



Efficiency of the financial markets during the COVID-19 crisis: time-varying parameters of fractional stable dynamics

Ayoub Ammy-Driss, Matthieu Garcin

► To cite this version:

Ayoub Ammy-Driss, Matthieu Garcin. Efficiency of the financial markets during the COVID-19 crisis: time-varying parameters of fractional stable dynamics. 2021. <hal-02903655v3>

HAL Id: hal-02903655

<https://hal.science/hal-02903655v3>

Preprint submitted on 25 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Efficiency of the financial markets during the COVID-19 crisis: time-varying parameters of fractional stable dynamics

Ayoub Ammy-Driss^a, Matthieu Garcin^{b,*}

November 24, 2021

Abstract

This paper investigates the impact of COVID-19 on financial markets. It focuses on the evolution of the market efficiency, using two efficiency indicators: the Hurst exponent and the memory parameter of a fractional Lévy-stable motion. The second approach combines, in the same model of dynamic, an alpha-stable distribution and a dependence structure between price returns. We provide a dynamic estimation method for the two efficiency indicators. This method introduces a free parameter, the discount factor, which we select so as to get the best alpha-stable density forecasts for observed price returns. The application to stock indices during the COVID-19 crisis shows a strong loss of efficiency for US indices. On the opposite, Asian and Australian indices seem less affected and the inefficiency of these markets during the COVID-19 crisis is even questionable.

Keywords – alpha-stable distribution, dynamic estimation, efficient market hypothesis, financial crisis, Hurst exponent

1 Introduction

The COVID-19 pandemic has strongly affected many persons, either for medical reasons or for the economic aftermath of the various prophylactic measures decided by governments, in particular lockdowns. The evolution of financial markets during the pandemic provides an illustration of the economic impact of these measures. According to several empirical studies, financial markets have indeed been strongly disturbed during this period [4, 6, 40, 70]. The question of the reaction of financial markets to a crisis is not specific to the COVID-19 pandemic. For example, we can cite a study of the impact of crisis in the 80s and the 90s on a model of dynamic for a stock market [21]. In general, these studies focus on variations of several statistics, such as jump intensity, implied volatility, parameters of factor models, divergence of price return densities, etc. To the best of our knowledge, no paper focuses on measuring the impact of the COVID-19 pandemic on market efficiency.

Market efficiency is the ability of market prices to reflect all the available information, so that no arbitrage is possible. In other words, if markets are efficient, price returns are not correlated with

*Corresponding author: matthieu.garcin@m4x.org.

^a ESILV, 92916 Paris La Défense, France.

^b Léonard de Vinci Pôle Universitaire, Research center, 92916 Paris La Défense, France.

each other and investors are not able to statistically determine what is more profitable between selling and buying a financial asset. Even if this dogma is sometimes questionable in calm periods, we wonder if it can resist to a crisis. We also wonder if we can observe regional disparities, hypothetically related to the magnitude of the outbreak in these regions, and how fast financial markets recover.

The Hurst exponent is a widespread indicator of market efficiency. The long-range feature associated to a Hurst exponent above $1/2$ is indeed traditionally interpreted as indicating predictability of a time series. It is used in finance [58, 16, 63, 8] but also in many other fields such as meteorology, for models of temperature [22]. It is often related to a model of dynamic with a specific fractal property, namely the fractional Brownian motion (fBm), introduced by Mandelbrot and van Ness [61]. The fractal property states that the variance of increments of duration τ is $\tau^{2H}\sigma^2$, where H is the Hurst exponent and σ^2 the variance of increments of duration 1. The fBm assumes that increments follow a Gaussian distribution. The fractal property of the fBm is then obtained by introducing a positive (respectively negative) correlation among increments if the Hurst exponent is above (resp. below) $1/2$. In the case where the Hurst exponent is $1/2$, the fBm is simply a standard Brownian motion (Bm), that is with independent increments. In the fBm framework, this value of the Hurst exponent is thus the only one consistent with the efficient market hypothesis (EMH).

However, the approach using the fBm assumes Gaussian price returns. It is not very realistic, since the presence of fat tails in the distribution of price returns is well documented [54, 9, 68, 37]. As an alternative to the Gaussian distribution, the alpha-stable distribution is appealing because it includes the Gaussian distribution as a particular case and entails fat tails, whose amplitude is directly related to α , a parameter of this distribution. When combining alpha-stable distributions and dependence between the successive price returns, we can get the fractional Lévy-stable motion (fLsm), which is thus a non-Gaussian extension of the fBm [78, 79, 80]. In this framework, the Hurst exponent is to be decomposed in $m + 1/\alpha$, where m is a memory parameter. If $m = 0$, adjacent price returns are independent and the market is efficient. We thus use this memory parameter as an alternative efficiency indicator.

The estimation of the fLsm partly relies on the estimation of alpha-stable distributions. Many estimation methods exist for this kind of distribution [2, 15]. In particular, we will focus on McCulloch's method, which is based on empirical quantiles [62]. In this paper, we are in addition interested in the dynamic estimation of this distribution, because we want to depict the chronology of the crisis, day after day. Several articles deal with the estimation of time-varying non-parametric densities [43, 40]. The case of parametric densities is simpler as it consists in estimating time-varying parameters. In both the non-parametric and parametric cases, the estimation at a given date takes into account the estimation at the previous date updated by the new observation. The balance between the previous estimation and the new observation is tuned by a discount factor. Several rules are possible for the selection of this free parameter. We focus on the minimization of a criterion related to arguments coming from the field of validation of density forecast [40, 26].

In the empirical part of the paper, we study the evolution of the two indicators of market efficiency, H and m for several stock indices, with a significance analysis. We have discovered that the Hurst exponent H detects less often a market inefficiency than does the memory parameter m of an fLsm. We foster the use of this parameter as efficiency indicator instead of the Hurst exponent. It indeed improves the standard Hurst approach insofar as it filters the kurtosis of price returns, which biases the Hurst indicator.

Besides the analysis of the impact of the COVID-19 on market efficiency, the innovative aspects of this paper include a selection rule for the discount factor of a dynamic parametric distribution, an estimation method of dynamic Hurst exponents, and the introduction of the memory parameter

of an fLsm as efficiency indicator.

The rest of the paper is organized as follows. Section 2 introduces the estimation of dynamic alpha-stable distributions, along with the selection rule of the discount factor. Section 3 provides some elements on market efficiency and details how indicators are built. Section 4 empirically studies the impact of COVID-19 on market efficiency. Section 5 concludes.

2 Alpha-stable distributions

The alpha-stable distribution is a generalization of the Gaussian distribution, appreciated for entailing fat tails. For this reason, it has been widely invoked in signal processing [77, 38, 39], with applications for example in medicine [76, 80] or in finance [54, 9, 68, 37].

We present below the estimation of a time-varying alpha-stable distribution, which we will apply later in this paper to financial time series. For this purpose, we present first the static estimation as well as the various representations of alpha-stable distributions. Regarding the dynamic distribution, a free parameter is to be selected. It is the discount factor. We propose a rule of selection in the last subsection.

2.1 Representation

Four parameters are used to depict a random variable following a stable distribution: $X \sim S_\alpha(\gamma, \beta, \mu)$. The parameter $\alpha \in (0, 2]$ is the one we will mostly be interested in. It determines the thickness of the tails. The parameter $\beta \in [-1, 1]$ is a skewness parameter. If $\alpha = 2$ and $\beta = 0$, we retrieve the Gaussian distribution. The last two parameters stand for the location ($\mu \in \mathbb{R}$) and the scale ($\gamma > 0$) of the distribution. We do not have any analytic expression for the probability density of X , but we can characterize the stable distribution by the mean of its characteristic function:

$$t \mapsto \mathbb{E}[\exp(itX)] = \begin{cases} \exp(i\mu t - \gamma^\alpha |t|^\alpha [1 - i\beta \text{sign}(t) \tan \frac{\pi\alpha}{2}]) & \text{if } \alpha \neq 1 \\ \exp(i\mu t - \gamma |t| [1 + i\beta \text{sign}(t) \frac{2}{\pi} \log |t|]) & \text{if } \alpha = 1. \end{cases}$$

We could use the Fourier transform to get the pdf from the characteristic function [76], but the above parameterization is not totally satisfactory insofar as the pdf is not continuous in the parameters, in particular when $\alpha = 1$ [67, 2]. Indeed, when $\beta > 0$, the density is shifted right when $\alpha < 1$ and left when $\alpha > 1$, with a shift toward $+\infty$ (respectively $-\infty$) when α tends toward 1 by below (resp. above) [67]. For applications to data and interpretation of the coefficients, this parameterization is thus to be avoided. For this reason, Nolan has proposed to use Zolotarev's (M) parameterization [84], which is also often called the S^0 parameterization. The characteristic function corresponding to $X \sim S_\alpha^0(\gamma, \beta, \mu_0)$ is [67, 2]:

$$t \mapsto \mathbb{E}[\exp(itX)] = \begin{cases} \exp(i\mu_0 t - \gamma^\alpha |t|^\alpha [1 + i\beta \text{sign}(t) \tan \frac{\pi\alpha}{2} (\gamma^{1-\alpha} |t|^{1-\alpha} - 1)]) & \text{if } \alpha \neq 1 \\ \exp(i\mu_0 t - \gamma |t| [1 + i\beta \text{sign}(t) \frac{2}{\pi} (\log |t| + \log \gamma)]) & \text{if } \alpha = 1. \end{cases} \quad (1)$$

This alternative parameterization is not far from the S_α one. The only difference is about the location parameter, which, in this new setting, corrects the shift exposed above for values of α close to 1:

$$\mu_0 = \begin{cases} \mu + \beta\gamma \tan \frac{\pi\alpha}{2} & \text{if } \alpha \neq 1 \\ \mu + \beta \frac{2}{\pi} \gamma \log \gamma & \text{if } \alpha = 1. \end{cases} \quad (2)$$

A Fourier transform makes it possible to get the pdf of a standard variable $S_\alpha^0(1, \beta, 0)$ [67]:

$$p(x; \alpha, \beta) = \frac{1}{\pi} \int_0^{+\infty} \cos(h(x, t; \alpha, \beta)) e^{-t^\alpha} dt,$$

with

$$h(x, t; \alpha, \beta) = \begin{cases} xt + \beta \tan \frac{\pi\alpha}{2} (t - t^\alpha) & \text{if } \alpha \neq 1 \\ xt + \beta \frac{2}{\pi} t \log t & \text{if } \alpha = 1. \end{cases}$$

We also obtain the pdf of a variable $X \sim S_\alpha^0(\gamma, \beta, \mu_0)$ with $\gamma \neq 1$ and $\mu_0 \neq 0$, by the mean of a translation, a scaling, and the substitution $s = \gamma t$, starting from the characteristic function provided in equation (1):

$$\begin{aligned} f(x; \alpha, \gamma, \beta, \mu_0) &= \frac{1}{\pi} \int_0^{+\infty} \cos \left(h \left(\frac{x - \mu_0}{\gamma}, \gamma t; \alpha, \beta \right) \right) e^{-(\gamma t)^\alpha} dt \\ &= \frac{1}{\gamma \pi} \int_0^{+\infty} \cos \left(h \left(\frac{x - \mu_0}{\gamma}, s; \alpha, \beta \right) \right) e^{-s^\alpha} ds \\ &= \frac{1}{\gamma} p \left(\frac{x - \mu_0}{\gamma}; \alpha, \beta \right). \end{aligned}$$

The formula of p contains an integral on an unbounded interval. For this reason, other formulations have been proposed [2, 57, 47, 15]. However, in the application to financial series, we find values of α far from 0, so that the truncated integral converges rapidly. We also get the corresponding cdf by numerically integrating $f(x; \alpha, \gamma, \beta, \mu_0)$.

2.2 Estimation

Many estimation methods of stable distributions exist [2, 15]. Some of them focus on the sole α parameter, using for instance a regression of extreme quantiles along with the extreme value theory [23]. Other methods make it possible to estimate all the parameters. In this class of methods, we can cite the estimation using L-moments, provided that $\alpha > 1$ [2, 44], empirical quantiles [62], the empirical characteristic function [74, 69, 51, 50, 80], and the maximum likelihood method [27]. This last method relies on the knowledge of the pdf, which may be approximated as exposed in the previous subsection, using Nolan's work [67].

We focus on the method using empirical quantiles, which has been introduced first by Fama and Roll, with the following assumptions: $\alpha > 1$, $\beta = 0$, and $\mu = 0$ [31]. This method is asymptotically biased. McCulloch proposed an extended version of the method, in which he corrected the asymptotic bias [62]. This version is also less restrictive with respect to the parameters, insofar as it only requires to have $\alpha \geq 0.6$.

The McCulloch's method consists in linking the four parameters to five empirical quantiles of probability levels 5%, 25%, 50%, 75%, and 95%. To this end, we first have to define two intermediate quantities:

$$v_\alpha = \frac{Q(0.95) - Q(0.05)}{Q(0.75) - Q(0.25)}$$

and

$$v_\beta = \frac{Q(0.95) + Q(0.05) - 2Q(0.50)}{Q(0.95) - Q(0.05)},$$

where $Q(p)$ is the theoretical quantile of probability p for a $S_\alpha(\gamma, \beta, \mu)$ variable [62, 2, 15]. Neither v_α nor v_β depend on γ and μ , so that α and β are functions of v_α and v_β : $\alpha = \phi_1(v_\alpha, v_\beta)$ and $\beta = \phi_2(v_\alpha, v_\beta)$. In practice, ϕ_1 and ϕ_2 are provided by tables [62]. Replacing the theoretical

quantiles $Q(p)$ by empirical quantiles $\widehat{Q}(p)$, we get estimators $\widehat{v}_\alpha$ for v_α and $\widehat{v}_\beta$ for v_β , as well as the following estimators for α and β :

$$\begin{cases} \widehat{\alpha} = \phi_1(\widehat{v}_\alpha, \widehat{v}_\beta) \\ \widehat{\beta} = \phi_2(\widehat{v}_\alpha, \widehat{v}_\beta). \end{cases}$$

The estimation of γ and μ relies on two other intermediate quantities which only depend on the already estimated α and β . For simplicity, we introduce the variable ζ defined by:

$$\zeta = \begin{cases} \mu + \beta\gamma \tan \frac{\pi\alpha}{2} & \text{if } \alpha \neq 1 \\ \mu & \text{if } \alpha = 1. \end{cases} \quad (3)$$

The intermediate quantities are:

$$v_\gamma = \frac{Q(0.75) - Q(0.25)}{\gamma}$$

and

$$v_\zeta = \frac{\zeta - Q(0.50)}{\gamma}.$$

They are such that $v_\gamma = \phi_3(\alpha, \beta)$ and $v_\zeta = \phi_4(\alpha, \beta)$. Their estimators $\widehat{v}_\gamma$ and $\widehat{v}_\zeta$ are obtained by replacing the quantiles by empirical ones. We thus have the following estimators for γ and ζ :

$$\begin{cases} \widehat{\gamma} = \frac{\widehat{Q}(0.75) - \widehat{Q}(0.25)}{\phi_3(\widehat{\alpha}, \widehat{\beta})} \\ \widehat{\zeta} = \widehat{\gamma} \phi_4(\widehat{\alpha}, \widehat{\beta}) + \widehat{Q}(0.50). \end{cases}$$

The deduction of the estimator of μ from $\widehat{\zeta}$ is straightforward using equation (3), as well as the version μ_0 of the location parameter in the parameterization S^0 using equation (2). In particular $\widehat{\mu}_0 = \widehat{\zeta}$ as soon as $\alpha \neq 1$.

2.3 Dynamic estimation

The time-varying adaptation of McCulloch's estimation technique amounts to estimating time-varying quantiles. Indeed, if we are able to infer dynamic quantiles $Q_t(p)$, the first two McCulloch statistics are defined in a dynamic fashion:

$$\begin{cases} v_{\alpha,t} &= \frac{Q_t(0.95) - Q_t(0.05)}{Q_t(0.75) - Q_t(0.25)} \\ v_{\beta,t} &= \frac{Q_t(0.95) + Q_t(0.05) - 2Q_t(0.50)}{Q_t(0.95) - Q_t(0.05)}, \end{cases}$$

so that we get time-varying α and β parameters. Time-varying quantiles make it also possible to define the last two McCulloch statistics and finally to fully estimate a dynamic alpha-stable probability distribution.

The subject of estimating dynamic quantiles is largely handled by the econometric literature. The favoured approach is based on quantile autoregression [49, 24, 41], like in the application to value-at-risk known as CAViaR [28]. A drawback of quantile regression is that different quantiles may cross: the monotonicity of quantiles is not necessarily preserved. The dynamic additive quantile, while keeping the autoregressive approach, deals with this limitation [42].

All the quantile autoregressions are based on a model of dynamic and we prefer to introduce a method in which we do not specify the evolution of quantiles with respect to previously estimated ones. The inspiration of such a model-free approach comes from the non-parametric statistics literature, in which we can estimate for example time-varying moments or even time-varying probability

densities [43]. In this perspective, we simply consider that the price returns X_t are alpha-stable random variables, each following its own alpha-stable distribution F_t , defined by parameters α_t , β_t , ω_t , and μ_t . In each distribution F_t , we only observe one variable X_t . This is not enough for estimating F_t . For this reason, we add another assumption on the dynamic of this time-varying distribution. We indeed consider that it evolves smoothly, so that price returns close in time will follow a close distribution.

Thanks to this assumption, we can use several price returns at times close to t in order to estimate F_t or its related quantiles $Q_t(p)$. We could estimate these quantiles $Q_t(p)$ by their empirical versions. For example, if one considers 100 observations, the empirical quantile of probability p is the $\lceil 100p \rceil$ -th smallest of the 100 observations. In other words, each of the 100 observations is associated with a probability $1/100$. However, price returns closer in time induce a lower bias in the estimation of $Q_t(p)$, so that the probability associated with these recent returns in the cdf related to Q_t should be higher. Thus, we use an exponentially-weighted quantile estimator (EWQ) and the corresponding discrete probability $\hat{P}_t^{EWQ}$, such that $\hat{P}_t^{EWQ}(x)$ is the probability for the price return at time t to be equal to x [65]. To this end, we estimate the EWQ-style discrete probability function by assigning to each observation X_i , for $i \in \llbracket 1, t \rrbracket$, a probability $p_{t,i}^\omega$ depending on a discount factor $\omega \in (0, 1)$, so that $\hat{P}_t^{EWQ}(X_i) = p_{t,i}^\omega$. The probability $p_{t,i}^\omega$ follows a standard form of exponential weight:

$$p_{t,i}^\omega = \frac{1 - \omega}{1 - \omega^t} \omega^{t-i}.$$

This expression ensures that the probabilities sum to 1, as required for probabilities:

$$\sum_{i=1}^t p_{t,i}^\omega = 1.$$

For $\omega \neq 1$ and big values of t , ω^t is close to zero and we have the approximation $p_{t,i}^\omega \approx (1 - \omega)\omega^{t-i}$. In addition to the discrete probability $\hat{P}_t^{EWQ}$, we can define the cdf $\hat{F}_t^{EWQ}$ associated with the EWQ:

$$\hat{F}_t^{EWQ}(x) = \sum_{i=1}^t \hat{P}_t^{EWQ}(X_i) \mathbb{1}_{X_i \leq x} = \sum_{i=1}^t p_{t,i}^\omega \mathbb{1}_{X_i \leq x}. \quad (4)$$

We can also write recursively this cdf [65]:

$$\hat{F}_t^{EWQ}(x) = \omega \hat{F}_{t-1}^{EWQ}(x) + (1 - \omega) \mathbb{1}_{X_t \leq x}.$$

More precisely, for each new observation X_t , the EWQ-style cdf associated with each past observations is discounted at the constant rate ω , whereas the new observation is provided with the highest probability, $1 - \omega$. The rationale is the following: the more recent the observation, the more likely its future occurrence.

These discrete probabilities make it possible to calculate easily time-varying quantiles, as generalized inverse functions of $\hat{F}_t^{EWQ}$. Compared to quantiles determined from dynamic kernel densities, which suffer from a high algorithmic complexity [83, 43], the EWQ approach seems computationally more performing. Indeed, for extreme quantiles, the algorithmic complexity of our method is less than linear with respect to the number of observations considered at each date.

We now describe precisely the procedure for calculating the EWQs. We begin with a first estimation, at date t_0 , of the quantile of probability p , $\hat{Q}_{t_0}^\omega(p)$, using the EWQ approach. To estimate the quantiles, we build the matrix $M_{t_0}^\omega$ containing the past observed price returns till t_0 , $\{X_i\}_{i \leq t_0}$,

sorted in descending order, along with their corresponding historical EWQ-style probability:

$$M_{t_0}^\omega = \begin{pmatrix} X_{\pi_{t_0}(t_0)} & p_{t_0, \pi_{t_0}(t_0)}^\omega \\ \vdots & \vdots \\ X_{\pi_{t_0}(i)} & p_{t_0, \pi_{t_0}(i)}^\omega \\ \vdots & \vdots \\ X_{\pi_{t_0}(1)} & p_{t_0, \pi_{t_0}(1)}^\omega \end{pmatrix},$$

where $X_{\pi_{t_0}(i)}$ is the i -th order statistic among the t_0 first observations and is obtained with the help of a permutation $\pi_{t_0}: \min_{i \leq t_0} \{X_i\} = X_{\pi_{t_0}(1)} \leq X_{\pi_{t_0}(2)} \leq \dots \leq X_{\pi_{t_0}(t_0)} = \max_{i \leq t_0} \{X_i\}$. Thanks to the definition of the generalized quantile and to equation (4), we get the following estimator for the quantile of probability p :

$$\widehat{Q}_{t_0}^\omega(p) = X_{\pi_{t_0} \left(\min \left\{ \tau \in \{1, \dots, t_0\} \mid \sum_{i=1}^{\tau} p_{t_0, \pi_{t_0}(i)}^\omega \geq p \right\} \right)}, \quad (5)$$

which is simply the lowest price return X such that the cumulated probability associated to lower returns, $\widehat{F}_{t_0}^{EWQ}(X)$, reaches p .

Iteratively, we can update this quantile in the following manner. We suppose we are given the probability distribution till time $t-1$ and that we want to estimate a quantile at time t . We first apply a probability decay to M_{t-1}^ω by a simple matrix product, so that we get a new matrix $\widetilde{M}_{t-1}^\omega$ containing sorted past observations till time $t-1$ along with their new probability at time t :

$$\widetilde{M}_{t-1}^\omega = M_{t-1}^\omega \begin{pmatrix} 1 & 0 \\ 0 & \omega \end{pmatrix}.$$

Then, we insert the new observation of time t in the matrix $\widetilde{M}_{t-1}^\omega$ thanks to a binary search in its first column. In the inserted line, we write the corresponding probability $1 - \omega$ in the second column. If we write $\mathcal{I}_t$ the position of the new observation, we get the new probability distribution matrix M_t^ω by:

$$M_t^\omega = \begin{pmatrix} X_{\pi_t(t)} & p_{t, \pi_t(t)}^\omega \\ \vdots & \vdots \\ X_{\pi_t(\mathcal{I}_t+1)} & p_{t, \pi_t(\mathcal{I}_t+1)}^\omega \\ X_{\pi_t(\mathcal{I}_t)} & p_{t, \pi_t(\mathcal{I}_t)}^\omega \\ X_{\pi_t(\mathcal{I}_t-1)} & p_{t, \pi_t(\mathcal{I}_t-1)}^\omega \\ \vdots & \vdots \\ X_{\pi_t(1)} & p_{t, \pi_t(1)}^\omega \end{pmatrix} = \begin{pmatrix} X_{\pi_{t-1}(t-1)} & \omega p_{t-1, \pi_{t-1}(t-1)}^\omega \\ \vdots & \vdots \\ X_{\pi_{t-1}(\mathcal{I}_t)} & \omega p_{t-1, \pi_{t-1}(\mathcal{I}_t)}^\omega \\ X_t & 1 - \omega \\ X_{\pi_{t-1}(\mathcal{I}_t-1)} & \omega p_{t-1, \pi_{t-1}(\mathcal{I}_t-1)}^\omega \\ \vdots & \vdots \\ X_{\pi_{t-1}(1)} & \omega p_{t-1, \pi_{t-1}(1)}^\omega \end{pmatrix}.$$

We can then calculate the empirical quantile in a manner similar to equation (5):

$$\widehat{Q}_t^\omega(p) = X_{\pi_t \left(\min \left\{ \tau \in \{1, \dots, t\} \mid \sum_{i=1}^{\tau} p_{t, \pi_t(i)}^\omega \geq p \right\} \right)}.$$

In order to diminish the algorithmic complexity of this method, if we are looking for a quantile above the probability 0.5, we prefer the following definition of the quantile, which is mathematically consistent with the one provided above:

$$\widehat{Q}_t^\omega(p) = X_{\pi_t \left(\min \left\{ \tau \in \{1, \dots, t\} \mid \sum_{i=\tau+1}^t p_{t, \pi_t(i)}^\omega \leq 1-p \right\} \right)}.$$

Finally, using McCulloch's method, these time-varying estimations of quantiles make it possible to estimate the dynamic parameters of the stable distribution, that we note $\hat{\alpha}_t^\omega$, $\hat{\beta}_t^\omega$, $\hat{\gamma}_t^\omega$, and $\hat{\mu}_t^\omega$. We stress the fact that the EWQ method is only intended to estimate non-parametric quantiles in order to infer time-varying parameters. In particular, in what follows, the time-varying cdf on which we focus is not the non-parametric EWQ-style cdf but the parametric alpha-stable cdf with the time-varying parameters estimated above.

2.4 Selection of the discount factor

The above dynamic estimation of quantiles and of the parameters of stable distributions relies on a free parameter, the discount factor ω , which depicts how fast the dynamic distribution evolves. If ω is close to 1, the distribution is almost constant. If ω is lower, the evolution of the distribution is faster and the description of the last observations will be more accurate. Nevertheless, this accuracy may be excessive and the evolution of the distribution may be non-significant. Indeed, in the extreme situation where ω is close to zero, the distribution will be very narrow and centered on the last observation, with a big divergence between two successive distributions. In order to find a good balance between accuracy and robustness, we decide to select the ω maximizing the ability of the density $\hat{f}_t^\omega$ estimated at date t to forecast the density of X_{t+1} , the price return at time $t + 1$.

Several definitions of what a good density forecast is are possible. Indeed, the true density at time $t + 1$ is never observed, so that we can only rely on one observation drawn in this density. In the non-parametric literature about time-varying densities, we find for instance a selection rule for the free parameter of the densities based on the maximization of a likelihood criterion [43]. We think that this criterion does not take properly into account the possibility of the occurrence of extreme events. The alternative solution we follow is based on an adaptation of a method coming from the literature of density forecast evaluation [40]. Indeed, even if $\hat{f}_t^\omega$ varies with t , so that we are provided with only one observation in this distribution, a simple transformation of each price return defines a distribution which remains the same through time. This transformation is the probability integral transform (PIT), usually introduced in the perspective of density forecast evaluation [26]:

$$Z_t^\omega = \hat{F}_{t-1}^\omega(X_t), \quad (6)$$

where $\hat{F}_{t-1}^\omega$ is the cdf corresponding to $\hat{f}_{t-1}^\omega$. In this literature, two conditions are required for the PITs: the Z_t^ω must follow a uniform distribution in $[0, 1]$ and they must be independent from each other.

The translation of these rules to the selection of free parameters in density estimation leads to two properties regarding ω [40]:

- ▷ uniformity of the PITs $Z_{t_0}^\omega, \dots, Z_T^\omega$: ω is to be selected so as to minimize the divergence between their empirical distribution and a uniform distribution,
- ▷ independence of the PITs: ω is to be selected so as to minimize the discrepancy, that is, for each subinterval of $[t_0, T]$ of size greater than a threshold ν , the divergence between the empirical distribution of the corresponding PITs and a uniform distribution is to be minimized.

In the perspective of the selection of a free parameter of a time-varying distribution, we can define the above divergence as a Kolmogorov-Smirnov statistic with an adaptation to compare directly divergences of distributions estimated on samples of different sizes [40]. Moreover, the minimal size ν of the subintervals considered is intended to be a threshold above which the asymptotic

framework required by the Kolmogorov-Smirnov statistic is satisfied. We consider $\nu = 22$ days, so that we expect the PITs to be uniform for scales larger than one month. As a consequence, the criterion to be minimized is:

$$d_\nu(Z_{t_0+1}^\omega, \dots, Z_T^\omega) = \max_{t_0+1 \leq s < s+\nu-1 \leq t \leq T} (\sqrt{t-s+1} \times k(Z_s^\omega, \dots, Z_t^\omega)),$$

in which k is the standard Kolmogorov-Smirnov statistic with respect to the uniform distribution:

$$k(Z_s^\omega, \dots, Z_t^\omega) = \max_{s \leq u \leq t} \left| \frac{u-s}{t-s} - Z_{\rho(u;s,t)}^\omega \right|, \quad (7)$$

where $v \in \llbracket a, b \rrbracket \mapsto \rho(v; a, b)$ is a permutation of $\llbracket a, b \rrbracket$ defining the new order $Z_{\rho(a;a,b)}^\omega \leq Z_{\rho(a+1;a,b)}^\omega \leq \dots \leq Z_{\rho(b-1;a,b)}^\omega \leq Z_{\rho(b;a,b)}^\omega$ [40]. In equation (7), $(u-s)/(t-s)$ is the empirical cdf of the PITs, whereas the sorted PIT $Z_{\rho(u;s,t)}^\omega$ is the theoretical uniform cdf. Finally, the optimal discount factor is defined as the solution of the following equation:

$$\omega^\star = \underset{0 < \omega < 1}{\operatorname{argmin}} d_\nu(Z_{t_0+1}^\omega, \dots, Z_T^\omega). \quad (8)$$

We can easily apply the above method to the case of an alpha-stable distribution. We only have to pay attention to the definition of the PIT in equation (6). Indeed, it relies on the cdf, which does not follow a simple expression for alpha-stable distributions. Nevertheless, as exposed in Subsection 2.1, numerical methods make it possible to calculate both the pdf and the cdf of an alpha-stable distribution. Therefore, the estimated cdf $x \mapsto \hat{F}_{t-1}^\omega(x)$ in equation (6) is the numerically evaluated cdf of an alpha-stable distribution with estimated parameters: $x \mapsto F(x; \hat{\alpha}_{t-1}^\omega, \hat{\gamma}_{t-1}^\omega, \hat{\beta}_{t-1}^\omega, (\hat{\mu}_0)_{t-1}^\omega)$.

3 Market efficiency

The market efficiency is a usual assumption in finance, which is in particular invoked when pricing derivatives. The EMH states that asset price series follow a random walk [30]. Investors and market makers update their expectations and their quotes at each instant, using the available information, so that it is not possible to beat the market. The EMH is convenient, because working with independent price returns makes the financial mathematics easier. But practitioners know that the EMH is not very realistic. The asset management industry in general aims at performing statistical arbitrages, whatever the transaction time scale, whether it is less than a second or more than a month. Maybe the asset manager will not win at each time, but in average he should win. Standard models relying on the EMH, such as the widespread geometric Brownian motion, are not consistent with the existence of statistical arbitrages.

We can also stress other unrealistic components of the EMH, such as the availability of the same information for every investor and market maker, as well as the rationality of all the agents, or the fact that we can substitute two assets provided that they are equally risky. Several alternatives to the EMH aim at describing more realistically the financial markets. It is the case of the fractal market hypothesis [73] and of the adaptive market hypothesis (AMH) [60]. The AMH is a good compromise between the standard EMH and the reality of statistical arbitrage. It states that a model can predict the market in average. But such a model can provide investors with performing forecasts during a limited time only, because other investors and market makers will progressively adapt their own models and decisions to this performing model. The AMH thus leads to a long-term efficiency of the market and allows statistical arbitrages for small time scales only.

Besides, it is worth noting that the price series of some assets are not far from the EMH. It is the case for major stock indices. But empirical studies show that these indices have fluctuating

efficiency and may encounter a loss of efficiency during financial crisis [3, 59]. Since the decisions of market makers and investors may not be the same if markets are efficient or if they aren't, it is important for them to determine from time series of prices whether the markets are efficient or not. We can cite many statistical indicators of market efficiency [75]. Some of them look for a predictability of price returns [19], using for example the amplitude of the parameters of a time-varying AR model [46, 66]. Other indicators measure a deviation from a random walk, using for example variance ratios [18] or a combination of several statistics such as fractal dimension and entropy [53, 52, 32]. But the most widespread indicator of market efficiency seems to be the Hurst exponent [58, 16, 63, 34].

The Hurst exponent is an indicator of long-range memory. Mandelbrot and van Ness introduced the fBm as a model consistent with a given Hurst exponent and extending the standard Bm by making the increments dependent on each other [61]. For this reason, the fBm is a popular model in finance, provided that practitioners do not want to comply with the EMH. In the fBm, the Hurst exponent H is also linked to the fractal property of the series, insofar as the variance of the increments of duration τ is τ^{2H} times the variance of increments of duration 1. Some estimators of the Hurst exponent use this fractal property instead of the long-range memory. However, given the fractal property estimated on the dataset, the fBm is not the only possible model. Other models may indeed have the same estimated Hurst exponent as the fBm but, relying on another specification, they may lead to other conclusions regarding the dependence of the increments and thus regarding the efficiency of the markets [1]. This fact has been documented for instance for foreign exchange rates, for which the stationarity of the time series biases the estimation of the Hurst exponent in the perspective of an fBm [35, 36]. Indicators of market efficiency should not bypass the diversity of models featuring a fractal property and consistent with the estimated Hurst exponent. Extensions of the fBm include some specificities for the Hurst exponent: it may vary deterministically through time, as in the multifractional Brownian motion [71, 11, 20, 34], it may be a random process, as in the multifractional process with random exponent [5, 14, 33, 37], or it may even be asymmetric [17, 81].

In what follows, after a presentation of the Hurst exponent in the perspective of the fBm, we will focus on a model in which increments are not necessarily Gaussian. In this framework, the fLsm is a natural extension of the fBm, in which increments follow an alpha-stable distribution [79, 82, 80, 37]. The fractal property of this process takes into account both the dependence between increments and the tail parameter of the alpha-stable distribution. This model thus makes it possible to disentangle the kurtosis of price returns and the efficiency of the market. It provides us with richer information than the sole Hurst exponent. This refinement is not superfluous. We will indeed see in the empirical section that the conclusions regarding the efficiency of stock indices are not the same if we consider the Hurst exponent or the fLsm-based indicator of market efficiency.

3.1 Hurst exponents

The Hurst exponent was originally introduced by Harold Edwin Hurst as an indicator of long-range memory that could be obtained thanks to the rescaled range (R/S) analysis [45]. Later, alternative estimation methods appeared, such as the detrended fluctuation analysis [72], or the absolute-moment method [11, 10, 12, 34], which is related to the notion of generalized Hurst exponent (GHE) [7, 25]. The absolute-moment method is a method mainly used by the community of statisticians of stochastic processes, because, contrary to the R/S analysis, it is strongly related to a stochastic process, namely the fBm.

The fBm is a generalization of the standard Bm. Increments of the fBm are Gaussian, since the

fBm is a fractional integral or a fractional derivative of a Bm, W_t [61]:

$$B_H(t) = \frac{\sigma}{\Gamma(H + \frac{1}{2})} \int_{-\infty}^{\infty} \left((t-s)_+^{H-1/2} - (-s)_+^{H-1/2} \right) dW(s).$$

The two parameters of the fBm are the Hurst exponent H , and the volatility parameter σ . If $H = 1/2$, the fBm is a Bm. If $H > 1/2$ (respectively $H < 1/2$), the fBm is the fractional integral (resp. fractional derivative) of order $H-1/2$ (resp. $1/2-H$) of a Bm; increments are thus positively (resp. negatively) correlated.

The absolute-moment method uses another definition of the fBm, which is consistent with the integral form provided above. Indeed, an fBm is also the only zero-mean Gaussian process, with zero at the origin, such that, for $s, t \geq 0$:

$$\mathbb{E}\{B_H(t)B_H(s)\} = \frac{\sigma^2}{2}(|t|^{2H} + |s|^{2H} - |t-s|^{2H}).$$

From this covariance, we get the self-similarity property:

$$\mathbb{E}\{|B_H(t) - B_H(s)|^k\} = \sigma^2 |t-s|^{kH},$$

for $k > 0$. This property states that the k -order absolute moment of the increments of duration $|t-s|$ is proportional to $|t-s|^{kH}$. Comparing two scales thus makes it possible to estimate H . If we focus on the two smallest scales, we indeed get the following estimator, for a time series of log-prices $X_1, X_2, \dots, X_t$:

$$\hat{H}_k = \frac{1}{k} \log_2 \left(\frac{(t-1) \sum_{i=1}^{t-2} |X_{i+2} - X_i|^k}{(t-2) \sum_{i=1}^{t-1} |X_{i+1} - X_i|^k} \right),$$

which converges almost surely toward H [11, 10].

We could use this estimator of H as an indicator of market efficiency. But, we are mostly interested here in evaluating the evolution of market efficiency through time. We need therefore a time-varying version of this estimator. This question is not new and the solution put forward in the literature is often based on the estimation in sliding windows [20, 13], possibly with a smoothing of the raw series of Hurst exponents as a post-processing [34]. But we prefer a method more consistent with the smoothing applied above for estimating the parameters of a distribution, that is with an exponential weighting, insofar as it overweights more recent observations and is thus a more relevant picture of the current state of the market. The closest method in the literature is a time-varying GHE using the exponentially-weighted moving average [64]. The main difference with the method we propose is that the GHE is always based on a linear regression of log-absolute moments of increments on several log time scales. Contrary to the GHE, we focus on only two scales, so that we get a simpler closed-form estimator:

$$\hat{H}_{k,t}^\omega = \frac{1}{k} \log_2 \left(\frac{(t-1) \sum_{i=1}^{t-2} \omega^{t-2-i} |X_{i+2} - X_i|^k}{(t-2) \sum_{i=1}^{t-1} \omega^{t-1-i} |X_{i+1} - X_i|^k} \right).$$

Beyond this simple formula, a more efficient implementation method is possible. Given the one-step statistic $M_{k,1,t} = \sum_{i=1}^{t-1} \omega^{t-1-i} |X_{i+1} - X_i|^k$ and the two-step statistic $M_{k,2,t} = \sum_{i=1}^{t-2} \omega^{t-2-i} |X_{i+2} - X_i|^k$, it consists of the following recurrence:

$$\begin{cases} M_{k,1,t+1} &= \omega M_{k,1,t} + |X_{t+1} - X_t|^k \\ M_{k,2,t+1} &= \omega M_{k,2,t} + |X_{t+1} - X_{t-1}|^k \\ \hat{H}_{k,t+1}^\omega &= \frac{1}{k} \log_2 \left(\frac{t M_{k,2,t+1}}{(t-1) M_{k,1,t+1}} \right). \end{cases} \quad (9)$$

We will discuss in Section 3.2 the selection of the parameter k in equation (9). The question of the optimal choice of the discount factor ω in the estimator $\hat{H}_{k,t}^\omega$ is also to be addressed. We could imagine to adapt the method exposed in section 2.4 to the case of an fBm. But it sounds more relevant to use the same discount factor for all the statistics of our work. We will thus simply use the discount factor chosen for estimating the time-varying stable distribution.

3.2 Fractional Lévy-stable motion

The fLsm is a generalization of the fBm in which increments follow an alpha-stable distribution, which admits the Gaussian distribution as a particular case. The fLsm is defined as the fractional integral or fractional derivative of a symmetric Lévy-stable motion [78, 79, 80]:

$$B_{H,\alpha}(t) = \frac{1}{C_{H,\alpha}} \int_{-\infty}^{\infty} \left((t-s)_+^{H-1/\alpha} - (-s)_+^{H-1/\alpha} \right) dL_{\alpha,\gamma}(s),$$

where $L_{\alpha,\gamma}(t)$ is a symmetric α -stable process of scale parameter γ and

$$C_{H,\alpha} = \left(\int_{\mathbb{R}} \left| (t-s)_+^{H-1/\alpha} - (-s)_+^{H-1/\alpha} \right|^\alpha ds \right)^{1/\alpha}.$$

In other words, increments follow a symmetric stable law and, if and only if $H - 1/\alpha > 0$, non-overlapping increments are positively dependent, that is with a positive codifference or a positive covariation [56]. The parameter H controls the scaling behaviour of the process, in the same manner as in the fBm, and the parameter α controls the thickness of the tails. The lower α , the fatter the tails. If $\alpha = 2$, increments are Gaussian and the fLsm is an fBm. When comparing an fBm and an fLsm, the fractal feature of the latter is not obtained only by adjusting the dependence of the increments but, in addition, by tuning the kurtosis of the underlying law. Therefore, we can write the Hurst exponent as the combination of a tail component, $1/\alpha$, and a memory parameter, m :

$$H = \frac{1}{\alpha} + m. \quad (10)$$

In this framework, the Hurst exponent is not the most relevant indicator of market efficiency and one should instead use m . As we have proposed time-varying estimators both for α and for H , we write the following time-varying estimator for m :

$$\hat{m}_t^\omega = \hat{H}_{k,t}^\omega - \frac{1}{\hat{\alpha}_t^\omega}, \quad (11)$$

where $\hat{\alpha}_t^\omega$ is the estimator of the parameter of the stable law defined in section 2.3. Efficient markets then correspond to $\hat{m}_t^\omega$ close to zero. The multifractional multistable motion is a model allowing this kind of dynamic with smoothly time-varying parameters α_t and γ_t . It is a localisable process in the sense that such a process admits at each time a local form, also called tangent process, which in our case is an fLsm [29, 55].

We also have to select an appropriate k for the estimator of H in equation (11). A widespread choice is $k = 2$, because it minimizes the variance of the estimator. However, for α -stable variables, the absolute moments are only defined for $k < \alpha$ [48]. We can mitigate this theoretical constraint by noting that the statistics of equation (9) are finite even for higher values of k . In the empirical part of this paper, we first follow the widespread choice $k = 2$. Then, we confirm the results with a lower value of k , namely $k = 0.5$, which is consistent with the minimal estimated α appearing in Figure 5.

4 Empirical study

We apply the above method to ten stock indices of various regions: USA (S&P 500, S&P 100), Europe (EURO STOXX 50, Euronext 100, DAX, CAC 40), Asia (Nikkei, KOSPI, SSE 180), and Australia (S&P/ASX 200). We have used data from Yahoo finance in the time interval between the 1st May 2015 and the 29th June 2020. The first date at which we estimate stable densities and the parameters of an fBm and an fLsm, that is t_0 , is the 1st November 2019. The period of study includes the financial crisis sparked by the COVID-19 pandemic.

We first determine for each stock index the optimal discount factor ω as in equation (8). The results are displayed in Table 1. We observe that the optimal discount factors are close to 0.95, whatever the index considered. For the rest of the empirical study, we consider a common discount factor ω^m , so that we can make fair comparisons between stock indices. We choose $\omega^m = 0.956$, which is the highest optimal discount factor measured for the various stock indices. This conservative choice limits the risk of providing a new observation with a too high weight in the dynamic estimators. A lower value could lead to spurious conclusions regarding the market efficiency for some stock indices.

Stock index	ω^*
S&P 500	0.946
S&P 100	0.939
EURO STOXX 50	0.949
Euronext 100	0.954
DAX	0.949
CAC 40	0.952
Nikkei 225	0.951
KOSPI	0.952
SSE 180	0.946
S&P/ASX 200	0.956
Mean value	0.949
Median value	0.950

Table 1: Optimal discount factor ω^* for several stock indices for stable densities between November 2019 and June 2020.

Using ω^m , we are able to determine the dynamic pdf of daily price returns. We display in Figure 1 these densities for four stock indices corresponding to regions with a different timing in the growth of the outbreak: S&P 500, the French CAC 40, the Chinese SSE 180, and the S&P/ASX 200 indices. For each of these indices, we plot the pdf before the crisis, in November 2019, at the peak of the crisis, and at the end of our sample, late June 2020. The peak of the crisis is not the same for all the indices. We define the peak date as the one leading to the maximal value for $|m|$. This date is in March for the four indices. Exact dates are provided in Table 3. For the four indices, the pdf at the peak shows fat tails and asymmetry. Late June, the pdf still has these features, except for SSE 180 index, for which the pdf is very similar to the one before the crisis, indicating a very fast recovery in China. The case of CAC 40 at the peak is also of interest, because the advent of fat tails and asymmetry is so abrupt that it does not crush the body of the pdf, contrary to other indices.

In this paper, we are mostly interested in determining whether the financial markets are efficient during a financial crisis. For this purpose, we have introduced two indicators. The first one is the widespread Hurst exponent, estimated here in a dynamic fashion as exposed in Section 3.1. But the Hurst exponent H provided above is an indicator of dependence between price returns only if




Figure 1: Estimated dynamic stable density of daily price returns for S&P 500 (top left), CAC 40 (top right), S&P/ASX 200 (bottom left), and SSE 180 (bottom right) indices. The dotted line is the density at the 5th November 2019, the black continuous line at the peak of the crisis, the grey line at the 29th June 2020.

these price returns are Gaussian. In a more general framework, if we consider the possibility of fat tails by the mean of an alpha-stable distribution, we define another efficiency indicator by the memory parameter m of an fLsm, as exposed in Section 3.2. In the first approach, the market is efficient for $H = 1/2$. In the second approach, which is more accurate because it takes into account the kurtosis of price returns thanks to the α parameter, the market is efficient for $m = 0$. The null hypothesis H_0 is the efficiency of each market. In order to know for which threshold of efficiency indicator we can reject H_0 with a given confidence p , we perform a simulation. We consider that the right price model corresponding to H_0 is a geometric Bm. We thus simulate a time series following this model. The length of the simulated series used for estimating the first parameters in t_0 is the same as in our financial dataset. We then simulate 4,000 other dates. For each of these dates, we estimate H , m , and α dynamically, using the discount factor ω^m . We consider that the bounds of the confidence interval with confidence level p , for the estimated H , m , and α , are the empirical quantiles of the corresponding parameters, estimated on the simulations, for probabilities $(1 - p)/2$ and $(1 + p)/2$.

We display the time-varying Hurst exponent in Figure 2 for the four focal stock indices of our study. Before the crisis, which begins in February or March with respect to the region, the Hurst exponent

is not significantly different from $1/2$, so that we cannot reject H_0 during this period. During the crisis, we observe very low Hurst exponents for the S&P 500 index which make it possible to reject H_0 with a confidence of more than 99%. But for the three other indices, the drawdown of the Hurst exponent is much less significant, in particular for the SSE 180 index. According to this approach, the market becomes clearly inefficient if we consider the S&P 500 index, whereas we cannot ascertain the inefficiency of the three other indices. It is also worth noting that after a peak downward, the Hurst exponent reaches abnormally high values for the French and the Australian indices. It suggests a persistence of the crisis: a short mean-reverting phenomenon followed by positively correlated price returns.




Figure 2: Daily evolution through time of the estimated Hurst exponent for S&P 500 (top left), CAC 40 (top right), S&P/ASX 200 (bottom left), and SSE 180 (bottom right) indices. The grey lines are the bounds of simulated confidence intervals with a confidence level of 95%, 99%, and 99.5%.

If we consider another indicator of market efficiency, namely the memory parameter m of an fLsm, the conclusions regarding the impact of COVID-19 on the market efficiency are not the same. In Figure 3, we observe first that m is in general negative, even before the crisis. It indicates a dominating mean-reversion phenomenon in the stock markets. But it is in fact often not significantly different from 0. When the crisis occurs, m goes downward and H_0 is rejected with a confidence higher than 99%, whatever the stock index. The most significant drawdown of m is again for the S&P 500 index. The duration of the significant inefficiency varies among the indices. The longer period is for the S&P 500 index. The S&P/ASX 200 has a very short period of inefficiency.

Using $k = 0.5$ instead of $k = 2$, as discussed in Section 3.2, leads to similar results. In particular, the




Figure 3: Daily evolution through time of the estimated m parameter of an fLsm for S&P 500 (top left), CAC 40 (top right), S&P/ASX 200 (bottom left), and SSE 180 (bottom right) indices. The grey lines are the bounds of simulated confidence intervals with a confidence level of 95%, 99%, and 99.5%.

loss of efficiency is very strong for the S&P 500 index, regardless of the indicator used, H or m . For the CAC 40 index, the Hurst exponent is never significantly below $1/2$, but m becomes significantly negative during the first lock-down. We display the corresponding evolutions in Figure 4.

It is also interesting to track another important parameter of the fLsm, namely α , which depicts the size of the tails of the distribution of price returns. When $\alpha = 2$, the price returns follow a Gaussian distribution. The lower α , the fatter the tails. For the American, French, and Australian indices, we observe in Figure 5 a negative impact of the crisis on α . It means that extreme events tend to occur more frequently and with a larger magnitude. This stylized fact is confirmed by another approach relying on non-parametric densities [40]. The evolution of α in the Chinese market is not similar to the three other indices and the values reached are less significantly different from 2. A progressive recovery toward high values of α is visible for the CAC 40 and the S&P/ASX 200 indices. For the S&P 500 index, we also observe an abrupt increase of α after the peak of the crisis, but it is of limited amplitude and α remains at a fairly low value.

We display in Tables 2 and 3 the range of values reached during the period by the two efficiency indicators, H and m , for the ten stock indices considered. These tables confirm that the greatest impact of COVID-19 on market efficiency occurred for US indices, whatever the efficiency indicator. We also note that inefficiency always leads to negative values of m . Indeed, the observed upper bounds in the period are never significantly different from 0. On the contrary, for the Hurst




Figure 4: Daily evolution through time of the estimated H (left) and m (right) parameters for S&P 500 (top) and CAC 40 (bottom) indices, with $k = 0.5$ instead of 2. The grey lines are the bounds of simulated confidence intervals with a confidence level of 95%, 99%, and 99.5%.

exponent, we find downward peaks below $1/2$ as well as upward peaks above $1/2$, so that it is unclear whether inefficiency leads to high or low values for H . In fact, the presence of fatter tails during the crisis biases this efficiency indicator. So, focusing on the m indicator, we find fairly synchronized peaks for indices of the same region: on 10th March 2020 in the USA, on 16th March 2020 in Europe. Other regions seem less affected: the maximal $|m|$ is indeed lower in Asia and Australia. However, the biased H indicator suggests a similar loss of efficiency in Europe and in Asia. This underpins again the relevance of the refinement of the efficiency indicator to take into account the kurtosis. The two efficiency indicators also lead to opposite conclusions when considering the situation at the end of the sample: according to H , markets are efficient again everywhere, whereas they are significantly inefficient in the USA and in Japan according to m .

5 Conclusion

We have shown to which extent the stock markets become inefficient during the COVID-19 crisis. The efficiency is clearly rejected in the case of the S&P index. On the contrary, the Hurst exponent does not make it possible to conclude about a loss of efficiency for the CAC 40, SSE 180, and S&P/ASX 200 indices. However, we have highlighted a limitation of this indicator: it does not take into account the dynamic kurtosis of price returns. Therefore, if we use the more appropriate

Stock index	Min H	Date of the min	Max H	Date of the max	H on T
S&P 500	0.122	2020-03-04	0.565	2020-02-25	0.452
S&P 100	0.098	2020-03-04	0.565	2020-02-25	0.427
EURO STOXX 50	0.288	2020-02-20	0.665	2020-03-12	0.488
Euronext 100	0.269	2020-02-20	0.680	2020-03-12	0.491
DAX	0.296	2020-02-20	0.678	2020-03-12	0.502
CAC 40	0.250	2020-02-20	0.676	2020-03-12	0.497
Nikkei 225	0.237	2020-01-08	0.677	2020-03-24	0.458
KOSPI	0.343	2020-01-23	0.757	2020-03-25	0.506
SSE 180	0.307	2020-03-05	0.576	2019-12-18	0.440
S&P/ASX 200	0.231	2020-02-20	0.686	2020-03-09	0.482

Table 2: Range of values reached by the efficiency indicator H between November 2019 and June 2020 for ten stock indices. T is the 29th June 2020. An efficient market corresponds to H close to $1/2$.

Stock index	Min m	Date of the min	Max m	Date of the max	m on T
S&P 500	-1.128	2020-03-10	0.065	2020-02-25	-0.355
S&P 100	-1.065	2020-03-10	0.065	2020-02-25	-0.464
EURO STOXX 50	-0.474	2020-03-18	0.043	2020-02-25	-0.188
Euronext 100	-0.603	2020-03-16	0.055	2020-05-15	-0.190
DAX	-0.576	2020-03-16	0.152	2020-03-09	-0.250
CAC 40	-0.690	2020-03-16	0.043	2020-02-25	-0.222
Nikkei 225	-0.417	2020-06-29	0.139	2020-03-19	-0.417
KOSPI	-0.355	2020-06-24	0.148	2020-03-24	-0.290
SSE 180	-0.460	2020-03-10	0.048	2019-11-06	-0.247
S&P/ASX 200	-0.504	2020-03-26	0.097	2020-02-25	-0.225

Table 3: Range of values reached by the efficiency indicator m between November 2019 and June 2020 for ten stock indices. T is the 29th June 2020. An efficient market corresponds to m close to 0.




Figure 5: Daily evolution through time of the estimated α parameter of an fLsm for S&P 500 (top left), CAC 40 (top right), S&P/ASX 200 (bottom left), and SSE 180 (bottom right) indices. The grey lines are the bounds of simulated confidence intervals with a confidence level of 95%, 99%, and 99.5%.

memory parameter of an fLsm, we observe the occurrence of an inefficiency period almost at the beginning of the crisis, even though it is less noticeable for the Chinese and the Australian indices.

We have also introduced in this paper the tools used for this analysis, namely the estimation of a dynamic stable distribution along with the estimation of the dynamic Hurst exponent and memory parameter of an fLsm. An important free parameter in this approach is the discount factor which is related to the speed at which the weight of past information decreases. We have used a selection rule based on the minimization of a criterion depicting the uniformity and the independence of the PITs, consistently with the literature on the validation of density forecasts.

Acknowledgement

The authors deeply thank Akin Arslan, Thomas Barrat, and Sarah Bouabdallah for their valuable help in the implementation of some of the methods described in this paper. MG thanks the participants of the Econphysics Colloquium 2021 in Lyon for useful comments.

References

- [1] F. Abergel and M. Politi. Optimizing a basket against the efficient market hypothesis. *Quantitative finance*, 13(1):13–23, 2013.
- [2] A. Alvarez and P. Olivares. Méthodes d’estimation pour des lois stables avec des applications en finance. *Journal de la société française de statistique*, 146(4):23–54, 2005.
- [3] P. Anagnostidis, C. Varsakelis, and C.J. Emmanouilides. Has the 2008 financial crisis affected stock market efficiency? The case of Eurozone. *Physica A: statistical mechanics and its applications*, 447:116–128, 2016.
- [4] K. Arias-Calluari, F. Alonso-Marroquin, M.N. Najafi, and M. Harré. Forecasting the effect of COVID-19 on the S&P500. *working paper*, 2020.
- [5] A. Ayache and M.S. Taqqu. Multifractal processes with random exponent. *Publicacions matemàtiques*, 49(2):459–486, 2005.
- [6] S.R. Baker, N. Bloom, S.J. Davis, K.J. Kost, M.C. Sammon, and T. Viratyosin. The unprecedented stock market impact of COVID-19. *working paper*, 2020.
- [7] A.L. Barabási and T. Vicsek. Multifractality of self-affine fractals. *Physical review A*, 44(4):2730–2733, 1991.
- [8] J.A. Batten, C. Ciner, B.M. Lucey, and P.G. Szilagyi. The structure of gold and silver spread returns. *Quantitative finance*, 13(4):561–570, 2013.
- [9] L. Belkacem, J. Lévy Véhel, and C. Walter. CAPM, risk and portfolio selection in α -stable markets. *Fractals*, 8(1):99–115, 2000.
- [10] A. Benassi, S. Cohen, and J. Istas. Identifying the multifractional function of a Gaussian process. *Statistics and probability letters*, 39(4):337–345, 1998.
- [11] A. Benassi, S. Jaffard, and D. Roux. Elliptic Gaussian random processes. *Revista matemática iberoamericana*, 13(1):19–90, 1997.
- [12] S. Bianchi. Pathwise identification of the memory function of multifractional Brownian motion with application to finance. *International journal of theoretical and applied finance*, 8(2):255–281, 2005.
- [13] S. Bianchi and A. Pantanella. Pointwise regularity exponents and market cross-correlations. *International review of business research papers*, 6(2):39–51, 2010.
- [14] S. Bianchi and A. Pantanella. Pointwise regularity exponents and well-behaved residuals in stock markets. *International journal of trade, economics and finance*, 2(1):52–60, 2011.
- [15] S. Borak, W. Härdle, and R. Weron. Stable distributions. In *Statistical tools for finance and insurance*, pages 21–44. Springer, Berlin, Heidelberg, 2005.
- [16] D.O. Cajueiro and B.M. Tabak. The Hurst exponent over time: testing the assertion that emerging markets are becoming more efficient. *Physica A: statistical mechanics and its applications*, 336(3-4):521–537, 2004.
- [17] G. Cao, J. Cao, and L. Xu. Asymmetric multifractal scaling behavior in the Chinese stock market: Based on asymmetric MF-DFA. *Physica A: statistical mechanics and its applications*, 392(4):797–807, 2013.

- [18] A. Charles and O. Darné. Variance-ratio tests of random walk: an overview. *Journal of economic surveys*, 23(3):503–527, 2009.
- [19] T. Chordia, R. Roll, and A. Subrahmanyam. Evidence on the speed of convergence to market efficiency. *Journal of financial economics*, 76(2):271–292, 2005.
- [20] J.-F. Coeurjolly. Identification of multifractional Brownian motion. *Bernoulli*, 11(6):987–1008, 2005.
- [21] J.A. Coutts, J. Roberts, and T.C. Mills. Parameter stability in the market model: tests and time varying parameter estimation with UK data. *Journal of the royal statistical society: series D (the statistician)*, 46(1):57–70, 1997.
- [22] J.K. Dagsvik, M. Fortuna, and S.H. Moen. How does temperature vary over time?: evidence on the stationary and fractal nature of temperature fluctuations. *Journal of the royal statistical society: series A (statistics in society)*, 183(3):883–908, 2020.
- [23] L. de Haan and T.T. Pereira. Estimating the index of a stable distribution. *Statistics & probability letters*, 41(1):39–55, 1999.
- [24] G. De Rossi and A.C. Harvey. Time-varying quantiles. *working paper*, 2006.
- [25] T. Di Matteo. Multi-scaling in finance. *Quantitative finance*, 7(1):21–36, 2007.
- [26] F.X. Diebold, T.A. Gunther, and A.S. Tay. Evaluating density forecasts, with applications to financial risk management. *International economic review*, 39:863–883, 1998.
- [27] W.H. DuMouchel. On the asymptotic normality of the maximum-likelihood estimate when sampling from a stable distribution. *Annals of statistics*, 1(5):948–957, 1973.
- [28] R.F. Engle and S. Manganelli. CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of business & economic statistics*, 22(4):367–381, 2004.
- [29] K.J. Falconer, R. Le Guével, and J. Lévy Véhel. Localizable moving average symmetric stable and multistable processes. *Stochastic models*, 25(4):648–672, 2009.
- [30] E.F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of finance*, 25(2):383–417, 1970.
- [31] E.F. Fama and R. Roll. Parameter estimates for symmetric stable distributions. *Journal of the American statistical association*, 66(334):331–338, 1971.
- [32] W. Feng, Y. Yang, Q. Yuan, C. Gu, and H. Yang. Evolution of scaling behaviors in currency exchange rate series. *Fractals*, 27(2):1950005, 2019.
- [33] M. Frezza. Modeling the time-changing dependence in stock markets. *Chaos, solitons & fractals*, 45(12):1510–1520, 2012.
- [34] M. Garcin. Estimation of time-dependent Hurst exponents with variational smoothing and application to forecasting foreign exchange rates. *Physica A: statistical mechanics and its applications*, 483:462–479, 2017.
- [35] M. Garcin. Hurst exponents and delampertized fractional Brownian motions. *International journal of theoretical and applied finance*, 22(5):1–26, 2019.
- [36] M. Garcin. A comparison of maximum likelihood and absolute moments for the estimation of Hurst exponents in a stationary framework. *working paper*, 2020.

- [37] M. Garcin. Fractal analysis of the multifractality of foreign exchange rates. *Mathematical methods in economics and finance*, 13-14(1):49–73, 2020.
- [38] M. Garcin and D. Guégan. Probability density of the empirical wavelet coefficients of a noisy chaos. *Physica D: nonlinear phenomena*, 276:28–47, 2014.
- [39] M. Garcin and D. Guégan. Wavelet shrinkage of a noisy dynamical system with non-linear noise impact. *Physica D: nonlinear phenomena*, 325:126–145, 2016.
- [40] M. Garcin, J. Klein, and S. Laaribi. Estimation of time-varying kernel densities and chronology of the impact of COVID-19 on financial markets. *working paper*, 2020.
- [41] R.H. Gerlach, C.W. Chen, and N.Y. Chan. Bayesian time-varying quantile forecasting for value-at-risk in financial markets. *Journal of business & economic statistics*, 29(4):481–492, 2011.
- [42] C. Gouriéroux and J. Jasiak. Dynamic quantile models. *Journal of econometrics*, 147(1):198–205, 2008.
- [43] A. Harvey and V. Oryshchenko. Kernel density estimation for time series data. *International journal of forecasting*, 28(1):3–14, 2012.
- [44] J.R. Hosking. L-moments: Analysis and estimation of distributions using linear combinations of order statistics. *Journal of the royal statistical society: series B (methodological)*, 52(1):105–124, 1990.
- [45] H.E. Hurst. Long-term storage capacity of reservoirs. *Transactions of the American society of civil engineering*, 116(1):770–808, 1951.
- [46] M. Ito, A. Noda, and T. Wada. International stock market efficiency: a non-Bayesian time-varying model approach. *Applied economics*, 46(23):2744–2754, 2014.
- [47] G. Julián-Moreno, J.E.L. de Vergara, I. González, L. de Pedro, J. Royuela-del Val, and F. Simmross-Wattenberg. Fast parallel α -stable distribution function evaluation and parameter estimation using OpenCL in GPGPUs. *Statistics and computing*, 27(5):1365–1382, 2017.
- [48] W. Karcher, E. Shmileva, and E. Spodarev. Extrapolation of stable random fields. *Journal of multivariate analysis*, 115:516–536, 2013.
- [49] R. Koenker and Z. Xiao. Quantile autoregression. *Journal of the American statistical association*, 101(475):980–990, 2006.
- [50] S.M. Kogon and D.B. Williams. Characteristic function based estimation of stable parameters. In R. Adler, R. Feldman, and M. Taqqu, editors, *A practical guide to heavy tails*, pages 311–335. Birkhauser, 1998.
- [51] I.A. Koutrouvelis. Regression-type estimation of the parameters of stable laws. *Journal of the American statistical association*, 75(372):918–928, 1980.
- [52] L. Kristoufek. On Bitcoin markets (in)efficiency and its evolution. *Physica A: statistical mechanics and its applications*, 503:257–262, 2018.
- [53] L. Kristoufek and M. Vosvrda. Measuring capital market efficiency: Global and local correlations structure. *Physica A: statistical mechanics and its applications*, 392(1):184–193, 2013.
- [54] P. Lambert and J.K. Lindsey. Analysing financial returns by using regression models based on non-symmetric stable distributions. *Journal of the royal statistical society: series C (applied statistics)*, 48(3):409–424, 1999.

- [55] R. Le Guével. An estimation of the stability and the localisability functions of multistable processes. *Electronic journal of statistics*, 7:1129–1166, 2013.
- [56] J.B. Levy and M.S. Taqqu. The long-range dependence of linear log-fractional stable motion. *Communications on stochastic analysis*, 5(1):187–210, 2011.
- [57] Y. Liang and W. Chen. A survey on computing Lévy stable distributions and a new MATLAB toolbox. *Signal processing*, 93(1):242–251, 2013.
- [58] F. Lillo and J.D. Farmer. The long memory of the efficient market. *Studies in nonlinear dynamics & econometrics*, 8(3):1–33, 2004.
- [59] K.P. Lim, R.D. Brooks, and J.H. Kim. Financial crisis and stock market efficiency: Empirical evidence from Asian countries. *International review of financial analysis*, 17(3):1571–591, 2008.
- [60] A.W. Lo. The adaptive markets hypothesis. *Journal of portfolio management*, 30(5):15–29, 2004.
- [61] B. Mandelbrot and J. van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM review*, 10(4):422–437, 1968.
- [62] J.H. McCulloch. Simple consistent estimators of stable distribution parameters. *Communications in statistics - simulation and computation*, 15(4):1109–1136, 1986.
- [63] S.K. Mitra. Is Hurst exponent value useful in forecasting financial time series? *Asian social science*, 8(8):111–120, 2012.
- [64] R. Morales, T. Di Matteo, R. Gramatica, and T. Aste. Dynamical generalized Hurst exponent as a tool to monitor unstable periods in financial time series. *Physica A: statistical mechanics and its applications*, 391(11):3180–3189, 2012.
- [65] V. Naumov and O. Martikainen. Exponentially weighted simultaneous estimation of several quantiles. *World academy of science, engineering and technology*, 8:563–568, 2007.
- [66] A. Noda. On the evolution of cryptocurrency market efficiency. *working paper*, 2020.
- [67] J.P. Nolan. An algorithm for evaluating stable densities in Zolotarev’s (M) parameterization. *Mathematical and computer modelling*, 29(10-12):229–233, 1999.
- [68] J.P. Nolan. Modeling financial data with stable distributions. In *Handbook of heavy tailed distributions in finance*, pages 105–130. North-Holland, 2003.
- [69] A.S. Paulson, E.W. Holcomb, and R.A. Leitch. The estimation of the parameters of the stable laws. *Biometrika*, 62(1):163–170, 1975.
- [70] B.M. Pavlyshenko. Regression approach for modeling COVID-19 spread and its impact on stock market. *working paper*, 2020.
- [71] R.F. Peltier and J. Lévy Véhel. Multifractional Brownian motion: definition and preliminary results. *working paper*, 1995.
- [72] C.K. Peng, S.V. Buldyrev, S. Havlin, M. Simons, H.E. Stanley, and A.L. Goldberger. Mosaic organization of DNA nucleotides. *Physical review E*, 49(2):1685–1689, 1994.
- [73] E. Peters. *Fractal market analysis: applying chaos theory to investment and analysis*. John Wiley & sons, New York, 1991.

- [74] S.J. Press. Estimation in univariate and multivariate stable distributions. *Journal of the American statistical association*, 67(340):842–846, 1972.
- [75] D.M. Rösch, A. Subrahmanyam, and M.A. Van Dijk. The dynamics of market efficiency. *The review of financial studies*, 30(4):1151–1187, 2017.
- [76] D. Salas-González, J.M. Górriz, J. Ramírez, Schloegl, E.W. M., Lang, and A. Ortiz. Parameterization of the distribution of white and grey matter in MRI using the α -stable distribution. *Computers in biology and medicine*, 43(5):559–567, 2013.
- [77] D. Salas-Gonzalez, E.E. Kuruoglu, and D.P. Ruiz. Modelling with mixture of symmetric stable distributions using Gibbs sampling. *Signal processing*, 90(3):774–783, 2010.
- [78] G. Samorodnitsky, , and M.S. Taqqu. *Stable non-Gaussian processes: stochastic models with infinite variance*. Chapman and Hall, New York, London, 1994.
- [79] S. Stoev and M.S. Taqqu. Simulation methods for linear fractional stable motion and FARIMA using the Fast Fourier Transform. *Fractals*, 12(1):95–121, 2004.
- [80] M. Šapina, M. Garcin, K. Kramarić, K. Milas, D. Brdarić, and M. Pirić. The Hurst exponent of heart rate variability in neonatal stress, based on a mean-reverting fractional Lévy stable motion. *Fluctuation and noise letters*, 19(3):2050026, 2020.
- [81] M. Šapina, M. Košmider, K. Kramarić, M. Garcin, M. Pirić, K. Milas, and D. Brdarić. Asymmetric detrended fluctuation analysis in neonatal stress. *Physiological measurement*, 39(8):085006, 2018.
- [82] A. Weron, K. Burnecki, S. Mercik, and K. Weron. Complete description of all self-similar models driven by Lévy stable noise. *Physical review E*, 71(1):016113, 2005.
- [83] K. Yu and M. Jones. Local linear quantile regression. *Journal of the American statistical association*, 93(441):228–237, 1998.
- [84] V.M. Zolotarev. *One-dimensional stable distributions*, volume 65. Translations of mathematical monographs, AMS, 1986.