 2: Descriptive statistics of NYSE stock returns (Section 5.1)

Ticker	Sample period	n	s^2	s^4	ARCH	P-value	0s	$\widehat{\pi}_0$
ARR	2007-12-03 to 2019-02-04	2812	1.93	17.54	70.35	0.00	352	0.13
BKT	2007-01-03 to 2019-02-04	3043	0.47	25.09	53.91	0.00	439	0.14
CO	2007-01-25 to 2019-02-04	3028	9.21	159.37	119.91	0.00	756	0.25
CPS	2010-05-25 to 2019-02-04	2189	3.35	31.33	2.65	0.10	298	0.14
CUBI	2012-02-21 to 2019-02-04	1750	4.51	29.59	0.01	0.92	223	0.13
DOOR	2009-07-24 to 2019-02-04	2399	5.03	84.62	0.06	0.80	831	0.35
EROS	2009-12-18 to 2019-02-04	2296	12.32	32.05	8.44	0.00	999	0.44
ESTE	2007-01-03 to 2019-02-04	3043	15.76	9.98	66.93	0.00	318	0.10
EVF	2007-01-03 to 2019-02-04	3043	1.30	24.12	79.96	0.00	310	0.10
FF	2008-07-14 to 2019-02-04	2659	6.34	13.97	80.44	0.00	481	0.18
FSB	2013-10-04 to 2019-02-04	1342	4.30	31.21	3.58	0.06	379	0.28
FTS	2007-01-05 to 2019-02-04	3041	1.79	22.08	0.29	0.59	1345	0.44
GNK	2014-07-15 to 2019-02-04	1148	23.28	16.00	30.78	0.00	126	0.11
GPRK	2007-08-29 to 2019-02-04	2878	9.70	44.94	2.08	0.15	1516	0.53
GTT	2007-01-03 to 2019-02-04	3043	22.96	33.03	18.22	0.00	965	0.32
ICL	2007-01-03 to 2019-02-04	3043	7.08	18.28	9.55	0.00	717	0.24
NOMD	2014-09-09 to 2019-02-04	1109	5.22	15.64	223.04	0.00	202	0.18
NVGS	2007-01-09 to 2019-02-04	3039	9.65	595.44	0.00	0.99	1629	0.54
OSB	2007-01-03 to 2019-02-04	3043	15.12	157.31	751.14	0.00	474	0.16
PARR	2012-09-05 to 2019-02-04	1613	17.90	168.14	338.76	0.00	184	0.11
TARO	2007-01-03 to 2019-02-04	3043	3.95	18.47	12.20	0.00	340	0.11
TIER	2014-02-27 to 2019-02-04	1243	6.57	155.46	0.04	0.84	343	0.28
TU	2007-01-16 to 2019-02-04	3035	1.76	28.33	0.25	0.62	1263	0.42
WCN	2007-01-04 to 2019-02-04	3042	3.34	24.29	7.93	0.00	368	0.12

n, the number of observations before lagging and differencing. s^2, sample variance of return. s^4, sample kurtosis of return. ARCH, Ljung and Box (1979) test statistic for first-order autocorrelation in ϵ_t^2, with P-value denoting the associated p-value. 0s, the number of zero returns. $\widehat{\pi}_0$, the proportion of zero returns.

Table 3: GARCH estimates (0-adjusted QMLE) of the daily NYSE returns (Section 5.1)

Ticker	$\widehat{\omega}$ $(s.e.)$	$\widehat{\alpha}$ $(s.e.)$	$\widehat{\beta}$ $(s.e.)$	$\widehat{\tau}$ $(s.e.)$	95% CI for τ Lower	Upper	$\chi^2(2)$ $(p-val)$
ARR	0.099 (0.0273)	0.274 (0.0848)	0.777 (0.0366)	−0.099 (0.0274)	-0.099	-0.045	3.784 (0.1508)
BKT	0.009 (0.0057)	0.105 (0.0285)	0.881 (0.0359)	0.016 (0.0134)	-0.009	0.043	1.240 (0.5379)
CO	3.127 (1.2138)	0.259 (0.1240)	0.473 (0.1648)	−0.885 (1.2726)	-3.127	1.609	1.936 (0.3798)
CPS	1.871 (0.9568)	0.375 (0.1646)	0.217 (0.2945)	1.702 (1.3963)	-1.034	4.439	0.054 (0.9735)
CUBI	0.278 (0.2201)	0.038 (0.0289)	0.872 (0.0816)	2.813 (2.1100)	-0.278	6.948	0.403 (0.8174)
DOOR	0.518 (0.3027)	0.059 (0.0566)	0.746 (0.1132)	10.000 (5.0839)	0.036	19.964	1.093 (0.5790)
EROS	0.308 (0.1425)	0.089 (0.0229)	0.880 (0.0251)	9.306 (2.6782)	4.057	14.556	1.452 (0.4838)
ESTE	0.250 (0.1766)	0.064 (0.0319)	0.928 (0.0353)	0.449 (0.4955)	-0.250	1.421	4.388 (0.1115)
EVF	0.022 (0.0101)	0.166 (0.0343)	0.843 (0.0287)	−0.013 (0.0354)	-0.022	0.057	0.444 (0.8009)
FF	1.017 (0.4266)	0.206 (0.0815)	0.673 (0.1034)	2.954 (1.5116)	-0.009	5.917	0.450 (0.7986)
FSB	0.166 (0.1410)	0.056 (0.0266)	0.903 (0.0499)	1.390 (0.8207)	-0.166	2.999	0.002 (0.9988)
FTS	0.015 (0.0151)	0.123 (0.0287)	0.862 (0.0322)	0.663 (0.1858)	0.298	1.027	0.556 (0.7574)
GNK	0.047 (0.0711)	0.042 (0.0136)	0.960 (0.0111)	0.190 (0.2076)	-0.047	0.597	9.899 (0.0071)
GPRK	1.066 (0.6481)	0.088 (0.0442)	0.799 (0.0945)	10.000 (4.5313)	1.119	18.881	2.980 (0.2254)
GTT	0.025 (0.0513)	0.057 (0.0165)	0.940 (0.0179)	3.648 (1.2871)	1.125	6.171	1.570 (0.4561)
ICL	0.000 (0.0172)	0.042 (0.0110)	0.957 (0.0121)	0.453 (0.1849)	0.091	0.816	0.975 (0.6142)
NOMD	0.126 (0.1259)	0.065 (0.0364)	0.897 (0.0591)	1.971 (1.0741)	-0.126	4.077	6.256 (0.0438)
NVGS	0.696 (0.9002)	0.048 (0.0450)	0.849 (0.1334)	10.000 (8.9374)	-0.696	27.517	1.830 (0.4005)
OSB	0.019 (0.0280)	0.055 (0.0114)	0.948 (0.0113)	0.445 (0.1929)	0.067	0.823	1.791 (0.4084)
PARR	0.068 (0.0429)	0.027 (0.0100)	0.963 (0.0135)	0.085 (0.2461)	-0.068	0.567	0.060 (0.9705)
TARO	1.123 (0.3247)	0.360 (0.0916)	0.400 (0.1146)	5.048 (1.6681)	1.778	8.317	0.808 (0.6676)
TIER	0.014 (0.0362)	0.041 (0.0170)	0.947 (0.0208)	1.515 (0.7548)	0.035	2.994	1.354 (0.5082)
TU	0.122 (0.0715)	0.122 (0.0373)	0.764 (0.0909)	2.593 (1.1298)	0.379	4.808	0.256 (0.8798)
WCN	0.012 (0.0112)	0.031 (0.0122)	0.963 (0.0135)	0.295 (0.1564)	-0.012	0.602	0.217 (0.8973)

0-adjusted QMLEs of $\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \tau 1_{\{\epsilon_{t-1}=0\}}$. $s.e.$, standard error of estimate. Upper bound of 95% CI for τ computed as $\widehat{\tau} + s.e.(\widehat{\tau}) \cdot 1.96$, where $s.e.(\widehat{\tau})$ is the standard error of $\widehat{\tau}$. Lower bound computed as $\max\{-\widehat{\omega}, \widehat{L}\}$, where $\widehat{L} = \widehat{\tau} - s.e.(\widehat{\tau}) \cdot 1.96$. To avoid explosive volatility-paths, the upper bound $\tau \leq 10$ is imposed during estimation. $\chi^2(2)$, the results from the portmanteau test of Section 3 of autocorrelation up to and including order 2 of η_t^2 (p-value in parentheses).

Table 4: Properties of $x_t = \widehat{\sigma}_{t,0adj} - \widehat{\sigma}_t$ for NYSE returns (Section 5.1)

Ticker	Avg x_t	P-value	Max x_t	Min x_t
ARR	0.005	0.434	0.394	-1.492
BKT	-0.001	0.220	0.025	-0.085
CO	-0.199	0.000	0.675	-20.862
CPS	-0.058	0.000	0.584	-1.193
CUBI	-0.266	0.000	0.552	-2.773
DOOR	-0.171	0.000	1.818	-3.792
EROS	-0.865	0.000	4.273	-19.983
ESTE	-0.064	0.000	0.434	-1.493
EVF	0.000	0.567	0.031	-0.065
FF	-0.132	0.000	1.262	-3.774
FSB	-0.582	0.000	0.560	-6.146
FTS	-0.149	0.000	0.445	-1.218
GNK	-0.302	0.000	0.023	-1.612
GPRK	-0.474	0.000	3.518	-16.605
GTT	-0.261	0.000	0.864	-6.134
ICL	-0.150	0.000	0.730	-1.497
NOMD	-0.390	0.000	0.837	-7.627
NVGS	-0.633	0.000	2.660	-20.711
OSB	-0.030	0.051	2.103	-3.251
PARR	-0.153	0.021	2.758	-13.040
TARO	-0.012	0.022	0.571	-1.106
TIER	-0.638	0.000	0.490	-5.629
TU	-0.065	0.000	0.918	-1.562
WCN	-0.103	0.000	0.234	-1.462

Avg x_t, average of x_t. P-value, the p-value of a two sided test with $E(x_t) = 0$ as null. The test is implemented via the regression $x_t = \mu + u_t$ with Newey and West (1987) standard errors, $H_0 : \mu = 0$ and $H_A : \mu \neq 0$.

Table 5: Intraday 5-minute USD/EUR returns (Section 5.2)

Descriptive statistics:

Sample	n	s^2	s^4	ARCH [$p-val$]	0s	$\widehat{\pi}_0$
2017-01-02 to 2018-12-31	147 347	7.84	89.1	113.1 [0.000]	29886	0.203

GARCH estimates:

Estimator	$\widehat{\omega}$ ($s.e.$)	$\widehat{\alpha}$ ($s.e.$)	$\widehat{\beta}$ ($s.e.$)	$\widehat{\tau}$ ($s.e.$)	95% CI for τ		$\chi^2(2)$ [$p-val$]
					Lower	Upper	
Standard	0.201 (0.027)	0.142 (0.008)	0.857 (0.007)	−0.161 (0.042)	−0.201	−0.078	
0-adjusted	0.429 (0.047)	0.143 (0.010)	0.855 (0.008)	−0.305 (0.075)	−0.429	−0.158	4.102 [0.129]
Moment	0.231	0.019	0.958				

Properties of $x_t = \widehat{\sigma}_{t,0adj} - \widehat{\sigma}_t$:

Avg x_t	P-value	Max x_t	Min x_t
−0.034	0.000	0.333	−10.143

n, number of returns. s^2, sample variance of return. s^4, sample kurtosis of return. ARCH, Ljung and Box (1979) test statistic for first-order autocorrelation in ϵ_t^2 (p-value in square brackets). 0s, number of zeros. $\widehat{\pi}_0$, proportion of zeros. $s.e.$, standard error of estimate. 95% CIs computed as $\widehat{\tau} \pm s.e.(\widehat{\tau}) \cdot 1.96$, where $s.e.(\widehat{\tau})$ is the standard error of $\widehat{\tau}$. Lower bound computed as $\max\{-\widehat{\omega}, \widehat{L}\}$, where $\widehat{L} = \widehat{\tau} - s.e.(\widehat{\tau}) \cdot 1.96$. $\chi^2(2)$, the result of the portmanteau test of Section 3 of autocorrelation up to and including order 2 of η_t^2 (p-value in square brackets). Moment, the modified moment-based estimator of Kristensen and Linton (2006), see the supplemental appendix. Avg x_t, average of x_t. P-value, the p-value of a two sided test with $E(x_t) = 0$ as null (implemented via the regression $x_t = \mu + u_t$ with Newey and West (1987) standard errors).

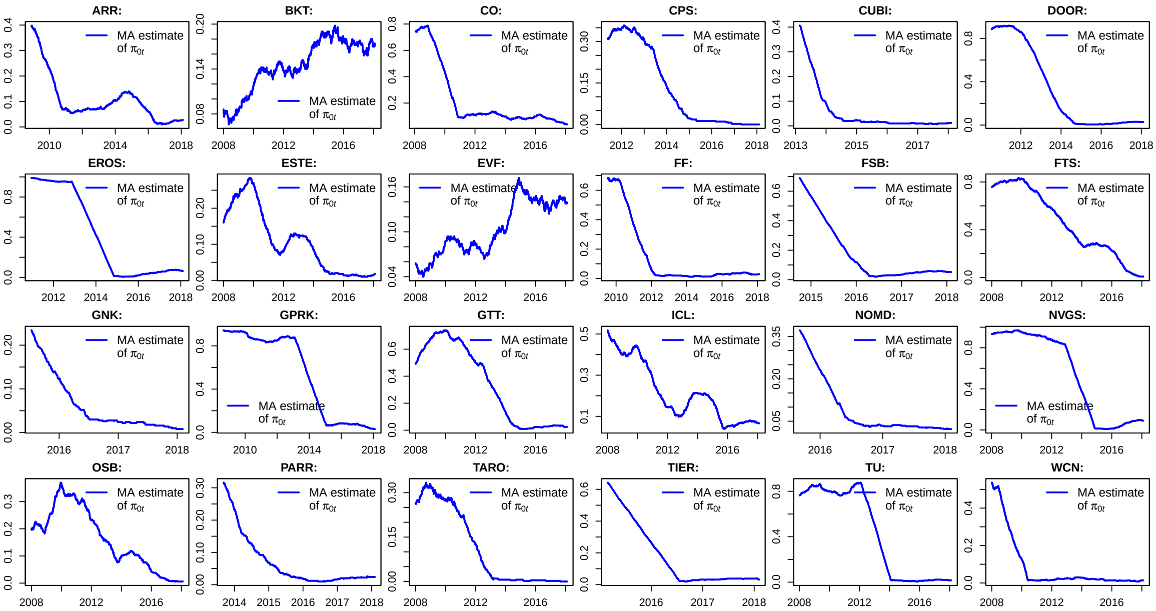

Figure 1: Moving Average (MA) estimates of the daily zero-probability π_{0t} for a subset of NYSE stocks (see Section 5). The moving average is computed as $\widehat{\pi}_{0t} = 500^{-1}\sum_{i=t-249}^{t+250}(1 - I_i)$, $t = 250, \ldots, n - 250$. Datasource: Bloomberg

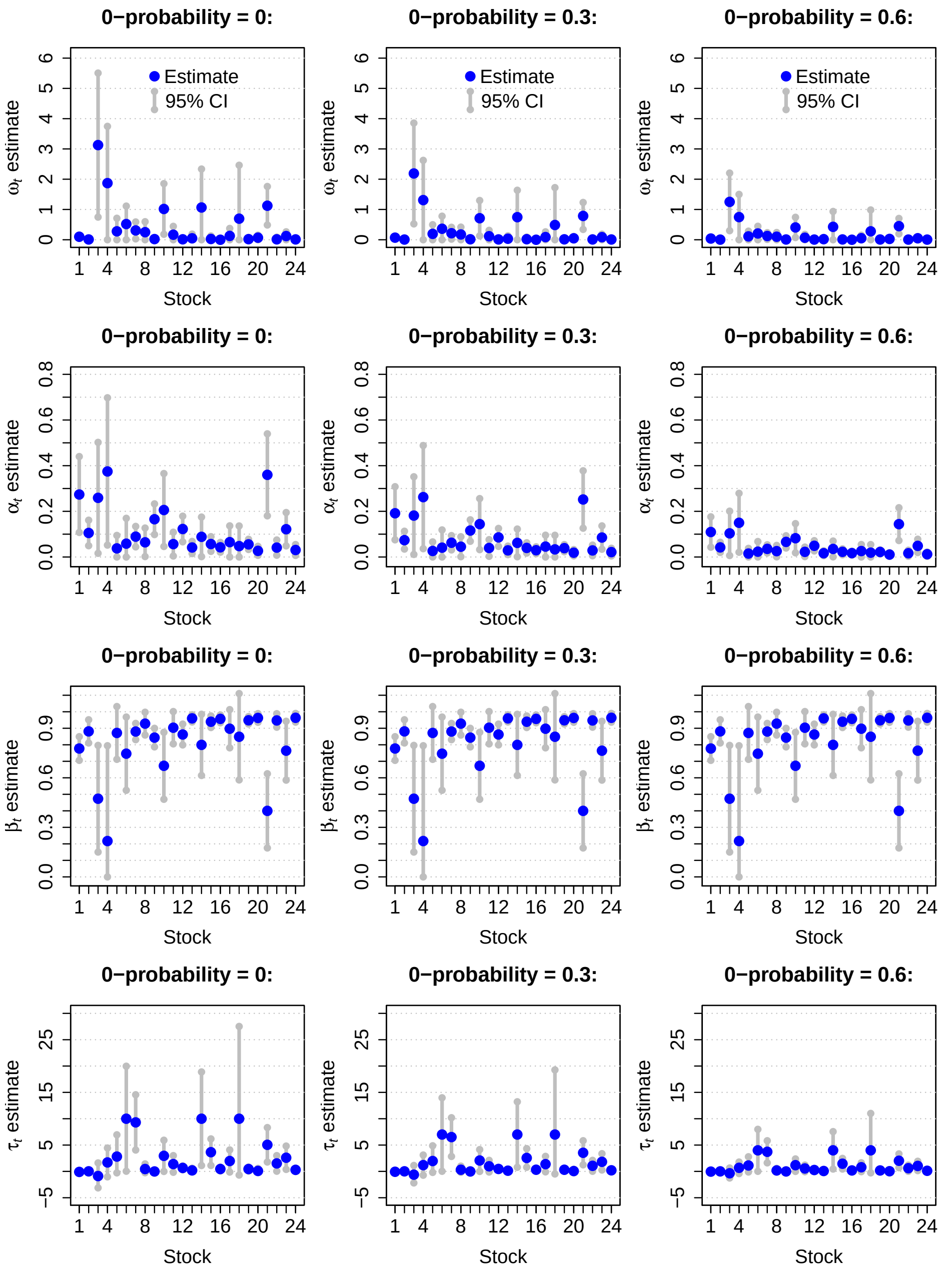

Figure 2: Estimates of the time-varying GARCH coefficients α_t, β_t and τ_t for different values of the 0-probability π_{0t}, together with 95% Confidence Intervals (CIs) (see Section 5.1). Datasource: Bloomberg

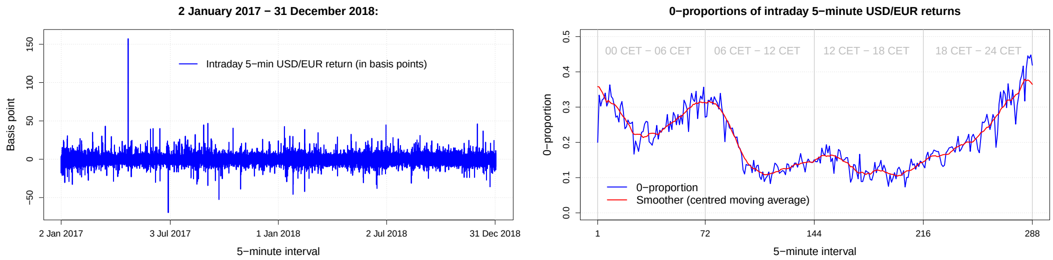

Figure 3: Intraday 5-minute USD/EUR log-returns in basis points (left graph) from 2 January 2017 to 31 December 2018 ($n = 147\ 347$), and the proportion of zero-returns in each intraday 5-minute interval (right graph). The smoother is a centred moving average of length 12 (see Section 5.2). Datasource: Forexite

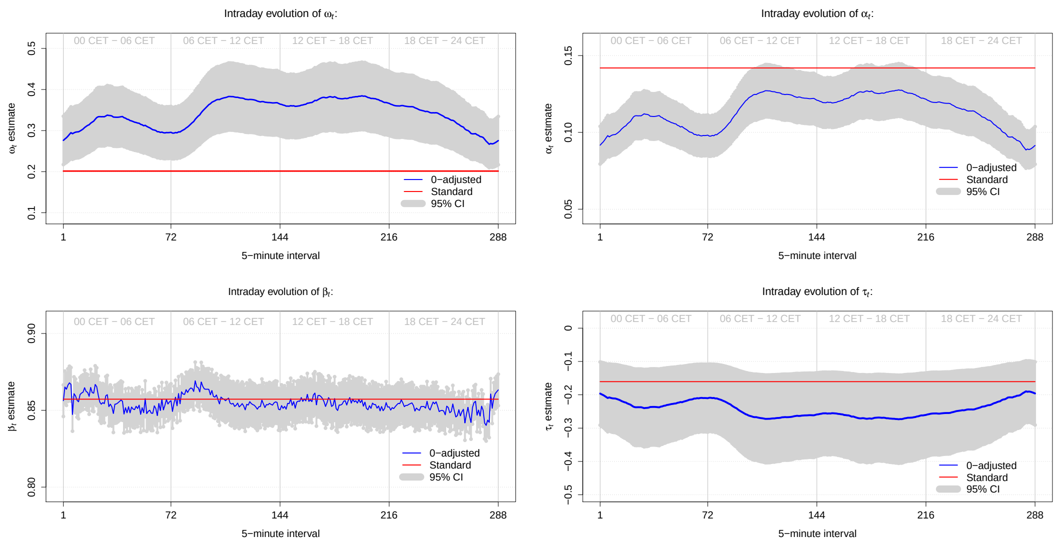

Figure 4: The time-varying intraday evolution of the GARCH coefficients in the conditional variance representation of USD/EUR 5-minute returns (see Section 5.2). Datasource: Forexite

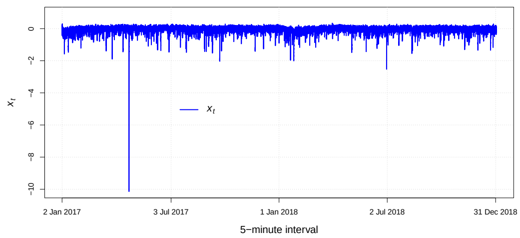

Figure 5: The difference $\widehat{\sigma}_{t,0adj} - \widehat{\sigma}_t$ for intraday USD/EUR 5-minute returns (see Section 5.2). Datasource: Forexite