The Zephyr K-Ratio

By Thomas Becker, Ph.D.
Zephyr Associates, Inc.

Introduction

In 1996, Lars Kestner introduced the K-Ratio as a complement to the Sharpe Ratio. With Version 8.1, Zephyr Associates makes the K-Ratio available to StyleADVISOR users. This article explains the use, the meaning, and the exact mathematical definition of the ratio.

In his 2003 book [1], Kestner made a modification to his original K-Ratio. For reasons that will be explained in Section 4 below, Zephyr strongly recommends using the original K-Ratio of 1996. This poses a naming problem: The term "K-Ratio" would, by default, refer to Kestner's own most recent version. To express Zephyr's preference for the original version, we refer to the ratio that we use as the "Zephyr K-Ratio." This should not be interpreted as Zephyr taking credit for the ratio. What we use is Kestner's original 1996 definition. We merely present an argument for preferring the original version over the modified one.

Section 1 gives a brief overview of the Zephyr K-Ratio and explains why an investor may want to complement the Sharpe Ratio with the Zephyr K-Ratio. No mathematical prerequisites are required.

Sections 2 goes into more detail about the Zephyr K-Ratio. It gives a largely visual explanation of the ratio's meaning. Still, no serious mathematics is required.

Section 3 presents the exact definition of the Zephyr K-Ratio. While no higher mathematics is required, it is assumed that the reader is rather mathematically inclined.

In Section 4, we give our argument for preferring Kestner's original K-Ratio over his 2003 modification. Exploring these arguments can actually be helpful to gain more insight into the meaning of the K-Ratio. We therefore recommend reading this section even if the reader is already convinced that the Zephyr K-Ratio is the one that should be used. The mathematical prerequisites for this Section are mild.

1. An End User's Perspective of the Zephyr K-Ratio

Figure 1 below shows the cumulative return of three portfolios. It is fair to say that these three portfolios are likely to have different appeal to different investors. Most investors would probably prefer the portfolio that is labeled "Consistent" over the other two: it increases in value over time, and it does so consistently. That is reflected in the fact that the cumulative return graph of this portfolio slopes upward, and, on a logarithmic scale, it very closely resembles a straight line. The Zephyr K-Ratio measures precisely that behavior: the more the cumulative return graph resembles a straight line, the higher the

K-Ratio, and the more the cumulative return graph tends upward, the higher the K-Ratio. For example, the portfolio "Consistent" in Figure 1, whose cumulative return graph is an almost perfect straight line, has a Zephyr K-Ratio of 180.13, as opposed to 6.43 for the other two portfolios.

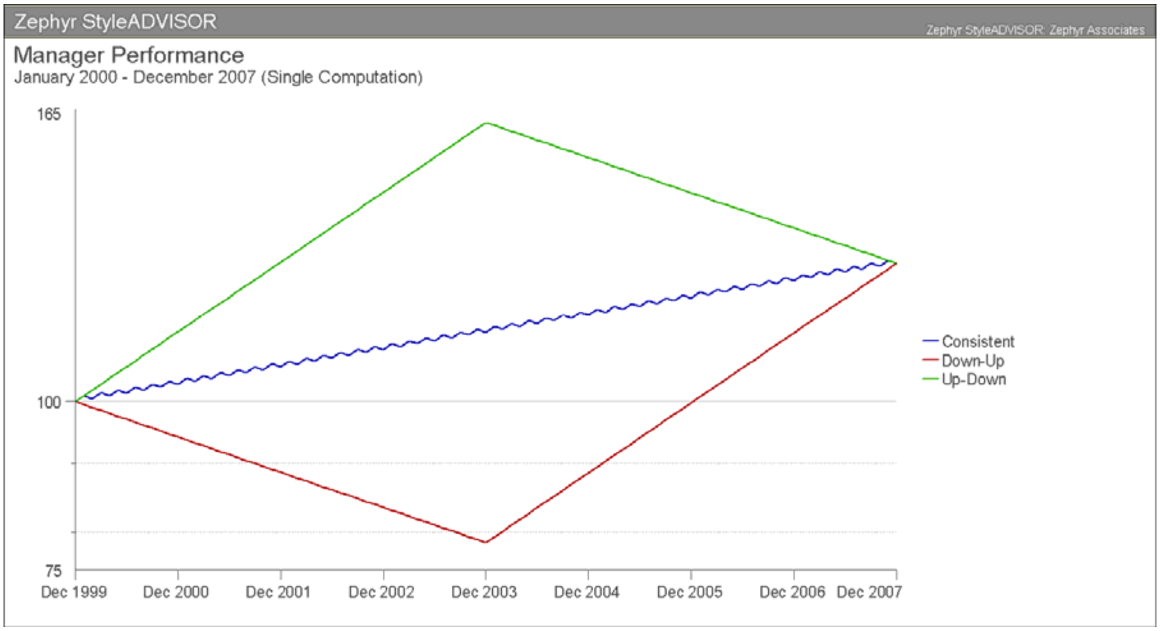

Figure 1

The Zephyr K-Ratio should be of particular interest to those who are currently using the Sharpe Ratio in their portfolio analyses. The Sharpe Ratio is the quotient of a measure of return and a measure of risk. It thus measures the amount of return per unit of risk. For the measure of return, the Sharpe ratio uses the portfolio's mean return over the riskless rate. The measure of risk is the standard deviation of the portfolio's return series.

Definition of the Sharpe Ratio

$$\text{Sharpe Ratio} = \frac{\text{portfolio's mean return above riskless investment}}{\text{standard deviation of portfolio's return series}}$$

For all its usefulness, the Sharpe Ratio has one limitation. The three portfolios whose cumulative return graphs are shown in Figure 1 all have the exact same Sharpe Ratio. That's because all three have the exact same set of monthly returns, namely, 48 returns of 1% and 48 returns of -.5%. The only difference is the order in which these monthly returns occur. The Sharpe Ratio does not capture that difference: it is entirely insensitive to the order of a portfolio's period returns.

The Zephyr K-Ratio aims to remedy this shortcoming of the Sharpe Ratio. Just like the Sharpe Ratio, the Zephyr K-Ratio is a quotient of a measure of reward and a measure of risk. However, the measures of reward and risk are chosen in such a way that they take into account the order of the period returns, and thus the shape of the cumulative return graph. To this end, a trend line is fitted to the logarithmic cumulative return graph[1]. *The*

[1] See Figure 7 below for an example of the trend line.

Zephyr K-Ratio then uses as its measure of risk is the degree to which the cumulative return graph deviates from that trend line. The measure of reward is the slope of that trend line.

Definition of the Zephyr K-Ratio

$$\text{Zephyr K-Ratio} = \frac{\text{slope of a trend line fitted to the portfolio's cumulative return series}}{\text{amount of deviation of the portfolio's cumulative return series from the trend line}}$$

Since the cumulative return graph of the portfolio named “Consistent” in Figure 1 resembles a straight line much more than the other two, it will hug its trend line more closely. This will result in a smaller denominator and thus a higher Zephyr K-Ratio. Due to the extreme nature of the example, the difference is actually dramatic: as mentioned earlier, “Consistent” has a Zephyr K-Ratio of 180.13 as opposed to 6.43 for the other two.

2. The Zephyr K-Ratio in More Detail

Figure 2 shows the cumulative return over time of a hypothetical manager, named Down-Up, for a time period of 8 years. Of the 96 monthly returns, the first 48 are all equal to -.5%. The remaining 48 are all equal to 1%. As can be seen from the chart, this return series roughly follows the behavior of the SP 500 over the same time period.

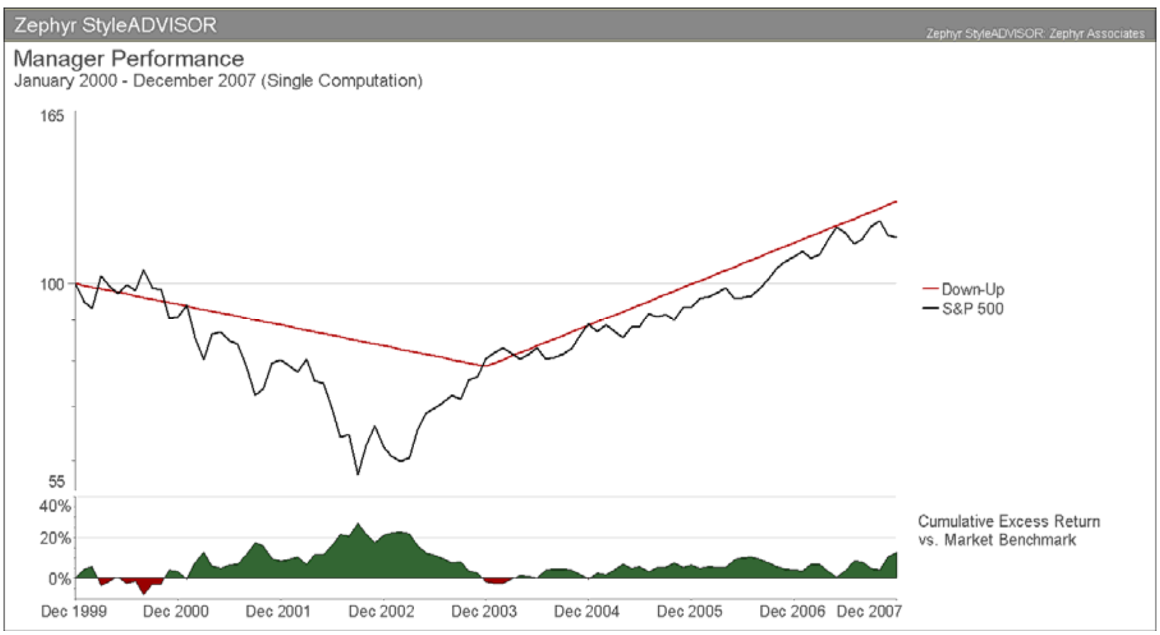

Figure 2

Next, we’ll consider two return series that differ from manager Down-Up only by the order in which the monthly returns occur. The first one, called Consistent, has a monthly return stream that alternates between 1% and -.5%. The second one, called Up-Down, is the mirror image of Down-Up: it has 48 returns of 1%, followed by 48 returns that are all -.5%.

Down-Up	48 returns of 1%, then 48 returns of -.5%
Consistent	alternating between 1% and -.5%
Up-Down	48 returns of 1%, then 48 returns of .5%

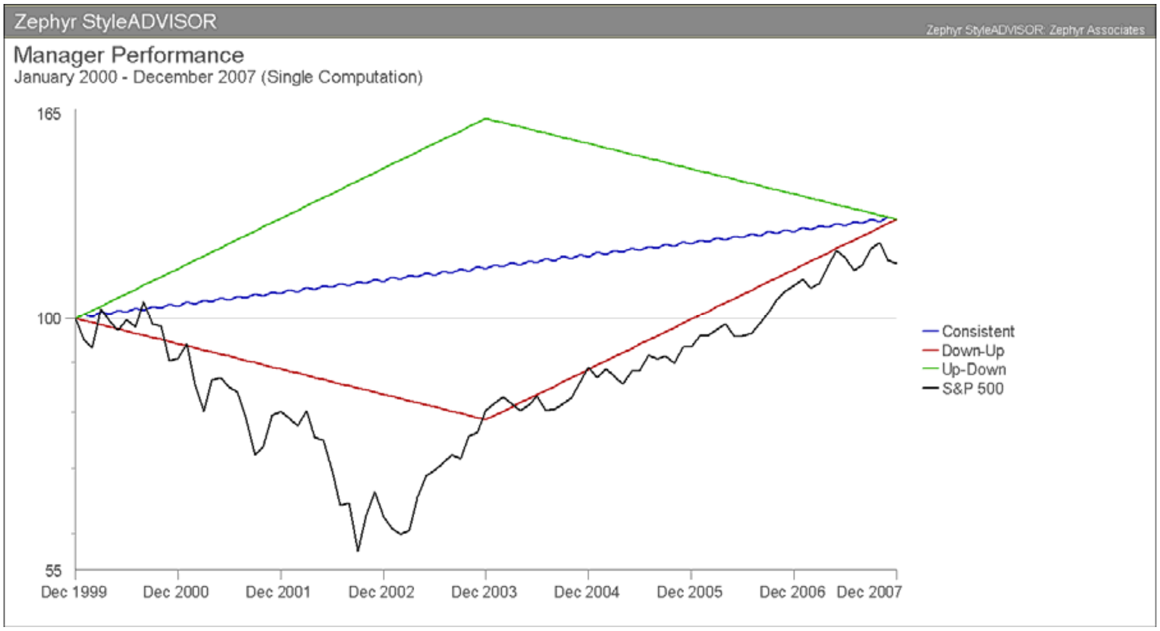

Figure 3

Figure 3 shows the resulting cumulative return over time. It is evident from the chart that many investors will not perceive these three managers as equally desirable. However, if we look at the classical absolute MPT statistics (absolute meaning "not versus a benchmark"), we see that these statistics all have the same value for the three managers:

Zephyr StyleADVISOR

Zephyr StyleADVISOR: Zephyr Associates

Custom Table

January 2000 - December 2007: Absolute MPT Statistics

	Return	Cumulative Return	Kurtosis	Standard Deviation	Sharpe Ratio	Omega (MAR = 3.00%)
Consistent	3.01%	26.75%	-2.04	2.61%	-0.11	1.01
Down-Up	3.01%	26.75%	-2.04	2.61%	-0.11	1.01
Up-Down	3.01%	26.75%	-2.04	2.61%	-0.11	1.01

Figure 4

This is due to the fact that these classical MPT statistics are inherently unable to distinguish between portfolios that differ only by the order in which the period returns occur. If we were to rate a set of managers by these statistics, then the three managers Consistent, Down-Up, and Up-Down would rate exactly the same, that is, they would be seen as indistinguishable and equally desirable. To see any difference between these managers at all within the framework of classical MPT, we'd have to look at statistics that are relative to a benchmark:

Zephyr StyleADVISOR

Zephyr StyleADVISOR: Zephyr Associates

Custom Table
January 2000 - December 2007: Relative MPT Statistics

	Excess Return vs. Market	Tracking Error vs.Market	Inf. Ratio vs. Market	Alpha vs. Market	Beta vs. Market
Consistent	1.35%	14.20%	0.09	3.07%	-0.01
Down-Up	1.35%	13.68%	0.10	2.97%	0.03
Up-Down	1.35%	14.38%	0.09	3.11%	-0.03

Figure 5

From Figure 5, we see that even the annualized excess return vs. the benchmark is the same for the three managers.

It would certainly be desirable to have a precise, mathematical way of describing the different behavior of our three return series without reference to a benchmark. Indeed, a number of statistics have been developed over the past decade or so that do just that. Here are a few examples:

Zephyr StyleADVISOR Zephyr StyleADVISOR: Zephyr Associates

Custom Table
January 2000 - December 2007: Hedge Fund Statistics

	Maximum Drawdown	Maximum Drawdown Begin Date	Maximum Drawdown End Date	Maximum Drawdown Length	Maximum Drawdown Recovery Date	Maximum Drawdown Recovery Length	Pain Index
Consistent	-0.50%	Dec 2007	Dec 2007	1	N/A	N/A	0.25%
Down-Up	-21.38%	Jan 2000	Dec 2003	48	Jan 2006	25	8.36%
Up-Down	-21.38%	Jan 2004	Dec 2007	48	N/A	N/A	5.67%

Figure 6

These statistics are often referred to as "Hedge fund statistics." They all deal with a specific aspect of a manager's cumulative return over time, namely, the occurrences of drawdowns and runups. The growing popularity of these statistics even outside the hedge fund world can be seen as an indication that the behavior of a manager's cumulative return series does matter to investors more than is acknowledged by classical MPT.

The Zephyr K-Ratio (as set forth by Lars Kestner in 1996) is the successful attempt to create a ratio that is analogous to the Sharpe Ratio insofar as it measures the reward per unit of risk. At the same time, the Zephyr K-Ratio is able to capture behavior that is caused by the order of returns. The Zephyr K-Ratio is based on the observation that the cumulative return graph of a completely riskless investment, that is, one with a constant return stream, is a straight line when plotted on a logarithmic scale. Therefore, it is plausible to measure the risk of a manager by the degree to which the logarithmic cumulative return chart deviates from a straight line. To this end, a trend line is fitted to the manager's cumulative return graph, as shown below[2]:

[2] Mathematically speaking, this trend line is the linear regression line for the cumulative return vs. time.

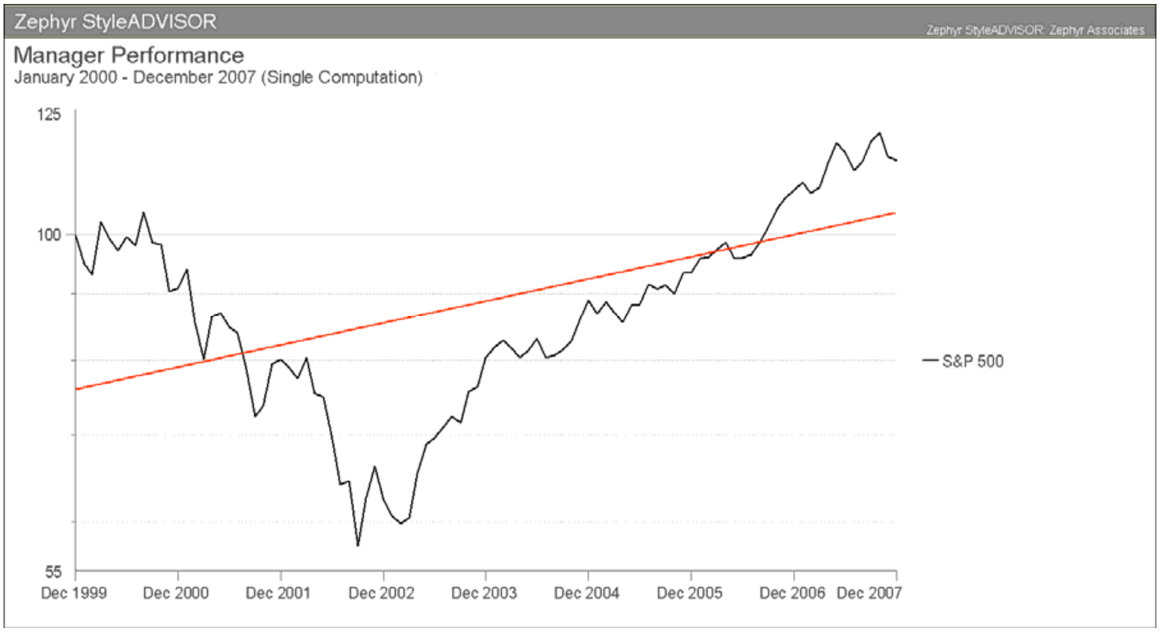

Figure 7

It is perhaps noteworthy that these trend lines, although not often encountered in portfolio theory, are often used to show long-term trends of certain markets such as the stock or bond market. One reason for the use of trend lines is that they take the endpoint sensitivity out of the cumulative return graph. In other words, a sharp upswing or downturn in a manager's returns near the beginning or end of the time period is not going to influence the trend line as much as it influences the actual cumulative return. One could thus argue that the slope of the trend line is a better measure of a manager's long-term return potential than the actual, realized cumulative return. This is indeed the approach of the Zephyr K-Ratio: the numerator of the Zephyr K-Ratio, that is, the measure of reward, is the slope of the trend line.

It remains to define the measure of risk, to be used as the denominator of the Zephyr K-Ratio. As mentioned earlier, the cumulative return graph of a perfectly riskless investment on a logarithmic scale is a straight line. In particular, the trend line of the cumulative return graph coincides with the cumulative return graph itself. Therefore, it is natural to measure the risk of an arbitrary portfolio by the degree to which its cumulative return chart deviates from the trend line. *In other words, the more ups and downs away from the trend line there are, and the farther away from the trend line they go, the more risky the investment is considered.* Fortunately, mathematicians have already developed a statistic that measures this behavior: in the theory of linear regression, this is known as the *standard error of the slope* of the trend line. For an intuitive understanding of the Zephyr K-Ratio, it is enough to know that the more closely the actual cumulative return graph hugs its trend line, the smaller the standard error of the slope will be.

We thus have the following definition of the Zephyr K-Ratio:

Definition of the Zephyr K-Ratio
Step 1: Fit a trend line (linear regression line) to the portfolio's logarithmic cumulative return graph.
Step 2: Define the Zephyr K-Ratio as

$$\text{K - Ratio} = \frac{\text{regression slope}}{\text{standard error of regression slope}} = \frac{\text{slope of the trend line}}{\text{amount of deviation from the trend line}}$$

The analogy to the Sharpe Ratio should be evident: the Zephyr K-Ratio measures a portfolio's reward per unit of risk. The difference is that here, the measures of risk and reward are chosen in such a way that they refer to the shape and slope of the cumulative return graph. The Sharpe Ratio, on the other hand, refers to the set of portfolio returns regardless of their order.

It turns out that the quotient that is the definition of the Zephyr K-Ratio, besides measuring reward per unit of risk, has a second meaning that has long been known to mathematicians. This second meaning has to do with hypothesis testing based on given sample data. Suppose the slope of the trend line is positive, that is, the trend line points upward. Then the Zephyr K-Ratio is a measure of the confidence with which the given data supports the following hypothesis: "The manager's cumulative return is, in the long run, actually positive."

In summary, we see that the Zephyr K-Ratio has three characteristics that distinguish it from the Sharpe Ratio:

1. The Zephyr K-Ratio's measure of reward uses a trend line to the logarithmic cumulative return graph, thus eliminating the end point sensitivity of a portfolio's cumulative return.
2. The Zephyr K-Ratio's measure of risk measures the deviation of a portfolio's logarithmic cumulative return from a straight line, thus reflecting the fact that movements away from a straight line are what investors perceive as risk.
3. The Zephyr K-Ratio incorporates a mathematically rigorous measure of confidence. The higher the Zephyr K-Ratio, the higher the confidence in a manager's ability to generate positive return over time.

Let us now return to our examples and look at the Zephyr K-Ratios for the series shown in Figure 3.

Series	slope of trend line (measure of reward)	standard error of trend line (measure of risk)	Zephyr K-Ratio (reward / risk)
Consistent	0.24689	.0013706	180.13
Down-Up	0.24689	0.038385	6.4319
Up-Down	0.24689	0.038385	6.4319
S&P 500	0.28014	0.057400	4.8805

Table 1

As was to be expected, the portfolio named Consistent, whose cumulative return graph almost equals a straight line, has a very large Zephyr K-Ratio. The portfolios Up-Down and Down-Up, by contrast, have a K-Ratio that resembles that of real world portfolios, such as the S&P 500. It is perhaps noteworthy that due to the symmetry between Up-Down and Down-Up, the Zephyr K-Ratio does not distinguish between these two. To capture the differences between these two in a single statistic, you'd have to look at things like drawdowns, recovery lengths, pain index, etc. What this goes to show is that while each of these statistics has its merits, no one number is the silver bullet that makes the rest superfluous.

3. The Mathematics of the Zephyr K-Ratio

The (in our opinion extraordinary) originality and ingenuity of Lars Kestner's K-Ratio lies in the idea of applying a linear regression line to a portfolio's cumulative return series and then using the slope of that line as a measure of reward and the standard error of the slope (that is, the degree to which the data deviates from a straight line) as a measure of risk:

$$\text{Zephyr K-Ratio} = \frac{\text{slope}}{\text{standard error of slope}}$$

After that, there is no need to invent any new mathematics: the exact mathematical definitions of the slope and its standard error have long existed in the theory of linear regression.

To calculate the Zephyr K-Ratio, one should first replace the dates on the horizontal axis of the portfolio's cumulative return graph with consecutive integers starting at 0. Any other equidistant series of numbers would of course work just as well, but a sequence of integers will give us naturally normalized values for slope and standard error. With these integers as independent *x*-values and the corresponding cumulative return values as dependent *y*-values, one can now calculate the slope of the regression line, that is, the numerator of the Zephyr K-Ratio, by the well-known formula

$$\text{slope} = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^{n}(x_i - \bar{x})^2} \qquad (1)$$

where $\bar{x}$ and $\bar{y}$ denote the mean of the *x*- and *y*-values, respectively. Statistics software packages will usually have this as a built-in function; for example, in Microsoft Excel it is the function SLOPE.

The standard error of the slope, that is, the denominator of the Zephyr K-Ratio, can be calculated from the *x*- and *y*-values by the formula

$$\text{standard error of slope} = \sqrt{\frac{\sum_{i=1}^{n}(y_i - \bar{y})^2 - \frac{\left[\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})\right]^2}{\sum_{i=1}^{n}(x_i - \bar{x})^2}}{(n-2)\cdot\sum_{i=1}^{n}(x_i - \bar{x})^2}} \tag{2}$$

The following equivalent formulas can be useful when using statistics software that does not have a function for the standard error of the slope, but has related statistics such as the standard error of the estimate.

$$\text{standard error of slope} = \frac{\text{standard error of estimate}}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}} \tag{3}$$

Formula (3) above is particularly convenient when using Microsoft Excel; in Excel, it translates into

STEYX(*y*-values, *x*-values) / SQRT(DEVSQ(*x*-values))

The standard error of the estimate can also be calculated directly from the *y*-values and their estimates:

$$\text{standard error of estimate} = \sqrt{\frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{(n-2)}} \tag{4}$$

Where $\hat{y}_i$ is the estimated value of y_i, that is,

$$\hat{y}_i = \text{intercept} + \text{slope}\cdot x_i \tag{5}$$

Formula (4) in conjunction with formula (3) is thus convenient if intercept and slope are available, but none of the more sophisticated statistics are.

Armed with the slope and the standard error of the slope, we obtain the Zephyr K-Ratio as

$$\text{Zephyr K-Ratio} = \frac{\text{slope}}{\text{standard error of slope}}$$

It so happens that when it is non-negative, this quantity is also the t-score for the rejection of the hypothesis, “The true regression line has slope less than or equal to 0.” In other words, the Zephyr K-Ratio has two meanings: on the one hand, it measures the portfolio's reward per unit of risk for a certain rather plausible definition of risk and reward. On the other hand, it is also a measure of the confidence with which we may assume that the true trend line for the portfolio's cumulative return is not flat or pointing downward.

The interpretation of the Zephyr K-Ratio as a t-score does of course imply a certain assumption of normality. The assumption is that the estimation error of the linear regression on the cumulative return data is normally distributed. It is a common assumption in Modern Portfolio Theory that the raw period returns are normally distributed. There has been ample research into that assumption; however, we are not aware of any research regarding the assumption that the estimation errors of the linear regression are normally distributed. One should perhaps not speculate before empirical data has been examined, but it seems plausible that normality of the estimation errors is no less likely than normality of the logarithmic period returns.

4. The Zephyr K-Ratio vs. the K-Ratio

The Zephyr K-Ratio as we have explained it in the previous three sections is the original K-Ratio that Lars Kestner introduced in 1996. In his 2003 book ([1]), Kestner made the following modification to the definition of the K-Ratio:

$$\text{Kestner's K - Ratio of } 2000 \text{ book} = \frac{\text{slope}}{(\text{standard error of slope}) \cdot (\text{number of data points})}$$

The following rationale for this modification is given on p. 87 of Kestner's book: "By dividing by the number of data points, we normalize the K-Ratio to be consistent regardless of the periodicity used to calculate its components." In this section, we will present our arguments for using the original version of the K-Ratio. As before, we will refer to the original version as the "Zephyr K-Ratio".

First off, when making any change to the Zephyr K-Ratio, one loses the property that the ratio, when it is non-negative, is equal to the t-score for rejecting the hypothesis "the slope of the trend return line is not positive." That property is certainly desirable to have; this is a reason to be wary of any suggestion to modify to the ratio, regardless of what it is.

We present two arguments for preferring the original K-Ratio over the modified version. The first one concerns the intended effect of the modification, namely, making portfolios comparable across different periodicities. We will argue that the original K-Ratio already guarantees that comparability; no modification is necessary in our opinion. The second argument concerns the effect of the modification on comparisons between portfolios with the same periodicity. We will argue that the modification has a detrimental effect in that situation.

To demonstrate the effect of the modification to the K-Ratio on data with different periodicity, we will convert the three portfolios and the market benchmark that we have been using throughout (see Figure 3) to quarterly data and inspect the way the two K-Ratios are affected by this change in periodicity. Below is the analog to Figure 3, with all series converted to quarterly.

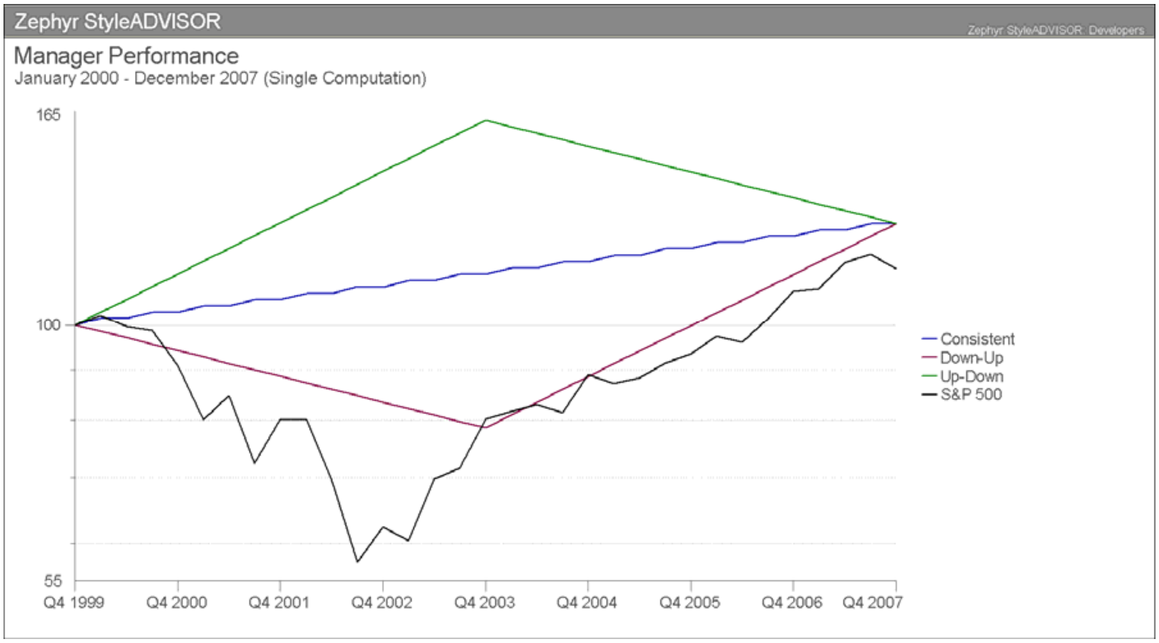

Figure 8

We see that this is the same picture as Figure 3, except that there is less granularity. More precisely, there are only a third of the data points, with no information on what is happening in between. Here are the corresponding K-Ratios:

	Zephyr K-Ratio		K-Ratio	
Series	monthly data	quarterly data	monthly data	quarterly data
Consistent	180.13	105.02	1.8570	3.1824
Down-Up	6.4319	3.6697	0.066308	0.11120
Up-Down	6.4319	3.6697	0.066308	0.11120
S&P 500	4.8805	2.6452	0.050314	0.080158

Table 2

We see that the Zephyr K-Ratio (Kestner's original 1996 version) decreases significantly as we pass from monthly to quarterly data. Moreover, this is not just any decrease: the change reflects, with the mathematical precision of Student's t-test, the loss of significance that is caused by the higher granularity of the quarterly data. We cannot see anything wrong with that. The modified, more recent K-Ratio, on the other hand, actually *increases* as we pass from monthly to quarterly data. The increase is significant insofar as the K-Ratio of the quarterly S&P 500 is higher than the monthly K-Ratio of the Up-Down portfolio, while within the same periodicity, the S&P 500 scores lower than Up-Down (see shaded cells in Table 2).

If the adjustment made for the modified K-Ratio were such that monthly and quarterly data now gave the same or roughly the same values, then we could see how an argument for the adjustment could be made. We would still prefer the Zephyr K-Ratio, where the loss of significance is reflected in the value of the ratio, but we could see how someone else would want the adjustment. However, an adjustment that actually shows significantly higher values for quarterly data as opposed to monthly data is something that we advise against.

Our second argument against modifying the original K-Ratio concerns its effect on comparisons between portfolios with the same periodicity. Remember, the intended effect of the modification concerns the comparison between portfolios of different periodicity. But it is clear that dividing by the number of data points has a non-trivial effect on comparisons between portfolios with the same periodicity as well, namely, when comparing portfolios with different amount of data. Clearly, it is preferable to only ever compare portfolios with the same amount of data. However, the Zephyr K-Ratio, because of its connection to the Student t-test, does allow comparisons across different amounts of data. Figure 9 shows an example.

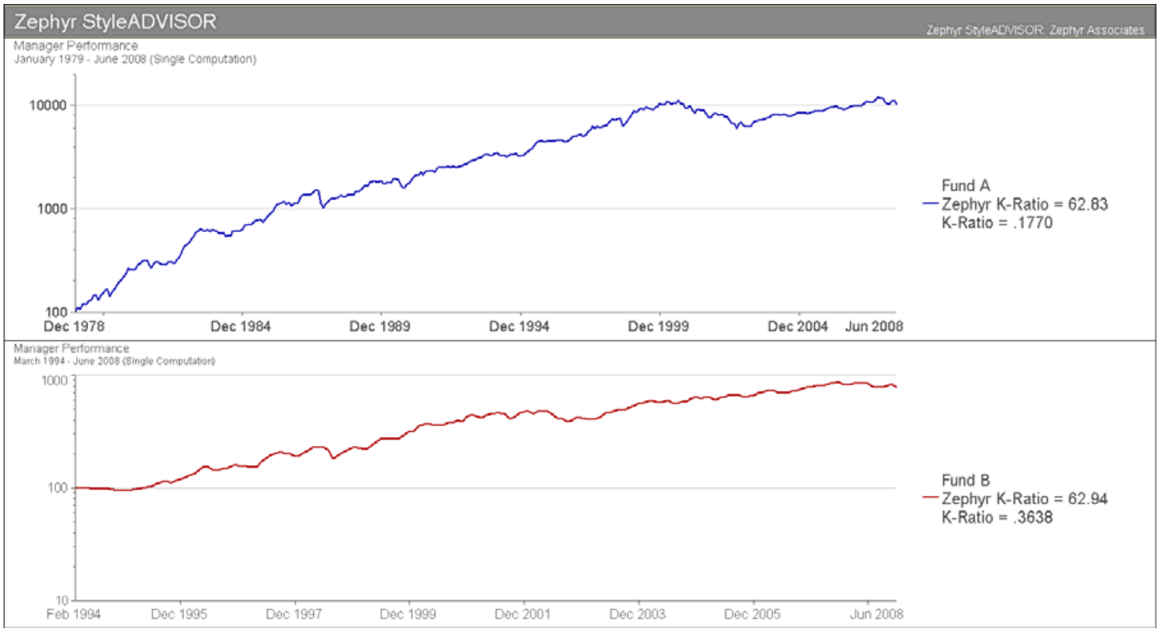

Figure 9

Here, Fund A has about twice as much data as Fund B. The slopes of the trend lines of the graphs of Fund A and Fund B are about equal (1.221 and 1.364, respectively). The graph of Fund A has visibly more deviation from a straight line than that of Fund B. However, because of the fact that Fund A has twice as much data to back up its trend line, the standard errors of the trend lines are about the same (.001944 and .002168, respectively). As a consequence, the Zephyr K-Ratio clocks in at almost the same value for both funds, namely, 62.83 and 62.94, respectively. All this is mathematically and statistically perfectly sound, and it conveys meaningful and useful information about the funds' cumulative return graphs.

The modified K-Ratios, as shown in the legends of Figure 9, on the other hand, are .1770 and .3638. In other words, Fund B now has more than twice the K-Ratio of Fund A. We simply do not see how one could explain such a dramatic change in the relationship between two monthly funds on account of an adjustment whose intention was to ensure comparability between portfolios whose data has different periodicity.

References

[1] Lars Kestner, *Quantitative Trading Strategies: Harnessing the Power of Quantitative Techniques to Create a Winning Trading Program, McGraw-Hill Traders Edge Series, 2003*