



Discrete hierarchy of sizes and performances in the exchange-traded fund universe



B. Vandermarliere^{a,b,*}, J. Ryckebusch^b, K. Schoors^b, P. Cauwels^c, D. Sornette^c

^a Department of General Economics, Ghent University, Belgium

^b Department of Physics and Astronomy, Ghent University, Belgium

^c Department of Management, Technology and Economics, ETH Zürich, Switzerland

HIGHLIGHTS

- The size distribution of exchange-traded funds exhibits a discrete hierarchy of sizes.
- This is shown with a spectral analysis of residuals and by studying the generalized derivative.
- Several explanatory mechanisms are offered for this discrete scale invariance.
- Larger ETFs exhibit a significantly stronger similarity compared with smaller ETFs.

ARTICLE INFO

Article history:

Received 31 August 2016

Available online 17 November 2016

Keywords:

Econophysics

Exchange-traded funds (ETFs)

Probability density function of ETF sizes

Discrete scale invariance

ETF size layers and performance measures

ABSTRACT

Using detailed statistical analyses of the size distribution of a universe of equity exchange-traded funds (ETFs), we discover a discrete hierarchy of sizes, which imprints a log-periodic structure on the probability distribution of ETF sizes that dominates the details of the asymptotic tail. This allows us to propose a classification of the studied universe of ETFs into seven size layers approximately organized according to a multiplicative ratio of 3.5 in their total market capitalization. Introducing a similarity metric generalizing the Herfindhal index, we find that the largest ETFs exhibit a significantly stronger intra-layer and inter-layer similarity compared with the smaller ETFs. Comparing the performance across the seven discerned ETF size layers, we find an inverse size effect, namely large ETFs perform significantly better than the small ones both in 2014 and 2015.

© 2016 Published by Elsevier B.V.

1. Introduction

An exchange-traded fund (ETF) can be thought of as a portfolio of stocks, commodities, or bonds, which is traded like stocks on stock exchanges. Exchange-traded funds have been made available as investment funds in the US in the early nineties and in Europe in the late nineties. Ever since, ETFs have emerged as a very important investment vehicle attracting ever increasing volumes of capital. Its attractiveness is partly due to the relatively low management and transaction costs involved, an element that is particularly important in times of low yields and low interest rates. Exchange-traded funds represent an increasingly important investment vehicle with potential hazards for systemic risk and possible dangerous menaces for the financial system [1–3]. For example, it has been shown that arbitrageurs can contribute to cross-sectional return co-movement via ETF arbitrage. The presence of a stock in ETFs increases return co-movement at both the fund and

* Correspondence to: Proeftuinstraat 86, B-9000 Ghent.

E-mail addresses: benjamin.vandermarliere@ugent.be (B. Vandermarliere), Jan.Ryckebusch@UGent.be (J. Ryckebusch), Koen.Schoors@UGent.be (K. Schoors), pcauwels@ethz.ch (P. Cauwels), dsornette@ethz.ch (D. Sornette).

the stock levels, where the effect is strongest among small and illiquid stocks [4]. These days, ETFs come in many different types of flavours [5]. For example, the degree of active management varies very much from one ETF to another.

The focus of this paper is on establishing a taxonomy of the equity ETF landscape on the basis of their size. From our discussion we exclude leveraged ETFs and ETFs holding bonds and commodities, mainly to not overly complicate the analysis. As our focus is on determining the robust and stylized features of the equity ETF landscape using size, we do not segregate by types of ETFs, for example in terms of managed versus active versus passive, or index tracking ETFs.

Size distributions often carry information about the underlying dynamics of a system. The analysis of the distribution of the equity ETF sizes described below discloses some features that suggest departures from a simple power-like tail. The occurrence of a fat tail in the distribution of ETF sizes does not really come as a surprise given the well-documented approximate Zipf-law distribution of firm capitalizations [6]. The fact, however, that there are strong indications that the tail is decorated with some log-periodic structure is remarkable. As this structure is connected with discrete scale invariance, one can infer some interesting constraints on the underlying dynamics of the equity ETF universe. Accordingly, we consider the disclosed log-periodic structure in the size distribution as a natural tool for classification of the universe of ETFs. The inferred classification of the ETFs in several size layers is used to study various economic indicators. We address questions like: ‘How similar are the various kinds of ETFs?’; ‘How do ETFs distribute their holdings over the wide landscape of possible holdings?’; and ‘Is there a connection between the ETF size and their performance?’. These questions are naturally motivated by the existence of the size effect, exploited in the famous Fama–French 3 factor model [7] that also addresses the fundamental issues of the relationships between diversification and performance.

The remainder of this paper is organized as follows. In Section 2 we present our empirical analysis of the equity ETF size distribution. We start off (Section 2.1) with providing details of the ETF size data used and with performing a maximum-likelihood fit to their distribution. This reveals indications for an interesting discrete hierarchical structure in the ETF size distribution that is discussed in more detail in Section 2.2. In order to put this structure on more solid grounds and to get better hold on the disclosed periodicity in the size distribution, in Section 2.3 we pursue a detailed analysis of the ETF size distribution using kernel density estimation and Lomb periodograms. In Section 2.4 we sketch some dynamical features of the ETF universe that may give rise to the observed hierarchical structure. We work out in detail how a model based on nonextensive (or, Tsallis) statistical mechanics, a current generalization of Boltzmann–Gibbs (BG) statistical mechanics, can give rise to the discerned oscillatory structures in the ETF size distribution. The basic premises of the proposed model is that the system consisting of all ETFs operates as an open system in a capital reservoir. The size of the ETF system is subject to capital exchange with the reservoir, whereby there is a mechanism of both preferential attachment and growth. In Section 3 we introduce a classification into seven layers of the equity ETFs based on the discerned log-periodic hierarchy. We also explore how the economic properties vary over the various size layers. Thereby, we investigate the intra-layer and inter-layer similarities (Section 3.1), the variations in the stock holding ubiquity and capitalization over the different layers (Section 3.2), and the connection between layer and performance (Section 3.3). Our conclusions are drawn in Section 4.

2. Analysis of the distribution of ETF sizes

2.1. Distribution of total net asset values of ETFs

At the end of 2014, we collected data for all exchange-traded funds (ETFs) labelled as equity ETFs from Thomson Reuters Eikon. This resulted in a set of 479 ETFs for which we obtained the total net assets and the entire composition of their portfolios. In total, this comprised 11,643 different assets and about 100,000 positions, for a total net assets over all ETFs of 1.399×10^{12} US\$. Fig. 1 includes the complementary cumulative distribution function (CCDF) of the total net assets of ETFs, i.e., the fraction of ETFs of total net assets larger than or equal to S . Also shown is the CCDF of the log-normal that best fits the data, as obtained by the maximum-likelihood method. The probability density function (PDF) of the log-normal law $\ln \mathcal{N}$ reads

$$\ln \mathcal{N}(\mu_L, \sigma_L^2) = \frac{1}{x\sqrt{2\pi\sigma_L^2}} e^{-\frac{(\ln x - \mu_L)^2}{2\sigma_L^2}}, \quad (1)$$

with μ_L the location and σ_L the scale parameter whose maximum-likelihood estimates are $\hat{\mu}_L = 18.7$ and $\hat{\sigma}_L = 2.24$. This corresponds to the mode (or most probable) ETF size of approximately 130×10^6 US\$ and a mean ETF size of 1.6×10^9 US\$. The much larger value of the mean compared to the mode reflects the existence of a very strong ‘‘fat tail’’ quantified by $\hat{\sigma}_L$.

When referring to fat tails, it is often convenient to use power law distributions. The tail of a log-normal distribution with large variance (as found here) is difficult to distinguish from a power law distribution (see e.g. Ref. [8] and Section 4.1.3 of Ref. [9]). Indeed, visually, the tail of the empirical CCDF shown in Fig. 1 seems roughly compatible with an asymptotic power law with an exponent of about 1 (Zipf’s law). Such an approximate asymptotic Zipf’s law has been documented for the distribution of firm sizes [6]. The fact that a similar approximate behaviour in the asymptotic tail is observed for the distribution of ETF sizes is not really a surprise as it can be expected from the presence of two joint and mutually reinforcing mechanisms. First, it is well known that the size of individual firms approximately obeys Zipf’s law [6,10–13]. This result is robust [14] and has been confirmed for different countries [10] and for several measures of firm size including

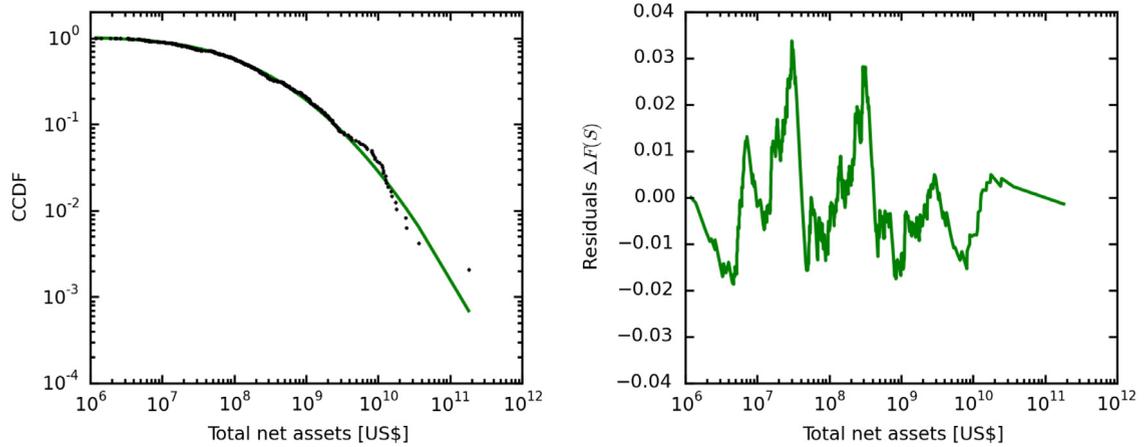


Fig. 1. The left figure shows the empirical complementary cumulative distribution function (CCDF) of total net assets as a function of total net assets (decimal log–log scale) for the set of 479 equity ETFs collected from Thomson Reuters Eikon in December 2014. The green full line is the maximum likelihood estimation of the lognormal distribution of Eq. (1) with $\hat{\mu}_l = 18.7$ and $\hat{\sigma}_l = 2.24$. The right figure shows the corresponding residuals – the difference between the lognormal fit and the data – as a function of total net assets. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

number of employees, profits, sales, value added, and market capitalizations. Therefore, randomly generated portfolios with weights roughly proportional to firm capitalizations will also have an asymptotic Zipf distribution in their tail, as a result of the generalized central limit theorem (see Section 4 of Ref. [9] for a pedagogical presentation). Second, Zipf’s law appears quite generically from the combination of three very robust ingredients, namely ETFs are born, they grow via proportional growth and then can also die or close. As outlined in Refs. [15, Chapter 10] and [16], mergers and acquisitions do not change significantly the overall picture. If the stochastic component of proportional growth is large, Zipf’s law is generically an excellent approximation of the tail [15,17].

Therefore, the observation of a fat tail that looks roughly like Zipf’s law is not of much significance. What is much more surprising is the existence of very large deviations from a smooth tail, as made apparent by the structure of the residuals $\Delta F(S)$ of the lognormal calibration also shown in Fig. 1. The pattern of these residuals clearly dominates the question of what is the asymptotic behaviour at large ETF sizes. A first preliminary conclusion is that there appears to be significant more texture to the tail of the CCDF than just a power law or log-normal tail. We now turn to the detailed quantitative analysis of these residuals.

2.2. Evidence of a discrete hierarchical texture in the distribution of ETF sizes by spectral analysis of the residuals

A visual inspection of the residuals $\Delta F(S)$ shown in Fig. 1 suggests a noisy oscillation. To ascertain the significance of this observation, we calculate the Lomb periodogram of these residuals, shown in Fig. 2. The use of the Lomb periodogram, instead of a Fourier transform, is required as a result of the non-even spacing of the pseudo-time variable, namely the logarithm of the total net assets S . Recall that the Lomb periodogram is a method for spectral analysis, which quantifies the contribution of each frequency to a given signal, based on the local least square fit of sine functions to the data [18]. In our case, the signal is the function $\Delta F(S)$ shown in the right panel of Fig. 1 expressed as a function of $\ln S$. A statistically significant oscillatory component would mean that $\Delta F(S)$ can be expressed as

$$\Delta F(S) = A + B \cos[\omega \ln S + \phi] + \mathcal{O}((\ln S)^2), \quad (2)$$

where (A, B, ϕ) are three constants and $\mathcal{O}((\ln S)^2)$ is a second-order residual function of amplitude much smaller than B .

It is important to note that ω is not an angular frequency in the usual sense, as it is the conjugate variable to $\ln S$ and not to S . In other words, as already mentioned, the Lomb spectral analysis is performed in terms of the variable $\ln S$. Thus, the presence of periodicity in the $\ln S$ variable means that the residual function $\Delta F(S)$ is log-periodic in the function S , i.e. it exhibits the symmetry of “discrete scale invariance” [19,20]. In particular, ω is dimensionless. Fig. 2 exhibits an extremely large peak at $\omega = 2.5 \pm 0.2$, which embodies the value of the scaling ratio $p_1 := \exp(2\pi/\omega_1) = 12.3$ for $\omega_1 = 2.5$ and quantifies the ratio of the geometrical series S_n at which the cosine in expression (2) is equal to 1 (i.e. $\omega \ln(S_n) + \phi = 2\pi n$, where n is an arbitrary integer). According to extensive simulations in the possible presence of heavy-tailed and correlated noise [21], one can ascertain that this peak at $\omega = 2.5 \pm 0.2$ is statistically highly significant. It expresses the existence of a discrete hierarchy of ETF sizes, roughly spaced according to the ratio $p_1 = 12.3$. Note also the existence of the two smaller peaks at $\omega_2 = 5.2 \pm 0.2 \approx 2\omega_1$ and $\omega_3 = 8.2 \pm 0.4 \approx 3\omega_1$. The presence of these harmonics strengthens the evidence for log-periodicity [22,23]. The peak at the lowest value $\omega = 0.78$ corresponds to an oscillation of about the size of the entire range of values, which can be expected just from cumulative noise effect [24] and we thus ignore it.

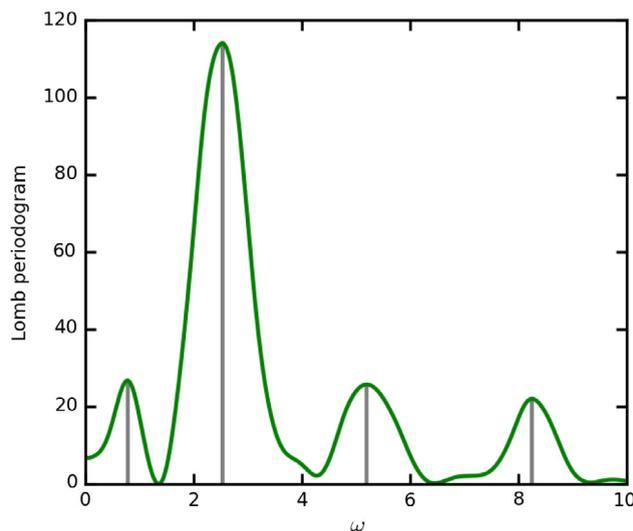


Fig. 2. Lomb periodogram of the residual function $\Delta F(S)$ shown in the right figure of Fig. 1. Here, ω is the conjugate variable to the logarithm of the ETF sizes. The occurrence of the three peaks at $\omega_1 = 2.5$ (large peak), $\omega_2 = 5.2 \approx 2\omega_1$ and $\omega_3 = 8.2 \approx 3\omega_1$ is interpreted in the text. As explained in the text, the peak at $\omega = 0.78$ is likely due to the conjunction of noise in the presence of a finite range of analysis.

2.3. Generalized derivative and Lomb periodogram of the PDF of ETF sizes

In science, and especially in statistics, it is challenging to prove the absolute reality of an empirical observation. But one can scrutinize the data with a variety of distinct and complementary methods, which altogether may provide confirming evidence of the claimed phenomenon and thus stronger trust in its genuine existence. Because the claim of discrete scale invariance and of a discrete hierarchical structure in the distribution of ETF sizes is rather unexpected and of possible economic importance, we present a detailed analysis of the observed log-periodicity using a completely different methodology, which follows precisely the procedure described in Refs. [25,26]. The procedure has three components: (i) the kernel density estimation (KDE) of the probability density function (PDF) (instead of using the CCDF) of the ETF sizes; (ii) the construction of the generalized (H, q) -derivative of the PDF, and (iii) the calculation of its Lomb periodogram.

Working with the PDF of ETF sizes has the advantage compared with the CCDF of being a local measure of the distribution, hence less prone to the influence of contamination by systematic biases. However, the PDF is more noisy and harder to estimate with limited data. A standard and robust estimation method consists in constructing its kernel density estimator, which is a kind of smoothed histogram. The Gaussian KDE of the PDF of the logarithms $\ln S$ of ETF sizes is defined as

$$\hat{f}_\sigma(\ln S) = \frac{1}{N} \sum_{i=1}^N \mathcal{N}(\ln S - \ln S_i, \sigma^2), \quad (3)$$

with $\mathcal{N}(0, \sigma^2)$ a zero-mean Gaussian distribution with variance σ^2 , and the sum is over the $N = 479$ data points $\ln S_i$. Further, in the context of KDE one refers to σ as the bandwidth. Fig. 3 shows the KDE of the PDF of ETF sizes for different bandwidths σ . The optimal bandwidth $\sigma_o = 0.22$ is determined with cross validation. Recall that, in cross validation, the model is first fit to part of the data, after which a quantitative metric is computed to determine how well this model fits the remaining data. Obviously, there are strong indications for oscillatory behaviour emerging from the KDE analysis of the PDF of ETF sizes. We use the generalized derivative of this function in order to gain a better insight into this oscillatory behaviour.

The generalized (H, q) -derivative of a function $f(x)$ is defined as [27,28]

$$D_q^H f(x) \equiv \frac{f(x) - f(qx)}{[(1-q)x]^H}, \quad (4)$$

and provides a robust metric of the trend or slope of a function. This is particularly useful to detect features in a noisy function, such as the PDF of ETF sizes studied here. Fig. 3 includes the D_q^H of the KDE of the PDF of ETF sizes for three bandwidths. As recommended in Refs. [27,28], we have scanned H from 0.5 to 0.9 in steps of 0.08, and q from 0.65 to 0.95 in steps of 0.06 and found that the results are robust. Accordingly, the displayed $D_{q=0.65}^{H=0.5} f(\ln S)$ results can be considered representative. One can observe three to four well formed oscillations in the logarithm of the ETF sizes S , quite similarly to the observations of the cumulative distribution approach.

In order to extract the strongest contributing frequencies, we have computed the Lomb periodogram of the generalized (H, q) -derivative of the Gaussian KDE of the PDF of the logarithms of the ETF sizes. We choose the kernel estimation with $\sigma_o/2$ as it is representative of the other estimators but exhibits the largest oscillatory amplitudes. The resulting periodograms

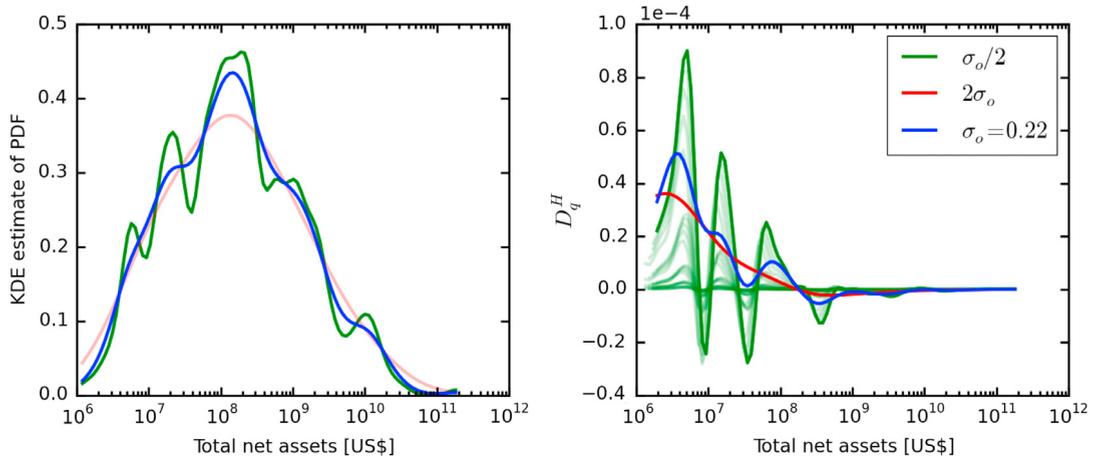


Fig. 3. The left figure shows the Gaussian KDE of the PDF of the logarithms of the ETF sizes for three different values of the bandwidth. The blue line is for the optimal bandwidth ($\sigma_o = 0.22$) determined using cross validation. The green and red line correspond with a bandwidth of $\sigma_o/2$ and $2\sigma_o$. The right figure shows the generalized derivative $D_q^{H=0.5}$ of the curves of the left figure. For $\sigma_o/2$ we also show D_q^H for different combinations of the values ($0.5 \leq H \leq 0.9, 0.65 \leq q \leq 0.95$). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

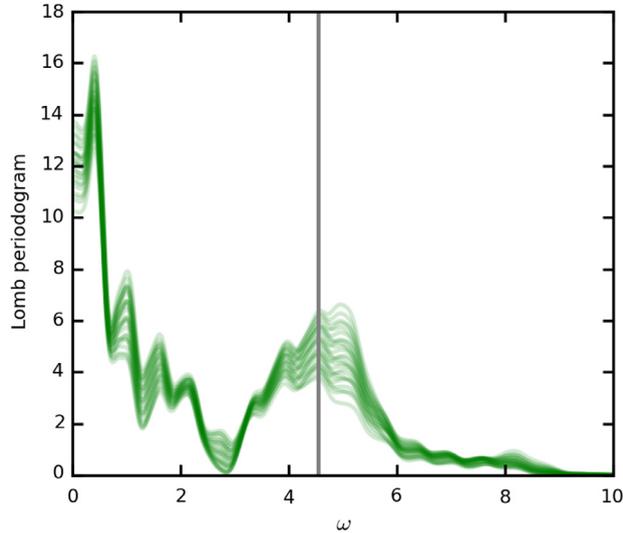


Fig. 4. Lomb periodogram of the generalized (H, q) -derivative of the Gaussian KDE of the PDF of the decimal logs of the ETF sizes with $\sigma_o/2$ and different values of the combination ($0.5 \leq H \leq 0.9, 0.65 \leq q \leq 0.95$). The black vertical line is at the centre value of $\omega = 4.6 \pm 0.6$, which corresponds to the scaling ratio for the ETF sizes S of $p = \exp(2\pi/\omega) = 3.9 \pm 0.4$.

are shown in Fig. 4. First, the peaks at low angular log-frequencies $\omega < 1.5$ represent oscillations with a wavelength of about the size of the entire range of values and hence can be ignored as explained above [24]. There is only one noticeable peak at a value of $\omega = 4.6 \pm 0.6$ that can be put in correspondence with the second harmonic $\omega_2 = 5.2 \pm 0.2$ previously reported. This angular log-frequency corresponds to a scaling ratio of $p_2 = \exp(2\pi/\omega_2) = 3.9 \pm 0.4$. Note that, when averaging the Lomb periodogram over the scanned H and q , the same estimate $\omega = 4.6 \pm 0.6$ for the unique significant peak is obtained, providing evidence that it has a real existence. There is no significant peak at ω_1 , likely as a result of the high-frequency noise associated with the construction of the PDF. Note that the general available theory of log-periodic functions indicates that different harmonics can have very different amplitudes that depend on subtle properties of the problem [29]. In other words, one should not be surprised that the different harmonics of log-periodicity express themselves with different amplitudes in distinct signals.

2.4. Mechanisms of discrete scale invariance in the PDF of ETF sizes

Sections 2.2 and 2.3 have presented statistically significant evidence of the existence of a discrete hierarchical structure in the distribution of ETF sizes, with preferred scaling ratios approximately equal to $p_1 \approx 12$ and $p_2 = \sqrt{p_1} \approx 3.5$.

Ref. [19] provides a review of the many mechanisms that can produce such a discrete hierarchy. While we cannot offer a definite mechanism and test for its relevance, the most likely candidates are the Kesten process [30,31] and aggregation/fragmentation dynamics [32]. The Kesten process can be used to describe the growth of portfolio as a result of the joint addition of new deposits and of stochastic proportional growth. Log-periodic PDFs emerge quite robustly if the multiplicative stochastic factors are not too broadly distributed. The aggregation/fragmentation dynamics could be also a limiting process for the formation of ETF portfolios, for which it can be shown that discrete scale invariance may emerge for quite general aggregation/fragmentation kernels.

Another explanation for the occurrence of distributions with a power law decorated by log-periodic oscillations finds its origin in the formalism of nonextensive statistical mechanics [33]. At any instant of time, the universe of ETFs behaves as an open system that seeks to find equilibrium with the whole of the capital market that acts as a reservoir. In the context of equilibrium statistical physics, this equilibrium process [34] gives rise to the well-known Boltzmann–Gibbs distribution of the sizes of the ETFs

$$p^{BG}(0 \leq S \leq \infty) = \frac{1}{T_0} \exp -\frac{S}{T_0}, \quad (5)$$

where the temperature T_0 acts as a typical scale parameter for the size of the ETFs. In this picture, all ETFs (independent of size) are subject to a similar stochastically driven capital exchange with the reservoir of the complete market. In other words, the universe of ETFs is embedded in the reservoir of investment products and the exchange between any element in the ETF universe and the reservoir can be parametrized by a single scale parameter T_0 . The larger T_0 the larger the average size of the ETFs. Obviously, Eq. (5) does not give rise to fat tails in the distribution of the ETF sizes and cannot be considered realistic given the observations of Section 2.1. A generalization, however, proceeds as follows. The Boltzmann–Gibbs exponential distribution (5) is a solution to the following equation

$$\frac{dP^{BG}(S)}{dS} = -\frac{1}{T_0} P^{BG}(S). \quad (6)$$

In nonextensive statistical mechanics, this equation is extended by adding a nonextensivity parameter n

$$\frac{dP(S)}{dS} = -\frac{1}{T(S)} P(S) = -\frac{1}{T_0 + \frac{S}{n}} P(S). \quad (7)$$

A particular solution to this equation is known as the Tsallis distribution

$$P(S) = \frac{n-1}{nT_0} \left(1 + \frac{S}{nT_0}\right)^{-n}. \quad (8)$$

The Tsallis distribution nicely interpolates between the Boltzmann–Gibbs exponential for ETF sizes S smaller than the scale parameter T_0 and a power law tail for $S \gg T_0$. The nonextensivity parameter n makes the temperature – in the current context a proxy for the typical amount of capital exchange of an ETF with the reservoir of investment products – dependent on the actual ETF size S . Loosely speaking, the parameter $\frac{1}{n}$ can be interpreted as a measure for the degree of preferential attachment [35], or the extent to which the rich ETFs get richer. In the current context, the n accounts for the fact that there is an increased linear tendency of an ETF to accrue money from the reservoir of investment products as it increases in size S . The quantity n determines the asymptotic behaviour of the distribution $P(S)$. The normalization condition $\int_0^\infty P(S) dS$ of the distribution (8) requires that $n > 1$. The Tsallis distribution, for example, provides an excellent fit to transverse momentum distributions in high-energy collisions with values of n of the order 6–8 [36]. Obviously, the limit $n \rightarrow \infty$ corresponds to a vanishing preferential attachment effect. The smaller n the larger the difference between the temperature associated with the small and the large ETFs. Small ETFs, that are defined as those with a current size smaller than the scale T_0 experience a temperature $T \approx T_0$ in their interaction with the capital market. Large ETFs, defined as ETFs larger than the scale T_0 , experience a temperature $T(S)$ that scales linearly with their size $T(S) = T_0 + \frac{S}{n}$.

Building on the connection between preferential attachment growth and nonextensive statistical mechanics [35] and following the derivations of Ref. [20] we now explain that the differential equation of the type (7) can give rise to distributions $P(S)$ that have a power law tail decorated with log-periodic oscillations if one adds an evolutionary aspect to the system. In finite difference form, the Eq. (7) can be written as

$$P(S + \delta S) = \frac{-n\delta S + nT_0 + S}{nT_0 + S} P(S) = \frac{-\delta S + T(S)}{T(S)} P(S), \quad (9)$$

where δS can be interpreted as a single-step small increment of the ETF size S . We now seek to find the solutions to the evolution equation (9) for a specific choice for the increment δS . As the changes δS can be anticipated to be proportional to the fluctuating temperature $T(S)$ one can introduce an additional scale parameter γ

$$\delta S \equiv \gamma n T(S) = \gamma n \left(T_0 + \frac{S}{n}\right) = \gamma n T_0 \left(1 + \frac{S}{nT_0}\right), \quad (10)$$

where γ can be made arbitrary small by imposing the condition $\gamma \ll \frac{1}{n}$ and recalling that $n > 1$. After inserting the expression (10) into (9) one finds that

$$P(S(1 + \gamma) + \gamma n T_0) = (1 - \gamma n)P(S). \tag{11}$$

In the asymptotic regime $S \gg T_0$, one finds

$$P(S(1 + \gamma)) \approx (1 - \gamma n)P(S) \quad (S \gg T_0), \tag{12}$$

an expression that for finite values of γ is directly recognized as the usual condition $P(\lambda S) = \mu P(S)$ for scale invariance of the function $P(S)$. It is well known [19,20] that the most general solution for the asymptotic part of the distribution is a linear combination of power laws with complex exponents $\alpha_k(\gamma, n)$

$$P(S) \approx \sum_{k \in \mathbb{N}} w_k S^{-\alpha_k(\gamma, n)} \quad (S \gg T_0), \tag{13}$$

with,

$$\alpha_{k \in \mathbb{N}}(\gamma, n) = -\frac{\ln(1 - n\gamma)}{\ln(1 + \gamma)} + \frac{2\pi i k}{\ln(1 + \gamma)}. \tag{14}$$

As is usually done, we retain only the terms in w_0 and w_1 and the real part of the function, to obtain

$$P(S) \sim S^{-n - \frac{n}{2}(n+1)\gamma + \phi(\gamma^2)} \left[w_0 + w_1 \cos\left(\frac{2\pi}{\ln(1 + \gamma)} \ln S\right) \right] \quad (S \gg T_0). \tag{15}$$

For large values of the ETF size S ($S \gg T_0$), the distribution $P(S)$ behaves as a power law decorated with a log-periodic oscillation of the type $\Delta F(S)$ defined in Eq. (2). This is compatible with the qualitative findings for the tail parts of the empirical distribution of ETF sizes (see Figs. 1 and 3). We stress that the log-periodic oscillation in the above distribution $P(S)$ is determined by the finite parameter γ that is connected with the time evolution of the system in accordance with multiplicative size increments δS that obey the relation (10). For infinitesimally small increments – that correspond with $\gamma \rightarrow 0$ – one has that

$$\lim_{\gamma \rightarrow 0} e^{-\alpha_k(\gamma, n)} = e^{-\alpha_0} \quad (\forall k), \tag{16}$$

and the asymptotic distribution of (15) reduces to the tail S^{-n} of the Tsallis distribution (8).

As a matter of fact, the proposed asymptotic solution (15) of the evolution equation, provides one with a prediction for the angular frequency of the oscillations in $\ln S$ after one time step

$$\omega_1 = \frac{2\pi}{\ln(1 + \gamma)}. \tag{17}$$

The measured distribution of ETF sizes is the result of many multiplicative evolution steps of the type (9) each with its finite characteristic scale parameter γ_t . The size of the ETF at the time instances t and $t - \Delta t$ are connected by an expression of the type (10)

$$\delta S_t = S_t - S_{t-\Delta t} = \gamma_t n \left(T_0 + \frac{S_{t-\Delta t}}{n} \right). \tag{18}$$

For the sake of simplicity, let us assume that there are κ time steps Δt and that all γ_t are equal: $\gamma_t = \gamma, \forall t$.

Proceeding in a fashion analogous to the above derivations and detailed in Ref. [20], one finds after κ time steps an asymptotic distribution $P(S)$ that is similar to the result of (15) apart from the following substitution in the angular frequency ω_κ of the $\cos(\omega \ln S)$ term

$$\omega_1 = \frac{2\pi}{\ln(1 + \gamma)} \implies \omega_\kappa = \frac{2\pi}{\kappa \ln(1 + \gamma)}. \tag{19}$$

This means that the angular frequency of the oscillations in $\ln S$ decrease with the finite number of time steps κ as $\frac{1}{\kappa}$. As a consistency check and referring to the observed oscillations in the size distributions of the ETFs in Fig. 4: for $\omega = 4.6$ one finds $\gamma = 0.014$ for $\kappa = 100$ and $\gamma = 0.0014$ for $\kappa = 1000$.

3. Analysis of the economic significance of the hierarchy of ETF sizes

We now explore whether the discrete hierarchical structure in the distribution of ETF sizes could be associated with economic properties of the ETF portfolios. Studying the return–risk properties of the hierarchy of ETF sizes is tantamount

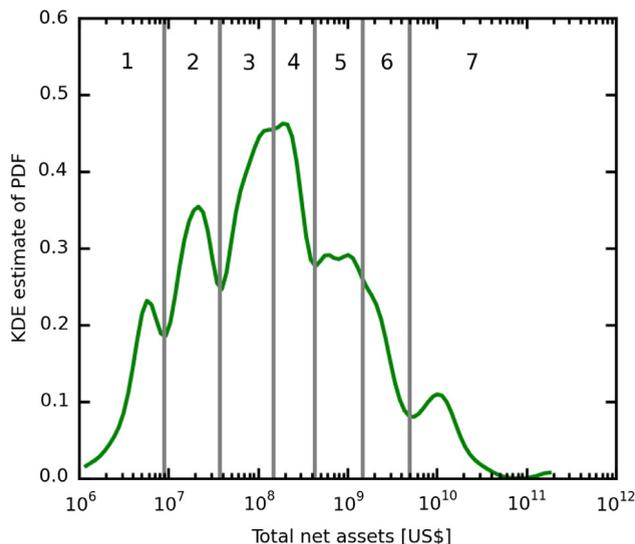


Fig. 5. Partitioning of the distribution of ETF sizes obtained with the Gaussian KDE of the PDF of the logarithms of the ETF sizes with $\sigma_o/2$, by identifying the minima and maxima that are separated by a factor close to $p_2 \approx 3.5$. The seven size layers are bracketed by the vertical lines.

to investigating the generalization of the size factor (also often referred to as SMB for “Small [market capitalization] Minus Big [market capitalization]”) of the three-factor Fama–French model [7]. Recall that the introduction of the SMB factor was motivated by the observation that small capitalization stocks have tended to do better than the market as a whole. The observation that the size effect is rather weak, especially in the last decade, has pushed Fama and French to extend their three-factor model to a five-factor model [37]. Therefore, we expect to find only weak signatures of the size hierarchy. Nevertheless, we propose that it is worthwhile to investigate a generalization of the dichotomy between small and big ETF sizes, by using the discrete hierarchy discovered above.

In order to construct the size layers partitioning the ETF universe, we use a specific geometric partitioning of the ETF universe based on the discovered scaling ratio $p_2 = \exp(2\pi/\omega_2) \approx \sqrt{p_1} \approx 3.5$, because it is present both in the analysis of the CCDF (Section 2.2) and of the PDF (Section 2.3) of ETF sizes. Moreover, it amounts to the simplest substructure to the dominant scaling ratio $p_1 \approx 12$ identified in Section 2.2. We partition the distribution of ETF sizes obtained with the Gaussian KDE of the PDF of the logarithms of the ETF sizes with $\sigma_o/2$, by identifying the minima and maxima that are separated by a factor close to $p_2 \approx 3.5$. The obtained set of seven size layers are represented in Fig. 5. Table 1 reports a number of properties for each size layer i , including the number of ETFs, the average number of holdings per ETF, the upper bound size (ub_i) and the ratio ub_i/ub_{i-1} . One can observe that the mean value of this ratio is 3.6, which is close to the scaling ratio of $p_2 = 3.5 \pm 0.2$, as expected. Note that the most probable ETF size of approximately 130 million US\$ determined in Section 2.1 falls close to the boundary between the third and fourth size layer. In contrast, the mean ETF size of 1.6 billion US\$ is close to the boundary between the fifth and sixth size layer.

3.1. Intra-layer and inter-layer similarity of stock holdings across ETF size layers

To investigate whether there is a connection between the different size layers and the portfolio composition of the ETFs, we compare the portfolio similarity of the different scales. The portfolio similarity $sim_{ee'}$ of ETFs e and e' is defined as

$$sim_{ee'} = sim_{e'e} = \frac{\sum_{i \in |h_e \cap h_{e'}|} w_{ei} w_{e'i}}{\sqrt{\sum_{k \in h_e} w_{ek}^2} \sqrt{\sum_{l \in h_{e'}} w_{e'l}^2}}, \quad (20)$$

where h_e and $h_{e'}$ are all the holdings of ETFs e and e' , and w_{ei} is the portfolio weight of holding i in ETF e .

Fig. 6 represents the matrix of average intra-layer and inter-layer similarities of ETFs, as defined by (20) across the seven size layers. Specifically, an entry $SIM(i, j)$ of this matrix is the average similarity $sim_{ee'}$ between the portfolios of all ETFs e in size layer i with all the ETFs e' in size layer j . Firstly, one sees that size layers consisting of larger ETFs are more self-similar. As the size layer number i increases, there is less diversity in the number of holdings used to construct the portfolios of the corresponding ETFs. The $SIM_i \equiv SIM(i, i)$ column of Table 1 reports these intra-size layer similarities. To sum up, two size layers of large ETFs are more similar than two size layers of smaller ETFs or than a size layer of small ETFs and a size layer of large ETFs. The first smallest size layer 1 breaks this regularity, which is kind of an oddity that can perhaps be associated with the very small sizes of these ETFs. The column SIM_M of Table 1 reports the average overlap of the ETFs in a given size

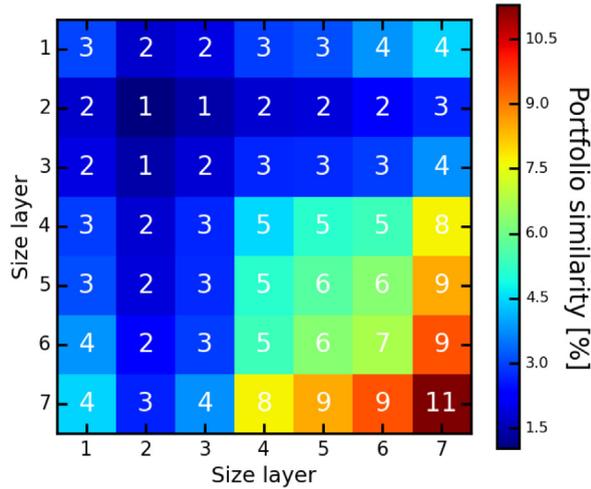


Fig. 6. Matrix of the average intra-layer and inter-layers similarities of ETFs across the seven size layers defined from Fig. 5 and detailed in Table 1. An entry $SIM(i, j)$ of this matrix is the average similarity expressed in percentages between the portfolios of all ETFs in size layer i with all the ETFs in size layer j . The similarity between two ETFs is defined by Eq. (20).

Table 1

For each identified ETF size layer $i = 1, 2, \dots, 7$, this table reports the corresponding interval of covered market capitalizations with the upper (ub_i) and lower bound (lb_i). For example, size layer 2 contains ETFs with capitalizations between 9×10^6 US\$ and 38×10^6 US\$. Further, for each size layer we provide the number of ETFs, the average number \bar{N}_h of holdings per ETF, the ratio ub_i/ub_{i-1} , the average overlap similarity over all ETF pairs in a band (SIM_i) and the average overlap with the market portfolio (SIM_M).

| Size layer $[lb_i, ub_i]$ (10^6 US\$) | #ETFs | \bar{N}_h | ub_i/ub_{i-1} | SIM_i (%) | SIM_M (%) |
|--|-------|-------------|-----------------|-------------|-------------|
| 1: $[0, 9]$ | 48 | 172 | | 3.03 | 13.04 |
| 2: $[9, 38]$ | 88 | 150 | 4.2 | 1.04 | 7.78 |
| 3: $[38, 150]$ | 109 | 183 | 3.9 | 1.82 | 8.54 |
| 4: $[150, 430]$ | 84 | 229 | 2.9 | 4.5 | 14.37 |
| 5: $[430, 1500]$ | 77 | 258 | 3.4 | 5.72 | 17.58 |
| 6: $[1500, 5000]$ | 43 | 281 | 3.3 | 6.74 | 19.44 |
| 7: $[5000, \infty]$ | 30 | 288 | | 11.29 | 26.13 |

layer with the market portfolio. To construct the market portfolio, we consider all the stocks held by ETFs. The weight of a stock in this portfolio is simply its market capitalization divided by the total market capitalization of all the stocks combined. Not surprisingly, one can observe that the larger size layers exhibit a stronger similarity to the market portfolio.

The greater intra-layer and inter-layer similarity of ETFs of large sizes is not surprising, as a large amount of capital to invest needs to find a large number of potential firms with not too large weights in order to limit market impact. As the universe of available stock investment is finite, and the set of attractive stocks is even more limited at any given time, it can be expected that the large ETFs exhibit significant overlaps in their holdings. For instance, who would not hold the largest firms such as Apple in their portfolio?

To quantify further this similarity in the holdings of the large ETFs, we define the two adjacency matrices \mathcal{M}_{bh}^{bin} and \mathcal{M}_{bh}^{frac} with dimensions $(N_b \times N_h)$. Here, $N_b = 7$ is the number of size layers and $N_h = 11,643$ is the total number of distinct holdings over the 479 equity ETFs considered. The fact that one ETF in the size layer b has a position in stock h is encoded by $\mathcal{M}_{bh}^{bin} = 1$, otherwise $\mathcal{M}_{bh}^{bin} = 0$. The second matrix is defined such that the element \mathcal{M}_{bh}^{frac} is equal to the fraction of ETFs in the size layer b that have a position in the holding h . The adjacency matrices \mathcal{M}_{bh}^{bin} and \mathcal{M}_{bh}^{frac} are shown in Fig. 7. The holdings are sorted from highest to lowest ubiquity in the 479 equity ETFs considered in our analysis. The adjacency matrices of Fig. 7 allow us to draw several conclusions. First, larger ETFs tend to use a smaller set of stocks to invest in. Second, larger ETFs tend to select increasingly from the same ubiquitous stocks. The ETFs in size layer 2, on the other hand, nicely sample from the entire space of holdings. Third, the first size layer 1 is an exception to this stylized picture and appears to be a smaller version of size layers 5 and 6.

3.2. Relationship between stock holding ubiquity and capitalization within ETF size layers

As already mentioned, an obvious explanation for why larger ETFs tend to be more similar is that they need to hold stocks with a larger market capitalization. Holding too many smaller stocks might prove too costly and not sufficient to absorb the capital in need of investment opportunities. The concentration and similarity of large ETFs may, in large part, just reflect the reduction in available large stocks to invest in. To investigate this hypothesis, we first study the relation between the

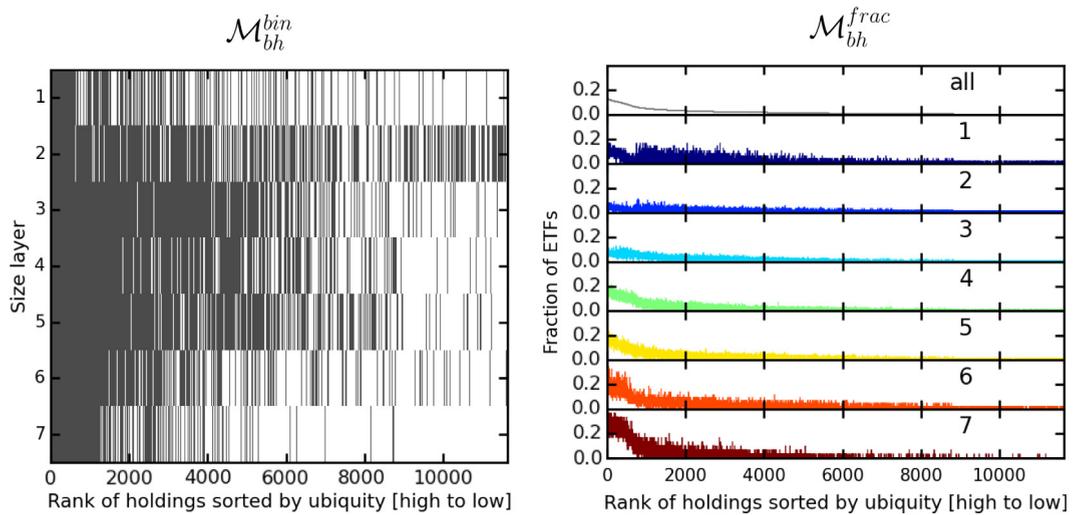


Fig. 7. The matrices \mathcal{M}_{bh}^{bin} (left) and \mathcal{M}_{bh}^{frac} (right) as defined in the text. The upper panel of \mathcal{M}_{bh}^{frac} shows the fraction of the equity ETFs in which a certain holding occurs. The 11,643 holdings are sorted according to their ubiquity with rank 1 corresponding to the most ubiquitous stock.

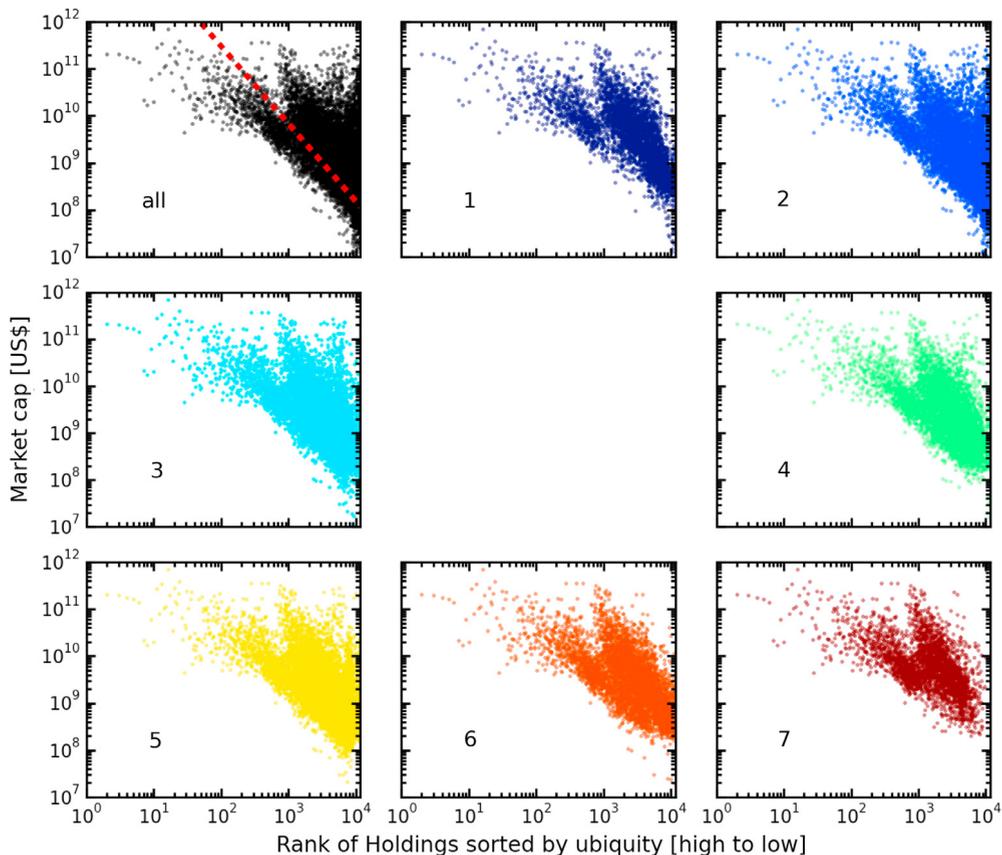


Fig. 8. For all the 479 ETFs and for each size layer separately, the market capitalization at the end of 2014 for all the holdings considered in Fig. 7. The holdings are sorted from highest to lowest ubiquity and the results are plotted on log–log axes (base 10).

ubiquity of a stock and its market capitalization. The results are shown in Fig. 8 for all holdings and per size layer respectively. The figure exemplifies that, up to the 500 largest firms (up to rank 500), the sizes of the corresponding firms are drawn from approximately the same distribution with a minimum size of about 5×10^9 US\$. In contrast, below rank ≈ 500 , one can observe a simple power law relationship relating the size of the smallest admissible firms with respect to their abundance in the universe of ETFs. The corresponding exponent of the power law is $\alpha = 1.7028 \pm 0.0002$.

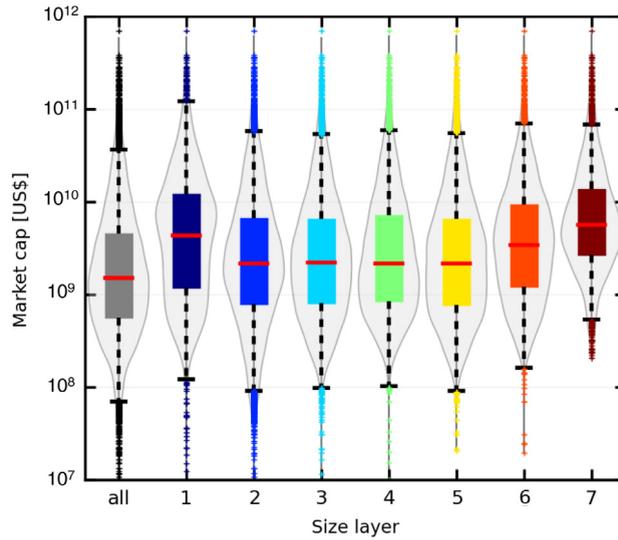


Fig. 9. For the entire sample of 479 equity ETFs and for each size layer separately, this figure shows the boxplots of the market capitalizations of all the stocks held by the ETFs. We note that the boxplot covering all the holdings is not a mere aggregation of all the size layers. In addition to presenting the boxplots, we show also the so-called violin plots that give the full distribution in thin lines along each vertical axis.

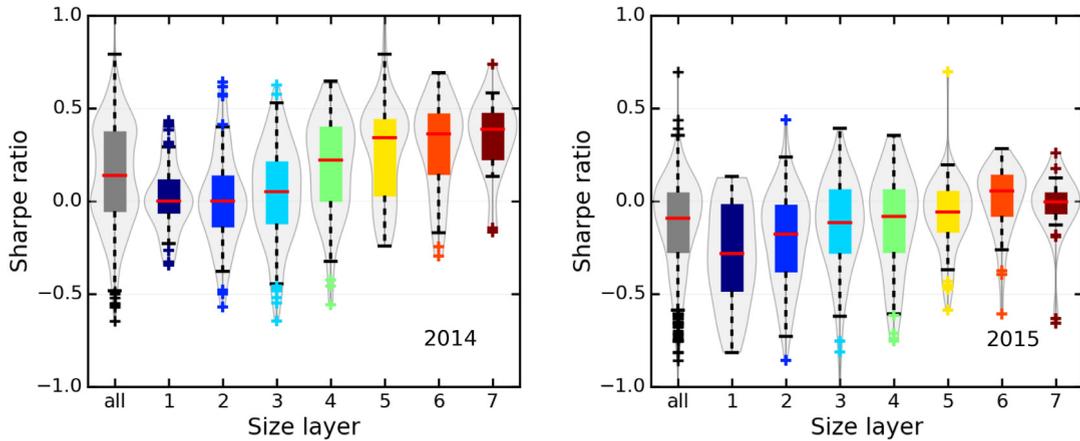


Fig. 10. Performance for the entire sample of ETFs and the different size layers in the years 2014 (left) and 2015 (right). We present the performance measure as boxplots complemented by violin plots giving the full distribution in thin lines along each vertical axis.

As expected, smaller stocks tend to be less ubiquitous in large ETFs. This effect is most obvious in the panel of size layer seven. On the contrary, we do not see that small ETFs only hold small cap stocks. Surprisingly, Fig. 8 uncovers some additional structure. There appear to be two clusters in the considered “market cap”-“holding ubiquity” matrix, and to guide the eye we have drawn a separation line in the panel including all stocks.

We have also investigated the distribution of the market capitalization of all the stocks that appear in the discerned size layers. Fig. 9 shows the market capitalizations of all the stocks that are held by the ETFs in the considered subset, be it the entire sample of ETFs or a particular size layer. The boxplot for all the holdings over all ETFs is not a mere aggregation of the boxplots of all size layers, as stocks can appear only once. For size layers 6 and 7, but remarkably also for size layer 1, the distribution is clearly shifted to stocks with a larger market capitalization. The violin plots of the full PDFs for the whole ETF market and each size layer reveal that the distributions have more structure than being uni-modal and symmetric. They exhibit significant skewness as well as bi-modality as is the case for size layers 1 and 5.

3.3. Investment performance across the seven ETF size layers

Fig. 10 presents a measure of performance of the ETFs in each size layer and in total for 2014 and 2015. Unfortunately, we do not have access to comparable data on ETFs for other calendar years, which prevents us from performing factor regressions as in Refs. [7,37]. As metric of performance, we use the Sharpe ratio, defined as the annualized mean return divided by annualized volatility (standard deviation of the returns). We take a reference risk free interest rate equal to 0. In

addition to presenting the standard boxplots, we show also the violin plots that give the full distribution in thin lines along each vertical axis.

In 2014, all median Sharpe ratios were positive, showing that the majority of ETFs generated positive returns. In addition, a significant inverse-size effect can be observed: the upper size layers 5–7 significantly over-perform the lowest size layers 1–3, with size layer 4 representing an intermediary case. In contrast, 2015 has been a difficult year for ETFs as well as for hedge-funds in general. Except for size layer 6, all other size layers have negative median Sharpe ratios. One can also observe an inverse size effect, in the sense that the upper size layers also perform better than the lower ones, but the difference is less pronounced than in 2014. The higher performance of large ETFs is reminiscent of the increasing returns to wealth inequality found, for example, for university endowments [38] and for households' portfolios [39]: larger endowments provide much larger returns as a result of better economies of scales, and of access to more investment opportunities and to more skilled managers. Similarly, larger households' wealth enable access to more diversified portfolios.

The violin plots of the full PDFs for the whole ETF market and each size layer in Fig. 10 reveal that the distributions have more structure than being uni-modal and symmetric. The discerned bi-modality of the PDF of Sharpe ratios for the whole market can be attributed to the distinct performance of two classes of ETFs, the less performing one represented mostly in size layers 2–4 and the more performing one populating the size layers 5–7.

4. Conclusion

We provided a novel detailed analysis of the size distribution of a universe of almost 500 equity ETFs and discovered a discrete hierarchy of sizes, which imprints a log-periodic structure on the probability distribution of ETF sizes that dominates the details of the asymptotic tail. We used the found discrete hierarchy to propose a classification of the whole universe of ETFs into seven size layers. Introducing a similarity metric, we found that the largest ETFs exhibit both stronger intra-layer similarity and stronger inter-layer similarity compared with the smaller ETFs. We have found strong indications that this reflects the obligation for large ETFs to spread their capitalization on a relatively more reduced set of large stocks. This lack of diversification in the classes of large ETFs seems to reinforce the concentration of stock capitalization known as Zipf's law. This concentration together with the similarity of holdings suggests potential vulnerability to systemic risks. We also provided comparative performance across the seven ETF size layers and found an inverse size effect, namely large ETFs perform significantly better than the small ones.

Acknowledgements

This work was supported by the Research Foundation Flanders (FWO-Flanders) (11R3114N). B. Vandermarliere was supported as an 'FWO-aspirant'.

References

- [1] I. Diaz-Rainey, G. Ibikunle, A taxonomy of the 'Dark Side' of financial innovation: The cases of high frequency trading and exchange traded funds, *Int. J. Entrepreneurship Innov. Manag.* 16 (1–2) (2012) 51–72.
- [2] Ayan Bhattacharya, Maureen O'Hara, Can ETFs increase market fragility? Effect of information linkages in ETF markets, 2016. Available at SSRN: <http://ssrn.com/abstract=2740699>.
- [3] Semyon Malamud, A dynamic equilibrium model of ETFs. Swiss Finance Institute Research Paper No. 15-37, 2015. Available at SSRN: <http://ssrn.com/abstract=2662433> or <http://dx.doi.org/10.2139/ssrn.2662433>.
- [4] Zhi Da, Sophie Shive, When the bellwether dances to noise: Evidence from exchange-traded funds, 2013. Available at SSRN 2158361.
- [5] https://en.wikipedia.org/wiki/Exchange-traded_fund#Types AND <http://etfdb.com/screener/>.
- [6] Robert L. Axtell, Zipf distribution of US firm sizes, *Science* 293 (2001) 1818–1820.
- [7] Eugene F. Fama, R. French Kenneth, Common risk factors in the returns on stocks and bonds, *J. Financ. Econ.* 33 (1993) 3–56.
- [8] Y. Malevergne, V. Pisarenko, D. Sornette, Testing the Pareto against the lognormal distributions with the uniformly most powerful unbiased test applied to the distribution of cities, *Phys. Rev. E* 83 (2011) 036111.
- [9] Didier Sornette, *Critical Phenomena in Natural Sciences (Chaos, Fractals, Self-Organization and Disorder: Concepts and Tools)*, Springer Series in Synergetics, Heidelberg, second ed., 2004.
- [10] Jeremy Ramsden, Gy. Kiss-Haypál, Company size distribution in different countries, *Physica A* 277 (1) (2000) 220–227.
- [11] Robert Axtell, Firm sizes: facts, formulae, fables and fantasies, in: Claudio Cioffi-Revilla (Ed.), *Power Laws in the Social Sciences*, Cambridge University Press, 2006.
- [12] Herbert A. Simon, Charles P. Bonini, The size distribution of business firms, *Amer. Econ. Rev.* 46 (1958) 607–617.
- [13] Orietta Marsili, Technology and the size distribution of firms: Evidence from Dutch manufacturing, *Rev. Ind. Organ.* 27 (2005) 303–328.
- [14] Yuji Ijiri, Herbert A. Simon, Skew Distributions and Sizes of Business Firms, in: *Studies in Mathematical and Managerial Economics (Book 24)*, North-Holland, Amsterdam, 1977.
- [15] A. Saichev, Y. Malevergne, D. Sornette, Theory of Zipf's Law and Beyond, in: *Lecture Notes in Economics and Mathematical Systems*, Springer, 2009, p. 632.
- [16] S. Lera, D. Sornette, Effects of mergers and acquisitions on firm size distributions, *J. Econ. Interact. Control* (2016).
- [17] Y. Malevergne, A. Saichev, D. Sornette, Zipf's law and maximum sustainable growth, *J. Econom. Dynam. Control* 37 (6) (2013) 1195–1212.
- [18] Jeffrey D. Scargle, *Studies in astronomical time series analysis. II-statistical aspects of spectral analysis of unevenly spaced data*, *Astrophys. J.* 263 (1982) 835–853.
- [19] D. Sornette, Discrete scale invariance and complex dimensions, *Phys. Rep.* 297 (5) (1998) 239–270. (extended version at <http://xxx.lanl.gov/abs/cond-mat/9707012>).
- [20] Grzegorz Wilk, Zbigniew Włodarczyk, Tsallis distribution decorated with log-periodic oscillation, *Entropy* 17 (1) (2015) 384–400.
- [21] W.-X. Zhou, D. Sornette, Statistical significance of periodicity and log-periodicity with heavy-tailed correlated noise, *Internat. J. Modern Phys. C* 13 (2) (2002) 137–170.

- [22] W.-X. Zhou, D. Sornette, Evidence of intermittent cascades from discrete hierarchical dissipation in turbulence, *Physica D* 165 (2002) 94–125.
- [23] W.-X. Zhou, D. Sornette, V. Pisarenko, New evidence of discrete scale invariance in the energy dissipation of three-dimensional turbulence: Correlation approach and direct spectral detection, *Internat. J. Modern Phys. C* 14 (4) (2003) 459–470.
- [24] Y. Huang, A. Johansen, M.W. Lee, H. Saleur, D. Sornette, Artifactual log-periodicity in finite-size data: Relevance for earthquake aftershocks, *J. Geophys. Res. (Solid Earth)* 105 (2000) 25451–25471.
- [25] W.-X. Zhou, D. Sornette, R.A. Hill, R.I.M. Dunbar, Discrete hierarchical organization of social group sizes, *Proc. R. Soc. Lond.* 272 (2005) 439–444.
- [26] Benedikt Fuchs, Didier Sornette, Stefan Thurner, Fractal multi-level organisation of human groups in a virtual world, *Sci. Rep.* 4 (2014).
- [27] Wei-Xing Zhou, Didier Sornette, Generalized q analysis of log-periodicity: Applications to critical ruptures, *Phys. Rev. E* 66 (4) (2002) 046111.
- [28] Wei-Xing Zhou, Didier Sornette, Nonparametric analyses of log-periodic precursors to financial crashes, *Internat. J. Modern Phys. C* 14 (08) (2003) 1107–1125.
- [29] S. Gluzman, D. Sornette, Log-periodic route to fractal functions, *Phys. Rev. E* 65 (2002) 036142.
- [30] Didier Sornette, Linear stochastic dynamics with nonlinear fractal properties, *Physica A* 250 (1) (1998) 295–314.
- [31] Per Jögi, Didier Sornette, Michael Blank, Fine structure and complex exponents in power-law distributions from random maps, *Phys. Rev. E* 57 (1) (1998) 120.
- [32] G. Ouillon, D. Sornette, A. Genter, C. Castaing, The imaginary part of rock jointing, *J. Physique I* 6 (8) (1996) 1127–1139.
- [33] C. Tsallis, Possible generalization of Boltzmann–Gibbs statistics, *J. Stat. Phys.* 52 (1–2) (1988) 479–487.
- [34] Adrian Dragulescu, Victor M. Yakovenko, Statistical mechanics of money, *Eur. Phys. J. B* 17 (2000) 723.
- [35] D.J.B. Soares, C. Tsallis, A.M. Mariz, L.R. da Silva, Preferential attachment growth model and nonextensive statistical mechanics, *Europhys. Lett.* 70 (2005) 70–76.
- [36] J. Cleymans, M.D. Azmi, Large transverse momenta and tsallis thermodynamics, *J. Phys. Conf. Ser.* 668 (1) (2016) 012050.
- [37] Eugene F. Fama, R. French Kenneth, A five-factor asset pricing model, *J. Financ. Econ.* 116 (2015) 1–22.
- [38] Thomas Piketty, *Capital in the Twenty-First Century*, Harvard University Press, 2014.
- [39] Claudio Campanale, Increasing returns to savings and wealth inequality, *Rev. Econ. Dynam.* 10 (2007) 646–675.