

# 10 Fallacies When Selecting CTAs

## Superficial aspects can be misleading

SEB GLOBAL QUANT TEAM

### Summary

In this article we are taking the perspective of an institutional investor who wants to invest in one or several commodity trading advisor funds (CTAs). Meetings we have had with investors over the last decade indicate that they do an initial quantitative screening, but then – after having met with the different CTA managers – put a lower weight on the hard facts (realised track record) and a higher weight on soft factors (all other kinds of information) when they make their final investment decision. This kind of behaviour either signals a belief that soft factors can forecast the future returns of the different CTA managers, and/or that investors make their investment decisions on the basis of emotions rather than objective facts.

We have scrutinised the following 10 fallacies in this article:

1. The Pitch Book Fallacy
2. The Slick Presenter Fallacy
3. The Big Team Fallacy
4. The Title Fallacy
5. The Long Experience Fallacy
6. The Brand Fallacy
7. The Technology Fallacy
8. The Trading Fallacy
9. The Transparency Fallacy
10. The Performance Fallacy

The research we carried out on 23 of the largest CTAs in the world resulted in the conclusion that neither team size nor experience is able to predict the risk-adjusted returns of different CTAs. After detailed discussions, we also drew the conclusion that none of the other soft factors are likely to contain any predictive power either: pitch books, communication skills of presenters, the number of people with PhD titles, the manager's brand, the technological appearance, the trading set-up or the degree of transparency provided. Finally, when it comes to performance figures, investors need to watch out for at least seven different pitfalls when managers present their past performance.

An objective like-for-like analysis of past performance may be the only remaining factor that potentially possesses some power to predict future performance. Investors may therefore want to spend more time analysing, understanding and adjusting track records, as well as classifying CTAs into sub-categories and running client portfolio simulations to find out which CTAs generate the largest improvement in the risk-adjusted returns for the client.

**Hans-Olov Bornemann**

*Portfolio manager and head of Global Quant Team*

### How we did the research

We started by looking at the constituents of three major CTA indices: BarclayHedge's BTOP50, the NewEdge CTA index and the Dow Jones Credit Suisse Managed Futures index. From this aggregated group of funds, we excluded funds:

- a) That tend to make qualitative (as opposed to quantitative) investment decisions;
- b) That do not offer public access to their performance data on Bloomberg;
- c) That did not exist prior to October 2006 or that had closed down by August 2013;
- d) That are copies of other funds.

We chose to use the launch date of our own fund, SEB Asset Selection (3 October 2006), as the start date for the research. The end date of August 2013 was chosen as we easily could use data from another study we had done with that end date. Following this methodology, we ended up with the following 23 CTA funds: Altis, Aspect, Boronia, Brummer & Partners Lynx, Campbell, Cantab, Eckhardt, Estlander, FTC, FX Concepts, Graham, IKOS, Lyxor Epsilon, Man AHL, Millburn, Nuwave, Ortus, Rivoli, Transtrend, SEB Asset Selection, SMN Diversified, Superfund and Winton.

Of course, it would have been great to have had at least 30 funds in the study, but rather than tweaking

the rules or changing the methodology, we have chosen to stay with the above 23 funds.

When it comes to return data, we have derived them from each fund's NAV per share series on Bloomberg (in dollars or euro). Sharpe<sup>1</sup> ratios take the currency-specific and period-specific risk-free rate into account. The number of years since inception and the number of relevant researchers per fund have been taken from publicly available sources in 2013. In some cases we have had to estimate these numbers. We stopped short of trying to classify each firm's employees with regard to their respective academic titles.

As regards the qualitative factors, we have chosen to discuss those fallacies from a behavioural finance perspective rather than making subjective assessments. Hopefully, those discussions can help fund selectors to avoid some of the pitfalls.

### 1. The Pitch Book Fallacy

It is time to get going with the first fallacy, the Pitch Book Fallacy. It is widely known that corporate finance people at the major investment banks are the masters of the universe when it comes to pitch book production. However, anybody who has met a larger CTA firm would probably agree that their pitch books are pretty good-looking too.

Most people would probably agree that neither the structure, lay-out nor the touch and feel of a pitch book has anything to do with an investment team's ability to generate excess returns. Anybody with an aesthetic sense and some basic knowledge of a presentation programme could come up with a very professional-looking presentation. Some managers believe it is worth the extra time and effort, others think that clients ignore the packaging.

The contents of presentations are, however, deemed to be relevant for judging a team's future alpha generation capabilities – that is why presentation materials are always used in meetings between clients and managers. At this point, let us make a distinction between hard contents and soft contents and define hard contents to be numbers and graphs directly or indirectly related to the fund's track record. In a corresponding way, we define soft contents to be the pages that describe the team's history, organisation, philosophy, processes and so on.

Let us discuss the soft contents first. Some investors believe that the soft contents give a more accurate forecast of a fund's future performance than the historical track record does. The only problem with this hypothesis is that most of the larger CTAs have investment processes and trading processes that are very similar to each other. Yet there is a fairly substantial dispersion in the funds' returns on a

year-to-year basis as well as the risk-adjusted returns over the longer term.

This is how we see the world: the risk-adjusted returns of CTA funds are not generated by the general and schematic investment, and trading processes pictured in the presentations or described in one-on-one meetings with investors. The risk-adjusted return of a specific fund is a direct function of the detailed specification of their alpha models. The model specifications are different between the different CTA funds and thus short-term and longer-term risk-adjusted returns will differ. However, since none of the largest CTA managers are willing to reveal their models in such detail to enable fund pickers to make a detailed comparison between funds, it is probably quite far-fetched to believe that you will find predictive power in the soft contents of presentations. You can quite easily tell, however, who is good at producing professional-looking hand-outs.

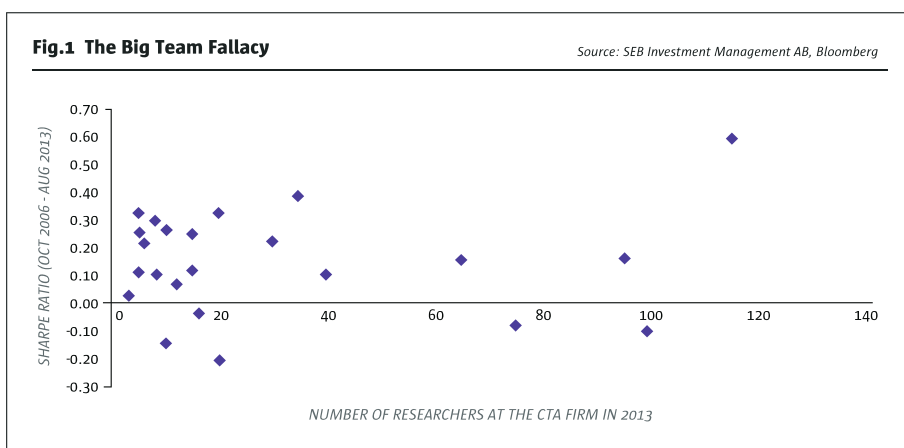
When it comes to the hard contents of presentations (the track record), let us come back to that a bit later. In spite of the mandatory consumer warnings, it may be the only relevant straw you can hold onto in the end.

## 2. The Slick Presenter Fallacy

The impression you get from a meeting is not just based on the presentation material used in the meeting. It is probably fair to say that the presenter's ability to connect with the client, to explain, to argue, to reason, to create a positive atmosphere, to establish a professional rapport and to convey energy is at least as important as the quality of the pitch book.

What a well versed presenter can do is to make you feel comfortable with the idea of handing over money to the asset manager. A good presenter is able to minimise any uncertainty that the investor may feel. Will the investment team be able to deliver good performance in the future? Will the asset management organisation be able to fulfil its risk control, compliance and other duties? Is the back and middle office staffed with competent people and equipped with good, efficient and safe systems? Is the fund valuation done by an independent and professional party? Are the assets of the fund held by a respectable custodian or not?

Of course, there are some hard facts (e.g., fund valuation being done by an independent and respected party or not, assets being held by a respected custodian or not) which the presenter cannot do too much about – i.e., typical yes-or-no questions. However, since most questions are of a more open nature, there is often enough flexibility for a good presenter to make a good impression.



## 3. The Big Team Fallacy

When it comes to the size of different investment teams, many fund pickers think that larger research teams should generate higher risk-adjusted returns than smaller teams. To most people, it seems like an obvious fact. The reasoning goes like this: “100 researchers must be able to beat a team of 10 researchers...” and the reasoning may continue like this: “of course, in a single year, randomness may disturb the picture, but over time, it must be true.”

Because it is such an obvious ‘truth’ and because it takes a bit of time to gather the relevant data to check the hypothesis, most people would not even bother to do the work. As opposed to the above-mentioned pitch book fallacy and slick presenter fallacy, it is substantially easier to test whether team-size has an effect on future returns or not.

In Fig.1 we have pictured the number of researchers per CTA manager in 2013 against the Sharpe ratio of each CTA fund over the period from October 2006 to August 2013. We have chosen to use risk-adjusted returns, but of course there are other performance measures that could be used as well.

For the full sample of 23 CTA managers, the correlation between the number of researchers and the risk-adjusted returns amounted to a mere 0.09. If you were to exclude the CTA funds with the highest and lowest Sharpe ratios, you get a negative correlation of -0.31, i.e., the more researchers you have, the lower your Sharpe ratio.

Irrespective of whether you choose to exclude some funds from the sample or not, the R-square (how well the values of the X-variable are able to explain the variation in the Y-variable) is very close to zero in both cases (only 1% if you include all observations and 10% if you eliminate the observations with the highest and lowest Sharpe ratios). From this analysis, it is not possible to conclude that larger research

teams are able to generate higher risk-adjusted returns than smaller research teams. Nor can you claim the opposite.

Thus, if you are looking for risk-adjusted returns, do not get fooled to believe that larger CTA teams will deliver better performance for you. In the CTA industry, team size seems totally uncorrelated to performance. This implies that larger teams (with the exception of one manager) seem to be unable to capitalise on their vast research resources.

At least two criticisms can be put forward to our analysis. First, the sample contains only 23 observations. The absolute minimum for statistical testing is generally thought to be 30 observations. We would like to encourage database companies who might have gathered this kind of data from a larger set of CTAs to publish their research on this topic.

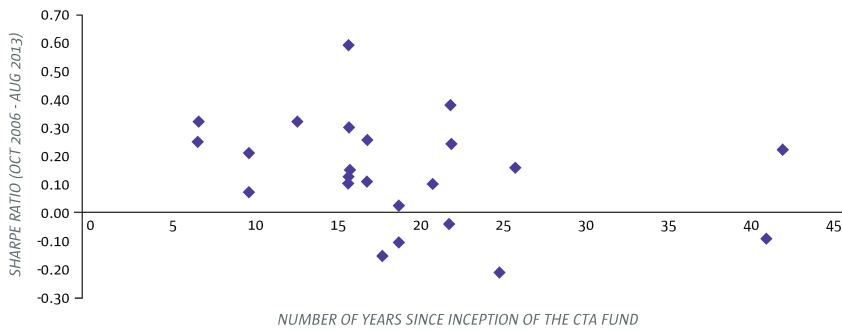
Second, observing the number of researchers at the end of the period rather than at the beginning of the period basically implies a test with perfect hindsight. Over this seven-year period, there has been a tendency that the most successful managers have been recruiting more researchers and the least successful have been making people redundant. In other words, if the number of employees had been measured ex ante, the correlation between team size and Sharpe ratio is likely to have been even lower (or more negative). We leave that for an independent researcher to study.

## 4. The Title Fallacy

Would it not be reasonable to assume that a person with a higher academic title has a better chance of delivering good performance than a person with a lower or no academic title? It seems like a reasonable hypothesis. Also, in the marketplace, a lot of investors seem to get impressed by academic titles, the PhD title in particular.

**Fig.2 The Long Experience Fallacy**

Source: SEB Investment Management AB, Bloomberg



The positive thing about people who have completed a higher academic education is that they have read a lot of academic research, learned to apply a scientific research methodology and have been equipped with fairly sophisticated theoretical models/tools for solving different kinds of problems. Also, one should not forget that their academic achievement reflects a high intellectual capacity and a willingness to work hard – very important ingredients for success in most jobs.

In a corresponding way as in the discussion about team size, human logic would say that a person with a higher education should have a greater chance of success than a person with a lower educational level. However, nobody seems to have bothered checking this hypothesis in reality. We refrained from calling the 23 CTA funds to ask for this kind of information as we thought they would not be willing to give it to us. However, we would encourage independent researchers to look into this. Our best guess is a correlation of +/-0.10 between educational level and risk-adjusted returns within the CTA industry. If that was the case, titles would be another fallacy that fund pickers may want to watch out for.

### 5. The Long Experience Fallacy

Another human logic is that people with a longer experience should be better at their job than people with a shorter experience. Again, implicit in this logic is the famous *ceteris paribus* assumption. The hypothesis can be tested by regressing the age of the respective funds to the funds' respective Sharpe ratios. We have done this in Fig.2.

Whether you include all observations or exclude the two observations with the highest and lowest Sharpe ratios, in both cases you get a correlation of about -0.30. In other words, the longer the fund had existed, the lower its risk-adjusted return turned out to be (on average). The human logic does not seem

to be particularly helpful when you are trying to spot the best performing CTA funds.

Is the explanation to this phenomenon that the best people may have left the company and started their own funds and that their replacements were not smart enough to keep up with the competition? Or, is it so that some teams – even if the same people remain on board – are unable to keep up with the new competition, in spite of the fact that they have had more time to develop their understanding of markets and to pursue more back-tests?

To be fair, one should not draw any far-reaching conclusions from this negative correlation. The R-squared only amounts to 9%, which is not a particularly high number. Moreover, considering that people may leave one CTA firm and join another or set up a competing firm, the age of the CTA programme is not necessarily indicative of how experienced a certain team is.

In any case, it is probably fair to say that one should not get impressed by the argument, “We are one of the most experienced teams in the industry.” There is no evidence in the real world indicating that CTA managers with 20+ years of experience are doing a better job than managers with seven to 10 years of experience.

### 6. The Brand Fallacy

Identifying and choosing the best CTA fund is not a particularly easy task. In such uncertain situations, there is a human tendency to prefer funds, teams and firms which you are more familiar with to those you have not known for an equally long time. The human logic goes something like this: “It is better to buy a fund which you know, because funds that you aren't very familiar with may contain all sorts of risks.”

The story about Goldman Sachs' quant team 2007-2011 is a good example of the brand fallacy. Nobody

could have imagined that such well performing funds, run by such talented people working for a firm with such a strong brand, could have contained such great risks and could have started to perform so poorly all of a sudden. Prior to the failure, LTCM enjoyed an equally stellar reputation in the hedge fund industry.

Brand awareness is very much a function of the number of years the brand has existed, you have known it, the way it has presented itself to investors/consultants, the way it has been interpreted by you and described by the press etc.

What is interesting with the brand discussion is that it is much easier for a CTA firm to build a certain brand than to deliver solid performance over time. Clients may want to keep this in mind, so that they do not end up buying the manager's brand instead of the manager's ability to generate performance.

### 7. The Technology Fallacy

Considering that people running CTA funds belong to one of the more geeky subcultures of society, clients may be fooled to believe that the geekiest people also are the best investment professionals.

Clearly, in certain areas of quantitative trading (high-frequency trading, HFT) the technological sophistication is indeed a key factor for achieving success. If you were consistently a millisecond slower than the fastest HFT manager when identifying alpha opportunities and trying to take advantage of them (for example a particular tick change in an instrument's bid or ask level), you would end up making no money at all, even if you were able to identify all kinds of lucrative alpha sources (do not forget, in the HFT field, the size of each opportunity is very small indeed).

In the CTA field, investors should be very cautious about drawing conclusions between apparent technological sophistication and the manager's ability to deliver performance. To a large degree, the technology show-off is only part of the marketing spiel.

### 8. The Trading Fallacy

Related to the technology fallacy is the trading fallacy, i.e., the marketing pitch that you can only deliver fantastic CTA performance if you have a state-of-the-art trading capability. This story has typically been pushed by the largest players in the industry, i.e., the ones who are struggling with huge assets under management.

You can look at trading from two different perspectives. First, you can have a situation where you have complete flexibility as to what instrument you trade and when you choose to trade it. If you

have done research in the high-frequency spectrum of the market and found that you are able to deliver respectable performance from such intra-day trading models, then you add these high-frequency models to the group of models that you are running in your CTA fund.

Any CTA team that has gone down this path has a) already chosen to implement or not to implement such high-frequency models (and they are already part of the observed track record) and b) discovered that the capacity of an intra-day model in any case is only a fraction of the capacity of a short-term model, which in turn is a fraction of the capacity of medium-term or longer-term models. In other words, there is no large CTA in the world that is getting a larger portion of its excess returns from high-frequency models. If there was, they would not be as highly correlated to the other CTAs as they are.

In summary, the more money you are managing in a CTA programme, the larger the transaction costs will be. This is true for any and all CTA firms. To the extent that a CTA firm claims that execution is incredibly important and that they are really good in that area, it should be very easy for them to prove these claims via a superior track record, i.e., that their execution costs are not eating up too much of the alpha in their models.

### 9. The Transparency Fallacy

Imagine a situation where a potential fund investor could get 100% access to all CTA managers' model specifications, back-testing systems, live-trading systems, quantitative and qualitative decision-making processes, IT development procedures, databases and other IT structures, security solutions, custodian set-ups, regulatory restrictions, detailed CVs of each and every employee and information on all other resources that are used at the CTA firms.

The investor would certainly get a much better understanding of how much or how little it takes to run a CTA fund. The client would be able to notice and record similarities and differences between IT architectures, organisations, processes, models, programming languages etc.

After having dug deeply into the programming code of each CTA manager's model, i.e., the heart of any CTA fund, it is quite likely that the investor will eventually ask him- or herself: "So what kind of fund performance will come out of all these different models in the end?"

At this point, the investor has two alternative ways to go: a) to compare the Sharpe ratios of the back-tests that each CTA manager has pursued, or b) to compare the historically realised Sharpe ratios, i.e.,

that have come out of their models in the real world (as opposed to the theoretical back-tests).

For this reason, you cannot compare Sharpe ratios of different CTA teams, unless you can secure that the back-tests have been conducted with similar and hopefully very limited amounts of data mining and curve-fitting. This implies that you are back to comparing and analysing past performance. In other words, after having completed your in-depth and all-encompassing research project at the different CTA managers, you are back to square one again analysing the CTA managers' track records.

Interestingly enough, instead of being leading indicators basically all of the above-mentioned factors seem to be lagging indicators with regard to the CTA fund's performance. That is, the better the fund's risk-adjusted performance has been, the larger the inflows and the assets under management, the greater the firm's revenues, the more people they have employed, the more they have overhauled their pitch books, and the more focus they have put on brand building. It should not be forgotten: the better the track record looks, the greater the expected payoff from spending money on marketing.

Given the above findings, more time should probably be spent on analysing the historically realised track records. By focusing your research on the hard numbers, you also avoid being deceived by the marketing pitches.

Analysing and comparing track records may not be as easy as some people think, however. It is time for the tenth fallacy.

### 10. The Performance Fallacy

Rule number one when comparing track records to each other is to compare them on a like-for-like basis. This implies a number of things:

- a) The first thing to keep in mind is that you should compare track records of different CTAs on the basis of the net fees (after potential rebates) that you as an investor would be paying in each case. Some CTAs publish institutional share classes with non-negotiable fees, other CTAs publish retail share classes on which institutional discounts are given.
- b) Second, you cannot directly compare the returns of CTA funds, if the time series represent different currency share classes. Since the pricing of currency forwards (which are used to hedge a fund's currency exposure from the base currency to the respective share class currencies) is a direct function of the interest rates that exist in the respective currencies, you should expect a higher return from the share class whose currency is enjoying a higher interest environment.

This phenomenon can easily be seen when you compare the performance of different share classes that belong to the same fund (and have the same fees). The difference in performance between two share classes over time should basically be equal to the interest rate differential between those two currencies over that time period.

Another way to explain this phenomenon is to look at the way a CTA portfolio is structured. First, there is the base portfolio, i.e., the part of the portfolio that invests the cash received from investors. A prudent CTA would invest this cash into the short-term government bill market and collect a risk-free rate of return on the investment. Second, there is an overlay portfolio consisting of futures and forwards (requiring collateral, but basically without financing need). Via their skills to forecast the direction of different markets, CTA managers are able to deliver some excess returns for investors in this part of the portfolio. When you are currency hedging a share class, you are not only getting rid of the currency exposure to the main/base currency, but you are also converting the underlying risk-free return from the main currency to the risk-free return that can be achieved in the share class currency.

Thus, when comparing funds, one should ideally use share classes that are hedged to the same currency. If it is not possible to find a share class of the desired currency, one could either try to adjust for the differential in risk-free rates during the period or at least be aware of the effect when analysing the results.

- c) Third, even if you choose to use data from the correct currency share classes when comparing two CTA funds, you still run the risk of comparing apples to oranges. We have seen fund selectors calculating and comparing funds on the basis of Sharpe Ratio Since Inception. Clearly, since funds tend to have different inception dates, you would compare numbers that have been calculated over different time periods and potentially draw the wrong conclusions.
- d) Fourth, a fund's net return is also dependent on its risk level. Thus, to assure a like-for-like comparison, the net returns need to be converted into risk-adjusted returns.

Another beauty about risk-adjusted returns like Sharpe ratios is that you can directly compare them even if they have been calculated using a variety of different currency share classes. The reason for this is that the calculation of the Sharpe ratio takes out the currency-specific risk-free rate and only contains the excess returns, which are currency-independent.

e) Fifth, even if the above mistakes may seem obvious to most people, there is yet another mistake you can make when trying to figure out which CTA fund to pick. Most people believe that they should pick the CTA manager with the highest stand-alone Sharpe ratio. When doing it in this way, you are making the assumption (probably unconsciously) that you would invest all of your money into a CTA fund. In reality, however, most clients prefer to keep some of their existing investments, e.g., equities/equity funds, bonds/bond funds, other hedge funds and so on, so that they get an overall diversified portfolio with a number of different exposures.

In such situations, investors should instead try to find the CTA fund that delivers the greatest value to the client's total portfolio. By calculating the Sharpe ratio (and other statistics you may be interested in) of the overall client portfolio (when including the respective CTA funds one at a time), you will be able to identify which CTA fund has the best fit to your specific portfolio. When doing these portfolio simulations, it is of course important to assure a like-for-like comparison, such that the volatility of the different CTA funds is normalised (put on an equal level). Since excess return is a function of the volatility in the fund, one can scale the excess return of different funds with the relative level of volatility.<sup>2</sup>

It should be noted that the goal of a CTA fund is to be a good 'Sharpe-booster' and 'tail-risk hedge' in the context of a client portfolio. Because of this goal, you should not expect CTA funds to have particularly high stand-alone Sharpe ratios (but you can probably expect the attractive combination of a positive skew and excess kurtosis). In fact, an abnormally high stand-alone Sharpe ratio should instead raise questions. Has the fund drifted away from the pure CTA strategy? If that is the case, the fund may not be able to protect the client portfolio in a bear market in the same way as it might have done in the past. Notice also that funds that may have boosted the client portfolio Sharpe ratio over a period of time, during which no bear markets developed, may look like a great CTA for a while. However, in years such as 2008 (equity bear market) and 2013 (bond bear market), it becomes pretty clear which funds actually possess the attractive tail-risk hedge and longer-term Sharpe-boosting features.

"CTA funds" come in many different flavours today:

i. CTA funds that have stayed true to the classical medium-term trend-following style, i.e., the strategy that seeks to boost client portfolio Sharpe ratios and protect client portfolios

during bear markets by being a statistical – not a perfect – tail-risk hedge;

- ii. Funds that are on their way to becoming multi-strategy hedge funds (probably doing so to maximise their stand-alone Sharpe ratio and to be able to handle larger assets under management);
- iii. The ones that have chosen to become longer-term trend-followers or who have complemented their CTA strategy with long-only exposures to bonds and/or equities (a style drift prompted by large AUM and/or a desire to maximise stand-alone Sharpe ratio);
- iv. Funds that have chosen to trade the futures markets on a higher frequency, i.e., high-frequency traders or short-term traders (who strive to be a complement to the medium-term trend followers and/or to protect client portfolios from shorter-term market corrections – these funds tend to have a fairly limited capacity, though).

In other words, investors who are searching for a classical CTA investment are advised to take a closer look at the different funds before making their decisions.

f) The sixth thing to keep in mind when comparing track records of CTA funds is the following: a like-for-like comparison also requires you to use the same data frequency for all funds. Theoretically, it should not matter if you calculate the annualised volatility on the basis of daily data, weekly data or monthly data. However, because CTA returns are not normally distributed and because they are not independent of each other, your estimates of the annualised volatility may differ quite considerably depending on whether you use e.g., daily or monthly data in your calculation. In other words, be sure to use the same data frequency for all CTA funds to get a fair comparison, even if some of them offer more frequent valuations than just monthly or weekly data.

g) Finally, when it comes to making forecasts about the future returns of CTA funds (as well as other hedge funds), some people just use the historical average return of the respective CTA fund as their forecast. Such forecasts may unfortunately be quite unrealistic. The reason is simple. Recall the fact that the net return of a CTA fund comes from two different sources, the risk-free rate on the base portfolio and the excess returns generated in the futures overlay.

Since the risk-free rate used to be clearly higher prior to 2008 than it has been thereafter, you

need to deduct the historical risk-free rate from the net return numbers and instead add the current risk-free rate (a much lower rate) in order to get a more sensible return forecast in today's interest environment. **THFJ**

## NOTES

1. In his article "The Sharpe Ratio" (1994) William Sharpe is first calculating the excess returns between the fund and the risk-free rate for each year (or another data frequency). Then he calculates a simple arithmetic average of the annual excess return numbers. The average return per annum is then put in relation to the standard deviation of the annual excess returns. This ratio is called the Sharpe ratio. William Sharpe is assuming a 1-period situation and is therefore using the arithmetic average for calculating returns per annum rather than the geometric average. Considering that investing is a multi-period phenomenon and that returns are 'base-dependent', we prefer to use the geometric methodology for calculating average returns per annum. A return of -50% in year 1 and +100% in year 2 would yield an arithmetic average return of  $(-50\% + 100\%) / 2 = +25\%$  per annum. With the same assumptions, the geometric average return will be  $((1 - 0.50) * (1 + 1.00))^{(1/2)} - 1 = 0\%$  per annum. Also, whereas William Sharpe first calculates the excess returns and then the average, we choose to calculate the geometric annualised returns for the fund and the risk-free rate first and then take the difference between the two annualised numbers. When it comes to the standard deviation of the excess returns, we approximate the volatility of the risk-free rate to be zero and just use the volatility of the funds to calculate our risk-adjusted return ratio. Rather than inventing a new name for our way of defining risk-adjusted return (which under normal situations will be quite similar to the original Sharpe ratio), we have chosen to refer to it as the Sharpe ratio.

2. This is only true as an approximation. The mathematics of scaling risk and excess returns are not so obvious. When you increase or reduce the size of the futures positions in a CTA fund to create a product with a different risk target, the excess return for a single day (i.e., the return beyond the risk-free rate or the return attributable to the base portfolio) will be directly proportionate to the excess return on a product that is run with another risk target. However, in a multi-period context, the effects of compounding start to distort the direct relationship between volatility and excess returns. Thus, in a multi-period context, scaling excess returns with the level of volatility should only be done as an approximation.