Algorithmic complexity theory and the relative efficiency of financial markets
— **Source link**

Ricardo Giglio, Raul Matsushita, Annibal Figueiredo, Iram Gleria ...+1 more authors

Institutions: Universidade Federal de Santa Catarina, University of Brasília, Federal University of Alagoas

Topics: Financial market, Order (exchange), Stock exchange and Insider trading

Related papers:

- On the Complexity of Finite Sequences

- Three approaches to the quantitative definition of information

- Introduction to Econophysics: Correlations and Complexity in Finance

- The Relative Efficiency of Stockmarkets

- Relationship between efficiency and predictability in stock price change

MPRA

Munich Personal RePEc Archive

Algorithmic complexity theory and the relative efficiency of financial markets

Giglio, Ricardo and Matsushita, Raul and Figueiredo, Annibal and Gleria, Iram and Da Silva, Sergio

Federal University of Santa Catarina

10 May 2008

Algorithmic complexity theory and the relative efficiency of financial markets

Ricardo Giglio[1], **Raul Matsushita**[2], **Annibal Figueiredo**[3],
Iram Gleria[4] **and Sergio Da Silva**[1]

[1]Department of Economics, Federal University of Santa Catarina, 88049970 Florianopolis SC, Brazil
[2]Department of Statistics, University of Brasilia, 70910900 Brasilia DF, Brazil
[3]Department of Physics, University of Brasilia, 70910900 Brasilia DF, Brazil
[4]Department of Physics, Federal University of Alagoas, 57072970 Maceio AL, Brazil

Abstract

Financial economists usually assess market efficiency in absolute terms. This is to be viewed as a shortcoming. One way of dealing with the relative efficiency of markets is to resort to the efficiency interpretation provided by algorithmic complexity theory. We employ such an approach in order to rank 36 stock exchanges, 37 individual company stocks, and 19 US dollar exchange rates in terms of their relative efficiency.

PACS numbers: 89.65.Gh; 89.20.–a

1. Introduction

Informational efficiency of financial markets is a central, hot topic in finance. It is meant that if price changes fully incorporate the information in possession of all the market participants such changes are unpredictable; the market in question is then said to be information efficient. Early empirical observations of the statistical properties of prices [1-6] detected a Brownian motion. In an efficient market populated by rational agents if the price is properly anticipated then it must fluctuate randomly [7]. The proof for this result employed the no-arbitrage condition. Such a stochastic process is a martingale that is, roughly, a probabilistic model of a fair game, one in which gains and losses cancel each other. In the presence of arbitrage, when informed traders move to

take advantage of their information, the price will move by an amount and in the direction that eliminates this advantage. As a result, there is an association between the unanticipated information obtained by the informed traders and the consequent movement of market prices. The uninformed traders can then infer from an observed price increase that some traders in the market have favorable information about the asset. Information is not wasted, and the price is a sufficient statistic for all the relevant information in possession of all the traders [8, 9]. So the efficient market theory [10] is the idea that prices in financial markets promptly adjust to reflect information.

After presenting an overview of market efficiency in their classic financial econometrics textbook, Campbell and coauthors [1] observed that (p. 24) the notion of *relative* efficiency, i.e. the efficiency of one market measured against another may be a more useful concept than the all-or-nothing (absolute) view taken by much of the traditional efficiency market literature. They made an analogy with physical systems that are usually given an efficiency rating based on the relative proportion of energy converted to work. Rating a piston engine 20% efficient means that on average 20% of the energy contained in the engine's fuel is used to turn the crankshaft, with the remaining 80% lost to other forms of work such as heat, light, or noise. It makes no sense to test statistically whether the engine is perfectly efficient. Similarly, the efficient market is an idealization that is unattainable, but that serves as a useful benchmark for measuring relative efficiency.

Indeed, one must regard the efficient market hypothesis as a limiting case. In practice, prices reflect only the information for which the acquisition costs cannot outweigh the benefits. There are also transaction costs. And information may not be widespread and reflected in the prices. Also, following the arrival of new information, market participants may diverge from one another in how they think it will impact prices; in other words, expectations are heterogeneous. Residual inefficiencies are always present in actual markets. These inefficiencies can introduce artificial patterns and then redundant information in real-world financial price series. Thus it is inappropriate to assess whether a given actual market is efficient or not. This is not a yes-no question; rather, efficiency should measure to what extent one market departs from the idealized efficient market. Relative efficiency is what really matters.

Algorithmic complexity theory makes a connection between the efficient market hypothesis and the unpredictable character of stock returns because a time series that has a dense amount of nonredundant information (such as that of the idealized efficient

market) exhibits statistical features that are almost indistinguishable from those observed in a time series that is random [12, 1-7]. As a result, measurements of the deviation from randomness provide a tool to assess the degree of efficiency of a given market. Because algorithmic complexity theory cannot discriminate between trading on noise and trading on information, it detects no difference between a time series conveying a large amount of nonredundant information and a pure random process.

We adopt such an approach. Doing so, we will be able to rank stock exchanges, individual company stocks, and currencies in terms of their relative efficiency. We will find, for instance, that the S&P 500 is 99.1% efficient whereas the Colombo Stock Exchange of Sri Lanka is only 10.5% efficient. This means that prices in the American stockmarket incorporate much more nonredundant information than its Sri Lankan counterpart does.

The absolute efficiency of stockmarkets has been investigated in a huge number of papers (for a review see Ref. [13]; an early survey is Ref. [10]), but we could track only three previous attempts similar to ours to deal with their relative efficiency. Shmilovici and colleagues [14] provide a test for the efficient market hypothesis (and not exactly for the relative efficiency of stockmarkets) that is based on the insight that the compression of the efficient market time series is not possible since there are no patterns. In that case, "stochastic complexity" is highest. The stochastic complexity of a time series is a measure of the number of binary digits needed to represent and reproduce the information in the time series. The authors use the Rissanen context tree algorithm to track patterns and then compress the series of 13 stock exchange indices as well as the stock prices of the companies listed on the Tel-Aviv 25. (Shmilovici and coauthors also claim that the approach in Chen and Tan [15] is one particular case of theirs.) Gabjin and coauthors [16] explicitly addressed the relative efficiency of currency markets using a tool called the approximate entropy statistic [17] aimed at tracking similar subsets in a time series. The approximate entropy statistic is strongly alignment-dependent, however, in that two parameters (lag interval and the similarity criterion) must be specified a priori [18]. Moreover, the statistic is not robust in terms of transitivity, i.e. a time series that is found to be more complex than another for some a priori lag and similarity parameters can become less complex for different values of the parameters [19]. Section 2 will show our distinct perspective applied to a larger database along with a simpler methodology that is based straightforwardly on the Lempel-Ziv ("deterministic") complexity index.

The rest of this paper is organized as follows. Section 2 discusses algorithmic complexity theory in more detail. Section 3 presents data and performs analysis. And Section 4 concludes.

2. Lempel-Ziv algorithmic complexity

Shannon entropy of information theory implies that a genuinely random series is the polar case where its expected information content is maximized, in which case there is maximum uncertainty and no redundancy in the series. The algorithmic (Kolmogorov) complexity of a string is given by the length of the shortest computer program that can produce the string. The shortest algorithm cannot be computable, however. Yet there are several ways to circumvent this problem. Lempel and Ziv [20] suggest a useful measure that does not rely on the shortest algorithm. (Rissanen context tree algorithm of stochastic complexity is another alternative.) Kaspar and Schuster [21] provide an easily calculable measure of the Lempel-Ziv index which runs as follows.

A program either inserts a new digit into the binary string $S = s_1, \ldots, s_n$ or copies the new digit to S. The program then reconstructs the entire string up to the digit $s_r < s_n$ that has been newly inserted. Digit s_r does not come from the substring $s_1, \ldots, s_{r-1}$; otherwise, s_r could simply be copied from $s_1, \ldots, s_{r-1}$. To learn whether the rest of S can be reconstructed by either simply copying or inserting new digits we take s_{r+1}, and then check whether this digit belongs to one of the substrings of S, in which case it can be obtained by simply copying it from S. If s_{r+1} can indeed be copied the routine goes on until a new digit (which once again needs to be inserted) appears. The number of newly inserted digits plus one (if the last copy step is not followed by inserting a digit) gives the complexity measure c of the string S.

As an illustration, consider the following three strings of 10 binary digits each.

A   0000000000

B   0101010101

C   0110001001

At first sight one might correctly guess that A is less random so that A is less complex than B, which in turn is less complex than C. The complexity index c agrees with such an intuition. For the string A one has only to insert the first zero and then rebuild the entire string by copying this digit; thus $c = 2$, where c is the number of steps necessary

to create a string. For the string B one has to additionally insert digit 1 and then copy the substring 01 to reconstruct the entire string; thus $c = 3$. For the string C one has to further insert 10 and 001, and then copy 001; thus $c = 5$.

The complexity of a string grows with its length. The genuinely random string asymptotically approaches its maximum complexity r as its length n grows following the rule $\lim_{n \to \infty} c = r = \frac{n}{\log_2 n}$ [21]. One may thus compute a positive finite normalized complexity index $LZ = \frac{c}{r}$ to get the complexity of a string relative to that of a genuinely random one. As the string approaches infinite $LZ \to 1$; however, very complex series in practical finite experiments usually have an LZ a little bit above one. The index also makes it possible to compare strings of distinct lengths as long as their lengths $\geq$ 1,000. Figure 1 shows a computer-generated pseudo-random string reaching the bulk of its convergence as it nears 1,000; from this threshold on there is slow asymptotical convergence toward an LZ index of one.

Here we consider sliding time windows, calculate the index for every window, and then get the average. For instance, for a time series of 2,000 datapoints and a chosen time window of 1,000 observations we first compute the LZ index of the window from 1 to 1,000, then the index of the window from 2 to 1,001, and so on, up to the index of the window from 1,001 to 2,000. Then we take the average of the indices.

As an illustration, Figure 2 shows three time series of 15,000 observations each, and the computed LZ indices of 14,000 sliding time windows of length 1,000. Figure 2a displays the index evolution of the series of computer-generated pseudo-random numbers (average LZ index = 1.062622). Figure 2b shows the index evolution of the series of the distances ("returns") between the first 15,001 adjacent prime numbers (average LZ index = 1.014342). And Figure 2c shows the index evolution of the series of natural logs of the distances between the first 15,001 adjacent primes (average LZ index = 1.025574). The distances between adjacent primes are believed to be genuinely random, and this agrees with our computed indices in Figure 2. Figure 3 shows the evolution of the LZ index for different parameter values of the logistic equation (1,000 iterations with the starting value set at 0.25). The solution to this equation is a series that depends on the value of its growth parameter. The series gets stable for low values of the parameter, which means low complexity. As the parameter grows the series behaves periodically, and then goes chaotic as the parameter approaches 4. This increased complexity agrees with the LZ index evolution in Figure 3.

Efficiency can be thought of as an absence of statistical dependence. The *LZ* complexity can also capture the type of statistical dependence, and is general enough to encompass Markov processes. In a sense, the *LZ* index provides a times series statistic more basic than those produced by autoregressive processes of any order. The Markov process describes the behavior of a random variable whose current observation depends on some subsets of past observations. For instance, a random process generated by a fair coin has no memory, so the probability of the next value conditional to the current one is $\frac{1}{2}$, i.e. the realizations are independent. Yet the sequence of, say, "sunny" or "rainy" days is still a random process, but the realizations are statistically dependent on each other. The probability of a sunny day tomorrow conditional to a sunny day today might be 0.75 rather than $\frac{1}{2}$.

If a time series exhibits statistical dependence, we expect it to present more patterns and accordingly to have lower complexity than an independent one. That the *LZ* complexity measure can also capture this feature can be seen in the example of an unfair coin, whose probability distribution of the current value is dependent on the previous one with probability p. If $p = \frac{1}{2}$ the process collapses to that of the fair coin with no memory (i.e. an independent process); if either $p = 0$ or $p = 1$ the current value is perfectly predictable from past observation, i.e. the process is totally dependent.

To set an experiment with an unfair coin we begin with an initial state that is independent from any event; thus, the process equals 0 if the pseudo-random number chosen is lower than $\frac{1}{2}$, and equals 1 otherwise. For all of the remaining realizations, the next value will repeat the current one with probability p.

If a current realization is 1, the next will be 1 again if the pseudo-random number generated is higher than p, and will be 0 otherwise. If the current realization is 0, the next will also be 0 if the pseudo-random number is higher than p, and will be 1 otherwise.

Now assume that the unfair coin repeats the current character with a probability $p = 0.2$. This means that if the current value is 1, the next will continue to be 1 with a probability $p = 0.2$. As an illustration, we generate two strings for such a Markovian process with $p = 0.2$ and $p = \frac{1}{2}$ respectively.

D   01111000001111111111

E   11001100100001010101

Figure 4 shows the *LZ* index of the sequences generated by the unfair coin for different probabilities. In the parabola-shaped curve, the index reaches its maximum when the process is statistically independent, and approaches its minimum when there is total dependency.

Some applications of the *LZ* index include the following. Li and coauthors [22] used the *LZ* index for DNA sequences to reconstruct the phylogenetic tree of several species of placental mammals among primates, rodents, and ferungulates. They concluded that primates and ferungulates are more closely related than rodents. Ferrario and colleagues [19] employed the *LZ* index to identify intrauterine growth restricted fetuses by analyzing the complexity of the heart rate variability signals. Shmulevich and Povel [23] suggested that the *LZ* index is able to identify temporal complexity in musical rhythms despite the fact that one needs long sequences in order to measure complexity.

3. Data and analysis

We collected seven years of daily data from July 2000 to July 2007 (2,000 observations) from 36 stock exchange indices (Table 1), 37 stock prices of companies listed on the NYSE, Nasdaq, and Bovespa (Table 2), and 19 US dollar exchange rates (Table 3). The source was Yahoo Finance and EconStats.

Analysis was performed with simple returns of the raw series. The return series were coded as ternary strings as follows [14]. Assuming a stability basin b for a return observation ρ_t, a datapoint d_t of the ternary string was coded as $d_t = 0$ if $\rho_t \leq -b$, $d_t = 1$ if $\rho_t \geq +b$, and $d_t = 2$ if $-b < \rho_t < +b$. The series would have become binary if we had shrunk the stability basin to the attractor zero, i.e. $b = 0$; yet we assumed $b = 0.0025$ following Shmilovici and colleagues. (We checked for the effects of changing b only to realize that the rankings did not alter too much; yet more research is needed to consider a more sophisticated analysis in the choice of b.) As an illustration, we take five daily percentage returns of the S&P 500 and compare them with $b = 0.25\%$. From 18 to 22 June 2007 the returns were, respectively, 0.652%, –0.1226%, 0.1737%, –1.381%, and 0.6407%. Thus the trading week was coded as 12201.

Figure 5 shows the evolution of the index using 1,000 sliding windows for (a) the computer-generated pseudo-random series (average *LZ* = 1.0180), (b) returns of the

Dow Jones (average LZ = 1.0201), (c) returns of the Shanghai Composite (average LZ index = 1.0032), and (d) returns of the Karachi 100 (average LZ index = 0.9918). Table 1 shows the average LZ index for the other stock exchanges. As can be seen, all the series seem to be very complex. They look more like the genuinely random series than the totally redundant, perfectly predictable series. (Check Figure 3 again to see that a periodic series has an LZ well below one.) Inspired by the experiment in Figure 1 we decided to consider $LZ = 1$ as a threshold in order to assess the relative efficiency of the series. We counted the number of occurrences where the LZ index was caught above one, and then considered that as a measure of relative efficiency. For the pseudo-random series the $LZ = 1$ threshold was surpassed 98.8% of the times; thus we say that it is 98.8% efficient.

The Dow Jones, Shanghai Composite, and Karachi 100 were found to be, respectively, 95.4%, 49.5%, and 23.7% efficient. Note that the Dow Jones series nears the pseudo-random series. Table 1 shows the measures for the other stock exchanges. As can be seen, the S&P 500 even beats the pseudo-random series. Thus it is safe to conclude that this American stockmarket is almost efficient. By contrast, the Colombo Stock Exchange was found to be only 10.5% efficient, which means that stock prices in that market convey some redundant information.

The procedure above was repeated for selected company stock prices (Table 2). Figure 6 shows the evolution of the LZ index using 1,000 sliding windows for (a) Coca-Cola (100% efficient), (b) Yahoo (99.65% efficient), (c) Vale (92.75% efficient), and (d) Aracruz (66.67% efficient). Table 3 shows the relative efficiency of selected US dollar exchange rates, whereas Figure 7 displays the LZ index evolution over the same sliding windows for the dollar price in terms of (a) pound sterling (average LZ index = 1.0223; 99.81% efficient), (b) euro (average LZ index = 1.0254; 99.45% efficient), (c) Brazilian *real* (average LZ index = 1.0156; 92.60% efficient), (d) Indian rupee (average LZ index = 0.9958; 43.54% efficient), and (e) Chinese yuan (average LZ index = 0.9266; 17.94% efficient). As for the latter example, China's currency remained pegged to the US dollar from 16 June 1994 to 21 July 2005; and this explains the initial low complexity of the dollar price in yuan terms in Figure 7. Matsushita and coauthors [24] showed that this meant that China played the "chaos game" over the period, which produced a Sierpinski triangle fractal in the yuan-dollar rate.

4. Conclusion

By considering data from 36 stockmarket indices, 37 individual company stock prices, and 19 US dollar exchange rates, this paper puts forward one way to assess the relative efficiency of financial markets. This is made possible thanks to the efficiency interpretation provided by algorithmic complexity theory. The latter makes a connection between the efficient market hypothesis and the unpredictable character of asset returns. The idealized efficient market generates a time series that has a dense amount of nonredundant information, and thus presents statistical features similar to a genuinely random time series.

Physical systems are usually given an efficiency rating based on the relative proportion of energy converted to work. We suggest an analogous efficiency rating based on the relative amount of nonredundant information conveyed by financial prices. The price of the idealized efficient market conveys information that is fully nonredundant; this market is then said to be 100% efficient.

Yet prices in real-world markets reflect only the information for which the acquisition costs cannot outweigh the benefits. Also, there are transaction costs, inside trading that cannot allow uninformed traders to extract information from prices, and heterogeneous expectations. Since such residual inefficiencies are always present in actual markets one should not expect some of them to be efficient in absolute terms. Yet by considering the random efficient market as a benchmark one can, for instance, say that the S&P 500 is 99.1% efficient whereas the Colombo Stock Exchange is only 10.5% efficient. This means that prices in the American stockmarket incorporate much more nonredundant information than its Sri Lankan counterpart does. And by saying that the Chinese currency is only 17.94% efficient we mean that the dollar price in yuan terms conveys some redundant information; indeed, artificial patterns were introduced by the recent peg of the currency to the US dollar.

Acknowledgments

Figure 1. As its length increases, a typical, computer-generated pseudo-random string seems to asymptotically converge to an *LZ* index of one.

Figure 2. *LZ* index evolution of (a) a series of computer-generated pseudo-random numbers (average *LZ* index = 1.062622), (b) a series of the distances between the first 15,001 adjacent prime numbers (average *LZ* index = 1.014342), and (c) a series of natural logs of the distances between the first 15,001 adjacent primes (average *LZ* index = 1.025574).

Figure 3. *LZ* index for increased values of the logistic growth parameter (1,000 iterations with the starting value set at 0.25). The series gets stable and then periodic for low values of the parameter (*LZ* complexity index well below one), and then goes chaotic as the parameter approaches 4.

Figure 4. *LZ* index and statistical dependence: experiment with an unfair coin. Maximum complexity is reached if the probability p is $\frac{1}{2}$, and lack of complexity obtains as the probability is either 0 or 1.

Figure 5. *LZ* index evolution over 1,000 sliding windows for (a) a computer-generated pseudo-random series (average *LZ* = 1.0180), (b) returns of the Dow Jones (average *LZ* = 1.0201), (c) returns of the Shanghai Composite (average *LZ* index = 1.0032), and (d) returns of the Karachi 100 (average *LZ* index = 0.9918).

Figure 6. Evolution of the *LZ* index using 1,000 sliding windows for (a) Coca-Cola (100% efficient), (b) Yahoo (99.65% efficient), (c) Vale (92.75% efficient), and (d) Aracruz (66.67% efficient).

Figure 7. *LZ* index evolution over 1,000 sliding windows for the dollar price in terms of (a) pound sterling (average *LZ* index = 1.0223; 99.81% efficient), (b) euro (average *LZ* index = 1.0254; 99.45% efficient), (c) Brazilian *real* (average *LZ* index = 1.0156; 92.60% efficient), (d) Indian rupee (average *LZ* index = 0.9958; 43.54% efficient), and (e) Chinese yuan (average *LZ* index = 0.9266; 17.94% efficient).

Table 1. The relative efficiency of selected stockmarket indices

| Stock Exchange	Country	Average LZ Index	Degree of Efficiency*, %
S&P 500	USA	1.0232	99.1
DAX 30	GER	1.0257	98.4
Nikkei 225	JPN	1.0432	98.2
All Ordinaries	AUS	1.0246	97.8
ATX	AUT	1.0173	97.4
Dow Jones	USA	1.0201	95.4
Korea Composite	KOR	1.0163	94.9
Tel-Aviv 100	ISR	1.0187	92.9
Hang Seng	HKG	1.0151	91.5
Straits Times	SIN	1.0153	90.3
CAC 40	FRA	1.0138	88.4
Helsinki General	FIN	1.0149	88.4
Kuala Lumpur SE	MAS	1.0158	88
FTSE 100	UK	1.0106	86.6
Prague X	CZE	1.0139	81
Bel 20	BEL	1.0118	80.4
IBC	VEN	1.0110	79.9
Madrid General	ESP	1.0201	79.3
Swiss Market	SUI	1.0101	78.4
Nasdaq Composite	USA	1.0080	75.4
Amsterdam EX	NED	1.0100	74.4
Bovespa	BRA	1.0127	67.8
IPC	MEX	1.0060	64
Merval	ARG	1.0050	62.9
Jakarta Composite	IDN	1.0054	62.1
Istanbul 100	TUR	1.0085	61.3
Moscow Times	RUS	1.0050	59.2
Copenhagen	DEN	1.0025	58.7
Athex Composite	GRE	1.0048	56.9
Bombay SE	IND	1.0010	53.3
Taiwan Weighted	TPE	1.0006	50.3
Shanghai Composite	CHN	1.0032	49.5
Philippines	PHI	0.9987	43.1
Lima General	PER	0.9903	37.9
Karachi 100	PAK	0.9918	23.7
Colombo SE	SRI	0.9795	10.5

* Hits above the threshold $LZ = 1$

Table 2. The relative efficiency of selected company stocks

Company	Stock Exchange	Average LZ Index	Degree of Efficiency*, %
Amazon	NYSE	1.0416	100
Coca-Cola	NYSE	1.0324	100
P&G	NYSE	1.0264	99.97
Intel	Nasdaq Composite	1.0292	99.92
eBay	Nasdaq Composite	1.0377	99.8
General Electric	NYSE	1.0274	99.66
Yahoo	Nasdaq Composite	1.0310	99.65
Texaco	NYSE	1.0264	99.46
Cisco	Nasdaq Composite	1.0357	99.44
Petrobras	Bovespa	1.0284	99.43
Pfizer	NYSE	1.0327	99.39
HP	NYSE	1.0298	99.38
Microsoft	Nasdaq Composite	1.0286	99.25
Goldman Sachs	NYSE	1.0311	98.78
J&J	NYSE	1.0275	98.73
Unilever	NYSE	1.0297	98.44
Nissan	Nasdaq Composite	1.0178	97.58
Merrill Lynch	NYSE	1.0279	97.33
JP Morgan	NYSE	1.0281	96.7
Oracle	Nasdaq Composite	1.0206	94.93
Citigroup	NYSE	1.0314	94.59
Vale	Bovespa	1.0193	92.75
Embraer	Bovespa	1.0258	91.59
Itau	Bovespa	1.0183	86.74
FedEx	NYSE	1.0186	86.6
Bradesco	Bovespa	1.0172	85.88
Exxon	NYSE	1.0161	85.56
Ford	NYSE	1.0152	84.26
Marcopolo	Bovespa	1.0072	77.36
Americanas	Bovespa	1.0136	76.68
Ipiranga	Bovespa	1.0111	76.55
Toyota	NYSE	1.0100	76.32
Wal-Mart	NYSE	1.0074	71.42
Ambev	Bovespa	1.0108	70.27
Aracruz	Bovespa	1.0048	66.67
Duratex	Bovespa	1.0048	65.17
Celesc	Bovespa	1.0005	50.03

$*$ Hits above the threshold $LZ = 1$

Table 3. The relative efficiency of selected US dollar exchange rates

Currency	Country	Average *LZ* Index	Degree of Efficiency*, %
Pound Sterling	UK	1.0223	99.81
Swedish Krona	SWE	1.0236	99.71
Norwegian Krone	NOR	1.0314	99.60
Euro	Eurozone	1.0253	99.45
New Zealand Dollar	NZL	1.0248	99.20
Swiss Franc	SUI	1.0169	99.12
Icelandic Krona	ISL	1.0184	97.48
Mexican Peso	MEX	1.0254	96.58
Brazilian Real	BRA	1.0156	92.60
Canadian Dollar	CAN	1.0219	90.07
South African Rand	RSA	1.0177	86.41
Japanese Yen	JPN	1.0153	85.51
Singapore Dollar	SIN	1.0074	66.48
Australian Dollar	AUS	1.004	63.41
Indian Rupee	IND	0.9957	43.54
Colombian Peso	COL	0.9913	21.98
Taiwan New Dollar	TPE	0.9794	21.17
Chinese Yuan	CHN	0.9265	17.94
Sri Lanka Rupee	SRI	0.9687	11.84

* Hits above the threshold $LZ = 1$

References

[1] Bachelier M L 1900 *Theorie de la Speculation* (Paris: Gauthier-Villars)

[2] Kendall M G 1953 The analysis of economic time series, part I: prices *Journal of the Royal Statistical Society A* **96** 11–25

[3] Roberts H V 1959 Stock market patterns and financial analysis: methodological suggestions *Journal of Finance* **14** 1–10

[4] Alexander S S 1961 Price movements in speculative markets: trends or random walks *Industrial Management Review* **2** 7–26

[5] Cootner P A (ed.) 1964 *The Random Character of Stock Market Prices* (Cambridge, MA: MIT Press)

[6] Fama E F 1965 The behavior of stock market prices *Journal of Business* **38** 34–105

[7] Samuelson P A 1965 Proof that properly anticipated prices fluctuate randomly *Industrial Management Review* **6** 41–50

[8] Hayek F A 1945 The use of knowledge in society *American Economic Review* **35** 519–30

[9] Grossman S J 1976 On the efficiency of competitive stock markets where traders have diverse information *Journal of Finance* **31** 573–85

[10] Fama E F 1970 Efficient capital markets: a review of the theory and empirical work *Journal of Finance* **25** 383–417

[11] Campbell J Y, Lo A W and MacKinlay A C 1997 *The Econometrics of Financial Markets* (Princeton: Princeton University Press)

[12] Mantegna R N and Stanley E 2000 *An Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge: Cambridge University Press)

[13] Beechey M, Gruen D and Vickery J 2000 The efficient market hypothesis: a survey *Reserve Bank of Australia Research Discussion Paper* 2000–01

[14] Shmilovici A, Alon-Brimer Y and Hauser S 2003 Using a stochastic complexity measure to check the efficient market hypothesis *Computational Economics* **22** 273–84

[15] Chen S H and Tan C W 1999 Estimating the complexity function of financial time series: an estimation based on predictive stochastic complexity *Journal of Management and Economics* **3** 3

[16] Gabjin O, Kim S and Eom C 2007 Market efficiency in foreign exchange markets *Physica A* **382** 209–12

[17] Pincus S 1991 Approximate entropy as a measure of system complexity *Proceedings of the National Academy of Sciences of the United States of America* **88** 2297–301

[18] Richman J S and Moorman J R 2000 Physiological time-series analysis using approximate entropy and sample entropy *American Journal of Physiology: Heart and Circulatory Physiology* **278** 2039–49

[19] Ferrario M, Signorini M and Magenes G 2005 Complexity analysis of the fetal heart rate for the identification of pathology in fetuses *Computers in Cardiology* **32** 989–92

[20] Lempel A and Ziv J 1976 On the complexity of finite sequences *IEEE Transactions on Information Theory* **22** 75–81

[21] Kaspar F and Schuster H G 1987 Easily calculable measure for the complexity of spatiotemporal patterns *Physical Review A* **36** 842–48

[22] Li B, Li Y B, He H B 2005 LZ complexity distance of DNA sequences and its application in phylogenetic tree reconstruction *Genomics, Proteomics & Bioinformatics* **3** 206–12

[23] Shmulevich I and Povel D J 2000 Measures of temporal pattern complexity *Journal of New Music Research* **29** 61–9

[24] Matsushita R, Figueiredo A, Gleria I and Da Silva S 2007 The Chinese chaos game *Physica A* **378** 427–42