ERGODICITY IN FINANCE

***Remarks on a study by Iván Bélyácz*[1]**

Edina Berlinger

ABSTRACT

In this paper, I provide a number of specific financial examples of when the ergodic hypothesis appears well-founded, and some when it does not. My conjecture is that tasks that demand only the prediction of risks (for example, derivative pricing or risk management) can be accomplished by studying historical data. If, however, an estimate of the expected returns is also needed (for example, in forecasting equity premiums or performance measurement), then conclusions drawn based on historical data become unreliable. Even a long time horizon does not necessarily increase stability, since often we cannot say more about processes even if we carry out our investigations on a historical scale; in other instances, however, we can indeed count on history repeating itself.[2]

JEL codes: G1, G2

Keywords: ergodicity, option pricing, bank risk management, equity premium puzzle, log-optimal portfolios, performance evaluation

1. INTRODUCTION

This paper was inspired by the study of ergodicity by *Iván Bélyácz* (2017) and the lively professional debate that followed. My goal is to present a number of specific financial problems in which issues of ergodicity may come to the forefront.

Bélyácz (2017) highlighted that (neo)classical economics, explicitly or not, takes as its starting point the ergodicity of economic processes. This essentially means that its adherents believe the probability distribution can be known through studying the past, so that economic actors must confront risk, rather than uncertainty, in the long run – risk that can be measured, priced, traded and hedged, and so managed well by quantitative methods. Bélyácz (2017, p. 51) unequivocally

1 Iván Bélyácz (2017): The debated role of ergodicity in (financial) economics. *Economy & Finance*, 4 (1), pp. 4–57.

2 The research was carried out with the support of the János Bolyai scholarship programme of the Hungarian Academy of Sciences (MTA).

disapproves of this approach: "We thus regard assumptio*ns about the relevance of ergodicity as false because we have doubts over the timelessness and immutability of economic processes.*"

Numerous definitions of ergodicity exist. As phrased by *Simonovits* (2017, p. 124), *"ergodicity is the stochastic generalization of the stability of a time-invariant deterministic dynamical system,"* meaning that a process is ergodic when the cross-sectional distribution in time converges towards a stable central limit. According to *Mellár* (2017, p. 98): *"Very simply put, ergodicity means that the evolution of economic processes over time is itself a periodically recurring, regular process, the characteristics of which can be identified and comprehended by mathematical and statistical means." Harcsa* (2017, p. 61), similarly to Bélyácz (2017), sees the essence of the ergodicity debate in the dichotomy of risk versus uncertainty. *Medvegyev* (2017, p. 110) states that a stochastic process is ergodic if the limit value in the infinite of the time average along individual trajectories exists. *Peters* (2011), and Peters and *Klein* (2013), use a stricter definition, whereby ergodicity requires equality of the averages calculated from the longitudinal (time) and cross-sectional (space) distributions. *Horst* (2008, p. 1) begins his article on the role of ergodicity in economics in the New Palgrave Dictionary thus: *"A stochastic system is called ergodic if it tends in probability to a limiting form that is independent of the initial conditions. Breakdown of ergodicity gives rise to path dependence. When path dependence occurs, 'history matters'."* Horst (2008) then goes on to present a number of economic models in which ergodicity does not hold true, for example in the case of endogenous preference formation, stochastic strategy revision in dynamic population games, or other dynamic social interactions.

The answer to the question of "ergodicity or nonergodicity" obviously depends on what kind of process we are considering and on what time horizon, as well as on the degree of accuracy we demand. Below I discuss these questions in relation to a number of specific financial applications.

2. STOCHASTIC PROCESSES AND ASSET PRICING

Let us assume that our boss gives us the task of forecasting some important financial indicator (exchange rate, yield, income etc.) for a specific point of time in future. In this case, we can essentially resort to one of two methods: on the one hand, in accordance with the principle of rational expectations, by processing all the available information and taking all complex interconnections into account, we can use the best theoretical model to determine the distribution of the random (cross-sectional) variable in space, allowing us to calculate the mean from this;

on the other hand, if a historical time series for the variable is given, then we can calculate the (longitudinal) average of the realizations in time. If a process is (in the strict sense) ergodic, then the two averages will be (asymptotically) equal. For this reason, if we can be sure that the examined process is ergodic, then we can complete the task assigned to us very simply and quickly since we only need to calculate the average of the observations, and the longer the available time series, the more accurate will be our forecast. If the process is nonergodic, however, then the time average can be entirely misleading.

Let us look at one of the simplest and most popular stochastic processes, geometric Brownian motion (GBM). If the probability variable x follows GBM, then for an infinitely short time dt, the change in x can be broken down into a trend and a random effect:

$$dx = \mu x dt + \sigma x dW, \qquad (1)$$

where μ is the trend variable, σ the volatility and dW the so-called Wiener component, which is an independent random variable with mean and standard variation of zero and $\sqrt{dt}$, respectively (*Hull*, 2015).

It is easy to see that the GBM process is nonergodic, since equality of the time and space averages is not fulfilled (Peters, 2011; Peters–Klein, 2013). The cross-sectional average of the exponential rate of growth is μ, while the limit value of the longitudinally averaged exponential rate of growth to infinity is only $\mu - \frac{\sigma^2}{2}$.

The distinction is vital, for example, if the trend μ = +5%, since then with 45% volatility the average of the realized log returns will be –5%. If an investor holds all their assets in this instrument, then despite the positive trend the value of their investment will approach zero over an infinite time horizon with probability of 1, since the expected value of the time-averaged exponential yield is negative. Due to the high volatility, in a given bad year the asset value may decrease to such an extent that even the positive expected return will be unable to compensate for this. Figure 1 shows a number of possible trajectories of this process and the certain –5% trajectory on a time horizon of 250 years with an initial value.

Figure 1
Some trajectories of geometric Brownian motion, μ = +5% and σ = 45%

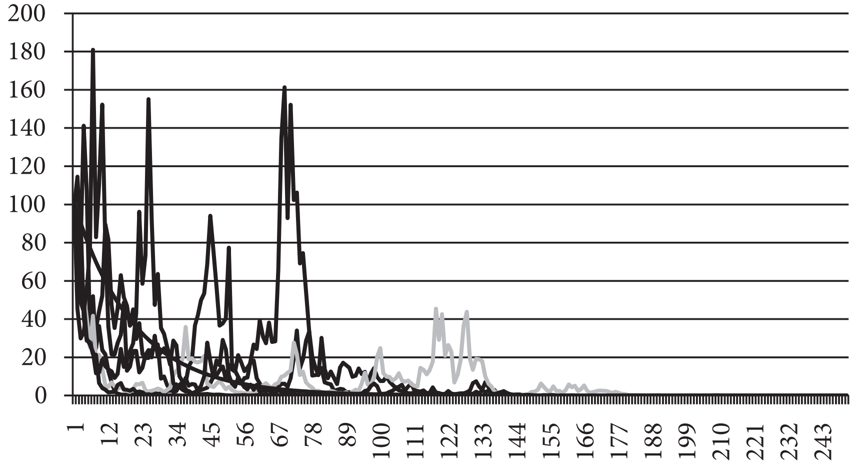

Source: own simulation

The GBM process is a good illustration of how little information historical observations carry in the case of nonergodic processes. For the first 70–80 years, volatility dominates, and in the case of most trajectories the declining path in the background is not apparent at all, so that we do not even notice that the time average overestimates the space average, i.e. the actual expected value.

In the GBM model, it is clear that if we regard the log return belonging to the Δt period, and not the price, as the probability variable, then this is at once both stationary and ergodic, since it follows a normal distribution with an expected value of $\left(\mu - \frac{\sigma^2}{2}\right)\Delta t$ and a standard deviation of $\sigma\sqrt{\Delta t}$, so that the time and space averages of the log returns converge and the cross-sectional distribution can be known based on historical realizations.

The renowned Black–Scholes option pricing formula is also based on the assumption that the price of the underlying asset follows GBM (*Black–Scholes*, 1973). To determine the value of the option, besides the spot price (x) of the underlying asset, we need to know "only" the future constant volatility (*If*), but not the price trend (μ). It is an empirical fact, however, that volatility is not constant, but prone to volatility clustering, meaning that quieter periods tend to alternate with more volatile ones. For this reason, in more advanced option pricing models, not only the price but also the volatility (and perhaps the interest rate) is stochastic; which then raises the question of how we should model the combined development of these factors, and whether the parameters of this multidimensional process – for example the volatility of volatility or cross-correlations – can be known based on historical realizations. Hull and *White* (1987) demonstrate that if the price and the square of the volatility each follows uncorrelated GBM, then this provides an explanation for the phenomenon of the "volatility smile" experienced in practice; in other words, that the single-factor Black-Scholes formula typically overprices ATM (at-the-money) options, and underprices deep ITM (in-the-money) and OTM (out-of-the-money) options. Using numerical methods, they also examine the case when the correlation coefficient between the two processes is constant, but not zero. The latest research goes even further, regarding the correlation, too, as a stochastic variable (see *Misik*, 2015). Of course this logic can be extended to infinity, with the building of increasingly sophisticated models, for which even the volatility of volatility of the correlations between correlations must be estimated. It can be seen that financial experts do not give up easily, and do not lose heart when they see volatility and correlation smiles; on the contrary, they redouble their efforts to wipe away those smiles and bring their increasingly complicated models closer to reality. All the while they adhere to the ergodic hypothesis, whereby the distributions are knowable, the models can be calibrated, and – with enough patience – the uncertainties can eventually be tamed into risks and channelled through customary risk management processes. One cannot fail to appreciate their audacity and virtuosity.

3. RISK MANAGEMENT

In the context of risk management in a bank, essentially it is necessary to pull off the feat, on a daily basis, of determining the composition of the current portfolio, identifying the risk factors that impact the value of the portfolio, modelling the stochastic behaviour of these factors, and – last but not least – generating the distribution of the entire portfolio's return for the next ten days. This distribution will be the basis for the risk management process, as it is only with knowledge of this that we can define the appropriate risk measure, which will eventually help us to determine how big the capital buffer and other reserves should be to cover potential losses with a high degree of probability. If the capital is inadequate, then either the portfolio must be restructured or additional capital must be brought in. Managing liquidity risk is also a part of the operation, taking into account prevailing liquidity trends. Forecasts must be constantly compared with actual data, and, if the discrepancy is too great, then the models must be redesigned. Risk appetite must be determined based on the organization's strategy, and risk limits distributed among the individual organizational units. The various units and their employees must be incentivized to follow the interests of the owners, clients and society (?), not their own. Internal settlement prices must be continuously recalculated, and profits or losses shared between stakeholders. Products must be designed in accordance with customers' demands, and these products priced. Markets must be acquired, competitors outrun, catastrophic situations survived, and snowballing regulatory expectations met. Thought must be given to what is meant by FinTech, to the nature of the bank of the future, and to the need to innovate. And the list goes on and on.

What can be said about the totality of tasks, in themselves and in combination, is that from both an ontological and cognitive perspective they are entirely hopeless. Disparate preferences render collective decision-making impossible (*Arrow*, 1963), multidimensional stochastic processes unknowable, and innovative leaps unforeseeable; there is no system of incentives that cannot be manipulated; risk appetite is hard to define (*Lamanda–Tamásné*, 2015), there is no rule of sharing appropriate to every reasonable axiom (*Csóka–Pintér*, 2016), unmanageable uncertainties are legion; and, furthermore, social context and historical determinism must also be taken into consideration. If the pricing model is too simple, then it cannot handle the complexity and will misprice; if too complicated, then it is not transparent and the suspicion arises that it is mere trickery. Nothing seems easier than to criticize the everyday practice of risk management. And yet bank regulators are even easier to criticize, since in reality they invite ridicule in their endeavours to keep track of the risks in the entire financial system, and in their attempts to intervene at the right moment and in the right manner in support of goals that are hard to put into words and often contradictory. It is particularly laughable when they end up struggling with the unintended consequences of their own interventions, or when they attempt to find excuses for why they did not foresee the crisis, why they did not uncover the fraud sooner, and so on and so forth.

Once we have finished laughing and feeling contempt for them, then we might offer some suggestions for what could be done instead of the present practice. What should we teach in universities? Theorems of impossibility and unknowability – or can we allow ourselves to offer, with eyes downcast, a few risk management models, which subsequently may or may not work? If we seriously believe that economic and financial processes are nonergodic, then most of the currently applied risk management techniques can be thrown out of the window. But what will take their place? "Expert" forecasts, birds in flight, coffee-grounds?

4. THE EQUITY PREMIUM PUZZLE

Mehra and *Prescott* (1985) were the first to raise the issue of the equity risk premium puzzle, which has occupied financial economists ever since. The essence of the puzzle is that over a relatively lengthy period on the U.S. equity market between 1889 and 1978, i.e. over a relatively long period in the leading economy of the world, the average risk premium was 6.18%, which is significantly higher than what we can explain based on standard microeconomic models. Although it is true that the standard deviation of the risk premium is also significant (16.67%), empirical experiments tell us that people are not really quite so risk-averse and that the majority would be content with appreciably less compensation. *Mankiw* and *Zeldes* (1991) showed that within the framework of a traditional CRRA (constant relative risk aversion) utility function, the risk premium and standard deviation observable on the market would be consistent with the assumption of a class of risk-averse investors for whom the two payoffs A and B below would be entirely equal in value:

A: 50% probability of a 50,000 dollar gain
50% probability of a 100,000 dollar gain

B: 100% probability of a 51,200 dollar gain

Since it is clear that investors are not that averse to risk, another explanation must be sought. When reality clashes with theory, the tension can be resolved in two ways: either we say that reality is not what it appears to be, or we replace the theory. Accordingly, some explanations will note that the risk premium was not so high in other countries and in other periods, or that returns appear so high only because of the survivorship bias; it can also be argued that an observation period of about 90 years is precious little to be able to determine from the significant volatility how great the long-term average return (premium) might be. The other school of thought, meanwhile, eagerly works to exchange the utility theory for the prospect theory, to incorporate other behavioural effects (e.g. narrow framing) into the decision-making model, or to take factors into consideration that are missing from the standard microeconomic models; for example, liquidity risk, the impact of the tax system, information asymmetry and transaction costs. Interest in the topic is well illustrated by more than 6,500 references to the original article according to Google Scholar, but unfortunately the "eureka moment" remains elusive for the time being and the matter has thus still not been put to rest.

And yet it is important to know whether an exceptionally high risk premium is a temporary phenomenon attributable to an unusual period, i.e. merely a fortuitous accident, or if we can expect the risk/return conditions observed in the past to remain in future. The fundamental parameter of every pricing formula is the expected (probable) risk premium. If we believe that the annual 6.18% surplus will come to us in the next 90 years as well (even if we don't actually understand why) – or if, in other words, we believe in ergodicity – then this will have numerous deep-reaching consequences: we will invest the greater part of our pension savings in risky assets; risk management will create greater value, and so financial experts will make a lot of money; company managers will concentrate stubbornly on short-term profit goals, since robust discounts will mean no one cares about the long-term impacts; and, likewise, we can also count on the attraction of long-term, responsible policies remaining low, and on the costs of crises being extremely high. It may be that the survival of the planet, and of humanity upon it, depends on what we think about the ergodicity of financial markets.

5. PORTFOLIO MANAGEMENT AND PERFORMANCE EVALUATION

A substantially less important question, which is nevertheless of interest to many, is how to get rich on the stock market. What is the optimal portfolio management strategy in the long term? Let us suppose that we can trade in a finite number of risky assets, and that we can restructure our portfolio on a daily basis based on historical observations. If the markets are efficient in at least a weak sense, then returns are independent in time and have no memory, and studying historical data offers us little to go by. If, in addition, asset returns derive from an identical distribution (i.e. are stationary), then there is no point in portfolio restructuring; indeed, this will even destroy value in the long term.

If, on the other hand, market returns are not independent in time, then it is possible that recurring patterns exist that a skilful trader can exploit. *Algoet* and *Cover* (1988) showed that if yields are ergodic (and stationary), then on an infinite time horizon a so-called log-optimal portfolio is the best choice (when the goal is to maximize the average log return of the portfolio to infinity). However, to determine this exactly, we would need to know the return-generating process itself. And unfortunately, the problem is precisely that, in general, we do not have this information at our disposal. It can be proven, however, that there exist so-called universally consistent strategies which asymptotically approach the average growth rate of the log-optimal portfolio in infinity (Algoet and Cover, 1988). The basic idea of these asymptotically log-optimal strategies is that we look for price patterns in the past similar to that most recently observed, and then examine what portfolio weighting would have been optimal immediately after the patterns emerged, restructuring our portfolio in the present accordingly. We then carry out this restructuring on a daily basis. Similarities in patterns can be defined in a number of ways, so that several kinds of log-optimal strategy can exist (*Györfi* et al., 2006). Empirical investigations reveal that log-optimal strategies work surprisingly well even on finite time horizons (of 10–20 years) (*Ormos* et al., 2009), which – albeit with reservations – we may interpret as a specific proof of the ergodicity of returns.

What remains a matter for debate, however, is whether active portfolio management is truly able to create value compared to the simple "buy the index and hold" strategy. We would think that measuring the performance of fund managers cannot be too difficult a task, since the performance is one-dimensional, measurable in money, and there is a great deal of reliable data at our disposal. Despite this, correct measurement of performance is almost impossible to accomplish. Based on *Bodie* et al. (2005), let us suppose that a fund's returns follow a stationary process, and that the fund manager really knows what they are doing since they are able to sustain a stable monthly 0.2% surplus yield, meaning 2%–3% on an annual level, which is an outstanding result. The fund's beta is 1.2, the individual standard deviation of monthly yields is 2%, the deviation in the monthly yield on the market portfolio is 6.5%, and the correlation is 0.97. If we carry out a hypothesis test using the usual statistical methods, then a 95% significance level will require an observation period of 384 months, i.e. 32 years. Someone might recommend that we use daily or even shorter returns for the measurement, but this makes no sense if the fund manager is working with an investment horizon of several months. In other words, unfortunately, the greater part of the fund manager's professional career will pass before they manage to figure out that they were not simply lucky, but genuinely possess exceptional abilities. Moreover, returns are not remotely stationary in practice. Indeed, paradoxically, the better the fund manager's timing skills, the less stationary will be the fund's returns. A skilful fund manager increases their exposure prior to an upswing, and reduces it before a recession, so that the standard deviation of the fund's yield will continually change parallel with the market trend. Under such circumstances, it is not even sure we will complete our measurement before the fund manager retires.

Without doubt there are also instances when, without any special need for measurements, we are able with our naked eye to determine with confidence that the performance of a given investment fund is outstanding. Figure 2 shows the relative performance of the special Fairfield Sentry investment fund.

Figure 2
Relative performance of the Fairfield Sentry investment fund, 1990–2007

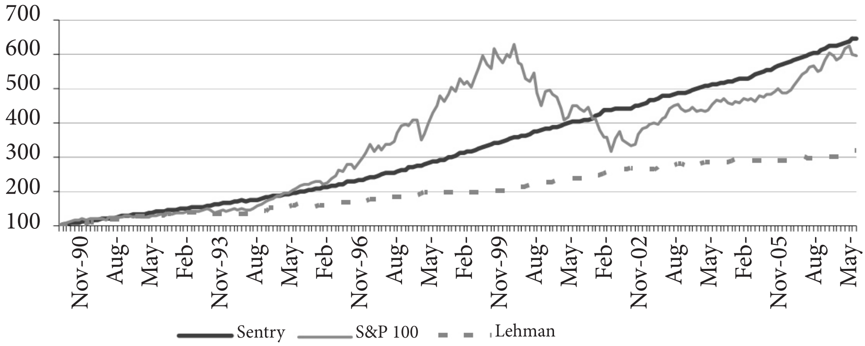

Source: Bloomberg

In the diagram we can see how our wealth would have changed over time if we had invested 100 dollars in November 1990 in the Fairfield Sentry fund, the Standard & Poor's 100 passive equity fund, or the Lehman bond fund. Fairfield Sentry generated approximately as great a yield as the equity index, with approximately the same volatility as the bond index. Investors in Fairfield Sentry would probably have been glad to acknowledge this diagram and to congratulate themselves on an excellent investment, without giving much thought to the ontological and cognitive aspects of ergodicity; but they would most certainly have been surprised when, at the end of 2008, the FBI began investigating the fund manager *Bernie Madoff*, exposing the biggest Ponzi scheme in world history, as a consequence of which they were suddenly obliged to register enormous losses. Could they have deduced this abrupt loss based on historical data, as part of the stochastic process itself (ergodicity), or was it an unforeseeable, external impact independent of all else (nonergodicity)? In the same way, should those who took out loans in Swiss francs have reckoned on unfavourable leaps in the exchange rate, even if they enjoyed lower borrowing costs without interruption in the preceding years? I think they should have: what seems too good to be true probably isn't true. Perhaps unlucky investors will learn from their mistakes and switch from adaptive expectations to rational expectations, but it is also possible that they have understood nothing, and that they can scarcely wait to rush headlong into another crackpot speculative bubble in order to regain their previous losses as quickly as possible.

6. LONG-TERM (IN)STABILITY

The time horizon of risk management in a bank ranges from 10 days (market risks) to one year (credit risks), the duration of derivatives is 1–2 years, the time horizon of active portfolio management and performance evaluation is sometimes 10 years, and that of the equity premium puzzle is approximately a century. But what can we say about economic processes over a much longer perspective in time?

Andrew Haldane, chief economist at the Bank of England, spoke in a lecture in 2015 about how nominal interest rates in developed economies have remained stuck at low levels unprecedented in human history (in Japan since 1995, and in the U.S., U.K. and the eurozone since 2009). Among the causes, he mentions the global low rate of growth, excess savings in the East, deficient investment in the West, worsening demographic trends and rising inequality. Since he does not expect these fundamentals to change conspicuously in future, Haldane breaks completely from historical trends (or ergodicity) in his forecast, predicting that interest rates will remain at low levels for a long time to come (Haldane, 2015).

Figure 3
Short and long-term interest rates in the past 5,000 years

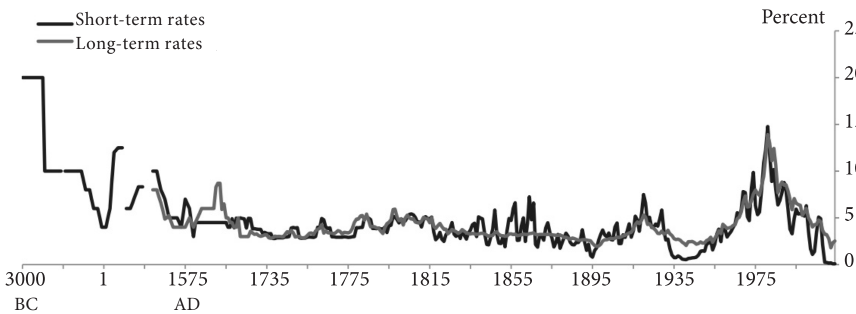

Note: The diagram shows the lowest documented interest rates before the 18th century (in the Babylonian, Greek, Roman and Byzantine empires, the Low Countries and Italy), and in the world's leading economies after the 18th century (U.K. and U.S.).

Source: Haldane (2015)

In contrast, *Piketty* (2015) believes that certain correlations, for example that the return on capital (*r*) is higher than the rate of economic growth (*g*), have remained stable for several centuries, and the capitalist economy, propelled by conformity to some internal rule, will sooner or later return to this path.

What economists broadly agree upon is that long-term economic growth is driven by technological development. For this reason, a fundamental question is what we think about the nature of the innovative process. Is it entirely a law unto itself and unpredictable (uncertainty), or does it obey discernible statistical regularities (risk)? Although the former approach is the more popular, many signs also point to the latter. For example, certain discoveries and inventions often come about simultaneously in time, but independently of one another; and certain processes fit precisely into a determined path (e.g. Moore's Law).[3] Meanwhile, *Oliveira* and *Barabási* [2005] reached the surprising finding in their empirical study that the dynamics of *Darwin*'s and *Einstein*'s written correspondence show the same power law distribution as that of internet e-mail communication of the present day (with only the exponent differing somewhat). In other words, it seems that technological advancement does not necessarily change the substantive features of human activity.

Others, however, happen to think that technological advancement overrides everything, that we are up to our necks in singularity, and from now on nothing will be the same as before. We are constantly hearing that the youth of today, Generation Z, are completely different from earlier generations in world history: their lives are spent on the World Wide Web, they have difficulty concentrating, they do not read, cannot remember data, cannot tolerate monotony, are afraid of commitment, and so forth. Where with this lead? The following quotations reveal that there is probably nothing new under the sun in this regard[4]:

- "Our young people (...) are badly brought up, caring nothing for authority and showing no respect for their elders. These days our sons (...) do not rise when an elderly person enters the room, they talk back to their parents and chatter instead of working. They are simply insufferable." (*Socrates*, 470–399 BC)
- "I no longer hold out any hope for the future of our country if today's youth comes to power tomorrow, because this youth is intolerable, knows no restraint, and are simply dreadful." (*Hesiod*, first half of 8th century BC)
- "The world has reached a crisis situation. Children no longer listen to their parents. The end of the world cannot be far." (Anonymous Egyptian monk, approx. 2,000 years ago)
- "The youth are rotten to the core. Young people are depraved and good for nothing. They will never be like the youth of old. Today's young people will be incapable of preserving our culture." (Babylonian clay tablet, approx. 3,000 years ago)

3 Moore's famous prediction was that the number of transistors in integrated circuits would keep doubling roughly every two years. Source: http://www.economist.com/node/3798505.

4 Source: http://www.szepi.hu/irodalom/pedagogia/tped_044.html

Unfortunately, there are also changes that may genuinely give us cause for concern. We surely cannot say, for example, that global warming is “nothing new under the sun.” It could easily be that the structural fissures observable in financial data are harbingers of truly serious ruptures in both the environment and the real economy.

7. SUMMARY

Through numerous greatly varying financial examples, we have thought about ergodicity in terms of the empirical knowability of stochastic processes. Our first observation was that it is not possible, based on the study of historical data, to determine exactly whether or not a process is ergodic, and at most we may have only an intuition. At the same time, it appears that it is not worth making generalized statements about the ergodicity of the economy. From time to time, it may certainly be useful to analyse time series, for example when pricing derivatives or during risk management; at other times, however, this may be completely misleading – as indicated, for example, by the equity premium puzzle or the difficulty of measuring a fund manager’s performance.

My conjecture is that any task that only requires the prediction of risks (for example, derivative pricing or risk management) is accomplishable without the need to dismiss the ergodic hypothesis. If, however, the expected returns also need to be estimated (for example, in forecasting equity premiums or performance measurement), then we are treading on more shaky ground. The reason for this may be that financial time series are highly prone to noise, and a great degree of volatility obscures the expected return. The shorter the examined time horizon, the more volatility outgrows the expected return by an order of magnitude, so that increasing the frequency of observation does not really help. If, on the other hand, we increase the length of the observation period, then we can be sure that sooner or later one or two breaks in the trend will occur. It is no accident that, on the basis of empirical experience, volatility can be reliably predicted, while forecasting of the expected return is entirely illusory.

Interestingly, even the long term does not necessarily enhance stability. Often we cannot say more about processes even if we carry out our investigations on a historical scale; for example, predictions of interest rates are little helped by looking back for 5,000 years, and neither can we rely on historical averages in forecasting temperature. In other instances, however, we can perhaps count on history repeating itself; the only trick then being to know which instances these are.

REFERENCES

Algoet, P. – Cover, T. (1988): Asymptotic Optimality and Asymptotic Equipartition Properties of Log-Optimum Investments. *Annals of Probability*, 16, pp. 876–898.

Arrow, K. J. (1963): *Social Choice and Individual Values*.Second Edition, New York and New Haven: John Wiley & Sons.

Béli, M. (2012): The Equity Risk Premium Puzzle in Hungary (Kockázati prémiumok rejtélye Magyarországon). *Hitelintézeti Szemle*, 11(5), pp. 403–441.

Bélyácz, I. (2017): The Debated Role of Ergodicity in (Financial) Economics. *Economy & Finance (Gazdaság* és *Pénzügy)*, 2017, 4(1), pp. 4–57.

Black, F. – Scholes, M. (1973): The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3), pp. 637–654.

Bodie, Z. – Kane, A. – Marcus, A. J. (2005): *Investments*. Published in Hungarian by Aula Publishing, Budapest.

Csóka, P. – Pintér, M. (2016): On the Impossibility of Fair Risk Allocation. *The B.E. Journal of Theoretical Economics*, 16(1), pp. 143–158.

Györfi, L. – Lugosi, G. – Udina, F. (2006): Nonparametric Kernel-based Sequential Investment Strategies. *Mathematical Finance*, 16, pp. 337–357.

Haldane, A.G. (2015): *Stuck: Speech by Andy Haldane*. Open University, Milton Keynes, http://www.bankofengland.co.uk/publications/Documents/speeches/2015/speech828.pdf

Harcsa, I. (2017): The Thesis of Ergodicity in Broader Contexts. *Economy & Finance (Gazdaság* és *Pénzügy)*, 2017, 4(1), pp. 58–68.

Horst, U. (2008): Ergodicity and Nonergodicity in Economics. In: *The New Palgrave Dictionary of Economics*, Second Edition (eds: Steven N. Durlauf, Lawrence E. Blume).

Hull J. C. – White, A. (1987): The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, (42), pp. 281–300.

Hull, J. C. (2015): *Options, Futures, and Other Derivatives*. Boston, MA: Pearson.

Lamanda, Zs. – Tamásné, V. Zs. (2015): Hungry for Risk: The Framework of Risk Appetite for Operational Risks (Kockázatra **éhezve**. A kockázati **étvágy** keretrendszere a működési kockázatoknál). *Pénzügyi Szemle*, 2015(2), pp. 217–230.

Mankiw, N. G. – Zeldes, S. P. (1991): The Consumption of Stockholders and Non-Stockholders. *Journal of Financial Economics*, (29), pp. 97–111.

Medvegyev, P. (2017): Some Observations on the Ergodicity of Economic Processes. *Economy & Finance (Gazdaság* és *Pénzügy)*, 2017, 4(2), pp. 110–119.

Mehra, R. – Prescott, E.C. (1985): The Equity Premium: A Puzzle. *Journal of Monetary Economics*, (15), pp. 145–161.

Mellár, T. (2017): Ergodicity and its Consequences: Comments on a Study by Iván Bélyácz. *Economy & Finance (Gazdaság* és *Pénzügy)*, 2017, 4(2), pp. 98–109.

Misik, S. (2015): Implicit Correlation on Foreign Currency Markets. *Economy & Finance (Gazdaság* és *Pénzügy)*, 2015, 2(4), pp. 288–305.

Oliveira, J. G – Barabási, A.L. (2005): Human Dynamics: Darwin and Einstein Correspondence Patterns. *Nature*, 437, 1251, DOI: 10.1038/4371251a.

Ormos, M. – Urbán, A. – Zoltán, T. (2009): An Empirical Study of Log-optimal Portfolios (Logoptimális portfóliók empirikus vizsgálata). *Közgazdasági Szemle*, Vol. LVI, January, pp. 1–18.

Peters, O. (2011): Optimal Leverage from Non-Ergodicity. *Quantitative Finance*, 11:11, pp. 1593–1602, DOI: 10.1080/14697688.2010.513338

Peters, O. – Klein, W. (2013): Ergodicity Breaking in Geometric Brownian Motion. *Phys. Rev. Lett.* 110, 100603, DOI: 10.1051/jp1:1992238

Piketty, T. (2015): *Capital in the Twenty-First Century.* Published in Hungarian by Kossuth Kiadó.

Simonovits, A. (2017): Ergodic Theory and Economics: Remarks on an Article by Iván Bélyácz. *Economy & Finance (Gazdaság* és *Pénzügy)*, 2017, 4(2), pp. 120–125.