

Maximum Entropy Production Principle for Stock Returns

Paweł Fiedor

African Institute of Financial Markets and Risk Management
University of Cape Town
Rondebosch, Cape Town, 7700, South Africa
Email: Pawel.F.Fiedor@IEEE.org

Abstract—In our previous studies we have investigated the structural complexity of time series describing stock returns on New York's and Warsaw's stock exchanges, by employing two estimators of Shannon's entropy rate based on Lempel–Ziv and Context Tree Weighting algorithms, which were originally used for data compression. Such structural complexity of the time series describing logarithmic stock returns can be used as a measure of the inherent (model-free) predictability of the underlying price formation processes, testing the Efficient–Market Hypothesis in practice. We have also correlated the estimated predictability with the profitability of standard trading algorithms, and found that these do not use the structure inherent in the stock returns to any significant degree. To find a way to use the structural complexity of the stock returns for the purpose of predictions we propose the Maximum Entropy Production Principle as applied to stock returns, and test it on the two mentioned markets, inquiring into whether it is possible to enhance prediction of stock returns based on the structural complexity of these and the mentioned principle.

I. INTRODUCTION

The concepts of entropy and entropy production in non-equilibrium processes formed the basis of a large part of modern physics, and have also created a core of ideological disputes about the evolution of the world in a more broadly defined scientific community. These discussions were started by the ground-breaking contributions made by Clausius, Boltzmann, Gibbs, and Onsager [1], [2], [3], [4]. Nowadays, the properties of entropy in various systems are widely known and studied [5], [6]. Much less developed are the concepts around the entropy production behaviour in non-equilibrium processes, even though these are crucial to understanding a wide variety of real-world systems [7], [8].

Most scientific disciplines strive to find an universal function, the extremum of which would determine the development of the studied system. Such functions have been successfully found only in some cases of relatively simple systems (the Fermat's principle, the principle of least action). Entropy has been seen by many researchers as a quantity describing the progress of non-equilibrium dissipative processes, which may be leading to such universal conclusions. Indeed in 1947 Prigogine proved the minimum entropy production principle. He subsequently applied it to the description of multiple non-equilibrium processes in physics, chemistry and biology [7], [9], [10], [11]. There still remains doubt as to whether one can formulate universal entropy principles, which would govern so diverse and dissimilar non-equilibrium processes, however.

978-1-4799-7492-4/15/\$31.00 ©2015 IEEE

Another principle, known as the Maximum Entropy Production Principle (MEPP), is much less known. The MEPP has been proposed and used by several scientists throughout the twentieth century, as they dealt with general theoretical issues of thermodynamics and statistical physics. This principle states that a non-equilibrium system develops so as to maximise its entropy production under present constraints. This principle has found many applications in physics, biology, and other fields [12], [13], [14], [15], [16]. We may argue for the existence of this principle in the following way: similarly to natural systems selecting the path out of available paths that minimises the potential or maximizes the entropy at the fastest rate given the constraints, financial systems choose the path that minimises the potential for profit at the fastest rate given the constraints (due to investors trying to grab these).

We may treat financial markets as non-equilibrium systems [17], [18]. Even though they are very complex systems indeed, it is fairly easy to keep track of entropy and entropy production within financial markets, as the price formation processes are describing the time-evolution of these systems. These processes leave a trace in the form of the time series describing prices at various time quanta, which are highly non-stationary, thus in practical applications time series containing logarithmic returns (percentage returns) are used [19], [20]. Having these time series we are able to estimate entropy production within such system (price formation process for a given financial instrument) by estimating the Kolmogorov–Sinai entropy [21], [22], or alternatively all positive Lapunov exponents. Since within the scope of financial markets Lapunov exponents tend not to have a finite value, we measure Kolmogorov–Sinai entropy by estimating Shannon's entropy rate based on discretised time series describing logarithmic returns for studied stocks [23], [6].

In our previous studies [24] we have analysed various markets in this way, and have found that as time quanta at which we measure the prices become shorter, the entropy production becomes smaller also. That is the price changes are more predictable in high frequency. We have also studied the relationship between the predictability (as quantified by entropy production) of a given price formation process and the profitability of standard trading algorithms operating on the same processes (stocks). As it happens, these standard trading algorithms do not perform better on the processes (stocks) which are more predictable [25], [24]. These methods do not use the structural complexity inherent in the time series describing prices, most likely due to their design assuming

there is no redundancy. To the best of our knowledge, the two afore-mentioned studies are the only ones dealing with the relation between predictability and profitability in the sense employed in this study, and both find no connection between the two (also see the introduction in Ref. [25]).

Since the standard methods of guessing future price changes do not use the information contained in entropy analysis of the price formation processes, we try to find out whether we can enhance the prediction simply by using the notion of entropy production itself. In doing so we introduce and partially test in practice the hypothesis of the Maximum Entropy Production Principle applied to financial markets. Of course, this principle states that the system maximises its entropy production under present constraints. There is as yet no single, unified theory of financial markets, and thus we do not know the constraints present. But we know that there are many such constraints, and we also know that financial markets are highly complex and very efficient systems, thus obviously a practical test will not find the system maximising entropy at all times. But if we can find that it does tend to follow the maximum entropy path more often than not, then it will provide a partial proof of this hypothesis. A full proof, as is often the case in studying complex systems, will be very hard if not impossible to obtain.

The paper is organised as follows. In Section 2 we present a method used to estimate entropy production in price formation processes, and the setup for testing the Maximum Entropy Production Principle. In Section 3 we apply the presented methodology to real data describing daily and high frequency stock returns from markets in Warsaw and New York. In Section 4 we discuss obtained results. In Section 5 we conclude the study and propose further research.

II. METHODS

To test the principle of maximum entropy production in financial markets we first need to define the concept of entropy production in the context of financial markets. We use the notion of entropy as defined in information theory, since price formation processes are information generating processes. Within this paradigm we can see the growth of Shannon's entropy with regards to the growth of the studied word length, that is Shannon's entropy rate, as an estimator of entropy production. Entropy rate, as defined by Shannon, is an estimator of Kolmogorov–Sinai entropy. The latter is an appropriate tool for accessing entropy for dynamical processes [21], [22].

Thus we need to define the entropy rate and a method of estimating it for practical purposes. The entropy rate, in information theory, is a term derivative to the notion of Shannon's entropy, which measures the amount of uncertainty in a random variable. Shannon's entropy of a random variable X is defined as

$$H(X) = - \sum_i p(x_i) \log_2 p(x_i) \quad (1)$$

summed over all possible outcomes $\{x_i\}$ with their respective probabilities of $p(x_i)$ [23].

The entropy rate is a generalisation of the notion of entropy for sequences of dependent random variables, and is thus a

good measurement of redundancy within studied sequences. For a stationary stochastic process $\mathcal{X} = \{X_i\}$ the entropy rate is defined either as

$$H(\mathcal{X}) = \lim_{n \rightarrow \infty} \frac{1}{n} H(X_1, X_2, \dots, X_n) \quad (2)$$

or as

$$H(\mathcal{X}) = \lim_{n \rightarrow \infty} H(X_n | X_1, X_2, \dots, X_{n-1}) \quad (3)$$

Entropy rate can be interpreted as a measure of the uncertainty in a quantity at time n having observed the complete history up to that point. Alternatively it can be seen as the amount of new information created in a unit of time in the studied process [6]. Entropy rate can thus be viewed as the maximum rate of information creation which can be processed as price changes for studied financial instruments [24]. This allows us to use Shannon's entropy rate as an estimator of entropy production within the studied price formation processes.

Entropy estimation within various systems constitutes a very broad and productive field. It has been especially active in the last two decades due to the advances in neurobiology, and the usefulness of information-theoretic constructs to studying many complex systems. Methods of estimating information-theoretic entropy rate can be divided into two separate groups [26]: maximum likelihood estimators, which analyse the empirical distribution of all words of a given length within the time series, and estimators based on data compression algorithms, particularly Lempel–Ziv [27], [28], [29] and Context Tree Weighting [30], [31] algorithms, both of which are precise even when operating on limited samples [32], [33]. The maximum likelihood estimators cannot deal with long-range dependencies, and as such we do not use them. In this study we use an estimator based on the Lempel–Ziv algorithm.

The Lempel–Ziv algorithm measures linear complexity of the sequence of symbols and has been first introduced by Jacob Ziv and Abraham Lempel in 1977 [29]. In practice, the algorithm counts the number of patterns in the studied time series, scanning it from left to right. The distribution of these can be used to find the entropy rate later. For example, the complexity of the below sequence

$$s = 101001010010111110$$

is equal to 8, as when scanning it from left to right we find eight distinct patterns [34]:

$$1|0|10|01|010|0101|11|110|$$

There have been numerous estimators of Shannon's entropy rate created on the basis of the Lempel–Ziv algorithm. In this study we follow earlier studies [25], [24] in using the estimator created by Kontoyiannis in 1998 [35]. It is widely used [31], [25] and has good statistical properties [35]. There is a large choice of slightly different variants to choose from [36], none of which have consistently better characteristics than the used estimator.

Formally, to calculate Shannon's entropy of a random variable X , the probability of each possible outcome $p(x_i)$ must be known. For practical purposes such probabilities are usually not known. Then entropy can be estimated by replacing probabilities with relative frequencies from observed

data. Estimating the entropy rate of a stochastic process is more complex than estimating entropy, as random variables in stochastic processes are usually interdependent. The mentioned Lempel–Ziv estimator of Shannon’s entropy rate is defined as:

$$\hat{H} = \frac{n \log_2 n}{\sum_i \Lambda_i} \quad (4)$$

where n is the length of the studied time series, and Λ_i denotes the length of the shortest substring starting from time i which has not yet been observed prior to time i (between times 1 and $i - 1$). It has been shown that for stationary ergodic processes $\hat{H}(X)$ converges to the entropy rate $H(X)$ with probability of 1 as n approaches infinity [35].

Having defined the methods for entropy rate estimation, we need to comment on the financial data that is to be fed into these algorithms. It is obvious that prices are the most important indicators of the conditions within financial markets. Other measures, for example volatility or realised volatility, may also be used, but are generally seen as less important. Since the time series describing prices themselves are problematic (non-stationary), we analyse time series containing logarithmic returns. Let us denote the most recent price of the studied financial instrument occurring on time t during the studied period by $p(t)$. Additionally, τ is the time horizon or price sampling frequency. Then for each stock the logarithmic returns are sampled

$$\rho_t = \ln(p(t)) - \ln(p(t - \tau)) \quad (5)$$

Such time series can describe logarithmic returns of the studied financial instruments at any time horizon, either daily (end of day prices) or intraday. For the purposes of the information-theoretic analysis we need to have discrete time series. Thus we discretise these time series by binning the values into Ω distinct states. The discrete logarithmic returns take values from an alphabet with cardinality Ω :

$$r_t \in \{1, \dots, \Omega\} \quad (6)$$

In our analysis we choose $\Omega = 4$, though in our previous studies we have shown that other values give reasonably similar results (see discussion in Ref. [24]). The four states represent four quartiles, and each state is assigned the same number of data points. This means that the model has no unnecessary parameters, which could affect the results and conclusions reached while using the data. Comparable experimental setups have been used in similar studies [25], [24] and proved to be very efficient [39].

Finally, we have to formally introduce the principle of maximum entropy production for financial markets, as well as the setup of the test for the principle of maximum entropy production in practice. In this study we understand entropy production in the sense of Shannon’s entropy rate, as defined above. Then if we study particular time series describing logarithmic returns in a window of length μ ($\{r_{t-\mu}, \dots, r_t\}$) we can, under the principle of maximum entropy production, say that the next price (r_{t+1}) will be assigned the state which maximises Shannon’s entropy rate, that is $H(\{r_{t-\mu}, \dots, r_{t+1}\})$. This approach does not predict the price exactly, but to the accuracy as specified in the discretisation step, that is differentiating among Ω different levels. In other words $\alpha, \beta \in \{1, \dots, \Omega\}$.

Formally the principle of maximum entropy production for stock returns can be described as:

$$\Psi = P(r_{t+1} = \alpha \mid \forall_{\beta \neq \alpha} H(\{r_{t-\mu}, \dots, r_t, \alpha\}) > H(\{r_{t-\mu}, \dots, r_t, \beta\})) = 1 \quad (7)$$

Given no constraints this approach should give 100% accuracy. But we know that financial markets are highly complex adaptive systems with many (unknown) constraints, and we have shown that they are nearly efficient [24], thus in practice we can only see whether the future (out of sample, predicted) state which maximises entropy production is significantly overrepresented in the historical data for various financial markets. Thus the principle of maximum entropy production for stock returns in practice could be described as:

$$\Psi \gg \frac{1}{\Omega} \quad (8)$$

In other words assuming discretisation into four states we want to study financial data with various values of parameter μ and find whether the percentage of times we can guess the next log return by maximising entropy rate given the history of length μ significantly exceeds 0.25, which we would get by guessing it randomly. We will attempt to test this in practice by estimating Ψ as the percentage of correct guesses when applying this procedure to historical data, and moving through the time series with moving window of length μ . This follows directly from the definition above, we calculate the probability (naively estimated through frequency from historical data) of the next price change being in the quartile which maximises entropy rate of the appropriate window. These will be further averaged over all studied stocks. If we can show that the above inequality holds in practice it would be the first step to proving that the principle of maximum entropy holds for financial markets. This in turn would be of great importance to market practitioners, and particularly analysts of financial markets. Though it is worth noting that we do not claim this principle in itself would prove a profitable trading strategy, only that it would hint an underlying principle according to which the markets work. In particular, the above does not incorporate the amount of loss each time a wrong prediction is made, neither does it incorporate the trading costs, or the effect of coarse-graining the predictions.

III. EXPERIMENTAL RESULTS

We start by introducing the datasets on which we test the principle of maximum entropy production, and the way in which we discretise them. In this study we use datasets from Warsaw’s and New York’s stock exchanges. This covers both emerging and mature markets, and the database used is large enough for the results not to be accidental.

We start with end of day time series for 360 securities traded on Warsaw Stock Exchange (GPW) (all securities which were listed for over 1000 consecutive days). This data has been downloaded from DM BOŚ database¹ and was up to date as of the 5th of July 2013. The data is transformed in the standard way for analysing price movements, so that the data points are the log ratios between consecutive daily closing prices, and those data points are, for the purpose of

¹<http://bossa.pl/notowania/metastock/>

the Lempel–Ziv estimator, discretised into four distinct states (for the discussion of this choice see Ref. [24]), ignored here due to length constraints). For high frequency price changes we use another set of data, that is intraday price changes for 707 securities listed on Warsaw Stock Exchange which had over 2500 price changes recorded (all securities with at least 2500 recorded price changes). Unlike in the daily data, in the intraday time series we have ignored data points where there was no price change, as those are not meaningful for data points separated by very small, and not uniform, time intervals. Most of the data points show no price change and calculating entropy in that manner would be pointless, thus in effect we estimate the predictability of the next price change. On the daily data no price change is meaningful however, and hence were not deleted from the data. Analysis with those removed from daily data gives similar results [24]. Note that we also ignore daily market close for the high frequency data, this will have negligible effect on the results. Please note that contrary to the standard procedures in some strands of economics, we do not remove drift or trend. That is because such a procedure adds the researchers own vision of what a trend is, and as such we would end up analysing our own construct, rather than the market. Such a procedure should be applied and compared with the results presented here in future research, which is omitted here due to length constraints.

We also study the daily price time series for 98 securities traded on New York Stock Exchange (NYSE100) (the 2 missing stocks were excluded due to missing data). The data has been downloaded from Google Finance database² and was up to date as of the 11th of November 2013, going 10 years back. The data is transformed in the same way as above. For high frequency price changes we use another set of data, that is intraday (1-minute intervals) price changes for the same 98 securities listed on NYSE 100. The data covers 15 days between the 21st October 2013 and the 8th of November 2013. Therefore the length of time series in both cases are comparable and sufficient for the used algorithms.

We note that since in the prediction we have to use two times series, one without the value that we predict ($\{r_{t-\mu}, \dots, r_t\}$) and one including this value ($\{r_{t-\mu}, \dots, r_{t+1}\}$). This raises a problem for the discretisation step. Since the length for these two samples is different the resulting quartiles may also be different. Additionally, if we want to check whether the prediction is correct we may run into an issue where the real values at time $t + 1$ is between quartiles, and we will have no good rule to decide on which of the bordering states should this value be assigned. For simplicity, in this study we have discretised the entire time series of length N , and not the subseries of length μ that are under study in each step as we move through the data. This is potentially problematic as we take into account future data. Nonetheless, such quartile discretisation is quite stable, and it also represents the analysts' concentration on what different levels of log returns mean in the long run. Thus we do not find this approach troubling. But to be cautious we need to consider a small correction to the principle of maximum entropy production for stock returns as defined above. If we use a global discretisation step then for some cases we may find that the predicted state for time $t + 1$ which maximises

entropy production may coincide with the state which is under-represented in the studied subseries. Then this state would be over-represented in the rest of the studied times series, increasing the likelihood of guessing it by random chance over the standard $1/\Omega$. Therefore, to be cautious, we need to compare the results to what we can reasonably get by randomly guessing the future value of logarithmic returns, taking into account the worst case scenario of described over-representation of the guessed state in the studied time series, for a given length of the time series, and a given value of the window μ . Then the principle of maximum entropy production for stock returns is defined as:

$$\Psi \gg \frac{N/\Omega}{N - \mu} \quad (9)$$

As can be seen, this correction is increasing with increasing ratio μ/Ω , thus we are able to find out whether this problem affects our analysis by observing how the results change with changing the parameter μ . In this study we use μ between 20 and 100. There is not much sense in estimating Shannon's entropy rate for time series shorter than 20, as the statistical noise will be significant for such calculations. On the other hand, we are not interested in using very long time series, thus there is no need to analyse values of μ higher than 100. For the above correction we will use N equal to the average within the studied set of stocks.

In figure 1 we show the results of testing the Maximum Entropy Production Principle for Warsaw Stock Exchange daily logarithmic returns (discretised into four quartiles) database of 360 stocks. In figure 1a we show an estimate of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all studied stocks with regards to the window length (μ) for which we employ the MEPP. This estimate is performed by moving through the whole studied series with the running window of size μ , and finding the percentage of times we guess the next price movement correctly (when using the MEPP). Within this figure we also show error bars of plus and minus one standard deviation. Further, the two lines denote the theoretic percentages of times we would expect to guess the next price movement by pure chance. The line below (horizontal) is unadjusted, or in other words $1/\Omega$, and the line above is adjusted for our setup as explained above. These results allow us to test the hypothesis as outlined in equations (8) and (9), for all studied values of the window length. These also allow us to see whether the window length affects the results, which can hint as to whether the adjustment mentioned above is necessary. In figure 1b we show a histogram of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all values of μ , together with vertical lines for unadjusted (on the left) and adjusted (on the right) theoretic percentages of times we would expect to guess the next price movement by chance. This allows us to clearly see how many of the studied stocks fall below the value expected for random guessing, helping to assess the stability of our results. In figure 1c we show a scatterplot of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all values of μ with regards to the Lempel–Ziv entropy rate (an estimator of the predictability of the price formation processes) for the whole studied time series. This scatterplot also features a linear regression fitted to the data, to present the correlation between

²<http://www.google.com/finance/>

the two variables. This scatterplot allows us to see whether predictions based on the MEPP are working better for more predictable price formation processes, which is not the case for standard methods of prediction used in trading algorithms. There is no clear relationship between the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all values of μ and the length of the whole studied time series, and as such these results have not been presented.

In figures 2, 3, and 4 we show the same results for Warsaw Stock Exchange intraday logarithmic returns (discretised into four quartiles) database of 707 stocks, New York Stock Exchange daily logarithmic returns (discretised into four quartiles) database of 98 stocks, and New York Stock Exchange intraday logarithmic returns (discretised into four quartiles) database of 98 stocks respectively.

In figure 5 we show the distributions of the estimates of the percentage of times the Maximum Entropy Production Principle (MEPP) correctly predicts the next price move (Ψ) averaged over all values of the parameter μ for logarithmic returns (discretised into four quartiles) for Warsaw Stock Exchange daily logarithmic returns (discretised into four quartiles) database of 360 stocks (figure 5a), Warsaw Stock Exchange intraday logarithmic returns (discretised into four quartiles) database of 707 stocks (figure 5b), New York Stock Exchange daily logarithmic returns (discretised into four quartiles) database of 98 stocks (figure 5c), and New York Stock Exchange intraday logarithmic returns (discretised into four quartiles) database of 98 stocks (figure 5d). These allow us to see where the mass of the distribution is, and whether the distributions have fat tails, which could be helpful for analysts and traders.

IV. DISCUSSION

First and foremost, we want to confirm or deny the hypothesis of the Maximum Entropy Production Principle for stock returns. Results presented in figures 1–4 support the hypothesis as outlined in equations (8) and (9). To decide whether we need the adjustment mentioned above, we observe how the results (Ψ) change with various values of window length μ and time series size N . We find that Ψ does not depend on either of these, with a small exception for low values of μ where statistical noise changes the results (although only slightly). If the scenario of described over-representation of the guessed state in the studied time series were affecting our analysis we would expect to see higher values of Ψ for higher values of μ and lower values of N , which is not the case. Therefore we conclude that it is safe to use the unadjusted hypothesis as defined in equation (8). At the same time we note that using the hypothesis as defined in equation (9) would leave the conclusions essentially the same. Having chosen the hypothesis to test against, we look at the average value of Ψ we have obtained in the four studied sets, and compare it with the values we would expect for random choice ($1/\Omega = 0.25$). For Warsaw Stock Exchange we have obtained average Ψ of 0.292 and 0.351 for daily and intraday logarithmic returns respectively. These are averaged over all studied stocks, and also all studied values of the window length μ . As mentioned above, an average over all values of μ is informative as there is practically no dispersion in these results, with standard deviations of 0.0345 and 0.0074 respectively. For New York

Stock Exchange we have obtained average Ψ of 0.27 and 0.281 for daily and intraday logarithmic returns respectively. In these there also is practically no dispersion, with standard deviations of 0.0142 and 0.0105 respectively. Additionally, for Warsaw Stock Exchange only an average of 3 out of 360 stocks show average Ψ below the 0.25 threshold (between 2 and 5 for various values of μ) for the daily log returns, and zero out of 707 stocks for the intraday log returns. For New York Stock Exchange only an average of 2 out of 98 stocks show average Ψ below the 0.25 threshold (between 1 and 4 for various values of μ) for the daily log returns, and zero out of 98 stocks for the intraday log returns. From these results we get a clear support for the hypothesis of the application of the MEPP to stock returns as defined in equation (8). We get only a handful of stocks for which the estimated Ψ is under the threshold. These results are particularly strong for intraday logarithmic returns, which is not surprising as these have been known to be more predictable. It is important that the result is robust with respect to the window length μ , as this makes the principle potentially useful in practice. We understand this phenomenon as stemming from the efficiency of the market being a product of *sub-efficiencies* among groups of traders with different investment horizons.

Second, we take a look at whether the Maximum Entropy Production Principle allows us to guess future price changes more accurately for these price formation processes which are more predictable in the sense of information theory (entropy rate for the whole studied series). In previous studies [25], [24], standard ways of predicting future price changes have been shown not to use the predictability (structural complexity) of the time series describing logarithmic returns. The two mentioned papers showed no effect of the entropy of the underlying price formation process on the performance of trading algorithms based on genetic programming and mean reversal strategy. We would like to find an approach which is capable of using this information. And, as can be seen in the figures described above, there is a strong negative correlation between Ψ and entropy rate for the studied time series describing logarithmic returns. In other words, there is a strong positive correlation between the explanatory power of the Maximum Entropy Production Principle and the predictability of the underlying price formation processes. The Pearson's correlation coefficient between the Ψ averaged over all studied values of μ and the Lempel–Ziv entropy rate for all studied stocks are equal to -0.927 and -0.931 for Warsaw Stock Exchange daily and intraday data, together with -0.752 and -0.765 for New York Stock Exchange daily and intraday data respectively. Thus we are able to conclude that the use of the MEPP as a tool for predicting future price changes of stocks uses the predictability of the price formation processes, contrary to popular methods used in trading algorithms, such as the mean reversal strategy.

Third, we take a look at the distributions of the estimates of the percentage of times the Maximum Entropy Production Principle correctly predicts the next price move (Ψ), averaged over all values of the window length μ , for logarithmic returns within the four studied sets. As can be seen in figure 5, almost the whole mass of the distribution is above 0.25 for all four studied sets of stocks, thus supporting our hypothesis of the Maximum Entropy Production Principle within financial markets. Additionally, we observe that for the large values of




Fig. 1. Maximum Entropy Production Principle (MEPP) testing results for Warsaw Stock Exchange daily logarithmic returns (discretised into four quartiles) database of 360 stocks: (a) Estimate of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all studied stocks vs the window length parameter (μ), together with error bars of one standard deviation, and two lines denoting values for random guessing the next price move (below unadjusted, above adjusted); (b) Histogram of Ψ averaged over all values of μ together with vertical lines for unadjusted (left) and adjusted (right) values for random guessing; (c) Scatterplot of Ψ averaged over all values of μ vs Lempel–Ziv entropy rate for the whole time series, with linear regression fitted. The values of Ψ are strongly positively correlated with the predictability of the price formation processes, but the results do not depend on the size of the studied time series.




Fig. 2. Maximum Entropy Production Principle (MEPP) testing results for Warsaw Stock Exchange intraday logarithmic returns (discretised into four quartiles) database of 707 stocks: (a) Estimate of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all studied stocks vs the window length parameter (μ), together with error bars of one standard deviation, and two lines denoting values for random guessing the next price move (below unadjusted, above adjusted); (b) Histogram of Ψ averaged over all values of μ together with vertical lines for unadjusted (left) and adjusted (right) values for random guessing; (c) Scatterplot of Ψ averaged over all values of μ vs Lempel–Ziv entropy rate for the whole time series, with linear regression fitted. The values of Ψ are strongly positively correlated with the predictability of the price formation processes, but the results do not depend on the size of the studied time series.

Ψ these distributions are characterised by fat tails, roughly approximating log-normal distribution. This is interesting and important, as the analysts and traders may concentrate on the handful of stocks with the largest predictability as defined by their Lempel–Ziv entropy rate, which are then likely to give the best results when applying the Maximum Entropy Production Principle for prediction purposes.

Finally, we also note that this method, similarly to earlier studies of structural complexity of stock returns, uses volatility levels to distinguish patterns. So while this method should give good results if Ω is larger than four, it will not give very good results when we apply it to stock returns discretised with $\Omega = 2$, that is only distinguishing positive and negative logarithmic returns. Within such framework there is not enough patterns within the data (which are mostly based on volatility clustering) for the predictability to be useful for analysis or predictions based on information theory.

V. CONCLUSIONS

In this paper we have introduced the concept of the Maximum Entropy Production Principle as applied to financial markets, and logarithmic stock returns in particular. We propose that such a principle governs the financial markets, but that it is bounded by a large number of unknown constraints, so that in practice it can only be observed as a statistical tendency, and not an immutable law. Thus our testing it, on the basis of

databases containing historical daily and intraday logarithmic returns for stock exchanges in Warsaw and New York, can only partially confirm the hypothesis of such principle. As is often the case with complex systems, a full proof of such a simple principle will be very hard to obtain. Nonetheless, the results confirmed our hypothesis and we have found that the application of such principle to the prediction of future stock returns gives consistently better returns than would be obtained by random chance. Additionally, the performance of such prediction, based on the Maximum Entropy Production Principle, is strongly dependent on the structural complexity, or model-free predictability, of the underlying price formation processes. This is important, as traditional trading algorithms do not possess such a characteristic. Together with the fact that the distribution of the performance of the Maximum Entropy Production Principle has fat tails, this allows analysts to pick stocks for which the underlying price formation processes are the most predictable, and concentrate their analysis or trading on these, knowing that it is more likely these will produce better results with the use of the analysed principle. In particular this means using intraday stock returns, as these are more predictable than their daily counterparts. Future studies should try to turn this principle into a viable trading strategy, and test it on other financial markets and instruments, such as currency exchange rates. Further studies should also try to develop the methodology in various ways, for example introducing the usage of permutation entropy for the purposes




Fig. 3. Maximum Entropy Production Principle (MEPP) testing results for New York Stock Exchange daily logarithmic returns (discretised into four quartiles) database of 98 stocks: (a) Estimate of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all studied stocks vs the window length parameter (μ), together with error bars of one standard deviation, and two lines denoting values for random guessing the next price move (below unadjusted, above adjusted); (b) Histogram of Ψ averaged over all values of μ together with vertical lines for unadjusted (left) and adjusted (right) values for random guessing; (c) Scatterplot of Ψ averaged over all values of μ vs Lempel–Ziv entropy rate for the whole time series, with linear regression fitted. The values of Ψ are strongly positively correlated with the predictability of the price formation processes, but the results do not depend on the size of the studied time series.




Fig. 4. Maximum Entropy Production Principle (MEPP) testing results for New York Stock Exchange intraday logarithmic returns (discretised into four quartiles) database of 98 stocks: (a) Estimate of the percentage of times the MEPP correctly predicts the next price move (Ψ) averaged over all studied stocks vs the window length parameter (μ), together with error bars of one standard deviation, and two lines denoting values for random guessing the next price move (below unadjusted, above adjusted); (b) Histogram of Ψ averaged over all values of μ together with vertical lines for unadjusted (left) and adjusted (right) values for random guessing; (c) Scatterplot of Ψ averaged over all values of μ vs Lempel–Ziv entropy rate for the whole time series, with linear regression fitted. The values of Ψ are strongly positively correlated with the predictability of the price formation processes, but the results do not depend on the size of the studied time series.

of such an analysis, or apply the same principle to time series describing logarithmic returns with normalised volatility. Studies should also be performed to establish whether it is possible to find the constraints to this principle on the markets, as mentioned above. It may be interesting to also compare the performance of MEPP with a simple method based on moving window averages for returns.

REFERENCES

- [1] R. Clausius, “Über eine veränderte Form des zweiten Hauptsatzes der mechanischen Wärmetheoriein,” *Annalen der Physik und Chemie*, vol. 93, no. 12, pp. 481–506, 1854.
- [2] L. Boltzmann, “Über die Mechanische Bedeutung des Zweiten Hauptsatzes der Wärmetheorie,” *Wiener Berichte*, vol. 53, pp. 195–220, 1866.
- [3] J. Gibbs, *Elementary Principles in Statistical Mechanics – Developed with Especial Reference to the Rational Foundation of Thermodynamics*. New York: C. Scribner’s Sons, 1902.
- [4] L. Onsager, “Reciprocal relations in irreversible processes,” *Physical Review*, vol. 37, no. 4, pp. 405–426, 1931.
- [5] J. Dugdale, *Entropy and its Physical Meaning*. London: Taylor and Francis, 1996.
- [6] T. Cover and J. Thomas, *Elements of Information Theory*. New York: John Wiley & Sons, 1991.
- [7] I. Prigogine, “Etude thermodynamique des phenomenes irreversibles,” Ph.D. dissertation, Universite libre de Bruxelles, 1947.
- [8] J. Keizer, *Statistical Thermodynamics of Nonequilibrium Processes*. New York: Springer-Verlag, 1987.
- [9] I. Prigogine, “Sur les fluctuations de l’equilibre chimique,” *Physica*, vol. 16, pp. 134–136, 1950.
- [10] F. Henin and I. Prigogine, “Entropy, dynamics, and molecular chaos,” *Proceedings of the National Academy of Sciences*, vol. 71, no. 7, pp. 2618–2622, 1974.
- [11] E. Gunzig, J. Gehehiau, and I. Prigogine, “Entropy and cosmology,” *Nature*, vol. 330, pp. 621–624, 1987.
- [12] L. Martyushev and V. Seleznev, “Maximum entropy production principle in physics, chemistry and biology,” *Physical Reports*, vol. 426, pp. 1–45, 2006.
- [13] H. Ozawa, A. Ohmura, R. Lorenz, and T. Pujol, “The second law of thermodynamics and the global climate system: A review of the maximum entropy production principle,” *Review of Geophysics*, vol. 41, pp. 1018–1042, 2003.
- [14] L. Martyushev, V. Seleznev, and I. Kuznetsova, “Application of the principle of maximum entropy production to the analysis of the morphological stability of a growing crystal,” *Journal of Experimental and Theoretical Physics*, vol. 91, pp. 132–143, 2000.
- [15] G. Beretta, “Maximum entropy production rate in quantum thermodynamics,” *Journal of Physics: Conference Series*, vol. 237, p. 012004, 2010.
- [16] S. Gheorghiu-Svirschevski, “Nonlinear quantum evolution with maximal entropy production,” *Physical Review A*, vol. 63, p. 022105, 2001.
- [17] L. Ingber, “Statistical mechanics of nonlinear nonequilibrium financial markets,” *Mathematical Modelling*, vol. 5, no. 6, pp. 343–361, 1984.
- [18] J. Voit, *The Statistical Mechanics of Financial Markets*. London: Springer, 2001.
- [19] R. Officer, “The Distribution of Stock Returns,” *Journal of the American Statistical Association*, vol. 340, no. 67, pp. 807–812, 1972.

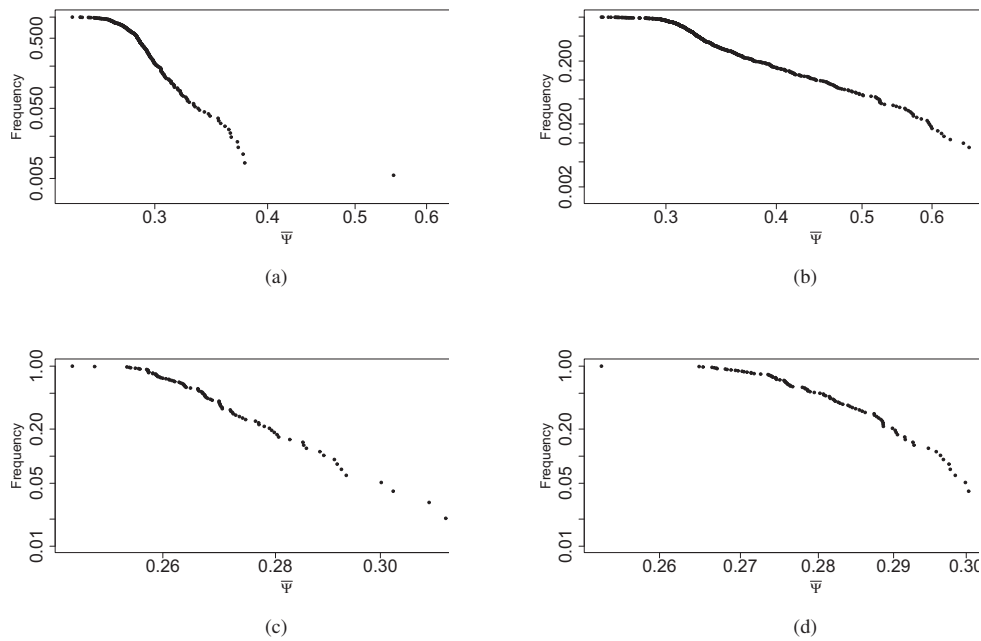


Fig. 5. The distributions of the estimates of the percentage of times the Maximum Entropy Production Principle (MEPP) correctly predicts the next price move (Ψ), averaged over all values of the window length μ , for logarithmic returns (discretised into four quartiles) for: (a) Warsaw Stock Exchange daily database of 360 stocks; (b) Warsaw Stock Exchange intraday database of 707 stocks; (c) New York Stock Exchange daily database of 98 stocks; (d) and New York Stock Exchange intraday database of 98 stocks. The distributions show signs of fat tails, particularly seen for the high values of Ψ . It is clear that almost the whole mass of the distributions is above 0.25, or the estimate for random guessing the next price change, thus supporting our hypothesis of the MEPP within financial markets.

- [20] C. Starica and C. Granger, "Nonstationarities in stock returns," *The Review of Economics and Statistics*, vol. 87, no. 3, pp. 503–522, 2005.
- [21] Y. Sinai, "On the Notion of Entropy of a Dynamical System," *Doklady Akademii Nauk SSSR*, vol. 124, pp. 768–771, 1959.
- [22] A. N. Kolmogorov, "Entropy per unit time as a metric invariant of automorphism," *Doklady Akademii Nauk SSSR*, vol. 124, pp. 754–755, 1959.
- [23] C. E. Shannon, "A Mathematical Theory of Communication," *Bell Systems Technical Journal*, vol. 27, pp. 379–423, 623–656, July, October 1948.
- [24] P. Fiedor, "Frequency Effects on Predictability of Stock Returns," in *Proceedings of the IEEE Computational Intelligence for Financial Engineering & Economics 2014*, A. Serguieva, D. Maringer, V. Palade, and R. J. Almeida, Eds. London: IEEE, 2014, pp. 247–254.
- [25] N. Navet and S.-H. Chen, "On Predictability and Profitability: Would GP Induced Trading Rules be Sensitive to the Observed Entropy of Time Series?" in *Natural Computing in Computational Finance*, ser. Studies in Computational Intelligence, T. Brabazon and M. O'Neill, Eds. Berlin Heidelberg: Springer, 2008, vol. 100.
- [26] Y. Gao, I. Kontoyiannis, and E. Bienenstock, "From the entropy to the statistical structure of spike trains," in *IEEE International Symposium on Information Theory*, Seattle, 2006, pp. 645–649.
- [27] M. Farah, M. Noordewier, S. Savari, L. Shepp, A. Wyner, and J. Ziv, "On the entropy of DNA: algorithms and measurements based on memory and rapid convergence," in *SODA'95: Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. San Francisco: Society of Industrial and Applied Mathematics, 1995, pp. 48–57.
- [28] I. Kontoyiannis, P. Algoet, Y. Suhov, and A. Wyner, "Nonparametric entropy estimation for stationary processes and random fields, with applications to English text," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 1319–1327, 1998.
- [29] A. Lempel and J. Ziv, "A Universal Algorithm for Sequential Data Compression," *IEEE Transactions on Information Theory*, vol. 23, no. 3, pp. 337–343, 1977.
- [30] F. Willems, Y. Shtarkov, and T. Tjalkens, "The Context-Tree Weighting Method: Basic Properties," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 653–664, 1995.
- [31] M. Kennel, J. Shlens, H. Abarbanel, and E. Chichilnisky, "Estimating entropy rates with Bayesian confidence intervals," *Neural Computation*, vol. 17, no. 7, pp. 1531–1576, 2005.
- [32] G. Louchard and W. Szpankowski, "On the average redundancy rate of the Lempel-Ziv code," *IEEE Transactions on Information Theory*, vol. 43, no. 1, pp. 2–8, 1997.
- [33] F. Leonardi, "Some upper bounds for the rate of convergence of penalized likelihood context tree estimators," *Brazilian Journal of Probability and Statistics*, vol. 24, no. 2, pp. 321–336, 2010.
- [34] A. Doganaksoy and F. Gologlu, "On Lempel-Ziv Complexity of Sequences," *Lecture Notes in Computer Science*, vol. 4086, pp. 180–189, 2006.
- [35] I. Kontoyiannis, "Asymptotically Optimal Lossy Lempel-Ziv Coding," in *IEEE International Symposium on Information Theory*. Cambridge: MIT, 1998.
- [36] Y. Gao, I. Kontoyiannis, and E. Bienenstock, "Estimating the entropy of binary time series: Methodology, some theory and a simulation study," *Entropy*, vol. 10, no. 2, pp. 71–99, 2008.
- [37] F. Willems, Y. Shtarkov, and T. Tjalkens, "Context weighting for general finite-context sources," *IEEE Transactions on Information Theory*, vol. 42, no. 5, pp. 1514–1520, 1996.
- [38] F. Willems, "The context-tree weighting method: Extensions," *IEEE Transactions on Information Theory*, vol. 44, no. 2, pp. 792–798, 1998.
- [39] R. Steuer, L. Molgedey, W. Ebeling, and M. Jiménez-Montaño, "Entropy and Optimal Partition for Data Analysis," *European Physical Journal B*, vol. 19, no. 2, pp. 265–269, 2001.