Measures of financial risk

S.Y. Novak
School of Science & Technology
Middlesex University, London

Abstract

The paper compares a number of available measures of financial risk and presents arguments in favor of dynamic measures. We argue that traditional measures are static, while dynamic measures lack statistical scrutiny. The main obstacle to building a body of empirical evidence in support of a dynamic risk measure is computational difficulty of identifying local extrema as price charts appear objects of fractal geometry. We overview approaches to financial risk measurement and formulate a number of open questions. The arguments are illustrated on examples of real data.

Key words and phrases: risk measures.

AMS 2000 Subject Classification: primary 60G70; secondary 60G55, 60E15.
JEL Classification: C53, G17.

1 Introduction

Accurate evaluation of risk is fundamental to financial well-being of financial institutions as well as individual investors. However, the issue is demanding. It involves sophisticated statistical analysis of market data.

The traditional approach to risk measurement is capable of forecasting the magnitude of a possible sharp market movement (see, e.g., Novak & Beirlant (2006)). However, the estimates appear static: they barely change with the inflow of new information, and may be unsuitable for active portfolio management. The "Technical Analysis" approach offers a truly dynamic risk measure m_{TA} (cf. Novak (2011), ch. 10). However, lacking statistical scrutiny affects its credibility. The main obstacle that prevents building a body of empirical evidence either in favor or against the use of m_{TA} is the computational difficulty caused by the fact that price charts appear objects of fractal geometry.

In the next section we briefly overview popular measures of risk, including Value-at-Risk (VaR) and a related measure called conditional VaR (CVaR) or Expected Shortfall. Section 3 discusses strengths and weaknesses of the traditional approach to risk measurement. In section 4 we briefly overview basic tools of the Technical Analysis approach. Section 5 presents properties of a dynamic risk measure. Section 6 concentrates on a number of open questions concerning computational and statistical issues related to dynamic risk measurement.

2 Traditional measures of risk

Let X denote the underlying financial/insurance data (e.g., the rate of return of a particular financial asset, the daily log-return of a share price, claim sizes, etc), and let

$$F(x) = \mathbb{P}(X \le x)$$

be the distribution function (d.f.) of $\mathcal{L}(X)$. Traditional measures of risk are the standard deviation

$$\sigma_X$$

and the beta

$$\beta_X = \mathrm{cov}(X; M)/\sigma_M^2 ,$$

where X is the rate of return of a particular financial asset, M is the rate of return of the market portfolio (often approximated by the S&P500 index) and σ_M^2 is the variance of M. In the case of ARCH/GARCH models one often deals with the conditional standard deviation σ_{X_n} of X_n given $(X_{n-1}, X_{n-2}, ...)$.

Among traditional measures of risk is Value-at-Risk (VaR). Up to a sign, Value-at-Risk is an extreme quantile:

$$m\%\text{-VaR} = -F^{-1}(m/100) .$$

Equivalently,

$$q\text{-VaR} = -F^{-1}(q) \qquad (0 < q < 1).$$

Recall that

$$F^{-1}(q) = \inf\{x : F(x) \ge q\} .$$

If F is continuous, then q-VaR is such a number that

$$\mathbb{P}(X \le -q\text{-VaR}) = q.$$

$m\%$-VaR indicates how far the quantity of interest (say, daily log-return X) can fall in approximately $m\%$ "worst" cases. For instance, if 1%-VaR equals 0.04, then in approximately 1% cases the quantity of interest, X, can be below -4%.

One often deals with log-returns $X_k = \ln(P_k/P_{k-1})$ instead of prices $\{P_k\}$, as log-returns are more likely to form a stationary sequence. If 1%-VaR for daily log-returns equals y, then roughly once in 100 days the value of a portfolio may be below e^{-y} times the previous day value.

VaR is probably the most popular measure of risk. Many banks routinely calculate VaR in order to monitor the current exposure of their portfolios to market risk. Knowing VaR allows a bank to decide how much money it needs to put aside to offset the risk of an undesirable market movement. For instance, Goldman Sachs deals with 5%-VaR; Citigroup, Credit Suisse First Boston, Deutsche Bank, JP Morgan Chase and Morgan Stanley use 1%-VaR (see, e.g., Wells et al. (2004), Gurrola-Perez & Murphy (2015)).

Closely related is another measure of risk known as Conditional VaR (CVaR), Expected Shortfall or BVaR (Artzner et al. (1999), Longin (2001), Pflug (2000), Rockafellar & Uryasev (1999)). It represents the average loss given there is a fall beyond VaR: assuming $\mathbb{E}|X| < \infty$,

$$\mathrm{CVaR} = -\mathbb{E}\{X | X \le -\mathrm{VaR}\}.$$

One often prefers dealing with positive numbers (for instance, we speak about "20.5% fall" of the S&P500 index on the "Black Monday" instead of "–20.5% rate of return"). If we switch from X to $-X$, then VaR is the upper quantile (the inverse of $F_c = 1 - F$):

$$q\text{-VaR} = F_c^{-1}(q)\,, \tag{1}$$

and

$$\text{CVaR} = \mathbb{E}\{X|X \geq \text{VaR}\} = \text{VaR} + \mathbb{E}\{X - \text{VaR}|X \geq \text{VaR}\} \tag{2}$$

(we sometimes omit prefix q-). Recall that

$$\mathbb{E}\{X - x|X \geq x\}$$

is the *mean excess function*, also known as the mean residual life function.

Properties of VaR and CVaR can be found, e.g., in Alexander (2008), Nadarajah & Chan (2015), Novak (2011), Pflug (2000).

Example 1. Let $\{X_n, n \geq 1\}$ be the ARCH(1) process with parameters $b > 0$ and $c \geq 0$:

$$X_n = \xi_n \sqrt{b + cX_{n-1}^2} \qquad (n \geq 2)\,, \tag{3}$$

where $\{\xi_i\}$ is a sequence of independent normal $\mathcal{N}(0;1)$ random variables. The conditional standard deviation of X_n given $(X_{n-1}, X_{n-2}, ...)$ is $\sigma_{X_n} = \sqrt{b + cX_{n-1}^2}$,

$$q\text{-VaR}_{X_n} \simeq -t_q\sqrt{b + cX_{n-1}^2}\ , \quad q\text{-CVaR}_{X_n} \simeq \exp(-t_q^2/2)\sigma_{X_n}/q\sqrt{2\pi}\ , \tag{4}$$

where $t_q = \Phi^{-1}(q)$ and Φ is the standard normal distribution function. The estimate of c is typically close to 0, making random variables σ_{X_n} , q-VaR$_{X_n}$ and q-CVaR$_{X_n}$ rather static. □

Note that σ, β, VaR and CVaR (as well as m_{TA} below) measure risk on the base of the past ("historical") data. A different approach based on "future values only" was suggested by Artzner et al. (1999): "The basic objects of our study shall be possible future values of positions or portfolios". Artzner et al. (1999) define a measure of risk related to the "acceptance set" A as

$$\rho_A(X) = \inf\{m\colon mr_o + X \in A\},$$

where X is "the final net worth of a position" and r_o is the rate of return of the risk-free asset.

3 Risk estimation

This section is devoted to the problem of estimating risk measures assuming a particular model. The topic of testing goodness-of-fit of particular models goes beyond the scope of this article.

One has to distinguish a quantity one wants to estimate (e.g., 1%-VaR) from a model (i.e., the assumptions on the class of distributions the unknown distribution $\mathcal{L}(X)$ belongs to).

If data is light-tailed (this, of course, needs to be checked), it hardly can exhibit extreme movements, and the assumption that $\mathcal{L}(X)$ has normal $\mathcal{N}(\mu;\sigma^2)$ distribution is not unreasonable. In the case of frequent data μ is typically negligent, and hence

$$q\text{-VaR}_{\mathcal{N}} \simeq -\sigma t_q\,,\; q\text{-CVaR}_{\mathcal{N}} \simeq \exp(-t_q^2/2)\sigma/q\sqrt{2\pi} \;\simeq q\text{-VaR}_{\mathcal{N}}\,. \tag{5}$$

According to (5), one only needs to estimate the standard deviation $\sigma \equiv \sigma_X$ (cf. Alexander (2001), §9.3).

Financial/insurance data often exhibits heavy tails (see Embrechts et al. (1997), Fama & Roll (1968), Longin (1996), Mandelbrot (1963)). This is particularly common to "frequent" data (e.g., daily log-returns of stock prices and stock indexes), while log-returns of less frequent data can exhibit light tails – well in line with the central limit theorem. Note that the stationary distribution of the ARCH(1) process (3) is heavy-tailed (see Goldie (1991), Embrechts et al. (1997), pp. 465-466).

The feature of heavy-tailed distributions is that a single observation can be of the same order of magnitude as the whole sum of sample elements: a single claim to an insurance company or a one-week market movement can cause a loss comparable to a one-year profit.

A number of procedures to check if the distribution is heavy-tailed are mentioned in Markovich (2007), Novak (2011).

The problem of evaluating risk from dependent heavy-tailed data attracts increasing attention of researchers (see Nadarajah & Chan (2015), Novak (2011) and references therein).

The distribution of a random variable X has a *heavy right tail* if

$$\mathbb{P}(X \geq x) = L(x)x^{-\alpha} \qquad (\alpha > 0), \tag{6}$$

where the (unknown) function L is slowly varying at ∞:

$$\lim_{x\to\infty} L(xt)/L(x) = 1 \qquad (\forall t > 0).$$

Number α in (6) is called the *tail index*. It is the main characteristic describing the tail of a heavy-tailed distribution.

If the distribution of a random variable X has a *heavy right tail* with tail index $\alpha > 1$, then

$$\mathbb{E}\{X|X \geq x\} \simeq \frac{\alpha}{\alpha-1}\,x$$

for large x (see, e.g., Embrechts et al. (1997), p. 162). For small q

$$q\text{-VaR} = \ell(q)q^{-1/\alpha}\,,$$

where ℓ is a slowly varying function (see Seneta (1976), Bingham et al. (1987)), and hence

$$q\text{-CVaR} \simeq \frac{\alpha}{\alpha-1}\,q\text{-VaR}$$

(cf. Longin (2001), equation (4), Novak (2011), p. 168).

Heavy-tailed distributions form a domain of attraction to one of the possible limit laws for the sample maximum (see Gnedenko (1943)). Interest to heavy-tailed distributions is inspired also by rich applications in hydrology, meteorology, etc. (see, e.g., Fraga Alves & Neves (2015)).

The problem of reliable estimation of the tail index α,

$$y_q := q\text{-VaR} \text{ and } z_q := q\text{-CVaR}$$

from heavy-tailed data is demanding (cf. Embrechts et al. (1997), Novak (2011)). The following estimators of VaR and CVaR are asymptotically normally distributed (see Novak (2002), (2011)) and appear reasonably accurate in examples of simulated data:

$$y_{n,q} \equiv y_{n,q}(x) = x\,(N_n(x)/qn)^{\hat{a}_n^{RE}}\,, \tag{7}$$

$$z_{n,q} \equiv z_{n,q}(x) = y_{n,q}/(1-\hat{a}_n^{RE})\,. \tag{8}$$

Here threshold x is a tuning parameter (it needs to be chosen), q is the given level and $\hat{a}_n^{RE}$ is the Ratio Estimator of index $a = 1/\alpha$:

$$\hat{a}_n^{RE} \equiv \hat{a}_n^{RE}(x) = \sum_{i=1}^{n} \ln(X_i/x)\mathbb{I}\{X_i > x\} \Big/ \sum_{i=1}^{n} \mathbb{I}\{X_i > x\} \tag{9}$$

(equivalently, $1/\hat{a}_n^{RE}$ is the Ratio Estimator of the tail index).

This approach is non-parametric. The comparison of finite-sample properties and asymptotic mean-squared errors of a number of tail index estimators in Novak (2011), p. 150, is in favor of the Ratio Estimator.

There are doubts that parametric models accurately describe real financial data: one usually cannot be sure if the unknown distribution comes from a chosen parametric family. The advantage of the non-parametric approach is that such a problem is void: the non-parametric class is so rich that one typically has no doubt the unknown distribution belongs to it. The disadvantage of the non-parametric inference is the presence of a tuning ("nuisance") parameter.

A procedure of choosing the tuning parameter for estimators (7), (8) and (9) has been suggested in Novak (2002), see also Novak (2011). It is data-driven, and suggests to
(i) plot an estimator $\hat{a}_n$ (e.g., $\hat{a}_n^{RE}$) as a function of the tuning parameter,
(ii) choose an interval $[x_-; x_+]$ formed by a significant number of sample elements, in which the plot demonstrates stability (the theoretical results in Novak (2002) state there should be such an interval of stability),
(iii) take the average value $\hat{a}$ of the estimator $\hat{a}_n$ in interval $[x_-; x_+]$ as the estimate of a (i.e., we choose the threshold $\hat{x}_n \in [x_-; x_+]$ such that $\hat{a}_n(\hat{x}_n) = \hat{a}$).
Since we take the average over an interval formed by a significant number of sample points, the procedure yields almost one and the same estimate despite the individual choice of the interval of stability: the variability with the choice of end–points is almost eliminated.

Example 2. *The "Black Monday" crash.* Forecasting the scale of possible extreme movements of financial markets is one of the main tasks of a risk manager. A particular question of this kind was raised in McNeil (1998): having the data over the period 01.01.1960 – 16.10.1987, was is possible to predict the magnitude of the next crash?

If data is heavy-tailed, then the standard deviation does not appear a proper tool to describe the risk associated with extreme movements of the portfolio returns even if a portfolio is optimal in the sense of the mean–variance portfolio theory.

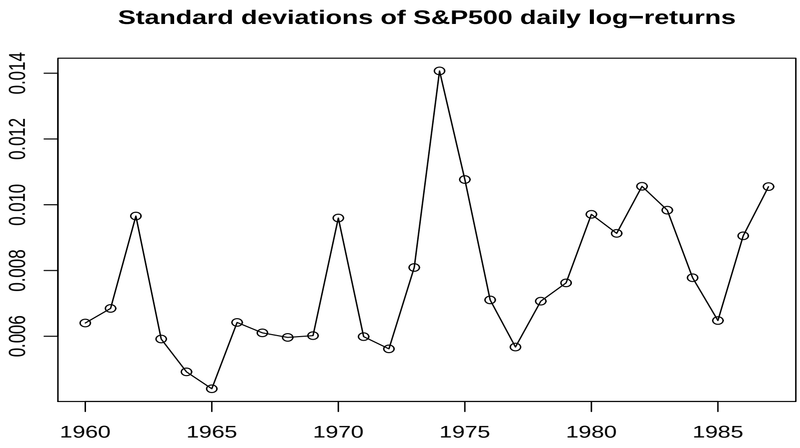

Figure 1: The standard deviations of daily log-returns of the S&P500 index in January 1960 – October 1987. The value of the standard deviation on 16.10.1987 is close to those on 16.10.1982, 16.10.1970 and 16.10.1962, and hardly can serve an indicator of a possible crash.

Fig. 1 presents the standard deviations of daily log-returns of the S&P500 index over the period 01.01.1960 – 16.10.1987 (the standard deviations were calculated using a one year of preceding data). The value of the standard deviation on 16.10.1987 is close to that on 16.10.1982, 16.10.1970 and 16.10.1962, and hardly can serve an indicator of a possible crash. Yet on Monday 19.10.1987 the S&P500 index fell by 20.5% — the worst daily fall of the index (see Fig. 2).

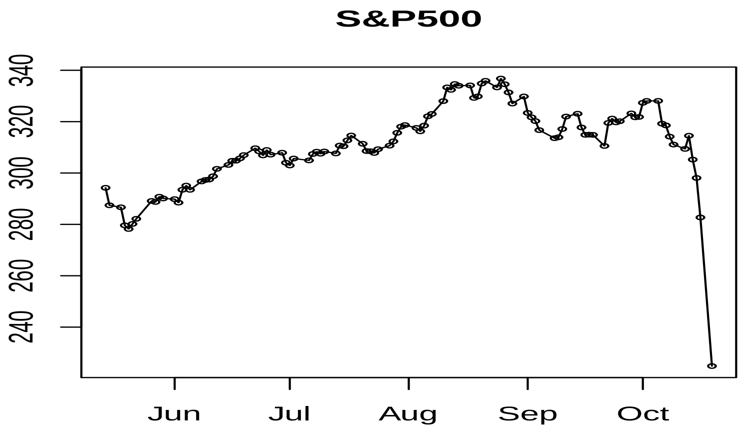

Figure 2: The S&P500 index in May 1987 – October 1987. On Monday the 19th, October 1987, the index fell by 20.5%.

Recall that if data is heavy-tailed then a single sample element, e.g., the loss over one particular day, can make a major contribution to the total loss over a considerable period of time. In particular, the "Black Monday" crash erased all the index had gained since March 1986.

Sometimes data exhibits such heavy tails that the variance is likely to be infinite. In all such situations Value-at-Risk and CVaR appear more suitable measures of risk.

We have applied estimators (7), (8) and the procedure of choosing the tuning parameter to daily log-returns of the S&P500 index over the period from 01.01.1960 till 16.10.1987 (formed by approximately 10000 trading days). The Dickey-Fuller test does not reject the hypothesis of stationarity. We use the Ratio Estimator in order to estimate the tail index (a comparison of a number of tail index estimators can be found in Markovich (2007) and Novak (2011)).

The plot of the Ratio Estimator $\hat{a}_n^{RE}(\cdot)$ is stable in the interval $[1.4; 5.4]$. The curve over that interval is formed by 285 points. There were 3319 falls during that period, while only 512 of them exceeded the 1% level. We conclude that interval $[1.4; 5.4]$ is formed by a significant number of sample elements. The average value of $\hat{a}_n^{RE}(x)$ as $x \in [1.4; 5.4]$ is

$$\hat{a} = 0.2515,$$

and threshold $\hat{x}_n$ is chosen so that $\hat{a}_n^{RE}(\hat{x}_n) = \hat{a}$. Thus, the tail index has been estimated at

$$1/\hat{a} = 3.97.$$

VaR and CVaR estimators (7), (8) with $q = 0.01\%$ and the threshold already chosen yield 18.1% for 0.01%-VaR and 24.2% for the corresponding CVaR. Hence the worst possible fall of the daily log-return of the S&P500 index in 40 years, according to the data available on the eve of the "Black Monday", was likely to be around 24.2%.

This is remarkably close to the value of the actual fall on 19.10.1987. The closing price of S&P500 on 16.10.87 was 282.94 (already 5% down on the previous day); the closing price of S&P500 on 19.10.1987 was 225.06, the log-return was equal to -0.229. $\square$

4 "Technical Analysis" of financial data

In this section we overview the approach to the analysis of financial data known as "Technical Analysis" (TA).

Measures of risk we discussed so far were static: they barely change with the inflow of new information and hence only suitable for long-term investment decisions; one hardly would use them for short-term investment decisions.

Measures like the standard deviation or VaR are static by the definition (it's not the standard deviation or VaR that changes but the values of our estimates of σ and VaR). Moreover, a long-term trend can be replaced with the trend in the opposite direction within few days, leaving static measures no time to react to the change (cf. the recent volatility of the crude oil prices). Thus, a short-term investor would prefer a dynamic risk measure.

It is widely believed that there are moments of time when investing in a particular financial instrument is less risky, and moments of time when the level of risk is high. One would prefer to invest when risk is low, and reduce or close a position when risk is high (cf., e.g., Choffray & de Mortanges (2015)). Determining such moments can be a key to successful investing, yet it's a difficult task. One way to locate such moments is by using a dynamic measure of risk.

We call a measure of risk dynamic if it changes considerably with the change of market data.

The first step towards developing a dynamic measure of risk was made with the introduction of the ARCH model (3). If the sequence $\{X_i, i \geq 1\}$ of random variables (say, daily log-returns) obeys (3), then the conditional variance

$$\sigma_n^2 \equiv \mathbb{E}\{(X_n - \mathbb{E}X_n)^2 | X_1, ..., X_{n-1}\} = b + cX_{n-1}^2$$

is a function of the previous observation.

A more general (not necessarily more accurate) is model ARCH(k), assuming

$$X_n = \xi_n \sqrt{b + c_1 X_{n-1}^2 + ... + c_k X_{n-k}^2} \qquad (n \geq k{+}1,\ k{\geq}1),$$

where $\{\xi_n, n \geq 1\}$ are independent normal $\mathcal{N}(0;1)$ random variables. Estimates of parameters $c, c_1, ..., c_k$ are usually small, meaning the influence of the past on the conditional variance is not dramatic.

A very different approach to risk measurement is based on TA tools.

Just like the traditional approach, which is sometimes called "Fundamental Analysis" (FA), Technical Analysis deals with a set of past (historical) data (e.g., daily closing share prices $\{P_n\}$ of a particular stock over the past two years).

Following the traditional approach, one starts by computing a handful of numbers (e.g., the rate of return

$$r_n = P_n / P_{n-1} - 1,$$

the standard deviation of $\{r_n\}$, the beta, the price to earnings ratio, etc), displays them and makes an executive decision (see, e.g., Elton & Gruber (1995), Luenberger (1997)). A practical advice on the use of the FA approach can be found, e.g., in Choffray & de Mortanges (2015).

Following TA, one also starts by computing a set of numbers (e.g., daily Open-High-Low-Close prices, moving averages, etc.). The main difference is that the results of calculation are displayed not in a numerical but in a graphical form (TA is sometimes called "charting").

Both approaches aim at spotting trends and points of entrance and exit. The idea behind TA is that a human eye can spot trends and regions of entrance and exit well before numerical measures like price to earnings ratio, VaR, etc, change considerably.

Background to the Technical Analysis. The Technical Analysis presumes that price movements form patterns, those price patterns can be studied, classified, etc, and an observer is capable of recognising a price pattern before the formation of a pattern is complete.

These assumptions mean the price movements are not considered completely random. Indeed, it is widely observed that prices often appear to exhibit trends, levels of support and resistance, etc. Not always an observer is capable of recognising a pattern the price is currently forming. However, an observer is free to invest only when s/he recognises a particular pattern. Moreover, the mere fact that many investors/speculators are aware of certain price patterns makes their reaction to corresponding price formations predictable, causing a plausible pattern even more likely to become a reality.

Note that TA assumptions contradict to ARCH and the Geometric Brownian Motion models, and suggest searching for models that allow for patterns of cyclical behavior. TA assumptions contradict also to the efficient market hypothesis (EMH). The arguments in favor and against

EMH and TA can be found, e.g., in Akram et al. (2006), Brock et al. (1992), Clarke et al. (2001), Elder (2002), Fama (1970), French (1980), Higson (2001), Irwin & Park (2007), Lo et al. (2000), Osler & Chang (1995), Poser (2003), Prechter & Parker (2007), Williams (1994).

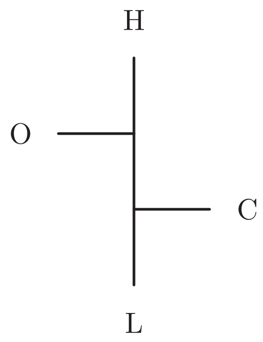

Figure 3: Open-High-Low-Close price bar. Fig. 5 below is an example of a price chart with Open-High-Low-Close price bars. Many web-cites offer displaying price charts as a sequence of price bars.

The present paper looks at TA tools from a different angle. The main point for us is that TA ideas lead to a dynamic measure of risk m_{TA} , which does change considerably as price changes. We overview few popular price patterns before discussing TA tools for dynamic risk measurement.

Basic Elliott waves. The price chart appears an object of fractal geometry. A common way to simplify the chaotic structure of the chart is by using price bars (Fig. 3).

Price charts often appear to exhibit patterns like those in Fig. 4: the general trend is up but it is interrupted by “corrections”. Such patterns are called Elliott waves.

The basic Elliott wave is a 5-leg zigzag (see Fig. 4 for an up-trend); it is typically followed by a 3-leg wave in the opposite direction. The use of straight lines in Fig. 4 is, of course, a simplification: each wave has “inside” a set of smaller scale waves, each “sub-wave” has again a set of even smaller waves inside, etc. (cf. monthly, weekly, daily, hourly, 10-min. and 1-min. price charts of the same stock or index).

Recognising an Elliott wave is not easy, and different observers can mark vertexes differently (cf. Fig. 5). Few empirically observed facts can be helpful in determining vertexes:

{**1**} The level of vertex 2 can be within 38% of the range of wave 1.

{**2**} The range of wave 3 (section 2–3) can be up to 262% of the range of wave 1. MACD histogram is positive during the 3^{rd} wave.

{**3**} Vertex 4 almost nether goes beyond the level of vertex 1. The range of wave 4 (section 3–4) is 38% – 50% of the range of wave 3. MACD histogram during the 4^{th} wave is negative.

{**4**} The range of the 5^{th} wave (section 4–5) is often within 62% – 100% of the 3rd one.

These are empirically observed facts (cf. Williams (1994)), there seems to be no statistical study behind these observations so far.

One can only be sure about locating a vertex only after the price has passed it. Making an investment decision on the basis of an assumption that the price is, say, in vicinity of vertex 2,

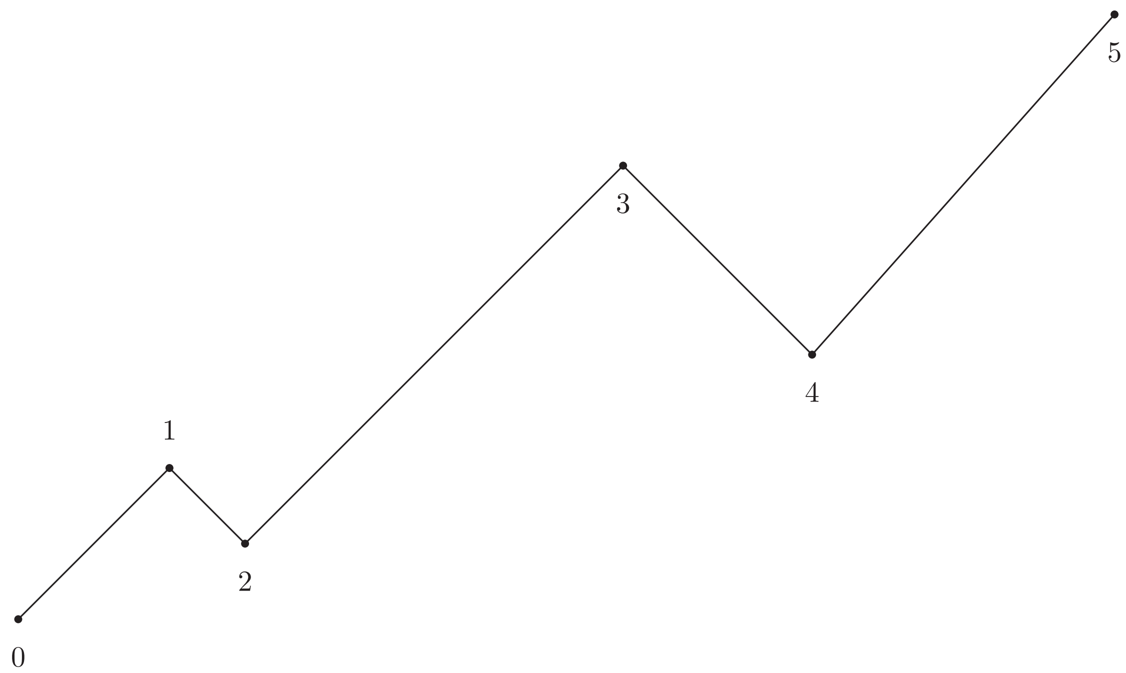

Figure 4: Basic Elliott wave for an up-trend (see also see Fig. 5). The use of straight lines is, of course, a simplification: each wave has "inside" a set of smaller scale waves, each "sub-wave" has again a set of even smaller waves inside, etc.

one understands an element of uncertainty is inevitable. However, dealing with imprecise objects is not uncommon in Statistics. Recall that if θ_n is an estimator of a certain quantity, θ, then $(1+1/n)\theta_n$ is often "equally good" estimator (e.g., the sample variance estimator

$$\hat{\sigma}_n^2 = n^{-1}\sum_{i=1}^{n} X_i^2 - \bar{X}^2 ,$$

where $\bar{X} = \sum_{i=1}^{n} X_i/n$, is a natural estimator constructed by the method of moments, while

$$\tilde{\sigma}_n^2 = \sum_{i=1}^{n}(X_i - \bar{X})^2/(n-1)$$

is an unbiased estimator).

Among TA signals/indicators one can mention
Moving Average (MA) — a basic indicator of a trend (MA(k) is the average of k past prices);
Moving Average Convergence-Divergence (MACD), one of the most popular TA indicators;
Relative Strength Index (RSI);
Breakout signal.
Formal definitions can be found, e.g., in Pring (2002), Williams (1994), Novak (2011), ch. 10. These indicators are functions of certain "tuning" parameters (e.g., the time scale).

5 Dynamic risk measurement

Measures of risk discussed in sections 1 – 3 (the standard deviation, VaR, CVaR) were static. In addition to static measures of risk suitable for long-term investment decisions, a short-term investor would like considering a dynamic measure of risk.

In this section we discuss a particular approach to dynamic risk measurement based on an idea of stop-loss from the Technical Analysis of financial data. Namely, we deal in this section with measure m_{TA} that measures a level of risk by the distance from the current price to the stop-loss level identified by the nearest local price minimum or maximum (see Novak (2011), ch. 10).

The definition of m_{TA}: for a holder of a "long" position it is the distance between the current price, P_n, and the last local price minimum below P_n:

$$m_{TA} = P_n - P_{n_*}\,, \tag{10}$$

where n_* is the time of the last local price minimum below P_n. For a holder of a "short" position n_* is the time of the last local price maximum above P_n and $m_{TA} = |P_n - P_{n_*}|$.

We denote by $m_{TA} \equiv m_{TA}(cP_n)$ the measure of risk when investing in c units of instrument P at time n (we sometimes omit P_n and c). The following example illustrates the definition.

Example 3. In the case of a basic Elliott wave the local extrema are the vertexes of the wave.

If the price is considered to be in wave 2, then for a holder of a "long" position m_{TA} is the distance between the current price and the level of vertex 0, which is the nearest local price minimum (Fig. 4); for a holder of a "short" position m_{TA} is the distance between the current price and the level of vertex 1.

Suppose a long position is opened when price is in wave 2. As the price moves down towards vertex 2, the paper loss increases while m_{TA} decreases. If the price falls below the level of vertex 0, the assumption that the price was in wave 2 is proved to be wrong, m_{TA} resets using the level of the last local minimum below the level of vertex 0, one would have also a breakout signal that the price was in the bear trend.

If the price is in wave 3 below the level of vertex 1, then for a holder of a "long" position m_{TA} is the distance between P_n and the level of vertex 2 (until the next local minimum is formed); for a holder of a "short" position m_{TA} is the distance between P_n and the level of vertex 1. After passing the level of vertex 1, m_{TA} is the distance between P_n and the level of the last local price minimum (typically, wave 3 is again a pattern with a number of local extrema, and m_{TA} is reset when a new local price minimum is formed and recognised); for a holder of a "short" position the risk jumps: m_{TA} is now the distance between P_n and the last local price maximum above the level of vertex 1.

If the price is in wave 4, then for a holder of a "long" position m_{TA} is the distance between P_n and the last local price minimum below P_n; for a holder of a "short" position m_{TA} is the distance between P_n and the level of vertex 3.

In wave 5 for a holder of a "long" position m_{TA} is the distance between P_n and the last local price minimum; for a holder of a "short" position the risk may be unlimited. $\square$

Note that a price chart is virtually an object of fractal geometry; a local extremum of a fractal cannot be identified. To be precise, a price chart is not exactly a fractal, and local extrema can be formally identified. However, by the local extremum a practitioner would mean a "recognisable" local minimum or maximum, cf. Fig. 4 and Fig. 5. Thus, m_{TA} is not uniquely defined (similarly, the definitions of patterns in TA are not mathematically strict). Just like TA patterns and signals, local extrema may be identified differently by different investors.

One approach to deal with this problem is to use data smoothing; the degree of smoothing is the "tuning" parameter. However, the choice of a tuning parameter becomes an issue.

Figure 5: Prices of the Akamai stock in Nov 2008 – Feb 2009. In early December an investor might have concluded that the price was in wave 2. With vertex 0 identified when the price was near \$9.50, the risk of holding a "long" position was around \$2 per share. The potential profit (if the position held till the end of wave 3) was around \$5.

Example 4. Consider, for instance, Fig. 5. In early December, when the price of Akamai stock was around \$11.50, an investor might have concluded the price was in wave 2. With vertex 0 identified when the price was near \$9.50, the risk of holding a "long" position was around \$2 per share, i.e., about 15%–20% of the value of the investment. The potential profit if the position held till the end of wave 3, according to property $\{2\}$, was around \$5.

Note that wave 4 can be identified as MACD histogram is negative in the first half of January 2009. As the price almost never falls below the level of vertex 3, an investor could have added a long position mid-January when the price was around \$13 with $m_{TA} < \$1$. We know the range of wave 5 is usually at least 62% of the range of section 0–3, i.e., one would expect the price to move up by \$4.2 – \$7; the ratio of the potential gain to the potential loss is greater than 4. The price did move up by about \$5; the S&P500 index was still declining during that period.

One could have used other TA indicators as well. For instance, Fig. 5 shows the MACD divergence signal: the price makes higher highs while the MACD histogram makes lower highs around the 10th of December, mid-December and early January. Expecting a short-term bear trend, an investor could have decided to open a short position in early January when MACD histogram made a local maximum. The measure of risk (the distance between the price $\approx \$15.50$ and the last local maximum under \$17) is $m_{TA} \approx \$1.5$. The chart shows the potential profit would be around \$3 per share. $\square$

Properties of $\mathbf{m_{TA}}$.

Property [**1**]. m_{TA} *is not symmetric.*

The level of risk for a holder of a long position is not equal, in general, to the level of risk for a holder of a short position. Note also that m_{TA} may grow as price moves in investor's favour.

Property [**2**]. $m_{TA} \geq 0$ and $m_{TA}(\lambda c P_n) = \lambda m_{TA}(c P_n)$.

Note that for a holder of a long position $m_{TA}(cP_n) \leq cP_n$.

For a holder of a short position $m_{TA} \leq \infty$ – not always the last local maximum can be identified (e.g., the price makes the new global maximum), meaning the amount of a possible loss for a holder of a short position is potentially unlimited.

Property [**3**]. m_{TA} *is typically sub-additive.*

Indeed, let X and Y be two instruments, and suppose that the last local minima of instruments X, Y and $X+Y$ has been identified within the period $[n_*, ..., n]$. Then for a holder of a long position

$$\begin{aligned} m_{TA}(X_n+Y_n) &= X_n + Y_n - \min_{n_* \leq i \leq n}(X_i + Y_i) \\ &\leq \Big(X_n - \min_{n_* \leq i \leq n} X_i\Big) + \Big(Y_n - \min_{n_* \leq i \leq n} Y_i\Big) = m_{TA}(X_n) + m_{TA}(Y_n)\,. \end{aligned}$$

If the last local maxima of instruments X, Y and $X+Y$ has been identified within the period $[n_*, ..., n]$, then for a holder of a short position

$$\begin{aligned} m_{TA}(X_n+Y_n) &= \max_{n_* \leq i \leq n}(X_i + Y_i) - X_n - Y_n \\ &\leq \Big(\max_{n_* \leq i \leq n} X_i - X_n\Big) + \Big(\max_{n_* \leq i \leq n} Y_i - Y_n\Big) = m_{TA}(X_n) + m_{TA}(Y_n)\,. \end{aligned}$$

Property [**4**]. m_{TA} *does not change if an investor adds cash to the portfolio.*

This property looks rather natural as only investing in the risky assets contributes to risk.

Note that Artzner et al. [4] have on the list of "desirable properties for measures of risk" the property of a measure to shrink by the amount of cash added. This is because Artzner et al. do not distinguish a measure of risk from a capital reserve/margin requirement: "the measures to the risk will be interpreted as the minimum extra cash the agent has to add to the risky position". The traditional approach (e.g., using σ, β, VaR, CVaR) does distinguish a measure of risk from a capital reserve/margin requirement. In other words, measuring risk is a separate act with respect to calculating a capital reserve/margin requirement.

Property [**5**]. Let X and Y be two instruments with prices $\{X_i\}$ and $\{Y_i\}$ respectively, and assume that

$$X_n = Y_n$$

(one can always consider rescaling $\{Y_i\}$, i.e., one compares investing in $\{X_i\}$ vs. investing in $\{zY_i\}$, where $z = X_n/Y_n$). If the last local minima of instruments X and Y has beed identified within the period $[n_*, ..., n]$ and

$$X_i \le Y_i \qquad (n_* \le i < n),$$

then for a holder of a long position

$$m_{TA}(X_n) \ge m_{TA}(Y_n). \tag{11}$$

If the last local maxima of instruments X and Y has beed identified within the period $[n_*, ..., n]$ and $X_i \ge Y_i$ as $n_* \le i < n$, then (11) holds for a holder of a short position. In other words, in both cases Y is less volatile than X.

Property [**6**]. m_{TA} *is not continuous.*

When a new local extremum is formed and recognised, it replaces the previous one in the calculation of m_{TA} (cf. Table 1).

Note that definition (10) of m_{TA} ignores the rate r_o of return of the risk-free asset (the same can be said about the definitions of a number of risk measures, e.g., the standard deviation, VaR and CVaR). One can adjust the definition of m_{TA} to take r_o into account: denote

$$m^*_{TA} = |P_n/(1+r_o)^{n-n_*} - P_{n_*}|\,,$$

where for a holder of a long position n_* is the time of the last local price minimum below P_n (for a holder of a short position n_* is the time of the last local price maximum above P_n). In particular, $m^*_{TA}(P^o_n) = 0$ if P^o_n is the risk-free asset.

The advantage of m_{TA} is that it changes considerably as the price changes. Traditional risk measures (estimated, say, from a one-year-long sample of preceding data) are rather static (see Example 5 below). Actually, they have to be static by the definition (it's not the standard deviation or VaR that changes but the values of our estimates of σ and VaR). By contrast, m_{TA} by the definition has to change as the price changes.

The disadvantage of m_{TA} is that it is not defined uniquely (like other tools of the Technical Analysis). That's because the nearest local minima/maxima of a fractal curve cannot be uniquely identified. One way to address the problem is to deal with price bars on a certain time scale (e.g., daily or weekly price bars). Note that some local extrema identified on a chart of daily price bars won't be present on weekly price charts. Moreover, some local extrema may be "too close" one to another. A practitioner may skip one of those "close" local extrema, effectively looking for minima/maxima of a group of "close" local extrema (i.e., zooming out on a particular chart area). For instance, in Example 5 we only consider local minima separated by two consecutive price bars. Thus, the time scale is the "tuning" parameter of m_{TA}. The presence of a tuning parameter is common in non-parametric estimation problems.

Example 2 (continued). We have calculated m_{TA} on the eve of the “Black Monday” (Table 1).

From mid-August 1987 the S&P500 index was mainly declining (see Fig. 2). Table 1 and Fig. 6 show steady decrease of the standard deviation of the S&P500 index since early September 1987 till 16.10.1987, suggesting lower level of risk (the conditional standard deviation was more dynamic, see Fig. 1 in McNeil & Frey (2000)).

The behaviour of m_{TA} is very different: Fig. 6 shows spikes of m_{TA} on the eve of the Black Monday. On 12.10.1987 the index has closed at 309.39, i.e., below the “last” local minimum identified on 21.09.1987 at 310.54 (“breakout” signal in TA terminology). This means m_{TA} has to be reset from $P_n - 310.54$ to $P_n - x_*$, where x_* is the last local minimum below P_n (x_* has been identified on 01.07.1987 at 302.94). Hence on 12.10.1987 $m_{TA} = 6.45$.

Date	S&P500	σ	m_{TA}
8.10.1987	314.16	28.55	3.62
9.10.1987	311.07	28.42	0.53
12.10.1987	309.39	28.27	6.45
13.10.1987	314.52	28.14	3.98
14.10.1987	305.23	27.99	2.29
15.10.1987	298.08	27.84	9.62
16.10.1987	282.7	27.67	4.49

Table 1: The S&P500 index, its standard deviations and m_{TA} on the eve of the Black Monday. While the standard deviation decreases, suggesting a reduction of the level of risk, m_{TA} exhibits spikes, indicating high level of risk.

On the previous trading day, 9.10.1987, the risk was $m_{TA} = 0.53$. The market has shown a 12-fold increase of the level of risk. The data was ringing a bell but would anyone hear it using the standard deviation? $\square$

While traditional measures of risk are too static, measure m_{TA} may appear too volatile. One may ask if there is anything in between?

As a possible candidate, we suggest the following measure:

$$m^+ \equiv m_c^+ = m_s + c m_{TA} \,. \tag{12}$$

Here m_s is a traditional (static) measure of risk (e.g., the standard deviation, VaR or CVaR) and $c \geq 0$ is a “tuning” parameter.

This approach suggests setting aside a somewhat larger amount of reserve, meaning smaller profits during quiet periods in return for higher financial stability during turbulent periods.

Example 5. In 2007 the S&P500 index had a level of resistance at 1406 (see Fig. 7). The presence of a level of resistance does not mean, of course, an increase of the level of risk – markets often bounce up from the levels of resistance (breaking through a level of resistance does however indicate a short-term downtrend and hence an increase of the level of risk).

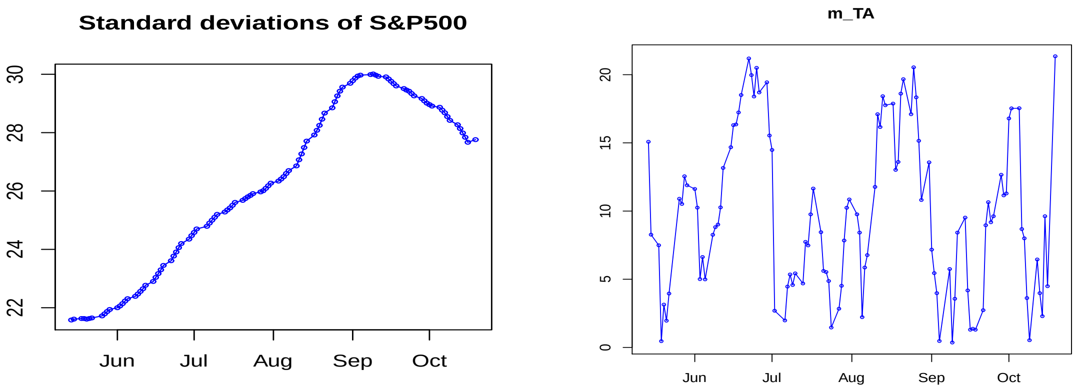

Figure 6: The standard deviations of the S&P500 index (left) and m_{TA} (right) in 14.05.1987–19.10.1987. The standard deviation decreases, suggesting lower risk; m_{TA} exhibits spikes, indicating high level of risk.

In 2007 the market did bounce up after touching the level of resistance. Eventually, after breaking through the level of resistance in January 2008, the index lost half of its value (see Fig. 8).

One would like to get a warning earlier, and measure m_s may come handy. It appears more suitable for the risk managers purposes, providing both a warning signal as well as a numerical measure of the increased level of risk.

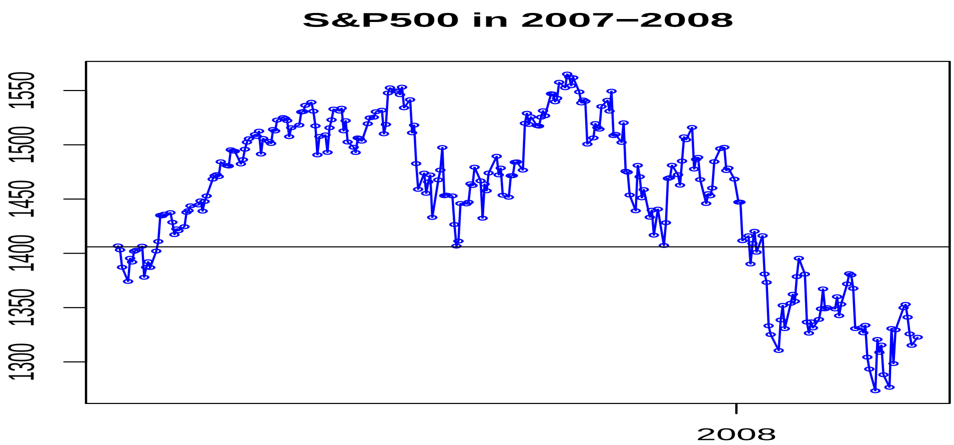

Figure 7: The S&P500 index in March 2007 – March 2008 had a level of resistance at 1406. After breaking through the level of resistance in January 2008 the index lost half of its value (see Fig. 8).

We have calculated m_{TA} for daily closing prices of the S&P500 index for the period from 27.11.2007 till 28.03.2008 (see Fig. 9). A local minimum is recognised if the daily closing price is smaller than the daily closing prices of two preceding and two consequent days.

Fig. 9 shows spikes of m_{TA} in January 2007; the highest value $m_{TA} = 135.07$ was achieved on 11.01.2008, indicating increased risk as well as providing a numerical measure of the level of risk. The S&P500 index was still at 1416.25 on 14.01.2008, providing time to act. After that it

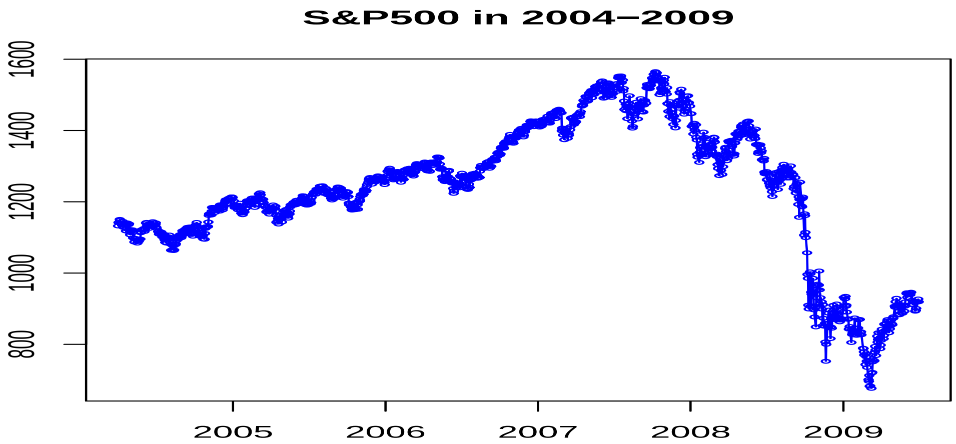

Figure 8: The S&P500 index in April 2004 – June 2009. After reaching 1565 in October 2007 the index fell to 678 in March 2009.

went on downtrend, briefly bounced back in May 2008 and then took a nose dive to reach 682.55 in March 2009 (see Fig. 8).

Fig. 9 presents also the combined measure (12) with the standard deviation as m_s and $c = 0.5$ for the period from 27.11.2007 till 31.03.2008. It is easy to see that the chart of measure m^+ is more dynamic than the chart of the standard deviation, yet less volatile than the chart of measure m_{TA} .

The peak values of m_{TA} and m^+ were achieved on 11.01.2008. The index closed at 1401.02 that day, breaking through the level of resistance. The preceding peak values were achieved on 08.01.2008, again with a break of the level of resistance. On both occasions the peak values of m^+ were more than twice the values of the recent local minima of m^+. The jumps of m_{TA} from through to peak in January 2008 were more than 100-fold. Thus, measures m_{TA} and m^+ signalled increased risk and provided a numerical measure of the level of risk.

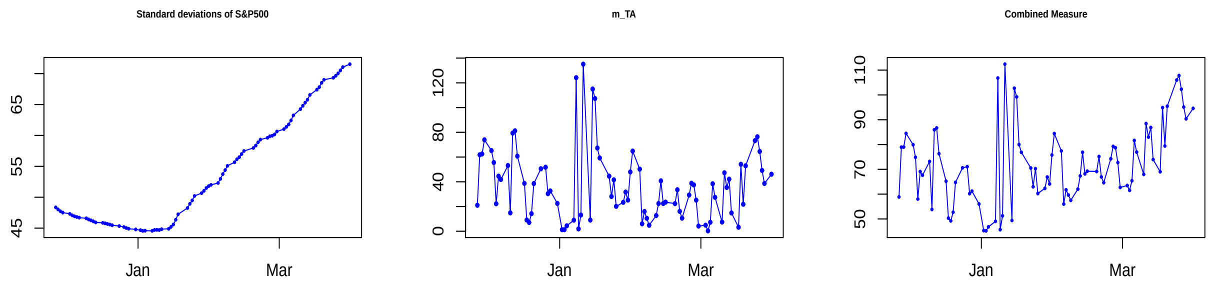

Figure 9: The standard deviation (left), measure m_{TA} (middle) and the combined measure m^+ (right) of the S&P500 index for the period November 2007 – March 2008.

What happened next suggests one would be right to act in January 2008 on the base of the indicator of increased risk provided by measures m_{TA} and m^+. □

6 Open problems and further research

TA approach is based on a large body of empirical observations (see, e.g., Elder (2002), Pring (2002), Williams (1994)) yet we are not aware of any comprehensive statistical study that would provide rigorous statistical background to the approach. Many questions appear to be left unanswered so far. We formulate below few natural questions.

(•) What can be said about the empirical distribution of the process $\{m_{TA}(P_n)\}$ for major stock indexes, commodities prices and major currency exchange rates?
The answer may be different for a holder of a long position and for a holder of a short position. It may be worth considering separately the periods of bull and bear trends.
(•) What is the distribution of the gain to risk ratio tested on major stock indexes, commodities prices and major currency exchange rates?
The answer may depend on the trigger signal, the stop/exit strategy and the chosen time scale (say, weekly or daily data).
(•) How often is a particular TA indicator profitable?
What is the empirical distribution of the rate of return when using a particular indicator?
Does the mean rate of return when using a particular indicator change considerably over time?

These questions are not easy to answer as the answers depend on the choice of the stop/exit strategy. For instance, an institutional investor often buys a stake and holds it, say, for two years. The natural question is:
(•) if an indicator is triggered, how often will the position be "in black" in two years time?

For a short-term investor the question can be formulated as follows:
(•) if a particular indicator has triggered opening a long position, the stop is at the level of the recent local minimum and the exit strategy is based on a specified signal, how often will the position be "in black"?

A considerable amount of empirical research is required to answer these questions as one would like to analyse a variety of markets as well as a variety of trigger signal/stop/exit strategy combinations.

The computational difficulty is caused by the fact that identifying local extrema of a near fractal is not straightforward: every price bar with a bar minimum lower than those of neighbouring bars is formally a local minimum. In fact, a local minimum of a fractal cannot be identified, and one needs a reasonable substitute to the local minimum. One approach to overcome the problem is to use data smoothing (e.g., moving average); variation with the degree of smoothing (the tuning parameter) may affect the results. Another approach is to use the minimum over an interval of size s, where s is a tuning parameter. The open problem is to suggest a procedure of choosing the tuning parameters. Besides, one would like to consider more than one time frame (e.g., daily and weekly price charts).

Conclusion. The paper overviews modern approaches to financial risk measurement, aiming to encourage interdisciplinary research on the topic of dynamic measures of risk.

We find that the traditional approach appears static from a short-term investor point of view.

The approach involving m_{TA} appears more suitable for dynamic risk measurement. Risk measure m_{TA} is truly dynamic (e.g., while the standard deviation was declining on the eve of the Black Monday, m_{TA} was indicating increased volatility). However, the definition of the dynamic risk measure involves local extrema of the price series. The computational difficulty of identifying local extrema is linked to the fact that price charts appear objects of fractal geometry. As a result, the approach involving m_{TA} lacks proper statistical scrutiny and requires considerable amount of empirical research.

Acknowledgments. The author is grateful to the Editor for helpful comments.

References

[1] Akram Q.F., Rime D. and Sarno L. (2006) Arbitrage in the foreign exchange market: turning on the microscope. — Swedish Institute for Financial Research, Research Report No 42.

[2] Alexander C. (2001) *Market models.* — New York: Wiley.

[3] Alexander C. (2008) *Value-at-Risk models.* — New York: Wiley.

[4] Artzner P., Delbaen F., Eber J.-M. and Heath D. (1999) Coherent measures of risk. — Math. Finance, v. 9, 203–228.

[5] Aronson D.R. (2006) Evidence-based Technical Analysis. — New Jersey: Wiley.

[6] Bingham N.H., Goldie C.M. and Teugels J.L. (1987) *Regular Variation.* — Cambridge: Cambridge University Press.

[7] Brock W., Lakonishok J and LeBaron B. (1992) Simple technical trading rules and the stochastic properties of stock returns. — J. Finance, v. 47, No 5, 1731–1764.

[8] Choffray J.-M. and de Mortanges C.P. (2015) Protecting assets under non-parametric market conditions. In: *Extreme events in finance* (F. Longin, ed.). Wiley.

[9] Clarke J., Jandik T. and Mandelker G. (2001) The efficient markets hypothesis. — In: *Expert financial planning: advice from industry leaders* (R.Arffa, ed.), 126–141. New York: Wiley.

[10] Embrechts P., Klüppelberg C. and Mikosch T. (1997) *Modelling extremal events for insurance and finance.* — Berlin: Springer.

[11] Elder A. (2002) *Come into my trading room.* — New York: Wiley.

[12] Elton E.J. and Gruber M.J. (1995) *Modern portfolio theory and investment analysis.* — New York: Wiley.

[13] Fama E.F. and Roll R. (1968) Some properties of symmetric stable distributions. — J. Amer. Statist. Assoc., v. 63, 817–836.

[14] Fama E.F. (1970) Efficient capital markets: a review of theory and empirical work. — J. Finance, v. 25, 383-417.

[15] Fraga Alves M.I. and Neves C. (2015) Extreme value theory: an introductory overview. — In: *Extreme value theory and its applications to finance and insurance* (F.Longin, ed.). Wiley.

[16] French K.R. (1980) Stock returns and the weekend effect. — J. Financ. Econom., v. 8, 55–69.

[17] Giraitis L., Leipus R. and Philipe A. (2006) A test for stationarity versus trends and unit roots for a wide class of dependent errors. — Econometric Theory, v. 22, No 6, 989–1029.

[18] Gnedenko B.V. (1943) Sur la distribution limite du terme maximum dune serie aleatoire. — Annals of Mathematics, 44, 423453.

[19] Goldie C.M. (1991) Implicit renewal theory and tails of solutions of random equations. — Ann. Appl. Probab., v. 1, 126–166.

[20] Gurrola-Perez P., Murphy D. (2015) Filtered historical simulation Value-at-Risk models and their competitors. – Working Paper No 525, Bank of England.

[21] Higson C. (2001) Did Enron's investors fool themselves? — Business Strat. Rev., v. 12, No 4, 1–6.

[22] Irwin S.H. and Park C.-H. (2007) What do we know about the profitability of Technical Analysis? — J. Economic Surveys, v. 21, No 4, 786–826.

[23] Lo A.W., Mamaysky H., Wang J. (2000) Foundations of Technical Analysis: computational algorithms, statistical inference and empirical implementation. — J. Finance, v. 55, No 4, 1705–1765.

[24] Longin F. (1996) The asymptotic distribution of extreme stock market returns. — J. Business, v. 69, No 3, 383–408.

[25] Longin F. (2001) Beyond the VaR. — J. Derivatives, No 8, 36–48.

[26] Luenberger D.G. (1997) *Investment science.* — Oxford: OUP. ISBN10: 0195108094

[27] Mandelbrot B.B. (1963) New methods in statistical economics. — J. Political Economy, v. 71, 421–440.

[28] Markovich N. (2007) *Nonparametric analysis of univariate heavy-tailed data.* — Chichester: Wiley.

[29] Markowitz H.M. (1952) Portfolio selection. — J. Finance, v. 7, No 2, 77–91.

[30] Matthys G. and Beirlant J. (2001) Extreme quantile estimation for heavy-tailed distributions. — Universitair Centrum voor Statistiek, Katholieke Universiteit Leuven: preprint.

[31] McNeil A.J. (1998) On extremes and crashes. — Risk, v. 11, p. 99–104.

[32] McNeil A.J. and Frey R. (2000) Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. — J. Empirical. Finance, v. 7, 271–300.

[33] Nadarajah S. and Chan S. (2015) Estimation methods for Value-at-Risk. — In: *Extreme value theory and its applications to finance and insurance* (F.Longin, ed.). Wiley.

[34] Novak S.Y. (2002) Inference on heavy tails from dependent data. — Siberian Adv. Math., v. 12, No 2, 73–96.

[35] Novak S.Y. and Beirlant J. (2006) The magnitude of a market crash can be predicted. — J. Banking & Finance, v. 30, 453–462.

[36] Novak S.Y., Dalla V. and Giraitis L. (2007) Evaluating currency risk in emerging markets. — Acta Appl. Math., v. 97, 163–175.

[37] Novak S.Y. (2011) *Extreme value methods with applications to finance.* — London: CRC. ISBN: 978-1-43983-574-6.

[38] Osler C. and Chang K. (1995) Head and shoulders: not just a flaky pattern. — Staff Report No 4, Federal Reserve Bank of New York.

[39] Pflug G.C. (2000) Some remarks on the Value–at–Risk and the conditional Value-at-Risk. — In: *Probabilistic constrained optimization*, 272–281. Kluwer: Netherlands.

[40] Poser S.W. (2003) *Applying Elliott wave theory profitably.* — New York: Wiley.

[41] Prechter R.R. and Parker W.D. (2007) The financial/economic dichotomy in social behavioral dynamics: the socionomic perspective. — J. Behavioral Finance, v. 8, No 2, 84–108.

[42] Pring M.J. (2002) Technical analysis explained. — New York: McGraw Hill. ISBN 0-07-138193-7.

[43] Rockafellar R.T. and Uryasev S. (1999) Optimization of conditional Value-at-Risk. — http://www.ise.ufl.edu/uryasev

[44] Seneta E. (1976) *Regularly varying functions.* — Lecture Notes Math., v. 508. Berlin: Springer.

[45] Wells D., Pretzlik C. and Wighton D. (2004) The balancing act that is Value-at-Risk. — Financial Times, 25.03.2004.

[46] Williams B.M. (1994) *Trading chaos: applying expert techniques to maximise your profits.* — New York: Wiley.

[47] Wilder J.W. (1978) *New concepts in technical trading systems.* — Trend Research. ISBN 0-89459-0278.

In: Extreme Events in Finance (F. Longin, ed.), 215–237. Handbook Series in Financial Engineering and Econometrics. Wiley, 2017. ISBN-10: 1118650190.