COLLECTIVE insight

Contents

Please email feedback and suggestions to:
collectiveinsight@mcubed.co.za

From the advisory committee

BY ROWAN WILLIAM-SHORT
African Harvest

WELCOME to the third issue of *Collective Insight*. The first two elicited favourable responses beyond our optimistic hopes, and we continue to receive requests for extra copies of both.

We chose the topic for this issue – "Use and abuse of numbers" – because each of us has encountered many striking examples of both in the South African investment arena. Some years ago, when calculators were first allowed in high school mathematics exams, it was feared that this step would render students' understanding of the concepts redundant. Nothing could be further from the truth. As a student of mine remarked, "JIJO!" (Junk In, Junk Out).

It now seems that the ready availability of powerful statistical tools is having a similar effect on the investment industry. The ease with which a multiple regression may be performed using basic spreadsheets has led to a plethora of absurd "findings" that would be debunked by even cursory examination of the underlying assumptions and, more especially, of how these have been violated.

The allure of templates to expedite equity analysis is another example. We all know that earnings growth expectations are computed as future earnings divided by past earnings, minus one – right?

I have, over 15 years, assembled a wicked collection of disastrous research. A common occurrence is the mindless application of the above formula. A company that previously lost 20c/share and now loses 40c/share has clearly lost twice as much – yet is shown to have grown earnings by +100% [(-40)/(-20) –1 = +1)]. By confusing the numeraire, commentators often report a currency to have declined by more than 100%.

Technical analysis, or charting, has its protagonists and its antagonists. What's not in doubt is that it's a field that unashamedly plagiarises mathematical and statistical nomenclature, seemingly to impart a greater degree of academic respectability, if not honesty. Slow and fast stochastics spring to mind. Their form, as used by chartists, bears no resemblance to their true meaning.

Many analysts are fiendishly accomplished at data mining, a pernicious technique of crunching numbers, putting aside the slew of disappointing results and boldly parading the minority of attention-grabbing "findings".

On the theme of data manipulation, our topic for this issue elicited a number of fascinating reports on asset managers' performance tables and on the hedge fund industry.

It's well documented that league tables have a remarkably low success rate in determining future investment performance. Nonetheless, they beguile many investors and therefore carry (misplaced) economic influence. Under these circumstances it's disappointing but not surprising that many sharp minds have been exercised on how to smooth wrinkles. In this issue you'll find very interesting material on this and on futile aspects of these tables. The burgeoning hedge fund industry is also investigated for the authenticity of some of its collective claims.

Contemplate the vast and baffling array of data that influences any share price and it becomes apparent that the human mind simply can't process all of it adequately. The availability of appropriate mathematical and statistical tools is essential. This issue of *Collective Insight* implores us all to apply some care in the use and interpretation of these techniques.

The Advisory Committee thanks everyone who submitted articles. We would also like to encourage all potential authors to start considering our next topic: Risk.

We welcome articles that can advance the debate around what risk actually means to participants in the investment process. What kind of risk measures or models provide the most useful insights? How can investors usefully employ the concept of risk budgeting in structuring their investment decisions? Publication date is set for November but please feel free to contact us at collectiveinsight@mcubed.co.za for further details on editorial content and sponsorship.

Torturing the statistics until they confess

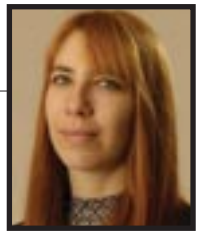

BY CHANTAL VALENTINE
Coronation

Coming up with a number is easy; interpreting it meaningfully isn't.

ANYONE who has spent any amount of time reading economic journals will have come across a number of instances where people with PhDs in econometrics and many years' modelling experience criticise complex research published by other people with PhDs in econometrics and many years' modelling experience.

Unfortunately, even far simpler representations of statistics can be misrepresented and criticised. This report deals with a few of the more basic issues, from the perspective of a user of economic statistics in South Africa.

There are two ways in which economic data tends to be presented to SA's investment community. The first is by pretty graphs, which show good "eyeball" relationships (only sometimes with any attempt to actually establish a statistical relationship between the variables presented). The second is a bombardment of seemingly impressive correlations between variables.

Playing with pictures

Charts showing seemingly good relationships between variables are the default method to present theories about the way economies and markets work. The two questions to ask when confronted with charts are: Why has the analyst chosen the timeframe that she/he has to present the data? And why has the analyst chosen to use particular Y-axis scales?

The first-mentioned in particular is subject to abuse. For example, I've seen a chart which "proves" that US Federal Reserve Bank cannot influence inflation, by illustrating high inflation rates in the US.

However, that chart starts at the first oil price shock in 1973 and ends before the first inflation hawk, Paul Volcker, took over at the Fed. It tells us absolutely nothing about how the Fed currently reacts to inflation.

Sometimes there are valid reasons to use timescales that are shorter than the available data – for example, if there's been a structural change in the variables being measured (indeed, that would be a good reason to exclude, for example, the above-mentioned time period when looking at how the Fed responds to inflation).

Often, however, it's the case simply that longer timescales either don't suit the analysis or do not "look" as good (which is why the back-up of a correlation coefficient is always useful).

Chart 1

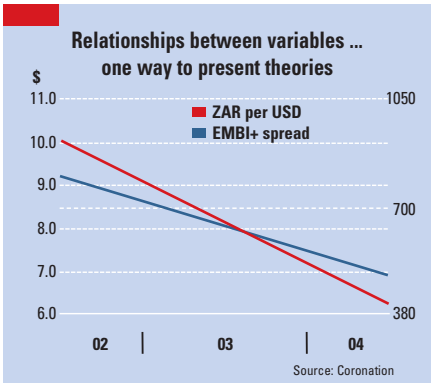

For example, Chart 1 shows the relationship between the rand/US dollar exchange rate and the JP Morgan Emerging Markets Bond index spread since mid-2002. The chart looks great. One might conclude that these two will continue to move in a step fashion and would imply that one can simply expect the relationship to continue, the second chart would warn that one needs to be aware of the type of circumstances that could see the relationship break down. The implication is very different.

Chart 2

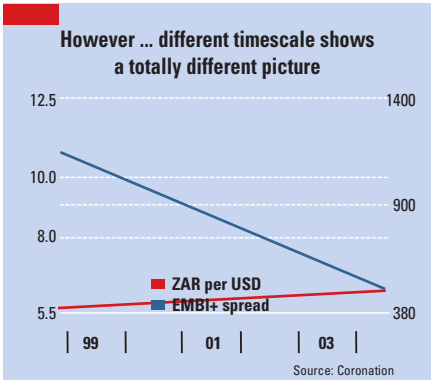

However, Chart 2, which shows the same relationship on a more extended timescale gives a very different picture. There are long periods of time when there's very little relationship between the two and periods when the correlation is even negative rather than positive.

Furthermore, on these different Y-axis scales even the good relationship of the past two years does not appear quite as good as it does in the chart above (of course, the statistical relationship remains the same).

It so happens that in this instance there are logical reasons to explain the periods of breakdown in the relationship. However, while the first chart watch carefully what global emerging market bonds are doing to help predict the rand.

Two charts of the same variables can lead to very different conclusions simply depending on the timescales chosen.

Going beyond "eyeballing" to examine the statistics, brings the differences home even more powerfully. In Chart 1 the correlation coefficient is an impressive 0,96. In Chart 2, it's –0,14! Not only is the relationship far weaker than that presented in the first chart – indeed, far too weak to be considered statistically useful – it now presents itself as a negative rather than a positive correlation. So, two charts of the same variables can lead to very different conclusions, simply depending on the timescales chosen.

Of course, sometimes timescales ▶▶ 6

3 can be too long. Structural changes in economies and markets change relationships. For example, in an environment of structurally lower inflation, a given GDP growth rate will be associated with a lower level of inflation than history may suggest.

Or one may be presented with data that one's told will return to a 40- or 50-year mean – even though there may have been significant structural changes over that time that will have led to a change in equilibrium value of the variable. (And it's sometimes claimed that a variable will revert to a mean when there have been no tests conducted to even establish whether or not it is, in fact, mean-reverting in the first place.)

Analysts use statistics much like a drunk man uses a lamp post: more for support than illumination.

Correlation causes. . . conclusions?

A good correlation by itself means nothing. A well-known study in Sweden found a high correlation between stork sightings and the number of births in certain regions of the country. Do we conclude that storks bring babies? Studies in Britain have found that asthma prevalence has risen as computer chip speed has improved. Do faster computers cause asthma?

Despite the fact that the first thing drummed into any economics student is that "correlation does not mean causation" there are still too many instances where a correlation is presented as some kind of "proof".

If the analyst cannot explain the relationship, one has to conclude that the correlation is spurious. That's data mining – simply looking for good statistical relationships without an underlying justification for the relationship. These analysts "use statistics much like a drunk man uses a lamp post, more for support than illumination".

Simply put, if one doesn't understand what the theoretical underpinning of a relationship is, there's little reason to be able to trust it to continue in the future.

There are often correlations that are good only because both variables are dependent on a third variable. A pertinent example is the relationship between the rand/US dollar exchange rate and that of the Australian dollar. The correlation is good – but not because the rand is dependent on movements in Antipodean currencies (or vice versa).

Rather, both countries are commodity exporters and their currencies are therefore influenced by commodity prices; and both are high-yielding currencies in a global context and therefore are influenced by global risk appetite.

Incidentally, there are also reasons for the occasional good relationship between the rand and the EMBI spread. Using the EMBI spread – or, perhaps, the Australian dollar – as a summary of, or proxy for, global factors that affect the rand is fine as long as it's understood that this is the case and that one's not being presented with some kind of "causal" link.

For causal relationships one should rather use – for example – commodity prices directly as an input into a rand model and understand why one's doing so.

Finally, models are sometimes presented with umpteen variables and amazingly high R-squareds. As more variables are added, that will raise the R-squared anyway – even if they make no economic sense in the equation.

Another problem – one that seems to be fairly common in some models presented in SA – is that some of the explanatory variables are themselves correlated (multicollinearity), which will bias the R-squared upwards.

For example, one might use the rand and inflation to explain interest rates – but there's obviously a relationship between the rand and inflation in the first instance. The textbooks tell us we can adjust for these factors; in practice it's not always clear-cut.

These issues aside, it's also important to understand what the purpose of the model is. Most are presented from the perspective "this is the best historical fit, therefore it's the best model to predict with".

First, only some of these are provided with an out-of-sample history to show that they are actually good at predicting. Second, one has to watch for "garbage in, garbage out": a historical relationship may be good but how trustworthy, for example, is the commodity price forecast used as an input to a rand forecast model?

Finally, a model showing a high R-squared may be excellent at explaining the past – but the more variables involved, the greater is the difficulty in using it to forecast, simply because one then has to provide forecasts for each of those variables to forecast the dependent variable. In practice, that leads to a high degree of forecast risk.

The bottom line is that it's very easy to manipulate charts and numbers. However, it's important to know if the statistics back up the proposed relationship and equally important to know if the relationship makes economic sense in the first place.

Correlation does not mean causation... if the analyst cannot explain the relationship it's probably spurious.

The analyst should be able to explain why he's chosen the variables and the timeframe and what kind of factors could cause a change in the relationship. Coming up with a number is easy; interpreting it meaningfully isn't. ■

CHANTAL VALENTINE

VALENTINE joined Coronation a year ago. She's responsible for formulating the macroeconomic view and fixed interest strategy, as well as contributing to asset allocation decisions.

She studied economics and finance. Prior to joining the asset management industry, she gained broad-based experience and has worked in academia, a mining house, a large bank and stockbroking over the past 14 years.

She was rated the top analyst in both economic trends (domestic) and fixed interest securities while on the sell-side. ■

Performance surveys may be hazardous to your wealth

BY ANN CABOT-ALLETZHAUSER
m-Cubed Asset Management

I AM a known activist. This isn't about giving peace a chance or saving the whales. It's about performance surveys and why they should be banned from all trustee meetings.

Let's be clear here. I'm not against interrogating and assessing manager performance as a principle. We need to gain insight into whether a manager's performance came from his/her deliberate bets (a sign of a good manager) or luck.

But that exercise is a far cry from taking every fund manager in the book and throwing them all into a rankings table on the pretext that we'll gain some useful information on which to base our manager selection (or, worse, rejection).

The travesty is that we seem susceptible. Show us a performance ranking and we cannot avoid coming to a subliminal, if not overt, conclusion that the guys on the top are "good" and the guys on the bottom are "bad".

March 1998

Top performing Managers – 10 years	Bottom performing managers – 10 years
Nibam	RMBAM
Sanlam	Allan Gray
SCMB	Foord and Meintjes
Norwich	

Source: Alexander Forbes

June 1998

Top performing Managers – 5 years	Bottom performing managers – 5 years
BoE	Foord and Meintjes
Fleming Martin	Allan Gray
Norwich	
Fedsure	

Source: Alexander Forbes

Too cynical am I? Jot down who you think South Africa's top fund managers are. Now, make a list of those on the bottom. Check your list against the performance rankings from 1998.

Think five years is too short? Let's take 10 years:

Note that the majority of the "winners" have since vanished. For many of these managers, their poor performance in the post-Crash environment resulted in assets leaving in droves, making them vulnerable to takeover by other managers.

As for yesterday's losers. . . well, does anyone remember who faced dire straits back then? Clearly subsequent performance resulted in selective memory.

▶▶ 7

7 **What has this exercise *not* achieved?**

- It hasn't proven who's the most skilful manager.
- Or who is the least skilful manager.
- Or which is the best asset management company.
- And certainly not which managers will perform the best in the future.

What it does illustrate is that there appears to be no "information content" in performance numbers relating to manager skill.

Let's try to understand why this is so.

Setting out the problem

In the Eighties the Magellan Fund, run by Peter Lynch, emerged as one of the top performing mutual funds of all time in the US, a remarkable fact because it was also one of the largest funds in the US.

Out of 13 years ended in 1989 it outperformed the S&P500 11 times. This should have convinced the sceptics that the man was a genius.

But Alan Marcus pointed out in his essay *The Magellan Fund and Market Efficiency*[1] that such a performance outcome provided no basis for such a conclusion.

Assume we were to hold a coin-flipping contest for 50 fund managers. We accept that coin-flipping is a purely random event. Still, the probability is greater than 40% that the winning "flipper" would flip heads 11 out of 13 times.

Was Lynch just a good "coin-flipper" then?

Marcus does identify other research he conducted that suggest that Lynch indeed did have skill, but none of his analysis was based on survey outcomes.

Challenge one: How many years of performance do we really need?

The correct answer is still hotly debated, but simply put, it's not three, five or even 10 years. The answer is that it's a function of the amount of times a manager gets to make an active investment decision. If we're evaluating a manager's asset allocation decisions alone, (typically one or two switches a year), we'd probably need more than 120 years.

By contrast, the graph seems to suggest that stock selection requires less time to assess. But that's deceptive. A good stock-picker doesn't just pick the right shares. Shares need to be properly weighted in a portfolio if the manager is going to "beat" their benchmark.

This is the exercise that massively complicates the skill assessment. Research by academics Clarke, De Silva and Thorley have argued that 91% of performance outcomes are a function of the mandate constraints imposed on long-only portfolios and only 9% is a function of the manager's stock selection skills.[2]

As such, no amount of years may be enough. The "noise" created by the mandate constraints effectively swamps the information content on manager skill.

Simplistically, one might argue that that the number of years needed to achieve a 95% confidence level is anywhere between 19 and 75 years. Not very practical.

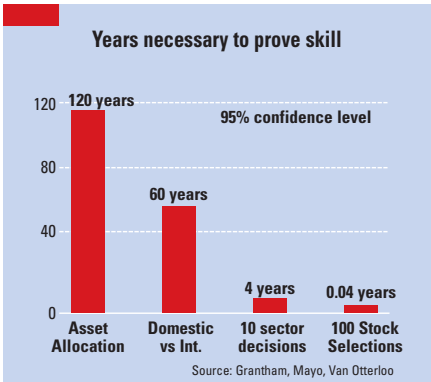

Challenge Two: What are performance surveys actually measuring?

If the information content on manager skill is so poor in performance surveys, then what explains the typically large variation of performances? The following summary is ordered in accordance to their contribution:

- Different strategic asset allocations.
- Different investment styles.
- Different appetites for risk.
- Different timeframes for investing.

In aggregate, these factors explain between 95% and 100% of performance variability.

Note that "skill" hasn't even made the list. It doesn't mean that it's absent – it's just that we'll probably never really know whether it's there...or not.

There's no question that survey providers are now moving quickly to ensure that more qualitative insights can be teased out of performance.

There's a growing awareness that if performance comparisons are going to have any meaning, then either apples must be compared to apples or statistical measures must be applied that render differences between mandates, risk and benchmarks less meaningful.

But either one of these routes is still fraught with difficulty.

Building a better 'apples to apples' approach

By including both the manager's benchmark performance alongside his/her total performance, we can now deduce how much added value the manager delivered against his/her own mandated target.

Even if we can't determine whether outperformance was a function of skill, we at least then know the manager delivered what was promised.

More importantly, though, if the benchmark performance is different from other managers' benchmarks, we know immediately that the managers' mandates are not really comparable.

But even surveys that include benchmark and active performance and separate managers into asset mix bands may still not be comparing like with like if one manager's mandate allows significant asset allocations variability and another manager is required to maintain a fixed asset allocation.

And unless the strategic allocations are identical, even the slightest variation in allocation in volatile markets can account for significant performance differences, masking skill completely.

Moreover, we still know nothing about whether we're comparing managers with similar investment styles or levels of aggression – two other factors, unrelated to skill, that explain differences in performance outcomes.

Without doubt, if the process of finding comparability reached its logical conclusion – every manager would get his/her own special category.

Is the "better statistic" route really a better option?

Statistical measures do claim to get around problems of comparability. Particularly popular additions are:

- Sharpe ratios.
- Sortino ratios.
- Downside risk measures.
- Information ratios.

The first three measures attempt to assess which manager was able to achieve the highest level of return or absolute performance at the lowest level of risk or destruction of wealth.

While these measures are interesting as a concept, (who wants to see their managers do otherwise is the theory),

the measures are only viable if all the managers in question are managing to these same objectives. That's highly unlikely.

Different timeframes will require different asset classes. Some managers will employ more volatile benchmarks than others (such as the JSE All-share to ensure that their strategic benchmark is consistent with their asset/liability modelling exercise.

The third statistic information ratio, calculates the ratio of the manager's active return above his/her own benchmark, against the risk he/she had to assume against that benchmark. This allows us to compare managers managing different benchmarks with different levels of aggressiveness.

As with the first three statistics, results can be enormously variable over time periods of less than five years. Compare a value manager's information ratio against a growth manager's information ratio in the period since the crash of 1998 for a case in point.

Because we have no meaningful equity benchmarks that reflect the differential performance of these two investment styles, the dramatically different information ratios say nothing about the skills of these two different sets of managers. They simply reflect which investment style worked over the period being measured. Once again, the issue of manager skill has eluded us.

The real abuse is in the application

Perhaps the most compelling reason to jettison performance surveys has more to do with how surveys invariably influence the manager selection and rejection process.

Research in both the US and UK has highlighted the fact that one of the most expensive decisions trustees can make is to replace a fund manager. Estimates range from 5% to 8% – a cost that a new manager would find exceedingly difficult to recoup. Given that we now know that survey rankings provide no insight into skill, it's disturbing that manager replacement is invariably failure to outperform their peers.

Consider this additional reality.

With the increase in focus on risk, most boards choose to diversify their pension fund managers. As Nobel Laureate Bill Sharpe pointed out more than 10 years ago though, the moment you use more than one manager, relative performance against the pack becomes totally meaningless.

Whether your managers perform differently enough from one another to make the diversification exercise meaningful is the critical decision.

Trustees would probably be better off just picking one of their managers from the top quartile and one from the bottom quartile. Makes you think, doesn't it?

1 Marcus, Alan J, *The Magellan Fund and Market Efficiency,* The Journal of Portfolio Management, 1990.

2 Clarke, Roger, De Silva, Harindra and Thorley, Steven; *Portfolio Constraints and the Fundamental Law of Active Management. Financial Analysts Journal*, September-October 2002. ■

ANNE CABOT-ALLETZHAUSER

CABOT-ALLETZHAUSER is Chief International Officer of m Cubed Asset Management. She's been an asset manager for 25 years, managing pension fund assets in North America, Japan, Britain, Europe and SA – her primary focus being risk management and quantitative modelling.

In SA she was responsible for the development of the multimanager management approach that's now m Cubed's hallmark. Cabot-Alletzhauser also oversees all its investment-related activities. ■

Miracle cure

Even Sharpe tools can give blunt results.

BY SHAUN LEVITAN
Alexander Forbes Asset Consultants

I RECENTLY read the results of an academic study that found that echinacea (a purple flower) is not successful in preventing the common cold. That's not the first study to arrive at the conclusion that there's no "quick fix" to preventing one of the oldest problems in the medical world.

However, I know of colleagues, friends and even family who swear by this herb's ability to keep them well through the winter months – despite most of them spending at least one day in bed recovering from flu.

In the investment arena we, like doctors, face many challenges. One challenge that receives a lot of attention is the selection of investments or investment managers. However, there are many of us who have embraced a quick-fix remedy to solve this important problem – we call it the Sharpe Ratio.

For those of you who have not yet been introduced to this "miracle cure", the Sharpe Ratio is a statistic that considers the excess performance of an investment over a risk-free rate of return and adjusts the figure by the risk of the investment. Importantly, the risk measure used is the standard deviation of the investment's returns.

There is a lot of intuition behind this number. After all, most investors desire higher returns than can be obtained by investing in so called risk-free investments, such as bank deposits and short-term Government bonds. However, most investors accept that additional returns can only be generated by assuming some risk. The Sharpe Ratio addresses this by measuring return in a context of the risk taken rather than simply as an absolute rate of return.

In isolation, the Sharpe Ratio has no meaning. An investor would not glean much information if all they knew about investment portfolio A was that it had a Sharpe Ratio of 1. However, if portfolio B had a Sharpe Ratio of 2 then, on a risk-adjusted basis, it should be chosen.

Different investments are marketed

10

9 as superior investments on the basis of having higher Sharpe Ratios and that's where the use and abuse begins. To get to the root of the problem, we need to go back to the original thinking behind the Sharpe Ratio.

A brief history lesson

In 1952, Markowitz revolutionised risk analysis with a paper on portfolio selection. He claimed that investors choose portfolios by only considering the portfolio's average return and risk (as measured by standard deviation). Therefore, a risk-adjusted return measure that used the portfolio's average return and its standard deviation had a strong theoretical foundation.

The Sharpe Ratio can be graphically depicted as seen in the diagram (right). The portfolio's expected return is plotted on the vertical axis and the portfolio's risk is plotted on the horizontal axis. The Ratio is the slope of the line between the risk-free rate (we assume 5%) and an investment in risk and return space.

The graph demonstrates that for the same increase in risk (as demonstrated by the yellow horizontal line), portfolio B delivers a larger increase in expected return. That's captured by the slope of the line (the Ratio) and, therefore, the portfolio with the steeper slope (higher Ratio) should be selected.

However, the underpinnings of Markowitz's theory require the return distributions of the assets to be normally distributed. Many asset class returns don't exhibit normal distributions. Empirical evidence demonstrates that equity return distributions exhibit a greater likelihood of extreme movements than is implied by a normal distribution.

Other studies have concluded that significant deviations from normality have been found in emerging markets, portfolios with derivatives, small companies, private equity and hedge funds.

In the same way that you need to do more than just look at a patient before diagnosis, in more than just the expected return and standard deviation need to be considered before choosing an investment.

But at least it measures risk?

The Ratio measures risk by considering the excess return of an investment over a risk-free rate and adjusting for the standard deviation of returns. The most controversial of its components is the use of the standard deviation as the risk measure.

It can be argued that the true risk to any investor is the risk of not meeting their objective. Some investors may not be concerned with variability of returns as much as the possibility of adverse variability – ie, it's the possibility of underperformance that may be more relevant for some. Standard deviation results in excessively positive deviations being penalised in the same way as excessively poor deviations.

Standard deviation also focuses on variability of returns relative to a portfolio's average return rather than an investor's target.

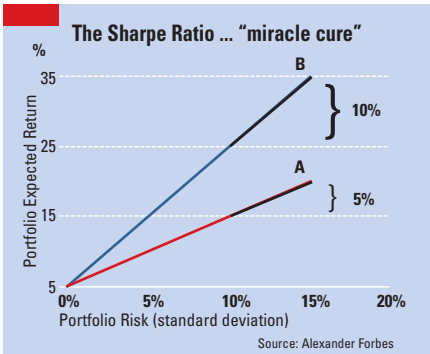

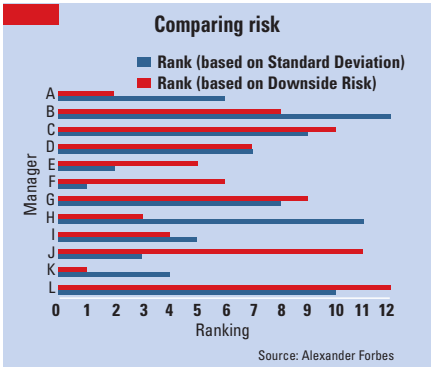

The graph above compares the "risk" of the current 12 managers of the Alexander Forbes Large Manager Watch Survey, as measured by their monthly returns for the year ended 30 April. Two definitions of risk are used: standard deviation of returns and downside risk only. A ranking of "1" indicates the lowest risk for a particular measure.

Interestingly, not one manager had the same ranking on both measures of risk. Manager H had a ranking of 11 (out of 12) for risk when standard deviation was the risk measure used, but climbed up to 3rd position when risk was measured relative to a target of positive returns. Manager H could therefore be suitable for a conservative investor seeking to preserve capital yet their risk-adjusted performance would look poor if the Sharpe Ratio was used.

Does that mean the Ratio should be discarded altogether? On the contrary. In the words of William Sharpe: "There's much to recommend a measure that at least takes into account both risk and expected return over any alternative that focuses only on the latter." Investors need to appreciate that each risk statistic measures a different aspect of risk. It's important to tread with care when selecting a risk measure.

"When I use a word, it means just what I choose it to mean – neither more or less" – Humpty-Dumpty in *Through the Looking Glass*, by Lewis Carroll.

Negative Sharpe Ratios

The analysis assumed that the Sharpe Ratios under consideration were positive. However, interpretation becomes complicated when a portfolio has a negative Ratio. That arises when the average return of a portfolio is lower than the risk-free rate. For example, consider portfolios C and D in the table below.

Portfolio C has a lower average return and higher risk than portfolio C and yet it has a Sharpe Ratio twice as large as Portfolio D. That presents problems for a naive interpretation of the Ratio. In this instance we're no longer comparing portfolios by their ability to generate superior risk-adjusted performance.

We need to refine our interpretation of the Ratio when it's negative. In such instances achieving a negative return with larger variability of returns is better, since the larger volatility implies that the probability of a positive return arising from this portfolio is higher.

Hedge funds

Hedge funds deliver relatively high expected returns and low standard deviations of returns but also have "attributes that are exactly opposite of those that investors desire – for example, the possibility of large losses". That means the Sharpe Ratio will systematically overstate true hedge fund performance relative to that of the standard market indices. Andrew Lo found that the Sharpe Ratio for a hedge fund can be overstated by as much as 65%.

That was the case with Long-Term Capital Management. In 1998 it boasted a Ratio more than nine times that of the

Portfolio return/Risk-free rate

	Average return	Standard deviation
Portfolio C	-20%	15%
Portfolio D	-12%	5%
Risk-free rate	5%	5%
Sharpe Ratio	-1,7	-3,4

long-term performance of the S&P500, due to the hedge fund demonstrating low but consistent returns.

The consistency results in a low standard deviation of returns and a very high Ratio. However, the use of gearing and derivatives by hedge funds means that there's a significant risk of a fund being completely wiped out. The Ratio doesn't capture that risk. Those funds employing such strategies without the loss in the return data would exhibit a very high Sharpe Ratio. The same phenomenon was evidenced by the IT boom – the possibility of a sector crash was not factored in.

Those investments that produce many small returns and the occasional large loss will appear to have a very high Ratio, just before the loss occurs.

More worrying is that Goetzman et al (2002) showed that it's possible to manipulate the Ratio using derivatives. As such there's a danger of relying on one number – especially when that number can be tweaked.

Not only is the Sharpe Ratio used to describe hedge fund performance in SA, it's also being used by international hedge fund indices like the CSFB/Tremont hedge fund indices. These surveys show the Ratio for each hedge fund participating in the survey. However, different hedge fund strategies have different characteristics regarding strategy, gearing, the use of derivatives and liquidity. Their return distributions will differ immensely and their Ratios will simply not be comparable.

CONCLUSION

So if there's no cure for the common cold, do doctors have no advice for us in preventing it? On the contrary; common sense suggests dressing warmly, eating healthily and exercising regularly. In much the same way, any investment decision should begin with a thorough qualitative analysis. Enhanced investment decisions can then be made by introducing a quantitative risk adjusted measure – as long as we understand its strengths and weaknesses.

"The statistical properties of hedge funds index returns and their implications for investors," Chris Brooks & Harry Kat, November 2001, Working Paper, Cass Business School, City University.
The statistics of Sharpe Ratios, Andrew W Lo, 2002, *Financial Analysts Journal*, 58(4):36-52. ■

SHAUN LEVITAN

LEVITAN is a senior consultant at Alexander Forbes Asset Consultants, where he heads asset liability modelling. His responsibilities include providing strategic advice to clients on asset allocation and objective setting, portfolio risk management research as well as employee share option professional services. ■

Return to normality

Watch out for those assumptions.

BY PETER URBANI
KnowRisk Consulting

MANY of the decisions made by investment professionals are based on the assumption that returns are normally distributed.

Most of the time that's sufficient to explain range of returns typically experienced. However, there are two instances where the normal distribution fails to adequately explain what's going on.

The first is when there's excess skewness and or kurtosis or peakedness in the distribution of returns. Such "lepto-kurtosis" (Greek for "Thin-arches") is commonly seen in financial time series. Similarly, positive or negative skewness is commonly seen in hedge funds and emerging markets, such as SA. The combined effect of the excess kurtosis and of negative skewness is to fatten the tails of the distributions. This means that losses when they occur tend to be larger than expected.

That brings us to the second failing of the normal distribution: namely, that it fails to adequately explain the size and frequency of extreme events.

That's because the ends of the normal distribution tail off too quickly to adequately describe the size and frequency of extreme events. Under the normal distribution, 68,3% of all observations should fall within one standard deviation of the mean, 95,4% within two standard deviations and 97,7% within three standard deviations. Thus under the normal distribution there's only a 0,3% probability of observing a greater than three standard deviation event.

If we consider the monthly returns from the JSE Securities Exchange's all-share index since 1925, we get a monthly return of just over 1% and a monthly standard deviation of around 5%.

Under the assumption of normality that means there's only a 0,3% probability of observing a monthly return of more than +16% or less than -14%. Out of the available 948 data points we should thus observe only about three periods when the returns were outside that range.

In fact, there have been 11 months when the monthly return on the all-share index has been less than -14%. That's almost nine times as often as predicted under the normal distribution.

That's not so evident when looking at the distribution of the JSE against the normal distribution on a normal scale. But when viewed on a log-scale the magnitude of the underestimation becomes more obvious.

The assumption of normality was made by Harry Markowitz in 1952, when he chose to use variance as the measure of risk. Variance is calculated as the sum of the squared differences from the mean or average. More typically, we take the square root of this number to get the standard deviation.

The normal, Guassian or Bell-shaped distribution is a symmetrical distribution that can be fully described by its mean and standard deviation and therein lies its problem.

Markowitz's reasons for using variance were largely pragmatic and due to the lack of computing power at the time, which required him to perform most of the necessary calculations by hand.

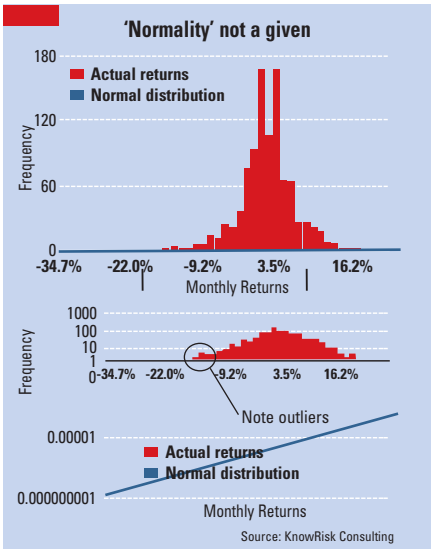

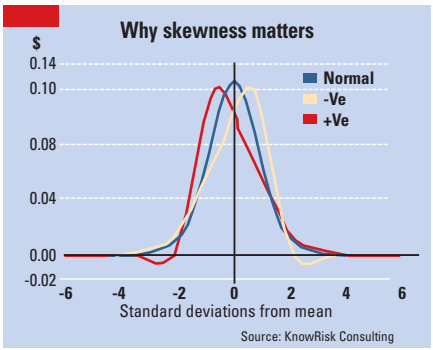

He later went on to say that semi-variance or downside deviation might be a more realistic representation of the way in which investors perceive risk – ie, that they feel more pain from the same level of losses than the pleasure they get from the same level of gains.

In a more recent paper Markowitz says that modern portfolio theory is a "normative" as opposed to a "positive" theory, in that it describes how investors ought to behave rather than how they necessarily do behave.

In the case of the JSE – and most markets – the mean and standard deviations alone are not enough to describe the shape of the distribution of returns. We need to also take the higher moments of skewness and kurtosis into account.

There are a number of ways to do so. For the purposes of this report I'll describe four such methods and compare the value at risk (VaR) and expected shortfall (ES) measures for each.

The first such "fat tailed" distribution we could use is the Student T distribution. In fact, its use is recommended whenever dealing with a sample rather than the full population of observations.

The second method is to match the moments of the distribution using a mixture of normal distributions.

The third is to modify the normal distribution to take account of the excess skewness and kurtosis using the Edgeworth expansion.

The fourth method is to use an extreme value type distribution. Here we use the generalised Pareto or Gumbel distribution.

We use the annualised statistics of the JSE all-share from 1925 to 2003 from the Firer and McLeod database. That gives the following mean, standard deviation, skewness and kurtosis.

We calculate the parametric VaR and ES using the four distributions and then compare the number of exceedences in the actual distribution to get the following results.

Remember the VaR at 99% is that minimum level of loss that we can expect 1% of the time. Given 948 observations we would expect to have 10 observations in excess of that amount.

The ES is the average loss that's

Table 1
JSE All Share Index (1925 - 2003)

Mean	13.13%
St Deviation	17.87%
Skewness	-0.65
Kurtosis	4.57
Degrees of freedom	3
Confidence level	**0.99**
Number of observations	948

Table 2

Value at Risk (VaR)	VaR	No. Obs Below	%
VaR Normal @ 99%	-28.44%	39	4.11
VaR Student T @ 99%	-33.72%	25	2.64
VaR Vol Mix @ 99%	-41.61%	16	1.69
VaR Cumbel @ 99%	-42.93%	16	1.69
VaR Cornish Fisher @ 99%	-48.24%	12	1.27

Table 3

Expected shortfall	ETL	No. Obs Below	%
ETL Normal @ 99%	-34.50%	25	2.64
ETL Student T @ 99%	-54.12%	5	0.53
ETL Vol Mix @ 99%	-55.16%	7	0.74
ETL Cumbel @ 99%	-56.45%	8	0.63
ETL Cornish Fisher @ 99%	-65.41%	3	0.32

realised when the VaR is exceeded. It's also called the conditional value at risk or extreme tail loss.

As can be seen, both the VaR and ES calculated using the assumption of normality significantly underestimate the number and size of tail or extreme events. Fortunately for the purposes of this report, we've used the 99% confidence level, which means that annual losses of up to 70% should only be experienced once in every 100 years. At 95%, or once in every 20 years, investors should expect a loss of up to 40% – still almost double that suggested by the normal distribution.

Table 4

	A SD	B Var	C SD
Mean	0.0	0.0	0.0
Var	1.0	1.0	1.0
SD	1.0	1.0	1.0
Skew	1.0	0.0	-1.0
Kurt	3.6	3.9	3.6
CL	0.99		
Value at Risk (VaR) for portfolio of R1m			
Assuming Normality			
	-R23,154	-R23,263	-R23,020
	-2.32%	-2.33%	-2.30%
After adjusting for Skewness & Kurtosis			
	-R23,165	-R25,297	-R37,053
VaR	-2.32%	-2.53%	-3.71%
ETL	-2.64%	-3.09%	-4.60%

It's disturbing to me that hedge funds are typically marketed on the strength of their lower than market risk as evidenced by their low standard deviations when in fact such funds are particularly poor candidates for having their risk measures using a normal measure. That's because they're supposed to have positive skewness (more up than down months) by design, yet in practice tend to be negatively skewed. Thus any symmetrical measure of risk used for a hedge fund, risks under- or over-stating its true level of risk.

In conclusion, investors should be wary of marketing hype and always remember the sting in the tail. They should take note of higher moments, particularly negative skewness and excess kurtosis, and should demand to see downside risk measures for all investments.

In the example of three hypothetical funds, all three have the same mean and standard deviation. However, fund A is positively skewed while fund C is negatively skewed.

If the VaR for each of these three investments is calculated under the assumption of normality using just the means and standard deviation, then the downside risk of the three funds is identical save for rounding errors.

However, if we take the excess skewness and kurtosis into account by using the modified VaR formula, it becomes clear that fund C has significantly more downside risk than either A or B.

That's because of the negative skewness of its distribution relative to that of the other two funds. As you move further out into the tail of the distribution, the degree of difference becomes even more marked, as evidenced by the extreme tail loss, also known as ES and conditional VaR of fund C. The JSE all-share index has long-run skewness and kurtosis of -0,65 and 4,57 respectively.

References:

The Legacy of Modern Portfolio Theory, Frank J Fabozzi, Francis Gupta and Harry Markowitz, *Institutional Investor Magazine*.

PETER URBANI

PETER Urbani is a risk consultant for KnowRisk Consulting. He was previously head of investment strategy at Fairheads Asset Managers and prior to that senior portfolio manager at Commercial Union Investment Management.

Behind the hedge fund numbers

Investors must understand what they are reading.

BY DR MICHAEL ROBERTSON
Edge Investments

Intrigued by stable returns and attractive diversification properties, investor demand for hedge fund exposure is on the increase. Yet some high profile hedge fund debacles have left investors concerned that the array of risk and return diagnostics don't describe the full extent of hedge fund risk.

Inside the indices

WITH a paucity of individual fund track records to analyse, investors and analysts have few data sources other than the growing range of hedge fund indices. These are widely used to extrapolate expected returns and produce risk estimates across the range of hedge fund strategies.

Unlike unit trust funds, hedge funds report voluntarily to one or more industry databases that compile indices from those funds that choose to be included. As a result, hedge fund indices represent only a general proxy for the industry rather than a comprehensive measure of it.

The first problem is that no databases really contain complete records of all operating funds in the industry. Even more problematic, hedge fund databases are unlikely to contain information on those funds that have ceased operation. The reason, for which there's no apparent remedial solution, is the lack of incentive for poor performing funds to report to the data vendors.

The result is a general overstatement of return from the upward survivorship bias in the index. The magnitude of the bias is impossible to determine since quantification requires the universe of hedge funds to begin with.

However, several international studies estimate that the bias that occurs from failing to have a complete record of the funds that have ceased operations, could mean that hedge fund indices may overstate hedge fund returns by as much as 2% to 3%/year.

To date that doesn't appear to have bothered too many data vendors. Anyone following the rapid development of the industry will have noticed a plethora of new efforts to produce and publish hedge fund performance indices. To further complicate matters, not only are these performance indices upwardly biased, but there's also wide variation in the index returns calculated from different databases.

Two separate international studies report the annual return for aggregate global hedge fund indices ranged from a low of 13,6% to a high of 18,1% for the period January 1995 to April 2001[1]. Another two studies find a difference of 1,5%/year over the 1994 to 1999 period between two widely used aggregate indices[2].

Additionally, returns on indices that appear to represent a common investment strategy were not highly correlated across different databases[3].

So investors should be extremely wary of any marketing justification that employs these unquantifiably biased indices to make its case.

Substandard deviation

The biased indices problem is further compounded by the fact that many of the conventional measures of risk have simply been transplanted from traditional asset management. For better or worse, doing so makes the unqualified assumption that what's useful for describing the characteristics of traditional investment portfolios applies automatically to hedge fund portfolios.

Investors should be extremely wary of any marketing justification that employs these unquantifiably biased indices

Among the most abused statistical measures in the hedge fund space is standard deviation. Intended to be a measure of dispersion, standard deviation is by far the most commonly used measure to describe the volatility of investment return. A small standard deviation corresponds to a low probability of extreme gains or losses, whereas a large standard deviation corresponds to high probability of extreme gains or losses.

The measure is most useful when the frequency of periodic returns represents a bell-curve shaped pattern. However, most hedge fund returns are not symmetrically distributed and display a skewed profile. Chart 1 illustrates the skewness in the return pattern of a sample of SA hedge funds, calculated using a monthly growing cross-sectional average of all funds in the database from July 1998 to March 2004.

The distribution of hedge fund returns in Chart 1 has, for example, a monthly standard deviation of 3,6%. If returns were symmetrically distributed this suggests that only in 5% of the months is return predicted to be less than -3,8% or more than 10,6%. However, observations in the tails of Chart 1 indicate that the number of extreme losses or gains, are greater than that implied by the standard deviation.

Chart 1

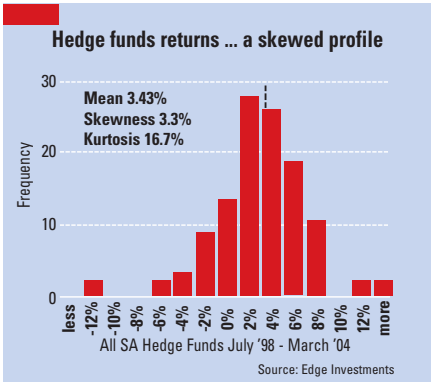

In addition, there is a further negative impact on the distribution of hedge fund returns introduced through the holding of illiquid securities, known as "thin trading bias".

Since some securities trade infrequently, they appear to have a low volatility. This results in the security being assigned a low numeric of measured risk, which often is not representative of the true economic risk inherent in the security. The bottom line is that measured standard deviation of return heavily underestimates hedge fund risk and is not fairly representative of the

range possible return outcomes.

While this statistical subtlety may not yet have occurred to investors, it ought to come as no surprise that the measured volatility of a vanilla hedge fund return series is always likely to be lower than that of an ordinary asset class. Using the inordinately low standard deviation of hedge fund returns as evidence, is to oversimplify the multidimensional risk picture that's intrinsic to hedge fund strategies.

Perhaps for reasons of familiarity or habit, the hedge fund industry in SA and overseas stubbornly persists with using symmetric measures of return dispersion to communicate risk to investors.

Down with drawdown analysis

If standard deviation fails to properly convey the magnitude of downside deviation in hedge funds, then perhaps alternate methods, such as the widely-quoted drawdown analysis, are better.

Drawdown analysis is a rudimentary peak-to-trough calculation that has intuitive appeal. Empirically easy to calculate, drawdown analysis is a theoretical lightweight in the risk stakes.

For example, there's very little that can be said about the economic risk of a hedge fund based on the empirical fact that three years ago it returned -4,1% in six weeks. For that reason, drawdown analysis has failed to attract the levels of serious research devoted to competing measures, such as volatility, value-at-risk or Sharpe Ratios.

Despite its simplicity of approach there are problems – not least of which is the existence of a weak proportional relationship between the magnitude of any given drawdown and the length of the track record of the fund. That's simply because a fund with a 20-year track record is more likely to have endured a far broader array of systematic dangers that a recently launched fund. So unless empirical drawdown figures are calculated over similar time periods it's less than fair to compare across funds.

The not-so-uncorrelated hedge fund

Another frequently cited hedge fund statistic borrowed from traditional asset management is the correlation coefficient. Widely quoted as evidence of the diversifying nature of hedge fund returns, correlation coefficients indicate the degree to which two return series are linearly related.

Since the measured returns of most hedge fund strategies typically have a low or negative correlation with the returns of conventional securities, the common conclusion is that aggregate portfolio volatility can be reduced through an allocation to hedge funds.

Not to be overlooked, especially when dealing with econometric time series, is the issue of spurious correlation – which arises when a correlation between two variables doesn't result from any direct relation between them but rather from their relation to other variables, either known or unknown.

If one could control this unknown variable, then the correlation is likely to be substantially different. For example, not many people noticed that all the CSFB/Tremont hedge fund strategy indices (except for the short bias strategy) showed a negative return in April 2004.

It's difficult to argue that there exists no hidden and shared systematic component of return among these various "uncorrelated" strategies. At least from a probability perspective the chances that all strategies have a negative return simultaneously is just too remote for there not to be a hidden common factor.

The lurking danger is that statistics – such as standard deviation, drawdown and correlation coefficients – are often quoted on hedge funds without qualification. Though the numeric relationships may be correct, the cause-and-effect nature is largely a matter of speculation. After all it's a little known fact that for the first half of 2004 the price of degreased Himalayan Yak wool, currently about R850/kg, was in fact highly correlated with German Bund futures traded in Chicago. So if you're a believer in the hard facts, then that's a diversification opportunity for sure!

Ultimately, an investor with a reasonable economic understanding of the source of a hedge fund's return is in a good position to decide whether the associated risks – however they're defined – are worthwhile.

We must accept though, that nothing's for nothing. Hedge fund strategies provide returns only because investors are being compensated for risks that

16

15 others would prefer to lay off or not take.

1. Chris Brooks & Harry Kat: *The Statistical Properties of Hedge Fund Index Returns and their Implications for Investors.*
2. William Fung & David Hsieh: *Hedge-Fund Benchmarks.*
3. Harry Kat: *10 Things That Investors Should Know About Hedge Funds.* ■

DR MICHAEL ROBERTSON

ROBERTSON was previously a risk consultant at Barra International responsible for supporting their equity risk and asset allocation models. He currently oversees all quantitative risk management procedures and research activities at Edge Investments.

In recent years he has produced six publications in scholarly journals on financial asset pricing. He holds a BCom in economics, an MBus.Sc and a PhD in finance from the University of Cape Town. ■

A tailored approach...

BY MARK VAN DER WALT
& ANNA NYE
Russell Investment Group

. . .to risk analysis in hedge funds can provide better insights.

NEWLY developed techniques that address unique problems with measuring hedge fund risk and return, are providing valuable new insights.

Specifically, by making adjustments for asymmetries introduced by performance fees and for the effects of stale pricing, a more realistic picture emerges of the true risk and return characteristics of hedge funds. So, while many of the claims for hedge funds are indeed hyped, investment strategies founded on careful research at a grassroots level should continue to prove fruitful.

Performance fees distort the picture

The first adjustment addresses the distorting effect of hedge fund fees. These fees are at least partly related to performance. For example, it's not unusual for them to run as high as 20% of returns above a hurdle rate, perhaps US Treasury bill returns, or in some cases even zero.

Such fees – which increase with positive returns but disappear when performance is poor – introduce a kink into hedge fund return patterns that makes analysing performance using traditional statistical tools difficult. This characteristic is exacerbated because, at least superficially, many hedge funds show markedly different behaviour in favourable markets than in unfavourable ones, even when the effect of fees is eliminated.

Three researchers – Asness, Krail and Liew – tested this and found evidence of such a kink in hedge fund return patterns. Their approach was to produce separate beta estimates for bull and bear markets. They found that hedge funds in most styles had much higher sensitivity to market index returns in bear than in bull markets.

Coping with stale pricing

The second major problem arises with stale prices – those that reflect past rather than current market conditions. Some hedge fund strategies trade in less liquid securities, making it difficult to base returns on up-to-the minute prices, such as those posted by markets where transactions are conducted almost continuously. Consequently, returns calculated today might not fully reflect the day's changes in market valuations.

So, while many of the claims for hedge funds are indeed hyped, investment strategies founded on careful research at a grassroots level should continue to prove fruitful.

To overcome stale pricing, the three researchers applied "lagged beta analysis" to the returns of a variety of hedge fund styles. This analysis identifies the correlation between returns of a hedge fund in one time period and those of a market index in earlier time periods. Should there be stale pricing problems, it will show up in the form of a significant relationship between current hedge fund returns and previous months' index returns. The result is a truer picture of the market exposures a given hedge fund style represents.

We extended the research by Asness, Krail and Liew in a number of ways. Looking at funds that fall into six common hedge fund strategy classifications, we found that lagged beta analysis produced evidence that stale pricing seems to be more prevalent in distressed securities funds and in equity long/short than in hedge funds classified as pursuing market neutral strategies.

In the case of the first two strategies, the strength of the relationship between the performance of the funds and the market increases strongly when prior and subsequent market performance (leads and lags) are brought into consideration. In the case of market neutral strategies the effect is less marked.

The observation of a kink in the performance pattern was puzzling for some time. We found that, even after removing the influence of fees on returns, a marked difference in semi-betas estimated for bull and bear market periods was shown at the individual manager level. So the kink was present. However, further research showed that these results were largely explained by three monthly returns: August to October 1998. These months span the Long-Term Capital Management meltdown, which endangered liquidity in many markets and necessitated massive intervention by the US Federal Reserve.

Removing this three-month period all but removed the bull and bear market asymmetry. Statistically, arguing that the remaining asymmetry is more than the result of chance would be difficult. ■

A diet of the mind

What's in a name.

BY DEON GOUWS
Sanlam Multi-Manager

"MATHEMATICIANS won the war. . . mathematicians broke the Japanese code and built the A-bomb. . . today we bequeath America's future into your able hands."

That description of the power of numbers – and those who work with them – is the opening comment from *A Beautiful Mind*, the film that won the best picture Oscar in 2002.

Based on Sylvia Nasar's biography of mathematician John Forbes Nash Jr, it provides an account of how problem student Nash develops a groundbreaking doctoral thesis in the field of game theory.

He subsequently starts suffering from delusional thinking, characteristic of persons who are diagnosed as schizophrenic. However, it all ends happily: Nash recovers but also receives the Nobel prize for his contribution to the world of economics.

The film shows how a young Nash tries to explain – and ultimately predict – everything around him by reducing it to mathematical theorems. He models a group of students playing touch football; he builds a formula based on the behaviour of a cluster of pigeons fighting over breadcrumbs.

Later on in the film – now in the delusional phase – Nash takes his mathematical bent to the next level. He pores over newspapers and magazines and contrives to find complicated patterns in a variety of headlines, articles and advertisements. He proceeds to interpret these as secret messages with strategic relevance to the Cold War.

Students playing touch football. . . pigeons fighting over breadcrumbs. . . secret messages hidden in magazine articles.

Perfect analogies for what the typical investment analyst seems to be doing on a daily basis: relating one number to another, looking for patterns, postulating some far-fetched causality.

To illustrate, consider the following selection of headlines from Bloomberg News, all culled in one 24-hour period.

US 10-year Treasury note declines on mortgage-related sales.

US Treasury notes rise on optimism for rates and less supply.

Treasuries fall; Fed signals it's moving close to rate increase.

US notes advance as Fed damps 18

speculation rates to increase.

I still have no idea about the real reason for the eventual move. But I do recall that many more contradicting opinions were posted, each seeming to depend simply on which particular number suited the author's book at a specific time of day. Figures don't lie – but liars can certainly figure.

However, finding some tenuous relationship between two sets of numbers does seem like the lesser of two evils when compared to another of the analyst's vices – the nondisclosure of relevant information. Financial accountants have long been guilty of this: if they can't accurately measure it, they won't want to report on it – regardless of how important "it" may actually be.

Derivatives exposure, share option schemes, intangible assets. . . the beancounters have always been better at finding excuses rather than solutions for these kinds of issues. The accounting fraternity does prefer to be precisely wrong rather than substantially correct.

But let us relate this principle back to the movie. The film was crowned with four Oscars, including one awarded to Akiva Goldsman for the Best Adapted Writing. And "adapt" he did: those filmgoers who read Nasar's book will know that several aspects of Nash's life were excluded from the film, including anti-Semitic comments, homosexual experiences, an illegitimate child and the divorce from his wife.

Would any of those who loved the film have changed their mind on Nash's "nice-guy status" if these issues had in fact been included? Perhaps. . . perhaps not. We'll never really know. But one can assume that those who made the film had their reasons for leaving them out, in much the same way that analysts will focus on certain numbers whilst omitting others – even though some of these may be fundamental to the investment case.

Ster-Kinekor Home Entertainment

The story of John Nash's life as portrayed by Russell Crowe in "A Beautiful Mind".

Research reports, like statistics, have therefore often been compared to a bikini: what it reveals is interesting but what it conceals is crucial. How many results are left out of research publications on the grounds of not being statistically significant, where something did perhaps *not* happen in accordance with *prior* expectations? Is there not some valuable information in such a counter-intuitive outcome?

There's an example of the significance of something that does in fact *not* happen in the film. Recall the scene where Nash meets the young niece of his imaginary university friend for the first time. The little girl soon sets off, running around among a flock of pigeons on the adjacent lawn.

And now for the crucial bit: did you notice that the pigeons continued to sit perfectly still through all of this?

In real life, pigeons clearly don't remain stationary when boisterous children try to chase them away. So did the film producers merely make a mistake with the computer-generated images of the birds? No. Those who made the film *wanted* us to notice; they were in fact giving us a most important clue – they were trying to help us realise (even if Nash himself didn't do so at the time) that there was indeed no kid running around among the pigeons.

If something didn't happen, it probably doesn't get reported. If something doesn't get reported, we won't consider its significance. If we don't consider something – if we just block it out – we're eventually bound to fall into the traps of apathy and ignorance. As John Nash says in the film: "It's a diet of the mind." We deserve to lose money to some of the more astute market participants in the process.

As far as diets of the mind are concerned, it's probably also worthwhile to note that any reference to the Nobel Prize in Economic Sciences (as the prize awarded to John Nash in 1994 is referred to) is not actually strictly accurate.

Even though such a reference may be the staple diet that the press have been feeding our minds with, there's in fact no such thing. More than 100 years ago Alfred Nobel specified that prizes should be awarded to individuals who "have conferred the greatest benefit on mankind" in five specific areas. Alas, economics was not one of them.

However, in 1968, the Bank of Sweden instituted the "Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel". Today the award enjoys similar status as the "real" Nobel Prizes. – those in the undoubtedly more noble fields of physics, chemistry, medicine, literature and peace. (Author's note: for purposes of simplicity, this report will also refer to the economics award as a Nobel prize).

From the beginning, those responsible for selecting the Nobel economics laureates tended to focus on rigorous number-crunchers. The citation for the second ever recipient, Paul Samuelson (1970), stated that the prize was for "the scientific work through which he has. . . actively contributed to raising the level of analysis in economic science".

Similar motivations were provided for many of the subsequent awards: "For the creation of econometric models. . ." (Lawrence Klein, 1980); "For his clarification of probability theory. . ." (Trygve Haavelmo, 1989); "For a new method to determine the value of derivatives. . ." (Robert Merton and Myron Scholes, 1997).

But regardless of how much you torture the numbers – and notwithstanding the fact that a proponent may have been crowned with a Nobel prize for doing so – the investment markets will not allow themselves to be prescribed to.

And never has there been a more dramatic example of that point than the case of Myron Scholes. It's now well documented how he lost the proverbial farm as a principal of the multibillion US dollar hedge fund business Long Term Capital Management less than nine months after receiving the 1997 award.

Another slightly ironic twist has been provided by Bill Sharpe, crowned with the Nobel Prize for Economic Science in 1990 in recognition of his pioneering work in quantifying the risk of an investment. Though his original work in the Sixties was good enough for a prize, Sharpe himself has moved on: his risk measures now focus more on return objectives rather than the traditional risk-free rate.

It was not until the new millennium that the panel of judges at the Bank of Sweden would eventually recognise that markets are in fact governed by more than just numbers: in 2002, Daniel Kahnemann received the award for his work in behavioural finance. Finally, there was recognition of the fact that emotions and biases count for at least as much as graphs and statistics do.

In closing, let's return to the film version of *A Beautiful Mind*. And let's remind ourselves of the very happy ending. John Nash accepts his Nobel prize in front of a packed audience. He proceeds to make an emotional speech, focusing on a lifetime of support from his loving wife. The tears flow.

Alas. We've been deceived once more. Nash did not make an acceptance speech. Nobel laureates never do.

However, according to the book, our hero did say a few words at a small champagne party. Sylvia Nasar paraphrases his concluding remarks as follows: "Nash said that he had won for game theory and that he felt that game theory was like string theory, a subject of great intrinsic intellectual interest that the world wishes to imagine can be of some utility."

If only the typical investment analyst had similar humility and intellectual integrity. ■

DEON GOUWS

GOUWS is a chartered accountant and a CFA charterholder; he also holds an MPhil degree from the University of Cambridge. He's the CEO and CIO of Sanlam Multi-Manager International, based in London.■