DIAMAN RATIO

D. Bernardi[*], R. Bertelli[†‡]

January 14, 2014

Abstract

The statistical indicators for the evaluation of the efficiency of a financial instrument are almost all based on the ratio between mean and variance or are, in any case, linked to the assumption that distribution of returns is normal. In most cases, the effective time-based sequence of the returns (hypothesis of independence) is not taken into consideration. In this work, after recalling the limits of the mean variance approach for the purposes of fund selection, a new indicator is proposed to measure the risk-adjusted performance. The DIAMAN Ratio considers the sequence of returns and is based on a definition of risk that is consistent with some well-established results from behavioural finance. The DIAMAN Ratio can be interpreted as an indicator of the persistence of returns: it analyses the strength of the trend (expected return) and the ability of the financial instrument to move around its own trend (risk).

[*]Founder and CEO, DIAMAN SIM S.p.A., Italy

[†]Associate Professor of Banking and Finance, Univeristy of Siena, Italy

[‡]We are grateful to F. Corielli and F. Canella for their suggestions and support.

1. Introduction: mean/variance approach and fund selection

"Most performance measures are computed using historic data but justified on the basis of predicted relationship. Practical implementations use ex post results while theoretical discussions focus on ex ante values. Implicitly or explicitly, it is assumed that historic results have at least some predictive ability"[1].

This work proposes an unusual approach to the use of the historic time series of performance in selecting financial portfolios, with particular focus on those characterised by active management. The historic time series of a fund's performance includes the work of the fund manager who constantly modifies the exposure to risk and the composition of the portfolio's risks. The distribution of the returns on the managed portfolio is, therefore, a "transformation" of the original return distribution of the underlying financial assets. Clearly, this transformation can be more or less intense, based on a variety of parameters. Period after period, the probability distribution from which the performance is "extracted", presents characteristics which vary considerably over time. We cannot exclude a relationship, even if it is unstable, between the form of distribution and past events, such as for example the determination of the portfolio's performance or of some of its constituents of the past period. In general, the power and the limits of the mean variance approach can be inferred by these simple words of Sharpe (1994)[1]: "we build on Markowitz's mean-variance paradigm, which assumes that the mean and the standard deviation of the distribution of one-period performance are sufficient statistics for evaluating the prospects of an investment portfolio".

As is well-known, relevant literature has heavily criticised the standard deviation as a measure (or rather as the only measure) of the portfolio risk, due both to endogenous reasons (the assumption of normality in the return distribution), and exogenous factors (investors and portfolio managers' risk perception according to behavioural finance).

In respect to the first aspect, see for example Eling – Schuhmacher (2006)[2]; for the second aspect Fisher - Statman (1999) [3], Estrada (2008) [4]. Moreover, the literature that focused on risk-adjusted performance measures in the field or of hedge funds or of "great investors" [5], has particularly emphasised the ability of the portfolio manager to "transform" the original return distribution of the underlying financial assets using appropriate skills and techniques, thus generating asymmetric return distributions which do not exist in nature [6]. In using risk-adjusted performance indicators for selection purposes, there is also another problem that shifts attention from the single-period form of return distribution to the result of a series of single-period performances linked over time (for a critique of Sharpe's index from this point of view, see Lo, 2002) [7]).

As is well-known, the fundamental hypothesis is that each determination is independent from past ones. This means that return distributions – regardless of the form - are independent. In other words, there is no serial correlation between returns. This is a very important hypothesis, but barely realistic. From the viewpoint of asset managers, it is absurd in certain respects. Aside from the numerous theoretical and practical "annualisation" problems of risk and risk-adjusted performance indicators, the hypothesis allows work to be done in the sphere of "normality", even when the single-period distributions are not "normal". In fact, thanks to the "central limit theorem", the sum of independent random variables (of any distribution) will be distributed normally - provided that a "certain" stability in terms of mean and variance is guaranteed. This brings us partially back to the starting point. And in any case, we can see that the independence of the return sequence is crucial. Indeed, by extending the time spans, it allows us to rely on a progressive log-normal distribution of the capital random variable, that results in a progressive "normality" of the return random variable (logarithmic). For a simple application, see Bertelli (1999) [8].

The topic of the serial correlation of returns was recently brought up again by some studies on the successes of the "momentum strategy" in portfolio management. These show that the independence of the return sequence is not (and has never been) a realistic hypothesis (for example [9]; [10]; [11]; [12]). In brief, the use of risk-adjusted performance indicators, for the purposes of fund selection, which assume a mean variance approach in a context: a) of stability of the probability density function of single-period returns; b) of mean and variance stability in distributions; c) of serial independence of returns, raises several doubts. This is especially so when the aim is to pick up significant and persistent differences in portfolio management ability. It is not only a general lack of confidence in the «predicting» abilities of historic performance series. Both the theorist and the practitioner feel uncomfortable due to a sensation of incompleteness and inadequacy in evaluation, that leads to constructing analytically detailed tables and graphs full of statistical indicators, but often with scarce operational value; or – worse – that are subject to brutal and conscious « data mining ».

2. An example of mean variance fund selection

The simple intuition leading to the proposal of a new risk-adjusted performance indicator arose due to the dissatisfaction with the rankings produced with the most classical and most commonly used (and also most criticized) indicator: the Sharpe Ratio. The origin of such unease can be seen by considering Figure 1.

Figure 1: **Fund Selection: an example**

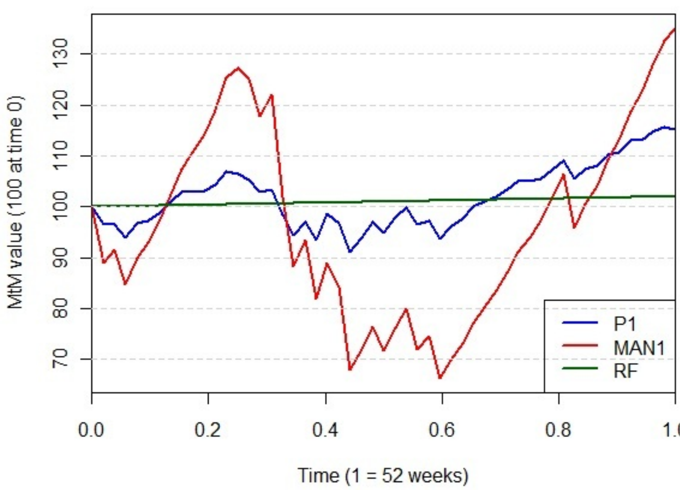

P1 is the MSCI USA in the last 52 weeks. RF is the capitalization of an interest rate of about 2%. MAN1 is the performance of a managed portfolio that has P1 as its operational benchmark. On the x-axis, the time in fractions of a year is shown. On the y-axis is the trend in the mark to market value reset to 100.

Did the operator do better than his benchmark? In terms of performance, absolutely. It is obvious that the best performance was «offset» by higher risk. The Sharpe Ratio suggests, however, that the manager was more efficient than the reference benchmark (see Table 1).

Table 1: **P1 vs MAN1 (Sharpe Ratio)**

	P1	**MAN1**
Ann. Return	15,23%	35,05%
Ann. Dev. St.	17,35%	49,30%
Sharpe Ratio	0,1094	0,1148

For certain aspects, by accepting the theoretical approach of the index, we are forced to agree with Sharpe. If we build a portfolio made up by 35% MAN1 and we invest the rest in monetary funds (2%), we obtain a portfolio with the same volatility as the index and a higher return. The greater volatility is more than compensated by the performance (Table 2).

Table 2: **P1 vs MAN1 (Sharpe Ratio)**

	P1	**MAN1**	**%MAN1**
Ann. Return	15,23%	35,05%	15,92%
Ann. Dev. St.	17,35%	49,30%	17,25%
Sharpe Ratio	0,1094	0,1148	0,1148

And yet, the asset manager would continue to deem it inappropriate to include MAN1 in the portfolio, even in light of this data.

This is because the problem is not volatility (or at least it is not only volatility), but the possibility that an unexpected and bad drawdown occurs. The return sequence has a bearing in the evaluation: « Rule no. 1: Never lose money. Rule No.2: Never forget rule No.1", that's what hedge fund managers keep saying. And in fact, the asset manager is supported by some asymmetric risk indicators, such as the Calmar Ratio (Young, 1991) and Sortino index (Sortino - van derMeer, 1991), that in fact raise some doubts (the data is collected in Table 3).

Table 3: **Some indicators, some statistics**

	P1	**MAN1**	**%MAN1**
Ann. Return	15,23%	35,05%	15,92%
Ann. Dev. St.	17,35%	49,30%	17,25%
Sharpe Ratio	0,1094	0,1148	0,1148
Maxx DD	-14,65%	-47,83%	-18,38%
Sortino Index	0,084463	0,071109	0,071528
Calmar Ratio	0,017966	0,016413	0,014952

It is not possible to construct a portfolio consisting of MAN 1 and monetary funds where performance is higher than the parity index for maxDD or downside deviation. Let's take a closer look. The perplexities are absolutely legitimate. We know MAN1 very well.

By analysing the performance of MAN1 according to the UP/DOWN Market Model, we realize that our « efficient » portfolio manager has a BETA UP equal to 1, a BETA DOWN equal to 4 and an ALPHA equal to 0.03 in relation to MSCI USA (P1). In brief, the manager is exposed to very significant drawdowns in the presence of repeated negative returns on the reference market. We can say that his ALPHA is pricy in terms of exposure to market risk.

3. The intuition of the Diaman Ratio

In fund selection the problem is to correctly define risk. Within a context of "loss aversion" ([15]; [16]; [17]), the symmetry of returns is not important, because an asymmetric evaluation is provided in any case and the return sequence matters, because each increase in losses has double the "weight" of each increase in profits. If the temporal return sequence matters, the definition of risk must start from the trend in invested capital, which is the consequence of the typical sequence of single-period returns. The trend in the mark-to-market value of the capital invested depends not only on the mean and variance of the returns, but also on their possible and variable self-correlation. In evaluating the action of a portfolio's manager, we must take in consideration his reaction in the case of unexpected losses and the potential opportunity to make unexpected profits (in both cases with losses and profits higher than the established goals).

In brief, the manager will tend to deviate as little as possible from an ideal trend in the managed assets, represented by the expected portfolio return or its objective performance. This view is absolutely compatible with the desires of the investor who accepts the risk only provided that he perceives the asset trend as sufficiently « stable ». The concept of « stability » in the trend of invested assets appears a correct representation of risk - or at least it takes on board an important aspect of what the investor perceives as risk in the financial investment.

Therefore, if the investor and operator have in mind an expected return, its risk-free projection (in other words the curve of the trend in the capitalized assets against the expected return over time) represents the ideal situation. The investment portfolio will be more appreciated, the more it is able to stay « on track», i.e. not deviate "too much" from the desired asset trend, on an equal forecast return basis. Obviously, a certain deviation compared to the expected return trend is accepted, provided that the expected trend leads to higher asset values.

This simple intuition is represented in Figure 2.

The three straight lines in Figure 2 represent the (logarithmic) risk-free trend in assets at the forecast return. It is obvious that a higher forecast return is better without risk. In order to pursue this objective, nonetheless, the investor must accept « deviations » compared to the ideal trend.

The DR's objective is to provide a measure of this type of risk and a « risk-adjusted performance », which can be readily compared with the risk-free rate.

The very simple choice has fallen on a linear regression model between the mark to market value of the invested assets and time. The angular coefficient of the straight line can be interpreted as expected return (β); the linear determination coefficient can be interpreted as risk (R^2).

The risk adjusted performance (the Diaman Ratio) is simply the result DR $= \beta \cdot R^2$. It should be noted that in the case of Figure 2, the three alternatives are as follows (table 4):

The DR gives indications of preference which are the opposite of the traditional Sharpe's index. In fact, in the two at-risk cases, the ability of the manager to remain « on track » is very high and constant. The greater expected performance of asset A makes it the preferred choice.

Figure 2: **Example**

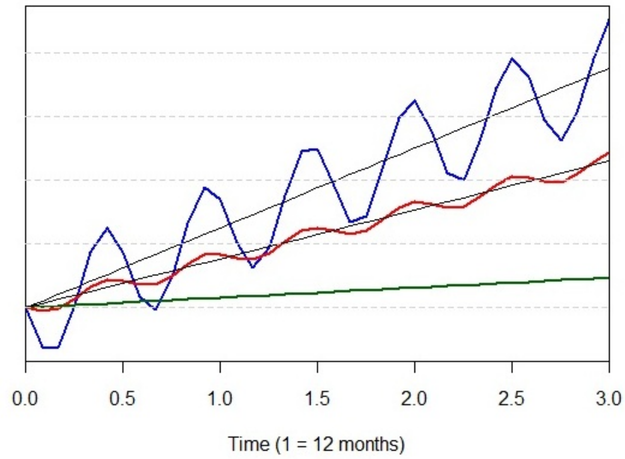

Table 4: **Sharpe Ratio and Diaman Ratio**

	A	**B**	**RF**
Mean	21,70%	11,62%	2,40%
St. Dev.	8,59%	1,43%	0,00%
Sharpe Ratio	0,1872	0,5364	
b	19,76%	11,85%	2,40%
R^2	0,927	0,994	1,000
Diaman Ratio	0,183	0,118	0,024

In the example presented in the previous paragraph (see Figure 2), by applying the DR, the asset manager would have not chosen MAN1, despite the higher Sharpe's index (Table 5).

Table 5: **P1 vs MAN1 (Diaman Ratio)**

	P1	**MAN1**	**%MAN1**
b	12,26%	5,89%	6,83%
R^2	0,3638	0,0080	0,0877
Diaman Ratio	4,46%	0,05%	0,60%

4. The Diaman Ratio

Let's suppose $P = (p_1, p_2, p_3, \dots, p_n)$ for the historical series of weekly logarithmic prices of a financial instrument and $t = (0, {}^1\!/\!_f, {}^2\!/\!_f, \dots, {}^{(n-1)}\!/\!_f)$ for the historic time series time where $f = 52$ and n is the length of the historical series.

The Diaman Ratio is calculated thus:

$$DR = \beta \cdot R^2$$

where

- β is the estimated value of the regression model $P_i = \beta \cdot t_i + \alpha + \varepsilon_i$
- R^2 is the calculation coefficient associated with the regression.

The use of logarithmic historic series is important for more accurate calculation, since the logarithm acts on the variability of the series and addresses the scale effect which is instead shown by linear series.

The estimated β is simply the annual logarithm growth rate of the historic series. To obtain the linear growth rate, calculate $e^{\beta} - 1$. This is even truer for the extreme case of a historic series growing at a constant rate. Indeed, if we calculate theβ of the historic series with these characteristics, the beta value will be equal to the growth rate value.

R^2 is included in the formula in order to keep in mind the ability of the predictor (time) to predict the values of the variable. If $R^2 = 1$, the series is growing monotone, with growth rate β, if $R^2 = 0$ the regression model is not well specified and therefore, there is too much variability around the estimatedβ.

Parameter f can also assume value 12 or 260 should the historic price series be monthly or daily. Parameter f has the purpose of allowing β to be read as the annual growth rate.

If we evaluate the DIAMAN Ratio in more detail, we can observe that it is nothing but the ratio between the co-variance of the two variables and the product of the variances.

$$DR = \beta \cdot R^2 = \frac{\sum (t - \bar{t})(p - \bar{p})}{\sum (t - \bar{t})^2} \cdot \frac{[\sum (t - \bar{t})(p - \bar{p})]^2}{\sum (t - \bar{t})^2 \cdot \sum (p - \bar{p})^2}$$

By simplifying the formula and considering that time is a growing determinate series, the ratio becomes:

$$DR = k_n \cdot \frac{[\sum (t - \bar{t})(p - \bar{p})]^3}{\sum (p - \bar{p})^2}$$

where

- $k_n = \frac{1}{[\sum (t - \bar{t})^2]^2} = \left(\frac{12 \cdot f^2}{(n-1) \cdot n \cdot (n+1)} \right)^2$
- $n =$ number of observations
- $f = 52$ for weekly time series

As can be seen from the formula, the innovative characteristic of this indicator is, therefore, the relation of the historic series to time, disregarding the concepts of return distribution, which present limits in terms of estimates and reliability.

Another characteristic is to separate the efficiency evaluation from the precise performance data, which is typically very unstable, in favour of a regression estimate that keeps indirect account of the variability in returns.

The third characteristic is to use an estimate parameter R^2 to evaluate the reliability of the result, since there could be periods in which the estimate is incorrect or barely reliable.

In conclusion, the last characteristic is that the DIAMAN Ratio can also be calculated to evaluate the historic risk-free series, since it also allows such series to be calculated. In this case $Rf = e^{DR} - 1 = e^{\beta} - 1$, since $R^2 = 1$. Therefore, if we wish to look at the analysed historic series compared to the risk-free series, we must perform the following operation: $DR_{NET} = DR - \ln(Rf + 1)$. Instead, if the rate changes in the analysed period, the historic series of the capitalized returns of risk-free rates can be considered and we can calculate: $DR_{NET} = DR_{SERIE} - DR_{RISKFREE}$.

5. Characteristics of the Diaman Ratio

Some practical characteristics of the indicator are analysed in this paragraph.

a) The DIAMAN Ratio is obtained by a regression of the historic price series against time, therefore, for the time span observed on an equal return and return volatility basis, the DIAMAN Ratio differentiates between the trajectories. The Sharpe Ratio, however, as calculated, does not take into consideration the temporal sequence of returns. In fact, it is possible to build n! historic series of n length with the same return and variance, but with totally different trajectories and therefore characteristics (for example the drawdown). Stating that all n! series can be considered similar in terms of efficiency would be a forced assumption.

For the table and graph below, we have taken an historic series and ordered the returns randomly, in two different ways, in order to obtain two trajectories which are very different from each other, but with the same mean and variance.

Figure 3: **Plot Series 1 and 2**

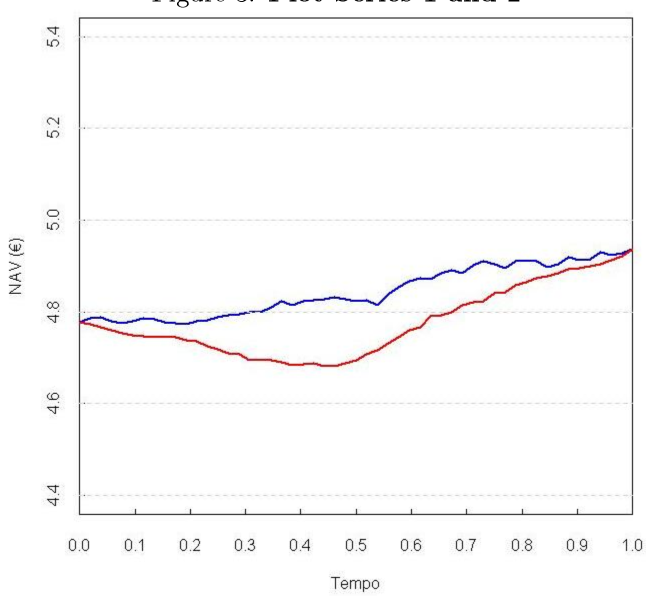

The result is that the Sharpe Ratio gives the same value for both historic series, while the DIAMAN Ratio is different for both series.

Table 6: **Example**

	Series 1	Series 2
Sharpe Ratio	2,608	2,608
Diaman Ratio	16,84%	10,68%

b) The DIAMAN Ratio is not linked to the risk-free rate, even if it can be easily introduced into the calculation (see previous paragraph). In fact, the values change due to the effect of a constant, but the classification remains the same. The introduction of the risk-free rate in calculating the Sharpe Ratio, on the other hand, leads to a non-unique evaluation of the final result. The use of a value instead of another for the risk-free rate may lead to different preferences within the same basket of instruments. The introduction of the risk-free rate in the ratio's numerator presents a fairly significant problem of arbitrariness that produces very significant changes in the indicator's results.

Table 7: **Fund Ranking**

			Sharpe Ratio		**Ranking**	
	Return	**St. Dev.**	**Rf = 2%**	**Rf = 4%**	**Rf = 2%**	**Rf = 4%**
Fund A	8,06%	3,10%	1,953	1,308	1	6
Fund B	27,44%	13,04%	1,951	1,798	2	1
Fund C	15,75%	7,08%	1,943	1,661	3	3
Fund D	4,66%	1,37%	1,940	0,484	4	10
Fund E	6,58%	2,39%	1,917	1,079	5	8
Fund F	11,76%	5,24%	1,862	1,480	6	4
Fund G	10,55%	4,61%	1,855	1,421	7	5
Fund H	6,45%	2,42%	1,837	1,011	8	9
Fund I	8,66%	3,71%	1,798	1,258	9	7
Fund L	40,74%	21,70%	1,785	1,693	10	2

The table calculates the Sharpe Ratios for 10 funds with Rf= 2% (first column) and Rf= 4% (second column). The second part of the table shows the changes in the classification of each fund.

c) The third characteristic concerns the quantity and quality of data. The change in the frequency of use of the data used to evaluate a series can lead to ranking results which are very different from each other using weekly data instead of daily data. The Sharpe Ratio is very sensitive to the calculation method, therefore if daily or weekly data is used rather than monthly data, the results change remarkably, as shown by the graphs below. The DIAMAN Ratio, on the other hand, despite changing in terms of results and ranking, is much more stable in the evaluation, although using short time spans (a year).

Figure 4: **Boxplot**

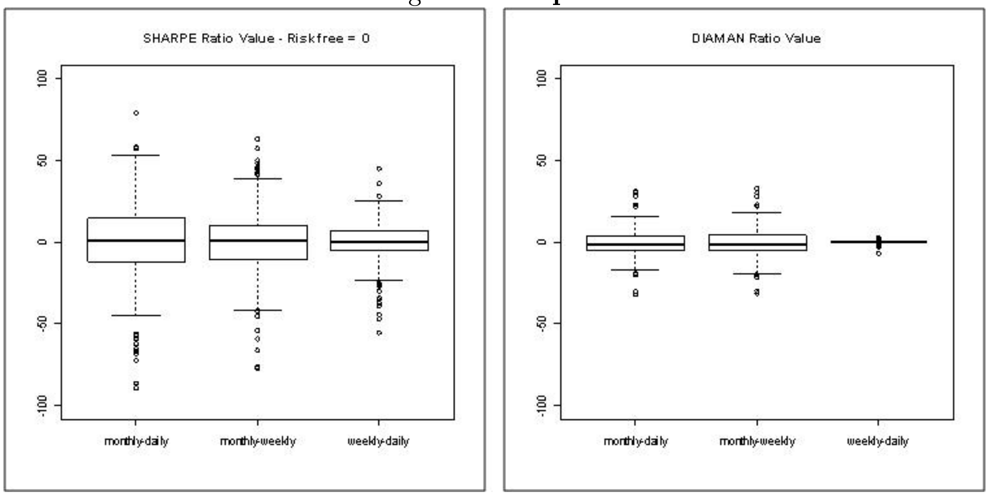

d) The DIAMAN Ratio is able to estimate both positive and negative trends, but has difficulties in the presence of changes of direction and non-linear historic series. This leads to the fact that when interpreting the value obtained, the indicator's values close to 0 should not be interpreted as better or worse compared to the positive or negative values, but as a non-reliable growth rate and, therefore, cannot be taken into consideration.

Figure 5: **Three series and related Diaman Ratio**

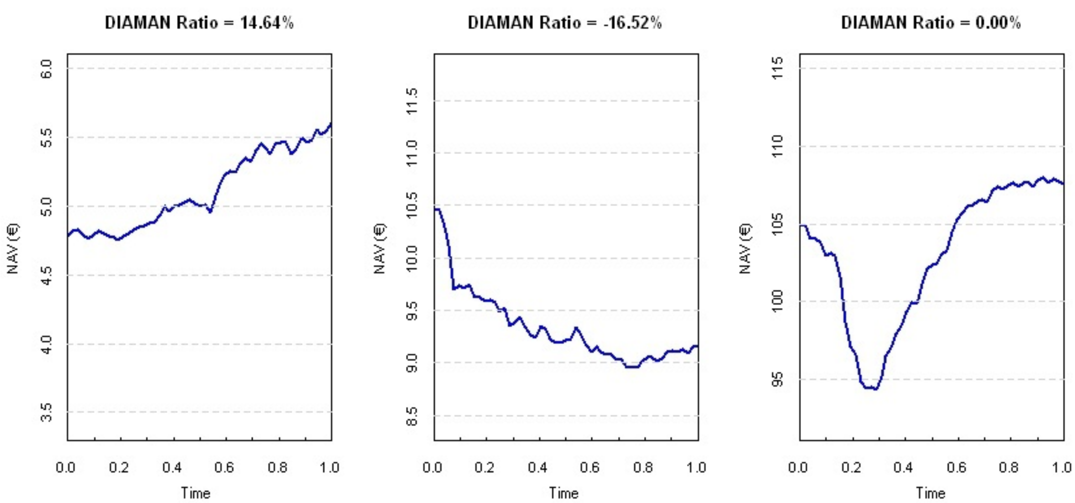

6. Practical case

The aim of the DIAMAN Ratio is to verify the continuity of the returns of a financial instrument. Absolute and total return funds set stable and constant performance over time, regardless of market conditions, as the typical objective of management. It is, therefore, an ideal style of management to assess the indicator's ability to select the best funds in the sector.

The database used for the analyses was kindly provided by FIDA Finanza Dati Analisi srl. The database consists of 497 flexible funds and bond funds with total or absolute return approaches. The timeframe analysed is 31/12/2003 – 31/08/2011. An analysis was made of both funds which were closed in the period and those that were opened; this means that the results are survivorship bias-free. The number of funds in the basket is therefore variable and goes from a minimum of around 200 funds to a maximum of around 400.

We tested the DIAMAN Ratio indicator by simulating a trading system. The main features of the trading system are as follows:

- Portfolio of 10 funds.
- Equal weighting of funds in the portfolio.
- Rebalancing of the portfolio on the first day of each month.
- The criterion for selecting funds is the maximum DIAMAN Ratio.
- Transaction fees are not considered since it is now common practice not to apply them to institutional operators.
- One day slippage to take account of the time between decision and operation.

Figure 6: **Back-test Performance vs Average Funds Basket**

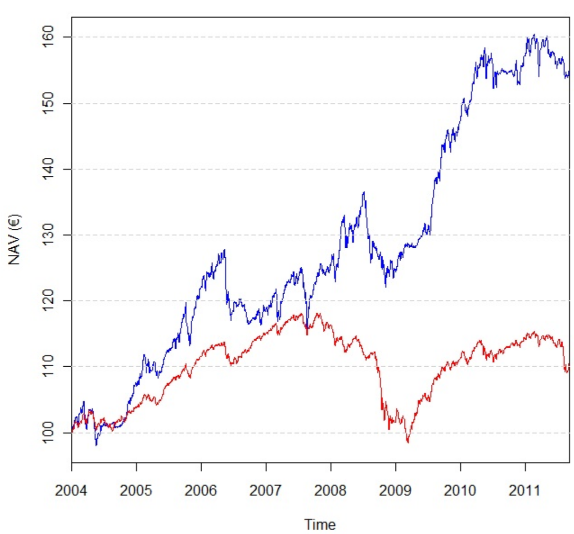

	Back-test	**Basket**
Return	54.65%	10.25%
Volatility*	6.53%	3.17%
MaxDD*	-10.64%	-16.65%
Ulcer Index*	4.01%	5.24%
Diaman Ratio*	5.27%	0.64%

* Method of calculation:

Volatility: standard deviation of return

MaxxDD: maximum loss recored

Ulcer: as described at http://en.wikipedia.org/wiki/Ulcer_Index

In order to have a better understanding of the validity of the indicator, the model analysed was not optimised but based only on arbitrary rules. The DIAMAN Ratio was calculated on the basis of a 6 months period, in other words on 130 observations.

Figure 6 shows the historical series of the trading system compared with the benchmark, in the sense of the average returns on the funds in the database. Besides the graph, we also provide a table summarising some performance and risk indicators of the example given.

From the results and the graph we can see that final extra-performance is equal to 44.41%. Globally volatility is greater, but it must also be considered that the portfolio consists of just 10 funds and, despite this, the maximum loss is lower than that of the basket (-10.64% compared to -16.65%). Also the ulcer index indicates greater consistency of return. The DIAMAN Ratio for the two series is 5.27% for the back–test and 0.64% for the fund average. To complete the previous analysis, we calculated the results of the back-tests as the parameter used to calculate the DIAMAN Ratio varied. Figure 7 shows the performance results and the maximum drawdown for the DIAMAN Ratio parameter values between 5 and 200 observations.

Figure 7: **Performance and MaxDD in relation to DR used**

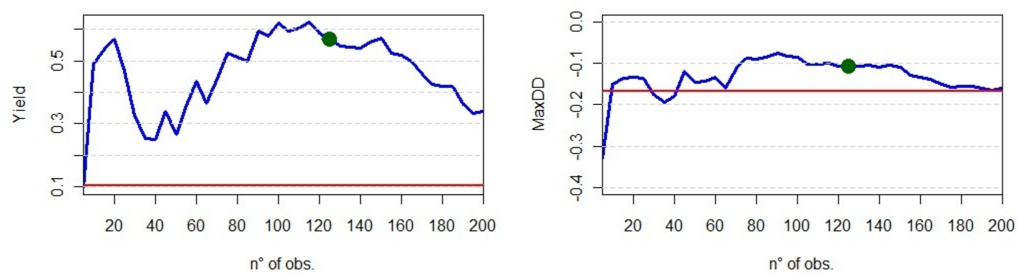

In the graphs the red line identifies the return (10.25%) and the maximum drawdown (-16.65%) recorded by the basket; the green point identifies the results relating to the back-test set out above; the blue line shows the results obtained as the parameter of the indicator varied. From the graph it is possible to see that, for indicator parameter values which range between 70 and 150, there is broad consistency of results both in terms of performance and in terms of risk. Such a result curve supports the hypothesis that the DIAMAN Ratio is a valid indicator to select funds which use an absolute return strategy.

7. Conclusion

We can state that the DIAMAN Ratio is a valid instrument to evaluate the efficiency of a financial instrument, especially to compare instruments in the same asset class, since it is able to evaluate correctly the returns' capacity to endure and the consistency in performance. For this reason, it is an extremely valid instrument, especially to evaluate absolute return instruments, in other words those which aim at a constant performance over time. Its unique character lies in the linkage of the returns over time and its simple composition allows more precise and unambiguous evaluation indications of a historic series to be obtained. With the necessary caveats, this indicator can also be used to estimate expected returns over time, and thus to optimize the portfolio, or to create tactical model portfolios.

References

[1] Sharpe W.F. (1994). “The Sharpe ratio”. *The Journal of Portfolio Management.*

[2] Eling M. and Schuhmacher F. (2006). “Does the choice of performance measure influence the evaluation of hedge Funds?”. *Working Papers on Risk Management and Insurance no. 29.* University of St. Gallen

[3] Fisher K. L. and Statman M. (1999). “A Behavioral framework for time diversification”. *Financial Analysts Journal.*

[4] Estrada J. (2009). “The Gain-Loss Spread: A New and Intuitive Measure of Risk”. *Journal of Applied Corporate Finance,* 104-114.

[5] Ziemba W. T. (2005). “The symmetric Downside Risk Sharpe ratio”. *The Journal of Portfolio Management.*

[6] Kat H. M. and Palaro H. P. (2005). “Who needs hedge funds? A Copula-Based Approach to Hedge Fund Return Replication”. *Alternative Investment Research Centre Working Paper No. 27.* Cass Business School research Paper, City University

[7] Lo A. W. (2002). “The statistics of Sharpe ratios”. *Finanncial Analists Journal.*

[8] Bertelli, R. (1999). “La gestione del risparmio previdenziale: il ruolo della banca-assicurazione”. *Materiali del Servizio Studi ASSBANK, n. 24*

[9] De Bondt W.F.M. and Thaler R. (1985). “Does the stock market overreact?”. *Journal of Finance, 40:793-80.*

[10] Asness C. S., Moskowitz T. J., Pedersen L. H. (2009). “Value and Momentum Everywhere”. *The Journal of Finance.*

[11] Fuertes A. M., Miffre J., Rallis G. (2010). “Tactical Allocation in Commodity Futures Markets: Combining Momentum and Term Structure Signals”. *Available at SSRN: http://ssrn.com/abstract=1127213*

[12] Dimson E., Marsh P., Staunton M. (2011). *Credit Suisse Global Investment Returns Sourcebook 2010.*

[13] Young T. W. (1991). “Calmar Ratio: A Smoother Tool”. *Futures, 20 (1), 40.*

[14] Sortino F. A. and van der Meer R. (1991). “Downside Risk”. *Journal of Portfolio Management, 17 (Spring),* 27–31.

[15] Kahneman D. and Tversky A. (1979). “Prospect theory: An analysis of decision under risk”. *Econometrica, 47,* 263–291.

[16] Shefrin H. and M. Statman, 1985. “The disposition to sell winners too early and ride losers too long: Theory and evidence”. *Journal of Finance, 40*, 777-791.

[17] Benartzi, Shlomo and Richard H. Thaler, 1995. "Myopic Loss-Aversion and the Equity Premium Puzzle". *Quarterly Journal of Economics 110.1*, pp. 73-92.