

# Commodity and Resource ETF Trading Patterns During the Financial Crisis

Torsten Heinrich\*

September 20, 2014

## Abstract

The current study estimates the parameters of a power law fit of the distribution of log returns of exchange traded funds (ETFs) before, during, and after the recent financial crisis. It is found, that there is considerable variation both between ETFs and between calm and turbulent phases. Exponents of the daily log return distribution are estimated to lie mostly between 3.0 and 5.0 depending on the ETF. In minute-by-minute trading data much lower power law exponents have been found concentrating between 3.0 and 4.0 and sometimes dropping to values close to or below 3.0. Further, there is evidence for changes in the distribution during times of turbulence (value of the exponent, improvement in the goodness of fit measures of the distribution). It can be hypothesized that effects such as infinite variance (for  $\alpha < 3$ ) or changes in the form of the distribution can occur, in turn affecting the predictability of the system which has implications for the possibility to control or regulate financial markets under such conditions.

**Keywords:** ETF, power law distributions, commodity prices, 2008 financial crisis, inverse cubic law of price changes

**JEL Codes:** C13, C14, G01, G12

## 1 Introduction

The current study uses daily price change data for 58 exchange traded funds (ETFs) for commodities and resources as well as higher resolution data for one of the ETFs (minute-by-minute data for the energy ETF xle) to estimate the distribution of the absolute values of the log returns  $r = |\log(\text{Price}_{t+1}/\text{Price}_t)|$ .<sup>1</sup> This distribution is, as will be detailed in the literature review below, commonly assumed to follow a power law shape  $p(x) = Cx^{-\alpha}$  with an exponent around  $\alpha = 3.0$ . Starting with this assumption, the samples are fitted to a power law distribution.

The purpose of this study is to investigate whether the distribution (that is, the estimated parameters) is stable across different ETFs and particularly over time. For the sample (which includes the time of the recent financial crisis) it is found, that there is considerable variation both between ETFs and between calm and turbulent phases.

Following the finding that log returns of financial assets may be power law distributed, the ramifications of this for the use of Gaussian statistics (as is common in economics and finance) have been vividly debated [2, 3, 4, 5, 1]. As higher moments do not exist for power law distributions (depending on the exponent), models relying on such moments (potentially including mean and variance) may lead to falwed results, particularly the underestimation of tail events (colorfully termed 'black swans' or 'dragon kings') such as financial crashes. As power law distributions are also taken as an indicator for self-organizing processes, the question comes to mind whether the processes are identical

---

\*University of Bremen; torsten.heinrich@uni-bremen.de

<sup>1</sup>Log returns are commonly used in finance since they are an approximation of raw returns  $r \approx 1 + e^r = \frac{\text{Price}_{t+1} - \text{Price}_t}{\text{Price}_t}$  but have certain desirable qualities, among others, they are additive in time and lognormally distributed raw returns would show up normally distributed in log returns form (which is a well-founded and convenient prior but turns out not to be supported by the data since log returns are still heavy tailed [1]).

(as indicated by the parameters of the distribution) for all financial assets including ETFs, if there are certain differences between assets, or over time, particularly against the background of a larger financial crisis. The present paper contributes to this discussion.

It has only recently been discussed which kind of power over the global commodity and resource markets is wielded by some very few large players [6]. For instance the global agricultural trade is controlled by 4 large trading houses (the so called 'ABCD companies'); with oil and gas as well as metal mining, the situation is very similar.

While this is not a trait that is unique to commodity trading - it is well known from the ownership structure of large global corporations [7] - the direct influence on commodity (and resource) prices arguably has a much larger potential to interfere with non-financial economic sectors and conceivably throw large parts of the socio-economic systems into turbulence.

Though the markets for resources and commodities (agriculture, metals, oil, energy, etc.) are controlled by large trading houses in recent years exchange-traded funds (ETFs) have made this investment opportunity available to a wider audience. While this does not remedy the possible danger resulting from more turbulent commodity and resource markets it helps to track the development of trading in these goods.

On the other hand the financial crisis 2007 to 2009 did not only cause the most severe economic downturn the entire developed world and beyond since the second world war, it was also one of the most destructive events in modern financial markets. The economic crisis that followed the financial crisis has probably not been resolved until today; most developed countries have since been in technical recession several times; the economy of a number of southern European countries is severely damaged and with the end of the US Federal Reserve Bank's quantitative easing policy a recent turn of the economic crisis hit the developing world. Further, it led - at least temporarily - to more careful investment decisions on the part of the financial market actors. One effect of this was an increased interest in investing in (and speculating with) resources and commodities. While parts of the developing world have always been susceptible to unforeseeable changes in resource and commodity prices caused by the financial market; even the food price crisis of 2010 may in part<sup>2</sup> have been a consequence of the recent financial crisis.

Seen within the context of such deep changes, it is difficult to believe that the statistical distribution of the price development in financial markets should have been completely unaffected. Studies [8, 9] of the price development in Asian (specifically Korean) financial markets during and after the Asian crisis of the 1990s suggest that there are indeed changes in the statistical properties of those distributions.

The current paper will proceed to investigate the ETF log return developments during the 2007 to 2009 financial crisis. Particularly, the question will be addressed if the distribution of price changes for time intervals of different scales (days, minutes) undergoes major changes. Section 2 will discuss the method and the data used in this study in the light of the literature; section 3 will present the results and discuss the consequences of the observed changes in the log returns. Finally, section 4 will draw conclusions.

## 2 How to Estimate Power Laws in Financial Market Data: Literature, Data, and Method

It has long been established that log returns of financial assets may follow power law distributions. The distinguishing property of this distribution is that it is scale free, self-similar in the sense that it is identical on different levels. Values that are  $h$  times higher are  $h^\alpha$  times less common, where  $\alpha > 0$  is called the power law exponent of the distribution, i.e.

$$h^\alpha p(hx) = p(x) \tag{1}$$

the solution to which is the canonical probability density function of the power law distribution

$$p(x) = Cx^{-\alpha} \tag{2}$$

---

<sup>2</sup>More obvious reasons were unfavorable weather and environmental conditions as well as the rise of biofuels.

where  $C$  is an positive factor (that guarantees that  $\int p(x) = 1$  as required for probability distribution functions). Originally observed by Mandelbrot [10, 11] in the cotton market in the 1950s, this has since been shown for many financial assets in many different time periods. Mandelbrot had notably found relatively low exponents of 1.7 to 2.0 [10]. During the 1990s, the exponents of price changes frequencies of different stock prices have been estimated by Jansen and de Vries [12] (who found exponents between 3.0 and 5.22) and by Lux [13] (finding exponents between 2.4 and 3.9) before Gopikrishnan et al. [14] proposed the 'inverse cubic law' stating that price changes are generally distributed according to a power law function with exponent  $\approx 3.0$ . While this has been confirmed a number of times [15, 16], some scholars have criticized the conjecture [17] and others have suggested that the distribution varies more than the 'inverse cubic law' suggests with some proposing that it may under certain conditions follow exponential distributions rather than power laws [8]. This has specifically been suggested for the case of Korean stocks during the Asian crisis [9]. Nevertheless it is widely accepted that power laws with exponents around 3.0 are a good approximation; this has also been the basis for many approaches to explain the emergence of financial power laws [18, 19, 20, 21]. No estimates exist so far regarding the assets investigated in this paper. There is a growing literature on ETFs, including contributions that are aware of the heavy tails and attempt to derive dedicated statistical methods - e.g. [22] but their focus is different from that of the current paper.<sup>3</sup> The more recent literature since the 2007-09 financial crisis has also further emphasized that apart from the power law distributions in finance there may be unanticipated outliers, termed 'black swans' or 'dragon kings' [23, 24], which might ex post also appear as a normal part of the empirical distribution but which make certain developments impossible to anticipate and evaluate ex ante [4, 5].

One of the crucial implications of the contention that the price changes (specifically log returns  $\log(p_{t+1}/p_t)$  but also other closely related measures like returns, differences, etc.) are power law distributed, is that a whole range of well-established statistical tools may not be applicable when dealing with this kind of data. While Gaussian statistics rely heavily on the moments (particularly variance) of samples, power law series are characterized by infiniteness of some moments.<sup>4</sup> The ramifications for financial market analysis and regulation policy would be extensive. This was bitterly lamented even before the financial crisis [2, 3], but the voices urging to take this point into account have become louder since. More recently this was argued more comprehensively in Taleb [5] and Mandelbrot [4]. Taleb [5] argues that the existence of the variance (i.e. if the exponent  $> 3$ ) does not make prediction easier if higher moments, particularly kurtosis, do still not exist. Of course, this is and remains contingent on the answer to the question whether the empirically observed distributions are indeed power laws or follow different distributions more favorable to Gaussian statistics (i.e. with finite moments). As seen below, power law and lognormal distributions are difficult to tell apart; in fact, the comparative statistical test based on the log-likelihood ratio which was employed in this study came out inconclusive for many of the data sets.

Related to this is the question of the explanation of financial power laws (if the distributions do actually follow a power law). As the scalefreeness property of power laws is quite unique, power laws are often taken as an indication of some kind of self-organizing process - both generally and specifically for financial assets. The general case is most illustratively argued in models of spacial organization [25] or implicit optimization of information packaging in languages widely discussed in the 1960s; for a recent overview see Mitzenmacher [26]. Gerig [1], in turn, gives the non-normal price change distribution (though he does not explicitly assume a power law) for financial assets as one of his examples for universal laws in economics. Self-organization mechanisms are also the basis for the different proposed models that explain the emergence of financial power laws. These include simple rational expectations models with trend trading<sup>5</sup>, particle dynamics models from physics, perlocation on networks (also adapted from physics), the scale free distribution of funds sizes [19], as well as herding behavior on the part of the financial market actors. A comprehensive overview is given by

<sup>3</sup>He et al. [22] for instance, are concerned with risk evaluation, not with the development of the industry as a whole and consequently not with evaluating and interpreting the shape of the log return distributions.

<sup>4</sup>All moments higher than  $\alpha + 1$  are infinite in power laws. That is, the variance becomes infinite for exponents smaller than 3.0; for exponents smaller than 2.0, the mean becomes infinite as well. Samples drawn from the distribution - including observations from real systems that follow the distribution - do, of course, always have a computable mean and variance (and statistical properties of higher-order) but with growing sample size they diverge towards infinity.

<sup>5</sup>This model only generates power law distributions if the requirement of the transitivity condition (that the present value of limit of assets at time  $t$  with  $t \rightarrow \infty$  be zero) is removed.

Lux [20].<sup>6</sup> Self-organization, in turn, implies that not only the shape and type of the distribution, but also its parameters are related to a potentially self-stabilizing mechanism.<sup>7</sup> This would entail the conclusion that either the parameters remain stable over time or the mechanism changes along with them. In the latter case, this would potentially also influence the possibility to predict, regulate, or control financial markets and is therefore a question of interest particularly against the background of financial crises.

The generally accepted statistical methods of the estimation and testing of power law distributions are those given by Clauset et al. [28]. The present paper follows these guidelines and discusses some central aspects of the method while referring to Clauset et al.'s original paper for a detailed introduction.

Daily trading data on 58 resource and commodity ETF's were analyzed<sup>8</sup> with the xle ETF being studied more closely for the time between 2006 and mid 2013.<sup>9</sup> The daily trading data is freely available from - among other sources - Yahoo Finance (<http://finance.yahoo.com>); higher resolution data, however, is generally proprietary, and has, in this case, been obtained from QuantQuote (<http://quantquote.com>).

Taking into account the discussion of the last 20 years and the scientific consensus on the existence of a power law distribution in financial price change (return) frequencies, the current paper assesses the development of the distribution in the context of the recent financial crisis. Thus starting with the hypothesis that the data series follow a power law distribution  $p(x) = Cx^{-\alpha}$  for observations  $x$  larger than some  $x_{min}$ , the power law exponents  $\alpha$  and the minimum values  $x_{min}$  were simultaneously computed using the method proposed by Clauset et al. [28]. The standard error of the exponents were bootstrapped from sets of  $n = 500$  artificial samples for each value.<sup>10</sup>

Following Clauset et al. [28] the goodness of fit is assessed in a double approach: First, a Kolmogorov-Smirnov test is used to ascertain the plausibility of the fitted power law distribution, second, the log-likelihood ratio is used to compare the goodness of the fit to two alternative distributions (lognormal and exponential distributions). The KS-test compares the Kolmogorov-Smirnov statistic<sup>11</sup> for the data set and a number of artificial samples drawn from the fitted power law distribution. The hypothesis (i.e. that the empirical sample may be drawn from fitted power law distribution) is rejected if more than 90% of the artificial samples are better representations of the distribution from which they were drawn than the empirical sample which was fitted to this distribution.<sup>12</sup> The comparative goodness of fit test (Vuong's test) compares the power law fit with a fit of the same sample (the  $x > x_{min}$  part) to an alternative distribution (here lognormal and exponential distributions), computes the log-likelihood ratio, and assesses the significance of the result using its Gaussian error (and rejects the significance if the latter is smaller than a threshold of 0.1).<sup>13</sup>

One of the major problems with the estimation of power laws is, as pointed out by many scholars

<sup>6</sup>None of the models is generally accepted, since, as Lux explains, with all of these models there are a number of problems, either related to unrealistic power law exponents or to the missing microfoundation of the model.

<sup>7</sup>Consequently, much of the literature seems to favor both stable parameters and an accordingly rigid mechanism (that would lead to, for instance Gabaix et al.'s 'inverse cubic law'). Other studies investigating unrelated power laws found for instance in the power law parameters of the size distribution of cities over the last millenium) no significant changes [27].

<sup>8</sup>They were studied far as they are available from Yahoo Finance. For some ETF's this is since the late 1990s or early 2000s while for others only the most recent years are available and many have only been established quite recently.

<sup>9</sup>It should be noted that although the fund existed since the 1990s, trading was very limited until ca. 2005, reaching a multiple of its earlier extent since 2006.

<sup>10</sup>The resulting values are very close to the theoretical standard errors without higher order correction term (which is difficult to compute exactly), given by Clauset et al. [28] as  $\sigma = \frac{\alpha-1}{\sqrt{n}} + O(\frac{1}{n})$  where  $O(\frac{1}{n})$  is the higher order correction.

<sup>11</sup>The maximum difference between the empirical and the fitted cumulative distribution.

<sup>12</sup>Consequently, the test yields a p-value  $p_{KS} \in [0, 1]$  with a rejection threshold  $p_{KS}^* = 0.1$ . The accuracy of the test was  $\varepsilon = 0.022$  in the current study. The accuracy follows the function  $\varepsilon = \frac{1}{2\sqrt{n}}$  for  $n$  synthetic power law distributions generated for comparison [28]. The current study uses  $n = 500$ , hence  $\varepsilon = 0.022$ .

<sup>13</sup>That is, the log-likelihood ratio is  $R = \log \left( \prod_i \frac{p_{powerlaw}(x_i)}{p_{alternative}(x_i)} \right)$ ; it is positive if the power law fit is better, negative if the alternative performs better. The Gaussian error is  $p_V = \frac{1}{\sqrt{2\pi\sigma^2}} |1 - \text{erf}(R)|$ ; a low Gaussian error (here  $p_V < 0.1$ ) indicates that the given log-likelihood ratio would have been unlikely to appear if actually the other distribution was a better fit than the one indicated by the log-likelihood ratio. Negative log-likelihood ratios with low Gaussian errors do consequently indicate that the power law fit should be rejected in favor of the respective alternative distribution.

who work in this field, that several other distributions have a roughly similar shape to power laws. This is especially true for exponential distributions

$$p(x) = \lambda e^{-\lambda x} \quad (3)$$

and log-normal distributions

$$p(x) = x^{-1} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\log x - \mu)^2}{2\sigma^2}}. \quad (4)$$

The equations given here are the probability density functions of the two distributions with  $\lambda$ ,  $\mu$ , and  $\sigma^2$  being positive parameters. The natural logarithms of log-normal distributed random numbers follow a Gaussian normal distribution with mean  $\mu$  and variance  $\sigma^2$ .

While the above discussed regularity for power law distributions requires them to take the form of a straight line in a log-log plot (both axes logarithmic)<sup>14</sup>, exponential distributions yield a straight line for linear x- and logarithmic p-axis which implies a downward curved shape in a plot with both axis on logarithmic scale.<sup>15</sup> This difference may, however, if the curvature is only slight, be difficult to detect.

Power law and lognormal distributions are even harder to distinguish than power law and exponential distributions. The above probability density function already suggests that the distribution (containing the factor  $x^{-1}$ ) follows a similar shape as a power law and the exponential distribution.<sup>16</sup>

Other distributions that are often proposed to explain heavy-tailed data include compound distributions which follow power laws only over a part of their domain while the tail follows e.g. an exponential distribution (power laws with exponential cutoff). These, however, are (1) even more difficult to fit and to distinguish from other distributions, (2) even more difficult to explain since there is no apparent reason for the break of the cutoff, and (3) have been criticized vigorously on the basis that samples drawn from power law distributions tend to underestimate the tail weight especially for small samples (since the 'dragon king' events are - though existent - unlikely to be part of the sample) [4, 5].

As discussed in the literature review above, the implications resulting from the distributions differ sharply: lognormal and the exponential distributions do have finite moments, while a power law distribution does not - with all resulting difficulties with respect to predictability and control.

In section 3 the current study proceeds to estimate and compare the parameters of power law fits for daily log returns of the different ETFs. It uses higher resolution data (minute-by-minute) to assess the time development of the parameters of one of the ETFs (xle) during and after the recent financial crisis. The question of the interpretation and possible explanations is taken up in the conclusion (section 4) again.

### 3 Results

The power law parameters estimated for the 58 ETFs included in the dataset as well as corresponding goodness of fit statistics are given in table 1. For 30 of the 58 ETFs, the fit is deemed plausible by the KS test and the power law hypothesis is at least not significantly rejected in favor of either lognormal or exponential distributions. (For some, including xle, which has been selected to be studied in more detail below, the power law fit is significantly better than the exponential distribution; comparisons to lognormal are almost all inconclusive.)

The estimated exponents range from ca. 3.0 ( $2.97 \pm 0.194$  and  $3.044 \pm 0.1$ ) to around 8.0 ( $7.927 \pm 1.661$ ). It should be noted, however, that the standard error is rather large for all those ETFs with

---

<sup>14</sup>Consider the transformation  $x' = e^x$  for the x-axis and  $p' = \log(p) + A$  for the p-axis where A is an arbitrary constant to shift the function up or down on the resulting logarithmic axis. This may be necessary since  $\log(p)$  maps 1 to 0 which may not be desirable for plotting statistical distributions. The above power law probability density function then takes the quasi-linear form  $p'(x') = \log C - \alpha x + A$  - obviously (and independently from A) a straight line.

<sup>15</sup>Consider again the transformations proposed in footnote 14. Logarithmic p- and linear x-axis yield a function  $p'(x) = \log \lambda - \lambda x + A$  while in a log-log plot, the downward curved function  $\log \lambda - \lambda e^x + A$  is obtained.

<sup>16</sup>The logarithmic form reveals linear ( $\log x$  and quadratic ( $\log^2 x$ ) components. The latter ensure that all moments of the distribution are finite while being hard to detect if they are small.




Figure 1: Power law tail exponents ( $\alpha$ , with standard error) and share of the distribution in the fitted tail ( $p(x \geq x_{min})$ ) for resource and commodity ETFs (plausible fits only) (agriculture related in green, precious metals in blue, other metals in cyan, uranium in violet, oil and natural gas in red, energy related in black).

high estimated exponents  $\alpha$  - those must therefore be treated with extreme caution (though they are not directly rejected by the employed goodness of fit measures).

Exponents and standard error are also shown in a plot against the share of the sample that was fitted ( $p(x > x_{min})$ ) in figure 1 (plausible fits only, those indicated in bold face in table 1). It might be concluded that ETFs of certain types are concentrated around certain values. For instance, for energy related ETFs a comparatively larger part of the sample seems to be part of the power law tail; oil and natural gas related ETF's seem to tend towards lower values in their exponents with precious metals being somewhat higher (may be interpreted as less volatile prices). Agriculture and base metal related ETF's seem to be less homogenous.

As is evident from some (large) error terms (and some small fitted shares of the sample), the sample size is rather small for most of the ETFs (since the ETFs emerged only in recent years). Consequently, daily return data would not allow good results if a separation of the sample into successive time frames was attempted in order to study the time development of the distributions. Therefore, for one of the ETFs with better and more plausible fits, the energy related ETF xle, higher resolution data (minute-by-minute trading data) was studied for a time frame from 2006 to mid 2013; month and day intervals were considered. The development of the log returns is shown in figure 2, the distribution for year intervals between 2006 and 2012 is shown in figures and

It was found that, first, the minute log returns followed distributions with much lower exponents (mostly between 3 and 4) than the day log returns ( $4.11 \pm 0.177$ ). Second, not for all of the intervals, the fitted distributions are plausible (based on the KS-test and the log-likelihood comparison to lognormal and exponential fits), also, for daily intervals, the fitted values vary widely.<sup>17</sup> The estimates are shown in figures 3 (month intervals, with standard errors) and 4 (day intervals) with plausible distributions (based on KS test and log-likelihood ratio compared to lognormal and exponential distributions) highlighted. Third, there are two phases in which the trading data undergo some turbulence (elevated absolute values of log returns), namely 2008-09 and again 2011. These times coincide with or are preceded by (in the 2008 case) a phase during which more of the intervals are deemed plausible (KS-

<sup>17</sup>Similar to the day log returns for this ETF, based on the log-likelihood ratio, exponential fits are strongly rejected in favor of the power law fit for many of the day and month intervals; however, no conclusive result emerges for the comparison with the the lognormal fit.

ETF	Resource/Commodity	$\alpha$	$x_{min}$	KS p-value	LLR vs. lognormal	LLR vs exponential
nucl	Nuclear Energy	2.902±0.117	0.018	0.002	-1.553 (0.12)	0.909 (0.363)
cane	Sugar Cane	<b>2.97±0.194</b>	<b>0.014</b>	<b>0.18</b>	<b>-0.971 (0.332)</b>	<b>0.661 (0.509)</b>
gex	Alternative Energy	<b>3.044±0.1</b>	<b>0.023</b>	<b>0.104</b>	<b>-1.636 (0.102)</b>	<b>1.868 (0.062)</b>
pkol	Coal	3.092±0.124	0.026	0.016	-1.723 (0.085)	0.584 (0.559)
kol	Coal	3.111±0.136	0.035	0.006	-2.433 (0.015)	-1.432 (0.152)
cafe	Coffee	3.154±0.164	0.016	0.0	-2.711 (0.007)	-3.366 (0.001)
cut	Timber	3.159±0.123	0.02	0.002	-2.363 (0.018)	-0.87 (0.384)
fan	Wind Energy	3.168±0.138	0.025	0.076	-1.385 (0.166)	0.875 (0.381)
crud	Crude Oil	<b>3.207±0.211</b>	<b>0.017</b>	<b>0.528</b>	<b>-1.1 (0.271)</b>	<b>0.381 (0.703)</b>
slx	Steel	3.302±0.145	0.036	0.012	-1.497 (0.134)	0.436 (0.663)
dba	Corn, Wheat, Soybeans, Sugar	3.356±0.13	0.016	0.032	-1.974 (0.048)	-0.204 (0.839)
tan	Solar Energy	3.48±0.206	0.051	0.004	0.413 (0.679)	1.263 (0.207)
jjt	Tin	<b>3.49±0.224</b>	<b>0.037</b>	<b>0.434</b>	<b>-0.966 (0.334)</b>	<b>0.304 (0.761)</b>
jjc	Copper	3.538±0.155	0.026	0.0	-2.866 (0.004)	-3.089 (0.002)
ixc	Energy	3.637±0.239	0.034	0.012	0.482 (0.63)	1.326 (0.185)
slv	Silver	3.675±0.167	0.031	0.0	0.409 (0.682)	1.366 (0.172)
sco	Crude Oil	3.726±0.217	0.058	0.014	0.419 (0.675)	1.218 (0.223)
gdx	Gold	<b>3.776±0.22</b>	<b>0.041</b>	<b>0.624</b>	<b>-0.591 (0.554)</b>	<b>1.226 (0.22)</b>
sivr	Silver	<b>3.798±0.222</b>	<b>0.028</b>	<b>0.1</b>	<b>-0.766 (0.444)</b>	<b>0.704 (0.481)</b>
ree	Rare Earths	3.831±0.306	0.121	0.026	-1.309 (0.19)	-1.111 (0.267)
corn	Corn	<b>3.861±0.313</b>	<b>0.027</b>	<b>0.284</b>	<b>-0.586 (0.558)</b>	<b>0.395 (0.693)</b>
nlr	Uranium, Nuclear Energy	<b>3.87±0.278</b>	<b>0.035</b>	<b>0.696</b>	<b>-0.481 (0.63)</b>	<b>1.144 (0.253)</b>
jjn	Nickel	3.89±0.201	0.036	0.01	-1.802 (0.072)	-1.167 (0.243)
uga	Gasoline	3.958±0.258	0.037	0.036	-1.807 (0.071)	-1.874 (0.061)
oil	Crude Oil	3.97±0.209	0.034	0.114	-1.787 (0.074)	-1.236 (0.216)
dbo	Crude Oil	4.074±0.218	0.029	0.062	-1.448 (0.148)	-0.525 (0.6)
remx	Rare Earths, 'Strategic Metals'	<b>4.088±0.365</b>	<b>0.032</b>	<b>0.208</b>	<b>-0.921 (0.357)</b>	<b>-0.412 (0.68)</b>
lit	Lithium	<b>4.094±0.354</b>	<b>0.025</b>	<b>0.488</b>	<b>-0.088 (0.93)</b>	<b>1.41 (0.158)</b>
iau	Gold	4.095±0.257	0.023	0.0	0.287 (0.774)	1.688 (0.091)
xle	Energy	<b>4.11±0.177</b>	<b>0.029</b>	<b>0.506</b>	<b>-0.011 (0.991)</b>	<b>2.671 (0.008)</b>
uso	Crude Oil	<b>4.129±0.235</b>	<b>0.035</b>	<b>0.138</b>	<b>-1.484 (0.138)</b>	<b>-0.886 (0.376)</b>
dbp	Precious Metals	<b>4.135±0.258</b>	<b>0.024</b>	<b>0.568</b>	<b>-0.37 (0.711)</b>	<b>1.344 (0.179)</b>
agq	Silver	4.144±0.332	0.075	0.006	0.358 (0.72)	1.585 (0.113)
bno	Crude Oil	4.191±0.305	0.022	0.042	-0.937 (0.349)	-0.327 (0.744)
ung	Natural Gas	4.203±0.202	0.038	0.0	0.626 (0.531)	1.647 (0.1)
nib	Cocoa	<b>4.212±0.247</b>	<b>0.029</b>	<b>0.148</b>	<b>-1.303 (0.193)</b>	<b>-0.448 (0.654)</b>
ld	Lead	4.252±0.263	0.042	0.04	-1.522 (0.128)	-1.108 (0.268)
dbe	Energy	4.293±0.271	0.03	0.042	-1.351 (0.177)	-0.85 (0.395)
ura	Uranium	<b>4.327±0.425</b>	<b>0.041</b>	<b>0.104</b>	<b>-0.22 (0.826)</b>	<b>0.706 (0.48)</b>
uhn	Diesel, Heating Oil	4.33±0.276	0.032	0.056	-1.597 (0.11)	-1.414 (0.157)
dgp	Gold	<b>4.338±0.352</b>	<b>0.053</b>	<b>0.308</b>	<b>-0.047 (0.963)</b>	<b>1.379 (0.168)</b>
jjv	Aluminium	<b>4.392±0.418</b>	<b>0.036</b>	<b>0.294</b>	<b>0.081 (0.935)</b>	<b>1.155 (0.248)</b>
sgg	Sugar	4.501±0.239	0.033	0.072	-0.778 (0.436)	0.553 (0.58)
pplt	Platinum	<b>4.507±0.393</b>	<b>0.02</b>	<b>0.308</b>	<b>-0.494 (0.621)</b>	<b>0.334 (0.739)</b>
pltm	Platinum	<b>4.51±0.453</b>	<b>0.036</b>	<b>0.244</b>	<b>-0.784 (0.433)</b>	<b>-0.258 (0.796)</b>
soyb	Soy Beans	<b>4.522±0.542</b>	<b>0.024</b>	<b>0.272</b>	<b>0.045 (0.964)</b>	<b>1.029 (0.303)</b>
gld	Gold	<b>4.624±0.328</b>	<b>0.026</b>	<b>0.492</b>	<b>-0.012 (0.99)</b>	<b>1.501 (0.133)</b>
nags	Natural Gas	<b>4.669±0.541</b>	<b>0.032</b>	<b>0.302</b>	<b>-0.019 (0.984)</b>	<b>0.968 (0.333)</b>
gaz	Natural Gas	<b>4.816±0.455</b>	<b>0.059</b>	<b>0.396</b>	<b>-0.326 (0.745)</b>	<b>0.62 (0.536)</b>
pall	Palladium	<b>5.016±0.449</b>	<b>0.034</b>	<b>0.58</b>	<b>-0.66 (0.509)</b>	<b>-0.055 (0.956)</b>
sil	Silver	<b>5.025±0.491</b>	<b>0.042</b>	<b>0.122</b>	<b>-0.296 (0.767)</b>	<b>0.501 (0.616)</b>
bal	Cotton	<b>5.107±0.431</b>	<b>0.043</b>	<b>0.296</b>	<b>-0.666 (0.505)</b>	<b>-0.095 (0.924)</b>
cow	Cattle, Hogs, Livestock	<b>5.495±0.455</b>	<b>0.018</b>	<b>0.138</b>	<b>-1.237 (0.216)</b>	<b>-1.475 (0.14)</b>
dbb	Base Metals	5.798±0.584	0.04	0.146	-1.295 (0.195)	-1.839 (0.066)
nini	Nickel	<b>5.886±2.469</b>	<b>0.086</b>	<b>0.212</b>	<b>0.057 (0.954)</b>	<b>0.613 (0.54)</b>
jo	Coffee	<b>6.881±0.915</b>	<b>0.045</b>	<b>0.424</b>	<b>-0.154 (0.878)</b>	<b>0.373 (0.71)</b>
weat	Wheat	7.599±1.86	0.04	0.076	0.061 (0.952)	0.641 (0.521)
cper	Copper	<b>7.927±1.661</b>	<b>0.031</b>	<b>0.168</b>	<b>0.169 (0.866)</b>	<b>0.629 (0.529)</b>

Table 1: Estimated parameters ( $\alpha$  and  $x_{min}$ ) for the power law fit of daily log returns in 58 ETFs as well as absolute (p-value of the Kolmogorov-Smirnov test) and comparative (log-likelihood ratio compared to lognormal and exponential distribution fit incl. Gaussian error probabilities in brackets) goodness of fit measures. Low KS p-value (here  $< 0.1$ ) indicates that the fitted power law is not plausible and should be rejected. Negative log-likelihood ratios indicate that the alternative distribution yields a better fit provided that the associated Gaussian error is low (here  $< 0.1$ ) (Vuong's test). ETF's are sorted by estimated  $\alpha$ ; Plausible fits (that are not rejected by either the KS test or the log-likelihood ratio compared to lognormal or exponential distributions) are highlighted in bold face.




Figure 2: Minute log returns, xle, 2006 to mid 2013.

test and log-likelihood ratio), and which could therefore be concluded to more closely follow a power law shape. More significantly (but unsurprisingly), in these phases, the entire distribution seems to shift to the right, the fitted parameters showing higher much  $x_{min}$ . Fourth, for month intervals the estimated exponent jumps to higher values (up to the value estimated for the day log return distribution, 4.1) during the turbulent phases. This may indicate that the bulk of the distribution shifts to higher values with the tail also being shortened (the drop to 0 occurring more sudden). It should be noted, that this would not be in line with the literature. In fact, if only a small part of the tail is fitted than the entire sample may be subject to understating the probability of high values (an effect widely described in the literature [4, 5]) resulting in the fit erroneously yielding a too steep slope (too high exponent).

To enable comparing the overall slope of the distributions (in log-log form) notwithstanding that the best fit often only encompasses a small part of the tail of the distribution, it has been attempted to fit the sample to a power law distribution forcing an  $x_{min}$  smaller than a maximum value of 0.012. In this case, the KS test results in rejection of the distribution for virtually all day intervals.<sup>18</sup> However, surprisingly, the log-likelihood ratio compared to lognormal and exponential distributions does not reject the power law distribution more often than for the above best fit (in fact, less often during the turbulent phase 2008-09). The parameters estimated for these day and month intervals also assumes a much less volatile and more consistent form. Finally, the so estimated slope falls steeply to around 2 before rising again to usual levels of around 3.5 in early 2009 (figure 5).

Though this does probably not indicate a drop of the tail exponent of the power law distribution (for which the estimated values shown in figure 4 remain the valid best fit), it indicates that a larger portion of the sample than just the power-law fitted tail may follow a distribution that remains closer in shape to a power law than either a lognormal or an exponential distribution. Taleb's and Mandelbrot's argument that samples tend to underrepresent power law tails is relevant to this.

In any case, it can be stated that the parameters estimated for the distribution do undergo fundamental and systematic changes during times of larger turbulence in the log returns (and in the price changes).

Finally, it should be noted that for most of the fits presented here, the log-likelihood ratio against a lognormal fit did not offer conclusive results, which means that the explanation and interpretation

<sup>18</sup>It should be noted that the KS test does not work very well on very small samples. The rejection of these fits by the KS test certainly indicates that it is not a perfect power law; it might, however, indicate a similar distribution that remains closer in shape to a power law than either a lognormal or an exponential distribution.






Figure 3: Power law distribution fit parameters ( $\alpha$  in red, with standard error,  $x_{min}$  in blue) as well as number of fitted observations (upper graph) in minute log return data, xle, monthly estimates. Plausible fits (not rejected by either KS test or log-likelihood ratio against lognormal and exponential fits) are marked.

of the distribution following a power law as outlined here (which would be in line with much of the literature) is only one of two possibilities. The other, that can not be ruled out based on the present findings, is that the distribution actually might also follow a lognormal distribution.

## 4 Conclusion

A number of approaches to estimate, explain, and model the emergence of the generally observed power-law structure with agent-based models (as well as the earlier non-agent-based bubble model) have been detailed in section 2 above. What is, however, not usually discussed in the debate about the modeling of financial power laws is the possibility of the change of exponents, the mechanisms leading to this, and possible other consequences resulting from such mechanisms. The current study found that the power laws the log return distributions follow do indeed change in times of turbulence. While it is apparent that the entire distribution is shifted to higher values and the tail appears steeper (fitted with higher exponents), indications exist (and have been presented) that larger parts of the distribution may actually be better fitted to a less steep power law, i.e. one with lower exponents and a higher probability for tail events ('dragon king' events).

Numerically, the power law exponents of daily log returns have been measured to fall in the approximate range between 3.0 and 5.0 for many of the investigated ETFs. For minute log returns of the ETF xle, the exponents seem to be considerably lower, falling in the range 3.0 to 4.0 (while the daily log return exponent for xle was estimated at  $4.11 \pm 0.177$ ). Certain differences between ETFs of different types have also been observed. While good share of the fits was found to be reasonably supported by goodness of fit tests (KS-test and log-likelihood ratio compared to lognormal and exponential fits), comparison to lognormal fits have mostly been inconclusive with both distributions being plausible hypotheses.

Therefore, the possibility that the investigated data might be lognormal distributed can not be ruled out, but conjecturing in line with the literature that the power law hypothesis holds, the consequences are not favorable for the prospect of predicting, regulating, or controlling trade and speculation in the goods (and respective ETFs) under investigation: As outlined by various scholars [4, 5, 1], not only do infinite moments present a problem, but also is the likelihood of tail events systematically underestimated. While Taleb's argument holds that predictability with Gaussian statistics as commonly

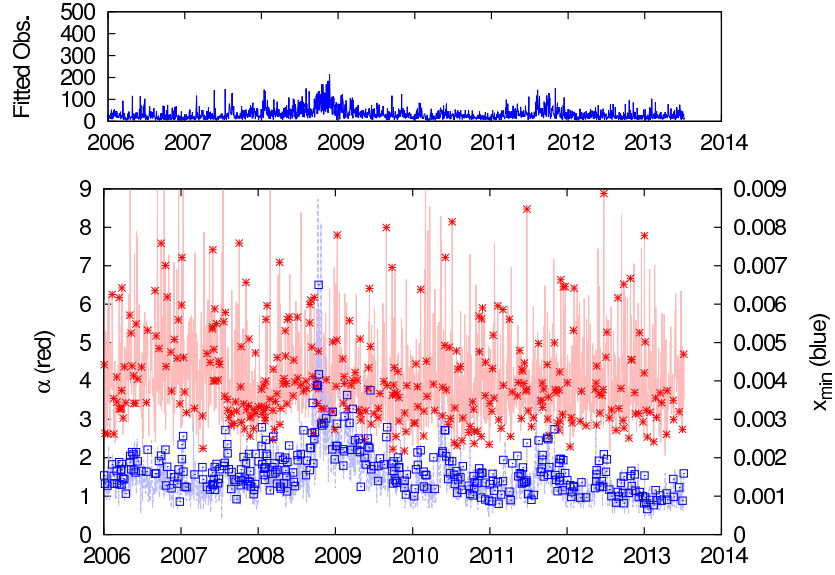


Figure 4: Power law distribution fit parameters ( $\alpha$  in red,  $x_{min}$  in blue) as well as number of fitted observations (upper graph) in minute log return data, xle, daily estimates. Plausible fits (not rejected by either KS test or log-likelihood ratio against lognormal and exponential fits) are marked.




Figure 5: Fit parameters ( $\alpha$  in red,  $x_{min}$  in blue) with a fit that forces  $x_{min} < 0.0012$ , as well as number of fitted observations (upper graph) in minute log return data, xle, daily estimates. Fits not rejected by the log-likelihood ratio against lognormal and exponential fits are marked.




Figure 6: Minute log return distribution, xle, year intervals, probability distribution (PDF) and complementary cumulative distribution (survival function, SF), 2006 - 2009




Figure 7: Minute log return distribution, xle, year intervals, probability distribution (PDF) and complementary cumulative distribution (survival function, SF), 2010 -2012

used in economics and finance is already severely impeded with infinite kurtosis (i.e. for  $\alpha < 5$ ), even lower exponents would tend to make it even more difficult. A consequence of the present research is that most ETFs seem to fall in this range, and even more so with short term log returns. (And an indication exists that this may even be aggravated in more turbulent phases.)

While this has been the case for a long time for financial assets and has arguably contributed to the repeated occurrence of completely unanticipated crashes, it may play an elevated role for resource and commodity trading in the near future. During 6 months of the second half of 2010, global food prices rose to previously unknown levels. The price of sugar rose by 76%, for sorghum, wheat, maize, soy and palm oil, as well as cotton the prices increased by more than 50% [29]. Particularly in the developing world, this led to an aggravation of poverty [29]. While the reason for the crisis were unfavorable weather conditions and expectations of poor harvests and crop failures as well as the displacement of food production by biofuels, speculation in agricultural products may have been a contributing factor. The markets are still dominated by large trading houses but the rise of ETFs led to commodities and resources being more and more also the target of other financial market actors. It is not likely that this will mitigate the problem.

## Acknowledgements

The author is indebted to two anonymous reviewers for many helpful comments and suggestions.

## References

- [1] Austin Gerig. Universal laws and economic phenomena. *Complexity*, 17(1):9–12, 2011.
- [2] Bill McKelvey and Pierpaolo Andriani. Why gaussian statistics are mostly wrong for strategic organization. *Strategic Organization*, 3(2):219–228, 2005.
- [3] Pierpaolo Andriani and Bill McKelvey. Beyond Gaussian averages: redirecting international business and management research toward extreme events and power laws. *Journal of International Business Studies*, 38:1212–1230, 2007.
- [4] Benoit B. Mandelbrot. “new methods of statistical economics,” revisited: Short versus long tails and gaussian versus power-law distributions. *Complexity*, 14(3):55–65, 2009.
- [5] Nassim Nicholas Taleb. Finiteness of variance is irrelevant in the practice of quantitative finance. *Complexity*, 14(3):66–76, 2009.
- [6] Sophia Murphy, David Burch, and Jennifer Clapp. Cereal secrets: The world’s largest grain traders and global agriculture. Technical report, Oxfam, 2012. Oxfam Research Report, <http://blogs.oxfam.ca/sites/blogs.oxfam.ca/files/rr-cereal-secrets-grain-traders-agriculture-03082012-en.pdf>, accessed 04/30/13.
- [7] Stefania Vitali and Stefano Battiston. The community structure of the global corporate network. Papers, arXiv.org, 2013.
- [8] A. Christian Silva, Richard E. Prange, and Victor M. Yakovenko. Exponential distribution of financial returns at mesoscopic time lags: a new stylized fact. *Physica A: Statistical Mechanics and its Applications*, 344(1–2):227 – 235, 2004.
- [9] Woo-Sung Jung, Seungbyung Chae, Jae-Suk Yang, and Hie-Tae Moon. Characteristics of the korean stock market correlations. *Physica A: Statistical Mechanics and its Applications*, 361(1):263 – 271, 2006.
- [10] B.B. Mandelbrot. Scaling in financial prices: I. tails and dependence. *Quantitative Finance*, 1(1):113–123, 2001.

- [11] Benoit B. Mandelbrot and Richard L. Hudson. *The (mis)behavior of markets*. Basic Books, New York, NY, 2004.
- [12] Dennis W Jansen and Casper G De Vries. On the frequency of large stock returns: Putting booms and busts into perspective. *The review of economics and statistics*, pages 18–24, 1991.
- [13] Thomas Lux. The stable paretian hypothesis and the frequency of large returns: an examination of major german stocks. *Applied financial economics*, 6(6):463–475, 1996.
- [14] Parameswaran Gopikrishnan, Martin Meyer, LA Nunes Amaral, and H Eugene Stanley. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B-Condensed Matter and Complex Systems*, 3(2):139–140, 1998.
- [15] Parameswaran Gopikrishnan, Vasiliki Plerou, Luis A Nunes Amaral, Martin Meyer, and H Eugene Stanley. Scaling of the distribution of fluctuations of financial market indices. *Physical Review E*, 60(5):5305, 1999.
- [16] Vasiliki Plerou, Parameswaran Gopikrishnan, Luis A Nunes Amaral, Martin Meyer, and H Eugene Stanley. Scaling of the distribution of price fluctuations of individual companies. *Physical Review E*, 60(6):6519, 1999.
- [17] J Doyne Farmer and Fabrizio Lillo. On the origin of power-law tails in price fluctuations. *Quantitative Finance*, 4(1):7–11, 2004.
- [18] Thomas Lux and Michele Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719):498–500, 1999.
- [19] Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423:267–270, 2003.
- [20] Thomas Lux. Financial power laws: Empirical evidence, models, and mechanism. Economics working paper / Christian-Albrechts-Universität Kiel, Department of Economics 2006,12, 2006. <http://hdl.handle.net/10419/3923>.
- [21] Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and Eugene Stanley. A unified econophysics explanation for the power-law exponents of stock market activity. *Physica A: Statistical Mechanics and its Applications*, 382(1):81–88, 2007.
- [22] Juan He, Jian Wang, Xianglin Jiang, Xiangfeng Chen, and Lei Chen. The long-term extreme price risk measure of portfolio in inventory financing: An application to dynamic impawn rate interval. *Complexity*, 2014.
- [23] Nassim N. Taleb and Charles S. Tapiero. Risk externalities and too big to fail. *Physica A: Statistical Mechanics and its Applications*, 389(17):3503 – 3507, 2010.
- [24] Didier Sornette. Dragon-kings, black swans and the prediction of crises, 2009. Swiss Finance Institute Research Paper No. 09–36. Available at arxiv: <http://arxiv.org/abs/0907.4290>.
- [25] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Physical Review A*, 38(1):364–374, July 1988.
- [26] Michael Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1(2):226–251, 2004.
- [27] Michael Batty. Visualizing space–time dynamics in scaling systems. *Complexity*, 16(2):51–63, 2010.
- [28] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51:661–703, November 2009.
- [29] Maros Ivanic, Will Martin, and Hassan Zaman. Estimating the short-run poverty impacts of the 2010–11 surge in food prices. *World Development*, 40(11):2302 – 2317, 2012.