Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Power-law correlations in finance-related Google searches, and their cross-correlations with volatility and traded volume: Evidence from the Dow Jones Industrial components

Ladislav Kristoufek*

Institute of Information Theory and Automation, Czech Academy of Sciences, Pod Vodarenskou Vezi 4, 182 08, Prague 8, Czech Republic Warwick Business School, University of Warwick, Coventry, West Midlands, CV4 7AL, United Kingdom

HIGHLIGHTS

- Online searches for the DJIA component stocks are power-law correlated.
- The uncovered persistence is very strong with Hurst exponents close to unity.
- Cross-correlations between the searches, traded volume and volatility vary strongly across stocks and scales.

ARTICLE INFO

Article history: Received 8 December 2014 Received in revised form 26 January 2015 Available online 11 February 2015

Keywords: Online searches Google Trends Long-term memory Cross-correlations Volatility Traded volume

ABSTRACT

We study power-law correlations properties of the Google search queries for Dow Jones Industrial Average (DJIA) component stocks. Examining the daily data of the searched terms with a combination of the rescaled range and rescaled variance tests together with the detrended fluctuation analysis, we show that the searches are in fact power-law correlated with Hurst exponents between 0.8 and 1.1. The general interest in the DJIA stocks is thus strongly persistent. We further reinvestigate the cross-correlation structure between the searches, traded volume and volatility of the component stocks using the detrended cross-correlation and detrending moving-average cross-correlation coefficients. Contrary to the universal power-law correlations structure of the related Google searches, the results suggest that there is no universal relationship between the online search queries and the analyzed financial measures. Even though we confirm positive correlations. In addition, the correlations vary quite strongly across scales.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Analysis of online activity of the internet users has proved its worth in various disciplines, most notably in psychology [1-5], ecology [6-9], epidemiology [10-14], medicine [15,16], linguistics [17], politology [18], sociology [19,20] and in a wide range of economics, marketing and finance [21-35]. In the economic and financial applications, the focus has been primarily put on the search queries on various search engines such as Google, Yahoo! and Baidu. Bank et al. [26] find

E-mail address: kristoufek@icloud.com.

http://dx.doi.org/10.1016/j.physa.2015.02.057 0378-4371/© 2015 Elsevier B.V. All rights reserved.







^{*} Correspondence to: Institute of Information Theory and Automation, Czech Academy of Sciences, Pod Vodarenskou Vezi 4, 182 08, Prague 8, Czech Republic.

connection between Google searches and liquidity at the German stock market. Bordino et al. [27] study traded volume of the NASDAQ-100 index component stocks and they report that it is correlated with the related searches of the Yahoo! engine. Vlastakis and Markellos [28] find positive correlation between internet search queries for NASDAQ and NYSE stocks and their traded volume and volatility. Dzielinski [29] introduces an uncertainty measure based on the financial online search queries. Preis et al. [31] show that Google searches for financial terms can be used for profitable trading strategies. Kristoufek [32] utilizes popularity of the Dow Jones stocks measured by Google search queries for portfolio diversification. Kristoufek [33] further studies dynamics between Google searches, Wikipedia page views and dynamics of the Bitcoin crypto-currency uncovering a strong relationship between these. Moat et al. [34] report that even Wikipedia page views can be utilized for the trading strategy construction. Curme et al. [35] cluster the online searches into groups and show that mainly politics and business oriented searches are connected to the stock market movements.

The most frequently reported relationship between the online searches, traded volume and volatility directs further to the dynamic characteristics of the online searches time series. As traded volume and volatility have been repeatedly studied for their power-law correlation structures [36–38], the same research line is at hand for the online searches as well. Potential long-term memory of the online activity has further implications for modeling and correct inspection of dynamics between the searches and other series. Here, we examine the correlation structure of the Google searches related to the Dow Jones Industrial Average (DJIA) index components. Daily Google searches data are utilized for the components of DJIA and as such, we present the first such study of the correlation structure of the online searches. To do so, we apply the rescaled range and rescaled variance tests to uncover the power-law correlations structure and we further proceed with the detrended fluctuation analysis of the search queries series. As it turns out that the DJIA-related Google queries are in fact power-law correlated, we reinvestigate a popular topic of cross-correlations between the searches, traded volume and volatility of the examined stocks. As we find the online searches to be power-law correlated and on the edge of (non)stationarity, we utilize the newly proposed correlation coefficients based on the detrended cross-correlation and detrending moving-average cross-correlation analyses.

The paper is organized as follows. In Section 2, we describe the used methodology, specifically the rescaled range and rescaled variance tests together with the moving block bootstrap significance criterion, and the detrended fluctuation analysis as well as the correlation coefficients. Section 3 introduces the dataset and presents the results. Section 4 concludes. We show that the Google searches related to the DJIA component stocks show scaling characteristic for the power-law correlated processes. This is supported by all utilized methods. General interest in the publicly traded companies thus shares similar properties to the variance and traded volume series — there are profound periods of high interest followed by long-lived periods of low interest. However, the search series always revert back to a long-term trend so that no explosive behavior is observed. After taking the long-term memory aspect of the online query series into consideration, the correlations between the searches, traded volume and volatility become quite unstable and no universal relationship is found. The initial long-term memory analysis thus proves to be crucial for a correct treatment of cross-correlations between the online searches and various possibly connected series.

2. Methodology

2.1. Long-term memory and its tests

Long-term memory (or alternatively long-range dependence and long-range correlations) is defined through a powerlaw decay of the auto-correlation function $\rho(k)$ which scales as $\rho(k) \propto k^{2H-2}$ for lag $k \rightarrow +\infty$ [39–41]. The series are then referred to as the power-law (auto-)correlated processes as well. The characteristic parameter of such processes is Hurst exponent *H*, or alternatively parameter α , which takes values between 0 and 1 for stationary processes. The breaking value of H = 0.5 characterizes a process with no long-term memory. Processes with H > 0.5 are usually referred to as persistent processes whereas the ones with H < 0.5 as anti-persistent processes. The former ones are reminiscent of locally trending processes which, however, keep their stationarity (for H < 1) and return to their mean value quickly enough. The latter ones are very erratic in behavior as they switch their direction more frequently than uncorrelated processes. Integrating the stationary long-range dependent processes once creates an additional category of processes which have interesting properties. For $1 \le H < 1.5$, we have non-stationary yet still mean reverting processes. The frontier of H = 1.5 marks a unit root process and H > 1.5 characterizes processes which are non-stationary and not mean reverting, i.e. explosive processes. The long-term memory property of time series has far-reaching consequences for the time series modeling and forecasting mainly due to its implication of a non-summable auto-correlation function [41]. Therefore, it is essential to distinguish between long-range dependence with its power-law correlations and short-range dependence with its exponential correlation structure. For this purpose, we utilize the modified rescaled range test and the rescaled variance test.

The modified rescaled range test [42] is an adjusted version of the original rescaled range analysis [39]. Both methods are based on scaling of the rescaled ranges with an increasing time series length. For the time series $\{x_t\}$ with t = 1, 2, ..., T, the testing statistic V_T is defined as

$$V_T = \frac{R}{S\sqrt{T}}$$

where *R* is a range of the profile of the analyzed series,

$$R = \max_{t=1,...,T} \left(\sum_{i=1}^{t} (x_i - \bar{x}) \right) - \min_{t=1,...,T} \left(\sum_{i=1}^{t} (x_i - \bar{x}) \right),$$

with \bar{x} being the time series average, S is a heteroskedasticity and autocorrelation consistent (HAC) estimator of the standard deviation of the original series, defined as

$$S^{2} = \widehat{\gamma}(0) + 2\sum_{k=1}^{q} \left(1 - \frac{k}{q+1}\right) \widehat{\gamma}(k), \tag{1}$$

with $\hat{\gamma}(k)$ being an estimated auto-correlation with lag *k* using the Bartlett-kernel weights. Note that $\hat{\gamma}(0)$ is an estimated variance. The crucial difference between the original and the modified version of the test stems in Eq. (1) which is constructed to control for a possible short-term memory bias. Selection of the parameter *q* then becomes crucial as an overshot parameter *q* can suppress even long-term memory whereas an undershot *q* parameter can direct to a misleadingly found long-term memory which in fact is only a strong short-term memory. We stick to an automatic selection criterion of the parameter as proposed by Lo [42]

$$q^* = \left\lfloor \left(\frac{3T}{2}\right)^{\frac{1}{3}} \left(\frac{2|\hat{\rho}(1)|}{1-\hat{\rho}(1)^2}\right)^{\frac{2}{3}} \right\rfloor$$
(2)

where $\hat{\rho}(1)$ is the sample first order autocorrelation and [] is the lower integer operator.

The rescaled variance test [43] is based on a very similar idea as the previous one but, as the name suggests, it is based on the profile variance rather than the profile range so that it is less sensitive to extreme values. The testing statistic M_T is then defined as

$$M_T = \frac{\operatorname{var}(X)}{TS^2}$$

where var(X) is the variance of the profile of the original series. To control for the short-term memory bias, the HAC standard deviation from Eq. (1) and the optimal q^* from Eq. (2) are used here as well.

Even though both V_T and M_T have well defined asymptotic critical values [42,43], we opt for an alternative approach utilizing the moving block bootstrap methodology [44,45] due to a finite sample, a very heterogeneous dynamics of the analyzed series as well as their distributional properties. In the procedure, surrogate series are formed by shuffling the blocks of a fixed size from the original series. This way, the short-term correlations and distributional properties are kept but the long-term correlations are shuffled away creating a distribution of the testing statistic under a more realistic null hypothesis. In our application, we fix the block size to 25 observations and we bootstrap 1000 surrogate series to obtain the testing statistic distribution under the null hypothesis.

2.2. Detrended fluctuation analysis

Detrended fluctuation analysis (DFA) [46–48] is the most popular and the most frequently applied time domain estimator of Hurst exponent. This is mainly due to the fact that DFA works under various settings such as non-stationarity and trends [48], periodic cycles and seasonalities [49], and heavy tails [50].

The procedure is based on the following steps. We work with the profile X(t) of the series $\{x_t\}$ with t = 1, ..., T defined

as

$$X(t) = \sum_{i=1}^{t} (x_i - \bar{x}).$$

The profile is divided into $T_s \equiv \lfloor T/s \rfloor$ non-overlapping windows of length *s* which is referred to as a scale. The time series length *T* may be non-divisible by *s* which creates an issue with the end of the series which would not be used in the procedure. For this purpose, the series is in addition divided into boxes from the end of the series so that we obtain $2T_s$ boxes of size *s*. In each of these boxes, we calculate a mean squared deviation from the linear time trend inside the box. This means that for the *j*th box of size *s*, we obtain

$$F^{2}(j,s) = \frac{1}{s} \sum_{i=1}^{s} (X(s[j-1]+i) - \widehat{X_{j}(i)})^{2}$$

where $X_j(i)$ is a linear fit of a time trend at position *i* in window *j*. In a similar manner, we obtain the fluctuation for the boxes formed from the end of the series as

$$F^{2}(j,s) = \frac{1}{s} \sum_{i=1}^{s} (X(T-s[j-T_{s}]+i) - \widehat{X_{j}(i)})^{2}.$$

We then construct a fluctuation for specific scale s as

$$F(s) = \left(\frac{1}{2T_s} \sum_{j=1}^{2T_s} [F^2(j, s)]\right)^{\frac{1}{2}}$$

and finally, we obtain Hurst exponent via the scaling law

$$F(s) \propto s^{H}$$
.

In the application, we estimate the exponent for scales between $s_{min} = 10$ and $s_{max} = 500 \approx T/5$. Moreover for better illustrational purposes, we base the estimation and the results on scales *s* which are powers of 10 to a single decimal point. Note, however, that the results do not change qualitatively for other specifications of scales and box splitting procedures and such approach is thus kept primarily for a straightforward presentation of the results.

2.3. DCCA and DMCA coefficients

The detrended cross-correlation coefficient $\rho_{DCCA}(s)$ for scale *s* as proposed by Zebende [51] is a combination of the detrended fluctuation analysis (DFA) [46–48] and the detrended cross-correlation analysis (DCCA) [52–54]. The coefficient is defined as

$$\rho_{DCCA}(s) = \frac{F_{DCCA}^2(s)}{F_{DFA,x}(s)F_{DFA,y}(s)},$$

where $F_{DCCA}^2(s)$ is a detrended covariance between profiles of series $\{x_t\}$ and $\{y_t\}$ based on a window of size *s*, and $F_{DFA,x}^2$ and $F_{DFA,y}^2$ are detrended variances of profiles of the separate series, respectively, for a window size *s*.¹ For time series of length *T*, the series is divided into non-overlapping boxes of length *s*. In each box, fluctuation functions are computed for linearly detrended series which are in turn averaged over all boxes of the same length. In the case when *T* is not divisible by *s*, the series is divided from the beginning as well as from the end and the averages are based on these sub-periods as in the case of DFA. More details about the methods and some alternative specifications can be found in Refs. [48,52,55–59].

The detrending moving-average cross-correlation coefficient $\rho_{DMCA}(\lambda)$ for scale λ has been introduced by Kristoufek [60] as an alternative to the DCCA coefficient. The method builds on a connection between the detrending moving average (DMA) procedure [61,62] and detrending moving-average cross-correlation analysis (DMCA) [63,64]. The coefficient is defined as

$$\rho_{DMCA}(\lambda) = \frac{F_{DMCA}^2(\lambda)}{F_{x,DMA}(\lambda)F_{y,DMA}(\lambda)},$$

where $F_{DMCA}^2(\lambda)$, $F_{DMA,x}^2(\lambda)$ and $F_{DMA,y}^2(\lambda)$ are a detrended covariance between profiles of the examined series and detrended variances of the separate series, respectively, with a moving average parameter λ . Fluctuation functions are based on series detrended by a centered moving average of length λ . Various specifications can be utilized for the detrending but the centered averaging has been shown to outperform the contenders [65]. Contrary to the DCCA coefficient, the DMCA coefficient is not based on a box-splitting and it is thus computationally more efficient. More details can be found in Refs. [60,62,63].

In a series of papers, Kristoufek [55,60] shows that the statistical properties of both methods depend strongly on longterm memory properties of the separate series. Moreover, reliability of the methods is not constant for different levels of correlation between the studied processes either. To control for such effects, we apply Theiler's Amplitude Adjusted Fourier Transform (TAAF) [66]. This method reconstructs the series with the same spectral as well as distributional properties as the original one. This way, we obtain two series with an unchanged auto-correlation and distributional structure which are, however, pairwise uncorrelated. Statistical significance of estimated correlations based on DCCA and DMCA can be then obtained and tested.

Specifically for each studied pair of processes, we obtain TAAF transformed series which are not cross-correlated but retain the auto-correlation and distribution properties of the original series. The DCCA and DMCA coefficients are then estimated for such series. As the series are not cross-correlated, the expected value of the coefficients is zero. However, variance of the estimates can be possibly high. Therefore, we estimate the coefficients on 1000 surrogate series to obtain a finite sample distribution under the null hypothesis of no cross-correlations between series which controls for both long-term auto-correlations and distributional properties.

3. Data and results

Google provides search query time series for specified terms from the year of 2004 onwards. However, the series are not reported as a pure number of searches for a given term but these are renormalized according to the Google algorithm which

(3)

¹ DCCA is a bivariate generalization of DFA presented in the previous section.

No.	Company full name	Company short name	Ticker	Search query
#1	3M Company	3M	MMM	3M
#2	Caterpillar Incorporated	Caterpillar	CAT	Caterpillar
#3	Coca-Cola Company	Coca Cola	КО	Coca Cola
#4	E. I. du Pont de Nemours and Company	Du Pont	DD	DuPont
#5	Exxon Mobil Company	Exxon Mobil	XOM	Exxon
#6	General Electric Company	General Electric	GE	GE
#7	Home Depot Incorporated	Home Depot	HD	Home Depot
#8	Intel Corporation	Intel	INTC	Intel
#9	International Business Machines	IBM	IBM	IBM
#10	J. P. Morgan Chase	J. P. Morgan	JPM	JP Morgan
#11	Johnson & Johnson	Johnson & Johnson	JNJ	Johnson Johnson
#12	McDonald's Corporation	McDonald's	MCD	McDonalds
#13	Merck & Co., Inc.	Merck	MRK	Merck
#14	Microsoft Corporation	Microsoft	MSFT	Microsoft
#15	Procter & Gamble Company	Procter & Gamble	PG	P&G
#16	The Boeing Company	Boeing	BA	Boeing
#17	United Technologies Corporation	United Technologies	UTX	UTC
#18	Walt Disney Company	Walt Disney	DIS	Disney

 Table 1

 Searched terms and DJIA component stocks.

can be in essence seen as rescaling the searches into the 0–100 interval so that the number represents the proportion of the specified searched term among all searched terms in time being kept between 0 and 100. Moreover, the obtained numbers are based on sampling from all searched terms so that these represent an estimated rescaled proportion. Even though such rescaling procedure can somewhat dilute the information content of the series, the empirical results summarized in the introductory section show otherwise.

The Google data can be downloaded freely from the Google Trends website (trends.google.com) at a weekly frequency. To obtain the data at a higher frequency, specifically the daily one, one needs to download the series in three-months sections and the series further need to be rescaled and chained together. We apply such procedure for the component stocks of the Dow Jones Industrial Average (DJIA) index between years 2004 and 2013 (apart from Exxon Mobil, J. P. Morgan and Procter & Gamble for which the series are several months shorter which will be evident later in the text) and thus obtaining 2516 observations for most series. The most severe issue with the Google queries data is its relative arbitrariness in defining the searched terms. Further, the sampling and thresholding procedure applied by Google for its search series guite frequently ends up with reporting incomplete series. If the specified term is not searched for frequently enough, the series is practically useless. We thus analyze only the component stocks which provide reliable search query series. Out of 30 DJIA stocks, we end up with 18 stocks for which the Google series are reliable without discontinuities. The analyzed stocks are summarized in Table 1. We have tried various combinations and specifications of the searched terms and we report the ones which provided the most complete series.

The Google searches for the analyzed stocks are illustrated in Figs. 1 and 2. These uncover that the searching frequencies for the component stocks are very heterogeneous. The trends are sharply decreasing (IBM, Merck, Microsoft), slowly decreasing (3M, Boeing, Du Pont, GE, Intel), or reversely increasing rapidly (McDonald's) or slowly (Caterpillar, Coca Cola, Exxon Mobil, Home Depot), or remain quite stable in time (Johnson & Johnson, J. P. Morgan, United Technologies, Walt Disney). Most of the series show strong seasonal patterns (hence the choice of the DCCA and DMCA techniques which are constructed for such series) mainly connected to the end of the year but also some stronger patterns as for Home Depot. The examined dataset thus provides a complex selection of various dynamic behaviors.

Before we get to the estimated values of Hurst exponent and thus to the type of memory in question, we first test whether the analyzed series are in fact power-law correlated. In Table 2, we present the testing statistics as well as the corresponding *p*-values for the rescaled range and rescaled variance tests as described in the previous section. Apart from two cases (Coca Cola and IBM), the power-law correlations are reported for all series (the null hypothesis is rejected by at least one of the tests at at least 10% level). It needs to be stressed that levels of the optimal *q* parameter climb high for all and very high for some cases, sometimes taking into consideration as much as 372 lags of the covariance function (here specifically for IBM). This only strengthens the claim that the analyzed Google series are long-term correlated. This is due to the fact that taking into account already tens of lags practically means considering long-term memory, even more so for hundreds of lags.

Table 2 also reports the estimated Hurst exponents which are further supported by Figs. 3 and 4. In the figures, we report an evident power-law scaling of the fluctuation functions according to Eq. (3). For all series, the scaling is very stable and the estimated Hurst exponents are thus reliable. Table 2 shows that Hurst exponents vary between 0.8 and 1.1. The Google searches are thus strongly persistent for all the analyzed series. Even though the memory is very strong for these series, Hurst exponents still remain below 1.5 which implies that the series stay mean reverting. In the DFA context, this means that even though the online queries series tend to wander away from the long-term trend, they always return to it and they never explode. The fact that the series remain on the edge of stationarity and non-stationarity (around H = 1) only highlights the need for a careful treatment of such series in multivariate settings which are standardly applied in the empirical literature.

L. Kristoufek / Physica A 428 (2015) 194-205



Fig. 1. Normalized Google searches (Part 1). Covered period ranges between 1.1.2004 and 31.12.2013 with a daily frequency.

To further illustrate the usefulness of the presented results, we reinvestigate the relationship between Google searches, traded volume and volatility. The traded volume for each component stock of the DJIA index is directly available at finance.yahoo.com as well as are the open, close, high and low prices. We utilize the provided information and construct volatility series using the Garman–Klass variance estimator [67] defined as

$$\widehat{\sigma}_{GK,t}^2 = \frac{(\log(H_t/L_t))^2}{2} - (2\log 2 - 1)(\log(C_t/O_t))^2 \tag{4}$$

where H_t and L_t are daily highs and lows, respectively, and C_t and O_t are daily closing and opening prices, respectively. The estimator possesses very good statistical properties and serves as an excellent choice without a need of using high-frequency data [68]. We study a logarithmic transformation of both $\hat{\sigma}^2_{GK}$ and the traded volume series which is a standard procedure in the applied literature. The transformation of the original variance series allows us to comment on both variance and volatility as the logarithmic variance becomes just twice the logarithmic volatility.

We examine the correlations between Google searches, traded volume and volatility at various scales using the DCCA and DMCA coefficients. For the DCCA coefficient, we study the correlation between the searches and traded volume, and between the searches and volatility for scales between 10 and 250 with a step of 10. For the DMCA coefficient, we use moving window lengths between 11 and 251 with a step of 10 as well. This way, we obtain comparable results using these methodologies.



Fig. 2. Normalized Google searches (Part 2).

Figs. 5 and 6 depict the results for variance and traded volume, respectively, for both methods. Only significant correlations with the *p*-value below 0.10 are reported, the insignificant ones are set to zero. We find several interesting results. First, the DMCA method reports more stable results with more significant coefficients. This is well in hand with the numerical results presented by Kristoufek [55,60]. Second, the correlations for traded volume are in general higher than the ones for volatility. Third, a majority of significant correlations occur at the lower scales. There thus seems to be rather short-term or medium-term relationship between the online searches and the examined financial indicators. In the long-term, only few correlations are identified as significant. Fourth, the level of correlations varies considerably across stock titles. There thus seems to be no universality in the relationship between the searches, and volatility and volume. Tables 3 and 4 further illustrate the heterogeneity of the results. There, we present the average DCCA and DMCA coefficients across scales together with their significance level. The above mentioned results are supported. First, the significance, level and sign of the correlations vary widely. Second, the DMCA procedure delivers more significant results. Third, the correlations are higher for volume than for volatility. Nevertheless, many of the significant correlations are still below a level of 0.05 and practically all the correlation coefficients fall between -0.2 and 0.2. The correlations are thus very weak even if found statistically significant.

There are still some interesting results mainly connected to various signs of the correlations. For example Microsoft shows some unorthodox behavior for volatility. A positive relationship is usually reported, whereas the search queries for Microsoft are negatively correlated with volatility. Conversely, traded volume shows a positive correlation. It thus seems

Table 2

Long-term memory tests and estimated Hurst exponent.

Company	# of obs.	V_T	p-value	M _T	p-value	q^*	H _{DFA}
3M	2516	2.9267	0.0000	0.0003	0.0000	46	0.9452
Boeing	2516	2.7083	0.0000	0.0003	0.0000	65	0.9086
Caterpillar	2516	2.0131	0.0120	0.0001	0.0030	33	1.0467
Coca Cola	2516	1.6006	0.1848	0.0001	0.2647	23	0.9016
Du Pont	2516	2.6857	0.0000	0.0002	0.0000	63	0.9276
Exxon Mobil	2265	2.8268	0.0000	0.0003	0.0000	37	0.9313
General Electric	2516	2.8344	0.0000	0.0003	0.0000	63	0.9652
Home Depot	2516	1.6500	0.0599	0.0001	0.1359	67	1.1557
IBM	2516	1.2279	0.6424	0.0001	0.1099	372	1.1478
Intel	2516	1.7875	0.0070	0.0001	0.0000	164	1.1486
Johnson & Johnson	2516	2.5457	0.0000	0.0002	0.0000	34	0.9964
J. P. Morgan	2393	2.4226	0.0000	0.0002	0.0000	26	0.8261
McDonald's	2516	2.0612	0.0000	0.0002	0.0000	127	1.0373
Merck	2516	2.3400	0.0000	0.0002	0.0000	88	0.9314
Microsoft	2516	1.5279	0.0729	0.0001	0.0060	206	1.0896
Procter & Gamble	2393	3.2807	0.0000	0.0005	0.0000	25	0.9468
United Technologies	2516	2.7528	0.0000	0.0002	0.0070	22	0.8747
Walt Disney	2516	3.0381	0.0000	0.0004	0.0000	38	1.0003

Table 3

Average DCCA and DMCA correlations between volatility and Google searches.

	$\overline{\rho_{DCCA}}$	$\widehat{\sigma}_{ ho_{DCCA}}$	p-value	$\overline{\rho_{DMCA}}$	$\widehat{\sigma}_{ ho_{DMCA}}$	p-value	significant	sign
3M	0.0657	0.0270	0.0149	0.0648	0.0112	0.0000	11	+
Boeing	-0.0372	0.0275	0.1771	-0.0281	0.0162	0.0831	×✓	_
Caterpillar	0.1087	0.0231	0.0000	0.1136	0.0086	0.0000	11	+
Coca Cola	0.0315	0.0214	0.1422	0.0369	0.0093	0.0001	×✓	+
Du Pont	0.1658	0.0194	0.0000	0.1590	0.0096	0.0000	<i>」」</i>	+
Exxon Mobil	0.0432	0.0170	0.0111	0.0411	0.0101	0.0000	<i>」」</i>	+
General Electric	0.1827	0.0173	0.0000	0.2181	0.0067	0.0000	<i>」」</i>	+
Home Depot	-0.0095	0.0273	0.7283	-0.0193	0.0080	0.0156	×✓	_
IBM	0.1093	0.0247	0.0000	0.1154	0.0047	0.0000	11	+
Intel	0.0287	0.0252	0.2536	0.0302	0.0125	0.0160	×✓	+
Johnson & Johnson	-0.0953	0.0308	0.0020	-0.1431	0.0227	0.0000	<i>」」</i>	-
J. P. Morgran	0.0282	0.0176	0.1093	0.0501	0.0094	0.0000	×✓	+
McDonald's	0.1509	0.0176	0.0000	0.1595	0.0073	0.0000	11	+
Merck	0.0216	0.0328	0.5090	0.0296	0.0228	0.1955	××	0
Microsoft	-0.1635	0.0323	0.0000	-0.1608	0.0250	0.0000	11	_
Procter & Gamble	-0.0118	0.0172	0.4944	-0.0141	0.0085	0.0960	×✓	_
United Technologies	-0.0813	0.0259	0.0017	-0.1076	0.0191	0.0000	11	_
Walt Disney	0.0841	0.0195	0.0000	0.0509	0.0052	0.0000	11	+

Table 4

Average DCCA and DMCA correlations between traded volume and Google searches.

-			-					
	$\overline{\rho_{DCCA}}$	$\widehat{\sigma}_{ ho_{DCCA}}$	p-value	$\overline{\rho_{DMCA}}$	$\widehat{\sigma}_{ ho_{DMCA}}$	p-value	significant	sign
3M	0.1599	0.0302	0.0000	0.1667	0.0179	0.0000	11	+
Boeing	-0.0280	0.0251	0.2647	-0.0468	0.0186	0.0117	×✓	_
Caterpillar	0.0785	0.0266	0.0032	0.0848	0.0144	0.0000	11	+
Coca Cola	0.0331	0.0230	0.1508	0.0436	0.0135	0.0013	×✓	+
Du Pont	0.1181	0.0184	0.0000	0.1051	0.0148	0.0000	11	+
Exxon Mobil	0.0936	0.0183	0.0000	0.0889	0.0073	0.0000	11	+
General Electric	0.1093	0.0218	0.0000	0.1994	0.0056	0.0000	11	+
Home Depot	0.0262	0.0341	0.4423	0.0382	0.0274	0.1638	××	0
IBM	0.1547	0.0151	0.0000	0.1407	0.0093	0.0000	11	+
Intel	0.0357	0.0254	0.1609	0.0345	0.0096	0.0003	×✓	+
Johnson & Johnson	0.0211	0.0254	0.4066	-0.0172	0.0173	0.3192	××	0
J. P. Morgran	0.1258	0.0195	0.0000	0.1008	0.0077	0.0000	11	+
McDonald's	0.2032	0.0185	0.0000	0.2168	0.0133	0.0000	11	+
Merck	0.0463	0.0447	0.3009	0.0328	0.0318	0.3022	××	0
Microsoft	0.0993	0.0248	0.0001	0.0937	0.0087	0.0000	11	+
Procter & Gamble	0.0724	0.0212	0.0007	0.0865	0.0199	0.0000	11	+
United Technologies	-0.0856	0.0238	0.0003	-0.0809	0.0156	0.0000	11	_
Walt Disney	0.2297	0.0437	0.0000	0.1317	0.0209	0.0000	<i>J J</i>	+

that general interest in Microsoft is mainly tied with positive news which stabilize the stock price rather than with negative news that would make the price more volatile. Similar dynamics is found for Johnson & Johnson. The only stock which gives



Fig. 3. DFA scaling of Google searches related to the DJIA component stocks (Part 1). Log-log representation shows a profound linear scaling characteristic for long-range correlated processes. Estimated Hurst exponents are summarized in Table 2.

insignificant results for both financial quantities is Merck. Other stocks show either positive and thus expected correlations or only weak negatives ones.

4. Discussion and conclusions

We have analyzed the power-law correlations in the online search queries for the DJIA stock components. By reconstructing the daily Google search series, we have been able to obtain enough observations for a valid analysis of long-range dependence. Using the combination of the rescaled range and rescaled variance tests and the detrended fluctuation analysis, we have shown that the online searches are indeed power-law correlated. Importantly, the level of long-term memory is very high with Hurst exponents around unity (between 0.8 and 1.1) for all the analyzed stocks. Such results suggest that the finance-related online searches have similar dynamic properties to stock variance and traded volume which are themselves power-law correlated. The information flow coming into the stock markets evidently enters the general public interest and keeps it for quite long periods and its dissipation is thus not immediate. The fact that the online searches and implied attention are usually attributed to retail and small investors, such information and attention dissipation fits into the picture of a small investor using the information for decision-making in a longer term. Such persistent dynamics of the series might also arise from indecisiveness of the small investors which would think twice before investing into a specific stock. Online searches then cluster and keep their level for longer time intervals. The results remain fascinatingly universal across the





'n

Scale

DCCA correlation between Google

searches and variance

0.10 0.00

-0.10

-0.20 -0 30

j

HHH 160

210

DMCA correlation between Google

searches and variance

Scale



Fig. 6. Correlation coefficients between Google searches and traded volume. The notation holds from Fig. 5.

analyzed stocks. Even though the global dynamics of the series is very heterogeneous with different speeds of trends or various volatility levels, they all remain strongly persistent with a smooth scaling of fluctuations.

In addition, we have studied the relationship between Google searches, traded volume and volatility using the recently proposed DCCA and DMCA coefficients. The results primarily suggest that there is no universal relationship between the online search queries and the analyzed financial measures. Even though we confirm positive correlation for a majority of pairs, there are several pairs with insignificant or even negative correlations. Further, the correlations vary quite strongly across scales. The online searches have thus retained their potential for financial modeling and various applications but our findings suggest that one needs to carefully study each stock or asset separately as the usefulness of the queries can fluctuate considerably. The reported results do not necessarily contradict some previous findings which find statistically significant connections [26–28,31] or time varying correlations [35]. However, we stress that there seems to be no universal and global relationship between the online searches and relevant financial variables (traded volume and volatility).

Our results open an interesting area of further research of the topic. First, the power-law properties of the correlation structure might be observed also in different types of search queries, not necessarily only for stocks or financial markets in general. This would show how information or information seeking dissipates in time and how such behavior connects to other specific phenomena of the relevant time series. Second, the results indicate that the online searches are strongly persistent and on the edge of (non-)stationarity. Such characteristic implies that the simple correlation studies reported in the literature should take this property into consideration as an inappropriate analysis of persistent data using tools for short-range dependent series can produce spurious and in turn misleading results. Third, knowing the basic dynamic properties of the series helps to construct the forecasting models which are of high interest for practitioners, specifically in risk management.

Acknowledgments

Support from the Czech Science Foundation under project No. 14-11402P and the Research Councils UK via Grant EP/K039830/1 is gratefully acknowledged. I would also like to thank Tomas Vakrman for help with a dataset collection. Google data are registered trademarks of Google Inc., used with permission.

References

- [1] M.J. McCarthy, Internet monitoring of suicide risk in the population, J. Affect. Disord. 122 (2010) 277–279.
- [2] A. Page, S.-S. Chang, D. Gunnell, Surveillance of Australian suicidal behaviour using the Internet? Aust. N. Z. J. Psychiatry 45 (2011) 1020–1022.
- [3] H. Sueki, Does the volume of Internet searches using suicide-related search terms influence the suicide death rate: data from 2004 to 2009 in Japan, Psychiatry Clin. Neurosci. 65 (2011) 392–394.
- [4] A.C. Yang, S.-J. Tsa, N.E. Huang, C.-K. Peng, Association of Internet search trends with suicide death in Taipei City, Taiwan, 2004–2009, J. Affect. Disord. 132 (2011) 179–184.
- [5] A. Hagihara, S. Miyazaki, T. Abe, Internet suicide searches and the incidence of suicide in young people in Japan, Eur. Arch. Psychiatry Clin. Neurosci. 262 (2012) 39–46.
- [6] M.L. McCallum, G.W. Bury, Google search patterns suggest declining interest in the environment, Biodiversity Cons. 22 (2013) 1355–1367.
- [7] M.L. McCallum, G.W. Bury, Public interest in the environment is falling: a response to Ficetola (2013), Biodiversity Cons. 23 (2014) 1057–1062.
 [8] G.F. Ficetola, Is interest toward the environment really declining? the complexity of analysing trends using internet search data, Biodiversity Cons. 22
- (2013) 2983–2988.
 [9] D. Verissimo, D.C. MacMillan, R.J. Smith, J. Crees, Z.G. Davies, Has climate change taken prominence over biodiversity conservation? Bioscience 64 (2014) 625–629.
- [10] Philip M. Polgreen, Yiling Chen, David M. Pennock, Forrest D. Nelson, Robert A. Weinstein, Using internet searches for influenza surveillance, Clin. Infect. Dis. 47 (11) (2008) 1443–1448.
- [11] Jeremy Ginsberg, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, Larry Brilliant, Detecting influenza epidemics using search engine query data, Nature 457 (7232) (2009) 1012–1014.
- [12] H.A. Carneiro, E. Mylonakis, Google trends: a web-based tool for real-time surveillance of disease outbreaks, Clin. Infect. Dis. 49 (2009) 1557–1564.
- [13] A. Seifter, A. Schwarzwalder, K. Geis, J. Aucott, The utility of "Google Trends" for epidemiological research: Lyme disease as an example, Geospatial Health 4 (2010) 135–137.

- [14] A.F. Dugas, Y.-H. Hsieh, S.R. Levin, J.M. Pines, D.P. Mareiniss, A. Mohareb, C.A. Gaudos, T.M. Perl, R.E. Rothman, Google Flu trends: correlation with emergency department influenza rates and crowding metrics, Clin. Infect. Dis. 54 (2012) 463–469.
- [15] F. Linkov, D.H. Bovbjerg, K.E. Freese, R. Ramanathan, G.M. Eid, W. Gourash, Bariatric surgery interest around the world: what Google Trends can teach us, Surg. Obes. Relat. Dis. 10 (2014) 533-539.
- [16] D.A. Telem, A.D. Pryor, Google Trends: is it a real tool to predict the future of bariatric surgery or merely a marketing landmine? Surg. Obes. Relat. Dis. 10 (2014) 538–539.
- [17] D. Mocanu, A. Baronchelli, N. Perra, B. Goncalves, Q. Zhang, A. Vespignani, The Twitter of Babel: mapping world languages through microblogging platforms, PLoS One 8 (2013) e61981.
- [18] P.T. Metaxas, E. Mustafaraj, Social media and the elections, Science 338 (2012) 472-473.
- [19] D. Rybski, S.V. Buldyrev, S. Havlin, F. Liljeros, H.A. Makse, Scaling laws of human interaction activity, Proc. Natl. Acad. Sci. USA 106 (2009) 12640–12645.
- [20] P.A. Grabowicz, J.J. Ramasco, B. Goncalves, V.M. Eguiluz, Entangling mobility and interactions in social media, PLoS One 9 (2014) e92196.
- [21] T. Preis, D. Reith, H.E. Stanley, Complex dynamics of our economic life on different scales: insights from search engine query data, Philos. Trans. R. Soc. A 368 (2010) 5707–5719.
- [22] Jordi Mondria, Thomas Wu, Yi Zhang, The determinants of international investment and attention allocation: using internet search query data, J. Int. Econ. 82 (1) (2010) 85–95.
- [23] S. Goel, J.M. Hofman, S. Lehaie, D.M. Pennock, D.J. Watts, Predicting consumer behavior with Web search, Proc. Natl. Acad. Sci. USA 7 (41) (2010) 17486–17490.
- [24] S. Vosen, T. Schmidt, Forecasting private consumption: survey-based indicators vs. google trends, J. Forecast. 30 (6) (2011) 565–578.
- [25] M.S. Drake, D.T. Roulstone, J.R. Thornock, Investor information demand: evidence from google searches around earnings announcements, J. Accoun. Res. 50 (4) (2012) 1001–1040.
- [26] M. Bank, M. Larch, G. Peter, Google search volume and its influence on liquidity and returns of German stocks, Financ. Mark, Portf. Manag. 25 (2011) 239–264.
- [27] I. Bordino, S. Battiston, G. Caldarelli, M. Cristelli, A. Ukkonen, W. Ingmar, Web search queries can predict stock market volumes, PLoS One 7 (2012) e40014.
- [28] N. Vlastakis, R.N. Markellos, Information demand and stock market volatility, J. Bank. Financ. 36 (2012) 1808–1821.
- [29] M. Dzielinski, Measuring economic uncertainty and its impact on the stock market, Finance Res. Lett. 9 (2012) 167–175.
- [30] T. Preis, H.S. Moat, H.E. Stanley, S.R. Bishop, Quantifying the advantage of looking forward, Sci. Rep. 2 (2012) 350.
- [31] T. Preis, H.S. Moat, H. Eugene Stanley, Quantifying trading behavior in financial markets using Google Trends, Sci. Rep. 3 (2013) 1684.
- [32] L. Kristoufek, Can Google Trends search queries contribute to risk diversification? Sci. Rep. 3 (2013) 2713.
- [33] L. Kristoufek, Bitcoin meets Google Trends and Wikipedia: quantifying the relationship between phenomena of the Internet era, Sci. Rep. 3 (2013) 3415.
- [34] H.S. Moat, C. Curme, A. Avakian, D.Y. Kenett, H.E. Stanley, T. Preis, Quantifying wikipedia usage patterns before stock market moves, Sci. Rep. 3 (2013) 1801.
- [35] C. Curme, T. Preis, H.E. Stanley, H.S. Moat, Quantifying the semantics of search behavior before stock market moves, Proc. Natl. Acad. Sci. USA 111 (2014) 11600–11605.
- [36] R. Cont, Empirical properties of asset returns: stylized facts and statistical issues, Quant. Finance 1 (2) (2001) 223-236.
- [37] S.-H. Poon, C. Granger, Forecasting volatility in financial markets: a review, J. Econ. Lit. 41 (2003) 478–539.
- [38] S.-H. Poon, C. Granger, Practical issues in forecasting volatility, Financ. Anal. J. 61 (1) (2005) 45-56.
- [39] H.E. Hurst, Long term storage capacity of reservoirs, Trans. Am. Soc. Eng. 116 (1951) 770–799.
- [40] J. Beran, Statistics for Long-Memory Processes, in: Monographs on Statistics and Applied Probability, vol. 61, Chapman and Hall, New York, 1994.
- [41] Gennady Samorodnitsky, Long range dependence, Found. Trends Stoch. Syst. 1 (3) (2006) 163–257.
- [42] A. Lo, Long-term memory in stock market prices, Econometrica 59 (5) (1991) 1279–1313.
- [43] L. Giraitis, P. Kokoszka, R. Leipus, G. Teyssiere, Rescaled variance and related tests for long memory in volatility and levels, J. Econometrics 112 (2003) 265–294.
- [44] B. Efron, R. Tibshirani, R.J. Tibshirani, An Introduction to the Bootstrap, Chapman & Hall, 1993.
- [45] V. Srinivas, K. Srinivasan, Post-blackening for modeling dependent annual streamflows, J. Hydrol. 230 (2000) 86–126.
- [46] C. Peng, S. Buldyrev, A. Goldberger, S. Havlin, M. Simons, H. Stanley, Finite-size effects on long-range correlations: implications for analyzing DNA sequences, Phys. Rev. E 47 (5) (1993) 3730–3733.
- [47] C. Peng, S. Buldyrev, S. Havlin, M. Simons, H. Stanley, A. Goldberger, Mosaic organization of DNA nucleotides, Phys. Rev. E 49 (2) (1994) 1685–1689.
- [48] J. Kantelhardt, S. Zschiegner, E. Koscielny-Bunde, A. Bunde, S. Havlin, E. Stanley, Multifractal detrended fluctuation analysis of nonstationary time series, Physica A 316 (1-4) (2002) 87-114.
- [49] K. Hu, P. Ivanov, Z. Chen, P. Carpena, H.E. Stanley, Effect of trends on detrended fluctuation analysis, Phys. Rev. E 64 (2001) 011114.
- [50] J. Barunik, L. Kristoufek, On Hurst exponent estimation under heavy-tailed distributions, Physica A 389 (18) (2010) 3844–3855.
- [51] G.F. Zebende, DCCA cross-correlation coefficient: quantifying level of cross-correlation, Physica A 390 (2011) 614–618.
- [52] B. Podobnik, H.E. Stanley, Detrended cross-correlation analysis: a new method for analyzing two nonstationary time series, Phys. Rev. Lett. 100 (2008) 084102.
- [53] W.-X. Zhou, Multifractal detrended cross-correlation analysis for two nonstationary signals, Phys. Rev. E 77 (6) (2008) 066211.
- [54] Z.-Q. Jiang, W.-X. Zhou, Multifractal detrending moving average cross-correlation analysis, Phys. Rev. E 84 (1) (2011) 016106.
- [55] L. Kristoufek, Measuring correlations between non-stationary series with DCCA coefficient, Physica A 402 (2014) 291–298.
- [56] L.-Y. He, S.-P. Chen, Nonlinear bivariate dependency of price-volume relationships in agricultural commodity futures markets: a perspective from multifractal detrended cross-correlation analysis, Physica A 390 (2011) 297–308.
- [57] L.-Y. He, S.-P. Chen, Multifractal detrended cross-correlation analysis of agricultural futures markets, Chaos Solitons Fractals 44 (2011) 355–361.
- [58] J. Cao, G. ad Cao, L. Xu, L.-Y. He, Detrended cross-correlation analysis approach for assessing asymmetric multifractal detrended cross-correlations and their application to the chinese financial market, Physica A 393 (2014) 460–469.
- [59] G.-J. Wang, C. Xie, L.-Y. He, S. Chen, Detrended minimum-variance hedge ratio: a new method for hedge ratio at different time scales, Physica A 405 (2014) 70–79.
- [60] L. Kristoufek, Detrending moving-average cross-correlation coefficient: measuring cross-correlations between non-stationary series, Physica A 406 (2014) 169–175.
- [61] N. Vandewalle, M. Ausloos, Crossing of two mobile averages: A method for measuring the roughness exponent, Phys. Rev. E 58 (1998) 6832–6834.
- [62] E. Alessio, A. Carbone, G. Castelli, V. Frappietro, Second-order moving average and scaling of stochastic time series, Eur. Phys. J. B 27 (2002) 197–200.
- [63] S. Arianos, A. Carbone, Cross-correlation of long-range correlated series, J. Stat. Mech. Theory Exp. 3 (2009) P03037.
- [64] L.-Y. He, S.-P. Chen, A new approach to quantify power-law cross-correlation and its application to commodity markets, Physica A 390 (2011) 3806–3814.
- [65] A. Carbone, G. Castelli, Scaling properties of long-range correlated noisy signals: application to financial markets, Proc. SPIE 5114 (2003) 406–414.
- [66] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, J.D. Farmer, Testing for nonlinearity in time series: the method of surrogate data, Physica D 58 (1992) 77–94.
- [67] M.B. Garman, M.J. Klass, On the estimation of security price volatilities from historical data, J. Business 53 (1980) 67–78.
- [68] R.Y. Chou, H. Chou, N. Liu, Handbook of Quantitative Finance and Risk Management, Springer US, 2010, pp. 1273–1281. Chapter Range volatility models and their applications in finance.